LIGO - T890035-00 - R

62

List of 3/15/89

REPORT ON LIGO CONTROL SYSTEMS AND DATA ANALYSIS

Computer and Data System Requirements for Phase A Facilities and Interferometers

Yekta Gürsel Andrew Jeffries Robert Spero

13 March 1989
Version 2.3

Abstract

The computer system design for Phase A Facilities and Interferometers are outlined in sufficient detail to identify the major elements and to show how they interconnect. Data storage and analysis issues that affect off-site operations are identified.

1 Definition and Purpose of the Data Systems

The LIGO comprises three functional entities: the Vacuum System and associated facilities, the interferometers and associated control systems and monitors, and the data control and analysis system. The data systems link the interferometer functions and log data from components of the vacuum system and environmental ("housekeeping") sensors. They provide control of the interferometer components (at a level of intimacy with detector operations yet to be determined), and serve as a conduit between the detector signals containing gravity wave information and the storage medium for those signals. The extraction of signals from the archived output of the detectors—probably done off-site—is the final role of the data systems.

2 Model Data and Control System

Our description of the data and control system is linked to the LIGO Mission, as embodied in part in Drawing 89L-307. We describe the on-site implementation of the major components of the data systems, identifying the functional role of computers, how they may be linked, and their approximate count. The conclusions of the groups working on Environmental Specifications and Support Facilities (in progress as of this writing) will influence the details of data systems implementation. Communications between buildings will be covered by another working group, meeting starting 16 March.

2.1 Design Principles

Our design is based on the following principles:

Modularity matching the vacuum system design Each of the Mission interferometers will have an associated data system. The separate data systems will be similar in purpose and function, and each will have the capability to operate independently of the others. The separate data systems are built from a few modules: control computer, communications link, and logging computer. Like modules will be largely interchangeable.

Use of existing standards As a design economy, commercially available and industry standard computers, interconnections, and operating systems will be used to a large extent. The special nature of some of the LIGO control and data requirements may require custom designs for some subsystems.

Reliability The control and data systems should be reliable and fault-tolerant to the level that they do not contribute significantly to LIGO down-time. This implies that software is thoroughly tested before being used to control interferometers or to store data, and that release control procedures are instituted. Hardware reliability can be enhanced by conservative design and by providing on-site spares for critical parts and subsystems.

Accessible user interface A minimum of training should be needed for staff and scientists to rapidly access information and to become acquainted with the operation of the interferometers.

2.2 Capabilities

The data systems will control the overall operation of the interferometers, aid in noise diagnostics, and mediate the storage of interferometer outputs and other signals needed to search for and analyze gravity waves. These functions impose the following capabilities:

Automatic control of interacting functions The simultaneous operation of several detectors will probably be too complex for control by human operators, even with most of the functions governed by semi-autonomous servos. The control and data systems will serve as concentrators for the control and monitoring of many—perhaps thousands—of signals necessary for interferometer operation. Data systems computers will help assure that interacting parts work together properly, even if they are separated by large distances. This requires that all control computers are linked, and that at least one computer has access to the state of all important elements of the interferometers.

Remote access to manual control of detector functions All mechanical positioning stages, electronic gain settings, and switches are to be accessible for manual adjustment via the computer system. This allows operation of the interferometers from an on-site but separate control room, minimizing disturbances from close human intervention.

Cross-correlation diagnostics The data system will have the capability to find unexpected correlations between the signals controlling or related to the detector operation (including the housekeeping data), and between these signals and the gravity wave signal. This will include quick-look analyses, and logging the signals for more thorough off-line analyses. Building the diagnostic software may be a considerable task, but it will result in a valuable tool for improving detector performance.

Archiving of primary detector signals The principal output of each interferometer is a single continuous round-the-clock signal representing the difference in two arm lengths. One or more such signals per interferometer will be recorded at gravity-wave frequency bandwidth.

Environmental, vacuum, and laser monitors The signals provided by the environmental monitoring, vacuum, and laser subsystems comprise the "housekeeping" data stream. The control computers should have ready access to these signals, and there should be provision for logging all of them, at least at low bandwidth.

Clock timing and data stamping To combine signals from detectors at separate sites, the time must be recorded at each site—along with the data stream—to an accuracy corresponding to less than a period of the highest-frequency capability of the detectors. If the detector bandwidth is 10 kHz, an error in Universal Time of 0.1 msec or less is tolerable. In the search for periodic sources, the clock is required to have a phase error of less than one radian of the gravity wave signal over the integration time.

Standard data format International standards will be established for the format of main (interferometer) and ancillary (housekeeping) data. The standards will be disseminated to assure the capability of free exchange of data.

2.3 Block Diagram of Model Control and Data System

The figure on the next page shows a possible configuration of the ligo data and control system. It consists of a number of modular units, each with a unique functional task.

Functional Diagram of Ligo Data and Control System

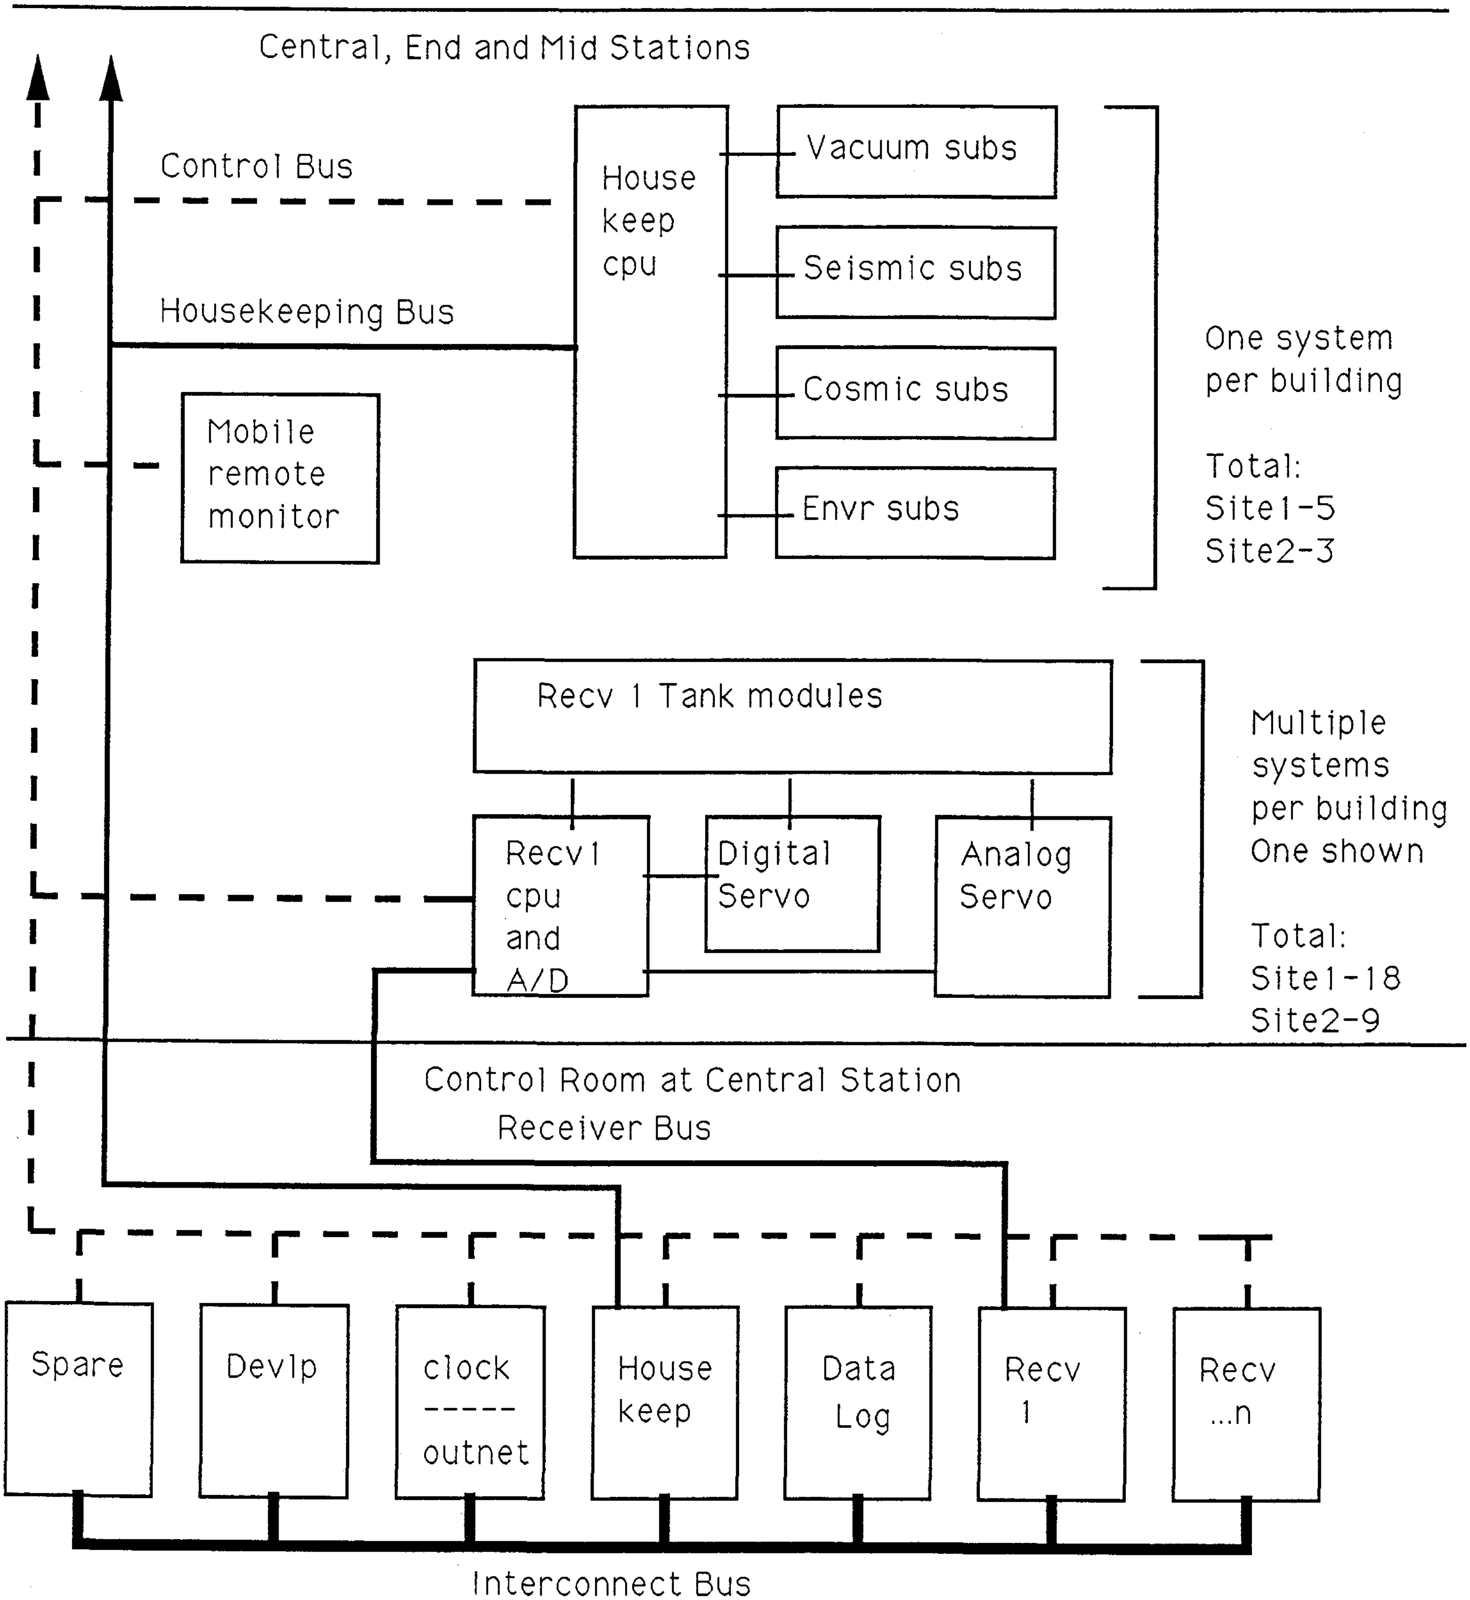

The housekeeping subsystem collects data from the vacuum control system and allows monitoring of such functions as tube and tank pressure at various locations, ion pump currents, and the state of vacuum valves. The seismic subsystem collects data from the vibration and strain gauges located in each of the major tanks. The cosmic ray shower detector is a veto signal with detectors located at each of the buildings. These systems detect showers and give timing and rough energy information. It is likely that this equipment would be standard camac modules and scintillation detectors. The environmental module collects information such as the acoustic sensor and temperatures, laser status and any other type of veto signals. All of these subsystems connect with the housekeeping module located near the tank vertex.

The receiver modules, consisting of both analog and digital electronics, are located close to the tanks of the signals they monitor (see drawing 89L-307). It is probable that some of the position and feedback control will be via digital servo at this location along with local digitization of other analog signals. The three boxes in the vicinity of each chamber (or cluster of chambers) might be contained in one relay rack. All of the electronics and motorized positioners within the chambers are controlled and monitored by these boxes. The estimate of 18 modules at site 1 assumes that one module may serve a cluster of chambers, reflecting an attempt to keep long analog cables to a minimum. There is capability for a mobile remote terminal for local diagnostics.

The various signal paths have been divided into three categories: a control bus that sends packets to and from the ancillary modules, a housekeeping bus that connects the realtime housekeeping module with the control housekeeping machine, and individual receiver buses that carry the detector output and status information to the corresponding receiver control computers.

Within the control room there is a set of nearly identical computers (Spare, Development, Housekeeping, $\text{Receiver}_1 \ldots \text{Receiver}_n$) that are interconnected via a standard network. These machines run standard operating systems, such as unix, and serve as the operator and scientist interface to the observatory. The development system allows new routines and other hardware to be tested before being placed in service. A spare module gives quick turnaround in the event of failure. The data logging module is dedicated to logging the interferometric and housekeeping data. The clock/outnet module consists of the specialized clocking system (distributed to the various a/d systems) and the offsite communications.

2.4 Number of Signals, Bandwidths, and Recording Requirements

We compile here the principal parameters of the signal-handling and controlling capabilities required of the data system. The signals correspond approximately to the "Electrical and Optical Feedthroughs ..." report, Version 1.1 (two of us, A. J. and R. S., contributed to that report). The following table is extracted from information in that report, and used here for counting data system signals. For simplicity, infrequently appearing functions such as Shark-1 are omitted.

Site 1 Best Guess Tank and Function Count VER 1.7							
Tank	Count	Shark-6	PD4	LD	Motor-3	Motor-6	Vib-mon
TM1	8	3	2	2	1	3	1
TM2	4	6	4	4	2	6	1
HAM	58	14	4	3	4	1	
BS12	1	18	14	11	8	10	1
BS8	1	10	9	9	4	5	1
SAT	6	4	5	4	2	3	1
MAN	6		6	6			
Total		912	353	286	272	139	20

The environmental monitoring, vacuum system, and laser system signal requirements are TBD (or if already determined, unknown to us). When these requirements become available, they will be added to the signal list. (*Note: the in-air laser stabilization and initial pointing systems—everything up to the entrance of the laser beam into the vacuum system—needs the attention of another working group.*)

The following tables list the analog and digital signals going to and from the data system control computers.

The "No. of lines" column in the analog signal table is our best guess of the signal count for four interferometers, as specified for Site 1 in the

Mission document. Two bandwidths are specified for each analog input line: "Monitor", which is the maximum bandwidth available for monitoring purposes in units of kHz, or kilo-samples/sec, and "Record", which is the bandwidth for archiving the data. The units of recording bandwidth, kBps, is kilo-bytes/sec—equal or double the sample rate in kHz, depending on whether the samples are one or two bytes wide. Signals that have a lower recording bandwidth than monitor bandwidth are assumed to be preprocessed by pulse-stretching and peak-detecting filters, so that impulsive events are not missed.

Many of the analog signals are likely to be mostly noise. By adding whitening filters and compression, the number of bits that need be recorded can be cut substantially (a calculation based on Gaussian noise suggests that three bits per sample is adequate). The tabulated numbers are the conservative, uncompressed bandwidths.

"Interf. ϕ" is the interferometer phase, the main interferometer signal; "Susp. ϕ " is the signal from the suspension-point interferometer. The "Critical Shark" signals are those associated with the test masses, beamsplitters, or critical steering mirrors. The entry 288 assumes (3 shark detectors / test mass) · (6 signals / shark-6) · (4 test masses / interferometer) · (4 interferometers). Each of these shark photodiode signals appears to the data system in analog form. The analog-to-digital conversion should be done close to the chamber housing the shark.

The "Noncrit Shark" signals are associated with suspended components that should not have first-order coupling to the detector noise. We assume that the control will be accomplished primarily by a tight analog loop; if the control computers were burdened with closing the servo loop, this would require that the analog input and output signal counts match the number of wires specified in the "Feedthrough" document. (*Note: Substituting control computers for conventional analog circuitry would permit a substantial reduction in the number of electronics modules as specified by the "Receiver Support" working group. This would impose additional requirements on the signal count, software complexity, processing speed, latency, and dynamic range of the computers, which we do not analyze here.*)

Our estimates of the recording bandwidth assumes that secondary signals, such as seismometers and microphones, will be fully analyzed in the search for periodic gravity waves. (All bandwidths except for those of the main interferometer output can be reduced for burst searches.) The aggregate analog recording bandwidth, 5.9 Mbytes/sec or 21 Gbytes/hr, is about 80% occupied with the critical shark and quadrant photodiode signals. If this bandwidth is a significant cost driver, the design can be refined to reduce the burden placed on the data system. For example, not all six degrees of freedom of the critical sharks are equally important; one way to cut back is to monitor only two degrees of freedom at full bandwidth. Considerable further reductions may be made after a trial period of running demonstrates that some signals are not important indicators to have in the archives.

The line counts for the Motor-3 and Motor-6 signals are tabulated by taking the single-wire count and dividing by 8, assuming the data are packed in bytes. Most of the motor lines are for shaft encoders. It is assumed the motors are controlled intermittently. For example, a motor might be adjusted for a few seconds once an hour to compensate for thermal expansion within a vacuum chamber. The motor shaft encoder and the command signals will be recorded to keep track of motor activity, resulting in an aggregate bandwidth much lower than for the analog signals. The shark count of 1824 assumes that the control computer can adjust the gain and bias of each shark degree of freedom, but does not directly control the coil current.

3 Issues in Analysis of Data

We imagine that the principal analysis task will be done at a separate computing facility, probably using university computers. A part of the regular operations at the sites will be to make daily copies of the critical data and to ship the copies (electronically, or, more likely, physically) to the data analysis facility. The attached report by Y. G., including supporting references, addresses the issue of how to analyze the data for the presence of gravity wave information. Considerable progress has been made in recent years in the development of algorithms to search for and analyze interferometer signals.

Most of the data analysis can be done with conventional, small, reasonably fast computers, such as current-generation RISC-based desktop machines. A possible exception is that specialized computing hardware, such as dedicated parallel processors, may be needed to search for periodic signals of unknown frequency and unknown position on the sky. The creation of efficient algorithms for searching long data streams (several months of continuous operation) for such signals is a current topic of research.

Analog Inputs to Data System VER 2.3							
Signal	No. of Lines	Bits	Monitor (kHz)			Record (kBps)	
			min	b.g.	max	/line	Total
Interf. ϕ	8	16	5	20	50	40	320
Microphone	6	8	5	20	50	20	120
Susp. ϕ	8	8	1	5	20	5	40
Crit. Shark	288	16	0.1	5	20	10	2880
Noncrit. Shark	5156	8	0.1	5	20	0.1	516
Crit. PD-4	365	8	1	5	20	5	1825
Noncrit. PD-4	1008	8	1	5	20	0.1	101
Vib-mon	108	16	0.1	5	20	0.5	108
Env. Mon.	TBD						
Vac. Sys.	TBD						
In-air laser	TBD						
Total							5910

Digital Inputs/Outputs of Data System VER 1.8		
Signal	Number of Lines	Bits/Line
Motor-3	470	8
Motor-6	532	8
Shark control	1824	8

BATCH START

STAPLE
OR
DIVIDER

Reference (3) (of YG memo on data.syst.)

SOURCES OF GRAVITATIONAL RADIATION: COALESCING BINARIES

B.F. Schutz

Department of Physics
University of Wales College of Cardiff
P.O. Box 78, Cardiff CF1 1XL, Wales, U.K.

ABSTRACT

Binaries consisting of compact objects -- neutron stars, or black holes -- seem likely to become the source of gravitational radiation most frequently observed in the late 1990's by the ground-based, laser-interferometric detectors now under development. The radiation comes from the orbital motion as the stars spiral together due to gravitational radiation reaction. The low frequencies observable from space make new classes of binaries detectable, such as coalescing massive black holes, and make it possible that several coalescencing systems with relatively long times to coalescence could be detected in our own galaxy. I examine the prospects for detecting such sources with various types of space-based detectors, and then consider the sorts of things we could learn from them, including new tests of general relativity.

1. INTRODUCTION TO COALESCING BINARIES

There are many classes of binary systems that emit detectable gravitational radiation in the low frequency regime suited to observations from space. (See Thorne /1/ for a recent review.) The class that is perhaps the hardest to predict, and which therefore could be the most rewarding to observe, is the coalescing binary. A good working definition of a coalescing binary is a system whose components are massive enough and close enough together for the loss of energy and angular momentum in gravitational waves to cause a significant evolution of the system during the period of observation. These systems were first considered as potential sources of gravitational radiation by Clark & Eardley /2/.

Our definition requires that the components be compact objects: white dwarfs, neutron stars, or black holes. Binaries consisting of main sequence stars are not relativistic enough to have a sufficiently short gravitational radiation evolution timescale. During all but the very late stages of the evolution of systems containing neutron stars and/or black holes, the orbital motion can be very accurately modelled as a Newtonian point-mass binary, with gravitational radiation reaction as given by the 'quadrupole formula' /3/. In most circumstances it is also permissible to take the orbital eccentricity of the model equal to zero, since gravitational radiation reduces the eccentricity faster than it does the orbital radius. This model gives the following basic equations describing the gravitational waves from a system consisting of two objects with total mass $M \times M_\odot$ and reduced mass $\mu \times M_\odot$, at a distance $r \times 10$ kpc, whose quadrupole radiation comes off at the frequency $f \times 0.1$ Hz. For the amplitude we have

$$h_{max} = 2.59 \times 10^{-21} \mu M^{2/3} f^{2/3} r^{-1} . \tag{1.1}$$

By h_{max} we mean the largest amplitude when the system and the detector are most favourably oriented with respect to each other. The typical r.m.s. amplitude, averaged over orientations of both the binary system and the detector, is somewhat smaller: for an interferometric detector it is down by √5. The timescale for decay of the orbit is given by

$$\tau = f/\dot{f} = 7.97 \times 10^{8} \mu^{-1} M^{-2/3} f^{-8/3} \text{ sec.} \tag{1.2}$$

Because the orbital decay accelerates so quickly, the actual lifetime is 3/8 of this timescale. Notice that, since both h and τ are measureable, one can take their product to find r: the masses M and μ drop out. This ability to measure r is almost unique in astronomy, and it forms the basis of a new method of determining Hubble's constant from ground-based gravitational wave observations at kiloHertz frequencies /4,5/. However, it is important to note that h itself can only be measured if

To appear in Advances in Space Research (talk presented at Cospar XXVII, July 1988)

the various orientation angles can be determined, and this normally requires 3 or 4 detectors to provide enough information. If space-based detections involve, say, only a single interferometer, then only an estimate of h and therefore of r will be possible.

It is clear from this that with reasonable sensitivity one could expect to detect any systems consisting of white dwarfs, neutron stars or solar-mass black holes if they are in our galaxy, and that they will be changing in the observation time if they have a frequency higher than 0.1 Hz or so. By rescaling the same equations to larger and more distant masses, we can see that we would get a similar amplitude (larger by a factor of about 1.71) and a similar timescale (smaller by the same factor) for a system consisting of two 10^6 $M_\odot$ black holes at a distance of 10 Gpc (cosmological redshift considerably larger than 1) radiating at a frequency of 10^{-4} Hz. (This ignores cosmological redshift and curvature effects, to which I will return below.)

The signal will only stand out from the noise over a long integration time. If the system is not changing in frequency significantly over the period of observation, then simple Fourier analysis will pick it out, provided that there is no confusion with other binary systems. The signal-to-noise ratio will be proportional to the square root of the observing period. If the frequency does change significantly, then matched filtering of the data will be necessary, and the signal-to-noise ratio will improve by the square root of τ. In principle this filtering can be done in the same way as is being planned for ground-based detection: see the articles in reference /6/. It is possible that the background 'noise' of other binary systems will be reduced by a larger factor by matched filtering than would true white noise sources, such as shot noise in an interferometer, since the binary signals are phase-coherent over long periods of time, and not necessarily very densely spaced in frequency.

2. KEY NUMBERS FOR COALESCING BINARIES

The formulas in the last section help us get a rough idea of the sorts of systems we can observe. There are a few useful numbers that discriminate among important regimes of detection.

Given an observing time $T \times 10^7$ sec, what is the frequency $f_{coalescing}$ of a system whose remaining lifetime equals T? This is a system that can be followed all the way to coalescence in the time of observation, using of course ground-based detectors if necessary for the higher frequencies in the late stages of coalescence. The critical frequency is:

$$f_{coalescing} = 0.52\, \mu^{-3/8}\, M^{-1/4}\, T^{-3/8} \text{ Hz.} \qquad (2.1)$$

For Galactic systems, this frequency is rather too high to be accessible to many space-based detectors. If we were generous and imagined that a space-based detector could identify a system that might coalesce in a much longer time, say 20 years, which could then be picked up by ground-based systems, what would the change in the critical frequency be? The weak dependence on T means that it only decreases to 0.11 Hz, still too high for presently planned detectors. For the cosmological system considered above, however, the situation is brighter: the critical frequency for two 10^6 $M_\odot$ black holes is 4.2×10^{-4} Hz, a very accessible frequency.

Another important question is, what is the dividing frequency $f_{changing}$ between systems that change their frequency measureably in the observation time $T \times 10^7$ sec and those that do not? The key point is that as T increases, one's frequency resolution also increases, so that one can distinguish frequencies separated by as little as 1/T. The critical frequency is, in the same notation,

$$f_{changing} = 4.07 \times 10^{-3}\, \mu^{-3/11}\, M^{-2/11}\, T^{-6/11} \text{ Hz.} \qquad (2.2)$$

This is the frequency at which the system would change by $\Delta f = 1/T$ in the observing time T. This is accessible from space, and it shows that Galactic coalescing binary systems can be identified, even if they cannot be followed to coalescence. (This equation is roughly equivalent to Eq. (35) of /7/.) We will discuss below the likelihood that the Galaxy contains any such systems.

Finally, what is the highest frequency $f_{interaction}$ of radiation we could expect a system to emit before it is so drastically altered by the close interactions of its components that it becomes difficult to model? A crude estimate, which omits tidal effects, is to take this upper limit to be the point at which the components would be in contact. This limit is determined by the radius of the less compact of the two components. For a system with the same parameters as before, and whose less compact star has a radius $R \times 10$ km, this is

$$f_{\mathrm{interaction}} = 1.3 \times 10^{3}\, M^{1/2}\, R^{-3/2}\ \mathrm{Hz}. \qquad (2.3)$$

Clearly, one does not need to worry about interactions *for* low-frequency observations of systems consisting of neutron stars and/or black holes, until the mass of the black hole exceeds about 10^6 $M_\odot$. Conversely, one can hope to see interaction effects if one observes massive black holes coalescing. I will return to the significance of this below. If the system contains a white dwarf, its evolution will be the same as for more compact objects of the same mass until it reaches a frequency no higher than about 0.2 Hz. Even if this cutoff is too high to be observed from space, one might be able to see tidal effects beginning at lower frequencies. This problem has been discussed in great detail in ref. /7/.

3. OBSERVABILITY IN VARIOUS DETECTORS

There are three space-based methods of detecting coalescing binaries that I will consider: Doppler tracking of interplanetary spacecraft (which is already taking place, *e.g.* /8,9/), the so-called 'Skyhook' /10/, and an interferometric beam detector /11/.

Doppler tracking involves looking for the changes caused by gravitational waves in the round-trip signal-transmission time between the Earth and a spacecraft. It is broadband, down to a bit below 10^{-4} Hz, and its past sensitivity to the amplitude h of a gravitational wave has been of the order of 10^{-14} for broadband bursts and 3×10^{-16} for continuous waves observed over a 4-month period. The planned Galileo mission will be likely to improve this by an order of magnitude, and foreseeable technical improvements could go as far as a further factor of 100. None of the long-lived coalescing binaries (the galactic sources) would be likely to reach this continuous-wave amplitude, unless it were among the Sun's nearest neighbors. For bursts, we must take into account the importance (mentioned above) of matched filtering of the output in improving the signal-to-noise ratio. This will be important to our analysis of the other detectors as well. For a source that is as narrow band as a coalescing binary, the effect of filtering on a signal of amplitude h is to enhance the signal-to-noise ratio to that which a broad-band burst of amplitude $h\sqrt{n}$ would have, where n is the number of cycles of the waveform while it remains in the bandwidth of the detector /1,6/. In our case n is roughly $f\tau$, where τ is the lifetime given in Eq.(1.2), and f is the lower limit on the freqency of the detector. Putting this together with the maximum amplitude given in Eq.(1.1), and demanding that the interaction frequency given in Eq.(2.3) be within the sensitivity range, we see that only coalescing black holes of mass less than 10^7 $M_\odot$ at distances less than 100 Mpc stand a chance of being detected by the Galileo mission. Such events are not impossible, but must be very rare in such a small fraction of the Hubble volume. However, future technical improvements might put such systems within reach even at interesting distances like 3-10 Gpc. Clearly, there is strong motivation to pursue such improvements.

The Skyhook would be a fairly narrow-band detector, sensitive between 10^{-1} and 10^{-2} Hz to bursts at a level of $h \sim 10^{-16}$ and to continuous waves at $h \sim 10^{-19}$ if it makes a continuous 4-month observing run. This sensitivity, roughly comparable in amplitude to but higher in frequency than the best that can be expected from Doppler tracking, is not likely to be adequate for galactic coalescing binaries, nor would it allow the Skyhook to see massive black holes coalescing in its frequency range unless their masses are below 10^6 $M_\odot$ and their distance considerably less than 500 Mpc. Again, the prospects for this seem slim.

A beam in space, however, could have a much better sensitivity: perhaps as low as 3×10^{-26} for continuous waves over 4 *months*, and 10^{-22} for bursts, over a bandwidth of 10^{-3} to 10^{-1} Hz. This would be more than adequate to detect galactic coalescing binaries and even strong ones in the Andromeda galaxy (M31). It would also see coalescences of binary black holes of mass 10^4 – 10^6 $M_\odot$ essentially anywhere in the observable universe. The beam in space is the instrument with the most promise for returning useful astrophysical information from gravitational wave observations.

4. EVENT RATES AND COSMOLOGICAL CONSIDERATIONS

Although we have good reason to suspect that solar-mass coalescing binary systems exist in reasonable numbers, it is very difficult from present astronomical observations to give a firm estimate of their numbers. The first attempt was by Clark, *et al* /12/. A more up-to-date review of the factors affecting the estimates and their uncertainties is given by Schutz /13/. The conclusion is that binaries consisting of two neutron stars coalesce at the rate of about 1 per galaxy per 10^4 years, with an uncertainty that is possibly as large as a *factor* of 100 either way. Similar rates would apply to binaries consisting of one neutron star and one white dwarf. The rate for two-white-dwarf binaries has been estimated by Evans, Iben, and Smarr /7/ to be perhaps as large as the rate of Type-I supernovae, which may be caused by the coalescence of two white dwarfs. This would suggest a rate of 0.15 per galaxy per year.

This leads to estimates of the likelihood of observing any given class of systems. For example, if the two-neutron-star rate is once per 10^5 years in our Galaxy, then we are unlikely to see such a system with an age less than 10^5 years, and therefore with a frequency greater than 4×10^{-3} Hz. But if the rate is 100 times larger, then the highest frequency we might observe is 2×10^{-2} Hz, well above $f_{changing}$, but well below $f_{coalescing}$. The number of systems per unit frequency increases at low frequencies as $f^{-11/3}$ /7/, so in this optimistic case we might see some 20-30 neutron-star systems above $f_{changing}$ that could be identified as coalescing binaries. White dwarf systems would be far commoner. The problem for detection here is the confusion caused by the large number of distant but detectable sources. This affects mainly the low-frequency systems. All systems above $f_{changing}$ ought to stand out from the continuum, even without allowing for the fact that their distinctive frequency change allows them to be filtered from the background /7/. There might be thousands of such systems detectable by a beam in space /7/. These would be indistinguishable from neutron-star systems of similar masses unless tidal effects can be detected; but it may well be that there will be little overlap in the distribution of masses between neutron-star binaries and white-dwarf binaries.

It is equally difficult to give an estimate of the coalescence rate of massive black holes at cosmological distances. It is popular to believe that active galactic nuclei contain supermassive black holes in the mass range 10^6 - 10^9 $M_\odot$, and there is mounting evidence that ordinary galaxies like our own may contain smaller black holes, perhaps 10^3 - 10^6 $M_\odot$. There are several plausible scenarios for the formation of these holes. One route might be the formation by gas-cloud collapse of black holes of 1000 solar masses, followed by the hierarchical coalescence of larger and larger black holes. If this were the case, then the formation of a 10^9 $M_\odot$ hole would involve a large number of observable coalescences, and would give event rates of several per year, even if only 1 galaxy in 1000 has a supermassive black hole in it. This is the most optimistic scenario, of course, and it may be that only gravitational wave observations from space will be able to shed light on the question of the formation of these giant black holes.

When studying coalescing binary systems that are at cosmological distances, it is important to take into account the effects of the redshift z. It can be shown that this is very simple: Eqs.(1.1) & (1.2) still apply, but the 'observed' masses μ and M are (1 + z) times the true values of the masses, as measured near the source, and the distance r is the luminosity distance. The frequency f in the equations is the observed frequency, which is also redshifted from the original frequency. If a given system is taken to larger and larger redshifts, and if the system can be followed all the way to coalescence, then the signal-to-noise ratio will depend on redshift to lowest nontrivial order as

$$S/N \propto z^{-1} \left[1 + \left(\frac{2 + 3q_0}{6}\right) z + \ldots \right]. \qquad (4.1)$$

Here q_0 is the usual deceleration parameter /14/. The 1/z dependence is just what one expects, but the effect of going to higher redshifts is to increase the signal-to-noise ratio above what one would expect: in this sense, coalescing binaries are easier to detect at larger redshifts.

5. WHAT WE WOULD LEARN FROM COALESCING BINARIES

From our discussion above of the sorts of systems that would be candidates for observation, it is clear that very fundamental astrophysical information is available in these waves. Let us start with the smaller, local systems first.

An important unknown is the number and distribution (in age and frequency) of coalescing binary systems in our Galaxy. By the time space observations with sufficient sensitivity take place, we should already have a good idea from ground-based detections what the typical neutron-star coalescence rate should be, averaged over some 10^6 galaxies or more. But it is possible that several distinct populations contribute to the rate: old systems (Populations II and III) that formed with or soon after the Galaxy and whose binaries are gradually decaying, and new systems (Population I) that are still forming precursor systems. By observing the spatial distribution of coalescing binary systems in our Galaxy, these two populations could be distinguished, and this would in turn tell us much about the differences between star formation now and in the remote past. (This will not be easy to do with only one detector in space, however.) Among the coalescing systems might be a few percent that consist of a neutron star and a black hole, and it would be interesting to see these.

Seeing white-dwarf binaries would be equally interesting. If their numbers correspond to a birthrate equal to the rate of supernovae of Type I, this would be strong evidence for the idea that they end as such supernovae. Observations of tidal effects in white dwarf binaries would also give considerable information about pre-supernova systems and about accretion and mass exchange in interacting binaries.

The cosmological systems would of course be even more dramatic. Not only would their detection shed light on the nature and formation of active galactic nuclei and quasi-stellar objects, but they enough to see the nonlinear effects of general relativity on the orbit of the holes (which should have been simulated accurately on computers by then), it will be possible to compare theory and observation in the very strongest field limit. This test is also possible from ground-based observations of two-black-hole binaries of moderate mass (10 - 20 $M_{\odot}$), but these may well not have been observed before sensitive space-based detectors fly.

In both cases there is astrophysical information in the low-frequency gravitational waves that is probably obtainable in no other way. Let us hope that the development of space-based detectors will make steady progress.

REFERENCES

1. K.S. Thorne, in: *300 Years of Gravitation*, eds. S.W. Hawking and W. Israel, Cambridge University Press, Cambridge, 1987.
2. J.P.A. Clark and D.M. Eardley, *Astrophys. J.*, 215, 315 (1977).
3. A. Krolak and B.F. Schutz, *Gen. Rel. Gravit.*, 19, 1163 (1987).
4. B.F. Schutz, *Nature*, 323, 310 (1986).
5. B.F. Schutz, in: Gravitational Collapse and Relativity, eds. H. Sato and T. Nakamura, World Scientific, Singapore, 1986.
6. B.F. Schutz, ed., *Gravitational Wave Data Analysis*, Reidel, Dordrecht, 1988.
7. C.R. Evans, I. Iben, Jr., and L. Smarr, Degenerate Dwarf Binaries as Promising Detectable Sources of Gravitational Radiation, California Institute of Technology preprint GRP-123, March 1987.
8. J.W. Armstrong, F.B. Estabrook, and H.D. Wahlquist, *Astrophys. J.*, 318, 536 (1987).
9. R.W. Hellings, P.S. Callahan, J.D. Anderson, and A.T. Moffet, *Phys. Rev.*, D23, 844 (1981).
10. V.B. Braginsky and K.S. Thorne, *Nature*, 316, 610 (1985).
11. J.E. Faller, P.L. Bender, J.L. Hall, D. Hils, R.T. Stebbins, and M.A. Vincent, this volume.
12. J.P.A. Clark, E.P.J. van den Heuvel, and W. Sutantyo, *Astron. Astrophys.*, 72, 120 (1979).
13. B.F. Schutz, in: /6/.
14. B.F. Schutz, *A First Course in General Relativity*, Cambridge University Press, Cambridge, 1985.

BATCH START

STAPLE
OR
DIVIDER

Reference (4) (of YG memo on data. syst.)

SOURCES OF GRAVITATIONAL WAVES*

Bernard F. Schutz
Department of Applied Mathematics and Astronomy
University College Cardiff
Cardiff, Wales, U.K.

November 15, 1988

Abstract

Sources of low-frequency gravitational radiation are reviewed from an astrophysical point of view. Cosmological sources include the formation of massive black holes in galactic nuclei, the capture by such holes of neutron stars, the coalescence of orbiting pairs of giant black holes, and various means of producing a stochastic background of gravitational waves in the early universe. Sources local to our Galaxy include various kinds of close binaries and coalescing binaries. Gravitational wave astronomy can provide information that no other form of observing can supply: in particular, the positive identification of a cosmological background originating in the early universe would be an event as significant as the detection of the cosmic microwave background was.

1 Introduction

Almost every speaker at this Workshop who has discussed methods of detecting gravitational waves from space has included a discussion of possible sources of gravitational waves at low frequencies. My aim here is not to repeat these discussions, but to put them in their astrophysical context: why is gravitational wave astronomy interesting? A good source for further reading is Thorne (1987).

In general terms, gravitational waves open up a qualitatively new window on the universe. The information they carry reflects the large-scale mass distribution of the source, on distance scales of the same order as the gravitational wavelength. By contrast, observable electromagnetic radiation is of much higher frequency, and comes from small regions: atomic size for visible wavelengths, for example. As a consequence, astrophysical modelling of large-scale structures requires assumptions that enable one to go from the small scale to the large: assumptions of local thermodynamic equilibrium, of homogeneity, of symmetry, and so on. Gravitational waves will enable more direct modelling of the source, and will be complementary to electromagnetic waves when both are available.

1.1 A brief look at sources of high-frequency radiation

It will help us to look briefly first at sources of high-frequency gravitational waves, even though they are of more relevance to ground- based detectors than to space-based ones. Some of them are closely related to low-frequency sources, and if they are detected from the ground they will provide further incentive for looking from space. For a review of ground-based detection, see Schutz (1988).

Gravitational collapse

Collapse to form neutron stars or black holes in the mass range $1-10M_\odot$ will radiate waves in the frequency range 1—10 kHz, with an amplitude that depends on how much asymmetry there is in the collapse. These collapses at least sometimes result in Type II supernova explosions. The rate at which Type II supernovae occur is relatively well known, but the fraction of collapse events that produce strong enough gravitational waves is not. Since the characteristic period of the waves is proportional to the light-travel time around the collapsed object, the dominant frequency scales as $1/M$. For sufficiently large M, this source will produce low-frequency waves detectable from space. (See the article by Stark in this volume.)

*Text of a talk presented to the NASA Workshop on Relativistic Gravitation Experiments in Space, Annapolis, MD, 28-30 June 1988

Coalescing binaries

This is one of the most promising sources of waves detectable from the ground, once broadband laser detectors reach their expected sensitivity. The famous 'Binary Pulsar' PSR 1913+16 is a precursor of such a system: in some 10^8 years it will have evolved through gravitational radiation reaction into an almost perfectly circular orbit with a period of 20 msec and a separation between the stars of about 150 km. At this point it will be a strong source of gravitational waves at 100 Hz, within the expected observing window of laser-interferometric detectors. During the next two seconds the stars will spiral together and coalesce; before they coalesce, they will have emitted some 500 or so cycles of radiation at ever increasing frequency. Because the signature of this radiation, or 'chirp', is unique and predictable, it is possible to filter weak signals out of the noise of an interferometer. Consequently, coalescing binaries can be seen some 25 times further away than moderately strong gravitational collapses (supernovae). The expected event rate is very uncertain. Again, the frequency of the waves is inversely proportional to the mass of the system, so binaries consisting of massive black holes could be detected from space. So, too, might the precursor systems when the stars are still well separated, as in the present Binary Pulsar. I will return to this source in Section 1.2 below.

Pulsars

Pulsars emit gravitational waves if they are non-axisymmetric. The frequency of the waves will be twice the rotation frequency of the star. We have little idea of what strength to expect from known pulsars, but it is unlikely that any slowly rotating former pulsar would be a strong source of gravitational waves at low frequencies.

Accreting neutron stars

Neutron stars in X-ray binaries can be spun up by accretion, possibly until they reach a rotation rate at which they encounter a non-axisymmetric rotational instability. As Wagoner (1984) has pointed out, further accretion will drive the instability, until it has sufficient amplitude so that the gravitational waves that are radiated carry away as much angular momentum as that which is being accreted. The system then becomes a steady source of gravitational waves. Several galactic X-ray sources are candidate sources. We do not know enough about the behavior of matter at neutron star densities to predict what the frequency of this radiation should be. If X-ray observations — such as those proposed for the XLA satellite (see the talk by Woods at this meeting) — detect low-amplitude variability in X-ray sources, ground-based detectors could search for the associated waves. Successful observations would be enormously important for neutron star (and hence for nuclear) physics. It is most unlikely that any of this radiation will be at frequencies below 10 Hz.

1.1.1 Stochastic background

There are many postulated sources of a measureable stochastic background at kiloHertz frequencies, all of them cosmological. Perhaps the most interesting are cosmic strings, which might have acted as seeds for galaxy formation. If they did, there is a firm prediction that the gravitational wave background they would have produced should have an energy density of 10^{-7} of the closure density (Vachaspati & Vilenkin 1985). There is no preferred frequency for this background, so the waves' spectrum should be scale-invariant. Detection of this backgound would provide strong evidence, not only for the string model of galaxy formation, but also for the particle-physics theories that lead to strings. See the talk by Matzner at this meeting for more details on backgrounds.

1.2 Coalescing Binaries in more detail

The interest in coalescing binaries of neutron stars or black holes is easier to understand if we write down the formulas for the amplitude h of the gravitational waves and the timescale τ for the coalescence of the system, in terms of the total mass M_T of the system, its reduced mass μ, the frequency f of the radiation, and the distance r to the system:

maximum h (when the system is viewed down the axis)

$$h_{max} = 3.6 \times 10^{-23}\left(\frac{M_T}{2.8M_\odot}\right)^{2/3}\left(\frac{\mu}{0.7M_\odot}\right)\left(\frac{f}{100\mathrm{Hz}}\right)^{2/3}\left(\frac{100\mathrm{Mpc}}{r}\right),$$

and

coalescence timescale

$$\tau := \frac{f}{\dot{f}} = 5.6\left(\frac{M_T}{2.8M_\odot}\right)^{-2/3}\left(\frac{\mu}{0.7M_\odot}\right)^{-1}\left(\frac{f}{100\mathrm{Hz}}\right)^{-8/3}\mathrm{sec}.$$

When viewed in other directions, the binary produces a wave amplitude that is h_{max} times angular factors. A network of four broadband detectors can determine these angular factors and thereby measure h_{max}.

Notice that the product $h_{max}\tau$ depends only on r: *coalescing binaries are standard candles!* It is extemely difficult in astronomy to find observable systems that can provide reliable distance measures. Coalescing binaries are of great interest for this reason. See the talk by Wahlquist at this meeting for further discussion of these binaries in the context of space-based observations.

For low-frequency observing, there are two frequencies which are useful to remember. If one expects to observe a system consisting of two $1.4M_\odot$ neutron stars for an observation period of 10^7sec, then the first important number is that a binary with an initial frequency of 0.5 Hz will just reach coalescence at the end of the observing period. This is in some sense the optimum frequency to search for coalescing systems at, since they are easiest to observe when they change the most in the observation period. If they are picked up at a lower frequency, they change less dramatically in 10^7sec. Unfortunately, frequencies of 0.1—1 Hz are the worst from the point of view of detector noise! The second number to keep in mind is that if a system with the assumed masses has $f < 7 \times 10^{-3}$Hz, then it will not change its frequency by a measureable amount during a 10^7sec observation. This frequency is roughly the dividing line between standard binaries and coalescing binaries, from an observational point of view.

2 Sources of Low-Frequency Gravitational Waves

There is a natural division of likely sources into two categories: cosmological sources, which are strong and distant; and galactic sources, which are local but weak.

2.1 Cosmological sources.

Formation of a giant black hole

Many astrophysicists believe that the most plausible explanation for quasars and active galactic nuclei is that they contain massive (10^6–$10^9M_\odot$) black holes that accrete gas and stars to fuel their activity. There is growing evidence that even so-called 'normal' galaxies like our own and Andromeda (M31) contain black holes of modest size (10^4–$10^6M_\odot$) in their nuclei (Blandford 1987). It is not clear how such holes form, but if they form by the rapid collapse of a cluster of stars or of a single supermassive star, then with a modest degree of non-symmetry in the collapse they could produce amplitudes $h \sim 10^{-16}$ - 10^{-18} in the low frequency range observable from space. If a detector had a spectral noise density of $10^{-20}\mathrm{Hz}^{-1/2}$ (see the talk by Bender at this meeting — this might be a conservative figure), then such events could have signal-to-noise ratios (S/N) of as much as 1000. This strong a signal would permit a detailed study of the event. If every galaxy has one such black hole formed in this way, then there could be one event per year in a detector. If no such events are seen, then either giant black holes do not exist or they form much more gradually or with good spherical symmetry.

Star falling into a giant black hole

If black holes power active galactic nuclei, they do so by swallowing stars and gas. Occasionally, neutron stars should fall into them. Neutron stars are compact enough not to be disrupted by tidal forces before reaching the horizon, so they will give a coherent gravitational wave burst with a frequency similar to that which the black hole gave off when it formed. Fairly reliable numerical calculations of this radiation exist (see Thorne 1987 for references), and they suggest that an event in the Virgo Cluster of galaxies

would give an amplitude $h \sim 10^{-21}$ and $S/N \sim 10$. The event rate, however, is very uncertain: although the Virgo Cluster contains over 1000 galaxies, their central black holes are quiescent and may by now have already consumed all the stars that are in orbits that take them near to the hole.

Coalescence of giant black holes.

If two black holes of mass $10^6M_\odot$ or more collide and coalesce, they will emit radiation which is at least as strong as we have suggested above for the formation of such holes. The waveform would have a characteristic signature, from which one could identify the event with some confidence. Such collisions could result from the merger of two galaxies that both contain black holes. Merged galaxies are not uncommon, especially in the centers of clusters; after the merger, dynamical friction could bring both holes to the center, where they would coalesce. Alternatively, it might be that giant black holes in the centers of galaxies themselves form, not be a single collapse, but by a sort of hierarchical merger of smaller black holes. Again, the event rate is very uncertain, but the events would be strong, $S/N \sim 1000$.

Stochastic background of cosmological origin

Gravitational waves having frequencies below 10^{-2}Hz today may be redshifted relics of waves emitted in much earlier phases of the Big Bang. See Matzner's talk at this meeting for a full discussion of the different mechanisms which might produce such waves. Among the most interesting observationally are inhomogeneities associated with inflation, which might produce a scale-invariant spectrum with a spectral density $\sim 10^{-21}\mathrm{Hz}^{-1/2}$ at 10^{-4} Hz; and early anisotropies, which might produce a 'line' of radiation at about 10^{-5}Hz, with spectral density $10^{-20}\mathrm{Hz}^{-1/2}$. If these backgrounds could be detected and identified by their spectrum, they would provide the most direct evidence that the early universe was dominated by the sort of particle physics effects that are fashionable but speculative in modern cosmological theory: inflation, spontaneous symmetry breaking, cosmic strings, and so on. The implications for cosmology and physics as a whole would be fully as significant as the discovery of the cosmic microwave background was 25 years ago. Clearly, this is one of the most important gravitational wave experiments possible from space. But it may not be easy, since as we will see below there are other backgrounds due to binary stars that could obscure any cosmological background.

2.2 Galactic sources

Coalescing binary precursors

If observing from space is confined to frequencies below 0.1 Hz, then our earlier discussion of coalescing binaries in Section 1.1 makes it clear that no solar-mass systems will be discovered that can be followed all the way to coalescence. However, it should be possible to see some precursors either as ordinary binaries (*i.e.*, below 7×10^{-3}Hz). The Binary Pulsar system itself will be just detectable at about 10^{-4}Hz if the spectral noise density of the detector is $10^{-20}\mathrm{Hz}^{-1/2}$. Because pulsar radiation is beamed, it is likely that there are similar systems even closer to us that we do not observe because their beams are pointed in the wrong direction. If the nearest is 2 kpc away, then it might give $S/N \sim 10$ if it is favorable oriented with respect to the detector. A precursor with a frequency of 10^{-3}Hz could be seen as far away as the Andromeda galaxy (M31) with $S/N \sim 10$. Since the number of precursor systems is very uncertain (see Schutz 1988), there is a good possibility that such a system, with a lifetime of only 10^4 years, would be seen.

Close white-dwarf binaries

Systems like this are associated with cataclysmic variables, Type I supernovae, and especially with models of the formation of isolated millisecond pulsars by the coalescence and subsequent collapse of the two white dwarfs. They are more numerous than neutron-star binaries, so the nearest may be considerably closer, with a $S/N \sim 100$ or more in a $10^{-20}\mathrm{Hz}^{-1/2}$ detector.

Individual binaries

A number of nearby binary systems are known which produce radiation that is strong enough to be observed by space-based detectors. See Thorne (1987) and references therein for a list. This is one of the few certain sources of gravitational waves at these frequencies.

Background noise from binaries

Another certain source is the vast number of ordinary binary sytems, whose radiation reaches us from random directions and at random frequency. A single space-based detector will have little directional resolution, so below about 10^3Hz it will be receiving waves from so many systems that they will be more closely spaced in frequency than the frequency resolution one can obtain in a 10^7sec observing run. (See Thorne 1987 or the talk by Bender at this meeting for details of the expected spectrum.) This background is of interest in its own right, since detecting it would give a measure of the distribution of periods in the binary population of the Galaxy. But it can also be a nuisance, obscuring other interesting sources. There are at least two possible ways to beat this noise. One is to obtain directional information about the gravitational waves, for example by flying two detectors. In any given solid angle, the confusion caused by the background will be reduced by the ratio of the solid angle to 4π. The second method is to make use of the fact that the 'noise' produced by these binaries is not true white noise: at any single frequency the amplitude is constant and the phase remains coherent over the observing period, since it is just the signal of a single binary system. This property may make it easier to filter for signals that do not have constant frequency, such as black-hole bursts or waves from relatively massive coalescencing binaries, since the 'noise' is not really stochastic.

3 Conclusion

There are a great variety of possible sources of gravitational waves at milliHertz frequencies. Some are rather speculative and some are essentially certain, considerably more certain in fact than any of the postulated sources detectable by ground-based detectors. Observations of or even good upper limits on some of these sources would contribute valuable information to astrophysical modelling of different types of binary star systems, neutron stars, quasars, active galaxies, and the early universe. In particular, the discovery of a gravitational wave background of cosmological origin would be of the greatest significance to astronomy and physics. Despite the great difficulties involved in building sensitive space-based detectors, the possible scientific returns make a strong case for going ahead with them.

Acknowledgements

It is a pleasure to acknowledge useful conversations with John Armstrong, Peter Bender, Ron Hellings, and Kip Thorne.

References

Blandford, R. 1987, in S. W. Hawking & W. Israel (eds.), *300 Years of Gravitation* (Cambridge University Press, Cambridge, U.K.), p. 277.
Schutz, B.F. 1988, *Gravitational Wave Data Analysis* (Reidel, Dordrecht).
Thorne, K.S. 1987, in S. W. Hawking & W. Israel (eds.), *300 Years of Gravitation* (Cambridge University Press, Cambridge, U.K.), p. 330.
Vachaspati, T. and Vilenkin, A., *Phys. Rev. D*, **31**, 3052.
Wagoner, R.V. 1984, *Astrophys. J.*, **278**, 345.

Reference (5) (of YG memo on data syst.)

GRAVITATIONAL RADIATION

Bernard F. Schutz
Department of Physics, University of Wales College of Cardiff
Cardiff, Wales, U.K.

Abstract

In this talk I will review the present status and plans of gravitational wave detector groups and the current thinking about important likely sources of gravitational waves. By the middle of next year, several groups should have instrumented cryogenic bar detectors for continuous observing; these have sufficient sensitivity to detect or set interesting limits on gravitational waves from supernova explosions in our Galaxy. The next generation of ultracryogenic bars is under development. Laser interferometric detectors can match present cryogenic bars in sensitivity, but not in long-timescale operation, and laser groups are concentrating on developing their plans for kilometer-scale detectors. Such plans are well advanced in several countries, and a number of projects may be approved by the next Texas Symposium. These large laser interferometers are being designed to reach a sensitivity of $h = 10^{-22}$ over a broad bandwidth, and this improvement over the old goal of 10^{-21} has brought about a revolution in our thinking about what sources are likely to be detected. Supernovae and other gravitational collapses are still important sources, and even moderately strong events (with $0.01 M_{\odot}c^2$ gravitational wave energy) might be detected at 60 Mpc. Even given the great uncertainty we presently have about the likely strength of radiation from such sources, prospects are very good for detecting at least a few per year. Prospects are even better for detecting coalescences of compact-object binaries. Because it is possible to apply pattern-matching techniques to the analysis of the data, it will be possible to detect neutron-star coalescences out as far as 1 Gpc, *i.e.* out to redshifts of 0.1 to 0.2. Moreover, coalescing binaries are *standard candles:* observations with a worldwide network of detectors can determine the distance to a coalescing binary directly. The astrophysical consequences of this are considerable: one can contemplate measurements of Hubble's constant, statistical studies of the distribution of stars at moderate cosmological distances, and the direct identification of black holes in coalescence events. Other sources are also interesting: all-sky searches for nearby unknown pulsars will be possible, and interesting limits should be set on a possible stochastic background of gravitational waves. The timescale for the development of laser interferometric detectors is rather long, however, so for the next five years at least, bar detectors, operating in conjunction with neutrino detectors, provide the best chances of a first reliably confirmed detection of gravitational waves. While detectors are being developed, there are a number of astrophysical questions about possible gravitational wave sources that need to be answered.

1 Introduction

Talks on gravitational radiation at previous Texas Symposia have spanned a wide range of subjects, from efforts to detect gravitational waves to the mathematical questions concerning the quadrupole formula and the equations of motion of radiating systems. In this talk I have chosen to concentrate on detectors and likely sources, partly because that is where my own interests lie at the moment, but mostly because of the great deal of activity these subjects have seen in recent years, which has led to many developments that are not as well known as they should be in the relativity and astrophysics communities.

In concentrating on detectors and sources, I am ignoring other active areas of the subject. The "quadrupole controversy" that began in the late 1970's (see *e.g.* Ehlers, *et al*, 1976 or the Texas Symposium talk on the subject ten years ago, Ehlers 1980) has largely abated, many independent approaches having confirmed the validity of the formula in the Newtonian limit (Damour 1987, Balbinski *et al* 1985). Nevertheless, there is a great deal of work on equations of motion, for example, especially with reference to relativistic binary systems (see Damour 1987).

Based on a talk given at the 14th Texas Symposium on Relativistic Astrophysics, Dallas, Texas, December 1988

Another area that I will not be able to cover is that of numerical relativity, which is crucially important to the reliable calculation of source strengths. A comprehensive review of that subject would require a full lecture of its own. The field is developing in two directions: (i) to re-do with better accuracy and more physics the earlier spherical and axisymmetric calculations of black-hole collisions and gravitational collapse; and (ii) to push on to fully three-dimensional calculations in general relativity. For more details see the proceedings of the recent workshop at the National Center for Supercomputer Applications (Evans, *et al*, 1989).

Instead, what I will be concentrating on are recent progress in detector development and the consequent broadening of our horizons with regard to what sources we believe we shall be able to see. For comprehensive reviews in greater depth, see Thorne (1987), Schutz (1989a), and Blair (1989).

2 Detectors, Present and Planned

2.1 Bar Detectors

Bar detectors have been under development since Weber's pioneering efforts that began in the late 1950's. Present bars operate with a sensitivity of a few times 10^{-18}, and bars now under development could push that to 10^{-20}. At that point they will be close to their quantum limit, and further progress will require either beating that limit (see Caves *et al* 1980) or new designs incorporating new materials, about which I will say a little below.

The present status of bar detector groups is summarized in Table (1). There are three kinds of bars in operation today:

- Room temperature bars. These have typical sensitivity of $\sim 3\times10^{-17}$ or worse. The only detectors taking data at the time of SN1987a were the room-temperature bars at Rome and Maryland. Pizzella will speak about those observations later in this meeting.
- Torsion pendulum. Detectors of this type have been constructed at Tokyo to look for radiation from the Crab Nebula pulsar. Consisting of two large masses connected by a torsion fiber, they can be tuned to a lower frequency than solid bars can reach. Current published limits on the gravitational radiation from the Crab are about 10^{-21}.
- Cryogenic bars. Cooled to 4.2 K, these can reach as low as 10^{-18}. There have been several coincidence experiments with these detectors, the most extensive of which was the recent three-way experiment of L.S.U.-Rome-Stanford, which set a limit of 7.5×10^{-18} on bursts during a limited period of time. Unfortunately, none of the groups developing these antennas was funded well enough for them to have two detectors, one to take data and the other as a development testbed. Consequently, when SN1987a went off, only the one in Perth was actually cooled down, and that one was not taking data because of excess noise problems. This situation is being rectified. By early 1989, the bars at Stanford, Rome, and Perth should be cooled down and taking data. The L.S.U. bar, which in late 1988 completed a six-month run at 4.2 K, may join them later. Figure (1) shows a short stretch of data from that run. The groups are linked by the observing network called GRAVNET, which will coordinate the analysis of data.

Cryogenic bars can be very frustrating to work on. It takes a month or more to cool one down, and if anything is wrong or needs changing, it may take a further month to warm it back up. Progress is therefore necessarily a matter of patient, steady development.

A typical summary of the data taken by a bar detector is in Fig. (1). The number of events per unit energy interval is plotted against their energy. The curve follows the expected thermal distribution for an effective noise temperature of 103 mK, except for a very few high-energy events. This is characteristic of both bar and laser interferometer detectors, that there is always a small number of non-Gaussian, unmodelled noise events. Coincidence experiments are necessary to exclude these. The noise temperature is much lower than the bar temperature because measurements are being made on the state of the mode over

The active bar-detector groups

Institution	Room-temperature	Torsion	Cryogenic	Ultracryogenic
Stanford U.	-	-	√	√
Louisiana State U.	-	-	√	√
U. of Maryland	√	-	√	√
U. of Rome	√	-	√	√
U. W. Australia (Perth)	-	-	√	-
Moscow State U.	√	-	-	-
Tokyo U.	-	√	-	-
China (& Beijing)	√	-	-	-

Table 1: Summary of the facilities developed at and planned for the active bar-detector groups.

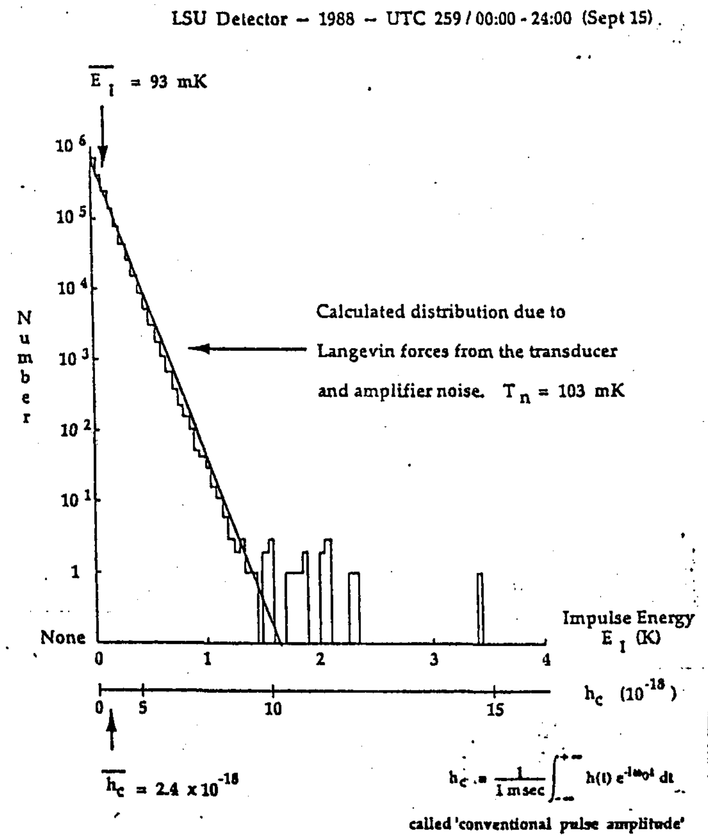

Figure 1: Data taken by the L.S.U. detector during one day of operation. See text for discussion.

timescales short compared to the time Q/ω it takes a mode with frequency ω to exchange thermal energy with the rest of the bar.

For the immediate future, several groups have plans for ultra-cryogenic bars, reaching temperatures of 100 mK or less. Stanford expects to take delivery of its new refrigerator in early 1989, and Rome began constructing their milliKelvin bar in 1988. Similar plans exist at Maryland and L.S.U. By 1991-2 we may see bars operating with sensitivity at the $10^{-19} - 10^{-20}$ level.

It is very important for astrophysics that these developments go ahead without damaging the efforts to keep the 4.2 K bars on the air. For at least the next 5 years, bars will be our only gravitational wave detection system, and there is at least a fair chance that there will be a gravitational collapse in our Galaxy in that time. Operating in coincidence with neutrino detectors, cryogenic bars would have an excellent chance of seeing gravitational radiation from such an event.

In the more distant future, the prospects for making measurements below the quantum limit, although in principle possible, do not seem encouraging at this time. But it may be possible to build bars whose quantum limit is considerably lower, near to 10^{-22} (W. Fairbank, private communication). These might be made of materials with a higher speed of sound, such as silicon carbide, allowing a larger size for a given resonant frequency. These might take bars into the same sensitivity regime as the planned laser interferometric detectors, although with much narrower bandwidth.

2.2 Laser Interferometric Detectors

The idea of using laser interferometers to detect gravitational radiation also goes back to Weber, who did not pursue it because of the limitations of the technology of the time. The first working interferometer for gravitational wave detection was built by R.L. Forward and associates at Hughes Research Laboratories (Moss, *et al*, 1971; Forward & Moss 1972; Forward 1978). Developments in laser and mirror technology, allied to clever ideas for the optical configuration of the detectors (such as various forms of recycling, as described by Drever 1983, Vinet *et al* 1988, and Meers 1988), have greatly enhanced the potential sensitivity of these instruments.

Current designs are based on arm lengths of 3–4 km. Roughly speaking, they would be built in two stages. The first stage would aim at a sensitivity of about 10^{-21} over a 1 kHz bandwidth. The second stage would implement more sophisticated optical techniques, leading to noise levels of $h \sim 10^{-22}$.

In addition, there is every expectation that the final-stage interferometers should be able to achieve good isolation from seismic noise down at least to 100 Hz, and possibly as low as 10 Hz. This low-frequency observing window, coupled with the enhanced sensitivity, has made it possible to contemplate detecting a much wider range of sources than one envisioned with bar detectors. The most exciting "new" source is the coalescing binary, which I will discuss in detail below.

The present status of laser interferometric detectors is summarized in Table (2). Working prototypes with arm lengths in the range 10–40 m have been constructed at Glasgow, Munich, Caltech, and recently in Japan. The sensitivity of the Glasgow detector is displayed in Figure (2). The flat spectrum is what detector groups hope to achieve over a wider bandwidth in full-scale detectors. The vertical units indicate the actual displacement being measured. To convert to an accuracy in h, multiply by the square root of the bandwidth of the expected signal (say 1000 Hz) and divide by the arm length (10 m). This gives a sensitivity of $h \sim 4 \times 10^{-18}$, comparable to the cryogenic bars. This is the best displacement sensitivity of any interferometer at present, but the Munich detector has better h sensitivity because of its longer arm length.

Special-purpose interferometers, designed to study particular technical problems in this area, have been built at M.I.T., Paris, and Pisa. In the U.S.S.R., plans are being made to convert a low-sensitivity interferometer built for geophysical purposes into a sensitive gravitational wave laser interferometer. Various detectors have made observing runs of limited duration, and a coincidence experiment lasting about one week is planned for early 1989 by Glasgow and Munich. There are no plans to run continuously, as the bars intend, because (i) the prototype interferometers are engineering testbeds for the large-scale detectors and cannot be left alone to take data for long periods, and (ii) laser interferometers are difficult systems to keep locked onto a fringe, and this must be done essentially by hand at the moment. Part of the design effort for large-scale detector will go into their control systems, but this considerable effort is not likely to be devoted to the prototypes. This means that until the large-scale projects are constructed and in operation, the only systems capable of sustained operation will be the bar detectors.

The active laser interferometer groups

Institution	Prototype	Special Purpose	Plan for Full-Scale Project
Calif. Inst. of Tech.	√	-	√
Mass. Inst. of Tech.	-	√	√
U. Glasgow	√	-	√
Max Planck Q. Opt. Munich	√	-	√
I.N.F.N. Pisa	-	√	√
C.N.R.S. Orsay—Paris	-	√	√
U. W. Australia (Perth)	-	-	√
Tokyo	√	-	√
India	-	-	√
U.S.S.R.	-	√	√

Table 2: Summary of the facilities developed at and planned for the active laser interferometric-detector groups.

Regarding future plans, the American project is presently (December 1988) engaged in a design exercise, and expects to submit its fully costed proposal for two detectors to the National Science Foundation in Autumn 1989. If approved, it could start in 1990. In other countries, the large cost of these instruments is a spur to international collaboration. In the U.K., there is a commitment of about 25% of the money required to build a detector, and discussions with possible partners in other countries are actively underway. A 1990 start to this project does also does not seem unrealistic, provided a suitable partner can be found. In Italy there is also a sum slightly larger than the British funding committed to the project, and a collaboration with France is likely. That detector would be built in Pisa. Other countries have entered the field recently: Australia, India, and Japan may well enter into partnerships with other countries to build detectors. From the scientific point of view, the more detectors the better. At least three and preferably four are required to reconstruct the gravitational wave completely, giving its intrinsic amplitude and its position on the sky. More detectors operating in coincidence will allow all the detectors to set lower thresholds against noise and thereby to see further away.

2.3 Space-based Detectors

I will not have much time to discuss detection of gravitational waves from space, but a few words are appropriate here. Methods presently used include tracking of interplanetary spacecraft and pulsar timing. Transponding to interplanetary spacecraft is sensitive to low-frequency gravitational waves, *i.e.* from 10^{-4} to 10^{-2}Hz. Typically this method is sensitive at the level $h \sim 10^{-13}$, but this may improve by a factor of 10 or more in the near future. (See articles in Hellings 1989.) Laser tracking could improve this to the level at which it would be likely to see the formation of giant $10^6 M_{\odot}$ black holes in the centers of galaxies. See Thorne (1987) for a recent review.

Pulsar timing is a technique for detecting a stochastic background of gravitational waves. The millisecond pulsar PSR 1937+214 is so stable that from the absence of big fluctuations in its apparent period one can set upper limits on the gravitational wave field the signals are passing through. This is most sensitive to waves with periods of the order of the time of observation. The present limit is that these waves contain an energy not more than 10^{-6} of the cosmological closure density (Thorne 1987), tantalizingly close to the density of 10^{-7} predicted by cosmic string theory if strings seed galaxy formation (Vachaspati & Vilenkin 1985). This limit should improve with time, but the extent to which it does will depend on whether Earth-based clocks can be made stable enough to measure (or limit) fluctuations in the pulsar's period. At the moment, the pulsar is just about as stable as any man-made clock. Discovery of a second or third very stable millisecond pulsar would dramatically change the situation, and would allow the pulsars to be used as an "interferometer" for the detection of very-low-frequency radiation.

A technique for the future that has been very seriously studied is the possibility of putting a laser interferometer in space. Consisting of three independent spacecraft arranged in an L-shape with arms

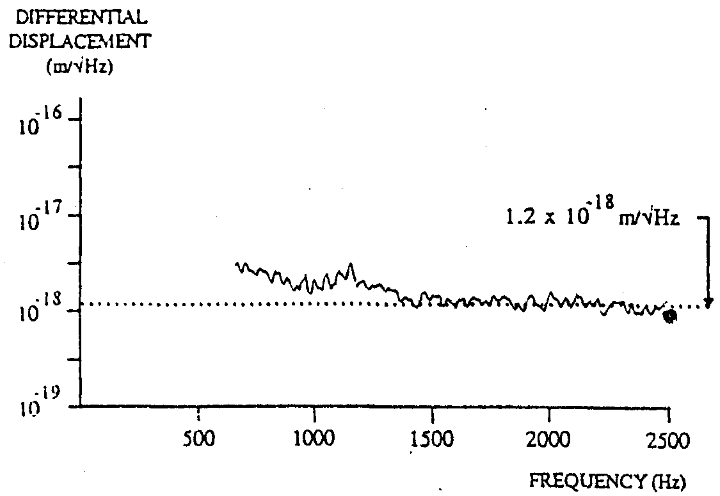

Figure 2: The noise performance of the Glasgow prototype in mid-1988.

of the order of 10^{-9} km, it might sit in the Earth's orbit about the Sun at some distance from the Earth. With realistic laser power, it could reach a sensitivity of 10^{-22} over a bandwidth of 10^{-4}–10^{-3} Hz. This would suffice to detect many galactic binaries, including the radiation from the Binary Pulsar PSR1913+16, as well as interesting cosmological sources, such as the formation of massive black holes in the centers of quasars.

3 Sources of Gravitational Waves

3.1 Supernovae and Gravitational Collapse

Supernovae have always been the source that detector builders aimed to detect; as the most violent event in our corner of the universe, it seemed the most likely candidate for detection. While coalescing binaries have to some extent displaced supernovae as the "most likely" source, at least for the laser interferometers, supernovae are still extremely important sources, particularly for the observing network of cryogenic bars in the next five years or more. Observations of supernovae, or of a gravitational collapse that is not accompanied by a strong electromagnetic outburst, would be extremely interesting for astrophysics. There is considerable uncertainty about the likely strength of gravitational waves from a supernova, and I will discuss this below. But first it is useful to see what sort of range detectors would have for a source of a given strength.

I will characterize the source by the total energy radiated in gravitational waves during approximately one millisecond; *for both $1M_\odot$ neutron stars and $10M_\odot$ black holes, this is the timescale of the biggest burst.* Some kinds of gravitational collapses may give strong wave-trains of much longer duration; provided that numerical calculations can provide us with predicted waveforms that we can use for pattern matching, such collapses may be even easier to detect than simple bursts of the same total energy. (This is because the radiation is likely to come out at a lower frequency: see Schutz 1989b for the reason why frequency matters.) I will consider two types of burst, which I will call strong and moderate.

- Strong bursts contain (by definition) about $0.1M_\odot c^2$ of energy in the gravitational waves. This is as much as any theorist might hope for from the formation of a neutron star, but it is a much more modest amount of energy to extract from a collapse that forms a $10M_\odot$ black hole. Such a burst would have an amplitude of

$$h \approx 6 \times 10^{-18} \text{ in our Galaxy,} \tag{1}$$

which could be seen by the cryogenic bars, or

$$h \approx 4 \times 10^{-21} \text{in the Virgo cluster,} \tag{2}$$

which would be visible to the first-stage laser interferometers.

- A moderate burst has, by my definition, ten times less energy: $0.01M_\odot c^2$. This is not an unreasonable amount of energy to come from a highly non-axisymmetric rotating collapse that forms a neutron star: the binding energy released is ten times larger, so this can easily be fit into the energy budget. If a $10M_\odot$ black hole is formed, the "moderate" burst is not very much more than axisymmetric collapses might generate. Its amplitude is down by roughly a factor of 3 from that of a strong burst, and in particular

$$h \approx 4 \times 10^{-22} \text{at a distance of 60 Mpc.} \tag{3}$$

This would be strong enough to be seen by laser interferometers when they achieve their design goals. Moreover, a network of detectors could measure locations on the sky to an accuracy of 1° or better (Gursel and Tinto 1989). This distance is at least three times larger than that to the Virgo cluster, which has commonly been regarded as the minimum distance for obtaining an interesting event rate. In this volume of space there are several starburst galaxies, where star formation and hence presumably supernovae are occurring at a much higher rate than in most galaxies. There might well be thousands of supernovae per year in the volume out to 60 Mpc. The nearest "starburst" galaxy, M82, is only 3 Mpc away, and has a supernova once every few years (Kronberg, *et al*, 1985).

Now it is time to return to the uncertainties surrounding supernovae as gravitational wave sources, which arise because a perfectly spherically symmetric collapse would produce no gravitational radiation at all. In order for a Type II supernova to produce even a moderate burst of radiation, the collapsing core must presumably become very asymmetric. The most plausible candidate for producing asymmetry is rotation. If the collapsing core contains enough angular momentum to make rotation dynamically dominant when it reaches neutron-star densities (and this is not much angular momentum — the Sun has much more than this), then the non-axisymmetric "bar mode" instability may be excited, and the resulting tumbling cigar-shaped core would be a strong source of gravitational waves.

There are two problems making this scenario quantitative. The first is that numerical calculations are not yet good enough to predict how much radiation will come from a core with a given amount of angular momentum. As computers improve, this problem will probably be solved, but it will not be easy: detailed nuclear physics and neutrino transport need to be included. A first step in this direction is described by Bludman in his talk at this Symposium. The second problem is that we don't know what initial conditions are likely: how much rotation does the pre-collapse degenerate core typically have? Observations of pulsars suggest that they may be formed with relatively slow rotation rates: the Crab pulsar is young and may never have been rotating near to its breakup velocity. This argues that there is little rotation in a collapsing core. On the other hand, SN1987a offers some suggestion that rotation may have been important in its collapse: polarization measurements (Cropper, *et al*, 1988) suggest that the expanding cloud has an elliptical shape, and the axis of the ellipse lines up with the direction to the mysterious "spot" that appeared in speckle photographs (Meikle, *et al*, 1987; Nisenson, *et al*, 1987). Rotation must be a strong candidate for the mechanism that produces these alignments.

There are other important questions that need more work before they can be answered. The radiation that can be expected from Type I supernovae is also uncertain: if the exploding star completely disintegrates, then none can be expected. But if occasionally a neutron "cinder" is left behind, then the original star probably has enough angular momentum to produce a strongly deformed core.

Another question is whether there are electromagnetically quiet collapses: do all gravitational collapses result in visible supernovae? The statistics of stellar births, deaths, supernovae, and pulsars are not good enough to exclude this possibility. On theoretical grounds, there is some reason to think that quiet collapses may be common. Current successful Type II supernova models are delicately balanced: the escaping neutrinos provide only just enough energy to power the shock outwards through the envelope and blow off the envelope. If one has a strongly rotating collapse, the lower core densities and longer timescale may weaken the shock and prevent it from blowing off the outer shell. There could therefore be an anticorrelation between electromagnetic and gravitational radiation intensities from a gravitational collapse: the strongest gravitational wave emitters may have the weakest explosions, and may therefore even go on to form black holes.

It may be that only gravitational wave observations of supernovae will settle some of these questions. But it would surprise me if the answers conspired to give no observations of bursts in laser interferometer detectors. Of all the supernovae out to 60 Mpc, if even only one percent produce moderate bursts, the network could register tens of events per year. And strong bursts could be seen to nearly 200 Mpc. There may be hundreds of black holes like Cygnus X-1 formed per year in that volume.

3.2 Gravitational Waves from SN1987a

It is one of the great frustrations of gravitational wave astronomers that the cryogenic bars were not taking data at the time of the supernova. What sort of waves would have been expected? Even if the burst was exceptionally strong, say $1.0 M_\odot c^2$ of energy in the radiation itself, its amplitude would have been only about $h \sim 5 \times 10^{-18}$, just within range of the cryogenic bars. The room-temperature bars at Rome and Maryland were sensitive at about 3×10^{-17}. No significant coincidences have been reported between them at the time of the supernova (as defined by the time of arrival of the bursts of neutrinos in any of the neutrino detectors).

However, it has been reported that there *were* excess correlations between events in the neutrino detectors and the excitation of the bar detectors 1.2 s before each neutrino event, over a two-hour period including the supernova, and Pizzella describes the latest results on that elsewhere in this volume. I will remark here only that if these events are gravitational waves from the Large Magellanic Cloud, then even at the 1σ level in a 3×10^{-17} detector, each event would represent the release of about $40 M_\odot c^2$ of energy in gravitational radiation. They therefore cannot be gravitational waves if general relativity is correct and their source is physically associated with the supernova. Other plausible theories of gravity can be excluded as well (de Rujula 1989).

3.3 Coalescing Compact-Object Binaries

In its final moments some 10^8 years from now, the "Binary Pulsar" PSR 1913+16 will be a pair of neutron stars in a perfectly circular orbit that is decaying rapidly because of gravitational radiation reaction. The orbit is well approximated as a simple Newtonian point-particle orbit with quadrupole radiation reaction (Peters & Mathews 1963). Tidal, post-Newtonian or other effects *do not begin to make it deviate* from this behavior until the gravitational waves have a frequency of some 500 Hz or so (Clark & Eardley 1977, Krolak & Schutz 1987).

As was first pointed out in a remarkable early paper by Forward & Berman (1967), these sources are good candidates for detection by broad-band detectors. The energy carried off by the waves from the orbit will be enormous: some $6 \times 10^{-3} M_\odot c^2$ as the radiation frequency rises from 100 to 200 Hz. Moreover, the very predictable wavetrain allows matched filtering to be applied to the data, which allows one to dig into the noise and see these systems at very great distances.

The maximum amplitude is along the axis of the orbital angular momentum. These gravitational waves have amplitude

$$h_{max} = 3.6 \times 10^{-23} \left(\frac{M_T}{2.8 M_\odot}\right)^{2/3} \left(\frac{\mu}{0.7 M_\odot}\right) \left(\frac{f}{100\mathrm{Hz}}\right)^{2/3} \left(\frac{100\mathrm{Mpc}}{r}\right), \tag{4}$$

where M_T is the total mass and μ the reduced mass of the binary. The amplitude in other directions is reduced by angular factors of order 1, and the *detected* amplitude is further reduced by antenna-pattern

effects. With a network of three or four laser interferometers it will be possible to unravel these angular factors and reconstruct h_{max} from the observations.

The orbit decays due to gravitational radiation reaction, and the orbital period decreases. The formal coalescence timescale is given by

$$\tau := \frac{f}{\dot{f}} = 5.6 \left(\frac{M_T}{2.8 M_\odot}\right)^{-2/3} \left(\frac{\mu}{0.7 M_\odot}\right)^{-1} \left(\frac{f}{100\mathrm{Hz}}\right)^{-8/3} sec. \tag{5}$$

The true time to coalescence is 3/8 of this. If we look at these two equations, we see that the way that the (generally unknown) stellar masses enter the relations is the same in both cases. Therefore, the product of the two observables is independent of the masses: *$h_{max}\tau$ depends only on r*, the distance to the binary (Schutz 1986). Coalescing binaries are therefore standard candles, and this makes them useful for many purposes, including the determination of Hubble's constant (Schutz 1986). I will return to this below.

Conversely, detection of the coalescing binary does not determine the individual masses of the stars. It only tells us the *mass parameter* of the system, defined by:

$$\mathcal{M} = m_1^{3/5} m_2^{3/5} / (m_1 + m_2)^{1/5}, \tag{6}$$

or equivalently by the more transparent formula,

$$\mathcal{M}^{5/3} = \mu M_T^{2/3}. \tag{7}$$

Only if the signal is strong enough to detect post-Newtonian orbital effects in the waveform will it be possible to determine the individual stellar masses (Krolak & Schutz 1987).

By using pattern-matching techniques, it should be possible for the second-stage laser interferometers to detect coalescing binaries out as far as 1 Gpc, where the cosmological redshift is 0.1–0.2 (Thorne 1987, Krolak 1989). A network of four detectors would detect 5–10% of all the sources in this volume (Tinto 1989).

As with gravitational collapse, there are considerable uncertainties about coalescing binaries as gravitational wave sources. The most important is their event rate. An early estimate suggested there would be 3 per year out to 100 Mpc (Clark, *et al* 1979). If we extrapolate this to our expected range of 1 Gpc, and allow for the fraction we expect to detect, we find that this gives an event rate in the network of detectors of 150–300 per year.

But the Clark, *et al*, event rate is based on a number of assumptions, not least of which is that the one coalescing binary precursor system that we do observe — the Binary Pulsar — is representative of its class. We would learn much more about the event rate if we could discover other precursors. In this connection, the recently discovered binary pulsar in the globular cluster 47 Tuc is potentially very important (Ables, *et al* 1988); but we must await confirmed measurements of the orbital elements and their rates of change. At the moment, I would estimate that it would be fair to say that the event rate is uncertain by a factor of 10-100 either way. Even the pessimistic lower bound would give a *few* events per year in a network of detectors. This is why I feel that coalescing binaries are the most likely source that laser interferometers will detect in the long run.

Another approach to understanding coalescing binaries is to study their evolutionary history. How many binary evolutionary paths could lead to systems with two compact objects close enough to have an orbital lifetime less than 10^{10} years? What fraction of them could contain black holes of $10 M_\odot$ or more? If the fraction of compact X-ray binaries in our Galaxy that have black holes is typical of coalescing binary precursors, then perhaps 1% of coalescences would involve a neutron star and a black hole, while 10^{-4} might involve two black holes. Since systems with two $14 M_\odot$ black holes could be seen roughly ten times further away than two-neutron-star systems, they could account for as much as ten percent of the total number of systems detected. Another question concerning the formation of coalescing binary precursors is whether conditions in evolved globular clusters lead to the formation of such systems more often than in the disk of the Galaxy. It is interesting that 47 Tuc is one of the most centrally condensed clusters, and it may have just undergone central core collapse and reheating by the formation of a few dozen binary stars. If a good fraction of globular clusters go through this cycle many times, how often will coalescing binary precursors be formed?

It might be possible to get further information about coalescing binaries if we knew what the coalescence event looked like in electromagnetic radiation. An interesting suggestion by Blinnikov *et al* (1984) is that in a system with unequal-mass neutron stars, the more massive star will strip the less massive one, and this will go on until the less massive one reaches the minimum mass of a neutron star, about $0.1 M_\odot$. At this point the small neutron star is unbound and will explode. What will this explosion look like? Will it look like a weak Type I supernova? Will it give a detectable X- or gamma-ray burst? What will happen to the primary star? If it accepts all the mass shed by the secondary, it may exceed the upper limit on the mass of a neutron star, and it may collapse to a black hole. Theoretical study of these questions could shed considerable light on the question of whether coalescing binaries will be detected frequently or at all.

If coalescing binaries are detected at the rate predicted by Clark, *et al*, then they become very interesting for astronomy:

- Nearby events (within 100 Mpc) could be used to determine Hubble's constant with an accuracy of a few percent. Our recent estimates of the accuracy of the angular positions that a network will be able to infer are much better than I assumed when I first suggested this method (Schutz 1986). It may take only a handful of events to pin down H_0 this way.
- If we detect a few hundred events per year from 500 Mpc and beyond, each with a distance and an angular position, we can do statistical studies of the stellar distribution on scales not heretofore possible. Gravitational waves are not obscured by any intervening matter, so they represent an ideal survey medium.
- Each coalescing binary event gives us a measure of the masses of the component stars, the mass parameter $\mathcal{M} = \mu^{3/5} M_T^{2/5}$. Statistical studies of this could help us infer the neutron star mass function.
- Coalescences involving black holes could be the most reliable method available for positive identification of black holes. Those involving two black holes can be modeled by computers, and therefore provide a test of general relativity.
- Coalescences of two black holes seen at $z \sim 1.0$ are bound to give us new statistical information about conditions at that epoch. They may also be candidates for gravitational lensing (Krolak & Schutz 1987). The different lensed components could not be resolved spatially, but the events would arrive at different times; the time-delay would therefore be known extremely accurately (to less than a second). The elliptical polarization of the generic coalescing binary gravitational wave would allow one to distinguish positive from inverted images.

Figure (3) summarizes the great variation in the effective ranges of expected gravitational wave detectors, when looking at supernovae and coalescing binaries.

As a sobering final thought on the subject, let us recall that it will be some time before abundant data on coalescing binaries becomes available: if the Clark *et al* event rate is correct, then only when laser interferometers reach a sensitivity of 10^{-22} can they begin to do serious observing. This may be ten years from now.

3.4 Other Sources: Pulsars and the Stochastic Background

Pulsars. Pulsars are particularly interesting sources to search for, because the gravitational radiation they may emit will not be beamed: the statistics of a "complete" search of the solar neighborhood would not be affected by the uncertainties due to beaming factors that afflict radio searches. Moreover, since the radiation would presumably be due to crustal deformations, old radio-quiet pulsars could still in principle be sources of gravitational waves. On the other hand, there is considerable doubt about the amplitude of any gravitational waves that may be emitted, and it may be that only by observing gravitational waves from pulsars or setting upper limits on their amplitudes will we be able to settle the question.

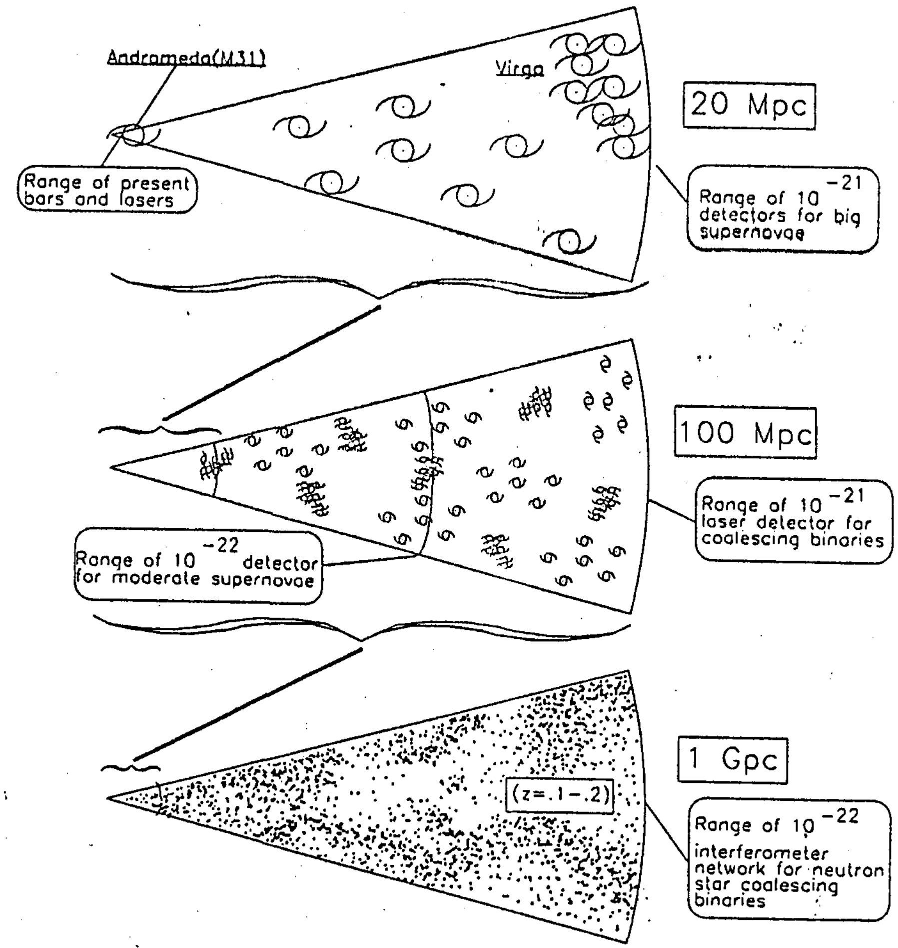

Figure 3: Wedge-shaped sections of the universe displayed on different scales to indicate the ranges of different detectors for supernovae and coalescing binaries. The observer sits at the apex of the cone.

Present observational limits come from the Tokyo experiment (Owa, *et al*, 1986), which can reach down to about $h \sim 3 \times 10^{-22}$ at 60 Hz. (The dominant quadrupole radiation will come out at twice the frequency of rotation of the pulsar, just as it does for binary systems.) This limit is far above an indirect one of 10^{-24}, which is the maximum allowed if the entire spindown of the Crab pulsar is due to its losing energy to gravitational waves. Interferometers in their ultimate configuration could in principle reach to 10^{-27} at 60 Hz, provided they can beat the seismic noise at such a low frequency.

Finding known pulsars is one thing: making an all-sky search in gravitational waves is another. The primary problem is that long integration times (typically 10^7 s) are required, and during such long times the Earth's various motions produce very significant Doppler shifts in the pulsar signal. Moreover, these shifts depend on the position of the pulsar in the sky. Removing these shifts is computationally prohibitive for a search by a single detector (Schutz 1989c). But a less-sensitive all-sky search is possible using cross-correlation of three or more interferometers. Gursel & Tinto (1989) have shown how to remove the effects of different instrumental polarizations in order to reach the maximum sensitivity. Their method works provided the detectors are sufficiently close that there is no significant Doppler shift of the signal in one relative to the other. For millisecond pulsars the detectors need to be within about 600 km of one another, but a world-wide network could search for pulsars efficiently below 100 Hz.

Wagoner stars. Wagoner (1984) pointed out that an accreting neutron star that is spun up to the angular velocity at which a non-axisymmetric mode of the star becomes unstable to gravitational radiation (a CFS instability: see Thorne 1987, Schutz 1987) will become a gravitational wave beacon, emitting gravitational waves at exactly the rate required to carry off any further accreted angular momentum. Such a source would be an X-ray source as well, and in fact the gravitational wave luminosity would be proportional to the X-ray luminosity. There are several Galactic X-ray sources in which this scenario is possible and whose luminosity is large enough to suggest that interferometers could see their gravitational waves, most particularly Sco X-1 (Thorne 1987). They could be searched for in the same way as for pulsars.

Stochastic background. A stochastic background can also be found by cross-correlation between detectors. The detectors should be no more than a reduced wavelength of the gravitational waves apart, for maximum sensitivity. There are many possible sources for such backgrounds: cosmic strings (Vachaspati & Vilenkin 1985; Brandenberger, *et al*, 1986); very massive objects — VMOs — formed in the early universe (Bond & Carr 1984); or phase transitions in the early universe (Thorne 1987). In particular, the cosmic string scenario makes a fairly definite prediction of the energy density of the background: if cosmic strings are seeds for galaxy formation, then the gravitational radiation they produce as they decay should contain at present about 10^{-7} of the closure density. Laser interferometers could reach to 10^{-9} of closure, so they could test this prediction at 100 Hz. It may also be tested soon by observations of very stable millisecond pulsars, but this would be at low frequencies (10^{-9} Hz) and it will depend on progress in the development of very stable laboratory clocks.

3.5 Low-Frequency Sources

Although I have concentrated on sources of gravitational waves at frequencies above about 50 Hz, which would be observable from the ground, there is considerable interest at present in the design of space-based detectors that can reach as low as $h = 10^{-22}$ at 10^{-3} Hz (Hellings 1989). There are a number of interesting sources at these frequencies, that make the development of such detectors desirable. These include:

- Formation and coalescence of black holes of mass greater than $10^6 M_\odot$ in the centers of galaxies.
- Binary systems in our Galaxy: there are many systems with periods of a few hours or less that could in principle be observed. Binaries consisting of two white dwarfs are particularly interesting (Evans, *et al*, 1987). Neutron-star binaries that will eventually coalesce could be identified at a much earlier stage of their evolution. In fact, there are so many binaries in this frequency range that they will become a nuisance, acting as a noise background against which it will be harder to detect bursts from massive black hole events in distant galaxies.

- Stars falling into nearby galactic black holes. Even nearby galaxies may have "dead" black holes in their nuclei, and neutron stars may occasionally fall into them. Such an encounter would emit a detectable burst of radiation.

For more details, see Thorne (1987), Schutz (1989b,d).

4 A Gravitational Wave Astronomer's Shopping List of Astrophysical Problems

There are many questions that lead to uncertainties in our predictions about gravitational wave sources that do not need to wait for gravitational wave observations before they can be solved. The following very personal "shopping list" is offered in the hope that it will encourage astrophysicists to give some needed attention to these problems during the time when the large laser interferometers are being constructed. The list is by no means exhaustive; it simply represents the problems that I personally would most like to see progress on during the next five years.

- Incorporate significant rotation into supernova collapse models. Although this is a long-term goal of numerical general relativity, present highly-developed spherical collapse codes that do the nuclear physics correctly could be modified within Newtonian gravity to examine rotation. Even an extension and updating of the work of Müller & Hillebrandt (1981) on axisymmetric collapse would be very useful. One would especially want to treat neutrino transport correctly (see Bludman in this volume). Given the well-known sensitivity of present collapse/supernova models to details of the nuclear physics, it would be interesting to see how realistic amounts of rotation might modify the picture.
- Are there electromagnetically quiet gravitational collapses? If one has confidence in at least a spherically symmetric collapse code, then it would be interesting to explore some sort of reasonable parameter space of initial conditions and physical uncertainties to see whether there is a class of stars that could collapse without a big explosion.
- Model the coalescence of two neutron stars from circular orbits. There are really three phases of coalescence that need to be better understood:
  - Model the interactions of the stars once they are too close to be treated as point masses. This could be attacked with analytic calculations and numerically. As a numerical problem, it is one of the most demanding goals one could set for a numerical relativity program. Head-on collisions of neutron stars have been treated by Evans (1989), and a remarkable first attempt within Newtonian gravity at modeling coalescence of two stars that start out in roughly circular orbits with their surfaces in contact has been reported by Nakamura (1989). As computers improve and more physics can be put in, one would like to see tidal effects and mass transfer modeled accurately. Is there a long "coasting" stage in which the orbital frequency is roughly constant as mass is transferred? Such models would help us to extend the filters we use for digging coalescence signals out of the noise beyond the point-mass stage, increasing the range of any detectors and improving the signal-to-noise ratio of any detection. This would have the further effect of improving the determination of the angular position and distance of any event.
  - Model the explosion of a mini-neutron star. The late stage of mass transfer between two neutron stars or a neutron star and a black hole may result in the stripping of a neutron star down to its minimum mass, at which point it would explode. What would such an explosion look like? Are there observational tests or searches one could perform?
  - Model the subsequent development of the primary star in the binary. If the secondary explodes because it has too little mass, does the primary in its turn collapse because it gains too much? What electromagnetic radiation would such a collapse produce? Given that the primary is probably rotating rapidly because of the accreted angular momentum, does its collapse produce

a long wavetrain of gravitational waves? If so, then numerical predictions of this wavetrain can be used to produce filters for the incoming data that will aid in detecting any collapses that may follow detected coalescing binary events.

- What is the coalescing binary event rate? It may be possible to provide much more confident predictions of this from studies of binary star evolution and from observational tests of the predictions of coalescence models, such as models for the explosion of mini-neutron stars.

In his Summary of the first Texas Symposium, Peter Bergman (1964) wrote that we could hope that gravitational waves will be detected, but only in the "rather distant future". Twenty-five years later, that hope is rather more concrete. If the funding authorities cooperate, and if our predictions regarding gravitational wave sources are not too far wrong, then in the *foreseeable* future — maybe at the 10^{th} Texas Symposium in 1998, maybe well before — we should have a talk, not just about gravitational wave detectors, but about gravitational wave *astronomy*. In the meantime, it is not just the gravitational wave experimentalists who will be working hard: the more theoretical understanding we have about likely gravitational wave sources, the easier it will be to dig their radiation out of the noise when the detectors come on line.

References

Ables, J.G., Jacka, C.E., McConnell, D., Hamilton, P.A., McCulloch, P.M., and Hall, P.J., 1988, *I.A.U. Circ*, no. 4602.

Balbinski, E.F.L., Detweiler, S., Lindblom, L., and Schutz, B.F., 1985, *Mon. Not. R. astr. Soc.*, 213, 553 - 561.

Bergman, P.L., 1964, in Robinson, I., Schild, A., and Schucking, E.L., eds., *Quasi-Stellar Sources and Gravitational Collapse* (University of Chicago Press, Chicago), 431 - 432.

Blair, D., ed., 1989, *Gravitational Radiation* (Cambridge University Press, Cambridge England).

Blinnikov, S.I., Novikov, I.D., Perevodchikova, T.V., and Polnarev, A.G., 1984, *Sov. Astron. Lett.*, 10, 177.

Bond, J. R., and Carr, B. J., 1984, *Mon. Not. R. astr. Soc.*, 207, 585.

Brandenburger, R.H., Albrecht, A., and Turok, N., 1986, *Nucl. Phys. B*, 277, 605.

Caves, C.M., Thorne, K.S., Drever, R.W.P., Sandberg, V.D., and Zimmerman, M., 1980, *Rev. Mod. Phys.*, 52, 341.

Clark, J. P. A., and Eardley, D. M., 1977, *Astrophys. J.*, 215, 315.

Clark, J. P. A., van den Heuvel, E. P. J., and Sutantyo, W., 1979, *Astron. & Astrophys.*, 72, 120.

Cropper, M., Bailey, J., McCowage, J., Cannon, R.D., Couch, W.J., Walsh, J.R., Strade, J.O., and Freeman, F., 1988, *Mon. Not. R. astr. Soc.*, 321, 695 - 722.

Damour, T., 1987, in Hawking, S.W., and Israel, W., eds., *300 Years of Gravitation* (Cambridge University Press, Cambridge, England).

de Rujula, A., 1989, in Blair, D., ed., *Proceedings of the Fifth Marcel Grossmann Meeting* (to appear).

Ehlers, J., 1980, *Ann. N. Y. Acad. Sci.*, 336, 279 - 293.

Evans, C.R.,, 1989, in Evans, C., Finn, L.S., and Hobill, D., eds., *Frontiers of Numerical Relativity* (Cambridge University Press, New York), in press.

Evans, C.R., Finn, L.S., and Hobill, D., eds., 1989, *Frontiers of Numerical Relativity* (Cambridge University Press, New York), in press.

Evans, C.R., Iben, Jr., I., and Smarr, L., 1987, *Astrophys. J.*, 323, 129 - 139.

Forward, R.L., 1978, *Phys. Rev. D*, 17, 379.

Forward, R.L., and Berman, D., 1967, *Phys. Rev. Lett.*, 18, 1071 - 1074.

Forward, R.L., and Moss, G.E., 1972, *Bull. Amer. Phys. Soc.*, 17, 1183.

Gursel, Y., and Tinto, M., 1989, *preprint*.

Hellings, R.W., 1989, *Proceedings of a NASA Workshop on Experimental Gravitation in Space* (NASA, to appear).

Krolak, A., 1989, in Schutz, B.F., ed., *Gravitational Wave Data Analysis* (Kluwer, Dordrecht), 59 - 69.

Krolak, A., and Schutz, B.F., 1987, *Gen. Rel. Gravit.*, 19, 1163 - 1171.

Kronberg, P.P, Biermann, P., and Schwab, F.R., 1985, *Astrophys. J.*, 291, 693 - 707.

Meikle, W.S.P., Matcher, S.J., and Morgan, B.L., 1987, *Nature*, 329, 608.

Moss, G.E., Miller, L.R., and Forward, R.L., 1971, *Appl. Optics*, 10, 2495.

Müller, E., and Hillebrandt, W., 1981, *Astron. Astrophys.*, 103, 358.

Nakamura, T., 1989, in Evans, C., Finn, L.S., and Hobill, D., eds., *Frontiers of Numerical Relativity* (Cambridge University Press, New York), in press.

Nisenson, P., Papaliolios, C., Karovska, M., and Noyes, R., 1987, *Astrophys J. Lett.*, 320, L15.

Owa, S., Fujimoto, M.-K. Hirakawa, H., Morimoto, K., Suzuki, T., and Tsubono, K., 1986, in Ruffini, R., ed., *Proceedings of the Fourth Marcel Grossmann Meeting* (North-Holland, Amsterdam).

Peters, P.C., and and Mathews, J., 1963, *Phys. Rev.*, 131, 435.

Schutz, B.F., 1986, *Nature*, 323, 310 - 311.

Schutz, B.F., 1987, in Carter, B., and Hartle, J., eds., *Gravitation in Astrophysics* (Plenum, New York).

Schutz, B.F., ed., 1989a, *Gravitational Wave Data Analysis* (Kluwer, Dordrecht).

Schutz, B.F., 1989b, *Class. Quant. Grav.*, (submitted).

Schutz, B.F., 1989c, in Blair, D., ed., *Gravitational Radiation* (Cambridge University Press, Cambridge England).

Schutz, B.F., 1989d, in Schutz, B.F., ed., *Gravitational Wave Data Analysis* (Kluwer, Dordrecht), 3 - 18.

Thorne, K.S., 1987, in Hawking, S.W., and Israel, W., eds., *300 Years of Gravitation* (Cambridge University Press, Cambridge, England), 330 - 458.

Tinto, M., 1989, in Schutz, B.F., ed., *Gravitational Wave Data Analysis* (Kluwer, Dordrecht), 299 - 313.

Vachaspati, T., and Vilenkin, A., 1985, *Phys. Rev. D*, 31, 3052.

Vinet, J.-Y., Meers, B., Man, C. Nary, and Brillet, A., 1988, *Phys. Rev. D*, 38, 433-447.

Wagoner, R. V., 1984, *Astrophys. J.*, 278, 345.

Reference (6) (gf YG memo on data syst.)

Data Processing, Analysis, and Storage for Interferometric Antennas

Bernard F. Schutz
Department of Physics
University of Wales College of Cardiff
Cardiff, Wales

January 10, 1989

Abstract

The problem of detecting gravitational wave sources with broad-band interferometric detectors is characterized by the need to store and process large amounts of data, using pattern-matching techniques to find weak radiation buried in the noise. I review the principal features of this problem: setting thresholds for reliable detection, using matched filtering to find coalescing binaries and determine their positions in the sky, the demands that real-time filtering makes on computing machinery, the difficulties of the pulsar-search problem, using cross-correlation to conduct an all-sky pulsar search, using a network of detectors to reconstruct the gravitational wave (the inverse problem), and difficulties of data storage and exchange. Many results are presented here for the first time.

1 Introduction

Laser-interferometric gravitational wave antennas face one of the most formidable data handling problems in all of physics. The problem is compounded of several parts: the data will be taken at reasonably high data rates (of the order of 10 kHz of 12-16 bit data); they may be accompanied by four times as much "housekeeping" data to ensure that the system is working appropriately; the data will be collected 24 hours a day for many years; the data need to be searched in real time for a variety of rare, weak events of short duration (1 second or less); the data from three or more detectors need to be cross-correlated with each other; and the data need to be archived in searchable form in case later information makes a re-analysis desirable. One detector might generate 300 Mbytes of data each hour. Even using optical discs or digital magnetic tapes with a capacity of 2 Gbytes, a network of 4 interferometers would generate almost 5000 discs or tapes per year. The gathering, exchange, analysis, and storage of these data will require international agreements on standards and protocols. The object of all of this effort will of course be to make astronomical observations. Because the detectors are nearly omni-directional, a network of at least three and preferably more detectors will be necessary to solve the "inverse problem", *i.e.* to reconstruct a gravitational wave event completely, from which the astronomical information can be inferred.

In this chapter I will discuss the mathematical techniques for analysing the data and solving the inverse problem, the technical problems of handling the data, and the possibilities for international cooperation, as they appear in late 1988. This discussion can only be a snapshot in time, and a personal one at that. The subject is one that can be expected to develop considerably in the next decade. I will orient the discussion toward ground-based interferometers, with the sensitivity and spectral range expected of the instruments that are planned to be built in the next decade. Much of the discussion naturally is equally applicable to present prototypes, but it is important to look ahead towards future detectors so that their data problems can be anticipated in their design. A large part of the section on data analysis also applies to space-based interferometers or to the analysis of ranging data for interplanetary spacecraft, although in these cases the volume of data is much lower because they operate as low-frequency detectors. I will

To appear in D. Blair, *Gravitational Radiation* (Cambridge University Press, Cambridge, England 1989).

also assume that the interferometers will operate in a reasonably broad-band mode, even though there exist a number of techniques for enhancing their sensitivity in narrower bandwidths. In the extreme narrow-banding case, in which the detectors have a bandwidth smaller than that of the waves, the data analysis problem resembles that for bar detectors, as discussed by Pizzella elsewhere in this volume.

2 The Analysis of the Data

2.1 Signals to look for

The likely sources of gravitational radiation are described by David Blair in the first part of this book. If a source is strong enough to stand out above the noise in the time-series of data coming off the machine, then simple threshold-crossing criteria can be used to isolate candidate events. If the event is too weak to be seen immediately, it may still be picked up by pattern-matching techniques, but the sensitivity to such events will depend upon how much information we have about the expected waveform. At the present time, we have little idea of what waveform to expect from bursts of radiation from gravitational collapse (supernovae or electromagnetically quiet collapses), so their detectability depends upon their being strong enough to stand up above the broad-band noise. (Future detailed numerical calculations of gravitational collapse may change this, of course.) On the other hand, we have detailed predictions for the waveforms from binary coalescence and from continuous-wave sources such as pulsars; these can be extracted from noisy data by various techniques, such as matched filtering. Pulsars with a known position may be found from the output of a single detector by sampling techniques. An all-sky search for unknown pulsars can be performed by cross-correlating the output of three or more detectors. There is also the possibility of a stochastic background of radiation; cross-correlation techniques between detectors can also search for this.

2.2 Analysis of the data from individual detectors

Bursts and continuous-wave signals can in principle be detected by looking at the output of one instrument. Of course, one would like to have coincident observations of the same waves in different detectors, for several reasons: to increase one's confidence that the event is real, to improve the signal-to-noise ratio of the detection, and to gain extra information with which to solve the inverse problem. It might be thought that the detection problem splits into two parts: first find the events in single detectors, then correlate them between detectors. In many cases this will work, but in some cases it will only be possible to detect signals in the first place by cross-correlating the output of different detectors. In this section I will address the problem of finding candidate events in single detectors. Cross-correlation will be treated later.

2.2.1 Finding broad-band bursts

A broad-band burst is an event whose energy is spread across the whole of the bandwidth of the detector (which I will take to be something like 100–2000 Hz), although considerable efforts are now being devoted to techniques for extending the bandwidth down to 40 Hz or less. To be detected it has to compete against all of the detector's noise, and the only way to identify it is to see it cross a pre-determined amplitude threshold in the time-series of data coming from the detector. The main burst of radiation from stellar core collapse will be like this. Numerical simulations of axisymmetric collapse (Evans 1986, Piran & Stark 1986) reveal, among other things, that after the main burst there is – at least if a black hole is formed – a "ringdown phase" in which the radiation is dominated by the fundamental quasi-normal mode of the black hole. This phase lends itself to some degree of pattern-recognition, such as that which I will describe for coalescing binaries in the next section. But it is unlikely that ringdown radiation will substantially improve the signal-to-noise ratio of a collapse burst, since it is damped out very quickly. Some simplified models of non-axisymmetric collapse (*e.g.* Ipser & Managan 1984) suggest that if angular momentum dominates and non-axisymmetric instabilities deform the collapsing object into a tumbling tri-axial shape, then a considerable part of the radiation will come out at a single slowly-changing frequency. If future three-dimensional numerical simulations of collapse bear this out, then this would also be a candidate for pattern-recognition. But one must bear in mind that even if we have good predictions of waveforms from simulations, there will be an intrinsic uncertainty due to our complete lack of knowledge of the initial conditions we might expect in a collapse, particularly regarding the angular momentum of the core. So it is not clear yet whether collapses will ever be easier to see than the time-series threshold criteria described next would indicate.

2.2.1.1 Simple threshold criteria. The idea of setting thresholds is to exclude "false alarms" — apparent events that are generated by the detector noise. Thresholds are set at a level which will guarantee that any collection of events above the threshold will be free from contamination from false alarms at some level. The "guarantee" is of course only statistical, and it relies on understanding the noise characteristics of the detector. I will assume here that the noise is white and Gaussian.

This should be a good first approximation, but there are at least two important refinements: first, detector noise is frequency-dependent, and when we consider coalescing binaries this will be important; and second, one must allow for unmodelled sources of noise that will occasionally produce large-amplitude "events" in individual detectors. This latter noise can be eliminated by demanding coincident observations in other detectors, provided we assume that it is independent of noise in the other detectors and that it is not Gaussian, in particular that there are fewer low-amplitude noise events for a given number of large-amplitude ones than we would expect of a Gaussian distribution. This implies that the cross-correlated noise between detectors will be dominated by the Gaussian component. These assumptions are usually made in data analysis, but it is important to check them as far as possible in a given set of data.

Assuming that the noise amplitude n in any sampled point has a Gaussian distribution with zero mean and standard deviation σ, the probability that its absolute value will exceed a threshold T (an event that we call a "false alarm" relative to the threshold T) is

$$p(|n|>T) = (\frac{2}{\pi})^{1/2} \int_T^\infty e^{-n^2/2\sigma^2} dn \approx (\frac{2}{\pi})^{1/2} (\frac{\sigma}{T} - \frac{\sigma^3}{T^3} + \ldots) e^{-T^2/2\sigma^2}. \quad (1)$$

In the asymptotic approximation given by the second line, the first term gives 10% accuracy for $T > 3.2\sigma$, and the first two terms give similar accuracy for $T > 2.5\sigma$. If we want the expected number of false alarms to be one in N_{obs} data points, then we must choose T such that

$$p(|n|>T) = 1/N_{obs}. \quad (2)$$

This is a straightforward transcendental equation to solve. For example, if we imagine looking for supernova bursts of a typical duration of 1 ms, then we might be sampling the noise in the output effectively 1000 times per second. (If we want to reconstruct the waveform we might the data at its raw sampled rate, say 4 kHz; but this would require a larger signal-to-noise ratio than simple detection, for which we could use the data sampled at or averaged over 1 ms intervals.) If we wish no more than one false alarm per year, then we must choose $T = 6.6\sigma$.

If we have two detectors, with independent noise but located on the same site, then we can dig deeper into the noise by accepting only *coincidences*, which occur when both detectors simultaneously cross their respective thresholds T_1 and T_2. Given noise levels σ_1 and σ_2, respectively, the criterion for the threshold is

$$p(|n|>T_1)p(|n|>T_2) = 1/N_{obs}. \quad (3)$$

For two identical detectors ($\sigma_1 = \sigma_2$), each making 1000 observations per second, the threshold T needs to be set at only $4.5\sigma_1$ to give one false alarm per year. Similarly, three identical detectors on the same site require $T = 3.6\sigma$ and four can be set at $T = 3.0\sigma$. The improvement from two to four detectors is a factor of 1.5 in sensitivity, or a factor of 3 in the volume of space that can be surveyed, and hence a similar improvement in the expected event rate. This favorable cost/benefit ratio — in this case, a factor of three improvement in event rate for a factor of two increase in expenditure — is characteristic of networks of gravitational wave detectors, and indeed of any astronomical detector network whose sensitivity is limited by internal noise uncorrelated between instruments.

2.2.1.2 Threshold criteria with time delays. I have qualified the discussion of multiple detectors so far by demanding that they be on the same site; the reason is that if they are separated, then allowing for the possible time delay between the arrival of a true signal in different detectors opens up a larger window of time in which noise can masquerade as signal. Suppose that two detectors are separated by such a distance that the maximum time delay between them is W measurement intervals. (For example, Glasgow and California are separated by about 25 ms, which we take to be effectively ± 25 measurement intervals for collapse events. This gives a total window size of 50 measurements.) Then in Eq. (3), the appropriate probability to use on the right-hand side is $1/N_{obs}/W$, since each possible "event" in one detector must be compared with W possible coincident ones in the other. For two identical detectors, a "typical" window $W = 40$ raises the threshold T to 4.9σ. This is an 8% decrease in sensitivity, or a 22% decrease in volume surveyed.

For three detectors, the situation begins to get more complex: as we will see later, if three detectors see an event that lasts considerably longer than their resolution time, there is a self-consistency check which may be used to reject spurious coincidences. (The check is that three detectors can determine the direction to the source, which must of course remain constant during the event.) For four detectors, even a few resolution times are enough to apply a self-consistency check. In principle, the quantitative effect of these corrections will depend on the signal-to-noise ratio of the event, since strong events can be checked for consistency more rigorously than weak events can. But the level of the threshold in turn will determine the minimum signal-to-noise ratio. A full study of this problem has not yet been made, and can probably only be undertaken in the light of a more thorough investigation of the inverse problem [see Sec. (3) below].

2.2.2 Extracting coalescing binary signals

Coalescing binaries are good examples of the type of signal that will probably only be seen by applying pattern-matching techniques: the raw amplitude from even the nearest likely source will be below the level of broad-band noise in the detector. Nevertheless, the signal is so predictable that interferometers should be able to see such systems ten times or more as distant as collapse sources. We will see that the signal depends on two parameters, so when we discuss the coincidence problem from the point of view of pattern-matching, we will have to consider the added uncertainty caused by this.

2.2.2.1 The coalescing binary wave form. The amplitude of the radiation from a coalescing binary depends on the masses of the stars and the frequency f of the radiation, which together determine how far apart the stars are. It is usual to assume that the stars are in circular orbits. This is a safe assumption if the binary system has existed in its present form long enough for its orbit to have shrunk substantially, since the timescale for the loss of eccentricity, $e/\dot{e}$, is 2/3 of the similar timescale for the decrease of the semimajor axis a. If the binary has only recently been formed, *e.g.* by tidal capture in a dense star cluster, then more general wave forms can be expected. The model assumes point particles in a Newtonian orbit, with energy dissipation due to quadrupolar gravitational radiation reaction. I will discuss corrections to this briefly below. The radiation amplitude when the radiation frequency is f is given by the function:

$$A_h(f) = 2.6 \times 10^{-23} (\frac{\mathcal{M}}{M_\odot})^{5/3} (\frac{f}{100\mathrm{Hz}})^{2/3} (\frac{100\mathrm{Mpc}}{r}), \quad (4)$$

where $\mathcal{M}$ is what I shall call the *mass parameter* of the binary system, defined for a system consisting of stars of masses m_1 and m_2 by the equation

$$\mathcal{M} = m_1^{3/5} m_2^{3/5} / (m_1 + m_2)^{1/5}, \quad (5)$$

or equivalently by the more transparent formula,

$$\mathcal{M}^{5/3} = \mu M_T^{2/3}, \quad (6)$$

where μ is the usual reduced mass and M_T the total mass of the system. A system consisting of two $1.4M_\odot$ stars has $\mathcal{M} = 1.22M_\odot$. The numerical value of h_f is actually the *maximum* observable value of the amplitude h, which obtains when the system is viewed down the axis of its angular momentum. One must insert angular factors in front of the expression to get the wave amplitude in other directions. If one averages over these angular factors *and* over the angular factors that describe the antenna pattern of an interferometer, one obtains an effective *mean amplitude* only 2/5 of the maximum (Thorne 1987, Krolak 1989).

The binary's orbital period changes as gravitational waves extract energy from the system. The frequency of the radiation is twice the orbital frequency, and its rate of change is

$$\frac{df}{dt} = 13(\frac{\mathcal{M}}{M_\odot})^{5/3}(\frac{f}{100\text{Hz}})^{11/3}\text{Hz/s}. \tag{7}$$

The maximum wave amplitude we expect, therefore, has the time-dependence

$$h_{max}(t) = A_h[f(t)]\cos(2\pi\int_{t_a}^{t} f(t')dt' + \Phi), \tag{8}$$

where t_a is an arbitrarily defined "arrival time" (conveniently taken to be, say, the time when the signal reaches an arbitrarily chosen frequency f_a), and Φ is the signal's phase at time t_a. This phase depends on where in their orbits the stars are when the frequency reaches f_a. The amplitude increases slowly with the frequency-dependence of A_h. Doing the frequency integral explicitly gives

$$f(t) = 100\text{Hz} \times [(\frac{f_a}{100\text{Hz}})^{-8/3} - 0.34\mathcal{M}^{5/3}(t - t_a)]^{-3/8}, \tag{9}$$

where t is measured in seconds. The phase integral is

$$2\pi\int_{t_a}^{t} f(t')dt' = 3000\mathcal{M}^{-5/3}\{(\frac{f_a}{100\text{Hz}})^{-5/3} - [(\frac{f_a}{100\text{Hz}})^{-8/3} - 0.34\mathcal{M}^{5/3}(t - t_a)]^{5/8}\}. \tag{10}$$

Putting this into Eq. (8) for $h_{max}(t)$ gives the desired formula, which we will use in the next section. Notice that coalescence in the two-point-particle model occurs when $f = \infty$. For a system whose radiation is at frequency f, the remaining lifetime until this occurs is

$$T_{coalescence}(f) = 3.0(\frac{M_\odot}{\mathcal{M}})^{5/3}(\frac{100\text{Hz}}{f})^{8/3}\text{s}. \tag{11}$$

This is 3/8 of the formal timescale $f/\dot{f}$ deducible from Eq. (7). Of course, for realistic stars the Newtonian point-particle approximation breaks down before this time, but if the stars are neutron stars or solar-mass black holes, corrections need be made only in the last second or less. Corrections due to post-Newtonian effects are the first to become important in this case, followed by tidal and mass-transfer effects. These have been considered in detail by Krolak & Schutz (1987) and Krolak (1989). If at least one of the stars is a white dwarf, tidal corrections will become important when $T_{coalescence}$ is still 1000 years or so, and f is tens of milliHertz; the system would only be observable from space (Evans, *et al*, 1987).

We shall need below not only the waveform $h(f)$, but also its Fourier transform. We shall denote the Fourier transform of any function $g(t)$ by $\tilde{g}(f)$, given by

$$\tilde{g}(f) = \int_{-\infty}^{\infty} g(t)e^{-2\pi ift}dt. \tag{12}$$

Provided that the frequency of the coalescing binary signal is changing relatively slowly (*i.e.*, that $T_{coalescence} \gg 1/f$), the method of stationary phase can used to approximate the transform of $h_{max}(t)$, $\tilde{h}_{max}f$ (Thorne 1987; Dhurandhar, *et al*, 1989). We shall only need its magnitude,

$$|\tilde{h}_{max}f| = 3.7 \times 10^{-24}(\frac{\mathcal{M}}{M_\odot})^{5/6}(\frac{100\text{Hz}}{f})^{7/6}(\frac{100\text{Mpc}}{r})\text{Hz}^{-1}. \tag{13}$$

This gives good agreement with results of numerical integrations performed by Schutz (1986). We shall use it in the next sections.

2.2.2.2 The mathematics of matched filtering: finding the signal. Matched filtering is a linear pattern-matching technique designed to extract signals from noise. For references on the theory outlined in this and subsequent sections, the reader may consult a number of books on signal analysis, such as Srinath & Rajasekaran (1979). To use matched filtering we have first to define some properties of the noise, $n(t)$. We expect that $n(t)$ will be a random variable, and we use angle brackets $< >$ to denote expectation values of functions of this noise. It is usually more convenient to deal with the noise as a function of frequency, as described by its Fourier transform $\tilde{n}(f)$. We shall assume that the noise has zero mean,

$$< n(t) > = < \tilde{n}(f) > = 0.$$

We shall also assume that the noise is *stationary*, *i.e.* that its statistical properties are independent of time. Then the *spectral density* of (amplitude) noise $S(f)$ is defined by the equation

$$< \tilde{n}(f)\tilde{n}^*(f') > = S(f)\delta(f - f'), \tag{14}$$

where a * denotes complex conjugation. This says two things: (*i*) the noise at different frequencies is uncorrelated; and (*ii*) the autocorrelation of the noise at a single frequency has variance $S(f)$, apart from the normalization provided by the delta function, which arises essentially because our formalism assumes that the noise stream is infinite in duration. (Texts on signal processing often define $S(f)$ in terms of a normalized Fourier transform of the autocorrelation function of a discretely sampled time-series of noise $n_j(t)$. The continuous limit of this definition is equivalent to ours.) Since $n(t)$ is real, $S(f)$ is real and an even function of f.

White noise has a constant spectrum, which means that $S(f)$ is independent of f. Interferometers have many sources of noise, as described in the article by Winkler in this volume. In this treatment I will consider only two: shot noise, which limits the sensitivity of a detector at most frequencies; and seismic noise, which I idealize as a "barrier" that makes a lower cutoff on the sensitivity of the detector at a frequency f_s. The shot noise is intrinsically white (that is, as a noise on the photodetector), but — depending on the configuration of the detector — the detector's sensitivity to gravitational waves depends on frequency, so the relevant noise is the photon white noise divided by the frequency response of the detector (called its *transfer function*). We call this function $S_h(f)$. I will assume that the detector is in the standard recycling configuration, so that (allowing for the seismic cutoff) we have

$$\begin{aligned} S_h(f) &= \tfrac{1}{2}\sigma_f^2(f_k)[1 + (f/f_k)^2] \quad \text{for} f > f_s, \\ &= \infty \quad \text{for} f < f_s. \end{aligned} \tag{15}$$

Here f_k is the so-called "knee" frequency, which may be chosen by the experimenter when recycling is implemented, and $\sigma_f(f_k)$ is the standard deviation of the frequency-domain noise at f_k. In the usual discussions of source strength *vs* detector noise [*e.g.* Thorne(1987), what is taken to be the detector noise as a function of frequency f is $\sigma_f(f)$, not $[S_h(f)]^{1/2}$, because it is assumed in those discussions that the knee frequency f_k will be optimized by the experimenter for the particular range of frequencies being studied, so that σ_f is representative of the noise that the experimenter would encounter. Later in this section we will see that the optimum value of f_k for observing coalescing binaries is $1.44f_s$.

Now, the fundamental theorem we need in order to extract the signal from the noise is the matched filtering theorem. If we have a signal $h(t)$ buried in noise $n(t)$, so that the output of our detector is

$$o(t) = h(t) + n(t),$$

and if the Fourier transform of the signal is $\tilde{h}(f)$, then the best linear way to determine whether the signal is present is to correlate the output with a *filter* $q(t)$, chosen in the optimum manner given below. The correlation is

$$c(t) = (o \circ q)(t)$$
$$= \int_{-\infty}^{\infty} o(t')q(t'+t)dt' \quad (16)$$
$$= \int_{-\infty}^{\infty} \tilde{o}(f)\tilde{q}^{*}(f)e^{2\pi i f t}df \quad (17)$$

The expectation value of $c(t)$ is the signal,

$$< c(t) >= (h \circ q)(t), \quad (18)$$

and the noise is the square root of the variance of the correlation,

$$< [c(t) - < c(t) >]^2 >= \int_{-\infty}^{\infty} S(f)\, |\tilde{q}(f)|^2\, df. \quad (19)$$

This gives a "raw" signal-to-noise ratio of

$$\frac{S}{N}(t) = \frac{(h \circ q)(t)}{[\int_{-\infty}^{\infty} S(f)\, |\tilde{q}(f)|^2\, df]^{1/2}}. \quad (20)$$

The idea of matching the filter to the signal comes from finding the filter $q(t)$ that maximizes this signal-to-noise ratio. It is not difficult to show that the optimal choice of filter for detecting the signal $h(t)$ is

$$\tilde{q}(f) = k\tilde{h}(f)/S_h(f), \quad (21)$$

where k is any constant. With this filter, if the output contains a signal, then $c(t)$ will reach a maximum at a time t that corresponds to the time in the output stream at which the signal reaches the point $t' = 0$ in the waveform $h(t')$. Of course, noise will distort the form of $c(t)$, but the expected amplitude signal-to-noise ratio S/N in $c(t)$ (ratio of maximum value to the standard deviation of the noise) is given by the key equation

$$\left(\frac{S}{N}\right)^2 = 2\int_0^{\infty} \frac{|\tilde{h}(f)|^2}{S_h(f)}df. \quad (22)$$

It can be shown that this is the largest S/N achievable with a linear filter. Moreover, given a waveform $h(t)$ that one wants to look for, and given a seismic cutoff frequency f_s, one can ask what value of the knee frequency f_k one should take in $S_h(f)$ in Eq. (22) to maximize S/N. For coalescing binaries, one can use the explicit expression for $\tilde{h}(f)$ given in Eq. (13) to show that this value, as mentioned earlier, is (Thorne 1987, Krolak 1989)

$$f_{k\,opt} = 1.44 f_s.$$

Naturally, in a real experiment one does not know if a signal is present or not. One then uses the size of S/N to decide on the likelihood of the correlation being the result of noise. A widely used criterion is the Neyman-Pearson test of significance (Davis 1989), based on the *likelihood ratio*, defined as the ratio of the probability that the signal is present to the probability that the signal is absent (false alarm). If the noise is Gaussian, then the Neyman-Pearson "best" criterion is just to calculate the chance of a false

alarm in the matched filter given by Eq. (21), exactly as described in Sec. (2.2.1.1) with x/σ replaced by S/N.

If the noise is not Gaussian (and we do of course expect that there will be a non-Gaussian component to the noise) then matched filtering does not necessarily give the best discrimination against false alarms. If experiments show that the noise is predominantly non-Gaussian, then we shall have to look at other filtering methods.

Searches for coalescing binaries can therefore be carried out by applying threshold criteria to the correlations produced by filtering. The false-alarm probabilities for detecting a coalescing binary have to be calculated with some care, however, because we must allow for the fact that we have in general to apply many independent filters, for different values of the mass parameter $\mathcal{M}$, and this increases the chance of a false alarm. I will consider the necessary corrections in Sec. (2.2.2.4) below.

In practice, one only samples the data stream at a finite rate, not continuously. It is clear from Eq. (22) that one must sample at least as fast as is required to determine $\tilde{h}(f)$ at all frequencies that contribute to the integral significantly: at least twice as fast as the largest required frequency in $\tilde{h}(f)$. For the coalescing binary, whose transform is given approximately by Eq. (13), the power spectrum $|\tilde{h}(f)|^2$ is falling off as $f^{-7/3}$, and the recycling shot noise multiplies a further factor of f^{-2} into this. Thus, when f rises to, say, 4 times f_s, the integrand in Eq. (22) will have fallen off to about 0.005 of its value at f_s. Truncating the integration here should be enough to guarantee that the filter comes within 1% of the optimum signal-to-noise ratio. This would require a sampling rate of $8f_s$, or 800 Hz if we take $f_s = 100$ Hz.

2.2.2.3 Determining the time-of-arrival of the signal. It is important for gravitational wave experiments that, by filtering the data stream, one not only determines the presence of a signal, but one also fixes its "time of arrival", defined as the time t_{arr} at which the signal reaches the $t' = 0$ point in the filter $h(t')$. The standard deviation in the measurement of t_{arr} is δt_{arr}, which is given by an equation similar to Eq. (22) (Srinath & Rajasekaran 1979; Dhurandhar, *et al*, 1989):

$$\frac{1}{\delta t_{arr}^2} = 8\pi^2 \int_0^{\infty} \frac{f^2\, |\tilde{h}(f)|^2}{S_h(f)}df = 2\int_0^{\infty} \frac{|\tilde{\dot{h}}(f)|^2}{S_h(f)}df, \quad (23)$$

where $\tilde{\dot{h}}(f)$ is the Fourier transform of the time derivative of $h(t)$. If either the signal or the detector's sensitivity is narrow-band about a frequency f_0, then a reasonable approximation to Eq. (23) is

$$\delta t_{arr} = \frac{1}{2\pi f_0}\frac{1}{S/N}, \quad (24)$$

where S/N is the optimum signal-to-noise ratio as computed from Eq. (22). This is a good approximation as long as S/N is reasonably large compared to 1. If we use Eq. (13) for $\tilde{h}(f)$ then it is not hard to show that, for coalescing binaries (Dhurandhar, *et al*, 1989)

$$\delta t_{arr} = 0.84(\frac{100\text{Hz}}{f_s})\frac{1}{S/N}\ \text{ms}. \quad (25)$$

For example, if the signal-to-noise ratio is 7 (the smallest for detection by a single detector) and the seismic limit is 100 Hz, then the timing accuracy would be 0.1 ms. If the signal-to-noise is as high as 30, which could occur a few times per year (see below), then the signal could be timed to 30 μs. Considering that the time it takes the wave to travel from one detector to another will typically be 15–20 ms, this timing accuracy would translate into good directional information. I will explain below how this can be done. However, in practice it will turn out that these numbers are too optimistic, perhaps by a factor of two. The reason is that the presence of noise will also make the determination of the mass parameter uncertain to some degree, and this turns out to affect the timing accuracy: Schutz (1986) has shown that a small change in the mass parameter can masquerade as a displacement in the time-of-arrival of the signal. It may also be that errors in the determination of the waveform phase Φ will do the same. This effect will have to be quantified before realistic estimates of the timing accuracy can be made.

It may seem paradoxical that, if detector physicists succeed in lowering the seismic barrier to, say, 50 Hz, the arrival-time-resolution given by Eq. (24) appears to get worse as f_s^{-1}! This is not a real worsening, of course: the increase in S/N due to the lower seismic cutoff (gaining as $f_s^{-7/6}$ if f_k remains optimized to f_s) more than compensates the $1/f_s$ factor, and the timing accuracy improves. Our remarks earlier about the importance of sampling at the correct rate apply here as well: if the sampling rate is smaller than twice the largest frequency at which the integrand in Eq. (23) contributes significantly, then in the numerical calculation the arrival time accuracy will be worse than optimum. This is an important lesson: *in choosing one's sampling speed one should ensure that one can get good accuracy in Eq. (23), whose integrand falls off less rapidly with frequency than that of Eq. (22).* If one does sample at an adequate rate, then it is possible to determine the time of arrival of a signal to much greater precision than the sampling time, provided the signal-to-noise ratio is much greater than unity. (See, for example, the numerical experiments reported by Gursel & Tinto 1988.) For a coalescing binary, taking timing accuracy into account does not significantly increase the sampling rate over that required for a good signal-to-noise ratio.

2.2.2.4 Threshold criteria for filtered signals. When searching a data stream for coalescing binary signals, we cannot presume ahead of time that we know what the mass parameter $\mathcal{M}$ will be: not all neutron stars may have mass $1.4M_\odot$, and some binaries may contain black holes of mass 15 or 20 $M_\odot$. We therefore will have to filter the data with a family of filters with $\mathcal{M}$ running through the range, say, $0.5 - 25M_\odot$.

How many filters should there be? This question has not yet received enough study. The calculations of Dhurandhar, *et al* (1989) show that two filters with mass parameters differing by a few percent have significantly reduced correlation, so the filters in the family should not be more widely spaced than this. However, it is not known whether they should be more closely spaced, to avoid missing weak signals. If we take successive filters to have mass parameters that increase by 2% at each step, then we need about 200 filters to span the range $(0.5, 25)$ in $\mathcal{M}$.

However, there is also another parameter in the filter, Eq. (8): the phase Φ, about which I have so far said little. When the wave arrives at the detector with frequency f_s, so that it is just becoming detectable, its phase may be anything: this depends on the binary's history. Filters with different phases must therefore be used. This question has not received enough study either, but numerical experiments by S. P. Lawrence (private communication) indicate that one might need about four differently phased filters for each mass parameter in order not to lose signal-to-noise in the search. This *increases* the number of filters to about 1000. In Sec. (2.2.2.6) we will look at the computing demands that this filtering makes on the data analysis system. In the present section we shall consider the signal-to-noise implications.

First it will be necessary to establish what the filtering equivalent of the sampling rate is, so that we can calculate the probability of, say, one false alarm per year. In our original calculation of the false-alarm probability, the sampling rate told us how many independent data points there were per year, on the assumption of white noise, which meant that each data point was statistically independent, no matter how rapidly samples were taken. In the present case, the output of the filter is the correlation given in Eq. (16). It has noise in it, but the noise is no longer white, having been filtered. The key number that we want here is the "decorrelation time", defined as the time interval τ_s between successive applications of the filter that will ensure that the outputs of the two filters are statistically independent. The analog here of the sampling rate in the burst problem is $1/\tau_s$, which I will call the effective sampling rate. This is the rate at which successive independent data points arrive.

To develop a criterion for statistical independence, we consider the autocorrelation function of the filter output when the detector output $o(t)$ is pure noise $n(t)$:

$$a(\tau) = \int_{-\infty}^{\infty} c(t)c(t+\tau)dt. \tag{26}$$

We shall take the decorrelation time to be the time τ_s such that $a(\tau)$ is small for all $\tau > \tau_s$. We can learn what this is by noting that it is not hard to show that the Fourier transform of $a(\tau)$ is, when the optimal filter given in Eq. (21) is used,

$$\tilde{a}(f) = \frac{|\tilde{h}(f)|^2}{S_h(f)}. \tag{27}$$

For coalescing binaries, we have already discussed some of the properties of this function at the end of Sec. (2.2.2.2). It is strongly peaked near f_s, and in particular the seismic barrier cuts it off rapidly below f_s. It follows that for times $\tau \gg 1/f_s$ the autocorrelation function is nearly zero: the effective sampling rate is about f_s. To play it safe, I will work with a rate twice this large, or an effective sampling time of 0.005 s. This gives effectively 6×10^9 samples — statistically independent filter outputs — per year.

Now, assuming that the noise is Gaussian, the calculation of the false-alarm probability for any size network looks similar to our earlier one in Sec. (2.2.1.2). What we have to allow for is that there will be some 1000 independent filters, each of which could give a false alarm. Of course, the false alarm occurs only if each detector registers an event in the *same* filter, so it is like doing 1000 independent experiments with no filter at all and a sampling time of 0.005 s, or one experiment with no filter and a sampling time of 5×10^{-6} s. This increases the number of points by a factor of 200 over the number we used in Sec. (2.2.1.1), but this factor makes only a modest difference in the level of the thresholds. For example, for one false alarm per year, and no correction for time-delay windows, the thresholds are: for one detector, 7.4; for two, 5.1; for three, 4.0; and for four, 3.4. For example, the three-detector threshold is 12% higher than for unfiltered data taken at 1 kHz.

These figures should not be taken as graven in stone: they illustrate the consequences of a particular set of assumptions. A better calculation of the noise properties of the filters is needed, and in any case one will have to ensure that the detector noise really obeys the statistics we have assumed.

2.2.2.5 Two ways of looking at the improvement matched filtering brings. The discussion of matched filtering so far has been fairly technical, with the emphasis on making reliable and precise estimates of the achievable signal-to-noise ratios and timing accuracy. In this section I will change the emphasis and try to develop approximate but instructive ways of looking at the business of matched filtering. The emphasis will be on understanding how matched filtering improves the sensitivity of an interferometer beyond its sensitivity to wide-band bursts. We will look at two points of view: comparing the sensitivity of the detector to broad-band and narrow-band signals that have either the same amplitude or the same total energy.

First let us consider two signals of the same amplitude h, one of which is a broad-band burst of radiation centered at f_0 and the other of which is a relatively narrow-band signal with n cycles at roughly the frequency f_1. The signals are observed with different recycling detectors optimized at their respective frequencies, f_0 and f_1, possibly contained in the same vacuum system, as is envisioned in some present designs. The broad-band signal has

$$\begin{aligned}\left(\frac{S}{N}\right)^2 &= 2\int_0^\infty \frac{|\tilde{h}(f)|^2}{S_h(f)}df. \\ &\approx \frac{2}{\sigma_f^2(f_0)}\int_0^\infty |\tilde{h}(f)|^2\,df \\ &\approx \frac{1}{\sigma_f^2(f_0)}\int_{-\infty}^\infty |h(t)|^2\,dt.\end{aligned} \tag{28}$$

Now, the integrand in Eq. (28) for a burst lasts typically only for a time $1/f_0$, so we have

$$S/N \approx \frac{h}{\sigma_f(f_0)f_0^{1/2}}. \tag{29}$$

For the narrow-band signal, we obtain again Eq. (28), but with f_0 replaced by f_1. Now, however, the signal lasts n cycles, a time n/f_1. This leads immediately to

$$S/N \approx \frac{hn^{1/2}}{\sigma_f(f_1)f_1^{1/2}}. \tag{30}$$

Comparing Eqs. (29) and (30), we see that a narrow-band signal has an advantage of $\sqrt{n}$ over a burst of the same amplitude and frequency, provided we have enough understanding of the signal to use matched filtering. For the coalescing binary one may approximate n by $f^2/\dot{f}$, and this can be large (of order 200). Coalescing binaries gain further when compared to supernova bursts because of their lower frequency: because σ_f depends on f as $f^{1/2}$, there is a further gain of a factor of f_0/f_1, which can be 7 or so. Therefore, a coalescing binary signal might have something like 100 times the S/N of a supernova burst *of the same amplitude!* This exaggerates somewhat the advantage that coalescing binaries have as a potential source of gravitational waves, since their intrinsic amplitudes may be smaller than those from supernovae, but it does show why they are such interesting sources.

The other way of looking at filtering is in energy terms. This is very instructive, because it shows "why" matched filtering works. We have just seen that a narrow-band signal with n cycles has a higher S/N than a broad-band burst of 1 cycle that has the same amplitude and frequency, by a factor of $\sqrt{n}$. But the *energy* in the narrow-band signal is n times that in the burst. This is because the energy flux in a gravitational wave is

$$\mathcal{F}_{gw} \approx \frac{4c^3}{\pi G}h^2 f^2, \tag{31}$$

and thus the total energy E in a signal passing through a detector during the time n/f that the burst lasts is given by the proportionality

$$E \propto h^2 f^2 (n/f) = nfh^2.$$

If we solve this expression for nh^2 and put it into Eq. (30), we find

$$\frac{S}{N} \propto \frac{E^{1/2}}{f\sigma_f(f)}. \tag{32}$$

Since this is independent of n, it applies to broad-band and narrow-band signals equally. It shows that if two signals send the same total energy through an interferometric detector, and if they have the same frequency, then they will have the same signal-to-noise ratio, again provided we have enough information to do the matched filtering where necessary. This provides a somewhat more realistic comparison of coalescing binaries and supernovae, since a coalescing binary radiates a substantial amount of energy in gravitational waves, of the order of 0.01 $M_\odot$. This is similar to the energy radiated by a moderate to strong gravitational collapse. The advantage that coalescing binaries have is that they emit their energy at a lower frequency. The factor of $f\sigma_f \propto f^{3/2}$ in Eq. (32) gives them an advantage of a factor of roughly 20 over a collapse generating the same energy at the same distance. If laser interferometric detectors achieve a broad-band sensitivity of 10^{-22}, as current designs suggest will be possible (see the article by Winkler in this volume), then they will be able to see moderate supernovae as far away as 50 Mpc. This volume includes several starburst galaxies, where the supernova rate may be much higher than average. They will therefore also be able to see coalescing binaries as far as 1 Gpc.

2.2.2.6 The technology of real-time filtering. *In this section I will discuss the technical feasibility* of performing matched filtering on a data stream in "real time", *i.e.* keeping up with the data as it comes out of a detector. Since coalescing binaries seem to make the most stringent demands, I will take them as fixing the requirements of the computing system. We have seen that we need a data stream sampled at a rate of about 1 kHz in order to obtain the best S/N and timing information, so I will use this data rate to discover the minimum requirements. It is likely that the actual sampling rates used in the experiments will be much higher, but they can easily be filtered down to 1 kHz before being analyzed. If the seismic cutoff is 100 Hz, then the duration of the signal, at least until tidal or post-Newtonian effects become

important, will be less than 2 seconds in almost all cases. This means that a filter need have no more than 2000 2-byte data points.

The quickest way of doing the correlations necessary for filtering is to use fast Fourier transforms (FFT's) to transform the filter and signal, multiply the signal transform by the complex conjugate of the filter transform, and invert the product to find the correlation. The correlation can then be tested for places where it exceeds pre-set thresholds, and the resulting candidate events can be subjected to further analysis later. This further analysis might involve: finding the best value of the mass parameter and phase parameter; filtering with filters matched to the post-Newtonian waveform to find other parameters that could determine the individual masses of the stars; looking for unmodelled effects, such as tides or mass transfer; looking for the final burst of gravitational radiation as the two stars coalesce; and of course processing lists of these events for comparison with the outputs of other detectors. Since the number of significant events is likely to be relatively small, the most demanding aspect of this scenario is likely to be the initial correlation with 1000 coalescing binary filters.

One way the processing might be done is as follows. The discrete correlation between a data set containing the N values $\{d_j,\ j=0,\ldots,N-1\}$ and a filter containing the N values $\{h_k,\ k=0,\ldots,N-1\}$ is usually given by the *circular* correlation formula:

$$c_k = (d \circ h)_k = \sum_{j=0}^{N-1} d_j h_{j+k},\ k=0,\ldots,N-1, \tag{33}$$

where we extend the filter by making it periodic:

$$h_{j+N} = h_j\ \forall j.$$

The circular correlation formula has a danger, because the data set and filter are not really periodic. In practice, this means that we should make the data set much longer than the (non-zero part of the) filter, so that only when the filter is "split" between the beginning and the end of the data set does the circular correlation give the wrong answer. Thus, even if each filter requires only $N_h \leq 2000$ points, it is more efficient to split the data set up into segments of length $N \gg N_h$ points, and to use a filter which has formally the same length, but the first $N - N_h$ of whose elements are zero. (I am grateful to Harry Ward for bringing this point to my attention.) The "padding" by zeros ensures that the periodicity of h corrupts only the last N_h elements of the correlation. This can be rectified by forgetting these elements and beginning the next data segment N_h elements before the end of the previous one: this overlap ensures that the first N_h elements of the next correlation replace the corrupt elements of the previous one with correct values. Since this procedure involves filtering some parts of the data set twice, it is desirable to make it a small fraction of the set, namely to make N_h small compared to N. This efficiency consideration is, however, balanced by the extra numerical work required to calculate long correlations, proportional to $N \ln N$. This arises as follows.

The fastest way to do long correlations on a general-purpose computer is to use Fourier transforms (or related Hartley transforms). For a discrete data set $\{d_j,\ j=1,\ldots,N-1\}$, the discrete (circular) Fourier transform (DFT) is the set $\{\tilde{d}_k,\ k=1,\ldots,N-1\}$ given by

$$\tilde{d}_k = \sum_{j=0}^{N-1} d_j e^{-2\pi ijk/N}, \tag{34}$$

with the inverse transform

$$d_j = \frac{1}{N}\sum_{k=0}^{N-1} \tilde{d}_k e^{2\pi ijk/N}. \tag{35}$$

Then the discrete version of the convolution theorem Eq. (17) is as follows. Given the (circular) correlation $\{c_j\}$ of two sets $\{d_j\}$ and $\{h_j\}$ as in Eq. (33), its DFT is

$$\tilde{c}_k = (\tilde{d}_k)^* \, \tilde{h}_k, \tag{36}$$

where an asterisk denotes complex conjugation.

Fast Fourier transform (FFT) algorithms may require typically $3N \log_2 N$ real floating-point operations (additions and multiplications) to compute the transform of a set of N real elements, provided N is an integer power of 2 (which can usually be arranged). (I neglect the overhead of integer arithmetic concerned with the index manipulations in such routines.) To compute the correlation of two such sets, then, would require three transforms — two to produce $\tilde{d}_k$ and $\tilde{h}_k$ and a third to invert the product $\tilde{c}_k$ — and the multiplication of the two original transforms, giving a total of $9N\log_2 N + 4N$ real floating-point operations. This is to be compared with the $2N^2 - N$ operations required to calculate the correlation directly from Eq. (33). As long as $N \geq 16$ it will be quicker to use FFTs.

In practice, one would compute once and store the DFT of all M filters, so that in real time the data would have to be transformed only once, and then M products of data and filter calculated and inverse-transformed. This would require $3N(M+1)\log_2 N + 4NM$ floating point operations.

We must now remind ourselves that in order to achieve the economies of the FFT algorithm, we must use the circular correlation, which has an extra cost associated with the overlaps we are required to take in successive data sets. For a given filter length (say $N_f < N$ non-zero points in the filter time-series), we can reduce the fractional size of these overlaps by making N larger, but this increases the cost of the FFT logarithmically in N. Is there an optimum ratio N_f/N? The total cost of analysing a data set containing a very large number $N_{tot} \gg N$ of elements, split up into segments of length N is

$$N_{fl\ pt\ ops} = \frac{N_{tot}}{N - N_f}[3N(M+1)\log_2 N + 4NM].$$

We want to minimize this with respect to variations in N holding N_f and M (the number of filters) fixed. It is more convenient to introduce the variable $x = N_f/N$, which measures the fractional overlap of successive data sets. In terms of x the expression is:

$$N_{fl\ pt\ ops}(x) = \frac{N_{tot}}{1-x}[3(M+1)\log_2 \frac{N_f}{x} + 4M]. \tag{37}$$

As long as the number of filters M is large, the optimum x will be independent of M: it will depend only on N_f, the "true" length of the filter. If we take N_f to be 2000, then the optimum x is .057; if $N_f = 1000$ then the best x is .061. But the minimum in $N_{fl\ pt\ ops}$ is a flat one, and one can increase the value of x quite a bit without compromising speed. This is important, because each stored filter transform must contain N points, so the larger we make x, the smaller will be our core memory requirements. This is illustrated in the following table, which gives x and $N_{fl\ pt\ ops}/NM$, the number of floating-point operations per data point per filter, as required by various strategies, always taking N_{tot} to be an integer power of 2.

From this it is clear that choosing an overlap between successive data sets of around 25% gives a CPU demand that is only slightly higher than optimum and reduces storage requirements to a minimum. Based on this, and assuming a data rate of 1000 2-byte samples per second with a 2-second filter length ($N_f = 2000$), it follows that doing 1000 filters in real time requires a computer capable of 60 Mflops (where 1 Mflop is 10^6 floating-point operations per second), and storage for 1000 filters, each of length 16k bytes. This is not far from the capabilities of present-day inexpensive (< \$100k) workstations with add-on array-processors, or of stand-alone arrays of transputers or other fast microprocessors. In five years it should be trivial.

There are many possible ways to speed up the calculation if CPU rates are a problem. It might be that special-purpose digital-signal-processing chips would be faster than general-purpose microprocessors for this problem. It might be possible to do the calculation in block-integer format rather than floating-point, with filters that consist of crude steps rather than accurate representations of the waveform (Dewey 1986). These should be analyzed further. Another possible CPU-saver is described in the next section.

N_f	N	x	$N_{fl\ pt\ ops}/N_{tot}M$
1000	2^{11}	.488	72
1000	2^{12}	.244	53
1000	2^{13}	.122	49
1000	2^{14}	.061	49
1000	2^{15}	.031	50
2000	2^{12}	.488	78
2000	2^{13}	.244	57
2000	2^{14}	.122	52
2000	2^{15}	.061	52
2000	2^{16}	.031	54

Table 1: The consequences of various strategies for applying filters of "true" length N_f, padded out with zeros to a length N, to very long data sets. See text, especially Eq. (37), for details.

2.2.2.7 Smith's interpolation method for coalescing binaries. An alternative strategy for coalescing binaries has been proposed and implemented by Smith (1987). This clever idea is based upon the following observation: if two coalescing systems of different mass parameters happen to have the same time of coalescence, then their signals' frequencies will remain strictly proportional to one another right up to the moment of coalescence. This follows from the fact that df/dt is proportional to a power of f, so that, as remarked after Eq. (11), there is a constant α independent of the masses such that $T_{coalescence} = \alpha f/\dot{f}$. If two signals with present frequencies f_1 and f_2 have the same $T_{coalescence}$, then it follows that

$$\frac{df_1}{df_2} = \frac{\dot{f}_1}{\dot{f}_2} = \frac{f_1}{f_2}.$$

Since if their times to coalescence are equal at one time then they are necessarily equal for all later times, this equation can be integrated to give $f_1/f_2 = \text{const}$.

Now suppose that the data stream is sampled at constant increments of the phase of signal 1, *i.e.* it is sampled at a rate that accelerates with the frequency f_1. Then if a Fourier transform is performed on the sampled points, the signal will appear just as a pure sinusoid, allowing it to be identified without sophisticated filtering. Moreover, and this is the key point, every other signal with the same time to coalescence will have been sampled at constant increments of its phase as well, since its frequency has been a constant times the first signal's frequency. So signals from *any binary coalescing at the same time*, no matter what its mass parameter, will be exposed by the single Fourier transform. Thus, one Fourier transform would seem to have done the work of all 1000 filters!

The situation is not quite that good, however, because a signal with a different coalescence time will not be visible in the transform of the points sampled in the manner just described. Therefore, data must be sampled over again at the increasing rate *ending at each possible time of coalescence of the binary*. If this is done, then every possible signal will be picked up.

One way of implementing this method would be to sample the detector output at a constant rate (*e.g.* 1000 Hz) and then interpolate to form the data sets that are given to the FFT routine. If we compare this interpolation method with the filtering described earlier, one trades the work of doing 1000 Fourier transforms on a stretch of data for the work of interpolating many times. The actual comparison depends on the number of operations required by the interpolation algorithm, but in general Smith's method with interpolation becomes more attractive as the number of filters one must use increases.

Another way of implementing Smith's method — and the way she herself used — would be to sample the detector output very fast, say at 10 kHz, and then to extract a data set at a slower rate (perhaps 500 – 1000 Hz) by selecting from the sampled points those points closest in time to the places one ideally would wish to sample. This is a far faster procedure than interpolating, and it seems to me that it would not necessarily be less accurate than a simple interpolation algorithm. I will call this *stroboscopic sampling*; we will meet it again when we discuss searches for pulsars. I do not know of any theoretical analysis of it; in particular, one would like to understand what it does to the noise background. The idea, at least in astronomy, seems to go back to Horowitz (1969), who devised it for optical searches for pulsars.

It may well be that for 1000 filters, Smith's method will be more efficient than filtering. However, it has at least two significant disadvantages over filtering:

1. It is restricted only to looking for the Newtonian coalescing binary signal: even any corrections (such as for post-Newtonian effects) will have to be searched for by filtering the sampled data sets, and the sets are essentially useless in searches for other kinds of signals that we may wish to filter from the data.

2. Signals with the same coalescence time but different mass parameters will enter the observing window (say, $f > 100$ Hz) at different times, and this presents a possible problem that was first pointed out by Harry Ward. If one decides to break the data stream into sets of length, say, 2–3 s, appropriate to coalescing $1.4M_{\odot}$ neutron stars starting at 100 Hz, then the set will be much too long for a signal from a binary system of two $14M_{\odot}$ black holes that will coalesce at the same time. The black hole system will have frequency 24 Hz when the data set begins, and will be buried in the low-frequency detector noise. When the data are transformed, this noise will be included in the transform, and the signal-to-noise ratio will accordingly be reduced. The matched filtering method does not suffer from this drawback, since it filters out the low-frequency noise. It might be possible to avoid this problem by pre-filtering the data stream before it is sampled or interpolated, removing the low-frequency noise (and signal).

Given our present uncertainties about sources, my own prejudice is to use filtering because of its inherent flexibility; but Smith's method may become important if filtering places too great demands on the computing system.

2.2.3 Looking for pulsars and other fixed-frequency sources

2.2.3.1 Why the data-analysis problem is difficult. There are many possible sources of gravitational radiation that essentially radiate at a fixed frequency. Pulsars, unstable accreting neutron stars (the Wagoner mechanism), and the possible long-term spindown of a newly-formed neutron star are examples. In some cases, such as nearby known pulsars, we will know ahead of time the frequency to look for and the position of the source. But most continuous sources may have unknown frequencies; indeed they will only be discovered through their gravitational waves. I will first discuss the detection problem for sources of known frequency, and then consider searches for unknown sources. Throughout this discussion, the word "pulsar" will stand for any continuous source. The most complete discussion of this problem of which I am aware is the Ph.D. thesis of Livas (1987).

If we were on an observing platform that had a fixed velocity relative to the stars, and therefore to any pulsar we might be looking for, then finding the signal would be just a matter of taking the Fourier transform of the data and looking for a peak at the known frequency. This is a special case of matched filtering, since the Fourier integral is the same as the correlation integral in Eq. (18) with the filter equal to a sinusoid with the frequency of the incoming wave. Therefore, the signal-to-noise ratio for an observation that lasts a time T_{obs} would increase as $T_{obs}^{1/2}$, just as in Eq. (30). However, the Earth rotates on its axis and moves about the Sun and Moon, and these motions would Doppler-spread the frequency and reduce its visibility against the noise.

How long do we have to look at a source before it becomes necessary to correct for the Earth's motion? If we consider only the Earth's rotation for the moment, then in a time T_{obs} the detector's velocity relative to the source changes by an amount $\Delta v = \Omega_{\oplus}^2 R_{\oplus} T_{obs}$, where $R_{\oplus}$ is the Earth's radius and $\Omega_{\oplus}$ its angular velocity of rotation. In a source of frequency f, this produces a change $\Delta f_{Doppler} = vf/c$. But the frequency resolution of an observation is $\Delta f_{observable} = 2/T_{obs}$. The Doppler effect begins to be important if $\Delta f_{Doppler} = \Delta f_{observable}$. Solving this for T_{obs} gives T_{max}, the maximum uncorrected observing time:

$$T_{max} = \left(\frac{2c}{\Omega_{\oplus}^2 f R_{\oplus}}\right)^{1/2} \approx 70\left(\frac{f}{1\mathrm{kHz}}\right)^{-1/2}\mathrm{min}. \tag{38}$$

Using the same formula for the effects of the Earth's orbit around the Sun gives a time roughly 2.4 times as long. The Earth's motion about the Earth-Moon barycenter also has a significant effect. Since any

serious observation is likely to last weeks or months, the Doppler effects of all these motions must be removed, even in searches for relatively low-frequency signals (10 Hz).

2.2.3.2 Angular resolution of an observation. The Doppler corrections one has to apply depend on the location of the source in the sky. Since the spin axis of the Earth is not parallel to its orbital angular momentum vector, there is no symmetry in the Doppler problem, and every location on the sky needs its own correction.

It is of interest to ask how close two points on the sky may be in order to have the same correction; this is the same as asking what the angular resolution of an observation might be. Let us first imagine for simplicity that our detector participates in only one rotational motion, with angular velocity Ω and radius R. If two sources are separated on the sky by an angle $\Delta\theta$ (in either azimuth or altitude), then the difference between the Doppler corrections for the two sources depends on the *difference* between the changes in the detector's velocities relative to the two sources. For small $\Delta\theta$ this is $\Delta v = \Delta\theta\Omega_{\oplus}^2 R_{\oplus} T_{obs}$. Using this velocity change, the argument is otherwise identical to that given in the previous section, provided that $T_{obs} < 1/\Omega$. For longer T_{obs}, the frequency shift begins to decrease again, so the maximum shift is given at least roughly by setting T_{obs} equal to $1/\Omega$ in this argument. The result is that

$$\Delta\theta = T_{max}^2 \max(\Omega^2, 1/T_{obs}^2). \tag{39}$$

When looking at a source with a frequency of 1 kHz, then for the Earth's rotation, and an observation lasting longer than a day, this gives

$$\Delta\theta_{rotation} = 0.1\left(\frac{f}{1\mathrm{kHz}}\right)^{-1}\mathrm{rad}, \tag{40}$$

which is about half a degree for a millisecond pulsar. For the Earth's motion about the Earth-Moon barycenter, and for an observation lasting longer than one month, the angular resolution improves to 0.01 rad, or approximately half a degree. For the Earth's orbit around the Sun, we have

$$\Delta\theta_{orbit} = 1\times10^{-6}\left(\frac{f}{1\mathrm{kHz}}\right)^{-1}\left(\frac{T_{obs}}{10^7\mathrm{s}}\right)^{-1}\mathrm{rad}, \tag{41}$$

which is about 0.2 arcsec for a millisecond pulsar. The Earth's orbital motion therefore affords the better angular resolution, but it also makes the most stringent demands on applying the corrections. In particular, uncertainties in the position of the pulsar being searched for, orbital motion of the pulsar in a binary system, proper motion of the pulsar (*e.g.*, a transverse velocity of 300 km/s at 200 pc), or unpredicted changes in the period [anything larger than an accumulated fractional change $\Delta f/f$ of $10^{-10}(f/1\mathrm{kHz})^{-1}$] will all require special techniques to compensate for the way they spread the frequency out over more than the frequency resolution of the observation.

2.2.3.3 The technology of performing long Fourier transforms. We shall see that there are several different strategies one can adopt to search for pulsars, whether known ahead of time or not, but all of them will involve performing Fourier transforms of large data sets. It will help us compare the efficiencies of different strategies if we first look at how this can be done.

If one imagines that the observation lasts 10^7 s with a sampling rate of 1 kHz, then one must perform an FFT with roughly 10^{10} data points. This requires roughly $3N\log_2 N$ operations for $N = 2^{34} = 1.7\times10^{10}$. This evaluates to 1.7×10^{12} operations per FFT. Given the 50 Mflops computer we required earlier for filtering for coalescing binaries, this would take about 10 hours. This is not unreasonable: over 200 FFTs could be computed in the time it took to do the observation.

The real difficulty with this is the memory requirement: FFT algorithms require access to the whole data set at once. To achieve these processing speeds, the whole data set would have to be held in fast memory, all 20 Gbytes of it. Unless there is a revolution in fast memory technology, it does not seem likely that this will be possible, at least not at an affordable level. One could imagine being able to store the data on a couple of 10-Gbyte read/write optical discs, and then using a mass-store-FFT algorithm, which uses clever paging of data in and out of store. This would still be very slow, but its exact speed would depend on the computer system.

An alternative method of calculating the Fourier transform would be to split the data set up into M chunks of length L, each chunk being small enough to fit into core. By performing FFTs on data sets of length L it is possible to calculate the contribution of each subset to the total transform. It is not hard to show that the work needed to construct the full transform from these individual sets is about M times the work needed to do it as a single set. With a memory limit of 200 Mbytes and a machine capable of 50 Mflops, it might be possible to do one or two Fourier transforms in the time it takes to do the observation. With the same memory in a machine capable of 1 Gflop, one could do 40 Fourier transforms in the same time. These are big numbers for memory and performance, but they may be within reach of the interferometer projects by the time they go on-line. The numbers become even more tractable if we are looking for a pulsar under 100 Hz: with a data rate of only 100 Hz, say, the work for a given number M of subsets goes down by a factor of about 11. It is clear that it is possible to trade off memory against CPU speed; the technology of the time will dictate how this trade-off is to be made.

If it proves impossible to compute the full transform exactly, there are approximate methods available, such as to subdivide the full set into M subsets as above, but then only to compute the power spectrum of each subset and to add the power spectra together. This reduces the frequency resolution by a factor of M, with a proportionate decrease in the spatial resolution and in the number of different positions that an observation might need to search. It also reduces the signal-to-noise ratio of the observation. It may be possible to use techniques developed for radio pulsar searches (Lyne 1989).

2.2.3.4 Detecting known pulsars. The earliest example of using a wide-band detector to search for a known pulsar is the experiment of Hough, *et al* (1983), which set an upper limit of $h < 8 \times 10^{-21}$ on radiation from the millisecond pulsar, PSR 1937+214. Future interferometers could better this limit by many orders of magnitude, but they will have to do long observing runs (some 10^7 s) to achieve maximum sensitivity. The analysis of the vast amount of data such experiments will generate poses greater problems for analysis than those we addressed for coalescing binaries.

Let us assume that we know the location and frequency of a pulsar, and we wish to detect its radiation. We need to make a correction for the Doppler effects from the known position, or from several contiguous positions if the position is not known accurately enough ahead of time. As for coalescing binaries, we can do this either by filtering, or by interpolation or stroboscopic sampling.

Let us consider filtering first. In this context a filter is just a sinusoidal signal Doppler-shifted to give the expected arrival time of any phase at our detector. If only one rotational motion of our detectors were present, and if the observation were to last several rotation periods, then only points separated in the polar direction would need separate filters: points separated in azimuth have waveforms that are simply shifted in time relative to one another, and so correlating the data in time with only one filter would take care of all such points.

However, our detectors participate in at least *three* rotational motions about different centers, and the observations will probably last only a fraction of a period of the most demanding motion, the solar orbital one. This means that filters lose one of their principal advantages: searching whole data sets for similar signals arriving at different times.

Filtering requires that at least three FFTs of long data sets must be performed: of the filter, of the sampled data, and of their product to find the correlation. Even for a well-known source, there will have to be several filters, because the phase of the wave as it arrives will not be predictable, nor will its polarization. The phase of the wave depends on exactly where the radiating "lump" on the pulsar is. A given detector will respond to the two independent polarizations differently as it moves in orbit around the Sun; the polarization will generally be elliptical, but the proportion of the two independent polarizations and the orientation of the spin axis are unknown. Each of these variables must be filtered for, and each filter needs two more FFTs (the data set needs to be transformed only once).

If the source's position and/or frequency are not known accurately, then even more filters will be required, each adding two further FFTs. Given the problems we saw we might have with FFTs, this could be a costly procedure.

Let us therefore turn to interpolation, or its even more efficient cousin, stroboscopic sampling [defined in Sec. (2.2.2.7)]. Given a position on the sky, one samples the data at intervals of time that represent the arrival of constant increases in the phase of the expected signal at the detector. A simple Fourier transform of the data set thus produced will find the signal, including its phase. The changing response of the rotating antenna to the polarization of the signal will cause a small modulation of the signal, spreading out the Fourier transform slightly.

Moreover, stroboscopic sampling would be particularly easy and fast in this case, since the sampling rate is always nearly constant in the Earth's rest frame. Consider a case where 10 kHz data are stroboscopically sampled at an average rate of 1 kHz. Normally one takes every tenth point, but occasionally one must skip to the ninth or eleventh point to keep in synchronism with a clock in the inertial frame. The calculation of these "leap points" is the only floating-point operation needed for this case.

Even if the frequency of the pulsar is not known well enough, there is no real problem: the exact frequency will come out of the Fourier transform. If the position is not known well enough, then different samplings must be made for different positions, even those separated only in azimuth. This might require more positions to be analyzed separately than with filters, where some of the azimuthal ambiguity can be resolved by a single filter. Each position requires a data sampling and a Fourier transform; this is to be compared with two Fourier transforms for each filter used.

Sampling has two advantages over filtering. First, stroboscopic sampling is much faster than an FFT of the full data set, and it makes no special demands on fast memory. Second, filters must be matched in phase and polarization to that of the signal, and even for known sources these parameters will not usually be known in advance. Therefore, several filters will be needed for each position and frequency, while the Fourier transform of a sampled data set will determine the phase and polarization directly. Against this is the possibility that filtering will take care of a number of azimuthally-separated positions, while sampling can handle only one at a time.

It may well be that, on balance, sampling will be the method of choice for this problem. But the choice will depend to some degree on the pulsar being looked for: on how accurately its position and frequency are known in advance, and on other factors such as its slowing down rate, its stability, its proper motion, or its possible membership of a binary system.

2.2.3.5 Searching for unknown pulsars. One of the most interesting and important observations that interferometers could make is to discover old nearby pulsars or other continuous wave sources. There may be thousands of spinning neutrons stars — old dead pulsars — for each currently active one. The nearest may be only tens of parsecs away. But we would have to conduct an all-sky, all-frequency search to find them. We shall see in this paragraph that a single detector cannot do this job with foreseeable computer technology. However, when we study cross-correlation methods among several detectors in Sec. (2.4.3) below, we will find that three detectors working together can perform an efficient all-sky search, albeit at lower sensitivity than one could achieve if the methods examined below could be made practical.

The *filtering method* described above for detecting known pulsars works only if the frequency of the radiation is known in advance. If we try to use filters to search all the sky, how many will we need? Let us again assume that the observation will last 10^7 seconds and that the data rate will be 1 kHz. The motion about the Sun is the most important, requiring different filters for positions spaced 10^{-6} radians apart [Eq. (41)]. In the polar direction, this means we need of order 3×10^6 angles. In the azimuthal direction, we might not need so many: as we noted above, if there were no other motions and if the observation lasted several years, then we would need only one filter for an entire 2π azimuthal strip. But there are other motions, so I will approximate their effect by demanding a resolution in the azimuthal direction corresponding to that required in the polar direction by the Earth's motion about the Earth-Moon barycenter, 1×10^{-2} radians [Eq. (40)]. This would give 600 angles azimuthally, leading to a total of 2×10^9 angular positions.

Let us also assume, somewhat optimistically, that we need four filters to cover different phases and four for polarizations. Since each filter is constructed for a given frequency, and we have as many frequency intervals in the Fourier transform as we have data points before the transform (10^{10}), we need 3×10^{20} filters to do an all-sky, all-frequency search with filters! Of course, pulsar searches are likely to take place in interferometers which are being run in a narrow-banding mode, such as resonant recycling. This enhances their sensitivity in a bandwidth that is some 1% or less of the full detector bandwidth. This would reduce the number of filters needed for any observation to something like 10^{18} or so, but this is still huge. And the cost is that pulsars outside the bandwidth of the detector cannot be found.

The alternative method, *stroboscopic sampling*, improves the situation, but not enough to be tractable. For this a different data set has to be extracted for each angular position, even in the azimuthal direction, requiring about 10^{13} data sets. Each one will, when Fourier transformed, search all frequencies, phases, and polarizations for us, so this method is better than filtering by a factor of 10^7 or more, but it still makes impossible demands of the computing machinery. Even a 1 Tflop computer which could hold the whole data set in store would take some 10^6 years to perform all the FFTs!

This number scales as $(T_{obs}f)^3$, so that observations at lower frequency or of shorter duration might be possible. For example, an observation at 100 Hz lasting only 1 day would take a 1 Gflop computer half a year to analyze with stroboscopic sampling. It becomes clear that sensitive all-sky, all-frequency searches using the data of only a single detector will be difficult, and perhaps impossible for long data runs. In Sec. (2.4.3) below we will see that using cross-correlations between different detectors improves the situation dramatically.

I do not believe we have heard the last word on this subject, however. The importance of the problem means that it may be rewarding to look at it more deeply to see if some special features of it can be exploited to make it tractable.

2.3 Combining lists of candidate events from different detectors

Until now I have kept the discussion to the analysis of one detector's data, but it is clear that for the best signal-to-noise ratio and for the extraction of complete astrophysical information, detectors must operate in coincidence. I will consider in this section the simplest method of coordinated observation: exchanging lists of events detected in individual detectors. I have elsewhere (Schutz 1989) called this the "threshold mode" of network data analysis, because each detector's criterion for an "event" is that its amplitude crosses a pre-set threshold.

We have seen in Sec. (2.2.1.2) how the thresholds can be determined. Once events have been identified by the on-line computer — either in the time-series of data directly or by filtering — it is important that the data from these events be brought together and analyzed as quickly as possible. If the event is a supernova, we have considerably less than a day before it might become bright enough to be seen optically, and optical astronomers need to be told of it as quickly as possible. If the event is a coalescing binary, there may be even more urgency: the absence of an envelope around a neutron star means that any radiation emitted may come out with much less delay than in a supernova. Since we know so little about what such events look like, it would be valuable to have optical telescopes and orbiting X-ray telescopes observe the region of the event as quickly as possible.

The rapid exchange of data is certainly possible: with modern electronic mail systems, it would be easy to arrange that the on-line computers could automatically circulate lists of events and associated data periodically, such as every hour. We should bear in mind that, if the threshold is set so that a network would have a four-way false alarm only once per year at a data rate of 1 kHz, then each detector will see a spurious noise-generated event three times per second! It will be impossible to distinguish the real events from the spurious until the lists of events from the various detectors are compared. The initial lists need not contain much data, so links over the usual telephone networks will be fast enough at this stage.

What sort of data must be exchanged? If the event is seen in a filter, the list should include the amplitude of the event, the parameters of the best-fit filter, and an agreed measure of the time the signal arrived at the detector (such as when a coalescing binary signal reached some fiducial frequency, *e.g.* 100 Hz). It will probably also be necessary to include calibration data, as the sensitivity of interferometers will probably change from time to time. If the signal has a high signal-to-noise ratio, then it may be desirable to include other information, such as its correlation amplitude with other filters, or even the raw unfiltered data containing the signal. The feasibility of this will depend upon the bandwidth of available communication channels.

If the event is a broadband burst seen in the time-series, then it will be even more important to exchange the raw data, along with timing and calibration information. If raw-data exchange is impossible, then at least some description of the event will be needed, such as when it first crossed the threshold, when it reached its maximum, and when it went below threshold.

Once likely coincidences among detectors have been identified, it will then be useful to request the on-line computers to send out more detailed information about the selected candidate coincidences. Since these requests will be rarer, it will not overburden the communications networks to exchange raw data and more complete calibration information for the times in question. If the events then still seem significant, they should be broadcast to other astronomers and analyzed more thoroughly at leisure.

2.4 Using cross-correlation to discover unpredicted sources

The threshold mode of analysis is unsuitable for some sources, such as continuous waves or weak events that we have not predicted well enough ahead of time to construct filters for. In these cases, the "correlation mode" is appropriate: using cross-correlations between the data streams of different detectors.

Cross-correlation has its own problems, however: its signal-to-noise relations are rather different from filtering, and the different polarizations of different detectors mean that signals in two different detectors from the same gravitational wave may not exactly correlate. In the next section I will give a general discussion of cross-correlation, addressing the behavior of noise and assuming that the two data streams contain the same signal. The problem of polarization has been solved by Gursel & Tinto (1989); I will discuss their approach in the subsequent section.

2.4.1 The mathematics of cross-correlation: enhancing unexpected or unknown signals

It is useful to think of cross-correlation as the use of one data stream as a filter to find things in the other data stream. Thus, if the first stream contains a signal that hasn't been predicted, one can still find it in the second. If we adopt this point of view, then we must face two important differences between matched filtering and cross-correlation as a means of enhancing signal-to-noise ratios. These are:

1. The "filter" is *noisy*. In fact, in the case of most interest, the signal is below the broad-band noise and the power in the filter is dominated by the noise. If we really had an instrument with an infinite bandwidth, then the noise power would be infinite and we would *never see the signal*. *In practice*, we will see below that we must filter the data down to a finite bandwidth before performing the correlation in order to achieve an acceptable signal-to-noise ratio.

2. The "filter" also contains the signal we wish to find, of course, but the amplitude of this part of the filter is not known *a priori*: it is the amplitude of the incoming signal. This means that if the incoming signal is reduced by half, the response of the filter to it will go down by a factor of *four*. We shall see that this leads to the biggest difference between matched filtering and cross-correlation when they are applied to long wavetrains: the enhancement of signal-to-noise in cross-correlation increases only as the fourth root of the observing time or the number of cycles in the signal, not as the square root we found in Eq. (30).

If we have two data streams o_1 and o_2 containing the same signal h but independent noise amplitudes n_1 and n_2,

$$o_1(t) = h(t) + n_1(t), o_2(t) = h(t) + n_2(t), \tag{42}$$

their cross-correlation is

$$o_1 \circ o_2 = h \circ h + n_1 \circ h + h \circ n_2 + n_1 \circ n_2. \tag{43}$$

The "signal" is the expectation of this, which is just $h \circ h$. The variance of the correlation, however, is a problem. The final term contributes

$$\begin{aligned} <|n_1 \circ n_2|^2> &= < \int \tilde{n}_1(f)\tilde{n}_1^*(f')\tilde{n}_2^*(f)\tilde{n}_2(f')e^{2\pi i(f-f')t}df df' > \\ &= \int S_1(f)S_2(f)\delta(f-f')\delta(f-f')e^{2\pi i(f-f')t}df df'. \end{aligned}$$

The presence of *two* delta functions in the integrand makes this expression infinite: if we allow all the noise in the detectors to be cross-correlated, then the variance of the correlation will swamp the signal. The solution is (i) to *filter* the output down to a suitable bandwidth B before correlating, and (ii) to perform the correlation only over a finite stretch of data lasting a time T. If we use a superscript F to denote the filtered version of a quantity, then the analog of $n_1 \circ n_2$ is

$$I_{12}(t) = \int_0^T n_1^F(t')n_2^F(t'+t)dt'. \tag{44}$$

Its variance is

$$<|I_{12}(t)|^2> = \int_0^T \int_0^T < n_1^F(y)n_1^{F*}(y')n_2^F(y+t)n_2^{F*}(y'+t) > dy dy'. \tag{45}$$

The key to evaluating this is the expectation

$$< n_1^F(t)n_1^F(t') > = 2\int_{f_1}^{f_2} S_1(f)\cos[2\pi f(t-t')]df, \tag{46}$$

where f_1 and f_2 are the lower and upper limits of the filtered frequency band ($f_2 = f_1 + B$), and where the factor of 2 arises because negative frequencies must be included in the filtered data as well as positive ones. It is a straightforward calculation to show that, assuming for simplicity that $S_i(f)$ has the constant value $\sigma_{i\,f}^2$ over the bandwidth, then for the most important case $2\pi BT \geq 1$,

$$<|I_{12}(t)|^2> \approx 8\sigma_{1\,f}^2\sigma_{2\,f}^2 BT. \tag{47}$$

This part of the noise is proportional to the bandwidth of the data and the duration of the correlation. The duration will usually be chosen so that the condition $2\pi BT \geq 1$ is satisfied, for otherwise the experiment would be too brief to detect any signal that fit within the bandwidth B. The remaining contributions to the variance of the cross-correlation come from the second and third terms of Eq. (43) (strictly, from their filtered and finite-time analogs). These are just like Eq. (19), and add to Eq. (47) a term equal to $(\sigma_{1\,f}^2 + \sigma_{2\,f}^2)\int_0^T |h^F(t)|^2\,dt$.

The case of most interest to us is where the "raw" signal $h^F(t)$ is smaller than the time-series noise in the bandwidth B in each detector, $n_i^F(t)$. Then the variance is dominated by Eq. (47) and we have the following expression for the signal-to-noise ratio of the cross-correlation

$$SN-\text{ratio} = \frac{\int_0^T |h^F(t)|^2\,dt}{[8\sigma_{1\,f}^2\sigma_{2\,f}^2 BT]^{1/2}}. \tag{48}$$

This has considerable resemblance to the filtering signal-to-noise ratio given in Eq. (20), and this justifies and makes precise our notion that cross-correlation can be thought of as using a noisy data stream as the filter. To convert Eq. (20) into Eq. (48), we must (i) replace the filter in the numerator with the signal h^F that is in the noisy "filter", and (ii) replace the filter power in the denominator with the noise power of the noisy filter, since we have assumed this power is the largest contributor to the noise.

However, Eq. (48) does not give us the signal-to-noise ratio for the gravitational wave signal, since its numerator is proportional to the *square* of the wave amplitude. This is the effect that we noted at the

beginning of this section, that the "filter" amplitude is proportional to the signal amplitude. The true amplitude signal-to-noise ratio is the square root of the expression in Eq. (48):

$$\frac{S}{N} = \frac{[\int_0^T |h^F(t)|^2\,dt]^{1/2}}{[8\sigma_{1\,f}^2\sigma_{2\,f}^2 BT]^{1/4}}. \tag{49}$$

There are two cases to consider here, long wavetrains and short pulses.

Long wavetrains. The best signal-to-noise is achieved if we match the observation time T to the duration of the signal or, in the case of pulsars, make T as long as possible. Let us assume for simplicity that the two detectors have the same noise amplitude, and let us denote by R the "raw" signal-to-noise ratio of the signal (its amplitude relative to the full detector noise in the bandwidth B),

$$R = \frac{h}{(2B\sigma_f^2)^{1/2}}.$$

Then we find

$$\frac{S}{N} \approx (\frac{BT}{8})^{1/4}R. \tag{50}$$

The signal-to-noise ratio increases only as the fourth root of the observation time. If we are looking at, say, the spindown of a newly formed pulsar, lasting 1 s, and we filter to a bandwidth of 1 kHz because we don't know where to look for the signal, then the enhancement factor $(BT/8)^{1/4}$ is only 3.3: short wavetrains are improved, but not dramatically. If we are looking at a pulsar, again in a broadband search with 1 kHz bandwidth, but in an observation lasting 10^7 s, then the enhancement of signal-to-noise is a factor of about 200. This is a lot less than can be achieved by filtering, but it is nonetheless significant. It shows that cross-correlation can be useful in pulsar searches. We will see below in detail how an all-sky search may be conducted by these methods.

Short pulses. Here one would set the bandwidth B equal to that of the pulse; if the pulse has duration roughly $T = 1/B$, and if again the two detectors have the same noise amplitude, then Eq. (49) gives a signal-to-noise ratio that is a factor of roughly $8^{1/4} \approx 1.7$ smaller than the optimum that filtering can achieve. For $TB \approx 1$ some of our approximations are breaking down, but it is reasonable that using this noisy filter would reduce the signal-to-noise by a factor of order 2. Since in this case filtering does not enhance the signal-to-noise ratio, neither does cross-correlation: if the signal is too weak to be seen above the broadband (bandwidth B) noise in one detector, it will not be found by cross-correlation.

I will conclude this section by pointing out another difference between filtering and cross-correlation. Since for signals below the broadband noise ($R < 1$), we do not know where the signal is in the data stream used as a filter, it follows that we cannot determine the time of arrival of the signal from the correlation, apart from a relatively crude determination based upon the presence or absence of correlations between given data sets of length T. The correlation also does not tell us the waveform and therefore it cannot determine the true amplitude of the signal. It can, however, determine the time-delays between the arrival of events at different detectors.

2.4.2 Cross-correlating differently polarized detectors

A more sophisticated approach to correlation has been devised by Gursel & Tinto (1989) in their approach to the inverse problem, which I will describe in more detail in Sec. (3) below. It works if there are at least three detectors in the network. I shall neglect noise for simplicity in describing the method. If we let θ and ϕ be the angles describing the position of the source on the sky and we use α_i, β_i, and χ_i to represent the latitude, longitude, and orientation of the i'th detector respectively, and if we have some to have is 500 km. This may be achievable within Europe, but it seems most unlikely that detectors in the USA will be built this close together. The data analysis is exactly the same as for two detectors on the same site.

3 Interpretation of the Data: the Inverse Problem

The inverse problem is the problem of how to reconstruct the gravitational wave from the observations made by a network of detectors. A single detector produces limited information about the wave; in particular, on its own it cannot give directional information and therefore it cannot say what the intrinsic amplitude is. With three detectors, however, one can reconstruct the wave entirely. In the last two or three years there has been considerable progress in understanding the inverse problem: see Tinto (1987), Dhurandhar & Tinto (1988), and Gursel & Tinto (1989). I will summarize the main ideas as I understand them at present, but this is an area in which much more development is likely soon.

3.1 Bursts seen in several detectors

A gravitational wave is described by two constants — the position angles of its source, (θ, ϕ) — and two functions of time — the amplitudes of the two independent polarizations $h_+(t)$ and $h_\times(t)$. Simple counting arguments give us an idea of how much we can learn from any given number of detectors. I will ignore the effects of noise at first.

For signals that stand out above the broadband noise:

- A single detector gives its response $r(t)$ and nothing else. Nothing exact can be said about the waves unless non-gravitational data can be used, as from optical or neutrino detections of the same event.
- Two detectors yield two responses and one approximate time delay between the arrival of the wave in one detector and its arrival in the other. Two functions of time and one constant should not be enough to solve the problem, and indeed they are not. The time delay is only an approximate one, because the two detectors will generally be responding to different linear combinations of $h_+(t)$ and $h_\times(t)$, so there will not be a perfect match between the responses of the two detectors, from which the time delay must be inferred. The time delay will confine the source to an error-band about a circle on the sky in a plane perpendicular to the line joining the detectors. The antenna patterns of the detectors can then be used to make some places on this circle more likely than others, but the unknown polarization of the wave will not allow great precision here. If the location of the source can be determined by other means, and if noise is not too large, then the two responses can determine the two amplitudes of the waves.
- Three detectors cross the threshold into precision astronomy, at least when the signals stand out against the broadband noise. Here we have three functions of time (the responses) and two constants (the time delays) as data, and this should suffice. As described in Sec. (2.4) above, correlations among the three detectors can pin down the location of the source and, if noise is not too important, the time-dependent amplitudes as well. In this case, there is redundant information in the data that can be used to test Einstein's predictions about the polarization of gravitational waves: pure transverse quadrupole.

If noise is so important that filtering is necessary, there is a completely different way of doing the counting. A given filter yields only two numbers: the maximum value of the correlation and the time the signal arrives (*i.e.* when it best matches the filter). We can only assume that the signal's waveform matches the "best" filter, so instead of two unknown time-dependent amplitudes we will have the response of the filter, the time of arrival, and a certain number of parameter constants that distinguish the observed waveform from others in its family.

Let us concentrate on coalescing binaries. The signal from a coalescing binary is an elliptically polarized, roughly sinusoidal waveform. The filters form a two-parameter family, characterized by the mass parameter $\mathcal{M}$ and the phase of the signal Φ, as in Eq. (8). The parameters we should be able to deduce are:

the amplitude h of the signal, the ellipticity e of its polarization ellipse (1 minus the ratio of the minor and major axes), an orientation angle ψ of the ellipse on the sky, and the binary's mass parameter $\mathcal{M}$. From these data we can not only determine the distance to the system, but also the inclination angle of the binary orbit to the line of sight (from e) and the orientation of the orbital plane on the sky (ψ).

The mass parameter $\mathcal{M}$ will be determined independently in each detector, and of course they will all agree if the event is real. Each detector in addition contributes the response of the filter, the phase parameter, and the time of arrival; these data must be used to deduce the five constants $\{\theta, \phi, h, e, \psi\}$. Here is how various numbers of detectors can use their data:[1]

- One detector does not have enough data, so it can only make average statements about the amplitude.
- Two detectors have four useful data: two responses, one phase difference, and one time delay. (Only the *differences* between the phases and times of arrival matter: the phase and time of arrival at the first detector are functions of the history of the source.) If the two detectors were identically polarized, the phase difference would necessarily be zero. A nonzero phase difference arises because the two principal polarizations in an elliptically polarized wave are 90° out of phase, so if the detectors respond to different combinations of these two polarizations, they will have different phases. With four data chasing five unknowns, the solution will presumably be a one-dimensional curve on the sky, but the problem has not yet been studied from this perspective.
- Three detectors have seven data: three responses, two phase differences and two time delays. The two time delays are sufficient to place the source at either of the intersections of two circles on the sky. For either location, the three responses determine h, e, and ψ. Presumably the phase differences would be consistent only with one of these positions, thereby solving the problem uniquely and incidentally providing the phase differences as a test of general relativity's model for the polarization of gravitational waves.

I have presented two different approaches to the inverse problem, but presumably in many situations the real situation will lie somewhere in-between. As far as I know, no work has been done on solving the inverse problem for intermediate cases, where noise is significant but not dominant.

3.2 The inverse problem for pulsars

We have already largely discussed how the inverse problem would be solved by cross-correlations. For observations of a pulsar (or other continuous source) already known from radio data, the position and frequency will be obtained from the filter or Fourier transform of stroboscopically sampled points that gives the best response: this will determine the position of the source, its polarization, and its frequency to very high accuracy. The amplitude of the waves will be determined to an accuracy that depends on the signal-to-noise ratio.

Notice that if a relatively low-frequency signal is detected by either filtering or sampling a much faster data stream (say 100 Hz samples from a 10 kHz data stream), then it would be possible to go back to the original data stream and refine the position and frequency by finding the filter that best fits a faster sampling of the data. In this way, positional accuracies approaching 10 milliarcseconds may be possible.

Once a signal has been identified, higher frequency harmonics should be searched for as well, since it is unlikely that the mass deformation giving rise to the gravitational waves will be such as to radiate at only one frequency. These harmonics will contribute detailed information about the star's structure.

Depending on the particular pulsar, there might be other information in the gravitational waves. It might be possible to see a proper motion, to measure the spindown rate or even its derivative, or (for pulsars nearer than 10–100 pc) to measure the parallax and hence the distance.

[1] This discussion is very different from previous ones I have given, *e.g.* Schutz (1989). In these I had not yet appreciated the importance of being able to determine the phase parameter independently of the time of arrival. This extra information makes it possible to solve the inverse problem with fewer detectors than I had previously believed.

definition of polarization of the waves so that we can describe any wave by its amplitudes h_+ and $h_\times$, then the response $r = \delta\ell/\ell$ of the i'th detector is a function of the form

$$r_i(t) = E_{+i}(\theta,\phi,\alpha_i,\beta_i,\chi_i)h_+[t-\tau_i(\theta,\phi)] + E_{\times i}(\theta,\phi,\alpha_i,\beta_i,\chi_i)h_\times[t-\tau_i(\theta,\phi)], \qquad (51)$$

where $\tau_i(\theta,\phi)$ is the time-delay between receiving a wave coming from the direction (θ,ϕ) at some standard location and at the position of the detector. We shall define the "standard location" by setting $\tau_1 = 0$. We need not be concerned here with the precise form of the functions E_{+i}, $E_{\times i}$, and τ_i, nor with the exact definitions of the various angles.

The response equations of the first two detectors may be solved for h_+ and $h_\times$ and substituted into the response equation for the third to predict its response, for an assumed direction to the source. Let this prediction be r_{3-pred}:

$$r_{3-pred}(t) = -[D_{23}r_1(t-\tau_3) + D_{31}r_2(t+\tau_2-\tau_3)]/D_{12}, \qquad (52)$$

where D_{ij} is the determinant

$$D_{ij} = E_{+i}E_{\times j} - E_{\times i}E_{+j}.$$

If there were no noise in the detectors, then for some choice of angles θ and ϕ there would be exact agreement between r_{3-pred} and the actual data from detector 3, r_{3-obs}. Given the noise, the best one can do is to find the angles that minimize the squared difference $d(\theta,\phi)$ between the predicted and observed responses during the interval of observation:

$$d(\theta,\phi) = \int_0^T |r_{3-obs}(t) - r_{3-pred}(t)|^2\, dt. \qquad (53)$$

Hidden in the integral for d are the correlation integrals we began with, *e.g.* $\int r_3(t)r_1(t-\tau_3)dt$. These will normally be the most time-consuming part of the computation of d for various angles, and should usually be done by FFT. Once the correlations have been computed for all possible time delays, they may be used to find the minimum of d over all angles; this will determine the position of the source. Notice that if the noise is small, this information can then be substituted back into Eq. (51) for the first two detectors to find $h_+(t)$ and $h_\times(t)$. This solves the inverse problem. But if the source is weaker than the noise, then this substitution will give mostly noise.

The information we have gained about the unpredicted source, even if it is weak, is that it is there: its position is known and its arrival time can be determined roughly by restricting the time-interval over which the correlation integrals are done and finding the interval during which one gets significant correlations. This is enough to alert other astronomers to look for something in the source's position.

2.4.3 Finding pulsars by cross-correlation

The Gursel-Tinto method can determine the direction to broadband burst sources, but it cannot discover them because, as we have already seen, cross-correlation does not enhance the signal-to-noise ratio of short bursts. But pulsars (and other continuous wave sources) are different: cross-correlation will help pull them out of the noise, and so we can expect that the three-detector analysis can be the basis of an all-sky search for such sources.

However, before we can be sure we can make this work, we have to ensure that the signals really do have the same time-dependence in each detector. Since detectors at different places on the Earth will have somewhat different motions and hence redshifts relative to a given source, it is not obvious that this will be possible. On the other hand, detectors on the same site will clearly always keep in phase with each other, so what we need to determine is the maximum separation over which detectors can remain correlated. I am grateful to J. Hough and H. Ward for clarifying my thinking on what follows.

The problem is closely related to that solved in Secs.(2.2.3.1) and (2.2.3.2). In a time T_{obs} we can tolerate a relative frequency shift between the two detectors of at most $2/T_{obs}$, and therefore a difference in velocity of $v/c = 2/fT_{obs}$ when observing waves of frequency f. If two detectors are separated by a distance R, then the Earth's rotation produces a maximum velocity difference $R\Omega_\oplus$. This leads to the constraint

$$R < \frac{2c}{\Omega_\oplus f T_{obs}} \approx 200\frac{1\text{kHz}}{f}\text{km}, \qquad (54)$$

where I have taken T_{obs} equal to 1/2 day, to get the maximum velocity difference that the Earth's rotation can generate. The other motions of the Earth have smaller angular velocity and therefore make less stringent demands on R.

It follows that a network of detectors that searches for millisecond pulsars by cross-correlation must be no more than a few hundred kilometers in size. A worldwide network could search for low-frequency pulsars (tens of Hz). Note that a pulsar signal from directly above the Earth's North Pole would experience no redshift in any detector at all. The constraint above gets weaker as $1/\sin\theta_p$, where θ_P is the angular position of the pulsar measured down from the pole.

The Gursel-Tinto correlation method completely changes the picture for an all-sky search for pulsars. It requires only a few Fourier transforms, and therefore can be done with an effort that is comparable to or less than that required for detecting known pulsars by filtering. How is this possible? Are we getting something for nothing?

The answer, of course, is no. The improvement in signal-to-noise is less than that for filtering, so the search cannot go as deep as individual pulsars can be found. Moreover, all the correlation will tell us about unknown pulsars is that there is a source at a particular position. Its frequency and other characteristics (such as whether it is in a binary system) do not come out. Nevertheless, the position is probably the most important datum: it allows intensive searches for pulsed electromagnetic radiation from that direction, and reduces the amount of work that one would have to do to try to discover the pulsar in the data of a single detector, using stroboscopic sampling.

2.5 Using cross-correlation to search for a stochastic background

Another very important observation that interferometers will make is to find or set limits upon a background of radiation. This is much easier to do than finding discrete sources of continuous radiation, because there is no direction-finding or frequency-searching to do. This problem has been discussed in detail by Michelson (1987).

The most sensitive search for a background would be with two detectors on the same site, with the same polarization. Current plans for some installations envision more than one interferometer in one vacuum system, which would permit a correlation search. One would have to take care that common external sources of noise are excluded, especially seismic and other ground disturbances, but if this can be done then the two detectors should respond identically to any random waves coming in, and should therefore have the maximum possible correlation for these waves. The correlation can be calculated either by direct multiplication of the sampled data points ($2N$ operations per time delay between the two data sets) or by Fourier transform methods as in Sec. (2.2.3.3) above. We are only interested in the zero-time-delay value of the correlation, but in order to test the reality of the observed correlation, one would have to compute points at other time delays, where the correlation is expected to fall off. (How rapidly it falls off with increasing time delay depends on the spectrum of the background.) The choice of technique — direct multiplication or Fourier transform — will depend on the number of time delays one wishes to compute and the capacity of one's computer.

If separated detectors are used, the essential physical point is that two separated detectors will still respond to waves in the same way if the waves have a wavelength λ much longer than the separation between the detectors. Conversely, if the separation between detectors is greater than $\lambda/2\pi$, there is a significant loss of correlation. It is important as well to try to orient the detectors as nearly as possible in the same polarization state. In order to search up to 1 kHz, the maximum separation one would like

4 Data storage and exchange

Although the amount of data generated by a four-detector network will be huge, I would argue strongly that our present ignorance of gravitational wave sources makes it important that the data should be archived in a form that is relatively unprocessed, and kept for as long a time as possible, certainly for several years. It may be that new and unexpected sources of gravitational waves will be found, which will make it desirable to go over old data and re-filter it. It may also be that new classes of events will be discovered by their electromagnetic radiation, possibly with some considerable delay after the event would have produced gravitational waves, and a retrospective search would be desirable. In any case, we have already seen that it will be important to exchange essentially raw data between sites for cross-correlation searches for unknown events. Once exchanged, it is presumably already in a form in which it can be stored.

4.1 Storage requirements

We have seen in the introduction that a network could generate 5000 optical discs or videotapes per year. Data compression techniques and especially the discarding of most of the housekeeping data at times when it merely indicated that the detector was working satisfactorily could reduce this substantially, perhaps by as much as a factor of ten. The cost of the storage media is not necessarily trivial. Even videotapes will run at \$2,000–\$20,000 per year, a cost that is not likely to decrease substantially; and optical discs of large capacity could cost 10–30 times as much at present prices (which will, hopefully, come down). Added to this is the cost of providing a suitable storage room, personnel to supervise the store, and equipment to make access to the data easy.

4.2 Exchanges of data among sites

We have already seen how important it will be to cross-correlate the raw data streams. At a data rate of some 100 kbytes per second, or even at 10 kbytes per second if the data are compressed as described above, this is well above the present bandwidths of international electronic mail networks. But these bandwidths are being constantly upgraded, and so in five years the situation may be considerably different: it may be possible, at reasonable cost, to exchange short high-bandwidth bursts of data regularly via optical-fiber-to-satellite-to-optical-fiber routes. Alternatively, a cheaper solution might be to exchange optical discs or videotapes physically, by airmail or international courier. This would be slower, but if lists of filtered events were exchanged on electronic mail networks, then there may be less urgency about exchanging the full data sets.

4.2.1 Protocols, analysis and archiving centers

It will be clear from our discussion that exchanging and jointly analysing data will require careful planning and coordination among all the groups. Discussions to this end are in a rudimentary stage now, but could soon be formalized more. Besides decisions on compatible hardware, software, data formats and modes of exchange, there are a number of "political" questions that need to be resolved before observations begin. We are dealing with data that the groups involved have spent literally decades of their scientific careers to be in a position to obtain, and the scientific importance of actual observations of gravitational waves will be momentous. Questions of fairness and proprietary rights to the data could be a source of considerable friction if they are not clearly decided ahead of time. A model for some of these decisions could be the protocols adopted by the GRAVNET network of bar antennas, described elsewhere in this volume. Other models might be international VLBI, or large particle-physics collaborations.

Some of the questions that need to be addressed are:

- how much data needs to be exchanged;
- what groups have the right to see and analyze the data of other groups and what form of acknowledgement they need to give when they use it;

- what powers of veto groups have over the use of their data, for example in publications by other people;
- how long the proprietary veto would last before the data become "public domain" (the funding agencies will presumably apply pressure to allow ready access to the data by other scientists after some reasonable interval of time);
- how long the data need to be archived.

Given the volume of data and the logistical complications of multi-way exchanges of it, it may be attractive to establish one or more *data analysis and archiving centers*, perhaps one in Europe and one in the USA. These would collect the data and store it, and perform the cross-correlations that can only be done with the full data sets on hand. They would be "facilities" for the use of the various groups, and would coordinate the distribution of data — not just the event lists for immediate coincidence analysis but also raw data — to anyone who wanted it and had a right to it. They could also be the locations for the massive computers that would be necessary to search for pulsars. They could serve other functions as well, such as maintaining wider databases on-line (bibliographies, astronomical catalogues, etc.), distributing newsletters and even scientific reports and papers electronically, or coördinating the development of common software for data reduction. All of this could in principle be made available to all the groups simply by allowing group members to log onto the centers' computers remotely.

5 Conclusions

In this review I have set out what I understand about the data analysis problem as of December 1988. Evidently, the field is covered very non-uniformly: coalescing binaries have received much more attention than pulsars or stochastic sources so far, and protocols for data exchange are something mainly for the future. The inverse problem is under intensive study right now.

Nevertheless, it is clear that questions of the type we have discussed here will influence in an important way decisions about the detectors: how many there will be, where they will be located, what their orientations will be, what weights one should apply to the various important parameters affecting their sensitivity (*e.g.*, length, seismic isolation, laser power) when deciding how to apportion limited budgets to attain the maximum sensitivity. Other questions that I have not addressed will also be important, particularly choosing the particular recycling configuration most suitable to searching for a given class of sources.

From the present perspective, it seems very likely that in ten years a number of large-scale interferometric detectors will be operating with a broadband sensitivity of around 10^{-22}. The data should contain plenty of coalescing binaries and at least a few supernovae; but the most exciting thing that we can look forward to is the unexpected: will this sensitivity suffice to discover completely unanticipated sources? The best way to ensure that it does is to make sure that our data-analysis algorithms and data-exchange protocols are adequate to the task: given the enormous efforts being made by the hardware groups to develop the detectors, and the considerable amount of money that will be required to build them, it is important that development of the data-analysis tools not be left too late. The data-analysis problem deserves much more work than it has so far received.

Acknowledgements

It is a pleasure to acknowledge the helpfulness of many colleagues, from whom I have learned much about data analysis and many of whom have contributed crucial insights into the problem. They include Mark Davis, Ron Drever, Jim Hough, Andrzej Krolak, Sean Lawrence, Alberto Lobo, Brian Meers, Eduardo Nahmad-Achar, Norna Robertson, Kip Thorne, Massimo Tinto, Philippe Tourrenc, Harry Ward, Rai Weiss, and Sheryl Smith Woodward. I am grateful to Drs. Hough, Robertson, Tinto, and Ward for correcting errors in a first draft of this chapter. Any remaining errors are, of course, my own.

This article was produced by the LaTeX formatting system.

References

Davis, M.H.A., 1989, in B.F. Schutz, ed., *Gravitational Wave Data Analysis* (Reidel, Amsterdam).

Dewey, D., 1986, in R. Ruffini, ed., *Proceedings of the Fourth Marcel Grossmann Meeting on General Relativity* (Elsevier, Amsterdam) 581.

Dhurandhar, S.V., Schutz, B.F., and Watkins, W.J., 1989, *in preparation*.

Dhurandhar, S.V., and Tinto, M., 1988, *Mon. Not. R. astr. Soc.*, **234**, 663–676.

Evans, C.R., 1986, "An Approach for Calculating Axisymmetric Gravitational Collapse", in J. M. Centrella, ed., *Dynamical Spacetimes and Numerical Relativity* (Cambridge University Press, Cambridge, England) 3–39.

Evans, C.R., Iben, Jr., I., and Smarr, L., 1987, *Astrophys. J.*, **323**, 129–139.

Gursel, Y., and Tinto, M., 1989, *preprint*.

Horowitz, P., 1969, *Rev. Sci. Inst.*, 40, 369–370.

Hough, J., Drever, R.W.P., Ward, F., Munley, A.J., Newton, G.P., Meers, B.J., Hoggan, S., and Kerr, G.A., 1983, *Nature*, 303, 216.

Ipser, J.R., and Managan, R.A., 1984, *Astrophys. J.*, 282, 287.

Krolak, A., 1989, "Coalescing Binaries to Post-Newtonian Order", in B.F. Schutz, ed., *Gravitational Wave Data Analysis* (Reidel, Amsterdam).

Krolak, A., and Schutz, B.F., 1987, *Gen. Rel. Gravit.*, 19, 1163–1171.

Livas, J.C., 1987, *Upper Limits for Gravitational Radiation from Some Astrophysical Sources* (Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, Mass.).

Lyne, A., 1989, in B.F. Schutz, ed., *Gravitational Wave Data Analysis* (Reidel, Amsterdam).

Michelson, P.F., 1987, *Mon. Not. R. astr. Soc.*, 227, 933–941.

Piran, T., and Stark, R.F., 1986, "Numerical Relativity, Rotating Gravitational Collapse, and Gravitational Radiation", in J. M. Centrella, ed., *Dynamical Spacetimes and Numerical Relativity* (Cambridge University Press, Cambridge, England) 40–73.

Schutz, B.F., 1986, "Hubble's Constant from Gravitational Wave Observations", in H. Sato, and T. Nakamura, eds., *Gravitational Collapse and Relativity* (World Scientific, Singapore) 350–368.

Schutz, B.F., 1989, "Data Analysis Requirements of Networks of Detectors", in B.F. Schutz, ed., *Gravitational Wave Data Analysis* (Reidel, Amsterdam).

Smith, S. ,1987, *Phys. Rev. D*, 36, 2901–2904.

Srinath, M.D., and Rajasekaran, P.K., 1979, *An Introduction to Statistical Signal Processing with Applications* (Wiley, New York).

Thorne, K.S., 1987, " Gravitational Radiation", in S.W. Hawking, and W. Israel, eds., *300 Years of Gravitation* (Cambridge University Press, Cambridge, England) 330–458.

Tinto, M., 1987, *Theoretical Aspects of Gravitational Wave Detection* (Ph.D. Thesis, University College Cardiff, Cardiff, U.K.).