Search for continuous gravitational waves from globular cluster NGC6544 using barycentric resampling

by

Santiago Caride

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in The University of Michigan
2015

Doctoral Committee:

Professor Keith Riles, Chair
Professor Hui Deng
Professor August Evrard
Professor David W. Gerdes
Professor Mateusz Ruszkowski

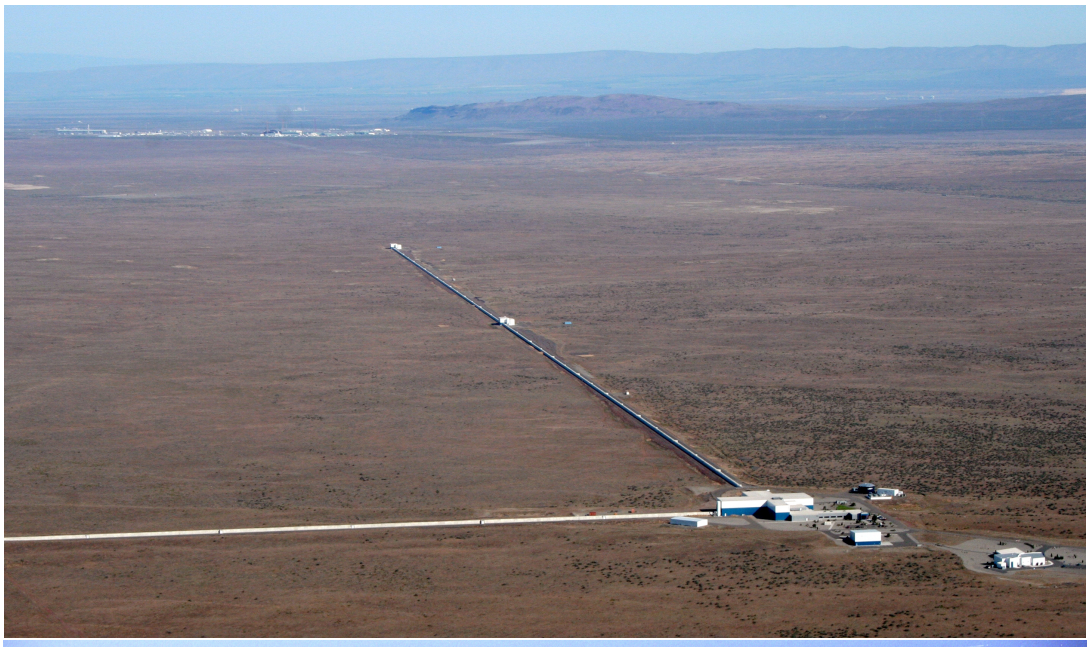

ACKNOWLEDGEMENTS

First of all, my Ph.D. adviser, Prof. Keith Riles; this project would not have been possible without your dedicated support. Your firm intuitive grasp of the fundamentals and keen attentive eye towards the details have both been immensely important in the development and refinement of this work and of this student. I am very grateful to have had the opportunity to work and study with you during my time at the University of Michigan. It's been a pleasure and I look forward to many more Michigan dinners at future LVCs!

Thanks are also due to others within the LIGO and VIRGO collaborations. Prof. Nelson Christensen, at Carleton College, was my gateway into LIGO as an undergraduate. Thanks for introducing me to the field, and to what it meant to be a scientist. Dr. Greg Mendell, at the Hanford observatory, set much time aside for me while I worked on my first project as a graduate student, both at the University of Michigan and during my summer at the Hanford detector. Thanks for your patience and generosity. Dr. Reinhard Prix, at the Albert Einstein Institute, worked with me to backport code modifications into the LIGO libraries after my initial solo efforts. Thanks for taking the time to bring me into the fold. Doctors Teviet Creighton and Ian Jones, my review committee; you've both improved my research considerably with your attention to detail, and improved my own understanding of the material with insightful questions about and evaluations of my work. Thanks for taking an hour (and more!) out of every week and dedicating it to me and my search. Finally,

although you are far too many to list by name, some measure of thanks is due to each and every individual in the LIGO and VIRGO collaborations for making this effort to detect gravitational waves a reality, and collectively giving me the opportunity to work with you to achieve that (fingers crossed!) first detection.

Thanks also to the students, past and present, that I've worked with in my time here at Michigan. Vladimir, you were on your way out when I made my way in, but still provided help and an important first impression of what a successful graduate student looked like. Evan, you've been a great colleague and friend both here in Ann Arbor and since your graduation. I've still got your TV, and I'll see you next time we sail around the world! Grant, who blazed the thesis trail before me, and Jax, who's hot on my heels: you've both been great fun to work and spend time with, and I look forward to our paths crossing again. And finally, to Ansel and Orion, the new guard in our Michigan group, I hope to have been of some small service and look forward to seeing you make your own (continuous) waves.

I'd be remiss not to thank my family for all of the love and support along the way. To my father, Ariel, for the initial inspiration to pursue science, and for all the early Sunday mornings along the way. To my mother, Ines, for keeping me honest and keeping me organized. Home is still my favorite place to be. To my siblings: Cecilia, thanks for putting up with me for these last two years; they've been by far my favorites in Ann Arbor. To Javier, for all the baseball and soccer and talk (even the philosophy.) To Martin, thanks for proving that at least one of us can be relatively normal; we'd all be so much worse off without you. And to Lauri: how's your forecast? Thanks for being a pressure release valve and giving me a few hours each week to just have fun.

Thanks to the friends who make life so much better: Dan (and Jeni, and Ryleigh!):

thanks for keeping me grounded in the real world. You've shaped me more than you know. Doctors Dylan, Rahul, and Arun, Neer-future Dr. Dutta, and Master Sathish– you've all walked this trail alongside me, even if a few steps ahead or behind. I look forward to flying back to Rochester for poker and board games for all eternity. To the rest of the Dilbert list, past and present: thanks for reading my daily ramblings. I'll keep buying calendars as long as you guys like.

This dissertation bears LIGO Document number DCC-P1500164 [1].

[1]For further questions, please contact the author: carides@umich.edu

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

The year 1905 is often called Albert Einstein's golden year. Those twelve months saw him publish his Special Theory of Relativity, an explanation of the photoelectric effect, and a treatise on Brownian motion. A decade later came the General Theory of Relativity [1], revolutionizing gravitational physics and suggesting the existence of gravitational waves. One century later, the Laser Interferometer Gravitational-wave Observatory (LIGO) stands on the brink of the first direct detection of these elusive signals.

Einstein himself was unsure if the waves predicted by his theory were detectable at all; over forty years of debate passed before, at a conference in 1957, a solution was presented by Felix Pirani [2]. Rather than focusing on the generation of gravitational waves, Pirani [3] described the effect of a gravitational wave's presence on a set of freely-falling objects, and showed that they would experience genuine motions with respect to one another [2]. Thus, gravitational waves were in principle detectable.

The first efforts to make such a detection were pioneered by Joseph Weber, who was in attendance at the 1957 conference [2]. Starting in 1960, he monitored the minute vibrations of massive aluminum cylinders, which came to be known as Weber bars [4]. Passing gravitational waves, he hoped, would set the bars into oscilla-

tion. Within the decade, Weber announced that his detectors were in fact detecting coincident signals, which he thought likely to be of astronomical origin [5]. Unfortunately, his results could not be replicated, and by the mid-1970s, the consensus was that gravitational waves were too weak to be detected by Weber's aluminum bar instruments [6].

The hunt, however, was not over. As early as 1962, the use of laser interferometers as gravitational wave directors had been proposed by Gertsenshtein and Pustovoit [2], and in 1972 Rainer Weiss essentially laid down the blueprint for what would become LIGO [7]. The interferometers were several orders of magnitude more sensitive than the aluminum bar detectors, and could be made even more sensitive by increasing the length of the arms. Hope within the gravitational wave community was renewed [6].

Gravitational-wave astronomy received another boost in 1975, when Hulse and Taylor reported their discovery of PSR1913+16 in radio waves [8]. This binary pulsar system followed a pattern of orbital decay precisely predicted by general relativity, and has continued to do so for decades [9] as shown in Figure 1.1. Hulse and Taylor received the 1993 Nobel for their discovery, and the gravitational wave community received a very strong piece of indirect evidence for the existence of their sought-after waves.

In 1989, a proposal to build two interferometric detectors with four-kilometer arms was submitted to the National Science Foundation. LIGO was funded in 1991, and two detectors were built within the decade; one in Hanford, Washington, and the other in Livingston, Louisiana. The two detectors reached their initial design sensitivities in 2005 [2]. This thesis centers on a search for continuous gravitational waves from neutron stars in the nearby globular cluster NGC6544, using data from these two detectors taken during LIGO's sixth science run (S6), which ran from July

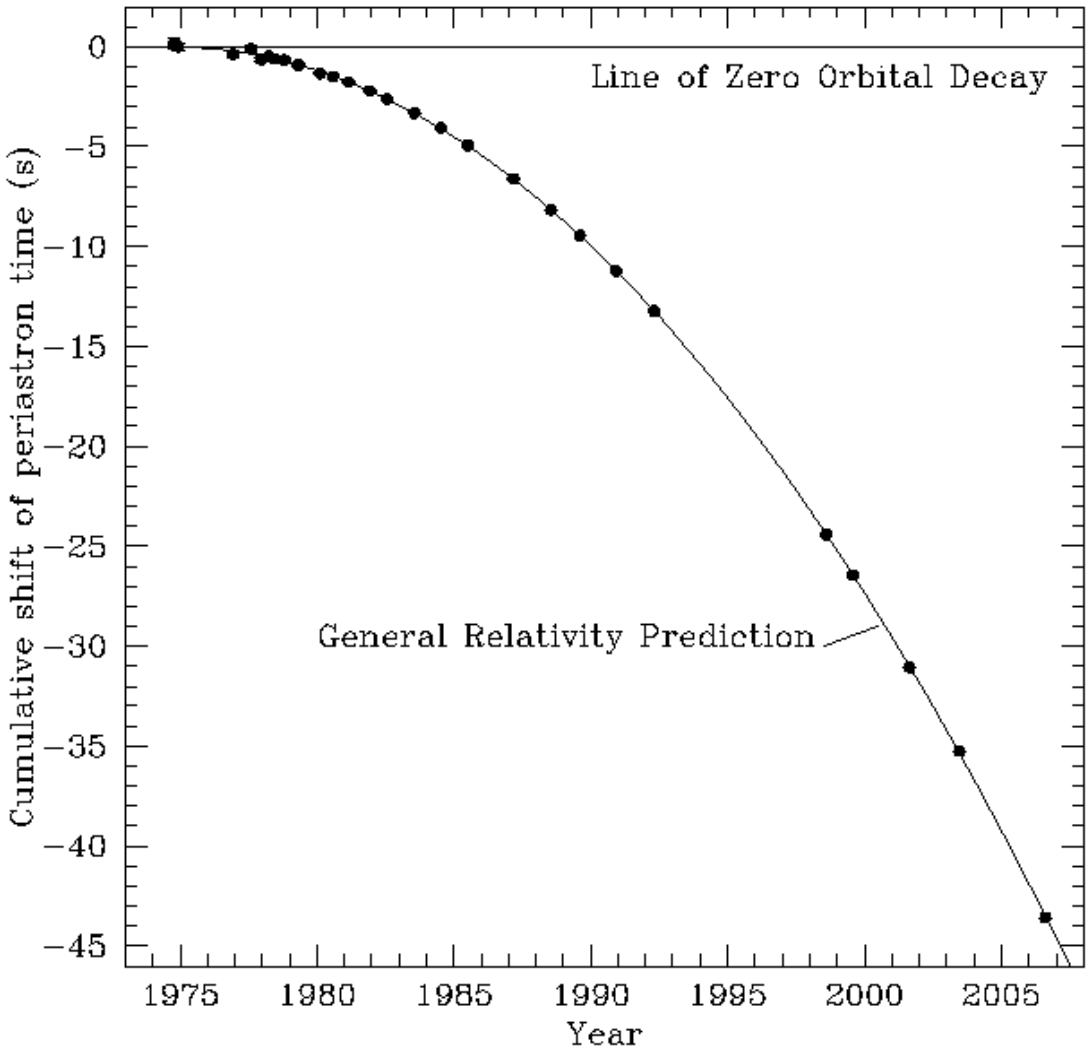

Figure 1.1: The orbital decay of the binary system PSR1913+16 has tracked general relativity's prediction since its discovery in 1975. This image is Figure 2 in [9].

2009 to October 2010.

Chapter 2 will describe the formalism of Einstein's General Relativity and the emergence of gravitational waves from the theory and briefly describe the major types of gravitational waveforms that LIGO scientists seek. Chapter 3 discusses how an interferometric detector responds to gravitational waves, and gives an overview of the specifications of the LIGO detectors. Chapter 4 describes in more detail continuous gravitational waves, the target of this thesis. The astrophysical sources and emission mechanisms of such waves are discussed, and the chapter concludes with a description of the globular cluster NGC6544 and the reasons it was selected as the target for this search. The fifth chapter describes a barycentric resampling technique which provides large computational gains, allowing for a search of a deeper parameter space with limited computational resources. Chapter 6 describes the methods and setup for the search conducted, and Chapter 7 is dedicated to the results of the search, including outlier followup and the setting of upper limits in the absence of a signal.

CHAPTER II

Gravitational Waves

Detecting gravitational waves is the central aim of this work. This chapter will describe gravitational waves and how they arise as a prediction of General Relativity. It will also provide a brief catalog of different types of gravitational waves, with their attendant characteristics, sources, and search techniques.

2.1 Waves in General Relativity

The spacetime interval ds is one of the central concepts of special relativity. The spacetime interval between any two neighboring points is defined by [6]

$$ds^2 \equiv -c^2dt^2 + dx^2 + dy^2 + dz^2 \tag{2.1}$$

where c is the speed of light. In more general terms, Eqn. 2.1 can be written

$$ds^2 = \eta_{\mu\nu}dx^\mu dx^\nu \tag{2.2}$$

where the Greek indices, ranging from 0 to 3, represent the spacetime coordinates t, x, y, and z, and $\eta_{\mu\nu}$ is the Minkowski metric, which is (in Cartesian spatial coordinates)

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{2.3}$$

The Minkowski metric represents the flat spacetime of special relativity, but general relativity uses the more general metric $g_{\mu\nu}$, which describes curved spacetime. In particular, we can look to the case of a small perturbation to a flat spacetime, for which we can use

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{2.4}$$

Here, $h_{\mu\nu}$ is the metric perturbation away from a flat Minkowski space. It carries all of the information about the space-time curvature–and therefore all the information about the gravitation.

Gravitation *is* curvature in General Relativity–that is, objects "feeling" the force of gravity are more correctly said to be feeling the curvature of space [10]. This idea is expressed mathematically by the Einstein field equation:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{2.5}$$

where $G_{\mu\nu}$ is the Einstein curvature tensor, G the gravitational constant, and $T_{\mu\nu}$ the stress-energy tensor. The stress-energy tensor is a symmetric tensor of order two which describes the "energy-like aspects of a system–energy density, pressure, stress, and so forth" according to Carroll [11], whose derivation we will follow. At its most

basic level, this equation is telling us that the curvature ($G_{\mu\nu}$) is described by the sources of mass density ($T_{\mu\nu}$)–spacetime curvature arises from matter.

The Einstein curvature tensor is defined by the relation

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R \tag{2.6}$$

where $g_{\mu\nu}$ is our aforementioned metric, $R_{\mu\nu}$ is the Ricci tensor, and R is the Ricci scalar. The Ricci scalar is defined by

$$R \equiv g^{\mu\nu} R_{\mu\nu} \tag{2.7}$$

where $g^{\mu\nu}$ is the inverse of the metric $g_{\mu\nu}$. In the case of the metric of Eq. 2.4, it is given to first order by

$$g^{\mu\nu} = \eta^{\mu\nu} - h^{\mu\nu}. \tag{2.8}$$

The Ricci tensor $R_{\mu\nu}$ is derived from the Riemann tensor $R^{\mu}_{\nu\rho\sigma}$; in particular

$$R_{\mu\nu} = R^{\alpha}_{\mu\alpha\nu} \tag{2.9}$$

where the Riemann tensor itself is defined as

$$R^{\mu}_{\nu\rho\sigma} \equiv \partial_\rho \Gamma^{\mu}_{\nu\sigma} - \partial_\sigma \Gamma^{\mu}_{\nu\rho} + \Gamma^{\mu}_{\alpha\rho} \Gamma^{\alpha}_{\nu\sigma} - \Gamma^{\mu}_{\alpha\sigma} \Gamma^{\alpha}_{\nu\rho}. \tag{2.10}$$

$\Gamma^{\rho}_{\mu\nu}$ is the Christoffel connection. A connection is a correction added to the partial derivative ∂_μ to make the resulting operation covariant [11]. The Christoffel

connection is the connection related to the metric $g_{\mu\nu}$, and which corrects geodesic paths [12]. It is defined as

$$\Gamma^{\rho}_{\mu\nu} \equiv \frac{1}{2} g^{\rho\sigma}(\partial_\mu g_{\sigma\nu} + \partial_\nu g_{\sigma\mu} - \partial_\sigma g_{\mu\nu}). \tag{2.11}$$

Substituting the weak perturbation metric (Eq.(2.4) into these equations gives, to first order in $h_{\mu\nu}$ (and remembering that since $h_{\mu\nu}$ is a weak perturbation, we can assume $|h_{\mu\nu}| \ll 1$):

$$\begin{aligned}
\Gamma^{\rho}_{\mu\nu} &= \frac{1}{2}(\eta^{\rho\sigma} - h^{\rho\sigma})[\partial_\mu(\eta_{\sigma\nu} + h_{\sigma\nu}) + \partial_\nu(\eta_{\sigma\mu} + h_{\sigma\mu}) - \partial_\sigma(\eta_{\mu\nu} + h_{\mu\nu})] \\
&= \frac{1}{2}(\eta^{\rho\sigma} - h^{\rho\sigma})(\partial_\mu h_{\sigma\nu} + \partial_\nu h_{\sigma\mu} - \partial_\sigma h_{\mu\nu}) \\
&= \frac{1}{2}(\partial_\mu h^{\rho}_{\nu} + \partial_\nu h^{\rho}_{\mu} - \partial^{\rho} h_{\mu\nu}) + \cdots ,
\end{aligned} \tag{2.12}$$

$$\begin{aligned}
R_{\mu\nu} &= \partial_\alpha \Gamma^{\alpha}_{\mu\nu} - \partial_\nu \Gamma^{\alpha}_{\mu\alpha} + \Gamma^{\alpha}_{\beta\alpha}\Gamma^{\beta}_{\mu\nu} - \Gamma^{\alpha}_{\beta\nu}\Gamma^{\beta}_{\mu\alpha} \\
&= \frac{1}{2}\partial_\alpha(\partial_\mu h^{\alpha}_{\nu} + \partial_\nu h^{\alpha}_{\mu} - \partial^{\alpha} h_{\mu\nu}) - \frac{1}{2}\partial_\nu(\partial_\mu h^{\alpha}_{\alpha} + \partial_\alpha h^{\alpha}_{\mu} - \partial^{\alpha} h_{\mu\alpha}) + \cdots \\
&= \frac{1}{2}(\partial_\mu \partial^{\alpha} h_{\nu\alpha} + \partial_\nu \partial^{\alpha} h_{\mu\alpha} - \Box h_{\mu\nu} - \partial_\mu \partial_\nu h) + \cdots ,
\end{aligned} \tag{2.13}$$

$$\begin{aligned}
R &= \frac{1}{2}(\eta^{\mu\nu} - h^{\mu\nu})(\partial_\alpha \partial_\mu h^{\alpha}_{\nu} + \partial^{\alpha}\partial_\nu h_{\mu\alpha} - \Box h_{\mu\nu} - \partial_\mu \partial_\nu h^{\alpha}_{\alpha}) \\
&= \partial^{\nu}\partial^{\alpha} h_{\nu\alpha} - \Box h + \cdots
\end{aligned} \tag{2.14}$$

Here $\partial_\alpha h^{\alpha}_{\mu} = \partial^{\alpha} h_{\mu\alpha}$, $\Box \equiv \partial_\alpha \partial^{\alpha} = \left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)$ (where ∇^2 is the derivative with respect to the spatial dimensions–$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}\right)$), $h = \eta^{\mu\nu} h_{\mu\nu}$, and the $\cdots$ indicate the neglected higher-order terms. (We will be evaluating every equation to first

order in $h_{\mu\nu}$; although henceforth we will not include these dots, the neglected terms exist.)

Substituting these last two equations into Eq.(2.6) yields

$$G_{\mu\nu} = \frac{1}{2}[\partial_\mu \partial^\alpha h_{\nu\alpha} + \partial_\nu \partial^\alpha h_{\mu\alpha} - \Box(h_{\mu\nu} - \eta_{\mu\nu} h) - \partial_\mu \partial_\nu h - \eta_{\mu\nu} \partial^\rho \partial^\sigma h_{\rho\sigma}]. \tag{2.15}$$

It is useful here to define

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} h \tag{2.16}$$

which has inversion

$$h_{\mu\nu} = \bar{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} \bar{h}. \tag{2.17}$$

Substituting Eq.(2.17) into Eq.(2.15) gives us a new form for the Einstein curvature tensor:

$$\begin{aligned} G_{\mu\nu} &= \frac{1}{2}[\partial_\mu \partial^\alpha (\bar{h}_{\nu\alpha} - \frac{1}{2}\eta_{\nu\alpha}\bar{h}) + \partial_\nu \partial^\alpha (\bar{h}_{\mu\alpha} - \frac{1}{2}\eta_{\mu\alpha}\bar{h}) \\ &\quad - \Box(\bar{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\bar{h} + \eta_{\mu\nu}\bar{h}) + \partial_\mu \partial_\nu \bar{h} - \eta_{\mu\nu}\partial^\rho \partial^\sigma (\bar{h}_{\rho\sigma} - \frac{1}{2}\eta_{\rho\sigma}\bar{h})] \\ &= \frac{1}{2}[\partial_\mu \partial^\alpha \bar{h}_{\nu\alpha} + \partial_\nu \partial^\alpha \bar{h}_{\mu\alpha} - \Box \bar{h}_{\mu\nu} - \eta_{\mu\nu}\partial^\rho \partial^\sigma \bar{h}_{\rho\sigma}]. \end{aligned} \tag{2.18}$$

We can make these equations especially clear by now selecting the transverse traceless ("TT") gauge, in which spacetime coordinates coincide with the world lines of freely falling test masses [13]. This is done by setting

$$\partial^\nu \bar{h}_{\mu\nu} = 0. \tag{2.19}$$

Using this choice of coordinates, Einstein's field equation becomes:

$$G_{\mu\nu} = -\frac{1}{2}\Box\bar{h}_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}, \tag{2.20}$$

or, equivalently,

$$\Box\bar{h}_{\mu\nu} = \left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \tag{2.21}$$

Lastly, we look to the weak-field limit–when we are far from sources of gravitation, and can take the stress-energy tensor $T_{\mu\nu} = 0$. Einstein's equation has become a homogeneous wave equation [14]:

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)h_{\mu\nu} = 0 \tag{2.22}$$

The elements of $h_{\mu\nu}$ then take the functional form $h(2\pi ft - \mathbf{k}\cdot\mathbf{x})$, which represents a plane wave propagating in the direction $\hat{k} \equiv \mathbf{k}/|\mathbf{k}|$ with speed c ($f = |\mathbf{k}|/2\pi c$). These are gravitational waves, moving at the speed of light c.

Since $h_{\mu\nu}$ is transverse and traceless, a gravitational wave propagating in the $\hat{z}$ direction will have $h_{\mu\nu}$ of the form [6]

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & a & b & 0 \\ 0 & b & -a & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{2.23}$$

It's apparent that this metric can be written as the sum of two components: $h = a\hat{h}_+ + b\hat{h}_\times$. The components are

$$\hat{h}_+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{2.24}$$

and

$$\hat{h}_\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{2.25}$$

These are the two orthogonal polarizations for gravitational waves, pronounced "h plus" and "h cross." The difference is best illustrated by the effect of a wave on a ring of particles, as seen in Figure 2.1. A plus-polarized wave propagating along the $\hat{z}$ axis will simultaneously compress the ring in $\hat{x}$ while stretching it in $\hat{y}$, and then alternately stretches the ring in $\hat{x}$ while compressing it in $\hat{y}$. A cross-polarized wave propagating along the $\hat{z}$ axis will act in a similar manner, but on axes rotated by 45°. The polarizations are invariant up to a sign change under a 90° rotation.

While electromagnetic radiation can be emitted by a variety of multipoles, gravitational waves are more restricted. Conservation of energy, conservation of linear momentum, and conservation of angular momentum forbid, respectively, monopole, dipole, and magnetic dipole radiation; gravitational radiation thus requires at least a nonzero quadrupole moment [13].

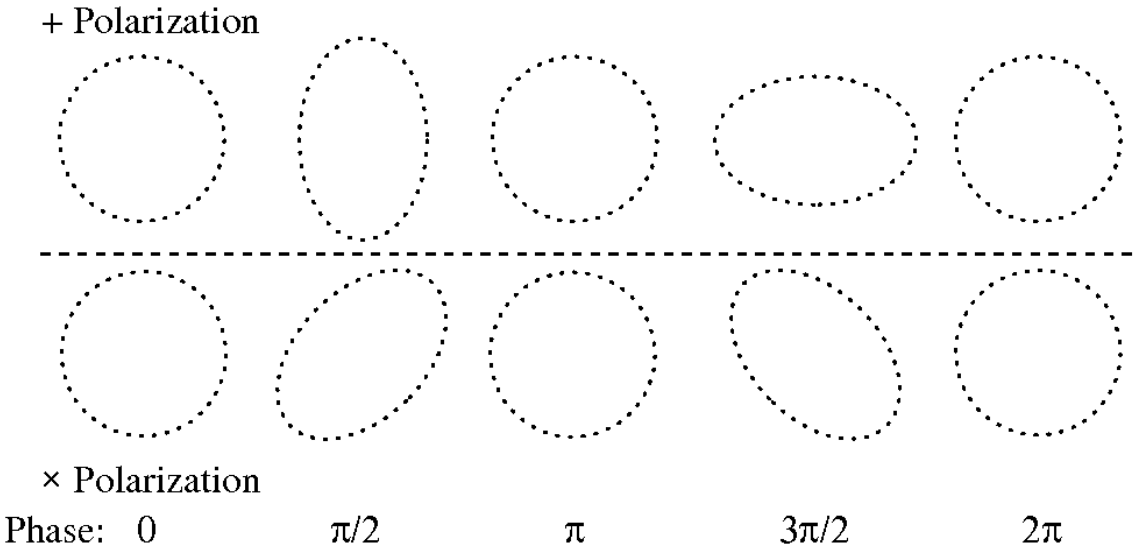

Figure 2.1: The effect of a gravitational wave traveling into the page ($\hat{z}$) on a ring of particles arrayed in the x-y plane. The top row shows a plus polarized wave, which first compresses space in $\hat{x}$ while stretching it in $\hat{y}$, and then vice-versa. The bottom row illustrates a cross-polarized wave, which has the same effect but rotated by 45° in the x-y plane.

2.2 Types of Gravitational Wave Sources

Gravitational waves come in a wide variety of categories and frequencies. Looking ahead to Chapter 3, I will here divide them into those that fall within the operating frequency range of ground-based interferometric detectors and those best accessible through other techniques.

2.2.1 The LIGO Band

Current ground-based interferometric detectors such as the Laser Interferometer Gravitational-wave Observatory (LIGO) are sensitive to gravitational waves over a broad range of frequencies, between 10 Hz and several kHz [13]. Within this wide band there are several different classes of signals to be sought after.

Compact binary coalescences (CBCs) are the most promising class of gravitational wave signatures LIGO hopes to detect. These are generated when two extremely

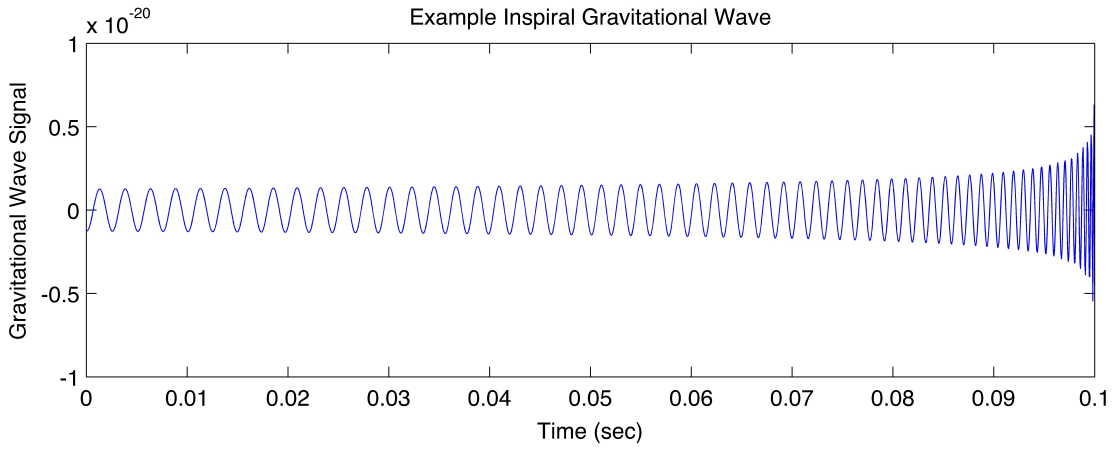

Figure 2.2: The signature "chirp" gravitational-wave signal of a compact binary coalescence. [Image: A. Stuver/LIGO]

massive objects (e.g., a pair of black holes or neutron stars, or a black hole-neutron star pair) inspiral towards each other, losing angular momentum and energy via gravitational radiation, before eventually colliding and merging to form a black hole. They are characterized by a "chirp" signal, a sinusoid increasing in both frequency and amplitude as the pair inspiral towards each other (see Figure 2.2); by a less-well-understood merging phase; then finally by a strong emission of gravitational waves known as "ringdown" as the combined object reaches its new stable state [15]. These signals are transient, present in the LIGO band of frequencies for durations on the order of minutes or seconds, depending on the masses of the inspiralling objects; at an Earth-based detector, the waves have characteristic amplitude [10]

$$h_0 \simeq 2 \times 10^{-21} \left(\frac{f}{1\text{ kHz}}\right) \left(\frac{r}{10\text{ Mpc}}\right)^{-1} \left(\frac{m_1}{1.4M_\odot}\right) \left(\frac{m_2}{1.4M_\odot}\right) \left(\frac{M}{2.8M_\odot}\right)^{-1/3} \tag{2.26}$$

where f is the gravitational wave frequency, r the distance to the source, m_1 and m_2 the respective masses of each inspiralling object, and M the total mass of the system.

The inspiral and ringdown phases are very well modeled, and advances in numerical relativity have improved the understanding of the merger phase as well [16,17,18]. A search for CBCs typically begins with a matched-filter step, where known waveforms are used as templates. The output of the matched-template filter is thresholded at a given signal-to-noise ratio (SNR), and events for which the SNR crosses the threshold are considered triggers for followup. Triggers are initially generated independently for each detector, then subjected to a coincidence requirement between detectors, and then further subjected to consistency tests [15]. If no signals are found, upper limits can be set on the rate of CBCs per search volume. The latest CBC searches, using data from the LIGO and VIRGO detectors, have set upper limits of 3.3×10^7 per cubic Megaparsec per year [19] for non-spinning binary black hole mergers, 1.3×10^4 per cubic Megaparsec per year [20] for binary neutron star mergers, and 3.1×10^5 per cubic Megaparsec per year [20] for black hole-neutron star mergers.

Burst signals are short unmodeled "bursts" of gravitational radiation, typically taken to be shorter than one second but with significant enough power to rise above the background noise floor of the detector. Sources of these "bursts" can be asymmetric core-collapse supernovae [21], or the ringdown portion of a CBC whose inspiral and merger stages are outside the LIGO band. Gamma-ray bursts (GRBs) and soft gamma repeaters (SGRs) can also be used as external triggers for burst searches [15]. Burst searches rely heavily on correlation and coincidence analyses between detectors, since the detectors can undergo "glitches," caused by environmental noise sources and instrumental artifacts, which mimic burst signals. Correlation need not be limited to gravitational-wave detectors–some of the most recent burst searches have correlated LIGO data with electromagnetic spectrum data from satellite observatories like Swift

and Fermi [22].

A stochastic background of gravitational waves could indicate the gravitational radiation produced just after the Big-Bang, which persists since gravitational waves interact weakly with matter [15]. Other cosmological sources could include phase transitions in the early universe, topological defects formed during these transitions, or cosmic strings; astrophysical sources can include the superposition of periodic signals from supernovae, spinning neutron stars, and low-mass X-ray binaries [13]. Stochastic signals are much weaker than those previously discussed, and searches for them involve integrations over long observation times and the use of environmental monitoring data to identify sources of instrumental noise which may be correlated across detectors. Recent stochastic searches have placed upper limits on Ω_{GW}, the energy density of the stochastic gravitational wave background, in four different frequency bands spanning 41.5 to 1726 Hz, with a lowest value of 5.6×10^{-6} for the frequency range between 41.5 to 169.25 Hz [23].

The final category of sources in the LIGO band, and the one of most importance to this work, are continuous-wave (CW) sources. Continuous gravitational waves are expected to be emitted by spinning neutron stars, observed as pulsars with spin frequencies between a few milli-Hz and about a kHz. Since an axisymmetric body has no quadrupole moment, spinning neutron stars emitting gravitational waves must have some axisymmetric asymmetry–a large "mountain" on their crust, or some asymmetry due to accretion processes. This class of sources will be discussed in detail in Chapter 4, and Chapters 6 and 7 will describe a search for gravitational waves from such an object in a nearby globular cluster.

2.2.2 Other techniques

Ground-based interferometry is only one technique for the detection of gravitational waves. Gravitational waves can arrive in a wide range of frequencies, many outside the LIGO frequency band described in the preceding section, which can be pursued with a variety of methods [13].

Waves with frequencies below 10^{-13} Hz, corresponding to wavelengths roughly the size of the universe, are expected to have been generated during the period of inflation closely following the Big Bang. These waves can be detected through the imprints they leave in the cosmic microwave background, like those searched for (but ultimately not found) in the recent BICEP2 experiment [24]. Massive binary black holes are expected to generate a background of waves between 10^{-9} and 10^{-7} Hz, which can be detected using time-of-arrival delays in the pulses of an array of millisecond pulsars [25]. Finally, the space between 10^{-5} Hz and 1 Hz is the domain of space-based interferometric detectors, such as the planned "evolved Laser Interferometer Space Antenna" (eLISA) [26]. Astrophysical sources of such signals include stellar-mass black hole and white dwarf binaries.

CHAPTER III

Detecting Gravitational Waves

3.1 Interferometry

An interferometer, broadly speaking, is a device which splits a beam of light into two and then recombines it. There are several types, with different schemes for splitting the light, altering its path length, and recombining the beams; we shall focus on the interferometer made famous by Michelson and Morley [27] in their attempt to measure the Earth's movement against the then-supposed luminiferous ether. This type is known, after its inventor, as a Michelson interferometer.

3.1.1 Michelson Interferometers

A Michelson interferometer has a source of coherent light, in most modern instances a laser, which is incident on a beamsplitter. A beamsplitter is a partially reflecting mirror, which allows some light to pass through it while reflecting the rest. Such a mirror placed at a 45° angle to the incident beam, for example, will result in two beams at right angles. These beam directions define the arms of the interferometer, and at the end of each is placed a fully reflecting mirror. Light reflected at the beamsplitter will travel up and down one arm; light passing through the beamsplitter will make a circuit of the other. The light beams, returning from the end mirrors, are superposed at the beam splitter, and the resulting combined beam is

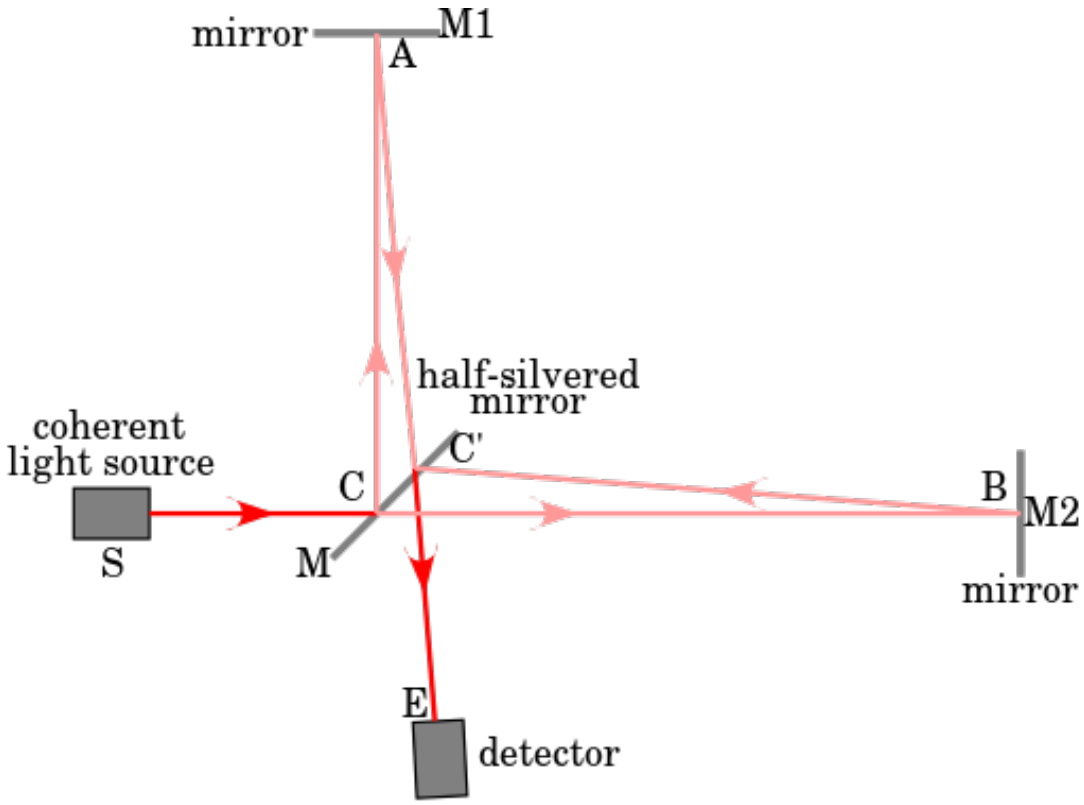

Figure 3.1: Schematic diagram of a Michelson interferometer. A coherent beam source is incident on a beamsplitter, which reflects some light while allowing the rest to pass through. Each "arm" of the interferometer is then reflected from a mirror and recombines at the beamsplitter. The recombined beam moves to a detector.

directed towards a photodetector. A schematic of such an interferometer is provided in Figure 3.1.

Michelson interferometers can be used as very precise length-sensing devices, because of the interference pattern produced when the two split beams recombine. We'll follow the analysis of Saulson [6], which itself is based on Haus [28].

If we take the input light (incident on the beam splitter) to be traveling in the x-direction, then its electric field is given by

$$E_{in} = E_0 e^{i(2\pi f t - kx)} \tag{3.1}$$

where E_0 is the amplitude of the light, f the frequency, t the time coordinate, and k the wavenumber $k = \frac{2\pi}{\lambda}$ with λ the wavelength. Assuming a beamsplitter that reflects half the light and transmits half the light, we can represent transmission by the amplitude transmission coefficient $t = i/\sqrt{2}$ and reflection by the amplitude reflection coefficient $r = 1/\sqrt{2}$, such that the field of the light which passes through the beam splitter and continues in the x-direction has field

$$E_t = i(E_0/\sqrt{2})e^{i(2\pi ft - kx)} \tag{3.2}$$

while that which is reflected into the other arm has field

$$E_r = (E_0/\sqrt{2})e^{i(2\pi ft - ky)} \tag{3.3}$$

The reflection at each end mirrors multiplies each wave by -1, and upon returning to the beamsplitter it is once again partly reflected and transmitted, with the same coefficients. Thus the field of the recombined beam moving towards the detector is

(3.4)

$$E_{out} = (i/2)E_0e^{i(2\pi ft - 2kL_x)} + (i/2)E_0e^{i(2\pi ft - 2kL_y)} = ie^{i(2\pi ft - k(L_x + L_y))}E_0 cos[k(L_x - L_y)]$$

where L_x and L_y are, respectively, the $\hat{x}$- and $\hat{y}$-arm lengths. Light returning from the mirrors which exits back towards the input source will have field

$$E_{refl} = -ie^{i(2\pi ft - k(L_x + L_y))}E_0 sin[k(L_x - L_y)] \tag{3.5}$$

The state of the light leaving the interferometer is thus determined by the difference in the arm lengths. When the lengths are equal, and the difference is zero, E_{refl}

goes to zero, and light of amplitude E_0 emerges towards the photodetector. On the other hand, for a phase difference of $\pi/2$, it is E_{out} which goes to zero while light of amplitude E_0 is reflected back at the input. If we measure the power, which is proportional to the square of the field, at the output detector, we will get

$$P_{out} = P_{in} cos^2[k(L_x - L_y)] = P_{in}(1 + cos[2(k(L_x - L_y))]) \tag{3.6}$$

A Michelson interferometer thus acts as a transducer–translating differences in light travel time into output optical power. Measuring the brightness of light at the output will give us the travel time difference up to an integer number of optical periods $\tau \equiv \lambda/c$.

3.1.2 Detection of GWs by Interferometers

The trick to using a Michelson interferometer as a gravitational-wave detector lies in making one key change to the experimental setup. Rather than fixing any of the optics to any single rigid structure, each mirror acts as a freely-falling mass. This way the optics will respond in a simple way to gravitational effects, especially when described in the transverse traceless gauge discussed in the previous chapter. The derivation here will once more follow [6].

Light travels at constant speed c in any inertial frame of reference. In the terms of relativity, this means that a ray of light connects sets of points separated by a spacetime interval ds equal to zero. Imagine, then, our Michelson interferometer with its arms aligned along the $\hat{x}$ and $\hat{y}$ axes, as in the previous section. Then we have

$$ds^2 = 0 = g_{\mu\nu} dx^\mu dx^\nu = (\eta_{\mu\nu} + h_{\mu\nu}) dx^\mu dx^\nu \tag{3.7}$$

Remembering the metric matrices from the previous chapter, this gives

$$ds^2 = 0 = -c^2dt^2 + (1 + h_{11}(2\pi ft - \mathbf{k}\cdot\mathbf{x}))dx^2 \tag{3.8}$$

along the $\hat{x}$-axis, and a similar expression for the $\hat{y}$-axis, but with h_{22} in place of h_{11}. We can then find the light travel time from the beam splitter to the end of the $\hat{x}$-arm by taking the square root of 3.8 and integrating:

$$\int_0^{\tau_{out}} dt = \frac{1}{c}\int_0^L \sqrt{1 + h_{11}}dx \tag{3.9}$$

Since we are working in the limit where h is a small perturbation of a flat spacetime, we can assume $|h_{11}| \ll 1$ and take a binomial expansion of the square root.

$$\int_0^{\tau_{out}} dt \approx \frac{1}{c}\int_0^L (1 + \frac{1}{2}h_{11}(2\pi ft - \mathbf{k}\cdot\mathbf{x}))dx \tag{3.10}$$

Similarly, the return trip from the end of the $\hat{x}$-arm to the beam splitter can be expressed as

$$\int_{\tau_{out}}^{\tau_{RT}} dt \approx -\frac{1}{c}\int_L^0 (1 + \frac{1}{2}h_{11}(2\pi ft - \mathbf{k}\cdot\mathbf{x}))dx \tag{3.11}$$

Combining the two gives the total round trip time τ_{RT}:

$$\tau_{RT} = \frac{2L}{c} + \frac{1}{2c}(\int_0^L h_{11}(2\pi ft - \mathbf{k}\cdot\mathbf{x})dx - \int_L^0 h_{11}(2\pi ft - \mathbf{k}\cdot\mathbf{x})dx) \tag{3.12}$$

For further simplicity, we'll assume a plus-polarized gravitational wave propagating along the $\hat{z}$ axis, with amplitude $h_{11} = h_{22} = h$. Then $\mathbf{k}\cdot\mathbf{x} = 0$. If we also assume

that the gravitational wave has a frequency f_{gw} such that $\tau_{RT} \ll 1/(2\pi f_{gw})$, then the metric perturbation is approximately constant during any given round trip and the perturbations to the light travel time in each arm will be equal and opposite. This then gives a total travel time difference of

$$\Delta\tau(t) = h(t)\frac{2L}{c} \tag{3.13}$$

This travel time difference can be expressed in terms of phase by dividing by the period of oscillation of the light ($\lambda/2\pi c$), such that

$$\Delta\phi(t) = h(t)\frac{2L}{c}\frac{2\pi c}{\lambda} = \frac{4\pi L}{\lambda}h(t) \tag{3.14}$$

Finally, note that since light travels at constant speed c, the travel time difference over a path-length difference ΔL can also be expressed as $2\Delta L/c$ (where the factor of two enters because light must travel twice the arm-length to complete one trip), which combined with 3.13 gives

$$h(t) = \frac{\Delta L}{L} \tag{3.15}$$

Given Eqn.(3.15), it is apparent that one of the keys to detecting small gravitational perturbations $h(t)$ is to make L as long as feasibly possible. The brute-force way to go about this is to create an interferometer with extremely long arms, but there are also other ways to increase the path length L. These and other features of the LIGO interferometers will be covered in the next section.

3.2 The LIGO Detectors

The LIGO collaboration operates a pair of Michelson-interferometer gravitational wave detectors. One is located in Hanford, Washington; the other in Livingston, Louisiana. This section will cover some of the challenges that arise when operating interferometric detectors, and some of the innovations that the LIGO detectors have adopted in response.

3.2.1 Fabry-Perot Cavities

The LIGO interferometers have arms four kilometers in length, which for typical values of $h(t)$ correspond to ΔL of approximately 4×10^{-18} meters [29]. To eliminate scattering of the laser beam by particulates, the interferometer arms are kept in vacuum at 10^{-9} torr. Further, each arm is made a Fabry-Perot cavity by the insertion of an input mirror on each arm (see Figure 3.2). A Fabry-Perot cavity consists of two aligned mirrors facing each other such that light bounces back and forth between them repeatedly. For cavity lengths $L = n\lambda/2$, where n is any integer, a Fabry-Perot cavity is in resonance–that is, the relative phase of all of the beams bouncing back and forth inside it is such that it transmits maximum power. This has the same effect as trapping light inside the arm and forcing it to make many trips before returning to the beam splitter; a Fabry-Perot cavity has an effective light "storage time" [6]

$$\tau_s = \frac{L}{c}\frac{\mathcal{F}}{\pi} \tag{3.16}$$

where $\mathcal{F}$ is the finesse of the cavity. The finesse of a cavity is a measure of the sharpness of its resonance, defined as the ratio between the width of a resonance peak at half-power and the free spectral range $\Delta f = c/2L$. The finesse is given by

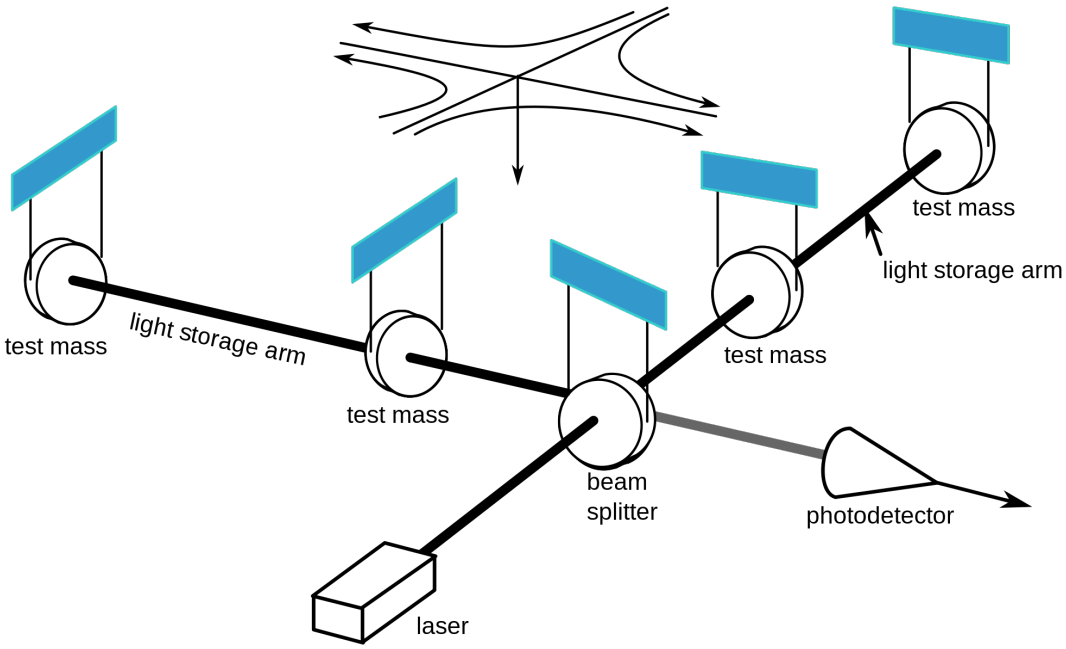

Figure 3.2: Schematic diagram of the LIGO Fabry-Perot interferometer. Suspended input and output test masses create a Fabry-Perot cavity in each arm. At top, a gravitational wave is incident from the $\hat{z}$-direction.

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} \tag{3.17}$$

where r_1 and r_2 are the amplitude reflectivities of the two mirrors which make up the cavity. Further, a Fabry-Perot cavity of finesse $\mathcal{F}$ gives $2\mathcal{F}/\pi$ times the phase shift in response to a gravitational wave, as long as the time scale for gravitational metric changes is long compared to $\mathcal{F}L/c$. Finally, in a Fabry-Perot cavity all of the many optical path lengths are superimposed spatially, as opposed to a folded design, so the optics can remain small in size, making it easier to manufacture them with the necessary optical properties. The LIGO interferometers lock their Fabry-Perot cavities to resonance using the Pound-Drever-Hall technique [30], and in their sixth science run used finesse $\mathcal{F} = 220$ and storage time $\tau_s = 0.95$ milliseconds [31].

3.2.2 Primary noise sources

The LIGO interferometers' noise curves are dominated by different sources at different frequencies. At low frequencies (below 50 Hz), the dominant source of disturbances is seismic–the motion of the ground coupling to the motion of the test masses. This is countered by both active (servo controls) and passive (multi-stage pendulums) seismic isolation systems on each of the test masses.[1] At medium frequencies (50 Hz - 200 Hz), the detectors are dominated by thermal noise. This is the Brownian motion of particles striking the test masses–countered by using materials with high quality factor Q. To this end, LIGO uses test masses of fused silica. At high frequencies the dominant source of error is shot noise–the counting error in the number of photons arriving at the output detector. Shot noise is inversely proportional to the input power, so it can be mitigated by increasing laser power, and LIGO makes use of a 1064 nm Nd:YAG laser operating at 10 Watts [29]. Increased laser power brings its own problems, however. Increased radiation pressure noise–movement induced in the test masses by the photons impacting them–can be compensated for by using heavier optics; LIGO mirrors weigh more than 10 kilograms. The heightened laser power also leads to thermal heating of the mirrors from photon absorption–despite having reflectivities of 99.997% [15]. This is compensated for with the aptly named Thermal Compensation System [32], which ensures the optics are heated to compensate distortions.

3.2.3 S6 noise sources

The search for gravitational waves described in Chapters 6 and 7 was performed in data taken during LIGO's sixth science run (S6), which spanned July 7, 2009

[1] Although installed at both sites now, during the sixth science run only the Livingston observatory had active isolation.

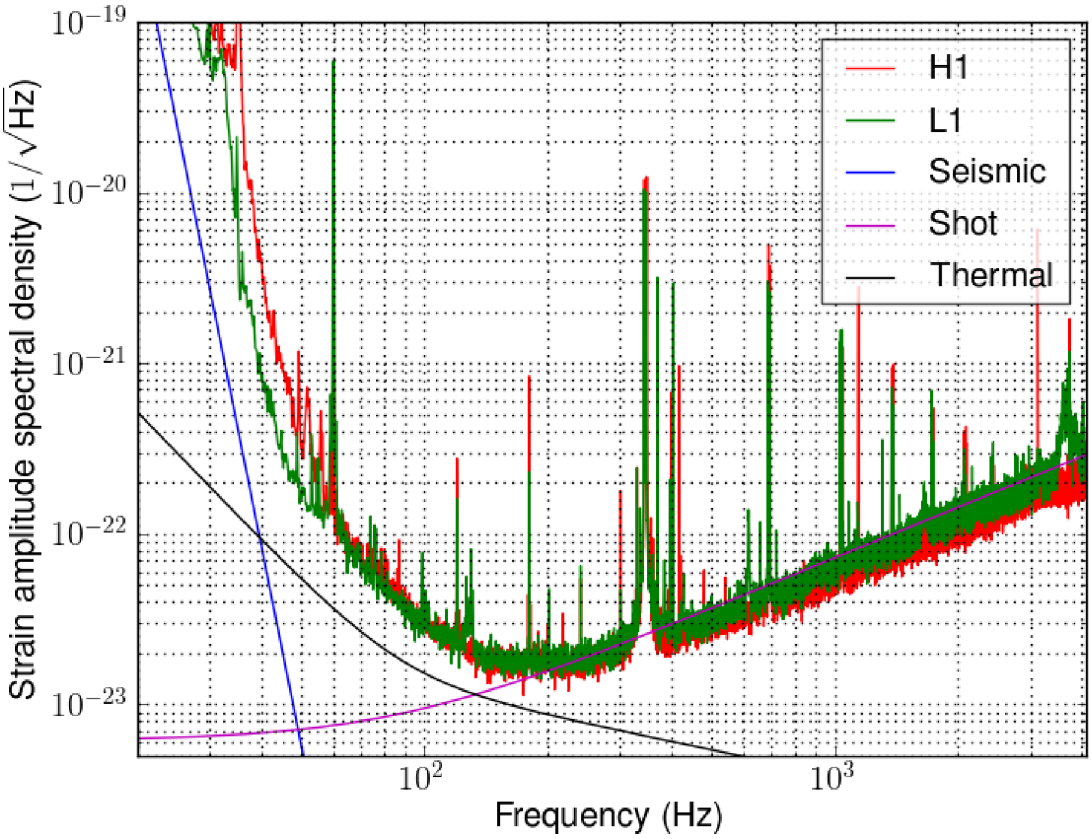

Figure 3.3: The noise curve for the LIGO detectors during S6. The green curve represents the Livingston (L1) detector; the red curve the Hanford (H1) detector. Included are the noise curves from the three dominant noise sources: seismic, thermal, and shot. Image credit: Macleod and Smith, 2013 [33].

to October 20, 2010. The LIGO noise curve as it existed during S6 is displayed in Figure 3.3, where strain sensitivity is plotted against frequency on logarithmic axes.

As can be seen from the figure, seismic noise is the main contributor at low frequencies, thermal noise at medium frequencies, and shot noise at high frequencies, as described in the preceding. At lower frequencies, there is also a discrepancy between the seismic noise curve and the LIGO sensitivity curve. This discrepancy is likely explained by what is known as Barkhausen noise [34]. The test masses are driven by electromagnetic actuator coils, and the low frequency currents applied to the magnets in those coils are upconverted into broadband actuation force noise below approximately 80 Hz [35]. This extra noise source will be mitigated in the next

generation (see next section) of detectors [31] by a quadruple-pendulum suspension system and electrostatic (as opposed to magnetic) actuation [36].

Also visible in the figure are several prominent noise lines: large disturbances at relatively well-defined stationary frequencies, which are very relevant for continuous-wave gravitational wave searches (described in Chapter 4) like the one we undertook (Chapters 6 and 7). Immediately apparent are the large disturbances at 60 Hz and its harmonics. These are due to the AC current supplied at 60 Hz by the U.S. power mains [37]. While large, these disturbances are relatively narrow in frequency; the search for the Crab pulsar was only moderately affected in LIGO's fifth [38] and sixth [39] science runs, when that pulsar's expected gravitational wave frequency was 59.77 Hz and 59.44 Hz, respectively. The large, broader disturbances between 330 Hz and 350 Hz and related harmonics are the detector's violin modes, the resonance frequencies for the suspended test masses; also visible near 400 Hz are a pair of calibration lines (one at each detector) used to measure interferometer response. Additionally, the sixth science run was plagued by a comb of 2-Hz and 16-Hz harmonics at both detectors; the specific mechanism causing these was never proven but they were known to be instrumental artifacts [37].

3.2.4 Advanced LIGO and the worldwide detector network

LIGO has recently completed an upgrade of all major components of its two detectors, which should result eventually in an order of magnitude increase in sensitivity [40]. Upgrades include fused silica suspensions welded to the fused silica optics, which combine with the aforementioned upgraded seismic isolation systems to hold LIGO optics in place to within 10^{-14} meters; an upgrade from 10-kg to 40-kg test masses to reduce radiation pressure noise; and an increase in laser power to 180 W. In addition, LIGO uses its multiple detectors to resolve the sky location of incom-

ing gravitational wave signals, and this capability too is receiving a boost as more interferometric detectors come online. The French-Italian VIRGO collaboration operates a 3-km detector in Northern Italy, which is also in the process of receiving upgrades [41], and a smaller 600-m interferometer in Hannover, Germany is pioneering quantum squeezing techniques [42]. Work is also under way on an underground, cryo-cooled interferometer in Japan [43], and plans are in place to install a 4-km LIGO detector, originally intended for Hanford, in India [44].

CHAPTER IV

CW GWs: Astrophysical Sources and Target Selection

Continuous-wave (CW) gravitational waves, described briefly at the end of Chapter 2, are long-lived quasi-monochromatic gravitational wave signals. They are characterized by a frequency f and its evolution ($\dot{f}, \ddot{f}, etc.$), their amplitude h, their polarization angle ψ and inclination angle ι, and a phase Ψ. CWs are expected to be weak, but long-lived, so that long stretches of observation time (days, months, years) can be integrated to search for them. The first section of this chapter will describe likely sources of these CW waves; the following section will discuss some common theorized mechanisms for their emission. The chapter will close by bringing these criteria together to explain the target selection for the search for gravitational waves described in Chapters 6 and 7.

4.1 Astrophysical Sources

Neutron stars are prime sources of gravitational waves. Chapter 2 discussed their roles as members of a compact binary coalescence, the target of a different branch of gravitational wave searches. Neutron stars in binary systems are also an area of ongoing development in gravitational wave searches–the pioneering TwoSpect [45] pipeline was only one method among several that recently took part in a mock data challenge [46] with simulated data from the Scorpius X-1 binary. This work focuses

on gravitational waves emitted by isolated spinning neutron stars.

4.1.1 Neutron stars

A neutron star is a compact object, roughly 10 kilometers in radius and $1.4M_{\odot}$ in mass [47]. These objects are the neutron cores of former stars, left after the stellar envelope is blown off in a supernova explosion [48]. Their existence was first postulated by Landau [49]; Baade and Zwicky first suggested their supernova origin [50]. Likely made nearly entirely of neutrons, these stars are supported against gravity by degeneracy pressure.

Since neutron stars are cold (i.e., they do not produce luminosity via fusion), in the simplest Newtonian case their structure can be expressed through three equations. The mass equation, which relates density to mass; the momentum equation, which expresses the hydrostatic equilibrium; and the potential equation, which describes the gravitational potential [47]. We can think of the star as a series of infinitesimaly thin shells of thickness dr at a radius r. Each shell will have mass dm given by its density ρ:

$$dm = 4\pi r^2 \rho dr. \tag{4.1}$$

Then the mass can be expressed as function of radius:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r), \tag{4.2}$$

or

$$\frac{dm(r)}{dr} = \frac{4\pi r^2 \epsilon(r)}{c^2} \tag{4.3}$$

where we have substituted the mass energy density $\epsilon = \rho c^2$. The total mass can then be found by integrating from the center to the stellar radius R, where R is defined as the location where the pressure vanishes.

$$M = \int_0^R \frac{dm}{dr} dr = \frac{4\pi}{c^2} \int_0^R \epsilon(r) r^2 dr \tag{4.4}$$

The hydrostatic equilibrium condition is that gravity must balance pressure. Considering again our infinitesimal shells, the change in pressure P across a shell bounded by the radii r and $r + dr$ is

$$P(r) - P(r + dr) = -\frac{dP}{dr} dr, \tag{4.5}$$

where the negative sign appears because the pressure decreases outward (recall that it vanishes at the stellar radius R). Pressure is force per area, and the gravitational force is determined by the enclosed mass, so we have

$$f_{grav} = \frac{Gm(r)}{r^2} dm = \frac{Gm(r)}{r^2} \frac{4\pi r^2 \epsilon}{c^2} dr \tag{4.6}$$

where G is the gravitational constant. The surface area of a spherical shell is $4\pi r^2$; balancing the pressure change with the gravitational force per area we have

$$\frac{dP}{dr} = -\frac{Gm(r)}{r^2 c^2} \epsilon. \tag{4.7}$$

Lastly, the radial change of the gravitational potential ϕ is given by

$$\frac{d\phi}{dr} = \frac{Gm(r)}{r^2} = -\frac{c^2}{\epsilon} \frac{dP}{dr}. \tag{4.8}$$

The generalization of these equations to the General Relativistic case are called the Tolman-Oppenheimer-Volkoff equations [51], and they can be written as a form that corresponds to the preceding Newtonian equations with correction factors [47].

$$\frac{dm}{dr} = \frac{4\pi r^2 \epsilon}{c^2} \tag{4.9}$$

$$\frac{dP}{dr} = -\frac{Gm(r)\epsilon}{c^2 r^2}\left(1+\frac{P}{\epsilon}\right)\left[1+\frac{4\pi r^3 P}{m(r)c^2}\right]\left[1-\frac{2Gm(r)}{c^2 r}\right]^{-1} \tag{4.10}$$

$$\frac{d\phi}{dr} = -\frac{c^2}{\epsilon}\frac{dP}{dr}\left(1+\frac{P}{\epsilon}\right)^{-1} \tag{4.11}$$

The mass equation is unchanged; the first two correction terms on Eq.(4.10) (as well as the correction factor on Eq.(4.11)) arise due to special relativistic effects. The last correction factor on Eq.(4.10) is due to General Relativity and as such can be used as a criterion for whether General Relativity is important, much like the Lorentz factor γ in special relativity. In particular, General Relativity will be unimportant provided that the condition

$$\frac{2GM}{c^2 R} \ll 1 \tag{4.12}$$

holds. By dimensional analysis, we can see that the quantity $2GM/c^2$ must have units of length; it is in fact the Schwarzschild radius. Thus Eq.(4.12) can be seen as the ratio of an object's Schwarzschild radius and its actual physical radius; if this ratio is small, General Relativity is unimportant. For example, this ratio for our Sun is $\approx 4 \times 10^{-6}$ [47].

A neutron star is thought to be composed of several layers. The surface of a neutron star has densities less than 10^6 g/cm^3 and contributes a negligible amount of mass to the star. It can be subdivided into an atmosphere, which shapes the spectra emitted by neutron stars, and an envelope, just below the atmosphere and important for thermal and transport properties. The surface is also a layer where extremely high magnetic fields are sometimes found. A relatively solid outer crust, one to two kilometers thick, covers the neutron superfluid core of the neutron star, which holds 99% of the neutron star's mass [47].

The crust is of particular interest to us, since it is here that a bump or "mountain" might arise and produce gravitational wave emission (see 4.2). Here we follow the model of Lattimer and Prakash [52]. The crust is 1-2 kilometers thick, and is usually separated into an outer and an inner crust. It is comprised of nuclei arranged into configurations of lowest energy, with density ρ increasing radially inward. In the outer crust, at relatively low ($\rho < 10^9 \mathrm{kg} \cdot \mathrm{m}^{-3}$) densities, the dominant nucleus is thought to be iron (^{56}Fe); as density increases, matter becomes ionized and the nuclei are embedded in a nearly uniform background of electrons [47]. The nuclei become more and more neutron rich; they may contain nucleon numbers of up to several hundred and arrange in crystalline lattices. At what is known as the neutron drip density, $\rho_{drip} = 4 \times 10^{14} \mathrm{kg} \cdot \mathrm{m}^{-3}$, the neutrons begin to "drip" out of the nuclei, immersing them in a sea of background neutrons [47]. The nuclear lattices change with increasing density from spherical shells to flat slabs of nuclei to voids embedded in nuclear matter to, when the density is so high that nucleons touch, a uniform nuclear fluid. At a density of $10^{17} \mathrm{kg} \cdot \mathrm{m}^{-3}$, the nuclei cease to exist and matter consists of a homogeneous neutron-proton fluid; this point defines the boundary of the outer core [47].

4.1.2 Pulsars

Some neutron stars emit electromagnetic radiation along the axes of their magnetic fields. If these neutron stars are spinning (and their dipole is misaligned with respect to their spin axis) then from the Earth this beamed radiation will appear as a pulse of radiation each time it sweeps across us, with a frequency equal to the rotational frequency of the star. This class of neutron stars, first observed in 1967 [53], are thus called pulsars.

Pulsars can be either isolated or in binary systems, as is true for the broader class of neutron stars to which they belong. "Supernova remnant associations" are pulsars associated with known supernova remnants, indicating a possibility that we may know which supernovae birthed them; "rotating radio transients" is the label given to pulsars which have a high degree of pulse-to-pulse variability, making them difficult to detect with Fourier techniques.

Figure 4.1 also groups the pulsar population into three subcategories [15]. In the upper right of the plot are the magnetars, pulsars with extremely high magnetic fields, greater than 10^{14} Gauss. The large bulk of the pulsar distribution is clustered roughly around a one-second period. They are no older than a few million years and have magnetic fields between 10^{11} and 10^{13} Gauss. In the lower left, with much smaller rotation periods, are the so-called millisecond pulsars (MSPs). This class of pulsar is thought to be much older, and are called "recycled" pulsars–their rotation period has slowed and then been spun back up, almost certainly by accretion from a companion star. There are isolated MSPs, for which it is possible the companion has been destroyed by the MSP [55] or kicked from binary orbit by the MSP's supernova. MSPs are of interest to us, since their rotation frequencies overlap with the frequency band of greatest sensitivity for LIGO.

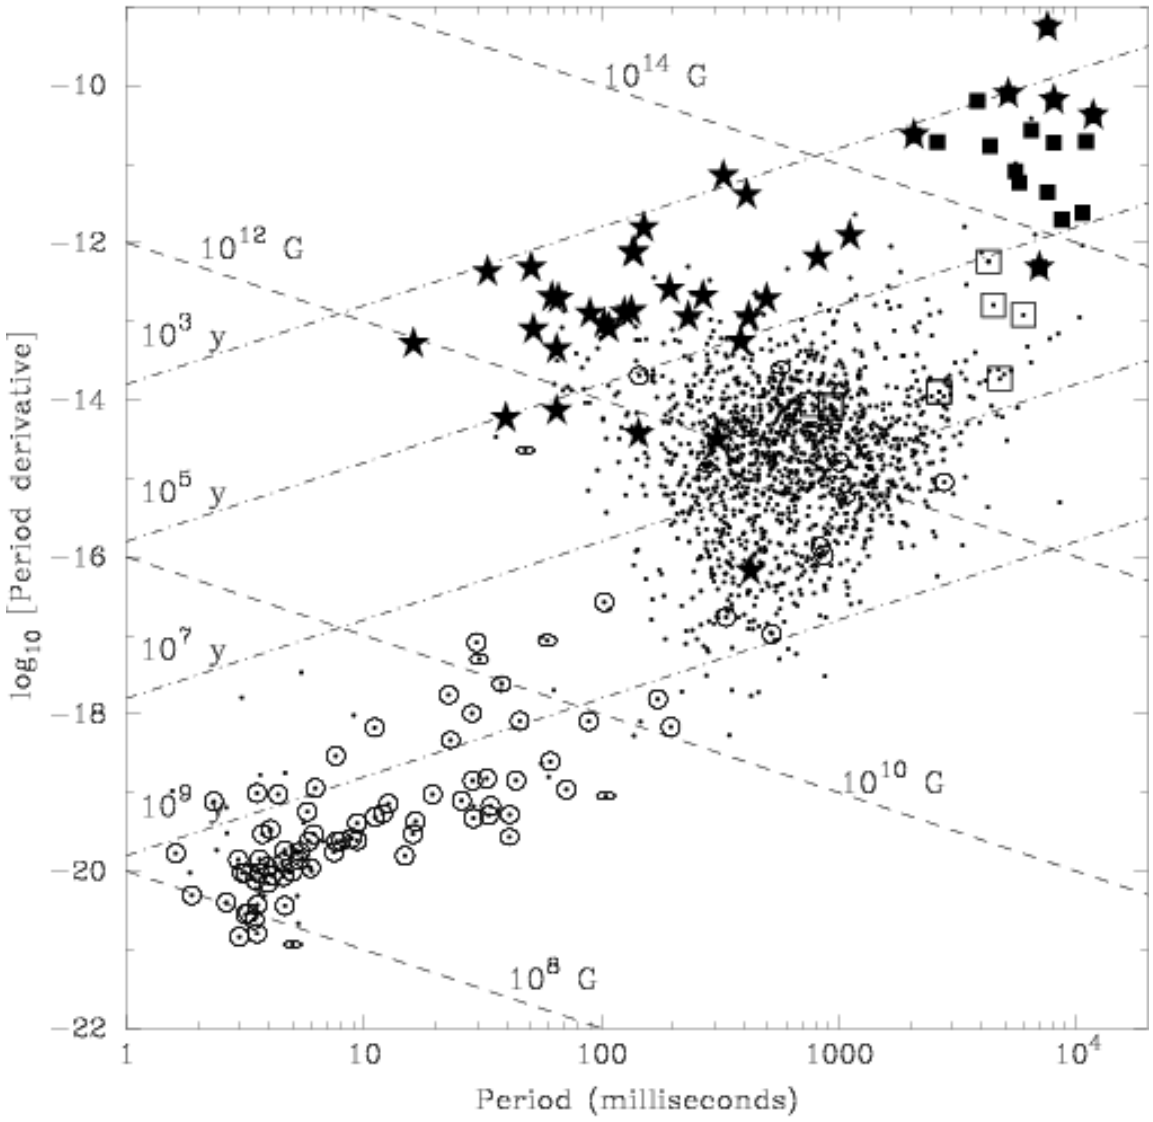

Figure 4.1: A period-period derivative plot showing the distribution of known pulsars. The plot shows isolated pulsars as dots, binary pulsars as circles/ellipses, magnetars as filled squares, rotating radio transients as open squares, and supernova remnant associations as filled stars. Also shown are lines of constant characteristic age (in years) and magnetic field strength (in Gauss). Plot taken from [54].

4.2 Emission mechanisms

There are several different mechanisms whereby a spinning neutron star can emit CW gravitational waves. We established in Chapter 2 that gravitational waves are quadrupolar and thus require some non-axisymmetric feature; this can be produced in the form of some non-axisymmetric mass distribution (a "mountain" or bump in the crust of a neutron star) or non-axisymmetric oscillations in the neutron star itself. This section will review these mechanisms and the gravitational wave strain each can be expected to produce.

4.2.1 Non-axisymmetric mass

One way to introduce non-axisymmetry is to imagine a bump or mountain on the crust of the spinning neutron star. The gravitational wave flux emitted by a spinning neutron star is [15]

$$L_{grav} = \frac{G}{5c^5} \langle \dddot{I}_{\mu\nu} \dddot{I}^{\mu\nu} \rangle \tag{4.13}$$

where the brackets indicate a time average, $I_{\mu\nu}$ is the moment of inertia tensor, G is the gravitational constant, and c is the speed of light. Imagine a spinning neutron star with a non-axisymmetric bump, rotating about the z-axis with angular frequency ω, and an observer located at an inclination angle ι with respect to the z-axis. The second derivative of the moment of inertia tensor of a rotating neutron star is then given by [56]

(4.14)

$$\ddot{I}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -16\pi^2 f^2 (I_{xx} - I_{yy}) \cos(4\pi\omega t) & -32\pi^2 f^2 (I_{xx} - I_{yy}) \sin(4\pi\omega t) & 0 \\ 0 & -32\pi^2 f^2 (I_{xx} - I_{yy}) \sin(4\pi\omega t) & 16\pi^2 f^2 (I_{xx} - I_{yy}) \cos(4\pi\omega t) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Combining Eq.(4.13) and Eq.(4.14), we have

$$L_{grav} = \frac{32G}{5c^5}\omega^6 (I_{xx} - I_{yy})^2 = \frac{32G}{5c^5}\omega^6 I_{zz}^2 \left(\frac{I_{xx} - I_{yy}}{I_{zz}}\right)^2. \tag{4.15}$$

We can see that there will be no associated gravitational wave emission without non-axisymmetry–if we have $I_{xx} = I_{yy}$, then $L_{grav} = 0$. We can define a measure of this non-axisymmetry, called the ellipticity ϵ:

$$\epsilon \equiv \frac{I_{xx} - I_{yy}}{I_{zz}}. \tag{4.16}$$

Then Eq.(4.15) becomes

$$L_{grav} = \frac{32G}{5c^5}\omega^6 I_{zz}^2 \epsilon^2. \tag{4.17}$$

The gravitational waves will actually be emitted at frequency $f = 2\left(\frac{\omega}{2\pi}\right)$. The quadrupolar nature of the waves makes them symmetric with respect to 180° rotations, and the neutron star is nearly transparent to gravitational waves. Thus for every rotation of the neutron star the gravitational waves will have swept past twice. The gravitational wave strain $h_{\mu\nu}$ of these waves is given by [14]

$$h_{\mu\nu} = \frac{2G}{c^4 d}\ddot{I}_{\mu\nu}, \tag{4.18}$$

where d is the distance to the neutron star. Putting together equations 4.14, 4.16, and 4.18, gives

$$\frac{32\pi^2\omega^2\epsilon I_{zz}}{d}\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\cos(4\pi\omega t)(1+\cos^2(\iota)) & -2\sin(4\pi\omega t)\cos(\iota) & 0 \\ 0 & -2\sin(4\pi\omega t)\cos(\iota) & \cos(4\pi\omega t)(1+\cos^2(\iota)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{4.19}$$

recalling that ι is the inclination angle between the observer and the z-axis. For the optimum case where $cos(\iota) = 1$, and using $f = 2\left(\frac{\omega}{2\pi}\right)$, the gravitational waves will have amplitude

$$h_0 = \frac{4\pi^2 G}{c^4}\frac{I_{zz} f^2}{d}\epsilon. \tag{4.20}$$

The maximum ellipticity ϵ_{max} is dependent on the star's magnetic field, composition, and structure, which are characterized in part by its breaking strain σ. The relationship between the two is [10]

(4.21)

$$\epsilon_{max} \approx 3.4\times10^{-7}\left(\frac{\sigma}{0.01}\right)\left(\frac{M_{NS}}{1.4M_\odot}\right)^{2.2}\left(\frac{R}{10km}\right)^{4.26}\left[1+0.7\left(\frac{M_{NS}}{1.4M_\odot}\right)\left(\frac{10km}{R}\right)\right]^{-1}$$

where M_{NS} is the mass of the neutron star and R its radius. Thus, for a maximum ellipticity of $\approx 2\times10^{-6}$, the gravitational wave strain will follow [10]

$$h_0 \leq 2\times10^{-24}\left(\frac{\epsilon}{2\times10^{-6}}\right)\left(\frac{I_{zz}}{10^{38}kg\cdot m^2}\right)\left(\frac{f}{1kHz}\right)^2\left(\frac{r}{1kpc}\right)^{-1} \tag{4.22}$$

4.2.2 Oscillations

Rapidly spinning neutron stars can support what are known as r-mode oscillations [57]. R-mode instabilities can be generated in a fluid star if the mode of oscillation is counter-rotating in the star's rotation frame, but rotates in the same sense as the star when seen by a distant observer. This occurs if the counter-rotation rate of the mode is lower than the rotation rate of the star. For example, consider as before a neutron star rotating with frequency ω. In the co-rotating frame, an r-mode oscillation is moving at $-\omega/3$. Seen from a non-rotating reference frame, this oscillation is moving at $2\omega/3$ in the direction of the rotation. Any gravitational waves emitted by this star would then carry away positive angular momentum, imparting negative angular momentum to the star. In the co-rotating frame, this negative momentum helps the instability to grow. This can lead to relatively long-lived r-mode oscillations, generating relatively stable gravitational waves [15]; the Coriolis force due to the star's rotation acts as a restoring force [58] and they are eventually damped by viscosity. The maximum amplitude of such oscillations is dependent on a neutron star's equation of state and the mechanisms by which it spins down; estimates are based on spin-down for isolated neutron stars, and x-ray flux for neutron stars in binary systems [10]. For low viscosity, saturation amplitudes can be large, leading to very rapid spindown, but for high viscosity, the spindown could be slow and well suited to the time scales of our search [59]. The most recent estimates range from a maximum as large as 0.14 for binary systems to as small as 5.1×10^{-6} for isolated neutron stars [60].

The frequency of gravitational waves from r-mode oscillations is related to the star's angular velocity ω via spherical harmonics. The frequency f of r-mode gravitational waves is given to lowest order by [15]

$$f = \frac{\Omega}{2\pi} = \frac{(l+2)(l-1)\omega}{2\pi(l+1)}, \tag{4.23}$$

so, for example, for the $l = 2$ harmonic, $f = 2\omega/3\pi = \frac{4}{3}f_{rot}$, where f_{rot} is the rotational frequency of the star. Higher order corrections depend on the star's equation of state and are highly dependent on the relativistic factor M/R, where M is a star's mass and R its radius [61].

4.2.3 Accretion

Some pulsars in binary systems will accrete matter from their companions. A pulsar's strong magnetic field can guide accreting matter to "hot spots" on the pulsar's surface, which can lead to crustal deformations–the bumps and mountains discussed earlier. In particular, accretion can maintain these in equilibrium against gravitationally driven annealing, leading to sustained gravitational waves [62]. For binary neutron stars it is also possible that the spin-up caused by accretion is balanced by a spin-down due to angular momentum carried away in emitted gravitational waves–this is known as the torque-balance limit [63]. The x-ray luminosity of a neutron star is a proxy for its accretion rate, so the gravitational wave amplitude from this torque-balance limit can be estimated from [10]

$$h_0 \approx 5 \times 10^{-27} \left(\frac{300\text{Hz}}{f_{rot}}\right)^{1/2} \left(\frac{F_X}{10^{-8}\text{erg} \cdot \text{cm}^{-2} \cdot \text{s}^{-1}}\right)^{1/2}. \tag{4.24}$$

where f_{rot} is again the spin frequency of the neutron star and F_X is the x-ray flux at the Earth. Assuming a torque-balance limit, the gravitational wave strain at the Earth for Scorpius X-1, the brightest known source of x-rays in the sky, is [64]

$$h_0 \approx 3 \times 10^{-26} \left(\frac{f_{GW}}{540\text{Hz}} \right)^{-1/2}. \tag{4.25}$$

where f_{GW} is the gravitational wave frequency, generally taken to be twice the start's rotation frequency.

In this way accretion is very welcome for CW gravitational waves, since it both breaks axisymmetries, producing non-zero quadrupole moments, and spins pulsars up, keeping their gravitational wave emissions in LIGO's most sensitive frequency band.

4.3 Target selection

This section will describe the criteria used to identify interesting targets for a CW gravitational wave search, and then describe the globular cluster NGC6544, which was selected for the search described in Chapters 6 and 7, in terms of these criteria.

4.3.1 Criteria

Multipole radiation in general induces a power law for spindown [47]:

$$\dot{\omega} = -K\omega^n \tag{4.26}$$

where K is a proportionality constant and the exponent n is known as the braking index and $n = 2m + 1$, where $m = 1$ indicates a dipole and $m = 2$ for a quadrupole. Thus $n = 3$ for a star emitting purely electromagnetic dipole radiation and $n = 5$ for a star emitting purely gravitational radiation. Differentiating Eq.(4.26) results in

$$\ddot{\omega} = -Kn\omega^{n-1}\dot{\omega}. \tag{4.27}$$

Multiplying both sides by ω then gives

$$\omega\ddot{\omega} = -Kn\omega^n\dot{\omega} = n\dot{\omega}^2 \tag{4.28}$$

and so

$$n = \frac{\omega\ddot{\omega}}{\dot{\omega}^2}. \tag{4.29}$$

The braking index n is thus a relation between the star's rotation frequency and its derivatives [15]. Assuming the braking index is stable (not time-dependent), we can also relate it to the age of the neutron star. Let $t_0 = 0$ be the time of the neutron star's birth, at which time it was spinning at some frequency ω_0. Then, integrating Eq.(4.26) from 0 to the present time t (with spin frequency ω), we have

$$t = \frac{1}{K(n-1)}\left(\frac{1}{\omega^{n-1}} - \frac{1}{\omega_0^{n-1}}\right). \tag{4.30}$$

But we know from Eq.(4.26) that $K = -\dot{\omega}/\omega^n$, and making that substitution into Eq.(4.30) gives

$$t = -\frac{1}{n-1}\left(\frac{\omega}{\dot{\omega}}\right)\left[1 - \left(\frac{\omega}{\omega_0}\right)^{n-1}\right]. \tag{4.31}$$

Now assume that the neutron star has been spinning down for a long time, such that $\omega_0 \gg \omega$, and the term in square brackets can be neglected. This t is then called the characteristic neutron star age τ:

$$\tau = -\frac{1}{n-1}\left(\frac{\omega}{\dot{\omega}}\right). \tag{4.32}$$

This ratio relation also holds for the gravitational wave frequency $f \propto f_{rot}$ (where $f_{rot} = \omega/2\pi$). The energy that a spinning neutron star is radiating away puts a bound on the flux in gravitational waves it can emit. Recall from earlier that we had, for the gravitational wave flux,

$$L_{gw} = \frac{32G}{5c^5} I_{zz}^2 \epsilon^2 (\pi f)^6 \leq -\left(\frac{dE}{dt}\right)_{rot} = -\frac{d}{dt}\left(\frac{\pi^2 I_{zz} f^2}{2}\right). \tag{4.33}$$

Solving for the ellipticity ϵ gives

$$\epsilon \leq \sqrt{\frac{5c^5}{32\pi^4 G I_{zz}} \frac{-\dot{f}}{f^5}} \tag{4.34}$$

which becomes, when we substitute in our age relation Eq.(4.32),

$$\epsilon \leq \sqrt{\frac{5c^5}{32\pi^4 G I_{zz}(n-1)\tau f^4}}. \tag{4.35}$$

We can then substitute this new age-based ϵ into Eq.(4.20) to get an age-based limit on the strain amplitude h_0.

$$h_{age} \leq \frac{1}{d}\sqrt{\frac{5GI_{zz}}{2c^3\tau(n-1)}} \tag{4.36}$$

Note that this limit does not depend on the frequency. So from this limit we can see that, even without knowing what frequencies we might search over, we should attempt to select young sources (small τ) that are close to the Earth (small d) to increase the strain amplitude of any potential signals.

4.3.2 NGC6544

This brings us to the target selected for the gravitational wave search which will be described fully in Chapters 6 and 7: the globular cluster NGC6544. The first question that must be asked is: why globular clusters?

The answer lies in the second parameter we can adjust–the characteristic age τ. This may seem counterintuitive, since globular clusters are known to be home to older, not younger, stars. However, the key here again is something touched on earlier in this chapter: the mechanisms which create the sources of gravitational waves. What might first induce an r-mode oscillation? Create a non-axisymmetric feature on a neutron star's crust? The driver, as discussed earlier, is accretion. Globular clusters have incredibly high stellar densities–luminosity density is a good proxy for neutron star density, and core-collapsed globular clusters can reach $10^6 L_\odot/pc^3$, and those surveyed by the Hubble Space Telescope have returned values as high as 10^8 [65]. With such high stellar densities, the probability that a neutron star has been recently perturbed–either by debris, or by another neutron star–goes up. If we imagine an event that spins up a neutron star, such as accretion or a collision, we can see that by the assumptions that went into Eq.(4.32) such an event effectively "resets" the neutron star's characteristic age. In this sense, globular clusters can host "young" stars.

Once we have settled on core-collapsed globular clusters, we'll want the closest ones possible, since the age-based limit on amplitude is proportional to $1/d$. The closest of all is NGC6397, a mere 2.0 kpc away, but the next-best is NGC6544 at 2.7 kpc. NGC6544 has one advantage over NGC6397, however: its sky location. The next chapter will discuss the doppler modulation of the frequency of a signal due to the detector's motion, i.e., the Earth's rotation about its axis and its revolution about

the sun. The Earth's rotation will smear out a prospective signal over a range of frequencies, and corrections must be applied to counter these Doppler effects. By the same token, these corrections will smear out signals which travel with the detector, such as the instrumental noise lines discussed in Chapter 3. The Doppler corrections will be largest for a target at low latitude, and the larger the corrections needed the larger the deleterious effect on instrumental noise lines will be. For this reason we favor a target at low latitude. NGC6544 is at close to $-25°$, where as NGC6397 is closer to $-54°$. For this reason we choose NGC6544, shown in Figure 4.2, as the target of our search for gravitational waves.

Figure 4.2: A Hubble image of the globular cluster NGC6544, the target of our search for gravitational waves.

CHAPTER V

Barycentric Resampling

Searches for continuous-wave (CW) gravitational-wave sources can be extremely computationally intensive. This chapter will introduce barycentric resampling, an algorithm used to greatly reduce the computational cost of such a search. The first section will discuss Doppler modulation and its effect on computational cost–the issue barycentric resampling resolves. Next follows a description of the signal model used in our search, and the final section will focus on the barycentric resampling method itself and the computational gains it accomplishes.

5.1 Doppler modulation

As discussed in the previous chapter, CW signals are quasi-monochromatic in nature–they are characterized most importantly by some nominal frequency f with very small time derivatives. The LIGO detectors, however, are not stationary with respect to a given source of CW waves. Located on the Earth's surface, they are affected diurnally and semi-diurnally by the planet's rotation on its axis, and to a much larger degree, annually by the planet's orbital motion about the Sun. Thus a signal from an astrophysical source, upon arrival at the detector, will have undergone Doppler modulation of its apparent frequency in the detector reference frame:

$$f_{Det} = f_0 \left(1 + \frac{\vec{v} \cdot \hat{r}}{c}\right). \tag{5.1}$$

Here f_{Det} is the frequency at the detector, f_0 is the source frequency, $\vec{v}$ is the detector velocity, $\hat{r}$ the unit vector pointing towards the source, and c the speed of light, since gravitational waves travel at this speed like their electromagnetic counterparts. This modulation requires correction in reconstructing the parameters of the gravitational wave signal arriving from the source.

A straightforward approach is to eliminate the earth's motion through a shift of reference frame to the solar system barycenter (SSB) frame. This is accomplished by shifting from detector time t to solar system barycenter time t_b, which is shifted by t_m: the time needed for a gravitational wavefront to travel from the detector to the solar system barycenter. The new time coordinate is thus

$$t_b = t + t_m = t + \frac{\hat{n_0} \cdot \vec{r}_d(t)}{c}, \tag{5.2}$$

where $\hat{n_0}$ is the unit vector pointing to the source in the SSB frame, and $\vec{r}_d(t)$ is the detector's time-dependent position in the same SSB frame. The phase of a gravitational wave signal in this frame is given by [66]

$$\Psi(t) = \varphi_0 + 2\pi \sum_{k=0}^{s} f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \hat{n_0} \cdot \vec{r}_d(t) \sum_{k=0}^{s} f_0^{(k)} \frac{t^k}{k!} \tag{5.3}$$

where φ_0 is the initial phase at $t = 0$, $f_0^{(k)}$ is the k^{th} derivative of the frequency f, and s is the highest order of the expansion. The third term on the right-hand side of Eq.(5.3) is the correction in the phase due to the detector's motion with respect to the source, and this correction has a dependence on the frequency of the signal.

Therein arises the large computational cost of a conventional barycentering method. In general, we will have from our detector some time series of data $x(t)$. To search for CW signals of some characteristic frequency f, it will be expedient to take a Fourier transform and perform our analysis in the frequency domain. In a directed search, like the one described in Chapters 6 and 7, we search at a single sky location but over a broad range of frequencies. But the Fourier transform of an SSB-shifted signal will look like

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{i\psi(f,\dot{f}^k,t)} dt \tag{5.4}$$

where the exponential term represents the frequency-dependent phase correction. This correction must then be calculated for every set of frequencies and frequency derivatives searched; this task quickly grows monumentally computationally expensive as the length of the time series increases. Barycentric resampling is the method we will use to overcome this computational limitation. Before turning to its detailed implementation, we will review the signal model for CW gravitational waves detected by LIGO interferometers.

5.2 The LIGO signal model

The gravitational wave strain detected at a LIGO interferometer can be written as [66]

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t) \tag{5.5}$$

where t is the time in the detector frame, $h_+(t)$ and $h_\times(t)$ are the respective strain amplitudes of the plus-polarized and cross-polarized components of the incoming

gravitational wave, and F_+ and $F_\times(t)$ are the respective plus- and cross-polarized beam-pattern functions of the interferometric detector, given by

$$F_+(t) = \sin(\zeta)[a(t)\cos(2\psi) + b(t)\sin(2\psi)] \tag{5.6}$$

$$F_\times(t) = \sin(\zeta)[b(t)\cos(2\psi) - a(t)\sin(2\psi)] \tag{5.7}$$

where ζ is the angle between the two detector arms (90° at both LIGO detectors), ψ is the polarization angle of the wave, and the detector response functions $a(t)$ and $b(t)$ are functions of source position (most often expressed in terms of right ascension α and declination δ), as well as detector position and orientation. Their full forms are given in [66].

These signals are likely to be buried among statistical noise, such that the time-series $x(t)$ measured at a given detector will be of the form

$$x(t) = n(t) + h(t) \tag{5.8}$$

where $n(t)$ represents the noise and $h(t)$ the (possibly-present) signal. In the ideal case where the noise is stationary and Gaussian, we can define a likelihood function Λ which we maximize over the parameters of the signal. For our signal the log-likelihood function $\ln(\Lambda)$ is given by [66]

$$\ln(\Lambda) = (x|h) - \frac{1}{2}(h|h), \tag{5.9}$$

where the inner product is defined by

$$(x|y) \equiv 4 * \mathbb{R}\left[\int_0^\infty \frac{\tilde{x}(f)\tilde{y*}(f)}{S_h(f)} df\right]. \tag{5.10}$$

Here $\mathbb{R}$ indicates the real part, $\tilde{x}$ indicates the Fourier transform and $*$ the complex conjugate, and $S_h(f)$ is the one-sided spectral density of the detector noise at the frequency f. In the approximation of stationary noise, $S_h(f)$ will be constant over an observation time T_0. Introducing the scalar product

$$(x||y) \equiv \frac{2}{T_0}\int_{-T_0/2}^{T_0/2} x(t)y(t)dt, \tag{5.11}$$

the log-likelihood function becomes approximately

$$\ln(\Lambda) \cong \frac{T_0}{S_h(f)}\left[(x||h) - \frac{1}{2}(h||h)\right]. \tag{5.12}$$

Maximizing this log-likelihood function over the extrinsic and unknown signal parameters results in the $\mathcal{F}$-statistic, which is defined by [66]

$$\mathcal{F} = \frac{4}{S_h(f)T_0}\frac{B\left|F_a\right|^2 + A\left|F_b\right|^2 - 2C\mathbb{R}(F_a F_b^*)}{D}. \tag{5.13}$$

Using the scalar product of Eq.(5.11) and the $a(t)$ and $b(t)$ of Eq.(5.6), $A = (a||a)$, $B = (b||b)$, $C = (a||b)$, and $D = A \cdot B - C^2$. F_a and F_b are integrals defined by

$$F_a(f) \equiv \int_{-T_0/2}^{T_0/2} a(t)x(t)e^{i\Phi_s(t)}dt \tag{5.14}$$

$$F_b(f) \equiv \int_{-T_0/2}^{T_0/2} b(t)x(t)e^{i\Phi_s(t)}dt \tag{5.15}$$

Here $x(t)$ is again the time-series of data from our detector, and $\Phi_s(t)$ is the ($k = 1, ..., s$) terms of Eq.(5.3)–that is, the portion of the phase correction due to the frequency derivatives, defined in the next section. In the following section we will see how this separation can help reduce computational cost.

5.3 The barycentric resampling method

Our earlier expression for the gravitational wave phase (Eq.(5.3)) can be split up into three parts: the initial phase φ_0, the $k = 0$ terms of its two sums, and the remainder. Recall that our time shift to the solar system barycenter had

$$t_m = \frac{\hat{n_0} \cdot \vec{r}_d(t)}{c} \tag{5.16}$$

and it is clear that for $k = 0$, Eq.(5.3) becomes

$$\begin{aligned} \Psi(t)^{(0)} &= \varphi_0 + 2\pi f_0^{(0)} \frac{t^1}{1!} + \frac{2\pi}{c} \hat{n_0} \cdot \vec{r}_d(t) f_0^{(0)} \frac{t^0}{0!} \\ &= \varphi_0 + 2\pi f t + 2\pi \frac{\hat{n_0} \cdot \vec{r}_d(t)}{c} f \\ &= \varphi_0 + 2\pi f (t + t_m). \end{aligned} \tag{5.17}$$

Then we can define the remaining summation terms as

$$\Phi_s(t) = 2\pi \sum_{k=1}^{s} f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \hat{n_0} \cdot \vec{r}_d(t) \sum_{k=1}^{s} f_0^{(k)} \frac{t^k}{k!}. \tag{5.18}$$

This is the Φ_s that appears in Eqs.(5.14) and (5.15). We can now make the barycentric time-shift substitution (Eq.(5.2)) into those two integrals. Taking the derivative with respect to t on both sides of the substitution gives

$$\frac{dt_b}{dt} = 1 + \frac{dt_m}{dt}, \tag{5.19}$$

but we know from Eq.(5.16) that

$$\frac{dt_m}{dt} = \frac{\hat{n_0} \cdot v_d(t)}{c}, \tag{5.20}$$

where $v_d(t)$ is the detector's velocity in the SSB frame–and dt_m/dt is therefore the Doppler shift of the source with respect to the detector in the SSB frame. The maximum Doppler shift experienced by an Earth-bound detector is of order 10^{-4}, which is to say $dt_m/dt \ll 1$, and therefore, by Eq.(5.19), $dt_b \approx dt$. This means that Eqs.(5.14) and (5.15) now become [15]

$$F_a(f) = \int_{-T_0/2}^{T_0/2} a(t_b)x(t_b)e^{-2\pi i f t_b}e^{i\Phi_s(t_b)}dt_b \tag{5.21}$$

$$F_b(f) = \int_{-T_0/2}^{T_0/2} b(t_b)x(t_b)e^{-2\pi i f t_b}e^{i\Phi_s(t_b)}dt_b \tag{5.22}$$

These are simply the Fourier transforms of the data stream $x(t)$ and the detector response, multiplied by a phase $e^{i\Phi_s(t_b)}$ which is independent of the nominal frequency, but dependent on its derivatives. This calculation can be performed much more computationally efficiently by fast FFT routines.

Processing a large frequency band at once may still be computationally infeasible, but we can also make use of heterodyning to split a target band into many smaller sub-bands, which can be analyzed in parallel on a distributed computing cluster. Take the Fourier transform of our time-series $x(t)$:

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt \tag{5.23}$$

Then the Fourier transform of a complex time-series $x_h(t) = x(t)e^{-2\pi i f_h t}$ is

$$\begin{aligned} \tilde{x}_h(f) &= \int_{-\infty}^{\infty} x(t) e^{-2\pi i f_h t} e^{-2\pi i f t} dt \\ &= \int_{-\infty}^{\infty} x(t) e^{-2\pi i (f+f_h) t} dt \\ &= \tilde{x}(f + f_h). \end{aligned} \tag{5.24}$$

By multiplying $x(t)$ by $e^{-2\pi i f_h t}$ we effectively shift all the frequencies in $x(t)$ by f_h.

The computational load can also be lightened by downsampling the data. In its simplest form, this consists of using only every j^{th} data point in order to downsample the time series by a factor of j. This will, however, change the Nyquist frequency of our time series, so care must be taken to avoid aliasing effects. This can be done by applying a low-pass filter with a sharp fall-off around the new Nyquist frequency, which will be

$$f_{Ny}^{new} = \frac{f_{Ny}^{old}}{j}. \tag{5.25}$$

The new time-series will have a sampling time $\Delta t = 1/(2f_{Ny}^{new})$. Consider, then, $F_a(f)$ for such a downsampled, heterodyned time-series ($F_b(f)$ is completely analogous):

$$F_a(f + f_h) = \int_{-T_0/2}^{T_0/2} a(t) x(t) e^{-2\pi i (f+f_h)(t+t_m)} e^{-i\Phi_s(t)} dt. \tag{5.26}$$

Lurking within the first exponent is our complex time-series $x_h(t)$:

$$x(t)e^{-2\pi i(f+f_h)(t+t_m)} = x_h(t)e^{-2\pi i f_h t_m}e^{-2\pi i f(t+t_m)} \tag{5.27}$$

We can get the integral back to its original form by then making the further substitution $z(t) = x_h(t)e^{-2\pi i f_h t_m}$. Now Eq.(5.26) has become

$$F_a(f+f_h) = \int_{-T_0/2}^{T_0/2} a(t)z(t)e^{-2\pi i f(t+t_m)}e^{-i\Phi_s(t)}dt \tag{5.28}$$

and we can now readily make the switch to barycenter time t_b.

$$F_a(f+f_h) = \int_{-T_0/2}^{T_0/2} a(t_b)z(t_b)e^{-2\pi i f(t_b)}e^{-i\Phi_s(t_b)}dt_b \tag{5.29}$$

In practice, our time-series will be discrete rather than continuous. The expression for a time series with N points is

$$F_a(f+f_h) = \sum_{k=1}^{N} a(t_b^k)z(t_b^k)e^{-2\pi i f t_b^k}e^{-i\Phi_s(t_b^k)}\Delta t_b \tag{5.30}$$

where t_b^k is the k^{th} point in the time-series as measured in the SSB frame, and $\Delta t_b = t_b^{k+1} - t_b^k$.

Thus, beginning with a time-series of data $z(t^k)$, linear in the detector frame, which is irregularly sampled in the solar system barycenter frame, we use $t_b^k = t^k + t_m$ to calculate $T^k(t_b^k)$, the set of detector times corresponding to a time-series linear in the solar system barycenter frame. Interpolating using our known time-series, we can calculate $z(T^k(t_b^k)$, which corresponds to the $z(t_b^k)$ in Eq.(5.30). The $a(t^k)$ and $b(t^k)$ are dealt with analogously, and the phase correction $e^{i\Phi_s(t_b^k)}$ can be calculated from its definition in Eq.(5.18). The factor of $\hat{n_0}\cdot\vec{r}_d(t)$ in Eq.(5.18) is already encoded in the time shift:

$$\hat{n_0} \cdot \vec{r}_d(t) = t_m \cdot c = (t_b^k - T^k(t_b^k)) \cdot c. \tag{5.31}$$

A completely analogous method will produce $F_b(f+f_h)$, and then from these two, the $\mathcal{F}$-statistic can be calculated (for more details of the implementation, consult [67]).

Figure 5.1 shows a comparison of the $2\mathcal{F}$ values calculated for a given signal under the standard Doppler demodulation ("LALDemod") and the new resampling method. It shows good agreement between the two methods, both in the presence of a signal (high $2\mathcal{F}$) and in its absence (low $2\mathcal{F}$). Another test of the resampling method was the recovery of hardware injections from S6.

A hardware injection is a fake signal inserted into LIGO data by physically moving the optics of the interferometer to simulate the arrival of a gravitational wave. It is defined in opposition to a software injection (see Chapter 6), which is a fake signal added to the data in post-processing. Ten pulsar-like hardware injections were inserted during the S6 science run, with varying parameters. The resampling code was used to search for a subset of these in the full S6 data alongside the standard search code. The results of one such search are shown in Figure 5.2. There is very good agreement between the results of the two methods.

How much faster is the resampling method? Assuming we have N data points, and N_{SFT} SFTs (an SFT (Short Fourier Transform) is a half-hour chunk of detector data, the standard data format used in LIGO CW searches), the number of operations required for computation is roughly $30 \cdot N_{SFT} \cdot N$ [15]. Resampling, on the other hand, needs approximately 30 operations to reach $F_a(t)$ and $F_b(t)$. A Fourier transform is of order $N \ln N$, so for the same number of data points and SFTs the total number of operations is $(30 + \ln N) \cdot N$. The ratio is then, to first order,

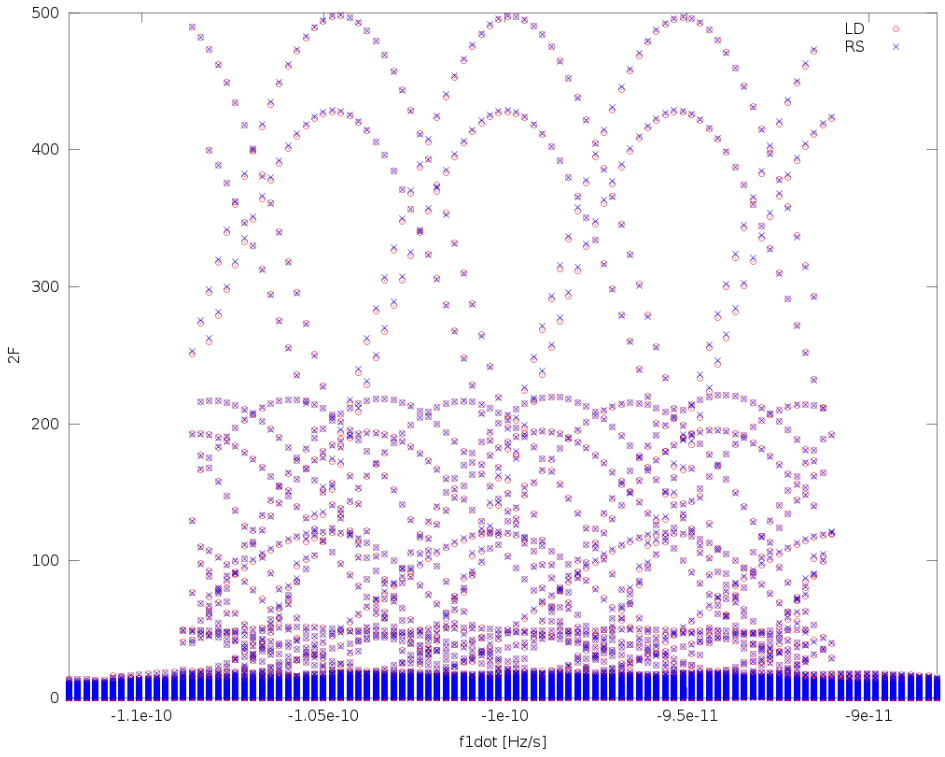

Figure 5.1: A comparison of $2\mathcal{F}$ values calculated versus spindown $\dot{f}$. The red circles are the values calculated using the standard Doppler demodulation code and the blue exes are the values calculated using the new resampling code. There are multiple values per $\dot{f}$ because multiple frequency bins were searched over, but the plot does not display this third dimension.

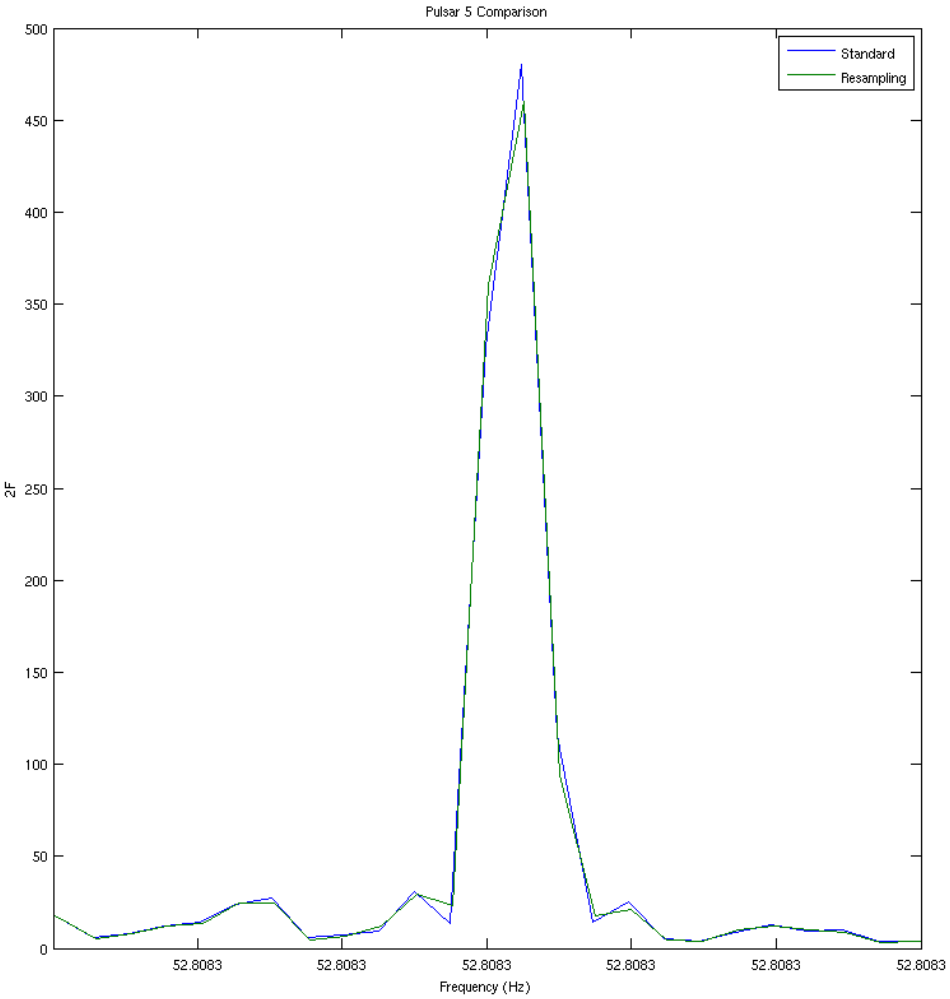

Figure 5.2: A comparison of $2\mathcal{F}$ values versus frequency. The blue curve represents the results calculated by the standard Doppler demodulation code; the green curve shows the results of the resampling code. The signal under examination was a hardware injection at frequency $f = 52.80832$ Hz.

$$\frac{N_{TOT}^{Standard}}{N_{TOT}^{Resamp}} \approx \frac{N_{SFT}}{\ln N}. \tag{5.32}$$

Thus we would expect resampling to be much quicker for long observation times. Figure 5.3 bears this out, showing the results of timing tests run using resampling and the standard demodulation technique described in the first section, over observation times ranging from 30 hours to 217 hours, which is comparable to the observation times of the search described in Chapters 6 and 7. The results are given in seconds per template per SFT. A template is a given combination of source frequency, frequency derivatives, and source sky location, which picks out a location in the parameter space of the search (see Chapter 6). For a time duration close to that of the search described in Chapters 6-7, the resampling method is more than an order of magnitude faster.

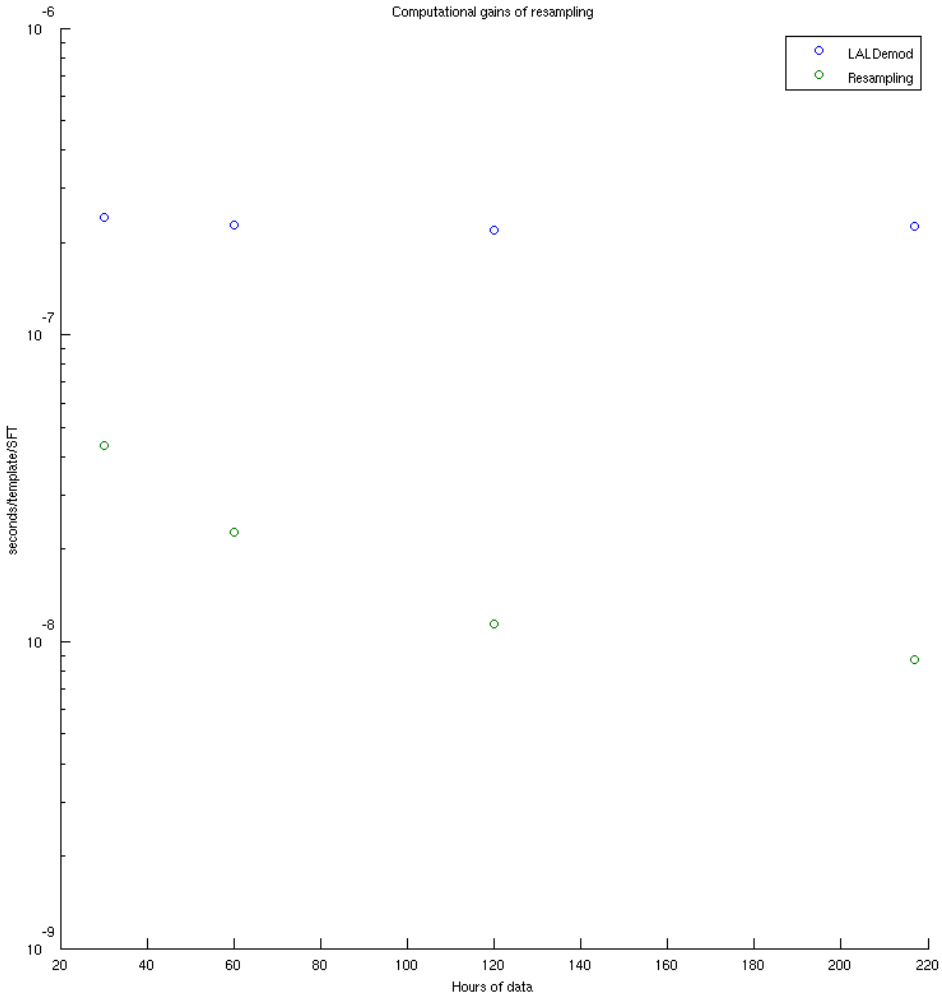

Figure 5.3: This plot compares the speed of the resampling and standard demodulation approaches on a quarter-hertz frequency band with varying observation times. The results are in seconds per template per SFT. A large observation time allows resampling to be much more computationally efficient.

CHAPTER VI

Search: Setup and Methods

This chapter and the chapter that follows discuss a search, carried out in data from LIGO's sixth science run (S6), for continuous-wave (CW) gravitational radiation from the nearby globular cluster NGC6544, which was introduced in Chapter 4. This chapter will describe the setup of the search and the methods used to execute it. We will describe an iterative method for arriving at a parameter space to search over, starting from an age-based limit on the expected strain of gravitational wave signals, and then describe our methods for selecting which stretches of data to use in the search. The final section will discuss our methods for conducting the search, finding and dealing with outliers, and setting upper limits on gravitational wave strength in the absence of a signal.

6.1 Parameter Space

The parameter space of a CW search refers to the range of parameters that the search will cover, such as sky locations, frequencies and their derivatives, and the so-called "nuisance parameters" (the inclination angle ι, polarization angle ψ, and initial phase ϕ_0). Each combination of these parameters produces a template for a gravitational-wave signal, and the data is then searched for a signal matching the template.

There are three broad classes of CW searches that are carried out within LIGO–targeted searches, all-sky searches, and directed searches. A targeted search is an effort to detect gravitational waves from a known source–for example, the Crab [38] or Vela [68] pulsars. In these searches the frequencies of the gravitational-wave signals are assumed to be twice the known rotational frequencies of the stars, which have known sky locations; in some cases, informed priors can be placed on the nuisance parameters as well. An all-sky search is in some sense the opposite. It assumes little, searching over a wide range in frequencies f and spindowns $\dot{f}$, and, as its name would suggest, surveys the entire sky as well. These searches are very computationally expensive and have been performed in LIGO data using the Einstein@Home distributed computing method [69, 70], as well as the PowerFlux search pipeline [71] for isolated neutron stars and the TwoSpect method [72] for stars in binary systems. Between these two extremes sit the directed searches. A directed search has a well-defined target, for example the galactic center [73] or Cassiopeia A [74], and so is concentrated on a single sky location. However, the frequency and other characteristics of the gravitational-wave signal are unknown, and the search is conducted over as wide an array of frequency and its derivatives as is computationally feasible and astrophysically motivated, while the $\mathcal{F}$-statistic described in Chapter 5 marginalizes over the "nuisance parameters." It is to this last category that our search belongs.

These searches are computationally limited. We would prefer to search over as wide a frequency parameter space as possible, but must make compromises based on the computational resources we can expect such a search to use. This search was run on the LDAS-CIT distributed computing cluster at the California Institute of Technology, and because of competition for computing resources, we chose to limit ourselves to 1000 CPU-months.

We began from the age-based limit on the gravitational wave strain h_0 introduced in Section 4.3:

$$h_{age} \leq \frac{1}{d}\sqrt{\frac{5GI_{zz}}{2c^3\tau(n-1)}}, \tag{6.1}$$

where d is the distance to the target, I_{zz} is the moment of inertia, τ is the (assumed) neutron star age, and n the braking index defined in Chapter 4. We use a distance to NGC6544 of 2.7 kpc, and the canonical moment of inertia 10^{38} kg $\cdot$m^2, derived from the mass ($1.4M_{\odot}$) and radius (10 km) for neutron stars given in Chapter 4. For choosing a search band, we assume a neutron star age of 300 years and a braking index $n = 5$, reflecting pure quadrupolar gravitational wave emission. This limit on the gravitational wave strain is then superimposed on the LIGO sensitivity curve for the S6 run, adjusted for an initial assumed observation time of two weeks, as shown in Figure 6.1. The LIGO sensitivity curve is given by using the noise power spectral density (PSD) harmonically averaged over all of S6 data for both LIGO interferometers. A running median with a 16-Hz window was further applied to smooth the curve.

The intersection of the age-based limit and the sensitivity curve gives us a frequency band over which we can expect LIGO sensitivity to beat the age-based limit for $n = 5$. From this frequency band, we then derive search bounds on the spindown $\dot{f}$ and second spindown $\ddot{f}$, as well. Recall from Chapter 4 the relation between the braking index n and the characteristic age of the neutron star, τ:

$$\tau = -\frac{1}{n-1}\left(\frac{f}{\dot{f}}\right). \tag{6.2}$$

We can rearrange Eq.(6.2) to give a bound on the spindown:

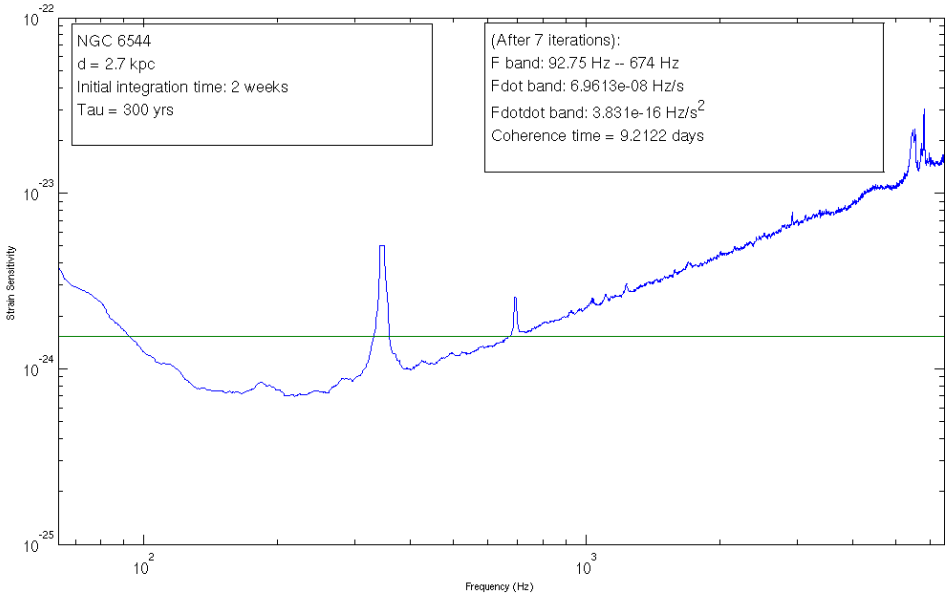

Figure 6.1: The result of our iterative process to arrive at a parameter space. The age-based limit on h_0 is the green horizontal line; the blue curve represents the upper limit of LIGO sensitivity for S6, and has been smoothed with a running median. The left text box lists the initial assumptions; the right text box the final parameter space and coherence time (observation time).

$$-\dot{f} = \frac{f}{(n-1)\tau}. \tag{6.3}$$

We do not know the frequency evolution of the target, so we take a conservative approach [15], covering both the pure magnetic dipole ($n = 3$) and gravitar ($n = 5$) cases, and look at all braking indices in the range $2 \leq n \leq 7$. This then establishes the range in spindown:

$$\frac{f}{\tau} \leq -\dot{f} \leq \frac{f}{6\tau}. \tag{6.4}$$

Similarly, from the definition $n = f \cdot \ddot{f}/\dot{f}^2$, and the same bounds on n, we can establish the range in second spindown to be

$$\frac{2\dot{f}^2}{f} \leq \ddot{f} \leq \frac{7\dot{f}^2}{f}. \tag{6.5}$$

To cover this parameter space with templates efficiently, we introduce the mismatch parameter m. The mismatch m represents the relative loss in our detection statistic (the $\mathcal{F}$-statistic) due to an offset ΔP in our template parameters from the true parameters of the signal. For our search, we take $m = 0.2$, meaning that we place templates such that, in the worst case, the closest template to the true signal will recover an $\mathcal{F}$-statistic value that is 80% of the $\mathcal{F}$-statistic that would be recovered by a template placed precisely on the true signal parameters. It is, in a sense, a measure of the maximum amount by which we are willing to be wrong in order to save computational cost.

Using this relative loss in signal strength due to an offset in signal parameters as an invariant "distance" measure, we can introduce a metric g over the parameter space. The metric's general form is complicated, taking into account all of a candidate signal's parameters, and detailed in [75]; in our case we can use a simplified metric since our directed search reduces our parameter space to frequency f and its time derivatives f^k, where k represents the k^{th} derivative. In this case we will need only the metric components g_{ff}, g_{ff^k}, and $g_{f^j f^k}$, where j and k are indices which run from 0 to the highest frequency derivative considered (in our case, 2). The resolution for each parameter is determined by the shift in that parameter required in order to shift the signal by one frequency bin [15]. The resolution scales as $\frac{1}{T^{k+1}}$ for the k^{th} derivative of frequency. These metric components, calculated using the method of [75], are then [15]:

$$g_{ff} = \frac{(\pi T)^2}{3}, \tag{6.6}$$

$$g_{ff^k} = \frac{(2\pi)^2 T^{k+2} f}{k!(k+2)(k+3)}, \tag{6.7}$$

and

$$g_{f^j f^k} = \frac{(2\pi f)^2 T^{j+k+2}}{j!k!(j+2)(k+2)(j+k+3)}. \tag{6.8}$$

where f is the frequency and T the observation time. The spacing between frequency templates, however, will be fixed due to our use of a Fast Fourier Transform algorithm in calculating the $\mathcal{F}$-statistic for multiple frequency templates. This means we cannot use the cross-diagonal terms (Eq.(6.7)) and the frequency parameter is projected out. In this projected metric, the template spacings are [15]:

$$df = 2\sqrt{\frac{m}{g_{ff}}}, \tag{6.9}$$

$$d\dot{f} = 2\sqrt{\frac{m}{\gamma_{\dot{f}\dot{f}}}}, \tag{6.10}$$

and

$$d\ddot{f} = 2\sqrt{\frac{m}{\gamma_{\ddot{f}\ddot{f}}}}, \tag{6.11}$$

where m is the mismatch, and $\gamma_{\alpha\beta} = \left(g_{\alpha\beta} - \frac{g_{f\alpha} g_{f\beta}}{g_{ff}}\right)$. Substituting our values for the metric components, we then have for our template spacings [15]

$$df = \frac{2\sqrt{3m}}{\pi}\frac{1}{T}, \tag{6.12}$$

$$d\dot{f} = \frac{12\sqrt{5m}}{\pi}\frac{1}{T^2}, \tag{6.13}$$

and

$$d\ddot{f} = \frac{20\sqrt{7m}}{\pi}\frac{1}{T^3}. \tag{6.14}$$

for a given observation time T.

We can then divide our known parameter ranges by the corresponding template spacings to obtain the total number of templates in terms of observation time. For example, for frequency we have

$$\frac{\Delta f}{df} = \frac{(\Delta f)\cdot \pi T}{2\sqrt{3m}}, \tag{6.15}$$

with similarly derived results for spindown and second spindown. We know our selected value for m, and from the timing tests discussed in Chapter 5 we already know our computational speed in seconds per template, so we can convert the total number of templates into a computation time, in terms of observation time T. We can then solve for the observation time T which satisfies our initial restraint (1000 CPU-months) on computing time. We then take this observation time T, make it our new initial estimate, and iterate the entire procedure again. As seen in Figure 6.1, after seven iterations this procedure converged to a roughly 583-Hz band (the actual search spanned from 92.5 Hz to 675 Hz) and an approximately 9.2-day observation time. The next step, then, was to find the best 9.2 days of data to use.[1]

[1]Note that restricting ourselves to the best nine-day observation time would change the sensitivity spectrum from

6.2 Data Selection

The search was carried out on data from LIGO's sixth science run (S6), which spanned July 7, 2009 to October 20, 2010; the data used was taken from LIGO's Hanford (H1) and Livingston (L1) detectors. Two different methods were used to determine which data would be searched, producing two different roughly nine-day streches; both were searched, allowing for the cross-referencing of search results between them.

The first method was to look for the most sensitive data from S6. This was done by taking nine week-long data samples from each detector spaced roughly 55 days apart, giving nine evenly spaced weeks throughout duration of S6. The data samples used are shown in Table 6.1. We chose four representative frequencies (100 Hz, 200 Hz, 400 Hz, 600 Hz) and generated strain noise power spectral densities (PSDs) in 1 Hz bands about these frequencies, using 0.01-Hz binning. The sensitivity h_{sens} was then taken to be

$$h^j_{sens} = \left(\frac{1}{\sqrt{(1/100) * \sum_{i=0}^{100} (S^i_h(f_i))^{-1}}} \right)_j \tag{6.16}$$

where $S^i_h(f)$ represents the PSD value of the i^{th} bin, at frequency f_i, and the index j runs from 1 through 4 and represents the four representative frequencies. The results are plotted in Figure 6.2; for all four frequencies the sensitivity improves with time. Based on this plot, the final weeks of S6 yielded the most sensitive data, and the data stretch chosen with this method was October 11-20, 2010 (GPS 970840605 – 971621841).

This method, however, does not take into account an important consideration:

the full S6 sensitivity curve used in this step; for practical reasons we do not reiterate the entire procedure after obtaining this best data stretch.

S6 sampling times			
Label	GPS Start	GPS End	Dates (UTC)
Week 1	931053000	931657800	Jul 8-15, 2009
Week 2	936053000	936657800	Sep 3-10, 2009
Week 3	941053000	941657800	Oct 31-Nov 7, 2009
Week 4	946053000	946657800	Dec 28, 2009-Jan 4, 2010
Week 5	951053000	951657800	Feb 24-Mar 3, 2010
Week 6	956053000	956657800	Apr 23-30, 2010
Week 7	961053000	961657800	Jun 20-27, 2010
Week 8	966053000	966657800	Aug 17-24, 2010
Week 9	971053000	971657800	Oct 14-21, 2010

Table 6.1: This table shows the weeks sampled to find the most sensitive S6 data. Times are given both in GPS and UTC calendar dates.

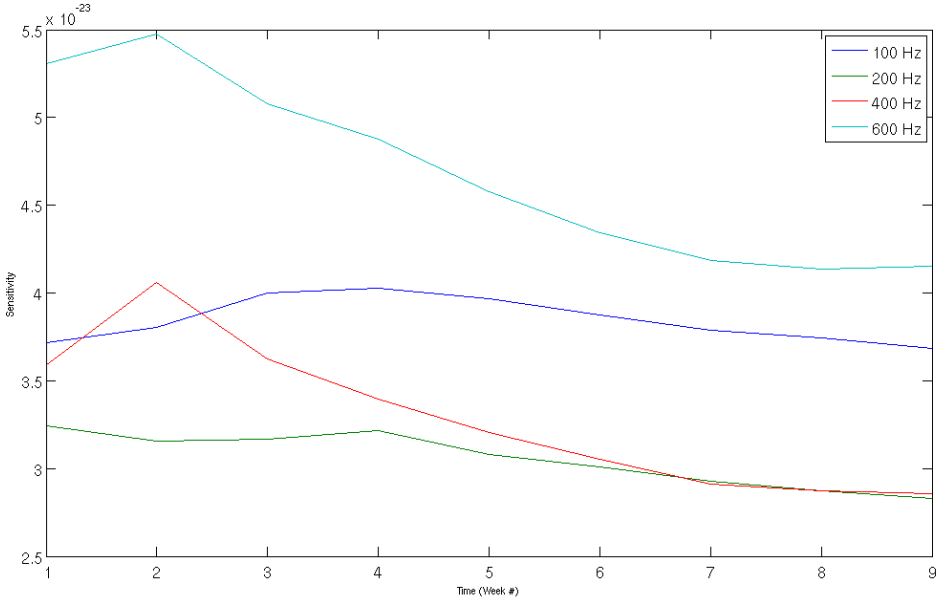

Figure 6.2: The comparative sensitivity at different points during S6. Each representative frequency is plotted with a different color. The final weeks of the science run are the most sensitive.

the detector livetime, i.e., the amount of data available from these nine days. The detectors are not always taking data, so despite an observation time of 217 hours, this data stretch produced only 202 half-hour SFTs (101 hours of data) at H1 and 172 half-hour SFTs (86 hours of data) at L1. An alternate method, which does account for livetime, is to maximize the figure of merit

$$\sum_{k,f} \frac{1}{S_h(f)} \tag{6.17}$$

where $S_h(f)$ represents the strain noise power spectral density at frequency f in the k^{th} SFT, and the sum is taken across all frequencies f in the search band and all k SFTs in a given 9.2-day observation time. This method favored a different data stretch: July 24–August 2, 2010 (GPS 964007133 – 964803598). This data stretch had slightly worse sensitivity than the first, but produced 368 H1 SFTs (184 hours of data) and 274 L1 SFTs (137 hours of data). Both data sets were searched independently.

6.3 Methods

The search was carried out with the program `ComputeFStatistic_v2`, a program used to calculate the multi-detector $\mathcal{F}$-statistic [76]. Use of the `useResamp=TRUE` flag variable enabled use of the barycentric resampling routine described in Chapter 5.

6.3.1 Search

The frequency band was broken up into 56 10-Hz bands, each containing 100 0.1-Hz subbands, as well as one 7.5-Hz band (75 subbands) and one 15-Hz band (150 subbands) to accommodate the boundaries. For each band an appropriate configuration file was used to create a directed acyclic graph (DAG) file with the

appropriate number of jobs (each job represents one 0.1-Hz subband) covering the desired frequency band with spindown and second spindown bands calculated using Eqs.(6.4) and (6.5). Each job is an instance of `ComputeFStatistic_v2`, searching over the designated parameter space, the selected SFTs, and returning three files.

The output file is a list of all templates with a joint detector $2\mathcal{F}$ ($2\mathcal{F}_J$) value ≥ 45, a threshold chosen to allow collection of templates from each band while preventing the output files from growing large enough to jeopardize the available disk space on the cluster. Each template is identified by a row in the output file with nine column entries, representing respectively the frequency f, right ascension α, declination δ, spindown $\dot{f}$, second spindown $\ddot{f}$, third spindown $\dddot{f}$ (set to zero by default for all cases; we did not search over third spindown), joint detector $2\mathcal{F}$-value $2\mathcal{F}_J$, H1 detector $2\mathcal{F}$-value $2\mathcal{F}_{\mathrm{H}}$, and L1 detector $2\mathcal{F}$-value $2\mathcal{F}_{\mathrm{L}}$. The histogram file is a histogram of all $2\mathcal{F}_J$ values returned, with no threshold applied. It can be used to provide a sanity check that the data follows a chi-squared distribution as expected [66, 76]. Finally, the "Loudest" file is a record of the template from each job (sky location and frequency variables) which produced the highest $2\mathcal{F}_J$ value; it also contains the $2\mathcal{F}$ values (both joint and individual) returned by that template.

6.3.2 Outliers

Results from a test band at 200.0 Hz, as well as a Mock Data Challenge (see Chapter 7 for more details on both), were used to establish a threshold value of $2\mathcal{F}_J$ as well as a smaller threshold on the individual detector $2\mathcal{F}$ values. The templates from the output files were subjected to these thresholds, as well as the veto condition that the $2\mathcal{F}_J$ value must be larger than both of the individual detector $2\mathcal{F}$ values; violation of this condition is caused by large $2\mathcal{F}$ values in one detector which are completely unsubstantiated in the other, and is a very good indicator of instrumental

noise lines as opposed to astrophysical sources.

From each output file, the loudest template, if any existed, that passed these thresholds and vetoes was added to a list of outliers. These outliers were then tested using time shifts and *extended looks*. In a time shift, the outliers from each data stretch (October and July-August) were evolved forwards or backwards in time, as appropriate, and searched for in the opposite data stretch, under the assumption that a true astrophysical signal should be present in both data sets since CW signals are implicitly long-lived. A set of 1000 software injections (fake signals with random frequency and nuisance variables generated with `Makefakedata_v4`, injected into detector data, and then searched for with `ComputeFStatistic_v2`) underwent the same treatment to provide a baseline threshold for signal detection.

In an extended look, each outlier was searched for in an expanded 20-day coherence time encompassing the original nine-day coherence time; the same assumption of signal continuity would predict a signal to roughly double in strength for a rough doubling of coherence time. These cases as well were tested with software injections to determine a threshold. In all cases, the searches were conducted over a parameter space envelope obtained by starting at the outlier frequency parameters $(f, \dot{f}, \ddot{f}) \pm 2$ bins, and evolving those ranges backwards or forwards in time using the extremum values of the next derivative (e.g., f evolved at maximum $\dot{f}$, $\dot{f}$ evolved at maximum $\ddot{f}$ to achieve the widest envelope possible. Figure 6.3 shows the envelopes contrasted with the uncertainty curves for the frequency parameters, given by

$$f(t) \pm \frac{1}{2}\Delta\dot{f} \cdot t, \tag{6.18}$$

where $\Delta\dot{f}$ is equal to half the bin width of a spindown bin in the initial search, t is the time, and $f(t)$ is the frequency at time t, and by

$$\dot{f}(t) \pm \frac{1}{6}\Delta\ddot{f} \cdot t, \tag{6.19}$$

where $\Delta\ddot{f}$ is half the bin width of a second spindown bin in the initial search. For frequency, the envelope is so conservative compared to the uncertainty curves that the highly magnified Figure 6.4 is provided to better illustrate the uncertainty curves. Assuming (as we do) a neutron star age of 300 years and a worst-case (for our parameter space) braking index n=2, the change in $\ddot{f}$ over a three-month lookback window (the difference between our October and July data sets) would be $2.17 \times 10^{-21}\mathrm{Hz/s}^2$. Our five-bin window (original bin $\pm$ two bins) covers $1.58 \times 10^{-16}\mathrm{Hz/s}^2$; it covers the uncertainty region easily.

Outliers detected with $2\mathcal{F}_J$ greater than the threshold established by the software injections were labeled candidates and received manual followup. The tests were not cumulative; an outlier needed only to survive one, not all of them, to be considered a candidate. The candidates and their followup are discussed in Chapter 7.

6.3.3 Upper Limits

The highest $2\mathcal{F}$ value in each 0.1-Hz subband, recorded in that job's "Loudest" file, is used by `lalapps_ComputeFStatAnalyticMonteCarloUpperLimit`, a semi-analytic program detailed in [13], in conjunction with the analyzed SFTs, to make a semi-analytic estimate of the 95% upper limit for that subband. A test 1-Hz band was chosen at 200.0 Hz, and 1000 software injections were performed in each 0.1-Hz subband, each given a signal strain equal to the semi-analytic estimate for the respective subband. An injection was considered detected if the search for the injection returned a $2\mathcal{F}$ value greater than that in the subband's "Loudest" file. For a 95% strain upper limit, we would expect 95% of signals generated at that signal strain to be detected;

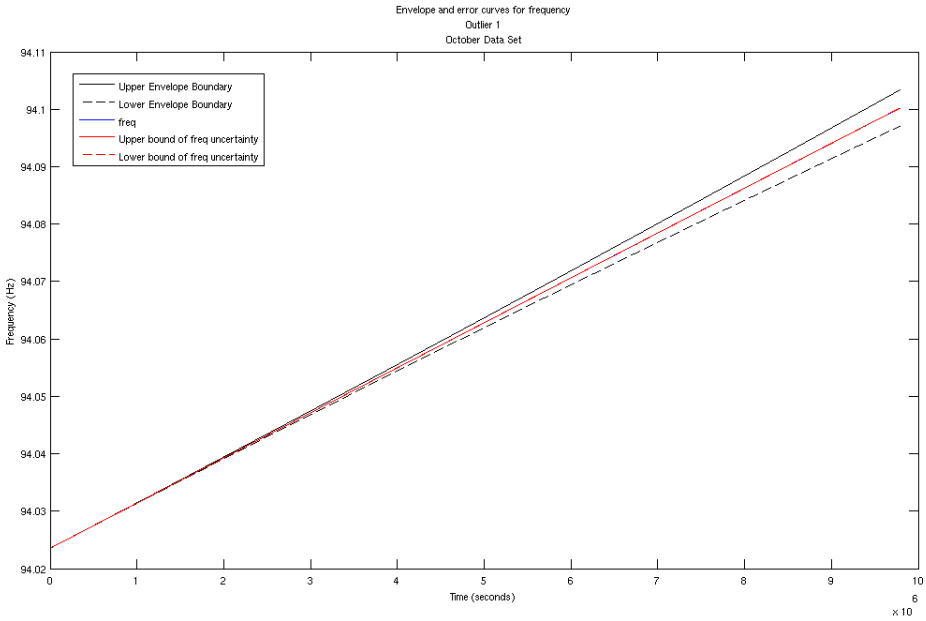

(a)

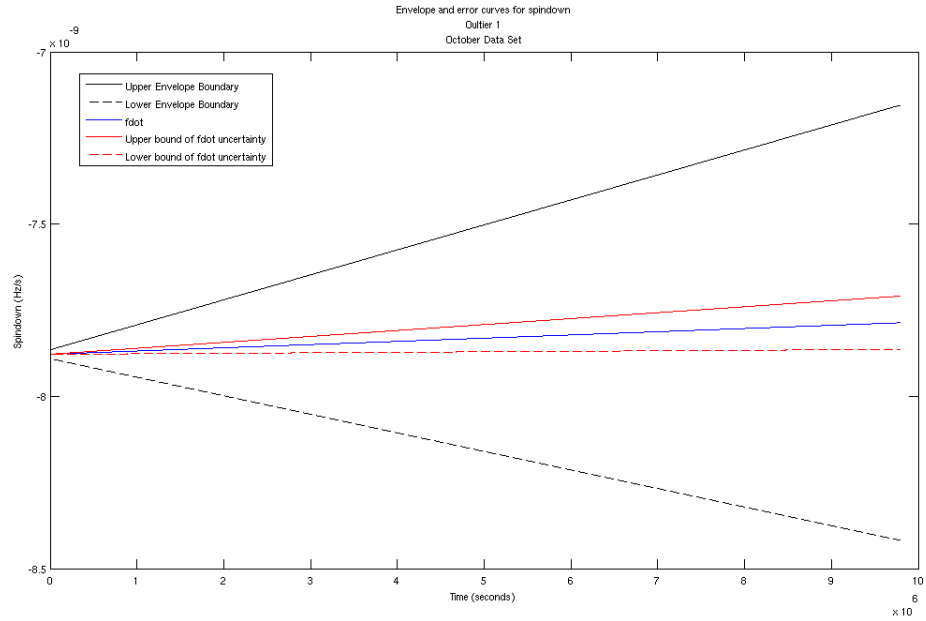

(b)

Figure 6.3: The envelopes (black lines) show the outer edges of the parameter space searched by a lookback search; the red lines show the uncertainty curves for the parameter in question; the blue line shows the time evolution of the parameter assuming no uncertainty. The top plot is for frequency, the bottom for spindown; both correspond to the first outlier from the October data set. It is clear from the plots that the envelopes searched will have the signal within them if it persists in time.

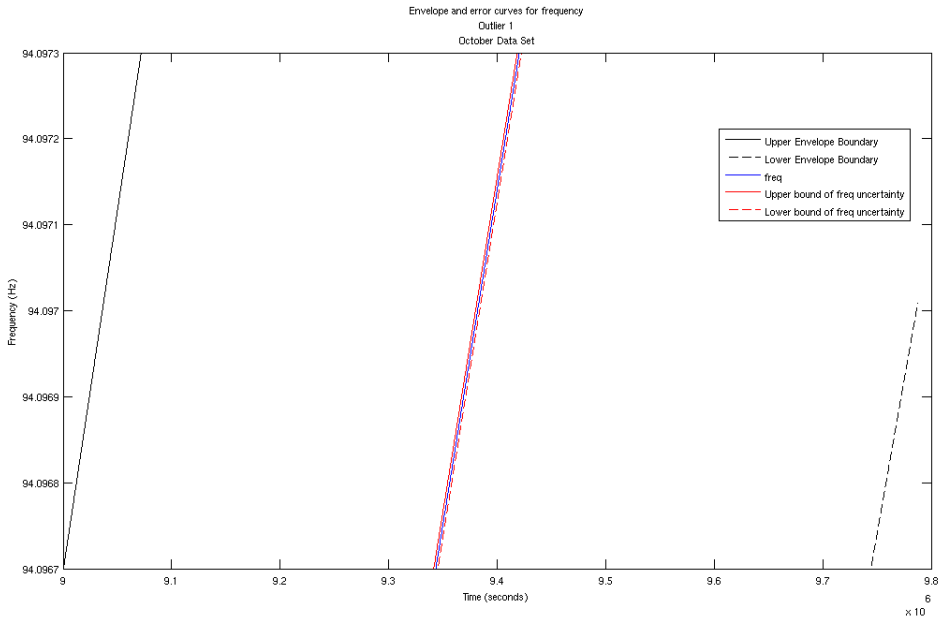

Figure 6.4: This is the frequency plot of Fig. 6.3, zoomed in far enough to make the frequency evolution curve distinguishable from the uncertainty curves. The black lines on either edge of the plot are the boundaries of the search envelope.

a test in Gaussian noise validated $\approx 97\%$ detection of injections, confirming that the semianalytic program returns a conservative 95% upper limit as claimed in [13]. Table 6.2 shows, however, the validation rate of each subband, and we see that in real detector data this validation does not quite hold. The semi-analytic estimates could not be taken at face value, and instead had to be refined through a program of software injections.

Every ten 0.1-Hz subbands are combined into one 1-Hz upper-limit band. The highest upper limit estimate (ULE) of the ten is taken to be the ULE for the 1-Hz band. In each upper limit band, eight sets of 125 injections each are made, with strength varying from $0.80 \cdot h_{nom}$ to $1.20 \cdot h_{nom}$, in increments of $0.05 \cdot h_{nom}$, where h_{nom} is the nominal predicted upper limit estimate. The injections are managed by a Python version of the Perl script used in [74], and considered recovered if they return a $2\mathcal{F}$ value greater than that found in their 0.1-Hz subband's "Loudest" file

Software Injections	
Band (Hz)	Validation Rate (%)
200.0-200.1	93.0
200.1-200.2	94.8
200.2-200.3	94.8
200.3-200.4	95.8
200.4-200.5	95.5
200.5-200.6	93.2
200.6-200.7	92.6
200.7-200.8	94.3
200.8-200.9	94.9
200.9-201.0	96.6

Table 6.2: This table shows the results of 1,000 software injections performed in each $0.1Hz$ sub-band of a test band between $200.0Hz$ and $201.0Hz$. An injection was considered validated if it returned a $2\mathcal{F}$ value greater than the $2\mathcal{F}$ value returned from its subband by the original search. All injections had strain h_0 equal to the 95% upper limit estimate derived for their respective subband by the semianalytic program `lalapps_ComputeFStatAnalyticMonteCarloUpperLimit` and randomly generated frequency and nuisance parameters.

from the original search. After a veto check to ensure that we are not seeing a false positive from the outlier which originally established the loudest $2\mathcal{F}$ in the search, the detection efficiency (injections recovered/total injections) is calculated for each set of 125 injections. These eight points are then plotted on a detection efficiency curve and fitted using a two-parameter sigmoid fit. The 95% upper limit can then be read off from the point where the curve crosses 95% detection efficiency; Figure 6.5 shows such a plot for a sample band. The results of this procedure and the significance of the upper limits obtained are discussed in the next chapter.

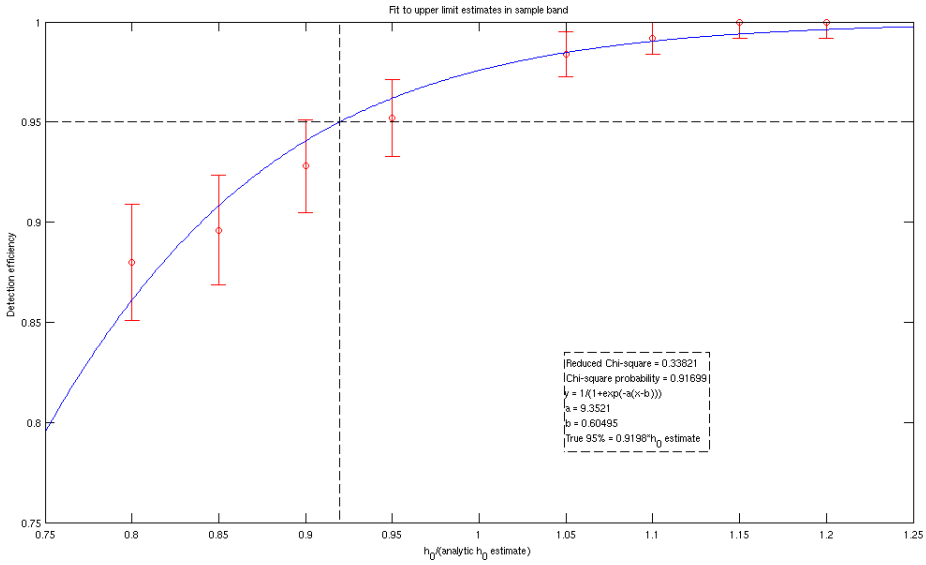

Figure 6.5: A plot demonstrating the upper limit validation technique for a sample band. The x-axis is in units of h_{nom}, the ULE for this upper-limit band; the eight points represent detection efficiencies for eight sets of 125 software injections. These eight points are then fit to a sigmoid curve (in blue); the 95% upper limit can then be read off from the point where the curve crosses 95% detection efficiency.

CHAPTER VII

Search: Results and Discussion

Results from this chapter are preliminary and have not yet been fully reviewed by the LIGO Scientific Collaboration.

This chapter presents the results of the search for gravitational waves described in the preceding chapter. The first section establishes the candidate selection criteria used for the search, through use of a test frequency band and a Mock Data Challenge (MDC). The second section then describes the use of this detection statistic to identify outliers on which to perform manual follow-up. The chapter concludes by reviewing the upper limits on gravitational-wave strain placed in the absence of a signal.

The analysis described here is expected to be the first published search for gravitational waves to employ the barycentric resampling method of [67], described in Chapter 5, which led to large gains in computational efficiency. The technique was first used in a search [15] for gravitational waves from the Calvera object [77], but the search was abandoned after pulsations were detected [78] from Calvera that placed it well outside the frequency band of sufficient LIGO sensitivity.

The search used approximately 900 000 core-hours to search approximately 4×10^{14} templates. (By comparison, the search for Cassiopeia A [74] used 420 000 core-hours

to search 7×10^{12} templates.) The gains in speed from resampling were put to use to search roughly 57 times the number of templates in slightly more than twice the computing time. A direct scaling approach would estimate that the search of [74] could have been completed in 15 750 core hours, for a rough $\times 27$ computational speedup.

7.1 Detection Statistic

As described in the previous chapter, the primary outputs of the search jobs were lists of outliers–candidate templates with $2\mathcal{F}$ values passing a set of thresholds and vetoes. These thresholds and vetoes collectively represent our candidate selection criteria.

The detection statistic used in the search was a combination of thresholds on each template's joint multidetector [76] $2\mathcal{F}$ value and individual detector [66] $2\mathcal{F}$ values, as well as the veto condition that a template's joint $2\mathcal{F}$ value should be larger than its individual detector $2\mathcal{F}$ values. The thresholds used were determined from a Mock Data Challenge and a test frequency band.

7.1.1 Mock Data Challenge

The Mock Data Challenge (MDC) consisted of a set of 1577 artificial CW signals injected into a set of real detector data from S6, which were then searched for using our resampled F-Statistic algorithm. The search was performed over a set of 707 SFTs generated for the Mock Data Challenge at CIT, for GPS times covering July 24 to August 3, 2010, which was identified by the same figure of merit described in the preceding chapter as the best ten days of S6 data. The injections were open–their frequency and nuisance parameters were known, so that searches could be relatively narrowly targeted. Searches were performed in a frequency range of ± 0.5 Hz from

the target frequency, split into ten 0.1-Hz sub-bands with a slight random dither applied so that the signal was randomly placed in the central 20% of the band.

The split into 0.1-Hz subbands combined with the knowledge that the true signal would be in the central 20% of the band meant that the injections would be located somewhere in the central two subbands. The subbands more than 0.2 Hz from the center, i.e., the first and last three subbands of each frequency band, were defined to be at the periphery of the searched frequency band and used for background estimation. A survey of the loudest joint multidetector $2\mathcal{F}$ value reported for each of these background subbands for a subset of the injections between 200 Hz and 240 Hz (used in a pilot run) gave a mean loudest joint multidetector $2\mathcal{F}$ value of ≈ 55 for subbands well away from injections. Given this background level, a multidetector threshold of $2\mathcal{F}_J = 60$ was chosen for candidate selection. A separate individual detector threshold of $2\mathcal{F} = 20$, derived from a test frequency band (see next section), was also applied. Candidates also had to pass the condition that the multidetector $2\mathcal{F}$ value be greater than the individual detector values to avoid being vetoed as instrumental artifacts, as explained in Chapter 6. Finally, candidates were subject to a veto if they fell outside of the central 20% of the searched frequency band and had frequencies within 1/16 Hz of either $N.25$ Hz or $N.75$ Hz, where N is an integer, to protect against blind injections which were also known to populate this data set.

Detection efficiencies (number of injections detected/total number of injections) were plotted against several measures of injection strength. Figures 7.1, 7.2, and 7.3 show these plots. The first figure shows a plot of detection efficiency versus the measure $h_0 \cdot \sqrt{T_{coh}/S_h(f)}$, where h_0 is the strain amplitude of the injection, T_{coh} is an observation time of 10 days, chosen as a representative observation time for these kinds of searches, and $S_h(f)$ is the power spectral noise density at the injection

frequency. This quantity thus measures the efficiency of a search versus the source strain strength of the injections, adjusted for typical observation times and real noise levels. The second graph shows detection efficiency versus another figure of merit:

$$h_{pc} = \sqrt{h_+^2 + h_\times^2} \cdot \sqrt{T_{coh}/S_h(f)}, \tag{7.1}$$

which incorporates the inclination angle ι through the relations

$$h_+ = \frac{1}{2} h_0 (1 + \cos^2 \iota) \tag{7.2}$$

and

$$h_\times = h_0 \cos \iota. \tag{7.3}$$

The figure of merit of Eq.(7.1) better measures the actual detected strength of an injected signal, making use of the quadrature sum of the strain amplitude for each polarization. Taking the inclination angle into account provides a more representative relationship between the injected strain amplitude and the gravitational-wave signal perceived by the detector, since a high inclination angle can weaken (or in the case of Eq.(7.3), eliminate) the strain amplitude for a given polarization. The third graph shows detection efficiency versus the sensitivity depth $\sqrt{S_h(f)}/h_0$, a measure of the ratio of background noise to signal strength. In all three cases the error bars are given assuming a binomial error $\Delta\epsilon$:

$$\Delta\epsilon = \sqrt{\frac{\epsilon(1-\epsilon)}{N_{bin}}} \tag{7.4}$$

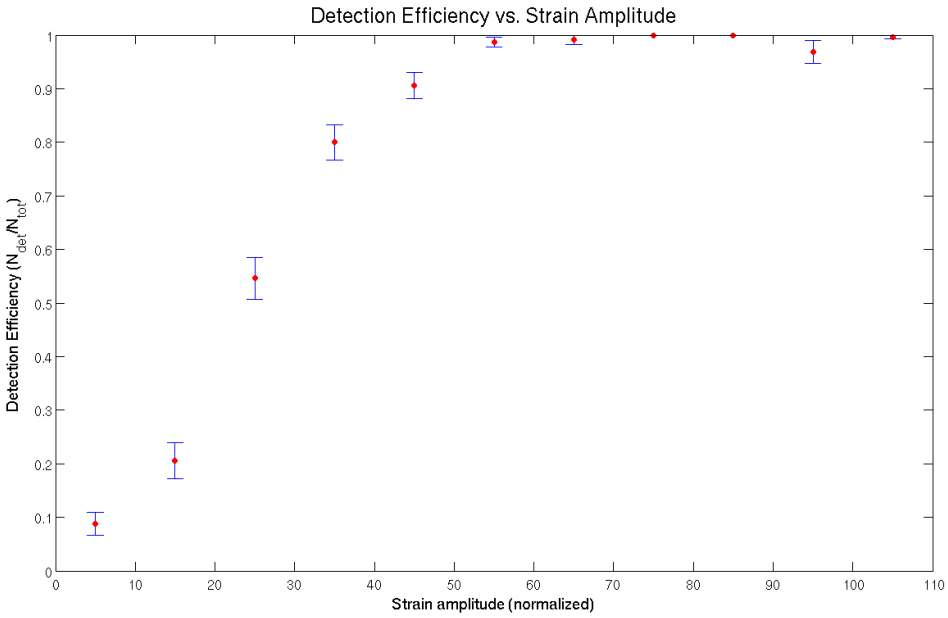

Figure 7.1: One measure of detection efficiency from the Mock Data Challenge. Detection efficiency ϵ is plotted with binomial error against the strain amplitude h_0, normalized to background power spectral density $S_h(f)$ and a ten-day observation time T_{coh}.

where ϵ is the detection efficiency N_{det}/N_{bin} (number detected per bin/total number of injections per bin).

The MDC established a baseline threshold for joint-detector $2\mathcal{F}$ values. We also desired to set a threshold on the individual detector $2\mathcal{F}$ values to protect against large noise fluctuations, and for this purpose, we conducted software injections in a test band.

7.1.2 Test Band

For this more detailed study, a frequency band between 200 Hz and 201 Hz was chosen; a band that sits in the "bucket"–the frequency band of maximum sensitivity for the LIGO detectors in their sixth science run (S6) configuration. The test band was split into 10 0.1-Hz subbands, and in addition to the one thousand software injections performed in each subband as described in Section 6.3, sets of one thousand

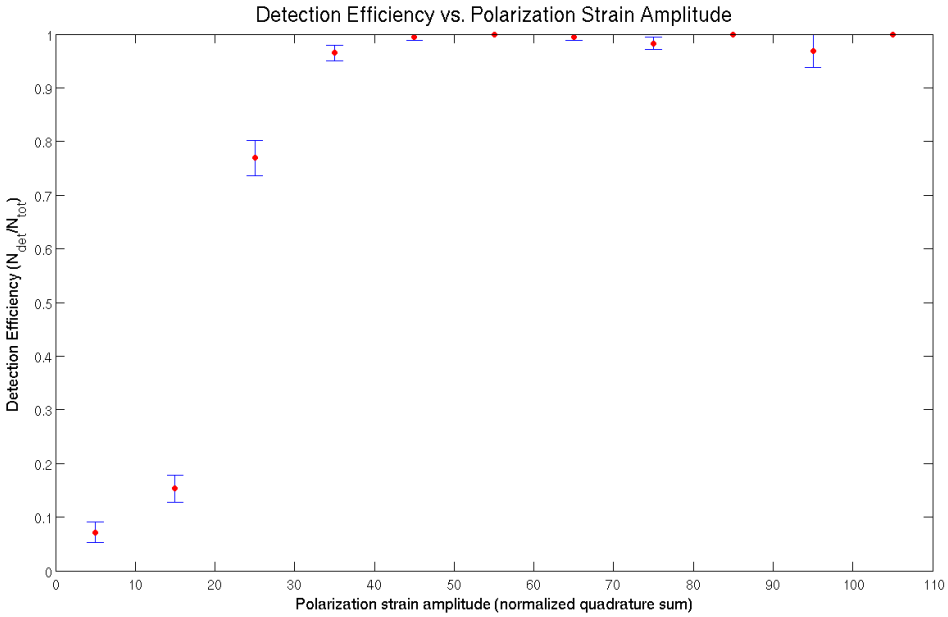

Figure 7.2: One measure of detection efficiency from the Mock Data Challenge. Detection efficiency ϵ is plotted with binomial error against the quadrature sum of polarization strain amplitudes h_+ and $h_\times$, normalized to background power spectral density $S_h(f)$ and a ten-day observation time T_{coh}.

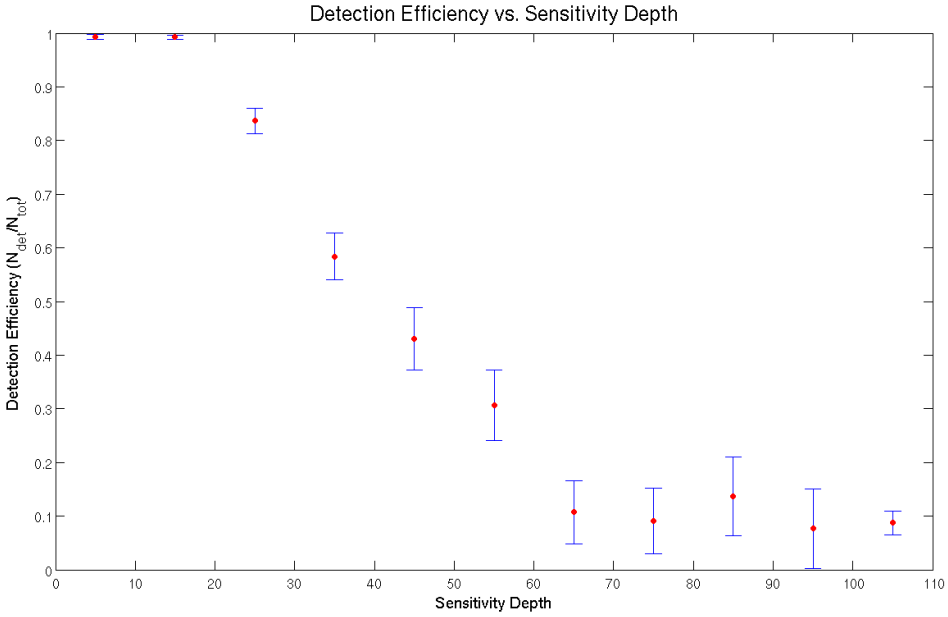

Figure 7.3: One measure of detection efficiency from the Mock Data Challenge. Detection efficiency ϵ is plotted with binomial error against sensitivity depth $\sqrt{S_h(f)}/h_0$, where h_0 is strain amplitude and $S_h(f)$ is the background power spectral density.

injections were performed in the subband between 200.0 Hz and 200.1 Hz at differing signal strengths to produce a detection efficiency curve like those from the MDC. Specifically, injections were performed at 100%, 80%, 60%, 40%, and 20% of the strain strength h_0 estimated for that subband by the semi-analytic program described in the previous chapter.

For each different signal strength, a range of individual detector $2\mathcal{F}$ thresholds were placed on the results. The maximum threshold was set at $2\mathcal{F} = 45$ because this was the minimum joint $2\mathcal{F}$ threshold placed on the output files; any higher threshold would eliminate all outliers (since our veto condition removed outliers with greater individual $2\mathcal{F}$ values than the joint value). The minimum threshold was set empirically at $2\mathcal{F} = 20$; outliers with joint $2\mathcal{F}$ greater than 45 were generally unaffected by thresholds set any lower. The individual detector threshold was meant to ensure a candidate signal was required to be present in both detectors; because of the differing orientations of each detector, however, it is possible for a signal to appear weaker in one detector than another [66]–one must be mindful of the false dismissal rate. The results of these sets of software injections are presented in Table 7.1 and plotted in Figure 7.4. An injection was considered detected if the loudest template that passed the individual detector threshold requirement reported a multidetector $2\mathcal{F}$ greater than that of the loudest outlier in that subband reported by the search. Each entry in the table represents the fraction of the 1000 injections that were detected. An individual detector threshold of $2\mathcal{F} = 20$ for each detector was chosen; this value was high enough to veto noise fluctuations but maintained a low false dismissal rate.

Efficiency vs. Trial for different single detector $2\mathcal{F}$ thresholds							
Strength	No threshold	$2\mathcal{F} = 20$	$2\mathcal{F} = 25$	$2\mathcal{F} = 30$	$2\mathcal{F} = 35$	$2\mathcal{F} = 40$	$2\mathcal{F} = 45$
100%	94.8%	93.1%	88.4%	81.3%	73.0%	65.5%	59.3%
80%	78.1%	76.1%	72.5%	63.8%	56.0%	49.6%	44.3%
60%	52.1%	50.4%	46.0%	38.0%	30.9%	23.1%	17.1%
40%	19.2%	17.6%	14.0%	8.0%	3.5%	0.9%	0.3%
20%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table 7.1: This table shows the results of sets of 1000 software injections. The first column gives the strength (as a percentage of the semi-analytic 95% upper limit estimate h_0) of the injected signals; the remaining columns give the percentage of injections (from a set of 1000) detected for a given individual detector $2\mathcal{F}$ threshold. Injections passing the individual detector threshold for all detectors and reporting a multidetector $2\mathcal{F}$ greater than that reported for their subband by the search were considered detected.

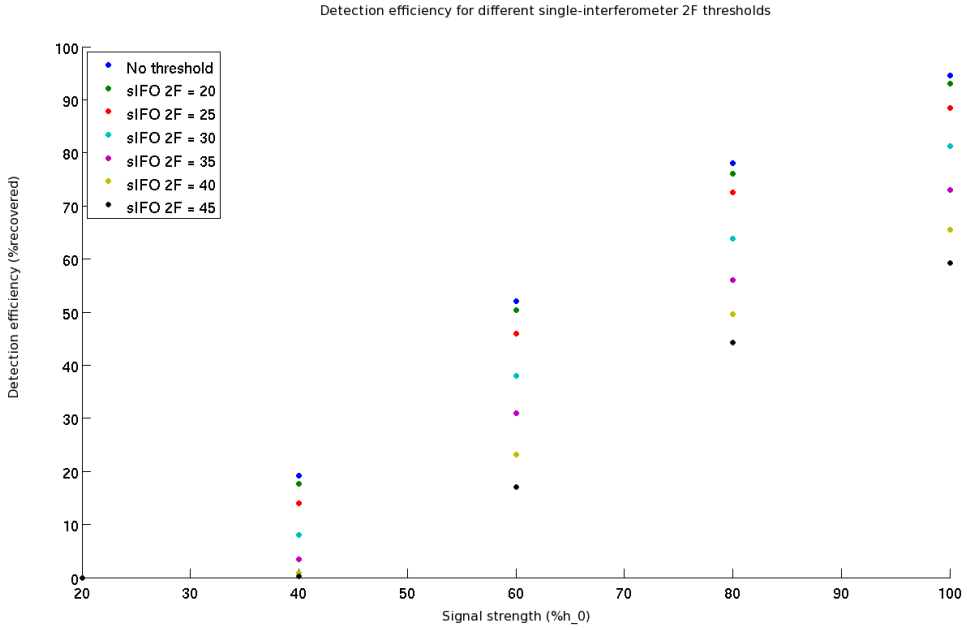

Figure 7.4: The results of Table 7.1 plotted as detection efficiency curves. Each individual detector $2\mathcal{F}$ threshold is signified by a different color marker.

7.2 Outlier Followup

Templates with $2\mathcal{F}$ values that met or exceeded the thresholds, and passed the vetoes, set up by the detection statistic were deemed outliers. This produced a list of 155 outliers from the July-August data stretch and a list of 168 outliers from the October data set. These outliers were subjected to the followup procedures described in the previous chapter–searched for in the opposite data stretch and in extended observation times in their original data stretch. These tests were not cumulative–an outlier need pass only one to be considered a "candidate" and passed to the next stage of followup. Even so, the combined 323 outliers produced only seven candidates, listed in Table 7.2.

These candidates were subject to manual followup. They were compared to strain histograms of run-averaged (i.e., over all of S6) spectra from each detector, to identify instrumental noise lines which could be responsible. Figure 7.5 shows one such strain histogram, which clearly shows a loud instrumental line in the Hanford (H1) detector (red) overlapping the outlier signal (blue) of the fourth candidate from Table 7.2. In five of the seven cases, the strain histograms gave clear evidence of an instrumental noise line responsible for the candidate, and in these cases records of prior detector characterization studies were consulted to provide explanations for the noise artifacts. In those cases the artifact is listed in Table 7.2 as well.

Two of the candidates (Outliers 79 and 131 from the October data set) were not clearly associated with known instrumental artifacts or noise lines. These were given another round of follow up, with a look-back and extended look (as described in the previous chapter) performed in June data. June was selected for being the farthest removed (in the time domain) available data of comparable sensitivity. Large

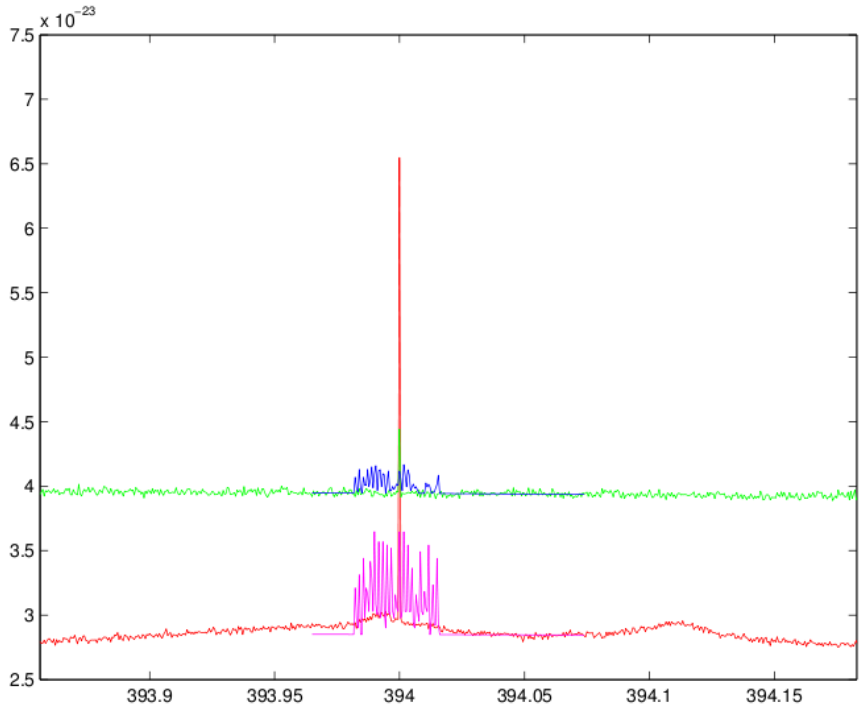

Figure 7.5: A run-averaged strain histogram for Outlier 77 (see Table 7.2). The red line shows a loud instrumental noise line in the Hanford detector overlapping the outlier signal, which is blue. The green line is the run-averaged spectrum of the Livingston (L1) detector, overlapping the Livingston outlier signal (magenta).

July-August Data					
Outlier	Search f (Hz)	Outlier f (Hz)	Search $2\mathcal{F}_J$	Followup $2\mathcal{F}_J$	Artifact, if any
27	192.4907	192.4956	612.969	300.712	Hardware Injection
74	392.2232	392.2315	189.903	173.787	Clock noise
77	394.0231	394.0307	228.268	197.300	Digital line
October data					
27	192.4195	192.4313	875.575	484.254	Hardware Injection
79	403.6424	403.8612	114.626	61.331	—
85	417.0394	417.1384	60.309	176.200	H1 Output Mode Cleaner Line
131	575.9658	576.5057	61.943	53.805	—

Table 7.2: This table lists the seven candidates which passed the first round of outlier followup. The columns give, respectively: the outlier's identifying number; the frequency of the outlier in the search; the frequency of the outlier in the followup data set in which it appeared; the $2\mathcal{F}$ value of the outlier in the search; the $2\mathcal{F}$ value of the outlier in the followup data set in which it appeared; the explanation, if any, provided by comparison with run-averaged strain histograms in conjunction with detector characterization records.

time separation leads to very different Doppler corrections needed to reconstruct an astrophysical source, corrections unlikely to reinforce an instrumental or environmental artifact. (The ideal interval is six months, which will give equal and opposite Doppler corrections, but April data was too insensitive compared to October data to be of comparable use.) Both outliers failed to pass the $2\mathcal{F}$ thresholds established by software injections in any of their June tests.

The search for Cassiopeia A [13] showed that the loudest $2\mathcal{F}$ value expected in the absence of signal depends on the number of templates searched; in that case, for $N = 6.2 \times 10^{12}$ statistically independent templates[1], the largest expected $2\mathcal{F}$ value was most likely to lie in the range $62 \leq 2\mathcal{F} \leq 75$. Furthermore, the range shifts upward with increasing N; an increase in N by two orders of magnitude led to an increase by roughly ten in the loudest expected $2\mathcal{F}$ value. As mentioned at the beginning of this chapter, our search used roughly two orders of magnitude more templates than the Cassiopeia A search, so our range of loudest expected $2\mathcal{F}$ would indeed be shifted upward. The $2\mathcal{F}$ values associated with Outliers 79 and 131 were

[1]The N_T templates used in our searches are not perfectly independent, but can be represented by N statistically independent templates where $N = 0.88N_T$. See section 8.7 of [13].

$2\mathcal{F}_J = 61.3$ and $2\mathcal{F}_J = 61.9$, respectively. The outliers' failure to pass the June tests and their marginal $2\mathcal{F}$ values led us to dismiss them as noise fluctuations.

Thus no gravitational wave signals were detected by our search. In the absence of a detection, we can set upper limits on the possible strength of gravitational waves in our data.

7.3 Upper Limit Validation

Upper limits were set for each 1-Hz band, using the procedure described in the previous chapter. A small number of bands had outliers so large that the semi-analytic method failed to converge to an estimate for h_0. Upper limits were therefore not set for these bands. For approximately 3% of the 1-Hz bands (41 of 583) in the October data set, the 95% upper limit had to be determined by extrapolating beyond their eight points, indicating that the initial semi-analytic estimate was significantly wrong (i.e., the true 95% upper limit did not lie between 80% and 120% of the initial 95% upper limit estimate). A further three bands set a 95% upper limit with uncertainty greater than $0.05 \cdot h_{nom}$, where h_{nom} is the initial semi-analytic estimate for h_0. For these 44 bands, the procedure was rerun with the extrapolated 95% upper limit as the new initial estimate around which the eight points were constructed. The same was done for 62 bands (59 extrapolated, three with large uncertainty) in the July-August data.

Figure 7.6 shows the semi-analytic upper limit estimates (ULEs) over the full band for the July-August data set; which was the more sensitive of the two because of its much greater livetime (642 SFTs vs. 374 SFTs for the October data set). The blue curve represents the expected sensitivity of the search for this data set, computed from the power spectral density at each frequency; there is good agreement with the

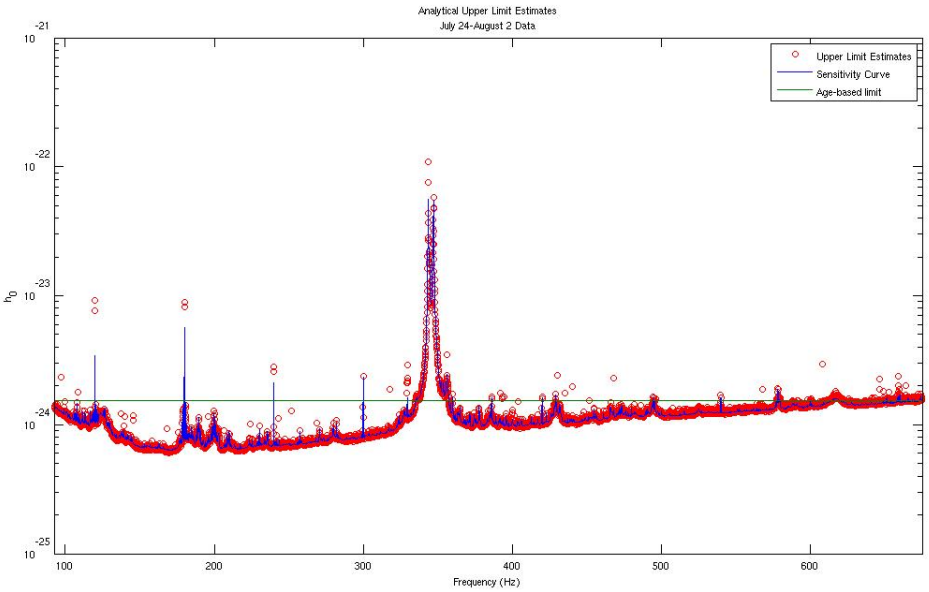

Figure 7.6: Upper limit estimates (red circles) compared to the sensitivity of the detector (blue curve) and the initial age-based limit on h_0 of Chapter 6 (green line).

ULEs. The green line represents the age-based limit derived in the previous chapter when first considering parameter space. Its intersection with the ULEs at either end of the plot is a confirmation that we correctly estimated the frequency band to search over. Everywhere the upper limits lie below the green curve, we have beaten the age-based spin-down limit for our source.

Figure 7.7 is a similar plot converting the upper limits on h_0 to upper limits on ellipticity (introduced in Section 4.2). The green curve represents the age-based limit on ellipticity, as calculated from Eq.(4.35) using the same assumptions (braking index $n = 5$, age $\tau = 300$ years) used in the parameter space calculations. We beat the age-based limit for our source everywhere the upper limits lie below the green curve.

Figure 7.8 is a histogram of the final 95% upper limits as a fraction of the semi-analytic ULEs. The distribution has a mean of 0.934, with standard deviation 0.061

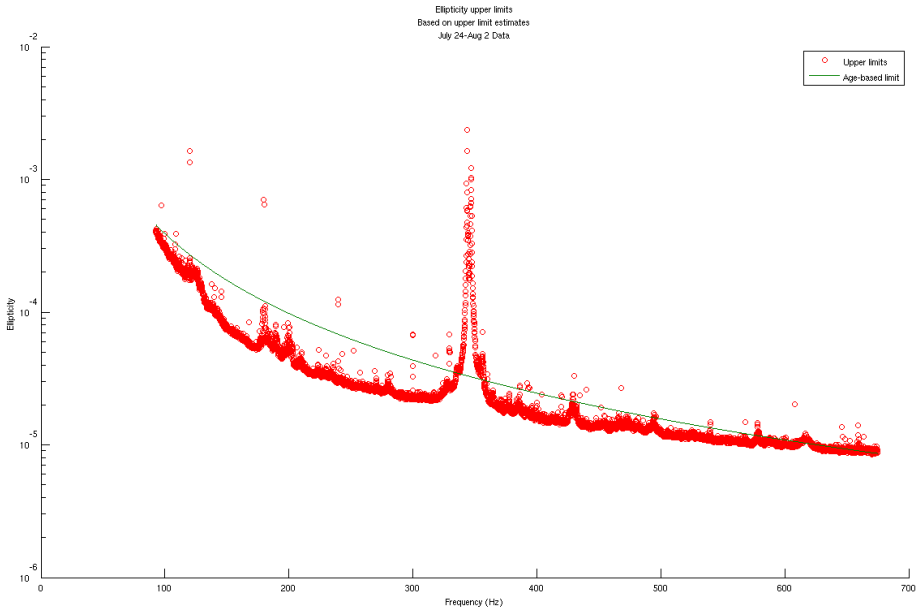

Figure 7.7: Upper limit estimates (red circles) compared to the initial age-based limit on ellipticity ϵ as calculated in Chapter 4 (green line).

(6.5%). The slight bias towards towards lower values is an expected consequence of our methodology, which systematically overestimates the upper limit by always using the largest per-0.1 Hz ULE for the combined 1-Hz upper limit bands.

It should be noted that this figure has had certain anomalous 1-Hz bands removed. 21 bands contained at least one, and sometimes as many as two, 0.1-Hz bands where no ULE was set because of outliers so large that the semi-analytic method failed to converge. This caused highly disturbed bands to go unaccounted for–loudest $2\mathcal{F}$ values were too low, creating an artificially low ULE for the 1-Hz band and an artificially high ratio when injections were made in the disturbed band. A further six bands had the opposite problem–they were contaminated by very large disturbances that were not large enough to cause the semi-analytic method to fail to converge, and as a result one 0.1-Hz band received a ULE much higher than its nine companions. This resulted in a ULE that was too high for the 1-Hz band, and artificially high

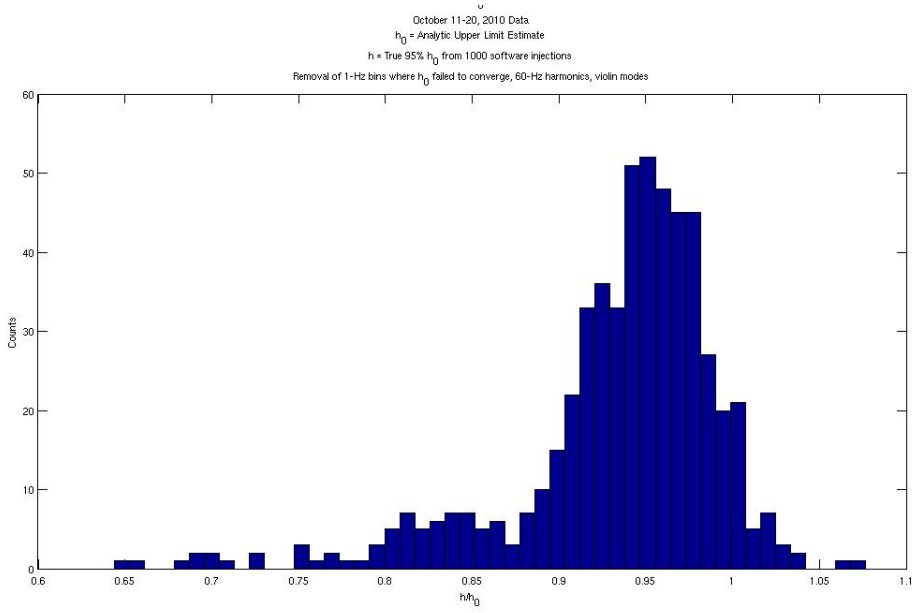

Figure 7.8: Histogram of true 95% upper limits as a fraction of inital semi-analytic upper limit estimates. The bias towards small ratios shows our upper limit estimates tend to be conservative, positing upper limits that are higher than the true value.

recovery rates pushed the 95% upper limit artificially low. Four of these six bands were disturbed by 60-Hz harmonics; the other two owed their disturbances to the violin modes of the detector. No upper limit was set in any of the above excluded bands, which in total represent less than 5% of the frequency band.

This search placed the first explicit upper limits on continuous gravitational wave strength from the nearby globular cluster NGC6544, and was indeed the first directed search for any globular clusters. The lowest upper limits were $h_0 = 6.7 \times 10^{-25}$ for the 1-Hz band starting at 173.0 Hz for the October data set, and $h_0 = 6.0 \times 10^{-25}$ for the 1-Hz band starting at 170 Hz for July-August data. These are comparable to the upper limits of 7×10^{-25} at 150 Hz obtained by the Cassiopeia A search [74]; the semi-coherent Galactic Center search [73] obtained 90% (as opposed to our 95%) confidence upper limits of 3.3×10^{-25} for frequencies near 150 Hz, but did not search over the second frequency derivative. A recent search over nine supernova remnants,

done without resampling, [79] set upper limits as low as 3.7×10^{-25} for the supernova remnant G93.3+6.9, but used a coherence time of over 23 days (and a frequency band of only 264 Hz). The same search set a comparable 95% upper limit of 6.4×10^{-25} for the supernova remnant G1.9+0.3, which has a sky position with comparable average sensitivity to that of NGC6544; that search used a comparable 9.1 days of data, but was limited to a 146 Hz search band, compared to our 583 Hz. Our search was carried out at substantially less computational cost because of barycentric resampling, and could thus search over a much larger parameter space.

The best upper limit on ellipticity, established using the July-August data set, was $\epsilon = 8.5 \times 10^{-6}$, for the 1-Hz band starting at 670 Hz. This is comparable to the best upper limit (4×10^{-5}) obtained by the Cassiopeia A search [13]; the Galactic Center search [73] set an upper limit of $\epsilon = 8.7 \times 10^{-6}$, albeit when assuming a moment of inertia three times the fiducial value ($10^{38}\mathrm{kg} \cdot \mathrm{m}^2$) that we used. The supernova remnant search [79] set a comparable upper limit on ellipticity at 7.6×10^{-5} for the supernova remnant G1.9+0.3.

Although too high for the maximum ellipticities for conventional neutron star models (see Section 4.2), there exist more exotic models for which this limit is in the range of interest: solid strange quark stars can support ellipticies up to 10^{-4}, and hybrid quark-baryon or meson-condensate stars might sustain up to 10^{-5} [80]. The search would have been capable of detecting a maximally deformed star described by these models; however, since we did not make a detection of a gravitational wave signal, we cannot make any statements about the equations of state of possible neutron stars in NGC6544. Their ellipticities could always simply be lower than the theoretical maxima.

CHAPTER VIII

Conclusion

This thesis presented the results of a search for continuous-wave (CW) gravitational waves from the globular cluster NGC6544 in data from the Hanford and Livingston detectors taken during LIGO's sixth science run (S6). We have derived gravitational waves from Einstein's equations, and the four families of LIGO searches were discussed. We described the LIGO interferometers, their response to gravitational waves, their fundamental sources of noise, and the steps taken to mitigate those sources in the next-generation Advanced LIGO detectors. We described in particular detail continuous-wave gravitational waves, the subject of our search: the composition of neutron stars, their likeliest astrophysical source; their emission mechanisms; the reasons globular cluster NGC6544 is an attractive target in a search for them. We explained the barycentric resampling method, and demonstrated its close to an order of magnitude computational speedup over standard demodulation techniques. Lastly, we described the search: the selection of data and the parameter space to be searched; methods for following up candidate events and setting and validating upper limits in the absence of a signal; and, finally, its results.

All candidate outliers were shown to be inconsistent with an astrophysical signal and/or due to instrumental artifacts, and no detection of gravitational waves could

be claimed. Upper limits were set in 1-Hz bands spanning the search frequency band between 92.5 Hz and 675 Hz. The lowest upper limit was $h_0 = 6.0\times10^{-25}$ for the 1-Hz band starting at 170 Hz, and the lowest upper limit on ellipticity was $\epsilon = 8.5\times10^{-6}$ for the 1-Hz band starting at 670 Hz, both in the data set spanning July 24 to August 3, 2010. These were comparable to other recent LIGO searches for gravitational waves from the youngest neutron star Cassiopeia A, the Galactic Center, and nine supernova remnants, but carried out at significantly less computational cost. The limits on ellipticity were too high for conventional neutron star models but in the range of maxima from more exotic equations of state; however without a detection we can make no statement about neutron star model favorability since ellipticities could simply be below the theoretical maxima.

As I write, the final preparations are being made for the first observing runs with the next-generation Advanced LIGO detectors. The extension of barycentric resampling to other targets of directed searches promises to lower computational costs across the board and allow for searches for more targets with the same resources, ushering in an era of quick and computationally cheap directed searches for CW gravitational waves in LIGO. The program developed here will be integral to those extensions. The use of semi-coherent methods, combining the results of multiple coherent observation times for greater sensitivity, will enable those searches to reach even greater sensitivities. Combined with the order of magnitude increase in the base sensitivity of the LIGO detectors in the Advanced Era, gravitational-wave astronomy is rapidly approaching a time when detections could go from a long-awaited holy grail to a relatively routine event. This is a very exciting time to be hunting.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] A. Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite 844-847.*, pages 844–847, 1915.

[2] P. Saulson. Gravitational Wave Detection: A Historical Perspective. Conference Presentation, April 2015. APS April 2015 Conference, Baltimore, MD.

[3] F. A. E. Pirani. On the Physical significance of the Riemann tensor. *Acta Physica Polonica*, 15:389–405, 1956.

[4] J. Weber. Detection and Generation of Gravitational Waves. *Physical Review*, 117:306–313, January 1960.

[5] J. Weber. Evidence for Discovery of Gravitational Radiation. *Physical Review Letters*, 22:1320–1324, June 1969.

[6] P. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing Company, Incorporated, 1994.

[7] R. Weiss. Electromagetically coupled broadband gravitational antenna. Technical Report 54, Research Laboratory of Electronics, MIT, 1972. Quarterly Progress Report.

[8] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophys. J Lett.*, 195:L51–L53, January 1975.

[9] J. M. Weisberg, D. J. Nice, and J. H. Taylor. Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16. *Astrophys. J*, 722:1030–1034, October 2010.

[10] E. Goetz. *Gravitational wave studies: detector calibration and an all-sky search for spinning neutron stars in binary systems*. PhD thesis, University of Michigan, 2010.

[11] S.M. Carroll. Lecture notes on general relativity. Technical Report gr-qc/9712019v1, arXiv, 1997-12-03.

[12] G. Meadors. *Directed searches for continuous gravitational waves from spinning neutron stars in binary systems*. PhD thesis, University of Michigan, 2014.

[13] K. Wette. *Gravitational waves from accreting neutron stars and Cassiopeia A*. PhD thesis, Australian National University, 2009.

[14] C.W. Misner, K.S. Thorne, and J.A Wheeler. *Gravitation*. W.H. Freeman, 1973.

[15] P. Patel. *Search for Gravitational Waves from a nearby neutron Star using barycentric resampling*. PhD thesis, California Institute of Technology, 2010.

[16] J.G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower. Comparisons of binary black hole merger waveforms. *Class. Quantum Grav.*, 24, 2007.

[17] L. Cadonati et al. Status of NINJA: the Numerical INJection Analysis project. *Class. Quantum Grav.*, 26:114008, 2009.

[18] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, K. Ackley, C. Adams, T. Adams, and et al. The NINJA-2 project: detecting and characterizing gravitational waveforms modelled using numerical binary black hole simulations. *Classical and Quantum Gravity*, 31(11):115004, June 2014.

[19] J. Aasi et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown in LIGO-Virgo data from 2009-2010. *Phys. Rev. D.*, 87(022002), 2013.

[20] J. Abadie et al. Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys. Rev. D*, 87(082002), 2012.

[21] C.D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Class. Quantum Grav.*, 26:063001, 2009.

[22] J. Abadie et al. Search for gravitational waves associated with gamma-ray bursts during LIGO Science Run 6 and Virgo Science Runs 2 and 3. *Astrophys. J*, 760(12), 2012.

[23] J. Aasi et al. Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009-2010 LIGO and Virgo Data. *Phys. Rev. Lett.*, 113:231101, 2014.

[24] Ade, P.A.R. et al. A Joint Analysis of BICEP2/Keck Array and Planck Data. *Phys. Rev. Lett.*, 114:101301, 2015.

[25] G. Hobbs et al. The international pulsar timing array project: using pulsars as a gravitational wave detector. *Class. Quantum Grav.*, 27:084013, 2010.

[26] P. Amaro-Seane et al. Low-frequency gravitational-wave science with eLISA/NGO. *Class. Quantum Grav.*, 29:124016, 2012.

[27] Michelson, A.A. and Morley, E.W. On the relative motion of the earth and the luminiferous ether. *Am. J. Sci.*, 34:333–345, 1887.

[28] H. Haus. *Waves and Fields in Optoelectronics*. Prentice-Hall, 1984.

[29] R. Weiss. The LIGO interferometers: how they work and how well they work. Technical Report G030024-00-D, Laser Interferometer Gravitational Wave Observatory, February 2003.

[30] E. Black. Notes on the pound-drever-hall technique. Technical Report T980045-00-D, Laser Interferometer Gravitational Wave Observatory, April 1998.

[31] K. Riles. Gravitational waves: Sources, detectors and searches. *Progress in Particle and Nuclear Physics*, 68:1–54, January 2013.

[32] R. Amin. Thermal compensation in LIGO. Technical Report G0900358, Laser Interferometer Gravitational Wave Observatory, April 2009.

[33] D. Macleod and J. Smith. Characterization of the ligo detectors during their sixth science run. Conference Presentation, March 2013. LVC March 2013 Conference, Bethesda, MD.

[34] P. J. Cote and L. V. Meisel. Self-organized criticality and the Barkhausen effect. *Physical Review Letters*, 67:1334–1337, September 1991.

[35] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, and et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901, July 2009.

[36] A. V. Cumming, A. S. Bell, L. Barsotti, M. A. Barton, G. Cagnoli, D. Cook, L. Cunningham, M. Evans, G. D. Hammond, G. M. Harry, A. Heptonstall, J. Hough, R. Jones, R. Kumar, R. Mittleman, N. A. Robertson, S. Rowan, B. Shapiro, K. A. Strain, K. Tokmakov, C. Torrie, and A. A. van Veggel. Design and development of the advanced LIGO monolithic fused silica suspension. *Classical and Quantum Gravity*, 29(3):035003, February 2012.

[37] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, and et al. Characterization of the LIGO detectors during their sixth science run. *Classical and Quantum Gravity*, 32(11):115012, June 2015.

[38] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, and et al. Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar. *Astrophys. J Lett.*, 683:L45–L49, August 2008.

[39] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, and et al. Gravitational Waves from Known Pulsars: Results from the Initial Detector Era. *Astrophys. J*, 785:119, April 2014.

[40] G. Harry. Advanced LIGO: The next generation of gravitational wave detectors. *Class. Quantum Grav.*, 27:084006, 2010.

[41] G. Losurdo. Advanced Virgo: where we are. Technical Report G1400356, VIRGO Collaboration and EGO, March 2014.

[42] R. Schnabel et al. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nat. Phys.*, 7:962965, 2011.

[43] Y. Asi et al. Interferometer design of the KAGRA gravitational wave detector. *Phys. Rev. D*, 88:043007, 2013.

[44] T. Souradeep. LIGO India Status. Technical Report G1500766, IndIGO Consortium, May 2015.

[45] E. Goetz and K. Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Class. Quant. Grav.*, 28(215006), 2011.

[46] C. Messenger et al. Gravitational waves from sco x-1: A comparison of search methods and prospects for detection with advanced detectors. Technical Report P1400217-v3, LIGO, 2015.

[47] M. Bruggen and S. Rosswog. *Introduction to High-Energy Astrophysics*. Cambridge University Press, 2007.

[48] D. Prialnik. *An Introduction to the Theory of Stellar Structure and Evolution*. Cambridge University Press, 2008.

[49] L. Landau. On the theory of stars. *Phys. Z. Sowjetunion*, 1(285), 1932.

[50] W. Baade and F. Zwicky. Cosmic rays from Super-novae. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 20, pages 259–263, may 1934.

[51] J. Oppenheimer and G. Volkoff. On massive neutron cores. *Physical Review*, 55(4):374–381, 1939.

[52] J. M. Lattimer and M. Prakash. The Physics of Neutron Stars. *Science*, 304:536–542, April 2004.

[53] A. Hewish et al. Observation of a rapidly pulsating radio source. *Nature*, 217(709), 1968.

[54] D. R. Lorimer. Blind surveys for radio pulsars and transients. In M. Burgay, N. D'Amico, P. Esposito, A. Pellizzoni, and A. Possenti, editors, *American Institute of Physics Conference Series*, volume 1357 of *American Institute of Physics Conference Series*, pages 11–18, August 2011.

[55] A. S. Fruchter, D. R. Stinebring, and J. H. Taylor. A millisecond pulsar in an eclipsing binary. *Nature*, 333:237–239, May 1988.

[56] M. Zimmermann and E. Szedenits, Jr. Gravitational waves from rotating and precessing rigid bodies - Simple models and applications to pulsars. *Phys. Rev. D*, 20:351–355, July 1979.

[57] Benjamin J. Owen, Lee Lindblom, Curt Cutler, Bernard F. Schutz, Alberto Vecchio, and Nils Andersson. Gravitational waves from hot young rapidly rotating neutron stars. *Phys. Rev. D*, 58(8):084020, Sep 1998.

[58] N. Andersson. TOPICAL REVIEW: Gravitational waves from instabilities in relativistic stars. *Classical and Quantum Gravity*, 20:105, April 2003.

[59] P. Arras, E. E. Flanagan, S. M. Morsink, A. K. Schenk, S. A. Teukolsky, and I. Wasserman. Saturation of the r-Mode Instability. *Astrophys. J*, 591:1129–1151, July 2003.

[60] B. J. Owen. How to adapt broad-band gravitational-wave searches for r-modes. *Phys. Rev. D*, 82(10):104002, November 2010.

[61] S. Yoshida, S. Yoshida, and Y. Eriguchi. R-mode oscillations of rapidly rotating barotropic stars in general relativity: analysis by the relativistic Cowling approximation. *Monthly Notices of the Royal Astronomical Society*, 356:217–224, January 2005.

[62] A. Melatos and D. J. B. Payne. Gravitational Radiation from an Accreting Millisecond Pulsar with a Magnetically Confined Mountain. *Astrophys. J*, 623:1044–1050, April 2005.

[63] L. Bildsten. Gravitational Radiation and Rotation of Accreting Neutron Stars. *The Astrophysical Journal Letters*, 501(1):L89, 1998.

[64] B. Abbott et al. Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Phys. Rev. D*, 76(8):082001, 2007.

[65] B. Owen. Globular cluster targets. https://wiki.ligo.org/CW/GlobularClusterTargets, July 2008.

[66] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58(6):063001, September 1998.

[67] P. Patel, X. Siemens, R. Dupuis, and J. Betzwieser. Implementation of barycentric resampling for continuous wave searches in gravitational wave data. *Phys. Rev. D*, 81(8):084032, April 2010.

[68] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, B. Allen, and et al. Beating the Spin-down Limit on Gravitational Wave Emission from the Vela Pulsar. *Astrophys. J*, 737:93, August 2011.

[69] B. Owen. Einstein@Home's search for spinning neutron stars. In *APS Meeting Abstracts*, page 12001, April 2007.

[70] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, and et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys. Rev. D*, 87(4):042001, February 2013.

[71] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, and et al. All-sky search for periodic gravitational waves in the full S5 LIGO data. *Phys. Rev. D*, 85(2):022001, January 2012.

[72] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, K. Ackley, C. Adams, T. Adams, and et al. First all-sky search for continuous gravitational waves from unknown sources in binary systems. *Phys. Rev. D*, 90(6):062010, September 2014.

[73] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, and et al. Directed search for continuous gravitational waves from the Galactic center. *Phys. Rev. D*, 88(10):102002, November 2013.

[74] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, C. Adams, R. Adhikari, P. Ajith, B. Allen, G. Allen, E. Amador Ceron, and et al. First Search for Gravitational Waves from the Youngest Known Neutron Star. *Astrophys. J*, 722:1504–1513, October 2010.

[75] R. Prix. Search for continuous gravitational waves: Metric of the multidetector F-statistic. *Phys. Rev. D*, 75(2):023004, January 2007.

[76] C. Cutler and B. F. Schutz. Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars. *Phys. Rev. D*, 72(6):063006, September 2005.

[77] R. E. Rutledge, D. B. Fox, and A. H. Shevchuk. Discovery of an Isolated Compact Object at High Galactic Latitude. *Astrophys. J*, 672:1137–1143, January 2008.

[78] S. Zane, F. Haberl, G. L. Israel, A. Pellizzoni, M. Burgay, R. P. Mignani, R. Turolla, A. Possenti, P. Esposito, D. Champion, R. P. Eatough, E. Barr, and M. Kramer. Discovery of 59 ms pulsations from 1RXS J141256.0+792204 (Calvera). *Monthly Notices of the Royal Astronomical Society*, 410:2428–2445, February 2011.

[79] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, T. Adams, and et al. Searches for continuous gravitational waves from nine young supernova remnants. *ArXiv e-prints*, December 2014.

[80] B. J. Owen. Maximum Elastic Deformations of Compact Stars with Exotic Equations of State. *Physical Review Letters*, 95(21):211101, November 2005.