


**From the Big Bang to Tumbleweeds:
Analysis of signals from relic gravitons, neutron stars, and
terrestrial gravitational noise in gravitational wave detectors**

Thesis by
Teviet Creighton

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2000
(Submitted October 22, 1999)

Acknowledgements

First and foremost, I would like to thank my advisor, Kip Thorne. His guidance and keen insight have been indispensable throughout my work.

I am also grateful to Patrick Brady, Curt Cutler, and Bernard Schutz, with whom I collaborated in my studies of neutron stars. My understanding of these issues also benefited greatly from the radio astronomy expertise of Stuart Anderson, and the physics and data analysis insights of Bruce Allen, Kent Blackburn, Ben Owen, Massimo Tinto, and Alan Wiseman. Scott Hughes has also been a source of helpful suggestions, both on data analysis in general and on gravity gradients in particular. Brad Sturtevant provided important guidance in my excursion into aeronautical physics, and Fred Raab braved thorn and thistle to collect the vital statistics of tumbleweeds.

Other friends and colleagues have also helped me, tangibly or intangibly, with my work at Caltech: Kashif Alvi, Christina Hood, Jim Mason, Patricia Purdue, Arthur Street, Michele Vallisneri, and most especially Pu Chen. Finally I would like to thank Jolien Creighton for leading me into this field, and for his ongoing help and collaboration.

This work was supported by NSF grants PHY-9424337 and AST-9731698, and by NASA grant NAG5-6840.

Abstract

This dissertation explores three separate issues in the field of gravitational-wave astronomy: optimal detection algorithms for quasi-periodic signals, gravitational-wave signatures of the equation of state in the early universe, and local Newtonian gravitational noise from nearby airborne masses as possible contaminants of the gravitational-wave signal.

Continuous quasi-periodic signals are waveforms that maintain phase coherence over times longer than practical observation times, although the phase may drift in a way that can be modeled with few parameters. Sensitivity to such signals is limited by the computational cost of the analysis, especially since the detection algorithm must search over many values of the parameters in the phase model; it is therefore crucial to develop computationally efficient search strategies. One such strategy is a *hierarchical stack search*: a technique combining coherent phase corrections on short stretches of data with incoherent frequency drift corrections among several such stretches. The procedure is repeated at least twice, with each pass increasing the confidence in any putative signal. This dissertation discusses how to choose parameter values and observation times for greatest sensitivity, and shows how several astrophysically interesting sources may be detectable by this method.

A background of gravitational waves originating in the Big Bang or a pre-Big-Bang collapsing universe will not thermalize in any cosmological epoch, but may be amplified by an intermediate epoch when the wavelengths were stretched outside the Hubble radius. The present-day spectral index is related simply and generically to the initial spectrum, and to the cosmological equation of state at the beginning and end of the intermediate epoch. This dissertation derives this relation, and compares it to related but more model-specific formulae in the current literature.

Finally, this dissertation considers two atmospheric sources of background Newtonian gravitational noise (infrasonic pressure waves and wind-advected density perturbations), and two sources of transient Newtonian gravitational signals (atmospheric shockwaves and massive airborne bodies, especially tumbleweeds). Neither background noise source will exceed the noise floor for advanced detectors, but sonic booms and wind-borne debris striking the detector can both produce detectable spurious signals through their gravitational effects. Possible corrective measures are discussed.

Contents

Acknowledgements	ii
Abstract	iii
Table of contents	iv
1 Introduction	1
1.1 General background	1
1.1.1 Interferometric detectors	2
1.1.2 Sources of gravitational waves	6
1.2 Specific background	10
1.2.1 Neutron stars	10
1.2.2 Primordial gravitational waves	12
1.2.3 Local gravitational fields	16
2 Phase-coherent search technique for periodic sources	18
2.1 Introduction	19
2.1.1 Gravitational waves from pulsars	19
2.1.2 Three classes of sources	20
2.1.3 The data analysis problem	23
2.1.4 Summary of results	24
2.1.5 Organization of this paper	26
2.2 Pulsar phenomenology	28
2.2.1 Spindown	28
2.2.2 Proper motions	30
2.2.3 Glitches	31
2.3 Gravitational waves from pulsars	31
2.3.1 Waveform	31
2.3.2 Signal at the detector	32
2.3.3 Parameter space	33
2.4 Data analysis technique	34
2.5 Parameter space metric	37
2.5.1 Mismatch	37

2.5.2	Metric and number of patches	39
2.6	Depth of an all sky search	41
2.6.1	Patch number versus observation time	42
2.6.2	Computational requirements	43
2.7	Computational requirements for a directed search	45
2.8	Future directions	47
3	Time-frequency and hierarchical search techniques for periodic sources	53
3.1	Introduction	54
3.1.1	Gravitational waveform	56
3.1.2	Parameter ranges	57
3.1.3	Search technique	58
3.1.4	Results	61
3.1.5	Organization of the paper	63
3.2	Mismatch	64
3.2.1	Directed search	67
3.2.2	Sky search	68
3.3	Thresholds and sensitivities	69
3.4	Stack-slide search	70
3.5	Hierarchical search: general remarks	74
3.5.1	Thresholds	74
3.5.2	Computational costs	75
3.6	Hierarchical search with stacking	76
3.7	Specialized searches	77
3.7.1	Galactic core pulsars	78
3.7.2	Newborn neutron stars	80
3.7.3	X-ray binaries	81
3.8	Future directions	85
3.9	Appendix A: Patch number formulae	86
3.10	Appendix B: Resampling error	87
4	Gravitational waves and the cosmological equation of state	93
4.1	Introduction	93
4.1.1	Organization of this paper	94
4.2	The cosmological and wave equations	94
4.2.1	Dependence on n	96
4.2.2	The initial and final spectra	98

4.3	Comparison with special cases	99
4.3.1	Standard inflation	99
4.3.2	Standard string-motivated rebound cosmologies	100
4.3.3	Nonstandard equations of state	101
4.4	Conclusions	102
5	Gravity gradient noise from airborne sources	103
5.1	Introduction	103
5.2	Atmospheric pressure waves	105
5.2.1	Ground absorption	110
5.3	Atmospheric temperature perturbations	110
5.3.1	Uniform airflow	113
5.3.2	Potential flow near the end station	114
5.3.3	Vortices	116
5.4	Shockwaves	119
5.4.1	Sonic booms	120
5.4.2	Vetoing shockwave signals	121
5.5	High-speed objects	123
5.6	Conclusions	127
5.7	Appendix: The temperature noise spectrum	128
	Bibliography	132

Chapter 1 Introduction

1.1 General background

Gravitational waves are a form of radiation predicted by Einstein's general theory of relativity (GR), and by most other theories of gravity that are consistent with the special theory of relativity. They are a generic result of the fact that changes in the gravitational field propagate at a finite speed, which, in GR at least, is the same as the characteristic speed c of light (or any other massless field). Simply put, imagine a configuration of masses A that changes over a timescale τ into a new configuration B . The changes in the associated gravitational field ripple outward at a speed c . In the *wave zone*, corresponding to distances from the field source that are $\gg c\tau$, what one sees are the field configurations of A and B , connected by *transverse* fields in an expanding shell of thickness $c\tau$. This is the gravitational wave pulse. At distances r much larger than the dimensions of the source, the expanding shell is spherical, and its field amplitude decreases as $1/r$. Meanwhile the force-producing static or background fields from that source decrease typically as $1/r^2$, so eventually the fields of the wave pulse or wave train come to dominate over the background gravitational field. In this way, dynamic objects at astronomical distances can produce detectable effects even though their direct gravitational influence is almost entirely negligible.

Of course, the above description applies not just to gravitational waves, but to any propagating massless field—in particular, to electromagnetic waves. Indeed, gravitational radiation is simply the gravitational analogue of electromagnetic waves. However, there are certain key differences between electromagnetic and gravitational fields, which have profound effects on the nature of their radiation.

Compared to gravitational fields, electromagnetic fields are much more strongly coupled to electromagnetic charges. However, these charges come in two signs, so that any macroscopic object will be essentially electrically neutral. Thus, electromagnetic radiation is typically an incoherent superposition of signals from the motions of microscopic charges. The radiation is therefore typically at high frequencies, and carries information about the thermodynamic state of the system. (This is only a general rule; centimetric and longer waves tend to be due to the bulk motions of charges in a system.) Also, since the waves couple strongly to charges, they are easily blocked by intervening material, and do not penetrate from the depths of dense emitters. On the plus side, their strong coupling makes them relatively easy to detect.

By contrast, gravitational fields are extremely weakly coupled to gravitational charges (i.e., masses), but these charges come in only one sign. Thus strong gravitational fields are produced by

macroscopic accumulations of mass, and gravitational waves are generated by the coherent motions of these masses. The radiation is therefore typically at low frequencies (below 10^4 Hz or so), and carries information about the bulk distribution and motion of mass in the system. The weak coupling of the field means that the waves are not significantly absorbed by any intervening matter; they carry information from regimes of density that are impenetrable to other forms of radiation. However, this weak coupling also makes them incredibly difficult to detect.

Another technical difficulty involved with detecting gravitational waves is that gravity, at least in GR, is thought to obey Einstein's general principle of relativity: that an isolated, freely-moving observer cannot discern the effects of gravity on the basis of purely local measurements. This is related to the principle of equivalence, in that an apparatus cannot measure a gravitational wave simply from the acceleration of a charge (test mass) in the detector, because gravity will affect the test mass and the detector equally. The only way to detect gravitational waves is by observing differential acceleration between separated test masses (or in extended bodies), where the differential acceleration grows with the distance between the test masses. The gravitational wave is thus a tidal field, as shown in Fig. 1.1. The characteristic amplitude of the wave is normally taken to be the double time integral of the tidal field, which is a dimensionless number. Physically it can be thought of as the fractional change in separation between two free test masses initially at rest. By contrast, electromagnetic waves (also shown in Fig. 1.1) are actual dipolar electric fields, which can produce local accelerations of electric charges within a detector.


Gravitational waves, like electromagnetic waves, come in two polarizations. Fig. 1.1 shows a right-circularly-polarized wave; left-circularly-polarized waves are also possible. More conventionally, one defines the polarization of a wave in terms of polarization states called $+$ and \times . Waves with pure $+$ polarization have tidal forces that stretch outward in the vertical direction and squeeze inward in the horizontal direction for half of the wave cycle, and then reverse for the other half of the wave cycle. Horizontal and vertical are defined with respect to some arbitrary reference frame transverse to the direction of propagation. Waves with pure \times polarization have the same stretching and squeezing behaviour, but in directions 45° off of the horizontal and vertical directions.

1.1.1 Interferometric detectors

In recent years, an international effort has been underway to build detectors capable of measuring the minute tidal fields generated at the Earth by a gravitational wave of astronomical origin. The two basic types of detectors are resonant bar detectors and interferometric detectors.

A resonant bar detector consists of an extended body, constructed of particular materials and suspended in such a way as to minimize the damping effects and thermal noise in its resonant modes. A passing gravitational wave induces differential tidal forces on the parts of the body, causing the body to vibrate. If the wave frequency is close to the resonant frequency of the body, the vibration

Electromagnetism



Gravity




Figure 1.1: Pictorial depiction of the difference between the dipolar electrical field of an electromagnetic wave, and the quadrupolar tidal field produced by a gravitational wave. Some points to note: 1. The gravitational radiation field is a tidal field, inducing *relative* accelerations on separated charges (test masses), whereas the electromagnetic radiation field produces measurable accelerations on individual charges. 2. The gravitational wave completes a cycle for every *half* cycle of the spinning masses. More generally, the fundamental frequency of a gravitational wave is twice the fundamental frequency of the source. Electromagnetic radiation, on the other hand, generally has the same fundamental frequency as the charge motion that produced it.

in the body will be strongly amplified. Motion sensors around the surface of the body measure this vibration. Bar detectors are thus narrowband detectors, being sensitive for the most part only to gravitational waves whose frequencies are confined to the resonance peak of the detector. Spherical bar detectors are equally sensitive to all polarizations of waves coming from any direction; other bar shapes will have preferred directions and sensitivities.

Perhaps a more promising class of detectors, and the one that I will focus on throughout this thesis, are interferometric detectors, which use laser interferometry to measure tiny perturbations in the separations of widely-separated test masses. A typical detector configuration is shown in Fig. 1.2. The test masses are suspended in a vacuum system and isolated from external vibrations. A laser beam that has been modulated at MHz frequencies is split and sent down each arm of the interferometer, where it bounces back and forth several times between the reflective test masses, then the light from the two arms is recombined and sent to the photosensor. The mirrors are normally positioned so that the recombined light interferes destructively (or nearly so). Any change in the relative path length between the two arms will change the phase of the recombined light, resulting in incomplete destructive interference. On timescales less than the storage time of light in each arm, then, the photosensor output will be a \sim MHz carrier wave whose amplitude can be thought of as an instantaneous measure of the relative lengths of the two arms. When a gravitational wave passes, the tidal field induces differential accelerations in the test masses as indicated on the figure. If the frequency of the wave is higher than the natural frequency of the test-mass suspension, the masses can be thought of as free masses. Then the fractional change in the arm length at any instant is just the dimensionless amplitude of the passing wave.

This assumes that the wave is of + polarization and propagating transverse to the interferometer plane; waves with other polarizations or coming from other directions will have reduced sensitivity, reducing the average sensitivity by a factor of $\sim \sqrt{5}$ in gravitational wave amplitude (p. 421 of [1]). This is one disadvantage of interferometric detectors over spherical bar detectors. However, interferometric detectors have the advantage that they can be constructed on large scales (up to several kilometres for terrestrial detectors), magnifying the differential motion produced by a given gravitational wave. Also, since the detector responds basically instantaneously to an incoming wave, it can detect waves over a large range of frequencies.

Currently there are several interferometric gravitational-wave observatories at various stages of construction or testing in various countries around the world. The TAMA project has built and is testing a 300 m interferometer in Japan. The GEO project is building a 600 m interferometer in Germany. The VIRGO project is constructing a 3 km interferometer in Italy. The LIGO project is building three kilometric interferometers in the United States: a 4 km and a 2 km detector in Washington, and a 4 km one Louisiana. Much of the discussion in chapters 2, 3, and 5 has been motivated by the anticipated sensitivities of the two 4 km LIGO interferometers. In addition to




Figure 1.2: Schematic of an interferometric gravitational-wave detector. Hanging test masses are indicated by number. Filled arrows indicate the direction of propagation of the laser beam. Open arrows indicate the differential motion induced in the test masses by a gravitational wave of + polarization moving transverse to the interferometer plane over half of its wave cycle.

these terrestrial interferometers, there is LISA, a proposed space detector, which would use test masses in separate spacecraft in solar orbit. In concept the space detector can be thought of as an interferometric detector, except that it uses phase-coherent active laser transceivers in place of actual mirrors, and does *not* use multiple reflections between the spacecraft to increase the optical path length.

The sensitivity of any interferometric detector is set by the measurement noise in the quantity $h(t) = \delta L(t)/L$, where L is the interferometer arm length and δL is the differential change in the arm lengths measured by the interferometer. The noise is normally quoted as a power spectral density $S_h(f)$, defined by $\langle \bar{h}(f)\bar{h}(f')^* \rangle = S_h(|f|)\delta(f - f')$, where $\bar{h}(f)$ is the Fourier transform of $h(t)$ in the absence of any signal, and $\langle \dots \rangle$ denotes an expectation over all instantiations of (stationary) noise. Sometimes the sensitivity is given as a characteristic noise amplitude $\sqrt{fS_h(f)}$. The noise power spectral density has dimensions of 1/(frequency), whereas the noise amplitude is dimensionless, giving it a certain aesthetic appeal. To go into detail about the various sources of noise that limit each interferometer is beyond the scope of this thesis, not to mention this introduction. Fig. 1.3 shows some anticipated noise spectra for interferometers in LIGO.

1.1.2 Sources of gravitational waves

As discussed earlier, gravitational and electromagnetic radiation are produced by qualitatively very different processes; it therefore makes sense to expect that strong gravitational emitters will not necessarily be strong electromagnetic emitters. Strong sources of electromagnetic radiation will typically be objects containing large numbers of charges in highly excited thermodynamic states, whereas strong sources of gravitational radiation will be massive objects with large bulk velocities. These properties are largely orthogonal: while there are some objects that exhibit both sets of properties, there is no reason to expect the two different populations of strong emitters to have much overlap. One can therefore expect gravitational-wave astronomy to reveal wholly new classes of previously undetected, perhaps even unpredicted, sources.

Nonetheless, current astrophysical theory can predict a number of potential sources of strong gravitational radiation, which serve as tangible objectives for the recent effort to build gravitational-wave detectors. The following is a brief synopsis of some of the predicted sources that are motivating the current observational effort. I group these sources into *burst sources* and *continuous emitters*.

Burst sources: These are sources that produce signals of finite duration, less than the observation time of the detector looking at it. The theoretical sensitivity of the instrument to such a source is fixed by the noise properties of the instrument and the amplitude of the signal from the source. Roughly, a source has a characteristic amplitude at a given frequency equal to the dimensionless amplitude h of its waveform at that frequency, multiplied by the square root of number of wave cycles




Figure 1.3: Characteristic noise spectra for interferometers in LIGO. The solid curve shows the expected noise spectrum for LIGO I, the first interferometers that will be used in LIGO for a gravitational-wave observation run, in 2002 [2]. The dashed curve shows a typical noise curve for LIGO II interferometers based on enhancements to be made in 2005 [3]; it should be noted that the LIGO II interferometers will also be tunable for improved narrow-band performance. The dotted curve shows an anticipated noise spectrum for “advanced” LIGO (LIGO III) interferometers, based on the original LIGO design proposal [4]. The ultimate sensitivities achievable at the LIGO facilities may be better by up to an order of magnitude.

that it spends in a unit logarithmic frequency interval about that frequency. If this substantially exceeds the dimensionless noise amplitude of the detector at that frequency, then the signal is detectable, provided that one scans the detector output with a template or source model that is matched in phase (and, somewhat less critically, in amplitude) to the incoming waveform. Less optimal search techniques will lower the effective sensitivity. Some proposed sources of burst-type signals are:

Binary inspirals. When two bodies orbit, they lose energy through gravitational radiation, causing them to spiral together. The last phases of inspiral before coalescence generate a characteristic “chirp” of gravitational waves increasing in amplitude and frequency. If the bodies are sufficiently massive and sufficiently compact (allowing for high orbital velocities and strong gravitational emissions), the final inspiral frequency will sweep through a detector’s pass-band within the observation time of the instrument, forming a burst signal. Some proposed sources in this category are: neutron star binaries, stellar-mass black hole binaries, binary systems of a neutron star and a stellar-mass black hole, neutron stars or stellar-mass black holes orbiting supermassive black holes in the galactic core, or possibly even two supermassive black holes inspiralling during a galactic merger.

Black hole mergers. When two black holes collide, they will produce a characteristic burst of gravitational radiation, followed by a damped ringdown as the final hole settles into a stationary state. Stellar-mass black holes will produce this burst at the frequencies of tens to hundreds of Hz accessible by ground-based detectors. Supermassive black hole mergers will be at the much lower frequencies (10^{-1} to 10^{-4} Hz) observed by space-based instruments. A great deal of computational effort is underway to model this process numerically, and produce suitable matched templates for signal analysis.

Supernovae. A supernova occurs when the degenerate core of a massive star or an accreting white dwarf exceeds its Chandrasekhar mass limit, causing it to collapse. If the object is spinning, or the collapse is not entirely symmetrical, it may produce a burst of gravitational waves at high frequencies (tens of Hz to several kHz). This is one of the few proposed sources that may produce a strong signal in both the gravitational and electromagnetic continua.

Newborn neutron stars. The product of a supernova explosion could in some cases be a hot, rapidly-rotating neutron star. This star may be highly unstable, and could radiate away most of its angular momentum over the course of a few months by means of gravitational radiation; during this time, the wave frequency would sweep from around a kHz to around a couple hundred Hz or so.

Continuous sources: These are sources whose emissions last much longer than the observation time of the instrument. The sensitivity of a detector to these sources therefore depends not only on the properties of the wave and the detector, but also on the length of time spent observing. The signal amplitude can be estimated as for burst sources, but, instead of using the number of cycles

spent by the wave in a unit logarithmic frequency, one should use the maximum number of cycles that the wave spends during the entire observation; this gives a sensitivity that generally improves as the square root of observing time. Continuous sources can include members of the class of burst sources if it is theoretically or computationally impractical to build a phase-matched template for the entire burst; sensitivity then depends on the square root of the length of the template. Some typical continuous emitters are:

Neutron stars. A rotating neutron star will emit gravitational radiation if it has a density or mass-current distribution that is not symmetric about its rotation axis. The waves will typically be at the second and higher harmonics of the rotation frequency. This source is discussed in more detail in section 1.2.1 and in chapters 2 and 3.

Compact binaries. At the low frequencies of space-based detectors, binary systems of white dwarfs, or of a white dwarf or neutron star around a supermassive black hole, may produce enough gravitational radiation to be detected, but not enough to complete their inspiral over the course of an observation. These sources would be observed as continuous emitters.

Unresolved binaries. White dwarf binaries are sufficiently common in our Galaxy that, at frequencies below about 10^{-4} Hz, there will be too many at any given frequency to resolve one from another over the course of an observation. The signals will then form a broadband stochastic background. Although such a continuous broadband signal is largely indistinguishable from the broadband noise in a given detector, it can be distinguished by correlating the output of two detectors having independent noise. In the correlated output, the amplitude of the common signal grows relative to the independent noise as the square root of the observation time.

Cosmic strings and domain walls. Another source of broadband stochastic background signals, these objects are the products of high-energy phase transitions in the early universe. When the fundamental particle fields in the early universe settled from a symmetric state into a lower energy vacuum state, they may have settled into different states in different parts of the universe. At the junctions of incompatible vacuum states there may be strings or sheets of the original high-energy “false” vacuum. As these objects merge or vibrate they will emit gravitational waves, which form a stochastic background in the universe today.

Primordial gravitons. Perhaps the ultimate achievement of gravitational-wave cosmology would be the detection of gravitational waves that come to us directly from the Big Bang, or even possibly from an imploding universe *before* the Big Bang. These waves are discussed in more detail in section 1.2.2 and chapter 4.

1.2 Specific background

This thesis deals with several widely-separated topics in the field of gravitational wave astronomy. Chapters 2 and 3 deal with data analysis schemes for gravitational waves from rotating neutron stars and similar objects. Chapter 4 discusses the spectrum of primordial gravitational radiation from the early universe. Chapter 5 deals with the noise and spurious signals produced in gravitational-wave detectors by the gravitational field of the Earth's atmosphere, or of airborne objects. Since these topics are quite separate, I provide in the following sections specific background material that will be assumed in subsequent chapters.

1.2.1 Neutron stars

A neutron star is a compact stellar object, having a mass typically in the range of one to two solar masses, but a diameter of 10 to 20 kilometres. Conveniently, these densities allow the stars to rotate rapidly, with periods measured in milliseconds. Pulsars are highly magnetized rotating neutron stars that produce detectable radio pulses as they rotate, by means of synchrotron radiation. Observations of radio pulsars show that many, indeed most, have frequencies in the 10–1000 Hz pass-bands of proposed interferometric detectors. Neutron stars also have intense gravitational fields, such that even neutron stars in other galaxies could produce detectable gravitational-wave signals, provided that the local gravity near the star fluctuates by some significant fraction as they rotate.

Of course the most appealing feature of neutron stars as potential gravitational-wave sources is that, unlike most other putative sources, neutron stars are known unequivocally to exist. Their main disadvantage is that their gravitational-wave amplitudes cannot be estimated, even to order of magnitude, with any certainty. In order to produce fluctuating gravitational tides as it spins, a neutron star must have a non-axisymmetric density distribution, or non-axial internal mass flows that can alter the gravitational field through gravitomagnetic forces (the gravitational analogue of magnetic fields). There is at present no compelling proof that axial symmetry will necessarily be violated to any significant degree.

Nonetheless, the fact that many radio pulsars exist in our Galaxy and have been detected makes them ideal first candidates for gravitational-wave searches. One simply has to pick a pulsar, determine its pulse timing, and correlate the gravitational-wave detector output with the pulse phase (and its harmonics). By this technique, the gravitational signal will grow linearly with observation time, while stochastic noise in the signal will grow only as the square root of observation time. After sufficiently long observation, one will either detect a signal, or at the very least begin to improve on the theoretical upper limits.

However, interesting as this is, a more ambitious goal is to search for signals from neutron stars that have not previously been detected, or whose rotation phase cannot be well characterized in

advance. This is an important task, since the strongest theoretical emitters of gravitational waves are likely to be in this category. For instance, the non-axial density perturbations are likely to be strongest in young neutron stars and to settle out as the star ages. Newborn neutron stars may even have unstable fluid modes, and be in the process of radiating away much of their angular momentum in the form of gravitational waves. In both cases, one might expect the youngest neutron stars to be in other galaxies, or to be shrouded in dust or a supernova remnant, both of which would make radio detection difficult. Accreting material from a companion might also produce nonaxisymmetric matter distributions or drive instabilities, and thus produce copious gravitational radiation; however, the strongest neutron star accretors do not seem to emit strong radio signals, so the gravitational-wave phase evolution will not be known in advance. Finally, there is the simple fact that if strong gravitational emitters are rare, one would do well not to restrict one's search sample to known radio pulsars.

Chapters 2 and 3 discuss the data analysis hurdles that need to be overcome in any search for pulsars or other incompletely-characterized coherent continuous signals; that is, a signal that can be treated as a modulated sinusoid of indefinite duration, where some or all of the frequency and modulation parameters are not known in advance. Simply put, the central problem is that search sensitivities are limited by the amount of computational processing power available for data analysis. Maintaining the advantage of longer coherent (i.e., phase-matched) signals not only requires accumulating larger data sets, but also requires more precise knowledge of the modulation. When the modulation is not fully characterized in advance, one must search over a range of parameters, with a precision that increases with the length of the observation. The result is that the computational power required to keep up with the data increases as some large power of the observation time. When computing resources are limited, one is forced to find an optimal tradeoff between increasing the observation time and maintaining close phase coherence, in order to maximize the sensitivity of the search.

Chapter 2 takes a first-cut look at the problem, using a simple brute-force search algorithm: a stretch of data is demodulated in the time domain according to a parameterized phase model, a power spectrum is then produced using a fast Fourier transform (FFT) and scanned for excess power in any frequency bin. The procedure is repeated over a mesh of parameter values in the parameter space. The modulation effects considered are intrinsic phase shifts due to spindown or gradual motion of the source (parameterized by the Taylor coefficients of the phase function), and the Doppler modulation due to the rotation and revolution of the Earth (parameterized by the position of the source on the sky). The analysis considers the best way to lay out search points in the parameter space, and the optimal tradeoff between the number of parameter values searched and the length of the data stretch. Even considering these optimizations, and given a computer capable of performing 10^{12} floating-point operations per second (1 Tflops), such a brute-force search

can only process about 18 days worth of data in a like amount of time if searching the entire sky for objects with relatively low intrinsic frequency drift (frequencies of $\lesssim 200\text{Hz}$ drifting on timescales of $\gtrsim 1000$ years). For the more dynamic behaviour expected of stronger gravitational emitters (e.g., for frequencies up to 1000Hz changing on timescales under 40 years), such a search algorithm is limited to processing stretches of data under a day in length.

Chapter 3 extends the methods and results of chapter 2 to more sophisticated search techniques, many of which are used by radio astronomers for radio pulsar searches. The first technique is a *stack search*, in which several power spectra from short stretches of data are stacked (added) incoherently to boost signal strength. This gives a lower signal than a single coherent Fourier transform of the entire dataset, but requires fewer computations, and thus can result in higher overall sensitivity in a computationally limited search. The second technique is a *hierarchical search*, in which an initial low-precision, short-observation-time search produces a list of candidate signals (most of which are due to noise), and then attempts to verify (or deny) these signals using a higher-precision, longer-observation-time search. Again, this allows computational resources to be allocated more efficiently, resulting in higher search sensitivities. Chapter 3 also considers two other types of signal modulation: Doppler modulation from a source in a close binary orbit, and stochastic frequency drift due to torques from material accreting onto a source.

The final achievable sensitivities depend strongly on the type of source one is searching for, and on how well one has constrained the parameter space of signal modulations in advance. In general, though, the threshold wave amplitude h_{th} that one can detect with 99% confidence using a 1 Tflops computer is within an order of magnitude of the threshold $h_{3/\text{yr}}$ achievable with 4 months of data if one had *precisely* characterized the phase evolution in advance. These sensitivities admit a number of possible detectable sources for LIGO II or LIGO III detectors, including Galactic core pulsars with dimensionless deformities on the order of 10^{-6} or higher, and newborn neutron stars shedding angular momentum via gravitational radiation at distances of 5 Mpc or so. Perhaps the most promising source would be an accreting low-mass x-ray binary such as Sco X-1, which, if emitting enough gravitational radiation to balance the angular momentum gained by accretion, would be radiating at five times the detectable threshold in narrowbanded LIGO II.

1.2.2 Primordial gravitational waves

Perhaps one of the most intriguing uses of gravitational wave astronomy is to study phenomena in regimes that are completely opaque to electromagnetic radiation, and even to neutrino observations, but not to gravitational radiation. Gravitational waves couple so weakly to matter that they are not absorbed to a significant extent by any concentration of matter that can occur in nature. While their propagation can be affected by the gravitational effects of large masses or by the overall curvature of the universe, these effects do not cause the gravitational radiation to thermalize; the information

content of the gravitational waveforms is essentially unaffected.

This feature of gravitational radiation is particularly important when considering primordial radiation from the early universe. The universe was opaque to electromagnetic radiation until the era of recombination, about a million years after the Big Bang, when the ionized plasma filling the universe cooled enough to recombine into neutral gas. This is the visible “surface” of the universe for electromagnetic radiation. By contrast, in the standard Big Bang model, primordial gravitational waves were created out of Planck-scale oscillations when the universe was only 10^{-43} seconds old and governed by the enigmatic laws of quantum gravity, and these waves have not been significantly absorbed by any process in the universe since that time. Gravitational waves therefore have the potential of directly transmitting information about the quantum-gravitational structure of the universe during the Planck era.

However, while primordial gravitational waves have not been absorbed or thermalized to any significant extent since their creation, they have been affected in predictable ways by the large-scale curvature of the expanding universe. This is in fact a prerequisite for the possibility of detectable primordial gravitational waves in the present day. Ordinarily when an oscillatory radiation field exists in an expanding universe, the amplitude of oscillation is decreased and the wavelength is increased in proportion with the increasing cosmological scale factor, normally denoted a . Quantum mechanically, this corresponds to conservation of the *number* of radiation quanta in any mode, while simultaneously redshifting the frequency (and hence lowering the energy) of each mode. If this were the only effect at work, then the number of gravitons in any mode today would be unchanged since the Planck era. The usual default assumption is that the universe began in a quantum-gravitational ground state, with only vacuum fluctuations in the gravitational field; this would result in no detectable gravitons in the present era.

Fortunately, the situation can change if the cosmological scale factor a increases faster than a linear function of cosmological proper time t . This is the criterion for an accelerated, or *inflationary*, universe. In an inflationary universe, the cosmological scale a (which determines the distance between points) is increasing as a higher power of t than the Hubble scale $a/(da/dt)$ (which determines the size of a causally connected region of spacetime). Eventually the wavelength of a given gravitational wave mode may get stretched outside the cosmological horizon, so that points in space one wavelength apart are no longer in causal contact. When this happens, the oscillation of the wave is suppressed, and what was previously a propagating fluctuation in the spacetime metric becomes a constant offset of the metric in each independent, causally disconnected region. The dimensionless amplitude of the gravitational wave quickly approaches a constant value, where it remains as long as the wave is outside the cosmological horizon.

After a (hypothetical) period of inflation, the universe enters an era of decelerated expansion, with a increasing more slowly than a linear function of t , and thus more slowly than the Hubble

scale. In this era, wave modes that were previously out of causal contact with themselves come back into causal contact, and begin oscillating again. At this point their amplitude begins to decrease as a^{-1} once again. However, the intermediate period when the wave amplitude was constant means that the amplitude has, on average, decreased more slowly than a^{-1} . This is called *parametric amplification* of the wave, since, in quantum mechanical terms, particle creation has occurred. It is *stimulated* particle creation if it amplifies an excited state of a mode (i.e., the mode originally contained real gravitons), or *spontaneous* particle creation if it amplifies the vacuum fluctuations (i.e., the ground-state waveform) of a mode. The latter is of particular interest to cosmologists, since it means that inflation could create a detectable graviton spectrum even if the initial state of the universe was a quantum-gravitational vacuum. The shape of the spectrum could then be used to characterize the expansion of the universe during the inflationary and post-inflationary eras, or even to test the assumption of a vacuum initial state. The most ambitious goal of such observations would be to probe the state of the universe at the very instant that space and time came into being.

Perhaps even more intriguing is the possibility of using gravitational waves to make observations of the universe *before* the Big Bang. Recent work in string theory has suggested that a collapsing universe might experience an era of strongly-coupled strings that could halt the collapse and cause the universe to rebound. Thus the “Big Bang,” i.e., the start of expansion, might be a string-energy-scale event rather than a Planck-energy-scale event. Since the string energy scale is thought to be at least one or two orders of magnitude lower than Planck energies, the universe would remain transparent to gravitational waves throughout the rebound.

Furthermore, during the collapse phase the cosmological scale factor a is expected to be experiencing accelerated contraction (i.e., as a lower power of $-t$ that is less than 1). This causes gravitational wave modes to be pushed outside of the cosmological horizon, as in the case of accelerated expansion. The parametric amplification, however, is dominated by coupling between the gravitational waves and the cosmological curvature around the time of the bounce (when the Hubble scale becomes undefined). This amplification is harder to see intuitively than in the case of inflation, but follows relatively simply from integrating Einstein’s field equations for a small perturbation on a homogeneous, isotropic cosmology. The result is that gravitational waves (whether they are real gravitons or vacuum oscillations) from the pre-Big-Bang universe will not only pass through the Big Bang unscathed, but may even be amplified to detectable levels in the process.

Chapter 4 discusses how one could use observations of a primordial gravitational wave background to measure certain properties of the universe during the Planck and inflationary eras, or during the pre-Big-Bang era, as the case may be. Simply put, the parametric amplification factor depends on the size of the scale factor a at the times when a particular gravitational wave was forced out of and when it reentered the Hubble radius. Since this occurred at different times for different wavelengths, the power-law index of the gravitational spectrum over some frequency band can be used to compute

some combination of the instantaneous power-law behaviours of a when waves of those frequencies exited and reentered the Hubble radius. This in turn can be used to study the properties of the dominant energy field in the universe at those times, since this field is what determines the rate of expansion or collapse. Or, if one knows the dominant energy field, and hence the expansion rate, one can factor out the amplification factor from the observed spectrum to obtain the initial spectrum of the Big Bang or pre-Big-Bang universe.

Unfortunately, the simplest and standard assumption of vacuum initial conditions, followed by exponential inflation, followed by evolution dominated by a relativistic radiation field, followed finally by pressureless matter-dominated evolution, gives a gravitational spectrum that is decreasing with frequency in a well-determined way. Since observations of the cosmic microwave background place an upper limit on the total amplitude of cosmological perturbations (from gravitational waves and other sources) at long wavelengths, one can extrapolate this to get an upper limit on the perturbations in the pass-bands of proposed detectors. For interferometric observatories such as LIGO the signal is a discouraging two orders of magnitude below the sensitivity of even advanced detectors, assuming a four-month observation, and increases only as the square root of the observation time. For LISA, which observes at lower frequencies, the sensitivity can approach the upper limit imposed by the microwave background [5]. However, this discouraging estimate has prompted many theorists to propose more elaborate cosmological models, including models with post-inflationary energy fields other than matter or relativistic radiation [6], and of course the string-motivated pre-Big-Bang cosmologies [7]. All of these models can have graviton spectra that increase with frequency over some band, and can, for certain choices of their parameters, produce a detectable background.

Given the large number of cosmological models in the literature, and the larger number of parameters that can influence the gravitational spectrum, the intent of my analysis in chapter 4 is not to show how one can produce a detectable background. Instead, the intent is to show how such a background, if detected, can measure certain general properties of the early universe, independent of the details of any particular cosmological model. Specifically, the spectral index (logarithmic slope) of such a background depends only on three parameters: the parameter $\gamma \equiv p/\rho$ of the equation of state at the time when the waves exited the Hubble radius ($\gamma = \gamma_i$) and when they reentered ($\gamma = \gamma_f$), and the average quantum occupation number $N(n)$ of a mode, where $n = a\omega$ is the conformal (expansion-invariant) frequency of the mode. Cosmological background spectra are usually measured by a dimensionless quantity $\Omega(\omega) = (\omega/\rho_c)(d\rho_{\text{gw}}/d\omega)$, where $d\rho_{\text{gw}}$ is the gravitational-wave energy density in a frequency interval $d\omega$ about the frequency ω , and ρ_c is the critical energy density required to make the universe spatially flat. The spectral index of this quantity produced by an inflationary universe is

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{\gamma_i + 1}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right),$$

and for a bounce cosmology it is

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{2\gamma_i}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right).$$

1.2.3 Local gravitational fields

Gravitational waves are not the only source of fluctuating gravitational fields at a detector site. The test masses in a gravitational-wave detector will also be influenced by the ordinary Newtonian gravitational forces of nearby moving masses. These forces are not gravitational radiation, but are the (effectively) instantaneous near-field forces at distances $r \ll c\tau$, where c is the speed of light (or gravitational) propagation, and τ is the timescale over which the field is varying. The strength of these forces decreases as r^{-2} , rather than r^{-1} as for gravitational radiation. Collectively these near-field forces are called *gravity gradients*, since they are simply gradients $\vec{\nabla}\Phi$ of the Newtonian potential field Φ . I should note that, in areas of research other than gravitational wave detection, the term “gravity gradient” is more often used to refer to the gradient in the gravitational field; that is, to the instantaneous local tidal field tensor $\nabla_a \nabla_b \Phi$. However, throughout this thesis I will continue to use the term to refer to the instantaneous local gravitational field vector $g_a = \nabla_a \Phi$, so as to distinguish it clearly from the tidal field tensor of the propagating gravitational wave.

For terrestrial detectors, the test masses are surrounded by other masses, both inside and outside of the detector facilities. Some of these masses may be in motion, and thus cause fluctuations in the gravitational force felt by the test mass. Although these masses and their motions are much smaller than the astrophysical sources of gravitational radiation, what they lack in luminosity they can make up for in proximity. The three basic sources of gravity gradients for terrestrial detectors are: density perturbations of the ground, density perturbations of the atmosphere, and individual massive objects. In each case, one is concerned primarily with time variations in these perturbations on the order of 0.3 seconds to fractions of a millisecond, corresponding to the 3 Hz to \sim kHz pass-bands of terrestrial gravitational-wave detectors. Fortunately, fluctuations in the local gravity tend to be strongly decreasing with frequency in this frequency range. Unfortunately, it is extremely difficult to isolate or shield the test masses from these forces; typically, the only way to reduce gravity gradient noise is to eliminate the source of the fluctuating field.

Density perturbations in the ground in the relevant frequency range are caused by a stochastic background of pressure and shear waves propagating through the upper layers of the Earth’s crust. Of primary concern for gravitational-wave detectors are waves in the upper few tens of metres of the crust: both pressure waves (P-waves), which generate the strongest density perturbations, and shear waves (S-waves), which can slightly raise and lower the level of the surface of the earth. These waves are not discussed in this thesis, but are dealt with in great detail by Hughes and Thorne [8].

Density perturbations in the atmosphere are primarily of two types: adiabatic pressure pertur-

bations (sound waves), which propagate at the speed of sound (about 330 m/s at the surface of the Earth), and isobaric temperature perturbations, which are carried past the detectors at the wind speed (anywhere from 0 m/s to 30 m/s or so). Sound waves can be caused by wind, weather, or human activity in the vicinity of the detector, whereas temperature perturbations are largely caused by solar heating and are then mixed into the surface airmass by turbulent convection and advection. These processes have been studied by micrometeorologists, and by optical astronomers who are concerned with the effects these perturbations have on optical seeing. Chapter 5 of this thesis applies some of the results of these studies to estimate the atmospheric gravity gradient noise in a LIGO-type interferometric detector, with the result that neither source of gravity gradient noise will dominate the noise floor even for advanced detectors. However, the gradients from temperature perturbations advected along turbulent vortices come close enough to the “advanced” LIGO (LIGO III) noise floor to warrant further study. Sonic impulses or shockwaves, such as those produced by supersonic projectiles or aircraft, can also produce transient signals in LIGO through their gravity gradients. Chapter 5 gives an analysis of these signals and their detectability, and concludes that the sonic boom of a supersonic aircraft is potentially a significant contaminant of the gravitational-wave signal; however, such a spurious signal should be easy to veto using external infrasound monitors.

Finally, there are a large number of individual objects moving in and around detector facilities whose gravity gradients could perturb the test masses. Inside the facility one must be concerned with the movements of people and equipment. This is discussed by Thorne and Winstein [9]. Fortunately, it is possible to regulate these human activities inside the facility to reduce the gravity gradient noise. Outside the facility, things are not so easy to control. Gravity gradients can be produced by anything from a passing rabbit or duck, to a rifle bullet, to a wind-blown tumbleweed. In some cases these objects may produce detectable signals through their gravity gradients. However, since these are transient signals, there is some hope that signal discrimination techniques will be able to veto these events, so as not to reduce the sensitivity to true gravitational waves. Chapter 5 concludes with an analysis of the most significant signals from outside objects. While most objects moving smoothly past the detector do not create significant signals in the detector pass-band, a massive object, such as a tumbleweed, that *collides* with the outer wall of the detector facility could produce a detectable signal. A fence $\sim 30\text{m}$ out from the detector test masses should reduce such events to undetectable amplitudes.

Chapter 2 Phase-coherent search technique for periodic sources

This chapter has appeared previously as the article “Searching for periodic sources with LIGO,” by Patrick R. Brady, Teviet Creighton, Curt Cutler, and Bernard F. Schutz [10].

Summary: We investigate the computational requirements for all-sky, all-frequency searches for gravitational waves from spinning neutron stars, using archived data from interferometric gravitational wave detectors such as LIGO. These sources are expected to be weak, so the optimal strategy involves coherent accumulation of signal-to-noise using Fourier transforms of long stretches of data (months to years). Earth-motion-induced Doppler shifts, and intrinsic pulsar spindown, will reduce the narrow-band signal-to-noise by spreading power across many frequency bins; therefore, it is necessary to correct for these effects before performing the Fourier transform. The corrections can be implemented by a parameterized model, in which one does a search over a discrete set of parameter values (*points* in the parameter space of corrections). We define a metric on this parameter space, which can be used to determine the optimal spacing between points in a search; the metric is used to compute the number of independent parameter-space points N_p that must be searched, as a function of observation time T . This method accounts automatically for correlations between the spindown and Doppler corrections. The number $N_p(T)$ depends on the maximum gravitational wave frequency and the minimum spindown age $\tau = f/\dot{f}$ that the search can detect. The signal-to-noise ratio required, in order to have 99% confidence of a detection, also depends on $N_p(T)$. We find that for an all-sky, all-frequency search lasting $T = 10^7$ s, this detection threshold is $h_c \approx (4-5)h_{3/\text{yr}}$, where $h_{3/\text{yr}}$ is the corresponding 99% confidence threshold if one knows in advance the pulsar position and spin period.

We define a coherent search, over some data stream of length T , to be one where we apply a correction, followed by an FFT of the data, for every independent point in the parameter space. Given realistic limits on computing power, and assuming that data analysis proceeds at the same rate as data acquisition (e.g., 10 days of data gets analyzed in ~ 10 days), we can place limitations on how much data can be searched coherently. In an all-sky search for pulsars having gravity-wave frequencies $f \leq 200\text{Hz}$ and spindown ages $\tau \geq 1000\text{Yrs}$, one can coherently search ~ 18 days of data on a teraflops computer. In contrast, a teraflops computer can only perform a ~ 0.8 -day coherent search for pulsars with frequencies $f \leq 1\text{kHz}$ and spindown ages as low as 40Yrs.

In addition to all-sky searches we consider coherent directed searches, where one knows in advance

the source position but not the period. (Nearby supernova remnants and the galactic center are obvious places to look.) We show that for such a search, one gains a factor of ~ 10 in observation time over the case of an all-sky search, given a 1Tflops computer.

The enormous computational burden involved in coherent searches indicates a need for alternative data analysis strategies. As an example we briefly discuss the implementation of a simple hierarchical search in the last section of the paper. Further work is required to determine the optimal approach.

2.1 Introduction

The direct observation of gravitational waves is a realistic goal for the kilometer-scale interferometers which are now under construction at various sites around the world [4, 11]. However, the battle to see these waves is not over when the detectors are constructed and running. Searching for gravitational wave signals in the interferometer output presents its own problems, not least of which is the sheer volume of data involved.

Potential sources of gravitational waves fall roughly into three classes: bursts, stochastic background, and continuous emitters.

Burst sources produce signals which last for times considerably shorter than available observation times. The chirp signals from compact coalescing binaries belong to this class. Since theoretical waveforms, valid during the inspiral phase of the binary evolution, have been accurately calculated using post-Newtonian methods [12], it is possible to search the data stream for chirps using matched filtering techniques. Detailed studies have been carried out to ascertain the optimal set of search templates [13, 14], and a preliminary investigation of search algorithms is now under way [15]. Detection of other, not so well understood, sources in this class—e.g., non-axisymmetric supernovae—has received limited attention [1].

Flanagan [16] has determined how to cross-correlate the output of two detectors in order to search for a stochastic background of gravitational radiation, which was implemented by Compton [17] and applied to data taken during a period of 100 hours by two prototype interferometers detectors in Glasgow and Garching [18]. In [5], Allen presents a detailed discussion of the potential significance of detecting a stochastic background. Compton's work, and simulations performed by Allen, have demonstrated that this kind of analysis requires minimal computational resources.

In this paper we consider some issues involved in searching for continuous wave sources. Throughout our discussion we use pulsars as a guide to develop a search strategy.

2.1.1 Gravitational waves from pulsars

Rapidly rotating neutron stars (pulsars) tend to be axisymmetric; however, they must break this symmetry in order to radiate gravitationally. The pulsar literature contains several mechanisms

which may lead to deformations of the star, or to precession of its rotation axis, and hence to gravitational wave emission. The characteristic amplitude¹ of gravitational waves from pulsars scales as

$$h_c \sim \frac{If^2\epsilon}{r} \quad (2.1)$$

where I is the moment of inertia of the pulsar, f is the gravitational wave frequency, ϵ is a measure of the deviation from axisymmetry and r is the distance to the pulsar.

Pulsars are thought to form in supernova explosions. The outer layers of the star crystallize as the newborn pulsar cools by neutrino emission. Estimates, based on the expected breaking strain of the crystal lattice, suggest that anisotropic stresses, which build up as the pulsar loses rotational energy, could lead to $\epsilon \lesssim 10^{-5}$; the exact value depends on the breaking strain of the neutron star crust as well as the neutron star's ‘geological history’, and could be several orders of magnitude smaller. Nonetheless, this upper limit makes pulsars a potentially interesting source for kilometer scale interferometers. Figure 2.1 shows some upper bounds on the amplitude due to these effects.

Large magnetic fields trapped inside the superfluid interior of a pulsar may also induce deformations of the star. This mechanism has been explored recently in [19], indicating that the effect is extremely small for standard neutron star models ($\epsilon \lesssim 10^{-9}$).

Another plausible mechanism for the emission of gravitational radiation in very rapidly spinning stars is the Chandrasekhar-Friedman-Schutz (CFS) instability, which is driven by gravitational radiation reaction [20, 21]. It is possible that newly-formed neutron stars may go through this instability spontaneously as they cool soon after formation. The radiation is emitted at a frequency determined by the frequency of the unstable normal mode, which may be less than the spin frequency.

Accretion is another way to excite neutron stars into emitting gravitational waves. Wagoner [22] proposed that accretion may drive the CFS instability. There is also the Zimmermann-Szedinitz mechanism [23] where the principal axes of the moment of inertia are driven away from the rotational axes by accretion from a companion star. Accretion can in principle produce relatively strong radiation, since the amplitude is related to the accretion rate rather than to structural effects in the star. However, accreting neutron stars will be in binary systems, and these present problems for detection that go beyond the ones we discuss in this paper. We hope to return to the problem of looking for radiation from orbiting neutron stars in a future publication.

2.1.2 Three classes of sources

Observed pulsars fall roughly into two groups: (i) young, isolated pulsars having periods of tens or hundreds of milliseconds, and (ii) older, millisecond pulsars. The young pulsars are most likely to deviate significantly from axisymmetry; however, they are generally observed to have low frequencies,

¹We adopt the definition of h_c provided in Eq. (50) of Thorne [1].

so that there is a competition between the frequency, f , and deviation from axisymmetry, ϵ , in Eq. (2.1). On the other hand, millisecond pulsars, whose waves are higher in frequency, tend to be quite old and well annealed into an axisymmetric configuration.

Radio observations can only probe a small portion of our galaxy in searching for pulsars. A significant effect reducing the depth of radio searches is dispersion of the signal by galactic matter between potential sources and the earth. Given current evolutionary scenarios for pulsars—that they are born in supernova explosions—it seems likely that most pulsars should be located in the galactic disk, and the youngest of these will also be shrouded in a supernova remnant, making them invisible to radio astronomers.

Blandford [24, 1] has pointed out that there could exist a class of pulsars which spin down primarily due to gravitational radiation reaction. For sources in this class the frequency scales as $f \propto \tau^{-1/4}$, where τ is the age of the pulsar. If the mean birth rate for such pulsars in our galaxy is τ_B^{-1} , the nearest one should be a distance $r = R\sqrt{\tau_B/\tau}$ from earth, where $R \simeq 10\text{kpc}$ is the radius of the galaxy. The intrinsic gravitational wave amplitude (that is, the amplitude h at some fixed distance) of a pulsar in this class is proportional to $\tau^{-1/2}$. Thus, the nearest source in this class would have a dimensionless amplitude h_c at the Earth

$$h_c \simeq 8 \times 10^{-25} \left(\frac{200\text{Yrs}}{\tau_B} \right)^{1/2}. \quad (2.2)$$

In arriving at this expression we have assumed that the age τ of typical pulsars in this class is much less than the age of our galaxy, so that the population has reached a steady state. This means that the gravitational ellipticity and the gravitational wave frequency must satisfy $\epsilon^2 \gg 1.2 \times 10^{-18} (\text{kHz}/f)^4$. Assuming the existence of such a class of pulsars, with $\tau_B \lesssim 2 \times 10^4$ Yrs, we see from Fig. 2.1 that there is a large region of parameter space that is both (i) detectable by the LIGO detector and (ii) physically reasonable, in the sense that $\epsilon < 10^{-5}$ and f lies in the range 200–1000 Hz.

Note that Blandford's argument can be slightly re-cast to yield an upper limit on the gravitational wave strength of *any* isolated pulsar—i.e., any pulsar whose radiated angular momentum is not being replenished by accretion. The age of an isolated pulsar must be shorter than the age computed assuming the spin-down is solely due to gravitational wave emission. Correspondingly, if we set τ_B equal to 40 yrs (corresponding to the birthrate for *all* pulsars), we get the following upper limit for measured gravitational wave amplitude of an isolated pulsar: $h_c < 2 \times 10^{-24}$. Of course, this is a statistical argument. This bound could certainly be violated by an isolated pulsar that just happens to be anomalously close to us.

It is important that any search strategy should be general enough to encompass all three of the above classes, allowing for the significant changes in frequency which may be inherent in the sources (see section 2.2).




Figure 2.1: Characteristic amplitudes h_c [see Eq. (2.11)] for several postulated periodic sources, compared with sensitivities $h_{3/\text{yr}}$ of the initial and advanced detectors in LIGO. ($h_{3/\text{yr}}$ corresponds to the amplitude h_c of the weakest source detectable with 99% confidence in $\frac{1}{3}\text{yr} = 10^7\text{s}$ integration time, if the frequency and phase of the signal, as measured at the detector, is known in advance.) Long-dashed lines show the expected signal strength as a function of frequency for pulsars at a distance of 10 kpc, assuming non-axisymmetries of $\epsilon = 10^{-5}$ and $\epsilon = 10^{-8}$, where ϵ is defined in section 2.3A. Upper limits are also plotted for the Crab and Vela pulsars, assuming their entire measured spindown is due to gravitational wave emission. The dotted lines indicate the strongest waves received at the earth for Blandford's hypothetical class of pulsars; each line corresponds to a particular birth rate.

2.1.3 The data analysis problem

The detection of continuous, nearly fixed frequency waves will be achieved by constructing power spectrum estimators and searching for statistically significant peaks at fixed frequencies. In practice, this is achieved by calculating the amplitude of the Fourier transform of the detector output given by applying a fast Fourier transform (FFT), a discrete approximation to the true Fourier transform:

$$\tilde{h}(f) = \frac{1}{\sqrt{T}} \int_0^T e^{2\pi ift} h(t) dt . \quad (2.3)$$

The main hope of detection lies in the fact that one may observe the sky for long time periods of time T . When such a data stretch is transformed to make the underlying signal monochromatic, the signal to noise ratio grows as \sqrt{T} in amplitude (or as T in the power spectrum). One will likely need to have integration times of several weeks or months in order for the expected signals from nearby sources to rise above the noise. However, such long data stretches pose a significant computational burden; using 10^7 s of data to look for signals with gravitational wave frequencies up to 500Hz requires calculating an FFT with $N \simeq 10^{10}$ data points. Calculation of a single such FFT would take about 1s on a 1Tflops computer, assuming that all 10^{10} points can be held in fast memory. Unfortunately, this is not the whole story.

The detection problem is complicated by the fact that the signal received at the detector is not perfectly monochromatic. Earth-bound detectors participate in complex motions which lead to significant Doppler shifts in frequency as the Earth rotates, and as it orbits around the sun (this orbit is significantly perturbed by the moon and the other planets). The time dependent accelerations broaden the spectral lines of fixed frequency sources spreading power into many Fourier bins about the observed frequency. In order to maintain the benefit of long observation times, it is therefore necessary to remove the effects of the detector motion from the data stream. This can be achieved by introducing an inertial (barycentered) time coordinate and carrying out the FFT with respect to it. The difficulty of doing this was first estimated by one of us [25]. However, we must also consider the additional complication that the signal may not be intrinsically monochromatic. If the signal exhibits intrinsic frequency drift, or modulation, due to the nature and location of the source — as is expected for pulsars which spin down with time — these effects can also be removed in the transformation to the new time coordinate.

Unfortunately, the demodulated time coordinate depends strongly on the direction from which the signal is expected, and on the intrinsic frequency evolution one assumes for the source. Thus, in searching for sources whose position and timing are not well known in advance, one must apply many different corrections to the data, performing a new FFT after each correction. Given the possibility that the strongest sources of continuous gravitational waves may be electromagnetically invisible or previously undiscovered, an *all sky, all frequency* search for such unknown sources is

of considerable interest. To obtain some idea of the magnitude of this task, consider searching the entire sky for signals with (fixed) frequencies up to 500Hz using 10^7 s worth of data. Assuming the entire data stream could be held in fast memory on a machine capable of 1Tflops, it would take 10^8 s to complete the search. Introducing intrinsic spindown effects into the search increases the computational cost, at fixed integration time, by many orders of magnitude. This computational cost is the central problem of searching for unknown pulsars in the output from gravitational wave detectors and is the focus of this paper.

2.1.4 Summary of results

We parameterize the space of pulsar signals by the position of the source on the sky $\{\theta, \phi\}$, entering through Doppler shifts due to the detector's motion, and by spindown parameters f_k which characterize the intrinsic frequency evolution. [See Eq. (2.13).] We constrain the range of possible values of the spindown parameters using the (spindown) age $\tau = f/\dot{f}$ of the youngest pulsar that a search can detect, thus $|f_k| \leq \tau^{-k}$. For the computationally-intensive search over all sky positions and spindown parameters, it is important to be able to calculate the smallest number of independent parameter values which must be sampled in order to cover the entire space of signals. We have accomplished this by introducing a distance measure and corresponding metric on the parameter space. The analysis is patterned after a similar one developed by Owen [14] for gravitational waves from inspiralling, compact binaries. Using our metric one can compute the volume of parameter space, thus inferring the number of independent points that must be sampled in order to cover the entire space. We define a *coherent search* to be one where we perform one demodulation and FFT of the data for every independent point in the parameter space. Besides telling us the computational requirements for a coherent search, the metric approach tells us how to place the points most efficiently in parameter space, in a similar way to that discussed by Owen.

We have found it useful to present the results based on several possible search strategies, which cover different regions of the parameter space. Accordingly, we define a pulsar to be *old* if its spindown age τ is greater than 10^3 Yrs and *young* if $\tau \gtrsim 40$ Yrs. A pulsar is considered to be *slow* if its gravitational wave frequency is $f \lesssim 200$ Hz and *fast* if $f \lesssim 10^3$ Hz.

A coherent all-sky search of 10^7 seconds of data for old, slow pulsars requires approximately 1.1×10^{10} independent points in the parameter space; only one spindown parameter is needed to account for intrinsic frequency evolution. In contrast, an all-sky search for fast, young pulsars in 10^7 seconds of data requires 8×10^{21} independent parameter space points to be sampled, using three spindown parameters to model intrinsic frequency evolution. Note that searches for old, fast pulsars (such as known millisecond radio pulsars) and young, slow pulsars (younger brothers of the Crab and Vela) are automatically subsumed under the latter search. These results mean the following. Assuming unlimited computer power and stationary, Gaussian statistics, a pulsar with

unknown position and period must have strain $h_c \approx 4.3h_{3/\text{yr}}$, if it is in our ‘old, slow’ category, and $h_c \approx 5.1h_{3/\text{yr}}$, if it is in our ‘young, fast’ category, to be detected with 99% confidence in a 10^7 second search. Here $h_{3/\text{yr}}$ is the strain required for detection with 99%-confidence in a 10^7 second integration, assuming the pulsar position and period are known in advance²:

$$h_{3/\text{yr}}(f) = 4.2 \sqrt{S_n(f) \times 10^{-7} \text{Hz}} . \quad (2.4)$$

Thus, when considering an all-sky, all-frequency pulsar search, the LIGO sensitivity curves shown in Fig.1 effectively overestimate the detector’s sensitivity by a factor of $\sim 4\text{--}5$, even in the limit of infinite computing power.

Our ability to perform searches for continuous waves will certainly be limited by the available computing resources. Assuming realistic computer power — say of order 10^{13} flops — we estimate that computing limitations will effectively reduce the sensitivity of the detector by another factor of ~ 2 , even for some reasonably optimized and efficient search strategy. However, more work will be needed to develop an optimized algorithm, and thus to refine this latter estimate.

While the concept of the metric is introduced in the framework of an all-sky search for unknown pulsars, it is clear that we may use the same approach to examine the depth of a search over limited regions of the parameter space. In particular, once the scope of a search is decided, the optimization procedure discussed in section 2.6 can be used to determine the observation time and grid spacing which maximizes the expected sensitivity of a search. As an example, we consider coherent *directed searches*, in which one assumes a specific sky position (such as a particular cluster or supernova remnant) and searches only over spindown parameters. Again, we present results for two concrete scenarios based on fast, young pulsars and old, slow pulsars. Similar considerations apply to directed searches as to all-sky searches; that is, the curves in Fig. 2.1 overestimate the detector sensitivity for 10^7 second integration. Table 1 summarizes the results for both cases.

We note that in each type of search, the number of parameter space points, and hence the computational requirements, were reduced significantly by the assumption that the points were placed with optimal spacings given by the metric formalism. Nevertheless, the bottom line is that limitations on computational resources will severely restrict the integration times that can be achieved. Assuming access to a Tflops of computing power (effective computational throughput, ignoring possible overheads due to interprocessor communication or data access), we find the following limits on coherent integration times: For young, fast pulsars we are limited to about 0.8 days for an all-sky search, and 18 days for a directed search. For older, slower pulsars, on the other hand, we are only limited to 9 days for an all-sky search, and nearly 160 days for a directed search. The threshold sensitivities that these strategies can achieve, relative to the noise curves in Fig. 2.1, are plotted as functions of

²This differs from equation (112) in [1] because we have specified 99% confidence, and we have used the correct exponential probability function for power.

Table 2.1: The number of independent parameter points $N_p(T, \mu_{\max} = 0.3)$ required for a coherent $T = 10^7$ s search, for four fiducial types of pulsar. We list the requirements both for all-sky searches and for directed searches (i.e., searches where the source position is known in advance). Also listed are the threshold values h_{th} of the characteristic strain h_c required to have 99% confidence of detection, assuming unlimited computer power. These threshold values are given by $h_{\text{th}}/h_{3/\text{yr}} = (1/1.90)\sqrt{\ln(50NN_p) - 1}$ where $N \equiv 2f_{\max}T$. Here $h_{3/\text{yr}}$ is the corresponding threshold, assuming the pulsar's position and period and are known in advance.

Search Parameters		N_p	$h_{\text{th}}/h_{3/\text{yr}}$	N_p	$h_{\text{th}}/h_{3/\text{yr}}$
f (Hz)	τ (Yrs)	(all-sky)	(all-sky)	(directed)	(directed)
< 200	$> 10^3$	1.1×10^{10}	3.7	3.7×10^6	3.3
$< 10^3$	$> 10^3$	1.3×10^{16}	4.2	1.2×10^8	3.5
< 200	> 40	1.7×10^{18}	4.3	8.5×10^{12}	3.9
$< 10^3$	> 40	8×10^{21}	4.6	1.4×10^{15}	4.1

computing power in Fig. 2.2.

2.1.5 Organization of this paper

In section 2.2 we outline the physics of pulsars which is relevant to the detection of continuous gravitational waves. The discussion is phenomenological and based almost entirely on pulsar data collected by radio astronomers. We focus attention on effects which may lead to significant frequency evolution over periods of several weeks of observation.

Then, in section 2.3, we introduce a parameterized model of the expected gravitational waveform, including modulating effects due to detector motion.

From this, we go on in section 2.4 to describe the basic technique used to search for signals, by constructing a demodulated time series. Livas [26], Jones [27] and Niebauer [28] have implemented variants of this basic search strategy over limited regions of parameter space (in particular they have not considered pulsar spin-down, and have restricted attention to small areas of the sky).

For the more computationally-intensive search over all sky positions and spindown parameters, it is important to be able to calculate the smallest number of independent parameter values which must be sampled in order to cover the entire space of signals. In section 2.5 we develop the metric formalism for calculating the number of independent points in parameter space.

In sections 2.6 and 2.7 we apply this formalism to determine the computational requirements of an all-sky search for unknown pulsars and a directed search, respectively.

Finally in section 2.8, we list some possible alternatives to a straightforward coherent search of the interferometer data. Detailed studies of the pros and cons of each are currently under investigation.




Figure 2.2: Relative amplitude sensitivities $h_{3/yr}/h_{th}$ achievable with given computational resources, for various coherent search strategies: (a) directed search for old ($\tau \geq 1000\text{Yrs}$), slow ($f \leq 200\text{Hz}$) pulsars, (b) all-sky search for old, slow pulsars, (c) directed search for young ($\tau \geq 40\text{Yrs}$) fast ($f \leq 1000\text{Hz}$) pulsars, and, (d) all-sky search for these same sources. For a given computational power, we have determined the optimum observation time as described in sections 2.6B and 2.7. Thus h_{th} is the expected sensitivity of the detector for the optimal observation time, and with 99% confidence, assuming only that the frequency bandwidth of the source is constrained in advance; see Eq. (2.56).

2.2 Pulsar phenomenology

Currently, the only expected sources of continuous, periodic gravitational waves in the LIGO band are pulsars. In this section, therefore, we review those properties of pulsars which may be important in the detection process. In general, the search technique we present later is capable of detecting any *nearly* monochromatic gravitational wave with sufficient amplitude. However, it is useful to have a concrete physical system in mind when considering the expected gravitational waveform.

That pulsars are rapidly rotating neutron stars is now well established [29]. Their high densities and strong gravitational fields allow them to withstand rotation rates of hundreds of times per second. Moreover, pulsar emission mechanisms require large magnetic fields, frozen into (co-rotating with) the neutron star. Indeed these large field strengths may produce non-axisymmetric deformations of the pulsar. However, the most remarkable feature of pulsars is the very precise periodicity of observed pulses.

There are more than 700 known pulsars, all at galactic distances, concentrated in the galactic plane. Based on the sensitivity limits of radio observations the total number of active pulsars in our galaxy is estimated to be more than 10^5 [30, 31].

2.2.1 Spindown

Pulsars lose rotational energy by electromagnetic braking, the emission of particles and, of course, emission of gravitational waves [32, 33]. Thus, the rotational frequency is not completely stable, but varies over a timescale τ which is of order the age of the pulsar. Typically, younger pulsars (with periods of tens of milliseconds) have the largest spindown rates. Figure 2.3 shows the distribution of rotational frequencies and spindown age, $\tau = f/(df/dt)$.

Current observations suggest that spindown is primarily due to electromagnetic braking; however, for detection purposes it is necessary to construct a sufficiently general model of the frequency evolution to cover all possibilities. For observing times t_{obs} much less than τ , the frequency drift is small and the rotational frequency³ can be modeled as a power series of the form

$$f(t) = (f_0/2) \left(1 + \sum_k f_k t^k \right). \quad (2.5)$$

If τ_{min} is the shortest timescale over which the frequency is expected to change by a factor of order unity, the coefficients satisfy

$$|f_k| \lesssim \tau_{\text{min}}^{-k}. \quad (2.6)$$

Clearly, for an observation time $t_{\text{obs}} \ll \tau_{\text{min}}$, the first few terms in this series will dominate.

³We choose to parameterize the frequency by what will be the gravitational wave frequency, f_0 , thus introducing the extra factor of 2 into this expression.




Figure 2.3: Gravitational wave frequency versus spindown age, $\tau = f/(df/dt)$, measured in years, for 540 pulsars which have measured period derivative. The figure clearly shows a large concentration of pulsars in the mid-left of diagram. Most of these are isolated pulsars. The standard evolutionary scenario suggests that pulsars move from higher frequencies and shorter spindowns left and up towards this main bunch. In contrast, many of the millisecond pulsars lying in the upper right of the figure are in binary systems, and it is widely believed that these are pulsars which have been spun up by mass accretion from the companion star.

Observations suggest that pulsars are born in supernova explosions with very short periods (perhaps several milliseconds), and subsequently spin down on timescales comparable to their age. Supernovae are observed in galaxies similar to our own at the rate of two or three per century, so we might expect $\tau_{\min} \sim 40$ years for pulsars in our galaxy. It is at this point that the distinction between various classes of pulsars becomes important. The known millisecond pulsars are old neutron stars which have been spun up to periods of only a few milliseconds, possibly by episodes of mass transfer from a companion star. As seen from Fig. 2.3, timing measurements of millisecond pulsars yield very long spindown timescales, $\tau_{\min} \gtrsim 10^7$ years.

2.2.2 Proper motions

Pulsars are generally high velocity objects [32], as can be inferred by the distance they move in their lifetimes. Proper motions cause Doppler shifts in the observed pulsar frequency. If the motion is uniform (constant velocity), it simply induces a constant frequency shift — an effect which is undetectable. However, acceleration and higher order derivatives of the source's motion will modulate the observed frequency.

Studies of millisecond pulsars in globular clusters have shown that acceleration in the cluster field can produce frequency drifts which are comparable in magnitude to the spindown effects [34, 35]. Once again, we expect these effects to be well modeled by a power series in t/τ_{cross} , where τ_{cross} is the time it takes the pulsar to cross the cluster. We expect that $\tau_{\min} \leq \tau_{\text{cross}}$ for these objects (since if not, the pulsar will already have escaped the cluster). Thus the frequency model adopted above should be sufficiently general to encompass the observational effects of proper motions of the sources.

A large proportion of millisecond pulsars are also in binary systems. Unfortunately, such pulsars participate in proper motions which vary over very short timescales (their orbital periods). The time-dependent Doppler effect due to these motions is *not* modeled accurately by a simple power series as in Eq. (2.5). They would require a more elaborate model involving as many as five unknown orbital parameters. Including these effects in a coherent, all-sky search strategy would be prohibitive (see section 2.6). In a search for gravitational waves from a known binary pulsar, however, it would be important to deal with this effect.

Proper motions can also affect a search if the star moves across more than one resolution element on the sky during an observation. For the lengths of observation periods envisioned here, this is unlikely to be a problem. In an observation lasting a year, however, a pulsar with a spatial velocity of $1 \times 10^3 \text{ km s}^{-1}$ at a distance of 300 pc will move by about half an arc-second, which is comparable to the resolution limit for our observations if the pulsar frequency is 1kHz.

2.2.3 Glitches

In addition to gradual frequency drifts due to spindown, some young pulsars exhibit occasional, abrupt increases in frequency. The physical mechanism behind these frequency *glitches* is not well understood, although the number of observations of glitch events is growing [30]. Given the stochastic nature of glitching, and the expectation that several months will elapse between major events, we will ignore glitching in this paper.

2.3 Gravitational waves from pulsars

In order to gain insight into the detection problem it is also important to understand the expected gravitational wave signal. Several mechanisms have been discussed in the literature which may produce non-axisymmetric deformations of a pulsar, and hence lead to gravitational wave generation [19, 20, 21, 23, 36, 37].

In general, a pulsar can radiate strongly at frequencies other than twice the rotation frequency. For example, a pulsar deformed by internal magnetic stresses, which are not aligned with a principal axis, can radiate at the rotation frequency and twice that frequency [38]. If the star precesses, it will radiate at three frequencies: the rotation frequency, and the rotation frequency plus and minus the precession frequency [23]. The important point, however, is that the signal at the detector is generally narrow band, exhibiting only slow frequency drift on observational timescales.

Therefore, in this section we outline the main features of the expected waveform and the corresponding strain measured at a detector for the case of crustal deformation; other scenarios give similar results except for the presence of more than one spectral component.

2.3.1 Waveform

Adopting a simple model of a distorted pulsar as a tri-axial ellipsoid, rotating about a principal axis with a frequency given by Eq. (2.5), one may compute the expected gravitational wave signal using the quadrupole formula. The two polarizations are

$$h_+ = h_0(1 + \cos^2 i) \cos\{2\pi f_0 [t + \sum f_k \frac{t^{k+1}}{k+1}]\}, \quad (2.7)$$

$$h_\times = 2h_0 \cos i \sin\{2\pi f_0 [t + \sum f_k \frac{t^{k+1}}{k+1}]\}, \quad (2.8)$$

where i is the angle between the rotation axis and the line of sight to the source. The dimensionless amplitude is

$$h_0 = \frac{2\pi^2 G}{c^4} \frac{I_{zz} f_0^2}{r} \epsilon, \quad (2.9)$$

where

$$\epsilon = \frac{I_{xx} - I_{yy}}{I_{zz}} \quad (2.10)$$

is the *gravitational ellipticity* of the pulsar. The distance to the source is r , and I_{jk} is its moment of inertia tensor.

The strength of potential sources is best discussed in terms of the characteristic amplitude h_c , defined in Eq. (50) of [1], and simply related to h_0 by

$$h_c = \sqrt{\frac{32}{15}} h_0. \quad (2.11)$$

For a typical $1.4M_\odot$ neutron star, having a radius of 10km and at a distance of 10kpc, the dimensionless amplitude is

$$h_c = 7.7 \times 10^{-25} \frac{\epsilon}{10^{-5}} \frac{I_{zz}}{10^{45} \text{ g cm}^2} \frac{10\text{kpc}}{r} \left(\frac{f_0}{1\text{kHz}} \right)^2. \quad (2.12)$$

The magnitude of the gravitational ellipticity, ϵ , represents the central uncertainty in any estimate of gravitational waves from pulsars. Models of neutron star structure generally include a crystalline outer layer, the *crust*, of the star surrounding a superfluid core. Since the moment of inertia of the crust represents only about 10% of the total moment of inertia and the superfluid core cannot support non-axisymmetric deformations, the tightest theoretical constraint, $\epsilon < 10^{-5}$, is set by the maximum strain that the neutron star crust may support [39, 1]. It has also been suggested that stresses induced by large magnetic fields might result in significant gravitational ellipticity. Recently, Bonazzola and Gourgoulhon [19] have considered this possibility, finding discouraging results; their calculations indicate $10^{-13} \lesssim \epsilon \lesssim 10^{-9}$ depending on the precise model they consider.

In any case, an upper bound on the gravitational ellipticity is $\epsilon \sim 10^{-5}$, although typical values may be significantly smaller.

2.3.2 Signal at the detector

Observing the gravitational waves using an earth-based interferometer introduces two further difficulties into the detection process: Doppler modulation of the observed gravitational wave frequency, and amplitude modulation due to the changing orientation of the detector.

For the purpose of detection, the Doppler modulation of the observed gravitational wave frequency, due to motion of the detector with respect to the solar system barycenter, is a large effect. Assuming the intrinsic frequency model (2.5) for the pulsar rotation, the gravitational wave fre-

quency measured at the detector is

$$f_{\text{gw}}(t) = f_0 \left(1 + \frac{\vec{v}}{c} \cdot \hat{n} \right) \left(1 + \sum_k f_k \left[t + \frac{\vec{x}}{c} \cdot \hat{n} \right]^k \right), \quad (2.13)$$

where $\vec{x}(t)$ is the detector position, $\vec{v}(t)$ is the detector velocity, and \hat{n} is the unit vector pointing to the pulsar, in some inertial frame. We generally choose this frame to be initially comoving with the Earth at $t = 0$. The frequency measured in this frame is identical to that measured at the solar system barycenter except for an unimportant constant shift in f_0 .

To understand the amplitude modulation we must introduce the Euler angles, $\{\Theta, \Phi, \Psi\}$, which specify the orientation of the gravitational wave frame with respect to the detector frame. The dimensionless strain at the detector is

$$h = F_+(\Theta, \Phi, \Psi)h_+ + F_\times(\Theta, \Phi, \Psi)h_\times \quad (2.14)$$

where F_+ and F_\times are the detector beam patterns given in Thorne [1]. In searching for continuous gravitational waves from a particular direction, the Euler angles become periodic function of sidereal time, thus resulting in an amplitude and phase modulation of the observed signal [1, 19, 26]. For observation times longer than one sidereal day, the amplitude modulation effectively averages the reception over all values of right ascension, and over a range of declination that depends on the precise position of the pulsar. In particular, the effect of this process is to allow detection of continuous waves from any direction, but at the cost of reducing the measured strain (see Fig. 2.4).

2.3.3 Parameter space

To facilitate later discussion it is useful to parameterize the gravitational waveform by a vector $\lambda = (\lambda^0, \vec{\lambda})$ such that

$$(\lambda^0, \lambda^1, \dots, \lambda^{s+2}) = (f_0, n_x, n_y, f_1, \dots, f_s). \quad (2.15)$$

Here s is the maximum number of spindown parameters included in the frequency model determined by Eq. (2.5). These vectors span an $s + 3$ dimensional space on which λ^α can be thought of as coordinates. (Note that $n_z^2 = 1 - n_x^2 - n_y^2$ is not an independent parameter.) In particular we denote the observed phase of the gravitational waveform by

$$\phi(t; \lambda) = 2\pi \int^t dt' f_{\text{gw}}(t'), \quad (2.16)$$

where $f_{\text{gw}}(t')$ is given by Eq. (2.13).

Initial interferometers in LIGO should have reasonable sensitivity to gravitational waves with

frequencies

$$f \geq 40\text{Hz} , \quad (2.17)$$

while advanced interferometers are expected to have improved sensitivity down to

$$f \geq 10\text{Hz} . \quad (2.18)$$

Moreover, theoretical constraints suggest that pulsars with spin periods significantly smaller than one millisecond are unlikely. This helps to constrain the highest frequency that one may wish to consider in an all sky search to be about 2kHz. According to the discussion in section 2.2, the spindown parameters satisfy

$$-\tau_{\min}^{-k} \leq f_k \leq \tau_{\min}^{-k} , \quad (2.19)$$

where τ_{\min} is the minimum spindown age of a pulsar to be searched for. Finally, n_x and n_y are restricted by the relation

$$n_x^2 + n_y^2 \leq 1 . \quad (2.20)$$

2.4 Data analysis technique

Radio astronomers are familiar with searching for nearly periodic sources in the output of their detectors [35, 40]. The technique employed by them is directly applicable to the problem at hand [26, 27].

In the detector frame the gravitational wave signal can be written as

$$h(t; \lambda) = \text{Re} \left[\mathcal{A} e^{-i\phi(t; \lambda)} \right] \quad (2.21)$$

where $\mathcal{A} = (h_{0+} + ih_{0\times})$, $h_{0+} = F_+(1 + \cos^2 i)h_0$ and $h_{0\times} = 2F_\times(\cos i)h_0$. The orbital phase $\phi(t; \lambda)$ is given by Eqs. (2.16) and (2.13). Introducing a canonical time

$$t_b[t; \vec{\lambda}] = \frac{\phi(t; \lambda)}{2\pi f_0} , \quad (2.22)$$

the above signal becomes monochromatic as a function of t_b . (The presence of the amplitude modulation complicates the following analysis without changing the conclusions significantly; therefore, we treat \mathcal{A} as constant in this and the next section⁴.) Figure 2.4 shows the normalized power spectrum computed from the signal as a function of t in Eq. (2.21) (with $f_k \equiv 0$), compared with the spectrum from the signal as a function of t_b . It is clear that the maximum power per frequency bin

⁴Amplitude modulation can be viewed as the convolution of the exactly periodic signal with some complicated window function. Thus, in reality, the power spectrum of a stretched signal will not be a monochromatic spike at a single frequency, but will be split into several discrete, narrow spikes spread over a bandwidth $\delta f \simeq 10^{-4}\text{Hz}$. After a preliminary detection, the amplitude modulation spikes would provide a discriminant against false signals [26].

is significantly reduced when frequency modulation is not accounted for.

Radio astronomers refer to this technique of introducing a canonical time coordinate as *stretching* the data. Since interferometer output will be sampled at approximately 16kHz, in a practical search for pulsars up to 2kHz gravitational wave frequency, the stretching can probably be achieved by resampling the data stream appropriately. This method, which is called *stroboscopic sampling* by Schutz [25], has the benefit of keeping the computational overhead introduced by the stretching process to a minimum. We will return to this issue in a later publication.

Now, a search of the detector output, $o(t)$, for gravitational waves from a known source is straightforward. One assumes specific parameter values $\vec{\xi}$ in the waveform (2.21), computes the demodulated time function $t_b[t; \vec{\xi}]$ using Eq. (2.22) and stretches the detector output accordingly, thus

$$o_b(t_b[t; \vec{\xi}]) = o(t) . \quad (2.23)$$

If the assumed parameters $\vec{\xi}$ are not too much different from the actual parameters $\vec{\lambda}$ of the signal, the stretched data will consist of a nearly monochromatic signal. One then takes the Fourier transform with respect to t_b ,

$$\bar{o}(f; \vec{\xi}) = \frac{1}{\sqrt{T_b^{\text{obs}}}} \int_0^{T_b^{\text{obs}}} e^{2\pi i f t_b} o_b(t_b) dt_b . \quad (2.24)$$

Here T_b^{obs} is length of the observation measured using t_b . The power spectrum is then searched for excess power. (The threshold is set by demanding some overall statistical significance for a detection; see section 2.6.) Notice that the gravitational wave frequency, $\lambda^0 = f_0$, is treated somewhat differently than the other parameters; the Fourier transform searches over all possible values in a single pass. Given a sampled data set containing N points, the entire process, from original data through to the power spectrum, requires of order $3N \log_2 N$ floating point operations (to first approximation).

If all the parameters are not known accurately in advance, it will be necessary to search over some of the remaining parameters $\vec{\lambda}$; a separate demodulation and FFT must be performed for each independent point in parameter space that one wishes to search. There are many possible refinements on this strategy which could reduce the computational cost of a search by circumventing certain stages of the procedure described here. We mention some of them in section 2.8, however, we focus attention on this baseline strategy in this paper.

One more issue that arises in the discussion of stretching is how it effects the noise in the detector. Throughout this paper we assume that the noise in the detector is a stationary, Gaussian process; however, when we stretch the output data stream the noise is no longer strictly stationary unless it is perfectly white. Real detectors will have colored noise, with correlations between points sampled at different times. Stretching the data modifies these correlations in a time dependent manner. In our




Figure 2.4: Power spectra for two simulated signals, each with gravitational wave frequency 5Hz, computed using approximately 10 days worth of data; they are normalized with respect to the maximum power achieved if the source were directly above an interferometer which remained stationary during the entire observation. The signal was assumed to come from declination 0° and right ascension 90° ; in fact, the amplitude modulation is only sensitive to changes in declination. The detector latitude was chosen to coincide with LIGO detector in Hanford Washington. The solid line corresponds to a Doppler and amplitude modulated gravitational wave signal. The dashed line is the same signal but with the Doppler modulation removed by stretching. The (unreasonably) low frequency was chosen for illustrative purposes, so that both curves could appear on the same scale. For realistic gravitational wave frequencies ($\sim 500\text{Hz}$) the Doppler modulated signal would be further reduced by roughly two orders of magnitude.

case this is a very small effect, having a characteristic timescale of several hours, and besides this the noise in real detectors may be intrinsically non-stationary on similar timescales due to instrumental effects. Correcting pulsar searches for such non-stationarity is an important problem, but one that we do not address here. We simply assume that $S_n(f)$, the power spectral density of the noise, can be estimated on short timescales and used in the conventional way for signal to noise estimates. Moreover, the effects of stretching on noise are only a consideration when the noise is not white; since stretching affects the power spectrum only within bands $\sim 10^{-1}$ Hz wide, the detector spectrum can usually be taken as white, unless we are near a strong feature in the noise spectrum. The precise nature of these effects is being explored by Tinto [41].

2.5 Parameter space metric

In general, neither the position of the pulsar nor its intrinsic spindown may be known in advance of detection. Therefore, the above process, or some variant on it, must be repeated for many different vectors $\vec{\xi}$ until the entire parameter space has been explored. How finely must one sample these parameters in order to minimize the risk of missing a signal? A similar question arises in the context of searching for signals from coalescing compact binaries using matched filtering; Owen [14] has introduced a general framework to provide an answer in that case. We adapt his method to the problem at hand by defining a distance function on our parameter space; the square of distance between two points in parameter space is proportional to the fractional loss in signal power due to imprecise matching of parameters. The number of discrete points which must be sampled can then be determined from the proper volume of the parameter space with respect to this metric.

2.5.1 Mismatch

The one-sided power spectral density (PSD) of the detector output, stretched with parameters $\vec{\xi}$, is

$$P_o(f) = 2 |\bar{o}(f; \vec{\xi})|^2 . \quad (2.25)$$

Now, suppose a detector output consists of a signal with parameters λ , and stationary, Gaussian noise $n(t)$ such that

$$o(t) = h(t; \lambda) + n(t) . \quad (2.26)$$

Thus, the expected PSD of the detector output, once again stretched with parameters $\vec{\xi}$, is

$$E[P_o(f)] = 2 |\bar{h}(f; \lambda, \Delta\vec{\lambda})|^2 + S_n(f) , \quad (2.27)$$

where $\Delta\vec{\lambda} = \vec{\xi} - \vec{\lambda}$, and $S_n(f)$ is the one-sided power spectral density of the detector noise. (As discussed at the end of the previous section, we ignore the small effects of stretching on the noise.) The notation $\tilde{h}(f; \lambda, \Delta\vec{\lambda})$ indicates the Fourier transform of a signal, with parameters λ , with respect to a time coordinate $t_b[t; \vec{\lambda} + \Delta\vec{\lambda}]$. We define the *mismatch* $m(\lambda, \Delta\lambda)$ to be the fractional reduction in signal power caused by stretching the data with the wrong parameters, *and* by sampling the spectrum at the wrong frequency; specifically,

$$m(\lambda, \Delta\lambda) = 1 - \frac{|\tilde{h}(f; \lambda, \Delta\vec{\lambda})|^2}{|\tilde{h}(f_0; \lambda, 0)|^2}. \quad (2.28)$$

Remember that $\lambda = (\lambda^0 = f_0, \vec{\lambda})$.

In the present circumstance, it is sufficient to consider a complex signal

$$h(t; \lambda) = \mathcal{A} e^{-2\pi i f_0 t_b[t; \vec{\lambda}]}, \quad (2.29)$$

where the amplitude \mathcal{A} is constant. The function $t_b[t; \vec{\lambda}]$, computed using Eqs. (2.22), (2.16) and (2.13), is explicitly written as

$$t_b[t; \vec{\lambda}] = t + \frac{\vec{x}}{c} \cdot \hat{n} + \sum_k \frac{f_k}{k+1} \left(t + \frac{\vec{x}}{c} \cdot \hat{n} \right)^{k+1}. \quad (2.30)$$

Now, the Fourier transform $\tilde{h}(f; \lambda, \Delta\vec{\lambda})$ is

$$\tilde{h}(f; \lambda, \Delta\vec{\lambda}) = \frac{\mathcal{A}}{\sqrt{T_b^{\text{obs}}}} \int_0^{T_b^{\text{obs}}} dt_b e^{i\Phi[t; \lambda, \Delta\vec{\lambda}]}, \quad (2.31)$$

where

$$\frac{\Phi[t; \lambda, \Delta\vec{\lambda}]}{2\pi} = \Delta\lambda^0 \hat{t}_b + f_0 (t_b[t; \vec{\lambda} + \Delta\vec{\lambda}] - t_b[t; \vec{\lambda}]) \quad (2.32)$$

and $\Delta\lambda^0 = f - f_0$. Here, t should be interpreted as a function of \hat{t}_b defined implicitly by $\hat{t}_b = t_b[t; \vec{\lambda} + \Delta\vec{\lambda}]$. Using Eqs. (2.30)-(2.32) it is easy to show that $m(\lambda, \Delta\lambda)$ has a local minimum of zero when $\Delta\lambda \equiv 0$;

$$m(\lambda, \Delta\lambda)|_{\Delta\lambda=0} = 0, \quad (2.33)$$

$$\partial_{\Delta\lambda^\alpha} m(\lambda, \Delta\lambda)|_{\Delta\lambda=0} = 0. \quad (2.34)$$

Thus, an expansion of the mismatch in powers of $\Delta\lambda$ is

$$m(\lambda, \Delta\lambda) = \sum_{\alpha, \beta} g_{\alpha\beta}(\lambda) \Delta\lambda^\alpha \Delta\lambda^\beta + \mathcal{O}(\Delta\lambda^3), \quad (2.35)$$

where

$$g_{\alpha\beta} = \frac{1}{2} \partial_{\Delta\lambda^\alpha} \partial_{\Delta\lambda^\beta} m(\lambda, \Delta\lambda)|_{\Delta\lambda=0} . \quad (2.36)$$

In this way the mismatch defines a local distance function on the signal parameter space, and, for small separations $\Delta\lambda$, $g_{\alpha\beta}$ is the metric of that distance function. Note that the metric formulation (2.35) will generally *overestimate* the mismatch for large separations, as demonstrated in Figure 2.5.

Calculations using this formalism are considerably simplified by partially evaluating the right-hand side of Eq. (2.36). The form of the signal (2.29) allows us to write

$$g_{\alpha\beta}(\lambda) = \langle \partial_{\Delta\lambda^\alpha} \Phi \partial_{\Delta\lambda^\beta} \Phi \rangle - \langle \partial_{\Delta\lambda^\alpha} \Phi \rangle \langle \partial_{\Delta\lambda^\beta} \Phi \rangle , \quad (2.37)$$

where Φ is given by Eq. (2.32), and where we use the notation

$$\langle \dots \rangle = \frac{1}{T_b^{\text{obs}}} \int_0^{T_b^{\text{obs}}} (\dots) dt_b \Big|_{\Delta\lambda=0} . \quad (2.38)$$

2.5.2 Metric and number of patches

Up until now, we have treated the frequency of the signal as one of the parameters, λ_0 , which must be matched. In our search technique, stretching and Fourier transforming the data yields an entire power spectrum, automatically sampling all possible frequencies. We would really like to know the number of times that this combination of procedures must be performed in a search. This requires knowledge of the mismatch $m(\lambda, \Delta\lambda)$ as a function of $\Delta\vec{\lambda}$, having already maximized the power (i.e., minimized m) over λ^0 . The result is the mismatch projected onto the $(s+2)$ -parameter subspace:

$$\mu = \min_{\lambda_0} m(\lambda, \Delta\lambda) = \sum_{ij} \gamma_{ij} \Delta\lambda^i \Delta\lambda^j \quad (2.39)$$

where

$$\gamma_{ij} = g_{ij} - \frac{g_{0i}g_{0j}}{g_{00}} , \quad (2.40)$$

and $i = 1, \dots, s+2$. We will generally refer to μ as the *projected mismatch*.

Technically, γ_{ij} should be computed from $g_{\alpha\beta}$ evaluated at the specific value of λ_0 at which the minimum projected mismatch occurred. However, since this number is unknown in advance of detection, we evaluate γ_{ij} for the largest frequency in the search space. In this way we never underestimate the projected mismatch.

In a search, the parameter space will be sampled at a lattice of points, chosen so that no location in the space has μ (given by Eq. (2.39) greater than some μ_{max} away from one of the points. This is equivalent to tiling the parameter space with *patches* of maximum extent $\mu_{\text{max}}^{1/2}$. The number of




Figure 2.5: Fractional reduction in measured signal power caused by demodulating with mismatched parameters (in this case, an error in the assumed declination of the source). The solid curve is the true power ratio, the dotted is that given by the quadratic approximation of the metric. Note that the widths of the curves agree well down to 70% power reduction ($m \sim 0.7$), beyond which the metric approximation significantly underestimates the range of parameters permitted for a specified power loss. The curves are computed for a sky position of 0° right ascension, 45° declination, and no spindown.

points we must sample at is therefore

$$N_p = \frac{\int_{\mathcal{P}} \sqrt{\det \|\gamma_{ij}\|} d^{s+2} \vec{\lambda}}{V_{\text{patch}}}, \quad (2.41)$$

where V_{patch} is the proper volume of a single patch, and $s + 2$ is the reduced dimensionality of the parameter space \mathcal{P} (excluding λ_0).

Optimally, one should use some form of spherical closest packing to cover the space with the fewest patches. Our solution uses hexagonal packing in two of the dimensions and cubic packing in all the others; in this way the volume of a single patch is

$$V_{\text{patch}} = \frac{3\sqrt{3}}{4} \left(\frac{4\mu_{\text{max}}}{s+2} \right)^{(s+2)/2} \quad (2.42)$$

Finally, we note that Eq. (2.41) may overestimate (conceivably, greatly overestimate) the number of points one must sample if the parameter space submanifold *folds in upon itself*, so that points that seem widely separated on the submanifold are actually close together in the embedding space of possible signals. While we have no reason to think this is occurring, we also have not seriously tried to investigate this possibility; it is a difficult, non-local question. Until this is resolved, it is perhaps safest to regard N_p given by Eq. (2.41) as an upper limit on the required number of sample points, though we suspect it is close to the actual number.

2.6 Depth of an all sky search

We are finally in a position to estimate the depth of a search for periodic sources using LIGO. The detector participates in two principal motions which cause significant Doppler modulations of the observed signal: daily rotation, and revolution of the Earth about the Sun. The latter is actually a complex superposition of an elliptical Keplerian orbit with a smaller orbit about the earth-moon barycenter, and is further perturbed by interactions with other planets. For now, however, we use a simplified model which treats both rotation and revolution as circular motions about separate axes inclined at an angle $\epsilon = 23^\circ 27'$ to each other. Although a simplification, this does remove any spurious symmetries from the model; thus, an actual search using the precise ephemeris of the earth in its demodulations should give comparable results. In this model, then, we write the velocity of the detector in a frame which is inertial to the solar system barycenter but initially comoving with the earth:

$$\begin{aligned} \vec{v} = & -(\Omega R_d \sin \Omega t - \Omega_A R_A \sin \Omega_A t) \vec{x} \\ & + (\Omega R_d \cos \Omega t - \Omega_A R_A \cos \epsilon [\cos \Omega_A t - 1]) \vec{y} \end{aligned} \quad (2.43)$$

$$-\Omega_A R_A \sin \epsilon [\cos \Omega_A t - 1] \vec{z}$$

where $R_d = 6.371 \times 10^8 (\cos l) \text{cm}$, l is the latitude of the detector, and $R_A = 1.496 \times 10^{13} \text{cm}$ is the distance from the earth to the sun. The angular velocities are $\Omega = 2\pi/(86400\text{s})$ and $\Omega_A = 2\pi/(3.155674 \times 10^7 \text{s})$. Our coordinate system measures \vec{x} towards the vernal equinox and \vec{z} towards the north celestial pole, and we arbitrarily choose to measure time starting at noon on the vernal equinox.

The number of spindown parameters f_k which must be included to account for all intrinsic frequency drift depends to a large extent on the type of pulsar one wishes to search for. We determined this number on a case by case basis, including all parameters which lead to a significant increase in the number of parameter space patches. Equivalently, the following geometric picture suggests a simple criterion for deciding when there is one spindown parameter too many included in the signal parameterization. Let λ^L be the last, ‘questionable’ spindown parameter f_s (so $L = s+2$). With respect to the natural metric γ_{ij} on parameter space, the unit-normal to surfaces of constant λ^L is just $\gamma^{iL}/(\gamma^{LL})^{1/2}$, where γ^{ij} is the inverse of γ_{ij} . The spindown parameter λ^L is unnecessary if the proper thickness of the parameter space in this normal direction nowhere exceeds half the proper grid spacing; that is, if

$$\max_p \left\{ 2\tau_{\min}^{-L+2}/(\gamma^{LL})^{1/2} \right\} < \sqrt{\mu_{\max}/L}. \quad (2.44)$$

In practice, one has included more spindown parameters than necessary if and only if $\min_p \gamma^{LL} > 4L\tau_{\min}^{-2L+2}/\mu_{\max}$.

2.6.1 Patch number versus observation time

It is extremely difficult to obtain a closed form expression for the metric, let alone its determinant. Therefore, we present results for two concrete scenarios which suggest themselves based on the discussion in section 2.2: (i) hypothetical sources with $f_0 \leq 1000\text{Hz}$, and spindown ages greater than $\tau = 40\text{Yrs}$; incidentally, this also includes the majority of known, millisecond pulsars; and (ii) slower sources ($f_0 \leq 200\text{Hz}$) having spindown ages in excess of $\tau = 1000\text{Yrs}$. The number of parameter space points which must be searched is plotted as a function of total observation time in Fig. 2.6. The numbers are normalized by a maximum projected mismatch $\mu_{\max} = 0.3$.

In considering an optimal choice of observation time, it is useful to construct an empirical fit to $N_p(t_{\text{obs}}, \mu_{\max})$. Notice first that all the parameters $\Delta\vec{\lambda}$ in Φ , given by Eq. (2.32), appear multiplied by the gravitational wave frequency f_0 ; thus, $N_p \propto (f_{\max})^{s+2}$ where f_{\max} is the maximum gravitational wave frequency to be searched for. Furthermore, provided the determinant of the metric is only weakly dependent on the values of the f_k , one may also extract a factor of $\tau^{-s(s+2)/2}$; our

investigations suggest the validity of this approach. In this way we arrive at the expression

$$N_p \simeq \max_{s \in \{0,1,\dots\}} [\mathcal{N}_s F_s(t_{\text{obs}})] , \quad (2.45)$$

where

$$\mathcal{N}_s = \left(\frac{f_{\text{max}}}{1\text{kHz}} \right)^{s+2} \left(\frac{40\text{Yrs}}{\tau} \right)^{s(s+1)/2} \left(\frac{0.3}{\mu_{\text{max}}} \right)^{(s+2)/2} \quad (2.46)$$

$$F_0(t_{\text{obs}}) = 6.9 \times 10^3 T^2 + 3.0 T^5 \quad (2.47)$$

$$F_1(t_{\text{obs}}) = \frac{1.9 \times 10^8 T^8 + 5.0 \times 10^4 T^{11}}{4.7 + T^6} \quad (2.48)$$

$$F_2(t_{\text{obs}}) = \frac{2.2 \times 10^7 T^{14}}{56.0 + T^9} , \quad (2.49)$$

and $T = t_{\text{obs}}/(8.64 \times 10^4\text{s})$ is the observation time measured in days. These formulae are normalized using only the data corresponding to Fig. 2.6(a), and subsequently compared with computed values for several frequencies and spindown ages τ . The analytic fit is in good agreement with the computed results for a variety of parameters; however, the fits generally break down for observation times less than one day. We stress that more spindown parameters may become important for observation times longer than 30 days.

Schutz [25] has previously estimated the number of points which must be searched in the absence of spindown corrections; he argued that this number scaled as T^4 for observation times longer than about a day. The difference between his previous estimate and the expression in Eq. (2.47), which shows that the number of points increases as T^5 , derives from an asymmetry between declination and right ascension which was not accounted for in his argument.

The benefit of the metric formulation is that it accounts for the significant correlations which exist between the intrinsic spindown and the earth-motion-induced Doppler modulations by using points which lie on the principal axes of the ellipsoids described by Eq. (2.39). Replacing the invariant volume integral in Eq. (2.41) by

$$\int_{\mathcal{P}} \sqrt{\prod_i \gamma_{ii}} d^{s+2} \vec{\lambda} \quad (2.50)$$

gives the number of points required for a search if, instead, one chooses them to lie on the $\{n_x, n_y, f_1, f_2, \dots\}$ coordinate grid. Figure 2.7 shows the total number of points computed using this method compared to the results obtained using the invariant volume integral. For sufficiently long integration times, the difference can be several orders of magnitude.

2.6.2 Computational requirements

The number of real samples of the interferometer output for an observation lasting t_{obs} seconds, and sampled at a frequency $2f_{\text{max}}$, where f_{max} is the maximum gravitational wave frequency being

searched for, is

$$N = 2f_{\max}t_{\text{obs}}. \quad (2.51)$$

For each $\vec{\lambda}$ that is used to stretch the detector output, a search then requires an FFT, calculation of the power, and some thresholding test for excess power. Assuming that the stretching and thresholding require negligible computations compared to performing the FFT and computing the power, the total number of floating point operations for a search is

$$N_{\text{op}} = 6f_{\max}t_{\text{obs}}N_p[\log_2(2f_{\max}t_{\text{obs}}) + 1/2], \quad (2.52)$$

where N_p is given by (2.45)-(2.49). The additive 1/2 inside the square brackets accounts for the three floating point operations per frequency bin which are required to compute the power from the Fourier transform.

A guideline for a feasible, long-term, search strategy is that data reduction should proceed at a rate comparable to data acquisition. Thus, the total computing power required for data reduction, in floating point operations per second (flops), is

$$P = \frac{N_{\text{op}}}{t_{\text{obs}}} = 6f_{\max}N_p(t_{\text{obs}}, \mu_{\max})[\log_2(2f_{\max}t_{\text{obs}}) + 1/2]. \quad (2.53)$$

For a prescribed maximum projected mismatch μ_{\max} , and maximum available computing power P_{\max} , this expression determines the maximum allowed coherent integration time. Alternatively, given the computing power available for data reduction, P_{\max} , it provides an implicit relation between μ_{\max} and the integration time.

The idea now is to choose μ_{\max} and t_{obs} so that we maximize the sensitivity of the search. In order to do this we must first obtain a threshold, above which we consider excess power to indicate the presence of a signal.

As discussed in section 2.4, we assume that the noise in the detector is a stationary, Gaussian random process with zero mean and PSD $S_n(f)$. In the absence of a signal, the power $P_o(f) = 2|\tilde{n}(f)|^2$ is exponentially distributed with probability density function

$$\frac{e^{-P_o(f)/S_n(f)}}{S_n(f)}. \quad (2.54)$$

We assume that there is independent noise in each of $f_{\max}t_{\text{obs}}$ frequency bins for a given demodulated power spectrum. In general the noise spectra obtained from neighboring parameter space points will not be statistically independent; however, one may expect that the correlations will be small when the mismatch between the points approaches unity. Therefore, we approximate the number of statistically independent noise spectra in our search to be $N_p(t_{\text{obs}}, \mu_{\max} = 0.3)$. In order that a

detection have overall statistical significance α , we must set our detection threshold so there is less than $1 - \alpha$ probability of *any* noise event exceeding that threshold. For a detection to occur, the power in the demodulated detector output must satisfy

$$\frac{P_o(f)}{S_n(f)} > \frac{\rho_c}{S_n(f)} = \ln \left[\frac{f_{\max} t_{\text{obs}} N_p(t_{\text{obs}}, \mu_{\max} = 0.3)}{1 - \alpha} \right], \quad (2.55)$$

where $P_o(f)$ was defined in Eq. (2.25), and ρ_c is the threshold power.

In other words, if the power at a given frequency exceeds ρ_c , we can infer that a signal is present; the expected power in the signal is then $\rho_c - S_n$. Thus, the minimum characteristic amplitude we can expect to detect is

$$h_{\text{th}} = \sqrt{\frac{(\rho_c/S_n - 1)S_n(f)}{\langle F_+^2(\Theta, \Phi, \Psi) \rangle (1 - \langle \mu \rangle) t_{\text{obs}}}}, \quad (2.56)$$

where $\langle F_+^2(\Theta, \Phi, \Psi) \rangle$ is the square of the detector response averaged over all possible source positions and wave polarizations. $\langle \mu \rangle$ is the expected mismatch for a source whose signal parameters $\vec{\lambda}$ lie within a given patch, assuming that all parameter values in that patch are equally likely. We note that the characteristic detector sensitivities $h_{3/\text{yr}}$ in Fig. 2.1 are obtained from this expression by setting $t_{\text{obs}} = 10^7$ seconds, $\langle \mu \rangle = 0$, and $f_{\max} t_{\text{obs}} N_p = 1$ in the expression for ρ_c ; this agrees with Eq. (2.4).

The optimal search strategy is to choose those values of t_{obs} and μ_{\max} which, for some specified computational power P_{\max} and detection confidence α , maximize our *sensitivity* Θ which is defined by

$$\Theta(t_{\text{obs}}, \mu_{\max}) \equiv \frac{1}{h_{\text{th}}} \propto \sqrt{\frac{[1 - \frac{s+2}{s+4}\mu_{\max}]t_{\text{obs}}}{\rho_c/S_n - 1}}, \quad (2.57)$$

where ρ_c/S_n is given by Eq. (2.55). Assuming an overall statistical significance of $\alpha = 0.99$, we have computed the optimal observation time t_{obs} and optimal maximum mismatch μ_{\max} , as functions of computing power, for the two searches considered in the previous subsection. The results are shown in Fig. 2.8.

2.7 Computational requirements for a directed search

In sections 2.5 and 2.6 we examined the computational requirements of an all-sky pulsar search. In this section we examine the computational requirements for a directed pulsar search, by which we mean a search where the position is known but the pulsar frequency and spin-down parameters are unknown. Obvious targets in this category are SN1987A, nearby supernova remnants that do not contain known radio pulsars, and the center of our galaxy. Such searches will clearly be among the first performed once the new generation of gravitational wave detectors begin to come on line.

Our treatment of directed pulsar searches closely parallels that of the all-sky search, so we

can be brief. Since the source position (n_x, n_y) is known, we can simply remove the Earth's motion from the data. Below we imagine that the signal has already been transformed to the solar system barycenter. Then the unknown parameters describing the pulsar waveform are

$$(\lambda^0, \lambda^1, \dots, \lambda^s) = (f_0, f_1, \dots, f_s), \quad (2.58)$$

where the f_i are the same as defined in Eq. (2.5) and s is just the number of spindown parameters included in the frequency model. We again calculate the metrics g_{ij} and γ_{ij} using Eqs. (2.37) and (2.40) respectively, and then calculate N_p using (2.41) (except the integral is now over s -dimensional parameter space). Assuming hexagonal packing in two dimensions and cubic packing in the others, the size of each patch is $V_{\text{patch}} = (3\sqrt{3}/4)(4\mu_{\text{max}}/s)^{s/2}$. (Except for $s = 1$, where $V_{\text{patch}} = 2\mu_{\text{max}}^{1/2}$.) We arrive at the expression

$$N_p \simeq \max_{s \in \{1, 2, \dots\}} [N_s G_s(t)], \quad (2.59)$$

where

$$N_s = \left(\frac{f_{\text{max}}}{1\text{kHz}} \right)^s \left(\frac{40\text{Yrs}}{\tau} \right)^{s(s+1)/2} \left(\frac{0.3}{\mu_{\text{max}}} \right)^{s/2} \quad (2.60)$$

$$G_1(t_{\text{obs}}) = 1.5 \times 10^3 T^2 \quad (2.61)$$

$$G_2(t_{\text{obs}}) = 6.97 \times 10^1 T^5 \quad (2.62)$$

$$G_3(t_{\text{obs}}) = 2.89 \times 10^{-4} T^9, \quad (2.63)$$

where $T = t_{\text{obs}}/(8.64 \times 10^4\text{s})$ is the observation time measured in days. Comparing these results with Eqs. (2.45)–(2.49), we see that for our fiducial parameter values ($f_{\text{max}} = 1\text{kHz}$, $\tau_{\text{min}} = 40\text{Yrs}$, $\mu_{\text{max}} = 0.3$) and observation times T of order a week, N_p is $\sim 10^5$ times larger for an all-sky search than for a directed search. Another way of putting this is: after using one's freedom to adjust the frequency and spin-down parameters in optimizing the fit, only $\sim 10^5$ distinguishable patches on the sky remain. Equivalently, a single directed search can cover an area of $\sim 10^{-4}$ steradians. Thus ~ 1000 week-long, directed searches would be sufficient to cover the galactic center region.

We can calculate the optimal μ_{max} and t_{obs} as a function of computing power for a directed search in the same way as we did for the all-sky directed search. (Except the factor $\frac{s+2}{s+4}$ in Eq. 2.57 becomes $\frac{s}{s+2}$ for the directed-search case.) The results are shown in Fig. 2.9, for our two fiducial types of pulsar. We see that knowing the source position in advance increases t_{obs} by only a factor of ~ 10 , for 1 Tflops computing power. The resulting gains in sensitivity can be seen in Fig. 2.2.

2.8 Future directions

Searching for unknown sources of continuous gravitational waves using LIGO, or other interferometers, will be an immense computational task. In this paper we have presented our current understanding of the problem. By applying techniques from differential geometry we have estimated the number of independent points in the parameter space which must be considered in all-sky and directed searches for sources which spin down on timescales short enough to produce observable effects; these numbers were used to compute the maximum achievable sensitivity for a coherent search (see Fig. 2.2). Furthermore, the metric formulation can be used to optimally place the parameter space points which must be sampled in a search.

Our analysis takes no account of bottlenecks in the analysis process due to data input/output and inter-processor communication. These are important issues which may impose further constraints on the maximum observation time; however, it seems premature to address such problems until we know the hardware that will be used to conduct searches for continuous waves.

Unfortunately, Fig. 2.8 shows that it will be impossible to search, in one step, 10^7 seconds worth of data over all sky positions. However, it is also unnecessary. We foresee implementing a hierarchical search strategy, in which a long data stream is searched in two (or more) stages, trading off sensitivity in the first stage for reduced computational requirements. Having determined a number of potential signals in the first stage—presumably at a threshold level which allows many false alarms due to random noise—these candidate events would be followed up in the second stage, using longer integration times. The longer integration times would be possible because the search would only have to be performed over much smaller regions of the parameter space, in the neighborhoods of the candidate signal parameters. In this way, one can achieve a greater sensitivity than a coherent search using the same computational resources.

Clearly one can imagine many different implementations of this rough strategy, and we have not yet determined the optimal one. Nevertheless, we have considered the simple example where the data is searched in two stages. Candidate signals from an all-sky search of a short stretch of data [$T^{(1)}$ seconds long] are followed up using longer Fourier transforms to achieve greater sensitivity. One can estimate $T^{(1)}$ using Fig. 2.8 and an assumption that roughly half of the total computing budget is used on the first stage; this turns out to be a valid assumption. A simple argument along these lines goes as follows. Consider a search for ‘young, fast’ pulsars that begins by coherently analyzing stretches of data that are all ~ 1 day long (possible with $\sim 4 \times 10^{12}$ flops, by Fig. 2.8). Imagine that in the second stage of the search one follows up all templates such that $P_0(f, \vec{\lambda}) > 4.6S_n(f)$, by seeing whether templates with roughly the same parameter values are exceeding this threshold every day. (Here $P_0(f, \vec{\lambda})$ is the power of the stretched data at frequency f , for stretch $\vec{\lambda}$. This threshold implies that one is following up only one out of every hundred templates.) It seems likely that this

second stage will not be more computationally intensive than the first. To exceed this threshold, a pulsar must have $h_c \gtrsim 12h_{3/\text{yr}}$. This is factor of roughly 3 better than if one restricted oneself to coherent searches considered above, but is a factor of 3 worse than the sensitivity one could achieve with unlimited computing power.

A refinement of this strategy would be one in which the first pass consists of several incoherently-added power spectra. That is, one slices the data into N sequential subsets, performs a full search (as described in this paper) for each subset, and adds up the power spectra of the resulting searches for each of the parameter sets. This technique has been used to good effect by radio astronomers searching for pulsars [35]. Since the addition of power spectra is incoherent, there is a loss of signal-to-noise ratio in the final summed power spectrum of $1/\sqrt{N}$ in relation to a full coherent search over the whole timescale. However, the computational savings involved allow one to search stretches of data which are much longer overall. For some optimal choice of N , this will result in higher sensitivities when one follows up candidate detections using coherent searches. D. Nicholson (private communication) has estimated that a 1Tflops computer could perform such a search of 10^7 s of data, over all sky positions but ignoring pulsar spindowns. A subsequent paper will present a concrete analysis of this and other hierarchical scenarios [42].





Figure 2.6: Number of independent points in parameter space as a function of total observation time, using a maximum projected mismatch $\mu_{\max} = 0.3$. The parameter ranges chosen were: (a) maximum gravitational wave frequency 1000Hz, minimum spindown age $\tau_{\min} = 40\text{Yrs}$ (hypothetical young pulsars); (b) maximum gravitational wave frequency 200Hz, minimum spindown age $\tau_{\min} = 10^3\text{Yrs}$ (observed, slow pulsars). The short-dashed curve represents the total number of patches ignoring all f_k . The long-dashed curve is the number of patches including only f_1 in the search. The dotted line is the number of patches including both f_1 and f_2 . Also shown is the empirical fit given in the text; it was normalized by the results shown in (a). In some regimes, searching over an additional spindown parameter would seem to reduce the number of patches; however, this actually only indicates regions where the parameter space extends less than one full patch width in the additional dimension. In such regimes one must properly discard the extra parameter from the search, forcing one to choose always the higher of the curves.




Figure 2.7: The total number of parameter-space points needed to search for pulsars having gravitational wave frequency up to 1kHz, and spindown age greater than $\tau = 40$ Yrs. The solid line is the number computed using the metric and properly accounting for correlations between various terms in the frequency evolution. The dotted line is the same number computed directly by assuming the points must lie on the grid of coordinates used to parameterize the signal. The benefits of using the metric to optimally place the points to be searched in parameter space is clear.




Figure 2.8: The optimum observation time (thick solid line), and maximal projected mismatch (dotted line) as functions of available computational power. Both graphs assume a threshold which gives an overall statistical significance of 99% to any detection (although the results should be insensitive to the precise value). Each of the graphs corresponds to: (a) the situation encountered when searching for periodic sources having gravitational wave frequencies up to 1000Hz, with minimum spindown ages $\tau_{\min} = 40$ Yrs. (b) The equivalent results for gravitational wave frequencies up to 200Hz, with minimum spindown ages $\tau_{\min} = 10^3$ Yrs. The transition region seen in figure (a) is due to the fact that a longer integration time would require searching over an additional spindown parameter, as seen in Fig. 2.6. In this region it is more efficient, as one adds computational power, to lower mismatch thresholds, rather than searching over the additional parameter.




Figure 2.9: The optimum observation time (thick solid line), and maximal projected mismatch (dotted line) as functions of available computational power for directed searches. Both graphs assume a threshold which gives an overall statistical significance of 99% to any detection (although the results are insensitive to the precise value). Each of the graphs corresponds to: (a) the situation encountered when searching for periodic sources having gravitational wave frequencies up to 1000Hz, with minimum spindown ages $\tau_{\min} = 40$ Yrs. (b) The equivalent results for gravitational wave frequencies up to 200Hz, with minimum spindown ages $\tau_{\min} = 10^3$ Yrs. The transition regions, where the optimum observation time does not increase, are due to the fact that a longer integration time would require searching over an additional spindown parameter.

Chapter 3 Time-frequency and hierarchical search techniques for periodic sources

This chapter is adapted from the paper “Searching for periodic sources with LIGO. II: Hierarchical searches,” by Patrick R. Brady and Teviet Creighton [42].

Summary: The detection of quasi-periodic sources of gravitational waves requires the accumulation of signal-to-noise over long observation times. This represents the most difficult data analysis problem facing experimenters with detectors like those at LIGO. If not removed, Earth-motion induced Doppler modulations and intrinsic variations of the gravitational-wave frequency make the signals impossible to detect. These effects can be corrected (removed) using a parameterized model for the frequency evolution. In a previous paper, we introduced such a model and computed the number of independent parameter space points for which corrections must be applied to the data stream in a coherent search. Since this number increases with the observation time, the sensitivity of a search for continuous gravitational-wave signals is computationally bound when data analysis proceeds at a similar rate to data acquisition. In this paper, we extend the formalism developed by Brady *et al.* [Phys. Rev. D **57**, 2101 (1998)], and we compute the number of independent corrections $N_p(\Delta T, N)$ required for incoherent search strategies. These strategies rely on the method of *stacked* power spectra—a demodulated time series is divided into N segments of length ΔT , each segment is Fourier transformed, a power spectrum is computed, and the N spectra are summed up. This method is incoherent; phase information is lost from segment to segment. Nevertheless, power from a signal with fixed frequency (in the corrected time series) is accumulated in a single frequency bin, and amplitude signal-to-noise accumulates as $\sim N^{1/4}$ (assuming the segment length ΔT is held fixed). For fixed available computing power, there are optimal values for N and ΔT which maximize the sensitivity of a search in which data analysis takes a total time $N\Delta T$. We estimate that the optimal sensitivity of an all-sky search that uses incoherent stacks is a factor of 2–4 better than achieved using coherent Fourier transforms, assuming the same available computing power; incoherent methods are computationally efficient at exploring large parameter spaces. We also consider a two-stage hierarchical search in which candidate events from a search using short data segments are followed up in a search using longer data segments. This hierarchical strategy yields a further 20–60% improvement in sensitivity in all-sky (or directed) searches for old (≥ 1000 yr) slow (≤ 200 Hz) pulsars, and for young (≥ 40 yr) fast (≤ 1000 Hz) pulsars. Assuming enhanced LIGO detectors (LIGO-II) and 10^{12} flops of effective computing power, we examine the sensitivity to sources

in three specialized classes. A limited area search for pulsars in the Galactic core would detect objects with gravitational ellipticities of $\epsilon \gtrsim 5 \times 10^{-6}$ at 200 Hz; such limits provide information about the strength of the crust in neutron stars. Gravitational waves emitted by unstable r -modes of newborn neutron stars would be detected out to distances of ~ 8 Mpc, if the r -modes saturate at a dimensionless amplitude of order unity and an optical supernova provides the position of the source on the sky. In searches targeting low-mass x-ray binary systems (in which accretion-driven spin up is balanced by gravitational-wave spin down), it is important to use information from electromagnetic observations to determine the orbital parameters as accurately as possible. An estimate of the difficulty of these searches suggests that objects with x-ray fluxes exceeding $2 \times 10^{-8} \text{erg cm}^{-2} \text{s}^{-1}$ would be detected using the enhanced interferometers in their broadband configuration. This puts Sco X-1 on the verge of detectability in a broadband search; the amplitude signal-to-noise would be increased by a factor of order ~ 5 –10 by operating the interferometer in a signal-recycled, narrow-band configuration. Further work is needed to determine the optimal search strategy when limited information is available about the frequency evolution of a source in a targeted search.

3.1 Introduction

The detection of gravitational waves from periodic sources is seemingly the most straightforward data analysis problem facing gravitational-wave astronomers. It is also the most computationally intensive. The long observation times required to detect these waves mean that Earth-motion induced Doppler effects, and intrinsic frequency drifts, degrade the signal-to-noise if not removed. Since these effects depend sensitively on the location and intrinsic properties of the source, searches for periodic (or quasi-periodic) sources will be limited primarily by the computational resources available for data analysis, rather than the duration of the signals or the lifetime of the instrument. For this reason, it is of paramount importance to explore different search strategies and to determine the optimal approach before the detectors go on line at the end of the century.

In a previous paper [10], hereafter referred to as Paper I, we presented a detailed discussion of issues that arise when one searches for these sources in the detector output. Using a parameterized model for the expected gravitational wave signal, we also presented a method to determine the number of independent parameter values that must be sampled in a search using coherent Fourier transforms (which accumulate the signal to noise in an optimal fashion). The results were presented in the context of single-sky-position directed searches, and all-sky searches, although the method outlined in Paper I is applicable to *any* search over a specified region of parameter space. Livas [26], Jones [27] and Niebauer *et al.* [28] have implemented variants of the coherent search technique without the benefit of the optimization advocated in Paper I.

In this paper, we discuss alternative search algorithms that can better detect quasi-periodic

gravitational waves using broadband detectors. These algorithms achieve better sensitivities than a coherent search with equivalent available computational resources. This improvement is accomplished by combining coherent Fourier transforms with incoherent addition of power spectra, and by using hierarchical searches that follow up the candidate detections from a first pass search.

The most likely sources of quasi-periodic gravitational waves in the frequency bands of terrestrial interferometric detectors are rapidly rotating neutron stars. We use these objects as guides when choosing the scope of the example searches considered below. Nevertheless, the search algorithms are sufficient to detect any source of continuous gravitational waves with slowly changing frequency.

A rotating neutron star will radiate gravitational waves if its mass distribution (or mass-current distribution) is not symmetric about its rotation axis. Several mechanisms that may produce non-axisymmetric deformations of a neutron star, and hence lead to gravitational wave generation, have been discussed in the literature [20, 21, 19, 37, 23, 22]. A neutron star with non-zero quadrupole moment produces gravitational waves at a frequency equal to twice its rotation frequency if it rotates about a principle axis. Equally strong gravitational waves can be emitted at other frequencies when the rotation axis is not aligned with a principal axis of the source [19, 38]. If the star also precesses, the gravitational waves will be produced at three frequencies: the rotation frequency, and the rotation frequency plus and minus the precession frequency [23].

For concreteness, we consider a model gravitational-wave signal with one spectral component. This is not a limitation of our analysis since the search strategy presented below is inherently broadband; it can be used to detect sources that emit gravitational waves at any frequency in the detector pass-band. Additional knowledge of the spectral characteristics of a signal might allow us to improve our sensitivity in the case when multiple spectral components have similar signal-to-noise ratio. In such a circumstance, a modified search algorithm would sum the power at all of the appropriate frequencies. In a background of Gaussian noise, the sensitivity would improve as (number of spectral lines)^{1/4} for only a moderate increase in computational cost.

Finally, we mention several other works that consider searching for quasi-periodic signals in the output of gravitational wave detectors. Data from the resonant bar detectors around the world has been used in searches for periodic sources. New *et al.* [43] have discussed issues in searching for gravitational waves from millisecond pulsars. Krolak [44] and Jaranowski *et al.* [45, 46] have considered using matched filtering to extract information about the continuous wave sources from the data stream. Finally, work is ongoing in the Albert Einstein Institute, Gölm, to investigate line-tracking algorithms based on the Hough transform [47, 48, 49]; this technique looks promising, although we must await results on the computational cost and statistical behavior before we can make a detailed comparison to the techniques described in this paper.

3.1.1 Gravitational waveform

The long observation times required to detect continuous sources of gravitational waves make it necessary to account for changes in the wave frequency. Physical processes responsible for these changes, and the associated timescales, were discussed in Paper I. In addition, the detector moves with respect to the solar system barycenter (which we take to be approximately an inertial frame), introducing Doppler modulations of the gravitational-wave frequency. To account for these two effects, we introduce a parameterized model for the gravitational-wave frequency $f(t; \lambda)$ and phase $\phi(t; \lambda) = 2\pi \int f(t; \lambda) dt$ measured at the detector:

$$f(t; \lambda) = f_0 \left(1 + \frac{\vec{v}}{c} \cdot \hat{n} \right) \left(1 + \sum_{k=1} f_k \left[t + \frac{\vec{x}}{c} \cdot \hat{n} \right]^k \right), \quad (3.1)$$

$$\phi(t; \lambda) = 2\pi f_0 \left(t + \frac{\vec{x}}{c} \cdot \hat{n} + \sum_{k=1} \frac{f_k}{k+1} \left[t + \frac{\vec{x}}{c} \cdot \hat{n} \right]^{k+1} \right). \quad (3.2)$$

Here f_0 is the initial, intrinsic gravitational-wave frequency, $\vec{x}(t)$ is the detector position, $\vec{v}(t)$ is the detector velocity, \hat{n} is a unit vector in the direction of the source, and f_k are arbitrary coefficients that we call *spindown parameters*. (We refer the reader to Paper I for a detailed discussion of this model and its physical origin.) The vector λ denotes the *search parameters* — the parameters of the frequency model that are (generally) unknown in advance. In the most general case that we consider below, the search parameters include frequency f_0 , the polar angles (θ, φ) used to specify \hat{n} , and the spindown parameters f_k :

$$\lambda = (\lambda^0, \lambda^1, \lambda^2, \lambda^3, \lambda^4, \dots) = (f_0, \theta, \varphi, f_1, f_2, \dots). \quad (3.3)$$

We note that the parameter $\lambda^0 = f_0$ defines an overall frequency scale, whereas the remaining parameters define the shape of the phase evolution. It is convenient to introduce the projected vector $\vec{\lambda} = (\lambda^1, \lambda^2, \lambda^3, \lambda^4, \dots)$ of shape parameters alone.

The strain measured at the interferometer is a linear combination of the $+$ and \times polarizations of the gravitational waves, and is given by the real part of

$$h(t; \lambda) = \mathcal{A} e^{-i[\phi(t; \lambda) + \Psi]}. \quad (3.4)$$

The time-dependent amplitude \mathcal{A} and phase Ψ depend on the detector response functions and the orientation of the source; they vary gradually over the course of a day (see references [19, 45]). In what follows, we treat \mathcal{A} and Ψ as constants. Our analysis may be generalized to include the additional phase modulation. This effectively increases the dimension of the parameter space, by one, and the number of points that must be sampled, by ~ 4 , which translates into a reduction in

relative sensitivity of $\sim 6\%$.

3.1.2 Parameter ranges

The computational difficulty of a search for quasi-periodic signals depends on the range of parameter values that are considered in the search. The intrinsic gravitational wave frequency f_0 ranges from (near) zero to some cutoff frequency f_{\max} . If gravitational waves are emitted at twice the rotation frequency, theoretical estimates [50, 51] suggest that

$$f_{\max} \lesssim 1.2 \text{ kHz to } 4 \text{ kHz} \quad (3.5)$$

depending on the equation of state adopted in the neutron star model. Observational evidence—the coincidence of the periods of PSR 1937+21 and PSR 1957+20—favors the lower bound on gravitational wave frequency $f_{\max} \simeq 1.2 \text{ kHz}$ [33]. The spindown parameters f_j are allowed to take any value in the range $|f_j| \leq (1/\tau_{\min})^j$ where $\tau_{\min} \sim f/\dot{f}$ is the characteristic timescale over which the frequency might be expected to change by a factor of order unity. Observations of radio pulsars provide rough guidance about the timescales τ_{\min} . In Paper I we considered two fiducial classes of sources that we denoted: (i) Young, fast pulsars, with $f_{\max} = 1000 \text{ Hz}$ and $\tau_{\min} = 40 \text{ yr}$, and (ii) old, slow pulsars, with $f_{\max} = 200 \text{ Hz}$ and $\tau_{\min} = 1000 \text{ yr}$. To facilitate direct comparison with the achievable sensitivities quoted in Paper I, we again use these two classes to present our results.

The two extremes of sky area to be searched are: (i) zero steradians for a *directed* search in which we know the source location in advance; e.g., a supernova remnant, and (ii) 4π steradians for an *all-sky* search. We consider both of these cases, as well as the intermediate case of a 0.004 steradian search about the galactic center.

It has been suggested recently that the gravitational wave frequency of new-born, rapidly-spinning neutron stars may evolve on a timescale of months rather than decades [52, 53, 54, 55]; an active r -mode instability can radiate away most of a neutron star's angular momentum in the form of gravitational waves within a year. Thus, new-born neutron stars may be loud enough to be detected in other galaxies, in which case optical detection of a supernova can serve as a trigger for a targeted search. Therefore, we consider the case of a directed search for sources with frequencies of $f_{\max}=200 \text{ Hz}$ and evolution timescales of $\tau_{\min}=1 \text{ yr}$.

A final class of sources that we consider are accreting neutron stars in binary systems. Several such binary systems have been identified via x-ray observations; the rotation frequencies of the accreting neutron stars are inferred to be $\sim 250\text{--}350 \text{ Hz}$ ($f_{\max}=700 \text{ Hz}$). Bildsten [56] has argued that these accreting objects in low mass x-ray binaries (LMXBs) may emit detectable amounts of gravitational radiation. Since the positions of these sources are well localized on the sky by their x-ray emissions, the earth-motion induced Doppler modulations of the gravitational waves can be

precisely determined. The difficulty with these sources is the unknown, or poorly-known, orbits of the neutron stars about their stellar companions, and the stochastic accretion-induced variations in their spin. We estimate the size of these effects, and outline a search algorithm in Sec. 3.7.3. These issues deserve further study in an effort to improve the search strategy.

3.1.3 Search technique

In searches for continuous gravitational waves, our sensitivity will be limited by the computational resources available, rather than the duration of the signal or the total amount of data. Therefore, the computational efficiency of a search technique is extremely important. For example, matched filtering (convolution of noise-whitened detector output with a noise-whitened template) may detect a signal with the greatest signal-to-noise ratio for any given stretch of data; however, it becomes computationally prohibitive to search over large parameter spaces with long data stretches. A sub-optimal, but more efficient, algorithm might achieve the best overall sensitivity for fixed computational resources.

We present two possible search strategies to accumulate signal to noise from the data stream. Central to both of these methods is the technique we adopt to demodulate the signal. We can remove the effects of Doppler and spindown modulations by defining a canonical time coordinate

$$t_b[t; \vec{\lambda}] = t + \frac{\vec{x}(t) \cdot \hat{n}}{c} + \sum_{k=1} \frac{f_k}{k+1} \left[t + \frac{\vec{x}(t) \cdot \hat{n}}{c} \right]^{k+1}, \quad (3.6)$$

with respect to which the signal, defined in Eq. (3.4), is perfectly sinusoidal:

$$h(t_b[t; \vec{\lambda}]) = \mathcal{A} e^{-2\pi i f_0 t_b[t; \vec{\lambda}] - i\Psi}. \quad (3.7)$$

(Remember, we treat \mathcal{A} and Ψ as constant in time.) The introduction of the new time coordinate can be achieved by re-sampling the data stream at equal intervals in t_b . The power spectrum, computed from the Fourier transform of the re-sampled data, will consist of a single monochromatic spike, whose amplitude (relative to broadband noise) increases in proportion to the length of the data stretch. In practice the data will be sampled in the detector frame, so that a sample may not occur at the desired value of t_b . Consequently, we advocate the use of nearest-neighbor (stroboscopic) resampling [25]. This method will not substantially reduce the signal to noise in a search provided the detector output is sampled at a sufficiently high frequency. (See Appendix B, Sec. 3.10.)

When the waveform shape parameters $\vec{\lambda}$ are not known in advance, one must search over a mesh of points in parameter space. The result of a phase corrected Fourier transform will be sufficiently monochromatic only if the true signal parameters lie close enough to one of the mesh points. In Sec. 3.2 we rigorously define what is meant by “close enough,” and show how to determine the

number of points for which corrections should be applied. This number depends on the adopted search strategy, and increases as a large power of the total observation time T_{data} . Note that the method of resampling followed by a Fourier transform has the benefit that a single Fourier transform automatically searches over all frequencies f_0 , leaving only the shape parameters $\vec{\lambda}$ to be searched explicitly. Other demodulation techniques, such as matched filtering, must apply separate corrections for each value of f_0 in addition to the $\vec{\lambda}$. This increases the computational cost dramatically.

A signal can also be accumulated *incoherently* from successive stretches of data by adding their power spectra [35]. However, even if each data stretch is demodulated to sufficient precision that the power from a signal is focused in a single Fourier frequency bin, residual errors in $\vec{\lambda}$ may cause the power to be at different frequencies between successive spectra. A more precise knowledge of the phase evolution is required to correct for this drift; i.e., a finer mesh in parameter space. Once a set of parameter corrections $\Delta\vec{\lambda}$ is assumed, it is relatively easy to correct for the frequency drift: successive power spectra are shifted in frequency by a correction factor Δf , where Δf is computed by differencing $f(t; \lambda)$, in Eq. (3.1), between the initial and corrected guesses for $\vec{\lambda}$, as a function of the start time of each data stretch. Once the spectra have been corrected by Δf , they can be added together. This accumulates signal-to-noise less efficiently than coherent phase corrections and FFT's, but is computationally cheaper.

Stack-slide search

The search techniques considered in this paper are variants on the following scheme. First, the data stream is divided into shorter lengths, called *stacks*. Each stack is phase corrected using a mesh of correction points sufficient to confine a putative signal to ~ 1 frequency bin in each stack. This procedure is complicated somewhat by the evolution of the spindown parameters with time. The following recipe determines the parameters used to phase correct each stack. Choose $\vec{\lambda}$ from the fine mesh, compute the *barycentred time* $t + \vec{x}(t) \cdot \hat{n}/c$ for the start of the entire search and the start of a given stack. Call the difference Δt_{start} . The appropriate spindown parameters to use for the given stack are those on the coarse grid which are closest to:

$$f_k = \frac{\sum_{n=k} \lambda^{(2+n)} \binom{n}{k} (\Delta t_{\text{start}})^{n-k}}{1 + \sum_{n=1} \lambda^{(2+n)} (\Delta t_{\text{start}})^n} . \quad (3.8)$$

The phase corrected stacks are FFT'd and the power spectra stored for as long as they are required in the next step. The individual power spectra are then shifted, relative to each other, to correct for residual frequency drift. The corrected power spectra are summed, and searched for spikes that




Figure 3.1: A flowchart representation of the *stacked slide* algorithm to search for sources of continuous gravitational waves. Notice that the computational cost of sampling the fine grid is reduced by sliding the power spectra, rather than re-computing an FFT for each point on the fine grid.

exceed some specified significance threshold¹. The procedure is summarized in the flowchart in Fig. 3.1.

Hierarchical search

We also consider a two-pass hierarchical search strategy. In this case, one performs an initial search of the data using a low threshold that allows for many false alarms. This is followed by a second pass, using longer stretches of data, but searching the parameter space only in the vicinity of the candidate detections of the first pass. This procedure is summarized in Fig. 3.2. The advantage of a hierarchical

¹The method of stacking power-spectra has been used by radio astronomers in deep searches for millisecond pulsars, although all corrections were applied to the data stream via resampling and not sliding the spectra. For more information on the implementation, see [35].




Figure 3.2: A flowchart representation of the *hierarchical* algorithm to search for sources of continuous gravitational waves. It should be noted that while this approach will almost certainly be incorporated into the eventual search algorithm for gravitational waves, the real benefit of such an approach will be to increase the confidence in a detection made using some other technique.

search are two-fold: (i) the low threshold on the first pass allows detection of low-amplitude signals that would otherwise be rejected, and (ii) the second pass can search longer data stretches on a limited computing budget, because of the reduced parameter space being searched, thus excluding false positives from the first pass. If the thresholds and mesh points are optimally chosen between the first and second passes, this technique achieves the best sensitivity of the strategies considered here and in Paper I for given computational resources.

3.1.4 Results

The sensitivity $\Theta = 1/h_{\text{th}}$ of a search is defined in Eqs. (3.38) and (3.39). The threshold strain amplitude h_{th} is defined such that there is a 1% *a priori* probability that detector noise alone will produce an event during the analysis, and therefore is the minimum characteristic strain detectable in the search. We compare our results for the sensitivity Θ to a canonical sensitivity determined by the search threshold $h_{3/\text{yr}} = 4.2\sqrt{S_n(f) \times 10^{-7}\text{Hz}}$, where $S_n(f)$ is the one-sided power spectral density

of the noise in the detector. This threshold is the characteristic amplitude of the weakest source detectable with 99% confidence in a coherent search of 10^7 seconds of data, if the frequency and phase evolution of the signal are known. The relative sensitivity Θ_{rel} is given by $\Theta_{\text{rel}} \equiv h_{3/\text{yr}}/h_{\text{th}}$; a relative sensitivity $\Theta_{\text{rel}} = 0.1$ for a search means that a signal must have a characteristic amplitude $h_c \gtrsim 10 \times h_{3/\text{yr}}$ to be detected in that search. Figure 3.3 shows $h_{3/\text{yr}}$ based on noise spectral estimates for three detector systems in LIGO: the initial detectors are expected to go on-line in the year 2000, with the first science run from 2002–2004; the upgrade to the enhanced detectors should begin in ~ 2004 , with subsequent upgrades leading to, and perhaps past, the advanced detector sensitivity. The expected amplitudes h_c of several putative sources are also shown; we use the definition of h_c given in Eq. (50) of Ref. [1], and Eq. (3.5) of Paper I. The strengths of gravitational waves from the Crab and Vela radio pulsars are upper limits assuming all the rotational energy is lost via gravitational waves. The estimates of waves from the r-mode instability are based on Owen *et al.* [55], and those from Sco X-1 are based on the recent analysis by Bildsten [56].

The results presented in this paper are based on three reasonable assumptions about any search strategy: (i) data analysis should proceed at roughly the same rate as data acquisition; (ii) finite computational resources will be available to perform the data analysis; (iii) in a search, the overall false alarm probability should be small (usually take to be 1%). Given these assumptions, there is an optimal choice for the duration T_{data} of data to be analyzed, and how many stacks it should be divided into. The optimal choice is that which *maximizes* the final sensitivity of a search subject to the constraints (i)–(iii)². To demonstrate the existence of this optimum point, we have plotted the relative sensitivity of a search for young, fast pulsars as a function of the number of stacks and the available computing power in Fig. 3.4. The optimal number of stacks can be read off the plot for fixed computing power. Note that the maximum sensitivity in this plot is quite flat, especially in the regime where one is most computationally bound. This may be extremely relevant when implementing these search techniques; data management issues may impose more severe constraints on the size and number of stacks than computational requirements do. This remains to be explored when the data analysis platforms have been chosen.

Figure 3.5(a) shows the optimal sensitivities that can be achieved, as a function of available computing power, using a stack-slide search. The results are presented for both fiducial classes of pulsars: old ($\tau \geq 1000$ yr) slow ($f \leq 200$ Hz) pulsars, and young ($\tau \geq 40$ yr) fast ($f \leq 1000$ Hz) pulsars. In each case, we have considered both directed and all-sky searches for the sources. The results should be compared with those of Paper I, in which we considered coherent searches without

²It may seem that the optimal duration of data to analyze is equal to the amount of data taken by the instrument. This is not necessarily true for a given algorithm (e.g., the stack-slide search, or the two-stage hierarchical search). Suppose, under the stated assumptions, we determine the optimal amount of data T_{data} to analyze using a given algorithm. Now, hold the computational resources fixed, but increase the amount of data by a factor of 10, so that we have 10 times as long to analyze it. Unfortunately, the computational cost increases by more than a factor of 10 because the number of parameter-space corrections increases faster than T_{data} . Thus, we cannot complete our analysis in the time it takes to acquire the data. Implications of this point are further discussed in Sec. 3.8.

stacking: the use of stacked searches gains a factor of $\sim 2\text{--}4$ in sensitivity.

The use of a two-pass hierarchical search can further improve sensitivities by balancing the computational requirements between the two passes. Figure 3.5(b) shows the sensitivities achievable when each pass uses a stack-slide strategy. The sensitivities achieved exceed those of one-pass stack-slide searches by $\sim 20\text{--}60\%$.

The computational requirements for all-sky, all frequency surveys are sufficiently daunting that we explore three restricted searches in Sec. 3.7: (i) a directed search for a newborn neutron star in the young ($\lesssim 1$ year old) remnant of an extra-galactic supernova, (ii) an area search of the galactic core for pulsars with $\tau \geq 100$ yr and $f \leq 500$ Hz, and (iii) a directed search for an accreting neutron star in a binary system (such as Sco X-1). Figure 3.6 shows the relative sensitivities attainable in such searches. With computational resources capable of 1 Tflops, we expect to see galactic core pulsars with enhanced LIGO if they have non-axisymmetric strains of $\epsilon \gtrsim 5 \times 10^{-6}$ at frequencies of ~ 200 Hz. Estimates of the characteristic strain of gravitational waves from an active r -mode instability in a newborn neutron star suggest that these sources will be detectable by the enhanced interferometers in LIGO out to distances ~ 8 Mpc; the rate of supernovae is ~ 0.6 per year within this distance. Finally, gravitational waves from accreting neutron stars in low-mass x-ray binary systems (LMXBs) may be detectable by enhanced interferometers in LIGO if we can obtain sufficient information about the binary orbit from electromagnetic observations. Sco X-1 is on the margins of detectability using the enhanced LIGO interferometers operating in broadband configuration. We estimate that the amplitude signal-to-noise from these sources could be improved by a factor of $\sim 5\text{--}10$ by operating the interferometer in a signal-recycled, narrow-band configuration.

3.1.5 Organization of the paper

In Sec. 3.2 we extend the metric formalism that was developed in Paper I to determine the number of parameter space points that must be sampled in a search that accumulates signal to noise by summing up power spectra. This method can then be used to compute the number of correction points needed in a stack-slide search. Approximate formulae, useful for estimating the computational cost of a search, are presented for the number of corrections needed in an all-sky search, and also in directed searches of a single sky position.

We discuss the issue of thresholding in Sec. 3.3. Then we present the computational cost estimates, and determine the optimal parameters for single-pass, stack-slide searches in Sec. 3.4.

Section 3.5 contains a general discussion of hierarchical searches for periodic sources using a single interferometer. Schutz [57] has emphasized the potential of hierarchical strategies in searches for periodic sources. The relationship between the threshold in the second stage of the search, and the threshold required in the first stage is discussed in detail. We also present the computational cost of each stage of the search. These results are used in Sec. 3.6 to determine the optimal search

parameters in hierarchical searches for our fiducial classes of sources.

We discuss three specialized searches in Sec. 3.7. In particular, we discuss issues that arise when the source of gravitational waves is in a binary system (e.g., an LMXB). Introducing a parameterized model of the binary orbit, we estimate the number of parameter space points that must be sampled in a search for the gravitational waves from one of the objects in the binary. When the gravitational-wave emitter is accreting material from its companion, we also allow for stochastic changes in frequency due to fluctuations in the accretion rate.

Detailed formulae for the number of points in parameter space when dealing with a stacked search are presented in Appendix A (Sec. 3.9). In Appendix B (Sec. 3.10) we discuss the loss in signal to noise that can occur when using nearest neighbor resampling to apply corrections to the detector output. If the data is sampled at 16 384 Hz, we demonstrate that this method will lose less than 1% of amplitude signal-to-noise for a signal with gravitational wave frequency ≤ 1000 Hz.

3.2 Mismatch

In a detection strategy that searches over a discrete mesh of points in parameter space, the search parameters and signal parameters will never be precisely matched. This mismatch will reduce the signal to noise since the signal will not be precisely monochromatic. It is desirable to quantify this loss, and to choose the grid spacing so that the loss is within acceptable limits. This can be achieved by defining a distance measure on the parameter space based on the fractional losses in detected signal power due to parameter mismatch. In Paper I we derived such a measure in the case where the search was performed using coherent Fourier transforms; this method was modeled after Owen's computation of a metric on the parameter space of coalescing binary waveforms [14]. We extend this approach to the case of incoherent searches, in which several power spectra are added incoherently, or *stacked*, and then searched for spikes.

Let $h(t; \lambda)$ be a hypothetical signal given by Eq. (3.4) with true signal parameters $\lambda = (f_0, \vec{\lambda})$; we use the complex form of $h(t; \lambda)$ without loss of generality. If the data containing this signal are corrected for some nearby set of shape parameters $\vec{\lambda} + \Delta\vec{\lambda}$, the signal will take the form

$$h_b(t; \lambda, \Delta\lambda) = \mathcal{A}e^{-i\{2\pi f_0 t_b[t; \vec{\lambda}] + \phi[t; \lambda] - \phi[t; (f_0, \vec{\lambda} + \Delta\vec{\lambda})]\}} , \quad (3.9)$$

where the subscript b is used to indicate the corrected waveform, and $\phi[t; \lambda]$ is defined in Eq. (3.2). In a stacked search, the data are divided into N segments of equal length ΔT seconds, each of these segments is Fourier transformed, and then a power spectrum $P_h(f; \lambda, \Delta\vec{\lambda})$ is computed according to the formula

$$H(f; \lambda, \Delta\vec{\lambda}) = 2 \sum_{k=1}^N |\tilde{h}_k(f; \lambda, \Delta\vec{\lambda})|^2 . \quad (3.10)$$

The Fourier transform of each individual segment is defined to be

$$\tilde{h}_k(f; \lambda, \Delta \vec{\lambda}) = \frac{A}{\sqrt{\Delta T}} \int_{(k-1)\Delta T}^{k\Delta T} e^{i\Delta\phi[t; \lambda, \Delta \lambda]} dt_b \quad (3.11)$$

where $\Delta\phi$ is given by

$$\Delta\phi[t; \lambda, \Delta \lambda] = 2\pi(f - f_0)t_b + \phi[t; (f_0, \vec{\lambda} + \Delta \vec{\lambda})] - \phi[t; \lambda]. \quad (3.12)$$

Here $\Delta \lambda = (f - f_0, \Delta \vec{\lambda})$ denotes the error in matching the modulation shape parameters *and* the error in sampling the resulting power spectrum at the wrong frequency. Both of these errors lead to a reduction in the detected power relative to the optimum case where the carrier frequency and the phase modulation are precisely matched.

The *mismatch* $m(\lambda, \Delta \lambda)$, which is the fractional reduction in power due to imperfect phase correction *and* sampling at the wrong Fourier carrier frequency, is defined to be

$$m(\lambda, \Delta \lambda) = 1 - \frac{H(f; \lambda, \Delta \vec{\lambda})}{H(f_0; \lambda, \vec{0})}. \quad (3.13)$$

Remember $\lambda = (\lambda^0, \vec{\lambda}) = (f_0, \lambda^1, \lambda^2 \dots)$. Substituting the expressions for H from Eq. (3.10) into Eq. (3.13), we find

$$m(\lambda, \Delta \lambda) = 1 - \frac{1}{N\mathcal{A}^2} \sum_{k=1}^N |\tilde{h}_k(f; \lambda, \Delta \vec{\lambda})|^2. \quad (3.14)$$

It is easily shown that $m(\lambda, \Delta \lambda)$ has a local minimum of zero when $\Delta \lambda = 0$. We therefore expand the mismatch in powers of $\Delta \lambda$ to find

$$m(\lambda, \Delta \lambda) = \sum_{\alpha, \beta} g_{\alpha\beta}(\lambda) \Delta \lambda^\alpha \Delta \lambda^\beta + \mathcal{O}(\Delta \lambda^3), \quad (3.15)$$

where (α, β) are summed over $0, 1, \dots, j, \dots$. The quantity $g_{\alpha\beta}$ is a local distance metric on the parameter space. This metric is explicitly given by

$$g_{\alpha\beta}(\lambda) \equiv \frac{1}{2} \partial_{\Delta \lambda^\alpha} \partial_{\Delta \lambda^\beta} m(\lambda, \Delta \lambda) |_{\Delta \lambda=0}, \quad (3.16)$$

where $\partial_{\Delta \lambda^\alpha}$ denotes a partial derivative with respect to $\Delta \lambda^\alpha$. It is convenient to express $g_{\alpha\beta}$ as a sum of metrics computed for the individual stacks, that is

$$g_{\alpha\beta}(\lambda) = \frac{1}{N} \sum_{k=1}^N g_{\alpha\beta}^{(k)}(\lambda) \quad (3.17)$$

where the individual stack metrics $g_{\alpha\beta}^{(k)}(\boldsymbol{\lambda})$ are explicitly given by

$$g_{\alpha\beta}^{(k)}(\boldsymbol{\lambda}) = \langle \partial_{\Delta\lambda^\alpha} \Delta\phi \partial_{\Delta\lambda^\beta} \Delta\phi \rangle_k - \langle \partial_{\Delta\lambda^\alpha} \Delta\phi \rangle_k \langle \partial_{\Delta\lambda^\beta} \Delta\phi \rangle_k . \quad (3.18)$$

The phase error $\Delta\phi$ is given in Eq. (3.12), and we use the notation

$$\langle \dots \rangle_k = \frac{1}{\Delta T} \int_{(k-1)\Delta T}^{k\Delta T} (\dots) dt_b \Big|_{\Delta\lambda=0} . \quad (3.19)$$

In a search, we will look for spikes in the power spectrum $H(f; \boldsymbol{\lambda}, \Delta\vec{\lambda})$ computed from the detector output; that is, we will look for local maxima in the frequency parameter f . The relevant measure of distance in the space of shape parameters $\vec{\lambda}$ is the fractional loss in power due to mismatched parameters $\Delta\vec{\lambda}$, but after maximizing over frequency. We therefore define the *projected mismatch* $\mu(\vec{\lambda}, \Delta\vec{\lambda})$ to be

$$\mu(\vec{\lambda}, \Delta\vec{\lambda}) = \min_f m(\boldsymbol{\lambda}, \Delta\boldsymbol{\lambda}) = \sum_{i,j} \gamma_{ij}(\vec{\lambda}) \Delta\lambda^i \Delta\lambda^j , \quad (3.20)$$

where

$$\gamma_{ij} = \left(g_{ij} - \frac{g_{i0}g_{j0}}{g_{00}} \right)_{\lambda^0=f_{\max}} \quad (3.21)$$

is the mismatch metric projected onto the subspace of shape parameters, and f_{\max} is the maximum frequency that we include in the search. The meaning of the minimization \min_f is clear from the definition of the mismatch in Eq. (3.14).

Technically, γ_{ij} should be computed from $g_{\alpha\beta}$ evaluated at the specific value of f at which the minimum projected mismatch occurred. Since this number is unknown in advance of detection, we evaluate γ_{ij} for the maximum frequency that we include in the search. In this way we never underestimate the projected mismatch.

The distance function, and in particular the metric in Eq. (3.21), can be used to determine the number of discrete mesh points that must be sampled in a search. Let \mathcal{P} be the space of all parameter values $\vec{\lambda}$ to be searched over, and define the maximal mismatch μ_{\max} to be the largest fractional loss of power that we are willing to tolerate from a putative source with parameters in \mathcal{P} . For the model waveform in Eqs. (3.1), (3.2), and (3.4) this parameter space is coordinatized by $\vec{\lambda} = (\theta, \phi, f_1, f_2, \dots)$ where θ, ϕ denote location of the source on the sky, and f_j are related to the time derivative of the intrinsic frequency of the source. Each correction point of the mesh is considered to be at the center of a cube with side $2\sqrt{\mu_{\max}/n}$, where n is the dimension of \mathcal{P} ; this insures that all points in \mathcal{P} are within a proper distance μ_{\max} of a discrete mesh point as measured

with the metric γ_{ij} . The number of patches required to fill the parameter space is

$$N_p(\Delta T, \mu_{\max}, N) = \frac{\int_{\mathcal{P}} \sqrt{\det |\gamma_{ij}|} d^n \lambda}{(2\sqrt{\mu_{\max}/n})^n}. \quad (3.22)$$

Since μ_{\max} is the maximum loss in detected power *after* the power spectra have been added, $N_p(\Delta T, \mu_{\max}, N)$ is the number of patches required to construct the *fine* mesh in the stacked search strategy described in Sec. 3.1.3. The coarse mesh in the stacked search strategy requires that the spikes in the individual power spectra be reduced by no more than μ_{\max} ; consequently, the number of points in such a mesh is simply $N_p(\Delta T, \mu_{\max}, 1)$.

3.2.1 Directed search

In most cases, the forms of Eqs. (3.2) and (3.6) are sufficiently complicated to defy analytical solution, especially since \vec{v} in Eq. (3.2) should properly be taken from the true ephemeris of the Earth during the period of observation. However, for the case of a directed search, that is a search in just a single sky direction, the phase correction is polynomial in t , and the metric can be computed analytically. To a good approximation, the metric is flat — the spacing of points in parameter space is independent of the value of the spindown parameters (f_1, f_2, \dots, f_s) . For a given number s of spindown parameters in a search, the right-hand side of Eq. (3.22) can be evaluated analytically. The result is expressed as a product $\mathcal{N}_s G_s$, where

$$\mathcal{N}_s = \frac{f_{\max}^s (\Delta T)^{s(s+3)/2}}{(\mu_{\max}/s)^{s/2} \tau_{\min}^{s(s+1)/2}} \quad (3.23)$$

depends on the maximum frequency f_{\max} (in Hz), the length of each stack ΔT (in seconds), the maximal mismatch μ_{\max} , and the minimum spindown age τ_{\min} (in seconds) considered in the search. The dependence on the number of stacks N is contained in $G_s(N)$, which are given by

$$G_0(N) = 1, \quad (3.24)$$

$$G_1(N) \approx 0.524N, \quad (3.25)$$

$$G_2(N) \approx 0.0708N^3, \quad (3.26)$$

$$G_3(N) \approx 0.00243N^6, \quad (3.27)$$

when $N \gg 4$. The detailed expressions for $G_s(N)$ are presented in Appendix A (Sec. 3.9). For up to 3 spindown terms in the search, the number of patches is then

$$N_p(\Delta T, \mu_{\max}, N) = \max_{s \in \{0,1,2,3\}} [\mathcal{N}_s G_s(N)]. \quad (3.28)$$

The maximization accounts for the situation where increasing s , the dimension of the parameter space, decreases the value of $\mathcal{N}_s G_s$ because the parameter space extends less than one patch width in the new spindown coordinate f_s ; one should not search over this coordinate.

3.2.2 Sky search

For signal modulations that are more complicated than simple power-law frequency drift, it is impossible to compute N_p analytically. In an actual search over sky positions as well as spindown, one should properly compute the mismatch metric numerically, using the exact ephemeris of the Earth in computing the detector position. In Paper I we computed $N_p(\Delta T, \mu_{\max}, 1)$ numerically, with the simplification that both the Earth's rotation and orbital motion were taken to be circular. However, in this paper we are concerned also with the dependence of N_p on the number of stacks N . This significantly complicates the calculation of the metric and its determinant, and makes it necessary to adopt some approximations in the calculation. Fortunately, the results of interest here are insensitive to small errors in N_p .

In Paper I we mentioned that there are strong correlations between sky position and spindown parameters. This requires the use of the full $s + 2$ dimensional metric. However, these correlations are due primarily to the Earth's orbital motion, which has a low-order Taylor approximation for times much less than a year. Therefore, we treat the number of patches as the product of the number of spindown patches times the number of sky positions \mathcal{M}_s , computed analytically using only the Earth's rotational motion. We note that this approximation is appropriate only for computing the number of patches; *when actually demodulating the signals, the true orbital motion would have to be included*. This approximation works well so long as the orbital *residuals* (the remaining orbital modulations after correction on this sky mesh) are much smaller than the spindown corrections being made at the same power in t . The residual orbital velocity at any power t^k is roughly

$$\sim \frac{\alpha_k}{k! \sqrt{\mathcal{M}_s}} \times \frac{r\Omega}{c} \times (\Omega t)^k, \quad (3.29)$$

where α_k is a number of order unity, \mathcal{M}_s is the number of sky patches, and $r = 1\text{AU}$ and $\Omega = 2\pi/\text{yr}$ are the Earth's orbital radius and angular velocity. When the range in this residual is comparable to or larger than the range in the corresponding spindown term $f_k t^k$, the "spindown" parameter space must be expanded to include the orbital residuals. The range in α_k is difficult to arrive at analytically. We have found that using a maximum value of ≈ 0.3 gives good agreement with the numerical results of Paper I (i.e., for $N = 1$), to within factors of ~ 2 .

One other approximation was made in computing the number of sky patches. We found that the measure $\sqrt{\det |\gamma_{ij}|}$ for the sky position metric is almost constant in the azimuth φ , and has a polar angle dependence that is dominantly of the form $\sin 2\theta$. When performing the integral over

sky positions, we approximated the measure by $\sqrt{|\det|\gamma_{ij}|} \simeq \text{constant} \times \sin 2\theta$; this approximation is accurate to about one part in 10^4 .

Given these approximations, the number of patches for a sky search is

$$N_p = \max_{s \in \{0,1,2,3\}} \left[\mathcal{M}_s \bar{\mathcal{N}}_s G_s \prod_{k=0}^s \left(1 + \frac{0.3r\Omega^{k+1}\tau_{\min}^k}{c k! \sqrt{\mathcal{M}_s}} \right) \right]. \quad (3.30)$$

The number of sky patches \mathcal{M}_s , in the $(s+2)$ -dimensional search, is given approximately by

$$\mathcal{M}_s \approx \frac{f_{\max}^2}{4\mu_{\max}/(s+2)} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (3.31)$$

where

$$A = 0.014, \quad (3.32)$$

$$B = 0.046(\Delta T/1 \text{ day})^2, \quad (3.33)$$

$$C = 0.18(\Delta T/1 \text{ day})^5 N^3. \quad (3.34)$$

This is a fit to the analytic result given in Appendix A (Sec. 3.9). The number of spindown patches $\bar{\mathcal{N}}_s G_s$ in the $(s+2)$ -dimensional search is

$$\bar{\mathcal{N}}_s G_s = \frac{s^{s/2}}{(s+2)^{s/2}} \mathcal{N}_s G_s \quad (3.35)$$

where \mathcal{N}_s and G_s are given in Eqs. (3.23)–(3.27), and the prefactor on the right corrects for the sky dimensions. The remaining product terms \prod in Eq. (3.30) represent the increase in the size of the spindown space in order to include the orbital residuals.

3.3 Thresholds and sensitivities

The thresholds for a search are determined under the assumption that the detector noise is a stationary, Gaussian random process with zero mean and power spectral density $S_n(f)$. In the absence of a signal, the power $P_n(f) = 2|\tilde{n}(f)|^2$ at each sampled frequency is exponentially distributed with probability density function $e^{-P_n/S_n}/S_n$. The statistic for stacked spectra is $\rho = \sum_1^N P_n(f)$. The cumulative probability distribution function for ρ , in the absence of a signal, is

$$\text{CDF}[\rho/S_n, N] = \int_0^{\rho/S_n} e^{-r} \frac{r^{N-1}}{(N-1)!} dr = \frac{\gamma(N, \rho/S_n)}{(N-1)!} \quad (3.36)$$

where $\gamma(N, \rho/S_n)$ is an incomplete gamma function.

A (candidate) detection occurs whenever ρ in some frequency bin exceeds a pre-specified threshold

ρ_c chosen so that the probability of a false trigger due to noise alone is small. There are $f_{\max}\Delta T$ Fourier bins in each spectrum, and $N_p(\Delta T, \mu_{\max}, N)$ spectra in the entire search. Therefore, we assume that a search consists of $N_p f_{\max}\Delta T$ independent trials of the statistic ρ , and compute the expected number of false events F to be

$$F = f_{\max}\Delta T N_p(\Delta T, \mu_{\max}, N) (1 - \text{CDF}[\rho_c/S_n, N]). \quad (3.37)$$

(In reality, there will be correlations between the statistic computed for different frequencies and different patches. Since this will reduce the number of independent trials, Eq. (3.37) overestimates the number of false events. This is a small effect that should not change the overall sensitivity of a search by much. It is only in the case that the number of trials is initially small that one should be concerned with this effect; unfortunately, we operate in the other extreme.) If $F \ll 1$, the number of false events F is approximately equal to the probability that an event is caused by noise in the detector. Consequently, $\alpha = 1 - F$ can be thought of as the confidence of detection. In a non-hierarchical search, the threshold ρ_c is set by specifying α and then inverting Eq. (3.37).

Finally, how does the threshold ρ_c affect the sensitivity of our search? We define a threshold amplitude h_{th} to be the minimum dimensionless signal amplitude that we expect to register as a detection in the search, that is

$$h_{\text{th}} = \sqrt{\frac{(\rho_c/N - S_n)}{\langle F_+^2(\Theta, \Phi, \Psi) \rangle (1 - \langle \mu \rangle) \Delta T}} \quad (3.38)$$

where $\langle F_+^2(\Theta, \Phi, \Psi) \rangle$ is the square of the detector response averaged over all possible source positions and orientations, and $\langle \mu \rangle = \mu_{\max}/3$ is the expected mismatch of a signal that is randomly located within a patch³. The *sensitivity* Θ of the search is then defined by

$$\Theta \equiv \frac{1}{h_{\text{th}}} \propto \sqrt{\frac{(1 - \mu_{\max}/3)\Delta T}{\rho_c/N - S_n}}. \quad (3.39)$$

For any given search strategy, the goal of optimization is to maximize the final sensitivity of the search, given limited computational power.

3.4 Stack-slide search

A stack-slide search is the simplest alternative to coherent searches we consider here. The main steps involved in the algorithm are shown in the flow chart of Fig. 3.1. In this section we estimate

³The average expected power loss for a source randomly placed within a cubical patch is $\langle \mu \rangle = \mu_{\max}/3$. In Paper I we quoted an average that was computed for ellipsoidal patches; this is not appropriate to the cubical grid that will likely be used in a real search.

the computational cost of each step, and determine the ultimate sensitivity of this technique.

The first step, before the search begins, is to specify the size of the parameter space to be searched (i.e., choose f_{\max} , τ , and a region of the sky), the computational power P that will be available to do the data analysis, and an acceptable false alarm probability. From these, one can determine optimal values for the maximal mismatch μ_{\max} for a patch, the number of stacks N , and the length ΔT of each stack, using the optimization scheme discussed at the end of this section. For now, we treat these as free parameters.

Coarse and fine grids are laid down on the parameter space with $N_{pc} = N_p(\Delta T, \mu_{\max}, 1)$ and $N_{pf} = N_p(\Delta T, \mu_{\max}, N)$ points, respectively. The data-stream is low-pass filtered to the upper cutoff frequency f_{\max} , and broken into N stretches of length ΔT .

Each of the steps above have negligible computational cost since they are done only once for the entire search. The subsequent steps, on the other hand, must be executed for each of the N_{pc} correction points.

Each stretch of data is re-sampled (at the Nyquist frequency $2f_{\max}$) and simultaneously demodulated by stroboscopic sampling for a set of demodulation parameters selected from the coarse grid. The result is N demodulated time series, each one consisting of $n = 2f_{\max}\Delta T$ samples. Since stroboscopic demodulation only shifts one in every few thousand data points (assuming a sampling rate at the detector of 16 384 Hz), the computational cost of the demodulation itself is negligible.

Each stretch of data is then Fourier transformed using a fast Fourier transform (FFT) algorithm with a computational cost of $3nN \log_2(n)$ floating point operations. Power spectra are computed for each Fourier series, costing 3 floating point operations per frequency bin, i.e., a total cost of $1.5nN$ floating point operations.

For demodulation parameters in the coarse grid, the power of a matched signal will be confined to ~ 1 Fourier bin in each power spectrum, but not necessarily the *same* bin in different spectra. To insure that power from a signal is accumulated by summing the N spectra, we must apply the following steps for each of the N_{pf} correction points in the *fine* mesh.

For a given set of parameters from the fine mesh, one determines which power spectra from the coarse grid are to be summed, using Eqs. (3.8). For each of the N spectra to be stacked, the frequency of a putative signal with initial frequency f_{\max} is computed using Eq. (3.1). These can be read from a look-up table generated when the meshes were laid out. Each spectrum is re-indexed so that the power from such a signal would be in the same frequency bin (we ignore the computational cost of this step), and the spectra are added [$0.5n(N - 1)$ floating point operations]. We automatically account for corrections at other frequencies by applying the fine grid corrections in this way. It may be possible to reduce the computational cost of this portion of the search by noting, for example, that we over count the fine grid corrections for signals with frequency $f_{\max}/2$ by a factor of 2^n where n is the dimension of the parameter space being explored. Since it is difficult to assess the

Table 3.1: The optimum stack length ΔT , number of stacks N , maximal projected mismatch μ_{\max} , and the relative sensitivity $\Theta_{\text{rel}} = h_{3/\text{yr}}/h_{\text{th}}$ as functions of available computational power for directed, stack-slide searches. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value). The optimization was performed numerically using simulated annealing, which accounts for some of the fluctuations in the observation times.

Young ($\tau_{\min} = 40\text{yr}$), fast ($f_{\max} = 1000\text{Hz}$) pulsars.

Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.52	157	0.49	0.13
3.16×10^{11}	1.15	56	0.43	0.13
1.00×10^{12}	1.24	60	0.41	0.14
3.16×10^{12}	1.51	55	0.43	0.15
1.00×10^{13}	1.65	58	0.43	0.16
3.16×10^{13}	1.79	62	0.43	0.17
1.00×10^{14}	2.01	63	0.42	0.18
3.16×10^{14}	2.31	62	0.42	0.19
1.00×10^{15}	2.57	64	0.43	0.20
3.16×10^{15}	2.93	64	0.43	0.21
1.00×10^{16}	3.17	69	0.44	0.22

Old ($\tau_{\min} = 10^3\text{yr}$), slow ($f_{\max} = 200\text{Hz}$) pulsars.

Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^7	1.15	236	0.49	0.23
3.16×10^7	1.58	206	0.49	0.25
1.00×10^8	2.04	199	0.50	0.28
3.16×10^8	2.63	192	0.49	0.31
1.00×10^9	3.54	172	0.49	0.34
3.16×10^9	4.29	183	0.49	0.38
1.00×10^{10}	5.26	192	0.49	0.42
3.16×10^{10}	6.52	197	0.49	0.47
1.00×10^{11}	12.62	95	0.47	0.52
3.16×10^{11}	19.41	59	0.43	0.55
1.00×10^{12}	21.28	62	0.42	0.58

feasibility of using this in a real search, we simply mention it so that it might be explored at the time of implementation.

The resulting stacked spectrum is scanned for peaks that exceed the threshold ρ_c . Since this has negligible computational cost, the number of floating point operations required for the entire search is

$$C = 3nNN_{pc} [\log_2(n) + 0.5 + N_{pf}(N - 1)/(6NN_{pc})]. \quad (3.40)$$

If data analysis proceeds at the same rate as data acquisition, the computational power P required to complete a search is $P = C/N\Delta T$ floating-point operations per second (Flops). Equation (3.40)

Table 3.2: Same as in Table 3.1, but for an all-sky search.

Young ($\tau_{\min} = 40\text{yr}$), fast ($f_{\max} = 1000\text{Hz}$) pulsars.				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.04	1232	0.50	0.07
3.16×10^{11}	0.06	890	0.49	0.07
1.00×10^{12}	0.08	651	0.49	0.08
3.16×10^{12}	0.11	475	0.49	0.08
1.00×10^{13}	0.16	329	0.50	0.08
3.16×10^{13}	0.21	250	0.48	0.09
1.00×10^{14}	0.30	166	0.50	0.09
3.16×10^{14}	0.36	152	0.49	0.10
1.00×10^{15}	0.37	179	0.50	0.10
3.16×10^{15}	0.41	184	0.50	0.11
1.00×10^{16}	0.51	154	0.50	0.11
Old ($\tau_{\min} = 10^3\text{yr}$), slow ($f_{\max} = 200\text{Hz}$) pulsars.				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	1.18	54	0.42	0.14
3.16×10^{11}	1.33	55	0.42	0.14
1.00×10^{12}	1.55	53	0.42	0.15
3.16×10^{12}	2.40	31	0.31	0.16
1.00×10^{13}	2.84	27	0.27	0.16
3.16×10^{13}	3.10	27	0.28	0.17
1.00×10^{14}	3.22	29	0.30	0.17
3.16×10^{14}	3.38	30	0.28	0.18
1.00×10^{15}	3.77	29	0.29	0.19
3.16×10^{15}	4.67	24	0.27	0.19
1.00×10^{16}	5.81	19	0.21	0.20

and the definition of $n = 2f_{\max}\Delta T$ imply that the computational power is

$$P = 6f_{\max}N_{pc} [\log_2(n) + 0.5 + N_{pf}(N - 1)/(6NN_{pc})]. \quad (3.41)$$

The final sensitivity Θ , defined in Eq. (3.39), of the search is determined once we know the function N_p , the frequency f_{\max} , the maximal mismatch μ_{\max} , and the confidence level $\alpha = 1 - F$. An optimized algorithm will maximize Θ as a function of μ_{\max} , N , and ΔT , subject to the constraints imposed by fixing the false alarm probability F , and the computational power P .

The results of the optimization procedure are given in Tables 3.1 and 3.2 for the fiducial classes of pulsar defined in Sec. 3.1.2. In each case we have set the probability of a false alarm threshold at $F = 0.01$ (indicating a 99% confidence that detector noise will not produce an event above threshold), and have determined optimal values of μ_{\max} , N , and ΔT for a range of values of the available computational power. Table 3.1 shows the results for directed searches for young, fast

($\tau_{\min}=40$ yr, $f_{\max}=1000$ Hz) and old, slow ($\tau_{\min}=1000$ yr, $f_{\max}=200$ Hz) pulsars, respectively. Table 3.2 shows the results for an all-sky search for the same two classes of source. The optimal sensitivities achieved by these searches are summarized in Fig. 3.5(a) of the Introduction.

It is worth pointing out here that we have assumed memory and I/O requirements to be irrelevant in determining the computational cost. Specifically, power spectra generated on the coarse grid are stored as long as they are needed for the sliding portion of the search. In practice, it may prove necessary to recompute power spectra, or to retrieve them from slow-access data storage. Such hardware- and implementation-specific details are beyond the scope of this paper.

3.5 Hierarchical search: general remarks

The basic hierarchical strategy involving a two pass search is represented schematically in Fig. 3.2. In the first pass, $N^{(1)}$ stacks of data of length $\Delta T^{(1)}$ are demodulated on a coarse and fine mesh of correction points computed for some mismatch level $\mu^{(1)}$, and then searched by stacked Fourier transforms. A threshold signal-to-noise level is chosen that will, in general, admit many false alarms. In the second stage, $N^{(2)}$ stacks of length $\Delta T^{(2)}$ are searched on a finer mesh of points computed at a mismatch level $\mu^{(2)}$, but only in the vicinity of those events that passed the first-stage threshold. The second stage will involve fewer correction points than the first, so the second-stage transforms can be made longer and more sensitive. The goal of optimization is to find some combination of $\Delta T^{(1)}$, $\Delta T^{(2)}$, $\mu^{(1)}$, $\mu^{(2)}$, $N^{(1)}$, and $N^{(2)}$ that maximizes the final sensitivity for fixed computational power P , and second pass false alarm probability $F^{(2)}$.

3.5.1 Thresholds

In the first pass of a hierarchical search, each of $N_f^{(1)} = f_{\max}\Delta T^{(1)}$ frequency bins in $N_{pf}^{(1)} = N_p(\Delta T^{(1)}, \mu^{(1)}, N^{(1)})$ stacked power spectra will be scanned for threshold crossing events. If (as we assume) all of these trials are statistically independent, the number of false events above the threshold $\rho^{(1)}$ will be

$$F^{(1)} = N_{pf}^{(1)} N_f^{(1)} (1 - \text{CDF}[\rho^{(1)}/S_n^{(1)}, N^{(1)}]). \quad (3.42)$$

We assume that the number of false events will significantly exceed the number of true signals in this pass; consequently, the number of events to be analyzed in the second pass will be $F^{(1)}$.

The second stage uses a coarse grid with $N_{pc}^{(2)} = N_p(\Delta T^{(2)}, \mu^{(2)}, 1)$ points, and a fine grid with $N_{pf}^{(2)} = N_p(\Delta T^{(2)}, \mu^{(2)}, N^{(2)})$ points. On average each false alarm will require $N_{pc}^{(2)}/N_{pf}^{(1)}$ coarse grid points, and $N_{pf}^{(2)}/N_{pf}^{(1)}$ fine grid points in the second stage. (When a second-pass mesh is coarser than the first pass's parameter determination, the corresponding ratio should be taken as unity.) Furthermore, since the first stage will identify the candidate signal's frequency to within ~ 2

frequency bins, the second-stage search should be over the $2\Delta T^{(2)}/\Delta T^{(1)}$ second-stage frequency bins that lie in this frequency range. Once again, we assume the noise in all frequency bins (and over all grid points) is independent, so the number of false events that exceed the threshold $\rho^{(2)}$ in the second stage is

$$\begin{aligned}
1 - \alpha &= F^{(2)} \\
&= 2F^{(1)} \frac{N_{pf}^{(2)} \Delta T^{(2)}}{N_{pf}^{(1)} \Delta T^{(1)}} (1 - \text{CDF}[\rho^{(2)}/S_n^{(2)}, N^{(2)}]) \\
&= 2f_{\max} \Delta T^{(2)} N_{pf}^{(2)} (1 - \text{CDF}[\rho^{(1)}/S_n^{(1)}, N^{(1)}]) \\
&\quad \times (1 - \text{CDF}[\rho^{(2)}/S_n^{(2)}, N^{(2)}]), \tag{3.43}
\end{aligned}$$

where α is our desired confidence level for the overall search.

The thresholds $\rho^{(1)}$ and $\rho^{(2)}$ cannot be assigned independently; rather, they should be chosen so that any *true* signal buried in the noise that would exceed (in expectation value) the second-stage threshold will have passed the first-stage threshold. In other words, it serves no purpose to set $\rho^{(2)}$ any lower than the weakest signal that would have exceeded $\rho^{(1)}$. A signal that is expected to pass the second-stage threshold exactly has an amplitude $|\tilde{h}^{(2)}|^2 = \rho^{(2)} - N^{(2)} S_n^{(2)}$. We define the *false dismissal probability* D to be the probability that such a signal will be falsely rejected in the first pass. Since the spectral power of a true signal increases with $N\Delta T$, the signal seen in the first pass has amplitude $|\tilde{h}^{(1)}|^2 = |\tilde{h}^{(2)}|^2 (N^{(1)} \Delta T^{(1)}) / (N^{(2)} \Delta T^{(2)})$, and the thresholds satisfy the relation

$$\begin{aligned}
D &= \text{CDF} \left[\frac{\rho^{(1)} - |\tilde{h}^{(1)}|^2}{S_n^{(1)}}, N^{(1)} \right] \\
&= \text{CDF} \left[\frac{\rho^{(1)}}{S_n^{(1)}} - \left(\frac{\rho^{(2)}}{S_n^{(2)}} - N^{(2)} \right) \frac{S_n^{(2)} N^{(1)} \Delta T^{(1)}}{S_n^{(1)} N^{(2)} \Delta T^{(2)}}, N^{(1)} \right]. \tag{3.44}
\end{aligned}$$

Now, for any choice of $\Delta T^{(1)}$, $\Delta T^{(2)}$, etc., the thresholds $\rho^{(1)}$ and $\rho^{(2)}$ are completely constrained by our choices of final confidence level α and false dismissal probability D . The false dismissal probability is fixed at $D = 0.01$ in our optimization; this is an acceptably low level, meaning that only one signal in a hundred is expected to be lost in this type of search.

3.5.2 Computational costs

The computational cost $C^{(1)}$ of the first stage of the search follows the same formula as for a simple non-hierarchical search, that is

$$\begin{aligned}
C^{(1)} &= 6f_{\max} N^{(1)} \Delta T^{(1)} N_{pc}^{(1)} [\log_2(2f_{\max} \Delta T^{(1)}) \\
&\quad + 0.5 + N_{pf}^{(1)} (N^{(1)} - 1) / (6N^{(1)} N_{pc}^{(1)})]. \tag{3.45}
\end{aligned}$$

For each of the $F^{(1)}$ first-stage triggers, the second stage requires $N_{pc}^{(2)}/N_{pf}^{(1)}$ (minimum 1) coarse grid corrections (each involving $N^{(2)}$ FFT's of length $\Delta T^{(2)}$), along with $N_{pf}^{(2)}/N_{pf}^{(1)}$ (minimum 1) frequency shifts and spectrum additions. Each of the coarse grid corrections requires the usual $2f_{\max}N^{(2)}\Delta T^{(2)}[3\log_2(2f_{\max}\Delta T^{(2)})+0.5]$ floating-point operations. The incoherent frequency shifts and spectrum additions require only $2(N^{(2)} - 1)\Delta T^{(2)}/\Delta T^{(1)}$ floating point operations since the frequency correction and power summation need only be applied over a bandwidth of ~ 2 first-pass frequency bins. The total cost of the second pass is therefore:

$$C^{(2)} = \frac{2F^{(1)}N^{(2)}\Delta T^{(2)}N_{pc}^{(2)}}{N_{pf}^{(1)}} \left\{ 3f_{\max}[\log_2(2f_{\max}\Delta T^{(2)}) + 0.5] + \frac{N_{pf}^{(2)}(N^{(2)} - 1)}{N^{(2)}N_{pc}^{(2)}\Delta T^{(1)}} \right\}. \quad (3.46)$$

We require that data analysis proceed at the rate of data acquisition. Since the amount of data used in the second-stage of the search will generally be greater than that used in the first, we require that the analysis be completed in $N^{(2)}\Delta T^{(2)}$ seconds. Thus the computational power is given by

$$P = (C^{(1)} + C^{(2)})/N^{(2)}\Delta T^{(2)}. \quad (3.47)$$

Our final sensitivity Θ is given by Eq. (3.39), using the observation time, mismatch level, and threshold of the *second* stage of the search. Optimization then consists of maximizing this function over the six parameters $\Delta T^{(1)}$, $\Delta T^{(2)}$, $\mu^{(1)}$, $\mu^{(2)}$, $N^{(1)}$, and $N^{(2)}$, for specified α , D , and P .

3.6 Hierarchical search with stacking

It turns out that the optimization described in the previous section is only weakly sensitive to the parameters $\mu^{(1)}$ and $\mu^{(2)}$; that is, even if we choose values for $\mu^{(1)}$ and $\mu^{(2)}$ quite different from the optimal ones, we can recover nearly all of the sensitivity by adjusting the other parameters for the same computational power P . In particular, if we arbitrarily fix $\mu^{(1)} = \mu^{(2)} = 0.3$ and re-optimize, we obtain sensitivities within 20% of the optimal.

This becomes very useful when we consider the generalized two-stage hierarchical search *with* stacking. Normally this would involve optimizing over six variables ($\mu^{(1),(2)}$, $N^{(1),(2)}$, and $\Delta T^{(1),(2)}$) with one constraint on P . However, by assuming that we can continue to set $\mu^{(1)} = \mu^{(2)} = 0.3$ with minimal loss of sensitivity, we can reduce our degrees of freedom back down to four minus one constraint.

The results of this optimization for our four canonical example searches are given in Tables 3.3 and 3.4. We have chosen a final confidence level $\alpha = 0.99$ and a false dismissal probability of $D = 0.01$.

Table 3.3: The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of directed, hierarchical searches. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max} = 0.3$. The last column gives the overall relative sensitivity $\Theta_{\text{rel}} = h_{3/\text{yr}}/h_{\text{th}}$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Young ($\tau_{\min} = 40\text{yr}$), fast ($f_{\max} = 1000\text{Hz}$) pulsars.						
Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	2.12	16	2.40	31	0.30	0.19
3.16×10^{11}	2.18	20	2.87	35	0.30	0.21
1.00×10^{12}	4.52	9	4.59	22	0.30	0.23
3.16×10^{12}	4.18	12	5.24	22	0.30	0.24
1.00×10^{13}	8.87	5	8.93	12	0.30	0.25
3.16×10^{13}	5.19	13	5.62	29	0.30	0.27
1.00×10^{14}	6.61	11	8.11	20	0.30	0.28
3.16×10^{14}	9.41	8	11.70	13	0.30	0.29
1.00×10^{15}	9.88	9	10.57	19	0.30	0.31
3.16×10^{15}	8.15	13	8.93	31	0.30	0.32
1.00×10^{16}	12.17	9	15.78	15	0.30	0.34

Old ($\tau_{\min} = 10^3\text{yr}$), slow ($f_{\max} = 200\text{Hz}$) pulsars.						
Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^7	2.37	58	2.39	185	0.30	0.36
2.51×10^7	3.43	47	3.46	132	0.30	0.39
6.31×10^7	4.31	47	4.69	110	0.30	0.42
1.58×10^8	6.66	32	7.07	76	0.30	0.46
3.98×10^8	4.77	81	5.69	152	0.30	0.50
1.00×10^9	10.24	31	11.69	62	0.30	0.55
2.51×10^9	9.75	48	9.80	115	0.30	0.59
6.31×10^9	14.35	34	15.00	75	0.30	0.64
1.58×10^{10}	17.71	32	20.65	60	0.30	0.69
3.98×10^{10}	19.78	36	23.01	69	0.30	0.74
1.00×10^{11}	25.44	32	28.76	64	0.30	0.81

Table 3.3 shows the results for directed searches for young, fast ($\tau_{\min}=40$ yr, $f_{\max}=1000$ Hz) and old, slow ($\tau_{\min}=1000$ yr, $f_{\max}=200$ Hz) pulsars, respectively. Table 3.4 shows the results for all-sky searches for the same two classes of source. The optimal sensitivities achieved by these searches are summarized in Fig. 3.5(b) in the Introduction.

3.7 Specialized searches

The strongest sources of continuous gravitational waves are likely to be the most difficult to detect since the frequency of the waves will be changing significantly as the source radiates angular momentum. As we have seen in the previous sections, an all sky search for these sources is unlikely to

Table 3.4: Same as in Table 3.3, but for an all-sky search.

Young ($\tau_{\min} = 40\text{yr}$), fast ($f_{\max} = 1000\text{Hz}$) pulsars.						
Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.09	169	0.10	462	0.30	0.09
3.16×10^{11}	0.12	160	0.12	451	0.30	0.09
1.00×10^{12}	0.16	105	0.16	349	0.30	0.10
3.16×10^{12}	0.20	83	0.26	204	0.30	0.11
1.00×10^{13}	0.28	56	0.34	174	0.30	0.11
3.16×10^{13}	0.50	30	0.56	97	0.30	0.12
1.00×10^{14}	0.60	31	0.78	72	0.30	0.13
3.16×10^{14}	0.94	21	1.07	55	0.30	0.14
1.00×10^{15}	1.17	22	1.30	47	0.30	0.15
3.16×10^{15}	1.33	23	1.77	40	0.30	0.16
1.00×10^{16}	2.64	11	2.75	26	0.30	0.17
Old ($\tau_{\min} = 10^3\text{yr}$), slow ($f_{\max} = 200\text{Hz}$) pulsars.						
Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	3.68	9	4.10	15	0.30	0.20
3.16×10^{11}	5.44	6	7.28	8	0.30	0.21
1.00×10^{12}	8.09	4	8.55	8	0.30	0.23
3.16×10^{12}	9.78	4	10.14	7	0.30	0.24
1.00×10^{13}	14.78	3	14.95	5	0.30	0.25
3.16×10^{13}	16.75	3	19.54	3	0.30	0.26
1.00×10^{14}	14.88	4	16.18	6	0.30	0.27
3.16×10^{14}	20.80	3	23.89	3	0.30	0.28
1.00×10^{15}	22.66	3	25.47	3	0.30	0.28
3.16×10^{15}	35.03	2	36.25	3	0.30	0.29
1.00×10^{16}	20.79	4	23.89	5	0.30	0.29

achieve the desired sensitivity with available computational resources. To reach better sensitivity levels, it will be useful to consider targeted searches for specific types of source. In this section, we consider three such searches: (i) neutron stars in the galactic core as an example of a limited area sky survey, (ii) newborn neutron stars triggered on optically observed extra-galactic supernovae, and (iii) low mass x-ray binary systems such as Sco X-1.

3.7.1 Galactic core pulsars

Area surveys of the sky will certainly begin with the region most likely to hold a large number of nearby sources. Based on population models of radio pulsars in our Galaxy [58], there should be many rapidly rotating neutron stars in the galactic bulge. As an example of a limited area search, we therefore consider the optimal strategy for searching an area of 0.004 steradians about the galactic core, for sources with frequencies $f \leq 500$ Hz and spindown ages $\tau \geq 100$ yr. The choice of a

Table 3.5: The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of an hierarchical search for pulsars located in a sky region of 0.004 steradians about the Galactic center, with $\tau \geq 100$ yr and $f \leq 500$ Hz. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max} = 0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	4.72	7	5.50	13	0.30	0.22
3.16×10^{11}	5.56	7	5.65	14	0.30	0.23
1.00×10^{12}	7.76	5	8.64	10	0.30	0.24
3.16×10^{12}	9.05	5	10.97	7	0.30	0.24
1.00×10^{13}	15.81	3	16.84	5	0.30	0.26
3.16×10^{13}	18.33	3	18.37	5	0.30	0.27
1.00×10^{14}	19.90	3	23.86	3	0.30	0.27
3.16×10^{14}	30.07	2	31.67	3	0.30	0.28
1.00×10^{15}	33.21	2	35.33	3	0.30	0.29
3.16×10^{15}	35.26	2	39.99	3	0.30	0.30
1.00×10^{16}	38.13	2	43.57	3	0.30	0.31

0.004 steradian search is arbitrary; it includes the entire molecular cloud complex at the core of the galaxy (~ 300 pc radius at a distance of ~ 8.5 kpc).

It is easy to include a correction factor, to allow for this limited area, in our calculation of the number of patches by reducing the ranges of the integral over \mathcal{P} in Eq. (3.22). Given the approximations in Sec. 3.2.2, this amounts to reducing N_p in Eq. (3.30) by

$$0.97 \times \left(\frac{0.004}{4\pi} \right), \quad (3.48)$$

where the multiplicative factor 0.97 is the correction for the difference in functional form between the mismatch metric and the angular area metric $d\Omega^2 = \sin^2 \theta d\theta d\phi$ in the direction of the galactic center (i.e., -28.9° declination).

The optimal choices of $N^{(1),(2)}$ and $\Delta T^{(1),(2)}$ for a hierarchical stacked search are shown in Table 3.5 as a function of available computing power; the relative sensitivity of this search is shown in Fig. 3.6 of the Introduction.

We note from Eq. (3.6) of Paper I that gravitational waves from rapidly rotating neutron stars might be expected to have a characteristic amplitude of

$$h_c = 2.3 \times 10^{-25} \frac{\epsilon}{10^{-5}} \frac{I_{zz}}{10^{45} \text{gcm}^2} \frac{8.5 \text{kpc}}{r} \left(\frac{f}{500 \text{Hz}} \right)^2, \quad (3.49)$$

where $\epsilon = (I_{xx} - I_{yy})/I_{zz}$ is the non-axisymmetric strain, I_{ij} is the moment of inertia tensor, r is the distance to the source, f is the gravitational wave frequency, and h_c has been averaged over the detector responses to various source inclinations [1]. Theoretical estimates of the strength of

the crystalline neutron star crust suggest that it can support static deformations of up to $\epsilon \sim 10^{-5}$, though most neutron stars probably support smaller deformations. From Figs. 3.3 and 3.6, we see that 1 Tflops of computing power should allow us to detect pulsars with strains as small as $\epsilon \sim 5 \times 10^{-6}$ at 8.5 kpc using enhanced LIGO detectors.

3.7.2 Newborn neutron stars

Several recent papers [52, 53, 54] have indicated that newly-formed fast-spinning neutron stars may be copious emitters of gravitational radiation. If the newborn neutron star is rotating sufficiently fast, its r -modes (axial-vector current oscillations whose restoring force is the Coriolis force) are unstable to gravitational radiation reaction. As the star cools, viscous interactions eventually damp the modes in isolated neutron stars. Numerical studies [55] indicate that neutron stars born with rotational frequencies above several hundred Hz will radiate away most of their angular momentum in the form of gravitational waves during their first year of life. Estimates of the viscous timescales, and the superfluid transition temperature, suggest that the r -modes are stabilized when the star cools below $\sim 10^9$ K and are rotating at ~ 100 – 200 Hz. During the evolutionary phase when most of the angular momentum is lost, the amplitude and spindown timescale are expected to be

$$h_c = 1.2 \times 10^{-24} \sqrt{\kappa} \left(\frac{f}{1\text{kHz}} \right)^3 \left(\frac{20\text{Mpc}}{r} \right) \quad (3.50)$$

$$\tau \approx \frac{580\text{s}}{\kappa} \left(\frac{1\text{kHz}}{f} \right)^6 \approx 6t. \quad (3.51)$$

These estimates are based on Eqs. (4.9) and (5.13) in Ref. [55]. (We note that the “characteristic amplitude” used in Ref. [55] is appropriate to estimate the strength of burst sources, and is different from our h_c .) Here κ is a dimensionless constant of order unity; it parameterizes our ignorance of the non-linear evolution of the r -mode instability. The distance to the neutron star is r , and t is the actual age of the star. Figure 3.3 shows h_c as a function of frequency with $\kappa = 1$ at distances $r = 2$ Mpc and 20 Mpc.

Sources outside our Galaxy are potentially detectable due to the high gravitational-luminosity of a newborn neutron star with an active r -mode instability. Nevertheless, it is a significant challenge to develop a feasible search strategy for these signals since the frequency evolves on such short timescales (compared to those considered above). One approach is to perform directed searches on optically observed supernova explosions. Although some supernovae may not be optically visible, and this instability may not operate in all newborn neutron stars, the computational benefits of targeting supernovae are substantial (if not essential). Based on the estimates in Ref. [55], most of the signal-to-noise is accumulated during the final stages of spindown. With limited computational resources, it seems best to limit the directed searches to frequencies $\lesssim 200$ Hz, when the spindown

Table 3.6: The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of an hierarchical for newborn neutron stars spinning down due to an active r -mode instability. We assume that a supernova has been identified and accurately located on the sky, so this is a directed search for an object with $\tau_{\min} = 1$ yr, and $f \leq 200$ Hz. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max} = 0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	1.08	4	1.16	6	0.30	0.08
3.16×10^{11}	1.22	4	1.25	7	0.30	0.09
1.00×10^{12}	1.22	5	1.26	8	0.30	0.09
3.16×10^{12}	1.67	4	2.16	5	0.30	0.10
1.00×10^{13}	2.51	3	2.79	3	0.30	0.10
3.16×10^{13}	2.63	3	3.39	3	0.30	0.11
1.00×10^{14}	2.28	4	2.32	6	0.30	0.11
3.16×10^{14}	4.58	2	4.97	3	0.30	0.11
1.00×10^{15}	3.94	3	4.04	3	0.30	0.12
3.16×10^{15}	5.94	2	6.07	3	0.30	0.12
1.00×10^{16}	4.63	3	5.00	3	0.30	0.13

timescale is ~ 1 yr. Table 3.6 shows the optimal search criteria in a hierarchical stacked search for neutron stars aged two months or older; the upper frequency cutoff is $f_{\max} = 200$ Hz and the minimum spindown timescale is $\tau_{\min} = 1$ yr. The sensitivities achievable in a search are shown in Fig. 3.6 of the Introduction.

Figure 3.6 shows that 1 Tflops of computing power will *not* suffice to detect newborn neutron stars as far away as the Virgo cluster (≈ 20 Mpc); however, such sources will be marginally detectable within ~ 8 Mpc by enhanced LIGO detectors. The NBG catalog [59] lists 165 galaxies within this distance (assuming a Hubble expansion of 75 km/s/Mpc, retarded by the Virgo cluster). From the Hubble types and luminosities of these galaxies, and the supernova event rates in [60], we estimate a total supernova rate of ~ 0.6 per year in this volume, of which $\sim 10\%$ would be of type Ia, $\sim 20\%$ of type Ib or Ic, and $\sim 70\%$ of type II. (We note that the total rate is consistent with values given in Ref. [61].) At present, it is not known what fraction of these will produce neutron stars with unstable r -modes.

3.7.3 X-ray binaries

A low-mass x-ray binary (LMXB) is a neutron star orbiting around a stellar companion from which it accretes matter. The accretion process deposits both energy and angular momentum onto the neutron star. The energy is radiated away as x-rays, while the angular momentum spins the star up. Bildsten [56] has suggested that the accretion could create non-axisymmetric temperature gradients in the star, resulting in a substantial mass quadrupole and gravitational wave emission. The star

spins up until the gravitational waves are strong enough to radiate away the angular momentum at the same rate as it is accreting; according to Bildsten's estimates the equilibrium occurs at a gravitational-wave frequency ~ 500 Hz. The characteristic gravitational-wave amplitudes from these sources would be

$$h_c \gtrsim 4 \times 10^{-27} \left(\frac{R}{10\text{km}} \right)^{3/4} \left(\frac{M}{1.4M_\odot} \right)^{-1/4} \times \left(\frac{F}{10^{-8}\text{erg cm}^{-2}\text{s}^{-1}} \right)^{1/2} \left(\frac{f}{600\text{Hz}} \right)^{-1/2}, \quad (3.52)$$

where R and M are the radius and mass of the neutron star, and F is the observed x-ray flux at the Earth.

The amplitude of the gravitational waves from these sources make them excellent candidates for targeted searches. If the source is an x-ray binary pulsar—an accreting neutron star whose rotation is observable in radio waves—then one can apply the exact phase correction deduced from the radio timing data to optimally detect the gravitational waves. (In this process, one must assume a relationship between the gravitational-wave and radio pulsation frequencies.) Unfortunately, radio pulsations have not been detected from the rapidly rotating neutron stars in all LMXBs (i.e., neutron stars that rotate hundreds of times a second). In the absence of direct radio observations, estimates of the neutron-star rotation rates are obtained from high-frequency periodic, or quasi-periodic, oscillations in the x-ray output during Type I x-ray bursts. (See Ref. [62] for a summary.) But this does not provide precise timing data for a coherent phase correction. To detect gravitational waves from these sources, one must search over the parameter space of Doppler modulations due to the neutron-star orbit around its companion, and fluctuations in the gravitational-wave frequency due to variable accretion rates. The Doppler effects of the gravitational-wave detector's motion can be computed exactly, because the sky position of the source is known.

In most cases, the orbital period of an x-ray binary can be deduced from periodicity in its x-ray or optical light curve. In some cases, the radial component of the orbital velocity can be computed by observing an optical emission line from the accretion disk, as was done with the bright x-ray binary Sco X-1 [63]. Such observations do not determine the phase modulation of the gravitational-wave signal with sufficient precision to make the search trivial; however, they do substantially constrain the parameter space of modulations.

In this subsection, we consider a directed search for an x-ray binary in which the orbital parameters are known up to an uncertainty δv in the radial velocity v_r of the neutron star, and an uncertainty $\delta\phi$ in the orbital phase. It is assumed that long-term photometric observations of the source can give the orbital period P to sufficient precision that we need not search over it explicitly.

We therefore parameterize the phase modulations as follows:

$$\phi(t; \lambda) = 2\pi f_0 \left(t + \frac{v_1 P}{2\pi c} \cos 2\pi t/P + \frac{v_2 P}{2\pi c} \sin 2\pi t/P \right), \quad (3.53)$$

where $\lambda = (f_0, v_1, v_2)$ are our search parameters, the gravitational-wave frequency f_0 is constrained to be $\leq f_{\max}$, and the pair (v_1, v_2) is constrained to lie within an annular arc of radius v_r , width δv , and arc angle $\delta\phi$.

Applying the formalism developed in section 3.2 to this problem gives essentially the same result as for a sky search over Earth-rotation-induced Doppler modulations if one converts time units by the ratio P/day . In the case of the Earth's rotation, a search over sky positions \hat{n} corresponds to a search over an area $\pi v_{\text{rot}}^2 \cos^2(\lambda)$ in the equatorial components of the source's velocity relative to the detector, whereas in the case of a binary orbit, the search is over a coordinate area $v_r \delta v \delta\phi$. So we can simply multiply Eqn. (3.30) by the ratio of these coordinate areas to obtain the number of grid points N_p in the parameter space:

$$N_p \approx \frac{(f_{\max} P)^2 v_r \delta v \delta\phi}{2\mu_{\max} c^2} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (3.54)$$

$$A = 0.5, \quad (3.55)$$

$$B = 1.6(\Delta T/P)^2, \quad (3.56)$$

$$C = 6.4(\Delta T/P)^5 N^3. \quad (3.57)$$

Accounting for the intrinsic phase variations of the spinning neutron star itself is problematic since changes in the accretion rate may lead to stochastic variations in the rotation frequency of the star. Consider a typical neutron star (with radius 10 km, mass $1.4M_\odot$ and moment of inertia 10^{45} g cm^2) in an accreting system. If the torque from accretion turns off, or doubles, for some reason, we can expect the gravitational-wave frequency to drift by

$$\Delta f \simeq 3.6 \times 10^{-8} \left(\frac{\dot{M}}{1.5 \times 10^{-9} M_\odot \text{ yr}^{-1}} \right) \left(\frac{t}{\text{day}} \right) \text{ Hz}, \quad (3.58)$$

over a time t , where \dot{M} is the rate of accretion onto the neutron star. For the sources of interest to us, accretion proceeds at (or near) the Eddington rate $\dot{M} \simeq 1.5 \times 10^{-9} M_\odot \text{ yr}^{-1}$. If we require that the frequency drifts by less than one Fourier bin $\Delta f \leq 1/\Delta T$ during a coherent observation, the observing time $t = \Delta T$ must satisfy

$$\Delta T \leq 18 \text{ days}. \quad (3.59)$$

The frequency will make a random walk as the accretion rate fluctuates about its mean value. This type of random walk cannot be modeled as a low-order polynomial in time. Nevertheless, the

Table 3.7: The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of an hierarchical search for sources of continuous gravitational waves that are in binary systems. We assume the orbit is characterized by two orthogonal velocity parameters that are known to within a total error of $17(\text{km/s})^2$, and that the frequency $f \leq 500$ Hz is experiencing a random walk typical of Eddington-rate accretion. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\text{max}} = 0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{max}	Θ_{rel}
1.00×10^{11}	7.12	11	8.60	53	0.30	0.40
3.16×10^{11}	7.33	14	8.51	57	0.30	0.40
1.00×10^{12}	4.72	49	8.31	58	0.30	0.40
3.16×10^{12}	5.13	50	8.47	58	0.30	0.41
1.00×10^{13}	5.84	57	8.27	56	0.30	0.41
3.16×10^{13}	6.14	56	8.46	57	0.30	0.41
1.00×10^{14}	6.78	55	8.38	57	0.30	0.42
3.16×10^{14}	6.69	48	9.17	57	0.30	0.42
1.00×10^{15}	7.40	49	8.93	58	0.30	0.42
3.16×10^{15}	7.53	55	8.99	56	0.30	0.42
1.00×10^{16}	7.00	51	9.71	58	0.30	0.42

stack-slide technique is well suited to search for these sources since corrections for the stochastic changes in frequency can be applied by shifting the stacks by +1, 0, or -1 frequency bins when required. In a search using N stacks, each of length ΔT , this kind of correction would be applied after a time t such that

$$3.6 \times 10^{-8} \text{Hz}(t/\text{day}) = 1/\Delta T. \quad (3.60)$$

The number of times that these corrections must be applied is then $N\Delta T/t$, and the number of distinct frequency evolutions traced out in this procedure is $3^{N\Delta T/t}$. Monte-Carlo simulations of stacked FFT's of signals undergoing random walks in frequency have shown that one can increase t by up to a factor of 4, i.e., allowing drifts of up to ± 2 frequency bins, while incurring only $\sim 20\%$ losses in the final summed power; however, we have not yet studied in detail how this combines with the mismatches generated from other demodulations, or how to search over all demodulations together in an optimal way. For now we assume a factor of $3^{0.003N(\Delta T/\text{day})^2}$ extra points in our search mesh for mismatches of $\mu_{\text{max}} \simeq 0.3$.

As an example, we consider a search for gravitational waves from the neutron star in Sco X-1. This system has an orbital period $P = 0.787313 \pm 0.000001$ days [64], a radial orbital velocity amplitude of $v_r = 58.2 \pm 3.0$ km/s, an orbital phase known to ± 0.10 radians [63], and an inferred gravitational-wave frequency $f_{\text{max}} \approx 500$ Hz [56]. We note that the uncertainty in P is basically

negligible over the < 18 day coherent integrations expected. The remaining uncertainties give:

$$N_p = \frac{10^5}{\mu_{\max}} 3^{0.003N(\Delta T/\text{day})^2} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (3.61)$$

$$A = 0.5, \quad (3.62)$$

$$B = 2.6(\Delta T/\text{day})^2, \quad (3.63)$$

$$C = 21(\Delta T/\text{day})^5 N^3, \quad (3.64)$$

where it is understood that $\Delta T \leq 18$ days, in order for the random-walk stack-slide corrections to achieve maximum sensitivity.

Table 3.7 shows the optimal search criteria for a hierarchical stacked search for the Sco X-1 pulsar under these assumptions. The sensitivities achievable in such a search are shown in Figure 3.6 of the Introduction. We see that 1 Tflops of computing power may be sufficient to detect this source using enhanced LIGO detectors if it is radiating most of the accreting angular momentum as gravitational waves. The sensitivity to these sources would be enhanced by a factor ~ 5 –10 if the interferometer is operated in a signal-recycled, narrow-band configuration during the search.

3.8 Future directions

We have presented in this manuscript the rudiments of a search algorithm for sources of continuous gravitational waves. We have completely neglected overheads associated with memory and I/O in our analysis. These are potentially serious issues that depend on the platform chosen for the analysis. To investigate them requires implementation of the schemes presented in this paper. This will test the algorithms with real data, and, hopefully, highlight the shortcomings of the computational cost estimates presented here. It will also allow direct comparison with other approaches such as the line-tracking method that is being explored by Papa et al. [47, 48, 49].

Several issues remain to be explored. The two-stage hierarchical searches discussed in this paper use only 60 days of data to achieve optimal sensitivity. At LIGO, it is planned to acquire ~ 1 yr of data in the first science run. Should we discard all but ~ 60 days of it? Clearly not. We have explored two possible search algorithms in this paper, but there are infinitely many algorithms to choose from. The goal is to find the one that uses *all* of the data to achieve optimal sensitivity in the sense that we have used it. The algorithm that achieves this goal will be more complex than the two-stage hierarchical algorithm — we suspect it will be a multi-stage hierarchical algorithm. At this point, the urgency is to implement any reasonable search algorithm. Experimental advances will probably outpace gains achieved by optimizing the algorithm to use all the data.

Further theoretical work is also needed to determine the parameter space that should be searched, especially in the case of active r -mode instabilities and radiating neutron stars in LMXBs. We have

performed preliminary estimates in this paper; however, the promising nature of these sources should make them high priorities in targeted searches.

Finally, it would be worthwhile to consider what advantage, if any, can be gained by using data from multiple interferometers at the initial detection stage of a search for continuous gravitational waves.

3.9 Appendix A: Patch number formulae

The approximate formulae given in Eqs. (3.25)–(3.27) are valid when $N \gg 4$. General expressions for the G_s can be derived by setting $\vec{x} = \vec{0}$ in Eqs. (3.2) and (3.6), and using Eqs. (3.12), (3.17), (3.18), (3.19), (3.21), and (3.22):

$$G_1(N) = \frac{\pi}{6\sqrt{5}} \sqrt{5N^2 - 4} \quad (3.65)$$

$$G_2(N) = \frac{\pi^2}{180\sqrt{105}} \sqrt{175N^6 - 840N^4 + 1100N^2 - 432} \quad (3.66)$$

$$G_3(N) = \frac{\pi^3}{75600\sqrt{105}} \sqrt{\begin{aligned} &3675N^{12} - 58800N^{10} + 363160N^8 - 1053360N^6 \\ &+ 1484336N^4 - 987840N^2 + 248832 \end{aligned}} \quad (3.67)$$

In Eq. (3.22) we approximate the metric as having constant determinant and evaluate it at the point of zero spindown; this introduces small errors of order $([N\Delta T]/\tau)^2$.

Equations (3.31)–(3.34) for the number of sky patches ignoring spindown and orbital motions provide an empirical fit to a numerical evaluation of the metric determinant. The determinant was found to have an approximate functional dependence $\sqrt{|\gamma_{ij}|} \sim |\sin 2\theta|$ with corrections of order $v/c \approx 10^{-4}$. Assuming this dependence to be exact, Eqs. (3.21) and (3.22) give

$$\mathcal{M}_s = \frac{2/\pi}{4\mu_{\max}/(s+2)} \sqrt{\det \left| g_{ij} - \frac{g_{0i}g_{0j}}{g_{00}} \right|_{\substack{f=f_{\max} \\ \theta=\pi/2}}} \quad (3.68)$$

Here $g_{\alpha\beta}$ is the mismatch metric, defined in Eqs. (3.17)–(3.19), computed using only the Earth-rotation-induced Doppler modulation. Since the Earth's rotation is a simple circular motion, and since we are evaluating the metric at a single point in parameter space, we can carry out the integrals in Eq. (3.19) analytically, to obtain

$$g_{\alpha\beta} = b_{\alpha\beta} - \frac{1}{N} \sum_{k=1}^N a_{k\alpha} a_{k\beta} \quad (3.69)$$

where

$$a_{k0} = \left(k - \frac{1}{2}\right) \Delta\Phi + \frac{v/c}{\Delta\Phi} \{\sin(k\Delta\Phi) - \sin([k-1]\Delta\Phi)\} \quad (3.70)$$

$$a_{k1} = \frac{f_{\max}v/c}{\Delta\Phi} \{\sin(k\Delta\Phi) - \sin([k-1]\Delta\Phi)\} \quad (3.71)$$

$$a_{k2} = \frac{f_{\max}v/c}{\Delta\Phi} \{\cos([k-1]\Delta\Phi) - \cos(k\Delta\Phi)\} \quad (3.72)$$

$$b_{00} = \frac{1}{12\Phi} \left\{ -24 + \frac{v}{c} + 6 \left(\frac{v}{c}\right)^2 \Phi + 4\Phi^3 + 24\frac{v}{c} \cos \Phi + 24\frac{v}{c} \Phi \sin \Phi + 3 \left(\frac{v}{c}\right)^2 \sin 2\Phi \right\} \quad (3.73)$$

$$b_{11} = \frac{f_{\max}^2(v/c)^2}{4\Phi} (2\Phi + \sin 2\Phi) \quad (3.74)$$

$$b_{22} = \frac{f_{\max}^2(v/c)^2}{4\Phi} (2\Phi - \sin 2\Phi) \quad (3.75)$$

$$b_{01} = b_{10} = \frac{f_{\max}v}{4\Phi} \left\{ -4 + 2\frac{v}{c}\Phi + 4 \cos \Phi + 4\Phi \sin \Phi + \frac{v}{c} \sin 2\Phi \right\} \quad (3.76)$$

$$b_{02} = b_{20} = \frac{f_{\max}v/c}{2\Phi} \left\{ -2\Phi \cos \Phi + \sin \Phi \left(2 + \frac{v}{c} \sin \Phi \right) \right\} \quad (3.77)$$

$$b_{12} = b_{21} = \frac{f_{\max}^2(v/c)^2}{2\Phi} \sin^2 \Phi . \quad (3.78)$$

Here $\Phi = 2\pi N\Delta T/(1 \text{ day})$ is the total angle over which the Earth rotates during the observation, $\Delta\Phi = \Phi/N$ is the angle rotated during each stretch of the data, and v is the maximum radial velocity relative to the detector at latitude $\lambda = 45^\circ$ of a point at a polar angle $\theta = \pi/4$ on the sky, that is

$$v = \frac{2\pi R_{\text{earth}} \cos \lambda \sin \theta}{1 \text{ day}} . \quad (3.79)$$

3.10 Appendix B: Resampling error

In this paper, we have assumed that coherent phase corrections are achieved through stroboscopic resampling: a demodulated time coordinate $t_b[t]$ is constructed, and the data stream $h(t)$ is sampled at equal intervals in t_b at the Nyquist rate for the highest frequency signal present, $f_{\text{Nyquist}} = 2f_{\max}$. However, since the data stream is initially sampled at some finite rate $f_s = Rf_{\text{Nyquist}}$ (where R is the *oversampling factor*), this can introduce errors: in general, there will not be a data point exactly at a given value of t_b , so the nearest (in time) datum must be substituted. Consequently, there will be residual phase errors $\Delta\Phi(t) \in [-\pi/2R, \pi/2R]$ caused by rounding to the nearest datum even if one chooses a phase model whose frequency and modulation parameters exactly match the signal. The phase of the resampled signal drifts until the timing error is $\geq 1/2f_s$, at which point one corrects the phase by sampling an adjacent datum, which shifts in time by $1/f_s$. These residual phase errors reduce the Fourier amplitude of the signal by a fraction

$$F = \left| \frac{1}{N} \sum_{k=1}^N e^{i\Delta\Phi(k/f_{\text{Nyquist}})} \right| , \quad (3.80)$$

where $N = f_{\text{Nyquist}}\Delta T$ is the total number of points in the resampled data stream.

The uncorrected signal will in general drift by many radians in phase, which is the reason why we must apply phase corrections in the first place. This means that $\Delta\Phi(t)$ will sweep through the range $[-\pi/2R, \pi/2R)$ many times over the course of the observation. So, regardless of the precise form of the phase evolution, we expect $\Delta\Phi(k/f_{\text{Nyquist}})$ to be essentially evenly distributed over this interval. Thus, replacing the sum with an expectation integral, we have

$$F \approx \left| \frac{R}{\pi} \int_{-\pi/2R}^{\pi/2R} e^{i\Phi} d\Phi \right| = \frac{\sin(\pi/2R)}{\pi/2R}. \quad (3.81)$$

The fractional losses in amplitude $1 - F$ for a few values of the oversampling factor R are given in Table 3.8.

The highest gravitational-wave frequencies we consider are 1000 Hz, requiring a resampling rate of $f_{\text{Nyquist}} = 2000$ Hz. Since LIGO will acquire data at a rate of 16 384Hz, corresponding to an oversampling factor of $R > 8$, we have a maximum signal loss due to resampling of $1 - F = 0.6\%$. Resampling errors will increase if the number of data samples is reduced by some factor *before* phase-correcting.

Table 3.8: The percentage reduction ($1 - F$) in amplitude of a signal as a function of the oversampling factor R . The LIGO interferometers will collect data at $f_s = 16\,384$ Hz, so that the data will be oversampled by $R \geq 4$ compared to the maximum gravitational wave frequency that we expect on physical grounds. In fact, it seems more likely that $R \simeq 8$ for real signals.

$R =$	2	3	4	5	6	7	8
$1 - F =$	10.0%	4.5%	2.6%	1.6%	1.1%	0.8%	0.6




Figure 3.3: Characteristic amplitudes h_c [see Eq. (3.5) in [10]] for several postulated periodic sources, compared with sensitivities $h_{3/\text{yr}}$ of the initial, enhanced (broad-band and narrow-band) and advanced detectors in LIGO. ($h_{3/\text{yr}}$ corresponds to the amplitude h_c of the weakest source detectable with 99% confidence in $\frac{1}{3}\text{yr} = 10^7\text{s}$ integration time, if the frequency and phase of the signal, as measured at the detector, are known in advance.) Long-dashed lines show the expected signal strength as a function of frequency for pulsars at a distance of 8.5 kpc assuming a gravitational ellipticity $\epsilon = 10^{-5}$ of the source (see Ref. [10]). Upper limits are plotted for the Crab and Vela pulsars, assuming their entire measured spindown is due to gravitational wave emission. The characteristic amplitude of waves from r -modes is also shown. These signals are not precisely periodic; rather, they chirp downward through a frequency band of ~ 200 Hz in 2×10^7 seconds. Finally, the strength of the gravitational waves from LMXBs, normalized to the observed x-ray flux from Sco X-1, is plotted under the assumption that gravitational waves are entirely responsible for their angular momentum loss.




Figure 3.4: A contour plot of the relative sensitivity defined in Sec. 3.1.4 as a function of available computing power, and the number of stacks in the search. The plot indicates the sensitivity that can be achieved using a stack-slide search for sources with $f \leq 1000$ Hz and $\tau \geq 40$ yr. The darkest shading represents the worst sensitivity. For fixed number of stacks the sensitivity improves with increasing computing power as expected. Notice that for fixed computing resources there is generally a point of optimal sensitivity; indeed the number of stacks at this optimal operating point should be compared with those given in Table 3.1. It is important to notice that the maximum falls off very slowly as the number of stacks increases.





Figure 3.5: Relative amplitude sensitivities $\Theta_{\text{rel}} = h_{3/\text{yr}}/h_{\text{th}}$ achievable with given computational resources, for (a) one pass stack-slide search strategies, and (b) two-pass hierarchical strategies (using the stack-slide algorithm in each pass). The results are presented for our fiducial classes of sources: (i) all-sky search for young ($\tau \geq 40$ yr), fast ($f \leq 1000$ Hz) pulsars, (ii) all-sky search for old ($\tau \geq 1000$ yr), slow ($f \leq 200$ Hz) pulsars, (iii) directed search for young, fast pulsars, and, (iv) directed search for old, slow pulsars. For a given computational power, we have determined the optimum observation time, and number of stacks as described in Secs. 3.4 and 3.6. Thus h_{th} is the expected sensitivity of the detector for an optimal stack-slide search, with 99% confidence. For comparison, the relative sensitivity achievable in a 10^7 second search with infinite computing power is shown in the bar on the right of each figure. When the relative sensitivity of a search exceeds this reference value, using finite computational resources, it indicates that the search uses more than 10^7 seconds of data.





Figure 3.6: Panel (a) represents the relative amplitude sensitivities $\Theta_{\text{rel}} = h_{3/\text{yr}}/h_{\text{th}}$ achievable with given computational resources, in three specialized searches: (i) A search for a new-born neutron star (whose direction is determined by observing an optical supernova) that is spinning down by gravitational wave emission via an active r -mode instability. We took $\tau \geq 1$ yr and $f \leq 200$ Hz. (ii) A search for pulsars in a region extending 0.004 steradians about the galactic core, with $\tau \geq 100$ yr and $f \leq 500$ Hz. (iii) A source in a binary orbit, e.g., Sco X-1. We assume the orbit is characterized by two orthogonal velocity parameters, known to within a total error of $17(\text{km/s})^2$; we further assume that the frequency $f \leq 500$ Hz experiences a random walk typical of Eddington-rate accretion. For each of these sources, panel (b) shows h_{th} for initial (upper lines), enhanced (middle lines), and advanced (lower lines) interferometers in LIGO, assuming 1 Tflops of computing power. Thus this is the characteristic amplitude of the weakest source that can be detected with 99% confidence using a two pass-hierarchical search strategy. We have also indicated the threshold sensitivity h_{th} for enhanced LIGO in narrow-band searches for LMXBs; the center frequency coincides with the expected wave frequency for Sco X-1.

Chapter 4 Gravitational waves and the cosmological equation of state

This chapter is adapted from the paper “Gravitational waves and the cosmological equation of state,” by Teviet Creighton [65].

Summary: As is well known, primordial gravitational waves may be amplified to detectable levels by parametric amplification during eras when their wavelengths are pushed outside the cosmological horizon; this can occur in both inflationary and “pre-big-bang” or “bounce” cosmologies. The spectrum of a gravitational wave background is expressed as a normalized spectral energy density $\Omega(\omega) \equiv (\omega/\rho_c)(d\rho_{\text{gw}}/d\omega)$, where ρ_c is the critical energy density that makes the universe spatially flat, and $d\rho_{\text{gw}}$ is the energy density of gravitational waves in a frequency band $d\omega$. The logarithmic slope of Ω is simply related to three properties of the early universe: (i) the gravitons’ mean initial quantum occupation number $N(n)$ ($\equiv 1/2$ for a vacuum state), where $n = \omega/a$ is the (invariant) conformal frequency of the mode and a is the cosmological scale factor, and (ii) & (iii) the parameter $\gamma \equiv p/\rho$ of the cosmological equation of state during the epoch when the waves left the horizon ($\gamma = \gamma_i$) and when they reentered ($\gamma = \gamma_f$). In the case of an inflationary cosmology, the spectral index is

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{\gamma_i + 1}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right),$$

and for bounce cosmologies it is

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{2\gamma_i}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right).$$

These expressions are compared against various more model-specific results given in the literature.

4.1 Introduction

The prospect of a detectable cosmological background of gravitational waves has opened up a new avenue for the investigation of fundamental physics and cosmology. Gravitational waves couple very weakly to matter, and are not blocked or thermalized during any cosmological epoch back to Planckian densities. The waves *do*, however, couple to large-scale cosmological spacetime curvature, and can therefore provide information about the evolution of this curvature, and hence about the fundamental physics of the matter fields driving this evolution.

A great deal of work has been done on this subject, deriving gravitational-wave spectra from various cosmological models, usually in hope of producing a model that would generate a detectable signature in planned gravitational-wave detectors. The most comprehensive of these is a recent paper by Gasperini [66], which presents a general prescription for constructing gravitational-wave spectra from fairly arbitrary cosmologies. The results in the present paper are consistent with Gasperini's results, although I derive them using a different formalism and consider a more generic cosmology.

The primary intent of this paper, however, is more the reverse of this: rather than deriving a spectrum from a particular cosmological model, this paper shows how generic properties of the early universe can be immediately deduced from an observation of a gravitational-wave background in any spectral band. To this end, I have chosen as my observable the spectral characteristic that is least susceptible to the tunings of particular cosmological models — namely, the spectral index — and have expressed it in terms of quantities that are most directly related to the underlying physics of the cosmological “fluid” — namely, the equations of state $p = \gamma\rho$. Furthermore, since the initial state of the gravitational waves (before amplification) is potentially one of the more interesting characteristics that might be deduced from observations, I have left it as a free parameter, rather than making the usual assumption of starting the waves in a quantum ground state.

4.1.1 Organization of this paper

In Sec. 4.2, I present a skeleton of the derivation of the main formulae of this paper. Much of the derivation has been done in one form or another in the published literature, and those familiar with the field will find nothing surprising, except perhaps the consideration of non-vacuum initial conditions in Sec. 4.2.2. Since the result is largely a generalization of the more model-specific formulae found in the literature, Sec. 4.3 compares this paper's formulae with those published previously. Sec. 4.4 presents some concluding remarks.

4.2 The cosmological and wave equations

I consider gravitational waves that are linear perturbations on the spatially flat FRW metric:

$$ds^2 = a^2(\eta)[-d\eta^2 + dx^2 + dy^2 + dz^2] + h_{\alpha\beta}dx^\alpha dx^\beta, \quad (4.1)$$

where the conformal time coordinate η is related to proper time t by $dt = a(\eta)d\eta$ (I have chosen units in which $c = 1$). The cosmology has a perfect fluid source that obeys the instantaneous equation of state $p = \gamma\rho$, where γ need not be constant but typically evolves slowly ($\gamma'/\gamma \ll a'/a$, where $' \equiv d/d\eta$). Causality considerations require that $\gamma \leq 1$; realistically, we can assume $\gamma < 1$, since not all of the energy in the universe will be in the maximally-stiff field. One can solve Einstein's

equations to zeroth order in h to obtain the exact solution

$$a(\eta) = a_0 \exp \left\{ \int_{\eta_0}^{\eta} \frac{d\eta_1}{a_0/a'_0 + \int_{\eta_0}^{\eta_1} d\eta_2 [1 + 3\gamma(\eta_2)]/2} \right\}, \quad (4.2)$$

which reduces to the usual power-law evolution $a(\eta) \sim \eta^{2/(1+3\gamma)}$ during epochs of (nearly) constant γ .

Following the notation in Eq. (2) of [67], I write the linear perturbations in the form:

$$h_{\alpha\beta} = \sum_{\mathbf{n},j} \frac{\mu_n(\eta)}{a(\eta)} U_{\mathbf{n}}(\mathbf{x}) e_{\alpha\beta}^{(j)}, \quad (4.3)$$

where $e_{\alpha\beta}^{(j)}$ is some basis of polarization tensors, $U_{\mathbf{n}}(\mathbf{x}) \propto e^{i\mathbf{n}\cdot\mathbf{x}}$ is a spatial harmonic function with a true (physical) wave number $\mathbf{k} = \mathbf{n}/a$ and frequency $\omega = n/a$ ($n = \|\mathbf{n}\|$), and $\mu_n(\eta)$ obeys the Schrödinger-like equation:

$$\mu_n'' + \left(n^2 - \frac{a''}{a} \right) \mu_n = 0. \quad (4.4)$$

Physically, the relative magnitudes of n^2 and the effective potential a''/a are related to whether the wave is inside or outside of the Hubble radius. Comparing \mathbf{k} to the Hubble radius $r_H = a/(da/dt)$, one has:

$$k^2 r_H^2 = \frac{n^2}{(a'/a)^2} = \frac{n^2}{a''/a} \times \frac{1-3\gamma}{2}. \quad (4.5)$$

Roughly speaking, a wave that is hitting the effective potential (a''/a increasing to meet n^2) is one that is exiting the Hubble radius, to order of magnitude. Similarly, a wave that is emerging from the effective potential is reentering the Hubble radius. The exception is for epochs when $\gamma \rightarrow 1/3$; see the remarks near the end of Sec. 4.2.1.

Cosmological amplification of waves occurs where initially oscillatory waves hit the effective potential during epochs of accelerated collapse ($a' < 0$, $\gamma > -1/3$) or accelerated expansion ($a' > 0$, $\gamma < -1/3$), when $|a''/a|$ is increasing, and then emerge into a post-inflationary universe ($a' > 0$, $\gamma > -1/3$) in which $|a''/a|$ is decreasing. A schematic of such an evolution is shown in Fig. 4.1. The periods when the wave initially hits the potential and when it finally emerges from it will be denoted by subscript i and f , respectively. The (complex) amplification factor of the emerging waves $\beta_n \equiv \mu_n(\eta_f)/\mu_n(\eta_i)$ is

$$\beta_n = \frac{1}{2i} e^{in(\eta_i - \eta_f)} \left[\frac{a_f}{a_i} \left(i + \frac{a'_f}{a_f n} \right) + \frac{a_i}{a_f} \left(i - \frac{a'_i}{a_i n} \right) + a_i a_f n \left(i - \frac{a'_i}{a_i n} \right) \left(i + \frac{a'_f}{a_f n} \right) \int_{\eta_i}^{\eta_f} \frac{d\eta}{a^2} \right]; \quad (4.6)$$

this is Eq. (11) of [67] with a sign error corrected, and reexpressed in the current notation.

If the evolution of $a(\eta)$ between a_i and a_f is monotonic, one can follow [67] and use the equation




Figure 4.1: Schematic of the evolution of the effective potential $|a''/a|$ (solid line) and Hubble scale $(a'/a)^2$ (dotted line) for inflation ($\eta < \eta_1$) followed by decelerating expansion ($\eta > \eta_1$). Also shown are the times η_i and η_f when a wave with squared conformal frequency n^2 (dashed line) hits and emerges from the potential barrier.

of state to reexpress the integral over η as an integral over a , eventually obtaining, for $a_f \gg a_i$:

$$|\beta_n| \approx \frac{1}{2} \frac{a_f}{a_i} \left| \left(i + \frac{a'_f}{a_f n} \right) \left(1 - \frac{2/3}{1 - \langle \gamma \rangle} \left[1 - \frac{i a_i n}{a'_i} \right] \right) \right|, \quad (4.7)$$

where $\langle \gamma \rangle$ is the average of γ from η_i to η_f with weighting factor $1/a^2$.

The derivation of Eq. (4.7) breaks down for the case of a bounce cosmology, where $a(\eta)$ is not monotonic. In this case, however, Eq. (4.6) is almost entirely dominated by the contribution of the integral $\int d\eta/a^2$ near the time of the bounce. Modeling the bounce as smooth with some finite concavity a''_0 around the minimum a_0 , we obtain:

$$|\beta_n| \approx \frac{1}{2} a_i a_f n \left| \left(i + \frac{a'_f}{a_f n} \right) \left(i - \frac{a'_i}{a_i n} \right) \right| \frac{1}{\sqrt{a''_0 a_0^3}}. \quad (4.8)$$

4.2.1 Dependence on n

The amplification factor β_n is at least implicitly dependent on n , since waves of different n will hit and emerge from the effective potential at different times η_i and η_f , and hence with different a_i , a_f . For periods of constant γ , we have $a \sim \eta^{2/(1+3\gamma)}$. The condition $n^2 = a''/a$ for hitting or emerging from the effective potential implies that $a_{i,f} \sim n^{-2/(1+3\gamma_{i,f})}$, and that $a'_{i,f}/(a_{i,f} n) \sim \sqrt{2/|1-3\gamma_{i,f}|} \sim \text{constant}$. If γ is changing gradually near $\eta_{i,f}$ these formulae remain true up to corrections of order $(\gamma'/\gamma)(a'/a)^{-1}$. Furthermore, the weighted average $\langle \gamma \rangle$ in Eq. (4.7) depends only very weakly on $\eta_{i,f}$, and hence on n , except for values of n^2 near the peak of the effective potential. These considerations, combined with Eqs. (4.7) and (4.8), give

$$|\beta_n| \sim n^{2/(1+3\gamma_i) - 2/(1+3\gamma_f)} \quad (4.9)$$




Figure 4.2: Schematic of the evolution of the effective potential and Hubble scale for inflation ($\eta < \eta_1$) followed by a radiation-dominated equation of state ($\eta > \eta_1$). Note that $\eta_f = \eta_1$ independent of n for a broad range of frequencies.

for inflationary cosmologies, and

$$|\beta_n| \sim n^{1+2/(1+3\gamma_i)+2/(1+3\gamma_f)} \quad (4.10)$$

for bounce cosmologies.

A special case is when $\gamma_f \approx 1/3$ (the equation of state for a relativistic or radiation-dominated fluid), for which the effective potential $a''/a \approx 0$. This is depicted in Fig. 4.2. If the transition to radiation dominance is rapid, then η_f , and hence a_f and a'_f , are independent of n for a wide range of n . However, waves that are well outside the Hubble radius at the moment of transition will have $a'_f/a_f \gg n$, as is clear from Eq. (4.5) and Fig. 4.2. So the appropriate terms in Eqs. (4.7) and (4.8) have behaviour

$$a_f \left| i + \frac{a'_f}{a_f n} \right| \sim \frac{1}{n}. \quad (4.11)$$

This is the same dependence as one would get by naively plugging $\gamma_f = 1/3$ into Eqs. (4.9) and (4.10), so these equations are correct even as $\gamma_f \rightarrow 1/3$. (This argument also applies for waves that *hit* the effective potential during radiation-dominated collapse.)

A similar situation occurs for waves that remain outside the Hubble radius throughout an intermediate period of $\gamma \approx 1/3$, when formally they have emerged from the effective potential and are oscillatory. This is depicted in Fig. 4.3. The effect of an intermediate period $\Delta\eta$ of zero effective potential is a term $\sim \sin(n\Delta\eta)/n$ in the final amplitude. If the waves remain well outside the Hubble radius throughout this era, then $n\Delta\eta \ll 1$, and this term is a constant with respect to n . Once again, Eqs. (4.9) and 4.10 remain unchanged. Both this and the preceding cases confirm one's physical intuition, that the final spectrum cannot resolve the details of sudden phase transitions that occur when the waves are much larger than the Hubble radius.

Eqs. (4.9), (4.10) also break down when $\gamma_{i,f} \approx -1/3$, but this condition will never arise. When




Figure 4.3: Schematic of the evolution of the effective potential and Hubble scale for inflation ($\eta < \eta_1$), followed by radiation dominance ($\eta_1 < \eta < \eta_2$), followed by generic decelerated expansion ($\eta > \eta_2$). Waves emerging from the potential barrier after η_2 will have been outside the Hubble radius throughout the period of radiation dominance.

$\gamma \approx -1/3$, one has $a \sim e^{\kappa\eta}$, so $a''/a \approx \text{constant}$, and waves will neither hit nor leave the effective potential.

4.2.2 The initial and final spectra

Most analyses of primordial gravitational waves assume that the initial state of the metric perturbations (during the Planck era or during pre-Big-Bang collapse) is a quantum-mechanical vacuum. While this is a reasonable assumption, one does not *know* the initial state in advance of observation. In fact, primordial gravitational waves are a potential means of observing the initial state of the cosmos. It is instructive, therefore, to leave the initial graviton spectrum as a free parameter.

I parameterize the initial spectrum using the mean quantum occupation number of graviton modes as a function of mode frequency, $N(n)$. In the semiclassical approach, the creation of gravitons can be treated as the classical amplification of the vacuum energy $\hbar\omega/2$ in each mode, so I normalize $N(n)$ to be $1/2$ for a mode in the ground state (no real gravitons), $3/2$ for the first excited state (one real graviton), and so on. This approach is valid provided the final state is classical ($|\beta_n| \gg 1$).

Cosmological perturbation spectra are normally expressed as a normalized energy density per logarithmic frequency interval, $\Omega \equiv (dE/dV d\ln\omega)/\rho_c$ (where ρ_c is the energy density required to make the universe spatially flat). There are $2\omega^2 d\omega dV$ modes in a frequency interval $d\omega$ in a volume dV , each with energy $\hbar\omega N(n)$ (where $n = \omega/a$), so the initial energy spectrum is:

$$\Omega_{\text{initial}}(n) = \frac{2n^4 N(n)}{a_{\text{initial}}^4 \rho_c}. \quad (4.12)$$

Note that a_{initial} is the (constant) scale factor of the universe at the time that one is evaluating the initial spectrum; it is not the same as a_i , which depends on n .

In considering the logarithmic slope of the spectrum I ignore overall amplitude factors, which are

highly model-specific. The energy spectrum at the present day is quadratic in $|\beta_n|$, so its logarithmic slope is:

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln \Omega_{\text{initial}}}{d \ln n} + 2 \frac{d \ln |\beta_n|}{d \ln n}, \quad (4.13)$$

where the evaluation is at $n = \omega/a_{\text{present}}$. Combining this with Eqs. (4.12), (4.9), and (4.10) gives the principal result of this paper:

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{\gamma_i + 1}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right) \quad (4.14)$$

for an inflationary cosmology, and:

$$\frac{d \ln \Omega}{d \ln \omega} = \frac{d \ln N}{d \ln n} + 2 \left(\frac{2\gamma_i}{\gamma_i + 1/3} \right) + 2 \left(\frac{\gamma_f - 1/3}{\gamma_f + 1/3} \right) \quad (4.15)$$

for a bounce-type cosmology.

4.3 Comparison with special cases

The formulae given above are consistent with results established previously in the literature. I first show this for the two “standard” cases of deSitter inflation and of cosmological rebound from a state of weakly coupled strings. I then compare with recent papers that consider more generic pre- and post-inflationary evolutions.

4.3.1 Standard inflation

The “standard” model for an inflationary cosmology consists of a period of deSitter expansion ($\gamma = -1$) due to the universe being in a false vacuum with nonzero energy density. In the post-inflationary era, the universe quickly reheats to a radiation-dominated equation of state ($\gamma = 1/3$), followed by cooling to a matter-dominated equation of state ($\gamma = 0$). If we assume the initial graviton state from the Planck era was a vacuum, then for waves that emerged into the radiation-dominated universe we have $N \equiv 1/2$, $\gamma_i = -1$, $\gamma_f = 1/3$, and

$$\frac{d \ln \Omega}{d \ln \omega} = 0. \quad (4.16)$$

This is, of course, the flat spectrum predicted by Harrison and Zel’dovich [68, 69].

The last two decades of the spectrum (10^{-16} – 10^{-18} Hz) consist of waves that emerged more recently, when the universe was matter-dominated ($\gamma_f = 0$). In this band the spectrum is tilted towards low frequencies, with a spectral index of -2 .

4.3.2 Standard string-motivated rebound cosmologies

The earliest, and still one of the most general, analysis of gravitational waves from a pre-Big-Bang universe was performed by Starobinski [70], who considered the present spectral index of gravitational waves in a cosmology with arbitrary asymptotic early- and late-time equations of state, and unknown evolution through the “bounce.” Eq. (2) of [70] can be shown to give the same spectral index as Eq. (4.15) with a flat initial graviton spectrum, given two notational differences. First, Starobinski uses a spectral energy density $\epsilon(\omega) \sim d\rho_{\text{gw}}/d\omega$, rather than $\Omega(\omega) \sim d\rho_{\text{gw}}/d\ln\omega$, so the spectral indices are related by $d\ln\Omega/d\ln\omega = d\ln\epsilon/d\ln\omega + 1$. Second, Starobinski’s equation of state parameter q is related to this paper’s γ by $q = 2/(3 + 3\gamma)$.

More recently, string theory has been used to provide a justification for the cosmological rebound, and to hypothesize a particular model of imploding universe that leads to a gravitational-wave spectrum increasing with frequency [7, 71]. These models are consistent with the current derivation provided one recognizes that the current formalism applies to the Einstein frame (where the metric represents intervals measured by physical rulers and clocks), rather than the string frame (where the metric geodesics describe the trajectories of weakly coupled strings). In the Einstein frame, a typical evolution consists of an epoch of decreasing lengthscale a and weakly coupled strings, followed by a period of strongly coupled strings that reverse the collapse, followed by relaxation of the string dilaton and an expanding radiation-dominated universe, eventually cooling to matter dominance. As usual, a vacuum initial state ($N \equiv 1/2$) is normally assumed. Most modes hit the effective potential during the epoch of weakly interacting strings, when the kinetic term of the dilaton dominates the energy background, giving a nearly maximally stiff equation of state ($\gamma_i \approx 1$). For waves that emerge during the radiation-dominated phase ($\gamma_f \approx 1/3$), Eq. (4.15) gives:

$$\frac{d\ln\Omega}{d\ln\omega} = 3. \quad (4.17)$$

Eq. (3.3) of [7] and Eq. (5.7) of [71] also give a spectral index of 3 for low frequencies. These papers also predict a logarithmic cutoff at high frequencies, $\Omega \sim \omega^3 \ln^2(\omega_s/\omega)$, where ω_s is the maximum frequency of waves that encountered the effective potential during the epoch of maximal stiffness (typically of order 10^{11} Hz, depending on the details of the strongly coupled string epoch). This effect occurs only for γ_i exactly equal to 1, which the current analysis doesn’t consider, and only affects the spectral index within a decade of ω_s in any case.

The low-frequency tail ($\omega < 10^{-16}$ Hz) of waves that emerge when $\gamma_f = 0$ will also be tilted towards high frequencies, but with a spectral index of 1.

4.3.3 Nonstandard equations of state

A recent paper [6] has explored the effects of an intermediate phase between the inflationary and radiation-dominated eras, in which the dominant cosmological fluid had an equation of state stiffer than radiation ($1/3 < \gamma_f < 1$). One consequence is that the background gravitational-wave spectrum would be tilted towards high frequencies for those wavelengths that emerged during this era. Assuming an initial vacuum ($N \equiv 1/2$) and deSitter inflation, Eq. (4.14) gives us:

$$\frac{d \ln \Omega}{d \ln \omega} = 2 \frac{\gamma_f - 1/3}{\gamma_f + 1/3}. \quad (4.18)$$

Eq. (3.32) of [6] gives the same dependence on γ .

Eq. (3.31) of [6] extends the analysis to the case of a maximally-stiff equation of state ($\gamma_f = 1$), which this paper does not consider. This gives a high-frequency logarithmic cutoff $\Omega \sim \omega \ln^2(\omega_1/\omega)$, where ω_1 is the highest frequency that encountered the amplifying potential. The cutoff does not affect the spectral index at frequencies more than a decade or so below ω_1 , which is likely in the MHz to GHz range (well above the pass-bands of proposed gravitational-wave detectors). This is analogous to the situation mentioned earlier, when $\gamma_i \rightarrow 1$ during collapse.

Schwarz [72] also considered non-radiation-dominated equations of state in the context of QCD phase transitions, which could cause temporary excursions from $\gamma_f = 1/3$, creating steps in the final gravitational-wave spectrum. Eq. (15) of [72] is identical to Eq. (4.18). Schwarz also points out that this spectral index is valid even if γ_f is changing rapidly, provided that the *total* change in γ_f is small ($\ll 1 + \gamma_f$).

Gasperini [66] has performed an even more general analysis of the primordial gravitational wave spectrum, considering generic evolution of $a(\eta)$ both in the early and late cosmological epochs. As usual, the initial state of the waves was taken to be a vacuum ($N \equiv 1/2$). Eqs. (4.14), (4.15) can then be written as:

$$\frac{d \ln \Omega}{d \ln \omega} = 4 + \frac{4}{1 + 3\gamma_i} - \frac{4}{1 + 3\gamma_f} \quad (4.19)$$

or

$$\frac{d \ln \Omega}{d \ln \omega} = 6 - \frac{4}{1 + 3\gamma_i} - \frac{4}{1 + 3\gamma_f} \quad (4.20)$$

for monotonic or bounce cosmologies, respectively. By comparison, Eq. (4.26) of [66] gives the spectral index as $4 - 2\nu_1 - 2\nu_{i+1}$, where $\nu = |\alpha - 1/2|$ and α is the exponent of the power law $a \sim \eta^\alpha$ during the epoch when the waves first strike the effective potential (for ν_{i+1}) and when they leave the effective potential (for ν_1). Gasperini's results are identical to Eqs. (4.19) and (4.20), provided one recognizes two things.

First, a phase of slowly varying α corresponds to a phase of slowly varying γ , with $\alpha = 2/(1+3\gamma)$. Inflationary cosmologies have $\gamma_i < -1/3$, hence $\alpha_{i+1} < 0$. Bounce cosmologies have $-1/3 < \gamma_i \leq 1$

and $\alpha_{i+1} \geq 1/2$. Either cosmology has $-1/3 < \gamma_f \leq 1$ and $\alpha_1 \geq 1/2$. (The condition $0 < \alpha < 1/2$ corresponds to the physically unrealistic case of $\gamma > 1$.)

Second, Gasperini assumes in Eq. (4.26) of [66] that all waves leave the effective potential at the same cosmological epoch, so that ν_1 is a constant across all frequencies (normally $\nu_1 = 1/2$ for radiation dominance). In fact, the spectral index (but not necessarily the normalization) given in Eq. (4.26) of [66] is valid for ν_1 varying across frequencies.

4.4 Conclusions

Eqs. (4.14) and (4.15) show how much can be determined about early cosmology from an observation of the spectral index of primordial gravitational waves. Although these formulae are quite simple, they still involve three independent parameters: γ_i , γ_f , and $N(n)$. To make definite statements about any one of them, one must make assumptions about the others. The usual procedure is to assume that we know $N(n)$ and γ_f , leaving γ_i to be deduced; however, it is quite possible that future advances in theory will fix both γ_i and γ_f , allowing gravitational-wave observations to probe directly the Planck-scale structure of the universe.

The formulae are also useful in showing which cosmological theories produce equivalent gravitational signatures. For instance, a monotonically inflating cosmology with $\gamma_{i_1} < -1/3$ will yield the same spectral index as a bounce-type cosmology with $\gamma_{i_2} = (\gamma_{i_1} + 1)/(3\gamma_{i_1} - 1)$ in the range $(-1/3, 1/3)$. Since these are physically valid ranges for γ_i in each case, it follows that a monotonic cosmology cannot be distinguished from a bounce cosmology given only the gravitational-wave spectral index.

The purpose of this paper was to present the simplest possible expressions relating a generic cosmological gravitational-wave signature to the physical processes that produced it, with a minimum of assumptions about the underlying cosmological model. I have deliberately chosen an approach of “maximum ignorance,” allowing, as much as possible, for the observations to dictate the cosmological parameters. Formulae such as these may prove valuable in an era when gravitational-wave observations begin to explore a new realm of physics, possibly revealing things beyond our most informed predictions.

Chapter 5 Gravity gradient noise from airborne sources

Summary: The test masses in gravitational-wave detectors will be sensitive not only to astrophysical gravitational waves, but also to the fluctuating Newtonian gravitational forces of moving masses in the ground and air around the detector. These effects are often referred to as gravity gradient noise. This paper considers the effects of gravity gradients from density perturbations in the atmosphere, and from massive airborne objects near the detector. These have been discussed previously by Saulson, who considered the effects of background acoustic pressure waves and of massive objects moving smoothly past the interferometer; the gravity gradients he predicted would be too small to be of serious concern even for advanced interferometric gravitational-wave detectors. In this paper I revisit these phenomena, considering transient atmospheric shocks, and estimating the effects of sound waves or objects colliding with the ground or buildings around the test masses. I also consider another source of atmospheric density fluctuations: temperature perturbations that are advected past the detector by the wind. I find that background acoustic noise and temperature fluctuations still produce gravity gradient noise that is below the noise floor even of advanced interferometric detectors, although temperature perturbations carried along non-laminar streamlines could produce noise that is within an order of magnitude of the projected noise floor at 10 Hz. A definitive study of this effect may require better models of the wind flow past a given instrument. I also find that transient shockwaves in the atmosphere could potentially produce large spurious signals, with signal-to-noise ratios in the hundreds in an advanced interferometric detector. These signals could be vetoed by means of acoustic sensors outside of the buildings. Massive wind-borne objects such as tumbleweeds could also produce gravity gradient signals with signal-to-noise ratios in the hundreds if they collide with the interferometer buildings, so it may be necessary to build fences preventing such objects from approaching within about 30m of the test masses.

5.1 Introduction

Interferometric detectors such as LIGO and VIRGO rely on exquisite sensitivity to the positions of hanging test masses in order to detect the perturbations of passing gravitational radiation. The sensitivity is so great that the measurements can also be affected by fluctuations in the local Newtonian gravitational field, which create tiny accelerations of the a mass. This noise source, known as gravity gradient noise or Newtonian gravitational noise, is caused by the near-field gravity of masses moving

near the interferometer, and is not to be confused with the far-field propagating gravitational waves that the instruments are intended to measure.

Gravity gradient noise has the potential to be quite insidious, since it cannot be shielded by improvements to the test-mass vibrational isolation. The only effective way to eliminate gravity gradient noise is to eliminate the moving masses that create the perturbing fields. Fortunately, though, the strongest perturbations to the local gravitational field are at frequencies well below the detectors' pass-bands. Since the proposed terrestrial interferometric detectors all have sensitivity cutoffs around 3 Hz or higher, we need only worry about noise sources that can perturb the local gravity field on timescales less than about 0.3 seconds. Most of the noise sources that I consider are motivated by the expected sensitivity of advanced LIGO interferometers, which were originally projected to have a low-frequency cutoff around 10 Hz, and instrumental noise of $S_h \sim 2 \times 10^{-45} \text{ Hz}^{-1}$ in the gravitational-wave signal output at that frequency [4]¹. Although changes in instrumentation technology will modify the ultimate sensitivity goals of LIGO, this "standard" advanced LIGO noise level is a good reference point when considering new noise sources. Also, as pointed out below, gravity gradient noise will make it difficult to push the detector noise much below this level at 10 Hz, regardless of improvements to the interferometers.

Saulson [73] was the first to estimate the effect of gravity gradient noise on terrestrial interferometric detectors, considering the effects of seismic waves passing through the earth and of sound waves in the air. In both cases he found the spectral density of noise in the interferometer path-length difference to be less than $10^{-39} \text{ m}^2 \text{ Hz}^{-1}$ around 10 Hz, corresponding to noise in the gravitational-wave signal at levels less than 10^{-46} Hz^{-1} for a 4 km interferometer. By comparison, this is significantly less than the noise floor of $\sim 2 \times 10^{-45} \text{ Hz}^{-1}$ that advanced LIGO interferometers expect to achieve at 10 Hz. More recently, a detailed analysis [8] has been made of seismic gravity gradient noise; this study indicated that seismic gravity gradient noise would be within a factor of 2 of the advanced LIGO noise floor at 10 Hz for most times, and could actually *exceed* this noise floor during seismically noisy times, making seismic gravity gradients a significant barrier to improvements in low-frequency sensitivity. It therefore seems prudent to revisit the issue of *atmospheric* gravity gradients as well.

In Sec. 5.2 I consider gravity gradients caused by atmospheric pressure perturbations—the same noise source considered by Saulson. Attention is paid, however, to the effects of the ground, and of buildings that reduce the pressure noise in the immediate vicinity of the interferometer test masses. However, I find that these tend only to weaken the gravity gradient noise in the pass-bands of interferometric detectors, reinforcing the conclusion that this noise source is not of any great concern.

A much stronger source of high-frequency density perturbations in the atmosphere is the presence

¹Specifically, I will be using the noise curve in Fig. 10 of [4], not Fig. 7, whose suspension thermal noise is a factor of 3 too small.

of temperature fluctuations, which are advected past a detector by the wind. In Sec. 5.3 I analyze this as a potential source of gravity gradient noise. However, I find that while small-scale temperature perturbations can produce high-frequency temperature fluctuations at any given point, they do not produce the same high-frequency fluctuations in the test mass position, since a given pocket of warm or cool air will affect the test mass gravitationally over the entire time that it is in the vicinity of the test mass, which is on the order of seconds. This produces a cutoff in the noise spectrum above a few tenths of a Hz. The presence of turbulent vortices around the interferometer buildings can increase the high-frequency component, but still probably not enough to show up in the gravitational-wave noise spectrum.

In Secs. 5.4 and 5.5 I turn away from sources of background noise to consider possible sources of transient gravity gradient signals that might be detected as spurious events in the gravitational-wave instruments. Sec. 5.4 extends the analysis in Sec. 5.2 to look at the effects of atmospheric shockwaves, such as might be generated by an explosion or supersonic aircraft. I find that sources such as these can indeed produce detectable signals that might be interpreted spuriously as gravitational-wave events. However, such signals would easily be vetoed using acoustic monitors outside the interferometer buildings.

Sec. 5.5 analyzes the gravity gradients produced by individual objects, such as wind-borne debris, moving around outside the interferometer buildings. Saulson considered this effect for the case of objects moving with fairly uniform velocity, but typically, in order to produce significant signal above 3 Hz, an object's motion must be changing on timescales of less than 0.3 seconds. In particular, I find that objects *colliding* with the interferometer buildings produce much stronger signals than objects simply passing by the buildings. As an example, tumbleweeds at the Hanford LIGO facility will be a steady source of spurious signals in advanced detectors if they are allowed to collide with the end stations. Preventing such signals requires shielding a region of a few tens of meters around the end station, screening any wind-borne debris that masses more than a few hundred grammes.

Sec. 5.6 presents some concluding remarks, including recommendations to the gravitational-wave experimental groups and possible directions for further research.

5.2 Atmospheric pressure waves

Pressure perturbations are the only type of atmospheric gravity gradient noise considered by Saulson [73]. The derivation in this section gives largely the same result as his.

Consider a plane pressure wave with frequency f propagating through a homogeneous airspace at some sound speed c . If the fractional pressure change $\delta p/p$ is small, it will induce an adiabatic density change $\delta\rho/\rho = \delta p/\gamma p$, where $\gamma \approx 1.4$ is the ratio of heat capacities at constant pressure and constant temperature for air at normal temperatures. The gravitational acceleration produced by

this wave in the direction of propagation \mathbf{e}_z is:

$$g_z(t) = \int \frac{G_z \delta \rho}{r^3} dV = \frac{2G\rho c}{\gamma p f} \delta p(t + 1/4f), \quad (5.1)$$

where $\rho \approx 1.3 \text{ kg m}^{-3}$ and $p \approx 10^5 \text{ N m}^{-2}$ are the ambient air density and pressure, and $\delta p(t)$ is the pressure perturbation measured at the same point as the acceleration is being measured. By symmetry, there is no acceleration transverse to the wave.

Now consider sound waves in the vicinity of the interferometer. First, since the interferometer is only sensitive to motions of the test mass parallel to the arms, the gravitational acceleration is reduced by a factor $\cos \theta$, where θ is the angle between the propagation direction and the interferometer arm.

Second, the interferometer test mass is inside a building, which can in principle be used to suppress noise within a distance r_{\min} of the test mass. Roughly speaking, this results in a high-frequency cutoff factor $C(2\pi f r_{\min}/c)$, where the function $C(x)$ depends on the precise shape of the building, the manner in which it reflects sound waves, and many other factors, but is normally close to 1 for $x \lesssim 1$. For instance, if one simply removes from the volume integral in Eq. (5.1) a cylinder with length and diameter both $2r_{\min}$ aligned with the z -axis, then $C(x) \sim 1$ for $x \lesssim 1$, but oscillates with an amplitude of ~ 0.3 for $x \gtrsim 1$. This function is shown in Fig. 5.1. The constant-amplitude oscillations of $C(x)$ for large x reflect the assumption that the sound wave has a coherence length much longer than the building size, so the field between the two ends of the excluded cylinder is generated almost entirely by the first half-wavelength beyond each cap. Realistically, the actual high-frequency behaviour of $C(x)$ will depend on how the sound waves bend and scatter around the building; however, this should not change the order of magnitude of $C(x)$. The behaviour shown in Fig. 5.1 is therefore probably a good estimate of the true cutoff function for $x \lesssim 1$, and a reasonable order of magnitude estimate for $x \gtrsim 1$. For the LIGO end stations, r_{\min} is of order 5 metres, giving $x \sim f/(10\text{Hz})$; the factor will not be too far off for the frequencies of greatest interest. More precise estimates would depend on the specific architectural details of a particular facility.

Third, the interferometer is on the ground, not in homogeneous empty space. For simplicity I assume that the waves are almost entirely reflected off the ground; I will justify this assumption in Sec. 5.2.1. In this case the gravity gradient in directions parallel to the ground contributed by the reflected wavefront is the same as if the wavefront were extended below ground, while the pressure perturbations measured by detectors near the ground (much less than a wavelength) will be doubled. The acceleration experienced by an interferometer test mass is therefore:

$$g_z(t) = \frac{G\rho c}{\gamma p f} \cos(\theta) C(2\pi f r_{\min}/c) \delta p(t + 1/4f). \quad (5.2)$$

The gravitational wave signal $h(t)$ in the interferometer is related to the acceleration of one of the




Figure 5.1: Plot of the factor $C(x)$ indicating the reduction in sonic gravity gradient noise due to setting the density perturbations to zero within a cylinder of diameter and length $2r_{\min}$ centered on the test mass (and aligned with the wave), versus $x = 2\pi r_{\min}/\lambda$ where λ is the wavelength. The oscillatory behaviour above $x = 1$ results from the ends of the cylinder coming in and out of phase with each other. A more accurate model of the scattering of sound waves off an interferometer end station building would modify the behaviour of $C(x)$ for $x > 1$, but probably not change its magnitude significantly.

test masses by $\ddot{h}(t) = g(t)/L$, or in frequency space $\bar{h}(f) = (2\pi f)^{-2}\bar{g}(f)/L$, where L is the length of the interferometer arm. Thus:

$$\bar{h}(f) = \frac{G\rho c}{4\pi^2\gamma p L f^3} \cos(\theta) C(2\pi f r_{\min}/c) i\tilde{\delta p}(f). \quad (5.3)$$

Assuming stationary noise, the one-sided spectral density $S_h(|f|)$ is given by $\langle \bar{h}(f)\bar{h}(f')^* \rangle = S_h(|f|)\delta(f-f')$, where $\langle \dots \rangle$ denotes an expectation over all random phases of all plane wave modes contributing to the noise, and $*$ denotes complex conjugation. Taking mode amplitudes and directions to be uncorrelated, this gives:

$$S_h(|f|) = \left[\frac{G\rho c}{4\pi^2\gamma L f^3} C(2\pi f r_{\min}/c) \right]^2 \langle \cos^2 \theta \rangle \frac{S_p(|f|)}{p^2}, \quad (5.4)$$

where $S_p(|f|)$ is the acoustic noise spectral density measured outside the building in the vicinity of a particular test mass. Since the two test masses in an arm are many wavelengths apart, their noise will be uncorrelated, and will thus add in noise power. Similarly, the noise from the two arms will add in power. (Actually this is a bit of an overestimate, since the noise in the motion of the test masses at the corner station will be somewhat correlated.) Noting that $\langle \cos^2 \theta \rangle = 1/3$, one finds that the total noise in the gravitational wave signal is:

$$S_h(|f|) = \left(\frac{G\rho c}{4\pi^2\gamma L} \right)^2 \frac{1}{3f^6 p^2} \sum_{i=1}^4 C(2\pi f r_{\min}^{(i)}/c)^2 S_p^{(i)}(|f|), \quad (5.5)$$

where i denotes a particular test mass in the interferometer, $r_{\min}^{(i)}$ is the dead air radius about the i th test mass, and $S_p^{(i)}(|f|)$ is the acoustic noise spectrum measured outside the building enclosing that test mass.

Infrasound noise spectra should be taken at the actual interferometer sites, but one can make estimates based on typical terrestrial atmospheric noise. An empirical study [74] collected 256 power spectra of 1–16 Hz infrasound data from a rural forest 50km from New York City over a period of months, and found that the average noise spectrum $S_p(f)$ was relatively flat at 6–16 Hz, though with widely varying amplitude: 25% of the spectra had noise under $\sim 100\text{nbar}^2/\text{Hz}$, 50% under $\sim 300\text{nbar}^2/\text{Hz}$, and 75% under $\sim 1000\text{nbar}^2/\text{Hz}$. I use Eq. (5.5) to compute the corresponding noise in the LIGO detector. Assuming that the end masses (with $r_{\min} \sim 5\text{m}$) dominate the contribution to the gravity gradient noise and contribute equally, the noise in the gravitational-wave signal around 10 Hz is:

$$S_h(|f|) \sim (6 \times 10^{-48} \text{Hz}^{-1}) \left(\frac{f}{10\text{Hz}} \right)^{-6} \left(\frac{S_p(|f|)}{1000\text{nbar}^2/\text{Hz}} \right) C(f/10\text{Hz}). \quad (5.6)$$

The results are plotted in Fig. 5.2, using infrasound power spectra read off of Fig. 3 of [74].

Even the third-quartile power spectrum is between two and three orders of magnitude below




Figure 5.2: Plot of dimensionless strain noise $\sqrt{fS_h(f)}$ versus frequency f for infrasonic atmospheric gravity gradients. The solid curve is the projected noise floor for advanced LIGO detectors; the dotted curves are the first, second, and third quartiles of noise produced by gravity gradients from ambient pressure waves. Data for the infrasonic noise power are taken from Fig. 3 of [74]. Clearly the pressure waves would not contribute significantly to LIGO noise in any frequency range.

the expected noise floor of $2 \times 10^{-45} \text{Hz}^{-1}$ at 10 Hz projected for advanced LIGO interferometers. Thus ambient infrasound is probably a negligible effect for determining the noise floor for most interferometric gravitational-wave detectors. Nonetheless the issue cannot be completely resolved without infrasonic noise data from the actual interferometer sites.

5.2.1 Ground absorption

In the derivation above I treated the sound waves as being reflected off the ground. Now consider what happens if a sound wave is absorbed by the ground. The energy flux in a travelling compression wave is $C^{3/2} \rho^{-1/2} \langle (\delta x/x)^2 \rangle$, where C is the compression modulus of the medium ($C_{\text{air}} = \gamma p_{\text{air}}$), and $\langle (\delta x/x)^2 \rangle$ is the average squared dimensionless compression factor over a wave cycle. The gravity gradient noise induced by such a wave goes as $c^2 \langle (\delta \rho)^2 \rangle$, where $\langle (\delta \rho)^2 \rangle \sim \rho \langle (\delta x/x)^2 \rangle$, and $c = \sqrt{C/\rho}$ is the wave speed in that medium. Therefore, if a sound wave is completely absorbed by the ground, the resulting ground motions will produce gravity gradient noise contributions in the ratio:

$$\frac{S_h^{(\text{ground})}}{S_h^{(\text{air})}} = \left(\frac{\rho_{\text{ground}}}{\rho_{\text{air}}} \right)^{3/2} \left(\frac{C_{\text{ground}}}{C_{\text{air}}} \right)^{-1/2}. \quad (5.7)$$

Now for $C_{\text{air}} = \gamma p_{\text{air}} = 1.4 \times 10^5 \text{N m}^{-2}$, $C_{\text{ground}} \sim 3 \times 10^8 \text{N m}^{-2}$, $\rho_{\text{air}} = 1.3 \text{kg m}^{-3}$, $\rho_{\text{ground}} \sim 1.8 \times 10^3 \text{kg m}^{-3}$, the ratio turns out to be of order 10^3 ; that is, if sound waves were completely absorbed by the ground, the resulting ground vibrations would produce gravity gradient noise levels about 1000 times greater than the atmospheric gravity gradients. However, it was shown in [8] that seismic gravity gradients are only just large enough to worry about. So the only way that atmospheric gravity gradients can be larger or of the same order as seismic gravity gradients is if the waves are mostly reflected off of the ground. This was one of the assumptions used in deriving Eq. (5.2).

5.3 Atmospheric temperature perturbations

The largest small-scale atmospheric density perturbations are caused not by pressure waves but by temperature perturbations. As heat is transported up through a convective atmospheric layer, convective turbulence mixes pockets of warm and cool air to form temperature perturbations on all lengthscales down to a few millimetres. On the timescales of interest (less than a second) these perturbations are effectively “frozen” into the airmass, while pressure differences disperse rapidly in the form of sound waves. Perturbations in the air density $\rho \propto p/T$ are therefore caused predominantly by the temperature perturbations, which are typically several orders of magnitude larger than the pressure perturbations. Although they are frozen into the airmass, these temperature perturbations can cause rapid time-varying density fluctuations $\delta \rho = -\rho \delta T/T$ as the wind carries them past a




Figure 5.3: Schematic of a pocket of air with temperature perturbation ΔT over a lengthscale l , being advected past a test mass k along a streamline S . The pocket has an instantaneous velocity $\mathbf{v}(t)$ and position $\mathbf{r}(t)$ relative to the test mass, and $x(t)$ is the projection of $\mathbf{r}(t)$ onto the axis of the interferometer arm. $\{S\}$ denotes a reference plane intersecting orthogonally with all streamlines S .

point in space. This is the primary source of “seeing” noise that affects optical astronomy.

The appendix to this chapter gives a rigorous mathematical derivation of the gravity gradient noise spectrum due to these temperature perturbations. This section gives a qualitative derivation that reproduces the final result to order of magnitude.

The gravity gradient signal at some frequency f is caused by pockets of warm or cool air with some lengthscale l being advected past the interferometer test mass at a speed v , where $l \sim v/2\pi f$. Consider a single such pocket of air with a temperature perturbation δT away from the ambient temperature T . The gravitational acceleration produced in the instrument as a function of time t is $g_x(t) = G\rho l^3(\delta T/T)x(t)r^{-3}(t)$, where ρ is the ambient air density, $r(t)$ is the distance of the air pocket from the test mass as it is blown past, and $x(t)$ is this distance projected onto the axis of the interferometer arm. This geometry is sketched in Fig. 5.3.

Now in general, the noise power spectral density in any quantity a due to a background of independent, uncorrelated events is $S_a(|f|) = (2/\Delta t)|\tilde{a}(f)|^2$, where $\tilde{a}(f)$ is the Fourier spectrum from a single event and Δt is the spacing between events. Assuming uncorrelated pockets of air, then independent pockets of air arrive along any given streamline at intervals $\Delta t \sim l/v$, and streamlines separated by more than l add noise incoherently (i.e., add linearly in power). This gives the following noise spectrum:

$$S_g(|f|) \sim \frac{2l}{v} \int_{\{S\}} \frac{dA}{l^2} \left(\frac{G\rho}{T} \right)^2 \delta T^2(l) |\tilde{G}_S(f)|^2, \quad (5.8)$$

where $\int_{\{S\}}$ denotes an integral over a plane crossing all streamlines S , and $\tilde{G}_S(f)$ is the Fourier transform of the function $G(t) = x(t)/r^3(t)$ taken along a given streamline. The quantity $\delta T^2(l)$ is the average squared temperature difference between points a distance l apart. Turbulent mixing theory, as well as actual micrometeorological measurements, predict a power-law behaviour for small separations: $\delta T^2(l) \sim c_T^2 l^p$, where p is typically $2/3$. This applies for horizontal separations l up to of order 50 times the height of a given air pocket above the ground [75]. For streamlines more than a metre or so above ground level, then, this behaviour for δT^2 should be good out to distances $l \sim 50\text{m}$, corresponding to frequencies $\gtrsim 0.2\text{Hz}(v/10\text{m s}^{-1})$.

Eq. (5.8) gives the gravity gradient noise on a given test mass in the interferometer. Denoting the test masses by the index $k = 1 \dots 4$, and assuming that each test mass contributes independently to the noise in the gravitational-wave signal h , one has $S_h(|f|) = (2\pi f)^{-4} L^{-2} \sum_k S_g(|f|)$. Combining this with Eq. (5.8) and the relation $l \sim v/2\pi f$ yields:

$$S_h(|f|) \sim 2 \left(\frac{G\rho}{LT} \right)^2 c_T^2 (2\pi f)^{-(p+7)} v^{p+4} \sum_k \int_{\{S\}} dA |\tilde{G}_{S,k}(f)|^2. \quad (5.9)$$

The more rigorous analysis in the appendix (Sec. 5.7) gives an expression with roughly the same form, but covers the factors of order unity, and also accounts for the fact that wind speed can vary along a streamline and between streamlines. The more accurate formula is:

$$S_h(|f|) = 2\pi^2 \left(\frac{G\rho}{LT} \right)^2 c_T^2 (2\pi f)^{-(p+7)} \sin(p\pi/2) \Gamma(p+2) \sum_k \int_{\{S\}} \tilde{F}_{S,k}(f)^* \tilde{G}_{S,k}(f) w dA, \quad (5.10)$$

where w is the wind speed of the streamline as it crosses the plane of integration, and $\tilde{F}_{S,k}(f)$ and $\tilde{G}_{S,k}(f)$ are Fourier transforms of functions $F_{S,k}(t)$ and $G_{S,k}(t)$ describing the motion of a point along a streamline S past a test mass k , of the form:

$$F(t) = \frac{x(t)}{r(t)^3} v(t)^{p+3}, \quad (5.11)$$

$$G(t) = \frac{x(t)}{r(t)^3}. \quad (5.12)$$

It is worth noting that the frequency structure of $S_h(|f|)$ depends on the time behaviour of the functions $x(t)$ and $r(t)$ describing the position of a point on a streamline relative to a test mass, which can be some distance away. The minimum distance r_{\min} from the test mass to the passing air is thus an additional important scale in the problem: if $x(t)$ and $r(t)$ change significantly only on timescales $\sim r_{\min}/v$, then the noise spectrum will be cut off at frequencies $\gtrsim v/2\pi r_{\min}$.

By comparison, the *temperature* noise spectrum $S_T(|f|)$ measured at a point depends only on the local properties of the atmosphere at that point. Applying similar order-of-magnitude arguments, one can write $S_T(|f|) \sim (2l/v) c_T^2 l^p |\tilde{H}(f)|^2$, where $H(t) \sim 1$ for $|t| \lesssim l/v$ and 0 otherwise. At high




Figure 5.4: Schematic of a uniform airflow streamline passing an interferometer corner station. θ is the angle between the airflow and the axis of the interferometer arm, r_0 is the distance from the test mass to the streamline at closest approach, $r_0 \cos \psi$ is the projection of this distance onto the ground, and $\mathbf{v} = \text{constant}$ is the velocity of the airflow.

frequencies $\gtrsim 0.2\text{Hz}(v/10\text{m s}^{-1})$ the system involves only one lengthscale l , giving the spectrum a power-law dependence:

$$S_T(|f|) \sim 2c_T^2 v^p (2\pi f)^{-(p+1)}. \quad (5.13)$$

This is the same, to order of magnitude, as the exact result given in Eq. (5.39).

5.3.1 Uniform airflow

The gravity gradient noise is easy to compute from Eq. (5.10) for the case of uniform airflow parallel to the ground with some constant velocity \mathbf{v} . Placing the reference plane orthogonal to \mathbf{v} and passing through the test mass, the equations of motion for a given streamline past the test mass become quite simple: $v(t) = w = v$, $x(t) = vt \cos \theta + r_0 \sin \theta \cos \psi$, $r(t) = (r_0^2 + v^2 t^2)^{1/2}$, where r_0 is the distance from the test mass to the nearest point on the streamline, and $r_0 \cos \psi$ is the projection of this distance onto the ground. The geometry of the airflow is shown in Fig. 5.4. It is easy to show that:

$$\tilde{G}(f) = -\frac{4\pi f}{v^2} [i \cos \theta K_0(2\pi f r_0/v) + \sin \theta \cos \psi K_1(2\pi f r_0/v)], \quad (5.14)$$

$$\tilde{F}(f) = v^{p+3} \tilde{G}(f). \quad (5.15)$$

I perform the integral over the above-ground half of the reference plane, out from some radius r_{\min} that is roughly the closest distance that the outside air can approach the test mass. This gives a

noise contribution from a single test mass equal to:

$$S_h(|f|) = 8\pi^2 \cos(\pi[p-1]/2)\Gamma(p+2) \left(\frac{G\rho r_{\min}}{TL}\right)^2 c_T^2 \left(\frac{v}{2\pi f}\right)^p (2\pi f)^{-5} \\ \times \left\{ \cos^2 \theta [K_0^2(x) - K_1^2(x)] + \frac{1}{2} \sin^2 \theta [K_1^2(x) - K_0(x)K_2(x)] \right\}, \quad (5.16)$$

where $x = 2\pi f r_{\min}/v$.

For typical values $v \sim 10\text{m/s}$ and $r_{\min} \sim 5\text{m}$, at frequencies above 10 Hz or so, one has $x \sim 30$ or more, well into the exponentially damped regime of the Bessel functions. Even for gale-force winds of 30m/s or so, the argument x of the Bessel functions will still be of order 10 or more. The asymptotic expansions of K_0 , K_1 , and K_2 give $K_0^2(x) - K_1^2(x) \sim K_1^2(x) - K_0(x)K_2(x) \sim \pi e^{-2x}/(2x^2)$. Also, I note that the total noise will be dominated by the contribution from the two end stations, which have smaller r_{\min} than the corner station. So the total noise in the interferometer for uniform airflow is:

$$S_h(|f|) = (2\pi)^3 \cos(\pi[p-1]/2)\Gamma(p+2) \left(\frac{G\rho}{TL}\right)^2 c_T^2 \left(\frac{v}{2\pi f}\right)^p (2\pi f)^{-5} e^{-4\pi f r_{\min}/v}. \quad (5.17)$$

I consider ‘‘typical’’ values of $p = 2/3$, $\rho = 1.3\text{kg m}^{-3}$, $T = 300\text{K}$, $L = 4000\text{m}$, and $r_{\min} = 5\text{m}$. The parameters v and c_T^2 can vary on a minute-by-minute basis, and should really be measured at the site of a given interferometer. However, $c_T^2 \sim 0.2\text{K}^2 \text{m}^{-2/3}$ is a typical daytime peak temperature fluctuation index [76], and $v \sim 20\text{m/s}$ might be typical of a fairly windy day. At frequencies around 10 Hz, this gives a noise spectrum of:

$$S_h(|f|) \sim (1.6 \times 10^{-40} \text{Hz}^{-1}) \left(\frac{f}{10\text{Hz}}\right)^{-23/3} 10^{-14f/(10\text{Hz})}. \quad (5.18)$$

The two dotted lines in Fig. 5.7 show the gravity gradient noise spectra computed from Eq. (5.18) for wind speeds of 10m/s and 30m/s.

Clearly the exponential cutoff makes this a negligible source of noise for LIGO or similar detectors. Physically this cutoff arises from the fact that the gravity from a particular temperature perturbation passing near the end station will affect the test mass coherently over the second or so that it takes to travel the width of the end station. Thus, even though the temperature noise spectrum has a high-frequency power law tail (reflecting the fact that temperature perturbations exist on all lengthscales), the gravity gradient signal will have this much sharper exponentially cut-off tail.

5.3.2 Potential flow near the end station

As described above, uniform airflow is not likely to produce much atmospheric gravity gradient noise in the pass-band of interferometric detectors, since the shortest timescale over which the gradients




Figure 5.5: Schematic of the airflow streamlines around the corner of an end station, assuming an incompressible vorticity-free flow. The dotted lines are the streamlines, r_{\min} is the minimum distance that the airflow can approach the test mass, R_m is the scale distance at which the flow velocity v approaches the free-streaming speed V , and X , Y , and R are coordinates used to describe the velocity field.

change is of order the wind crossing time of the interferometer buildings. However, if an air pocket could be made to accelerate over shorter timescales, it might produce a stronger gravity gradient signal at high frequencies.

One possibility is the acceleration of the air as it is forced up and around the wall of an end station: streamlines that approach the ground-wall corner of the end station can have curvature scales much shorter than the building size. Treating the flow as incompressible and vorticity-free, the resulting velocity field near the corner is (p. 27 of [77]):

$$v_x = -2AX, \quad v_z = 2AZ, \quad (5.19)$$

where $A = v/2R$ is a constant, and R , X , and Z are measured from the corner. This approximation is clearly only good near the corner, since it gives velocity increasing monotonically with radius; one would expect v to approach the free-streaming airspeed V at a distances R_m of order the building height. Fig. 5.5 shows a schematic of the airflow near the corner.

It is clear from Eq. (5.19) that, although the streamlines are sharply curved near the corner of the end station, the advection speed is smaller in direct proportion. The shortest timescale over which the motion can change is of order R_m/V for all streamlines. Thus the streamlines close to the corner

will contribute no more high-frequency noise than the streamlines further out, at distances of order R_m , and the spectrum should not differ greatly from the one for uniform flow, Eqs. (5.17), (5.18).

5.3.3 Vortices

Perhaps the most serious contenders for high-frequency atmospheric gravity gradient noise are circulating vortices of air near the end stations. This is somewhat stretching the assumptions of the formalism I have established, since I had previously separated the effects of the homogeneous turbulence (which establishes the temperature perturbations on various lengthscales), and wind flow (which carries these perturbations past the instrument). However, the results in this section are only expected to be good to order of magnitude anyway, in the absence of detailed hydrodynamic analysis of airflow past a particular interferometer building. I therefore apply the above formalism to a simple model for a turbulent-like flow past an interferometer end station.

For simplicity, I specialize to airflow along the axis of an interferometer arm, since these give the largest gravitational accelerations. A simple model for the turbulent flow is to take a uniform flow and then add a cycloidal motion to it. This gives $x(t) \sim vt - R \sin(vt/R)$, $z(t) \sim r_0 - R \cos(vt/R)$, $r(t) = \sqrt{z^2(t) + x^2(t)}$, where R is the radius of the cycloidal motion, and $r_0 \sim r_{\min}$ is the distance from the test mass to the unperturbed streamline. I treat the speed v along the streamline as a constant. If R is also a constant, then one would expect the Fourier transform of $G(t) = x(t)/r^3(t)$ [Eq. (5.12)] to have a spike at frequency $f = v/2\pi R$, with a width of order $\sim v/r_0$. However, to give a somewhat more realistic behaviour for R , I treat it as growing from zero at the leading edge of the end station to some scale value R_0 over the half-length $\sim r_{\min}$ of the end station: $R(t) = R_0 \sqrt{vt/r_{\min}}$. The square-root dependence mimics the less-than-linear growth of the thickness of a boundary layer. Although quite crude, this model covers the essential features of a turbulent flow past the building: a uniform translation, accompanied by circulating motions over a range of radii with a scale set by R_0 , with both the uniform flow rate and the circulation speed set by the free-streaming airspeed v . A typical streamline of this type is shown in Fig. 5.6.

The Fourier transform of $G(t)$ is too complicated to perform analytically, but is simple enough to compute numerically using a fast Fourier transform. The resulting $\tilde{G}(f)$ has the usual exponential cutoff with frequency scale $v/2\pi r_{\min}$, as for a smooth streamline, but then rises to a second peak value of $\sim (4/v) \sqrt{R_0^3/r_{\min}^5}$ at a frequency $f \sim 0.06v/R_0$ and decreases from there as f^{-3} . This high-frequency tail is the power law that one would expect from the cusps on the bottom of the cycloid. Since I treat v as constant, $\tilde{F}(f) = v^{p+3} \tilde{G}(f)$. The cross-sectional area of streamlines that contribute significantly around this peak frequency is of order $\sim 2r_{\min}R_0$. Plugging these into




Figure 5.6: Schematic of a vortex-like streamlines across the top of an end station. The quantity r_{\min} is the minimum distance that the airflow can approach the test mass, and R_0 is the typical radius of circulation at this distance.

Eq. (5.10) one gets a noise spectral density of

$$S_{h(\max)} \sim 1.5 \times 10^6 \left(\frac{G\rho}{LT} \right)^2 c_T^2 R_{\max}^{35/3} r_{\min}^{-4} v^{-5} \quad (5.20)$$

at the peak frequency of $f \sim 0.06v/R_0$, where I have assumed $p = 2/3$.

For wind speeds around 10m/s, a cycloid radius $R_0 \sim 0.06m(v/10m/s)$ puts the noise peak at the 10 Hz seismic noise wall for advanced LIGO detectors. The atmospheric gravitational noise contribution from a single end station is then:

$$S_{h(\max)} \sim (1.3 \times 10^{-49} \text{Hz}^{-1}) \left(\frac{c_T^2}{0.2 \text{Km}^{-2/3}} \right) \left(\frac{v}{10 \text{m/s}} \right)^{20/3} \left(\frac{r_{\min}}{5 \text{m}} \right)^{-4}. \quad (5.21)$$

This is over five orders of magnitude below the expected advanced LIGO noise floor of $2 \times 10^{-45} \text{Hz}^{-1}$ at 10 Hz. Gale-force winds ($v \sim 30 \text{m/s}$) will bring this up to $2 \times 10^{-46} \text{Hz}^{-1}$, still an order of magnitude below the advanced LIGO noise curve. The dashed lines in Fig. 5.7 show the actual data from the numerical Fourier transforms for these two cases.

Since even the worst-case estimate is still an order of magnitude below the advanced LIGO noise floor (as well as the seismic gravity gradient noise floor in [8]), it seems unlikely that turbulent vortices will be sufficient to raise the atmospheric gravitational noise to significant levels, even given the approximations made in this analysis. However, to settle this matter definitively would require a much more sophisticated numerical analysis of the temperature perturbations and airflow past the buildings of a particular facility.




Figure 5.7: Plot of dimensionless noise amplitude $\sqrt{f S_h(f)}$ versus frequency f . The solid line is the projected noise floor for advanced LIGO detectors. The two dotted lines are the gravity gradient noise levels caused by temperature perturbations advected along smooth streamlines at 10m/s (left curve) and 30m/s (right curve). The dashed lines are the noise levels caused by temperature perturbations advected along cycloidal vortices, as described in Sec. 5.3.3, at 10m/s (left curve) and 30m/s (right curve); in each case the size scales R_0 of the vortices have been tuned to maximize the noise at 10 Hz. Even given these fine-tunings, the gravity gradient noise curve in the worst-case scenario is still nearly an order of magnitude (in amplitude) below the advanced LIGO sensitivity.

5.4 Shockwaves

Although atmospheric pressure waves are unlikely to be a significant source of gravity gradient noise in interferometric gravitational-wave detectors, the sudden pressure changes caused by atmospheric shockwaves could potentially produce detectable transient signals in the detector, if such shocks occur in the vicinity of the detectors. Shocks are specifically a matter of concern because they can produce significant pressure changes over timescales less than 0.1 s, corresponding to the lower end of the pass-bands of most interferometric detectors.

Consider a shock that produces a sudden jump in air pressure in the vicinity of one of the interferometer test masses: $\delta p(t) = \Delta p \Theta(t - t_0)$. It is a simple matter to take the Fourier transform and apply Eq. (5.3) to obtain:

$$\bar{h}(f) = \frac{G\rho c}{8\pi^3\gamma L f^4} \frac{\Delta p}{p} i e^{2\pi f t_0} \cos(\theta) C(2\pi f r_{\min}/c). \quad (5.22)$$

Here, θ is the angle between the interferometer arm and the normal to the shock front. If the shock has a finite rise time Δt , we can mimic this analytically by convolving in the time domain with a Gaussian of width $\sigma \sim \Delta t$. In the frequency domain this multiplies our amplitude by a Gaussian of width $\sigma \sim 1/\Delta t$, giving an exponential cutoff at frequencies above $1/\Delta t$. Typical shocks from, for instance, supersonic objects have rise times on the order of a few milliseconds [78], corresponding to cutoff frequencies of a few hundred Hz. However, one expects the dominant contribution of the signal to come from much lower frequencies, before the building-size cutoff factor C kicks in.

Shocks are transient phenomena that will produce *signals* in the detector, rather than raising the noise floor. What one would like to know is what signal-to-noise ratio the shock will produce. In general this depends on what filters one is using to search for signals, and how well these filters overlap with the signal produced by a shock. However, the signal from a shock is likely to overlap quite well with templates designed to search for generic impulsive phenomena; such templates are likely to be used in advanced interferometers as control over nonstationary instrumental noise improves. Thus it is reasonable to consider the signal-to-noise ratio ρ that a matched filter would give:

$$\rho^2 = \int_0^\infty \frac{4\bar{h}(f)\bar{h}^*(f)}{S_h(f)} df = \int_{-\infty}^\infty \frac{|2f\bar{h}(f)|^2}{fS_h(f)} d\ln f. \quad (5.23)$$

This shows, roughly speaking, that the relative magnitude of the dimensionless signal amplitude $2|f\bar{h}(f)|$ and the dimensionless noise amplitude $\sqrt{fS_h(f)}$ over a logarithmic frequency interval gives a good indication of the signal-to-noise ratio produced in the detector.

5.4.1 Sonic booms

Sonic booms caused by supersonic bodies are one example of atmospheric shocks that might affect interferometric gravitational-wave detectors. Direct shockwaves from a supersonic aircraft will typically hit the ground in a “carpet” about 15–20 km wide under the aircraft’s flight path. Outside this carpet, the temperature gradient near the ground will completely reflect the shockwave before it touches down. However, shockwaves will also reflect downward off of the temperature inversion in the stratosphere and thermosphere, forming secondary and higher-order “carpets” out to many hundreds of kilometres. [79]. The presence or absence of these higher-order waves can depend quite sensitively on conditions in the upper atmosphere.

A detailed prediction of these effects is beyond the scope of this paper. However, to give an indication of their potential seriousness, I will consider what would happen if a supersonic aircraft were actually to overfly the instrument at a height of several kilometres.

The characteristic profile of a sonic boom is a symmetric N-wave, consisting of a shock that increases the pressure by an amount Δp , followed by a smooth decrease in pressure of $2\Delta p$ over a time Δt , followed by a second rising shock Δp to restore the ambient pressure. According to Eq. (9.78) of [80], the strength of the shocks is:

$$\frac{\Delta p}{p} \approx \frac{2^{1/4}\gamma}{(\gamma+1)^{1/2}}(M^2-1)^{1/8}\kappa\delta l^{3/4}r^{-3/4}, \quad (5.24)$$

where γ is the adiabatic coefficient of air (1.4 at normal temperatures), M is the Mach number of the aircraft (its speed divided by the sound speed), κ is a dimensionless form factor that depends on the shape of the aircraft (typically around 1), l is the length of the aircraft, δ is the ratio of the aircraft’s typical thickness to its length, and r is the closest distance that the aircraft came to the point of measurement. Between the two shocks, the rate of change of pressure in a direction \mathbf{e}_x parallel to the line of flight is given, in Eq. (9.80) of [80], as:

$$\frac{1}{p} \frac{dp}{dx} \approx \frac{\gamma}{\gamma+1} \frac{(M^2-1)^{1/2}}{M^2} \frac{1}{r}. \quad (5.25)$$

The shock fronts move outward at the sound speed c in the direction orthogonal to their surface, while the entire cone travels in the \mathbf{e}_x direction along with the aircraft at a speed Mc . The total change of pressure between the two shocks is $2\Delta p$. From these facts and Eqs. (5.24) and (5.25), one can show that the time between the two shocks is:

$$\Delta t \approx 2^{5/4}(\gamma+1)^{1/2} \frac{M}{(M^2-1)^{3/8}} \kappa\delta \frac{l^{3/4}r^{1/4}}{c}. \quad (5.26)$$

On frequency scales higher than $1/\Delta t$ (typically a few Hz for a supersonic aircraft a few kilometres

away), the sonic boom looks like simple Heaviside shocks, giving $\delta\tilde{p}(f) \sim 1/f$. At lower frequencies, though, the entire N-wave looks like the derivative of a δ -distribution, giving $\delta\tilde{p}(f) \sim f$. Performing the Fourier transform analytically and plugging into Eq. (5.3), one obtains:

$$\tilde{h}(f) = \frac{G\rho c}{4\pi^3\gamma L} \frac{1}{f^4} \frac{\Delta p}{p} \cos(\theta) C(2\pi f r_{\min}/c) \left[\frac{\sin(\pi f \Delta t)}{\pi f \Delta t} - \cos(\pi f \Delta t) \right] e^{2\pi i f t_0}, \quad (5.27)$$

where t_0 is the time when the midpoint of the N-wave crosses the detector. As expected, the amplitude goes roughly as f^{-4} , except for frequencies less than $1/\Delta t$, where it goes as f^{-1} .

Now let us plug in some typical numbers. The numbers $G = 6.67 \times 10^{-11} \text{m}^3 \text{kg}^{-1} \text{s}^{-2}$, $\rho = 1.3 \text{kg m}^{-3}$, $\gamma = 1.4$, $c = 332 \text{m s}^{-1}$, and $L = 4000 \text{m}$ can be treated as constant. A supersonic jet aircraft might have a length of $l = 10 \text{m}$, a typical diameter of $\delta l = 2 \text{m}$, and be traveling at Mach $M = 1.5$ at a distance of $r = 10 \text{km}$ or so. Let $\cos\theta$ be 1 for an upper limit. Then $\Delta t \sim 0.2 \text{s}$, and for frequencies $f \gtrsim 10 \text{Hz}$ the dimensionless signal amplitude is:

$$2|f\tilde{h}(f)| \sim 1.4 \times 10^{-19} (M^2 - 1)^{1/8} C(2\pi f r_{\min}/c) \left(\frac{\delta}{0.2} \right) \left(\frac{l}{10\text{m}} \right)^{3/4} \left(\frac{r}{10\text{km}} \right)^{-3/4} \left(\frac{f}{10\text{Hz}} \right)^{-3}. \quad (5.28)$$

This is three orders of magnitude above the expected noise floor of $\sqrt{f S_h(f)} \sim 1.4 \times 10^{-22}$ at 10 Hz for advanced LIGO interferometers!

By contrast, consider a .30-calibre rifle bullet ($l \approx 0.025 \text{m}$, $\delta \approx 0.3$) passing at Mach 3 within 10m of an interferometer test mass. (This stretches the assumption of a plane-wave shock front at the test mass, but the order of magnitude should be correct.) The bullet produces a much stronger double shock, but with a time interval $\Delta t \sim 0.5 \text{ms}$, so $1/\Delta t \sim 2 \text{kHz}$. The low-frequency tail of this signal will have dimensionless amplitude:

$$2|f\tilde{h}(f)| \sim 1.8 \times 10^{-23} \frac{M^2(M^2 - 1)^{-5/8}}{2.5} C(2\pi f r_{\min}/c) \left(\frac{\delta}{0.3} \right)^3 \left(\frac{l}{0.025\text{m}} \right)^{9/4} \times \left(\frac{r}{10\text{m}} \right)^{-1/4} \left(\frac{f}{10\text{Hz}} \right)^{-1}. \quad (5.29)$$

This is nearly an order of magnitude *below* the dimensionless noise amplitude in advanced LIGO, and therefore too small to be of any serious concern.

Fig. 5.8 shows more complete gravity gradient signal spectra computed using Eq. (5.27) with the above parameters for a supersonic aircraft and rifle bullet. These are plotted along with the anticipated dimensionless noise amplitude for advanced LIGO detectors.

5.4.2 Vetoing shockwave signals

While atmospheric shockwaves are a potential source of spurious signals in gravitational-wave detectors, they are easy to veto using environmental sensors. One need simply place infrasound mi-




Figure 5.8: Plot comparing dimensionless signal strengths $2|f\bar{h}(f)|$ of gravity gradients from sonic booms with the dimensionless interferometer noise amplitude $\sqrt{fS_h(f)}$, as a function of frequency. The solid line is the projected noise amplitude for advanced LIGO detectors. The dotted line is the signal from a .30 calibre bullet travelling at Mach 3 within 10 m of the end station; the dashed line is the signal from a 10 m long aircraft travelling at Mach 1.5 at a distance of 10 km. The aircraft's sonic boom can produce a detectable gravity gradient signal.

crophones outside the buildings and test-mass vacuum enclosures. If these sensors detect a pressure change of more than a millibar over timescales of 50–100 milliseconds, then one might expect spurious signals with dimensionless amplitude of $\sim 10^{-22}$ in the 10–20 Hz frequency range. The stretch of data containing the potential spurion can then be discarded.

Alternatively, if the same shock profile is detected in an array of at least three sound sensors, then one can determine the direction of propagation of the shock and predict the actual induced test-mass motions. The spurious signal could then be subtracted out of the data stream. This is a much trickier procedure, and would only be necessary in the unlikely event that significant amounts of data were being corrupted.

In either case, it is clear that infrasound sensors will be important environmental monitors for advanced interferometric detectors.

5.5 High-speed objects

Another potential source of spurious signals in the interferometer is the gravity gradient caused by the motion of an individual massive object past the interferometer end station, or the collision of such an object with the end station. The latter is particularly serious, since the sudden deceleration of the object can produce a signal at high frequencies. The issue of human-generated gravity gradient noise has been addressed in [9], but there are other sources outside the facility that must be considered, such as stray bullets and wind-borne debris. In particular, the Hanford LIGO facility is plagued by tumbleweeds, which can produce non-negligible gravity gradient signals.

The general formula for the spurious gravitational-wave signal produced by a moving object is:

$$\tilde{h}(f) = \frac{GM}{4\pi^2 L f^2} \int_{-\infty}^{\infty} \frac{x(t)}{r^3(t)} e^{2\pi i f t} dt, \quad (5.30)$$

where M is the mass of the object, $r(t)$ is its distance from the test mass as a function of time, and $x(t)$ is its distance from the end mass in the direction parallel to the interferometer arm. For an object travelling parallel to the ground in a straight line at speed v , Eq. (5.30) becomes [as in Eq. (5.14)]:

$$\tilde{h}(f) = \frac{GM}{Lv^2 \pi f} \left[K_0 \left(\frac{2\pi f r_{\min}}{v} \right) \cos \theta - i K_1 \left(\frac{2\pi f r_{\min}}{v} \right) \sin \theta \cos \psi \right] e^{2\pi i f t_0}, \quad (5.31)$$

where θ is the angle between the line of motion and the interferometer arm, r_{\min} is the distance of closest approach between the object and the test mass, ψ is the angle projecting this distance onto the ground, and t_0 is the time of closest approach. K_0 and K_1 are modified Bessel functions of the second kind of order 0 and 1, respectively.

Under moderately windy conditions (wind speeds up to 15m/s or so), tumbleweeds at the Hanford

LIGO facility will bounce along the ground at 5–10m/s. In stronger winds, the tumbleweeds become airborne, with speeds approaching the wind speed; they can fly right over the LIGO buildings, or impact with considerable force. The same may be true of wind-borne debris at other interferometer facilities. However, the value of r_{\min} is usually at least 5m, so for frequencies above 10 Hz one has $2\pi fr_{\min}/v \gtrsim 10$ even for very strong winds ($v \sim 30\text{m/s}$). In this regime the Bessel functions are exponentially damped, so these objects will not produce significant gravity gradient signals simply by blowing past the instrument. A rifle bullet, on the other hand, might be moving around 1000m/s, putting us in the small-argument regime of the Bessel functions, where $K_0(x) \sim -\ln(x)$ and $K_1(x) \sim x^{-1}$. Taking the most dangerous geometry $\theta = \pi/2$, $\psi = 0$, and assuming a bullet mass of around 5 grammes, this gives a dimensionless signal amplitude near 10 Hz of:

$$2|f\tilde{h}(f)| \sim 1.6 \times 10^{-22} \left(\frac{M}{5\text{g}}\right) \left(\frac{1000\text{m/s}}{v}\right) \left(\frac{5\text{m}}{r_{\min}}\right) \left(\frac{10\text{Hz}}{f}\right). \quad (5.32)$$

This gives a signal-to-noise ratio of about 1 at 10 Hz, a bit below the detectable threshold. In fact, even if one fine-tunes the bullet speed v , the largest signal-to-noise ratio that one can get at 10 Hz is about 2, for a speed of around 250m/s. Since events with signals less than about 5 times the noise will probably be ignored in any case, one can conclude that objects flying past a test mass are not likely to be serious sources spurious events.

If an object does not pass smoothly by the interferometer but instead collides with an end station, the signal at ~ 10 Hz can be large even for slow-moving objects: it is the deceleration time, not the end-station-crossing time, that sets the frequency scale of the signal. Suppose an object collides end-on with the end station at a speed v , coming to a stop within a distance d at constant acceleration. Let $t = 0$ denote the time that the object comes to rest. The motion of the object is then given by:

$$r(t) = x(t) = \begin{cases} r_{\min} - d - vt & t \leq -2d/v \\ r_{\min} + (vt)^2/4d & -2d/v \leq t \leq 0 \\ r_{\min} & t \geq 0 \end{cases}. \quad (5.33)$$

The Fourier transform of $1/x^2(t)$ is tricky to do analytically, so I have relied on numerical fast Fourier transforms, and then made approximate analytic fits to the result. However, one can qualitatively predict the shape of the signal in frequency space. The function $1/x^2(t)$ starts out near zero and then slowly rises over a timescale r_{\min}/v to a value r_{\min}^2 , then quickly levels off at that value over a timescale d/v . So on frequency scales $\ll v/r_{\min}$, $1/x^2(t)$ looks like a step function, whose Fourier transform goes as $1/f$. On frequency scales $\geq v/r_{\min}$ but $\ll v/d$, one sees the deceleration as a cusp (discontinuous first derivative), giving a Fourier transform that goes as $1/f^2$. On frequency scales $\geq v/d$, the deceleration appears smooth, but the onset of deceleration is sudden, giving a

$1/f^3$ behaviour. The following gives a good fit to the numerical Fourier transform:

$$\left| \int_{-\infty}^{\infty} \frac{e^{2\pi i f t}}{x^2(t)} dt \right| \sim \frac{1}{v r_{\min}} \left[5.9 \left(\frac{f r_{\min}}{v} \right) + 14 \left(\frac{f r_{\min}}{v} \right)^2 + 59 \left(\frac{f^3 r_{\min}^2 d}{v^3} \right) \right]^{-1}. \quad (5.34)$$

The two breakpoints separating the three branches are at frequencies $f_1 = 0.4v/r_{\min}$ and $f_2 = 0.24v/d$. More precisely, since the deceleration occurs over a well-defined time $2d/v$, the third branch of the Fourier transform is oscillatory with nodes every $v/2d$ in frequency space; the functional fit in Eq. (5.34) is an envelope containing these oscillations.

For a 5 g bullet striking an end station at 1000m/s, the signal at 10 Hz is dominated by the low-frequency tail regardless of stopping distance. The dimensionless signal amplitude is then:

$$2|f\tilde{h}(f)| \sim 3 \times 10^{-22} \left(\frac{M}{5\text{g}} \right) \left(\frac{10\text{Hz}}{f} \right)^2 \left(\frac{5\text{m}}{r_{\min}} \right)^2. \quad (5.35)$$

The signal amplitude is about the same as for a passing bullet, although the dependence on f (and r_{\min} and v) is different. For a tumbleweed or other wind-borne object, by contrast, the sudden deceleration can create significant high-frequency noise. A typical tumbleweed at Hanford has a mass of $M = 0.1\text{kg}$ and a diameter of 0.4m, and can be compressed by about half that amount ($d = 0.2\text{m}$). Larger weeds can be twice as large in diameter, putting their masses in the 0.5–1kg range [81]. For moderate to high wind speeds ($v = 10\text{--}30\text{m/s}$), the signal at frequencies above 10 Hz is dominated by the second branch of Eq. (5.34), giving a dimensionless signal amplitude of:

$$2|f\tilde{h}(f)| \sim 5 \times 10^{-21} \left(\frac{M}{1\text{kg}} \right) \left(\frac{v}{10\text{m/s}} \right) \left(\frac{10\text{Hz}}{f} \right)^3 \left(\frac{5\text{m}}{r_{\min}} \right)^3. \quad (5.36)$$

Thus even a “typical” 0.1kg weed at these speeds will produce a signal-to-noise ratio of around 4 at 10 Hz in advanced LIGO interferometers, which is in danger of being interpreted as a real gravitational-wave event. In a more extreme case, a 1kg tumbleweed blown airborne by a strong 30m/s wind will produce a signal 100 times higher than the noise at 10 Hz, which is easily detectable.

Fig. 5.9 shows the signal spectra for the objects discussed above, plotted against the dimensionless noise amplitude expected in advanced LIGO detectors. Since tumbleweeds are a potential source of spurious detectable events, one should consider ways to reduce the tumbleweed gravity gradient noise. Fortunately, the signal goes as r_{\min}^{-3} , so a simple fence preventing the weeds from approaching the end station should be sufficient. A fence 30 m out from the end station will reduce the signal-to-noise ratio to 1 for tumbleweed masses up to 1kg and speeds up to 30m/s, reducing the risk of spurious events.




Figure 5.9: Plot comparing dimensionless signal strengths $2|f\tilde{h}(f)|$ of gravity gradients from airborne objects with the dimensionless interferometer noise amplitude $\sqrt{fS_h(f)}$, as a function of frequency. The solid line is the projected noise amplitude for advanced LIGO detectors. The short-dashed line is the signal from a .30 calibre bullet passing next to the interferometer end station at 1000 m/s, while the dotted line is the signal from that same bullet colliding with the end station. The long-dashed line is the signal from a large 1 kg tumbleweed passing next to the end station at 30 m/s, while the dot-dashed line is the signal from that same tumbleweed colliding with the end station, assuming that it compresses by 20 cm on impact. The tumbleweed mass and speed are both near the upper end of the expected range, representing the greatest danger of spurious signals in the detector; evidently these tumbleweed signals *can* be detected quite easily by advanced LIGO detectors.

5.6 Conclusions

This paper has studied two sources of background gravity gradient noise, from infrasonic atmospheric pressure waves and from wind-advected temperature perturbations, in order to determine whether they constitute a limiting noise floor for interferometric gravitational-wave detectors—in particular, for the “advanced” LIGO detectors projected in [4]. The paper also analyzed two sources of gravity gradient signals, from transient atmospheric shockwaves and from massive airborne bodies, to determine whether they would constitute detectable spurious events in these interferometers. The following summarizes the results and suggests possible further work that may need to be done.

Current estimates suggest that infrasonic pressure waves will not be a significant source of gravity gradient noise, being over two orders of magnitude below the advanced LIGO noise floor at 10 Hz. Nonetheless, these estimates are not based on actual noise measurements at an interferometer site, so infrasound measurements at these sites are recommended to confirm them. Further empirical studies might also analyze the specific effects of building shapes and of infrasound coherence lengths on the noise spectrum, although these refinements would likely only serve to reduce noise estimates above 15 Hz or so.

Wind-advected temperature perturbations, although the dominant source of atmospheric density fluctuations, do not produce significant high-frequency gravity gradient noise, due to the long times that any particular pocket of warm or cool air spends in the vicinity of an interferometer test mass. A possible exception is when the airflow forms vortices around the interferometer buildings, since this will produce a noise spectrum peaked around the typical vortex circulation frequencies near the test mass. The current crude analysis of these effects suggests that the noise is still an order of magnitude below the advanced LIGO noise floor at 10 Hz even in the worst-case scenarios, but the model could be improved significantly. Numerical models of the airflow and of temperature perturbations near an interferometer building may be required to settle this issue definitively.

Gravity gradients from atmospheric shockwaves are potentially serious sources of spurious signals in interferometric gravitational-wave detectors. For instance, the sonic boom from a supersonic aircraft overflying an advanced LIGO detector could produce signal-to-noise ratios of several hundred. Although such overflights are expected to be rare or nonexistent, they point out the potential seriousness of shocks from weaker or more distant sources, even if the signals are several orders of magnitude smaller. It is therefore *strongly* recommended that advanced interferometric detectors include infrasonic detectors as environmental monitors. Such sensors could easily be employed to veto spurious atmospheric gravity-gradient events.

Gravity gradients from wind-borne objects such as tumbleweeds are another possible source of spurious events in gravitational-wave detectors, if these objects are allowed to collide with the buildings housing the interferometers. Fences or other structures should be used to keep these

objects at least 30 metres from the test masses, in order to eliminate the risk of spurious signals.

Obviously the number of things that can affect interferometric detectors through gravity gradient forces is immense; I have considered here only the few sources that I considered the most worrisome. I encourage other researchers to consider the implications of this often-neglected effect.

5.7 Appendix: The temperature noise spectrum

This appendix presents a more rigorous mathematical derivation of Eq. (5.10) used in Sec. 5.3.

Consider a time-varying field of temperature perturbations $\delta T(\mathbf{r}, t)$ about some average temperature T . This produces a gravitational perturbation $g_x(t) = \int dV G \rho x r^{-3} (\delta T/T)$, where ρ is the average air density, and the x -axis is along the interferometer arm. The spectral density of gravity gradient noise $S_g(|f|)$ is given by twice the Fourier transform of the gravity autocorrelation $C_g(\tau) = \langle g(t)g(t + \tau) \rangle$. Thus:

$$S_g(|f|) = 2 \left(\frac{G\rho}{T} \right)^2 \int_{-\infty}^{\infty} d\tau \int dV \int dV' \frac{xx'}{r^3(r')^3} \langle \delta T(\mathbf{r}, t) \delta T(\mathbf{r}', t + \tau) \rangle e^{2\pi i f \tau}. \quad (5.37)$$

The temperature noise measured at a point \mathbf{r}_0 , on the other hand, is given simply by:

$$S_T(|f|) = \int_{-\infty}^{\infty} d\tau \langle \delta T(\mathbf{r}_0, t) \delta T(\mathbf{r}_0, t + \tau) \rangle e^{2\pi i f \tau}. \quad (5.38)$$

On sufficiently small scales, the temperature perturbations in the Earth's turbulent boundary layer can be treated as homogeneous and isotropic. The expected squared temperature difference between two points is then a function only of their separation: $\langle [\delta T(\mathbf{r}) - \delta T(\mathbf{r} + \Delta\mathbf{r})]^2 \rangle = D_T(\|\Delta\mathbf{r}\|)$. The function $D_T(\Delta r)$ is called the *temperature structure function* of the atmosphere, and for small Δr reduces to a power law $D_T(r) = c_T^2 \Delta r^p$. If a wind with speed v blows these perturbations past a measuring station, the temperature autocorrelation is $\langle \delta T(\mathbf{r}_0, t) \delta T(\mathbf{r}_0, t + \tau) \rangle = \sigma_T^2 - (1/2)c_T^2 (v\tau)^p$ for small τ , where σ_T^2 is the mean squared temperature fluctuation. This results in a high-frequency power law tail:

$$S_T(|f|) = c_T^2 v^p (2\pi f)^{-(p+1)} \Gamma(p+1) \sin(p\pi/2). \quad (5.39)$$

Turbulent mixing theory, as well as micrometeorological measurements of $S_T(|f|)$, show that the value of p is normally $2/3$, characteristic of a type of turbulence known as Kolmogorov turbulence. See, for example, [75] for discussion of this type of turbulent mixing, also [76] and references therein.

We are interested in $S_g(|f|)$, which is somewhat trickier to calculate than $S_T(|f|)$, since it involves a correlation between points separated in space as well as time. However, chaotic turbulence will almost certainly destroy high-frequency correlations between widely separated points, so the high-frequency behaviour of $S_g(|f|)$ will come from correlations between nearby points. That is, the

high-frequency support of $\langle \delta T(\mathbf{r}, t) \delta T(\mathbf{r}', t + \tau) \rangle$ will come from those points \mathbf{r}' at time $t + \tau$ whose fluid elements were near \mathbf{r} at time t .

Consider two fluid elements moving along paths S and S' passing through \mathbf{r} and \mathbf{r}' , respectively. This is shown schematically in Fig. 5.10. Let r_0 be the distance from \mathbf{r} to the nearest point on S' , and let τ_0 be the time it would take a pocket of air at \mathbf{r}' to be carried to this point on S' . In order for these points \mathbf{r} and \mathbf{r}' to contribute to the high-frequency component of the spectrum, the distance r_0 must be fairly small, of the order v/f where v is the wind speed past \mathbf{r} . I treat the streamlines as relatively straight over these scales, in which case the temperature perturbations at (\mathbf{r}, t) and $(\mathbf{r}', t + \tau)$ correspond to physical pockets of air separated by a distance $\sqrt{r_0^2 + v(\mathbf{r})^2(\tau - \tau_0)^2}$, for τ near τ_0 . Assuming that the τ -dependence of the correlation function is due entirely to this advection, the correlation function can be written explicitly as:

$$\langle \delta T(\mathbf{r}, t) \delta T(\mathbf{r}', t + \tau) \rangle = \sigma_T^2 - \frac{c_T^2}{2} [r_0^2 + v^2(\tau - \tau_0)^2]^{p/2}, \quad (5.40)$$

where typically $p \approx 2/3$.

This term contains the entire τ -dependence of the gravity perturbation in Eq. (5.37), so the first integral I do is the Fourier integral over τ . This integral has the form $\int_{-\infty}^{\infty} (\beta^2 + x^2)^{\nu-1/2} e^{iax} dx = 2\pi^{-1/2} (2\beta/a)^\nu \cos(\pi\nu) \Gamma(\nu + 1/2) K_{-\nu}(a\beta)$. Formally the integral diverges for $\nu \geq 1/2$, but the closed-form expression remains approximately correct for large a provided the integrand is cut off smoothly for large $x \gg 1/a$. Physically this corresponds to the fact that a smooth, large-scale cutoff of the temperature correlations in Eq. (5.40) will not affect the high-frequency component of the temperature noise. For horizontal winds near the ground the spatial correlation function is cut off on horizontal distance scales of ~ 50 times the fluid elements' altitude z above ground, giving a low-frequency cutoff around $\sim 0.02v/z$ (Fig. 1.A4 of [75]). Typically this will be below the relevant frequency range for interferometric detectors; I ignore it to obtain pessimistic (upper-limit) noise estimates. The high-frequency noise tail is then:

$$S_g(|f|) = \left(\frac{G\rho}{T} \right)^2 \int dV \int dV' \left[\frac{\mathbf{x}\mathbf{x}'}{r^3(r')^3} c_T^2 v^p (\sqrt{2\pi}f)^{-(p+1)} a(p) \left(\frac{2\pi f r_0}{v} \right)^{(p+1)/2} K_{(p+1)/2}(2\pi f r_0/v) e^{2\pi i f \tau_0} \right], \quad (5.41)$$

where $a(p) = -2\pi^{-1/2} \cos(\pi[p+1]/2) \Gamma([p+2]/2) \approx 0.873$ for $p = 2/3$.

Next is the integral over dV' . The exponential decay of the Bessel function $K_{(p+1)/2}$ restricts the support of this integral to values $r_0 \lesssim v/2\pi f$, representing the fact that high-frequency fluctuations can only arise from the rapid change of the spatial correlation function over small lengthscales. This range in r_0 defines a narrow bundle of streamlines S' about the streamline S passing through \mathbf{r} . I

assume that the size of this bundle is less than the distance from the test mass to the bundle, so that the values of x' and r' on a given S' can be replaced with the nearby values on S . Now for points near \mathbf{r} the volume element dV' can be written in terms of the new parameters r_0 and τ_0 as $dV' = 2\pi r_0 dr_0 v(\mathbf{r})d\tau_0$. Since the airflow, being very subsonic, is nearly incompressible, the volume element retains this form for all points along the bundle: if $v(\mathbf{r}')$ decreases below $v(\mathbf{r})$, for instance, the length element $v(\mathbf{r}')d\tau_0$ will decrease, but the cross-section of the bundle (i.e., the relevant range of $2\pi r_0 dr_0$) will increase to compensate. Plugging in this volume element and integrating over r_0 , one obtains:

$$S_g(|f|) = \left(\frac{G\rho}{T}\right)^2 \int dV \int v d\tau_0 \left[\frac{xx'}{r^3(r')^3} c_T^2 v^p (\pi f)^{-(p+1)} a(p) \times 2\pi \Gamma([p+3]/2) \left(\frac{v}{2\pi f}\right)^2 e^{2\pi i f \tau_0} \right]. \quad (5.42)$$

Let t' be a new time coordinate denoting the time it takes for an air pocket to reach \mathbf{r}' from some fixed reference plane that crosses all streamlines orthogonally, and t be the corresponding coordinate for the point \mathbf{r} . Then τ_0 is just $t' - t$, and the volume element dV can be written as $w dt dA$, where dA is a cross-sectional area element on the reference plane, and w is the wind speed across that area element. Plugging this in, and ignoring the spatial separation between S and S' , one obtains:

$$S_g(|f|) = \left(\frac{G\rho}{T}\right)^2 c_T^2 (\pi f)^{-(p+3)} a(p) (\pi/2) \Gamma([p+3]/2) \times \int_{\{S\}} w dA \left(\int_{-\infty}^{\infty} \frac{x}{r^3} v^{p+3} e^{-2\pi i f t} dt \right) \left(\int_{-\infty}^{\infty} \frac{x'}{(r')^3} e^{2\pi i f t'} dt' \right), \quad (5.43)$$

where I have used $\int_{\{S\}}$ to denote an integral over the entire reference plane; i.e., over all streamlines S .

The noise in the gravitational-wave signal $h(t)$ due to the gradients at a given test mass is $S_h(|f|) = (2\pi f)^{-4} S_g(|f|)/L^2$, where L is the interferometer arm length. The noise at each test mass adds incoherently to the total signal. Combining these with Eq. (5.43) yields the result given in Eq. (5.10).




Figure 5.10: Schematic showing the parameters used to describe streamlines of air flowing past an interferometer test mass. $\{S\}$ denotes a plane intersecting all streamlines; S and S' are two such streamlines passing through the points \mathbf{r} and \mathbf{r}' , respectively, where the coordinate origin is centered on the test mass k . The distance between the streamlines, measured at \mathbf{r} , is r_0 . The positions of the points \mathbf{r} and \mathbf{r}' along the streamlines S and S' are parameterized by the times t and t' that it would take for a pocket of air to move from $\{S\}$ to \mathbf{r} or \mathbf{r}' ; τ_0 is the difference $t - t'$. $\mathbf{v}(\mathbf{r})$ is the wind velocity along S through the point \mathbf{r} ; \mathbf{w} is the wind velocity through the plane $\{S\}$.

Bibliography

- [1] K. S. Thorne, in *Three hundred years of gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987), Chap. 9, pp. 330–458.
- [2] B. Allen, *GRASP: a data analysis package for gravitational wave detection*, 1.7.0 ed., Department of Physics, P.O. Box 413, University of Wisconsin, Milwaukee WI 53201, USA, 1998, <http://www.lsc-group.phys.uwm.edu/ballen/grasp-distribution/>.
- [3] E. Gustafson, D. Shoemaker, K. Strain, and R. Weiss, Technical Report No. T990080-00-D, LIGO Scientific Collaboration, (unpublished), <http://www.ligo.caltech.edu/ligo2/>.
- [4] A. Abramovici *et al.*, *Science* **256**, 325 (1992).
- [5] B. Allen, in *Proceedings of the Les Houches School on Astrophysical Sources of Gravitational Waves*, edited by J. A. Marck and J. P. Lasota (Cambridge University Press, Cambridge, 1996), gr-qc/9604033.
- [6] M. Giovannini, *Phys. Rev. D* **58**, 083504 (1998), hep-th/9806329.
- [7] R. Brustein, M. Gasperini, M. Giovannini, and G. Veneziano, *Phys. Lett. B* **361**, 45 (1995), hep-th/9507017.
- [8] S. A. Hughes and K. S. Thorne, *Phys. Rev. D.* **58**, 122002 (1998).
- [9] K. S. Thorne and C. J. Winstein, gr-qc/9810016 (unpublished).
- [10] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, *Phys. Rev. D.* **57**, 2101 (1998), gr-qc/9702050.
- [11] C. Bradaschia *et al.*, *Nucl. Inst. A.* **289**, 518 (1990).
- [12] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, *Class. Quant.* **13**, 575 (1996).
- [13] T. A. Apostolatos, *Phys. Rev. D.* **54**, 2421 (1996).
- [14] B. Owen, *Phys. Rev. D.* **53**, 6749 (1996).
- [15] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, *Phys. Rev. D.* **53**, 3033 (1996).
- [16] E. E. Flanagan, *Phys. Rev. D.* **48**, 2389 (1993).

- [17] K. Compton, Ph.D. thesis, University of Wales, Cardiff, 1996.
- [18] D. Nicholson *et al.*, *Physics Letters A* **218**, 175 (1996).
- [19] S. Bonazzola and E.ourgoulhon, *Astron. Astr.* **312**, 675 (1996).
- [20] S. Chandrasekhar, *Phys. Rev. Let.* **24**, 611 (1970).
- [21] J. L. Friedman and B. F. Schutz, *Astrophys. J.* **222**, 281 (1978).
- [22] R. V. Wagoner, *Astrophys. J.* **278**, 345 (1984).
- [23] M. Zimmermann and E. Szedenits, Jr., *Phys. Rev. D.* **20**, 351 (1979).
- [24] R. D. Blandford, unpublished, 1984.
- [25] B. F. Schutz, in *The Detection of Gravitational Waves*, edited by D. G. Blair (Cambridge University Press, Cambridge, 1991), Chap. 16, pp. 406–451.
- [26] J. C. Livas, Ph.D. thesis, Massachusetts Institute of Technology, 1987.
- [27] G. S. Jones, Ph.D. thesis, University of Whales, 1995.
- [28] T. M. Niebauer *et al.*, *Phys. Rev. D* **47**, 3106 (1993).
- [29] S. L. Shapiro and S. A. Teukolsky, *Black Holes, White Dwarfs and Neutron Stars* (Wiley, New York, 1983), Chap. 10.
- [30] A. G. Lyne, *Phil. Trans. R. Soc. Lond. A* **341**, 29 (1992).
- [31] R. Narayan and J. P. Ostriker, *Astrophys. J.* **352**, 222 (1990).
- [32] R. N. Manchester, *Phil. Trans. R. Soc. Lond. A* **341**, 3 (1992).
- [33] S. R. Kulkarni, *Phil. Trans. R. Soc. Lond. A* **341**, 77 (1992).
- [34] A. Wolszczan *et al.*, *Nature* **337**, 531 (1989).
- [35] S. B. Anderson, Ph.D. thesis, California Institute of Technology, Pasadena, California, 1993.
- [36] M. Zimmermann, *Nature* **271**, 524 (1978).
- [37] M. Zimmermann, *Phys. Rev. D.* **21**, 891 (1980).
- [38] D. V. Gal'tsov, V. P. Tsvetkov, and A. N. Tsurulev, *Sov. Phys. —JETP* **59**, 472 (1984).
- [39] J. Shaham, *Astrophys. J.* **214**, 251 (1977).
- [40] J. Middleditch, Ph.D. thesis, University of California, Berkeley, 1975.

- [41] M. Tinto, private communication, 1997.
- [42] P. R. Brady and T. Creighton, gr-qc/9812014 (unpublished).
- [43] K. C. B. New, G. Chanmugam, W. W. Johnson, and J. E. Tohline, *Astrophys. J.* **450**, 757 (1995).
- [44] A. Krolak, gr-qc/9803055 (unpublished).
- [45] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D.* **58**, 063001 (1998), gr-qc/9804014.
- [46] P. Jaranowski and A. Krolak, gr-qc/9809046 (unpublished).
- [47] M. A. Papa, B. F. Schutz, S. Frasca, and P. Astone, Detection of continuous gravitational wave signals: pattern tracking with the Hough transform, To appear in Proceedings of the LISA Symposium, 1998.
- [48] M. A. Papa, P. Astone, S. Frasca, and B. F. Schutz, Searching for continuous waves by line identification, To appear in the proceedings of Gravitational Wave Data Analysis Workshop, 1997.
- [49] M. A. Papa, private communication, 1998.
- [50] S. L. Shapiro, S. A. Teukolsky, and I. Wasserman, *Astrophys. J.* **272**, 702 (1984).
- [51] J. L. Friedman, J. R. Ipser, and L. Parker, *Nature* **312**, 255 (1984).
- [52] N. Andersson, *Astrophys. J.* **502**, 708 (1998), gr-qc/9706075.
- [53] J. L. Friedman and S. M. Morsink, *Astrophys. J.* **502**, 714 (1998), gr-qc/9706073.
- [54] L. Lindblom, B. J. Owen, and S. M. Morsink, *Phys. Rev. Lett.* **80**, 4843 (1998), gr-qc/9803053.
- [55] B. J. Owen *et al.*, *Phys. Rev. D.* **58**, 084020 (1998), gr-qc/9804044.
- [56] L. Bildsten, *Astrophys. J.* **501**, L89 (1998), astro-ph/9804325.
- [57] B. F. Schutz, Sources of radiation from neutron stars, talk given at Gravitational Wave Data Analysis Workshop, MIT, 1996, gr-qc/9802020.
- [58] S. J. Curran and D. R. Lorimer, *Mon. Not. R. Astron. Soc.* **276**, 347 (1995).
- [59] R. B. Tully, *Nearby Galaxy Catalog* (Cambridge University Press, Cambridge, 1988).
- [60] S. van den Bergh and R. D. McClure, *Astrophys. J.* **425**, 205 (1994).
- [61] R. Talbot, *Astrophys. J.* **205**, 535 (1976).

- [62] M. van der Klis, in *The Many Faces of Neutron Stars*, edited by A. Alpar, L. Buccheri, and J. van Paradijs (Kluwer, Dordrecht, 1998 (in press)).
- [63] A. P. Cowley and D. Crampton, *Ap. J. Letters* **201**, L65 (1975).
- [64] E. W. Gottlieb, E. L. Wright, and W. Liller, *Ap. J. Letters* **195**, L33 (1975).
- [65] T. Creighton, gr-qc/9907045 (unpublished).
- [66] M. Gasperini, in *New developments in string gravity and physics at the Plack energy scale*, edited by N. Sanchez (World Scientific, Singapore, 1996), hep-th/9607146.
- [67] L. P. Grishchuk and M. Solokhin, *Phys. Rev. D* **43**, 2566 (1991).
- [68] E. R. Harrison, *Phys. Rev. D* **1**, 2726 (1970).
- [69] Ya. B. Zel'dovich, *Mon. Not. Roy. Astron. Soc.* **160**, 1p (1972).
- [70] A. A. Starobinski, *JETP Lett.* **30**, 682 (1979).
- [71] M. Gasperini, in *Proceedings, 3rd Colloque Cosmologie*, edited by H. de Vega and N. Sanchez (World Scientific, Singapore, 1995), hep-th/9506140.
- [72] D. J. Schwarz, *Mod. Phys. Lett. A* **13**, 2771 (1998).
- [73] P. R. Saulson, *Phys. Rev. D.* **30**, 732 (1984).
- [74] E. S. Posmentier, *J. Geophys. Res.* **79**, 1755 (1974).
- [75] N. E. Busch, in *Workshop on Micrometeorology*, edited by D. A. Haugen (American Meteorological Society, Boston, 1972), Chap. 1, pp. 1-65.
- [76] C. E. Coulman, *Ann. Rev. Astron. Astrophys.* **23**, 19 (1985).
- [77] L. D. Landau and E. M. Lifshitz, *Fluid Mechanics* (Pergamon Press, New York, 1959).
- [78] W. D. Hayes, *Ann. Rev. Fluid Mech.* **3**, 269 (1971).
- [79] N. K. Balachandran, W. L. Donn, and D. H. Rind, *Science* **197**, 47 (1977).
- [80] G. B. Witham, *Linear and Nonlinear Waves* (Wiley, New York, 1974).
- [81] F. Raab, private communication, 1999.