Experiments Relevant to the Development of Laser Interferometric Gravitational Wave Detectors

by

Norman Lewis Mackenzie

Department of Physics and Astronomy
University of Glasgow.

Presented as a thesis for the degree of
Ph.D. in the University of Glasgow.
September 1989.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

Contents

Acknowledgements

I am grateful for the interest and support shown by Prof. E. Laing and Prof. I. Hughes during the period of this work.

Thanks are due to Jim Hough for his guidance and for his advice during the writing of this thesis, and to Harry Ward and Brian Meers for many helpful discussions.

I would like to thank the other members of the gravitational wave group—Gavin Newton, Norna Robertson, Stuart Hoggan, Graham Kerr, John Mangan, David Robertson, Kenneth Strain, Caroline Cantley and Allan Carmichael—for their help and encouragement during this work.

For help with computing problems I am grateful to Allan Flavell and Paul Rosenberg.

Technical support was provided by Jim Pittillo, Angus McKellar, Allan Latta, John Jarvis and David Edwards.

I would like to give special thanks to my parents for their support and patience.

During the time that this work was carried out, I was in receipt of financial support from the SERC.

This thesis was produced using TEX and the LATEX formatting package.

Summary
Experiments Relevant to the Development of Laser Interferometric Gravitational Wave Detectors.

The development of gravitational wave detectors has been in progress for approximately twenty-five years. As yet there has been no clear evidence for the successful detection of such propagating fluctuations in the curvature of spacetime, but the prospects seem good that detectors of sufficient sensitivity to detect gravitational waves of astrophysical origin can be constructed in the near future. The most promising form of detector is the long baseline laser interferometer, and prototypes are being developed at a number of sites around the world. A 10 metre prototype is currently being developed in Glasgow. This thesis is an account of work based on the Glasgow prototype.

After an elementary introduction to the theoretical foundations of gravitational waves, various sources of gravitational radiation, the nature of their emitted signal and their strengths are considered. Suitable detectors and their possible sensitivities are reviewed. Noise sources which could limit the sensitivity of laser interferometer detectors and the constraints which these place on the design of the detector are discussed.

Since the test masses in an interferometer detector must be freely suspended as pendulums, yet their orientation must be accurately controlled to maintain correct alignment of the optical cavities forming the interferometer, an active orientation control system was developed and installed on the Glasgow prototype. This system

provides a high degree of positional and angular stabilisation at low frequencies while leaving the test mass essentially free at high frequencies. Some of the potential limitations and noise sources are noted and their magnitudes calculated.

A digital recording system was designed and used to record data from the prototype detector at Glasgow. The effects of the detector's response are analysed and techniques to recover the gravitational wave signal from the recorded data are described. The analysis of some data recorded with this system is then reported. The pulse statistics of the interferometer are analysed and the implications for searches for millisecond pulses of gravitational waves are discussed.

The results of a search for periodic signals in the detector output are presented. Various sources of contamination which may be present in the detector output are identified, limitations of the recorded data are noted, and techniques which may be used to reduce the importance of these effects are described.

Preface

Since 1978, a 10 m prototype laser interferometer gravitational wave detector has been under development at the University of Glasgow. This thesis describes certain aspects of this development and work carried out between October 1985 and October 1988 to quantify the performance of the detector and so aid its development.

Chapter 1 gives a brief introduction to the gravitational wave theory which motivates the development of such prototype detectors, and describes the expected sources and proposed detector techniques. This chapter is based on the published literature.

Chapter 2 discusses the Glasgow prototype detector and noise sources in laser interferometer detectors. Most of this work is derived from the published literature, including the publications of the Glasgow group.

In Chapter 3 an orientation control system, which may be used to align the optical cavities in the detector, is described. This works by applying torques to a mass suspended as a pendulum to which the cavity mirrors are fixed. This is based on a system developed on the prototype detector at Garching. The work was carried out at the instigation of J. Hough in collaboration with B. J. Meers.

Chapter 4 describes the design and development of the software and hardware of a data acquisition system for the prototype detector. The use of this system to record data from the prototype shortly after the identification of SN1987a is discussed along with the characteristics of the recorded data. Techniques for the removal of spurious detector signals are detailed.

The analysis of this recorded data is described in Chapter 5 and Chapter 6. In Chapter 5 the characteristics of the time series and the implications for searches for pulses of gravitational waves are discussed. In Chapter 6 techniques for the

detection of periodic gravitational waves are reviewed. The practical limitations of the recorded data are detailed and techniques for reducing or removing these are described. The work in Chapter 4, Chapter 5 and Chapter 6 was suggested by J. Hough.

In Chapter 7 the prospects for the proposed long baseline instruments are reviewed.

Appendix A describes measurements to determine the vibration isolation provided by a stack of alternate layers of lead and rubber.

Since the development of the data acquisition system described in Chapter 4 another system has been developed based on more modern hardware. David Robertson and Kenneth Strain were responsible for the filter hardware development and software development respectively. After initial tests showed that the software could not sustain the desired recording bandwidth—due to delays in the transfer of sampled data to the controller—a scheme for packing the data before transfer was proposed, developed and tested by the author. Appendix B lists this program.

Chapter 1

The Generation and Detection of Gravitational Waves

1.1 Introduction

The existence of gravitational waves is a prediction of the general theory of relativity and indeed all other theories of gravity which have local Lorentz invariance of their gravitational laws. These waves should carry energy away from mechanical systems which move in a suitable way and it should be possible, in principle, to detect this radiation. In order to detect these waves and directly verify their existence, it is necessary to predict the behaviour of those astrophysical sources which are expected to emit significant quantities of gravitational radiation and consider the nature of the wave propagation and methods of detection. This chapter will briefly introduce the theoretical predictions which motivate the present experimental effort to detect gravitational waves and will describe some of the expected sources and the effect of a passing gravitational wave on a suitable detector.

1.2 Einstein's Prediction of Gravitational Waves

In 1916, in a paper presented to the Royal Prussian Academy of Sciences, Einstein showed that the field equations of general relativity can be made linear and a solution obtained when a weak field approximation is introduced [Einstein 1916]. From this solution, Einstein was able to show that the gravitational field propagates through space with the speed of light, and he derived an expression for the energy which would be lost by a system of masses to gravitational radiation. In

a later paper presented to the Academy [Einstein 1918], Einstein derived an expression for the power absorbed from a passing gravitational wave by a mechanical system. These two papers therefore provided a description of the generation, propagation and detection of gravitational waves in the weak field limit. Since this weak field approximation holds in many cases of astrophysical interest, this linearised treatment of propagation and detection is often sufficient. However, the weak field approximation is not valid when considering emission from systems where gravity does significant work, for example binary star systems. In such cases, a more complete treatment is required.

1.3 The Linearised Theory of Gravity

In the special theory of relativity, events are represented by points in a flat, four dimensional spacetime known as Lorentz spacetime. The coordinates of this spacetime are denoted by x_α, where $\alpha = 0, 1, 2, 3$. The time coordinate is usually identified with x_0, with x_1, x_2 and x_3 representing the spatial coordinates. The interval between events is defined to be,

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu, \tag{1.1}$$

where $\eta_{\mu\nu}$ is the metric tensor of Lorentz spacetime, that is, taking c=1,

$$\eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{1.2}$$

In order to include gravity into this relativistic theory, Einstein, motivated by the equivalence of inertial and gravitational mass, introduced the concept of a curved spacetime in which freely falling particles follow geodesic paths. The geometry of this spacetime is determined by the distribution of mass-energy within the spacetime. Whereas the special theory of relativity uses a metric tensor which is

constant throughout all of the spacetime, in general relativity the interval between neighbouring events is determined by a metric tensor, $g_{\mu\nu}$, which is a function of the position in the spacetime and which describes the geometry of the spacetime. Since the spacetime is no longer flat it is useful to define a tensor, the Riemann curvature tensor, $R^{\alpha}{}_{\beta\mu\nu}$, which gives a coordinate independent measure of the curvature of spacetime at any event. This curvature tensor is a function only of the metric tensor and its derivatives and describes the manner in which the interval between geodesics changes, that is gravitational tidal effects.

The field equations of general relativity can be written in the form

$$G^{\mu\nu} = 8\pi T^{\mu\nu}, \tag{1.3}$$

where $G^{\mu\nu}$ is Einstein's tensor, and $T^{\mu\nu}$ is the stress-energy tensor.

Einstein's tensor is constructed from the Riemann curvature tensor and the metric and has the property of being divergence free, that is $G^{\mu\nu}_{;\nu} = 0$, implying the conservation of energy-momentum. The field equations therefore equate a tensor which is a function of the geometry of the spacetime with a tensor describing the distribution of energy momentum so that the distribution of energy-momentum affects the geometry of the spacetime. All forms of mass-energy are acted upon by gravity in the same way, that is they follow geodesic paths in the curved spacetime, and they all contribute to the curvature in an identical manner. This means that the particles which mediate the gravitational force, gravitons, will themselves contribute to the gravitational curvature and so the field equations of general relativity are non-linear. However, these non-linear field equations can be made linear to first order by introducing a weak field approximation. In this linearised theory the metric tensor can be written in the form,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \tag{1.4}$$

where $|h_{\mu\nu}| \ll 1$, and $\eta_{\mu\nu}$ is the metric of Lorentz spacetime.

It can be shown that, in this linearised theory, $h_{\mu\nu}$, and indeed all physical fields such as that represented by the curvature tensor, can be regarded as tensor

fields in Lorentz spacetime. Since we are considering weak fields we may raise and lower indices with $\eta_{\mu\nu}$ so that

$$h^{\mu}{}_{\beta} = \eta^{\mu\alpha} h_{\alpha\beta},$$

and by making the definitions,

$$h = h^{\alpha}{}_{\alpha},$$

$$\overline{h}_{\alpha\beta} = h_{\alpha\beta} - \frac{1}{2}\eta_{\alpha\beta} h,$$

and by choosing to work with a set of coordinates where

$$\overline{h}^{\mu\nu}_{,\nu} = 0,$$

it is possible to derive an expression for $G^{\mu\nu}$ in terms of $h^{\mu\nu}$. The weak-field equations derived in this manner are,

$$\Box \overline{h}^{\mu\nu} = -16\pi T^{\mu\nu}. \tag{1.5}$$

1.4 The Propagation of Gravitational Waves

The linearised field equations in vacuum,

$$\left(-\frac{\partial^2}{\partial t^2} + \nabla^2\right) \overline{h}_{\mu\nu} = 0, \tag{1.6}$$

have a simple, monochromatic, plane wave solution

$$\overline{h}_{\mu\nu} = A_{\mu\nu} \exp\left(i k_{\alpha} x^{\alpha}\right), \tag{1.7}$$

where $\mathbf{k}$ is a null vector and $\mathbf{A}$ is orthogonal to $\mathbf{k}$. This solution describes a wave of frequency $\omega = k_0 = (k_1^2 + k_2^2 + k_3^2)^{\frac{1}{2}}$ propagating with the speed of light. By choosing a specific gauge it can be shown that the amplitude $A_{\mu\nu}$ has two independent components. In particular, a gauge can be chosen such that $A_{\mu\nu}$ is

transverse to the direction of propagation and has a trace equal to zero. In this transverse, traceless, or TT, gauge the amplitude tensor takes the form

$$A_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{1.8}$$

for a wave travelling in the z direction. The two amplitudes, A_{xx} and A_{xy}, correspond to two independent polarisation states of the wave. The physical significance of these polarisation states will be discussed in the next section.

1.5 The Detection of Gravitational Waves

A single freely falling observer would not be able to detect any effect due to the wave; the wave produces no local effect. In fact in the TT gauge any freely falling observer which is initially at rest will remain at the same TT coordinate position, even in the presence of a gravitational wave. However, since the metric in the linearised theory is given by $h_{\mu\nu} + \eta_{\mu\nu}$, it is apparent that the geodesic separation of test masses in free fall will vary in a manner determined by the gravitational wave amplitudes. If we consider a system of two particles, one at the origin of a TT coordinate system and the other at the position $x = \varepsilon$, $y = z = 0$, then the proper distance between them is

$$l = \int_0^{\varepsilon} \left| ds^2 \right|^{\frac{1}{2}}. \tag{1.9}$$

For the particular case when $h_{xy} = 0$ and $h_{xx} = h$ we see that

$$l = \int_0^{\varepsilon} ((1+h)dx^2)^{\frac{1}{2}} \tag{1.10}$$

$$\simeq \left(1 + \frac{1}{2}h\right)\varepsilon \tag{1.11}$$

Similarly for two particles, one at the origin and the other at $x = z = 0$, $y = \varepsilon$, it is easily shown that the proper distance is given by

$$l = \left(1 - \frac{1}{2}h\right)\varepsilon \tag{1.12}$$

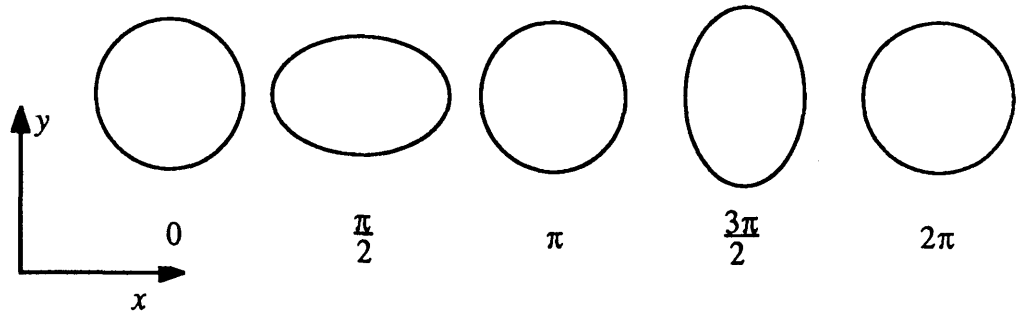

Figure 1.1: The effect of a gravitational wave with $A_{xy} = 0$ and $A_{xx} = A$ on a ring of test particles.

The apparent change in the separation, ε, of two test particles is therefore given by

$$dl = \frac{1}{2}h\varepsilon. \tag{1.13}$$

The effect of the wave can therefore be regarded as a strain in space which is proportional to the wave's amplitude. If such a gravitational wave encounters a circular ring of test particles which lie in a plane perpendicular to the wave's direction of propagation, then the ring's shape will be deformed as the wave passes. In the case described we have a linearly polarised wave, and the ring will deform into an ellipse which oscillates in the manner shown in Figure 1.1. The other, independent, linearly polarised wave, with $h_{xx} = 0$, $h_{xy} = h$ produces ellipses whose axes are rotated by an angle of 45°, as shown in Figure 1.2. These polarisation states can combine to produce circular and elliptical polarisation states. Since free particles at rest will remain at the same TT coordinate it was originally thought that there was no physical reality to gravitational waves, however, a thought experiment proposed by Bondi [Bondi 1957] demonstrated that these gravitational waves carry stress-energy. If two test masses are constrained to slide along a rod, then the effect of a suitably polarised wave will be to drive the test masses backwards and forwards along the rod. In the presence of friction this will generate heat; therefore energy has been extracted from the wave. Since the wave produces no local effects it is not possible to localise this stress-energy, though it is possible

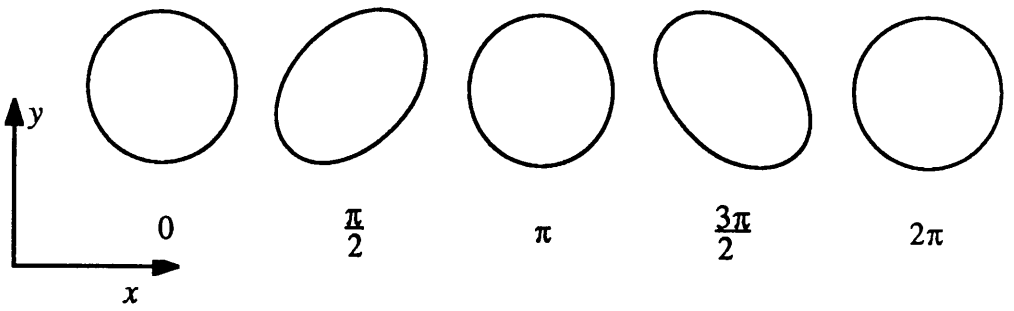

Figure 1.2: The effect of a gravitational wave with $A_{xx} = 0$ and $A_{xy} = A$ on a ring of test particles.

to calculate the stress-energy averaged over several wavelengths of the wave and this is given by [Misner *et al.* 1973]

$$T^{GW}_{\alpha\beta} = \frac{1}{32\pi} \sum_{i,j}^{3} \left\langle h^{TT}_{ij,\alpha}\, h^{TT}_{ij,\beta} \right\rangle , \tag{1.14}$$

where $\langle \ldots \rangle$ denotes the average over several wavelengths.

1.6 The Generation of Gravitational Waves

For electromagnetic radiation the fundamental, and most efficient, radiation mechanism is dipole radiation. There is no such mechanism for gravitational radiation since there is no equivalent of positive and negative charge. The conservation of momentum precludes the existence of gravitational dipole radiation. However, just as gravitational radiation can be detected by its effect on the quadrupole moment of a system of masses, so radiation should be emitted by any system of masses whose mass quadrupole moment changes with time. Although the gravitational radiation emitted by any source can in principle be calculated directly from the field equations, for real sources such a calculation requires considerable computational effort and the field of numerical relativity has only recently become possible with the advent of supercomputers. Fortunately it can be shown (see for example

Misner *et al.* 1973) that in the so called slow motion approximation, where all the velocities in the source are very much less than the speed of light, the gravitational wave emission is given approximately by

$$h_{jk}^{TT} = \frac{2}{r}\frac{\partial^2}{\partial t^2}\mathcal{I}_{jk}^{TT}(t-r), \tag{1.15}$$

where

$$\mathcal{I}_{jk} = \int \rho(t)\left[x^j x^k - \frac{1}{3}r^2\delta_{jk}\right] d^3x \tag{1.16}$$

is the mass quadrupole moment of the source integrated over the volume of the source, and r is the distance to the source.

The total power emitted in the form of gravitational waves is given by

$$L_{GW} = \frac{dE}{dt} = \frac{1}{5}\frac{G}{c^5}\left\langle \mathcal{I}_{jk}\,\mathcal{I}_{jk}\right\rangle . \tag{1.17}$$

This quadrupole formalism can be used to estimate the radiation emitted by a wide variety of sources. For example, consider the case of a bar length L, of mass M, rotating around its centre with an angular velocity ω. The power emitted by such a source in the form of gravitational radiation will be

$$\frac{4GM^2L^4\omega^6}{45c^5}. \tag{1.18}$$

Even if such a source could be constructed with $L = 10\text{m}$, $M = 10^5\text{kg}$ and $\omega = 1000\text{rad s}^{-1}$ the power output would be only 10^{-22}W. It can be seen then, that since the coupling of this gravitational energy to matter is equally weak, it is unlikely that any terrestrial source could generate sufficient gravitational radiation to be detected.

1.7 Sources of Gravitational Radiation

As discussed above, any system of masses whose quadrupole moment changes with time will emit gravitational radiation. However the amplitude of this radiation is so weak that with present technology it is not possible to construct a laboratory generator of gravitational waves whose output could be detected. However

it is expected that many violent astrophysical phenomena will be sources of gravitational waves with an amplitude large enough to be detectable with the long baseline interferometers which are currently being designed [Hough *et al.* 1986, Winkler *et al.* 1986, Linsay *et al.* 1983]. Indeed observations of the orbital parameters of the binary pulsar PSR1913+16 have already given indirect evidence that gravitational waves are being emitted [Taylor *et al.* 1979]. The direct detection of waves emitted from such sources and the determination of the parameters of the wave would provide a rigorous test of general relativity and other relativistic theories of gravity. After this initial detection, the analysis of gravitational wave signals would provide information about the large scale motion of many interesting astrophysical systems.

Gravitational wave sources fall into two main classes: those which produce short bursts of gravitational radiation and those which produce nearly monochromatic gravitational waves. In addition it is possible that there is a stochastic background of gravitational radiation. The following sections will introduce some of these sources and review current estimates of their gravitational wave amplitude.

1.7.1 Supernovae

Supernovae are possible sources of short bursts of gravitational radiation. They fall into two distinct categories, Type I and Type II.

Type I supernovae are thought to occur in binary systems in which a white dwarf accretes matter from a companion star. The exact mechanism for the violent outburst of energy is uncertain, but is thought to be due to either nuclear detonation of the white dwarf or the gravitational collapse of the white dwarf. Out to a distance of 15 Mpc, the distance to the Virgo Cluster of galaxies, a few Type I supernovae are seen per year.

Type II supernovae are thought to occur when the nuclear reactions in the cores of highly evolved massive stars are unable to provide sufficient radiation pressure to prevent rapid gravitational collapse. The core of the star collapses in $\sim$1 sec to a point where repulsion between the nucleons sets in and the cores bounces

outwards, generating a shock wave which ejects the envelope of the star. The time scale of the bounce is ~1msec. A neutron star or black hole will be left behind. It is expected that there will be a few Type II supernovae per year out to the Virgo Cluster.

Collapsing stars will only produce significant gravitational radiation if the collapse is not spherically symmetric. If a Type I supernovae is indeed a collapsing white dwarf, it is likely that it will have been spun up by matter accreted from its companion and that centrifugal forces will induce a non-spherical collapse. The degree of asymmetry in Type II supernovae is uncertain. In fact, most numerical simulations of stellar collapse have assumed that the progenitor star is spherically symmetric, remaining so during the collapse, since this assumption greatly simplifies the calculation. Saenz and Shapiro have considered the case where the core of the star is ellipsoidal [Saenz and Shapiro 1978, Saenz and Shapiro 1979, Saenz and Shapiro 1981]. They found that the core bounced repeatedly, becoming increasingly asymmetric, and that this effect greatly increased the efficiency of the gravitational wave emission. Their numerical solutions also predicted the waveforms of the gravitational waves emitted, typical examples are shown in Figure 1.3. One common feature of these waveforms is a short pulse lasting about 1ms. In addition some of the waveforms have an oscillatory nature corresponding to the core bounces and this feature may increase the probability of their detection. Depending on the model chosen, either a fast, free fall collapse or a slow collapse which is resisted by thermal pressure, Saenz and Shapiro's results predict that the energy emitted as gravitational waves is between

$$\Delta E \sim (10^{-4} \rightarrow 10^{-2}) Mc^2 \tag{1.19}$$

for a fast collapse (where M is the mass of the star) and

$$\Delta E \sim (10^{-6} \rightarrow 10^{-4}) Mc^2 \tag{1.20}$$

for slow collapse. However, since the frequency of the emitted radiation also changes, the wave amplitude in either case should be $h \sim 10^{-22}$ for sources out to the Virgo Cluster, and Saenz and Shapiro estimate that there should be one such event per month.

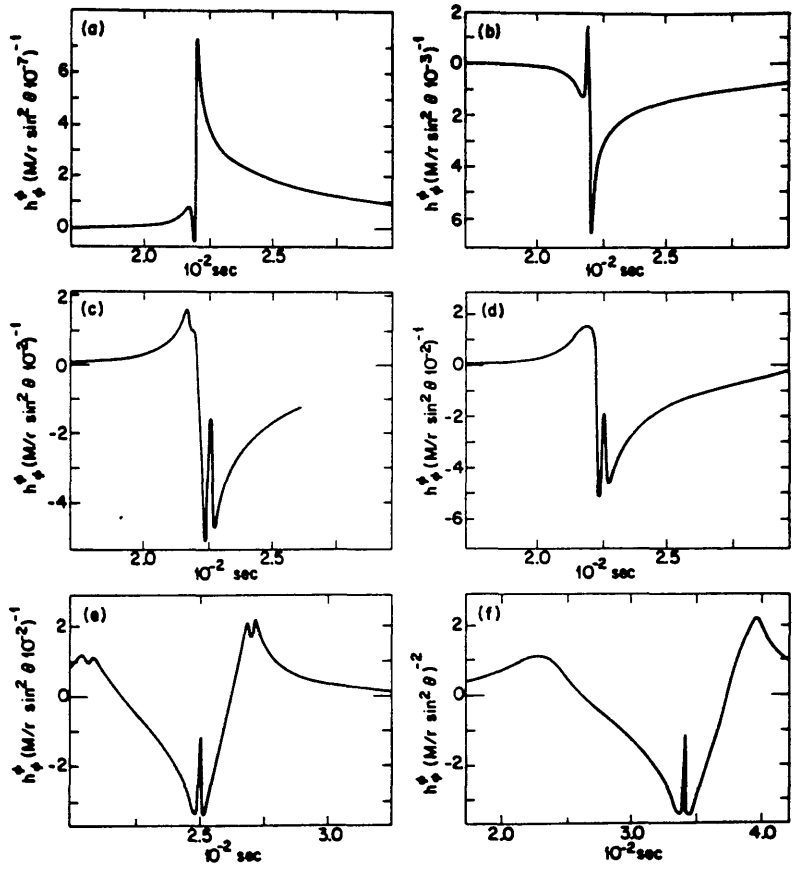

Figure 1.3: Gravitational waveforms predicted for the gravitational collapse of a Type II supernova (from Saenz and Shapiro 1978).

The gravitational radiation produced would be greatly increased if the core of the collapsing star became unstable, perhaps even splitting into two orbiting components. In such a case the efficiency of production could approach

$$\Delta E \sim 10^{-2} Mc^2. \tag{1.21}$$

1.7.2 Black Hole Formation

The creation of a black hole is expected to radiate considerable quantities of energy in the form of gravitational waves. Black holes may be formed by a number of astrophysical processes. For example, low mass black holes could be the final result of the gravitational collapse of stars which are too massive to form neutron stars. More massive black holes may be formed when stars collide and coalesce (a possibility in dense galactic nuclei), when accretion of material onto a neutron star increases its mass above the Oppenheimer-Volkoff mass, or even by the collapse of very massive stars. Event rates for black holes formed by stellar collapse have been estimated [Thorne 1987] at less than about 1 collapse per year out to the Virgo cluster. The formation of super massive black holes, $M \sim 10^6 M_\odot$, is believed to be much less than about 1 per year out to the Hubble distance.

When a black hole is formed it is expected to vibrate. The vibrations of the quadrupolar modes will be damped by gravitational radiation and so this radiation will have a characteristic damped oscillatory form from which information about the black hole can be obtained. The waveforms associated with specific black hole models have been calculated numerically [Stark and Piran 1985]. A typical example of such a waveform is shown in Figure 1.4. It can be seen that a short burst of gravitational waves, of order 1 millisecond, is produced. The strain amplitude at the earth of a source at distance r will be

$$h \sim 10^{-20} \left(\frac{\varepsilon}{0.01}\right)^{\frac{1}{2}} \left(\frac{10^3\,\mathrm{Hz}}{f_c}\right) \left(\frac{10\,\mathrm{Mpc}}{r}\right) \tag{1.22}$$

where f_c is the characteristic frequency and ε is the emission efficiency.

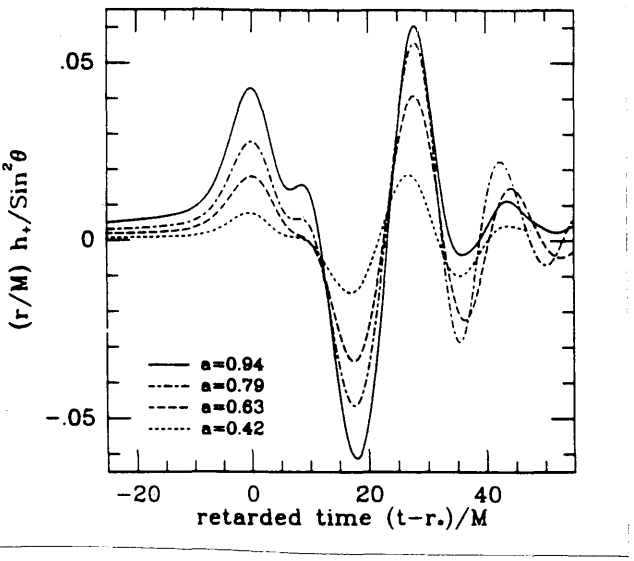

Figure 1.4: Calculated waveforms from a model of black hole formation for various angular momenta (from Stark and Piran 1985). M in solar masses. $M_\odot \simeq 4.93\mu$sec.

1.7.3 Stellar Collisions

Gravitational radiation should be produced by collisions of astrophysical objects. For example two black holes may collide head on, and as the quadrupole moment of the system changes gravitational waves are emitted. Numerical solutions of this problem [Smarr 1977] suggest that the energy carried by the emitted radiation will be of order

$$\Delta E = 0.001 Mc^2, \tag{1.23}$$

where M is the total mass of the system. If the encounter is not exactly head on then the collision should be more efficient.

At the moment of collision, shock waves may be set up in, for example, two neutron stars, and as the stars are accelerated by non-gravitational forces the energy conversion to gravitational radiation may be as high as $\Delta E = 0.01 Mc^2$. Estimates of the event rate of such sources suggest that there may be 1 such event per year out to 40Mpc with amplitudes of order $h \sim 10^{-22}$ [Clark *et al.* 1979]

1.7.4 Infall of Matter into Black Holes

The very large, $M \sim 10^5 M_\odot$, black holes which are believed to exist at the centre of galaxies will increase in size as stars are captured by the hole. As stars enter the hole the system will emit a burst of low frequency gravitational waves, typically around 10^{-4} Hz. For an object of mass m falling into a black hole of mass M the efficiency of energy conversion to gravitational radiation may be [Davis *et al.* 1971]

$$\Delta E \sim \frac{0.01 m^2 c^2}{M}. \tag{1.24}$$

By allowing for a rotating black hole, a more likely case, Detweiler was able to show [Detweiler 1977] that the efficiency could be as high as

$$\Delta E \sim \frac{1.5 m^2 c^2}{M}. \tag{1.25}$$

Although such sources are at frequencies which are too low to be seen by the large, ground based laser interferometer detectors currently being planned, they may be detectable by low frequency detectors in space which are being considered for development.

1.7.5 Neutron Starquakes

Pulsar slowdown rates often show glitches during which the pulsar period suddenly decreases. These have been interpreted as starquakes in the crust of the neutron star [Pines *et al.* 1974]. As the star slows down, centrifugal forces may no longer be able to support the crust and stresses may build up until they are relieved by a "starquake". This event may produce gravitational radiation since the quadrupole vibrational modes of the pulsar may become excited. Thorne has calculated that as much as 10^{38}J may be present in such a mode, at a frequency in the range 10^3Hz$\rightarrow$ 10^4Hz. Such a mode would have a reasonably high Q, perhaps corresponding to a damping time, t_*, of 1sec, and could produce a wave amplitude given by [Thorne 1978],

$$h \sim 10^{-23} \left(\frac{\Delta E}{10^{38}\mathrm{J}}\right) \left(\frac{3000\mathrm{Hz}}{f_g}\right) \left(\frac{1\mathrm{s}}{t_*}\right)^{\frac{1}{2}} \left(\frac{3\mathrm{kpc}}{r}\right). \tag{1.26}$$

1.7.6 Binary Stars

It is estimated that half of all stars are members of binary or multiple star systems. Since the quadrupole moment of binary systems varies with time these systems will produce periodic gravitational radiation at frequencies that are multiples of the orbital frequency [Peters and Mathews 1963], the distribution of the gravitationally emitted power being dependent on the ellipticity of the orbit. However the orbital frequency of these systems will usually lie well below the range of terrestrial detectors, typically $f < 10^{-3}$. Indeed, for main sequence stars there will be an upper limit to the orbital frequency since an orbital radius less than a few stellar radii will induce tidal instabilities in the system. Compact object binaries should have a higher upper frequency limit since the tidal disturbances will be smaller but there are considerably fewer of these systems and so they will in general lie at greater distances from the detector than normal binary systems. For this reason it is likely that doubly compact systems will be most easily seen at the end of their evolution when the gravitational wave emission is maximum (see section 1.7.9). The behaviour of such sources is, however, almost certain. For example the binary system μSco should be emitting gravitational waves with an amplitude at the earth of $h \sim 10^{-20}$ at a frequency of order 10^{-5}Hz [Douglass and Braginsky 1979]

1.7.7 Rotating Neutron Stars

A non-axisymmetric rotating neutron star or an axisymmetric one rotating about an axis that is not the symmetry axis, will produce gravitational waves. For example, a pulsar may have a rigid crust which is able to support a small bump less than a few metres high. This asymmetry would cause the emission of gravitational waves at twice the rotational frequency of the pulsar with an amplitude of order (in this simple case the efficiency, ε, is equal to the ellipticity) [Hough *et al.* 1986]

$$h = 10^{-21}\,\varepsilon \left(\frac{f}{100\text{Hz}}\right)^2 \left(\frac{10\text{kpc}}{r}\right). \tag{1.27}$$

For a pulsar whose symmetry axis precesses around its rotational axis, the gravitational radiation will be produced at frequencies which differ from the observed

pulsar frequency by the precession frequency [Zimmermann 1980]. The typical eccentricity of pulsars and hence the emitted gravitational wave amplitude is uncertain. However, since the emission of gravitational radiation will decrease the rotational energy of the pulsar, the pulsar's period will increase and it is therefore possible to place upper limits on a pulsar's eccentricity by measuring its slowdown rate by assuming that this slowdown is due entirely to the gravitational wave emission. In this way upper limits can be set for the Crab and Vela pulsars of $\varepsilon = 6 \times 10^{-4}$ and $\varepsilon = 4 \times 10^{-3}$ respectively [Thorne 1987]. It has been suggested by Zimmerman that there is evidence that the Crab and Vela pulsars possess a slight Chandler wobble due to misalignment of the symmetry and rotational axes and that more accurate values could be $\varepsilon = 5 \times 10^{-5}$ for the Vela pulsar and $\varepsilon = 2 \times 10^{-6}$ for the Crab pulsar [Zimmermann 1978], corresponding to gravitational wave amplitudes of $h = 5 \times 10^{-26}$ for the Vela pulsar and $h = 9 \times 10^{-27}$ for the Crab pulsar [Zimmermann 1978]. Indeed, the occurrence of glitches in the slowdown rate of pulsars suggests that it is quite possible that some pulsars may have a crust which is slightly irregular. In any event a large laser interferometer detector in the narrow band optical configuration which is described in section 1.8.3 should be able to place an upper limit of $h \sim 2 \times 10^{-29}$ on the gravitational wave emission of any millisecond pulsars after one year of observation.

1.7.8 Spinning, Accreting Neutron Stars

In a binary system comprising a neutron star and a non-compact star the neutron star may accrete matter from its companion. The accretion will cause the emission of X-rays which will ablate the companion further. The accretion will also act to spin up the neutron star—the orbital angular momentum of the companion is transferred to the rotational angular momentum of the neutron star—and it has been suggested that this is the mechanism for the creation of millisecond pulsars with very low spindown rates [Alpar *et al.* 1982]. When the rotational period of the star falls below a critical value of about 1ms, a value dependent on the structure and internal viscosity of the star, an instability may set up waves in the outer layers

of the star which rotate in the opposite direction to the star's rotation. These will cause the neutron star to emit gravitational waves at a frequency less than the rotational frequency [Chandrasekhar 1970, Friedmann and Schutz 1978], though it may be possible to determine this frequency from observations of the modulations of the emitted X-rays. Indeed the X-ray flux produced by the accretion gives an estimate of the gravitational wave amplitude of [Wagoner 1984]

$$h \sim 2.4 \times 10^{-28} \left(\frac{300\text{Hz}}{f}\right) \left(\frac{F_x}{10^{-17}\text{J/cm}^2/\text{sec}}\right)^{\frac{1}{2}}, \tag{1.28}$$

where F_x is the X-ray flux of the CFS unstable source, observed at the earth.

The recently discovered eclipsing millisecond pulsar PSR1957+20 is a good candidate for such a source [Fruchter *et al.* 1988]. In this binary system, about 800pc from the earth, a pulsar with a 1.6ms period is eclipsed every 9.2 hours by a large diffuse companion of only $0.2M_\odot$. This pulsar's behaviour is consistent with this theory of millisecond pulsar creation and it appears that X-ray emission from the pulsar is ablating the companion, which is believed to be a white dwarf. The material released from the companion then spins up the pulsar. Considering the difficulty in discovering rapidly fluctuating radio sources which are Doppler shifted, the discovery of this system and other millisecond pulsars, some of which are also in binary systems, such as that in the Globular cluster M4, suggests that this class of source may give interesting event rates for a laser interferometer detector.

1.7.9 Coalescing Compact Binaries

This class of source is currently believed to be the most promising for detection. Binary systems, considered above, will lose energy to gravitational radiation. This radiation reaction will cause the orbit to decay. The separation and orbital period of the stars will decrease, therefore increasing the gravitational radiation produced and reducing the eccentricity of the orbit. For a system of compact objects, this orbital decay can continue until the stars are separated by only 100km and the orbital frequency is several hundred hertz before tidal effects disrupt the behavior

of the system. In such a case a characteristic chirp gravitational waveform will be produced, the amplitude and frequency increasing with time as given by

$$\langle h \rangle = 10^{-23} m_T^{2/3} \mu f_{100}^{2/3} r_{100}^{-1}, \tag{1.29}$$

$$\tau = \frac{f}{\dot{f}} = 7.8 m_T^{-2/3} \mu^{-1} f_{100}^{8/3}, \tag{1.30}$$

where $f_{100} = f/100\text{Hz}$, m_T is the total mass of the system, μ is the reduced mass of the system, $r_{100} = r/100\text{kpc}$, and $\langle \ldots \rangle$ represents an average over the detector's sensitivity pattern and over the polarisation states of the incoming waves [Schutz 1986]. An example of such a waveform for two $3.0 M_\odot$ neutron stars is shown in Figure 1.5. Out of about 400 known pulsars only one system of two neutron stars is currently known, the binary pulsar PSR1913+16, which should coalesce in 1.4×10^8 years. Assuming that a pulsar is created in the Galaxy every 40 years and that 1 in 400 of all pulsars are in binary systems, it is possible to calculate a creation rate of one such system every 16000 years [Schutz 1988]. If it is further assumed that the number of compact binary systems is at present unchanging with time, so that the birth rate of binary systems is equal to the rate of coalescence, then out to 100Mpc, a volume containing $\sim 10^5$ galaxies, there should be 3 coalescences per year. Obviously this calculation is highly uncertain due to the small number statistics involved but the figure of 1 in 400 is in reasonable agreement with the observed proportions of binary and compact stars. Using optimal filtering techniques, which will be discussed in a later chapter, it should be possible to detect these waveforms in the output of wideband gravitational wave detectors. It has been noted by Schutz [Schutz 1986] that these sources can act as distance indicators, and that they could therefore be used to provide an accurate value for the Hubble constant. If we assume that the event rates which are currently predicted are correct and that the large scale interferometers will achieve their design sensitivity then even if the sources cannot be identified optically, the fact that galaxies cluster strongly would enable a statistical calculation to be made which should determine the Hubble constant to an accuracy of 10% in about one year.

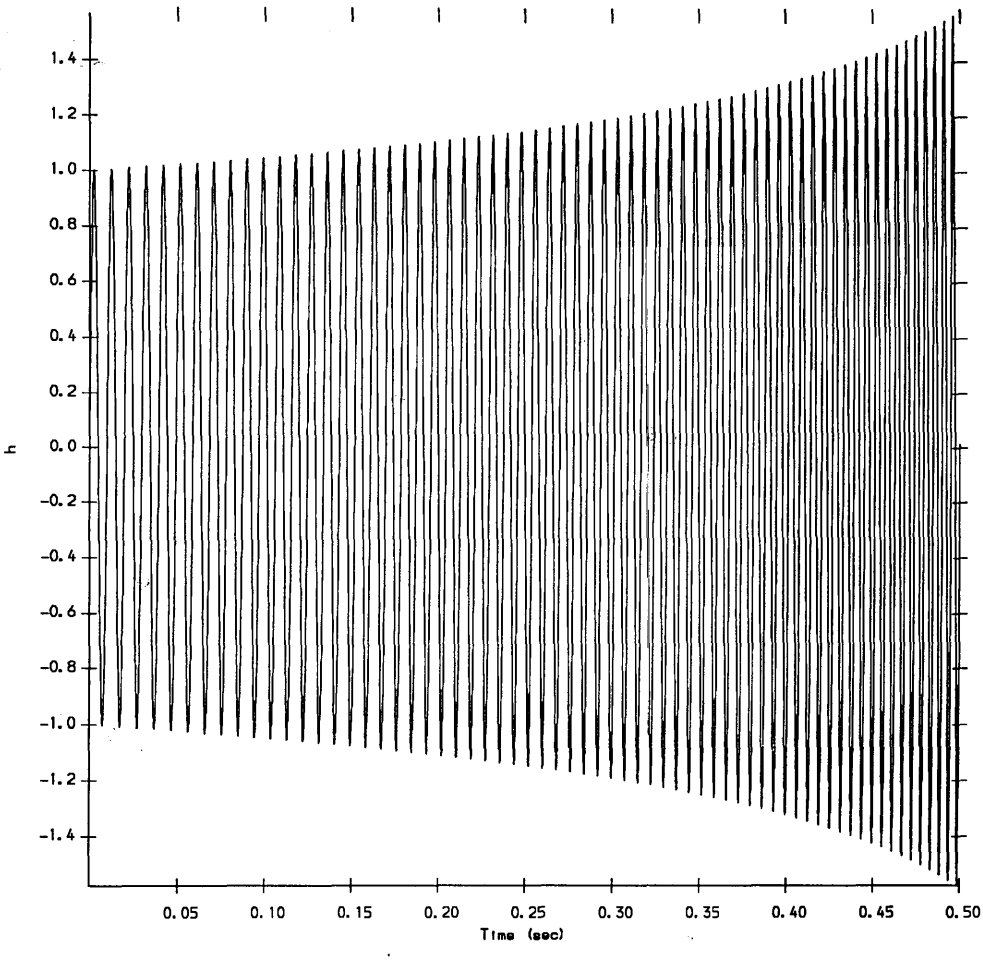

Figure 1.5: "Chirp" waveform produced by the coalescence of two 3.0 solar mass neutron stars.

1.7.10 Stochastic Background

The preceding sections have described possible gravitational wave sources which are associated with particular astrophysical objects. However there may be a background of gravitational waves which is due to the superposition of many sources or which has been produced early in the history of the universe. Such a background could only be detected by a cross-correlation of the outputs of two or more detectors. A number of possible sources of a stochastic background have been proposed. These include

Binary Stars.

For frequencies less than $\sim 10^{-2}$Hz the combined emissions of all the binary stars in the Galaxy will produce a gravitational wave background which should be detectable by planned low frequency detectors.

Primordial gravitational waves.

At times earlier than the Planck time, 10^{-43}s, gravitons were strongly coupled to matter, later becoming decoupled. If the universe went through a stage of exponential inflation, then small gravitational wave fluctuations would have been amplified, perhaps to interesting levels. Unfortunately different models produce vastly differing gravitational wave densities and so it is not possible to estimate accurately the expected gravitational wave amplitudes.

Cosmic strings.

Some Grand Unified Theories predict the existence of cosmic strings. These one dimensional objects would form closed loops which would begin to vibrate, radiating gravitational waves over a very wide frequency range. The amplitudes of this radiation which are predicted by current string theories should be detectable by many of the proposed detectors. In fact pulsar timing observations are already beginning to place interesting constraints on cosmic string theory.

Population III stars.

It has been proposed [Carr 1986] that there may have been a population of massive pregalactic stars which evolved very rapidly. Their supernovae would have contributed to a background of gravitational radiation. This class of source is highly uncertain.

1.8 Practical Detectors of Gravitational Radiation

The rich variety of possible sources of gravitational waves has motivated a considerable effort to design and construct detectors with sufficient sensitivity to detect the emitted waves. These detectors are designed to be sensitive at a variety of frequency ranges. Some of the more promising detectors will now be discussed.

1.8.1 Seismic Measurements

One of the first searches for gravitational waves was performed by monitoring the quadrupole distortions of the earth's normal modes [Forward *et al.* 1961]. A gravimeter on the surface of the earth would record changes in the local acceleration of gravity as a gravitational wave passes. A network of such detectors can monitor the earth's normal modes. A recent analysis of seismic data [Boughn and Kuhn 1984] placed a limit on the background at 0.31mHz of $h < 3 \times 10^{-15}$. An alternative method is to detect the gravitational radiation induced strain in the earth's crust by interferometrically monitoring the separation of piers, which are fixed to the ground. An analysis of the strain induced in the earth's crust by a gravitational wave was given by Dyson [Dyson 1969]. In this manner Levine and Stebbins were able to place a limit of $h < 3 \times 10^{-17}$ on the gravitational radiation from the Crab pulsar at frequencies of 30.2Hz and 60.4Hz [Levine and Stebbins 1972].

More recently, Braginsky and his collaborators have suggested using blocks of the earth's crust as detector elements [Braginsky *et al.* 1985]. They point out that the power spectrum of seismic disturbances usually has a minimum in the 0.03Hz region, due to the characteristic size of blocks of the earths crust. They claim that

by monitoring the strain in several individual blocks of the earth's crust, pulses with $h \sim 10^{-16}$ could be detected.

1.8.2 Resonant Bar Detectors

Bar detectors were initially proposed and developed by Joseph Weber at the University of Maryland [Weber 1960] and are still being used and developed by the majority of the groups which are active today. The first generation of these detectors typically consisted of a large cylinder of aluminium, around 1.5 tons, which was suspended inside a vacuum chamber in such a way that the bar was free to oscillate in its fundamental mode. The effect of a passing gravitational wave is to strain the bar; in fact at its resonant frequency the ends of a bar behave almost like free test masses and so the strain in the bar is almost equal to the strain in space that is generated by the gravitational wave. This strain was then measured using electro-mechanical transducers, for example piezoelectric transducers bonded around the middle of the bar or capacitive transducers monitoring the ends of the bar, whose output was amplified and recorded. The main noise sources in such an experiment are due to the thermal motion of the bar, thermal noise in the transducer and amplifier noise. The mean thermal energy in the bar's fundamental mode will be kT, where k is Boltzmann's constant and T is the bar's temperature. However if the mode has a high Q, that is if the coupling of the mode to its environment is low, then the mean change in the energy of the mode between measurements separated by times shorter than the damping time of the bar's mode will be approximately

$$\frac{kT\omega_m t_{meas}}{Q}, \tag{1.31}$$

where t_{meas} is the time interval between measurements and ω_m is the angular frequency of the bar's fundamental mode. Viewing the bar as a harmonic oscillator we see that the corresponding displacement of the ends of the bar is

$$x \sim \left(\frac{kTt_{meas}}{Mw_m Q}\right)^{\frac{1}{2}}. \tag{1.32}$$

The effect of thermal noise in the bar can therefore be reduced by cooling the bar, reducing its coupling to its environment—higher Q—or by increasing the mass or length of the bar. In addition the limitation to the sensitivity set by the bar's thermal motion could be reduced by reducing the time between measurements to be equal to the duration of the gravitational radiation pulses expected. However this last method requires that the transducer and amplifier have a sufficient bandwidth and increasing this bandwidth increases the effect of the transducer and amplifier noise. Therefore a bandwidth, and hence t_{meas}, can be chosen which optimises the sensitivity of the bar.

In 1969 Weber reported that two such detectors separated by 1000km were detecting coincident events with a gravitational wave amplitude in excess of 10^{-16} at the rate of about 100 such events per year and he claimed that some of these coincidences should occur due to random effects only once every 48 years [Weber 1969]. Moreover, the events appeared more likely when the bars were most sensitive to sources in the direction of the Galactic centre. Theoretical calculations were unable to explain the huge energy production implied by Weber's results. A number of groups, including one at the University of Glasgow, then built similar detectors in an attempt to confirm Weber's results. Although their detectors were equally sensitive, indeed some were more sensitive, these groups failed to detect any such events and it is now generally believed that Weber's events were not due to gravitational waves.

The detectors developed at Glasgow were different in design from those used elsewhere, using two separate masses joined by piezo-electric transducers. This arrangement increased the output from the transducer for a given motion of the bar, thereby increasing the tolerable sensor noise and the measurement bandwidth, at the expense of a reduced Q. With two such bars, operating in coincidence, a limit on the flux at the earth due to a stochastic background of less than $6.8 \times 10^3\,\mathrm{Wm^{-2}Hz^{-1}}$ was placed in the kilohertz region. A later experiment using this form of bar detector placed a limit of $h \sim 0.8^{+1.5}_{-0.8} \times 10^{-20}$ on the gravitational wave emission from the millisecond pulsar PSR 1937 + 214.

Recent improvements in the sensitivity of bar detectors have relied on cool-

ing the bars to liquid helium temperatures and using much improved transducers and amplifiers. For example, the cooled bar at Stanford University has a sensitivity to millisecond pulses of $h \sim 2 \times 10^{-18}$ [Bassan *et al.* 1983]. It consists of a 4.8tonne bar of an aluminium alloy, having a Q of order 5×10^7, cooled to 4.2K. A superconducting transducer and a SQUID are used to reduce the sensor noise.

The sensitivity of such bar detectors will eventually be limited by quantum mechanical effects, for ultimately a gravitational wave encountering a bar must produce one extra phonon in the bar's mode to be observable. This corresponds to a limit to the sensitivity of [Thorne *et al.* 1979]

$$\Delta x \sim \left(\frac{\hbar}{m\omega}\right)^{\frac{1}{2}}, \tag{1.33}$$

or a gravitational wave sensitivity of order 10^{-21} for a typical bar.

A number of authors have proposed methods for avoiding this quantum limit [Braginsky 1977, Unruh 1978, Caves *et al.* 1980] in the next generation of bar detectors. These "quantum non-demolition" or "back action evasion" detectors would operate by making a measurement of some physical observable of the bar in such a way that the measurement does not have any effect on later measurements of that observable. A number of variants of this technique have been proposed [Braginsky 1977, Caves *et al.* 1980]. For example, Braginsky has proposed a "stroboscopic" method of back action evasion. If a short, accurate, measurement of the position of a harmonic oscillator is made then the momentum of the oscillator and hence the amplitude of the oscillation will be disturbed. However the frequency of the oscillation is unchanged and so the oscillator must return to the measured position after one cycle. Therefore the amplitude and phase of the harmonic oscillator have been disturbed in a way that has no effect on successive measurements a half integer number of oscillations later. A gravitational wave, with suitable polarisation and phase, would however change the observed amplitude, while two detectors in parallel could determine both the amplitude and phase of the gravitational wave.

Much effort is currently being invested in the development of such detectors. However the experimental difficulties are formidable, and the effects of thermal

noise and sensor noise must also be reduced below the levels required for such measurements. Therefore it appears that it will be difficult to improve the sensitivity of resonant bar detectors to the levels where astrophysical gravitational radiation is observed regularly.

1.8.3 Laser Interferometers

The most promising ground based detectors are laser interferometers. The first interferometer detector was designed and constructed by Forward and his collaborators at the Hughes Research labs in Malibu, achieving a sensitivity of $2\times10^{-16}/\sqrt{\text{Hz}}$ at frequencies around 1kHz [Forward 1978]. Since then a number of other groups have continued the development of these detectors.

The principle behind such detectors is relatively simple. If a mass is suspended as a pendulum, then for fast displacements in the horizontal plane, i.e. those at frequencies large compared to the frequency of the pendulum, the mass essentially behaves as if it were freely falling. The separation of two such nearly free pendulums will therefore vary depending on the presence of gravitational waves. In principle, if a sufficiently accurate reference length is available, measurement of the separation of these two pendulums is sufficient to detect gravitational waves. However, the estimates of the strength of astrophysical sources suggest that the separation will vary by less than about 1 part in 10^{22} over millisecond timescales which places extreme demands on the stability of the reference length. For this reason practical laser interferometers will use two orthogonal arms and compare the relative lengths of these arms using a Michelson interferometer arrangement. Since a gravitational wave travelling in a direction perpendicular to the plane of the two arms, and with suitable polarisation, will extend one arm while shortening the other, this method has the advantage of doubling the effect of the wave compared to that seen by monitoring one arm relative to an absolute reference length, so that $dx = hl$, where dx is the relative length difference between the two arms, h is the gravitational wave amplitude and l is the length of the detector's arms. The idealised arrangement in Figure 1.6 has three test masses forming two

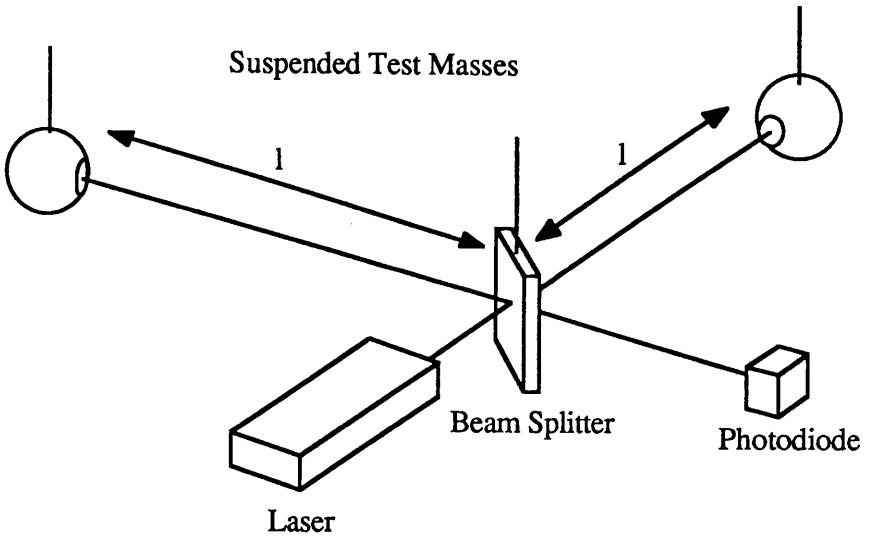

Figure 1.6: Schematic diagram of a Laser Interferometer Gravitational Wave Detector.

arms of the interferometer, each arm of length l. The detector's sensitivity may be limited by noise sources which disturb the masses, variations in the apparent optical path and by limitations in the measurement of this apparent separation. Some of the noise sources which may directly affect the position of the masses are thermal noise in the masses and pendulums, Brownian motion caused by the residual gases in the vacuum system, low frequency seismic noise and radiation pressure from the laser light acting on the masses. The optical path may be changed by refractive index fluctuations and by light being scattered within the interferometer. Examples of measurement limitations are shot noise in the photocurrent induced by the detected light, laser intensity and frequency fluctuations and beam geometry fluctuations. These noise sources will be discussed more fully in the following chapter, and it will be shown that for the detectors currently being planned the dominant noise source above 100Hz should be due to the photocurrent shot noise and that the signal to noise ratio of the detector is maximised by arranging that measurement is made of a dark fringe. The output of this detector is therefore of the form

$$I = \frac{I_0}{2}(1 - \cos\phi) \tag{1.34}$$

where $\phi - \pi$ is the phase difference between the two interfering beams and I_0 is the input laser intensity. The performance of an interferometer detector which is shot noise limited may be improved by using more sophisticated optical arrangements which recycle the light which would otherwise return towards the laser and be wasted. These techniques will be discussed in later sections.

It can be seen that, since relative length changes in the two arms are being measured and a gravitational wave signal produces a strain in space, the sensitivity could be improved by increasing the arm length up to a maximum of $\lambda_g/4$,where λ_g is the wavelength of the gravitational wave being searched for. If the length of the arm is increased further, then since the light spends more than half a gravitational wave period in the arm and a mean gravitational wave amplitude over the light travel time is measured, the sensitivity begins to fall again. For 1 kHz gravitational wave signals the optimum arm length is $\sim$ 100km. Since the two arms of the interferometer must be housed in a vacuum, such a detector would be prohibitively expensive to construct and site on the earth, but it is possible to simulate a long arm length by many passes of a shorter arm. Two forms of laser interferometer which use optical arrangements that effectively increase the arm length have been proposed and are currently being developed; Delay Line and Fabry-Perot detectors.

Delay Line Detectors

Delay line detectors were initially developed by Weiss [Weiss 1972] and are currently being developed on a 30 metre prototype at the Max-Planck-Institut für Quantenoptik in Garching. Initially the Glasgow group also tested a 1m prototype. In such a detector the laser light entering an arm is reflected back and forth many times between mirrors fixed to the test masses before it is allowed to return towards the beam-splitter as shown in Figure 1.7. A small relative displacement δx of the test masses therefore causes a change in the optical path which is $2N\delta x$, where N is the number of reflections of the beam from the end test mass.

In operation, the phase difference between the returning beams is measured using an RF phase modulation technique. The interferometer output is main-

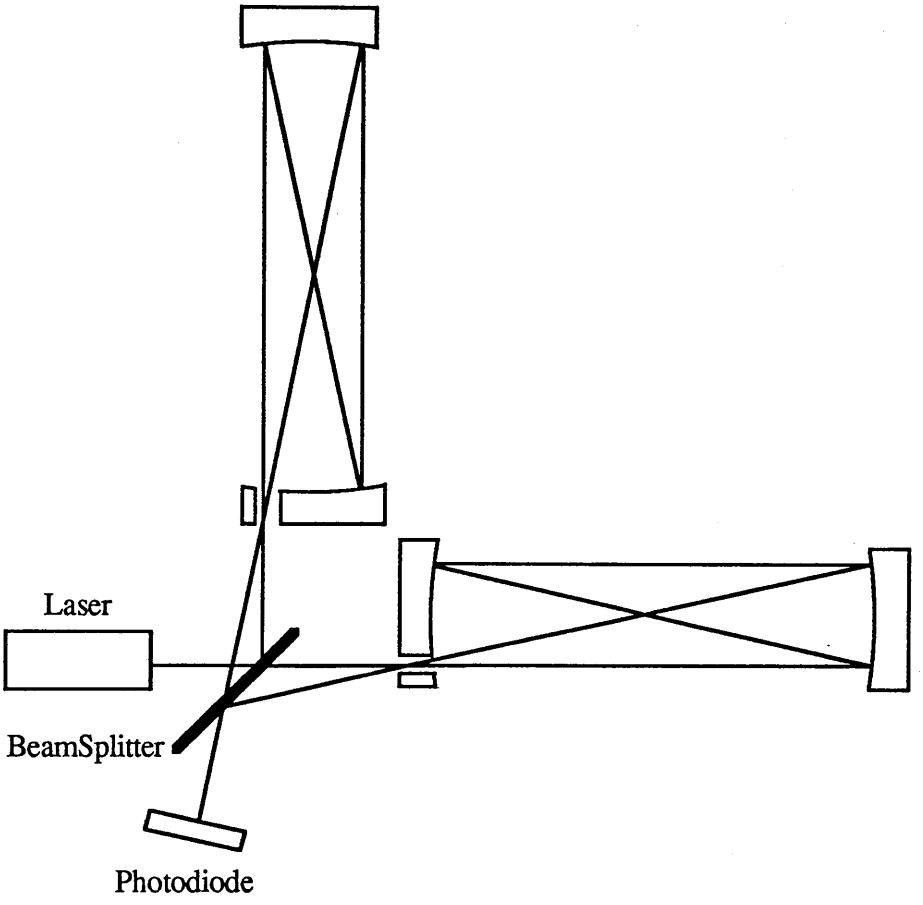

Figure 1.7: Delay Line arrangement for increasing the effective arm length.

tained on a null by adjusting the phase of the light from one of the arms using an electro-optic phase modulator. The correction signal which is fed back to the modulator therefore gives a direct measure of the relative lengths of the two arms. Such a feedback system can be made broadband and a high dynamic range can be achieved allowing the interferometer to operate stably over reasonably long timescales. For example it has been reported that the 30 metre prototype at Garching will remain locked for periods up to an hour and reacquisition is automatic [Shoemaker *et al.* 1988].

In principle such a detector can be made insensitive to frequency fluctuations of the laser light, since no phase difference between the interfering light beams will be generated if the lengths of the two arms are equal. However, if there is scattering of light within the arms of the delay line, so that light reaching the photodiode has a component that has travelled by a shorter or longer path, the frequency fluctuations can place a severe limitation on the sensitivity of the detector. This effect was identified on the Garching prototype and on the trial 1m detector at Glasgow. Techniques to reduce the effect of this noise source by frequency stabilising the laser light and by phase modulating the light have been developed, both at Garching [Schilling *et al.* 1981] and at MIT [Dewey 1986]. The group at Glasgow decided instead to develop Fabry-Perot detectors in order to reduce the problems associated with scattered light.

Fabry-Perot Detectors

The optical cavity system was initially developed by the Glasgow group on a 10 metre prototype and has since been implemented at Caltech on a 40 metre system. Each arm of the interferometer consists of a Fabry-Perot cavity which is held near resonance with the input laser light (see Figure 1.8). Small fluctuations in the cavity length produce large phase changes in the light leaking from the cavity so that interference of the light from the two arms will produce an output which varies with the relative lengths of the arms. Such an arrangement requires rather more sophisticated control systems than the delay line system, as the cavities must both

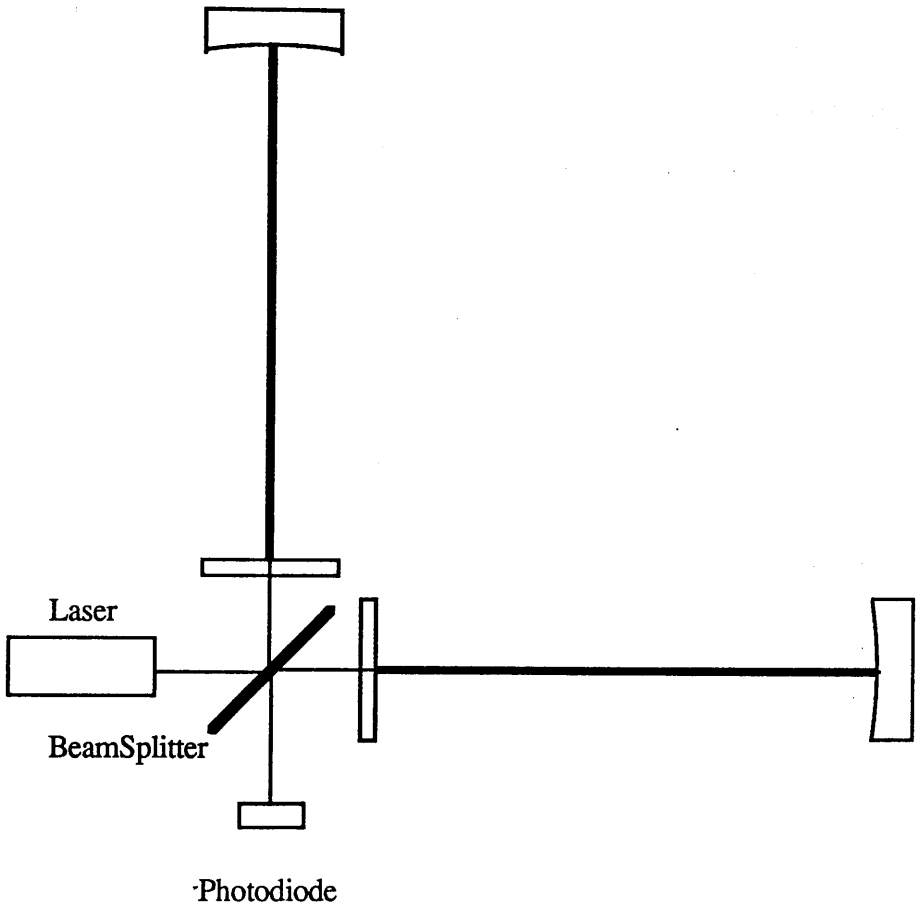

Figure 1.8: Fabry-Perot arrangement for increasing the effective arm length.

be held near resonance with the laser light in order to retain a linear length to phase response and to reduce the effect of intensity noise—this will be discussed more fully in the next chapter. In practice a fraction of the light returning from each cavity would be split off and directed to a monitor photodiode to keep the arms near resonance. Since the frequency of commercial lasers typically varies by $10^3 \mathrm{Hz}/\sqrt{\mathrm{Hz}}$ around 1000 Hz and, for example, the laser frequency should be stable to $10^{-5}\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ in order to remain sufficiently close to resonance with the cavities of the 10 metre prototype, the laser frequency is locked to one arm of the interferometer and the other arm's length is adjusted to keep it near resonance with this highly stabilised laser light. (In Delay-Line systems the null output condition can be maintained through the use of electro-optic modulators. In a Fabry-Perot system the length one of the cavities must be directly controlled to maintain resonance, since losses associated with a modulator in the cavity would severely limit the sensitivity of the detector). The common mode signals due to any remaining laser frequency fluctuations may be removed by interfering the light from the two arms or by monitoring the cavities separately and differencing these two signals to remove the common mode component.

Such a system has the advantage that light which is scattered in the cavity is much less likely to re-enter the system, since the beams lie on top of each other. In addition the mirrors can be kept considerably smaller than than those required by delay line systems. As mentioned before, Fabry-Perot detectors require rather more sophisticated servo systems than delay line detectors to maintain optimum performance, however work on the Glasgow prototype has shown that systems of suitably high performance can be implemented.

Standard Recycling

To obtain optimum performance, interferometers must be operated on a null fringe. If low loss mirrors are used then most of the output light is returned towards the laser, and is wasted. Drever has suggested that if a mirror is placed between the laser and the beam-splitter, as shown in Figure 1.9, then, by careful phase

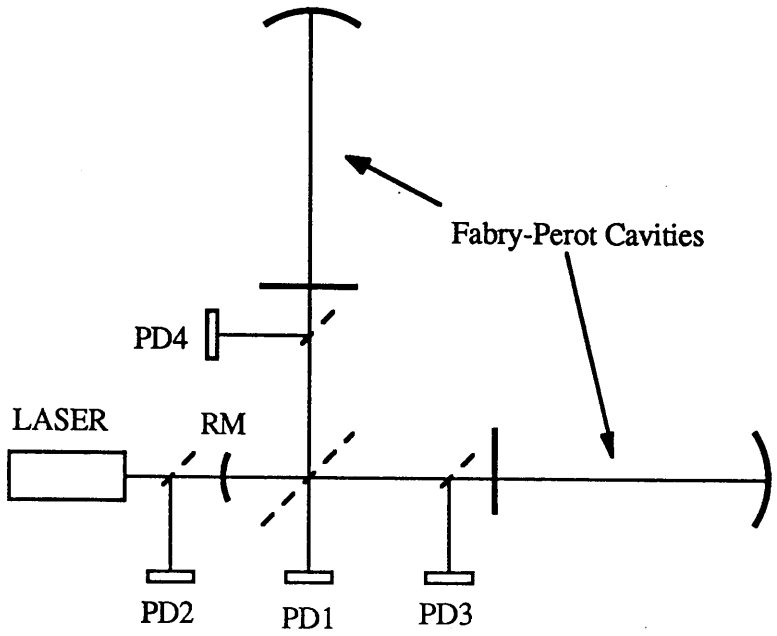

Figure 1.9: Possible arrangement for standard recycling. RM is the recycling mirror. PD1 is a photodiode which monitors the output signal. PD2 to PD4 are used to keep the cavities on resonance. Dotted lines represent beamsplitters.

matching, this light can be made to enter the interferometer again, thus increasing the effective laser input power and so reducing the shot noise limit [Drever 1983]. Losses in the optical components will determine the effective number of passes through the interferometer the light may make. The interferometer then becomes an optical cavity which can be held on resonance. If the main losses are due to light leaking through or absorbed by the end mirrors and these mirrors have a power loss coefficient A^2, then the sensitivity is improved by a factor [Meers 1988]

$$\left(\frac{\pi \nu_g}{4\nu_0 A^2}\right)^{-\frac{1}{2}}, \tag{1.35}$$

where ν_g is the gravitational wave frequency and ν_0 is the free spectral range of the optical cavities.

Resonant Recycling

Another optical scheme proposed by Drever [Drever 1983] allows a broadband laser interferometer to be made narrowband, with a corresponding increase in the sensitivity. In the presence of a periodic gravitational wave the lengths of the arms

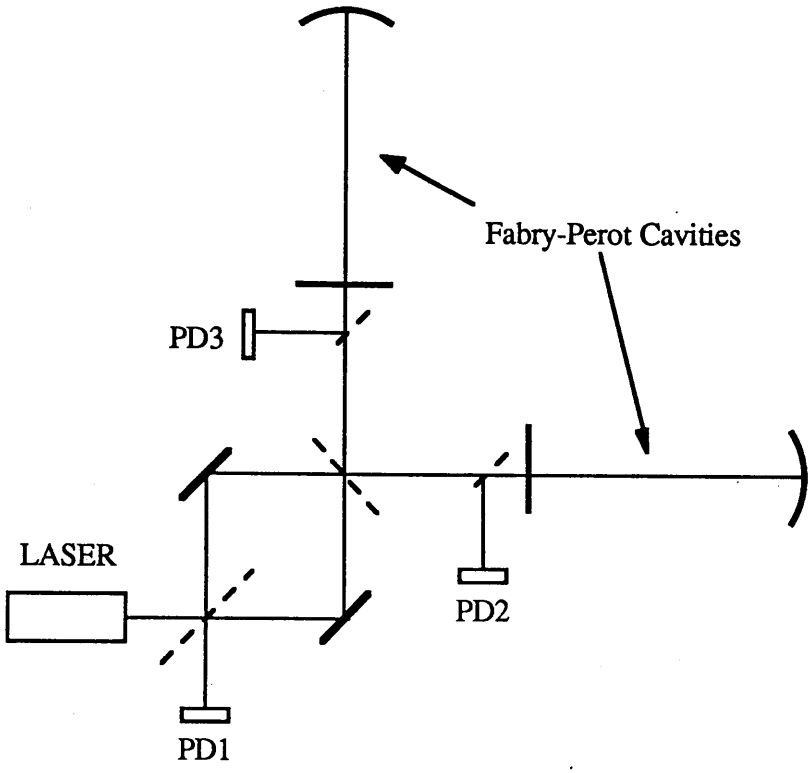

Figure 1.10: Possible optical arrangement for resonant recycling. PD1 is the photodiode which monitors the gravitational wave signal. PD2 and PD3 are used to keep the cavities on resonance. Dotted lines represent beamsplitters.

of the detector will oscillate. For both the Fabry-Perot and Delay Line forms, light entering an arm will effectively remain in the arm for some time, t_{stor}, before returning towards the beamsplitter. This storage time can be adjusted by varying the number of bounces in the delay-line case or by adjusting the reflectivity of the cavity mirrors in the Fabry-Perot case. Therefore if two light beams enter the arms just as one becomes longer than average and the other shorter, then the light which entered the shorter arm will return to the beamsplitter before the light which entered the other, longer, arm. If most of the light returning from each arm can then be made to enter the other arm, and the storage time is equal to one half of a gravitational wave period, then the phase delay between these two beams will increase linearly with time. The total integration time will be limited by the light losses within the detector. A suitable optical configuration is shown in Figure 1.10.

The improvement in the sensitivity achieved should be a factor [Meers 1988],

$$\frac{\pi \nu_g}{4\nu_0 A^2}, \tag{1.36}$$

while the effective bandwidth around ν_g is given by [Meers 1988]

$$\frac{2\nu_0 A^2}{\pi}, \tag{1.37}$$

where the terms are as defined in the previous section.

1.8.4 Laser Interferometers in Space

In order to avoid the low frequency, seismic, limit on the sensitivity of earth based laser interferometers, a number of authors have suggested and analysed designs for interferometers in space [Weiss *et al.* 1976, Faller *et al.* 1985]. Such a detector might consist of 3 masses in orbit around the sun and may have arm lengths of 10^6km and would therefore be sensitive to gravitational waves with frequencies below 0.1 Hz. Disturbances in the orbits of the masses, caused, for example, by the solar wind, will degrade the sensitivity at frequencies below 10^{-4}Hz. Sensitivities in this frequency range may be of order $10^{-20}/\sqrt{\text{Hz}}$.

1.8.5 Pulsar Timing

Certain pulsars, which are believed to have been spun-up by the accretion of material ablated from a companion star and therefore highly symmetric, have been discovered which display extremely stable pulse periods of the order of a few milliseconds. These pulsars therefore constitute a distant, highly stable clock and observation of the variation of the pulsar period relative to atomic clocks on the earth can set upper limits on the amplitude of very low frequency gravitational waves which are passing either the pulsar or the earth. Using this technique the limit placed on the gravitational wave amplitude for waves with a period of order 1 year is $3 \times 10^{-14}/\sqrt{\text{Hz}}$ [Davis *et al.* 1985]. This is presently limited by the accuracy of atomic clocks; the pulsars display periods which are more stable than the clocks.

1.8.6 Doppler Tracking of Spacecraft

Just as the signals from pulsars would be Doppler shifted by gravitational waves, so the signals received and transmitted by interplanetary space probes will be Doppler shifted. If a gravitational radiation pulse passing through the solar system passes the earth then transmissions from a space probe which are being received and signals which are being transmitted will be Doppler shifted. As the wave travels on and passes the probe then radio waves being received and retransmitted will be Doppler shifted. The receiver on the earth will therefore see a signal whose frequency varies at three distinct times. Indeed special signal processing can search for such waveforms and some information about the wave's direction of propagation can be obtained by the relative times of the three events. Using this technique with the Voyager I spacecraft, Hellings *et al.* have been able to place a limit on the gravitational wave density in the frequency range $2 \times 10^{-4}\text{Hz} < f < 10^{-3}\text{Hz}$ of $3 \times 10^{-14}/\sqrt{\text{Hz}}$ and found no evidence for bursts with periods of order 1000 seconds with amplitudes greater than 3×10^{-14}. The technique is currently limited by timing errors in the clock used as a reference for transmission and reception and by scintillation effects in the interplanetary medium and in the troposphere. These effects may be reduced on future experiments by operating at two different radio frequencies [Smarr *et al.* 1983].

1.8.7 Skyhook

A detector which has recently been proposed by Braginsky and Thorne is Skyhook. This would consist of two masses joined by a thin cable with a spring in the middle. The arrangement would orbit the earth and be stretched radially by the earth's gravitational field. Gravitational waves would cause variations in the extension of the spring which could be measured with a suitable sensor. The frequency range of such a detector would be 0.1-0.01 Hz. The noise behaviour of such a detector has been analysed by Braginsky and Thorne who have calculated that it would have a sensitivity of the order of $3 \times 10^{-17}/f^{\frac{1}{2}}/\sqrt{\text{Hz}}$, which is similar to that which might be obtained by monitoring blocks of the earth's crust [Braginsky and Thorne 1985].

1.9 Conclusions

It has been shown in this chapter that gravitational radiation should be emitted by a wide range of known astrophysical objects. The sources which have been predicted, and which should be detectable by the large scale laser interferometers, fall mainly into two classes; those which produce short, millisecond bursts, and those which emit nearly monochromatic waves. In addition coalescing binaries will produce a characteristic gravitational wave signal. It is necessary then to optimise detectors for these sources and to develop suitable data analysis algorithms for their detection. It is quite possible, however, that sources of significant levels of gravitational radiation exist which have not be observed directly by electromagnetic detectors. Gravitational wave detectors may therefore prove invaluable astrophysical tools for the investigation of the bulk motions of matter in the universe.

Chapter 2

Some Aspects of Interferometer Design

2.1 Introduction

The principles of laser interferometric gravitational wave detectors have been introduced in Chapter 1. The design and operation of the Glasgow 10 metre prototype will now be described. The limitations to the sensitivity of the prototype, set by various noise sources, will be considered along with the implications for the planned long baseline detectors. It will be seen that some of the limitations are of a fundamental nature and limit the sensitivity which can ultimately be achieved, while others are practical limitations which can be avoided by careful design, and it will be shown that many of the noise sources have an effect which is to displace the position of the test masses. Since gravitational waves are a strain in space, the importance of such noise sources can be reduced by increasing the separation between the test masses. This is the motivation behind the planned long baseline interferometers with arm lengths of 3 km or more.

2.2 The Prototype Detector

The Glasgow prototype is a Fabry-Perot design using two perpendicular 10 metre optical cavities. The laser is frequency locked to one cavity, while the second cavity is locked to the stabilised laser frequency. Therefore one cavity is used to measure the length of the other. Four vacuum tanks sit at the corners of a square with sides of 10 metres. The tanks are linked by four vacuum pipes which lie along the

sides of the square. The two test masses which hold the cavity end mirrors are suspended in two tanks at diametrically opposite corners of the square. The beam splitter and the test masses holding the cavity input mirrors are suspended in one of the remaining tanks; at present the fourth tank is unused. The vacuum system is supported by a steel framework which rests on four stacks of alternate layers of lead and rubber placed on the ground. These stacks provide isolation from ground motion at high frequencies. The beam splitter and the quarter wave plates are mounted on an aluminium plate which is suspended from an aluminium top plate. The top plate rests on a stack of alternate layers of rubber and lead. This stack has an upper normal mode at a frequency of order 10 Hz and provides isolation of the suspension points from motion of the vacuum chamber at frequencies greater than this. The cavity input mirrors are high reflectivity, multi-layer dielectric mirrors which are optically contacted to cylinders of fused silica which are also suspended as pendulums from the aluminium top plate with a period of around 1 second. Control of the orientation of these masses and damping of motions at the resonant frequency of the pendulum is achieved by movement of the points of suspension of the suspension wires. Since the fused silica is not of optical quality, each mass has a hole drilled along its axis through which the light may pass. By separately suspending the cavity mirrors the effects of resonances in the large central mass are reduced. The two cavity end mirrors are also multi-layer dielectric mirrors and they are fixed to the fused silica test masses and suspended in a similar manner to the input mirrors from plates which are supported by stacks. One of these end masses, that in the primary cavity, is damped to the ground and its orientation is controlled by the system described in Chapter 3, while the orientation of the secondary cavity end mass is adjusted in the same manner as the two cavity input mirrors, that is by adjusting the points of suspension. The pointing of each of the cavity test masses can be adjusted by the operator to achieve optimum alignment.

The laser is a commercial argon ion laser which has been extensively modified at Glasgow and which provides continuous, single mode light with a wavelength of 514 nm. The laser cavity mirrors are mounted on a resonator which is constructed of INVAR, an alloy with a low coefficient of thermal expansion, and which is

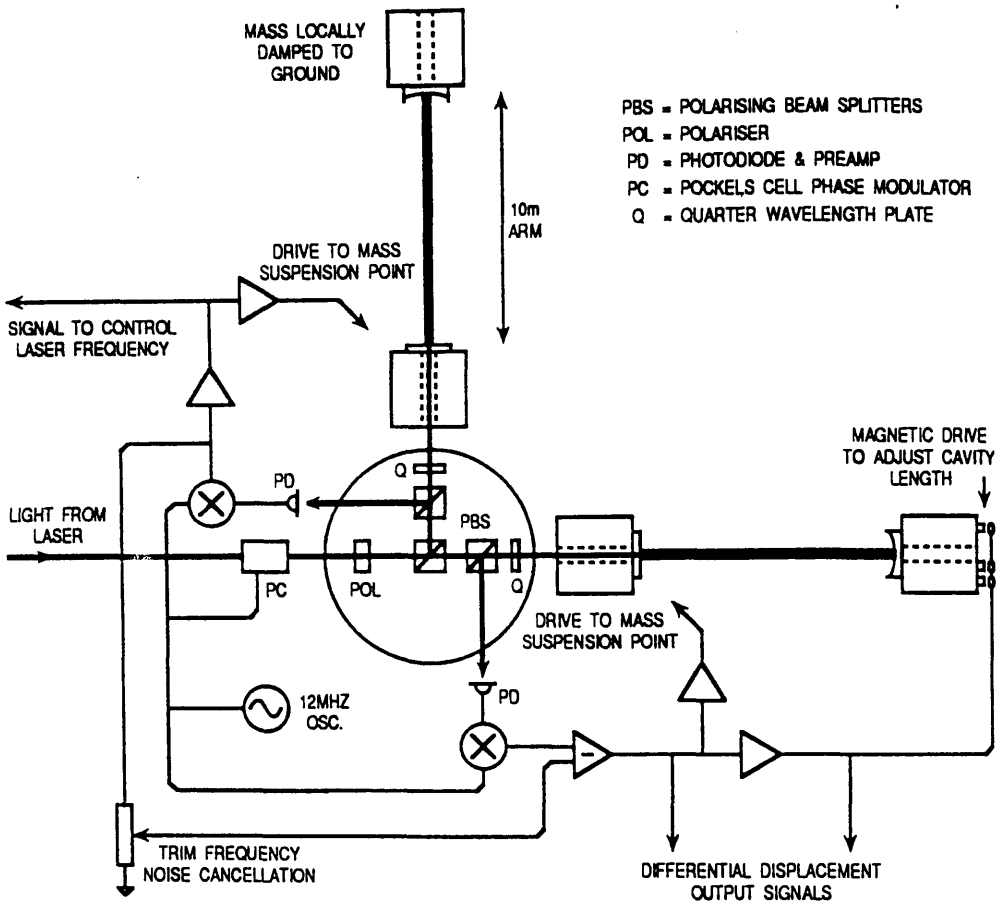
MASS LOCALLY DAMPED TO GROUND
10m ARM
PBS = POLARISING BEAM SPLITTERS
POL = POLARISER
PD = PHOTODIODE & PREAMP
PC = POCKELS CELL PHASE MODULATOR
Q = QUARTER WAVELENGTH PLATE
DRIVE TO MASS SUSPENSION POINT
SIGNAL TO CONTROL LASER FREQUENCY
MAGNETIC DRIVE TO ADJUST CAVITY LENGTH
PD
Q
LIGHT FROM LASER
PBS
PC
POL
Q
DRIVE TO MASS SUSPENSION POINT
PD
12MHZ OSC.
TRIM FREQUENCY NOISE CANCELLATION
DIFFERENTIAL DISPLACEMENT OUTPUT SIGNALS

isolated from the laser tube to reduce the coupling of vibrations produced by the motion of the laser cooling water. The cavity mirrors are mounted on piezo-electric transducers which allow control of the cavity length and hence the laser frequency. Light emerging from the laser is directed through optical isolators to prevent light reflected from subsequent optical components re-entering the laser cavity, and then passes through auxiliary optical components which are used to phase modulate the light and to apply fast phase corrections to the light to control the frequency of the emerging light. The light then passes through a monomode optical fibre which is used to suppress lateral motion of the laser beam, through mode matching optics and then enters the vacuum system through a porthole in the tank holding the beam splitter. Before entering the vacuum tank, some of the light is split off and directed towards a photodiode whose output signal is used to intensity stabilise the laser light entering the detector. When the light enters the tank it is passed through a polariser so that the light is linearly polarised with a well defined direction of polarisation before it reaches the beamsplitter. The beam splitter is a polarising beam splitter and it is arranged that 50% of the light is directed into each cavity. The cavity input mirrors are flat while the end mirrors are spherical with a radius of curvature approximately equal to 1.5 times the cavity length. This cavity configuration reduces the effects of coupling to high order radial modes and since the waist of the cavity lies at the flat input mirror allows smaller optical components, such as the beamsplitter, to be used. Before the light enters a cavity it is passed through a quarter wave plate. The light in each cavity is therefore circularly polarised. Light returning from a cavity passes once more through the quarter wave plate and returns towards the beamsplitter as linearly polarised light at 90° to the incoming light to that cavity. If the light from each cavity was split at the beamsplitter and these components interfered, the intensities of the light leaving the two output ports of the polarising beam splitter would depend on the phase difference between the two beams and from this a measure of the relative length difference between the two arms could be obtained.

At present the light from the two cavities is not recombined and interfered to obtain the gravity wave signal directly. Instead, the light from each cavity is

directed to its own photodiode using polarising beam splitters. The two photodiodes, which have a high frequency capability in order to detect the RF modulation signals, monitor the resonance condition of the cavities. The frequency of the laser is controlled to maintain resonance with one cavity, the primary cavity, using an RF reflection locking scheme, [Drever *et al.* 1985], and is therefore extremely stable over short timescales. Low frequency signals due to the pendulum motion of the cavity masses are fed back to the cavity mass holding the primary cavity input mirror. This reduces slow drifts in the laser frequency which would otherwise be induced and which would have to be tracked by the secondary cavity. In this reflection locking scheme, the light is phase modulated before reaching the cavity, producing sidebands on the laser light. These sidebands are not resonant with the cavity and so are suppressed. The light returning from the cavity consists therefore of two components; a directly reflected part which retains the sidebands and a cavity leakage part with the sidebands removed. The carrier frequency components of these two signals are out of phase when the cavity and laser light are resonant. The detected light will have a component at the modulation frequency whose amplitude is a direct measure of the phase difference between the light from the laser and the light returning from the cavity. Demodulation of this signal therefore gives a signal representing the cavity offset. The length of the other cavity, the secondary cavity, is adjusted to maintain resonance with the stabilised laser light using the same reflection locking scheme.

This technique of monitoring the cavities separately is simpler to implement than recombining and interfering the light from the two cavities, and the loss in sensitivity can ideally be only a factor of $\sqrt{2}$, due to increased photon shot noise caused by measuring two signals independently rather than measuring their sum. Laser frequency fluctuations appear as a common mode signal in the outputs of the two cavities and these can be reduced by adding the two cavity outputs with a suitable phase correction.

With this arrangement, a sensitivity of $1.7 \times 10^{-19}/\sqrt{\text{Hz}}$ has been achieved in a bandwidth of 1 kHz around 1 kHz. For millisecond pulses, this is approaching the sensitivity of the best of the bar detectors.

2.3 Fundamental Limitations to the Sensitivity of Interferometric Detectors

The noise sources which will be described in this section are of a fundamental nature and are set by the configuration of the experiment, for example the masses of the test masses and the input laser power. These noise sources set limits on such parameters as the laser input power required and the masses of the test masses.

2.3.1 Photon Shot Noise

The most important noise source in the proposed detectors is the shot noise introduced by photon counting statistics. It is simple to show that the effect of shot noise in the simple Michelson interferometer discussed in Chapter 1 is minimised by operating on a dark fringe, for if the output of this detector is of the form

$$I = \frac{I_0}{2}(1 - \cos\phi), \tag{2.1}$$

where $\phi - \pi$ is the phase difference between the two interfering beams and I_0 is the input laser power, then the number of photo-electrons produced in a time τ by light with frequency ν striking a photodiode with efficiency ϵ is

$$N = \frac{\tau I_0 \epsilon}{4\pi\hbar\nu}(1 - \cos\phi). \tag{2.2}$$

Therefore the change in the number of photo-electrons corresponding to a phase change $d\phi$ is

$$dN = \frac{\tau I_0 \epsilon}{4\pi\hbar\nu}\sin\phi\, d\phi, \tag{2.3}$$

and for a simple Michelson system interferometer this phase change relates to relative length differences in the two arms as

$$d\phi = \frac{4\pi dx}{\lambda}, \tag{2.4}$$

where dx is the difference in the optical path of the two beams and λ is the light's wavelength. Since the photo-electrons produced by the photodiode will

obey Poisson statistics, the noise due to the statistical fluctuations is

$$dN' = \sqrt{\frac{\tau I_0 \epsilon}{4\pi\hbar\nu}(1-\cos\phi)} \tag{2.5}$$

so that the corresponding limit to the positional fluctuations that can be detected over a time τ is set by the condition $dN > dN'$, that is,

$$dx > \left(\frac{\hbar\lambda c}{8\pi I_0 \epsilon \tau \cos^2(\phi/2)}\right)^{\frac{1}{2}}. \tag{2.6}$$

Therefore the effect of shot noise can be minimised by operating with $\phi = 0$, that is on a dark fringe, giving the sensitivity to pulses of duration τ of

$$dx > \left(\frac{\hbar\lambda c}{8\pi I_0 \epsilon \tau}\right)^{\frac{1}{2}}. \tag{2.7}$$

For a non-ideal detector, however, two other effects must be considered to obtain the correct shot noise limit. The first is that losses in the mirrors will result in a fringe visibility of less than one and hence a reduction in the sensitivity of the detector, while the second effect is that the increase in the effective arm length due to multiple passes of the laser light within the cavity or delay line improves the sensitivity. When these effects are taken into consideration the limit on the sensitivity of a Fabry-Perot interferometer, with identical cavities, set by shot noise is [Kerr 1986]

$$dx > \left(\frac{\hbar\pi\lambda c}{8\tau\epsilon F^2 I_0}\right)^{\frac{1}{2}} \times f(V, \beta), \tag{2.8}$$

where $f(V, \beta)$ is a factor which depends on the visibility, V, and on the modulation index, β, of the laser light and where

$$F = \frac{\pi\sqrt{r_1 r_2}}{(1 - r_1 r_2)} \tag{2.9}$$

is the finesse of the cavities, with r_1 and r_2 the amplitude reflection coefficients of the two mirrors forming a cavity. It is also assumed that absorption in the

input mirrors is negligible. The factor $f(V, \beta)$ should be no more than about four and should approach unity for high visibilities. As the finesse is increased the storage time of the cavity also increases. Since the sensitivity of the detector is degraded for pulses shorter than the cavity storage time, a value of the finesse should therefore be chosen which minimises photon shot noise for pulses longer than some minimum pulse time.

From Equation 2.8 it can also be seen that the effects of shot noise can be reduced by operating with higher light powers. Indeed the large detectors which are planned will use lasers which provide up to 200 W of continuous single mode light. This is limited by practical considerations such as heating of mirrors and other optical components due to absorption. It should also be noted that methods of improving the sensitivity beyond that set by shot noise have been proposed in which squeezed light is injected into the output port of the interferometer [Caves 1981]. This technique would decrease the error due to photon counting statistics at the expense of increased radiation pressure fluctuations.

2.3.2 Heisenberg Uncertainty Principle

One obvious limitation to the sensitivity of the detector is set by the uncertainty principle. Any attempt to measure the position of the mass to an accuracy of Δx will result in the momentum of the mass being uncertain to greater than $\Delta p \geq \hbar/\Delta x$. Therefore the limit to the sensitivity of a detector to pulses of duration τ set by the uncertainty principle is [Caves 1980]

$$dx > \left(\frac{2\tau\hbar}{m}\right)^{\frac{1}{2}}, \tag{2.10}$$

where the pulse corresponds to a bandwidth of $1/\tau$. In principle, quantum non-demolition techniques, similar to those discussed in Chapter 1 for bar detectors, could be used to increase the sensitivity beyond this limit, but as yet no scheme has been proposed which would be practicable.

2.3.3 Radiation Pressure

As the light power is increased the effects of shot noise decrease and the position of the mass is more accurately determined. However the effect of the random fluctuations in the number of photons striking the mirror disturbs the momentum of the mass more and more as the light power is increased. It is possible to regard the beam splitter as scattering each incident photon independently into one or other of the two arms [Edelstein *et al.* 1978], producing an anticorrelated effect between the two arms. The differential length change induced on the cavities is therefore equivalent to that produced by two independent light sources illuminating the cavities. Caves has shown [Caves 1980] that for some light power

$$I_{opt} \sim \frac{mc\lambda}{2\tau^2}, \tag{2.11}$$

the effects of shot noise and radiation pressure in a simple Michelson interferometer can be minimised, and the limit to the sensitivity set by the Uncertainty Principle can be reached. For a Fabry-Perot system the optimum input light power is given by [Kerr 1986]

$$I_{opt} \sim \frac{mc\lambda\pi}{4F^2\tau^2}, \tag{2.12}$$

which is roughly equal to 18 W for the current 10 metre prototype, and should be about 1 MW for the planned long baseline detectors. However, as mentioned above, practical limitations will limit the input light power to less than about 200 W and so the intra-cavity light power with recycling will be only of order 20 kW and radiation pressure noise should not be important in the proposed long baseline detectors unless the limit set by the uncertainty principle can be reached by the use of squeezed light.

2.4 Practical Limitations to the Sensitivity of Interferometric Detectors

This section will describe a number of noise sources which can in principle be reduced to arbitrarily small levels by careful design. However some of the design

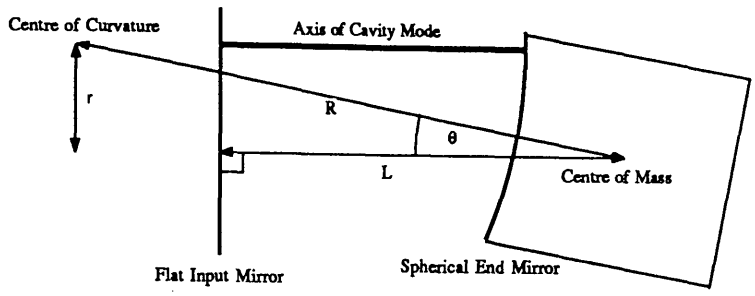

Figure 2.1: The effect of rotation of the masses on the cavity length.

features may be very difficult to achieve in a practical system and so it is important to calculate the magnitude of the effects and how they can be minimised most effectively.

2.4.1 Test Mass Orientation Fluctuations and Beam Geometry Fluctuations

The optical cavities used in the current prototype and to be used in the larger detectors have a flat input mirror and an end mirror with a radius of curvature equal to one and a half times the cavity length. If the orientation of the mirrors changes, then it is simple to show that the length of the cavity axis will also change. Consideration of Figure 2.1 shows that if the cavity axis and the centre of curvature of the mirror mounted on the end test mass are misaligned by a distance r, then the cavity length changes by

$$\delta l = r\delta\theta - \frac{r\,\theta^2\,\delta\theta}{3!} + \text{higher order terms}, \tag{2.13}$$

$$\sim r\delta\theta \tag{2.14}$$

due to variations θ in the orientation of the end test mass with respect to the test mass holding the flat, input mirror. This places limits on the fast changes in orientation of the test masses which can be allowed of

$$\delta\theta \sim 3\times 10^{-18}\left(\frac{l}{3\,\mathrm{km}}\right)\left(\frac{1\,\mathrm{mm}}{r}\right)\mathrm{rad}/\sqrt{\mathrm{Hz}} \tag{2.15}$$

for a detector sensitivity of $10^{-24}/\sqrt{\text{Hz}}$. Such a level of orientation noise is unlikely to be a problem if the source of the torque on the mass is an external disturbance transmitted through the suspension wires or directly applied to the mass, since such forces are likely to produce translations of the mass along the cavity axis which set a more severe limit on the detector's sensitivity and it is believed that such disturbances can be adequately filtered.

The cavity axis must also remain well aligned with respect to the laser beam, both in translation and rotation, in order to maximise the detector's sensitivity. The input laser beam will be in a Gaussian mode and it is important that as much of this light as possible is coupled into the fundamental transverse mode of the cavity. Misalignments will cause light to couple into higher order modes which do not contribute to the output signal. The amplitude of the fundamental mode in a cavity is given by

$$A = A_0 \exp\left(-(\theta/\theta_c)^2\right), \tag{2.16}$$

where $\theta_c = (2\lambda)/(\pi w)$ and w is the beam size given by $l^2\lambda^2/2\pi^2$ for the cavity configuration which is planned [Hough *et al.* 1986]. These misalignments will therefore reduce the fringe visibility and cause a loss of signal. It has been calculated that in order to retain 95% of the signal, an angular stability of order 10^{-7} rad must be achieved [Hough *et al.* 1986].

In addition to cavity orientation fluctuations, variations in the orientation and position of the input laser beam can couple into the output of the detector. For example if the beam width changes or if the beam is rotated or translated from the optimum configuration then the mode matching to the two cavities will vary with a corresponding reduction in the output signal to noise ratio.

An additional effect is that if two beams meet at an angle α and move laterally with respect to each other by an amount δx a phase difference of $2\pi\alpha\delta x/\lambda$ will be induced between the beams. This places limits on the tolerable positional fluctuations which can be permitted for the light returning towards the beam splitter which have been estimated as being of order $10^{-12}\,\text{m}/\sqrt{\text{Hz}}$ [Hough *et al.* 1986]. For a detector with arms of 3 km, this corresponds to an angular stability of the test masses of $3\times10^{-16}\text{rad}/\sqrt{\text{Hz}}$. This light will have four components—the light which

has been directly reflected from the two cavities, and the light leaking from each cavity—and all four components must be accurately controlled. Such lateral motion of the beams could be produced by orientation changes in the test masses and by the natural fluctuations in the light from the laser.

The light returning to the beam splitter from the two cavities must also be well aligned, with little distortion of the wavefronts, if the light is to interfere sufficiently well for recycling to be efficient (although the recycling cavity will act as a filter against beam motion).

Natural positional fluctuations in the laser light can be reduced by passing the light through either a mode cleaning cavity or a monomode optical fibre. For small angular deviations from some reference axis, the laser light can be thought of consisting of a superposition of transverse modes along the reference axis. Tilting the laser beam couples more light into the higher order transverse modes and the action of the fibre or cavity is to suppress these higher order modes. Active control of beam geometry fluctuations has also been demonstrated at Glasgow [Meers 1983]. Using these techniques the required degree of stabilisation of the laser beam should be possible.

2.4.2 Laser Intensity Noise

If a cavity and laser are held exactly on resonance then the system is insensitive to laser intensity fluctuations in the frequency range of the detector, for a null measurement is being made and the use of RF modulation effectively averages these low frequency intensity fluctuations to zero. This is important because intensity noise in lasers at low frequencies is usually rather high—due for example to vibrations of the laser induced by the flow of cooling water and plasma effects—but above 10 MHz argon ion lasers are usually shot noise limited. If the cavity and laser are not exactly on resonance, however, then a measurement of the offset is being made. The measurement system has a gain which is proportional to the laser intensity so that if the required accuracy in the measurement of the cavity offset is dx and the offset is δx, the laser intensity must be stabilised to better

than

$$\frac{\delta I}{I} \sim \frac{dx}{\delta x}. \tag{2.17}$$

For an arm length l and a desired sensitivity of h the required accuracy dx is simply given by

$$dx = hl. \tag{2.18}$$

Therefore we require

$$\frac{\delta I}{I} < \frac{hl}{\delta x}. \tag{2.19}$$

If we assume that the offset, δx, can be kept as small as $10^{-6}\mathrm{m}/G$, where G is the loop gain of the control system which may be made as high as 10^7 using conditionally stable feedback systems, and 10^{-6}m is the natural offset of the cavity from the correct locking point, then to detect millisecond pulses of amplitude 10^{-22}, an intensity stabilisation $\delta I/I$ of 10^{-7} over millisecond timescales is required. This offset will be due mainly to low frequency ground noise and the natural limit to changes in the cavity length of $\lambda/2$. It should be possible to reduce very low frequency drifts, caused, for example, by thermal expansion of the detector, below this limit since feedback systems with very high gains are possible at these low frequencies. Such a degree of intensity stabilisation has been achieved at Glasgow with a control system that uses an electro-optic modulator placed between polarisers as the control element [Mangan 1988].

The effects of radiation pressure in the two cavities may also depend on intensity fluctuations. For if the intensities in the two cavities are different, or if the test masses comprising the two cavities are of different mass, then light intensity changes may induce differential length changes, placing a limit to the sensitivity of a detector of [Hough *et al.* 1986],

$$h > 6 \times 10^{-24}\left(\frac{f}{1\mathrm{kHz}}\right)\left(\frac{M}{10\mathrm{kg}}\right)\left(\frac{\varepsilon_I + \varepsilon_m}{0.1}\right)\left(\frac{\delta I/I}{10^{-7}}\right)\left(\frac{I}{200\mathrm{W}}\right)\left(\frac{\tau}{1\mathrm{ms}}\right)\left(\frac{l}{3\,\mathrm{km}}\right), \tag{2.20}$$

where ε_I is the fractional difference in intensity between the two arms, ε_m is the fractional difference in mass of the test masses, l is the detector arm length, M is the mass of the test masses and f and τ are the characteristic frequency and duration of the gravitational waves being searched for. The effects of offsets place

rather more stringent limits on the tolerable intensity fluctuations with the result that radiation pressure effects should not be a problem in long baseline detectors.

2.4.3 Laser frequency fluctuations

As has been mentioned the cavities in a Fabry-Perot detector must be held close to resonance with the laser light. If the two cavities are identical then small fluctuations of the laser's frequency will be cancelled by interfering the beams from the cavities. However, due to variations in the cavities' mode matching, visibility and storage times the cancellation will be incomplete and small frequency fluctuations will produce an output signal. If the common mode rejection ratio of the cavities which can be maintained is D, then the required frequency stability is

$$\frac{d\nu}{\nu} = hD \tag{2.21}$$

Experiments performed at Glasgow [Kerr 1986] suggest that a value of D=100 may be possible, placing a limit of 10^{-21} on the fractional frequency fluctuations permissible over timescales of a few milliseconds. The frequency of the laser used on the present Glasgow prototype is currently stabilised to the length of the primary cavity to approximately 1 part in 10^{19} over millisecond timescales, a limit set by the effects of shot noise on the detected light, so such stabilisations should be possible with the light powers which are planned for long baseline detectors.

2.4.4 Thermal Noise

A mass suspended as a pendulum in thermal equilibrium will exchange heat with its surroundings with the result that, on average, an energy of $kT/2$ will be associated with each degree of freedom of the pendulum motion. Similarly, the test mass itself will be in equilibrium with its surroundings and so its internal resonances will each have a mean energy of kT. The test masses will therefore be translated, rotated and distorted by thermal noise and it is important to ensure that the effects are unimportant at the frequencies of interest.

Consider a test mass, mass m, suspended as a pendulum with angular frequency ω_0 and quality factor Q. Most of the thermal energy will lie in a bandwidth of ω_0/Q centered around ω_0. Above this resonant frequency the amplitude of the thermal fluctuations will fall as $1/f^2$. It can be shown that the amplitude of these fluctuations at a frequency f above the resonant frequency is given by

$$dx \sim \left(\frac{kT\omega_0}{8\pi^4 Qmf^4}\right)^{\frac{1}{2}} \mathrm{m}/\sqrt{\mathrm{Hz}}. \tag{2.22}$$

For four such masses forming the two arms of an interferometer the limit placed on the sensitivity of the detector to pulses of duration τ is [Hough *et al.* 1986]

$$h = \frac{1}{l}\left(\frac{kT\omega_0\tau^3}{3\pi^3 Qm}\right)^{\frac{1}{2}}. \tag{2.23}$$

A similar argument shows that if the test mass has a moment of inertia I, and the mass rotates about its equilibrium position with an angular frequency ω_0 and quality factor Q, then the amplitude of the rotation of the test mass at a frequency f above the resonant frequency is given by

$$d\theta \sim \left(\frac{kT\omega_0}{8\pi^4 QIf^4}\right)^{\frac{1}{2}} \mathrm{rad}/\sqrt{\mathrm{Hz}}. \tag{2.24}$$

This rotation will couple into an apparent change in the length of the cavity as described in Section 2.4.1 corresponding to an apparent strain sensitivity to pulses of duration τ of

$$h = \frac{r}{l}\left(\frac{kT\omega_0\tau^3}{3\pi^3 QI}\right)^{\frac{1}{2}}, \tag{2.25}$$

where r is as defined in Section 2.4.1.

If we assume that the test masses are solid fused silica of mass m and that they take the form of a right cylinder with equal length and diameter in order to maximise the frequency of the lowest internal resonance, then the moment of inertia will be equal to $I = 2.2 \times 10^{-3} m^{\frac{5}{3}}$. The limit to the sensitivity of the detector set by thermally induced rotations will therefore be

$$h \sim 10^{-24}\left(\frac{r}{1\,\mathrm{mm}}\right)\left(\frac{3\,\mathrm{km}}{l}\right)\left(\frac{100\,\mathrm{Hz}}{f}\right)^2\left(\frac{100\,\mathrm{kg}}{m}\right)^{\frac{5}{6}}\left(\frac{10^4\,\mathrm{sec}}{\tau_{damp}}\right)^{\frac{1}{2}}/\sqrt{\mathrm{Hz}}. \tag{2.26}$$

where τ_{damp} is the time taken for the amplitude of the rotational oscillations of the test masses to be damped to $1/e$ of their starting value. Therefore rotational thermal noise should not be a problem for long baseline interferometers if the natural damping of the rotational mode of the test masses can be kept small.

Thermal excitation of the resonant modes of the test masses may also limit the sensitivity of a detector. In this case the resonant frequencies should be well above the range of signal frequencies and the apparent mirror motion will be given by

$$dx \sim \left(\frac{4kT}{\omega_0 mQ}\right)^{\frac{1}{2}} / \sqrt{\text{Hz}}. \tag{2.27}$$

Masses with Q's of order 10^6 should be possible, so that with an arm length of 3 km and 400 kg masses the resonant frequency of the test masses must be above about 3 kHz. This should be possible with fused silica or aluminium masses.

2.4.5 Seismic Noise

Seismic noise can directly disturb the positions of the test masses. Apart from large amplitude earth tremors there is a constant background of micro-seisms, caused for example by wind, traffic and ocean waves. The amplitude of seismic noise for a quiet site typically has the form [Weiss 1972]

$$dx \sim \left(\frac{10^{-7}}{f^2}\right) \text{m}/\sqrt{\text{Hz}}. \tag{2.28}$$

By suspending the masses as pendulums, isolation is provided above the resonant frequency of the pendulum. In fact for a pendulum whose suspension point is subject to the disturbance given by Equation 2.28 the motion of the pendulum mass at frequencies above the resonant frequency is given by

$$dx \sim \left(\frac{10^{-7} f_0}{f^4}\right) \text{m}/\sqrt{\text{Hz}}. \tag{2.29}$$

where f_0 is the pendulum's resonant frequency. For an arm length of 3 km and a pendulum with a resonant frequency of 1 Hz this places a limit on the sensitivity of the detector at 100 Hz of $6 \times 10^{-19}/\sqrt{\text{Hz}}$. Further isolation of order 10^5 of the point

of suspension is therefore required at 100 Hz to obtain the desired sensitivity for these detectors. In the detectors which have been constructed to date, additional isolation has usually been obtained by supporting the suspension point on a stack of alternate layers of rubber and lead. Such a stack acts as low pass filter and the degree of isolation obtained can be considerable. The isolation of stacks is however more likely to be limited by internal resonances of the material constituting the stack. The behaviour of such a stack is discussed in Appendix A.

If it is desired to operate the long baseline detectors at frequencies of order 10 Hz then such passive isolation may not be sufficient. However active methods of isolation have been considered and successfully implemented [Robertson *et al.* 1982]. Seismic noise will still ultimately limit the sensitivity at low frequencies.

2.4.6 Refractive Index Fluctuations

If the pressure in the vacuum system is not kept sufficiently low then the refractive index of the gas in the tank will introduce a change in the apparent lengths of the two cavities. Fluctuations of the refractive index will therefore produce a phase difference between the two arms which is indistinguishable from a relative length change. The magnitude of this effect determines the vacuum pressure which must be attained in the detector arms. (In practice, effects such as Brownian motion of the test masses set a less stringent limit on the pressure.) It can be shown that the limit to the sensitivity of a Fabry-Perot detector is [Hough *et al.* 1986]

$$h = 4\beta\left(\frac{2kTP}{\pi cal}\right)^{\frac{1}{2}}/\sqrt{\text{Hz}}, \tag{2.30}$$

where P is the pressure, c is the mean molecular velocity, a is the beam radius and l is the arm length of the detector. The value of β is that as defined in the relation $n = 1 + \beta P$. For hydrogen, which is likely to be the main component of gases which leak from the surfaces of the vacuum system, $\beta \sim 10^{-9}/\text{Pa}$. The limit placed on the pressure for a sensitivity of $10^{-24}/\sqrt{\text{Hz}}$ is therefore of order 10^{-6} torr. Allowing for the partial pressure of other molecules, a pressure of 10^{-8} torr should be sufficient to achieve the desired sensitivity.

Gas in the tanks will also scatter light from the beam and this loss mechanism may prevent optimum operation of the detector while recycling. If light which is scattered reflects off the walls of the vacuum tank then it is possible for the detector to become sensitive to motion of the vacuum tank (see Section 2.4.7). The losses corresponding to the pressure limit set by refractive index fluctuations, however, should be negligible.

2.4.7 Scattered Light

If light is scattered from the fundamental mode of the cavity, is reflected from the walls of the vacuum pipes and returns to the beamsplitter to contribute to the output signal then the detector will be sensitive to motion of the reflecting surface which will phase modulate this scattered component. Frequency fluctuations of the laser light will also effectively phase modulate the scattered light. To prevent these effects limiting the sensitivity of the detector it will be necessary to place baffles in the vacuum system to trap any scattered light and to frequency stabilise the light. Fabry-Perot systems are well suited to high degrees of frequency stabilisation. Another technique which has been used to reduce the effects of scattered light in delay line systems is to phase modulate the light so that there will be little coherence between scattered light which has travelled by different paths [Schnupp 1985].

2.4.8 Mirror Quality

In order to achieve efficient recycling, the light lost from the detector on each cycle round the interferometer must be kept small and this places constraints on the losses caused by mirrors that can be tolerated. The cavity end mirrors must have a reflectivity which is very close to 1 while the input mirrors, which must be partially transmitting must have very low absorption and scattering coefficients. Multi-layer dielectric mirrors are being developed for laser gyros, having losses of order 2×10^{-5} and so mirrors of suitably high quality should be available. The mirrors must also be large enough that diffraction losses are kept sufficiently small,

for a loss of less than 10^{-5} this corresponds to a mirror radius which is 2.4 times the beam radius, and they must have an accurate surface figure so that wavefront distortion is kept small. However the mirror parameters set by these considerations are within the limits of current manufacturing techniques.

2.5 Conclusions

Although there are many possible noise sources in interferometric detectors, it appears that detectors with sensitivities to millisecond pulses of around 10^{-22} should be possible. This should enable the detection of many of the sources described in the previous chapter. The noise sources which should dominate in the long baseline detectors—shot noise at high frequencies and thermal noise associated with the pendulum motion of the masses at low frequencies—should produce noise which is Gaussian in nature. Gravitational wave signals must therefore be detectable above a background of Gaussian noise. Although the practical difficulties involved in the development of such detectors are formidable, the results of work on the prototype detectors at Glasgow, Garching, Caltech and MIT suggest that these difficulties may be overcome.

Chapter 3

An Orientation Control System for a Nearly-Free Test Mass

3.1 Introduction

It was shown in the previous chapter that the orientation of the test masses in a laser interferometer must be accurately controlled for optimum operation. In particular the mass's orientation and low frequency motion must be controlled if excess noise is not to limit the detector's output. This chapter describes a system that has been implemented at Glasgow on the 10 metre prototype detector to control the orientation and position of a test mass at low frequencies while leaving the test mass free at the higher frequencies at which gravitational wave signals are expected.

It is envisaged that the large, kilometer scale, detectors will be able to operate at frequencies as low as a few tens of hertz and so the test masses must be free above this lower limit. However the orientation of each test mass must be controlled to a sufficient accuracy that the cavity and laser axes are properly aligned. It is therefore important to have servo systems which can control the slow variations without affecting the response of the detector. The present Glasgow prototype has a lower frequency limit of a few hundred hertz and is therefore not so demanding on the servo system.

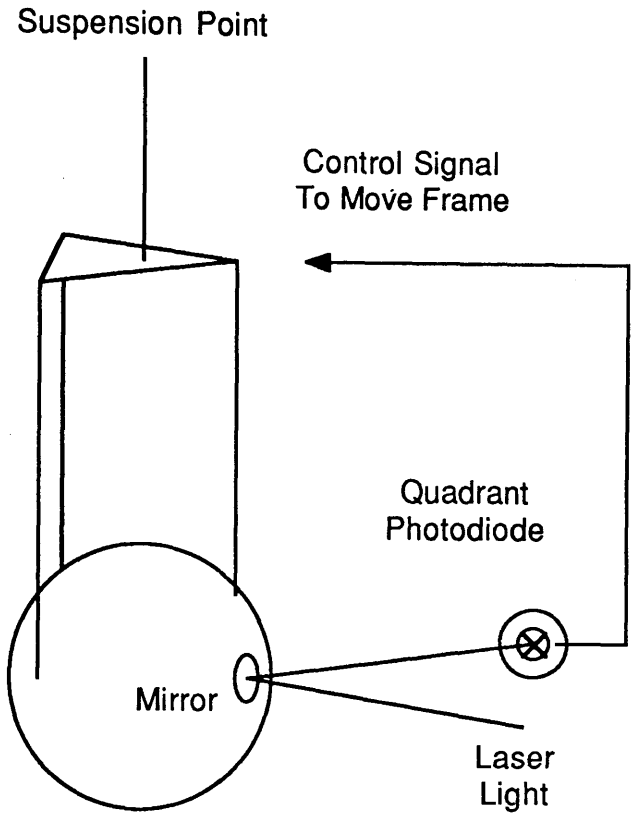

Figure 3.1: Schematic diagram of an orientation control system used at Glasgow. The mass's orientation is sensed optically and controlled by moving the suspension wires.

3.2 Orientation Control Systems

Various methods for controlling the orientation of the test masses in the prototype laser interferometer detectors have been developed. One method which has been installed at Glasgow [Ford 1981], and is still used to control some of the test masses, applies feedback to the mass through motions of the suspension wires. The orientation of the mass is monitored by reflecting light, generated by a HeNe laser, from a mirror fixed to the mass onto a position sensitive photodiode (Figure 3.1). Suitable feedback signals, derived from this photodiode, can then be applied to moving coil and piezoelectric transducers, attached to a light frame which is suspended from a seismically isolated top plate, and from which the mass is in turn

suspended. This method has been effective in controlling the orientation fluctuations of the mass to 10^{-5} rad/$\sqrt{\text{Hz}}$, however it has several limitations. For example, the monitor signal derived from the photodiode may be limited by the directional stability of the laser beam itself. The lasers used at Glasgow are mounted outside the vacuum system and so air currents can affect the directional stability of the beam, as can the presence of additional transverse modes in the lasers. In addition the control structure used to displace the suspension wires is, of necessity, rather complicated. It introduces resonances, typically between 10 and 100 Hz, into the response of the pendulum to motion of the suspension point and may act to damp the pendulum mode. It is therefore an advantage to have as simple a suspension system as possible and this suggests that the controlling forces must be applied directly to the test mass.

Another method which has been proposed for the long baseline instruments and which has been used on the MIT 1 metre prototype [Shoemaker 1987] is to apply electrostatic control forces to the mass. This has the advantage that it is relatively easy to shield the test mass from external electrostatic fields. However, in order to generate the required forces, it is necessary to apply high voltages to the electrostatic plates—with the possible problem of electrical discharge. If the test mass is not to be given an isolated static charge but is to be grounded, then the test mass must be subjected to a bias force in order to obtain a bidirectional change in the force. In addition, since the capacitance of the plates will change with separation, the plates must be driven by a constant charge source to apply a correct force feedback. Such sources must have high output impedances, of order $10^9\Omega$, in order that the charge on the capacitor plates does not discharge over timescales shorter than the period of the pendulum to be stabilised and such a current source requires careful design to achieve an adequate noise performance.

3.3 The Coil and Magnet System.

It was decided at Glasgow to implement a system where the correcting forces are generated between small permanent magnets mounted on the test masses and

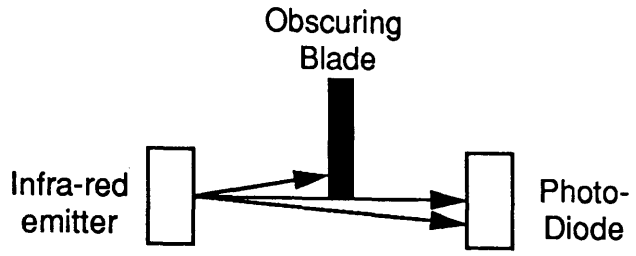

Figure 3.2: The principle of the shadow meter.

coils positioned nearby. Such a system had been implemented on the delay line detector at the Max-Planck-Institut [Shoemaker 1987]. In this system the position and orientation of the mass are monitored using "shadow meters" in which a blade fixed to the mass partially obscures the path between a light source and a light detector (see Figure 3.2). Small variations in the position of the mass modulate the intensity of the detected light and a signal representing the relative positions of the blade and the emitter/detector pair can therefore be derived. Small permanent magnets are fixed to the test mass and so, by varying the current in nearby coils, forces can be applied to control the orientation and position of the test mass. The test mass consists of a cylinder of fused silica, about 12 cm in length and 12 cm in diameter and with a mass of 2.2 kg. This is suspended as a pendulum with a period of about 1 second. Three magnets were fixed to the rear face of the cylinder, at the corners of an equilateral triangle, so that they lay 40 mm from the centre of the face. Small aluminium blades were fixed to the magnets and these were used to obscure the light in the shadow meter. Each coil and its corresponding shadow meter were combined in a unit which was designed at Garching (Figure 3.3). The three units were fixed to an aluminium back plate which in turn was fixed to the vacuum chamber. The positions of the coil units in the direction perpendicular to the back plate could be adjusted to ensure that when the system was locked, and the cavity was properly aligned, the blades attached to the mass lay in the correct position relative to the shadow meters (Figure 3.4).

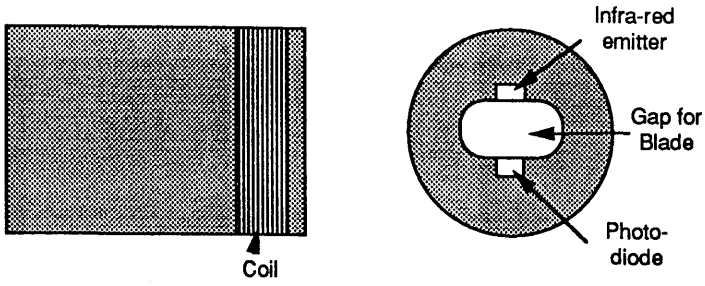

Figure 3.3: The shadow meter/coil assembly.

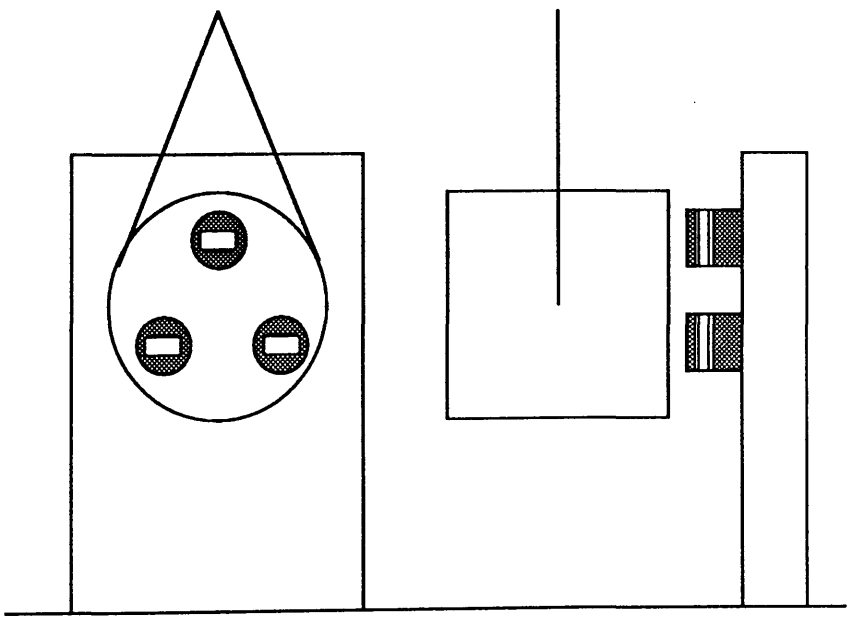

Figure 3.4: The arrangement of the orientation sensors and feedback elements on the prototype detector.

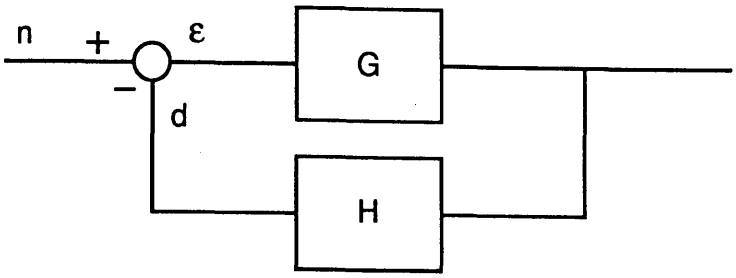

Figure 3.5: Schematic, block diagram of a negative feedback circuit.

Figure 3.5 shows the block diagram of a negative feedback system in which a quantity ϵ is controlled. G and H represent control elements with the transfer functions $G(s)$ and $H(s)$ with $s = 2\pi f\sqrt{-1}$, d is the contribution to ϵ provided by the control system and n represents a noise input, corresponding for example to the natural fluctuations of the quantity ϵ. We can obtain the relation

$$\epsilon = \frac{n}{1+GH}, \tag{3.1}$$

and so we see that the effect of the control system is to reduce the natural fluctuations of the system by a factor $1 + G(s)H(s)$. However, such a system may not be stable, that is its impulse response may not approach zero as time approaches infinity. For example if $G(s)H(s) = -1$ for some s then the value of ϵ becomes infinite at that frequency. We must therefore consider the conditions for stability of the system. It can be shown (for example [Di Stefano *et al.* 1976]) that the system will be stable if the roots of $1 + G(s)H(s)$ all lie in the left-half of the complex plane. This is equivalent to the requirement that the phase delay of the loop gain GH is less than 180° when $|GH|$ is unity.

The transfer function relating the motion of a lightly damped pendulum of mass m, resonant frequency ω_0 and damping coefficient γ to a force is given by

$$\frac{x}{F} = \frac{1}{m(s^2 + \gamma s + \omega_0^2)}. \tag{3.2}$$

For high frequencies, that is those above the resonant frequency, this corresponds to a phase delay of very nearly 180°. Therefore if the position x of the mass is

measured and a force proportional to this displacement is applied to the mass such that the unity gain frequency of this loop is above the resonant frequency of the pendulum, the servo system is marginally stable so that its impulse response tends to zero rather slowly. It is necessary to introduce a lead compensator, that is a differentiator, to reduce the phase delay below 180°. If the phase delay is sufficiently reduced then the impulse response of the system will be critically damped.

An equivalent view of this system is that the action of the position dependent force is to increase the resonant frequency of the pendulum, so that the impulse response is simply that of a lightly damped pendulum of the new resonant frequency, while the lead compensation, which generates a force proportional to the derivative of the position, that is the velocity, introduces a damping term which can critically damp the impulse response.

For the control system discussed here, three independent motions will be damped, namely the translation of the mass perpendicular to the plane defined by the three shadow meters, rotations of the test mass about a vertical axis and rotations about a horizontal axis parallel to the plane defined by the shadow meters. If each of the three units has an identical loop gain then the control of these motions could be simply analysed. For example translation of the mass would generate a common mode signal in the three coil units, and rotations and tilts would generate differential signals in the three units. For units which are not perfectly balanced, however, such an analysis becomes rather more difficult, as translations and rotations will be coupled.

The circuit shown in Figure 3.6 was developed to implement this control system. The first stage of amplification simply generates a voltage proportional to the photo-current from the sensor. The second stage includes the differentiator for damping, while the third amplifier has a high current output and is used to drive the coils. Since the pointing of the mass must be adjusted to align the cavity, the ability to steer the masses by applying offset voltages to the input amplifiers was also added.

In addition to the three units which were used to control the orientation, a fourth unit was used to damp horizontal motions of the mass, transverse to the

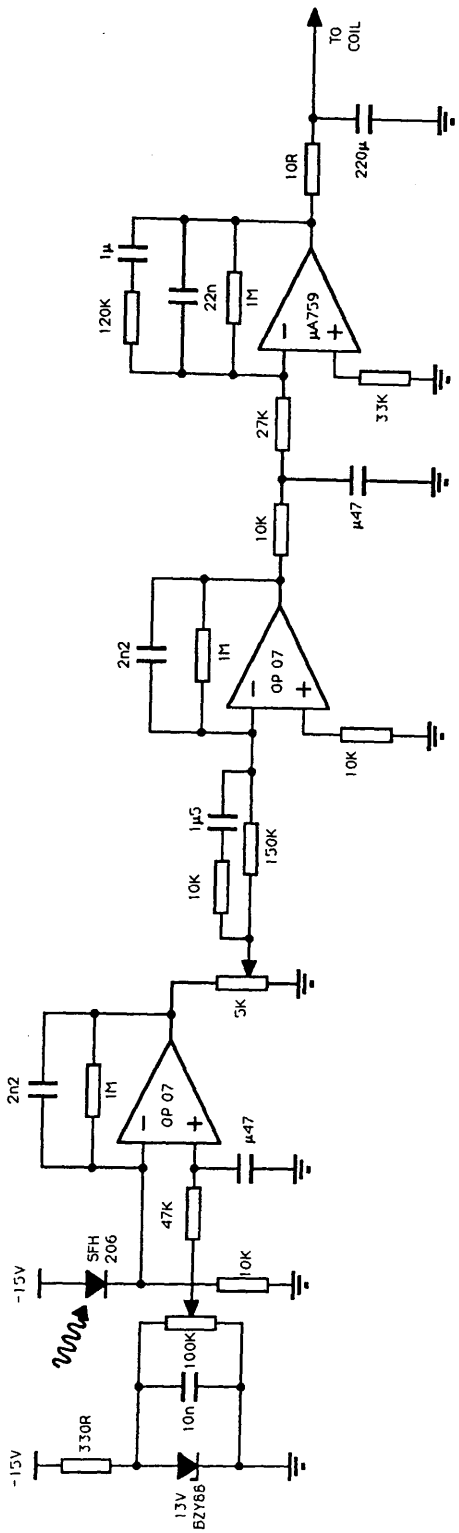

Figure 3.6: The circuit for the control system electronics.
(Based on a design from MPQ, Garching)

cavity axis.

3.4 Noise Sources

This system has several potential limitations which will now be discussed.

3.4.1 Motion of the Coils.

For low frequencies, where the loop gain of the system should be relatively high, the position of test mass is effectively fixed relative to the coil units. Any motion of the coil units, caused for example by inadequate isolation from seismic noise, will couple directly to the mass and so it is necessary to isolate the coil units from ground motion.

For higher frequencies, above the unity gain point, the gain of the feedback system is reduced so that the mass is effectively free from the control of the photodiode. However the motions of the coil may still couple into a force on the test mass if the separation of the magnets and coils is not optimum. The force along the x axis on a magnetic dipole, dipole moment m_D which is aligned with the axis and which is situated at point P in a magnetic field B is given by

$$F = m_D \left(\frac{dB}{dx}\right)_P . \tag{3.3}$$

For a circular current loop of radius a carrying a current I (Figure 3.7) the magnetic field along the axis of the coil is given by

$$B_z = \frac{\mu_0 I a^2}{2(a^2 + x^2)^{3/2}} \tag{3.4}$$

giving a force on a magnetic dipole with magnitude

$$F = \frac{3 m_D z \mu_0 I a^2}{2(a^2 + x^2)^{5/2}}. \tag{3.5}$$

The force generated by such a system should ideally be unaffected by small relative motions of the coils and magnets, that is $dF/dx = 0$, in order that motion of the

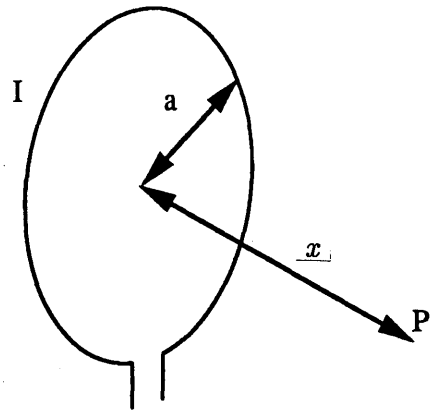

Figure 3.7:

coil units does not couple directly into motion of the test mass. For this simple case such a condition is satisfied when $x = a/2$. Therefore it is possible to set the separation of the dipole and coil so that the force on the dipole is independent to first order on variations of this separation. However, if the separation is not optimal then motion of the coils will vary the force on the mass.

To estimate this effect, consider the arrangement shown in Figure 3.8. A uniformly magnetised bar can be modelled as a solenoid of equivalent dimensions [Lorrain and Corson 1970]. Real permanent magnets will not be truly uniformly magnetised due to end effects in the material, however such effects should be small. The force versus separation relation for such an arrangement can be derived by integrating the forces between each element of the coil and magnet.

$$F_x = \int_0^{2\pi} d\theta \int_0^{2\pi} d\phi \int_a^b ds \int_0^c dp \int_0^e dq \; k \frac{fs(p+d+q)\cos(\theta-\phi)}{(f^2+s^2+(p+d+q)^2-2sf\cos(\theta-\phi))^{3/2}} \tag{3.6}$$

where k is a constant of coupling, and the other terms are as shown in Figure 3.8.

For the arrangement used at Glasgow, with magnets 9 mm in diameter and 3 mm thick and with coils with inner and outer diameters of 20 mm and 30 mm respectively and a length of 6 mm, numerical integration gives the result shown in Figure 3.9. Experimental values for the force/distance relation were derived by fixing a magnet to a mass suspended as a pendulum and applying a force with a

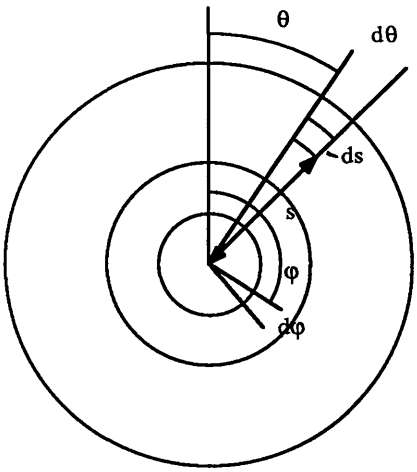
θ
dθ
ds
s
φ
dφ

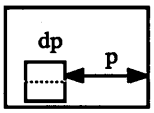
dp
p

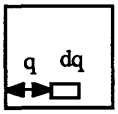
q
dq

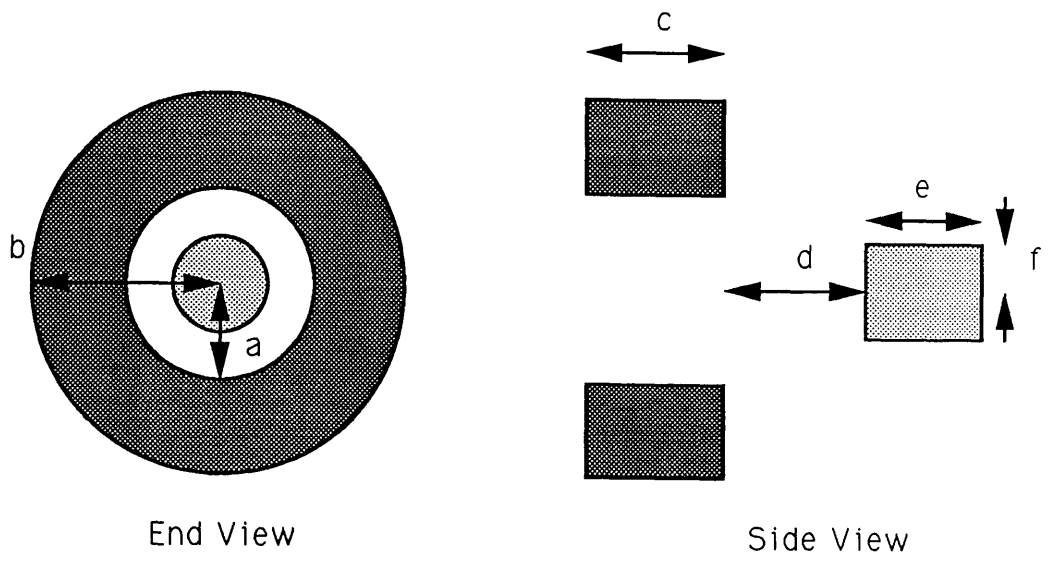

Figure 3.8: The elements of the model of a uniformly magnetised bar in the field of a solenoid.

coil unit. Measuring the pendulum deflection and the magnet/coil separation with a travelling microscope gave the points plotted in Figure 3.9. It can be seen that the experimental and theoretical results give maximum force at about 1.5 mm.

If the magnet/coil separation is as much as 1 mm away from the optimum position then Figure 3.9 shows that the value of dF/dx is of order $40F_0$ N/m, where F_0 is the force which is being applied to the mass. The coils each generate a force of order 0.24 N/A on the magnets, so that if we consider a case where each of the three output amplifiers generates an output current of 100 mA, then a motion of the coils at frequency f, above the resonant frequency of the pendulum, with amplitude A_0 will generate an equivalent motion of the test masses of

$$dx_{mass} \sim \frac{2.88}{m\omega^2} \times A_0. \tag{3.7}$$

From this it can be seen that to achieve a displacement noise of 10^{-18} m/$\sqrt{\text{Hz}}$ at 1 kHz requires a coil motion at 1 kHz of less than 3×10^{-11} m/$\sqrt{\text{Hz}}$. From Equation 2.28 it can be seen that this should not be a problem at 1 kHz even if the coils are not isolated from the ground.

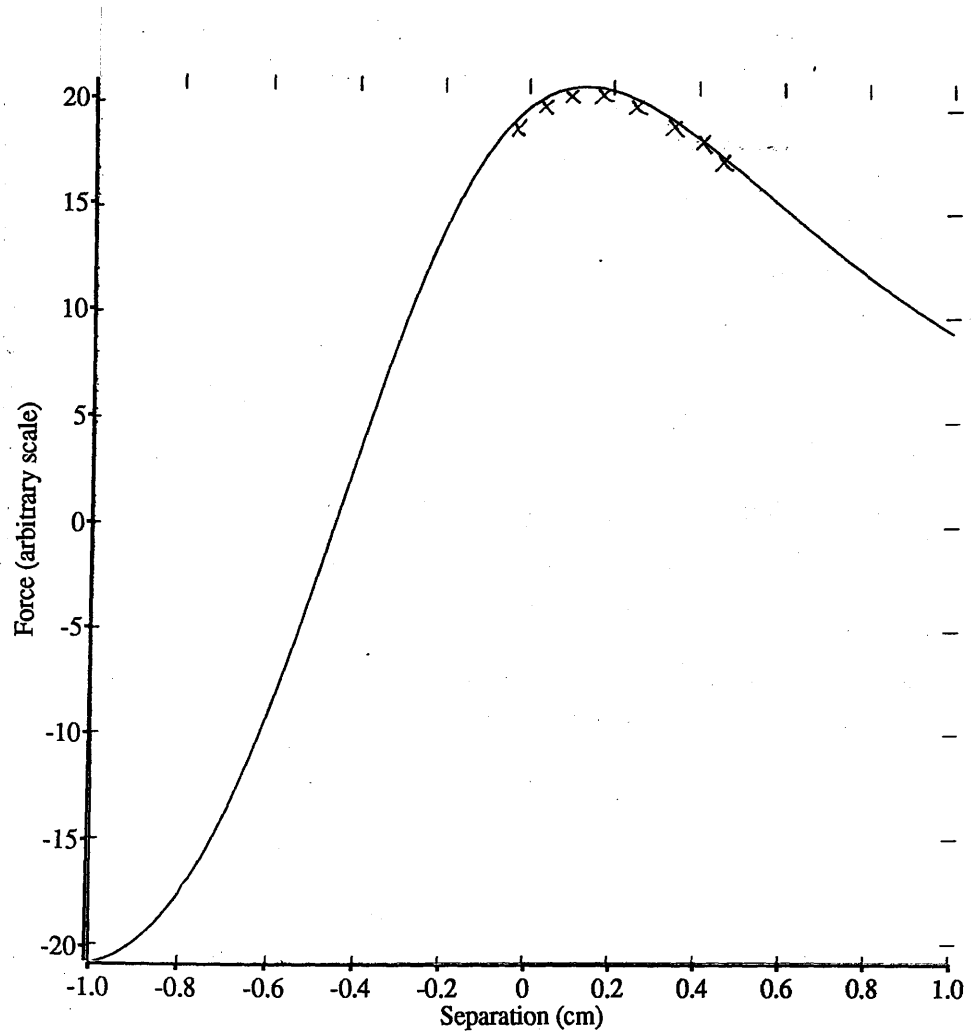

Figure 3.9: Variation of force with coil/magnet separation.

Motion of the coils may also couple to motion of the mass through eddy-currents induced in the coil by motion of the coil units relative to the magnets. For the present case, however, such an effect is small, and could if necessary be reduced by using a current source rather than a voltage source to drive the coils.

3.4.2 Detector and Amplifier Noise

The fundamental limitation to the stabilisation that can be achieved in this system should be set by shot noise in the photodiode's photocurrent, though in practice the light from the emitter is not shot noise limited. In fact, since the system only operates at frequencies of order 10 Hz, flicker noise in the emitted light is likely to be a limitation on the system. The root mean square noise level of the output of the first amplifier stage due to flicker noise and amplifier input noise is about $10\,\mu$V. In the linear region of the shadow meter's response a motion of 10^{-4} m typically generates a change in the output voltage of 10 V so that the corresponding limit to the orientation fluctuations which can be measured is therefore better than $10^{-8}\,\mathrm{rads_{rms}}$. This should be sufficient for the present system.

Amplifier noise may introduce current fluctuations at high frequencies which generate forces which are sufficiently strong to displace the mass and limit the sensitivity of the detector. It is therefore essential to reduce the gain of the amplifiers above the unity gain frequency of the servo system. In particular, since the output noise of an operational amplifier cannot be less that its input noise and the operational amplifier used in the final stage of the circuit of Figure 3.6 has an input voltage noise at 1 kHz of $60\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$, which if applied to the coils would generate a motion at 1 kHz of order $4 \times 10^{-18}\mathrm{m}/\sqrt{\mathrm{Hz}}$, it is essential to introduce low pass filtering after the final amplifier as shown in Figure 3.6.

3.4.3 Magnetic Field Noise

Background magnetic field fluctuations may couple forces to the test masses in two ways. Firstly by direct coupling of the magnets to magnetic field gradients, and secondly by reradiated fields generated by induced currents in the driving coils.

If the magnets have dipole moments which are well matched then such a magnetic field is unlikely to produce significant rotation, and any rotation produced at low frequencies would be reduced by the control of the feedback system. However at signal frequencies translations of the mass may limit the sensitivity of the detector.

For an electromagnetic plane wave $|\frac{dB}{dx}| \sim \frac{1}{c}|\frac{dB}{dt}|$, and the voltage induced across a coil of radius a and with N turns is $N\pi a^2 \frac{dB}{dt}$. A current in the coils of order $\frac{N\pi a^2}{Z}\frac{dB}{dt}$ is induced in the coils, where Z is the impedance of the coil and the amplifier output stage. (Here we assume that the reradiated field is smaller than the incident field.) If we make the assumption that the field of such a multiturn coil can be adequately approximated by a single turn coil with a current N times that of the multiturn coil, then we see by differentiating Equation 3.4 that, at the maximum of the reradiated field from a coil with $N = 500$ and if the impedance at 1 kHz is 70Ω

$$\left|\frac{dB}{dx}\right| \sim 2 \times 10^{-2} \left|\frac{dB}{dt}\right|, \tag{3.8}$$

which is approximately 10^7 times greater than that of the incident field. Therefore, provided that the induced voltage due to the background magnetic field variations, as measured across the coils, does not correspond to that which would produce significant translations of the test mass, background magnetic fields from distant sources, for example geomagnetic variations, should not be a problem. Measurements of the induced voltage across the coil show no evidence for such pick-up at significant levels.

For the case of nearby sources of stray magnetic fields, the value of dB/dx can be much greater than for a plane wave. However such a source is likely to be man-made and can therefore, in principle, be removed from the vicinity of the test mass. The most probable sources of such fields are electrical components, which are likely to emit periodic signals at harmonics of the line frequency.

3.4.4 Thermal Drift

Over long timescales the temperature of the emitters and detectors in the position sensors will vary. Any temperature dependency of the emitters or detectors will

cause a change in the apparent position of the blade relative to the emitter/detector pair. The orientation of the test masses will therefore drift with temperature. On the current prototype the alignment is performed by the operator and so must be checked periodically. However schemes for automatically aligning the cavities to the input laser beam have been designed and implemented at Glasgow. For example the orientation of the mass may be modulated and the visibility signal demodulated to give a correction signal. Slow thermal drifts should not be a serious problem, and may be reduced by controlling the diode current.

3.5 Results

The spectrum of the orientation fluctuations of the system as measured at the monitor point is shown in Figure 3.10. This voltage at the monitor point typically varies by $50\mu V_{rms}$ corresponding to a root mean square angular fluctuation of less than 2×10^{-8} rad. Since this is measured relative to the reference plane defined by the positions of the coils, this alone does not allow any limit to be set on the absolute accuracy of the orientation stability. As an independent check on the absolute stability, the orientation was also measured by reflecting a laser beam from a mirror fixed to the test mass onto a quadrant photodiode. In fact the accuracy of this measurement is limited by the natural positional fluctuations of the laser beam, and so only an upper limit to the angular fluctuations can be set. The spectrum of the angular fluctuations measured in this way is shown on Figure 3.11. The angular stability measured in this way is $< 10^{-7} \mathrm{rad}_{rms}$.

3.6 Conclusions

The system described in this chapter allows the control of the low frequency ($\leq$10 Hz) positional changes of the test mass to an accuracy of order 3×10^{-9} m, and the orientation fluctuations of the test mass over timescales of a few minutes are less than $10^{-7}\,\mathrm{rad}_{rms}$. Such a degree of positional stabilisation can greatly simplify the operation of the detector, since the laser frequency and the secondary

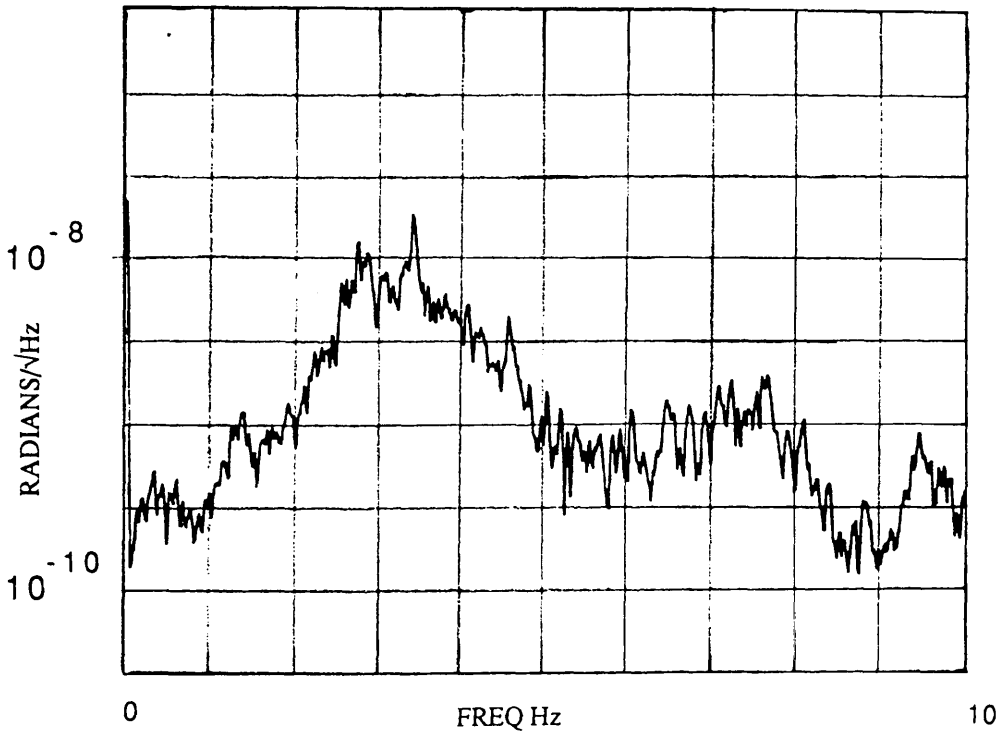

Figure 3.10: Spectrum of the closed loop monitor point signal.

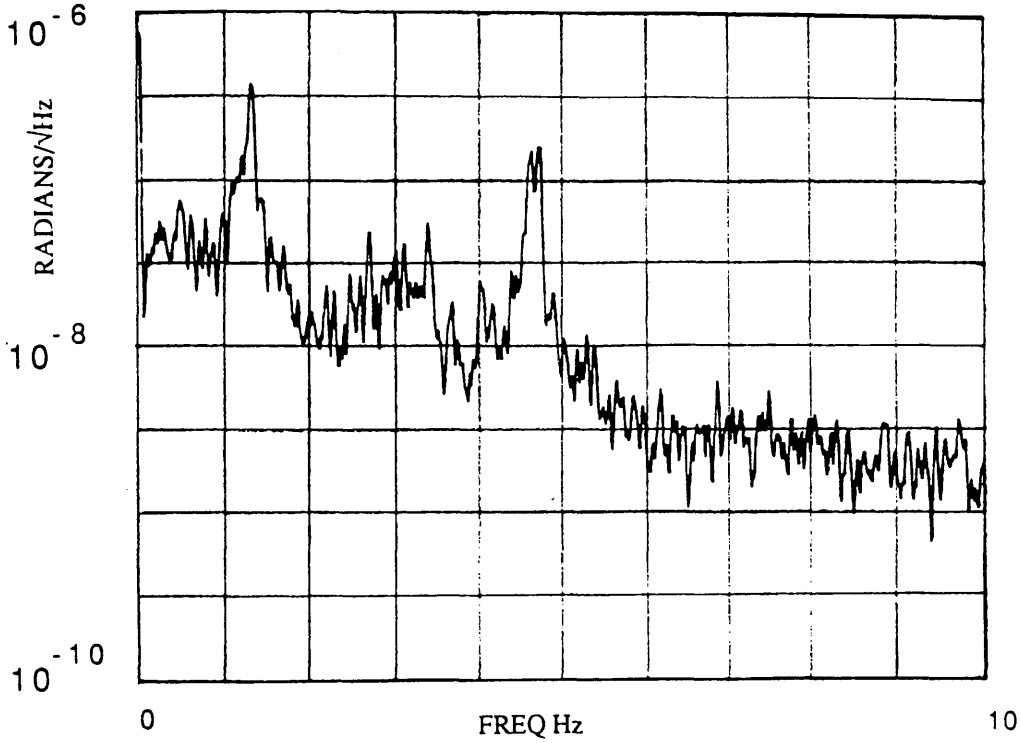

Figure 3.11: Spectrum of the angular fluctuations measured with the laser and quadrant photodiode.

cavity length are not forced to track large changes in the length of the primary cavity at the pendulum resonant frequency. The orientation stability is more than adequate to ensure that the present prototype's sensitivity is not limited due to cavity misalignments, since the alignment can be adjusted to an accuracy of the same order as the stability achieved. The noise performance at higher frequency is sufficient that the detector's sensitivity is not limited by motion of this test mass.

Such a control system may be used on the large scale detectors, though it is likely that such a system would require more isolation of the coils from ground motion at gravitational wave signal frequencies in order that field gradients and damping of the test mass due to eddy-currents will not limit the sensitivity. Background magnetic fields may also become more of a problem at such sensitivities. However, improved seismic isolation and the use of test masses with a large moment of inertia, suspended so that the resonant frequencies of the rotation and tilt modes of the mass are as low as possible, can reduce the natural orientation fluctuations of the mass, so easing the requirements on the control system.

Chapter 4

A Data Acquisition System for the Glasgow Prototype

4.1 Introduction

The astrophysical sources which were described in Chapter 1 fall into two classes; those which should be detectable online, such as short pulses of gravitational waves from supernovae which should be detectable in suitably bandpassed data or the short, approximately 1 second, oscillating bursts expected from coalescing compact binary systems which will be detectable using matched filtering techniques; and those which will only be detectable by intensive offline processing or by comparing the outputs of at least two detectors, such as nearly monochromatic gravitational waves and a stochastic background. For this second class of source it is necessary to record the output of these detectors continuously, while for the first there is an advantage in having a recording system which will preselect signals which may be due to gravitational waves. This chapter describes a digital recording system which has been developed and used to record data from the prototype detector at Glasgow.

4.2 Sampled Signals

A digital recording system is unable to record a continuous waveform, and so it is necessary to record discrete samples of the data. It can be shown (see for example [Bracewell 1986]) that a bandlimited signal can be recorded without loss of accuracy by a series of discrete samples of the signal separated by a sufficiently

small time interval. More precisely, the sampling theorem states that if the Fourier transform, $F(s)$, of a function $f(t)$, defined by

$$F(s) = \int_{-\infty}^{\infty} f(t) \exp(-i2\pi st)\, dt \tag{4.1}$$

and with the inverse transform defined by

$$f(t) = \int_{-\infty}^{\infty} F(s) \exp(i2\pi st)\, ds, \tag{4.2}$$

is zero for all frequencies greater than some frequency s_c, then the continuous function $f(t)$ can be uniquely determined from a knowledge of its sampled values

$$\hat{f}(t) = \sum_{n=-\infty}^{\infty} f(nT)\delta(t - nT), \tag{4.3}$$

where $T \leq 1/2s_c$ and the δ distribution is defined by

$$\delta(x) = 0 \text{ for } x \neq 0, \tag{4.4}$$

$$\int_{-\infty}^{\infty} \delta(x)dx = 1$$

(The function $f(t)$ must meet some conditions on its continuity for its Fourier transform to exist and it must be finite. However all physical signals will meet these conditions.) The conditions of the sampling theorem can be appreciated by considering a bandlimited signal and its Fourier transform as shown in Figure 4.1. Sampling of this signal is therefore equivalent to multiplying the time series by a series of impulses which are separated by the sampling interval. The convolution of two functions $f(t)$ and $g(t)$ is

$$h(t) = \int_{-\infty}^{\infty} f(u)g(t - u)du, \tag{4.5}$$

which will be denoted by the following notation

$$h(t) = f(t) * g(t). \tag{4.6}$$

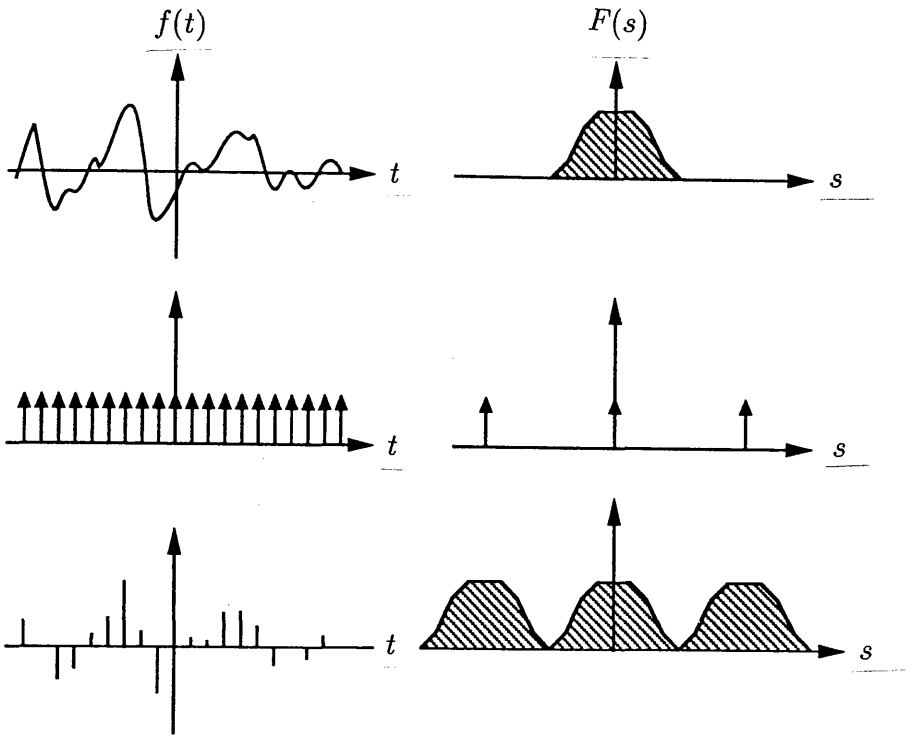

Figure 4.1: The effect of sampling on the spectrum of a bandlimited signal.

From these definitions it is simple to obtain the time-convolution theorem, namely that if two functions $f(t)$ and $g(t)$ have the Fourier transforms $F(s)$ and $G(s)$ respectively then $f(t) * g(t)$ has the Fourier transform $F(s)G(s)$. The frequency-convolution theorem can also be derived, namely that $f(t)g(t)$ has the Fourier transform $F(s) * G(s)$. From the convolution theorem we can see that the sampled signal has a Fourier transform which is the convolution of the signal's transform and the transform of a series of impulses. Since the Fourier transform of a series of impulses separated by T in the time domain is a series of impulses separated by $1/T$ in the frequency domain the effect of the convolution is to produce copies of the original transform separated by $1/T$. For a signal which is bandlimited to sufficiently low frequencies we can see that there is no overlap between any of the copies of the original transform and so in principle this sampled signal could be filtered to obtain the original signal. This is the requirement of the sampling theorem.

In practice it is not possible to obtain a signal which is perfectly bandlimited. In this case the signal and its Fourier transform may be as shown in Figure 4.2.

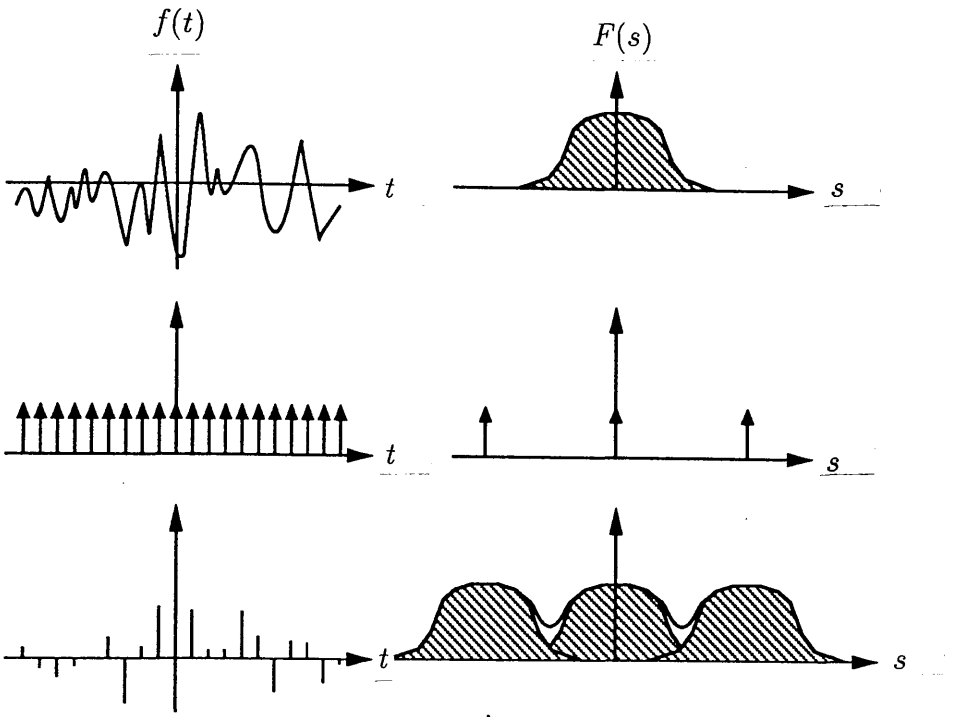

Figure 4.2: The effect of sampling on the spectrum of a signal which does not meet the requirements of the sampling theorem.

Sampling the signal introduces overlap between the copies of the signal's spectrum, so that spectral components at different frequencies are added in an irreversible way. The original signal cannot be retrieved from the sampled values by filtering. This effect is called aliasing. Therefore to record a signal containing components up to a frequency s_u it is necessary to low pass filter the signal to frequencies below s_u and sample the resulting signal at intervals of somewhat less than $1/2s_u$. This introduces a guard band at high frequencies, which allows for the transitional range of the filter.

The Discrete Fourier Transform

For a finite length of sampled data, so that we have N points $x_0, x_1, \ldots, x_{N-1}$, the Fourier transform is not appropriate. Rather the discrete Fourier transform (or

DFT), defined by

$$z_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j \times \exp\left(-i\frac{2\pi jk}{N}\right) \text{ for } k = 0, 1, \ldots, N-1, \qquad (4.7)$$

should be used. We therefore obtain information on the complex amplitude of frequency components at N distinct frequencies.

Since a finite time series of N samples is that which would be produced if an infinite time series was multiplied by a window function which is one for samples from 0 to $N-1$, and zero everywhere else, each of these N frequency bins has an effective bandwidth which has the form of a sinc function where $\mathrm{sinc}(x) = \sin(\pi x)/\pi x$ centred around that bin's frequency. For real data, x_j, it can be shown that there is redundancy in the z_k and in fact there are only $N/2$ independent complex amplitudes. The other half of the z_k effectively correspond to the values for negative frequencies, and can be disregarded for most purposes.

4.3 Hardware

The recording system was based on a Data General Nova 1220, a general purpose minicomputer. Two industry standard, 800 bpi, 9 track, 1/2 inch magnetic tape drives, with a linear recording speed of 37.5 inches per second, and an analogue to digital converter unit were connected to this computer. Transfer of data between the computer and the tape drives and ADC unit was performed using the Nova's Direct Memory Access data channel transfers.

The maximum data recording rate of the tape drives is 30000 bytes per second, a figure reduced slightly by inter-block gaps between each record on the tape. If run continuously at this maximum rate each 2400 foot tape would last for just under 13 minutes and would store almost 21 Mbytes.

Programs running on the Nova could initiate a transfer of a block of data from any area of the computer's memory on to tape. Having received the command, the interface would transfer the data in the background while the main program continued. Two flags are maintained by the interface; BUSY and DONE. When

the transfer of data to tape is in progress the BUSY flag is set. When the transfer is complete BUSY is cleared and DONE is set. These flags can therefore be examined under program control to monitor the progress of a transfer and to ensure that a new transfer request is not made while the previous transfer is in progress.

4.3.1 The analogue to digital converter

The analogue data system connected to the Nova was one that had been designed and developed for previous experiments at Glasgow. The analogue to digital conversion was performed by two 8-bit analogue to digital converters, DATEL ADC - EH8B1. These converters are of the successive approximation type and have a conversion time of $4\,\mu s$. They were configured so that their useful voltage input range was ± 5 volts. Their linearity and accuracy are specified as $\pm 1/2$ LSB and $\pm 0.2\%$ of FS $\pm 1/2$LSB respectively. The converters were preceded by sample and hold circuits designed at Glasgow which were triggered by an external digital clock signal. The data from both converters were read simultaneously into one 16-bit computer word. The use of two independent converters ensured that crosstalk which might be present if one converter was used to record two multiplexed signals might be avoided. The analogue front end can be instructed by a program running in the Nova to take a given number of samples and transfer them automatically into the Nova's memory. Like the magnetic tape interface the analogue data interface maintains two flags BUSY and DONE. When the correct number of samples has been transferred to the host the DONE flag is set. If the program starts another set of conversions within one sampling clock interval of detecting the transition of the DONE flag then no samples are missed between blocks.

Quantisation Error

Since an ADC represents continuously variable voltages by discrete levels, the sampled values are effectively equal to the input voltage plus some error term. For random, large amplitude signals, this noise term should have an amplitude distribution function which is uniform. Each sample of this uniform distribution

will be independent from those preceding and so the error signal will be white. Therefore the effect of quantisation error will be to add a white noise source to the voltage signal and thereby limit the dynamic range of the recorded signal. The variance of this quantisation noise is given by $Q^2/12$ where Q is the separation between successive levels of the ADC. The signals which are to be recorded from gravitational wave detectors will usually be random signals whose amplitude will have a Gaussian probability distribution and so to avoid clipping problems from the finite amplitude range of the ADC the signal level should be set to some fraction of the full range of the ADC. If the signal size is adjusted so that the rms amplitude of the signal is one f^{th} of the full range of the ADC the limit to the dynamic range is given by

$$\text{SNR} = 10\log_{10}\left(\frac{12}{f^2 2^{-2(b-1)}}\right) \tag{4.8}$$

where b is the number of bits used in the ADC, including the sign bit. For 8-bit converters and a signal whose rms level is one fifth of the full scale, this corresponds to a dynamic range of 39 dB.

For a sinusoidal signal in the absence of noise, the quantisation error should be regarded not as a noise source, but rather as a distortion. Indeed, if two sinusoidal frequencies of different frequencies are added then sampled, the power spectrum of the sampled signal may have significant power at the sum and difference frequencies (and many other higher order harmonics), that is the signals have effectively been mixed. To prevent this effect causing spurious peaks in the power spectrum it is important to remove large amplitude, narrow band signals, such as may be caused by wire resonances, before sampling, for even if these peaks are not in the frequency range of interest, the mixed frequencies are likely to be. This point is discussed further in Chapter 6.

4.3.2 Timing Accuracy and Clock Jitter

To compare the outputs of two separate detectors it is vital to record accurate timing information with the data. Since laser interferometer detectors are wide-band, the expected pulses should be of millisecond duration, but since the delay

in arrival times between detectors in Europe and North America may be as much as 20 milliseconds, a timing accuracy of a few milliseconds should be adequate for a simple coincidence experiment such as this. The timing system at Glasgow is based on the 60 kHz transmissions of the MSF standard frequency. This stable carrier is pulsed on and off once every second and is pulse width modulated with digitally encoded timing information. The time as recorded at Glasgow will differ from that of the transmitting station due to propagation delay and due to delay in the receiver and decoder. The propagation delay from Rugby, where the MSF transmitter is based, and Glasgow is of the order of 1.5 ms. This is almost constant but varies slightly due to atmospheric effects. (Subsequent comparison of the received MSF signal with a Rubidium clock which is accurate to 1 part in 10^{12} has shown that the propagation variation is usually less than about 10 μsec, changing over timescales greater that about ten minutes.) The receiver at Glasgow detects the falling edge of an MSF pulse for timing synchronisation, providing a digital pulse at the beginning of each minute. However the receiver incorporates a narrow, 60 kHz filter to prevent interference from powerful radio transmissions at nearby frequencies. The rise and fall time of the pulses is therefore limited by the bandwidth of this filter, with the result that the time is retarded by an additional 3 to 4 ms, a figure which may vary according to the signal level and the threshold set. The overall time delay is therefore approximately 4.5 ms to 5.5 ms. The time as recorded at Glasgow must therefore be corrected for this delay.

The sampling clock was derived from an oven controlled 6 MHz quartz crystal oscillator which was phase locked to the 60 kHz MSF transmissions. The MSF transmissions have an accuracy of 1 part in 10^9 and propagation delay variations cause an estimated variation in the received signal of 1 part in 10^8. A signal which is represented by samples which are not uniformly spaced will effectively be phase modulated. The effect will be worse at higher frequencies (it is equivalent to a Doppler shift of the signal) with the result that monochromatic signals will have their spectrum smeared out in the frequency domain. Similarly any drift in sampling frequency will be mirrored by an apparent drift in the frequency of sinusoidal signals. However, for data lengths of order 1 hour the clock derived

from the MSF transmissions is sufficiently stable that timing variation effects are negligible compared to the effects of windowing of the data due to losses of lock of the interferometer.

4.4 Event Recording Program

Wideband gravitational wave detectors will produce large quantities of output which will require storage and analysis. For example, if a single signal is sampled at the rate of 10000 samples per second and with 16-bit resolution then about 1.6 Gbytes of storage capacity are required each day of operation to record this signal. In addition it will be necessary to record a number of control and veto signals along with the detector output, so that it will probably be necessary to store approximately 20 Gbytes of data every day of operation. High density recording media such as optical disk drives and magnetic tape systems based on video-recorder technology should provide sufficient capacity. However, the access times for optical disk drives may be too long for online recording at the data rates envisaged while tape based systems only allow sequential access to the data. (Some tape systems, such as the EXABYTE 8 mm videotape system, will allow fast searching operations for specially written file marks. It remains to be determined whether such marks can be written periodically between blocks of data in real time.) If a search for bursts of gravitational waves is being made, it is only necessary to record short sections of data that contain signals that are candidates for gravitational wave pulses. The selection of such events could be made on the basis of reasonably simple filters. From this recorded data it would be possible to use more complicated filters to determine the parameters of the signal more accurately.

Prompt detection of candidate events will be vital to allow searches to be made for optical displays associated with the event producing the gravitational wave, for example a supernova or a coalescing binary. It is essential, therefore, that the data is analysed online to detect bursts of gravitational waves and coalescing binary waveforms.

A program was developed to monitor the analogue input channels and to record data on to magnetic tape along with timing information whenever some amplitude threshold was exceeded. The program monitored two analogue signals, sampled at 6 kHz each, and if any sample in either of the two signals exceeded a preset threshold for that signal, then a section of data from both channels at the time of the event would be saved to magnetic tape, preceded and followed by at least 340 ms of data. The program consists of a few main elements: a clock routine; a loop to fill buffers with sampled analogue data and to monitor this incoming analogue data for candidate gravitational wave events; and a routine to save buffers to tape.

4.4.1 Internal Clock

To have accurate timing information available to the program, it was decided to maintain a Real Time Clock in the Nova which was controlled by the MSF transmissions. A digital Input/Output board was fitted to the Nova which could be configured to interrupt the processor whenever one of 8 digital input lines went through a transition of a given polarity. One of these lines was connected to a 1 kHz signal which was derived by dividing down the 6 MHz signal which was phase locked to the MSF transmissions. Another line was connected to the minute pulse output of the MSF receiver. By counting the transitions of the 1 kHz line the computer could maintain a Real Time Clock with an accuracy of 1 ms. Synchronisation was achieved via the minute pulse. On starting, the program prompted the user for the date and the time of the next minute pulse. The minute pulse interrupt was then enabled and the program entered a loop. On receiving the interrupt, a routine was entered which disabled the minute interrupt, enabled the 1 kHz interrupt and began to record data. On each subsequent interrupt produced by the 1 kHz line a routine would be entered which updated the clock before returning to the main program. The time was stored in four 16-bit words in the following format;

- WORD1 The year
- WORD2 The month and day, month in high byte, day in low byte.
- WORD3 The hour and minute, hour in high byte, minute in low byte.

- WORD4 Seconds and milliseconds (stored as a count from 0 to 59999).

One potential problem with this clock routine is that the clock may be updated while the time is being read by the main program. If the seconds word is read from the clock and the time is updated to the next minute before the remaining words are read then the time as read from the clock may be out, typically by about one minute. However the probability of such a clash occurring is rather small, and since at least three consecutive blocks, each of 340 ms, are saved to tape and only one of these blocks could possibly be affected by this problem, the time as recorded on the tape can always be determined unambiguously. It was therefore not felt necessary to develop a more sophisticated clock routine. Such a routine could have kept two copies of the time, allowing the program to read one copy (by changing a pointer for example) while updating the other copy.

4.4.2 Main Loop

The program maintained three buffers, which we will refer to as BUF1, BUF2 and BUF3, that were filled circularly. The buffers were able to hold 2040 words of sampled data. As one buffer was filling the previous buffer was checked for samples which exceeded the threshold and, if such a sample occurred, this buffer was saved to magnetic tape along with those immediately preceding and following it. At the start of each buffer were 4 words of timing information, taken from the MSF derived real time clock just before the first sample of the buffer. In addition, a flag is associated with each buffer which is set whenever the buffer is saved, and cleared whenever new data is read into the buffer. The routine which was called to save particular buffers to magnetic tape could therefore check whether the buffer had been saved previously. In this way a buffer was never saved to tape more than once.

Having set up the clock, the program initially recorded some blocks of data and calculated arbitrary thresholds for both of the ADC channels based on their signal levels. Both of the ADC's were therefore monitored independently. It was not possible to periodically recalculate these thresholds, for this calculation could not

be performed in real time and so recalculation would have led to dead time. (This was not a major limitation of the system since with the addition of some simple hardware it would be possible to obtain an analogue signal which represented the rms value of the input signals. A minor addition to the program could have digitised this value and used this to calculate the threshold.) A transfer of analogue data into BUF1 was then initiated and on its completion a transfer of data into BUF2 was begun. On the completion of this second transfer the main loop was entered. The main loop began a transfer into BUF3 and as BUF3 was filling, checked BUF2 for the existence of any samples above the threshold. If none were found then the program would wait until BUF3 had filled, begin a transfer into BUF1 and check BUF3, followed by a transfer into BUF2 while BUF1 was being checked, before returning to the beginning of the loop.

If a sample exceeding the threshold was found, for example in BUF2, then the program would call the tape writing routine to save BUF1 to magnetic tape, wait until this write operation was complete and then call the write routine to save BUF2; after starting the write operation, the program can continue while the transfer takes place. At this point the program checked that BUF3 was not yet completely full, for if it had been samples may have been lost before the next buffer of sampled data was begun. If BUF3 was full on returning from starting to write BUF2 to tape then an error was flagged, otherwise the program waited until BUF3 filled, then began a transfer of sampled data into BUF1, waited until BUF2 had been saved to tape and then saved BUF3. While BUF3 was being saved the main loop could be re-entered at the point where BUF3 is checked. If any signals exceeding the threshold were found, this saving procedure was repeated with BUF3 replacing BUF2, BUF1 replacing BUF3 and BUF2 replacing BUF1. Since the tape writing routine ensures that no time is lost by rewriting a buffer to tape if it has already been saved, the program can run continuously with no dead time. In fact the program can be used to record data continuously simply by setting a threshold of zero.

4.5 Recording Runs after SN1987A

Shortly after the optical detection of supernova SN1987A in the Large Magellanic Cloud it was decided to record a few hours data from the three prototype laser interferometer detectors at Glasgow, Caltech and MIT. It is highly unlikely that gravitational waves of sufficient energy to be detectable would be being emitted by the supernova remnant. However, this data gathering run was the first coincidence run using interferometer detectors and it was hoped that the analysis of the data would provide useful feedback on the design of the detector. In total $5\frac{1}{2}$ hours of data were recorded at Glasgow, between the 5^{th} and 11^{th} of March 1987. Of these, 4 hours were recorded in coincidence with the prototype detector at Caltech with $2\frac{1}{2}$ hours of these recorded in coincidence with MIT also. The data was recorded using the event recording program described above, operated in continuous mode by setting a threshold of zero. The two signals recorded were derived from the laser frequency control system and the secondary cavity arm length control system. Unfortunately hardware failures in the data acquisition system shortly before the recording runs introduced two limitations. The first was that the second least significant bit of the ADC connected to the secondary feedback signal became fixed in a low state. This reduces the resolution of the secondary signal ADC to 6-bits, giving a dynamic range for a signal recorded at one fifth the full range of approximately 26 dB. The second limitation was caused by hardware problems with the disk drives on which the software was stored. It was necessary to use software which had been stored on magnetic tape and regenerate the executable files using a magnetic tape operating system. Unfortunately this software did not include a version of the buffer writing routine which allowed for the four extra words of timing information, so that four words were lost from the end of every buffer. Since each 16-bit word consists of 8-bit samples from the two ADC's, the result was that 4 samples were lost every 2040 samples from each of the two recorded signals.

Each of the two signals was sampled at 6 kHz. This gives a time between tape changes of just under 30 minutes. The signals were filtered to avoid aliasing of high frequency signals into the bandwidth of interest. The filters used were commercial

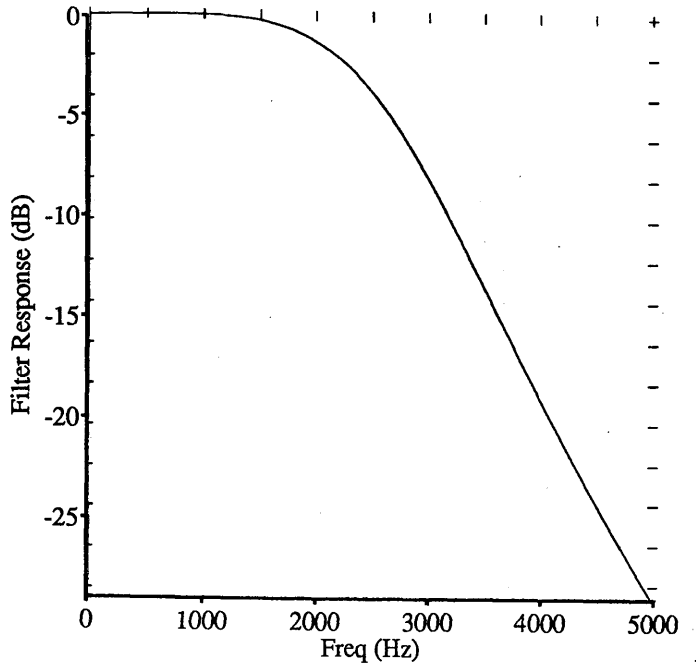

Figure 4.3: The amplitude response of the active anti-aliasing filters.

active filters with a response equivalent to a 6-pole Butterworth filter with a −3 dB frequency of 2.5 kHz. The transfer function of these filters is shown in Figure 4.3. This filter provides about 11 dB reduction of signals around 3.5 kHz compared to the signals at 2.5 kHz and so further passive filtering was added, and since the output noise of the detector rises sharply at low frequencies and large amplitude resonances at about 400 Hz are present, high pass filtering above a few hundred hertz was also introduced. The transfer function of the bandpass filtering is plotted in Figure 4.4. When combined with the natural fall off of the signal at frequencies above 1200 Hz this is sufficient to ensure that the error introduced below 2.5 kHz is less than a few percent.

4.5.1 The Laser Frequency Control System

The laser in the Glasgow prototype is frequency locked to maintain resonance with the primary cavity. While subtraction of the common mode signal in the two cavities due to residual laser frequency noise is performed, any frequency variations remaining, such as might be produced by a resonance of the laser cavity or by a

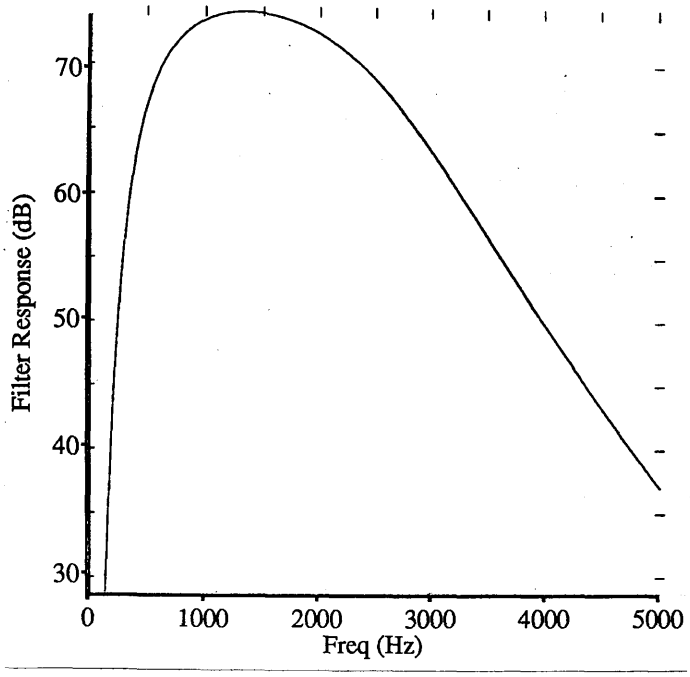

Figure 4.4: The amplitude response of the bandpass filtering applied to the secondary control system output.

laser oscillation will produce an apparent gravitational wave signal. For this reason it was decided to record the error point of the laser frequency stabilisation system. Periodic laser frequency fluctuations which would generate an apparent periodic gravitational wave signal should appear on this primary cavity error point, so that comparison of the spectrum of this signal with that of the gravitational wave signal would enable such spurious signals to be detected. However the loop gain of the laser frequency control system at signal frequencies is high, of order 10^{10}, and so the error point signal is rather small and the recorded primary signal is dominated by noise, predominantly due to line pick-up. While periodic variations may be detectable in the power spectrum, this signal is not able to veto broadband bursts, which may be caused for example by higher order laser modes changing the apparent laser frequency or by saturation of the amplifiers, and which will be indistinguishable from a gravitational wave signal.

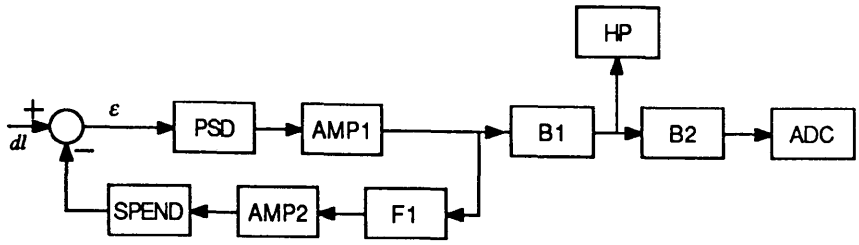

Figure 4.5: Schematic diagram of the secondary cavity length control system. (The elements are described in Section 4.5.2.)

The Secondary Cavity Control System

When the data acquisition system was developed the feedback system used for controlling the secondary arm length had the form shown in Figure 4.5. The resonance condition of the cavity was monitored by an RF reflection locking scheme which gives an output that, for small offsets, is proportional to the phase difference between the light from the cavity and the light from the laser. This phase sensitive detector therefore gives an output which is a direct measure of the offset of the cavity from resonance. This signal and the corresponding signal from the primary cavity were summed with a suitable phase correction to cancel the effects of residual laser frequency noise and the resulting signal was amplified and filtered. Forces could be applied to one of the test masses in the secondary cavity to which was bonded a multiturn coil. It was arranged that this coil sat in the gap of a loudspeaker magnet. By passing electrical current through the coil, correcting forces could be applied to the test mass to adjust the secondary cavity length to maintain resonance of the laser and the secondary cavity. By suitably filtering the relative length difference signal, a stable feedback system was possible with a unity gain frequency of order 1 kHz.

Ideally the signal recorded would be a direct measure of the relative length difference between the two arms. While such a signal can be obtained relatively easily if the unity gain frequency of the control system is well above the range of signal frequencies—for example in a Delay Line detector the null output condition

can be maintained by a control system which uses an Electro-Optic Modulator and which is inherently wideband— this signal is rather difficult to obtain in the case of the present prototype detector since the length of the secondary cavity is adjusted to maintain resonance with the laser light by controlling the position of the end test mass. Since the control system used for this task at Glasgow has a unity gain frequency of order 1 kHz, the relative length fluctuations at frequencies below 1 kHz are approximately reduced by a factor of the loop gain while relative length fluctuations at frequencies well above the unity gain frequency are unchanged by the control system. Signals at frequencies close to the unity gain frequency will be transformed in a way that depends critically on the phase margin of the feedback system. Therefore the recorded signal will effectively be the relative length difference signal passed through some linear filter whose transfer function depends on the characteristics of the secondary length feedback system.

In principle, knowing the transfer function of each element of the servo system, the recorded signal can be deconvolved to obtain the relative length signal. However, practical problems place a limit on the accuracy which can be achieved. For example the loop gain of the servo system depends on a number of experimental variables such as the laser intensity, the fringe visibility and the cavity storage times. While the cavity storage time should remain constant over the observation periods, since it is dependent mainly on the parameters of the mirrors, and the laser light is intensity stabilised, the values of these parameters are not usually known and visibility variations can be expected to change the loop gain of the system on timescales of a few minutes, varying for example as the alignment of the cavities changes. The servo system is conditionally stable, in order to achieve the high loop gains required at low frequencies to maintain the cavity sufficiently close to resonance. In practice the unity gain frequency must be kept at around 1 kHz to obtain a stable system, and so the control system electronics include variable components which can be adjusted to control the loop gain. The response of these electronic circuits is therefore also unknown. It can be seen then that there is no simple way to deduce the gravitational wave output from the detector's output alone and that calibration across the bandwidth of interest is necessary.

Calibration

During the recording runs after SN1987A, calibration was provided by modulating the length of the primary cavity. These modulations induce known relative length changes which appear on the output of the detector, allowing the response of the detector to be determined. The modulation signal was derived by triply differentiating a square-wave voltage waveform which was then applied to the coils of the orientation control system described in Chapter 3. Since the amplitude of a harmonic of a square waveform is inversely proportional to the frequency of the harmonic, one stage of differentiation should produce a series of harmonics of equal amplitudes. The other two stages of differentiation compensate for the response of the pendulum to forces above the resonant frequency of the pendulum. A square wave of 250 Hz was used which therefore provided calibration signals at odd harmonics of 250 Hz. The calibration peaks must have an amplitude which is great enough that they are easily detectable above the background noise and their amplitude must be relatively insensitive to this noise. For example, if data containing calibration peaks is transformed and a peak is 30 dB above the background noise then its amplitude is varied by only about 3%. However if the calibration peaks are always present in the recorded data their amplitude will always be greater than the noise and it would be necessary to remove them from the data before any examination of the time series could be made for burst signals. For this reason the calibration peaks must be applied periodically for relatively short intervals. During the recording runs the calibration peaks were applied to the detector for one second every 100 seconds. A typical spectrum of the output from the detector is shown in Figure 4.6 showing the calibration signals which were applied during the recording runs.

Artificial Chirps

Signals from coalescing binary systems are currently thought to be the most likely to be detected by long baseline interferometers. It was shown in Chapter 1 that these systems produce "chirp" waveforms which consist of a large number of os-

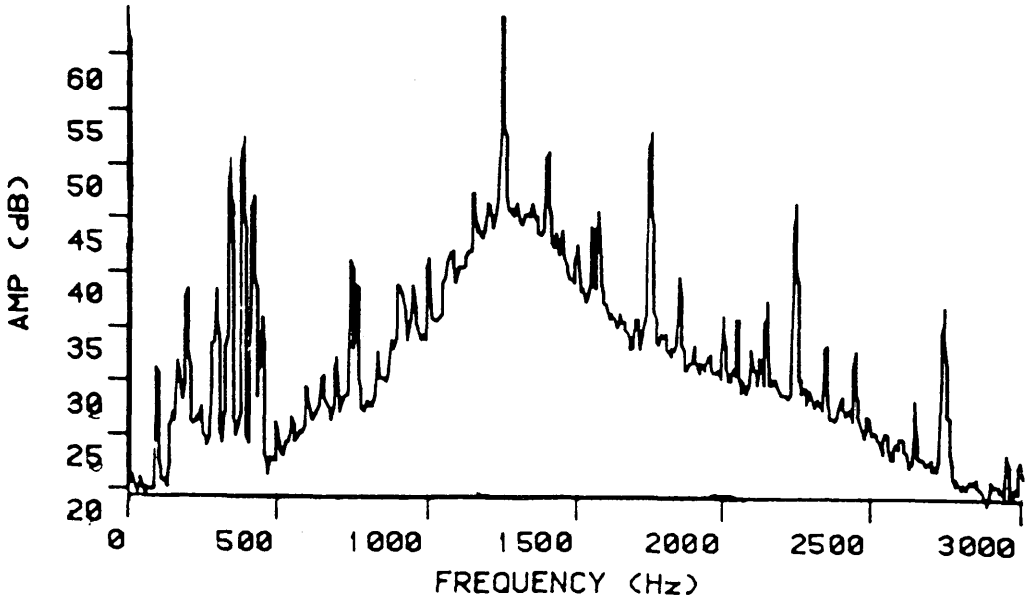

Figure 4.6: Spectrum of the detector output.

cillations, typically ≥ 100, whose frequency and amplitude increases with time. In order that tests of algorithms to detect these signals might be performed, such signals were generated using a digital computer and the waveform output using a digital to analogue converter. These signals modulated the length of the primary cavity in the same manner as the calibration signals at a range of amplitudes. Each signal corresponded to that emitted by a system of two $1.0\,M_{\odot}$ stars, and was speeded up by a factor of 3.75 to produce a signal in the bandwidth of the current prototype. (This speed up can be regarded as a Doppler shift of the signal and it can be seen from Equation 1.29 that such a shifted waveform is equivalent to an unshifted waveform with a different mass parameter.) The amplitude of the chirp waveform was varied during the recording runs. The waveform produced by the chirp generator was also added to the signal from the laser frequency control system and recorded. This allowed a comparison between the position of the signal in the time series as detected by suitable matched filters and the true position to be made.

4.5.2 Deconvolution of the Detector's Response

To retrieve the true gravitational wave signal from the recorded signal it will be necessary to filter the digital signal. Digital filters fall into two classes: non-recursive, or Finite Impulse Response (FIR), filters, which generate an output series by calculating weighted sums of an input series so that they perform a convolution; and recursive, or Infinite Impulse Response (IIR), filters which generate the output series by calculating a weighted sum of the input series and previous values of the output series. In general, recursive filters are more efficient than non-recursive filters when the convolution is performed directly. However the synthesis of recursive digital filters for an arbitrary transfer function can involve considerable computation, while non-recursive filters can be constructed relatively easily in the frequency domain. A non-recursive filter can also be implemented very efficiently using the convolution properties of the Fourier transform and the FFT. It was therefore decided to implement a non-recursive digital filter to deconvolve the recorded data. Before such a filter can be constructed we must determine the desired transfer function. This should be the inverse of the combined transfer function of the secondary cavity control system and the recording system.

From the calibration peaks we can only derive the response of the detector and recording system at 5 frequencies which are spaced by 500 Hz. Since we have more than five variable elements in the secondary cavity control system, it is not possible to obtain a unique solution to this this problem. Additional problems occur due to the presence of delays in the control system which cannot be accurately modelled. For example, these include propagation delays in electronic components and motion of the magnet due to reaction forces. Also, since the mirror and coil were on opposite faces of the test mass and the coil was fixed to a cylindrical former a few centimetres long, there will be a delay between motions of the coil and motions of the mirror. (In principle, some of these effects could be determined experimentally.)

Ignoring most of these effects, the transfer function of the detector can be calculated. The phase sensitive detector (PSD) produces a voltage that depends on the frequency difference between the laser light and the cavity eigenmode and can be modelled as a linear system with a single pole roll-off having a time constant which

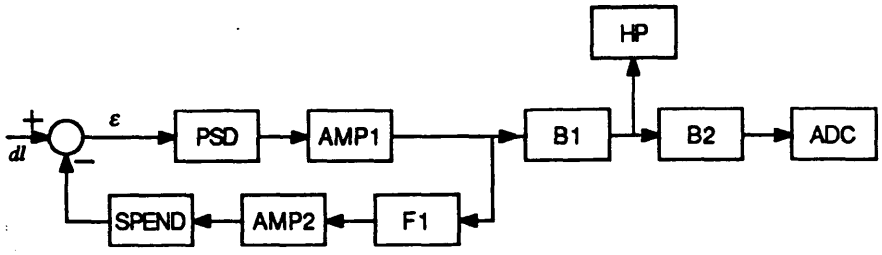

Figure 4.5: Schematic diagram of the secondary cavity length control system. (The elements are described in Section 4.5.2.)

is simply related to the cavity storage time. The gain of the PSD is dependent on the laser light intensity, the visibility of the fringes and the cavity linewidth. The linewidth is also simply related to the cavity storage time. Therefore we can adopt a model where the PSD has a transfer function relating the voltage output to the relative length difference between the two cavities of the form

$$\mathrm{PSD}(s) = \frac{k}{1 + s\tau} \tag{4.9}$$

where k is a parameter specifying the gain, of order 10^{12} volts/metre, τ is the time constant associated with the cavity storage time, typically around 100 μs and $s = i\omega$ where $i = \sqrt{-1}$ and ω is the angular frequency.

The output of the PSD was amplified and filtered as shown in Figure 4.5. AMP1 is a voltage amplifier based on a commercial operational amplifier. AMP2 is a voltage controlled current source. F1 represents passive filters which can be adjusted to control the stability of the control system. B1 and B2 are bandpass filters and amplifiers and HP represents a commercial spectrum analyser which was used to monitor the output signal during the recording runs.

The response of the test mass to the forces generated by the output current of AMP2 was modelled simply as a driven, damped, harmonic oscillator. So that

$$\mathrm{SPEND}(s) = \frac{C/m}{s^2 + \gamma s + \omega_0^2}, \tag{4.10}$$

where C is a constant defining the coupling strength, γ is the damping coefficient, ω_0 is the natural resonant frequency and m is the mass of the test mass. Allowance for a time delay due to propagation in the test mass was also modelled.

From the equations for the transfer functions of these elements an expression relating the recorded signal to the gravitational wave signal can be obtained. From the calibration signals the value of this expression can be obtained at the five calibration frequencies. Since the gain of the system formed by the detector and recording system changes with time, each period when the calibration peaks were applied gives different values for the transfer function. By varying the values of the unknown terms, it should be possible to vary the transfer function to agree with

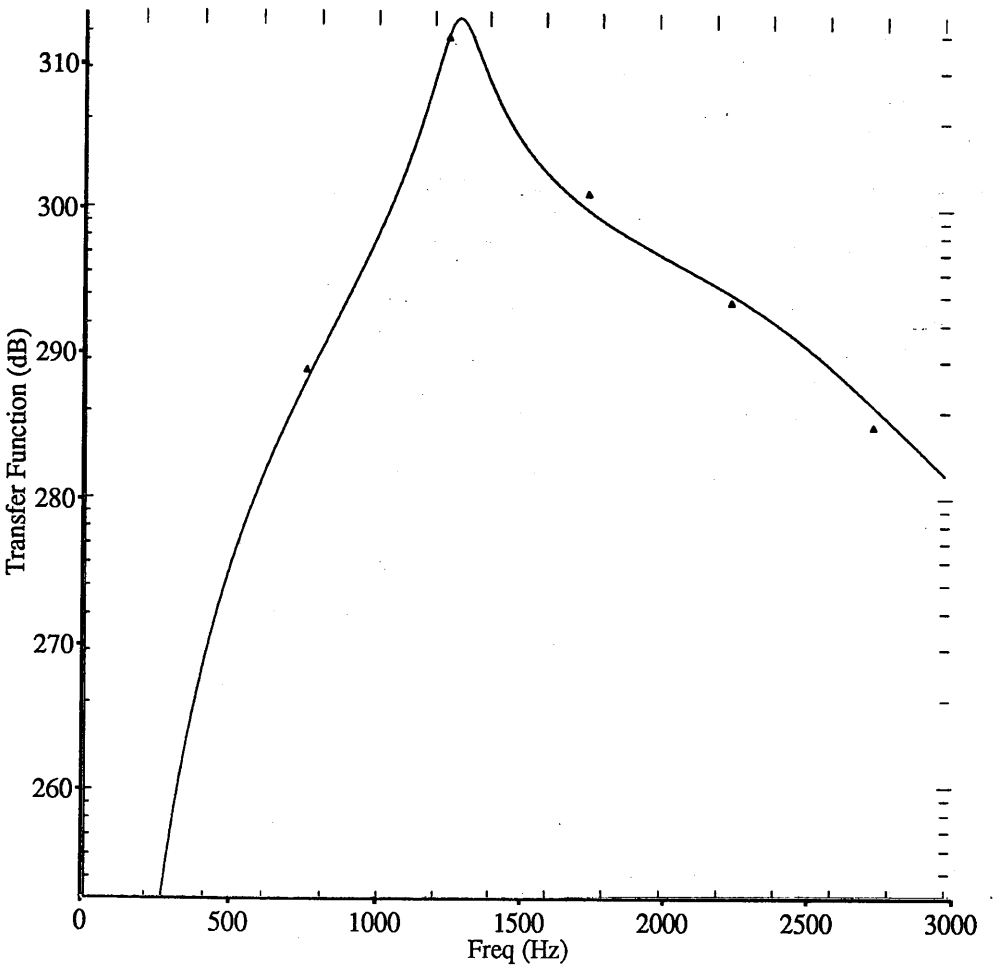

Figure 4.7: Results of the model for the recording system's transfer function. Also shown are the five values derived from one set of calibration peaks.

these values. However this problem is rather unstable and because of the problems discussed above and due to component tolerances the solution can be at best an approximation. An example of the transfer function which can be obtained in this way is shown in Figure 4.7 along with one set of calibration peaks.

The model can simulate the transfer function of the system fairly well at frequencies below 800 Hz and above 1500 Hz, but at frequencies around the unity gain frequency of $\sim$ 1100 Hz the system response depends critically on the phase delay and the model loses accuracy. In addition, since there are more parameters than known values, it is difficult to automate the determination of the optimum values

of the parameters.

One advantage of the application of the chirp waveforms, however, is that they can be used to calibrate the detector's amplitude response up to 1 kHz and the phase response up to 2.5 kHz. From a section of data containing such a chirp we can derive the transfer function of the detector and recording system by calculating the ratio of the discrete Fourier transforms of the recorded secondary cavity signal and the length changes of the primary cavity produced by the chirp waveform. To reduce the effects of noise this transfer function should be smoothed. This frequency domain specification can therefore be used to construct a digital filter.

Filtering a time series with a filter whose transfer function is $H(f)$ is equivalent to performing the convolution of the time series with the impulse response of the filter. We therefore wish to calculate the discrete convolution of the time series, $f(t)$, with the filter impulse response, $g(t)$. The impulse response of the filter should be windowed, for example with a Hanning filter, to reduce end effects [Brigham 1988]. (Each of the discrete frequency bins of the DFT has a response which is a sinc function where $\mathrm{sinc}(x)=\sin \pi x/\pi x$ and so has an effective width of one bin but a response that falls only as $1/x$. The effect of the window function is to increase the roll-off of each bin, and so improve the filter's response.) We therefore wish to calculate the discrete convolution

$$h(k) = \sum_{j=0}^{M-1} f(k-j)g(j), \tag{4.11}$$

where M is the number of points in the impulse response.

Convolution of two time series can be performed by multiplying their Fourier transforms and inverse transforming the result. Such a procedure is more efficient than performing the convolution directly because of the existence of efficient algorithms for the DFT. However the DFT will calculate the circular convolution of two time series; that is each of the indices in Equation 4.11 is taken MOD(M). This effect can be avoided by padding out the filter with zeros which cause no effect when they wrap around due to the circular convolution.

In fact for a data length N and a filter of length M padded to N with zeros, the inverse transform of the product of the discrete Fourier transforms of the signal

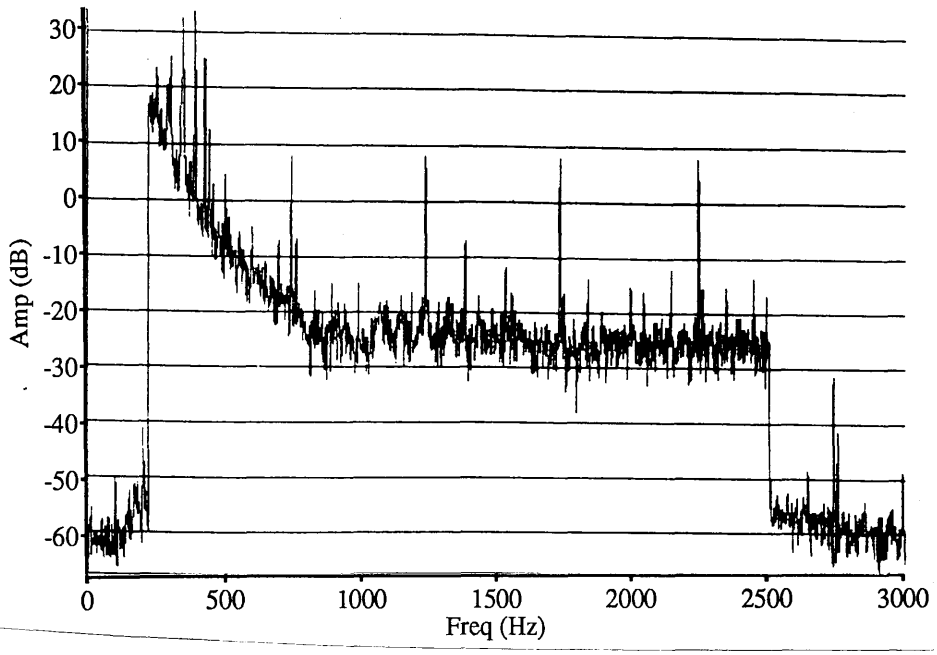

Figure 4.8: Spectrum of the output of the digital filter described in Section 4.5.2. (A bandpass from 250 Hz to 2500 Hz has also been applied.)

and filter contains $N - M + 1$ true convolution points. Therefore the time series to be filtered should be sectioned into blocks of length N, separated by $N - M + 1$, so that there is overlap between sections. The discrete fourier transform of some data which was filtered in this way using a filter of length M=4096, tapered with a window function $0.5(1 - \cos(2\pi j/M))$, and a transform of length N=32768, is shown in Figure 4.8. It is therefore possible to derive the gravitational wave signal from the recorded data using such a digital filter if the detector's transfer function can be derived.

4.5.3 Searches for Loss of Lock

Occasionally the feedback systems controlling the laser frequency and the secondary cavity's length are unable to maintain lock. This can be caused for example by thermal expansion of the laser's resonator or simply by large impulses applied

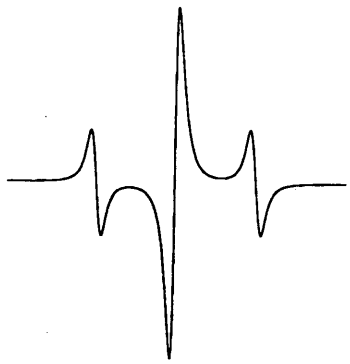

Figure 4.9: Schematic diagram of the output of the PSD against relative length differences between the two cavities.

suddenly to the test masses. The detector may regain lock automatically but more usually operator intervention will be required. While the detector is unlocked no useful data can be recorded, instead spurious signals are recorded.

When the detector loses lock the laser frequency and the secondary arm length will vary rapidly with respect to each other. The output of the RF reflection locking scheme as a function of the offset from resonance is as shown in Figure 4.9 [Drever *et al.* 1985]. Since resonance can occur at frequencies which are a free spectral range of the cavity apart, corresponding to an arm length change of 2.57×10^{-7} m the offset will sweep rapidly through a number of these resonances. A typical section of unlocked data is shown in Figure 4.10 while a typical section of locked data is shown in Figure 4.11. The most notable characteristic of the unlocked data is that the output is close to zero for many continuous samples. The status of the detector can therefore be determined from the recorded data simply by searching for this feature.

During the data runs after Supernova 1987A, the detector was monitored, and the times of losses and reacquisiton of lock were noted. A program was written to search this data for such consecutive low amplitude samples. An arbitrary choice was made that a loss of lock would be flagged when 30 consecutive samples less than $\sigma/2$ were detected, where σ is a typical value for the standard deviation of the noise when the detector is locked. If the samples were uncorrelated then this would have a probability of occurring by chance in any group of 30 samples

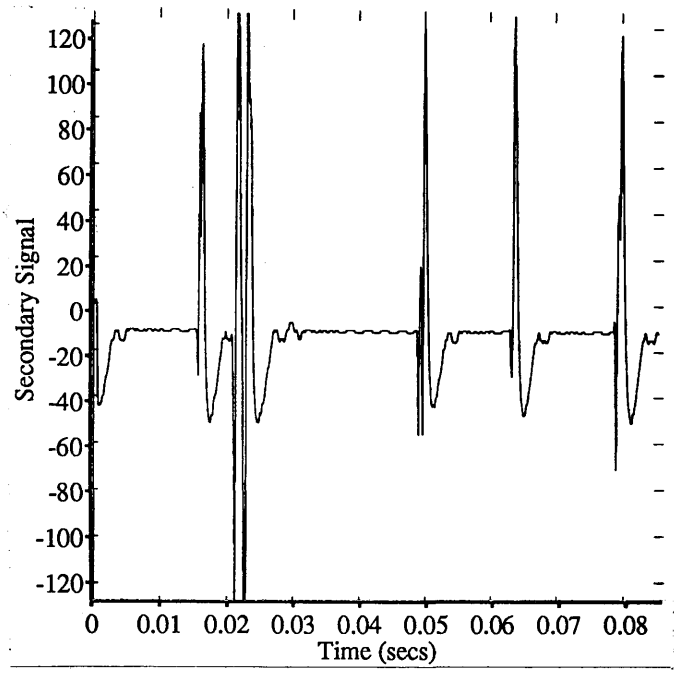

Figure 4.10: A section of the detector's output when the detector was unlocked.

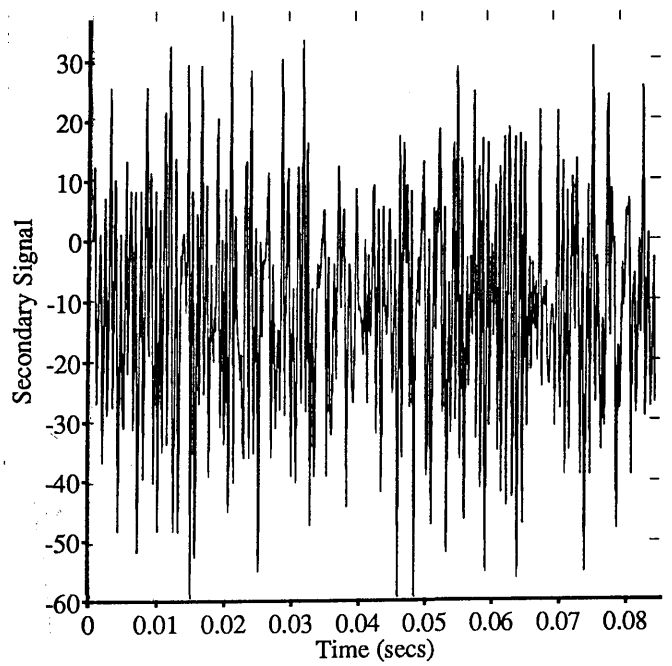

Figure 4.11: A section of the detector's output when the detector was locked.

of 3×10^{-13}. However correlation between samples introduced by the spectral response of the detector and recording system may reduce the probability of such spurious detections of loss of lock below this level.

Similarly reacquisition was flagged whenever a burst of 10 consecutive low amplitude samples had not occurred for 4000 samples. This allows a certain amount of dead time for the detector to settle after reacquisiton of lock. For safety some dead time was also allowed before a loss of lock.

When used to analyse the data recorded during the observations after Supernova 1987A the program successfully detected those losses of lock that had been recorded during the data runs with no spurious losses reported. This simple technique can therefore be used to detect periods of loss of lock, so that spurious signals can be removed from consideration.

4.6 Conclusions

Since the present sensitivity of the prototype laser interferometric gravitational wave detectors is not sufficient to give a reasonable chance of detecting emissions from astrophysical sources, there is currently little effort to perform searches for gravitational waves. Even after a supernova such as SN1987A there is currently little chance of detecting gravitational waves. However, recording and analysis of data allows the techniques which will be used on the larger, more sensitive, instruments to be developed and tested and provides a useful check on the performance of the prototype detectors.

The output from the current prototype detector does not give the exact gravitational wave signal, and because of varying elements in the control system and unknown tolerances it is difficult to derive this from this signal alone. The variable elements, however, fall into two types; electronic components which are adjusted by the operator to keep the control system stable and variations associated with the laser and cavities which will vary the two parameters in Equation 4.9. The problem of unknown values of the electronic components can be eliminated by recording both the input signal and the output signal to these amplifiers and so

obtain the transfer function experimentally. If the response of the mass to current in the coil can be determined experimentally then it should be possible to reduce the problem to determining the two parameters in Equation 4.9. This would in principle be possible using calibration signals at two frequencies, one above and one below the cavity corner frequency.

The following two chapters will describe some analysis of the recorded data which has been performed and which quantifies the noise performance of the Glasgow prototype.

Chapter 5

The Noise Statistics of the 10 metre Prototype

5.1 Introduction

Some of the sources that should produce gravity wave bursts were described in Chapter 1. It was shown that most of these sources will produce pulses with a duration of about 1 millisecond. Such pulses must be detected against a background of noise generated by processes within the detector. This chapter will describe the methods which are used for pulse searches and will give details of the performance of the 10 metre prototype detector at Glasgow during the data runs following SN1987A.

5.2 Expected Pulse Height Distribution

Many of the noise sources in laser interferometer gravitational wave detectors are assumed to be Gaussian in nature, that is their amplitude has a Gaussian distribution. If $p(V)dV$ is the probability that the instantaneous amplitude lies between V and $V + dV$ and the rms amplitude is V_{rms} then,

$$p(V) = \frac{1}{V_{rms}\sqrt{2\pi}} \exp\left(-\frac{V^2}{2V_{rms}^2}\right). \tag{5.1}$$

Noise sources that are Gaussian include shot noise, thermal noise, radiation pressure, laser intensity noise and refractive index fluctuations. In fact any noise source which consists of the sum of many independent, random elements will be a good

approximation to a Gaussian noise source. However it is possible that other noise sources will be present in the detector which are not Gaussian, for example seismic disturbances and mechanical coupling of external vibrations, laser frequency fluctuations due to mode-hopping in the laser and acoustic pick-up. Additionally, the servo electronics may occasionally saturate, giving large amplitude deviations on the output of the detector. It is hoped that interferometer detectors will be limited by a Gaussian noise source, shot noise, and since the statistics of this distribution are well understood the significance of particular events in the output of the detector can be calculated. Occasional non-Gaussian events may be tolerable if two independent detectors are working in coincidence and each can therefore mask the other's spurious signals. However, if the occurrence of these non-Gaussian events is too frequent, the sensitivity at which a strong claim to have detected gravitational waves can be made will be decreased. It is therefore important to measure the performance of the detector to ensure that the assumptions which are made about the probability distribution are valid. In addition such measurements may also aid the identification of such non-Gaussian noise sources, leading to their elimination, either by redesign of the relevant section of the detector or by recording suitable veto signals which can be used to mask out spurious events.

5.3 Pulse Searches

It is possible to show that for any signal in the presence of Gaussian white noise, the linear filter which maximises the received signal to noise power ratio is the matched filter [Brigham 1988]. This filter generates an output time series which is the correlation of the input time series with the signal waveform. The correlation of two signals, $f(t)$ and $g(t)$, is defined to be

$$z(t) = \int_{-\infty}^{\infty} f(u)g(t+u)du. \tag{5.2}$$

Comparison of Equations 4.5 and 5.2 shows that the correlation is just the convolution of $f(t)$ with the time reversed $g(t)$, so that application of the convolution theorem gives the result

$$Z(s) = F(s)G^*(s) \tag{5.3}$$

where $Z(s)$, $F(s)$ and $G(s)$ are the Fourier transforms of $z(t)$, $f(t)$ and $g(t)$ respectively, and $*$ denotes the complex conjugate. The optimum filter is therefore one which has a transfer function equal to the Fourier transform of the time reversed signal.

A number of idealised pulse shapes and the amplitudes of their Fourier transforms are shown in Figure 5.1. (For a symmetrical pulse the phase response will be linear with frequency.) For millisecond pulses we may therefore assume that the filter should have a Fourier transform which is centred about 1 kHz and which has a bandwidth of 1 kHz.

5.4 Recorded Data

The response of the detector and recording system used during the recording runs after SN1987A is approximately as shown in Figure 4.7. It therefore has the feature that the gain peaks around 1 kHz, and the phase versus frequency relation is almost linear around 1 kHz. The recorded data therefore represents the output of a filter which is a reasonable approximation to a filter for millisecond pulses and so we may examine the recorded data for millisecond pulses directly. Obviously this filter is not ideal. However, unless an assumption about the exact pulse shape to be searched for is made, there is little point in improving the filter response since it should be adequate for the detection of non-Gaussian events. In practice, a series of such filters would be required for a range of different pulse shapes.

The output of a filter with a bandwidth of 1 kHz will be correlated over a timescale of 1 ms, and so the recorded time series should be sampled at a rate above 2 kHz. (There is some advantage in sampling at higher data rates to be sure of sampling close to the peak of a millisecond pulse. These samples will not be independent however.)

A limitation of this system is that the amplitude of the recorded signal will change due to variations in the transfer function of the secondary cavity control system. (Typically by $\pm$ 25% at 1250 Hz.) During the recording runs the response of the detector was monitored and adjusted manually to maintain a stable control

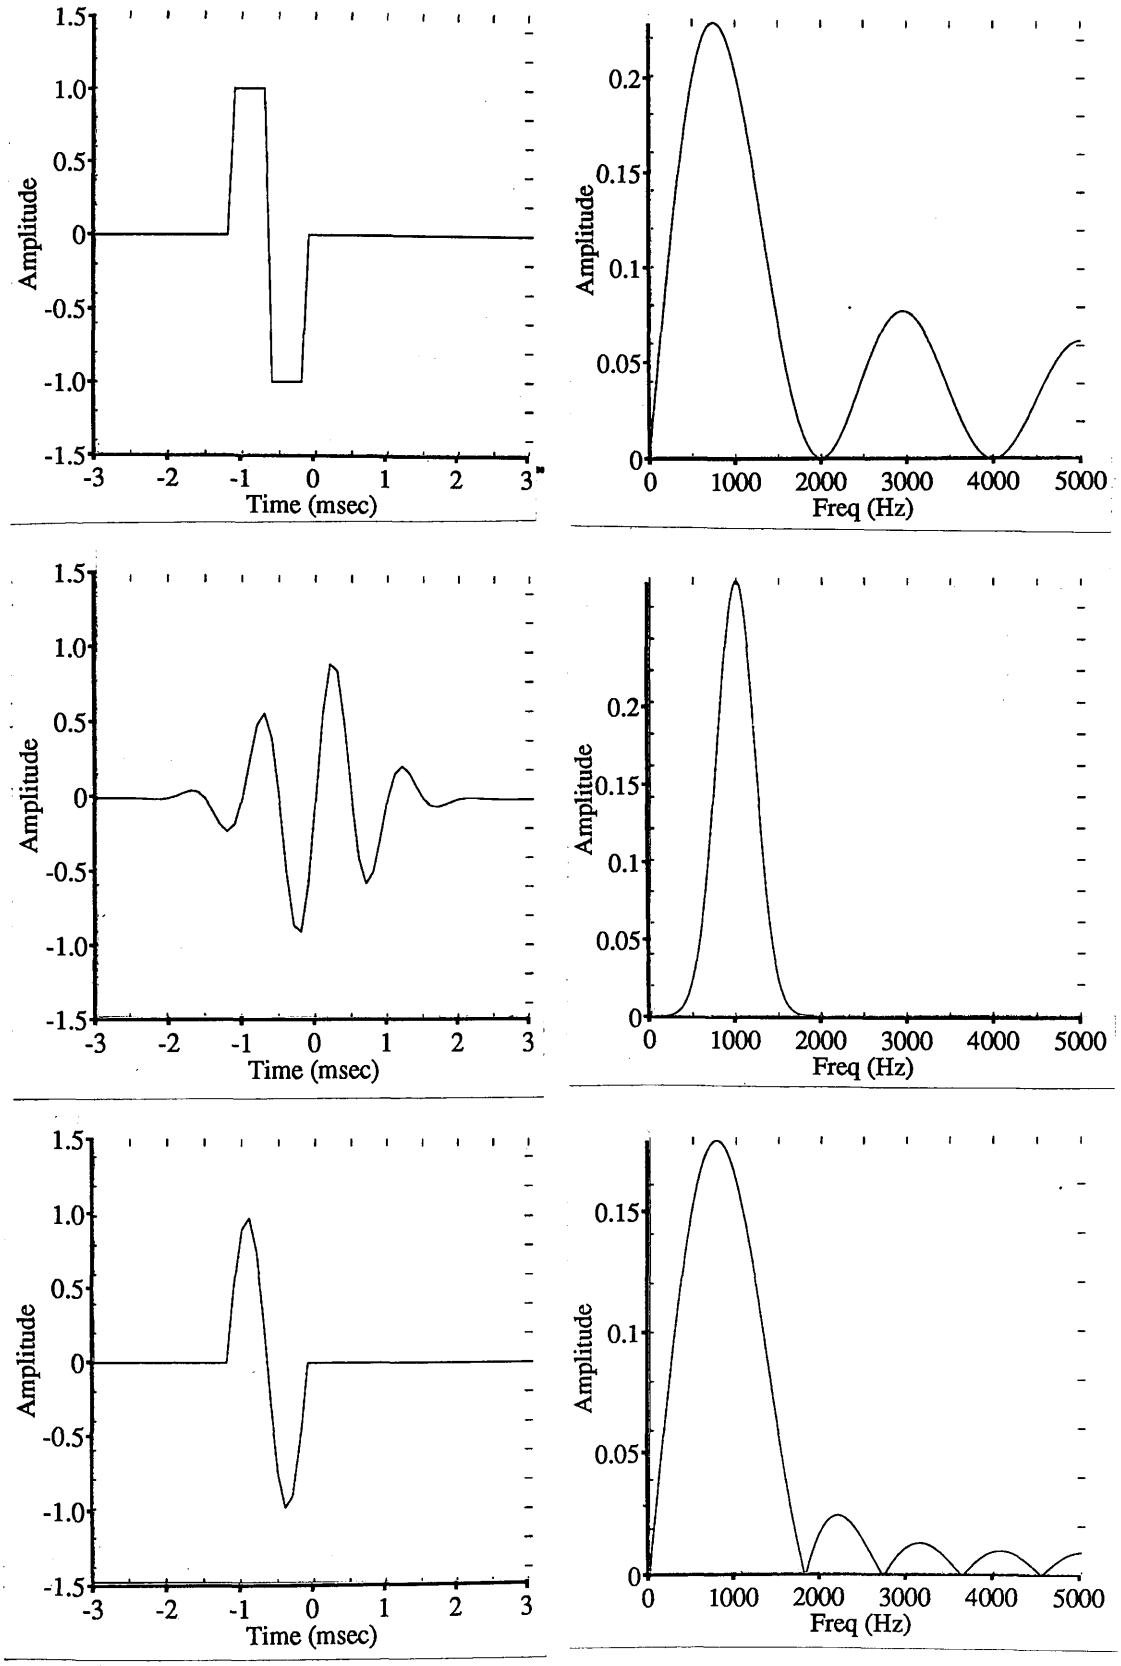

Figure 5.1: Some idealised pulse waveforms and their Fourier transforms.

system, as close to being critically damped as possible, and so the shape of transfer function is relatively unchanging with time. However, the recorded signal is derived from one point in this control system, and although an increase in the gain of elements before this point can be offset by changing the gain of elements following this point to maintain the same servo response, the recorded signal will still vary in amplitude. Since most of these variable elements simply change the gain over the frequency range of interest, we shall assume that the variation in the recorded signal with time is a frequency independent gain change. Such an assumption greatly simplifies the problem of coping with the time dependence of the recorded signal, since we may simply calculate the rms value of the signal over short timescales and normalise the data. If we calculate the rms value using 20000 independent samples, that is 20 seconds of data, then the statistical error in the estimate of the rms value is of order 1 part in 200, which is more than adequate, and the estimate is insensitive to large amplitude pulses. The value of 20 seconds is also small enough that this estimate should be able to track slow variations in the gain with time, due for example to cavity alignment variations. In this way we can generate a time series which should represent, to adequate accuracy, the output of a filter for millisecond pulses.

5.5 Pulse Height Distribution

The amplitude distribution of the normalised data can be compared with that expected for Gaussian noise. This is most easily achieved by noting that for $x = V^2$ and normalised data ($V_{rms} = 1$), we have, from Equation 5.1,

$$p(x) = \frac{1}{\sqrt{2\pi x}} \exp\left(-\frac{x}{2}\right), \tag{5.4}$$

or $-2\ln\left(p(x)\sqrt{2\pi x}\right) = x$, so that a graph of $\ln\left(p(x)\sqrt{x}\right)$ against x should be a straight line.

A histogram of the amplitude distribution of ~ 30 minutes of data (avoiding periods when calibration signals were applied, when the detector was unlocked and when the recorded signal saturated) is plotted in this way in Figure 5.2. The

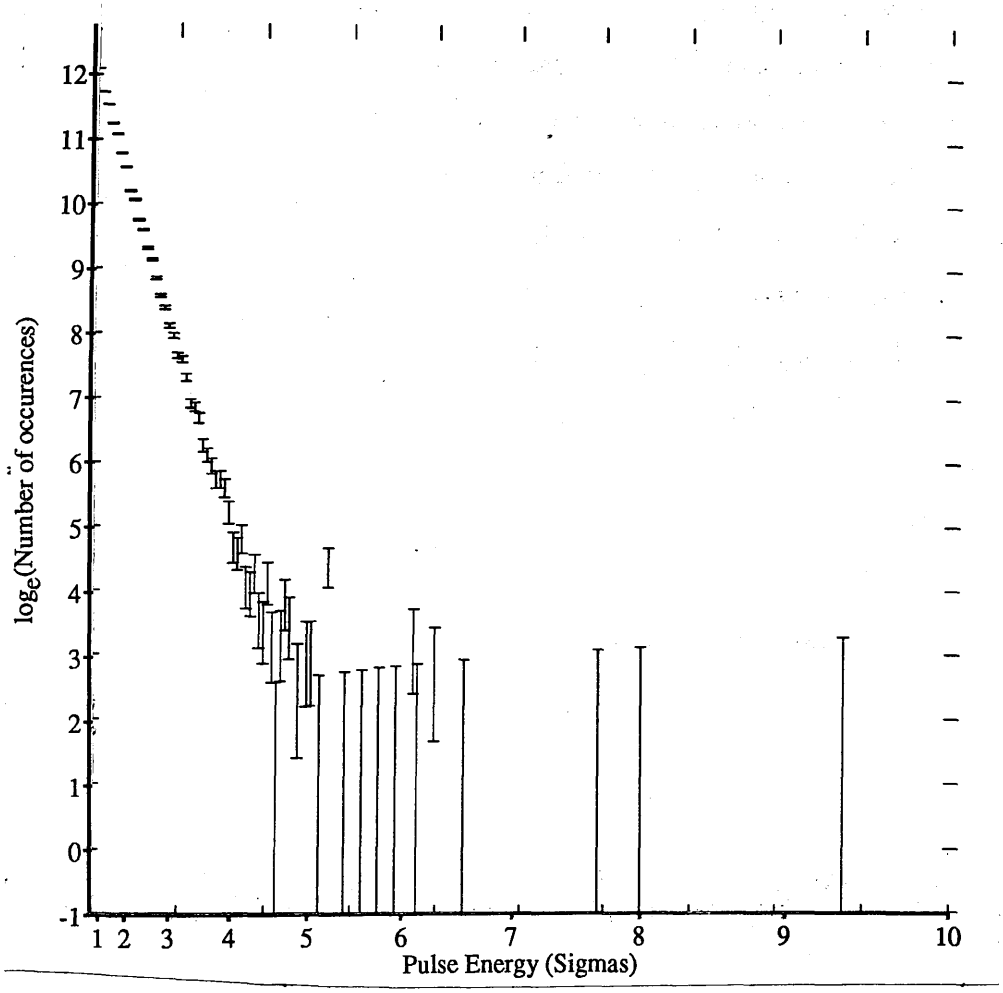

Figure 5.2: The pulse height distribution of the detector's output.)

distribution appears to be Gaussian up to about 5σ, above which a non-Gaussian tail is evident. The samples which are greater than 5σ are due to 14 discrete pulses. When these high amplitude events were examined, it was found that 8 were due to short periods (a few seconds) when the amplitude of the recorded signal is larger than usual. These pulses are not very much larger than those surrounding them and so can be disregarded as being due to fast gain variations which are not followed by the normalisation routine. It was found that 8 of the 14 pulses occur just before a loss of lock and 5 pulses occur within 5 seconds of each other. Since it was believed that the losses of lock were due to saturation of the amplifier driving the piezo-electric transducers controlling the laser cavity length, which were unable to track the thermal expansion of the laser resonator over long periods, it is likely that these pulses are also due to brief periods of amplifier saturation. If we disregard the clusters of pulses, and those within 10 seconds of a loss or regain of lock (when the detector will still be settling down), we are left with 3 pulses each with an amplitude greater than 6σ.

We can consider the time series as being due to a Gaussian background with non-Gaussian pulses added. The underlying noise statistics of the interferometer can be examined more easily by removing these samples from the histogram. Since there are no samples with amplitudes between 5σ and 6.0σ except for those produced by ringing of the larger amplitude pulses, these non-Gaussian pulses are easily identified. The underlying noise distribution is shown in Figure 5.3, and it can be seen that it is indeed Gaussian. This suggests that there is no evidence for non-Gaussian noise which produces pulses less than 5σ. And that the number of non-Gaussian events above a threshold is constant for thresholds up to 6σ.

The occurrence of such non-Gaussian pulses is not a major limitation of the prototype detector, for if, as believed, most of the pulses, if not all, are the result of amplifier saturation, then either the amplifiers can be given higher dynamic range to avoid saturation, or the amplifier output can be monitored and recorded as a veto on the detector's output. In fact, provided the source of a pulse can be identified, it can be disregarded as a candidate gravitational wave pulse with little loss to the efficiency of the detector.

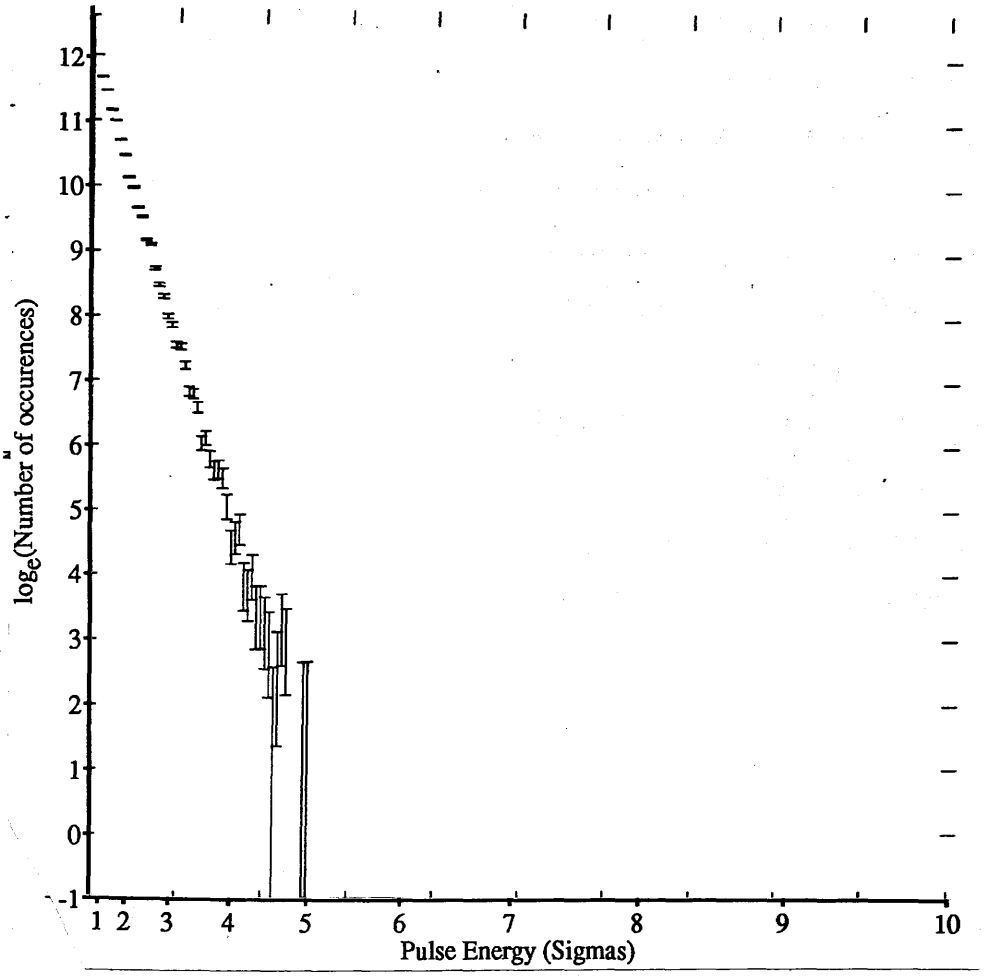

Figure 5.3: The pulse height distribution of the detector's output when large amplitude pulses have been removed.

5.6 Conclusions

The noise statistics of the prototype detector, demonstrate that the output is indeed Gaussian with the exception of approximately 10 events per hour, most of which are believed to be due to amplifier saturation and therefore, in principle, avoidable. Even if such an event rate of pulses cannot be reduced, a coincidence experiment between two or more such detectors should allow discrimination between these detector effects and gravitational wave signals. Consider, for example, two identical detectors on the same site. If the probability of any millisecond sample being a high amplitude, non-Gaussian event is P then the probability of a coincidence between the two detectors is P^2. From the event rate of the prototype we can conclude that $P \sim 3 \times 10^{-6}$, and so the probability for coincidence would be of order 10^{-11} which is about one coincidence every three years. For three detectors this coincidence rate falls to a level which can be safely ignored. Therefore the noise performance of the prototype during the recording runs after SN1987A was such that non-Gaussian events should not have seriously limited the sensitivity of a search for millisecond pulses in a coincidence search. For the data described there is no evidence for millisecond gravitational wave pulses with an amplitude greater than 3.8×10^{-17}.

Chapter 6

Periodic Signal Detection and the Frequency Domain Performance of the Prototype Detector

6.1 Introduction

In Chapter 1 it was shown that there is a strong possibility of detecting periodic gravitational waves with long baseline interferometric detectors. The broadband nature of interferometric detectors is important for searches for unknown sources of periodic gravitational waves, while the ability to place the detector in a narrow band configuration with considerably enhanced sensitivity, by recycling the light resonantly, increases the probability of detecting gravitational waves from known pulsars, such as the Crab and Vela pulsars or from neutron stars which are CFS unstable and for whom a modulated X-ray flux has been measured.

It is therefore important to ensure that the frequency domain performance of the detector is not limited by instrumental effects. For example, instrumental resonances may be present, and non-linear processes in the detector may introduce a range of spurious harmonics. This chapter describes the spectral analysis of some of the data recorded between 8.30am and 9.00am on the 11^{th} March 1987, after SN1987A was discovered.

6.2 The Detection of Periodic Signals

It was pointed out in the previous chapter that the linear filter which optimises the received signal to noise power is the matched, or correlation, filter. Therefore to detect a sinusoidal signal we should use a filter which has an impulse response which is a sinusoid of the same frequency. If the frequency is unknown *a priori,* then we must use a series of such filters at different frequencies. In any finite frequency range there are an infinite number of such frequencies. However, for a finite data length these matched filters need only have a bandwidth equal to $1/T$ Hz, where T is the length of the input time series in seconds, and be spaced $1/T$ Hz apart. By using two independent quadratures it is necessary to calculate the filter output for only one lag. This is just the calculation performed by the Discrete Fourier Transform (DFT). This calculation can therefore be efficiently performed using an FFT implementation of DFT.

One limitation of this technique is the picket-fence effect. Since each frequency bin of the DFT has a response to signals at nearby frequencies which is given by a sinc function, a signal at a frequency which lies exactly halfway between two consecutive bins of the DFT will contribute $A_0\text{sinc}(1/2)$ to the amplitude of each of these two bins, where A_0 is amplitude of the signal. If we simply look at the amplitudes of each bin in the DFT to discover signals then such a signal will contribute a maximum of $2A_0/\pi$, to any one bin in the transform, and so the improvement in signal to noise power ratio for such a split signal is less than one half the maximum possible value. This effect can be reduced in a number of ways. For example, the bandwidth of each frequency bin of the DFT can be reduced by performing a DFT of a longer data set, so that if the data was padded with zeros, say to twice its original length, and then transformed, the frequency resolution would be improved by a similar factor, two in this case. Since the sinc function with which the data's transform has been convolved is still that of the shorter data length, the minimum contribution of any signal is now $\text{sinc}(1/4)$. Alternatively, the time series can be multiplied by some window function, so that the convolution function is no longer a sinc function. A number of such window functions are

commonly used each with different properties. Such a window function effectively averages the signal in successive bins of the DFT. In fact, since it is simple to prove that a frequency peak which lies between two bins of the DFT contributes complex amplitudes to these bins which are exactly out of phase, subtraction of these two complex amplitudes therefore adds the signal coherently. This method should, in principle, reduce the signal to noise power loss factor to $\sim 10/\pi^2$, almost equivalent to that obtained by padding with zeros. (Correlations between consecutive bins of the DFT, introduced by window functions applied to the data output, may however introduce excess noise in this method.)

Since the power in any bin of the DFT due to noise is proportional to the bandwidth of the bin, while the signal, being narrowband, should contribute at least about half its power into one bin alone, the signal to noise power ratio is proportional to T and can therefore be improved by integrating for as long a period as possible. However, practical computational limits and signal purity considerations limit the useful integration period of a simple DFT.

6.3 The Effect of the Observer's Motion on Periodic Signals

The motion of the observer introduces two effects which change the nature of a purely monochromatic signal. Firstly, the motion of the observer introduces Doppler shifts in the observed frequency of the signal. Secondly the rotation of the observer causes amplitude modulation of the detected signal since the sensitivity of a laser interferometer is not equal in all directions or to all polarisations. These two effects will now be analysed in more detail.

6.3.1 The Doppler Effect

The Doppler effect will shift a spectral line at a frequency ν to an observed frequency ν_O given by

$$\nu_O = \frac{\nu}{1+z} \tag{6.1}$$

where z is the observed red-shift. The red-shift will be given by the velocity of the observer along a line between the observer and the source. This velocity, the topocentric radial velocity V_r^T, is the sum of two terms, the radial velocity of the observer with respect to the earth's centre, given by $\mathbf{v}' \cdot \mathbf{s}$ where $\mathbf{v}'$ is the velocity of the observer with respect to the earth's centre and $\mathbf{s}$ is a unit vector in the direction of the source, and the radial velocity of the earth's centre relative to the barycentre of the solar system, given by $\mathbf{v} \cdot \mathbf{s}$ where $\mathbf{v}$ is the velocity of the geocentre relative to the barycentre of the solar system. (Since the radial velocity of sources, which are not in binary systems, relative to the barycentre should be effectively constant over the timescales considered for integration of the signal the spectral lines of such a source will be shifted by a constant amount which can be neglected in the present discussion.) The red shift is therefore given by (adopting the convention that recession is positive)

$$z = \frac{V_r^T}{c} = -\frac{\mathbf{v}' \cdot \mathbf{s} + \mathbf{v} \cdot \mathbf{s}}{c}, \tag{6.2}$$

where c is the speed of light.

The value of $\mathbf{v}' \cdot \mathbf{s}$ is given approximately by [Green 1985]

$$\mathbf{v}' \cdot \mathbf{s} \sim -0.465 \cos\phi' \cos\delta \sin H \mathrm{km\,s^{-1}}, \tag{6.3}$$

where ϕ' is the geocentric latitude of the observer and δ and H are the source's declination and hour angle respectively, with $H = LST - RA$ where LST is the local sidereal time and RA is the right ascension of the source.

The value of $\mathbf{v} \cdot \mathbf{s}$ can be derived from tabulated values in the Astronomical Almanac. However an approximation which neglects the difference between the centre of the sun and the barycentre of the solar system gives a reasonable approximation for the variable part $\mathbf{v} \cdot \mathbf{s}$ of [Green 1985]

$$\mathbf{v} \cdot \mathbf{s} = -29.79 \cos\beta \sin(\lambda_\odot - \lambda), \tag{6.4}$$

where β and λ are the source's ecliptic latitude and longitude respectively and $\lambda_\odot$ is the true ecliptic longitude of the sun. For a given time β and λ can be determined

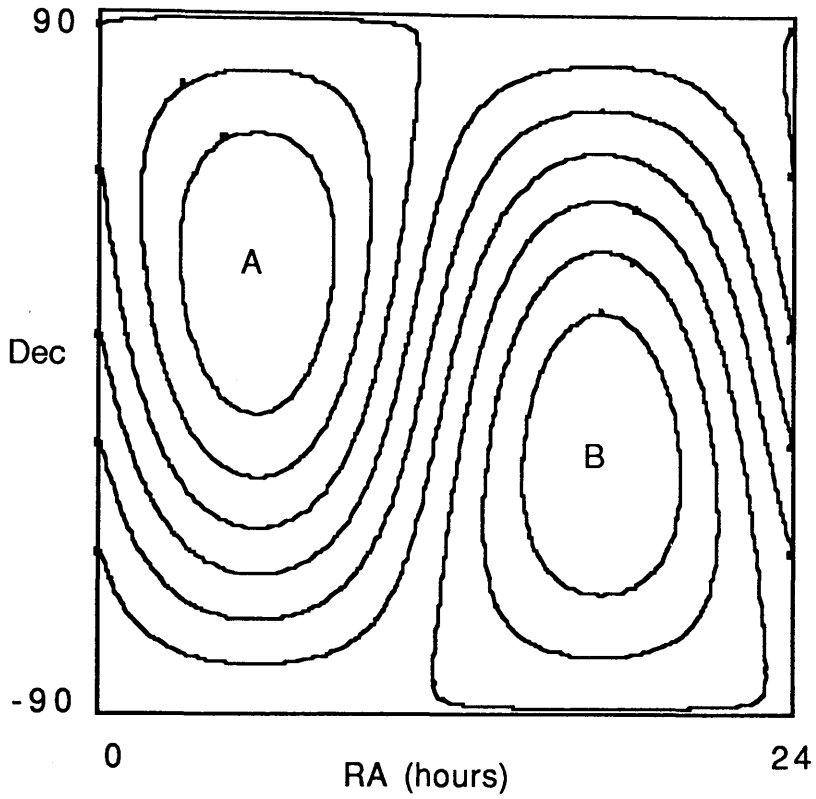

Figure 6.1: Red shift z as a function of source position for 08.45 GMT on 11/3/87. The contours are equally spaced. Points A and B correspond to $\pm 9.8 \times 10^{-5}$.

from the sources's RA and δ and the position of the sun can be determined so that the red shift can be calculated. For the 11^{th} March 1987 at 8.45 GMT, the red shift as a function of the source's position is plotted in Figure 6.1, with the rate of change shown in Figure 6.2.

If we make the simplifying assumption that the rate of change of the observed frequency is the maximum possible, then we find that

$$\frac{df}{dt} \sim f_0 \times 8.3 \times 10^{-11}\text{Hz/sec}, \tag{6.5}$$

where f_0 is the frequency of the signal. A DFT of a time series of length T seconds, gives information at frequencies separated by $1/T$, each with a bandwidth of $1/T$. If the frequency drift of the signal in time T is less than the bandwidth of one of

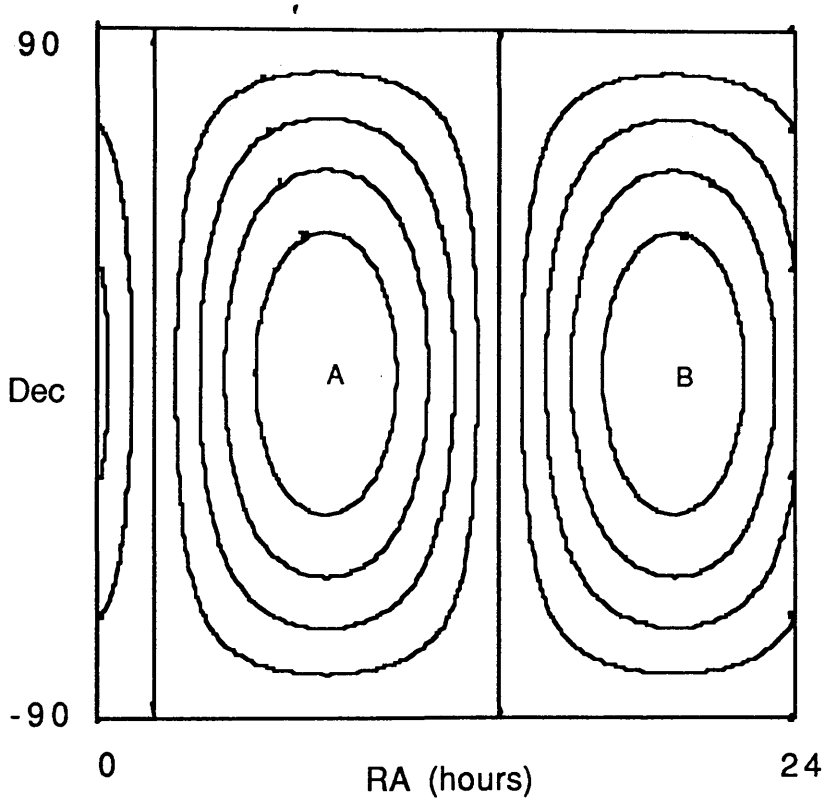

Figure 6.2: The rate of change in red shift as a function of source position for 08.45 GMT on 11/3/87. The contours are equally spaced. Points A and B correspond to $\pm 1.1 \times 10^{-7}$.

these frequency bins, then there will be little loss of signal to noise power ratio due to the frequency drift. For integration times longer than this, however, the frequency will drift through a number of bins and so rather more complicated, source dependent methods are required to obtain the optimum signal to noise ratio.

The maximum integration time that can be safely taken in this simple case is of order

$$T_{max}^2 \sim \left(f_0 \frac{df}{dt} \right)^{-1}, \tag{6.6}$$

so that for signals up to 2500 Hz we have $T_{max} \sim 36$ minutes. Since for efficiency, we wish to use a radix-2 transform, we will take 2^{23} points as the maximum length of transform that can be performed. For a sampling rate of 6000 Hz this corresponds to a period of 23.3 minutes.

6.3.2 Antenna Sensitivity Pattern

As the earth rotates, the direction of a source on the sky relative to the arms of the detector will change giving rise to amplitude modulation of the signal. Since gravitational waves have two independent polarisation states, and gravitational wave detectors have a response which is polarisation dependent, the response of the detector to a given direction will also depend on the polarisation of the wave. The time dependence of this effect will now be calculated. The difference in arm lengths induced by a gravitational wave is given by

$$\delta x\,(t) = h\,(t)\,l, \tag{6.7}$$

where l is the length of the detector's arms and $h\,(t)$ is an effective gravitational wave amplitude. The effects of polarisation and source direction are incorporated in the term $h\,(t)$.

The sensitivity pattern of a laser interferometer gravitational wave detector has been calculated by a number of authors [Forward 1978, Thorne 1987], and it has been shown that if the detector and source are oriented as shown in Figure 6.3

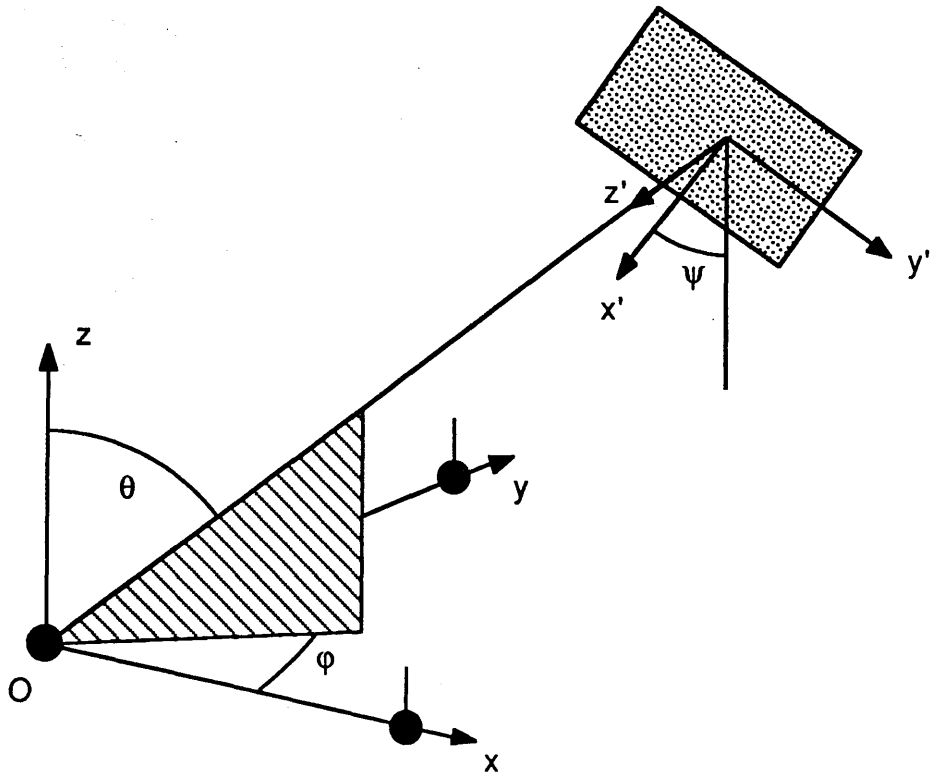

Figure 6.3: The relative orientations of the detector and source.
OX and OY represent the two arms of the detector with Z the direction of the zenith. Z' is the direction of propagation of the gravitational wave, with X' and Y' the polarisation axes.

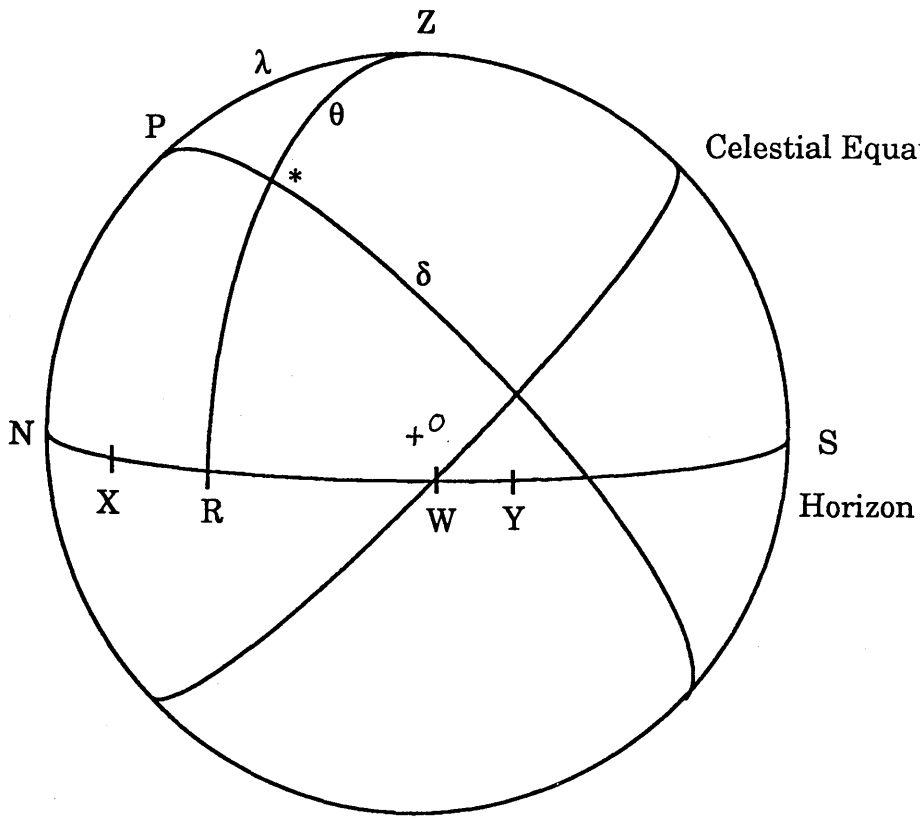

Figure 6.4: Celestial sphere of observer O.
The symbol '*' represents the pulsar.

then $h(t)$ takes the form

$$h(t) = F_+(\theta, \phi, \psi)\, h_+(t) + F_\times(\theta, \phi, \psi)\, h_\times(t). \tag{6.8}$$

where the beam-pattern factors F_+ and $F_\times$ are given by

$$F_+(\theta, \phi, \psi) = \frac{1}{2}\left(1 + \cos^2\theta\right)\cos 2\phi \cos 2\psi - \cos\theta \sin 2\phi \sin 2\psi \tag{6.9}$$

and

$$F_\times(\theta, \phi, \psi) = \frac{1}{2}\left(1 + \cos^2\theta\right)\cos 2\phi \sin 2\psi + \cos\theta \sin 2\phi \cos 2\psi. \tag{6.10}$$

Consider the celestial sphere of Figure 6.4 and the spherical triangle of Figure 6.5. The point Z represent the zenith of an observer O at a site with latitude λ. The

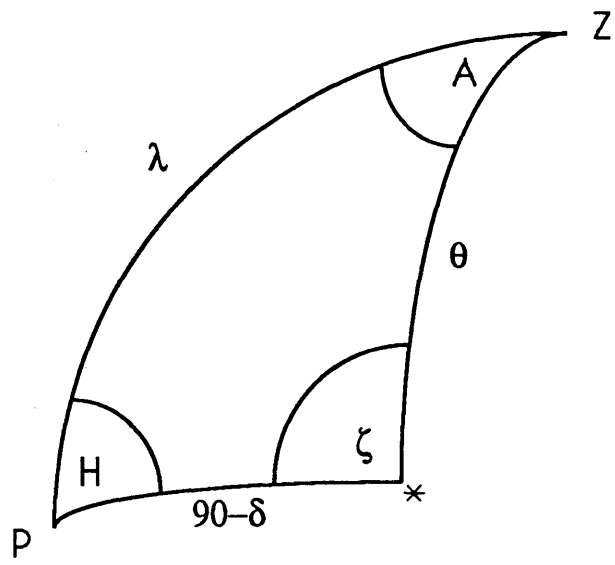

Figure 6.5: Spherical triangle of Figure 6.4.

points N, S and W are the points on the observer's horizon in the directions North, South and West respectively. The following identifications can be made;

$$\theta = Z*\tag{6.11}$$

$$\phi = XR = A - NX.\tag{6.12}$$

Equations 6.9 and 6.10 determine polarisation effects with an angle ψ relative to the constant–ϕ plane. However this angle will change as the earth rotates and in order to take account of the time dependence of this polarisation effect it is useful to define an unchanging direction on the celestial sphere. Such a direction will remain fixed relative to the source. We make the arbitrary choice of the great circle through points P and *. This defines a plane of constant–RA. The plane of constant–ϕ is given by the great circle through points Z and *.

With these definitions, and knowing H, λ and δ, we can use the standard equations of spherical trigonometry to obtain the relations

$$\cos\theta = \cos\lambda\sin\delta + \sin\lambda\cos\delta\cos H\tag{6.13}$$

$$\cos A = \frac{\sin\delta\sin\lambda - \cos\delta\cos\lambda\cos H}{\sin\theta} \tag{6.14}$$

$$\sin A = \frac{\sin H\cos\delta}{\sin\theta} \tag{6.15}$$

$$\sin\zeta = \frac{\sin H\sin\lambda}{\sin\theta} \tag{6.16}$$

$$\cos\zeta = \frac{\cos\lambda\cos\delta - \sin\lambda\sin\delta\cos H}{\sin\theta} \tag{6.17}$$

(since $0 \leq \theta \leq \pi$ there is no ambiguity about the sign of $\sin\theta$). H is the hour angle of the source and this can be calculated for a particular time. Therefore the values of F_+ and $F_\times$ can be determined as a function of time.

For the 11^{th} of March 1987 the sensitivity of the Glasgow prototype detector to a source in the direction of the Large Magellanic Cloud is plotted in Figures 6.6 and 6.7, corresponding to purely h_+ and purely $h_\times$ gravitational waves respectively.

This amplitude modulation produces sidebands on the observed signal. However, for periods of order 1 hour, the DFT is insufficient to resolve these sidebands and there will be little change in the received signal. Therefore the time dependence of the detector's response can be ignored and the effect of the antenna's sensitivity pattern is just to reduce the sensitivity to sources in particular directions and with particular polarisations.

6.4 FFT Program

The DFT defined by Equation 4.2, can be regarded as the multiplication of a vector of length N with an $N \times N$ matrix. The result, a vector of length N, is the DFT. This operation requires N^2 multiplications. However it can be shown [Bracewell 1986] that the $N \times N$ matrix can be factorised to reduce the total number of multiplications. This is the basis of the Fast Fourier Transform (FFT). The FFT algorithm requires a number of multiplications proportional to $N \log_2 N$, (Algorithms exist which reduce the number of multiplications further at the expense of an increased number of additions and subtractions. For modern computers with

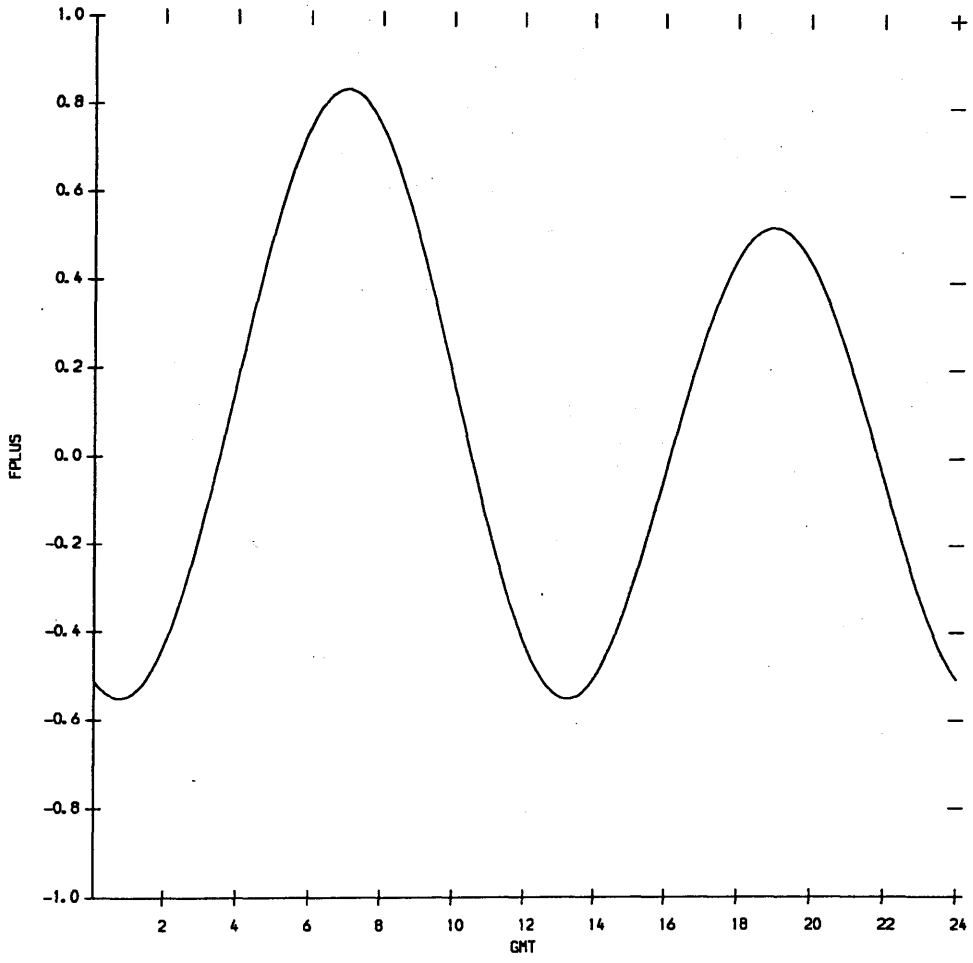

Figure 6.6: Sensitivity of Glasgow prototype to h_+ polariation against time for 11/3/87.

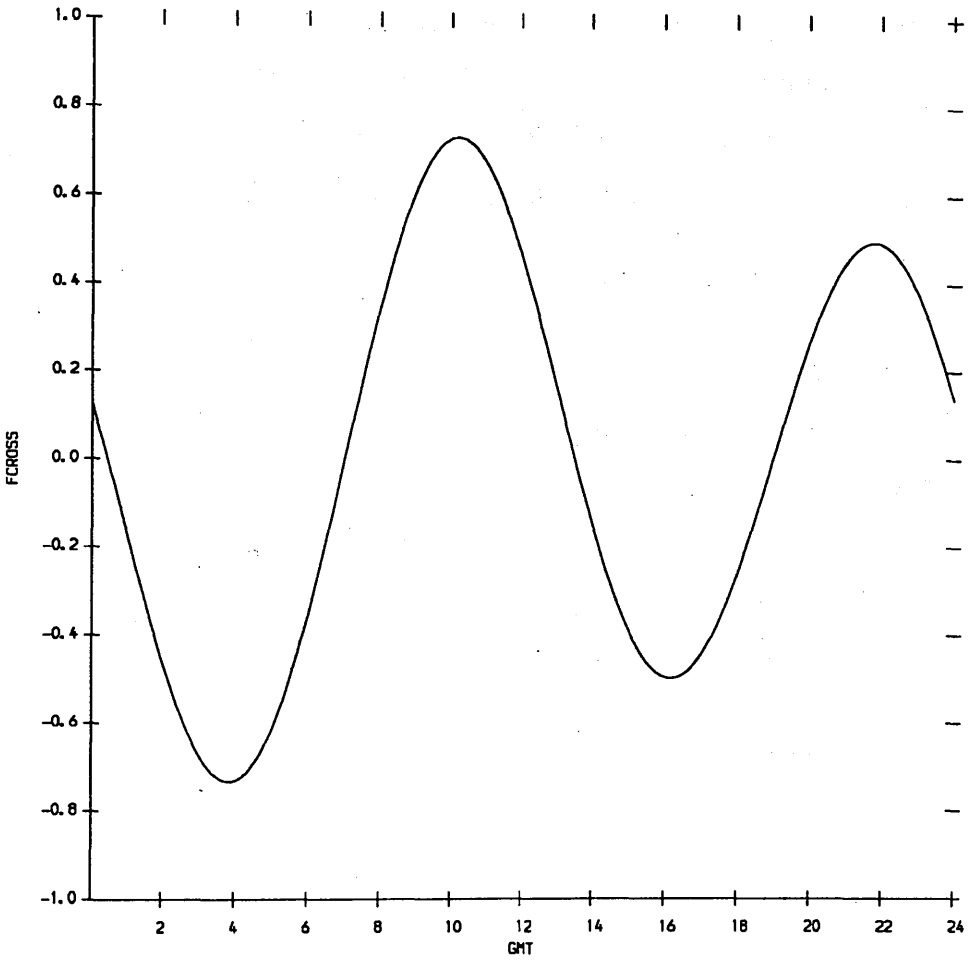

Figure 6.7: Sensitivity of Glasgow prototype to h_x polarisation against time for 11/3/87.

hardware floating point multiplication which takes no longer than floating point addition and subtraction, there is little advantage in using these algorithms.)

For large FFT's it may not be possible to store all of the data in memory at once. In such a case an algorithm which swaps data in and out of mass store as required must be used. The computer used to perform the FFT was an IBM 3090-150E mainframe with a vector processor, running the VM/CMS SP operating system. This operating system has a maximum addressing range of 16 Mbytes and so the largest array which could be stored is 8 Mbytes. A transform of 2^{23} points, stored as DOUBLE PRECISION 8 byte floating point numbers requires storage of 64 Mbytes and so such a mass store FFT had to be used.

The FFT routine used was a standard mass store routine, ACM Algorithm 545 [Fraser 1979]. For 2^{23} points it required 149 seconds of processor time. However, since data must be repeatedly swapped into and out of memory, the FFT is I/O limited and therefore takes approximately 40 minutes to perform.

Before the FFT was performed the data was preconditioned. The DC component was removed, to prevent convolution effects with a large amplitude zero frequency signal, and losses of lock, calibration peaks and periods when the ADC saturated were also windowed from the time series. The effects of some of these and other modifications to the time series will now be considered.

6.5 Power Line Pick-Up

One source of contamination in the spectrum of the recorded signal is interference at harmonics of the 50 Hz power line frequency. This frequency will drift with time; an example of the drift with time of the frequency of the fundamental is plotted in Figure 6.8. Higher harmonics are present in the output of the detector, and these contribute power to a wide range of frequencies. For example the 60^{th} harmonic, at 3000 Hz, varies by about ± 7 Hz. Ideally such interference would be removed, by improving the shielding and isolation of the components of the detector. However, if it can not be removed completely then it may be desirable to filter the recorded data with a filter which notches out the mains harmonics and which varies with

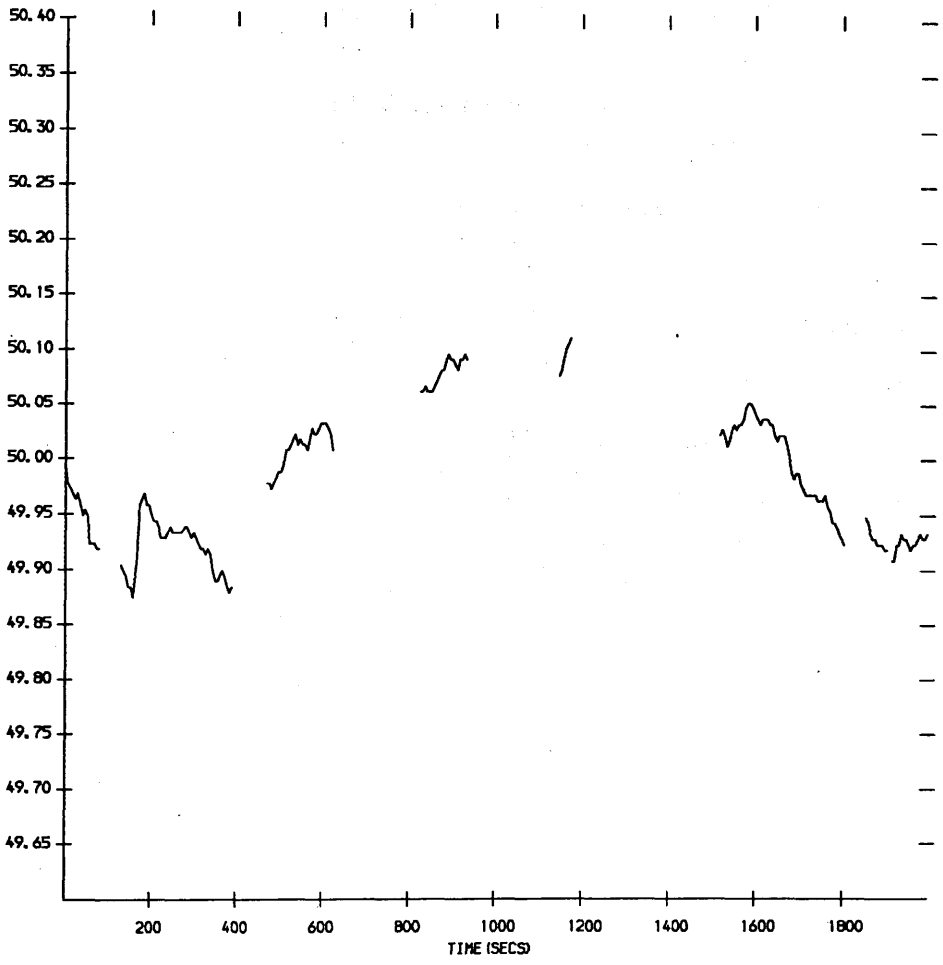

Figure 6.8: Typical variation in the frequency of the AC power line with time (Derived from the recorded secondary signal. Gaps are due to losses of lock.) Y axis in hertz.

time to track these frequencies. If the filter notches can be made narrower than the typical frequency deviation of the harmonics then, as these notches track the frequency drift, little background noise should be lost from any frequency. In any event, as the frequency drifts in and out of bins, adding incoherently over long timescales, the received signal to mains harmonic power in a frequency bin will be proportional to integration time for sufficiently long integration times, that is those above the characteristic timescale of frequency drift.

6.6 Effect of Lost Samples

As discussed in Chapter 4, 4 samples every 2040 samples were lost. If we zero these samples then this has the effect of multiplying the time series with a window function which is a repeating pattern of zero for 4 samples and one for 2036 samples. Therefore every peak in the transform of the data will be convolved with the transform of the window function, which has the form of a series of impulses at frequencies which are integer multiples of $6000/2040 = 1/0.34\,\mathrm{Hz}$ with amplitudes given by

$$\frac{2036}{2040}\mathrm{sinc}(2036 \times 6000 \times f/2040^2) \tag{6.18}$$

Large amplitude, narrow band signals will therefore produce harmonics which may be above the level of the background noise. For example the transform of the data at low frequencies is shown in Figure 6.9, where it can be seen that harmonics spaced $1/0.340\,\mathrm{Hz}$ apart are present. These are due to the convolution of the large amplitude harmonics which are the result of resonances, around 400 Hz, of the pendulum suspension wires. Since the amplitude of the zero frequency component of the transform of the window function is much larger than any of the other components, by about 50 dB, it is possible to remove these harmonics to first order, by deconvolving the spectrum. In this case care must be made to allow for the negative frequency components of the wire resonances in the Fourier transform which, when convolved with the transform of the window function, produce positive frequency peaks. The effects of leakage of the the peaks produced by the convolution into the bins of the DFT must also be considered. When this is done,

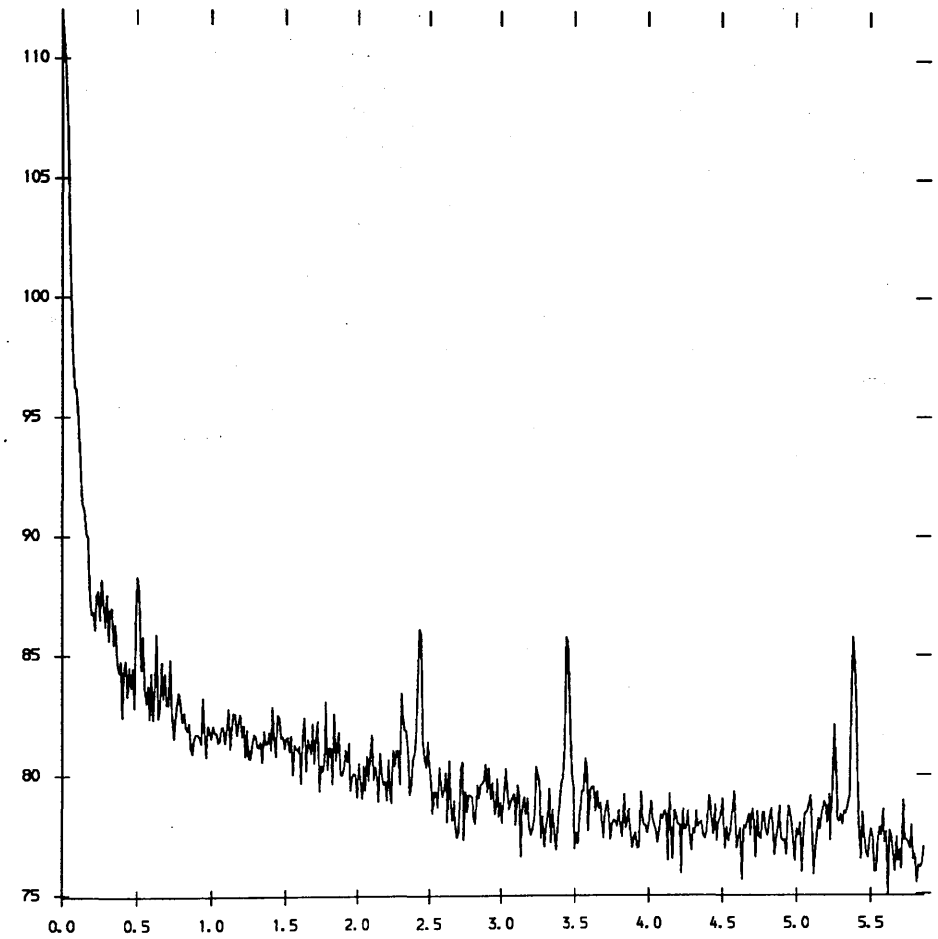

Figure 6.9: Section of the transform of the data. (Losses of lock, calibration signals and saturation all windowed out.) A low frequency section is shown since at these frequencies the spectrum is quantisation noise limited, so the amplitude at these frequencies is as low as possible and the peaks show most clearly. Y axis in dB's, X axis in hertz.

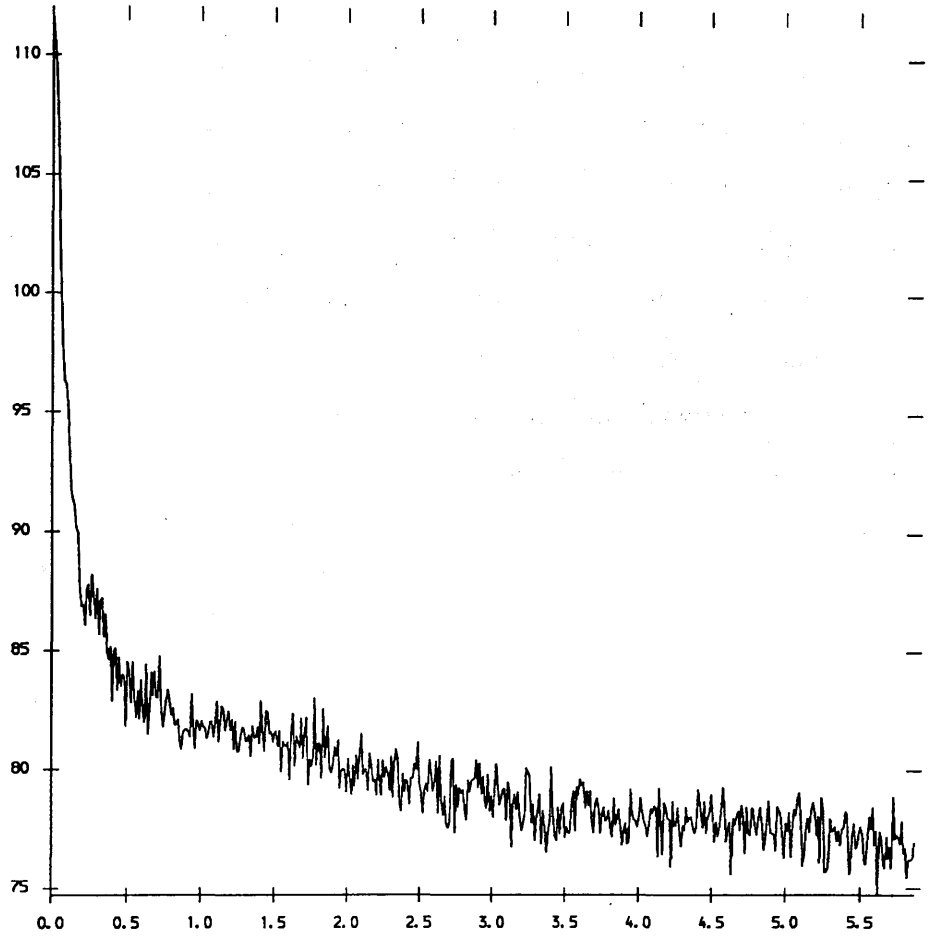

Figure 6.10: The result of deconvolving the data of Figure 6.9 for the peaks introduced by the lost samples. Y axis in dB's, X axis in hertz.

the deconvolved spectrum is as shown in Figure 6.10. The procedure can also be carried out (perhaps more easily) in the time domain. Essentially the large amplitude narrow band signals from the wire resonances would be interpolated for the lost samples.

6.7 Quantisation Error

It was noted in Section 4.3.1 that for some signals the effects of quantisation error should be regarded as a distortion. Since the hardware fault described in

Section 4.5 limited the resolution of the recorded secondary cavity signal to 6-bits it is possible that this effect will be present in the recorded data. To determine the amplitude of this effect a time series was generated to simulate the quantisation noise. A series of random samples from a Gaussian distribution was generated using a standard pseudo-random number generator (NAG routine G05DDF). This time series was then filtered by a digital filter with a frequency response approximately equal to the spectrum of the recorded signal. To this time series were added sinusoidal signals at frequencies and amplitudes corresponding to the resonances of the wires supporting the test masses. This signal was then quantised to the same resolution as the recorded data, but the quantisation error part alone was kept. Since the Fourier transform is linear, the transform of this data is equal to the change which quantisation introduces to the Fourier transform of the data.

The Fourier transform generated in this way did indeed show evidence for mixing of the wire resonance frequencies; a section of the transform is shown in Figure 6.11. For comparison the same section of the Fourier transform of the recorded data is shown in Figure 6.12. It is evident that many of the peaks in the spectrum may be due to this quantisation effect.

The amplitudes of these spurious peaks could be reduced by windowing out those parts of the time series where the amplitudes of the wire resonances are higher than normal. For the section of data considered here, the wire resonances are usually of low amplitude. Periodically however, they become excited and since the resonances are lightly damped they take of order 30 seconds to damp down. When these periods are removed by windowing, the resulting data has a Fourier Transform as shown in Figure 6.13. It can be seen that these spurious peaks have been removed. (It can also be seen that the effects of the convolution of the Fourier transform with the peaks corresponding to the wire resonances have been reduced to a level where they may be disregarded.)

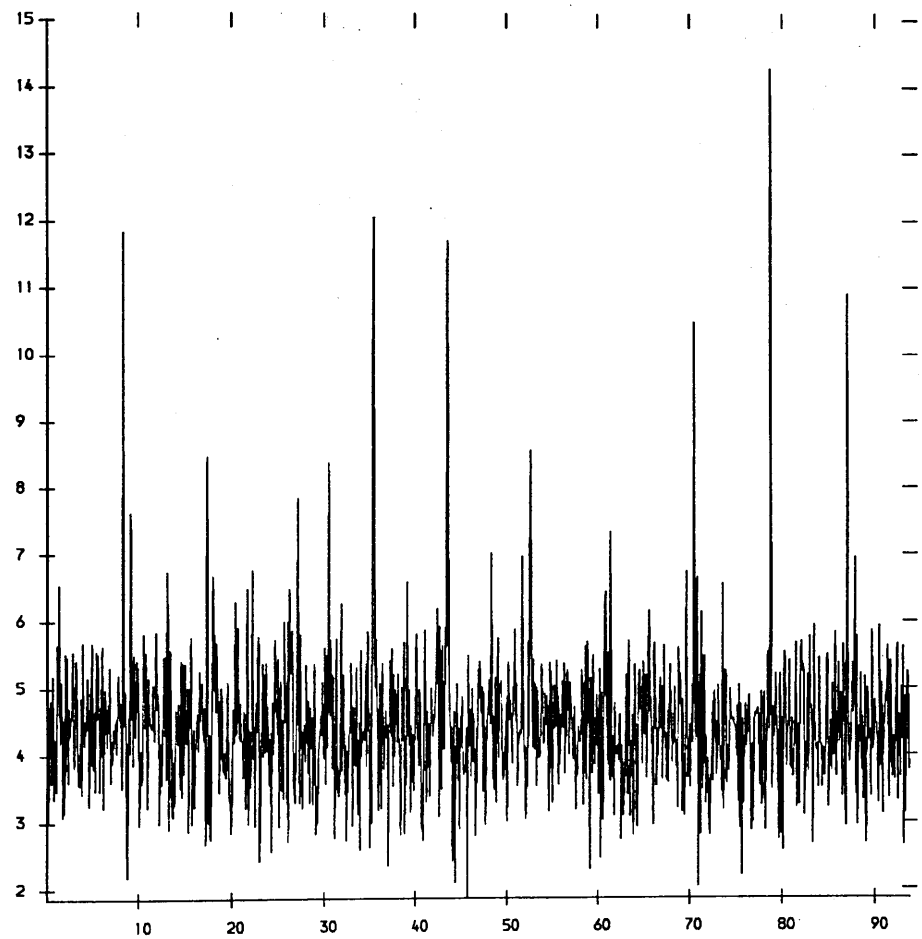

Figure 6.11: Simulation of the effect of quantisation noise on the recorded signal's spectrum. Y axis in dB's, X axis in hertz.

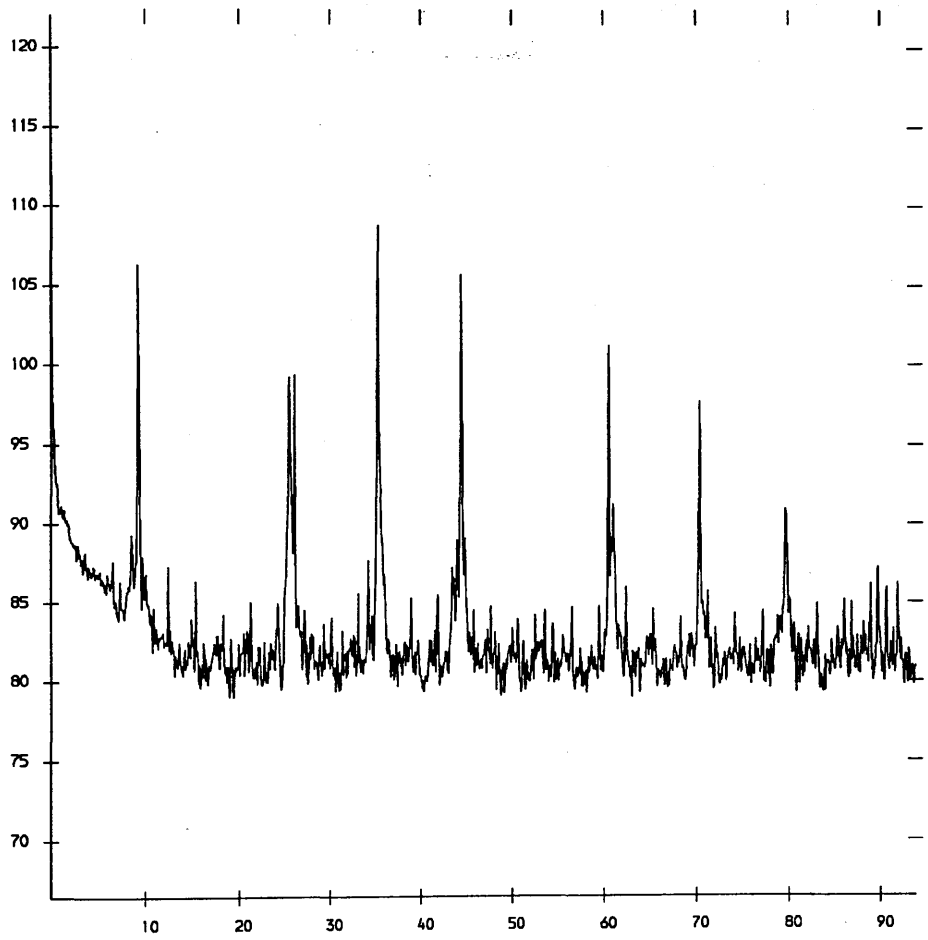

Figure 6.12: Low frequency section of the transform of the recorded signal (same frequency range as Figure 6.11)., showing quantisation noise. Y axis in dB's, X axis in hertz.

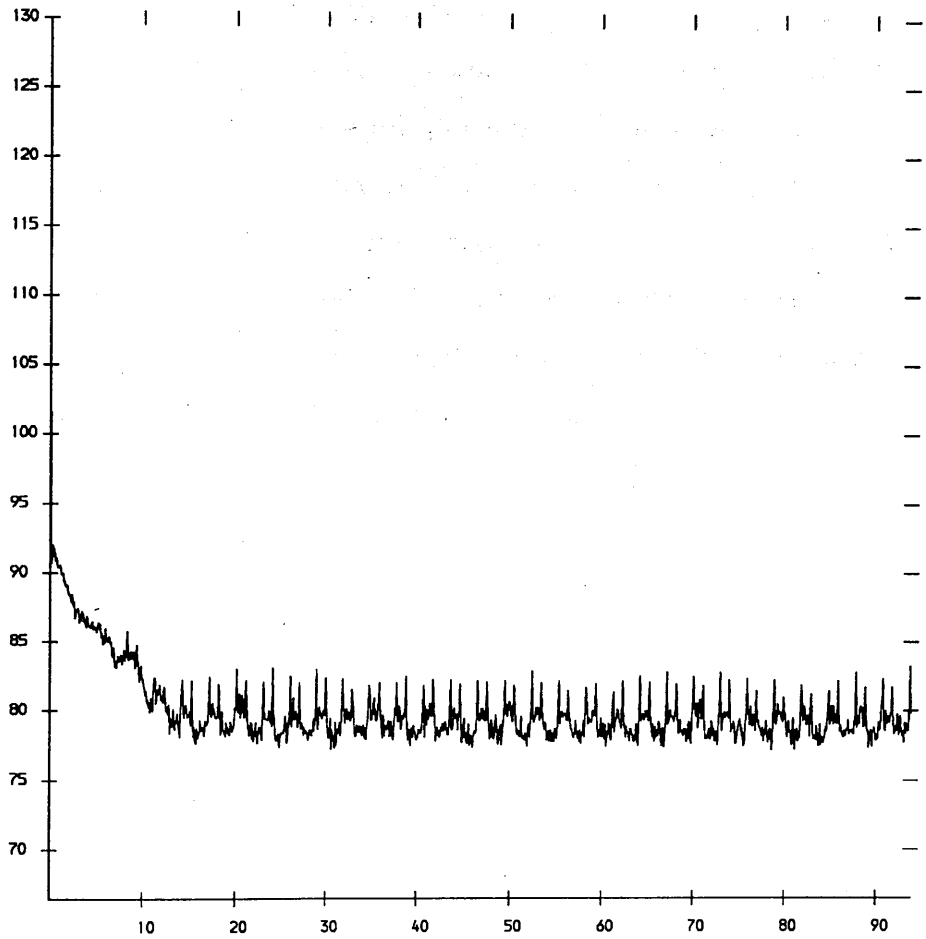

Figure 6.13: The transform of data which has had all periods of wire resonances removed. Equally spaced peaks are due to the windowing of the wire resonances by the lost samples.

6.8 Time Dependence of the Detector Transfer Function

Since the transfer function of the detector and recording system changes with time over intervals shorter than the desired transform length, it is not possible to derive the true power spectrum of the effective gravitational wave signal from a transform of the recorded signal. In principle, if the transfer function of the detector can be derived, the data could be filtered to derive the true gravitational wave signal. This, however, requires a filter which is also time dependent and therefore changes over timescales of order a minute. Above the unity gain frequency of the secondary cavity control system the level of the recorded signal is reasonably constant, typically varying by 10%, and even at the unity gain frequency the variation is typically only 25%. Therefore we shall consider only frequencies above 1.5 kHz for a search for periodic gravitational wave signals. For a search for possible instrumental effects we may search the entire transform noting that the effective gravitational wave amplitude of any peaks detected in this frequency range will be uncertain to about 10%.

6.9 Windowing of Calibration Peaks

The addition of calibration peaks introduces large amplitude harmonics at frequencies in the range of interest. Since these signals were applied for one second, each peak will have a bandwidth of order 1 Hz. While the signal bandwidth lost is small the calibration peaks should be removed prior to the transform to minimise the possibility of convolution effects in the frequency domain introducing spurious harmonics.

The window function chosen was one that was zero while the calibration peak was present and tapered over one second on each side of the calibration peak. For

a calibration peak starting at $t = 0$ the window function was given by

$$\begin{aligned} w &= 1 && \text{for } t < -1\,\text{sec} \\ w &= \tfrac{1}{2}(1 - cos(\pi t)) && \text{for } -1 \geq t < 0\,\text{sec} \\ w &= 0 && \text{for } 0 \leq t < 1\,\text{sec} \\ w &= \tfrac{1}{2}(1 + cos(\pi t)) && \text{for } 1 \leq t < 2\,\text{sec} \\ w &= 1 && \text{for } t \geq 2\,\text{sec} \end{aligned} \tag{6.19}$$

The taper is included to avoid sharp discontinuities, which could cause that the transform of the window function to have significant power at high frequencies; the choice of a one second cosine taper is arbitrary, The effect of high frequency terms in the convolution could be reduced by taking a longer taper at the expense of reducing the signal power in any peak.

The effect of this window function is to reduce the power in any frequency peak due to noise to about 98% of its true value. Similarly the amplitude of any periodic signal will be reduced to 98% of its true value, giving a degradation in the received signal to noise power ratio of 98%. Since this is less than the error caused by the dynamic range limitation of the 6-bit conversion the effects of this windowing can be safely ignored.

6.10 Losses of Lock

The effect of losses of lock can be considered equivalent to multiplying a continuous detector output by a window function which is zero whenever the detector is unlocked and one whenever it is locked. To the time series produced in this way is added a signal which is equivalent to the spurious pulses generated by the secondary cavity and laser passing through resonances as described in Chapter 4 while the detector is unlocked. The Fourier transform of the window function will have a zero frequency term whose amplitude is equal to the fraction of the total time of the data set that the detector is locked. Since the times of loss and regain of lock are essentially random and widely spaced, the transform of the window function is unlikely to contain any large amplitude harmonics and so the main

problem caused by losses of lock is simply one of signal to noise power ratio loss. Since pulses are generated while the detector is unlocked, it is important to remove these spurious signals. This can be done simply by replacing samples taken while the detector was unlocked by the mean value, which in this case is zero, since the DC component of the signal was removed before the transform. For the data analysed here, a cosine taper of one second was used to reduce end effects due to losses of lock.

6.11 Statistics of the Discrete Fourier Transform

Since the DFT of independent Gaussian numbers calculates weighted sums of these numbers, each value in the output of the DFT is itself a Gaussian random number and in particular the real and imaginary parts of a particular frequency bin are independent. Therefore the power in each bin will be a sample from a chi-square distribution of order two and the amplitude will be a sample from a Rayleigh distribution [Woodward 1953]. If a search for periodic gravitational waves is to be made we must select an arbitrary threshold level for the power in the bin. This threshold must be large enough that the probability of statistical fluctuations causing a bin to exceed the threshold is small. For this experiment we may arbitrarily choose a threshold that satisfies the condition that the probability of one or more bins in the entire transform being above the threshold, P_{all}, is 1 in 20, so that we are 95% confident that there will be no such bin due to statistical fluctuations.

If we assume that each bin has the same probability of crossing the threshold, given by P_{one}, then

$$(1 - P_{one})^N = 1 - P_{all}, \tag{6.20}$$

where N is the number of frequency bins. (For non-white noise and a single threshold this will not be true; however as will be shown later a threshold which is a function of frequency can be used to ensure this condition.) Since the amplitude of each bin is a sample from a Rayleigh distribution, the probability of the amplitude exceeding a value r is

$$P(>r) = \exp\left(-\frac{r^2}{2\sigma^2}\right), \tag{6.21}$$

where σ is the standard deviation of the Gaussian distributed real and imaginary parts of the discrete Fourier transform. For $N = 2^{23}$ and $P_{all} = 0.05$ the desired threshold is 6σ. Since the effects of noise in the frequency bin containing such a signal can both increase and decrease the amplitude in the DFT, we can be ~50% certain of detecting a 6σ peak with this threshold and 95% sure of detecting 8σ peaks.

Since the spectrum is not white, we must use a threshold equal to six times the standard deviation at a particular frequency. We must therefore estimate the local standard deviation of the noise and this is done by calculating the value for a given number of consecutive bins. If too many bins are used, the average will be affected by the variation with frequency, whereas if too few bins are used, the estimate will have a poor statistical stability and an individual peak with high power may bias the estimate. If N samples, $x_1, x_2, \ldots, x_N$ of a Gaussian distribution with zero mean are taken, then the sample variance can be taken as

$$S^2 = \frac{1}{N}\sum_{i=1}^{n} x_i^2. \tag{6.22}$$

The error in this estimate of the standard deviation S will be approximately $S/\sqrt{2N}$. If we use 1024 consecutive frequency bins then we have $N = 2048$ and the error in the estimate of S is only about 2%. To generate a similar error a single frequency peak would require an amplitude equal to about $10S$, so this estimate is reasonably robust. The average is being performed over only 0.73 Hz and so will be relatively unaffected by the non-white nature of the spectrum.

When the discrete Fourier transform generated as described in Section 6.4 is analysed in this way it is found that the amplitude distribution has the form shown in Figure 6.14, which also shows the expected theoretical distribution. Apart from a few bins with amplitudes above the 6σ level, the distribution follows the expected form quite closely. Those bins which are over 6σ are candidates for gravitational wave signals with the confidence levels calculated above. However, these peaks were identified and it was found that they were close to 1 kHz and 2 kHz. In fact they appear to have been due to a fundamental signal at 1 kHz, its second harmonic and sidebands at displacements of $\pm$10 Hz, $\pm$20 Hz and $\pm$30 Hz. Since

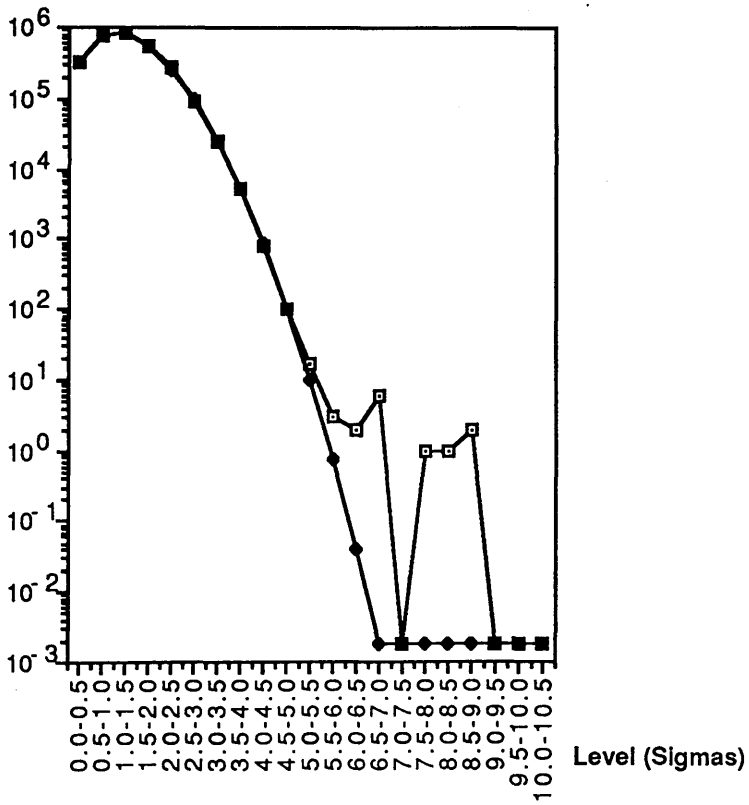

Figure 6.14: The theoretical (lower) and observed (upper) fourier transform statistics.

the NOVA computer has a hardware real time clock which generates frequencies of 1 kHz, 2 kHz and 10 Hz and since, to the resolution of the FFT (0.715 mHz), the harmonics in the recorded data are exactly related by integers it is likely that these harmonics are present due to pick-up from the clock signals. When these harmonics are eliminated, we find that there are no frequency bins whose amplitudes exceed 6σ and that there is therefore no evidence for periodic gravitational wave signals. Since the sensitivity of the detector was of order $2.5 \times 10^{-19}/\sqrt{\text{Hz}}$ between 1.5 kHz and 2.5 kHz during the data runs, and the detector had a duty cycle of order 60% for the data analysed, a limit on the strength of a suitably polarised source in the direction of the LMC of 6.9×10^{-20} can be placed in this frequency range at the 95% confidence level (8σ).

6.12 Conclusions

The results presented in this chapter demonstrate that the frequency domain performance of the prototype at the time of the recording runs is as expected. There is no evidence for any detector resonances in the frequency range above 1 kHz. Similarly there is no evidence for periodic gravitational waves at an amplitude threshold where we would expect that statistical fluctuations would cause only one frequency bin to cross the threshold in every twenty transforms.

Chapter 7

Future Prospects

Since 1960, when Weber first proposed the search for astrophysical gravitational waves using resonant bar detectors, the sensitivities of bar detectors have improved until they now have sensitivities approaching 10^{-18} for millisecond pulses. Coincidence runs between such bars, lasting for a few months, have been carried out; achieving sensitivities of 5×10^{-18}. However there has been no evidence for the detection of gravitational wave pulses and it is clear that much better sensitivities are required. Further advances in the sensitivity of bars will become increasingly difficult as the quantum limit is approached, so that it is likely that it will prove difficult to increase the sensitivities of bar detectors beyond a millisecond pulse sensitivity of 10^{-21}. The event rate of signals at such a sensitivity will be rather low.

The Glasgow prototype interferometer detector currently has a sensitivity to millisecond pulses of around 5×10^{-18}—approaching that of the best bars—and prototype detectors in Germany and America have similar sensitivities. A number of groups are currently proposing to build interferometers with arms of 3 km or more. Two detectors are currently planned in Europe—one is a collaboration between the Glasgow group and the group in Garching and the other is a collaboration between groups in France and Italy. An American gravitational wave observatory, consisting of two 4 km detectors, is also planned. If current estimates of source strengths are valid, the proposed long baseline instruments should allow the detection of millisecond pulses of amplitude 10^{-22}, at the rate of several events per year. With four or more detectors, it should be possible to determine the

polarisation state of the wave and the position of the source.

When these long baseline interferometer detectors come into operation, searches for a wide variety of gravitational wave signals will be possible. In addition to the bursts of gravitational waves mention above, periodic gravitational waves and signals from coalescing binaries should be detectable with these detectors. It is likely, at least initially, that searches for periodic sources of gravitational waves will be confined to known pulsars. This will enable the detector to be placed in a narrowband configuration to optimise its sensitivity at the expected pulsar frequency. The data analysis is also considerably simplified by considering only known sources, since their positions are also known and the motion of the observer can be allowed for. A number of millisecond pulsars in binary systems have been found recently, which may be detectable if the emission frequency is known.

Perhaps the sources of gravitational waves most likely to be detected first, however, are coalescing binary systems. These have the advantage that they emit many cycles of large amplitude gravitational waves over a short interval. Their short duration ensures that Doppler effects need not be considered, while their characteristic form is unlikely to be simulated by any detector noise source. Simple matched filtering techniques should therefore be sufficient to detect such waveforms. As an example of the signal to noise improvement which can be achieved using these techniques, Figure 7.1 shows the output of the prototype detector while one of the chirps described in Chapter 4 was being applied. The result of filtering this data with the correct matched filter is shown in Figure 7.2. The development of signal processing algorithms for detecting the predicted waveforms is also underway.

These detectors are many years in the future. However, current work is demonstrating that many of the techniques which may be needed to achieve full sensitivity, are practicable. At present, effort in the research groups developing interferometers is being concentrated on the development of improved optical techniques, improved seismic isolation and on the development of high power, continuous wave lasers. For example, light recycling has been demonstrated both in Garching and in Orsay, Paris, and a multi-stage pendulum, which should provide enough isolation from seismic noise for detectors to operate down to 10 Hz, is under development

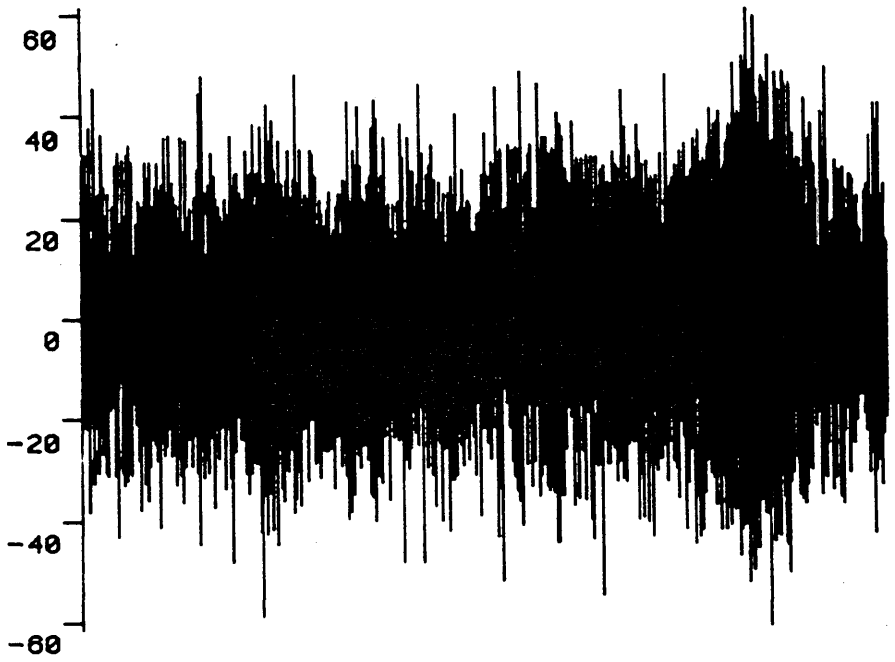

Figure 7.1: Sample of detector output containing a synthesised chirp. (X axis is 1.4 seconds.)

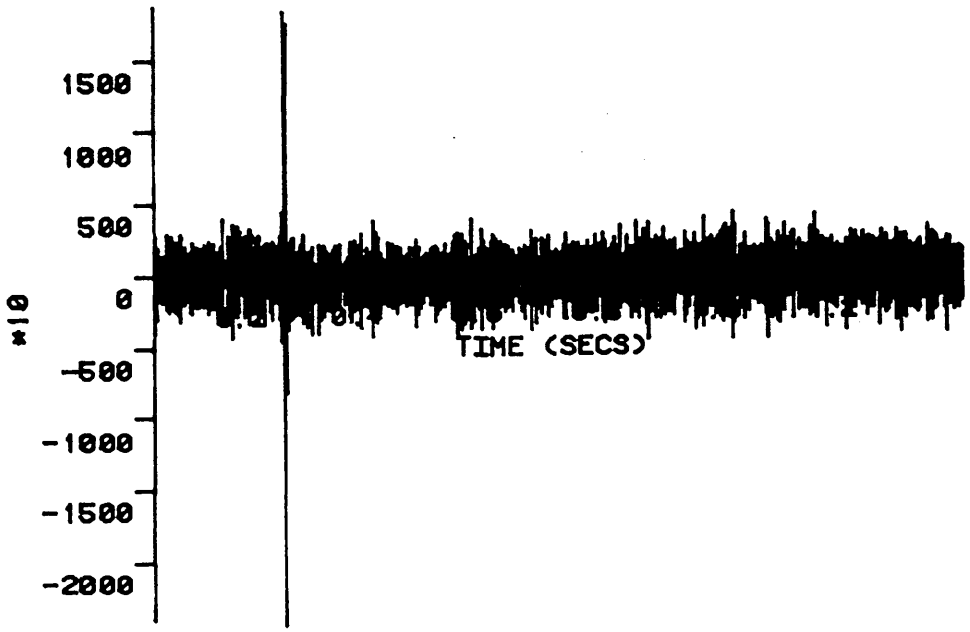

Figure 7.2: The result of filtering the data of Figure 7.1 with a suitable matched filter.

in Italy. The application of squeezing techniques to interferometers is also under investigation.

The technological difficulties involved in constructing such detectors are not insignificant, but neither do they appear insuperable and the potential rewards are high. Gravitational waves will carry information about the bulk motion of matter which cannot be derived in any other way, and their direct detection will therefore provide a new observational tool for the astronomer.

Appendix A

Lead and Rubber Stacks

Stacks of alternate layers of lead and rubber are used in most forms of gravitational wave detector to provide isolation between the apparatus and seismic noise. Such stacks can, in principle provide a very high degree of isolation in the frequency range of earth based detectors. However it is possible that cross-coupling of the modes of the stack and resonances of the components of the stack will severely limit the isolation provided by these stacks. A measurement of the response of such a stack was therefore made. A controlled vibration was introduced at the bottom of a stack comprising 5 lead blocks separated by 4 sheets of rubber. The stack was positioned on an aluminium plate which was supported by rubber pillars. A vibrator, fixed to the plate, could be driven by a voltage noise source. The magnitude of the vibrations at different levels of the stack was measured with a small commercial accelerometer. To reduce the effects of acoustic coupling between the layers of the stack, the apparatus was place inside a vacuum chamber at 5×10^{-2} torr.

A mass m which rests on a layer of rubber of thickness t, surface area A and with shear modulus μ (Figure A.1) can be viewed as a damped harmonic oscillator. If most of the damping of the motion of the mass is provided by the rubber then the damping of the mass can be modelled as proportional to the relative velocities of the two surfaces of the rubber, that is $\propto (\dot{x}_1 - \dot{x}_0)$. The transfer function of motion from the ground to such a mass is given by

$$\frac{\gamma s + \omega_0^2}{s^2 + \gamma s + \omega_0^2}, \tag{A.1}$$

where γ is the damping coefficient and $\omega_0^2 = \frac{\mu A}{tm}$ is the resonant frequency of the

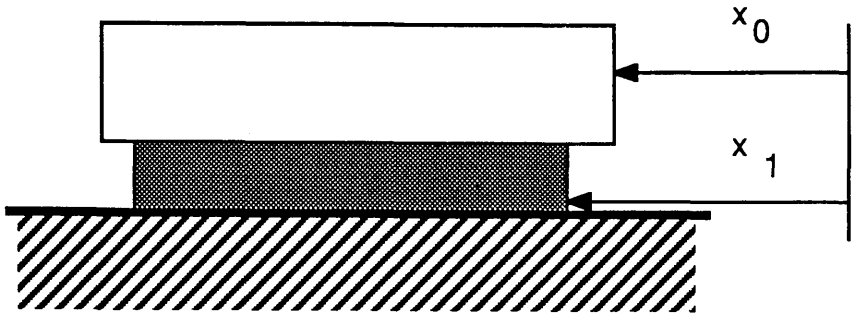

Figure A.1:

mass on the rubber sheet (typically 6 Hz for the masses and rubber sheets tested). The degree of isolation provided by such a layer therefore increases with frequency. For a heavily damped mass, $\gamma \sim \omega_0$, the isolation above the resonant frequency improves as $1/\omega$ while for no damping the isolation improves as $1/\omega^2$. For the case where the damping is light, the isolation improves as $1/\omega^2$ until a frequency ω_0/γ above which the isolation only improves as $1/\omega$. By cascading N such layers it is hoped to obtain isolations which improve as ω^{2N} for the lightly damped case and ω^N for the heavily damped case. However, it can be seen that the isolation achieved will depend on the relative masses of the components of the stack, so that such a high degree of isolation may not be achieved in practice. In addition, internal resonances in the masses of the stack may limit the isolation obtained.

Initial measurements revealed that the isolation provided by the stack was limited by a resonance at 880 Hz. This was identified with the fundamental transverse mode of the lead blocks, and its effects were reduced by replacing the sheets of rubber with small rubber spacers, positioned at the nodes of this resonance. The transmission of vibrations from one level of the stack to the next higher level for this modified arrangement was found to fall as 12 dB per octave, as expected for low damping. The addition of some damping material between the blocks reduced this figure to 6 dB per octave, confirming the theoretical prediction. However the frequency at which the fall-off began was around 100 Hz, substantially higher than expected. This is probably due to a high order resonance of the stack. At higher frequencies (around 1 kHz) vibrations transmitted to the accelerometer through the connecting cables placed a limit on the measured isolation, which was at least

40 dB over a frequency range of 400 Hz to 2.5 kHz.

If, however, the assumption is made that the isolation does indeed continue to improve at 12 dB per octave for each layer of the stack, then it can be calculated that, at 400 Hz, the isolation would be about 96 dB, rising to 144 dB at 800 Hz. Such high isolations are unlikely to be achieved, due to resonances within the individual components of the stack, such as the lead block resonance described above. For the current prototype only modest seismic isolation is required and these stacks are adequate. For the long baseline instruments it is likely that active isolation systems will be required.

Appendix B

1401 Data Acquisition Program

Since the data acquisition system described in Chapter 4 was completed and used to record the data which has been analysed in Chapters 5 and 6, a new data acquisition system has been developed. This is based on a 25 MHz COMPAQ 386, a fast, IBMPC compatible, microcomputer. The microcomputer acts as the Host for, and controls a commercial ADC unit, the Cambridge Electronic Design 1401. The 1401 has the ability to record 16 multiplexed 12-bit analogue channels and 16 bits of digital data and is controlled by a 4 MHz 65C02 microprocessor. Sampling is controlled by a Z8 coprocessor. Data is stored on an EXABYTE cartridge tape subsystem with a capacity of 2 GBytes which is linked to the COMPAQ *via* a SCSI interface. The 1401 can be programmed to transfer sampled data directly into the host computer's memory using DMA transfers. A program running in the host then transfers this data to the EXABYTE tape unit. This host program is essentially a standard program supplied by Cambridge Electronic Design which has received minor modifications to allow timing information to be saved along with the data. It was decided to use the 1401 to record 5 12-bit analogue signals, one 8-bit analogue signal and an 8-bit digital byte, each sampled at 10000 samples per second. By using the completion of the 8-bit channel sample to time the reading of the input port, this allows the MSF broadcast 60 kHz transmissions to be used as a sampling clock. Since the 12-bit samples are read from the ADC chip as 2 bytes the total data rate was therefore 120000 bytes per second. Experiments showed that this transfer rate could not be maintained reliably, due to delays in transfering data from the 1401 into the COMPAQ's memory. This was probably caused by

delays in resetting the DMA transfer counters at the end of each 64Kbyte transfer. However it was determined that the system would work reliably at a transfer rate of 100000 bytes per second and since 5 12-bit samples are recorded there were 20 bits which were being transferred which contained no useful information. It was decided to develop a 1401 program which would compress the data before transferring it to the COMPAQ so that the desired signals could be recorded reliably. The following program was designed to perform such a compression. During each cycle of the main loop of the program, eleven bytes are read from the ADC chip along with one byte from the digital input chip. These are compressed into ten bytes which are then written to the host interface port. The program was written so that the output bytes could be sent to the host interface separated by at least 20μs. Tests of the program showed that it would work reliably when the sampling clock was increased from 60 kHz to at least 70 kHz. Above this rate the data acqusition system was limited by the rate of the transfer of the data to tape.

```
;Command for the CED 1401 with the IBM interface to read
;5 12-bit ADC channels, 1 8-bit ADC channel and 1 8-bit
;digital input byte at 10kHz each.
;
;ADCHN   - The ADC channel register.
;ADCSR   - The ADC Control/Status Register.
;ADDAT   - The ADC high byte.
;ADLOW   - The ADC low byte.
;BLKSET  - Monitor routine to initialise Block Transfers.
;C4CSR   - Clock four's Control/Status Register.
;DBCSR   - Bi-directional digital port Control/Status Register.
;DBDAT   - Bi-directional digital port.
;DICSR   - Digital input Control/Status Register.
;DIDAT   - Digital input port.
;EVDIS   - External event input disable register.
;EVPOL   - External event and ADC Ext conv.  register.
;HOCSR   - The host Control/Status Register.
;HODIN   - The host input port.
;HODOT   - The host output port.
;RDITEM  - Monitor routine which returns command parameters.
;WRZ8    - Monitor routine which writes A to the Z8 mode register.
;Z8BUSY  - Monitor routine which waits for Z8 to go not busy.
;Z8MOD0  - Monitor routine to put Z8 into mode 0.
;        - (In mode 0 the Z8 does nothing.)
;
REPEAT = C4BP               ;Define an unused address.
BEGIN   .COM  SAMPL60
        .BYTE 0,1           ;Revision level.
        .WORD GO-BEGIN      ;Offset of start of program.
        .WORD 0             ;No initialise entry point.
        .WORD 0,0,0,0       ;Spare words.
```

```
        .BYTE 2,2           ;Must be one parameter.  (2 with name.)
        .BYTE %30           ;It's an unsigned 16-bit integer.
;
GO:     LDX   #%10          ;Get argument passed
        JSR   RDITEM        ;into Y and A.
        STY   REPEAT        ;Save as repeat count, lowbyte
        STA   REPEAT+1      ;and high byte.
        STZ   C4CSR         ;Disable Clock 4.
        STZ   ADCSR         ;Reset ADC chip.
        JSR   Z8MOD0        ;Put Z8 into mode 0 (idle).
        LDA   #%00          ;Positive edge triggering
        STA   EVPOL         ;for events.
        LDA   #%1F          ;Disable all events except
        STA   EVDIS         ;the ADC external convert.
        BIT   ADDAT         ;Clear any data from ADC.
        JSR   SETTRN        ;Set up DMA transfers.
        LDA   #%F0          ;These 4 lines are a bit of CED magic.
        JSR   SNDBYT        ;HOSIZE
        LDA   #%FF
        JSR   SNDBYT        ;HOSIZE+1
        LDA   REPEAT        ;Send back the repeat count
        JSR   SNDBYT        ;so that the Host DMA
        LDA   REPEAT+1      ;will know when to stop
        JSR   SNDBYT        ;the data transfers.
        LDA   #1            ;Put the Z8 into mode one
        JSR   WRZ8          ;(cyclical list of channels).
        LDA   #5            ;There will be six channels.
        JSR   WRZ8
        LDA   #0            ;Channel 0.
        JSR   WRZ8
        LDA   #1            ;Channel 1.
```

```
        JSR   WRZ8
        LDA   #2          ;Channel 2.
        JSR   WRZ8
        LDA   #3          ;Channel 3.
        JSR   WRZ8
        LDA   #4          ;Channel 4.
        JSR   WRZ8
        LDA   #5          ;Channel 5.
        JSR   WRZ8
        JSR   Z8BUSY      ;Wait for the Z8 to finish.
START1: BIT   HOCSR       ;Wait for a trigger char
        BVC   START1      ;to say host is ready.
        LDA   HODIN       ;Clear the character.
        STZ   DICSR       ;Set up digital input port A.
        SEI               ;Disable Interrupts
        LDY   #4          ;Set up digital port DBDAT so we see
        STY   DBCSR       ;data.  (Direction set at power on.)
WAIT:   BIT   DBDAT       ;Wait in this loop for a zero on
        BMI   WAIT        ;bit 7 and then a one on bit 7
WAIT2:  BIT   DBDAT       ;so that we definitely catch a
        BPL   WAIT2       ;transition (to 2.5usec).
        STZ   ADCHN       ;Start Z8 co-processor
        LDX   #8          ;Disable ADC interrupts and
        STX   ADCSR       ;enable external clock conversions.
L0:     BIT   ADCSR       ;Wait for ADC chan 0.
        BPL   L0
        LDA   DIDAT       ;Read digital input port.
        LDX   ADDAT       ;Read ADC chan 0 high byte.
L1:     BIT   HOCSR       ;Wait for host to become free.
        BPL   L1
        STA   HODOT       ;Send digital byte.
```

```
L2:     BIT   ADCSR     ;Wait for ADC chan 1.
        BPL   L2
L3:     BIT   HOCSR     ;Wait for host to become free.
        BPL   L3
        STX   HODOT     ;Send ADC chan 0 high byte.
        LDA   ADLOW     ;Read ADC chan 1 low byte.
        LDX   ADDAT     ;Read ADC chan 1 high byte.
        LSRA            ;Shift ADC chan 1 nibble
        LSRA            ;from bits 4-7 to bits
        LSRA            ;0-3.
        LSRA
L4:     BIT   HOCSR     ;Wait for host to become free.
        BPL   L4
        STX   HODOT     ;Send ADC chan 1 high byte.
L5:     BIT   ADCSR     ;Wait for ADC chan 2.
        BPL   L5
        ORA   ADLOW     ;Combine low nibbles of chans 1 and 2.
        LDX   ADDAT     ;Read ADC chan 2 high byte.
L6:     BIT   HOCSR     ;Wait for host to become free.
        BPL   L6
        STX   HODOT     ;Send ADC chan 2 high byte.
L7:     BIT   HOCSR     ;Wait for host to become free.
        BPL   L7
        STA   HODOT     ;Send combined nibbles.
L8:     BIT   ADCSR     ;Wait for ADC chan 3.
        BPL   L8
        LDX   ADLOW     ;Read ADC chan 3 low byte.
        LDA   ADDAT     ;Read ADC chan 3 high byte.
L9:     BIT   HOCSR     ;Wait for host to become free.
        BPL   L9
        STA   HODOT     ;Send ADC chan 3 high byte.
```

```
L10:    BIT   ADCSR     ;Wait for ADC chan 4.
        BPL   L10
L11:    BIT   HOCSR     ;Wait for host to become free.
        BPL   L11
        STX   HODOT     ;Send ADC chan 3 low byte.
        LDA   ADLOW     ;Read ADC chan 4 low byte.
        LDX   ADDAT     ;Read ADC chan 4 high byte.
        LSRA            ;Shift ADC chan 4 nibble
        LSRA            ;from bits 4-7 to bits
        LSRA            ;0-3.
        LSRA
L12:    BIT   HOCSR     ;Wait for host to become free.
        BPL   L12
        STX   HODOT     ;Send ADC chan 4 high byte.
L13:    BIT   ADCSR     ;Wait for ADC chan 5.
        BPL   L13
        ORA   ADLOW     ;Combine nibbles of chans 4 and 5.
        LDX   ADDAT     ;Read ADC chan 5 high byte.
L14:    BIT   HOCSR     ;Wait for host to become free.
        BPL   L14
        STX   HODOT     ;Send ADC chan 5 high byte.
L15:    BIT   HOCSR     ;Wait for host to become free.
        BPL   L15
        STA   HODOT     ;Send combined nibbles.
        JMP   L0        ;Loop round to repeat whole thing.
                        ;Rely on reset from host to get
                        ;us out of the loop.
SNDBYT: BIT   HOCSR     ;Wait for host to become free.
        BPL   SNDBYT
        STA   HODOT     ;Send the byte passed in A.
        RTS
```

```
SETTRN:  LDA   #0          ;Stack the three bytes that
         PHA               ;are required by the monitor
         PHA               ;routine BLKSET.
         LDA   #4          ;This is a circular mode transfer.
         PHA               ;(must JMP since parameters on stack.)
         JMP   BLKSET      ;Tell host about transfer and return.
```

Bibliography

[Alpar *et al.* 1982] Alpar, M.A., Cheng, A.F., Ruderman, M.A. and Shaham, J. (1982). *Nature,* **300,** 728.

[Bassan *et al.* 1983] Bassan, M., Fairbank, W.M., Mapoles, E., McAshan, M.S., Michelson, P.F., Moskowitz, B., Ralls, K. and Taber, R.C. (1983). In *Proceedings of the Third Marcel Grossmann Meeting on General Relativity,* ed. H. Ning, North-Holland, Amsterdam.

[Bondi 1957] Bondi, H. (1957). *Nature,* **179,** 1072.

[Boughn and Kuhn 1984] Boughn, S.P. and Kuhn, J.R. (1984). *Astrophysical Journal,* **286,** 387.

[Bracewell 1986] Bracewell, R.N. (1986). *The Fourier Transform and its Applications,* Second Edition, McGraw-Hill: New York.

[Braginsky 1977] Braginsky, V.B. (1977). In 8^{th} *International Conference on General Relativity and Gravitation,* August 7–12, 1977. Waterloo, Ontario, Canada.

[Braginsky *et al.* 1985] Braginsky, V.B., Gusev, A.V., Mitrofanov, V.P., Rudenko, V.N. and Yakimov, V.N. (1985). *Uspekhi Fizicheskikh Nauk,* **147,** 422.

[Braginsky and Thorne 1985] Braginsky, V.B. and Thorne, K.S. (1985). *Nature,* **316,** 610.

[Brigham 1988] Brigham, E.O. (1988). *The Fast Fourier Transform and its Applications,* Prentice-Hall: New Jersey.

[Carr 1986] Carr, B.J. (1986). In *Inner Space/Outer Space.* ed. E.W. Kolb, University of Chicago Press.

[Caves 1980] Caves, C.M. (1980). *Physical Review Letters,* **45,** 75.

[Caves 1981] Caves, C.M. (1981). *Physical Review D,* **23,** 8.

[Caves *et al.* 1980] Caves, C.M., Thorne, K.S., Drever, R.W.P., Sandberg, V.D. and Zimmermann, M. (1980). *Reviews of Modern Physics,* **52,** 341.

[Clark *et al.* 1979] Clark, J.P.A., van den Heuvel, E.P.J. and Sutantyo,W. (1979). *Astronomy and Astrophysics,* **72,** 120.

[Chandrasekhar 1970] Chandrasekhar, S. (1970). *Physical Review Letters,* **24,** 611.

[Davis *et al.* 1971] Davis, M., Ruffini, R., Press, W.H. and Price, R.H. (1971). *Physical Review Letters,* **27,** 1466.

[Davis *et al.* 1985] Davis, M.M., Taylor, J.H., Wiesberg, J.M. and Backer, D.C. (1985). *Nature,* **315,** 547.

[Detweiler 1977] Detweiler, S.L. (1977). In 8^{th} *International Conference on General Relativity and Gravitation,* August 7–12, 1977. Waterloo, Ontario, Canada.

[Dewey 1986] Dewey, D. (1986). Ph.D. thesis, Massachussetts Institute of Technology.

[Di Stefano *et al.* 1976] Di Stefano III, J.J., Stubberud, A.R. and Williams, I.J. (1976). *Feedback and Control Systems,* McGraw-Hill: New York.

[Douglass and Braginsky 1979] Douglass, D.H. and Braginsky, V.B. (1979). In *General Relativity, An Einstein Centenary Survey,* eds. S.W. Hawking and W. Israel, Cambridge Universtity Press.

[Drever 1983] Drever, R.W.P. (1983). In *Gravitational Radiation,* eds. N. Deruelle and T. Piran, North Holland: Amsterdam.

[Drever *et al.* 1985] Drever, R.W.P., Hall, J.L., Kowalski, F.V., Hough, J., Ford, G.M., Munley, A.J. and Ward, H. (1983). *Applied Physics B,* **31,** 97.

[Dyson 1969] Dyson, F.J. (1969). *Astrophysical Journal,* **156,** 529.

[Edelstein *et al.* 1978] Edelstein, W.A, Hough, J., Pugh, J.R. and Martin, W. (1978). *Journal of Physics E: Scientific Instruments,* **11,** 710.

[Einstein 1916] Einstein, A. (1916). *Sitzungsber Preuss. Akad. Wiss.*, 688.

[Einstein 1918] Einstein, A. (1918). *Sitzungsber Preuss. Akad. Wiss.*, 154.

[Faller *et al.* 1985] Faller, J.E., Bender, P.L., Hall, J.L., Hils, D. and Vincent, M.A. (1985). In *Procedings of the Colloquium 'Kilometric Optical Arrays in Space'*, Cargese (Corsica) 23–5 October 1984. ESA SP-226.

[Ford 1981] Ford, G.M. (1981). Ph.D. thesis, University of Glasgow.

[Forward 1978] Forward, R.L. (1978). *Physical Review D,* **17,** 379.

[Forward *et al.* 1961] Forward, R.L., Zipoy, D., Weber, J., Smith, S. and Benioff, H. (1961). *Nature,* **189,** 473.

[Fraser 1979] Fraser, D. (1979). *ACM Transactions on Mathematical Software,* Vol. 5, No. 4.

[Friedmann and Schutz 1978] Friedman, J.L. and Schutz, B.F. (1978). *Astrophysical Journal,* **222,** 281.

[Fruchter *et al.* 1988] Fruchter, A.S., Stinebring, D.R. and Taylor, J.H. (1988). *Nature,* **333,** 237.

[Green 1985] Green, R.M. (1985). *Spherical Astronomy,* Cambridge University Press.

[Hellings *et al.* 1981] Hellings, R.W., Callahan, P.S., Anderson, J.D. and Moffett, A.T. (1981). *Physical Review D,* **23,** 844.

[Hough *et al.* 1986] Hough, J., Meers, B.J., Newton, G.P., Robertson, N.A., Ward, H., Schutz, B.F. and Drever, R.W.P. (1986). A British Long Baseline Gravitational Wave Observatory, unpublished report submitted by Glasgow University to the SERC. GWD/RAL/86-001

[Kerr 1986] Kerr, G.A. (1986). Ph.D. Thesis, University of Glasgow.

[Lorrain and Corson 1970] Lorrain, P. and Corson, D.R. (1970) *Electromagnetic Fields and Waves,* Second Edition, W. H. Freeman: San Francisco.

[Levine and Stebbins 1972] Levine, J. and Stebbins, R. (1972). *Physical Review D,* **5,** 1465.

[Linsay *et al.* 1983] Linsay, P., Saulson, P., Weiss, R. and Whitcomb, S. (1983). A Study of a Long Baseline Gravitational Wave Antenna System; report prepared for The National Science Foundation. MIT: Cambridge, Massachusetts (unpublished).

[Mangan 1988] Mangan, J.B. (1988) Ph.D. thesis, University of Glasgow.

[Meers 1983] Meers, B.J. (1983). Ph.D. thesis, University of Glasgow.

[Meers 1988] Meers, B.J. (1988). *Physical Review D,* **38,** 2317.

[Misner *et al.* 1973] Misner, C.W., Thorne, K.S. and Wheeler, J.A. (1973). *Gravitation.* W. H. Freeman: San Francisco.

[Peters and Mathews 1963] Peters, P.C. and Mathews, J. (1963). *Physical Review,* **131,** 435.

[Pines *et al.* 1974] Pines, D., Shaham, J. and Ruderman, M.A. (1974). In *IAU Symposium No 53, Physics of Dense Matter,* ed. C. Hansen, p189. Reidel: Dordrecht.

[Robertson *et al.* 1982] Robertson, N.A., Drever, R.W.P., Kerr, I. and Hough, J. (1982). *Journal of Physics E: Scientific Instruments,* **15,** 1101.

[Saenz and Shapiro 1978] Saenz, R.A. and Shapiro, S.L. (1978). *Astrophysical Journal,* **221,** 286.

[Saenz and Shapiro 1979] Saenz, R.A. and Shapiro, S.L. (1979). *Astrophysical Journal,* **229,** 1107.

[Saenz and Shapiro 1981] Saenz, R.A. and Shapiro, S.L. (1981). *Astrophysical Journal,* **244,** 1033.

[Schilling *et al.* 1981] Schilling R., Schnupp, L., Winkler, W., Billing, H., Maischberger, K. and Rüdiger, A. (1981). *Journal of Physics E: Scientific Instruments,* **14,** 65.

[Schnupp 1985] Schnupp, L., Winkler, W., Maischberger, K., Rüdiger, A. and Schilling, R. (1985). *Journal of Physics E: Scientific Instruments,* **18,** 482.

[Schutz 1986] Schutz, B.F. (1986). *Nature,* **323,** 310.

[Schutz 1988] Schutz, B.F. (1988) Talk presented at the International Symposium on Experimental Physics, Guangzou, China, 3–8 August 1987.

[Shoemaker 1987] Shoemaker, D.H. (1987). Ph.D. thesis, Université de Paris-Sud.

[Shoemaker *et al.* 1988] Shoemaker, D.H., Schilling, R., Schnupp, L., Winkler, W., Maischberger, K. and Rüdiger, A. (1988) *Physical Review D,* **38,** 423.

[Smarr 1977] Smarr, L. (1977). *Annals of the New York Academy of Sciences,* **302,** 569.

[Smarr *et al.* 1983] Smarr, L., Vessot, R.F.C., Lundquist, C.A. and Decher, R. (1983). *General Relativity and Gravitation,* **15,** 129.

[Stark and Piran 1985] Stark, R.F. and Piran, T. (1985). *Physical Review Letters,* **55,** 891.

[Taylor *et al.* 1979] Taylor, J.H., Fowler, L.A. and McCulloch, P.M. (1979). *Nature,* **277,** 437.

[Thorne 1978] Thorne, K.S. (1978). In *Theoretical Principles in Astrophysics and Relativity,* eds. N. Lebovitz, W. Reid and P. Vandervoot. University of Chicago Press.

[Thorne 1987] Thorne, K.S. (1987). In *300 Years of Gravitation,* eds. S.W. Hawking and W. Israel, Cambridge Universtity Press.

[Thorne *et al.* 1979] Thorne, K.S., Caves, C.M., Sandberg, V.D. Zimmermann, M. and Drever, R.W.P., (1979). In *Sources of Gravitational Radiation,* ed. Smarr, L., Cambridge University Press.

[Unruh 1978] Unruh, W.G. (1978). *Physical Review D*, **18**, 1764.

[Wagoner 1984] Wagoner, R.V. (1984). *Astrophysical Journal*, **278**, 345.

[Weber 1960] Weber, J. (1960). *Physical Review*, **117**, 306.

[Weber 1969] Weber, J. (1969). *Physical Review Letters*, **22**, 1302.

[Weiss 1972] Weiss, R. (1972). *Quarterly Progress Report of the Research Laboratory of Electronics of the Massachusetts Institute of Technology*, **105**, 54.

[Weiss *et al.* 1976] Weiss, R., Bender, P.L., Misner, C.W. and Pound, R.V. (1976). *Report of the Sub-Panel on Relativity and Gravitation, Management and Operations Working Group for Shuttle Astronomy.* NASA: Washington, DC.

[Winkler *et al.* 1986] Winkler, W., Maischberger, K., Rüdiger, A., Schilling, R., Schnupp, L. and Shoemaker, D. (1986). In *Proceedings of the Forth Marcel Grossman Meeting on Recent Developments of General Relativity.* ed. R. Ruffini. North Holland: Amsterdam.

[Woodward 1953] Woodward, P.M. (1953). *Probability and Information Theory, with Applications to Radar*, Pergamon: London.

[Zimmermann 1978] Zimmermann, M. (1978) *Nature*, **271**, 524.

[Zimmermann 1980] Zimmermann, M. (1980) *Physical Review D*, **21**, 891.