Search for Gravitational Waves from a nearby neutron star using barycentric resampling

Thesis by
Pinkesh K. Patel

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

California Institute of Technology
Pasadena, California

2011

(Defended Aug 26, 2010)

To my Mum and Dad...

Preface

The work presented in this thesis was carried out within the LIGO Scientific Collaboration (LSC). The methods and results presented here are under review and are potentially subject to change. The opinions expressed here are those of the author and not necessarily those of the LSC.

The author gratefully acknowledges the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, which provided support for this work.

Acknowledgements

It is the memories that remain with you, even after one can't remember the exact moment, the faces of the people involved or even the year in which they happened. How does one thank all the people involved in creating all the wonderful memories I have had in the past five years? If I miss someone, know that I will remember all the sweet or sour memories that I have shared with you. Sour memories also have their place and make me the person I am. So thank you to all the people who thought they were unkind to me or said things I didn't want to hear. You also shaped me, just as much as the people I hold dear.

I would like to start by thanking the California Institute of Technology, lovingly known as Caltech. It is the best university on the planet. Caltech and its staff are exemplary and I cannot thank them enough. The only sad part about graduating is that I will be leaving Caltech. I will miss its Millikan pond, its scientifically minded grafitti, its wonderful easy-going attitude and its efficient and always smiling secretaries, among other things.

Alan Weinstein is the best advisor that I could have hoped for. He has given me all the freedom that even a postdoc may not hope to have. He has been a wonderful critic of my work and has always had my best interests at heart. He has stood by me when needed and fought for me (he has the blows to show for it). He has stayed up at 5 AM to hear me present in Europe, and has read every word that I have written. Thank you Alan.

The LIGO Scientific Collaboration has been a source of wonderful mentors to me. I would like to thank the Continuous Waves working group and the staff at LIGO Hanford for all their efforts in making me a scientist. Three people besides Alan stand out in mentoring me - Rejean Dupuis, Xavier Siemens and Joseph Betzweiser. Thank you guys, you have helped me learn a lot of tricks and made me a better person. I especially remember the day after a customary CW telecon, in which I goofed up something and got scolded by Rejean and Xavier sitting on the bridge of the Millikan pond. Or the night at which I called Joe, thinking that

I had a detection and a ticket to Stockholm for Rai Weiss. I would like to thank Xavier Siemens, David Hitlin, Christopher Hirata and Curt Cutler for serving as my thesis committee.

I have been fortunate to have had some fantastic friends over the past five years. Some of the memories that stand out in my mind are midnight arguments with Devdutt, Setu and Naresh, my first year roommates; going snowboarding with Evan and Berit; driving to glacier national park with Justin and Adam; playing settlers with Jake and Phil; sliding on a trashbin can down the Philadelphia Museum of Art steps with Nachiket; the council “jokes” with Krish and Sujit; filming the Oasis movie with Vikram, Shankar, Chaitanya and Tejaswi; playing mafia with Setu, Mansi, Prabha, Uday, Teja and Mayank; cooking with Harmony, JC and all my roommates. All of these and more memories will remain with me. I thank all my friends for the wonderful times that we have all had.

Vaibhav and Sonali have been the elder brother and sister that I have never had. Vaibhav has made me laugh more than anyone else and if I follow only a fraction of all the advice he has given me over the years, I will be wildly successful. Sonali has been a moral compass of sorts and the most fun person to be with. Zeesh has been the perfect companion on my journey through Caltech. He and I studied together for the qualifiers, shared a house for two years and have had numerous good times. Varun has been my partner in crime at Caltech. He has been a close confidant and a stress release valve. I have had a lot of fun going to different national parks and generally driving around with Naresh and Shriharsh. I would also thank Sowmya and Vijay for always being there for me and for all the times that we have cooked together.

Last, but by no means least, I would like to thank my parents, Kirit Patel and Jayshree Patel. They are the ones I aspire towards and idolize. My father and mother both came from a poor family and made a name for themselves through hard work and dedication and I can only hope to emulate them. They have encouraged me to do whatever I want and have been the best friends I have had. I hope our friendship lasts a long time.

Search for Gravitational Waves from a nearby neutron star using barycentric resampling

by

Pinkesh K. Patel

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

Abstract

Rapidly spinning neutron stars in our Galactic neighborhood are promising sources of quasi-monochromatic continuous gravitational waves observable by the current LIGO detectors. I describe a search done on the LIGO S5 data, looking for an isolated neutron star hypothesized to be at a distance of about 100 parsecs. This kind of search is computationally bound and is made possible by the implementation of barycentric resampling, which is described here as well. I also describe the work done at the Hanford LIGO site, while taking data for the Astrowatch program.

Contents

List of Figures

List of Tables

GW Gravitational Waves

CW continuous waves

LSC LIGO Scientific Collaboration

LIGO Laser Interferometer Gravitational-Wave Observatory

CasA Cassiopeia A

LMXB Low Mass X-Ray Binary

SNR Signal to Noise Ratio

NS Neutron Star

Chapter 1

Thesis overview and introduction to Gravitational Waves

This thesis describes a search done for gravitational waves (GW) from a nearby neutron star using LIGO (Laser Interferometer Gravitational-Wave Observatory) data. This suspected neutron star is called Calvera, which is hypothesized to be between 80-260 parsecs from Earth. The search for GWs from Calvera falls in between the traditional all-sky searches for GWs from neutron stars and targeted searches for known pulsars. In the case of Calvera, the sky location is known, but there is no information about its age, its spin frequency or frequency derivatives. Certain assumptions are made about the age and distance of the suspected neutron star and a broadband frequency and frequency derivative search is conducted. Such broadband searches are computationally bound and this thesis describes the technical challenges involved in conducting such a search. It also describes the implementation of barycentric resampling, without which this search would have been impossible to conduct, given current computational resources.

This chapter briefly introduces GWs and their effects on a detector like LIGO. Also discussed in this chapter are the potential sources of GWs which include compact binary coalescences (CBC) of neutron star-neutron star, neutron star-black hole and black hole-black hole systems, bursts of gravitational radiation from sources like core collapse supernovae, a gravitational wave stochastic background and continuous quasi-periodic GWs (CW). The LIGO detectors are introduced in Chapter 2. The interferometers are described from an operations perspective and a procedure for maintaining them in data collection mode is also described. Chapter 3 describes neutron stars and CW sources. Chapter 4 is a summary of techniques used to search for GWs from CW sources. Chapter 5 focuses on one of these techniques and the implementation of barycentric resampling. Calvera is introduced in Chapter 6 and a search for CWs from it using barycentric resampling is presented in Chapter 7. Conclusions and future work prospects are discussed in Chapter 8.

1.1 General Relativity and Gravitational Waves

General Relativity (GR) is one of the pillars of modern physics. It was developed by Albert Einstein in 1915. It is a theory that describes gravity as a curvature in the four dimensional spacetime. It predicts the existence of exotic objects like neutron stars and black holes and it forms the basis of our understanding of cosmology. One of its predictions is the existence of gravitational waves (GW), which are a strain of spacetime propogating at the speed of light. These GWs are practically unaffected by intervening matter between the source and us and are exceedingly weak by the time they reach us. They have been indirectly observed by noting the evolution of the orbits of the binary pulsar system PSR 1913+16, discovered by Russell Hulse and Joseph Taylor. For this discovery, they received the Nobel Prize in Physics in 1993. The evolution of the orbits matches the theoretical predictions of GR, which predicts the loss of angular momentum carried away by GWs, to within 0.3% [1]. However, given the weak nature of GWs, their direct detection has been a challenge. LIGO, the Laser Interferometer Gravitational–Wave Observatory aims to detect and study GWs.

The major result in GR is Einstein's equation, which in geometrized units of $G = c = 1$ is written as [2]

$$G_{\alpha\beta} = 8\pi T_{\alpha\beta}, \tag{1.1}$$

where $G_{\alpha\beta}$ is the Einstein tensor and $T_{\alpha\beta}$ is the stress energy tensor. The indices α

and β run from 0 to 3 and correspond to the spacetime coordinates of time and the three spatial axes. The Einstein tensor which also contains the metric, specifies the geometry of spacetime. The stress-energy tensor $T_{\alpha\beta}$ specifies the source terms from all forms of matter and energy. The stress-energy tensor, which is symmetric, can be divided into three parts, with T_{00} specifying the energy density, $T_{0\mu} = T_{\mu 0}$ being the energy flux or the momentum densities and the remaining 3×3 matrix being the stress tensor. While equation 1.1 looks compact and elegant, it is non-linear and has not been completely solved analytically.

The Einstein equation 1.1 can be solved under certain assumptions. One such assumption is the linearized weak field limit of gravitational waves far from a source such that $T_{\alpha\beta} = 0$. The metric can then be written as a small perturbation $h_{\alpha\beta}(x)$ of the metric of flat spacetime $\eta_{\alpha\beta}$, also known as the Minkowski metric [2]. This means

$$g_{\alpha\beta}(x) = \eta_{\alpha\beta} + h_{\alpha\beta}(x), \tag{1.2}$$

where the Minkowski metric can be written as

$$\eta_{\alpha\beta} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{1.3}$$

Here the definition is that the coordinates are in the order of (t, x, y, z), with t being time and the remaining being the three spatial coordinates. By choosing a transverse-traceless gauge, Einstein's equation can be written as a flat space wave equation

$$\left(-\frac{\partial^2}{\partial t^2} + \vec{\nabla}^2 \right) h_{\alpha\beta}(x) = 0 \tag{1.4}$$

where $\vec{\nabla}^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$. The generic solution to equation 1.4 is the general wave solution

$$h_{\alpha\beta}(x) = A_{\alpha\beta} e^{i\mathbf{k}\cdot\mathbf{x}}, \tag{1.5}$$

where, $A_{\alpha\beta}$ is a 4×4 amplitude matrix and $\mathbf{k}$ is a four dimensional wave vector, such that $\mathbf{k} \cdot \mathbf{x}$ is

$$\mathbf{k} \cdot \mathbf{x} = \omega t + \vec{k} \cdot \vec{x}. \tag{1.6}$$

Using equations 1.4 and 1.5, the fact that we have chosen a transverse-traceless gauge and orienting the propogation of the wave along the z axis, the amplitude matrix becomes

$$A_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{1.7}$$

Using equations 1.7 and 1.4 we can write the gravitational wave strain amplitude $h_{\alpha\beta}$ as

$$h_{\alpha\beta}(t,z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{+} & h_{\times} & 0 \\ 0 & h_{\times} & -h_{+} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} e^{i\omega(t-z)}. \tag{1.8}$$

Here h_{+} and $h_{\times}$ are the “plus” and “cross” polarizations of the gravitational wave. There are only two polarizations that exist according to GR. A schematic of their effect on a ring of particles if a gravitational wave incident into the page and consisting of only one of the two polzarizations can be seen in Figure 1.1. As can be seen in Figure 1.1, gravitational wave polarizations are invariant under a π radian rotation. In the language of quantum gravity this means that the exchange particle for gravity, the graviton has a spin of $\frac{2\pi}{\pi} = 2$.

Electromagnetic radiation can be produced by accelerating charge dipoles, charge-current dipoles and other multipoles. However in the case of gravitational radiation, the conservation of mass prevents radiation due to an accelerating mass monopole. The emission of gravitational radiation via an accelerated mass dipole is forbidden by conservation of linear momentum. Conservation of angular momentum similarly prevents radiation due to the acceleration of current dipoles.

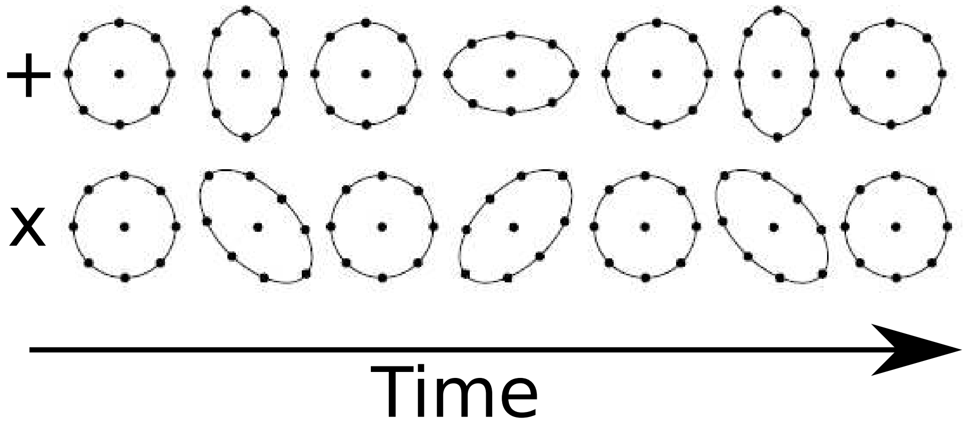

Figure 1.1: The effect of a gravitational wave consisting of only one of the two polarizations “+” and “×” on a ring of particles. The lines are just a guide to the eye. The GW would be passing perpendicularly through the page and the figure shows snapshots in time as the wave passes. It shows one whole wavelength for each polarization.

Thus the first multipole that can produce gravitational radiation is from the acceleration of a mass quadrupole. Such accelerating mass quadrupoles are common in nature; examples include orbiting binary stars, rotating non-axisymmetric stars etc.

1.2 Gravitational Wave Sources

The plausible gravitational wave sources that have been studied so far can be divided into four broad categories. Compact Binary Coalescences (CBC) have a transient, but well modeled waveform, usually divided into three parts called the Inspiral, Merger and Ringdown phases. Unmodeled transient waveforms are known as gravitational wave bursts. Among the long-lived forms are the modeled sinuosoidal waveforms emitted by spinning neutron stars, known as continuous waves (CW) and the stochastic gravitational wave background from the Big Bang or from a large number of weak astrophysical sources. Thus far, no searches for

GWs from these searches have yielded plausible detections.

1.2.1 Compact Binary Coalescences

Compact binary coalescences are perhaps the most highly anticipated sources of gravitational waves detectable by LIGO. They consist of three distinct phases of gravitational wave emission. The first stage consists of the inspiralling of two compact objects like two neutron stars, one neutron star and one black hole or two black holes. These compact objects lose energy and angular momentum in the form of gravitational radiation and inspiral in, eventually merging. The waveform is well understood for the inspiral phase and it depends on the masses of the two objects inspiralling. The merger waveform is however not well understood, but it is under active research in the field of numerical relativity. After the merger phase, the distorted black hole relaxes to a stationary Kerr state through a strong emission of gravitational waves, which is known as a ringdown.

An inspiral produces a signal that is commonly known as a chirp, which is a sinusoid increasing in amplitude and frequency until the merger phase begins. Since the frequency of the signal increases as a function of time, inspiral signals enter the LIGO band of interest from about 40 Hz onwards only in the last few minutes or seconds (depending on the masses) before the merger. This makes these signals transient in nature. A binary neutron star system consisting of two typical neutron stars of mass $1.4M_\odot$ would undergo about 1630 cycles in the LIGO band before merging in 26 seconds. A typical chirp signal can be seen in Figure 1.2.

The analysis of LIGO data for CBC waveforms includes a match filtering step, where analytical or phenomenological waveforms are used as filter templates. The output of the filter is then thresholded. Events whose SNR crosses the threshold are considered triggers for further analysis. The threshold is set low enough to let a potential signal through, but keep the number of triggers to a manageable value. The triggers are first generated independently for each detector. This is followed by a coincidence analysis, with appropriate sized windows allowing for light travel time between the the detectors involved. Such coincidence analysis helps reduce

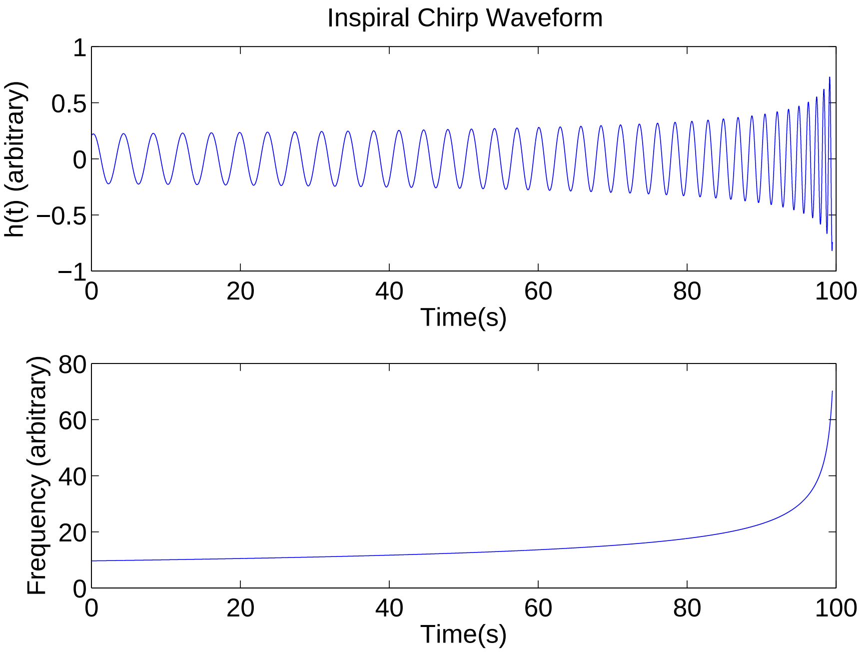

Figure 1.2: An example of an inspiral gravitational wave signal, emitted by a CBC. This plot is in arbitrary units and only to give a qualitative idea.

the number of total triggers by suppressing the background. These triggers are then subjected to consistency checks like the χ^2 [3] and r^2 [4] tests. CBC searches conducted by the LSC are optimized for detection of a signal (i.e. the cuts are loose). If no signals are found, then upper limits are set on the rate of CBCs per Milky Way Equivalent Galaxy or per sensitive search volume in Mpc^3. The searches conducted so far for CBCs can be found in [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15].

1.2.2 Bursts

A second category of transient searches involve unmodeled "bursts" of gravitational radiation. These bursts are usually assumed to be shorter than about 1 second, but with significant enough power to rise above the noise fluctuations in the detector. The possible astrophysical sources include asymmetric core collapse supernovae [16], ringdowns of black hole-black hole mergers [13], whose inspiral phases are not within the LIGO band and thus only the high frequency ringdowns can be seen, and disturbances along hypothesized structures known as cosmic strings. Some other astrophysical phenomena are used as external triggers for burst searches such as extra-galactic Gamma-ray bursts (GRBs) and galactic Soft-Gamma repeaters (SGRs), both of which are highly energetic events and which could potentially emit prodigious amounts of gravitational radiation. Short GRBs (lasting only a few seconds to minutes) are suspected to be neutron star-neutron star or neutron star-black hole mergers. SGRs are believed to be neutron stars with huge magnetic fields which flare up from time to time and emit prodigious amounts of electromagnetic radiation. A discussion of an SGR search in which I contributed with vetoes will be discussed in Chapter 2. Besides all this, burst searches aim to be inclusive in case a completely unmodeled and ill-understood source is emitting gravitational radiation.

The search techniques for burst searches always include a correlation between detectors, since the LIGO and Virgo detectors undergo numerous “glitches”, which can mimic a burst signal. These glitches and other noise fluctuations are however unlikely to happen within the light travel time between any two detectors and thus coincidence is an efficient technique to reduce the background noise. The LIGO Scientific Collaboration (LSC) has developed many different search pipelines, which include techniques like excess power, cross-correlation and coherent methods. Some of the searches done by the LSC for burst sources can be found in [17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33].

1.2.3 Stochastic Background

Since gravitational waves couple weakly with matter, most of the gravitational radiation produced just after the Big-Bang would still exist as a stochastic background. Detecting these GWs would provide unique information on the earliest moments of the universe. Other sources of stochastic GWs include phase-transitions in the early universe, cosmic strings, topological defects formed during the symmetry-breaking phase transitions in the early universe, or other unresolved astrophysical sources like CBCs, neutron stars or supernovae. Most models of the spectra expected for a stochastic background approximate it by a power law such that $\Omega_{\mathrm{GW}}(f) \propto f^{\alpha}$, where α is different for different mechanisms. The standard "slow roll inflation" cosmological model predicts predicts a flat spectrum and thus an $\alpha = 0$.

GW stochastic background is expected to be so weak that it would not rise above the noise floor of the detector, thus requiring an integration over time to increase the SNR. If two independent detectors are used for a cross-correlation measurement, then the signal would increase as a function of the integration time T, while the noise would only increase as $\sqrt{T}$. However in the case of some detectors, the noise is not independent, for example the two co-located detectors at Hanford (H1 and H2, to be discussed in Chapter 2) share a lot of noise sources. Even in the case of detectors which are seperated geographically like L1 and H1, there is a price to pay in the form of a reduction of the signal due to the separation time delay between the two detectors and the misalignment of their arms. The LSC's stochastic background searches can be found in [34, 35, 36, 37, 38, 39].

1.2.4 Continuous Waves

Continuous GWs are expected to be emitted by neutron stars, which are observed as pulsars spinning at frequencies ranging from a few milli-Hz to about a kHz. Spinning neutron stars would emit gravitational radiation if there were some non-axisymmetric asymmetry in their crust, or due to some fluidic asymmetries or due to accretion of matter. These and other mechanisms of emission are discussed in Chapter 3. The detection of continuous GWs involves the integration of long stretches of data to dig out signals buried deep in the noise. The search methods are discussed in detail in Chapter 4.

Chapter 2

LIGO Detectors and their Operation

This chapter discusses the basics of the LIGO interferometric detectors and their operation. It will discuss how an interferometer is brought into a resonance lock and maintained there. A brief discussion of the data readout scheme then follows. After a short history of LIGO's science runs, Astrowatch is introduced and my participation in it is discussed. Improvements and detector characterization work done during Astrowatch follow.

2.1 LIGO overview

The LIGO detectors were constructed in the 1990's by the California Institute of Technology and the Massachusetts Institute of Technology, with funding from the National Science Foundation (NSF). Three kilometer scale detectors were built with one 4-kilometer interferometer each in Hanford, Washington (H1) and Livingston, Louisiana (L1). At the Hanford site, one 2-kilometer interferometer called H2 was also built.

A timeline showing the progress of LIGO over the years is shown in Figure 2.1. This figure shows the three stages of LIGO, Initial LIGO which ran from about 2000 to the end of 2007, enchanced LIGO from 2009 to the present and the construction and funding of advanced LIGO which is expected to start operation in 2014. Figure 2.1 also shows the strain sensitivities achieved over the years in the most sensitive part of the LIGO spectrum. The detectors have been operated

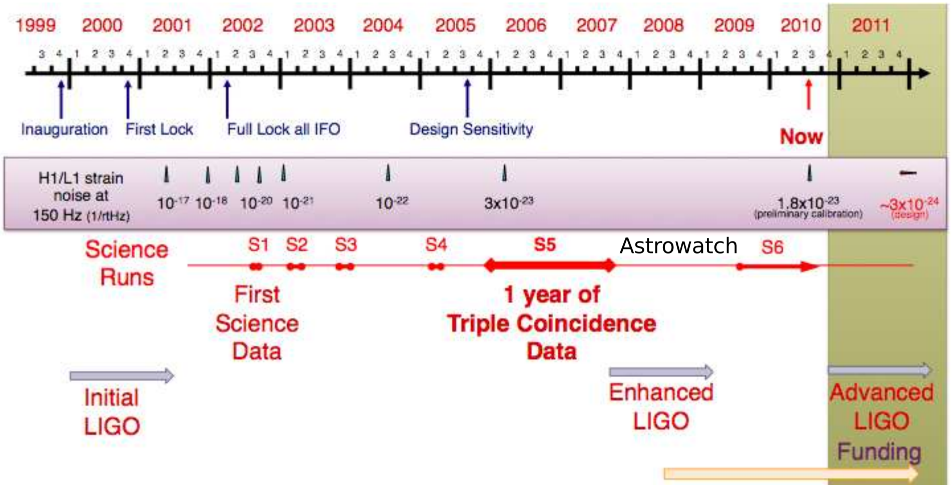

Figure 2.1: A timeline of LIGO. The figure shows the three stages of LIGO, Initial LIGO from 2000 to the end of 2007 (science runs S1 through S5), followed by Enhanced LIGO from 2009 to the present (science run S6) and Advanced LIGO, which is currently under construction. Astrowatch took place in 2008 between S5 and S6. The strain senstivities achieved over the years in the most sensitive part of the LIGO spectrum are also shown.

in an observational mode from time to time, during which they are run for data collection for a significant amount of time. There have been six such official science runs that have been called S1 through S6. The senstivities achieved during S5 met the initial LIGO design sensitivity requirement and these are shown in Figure 2.2. S5 was the longest science run with approximately one year of triple coincidence time (time during which all three LIGO detectors were operational). S5 ran for a calendar time of about 2 years from November 2005 to October 2007. At the end of S5, the H1 and L1 detectors were taken offline and upgraded to enhanced LIGO during 2008. During that time H2 was run under an Astrowatch mode. For more details please see Section 2.3.

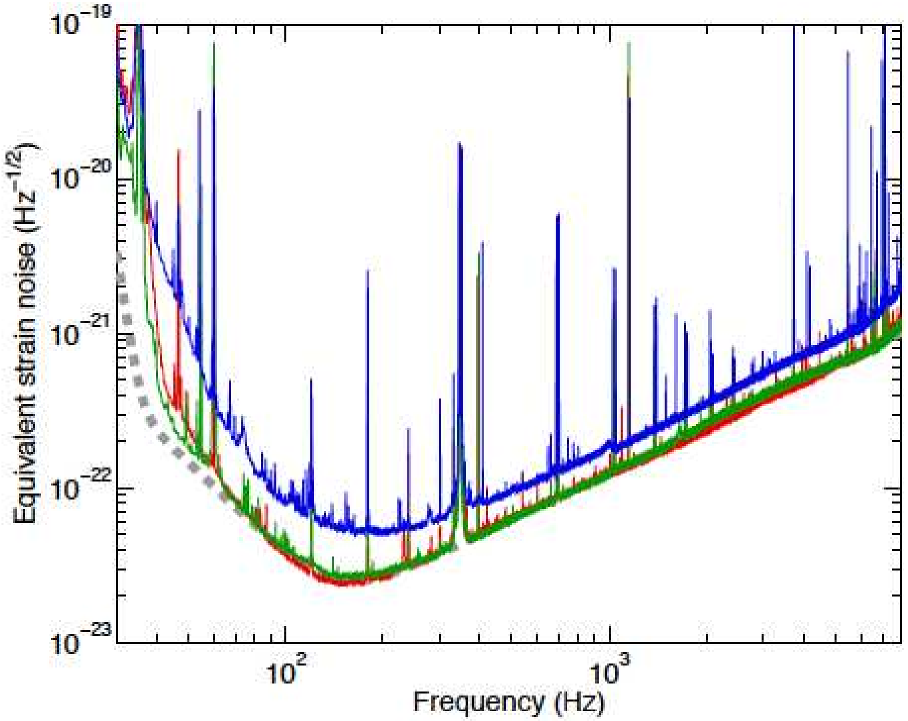

Figure 2.2: Strain sensitivities, expressed as amplitude spectral densities of detector noise converted to equivalent GW strain. The vertical axis denotes the rms strain noise in 1 Hz of bandwidth. Shown are typical high sensitivity spectra for each of the three interferometers (red: H1; blue: H2; green: L1), along with the design goal for the 4-km detectors (dashed grey). Figure and caption courtesy [40].

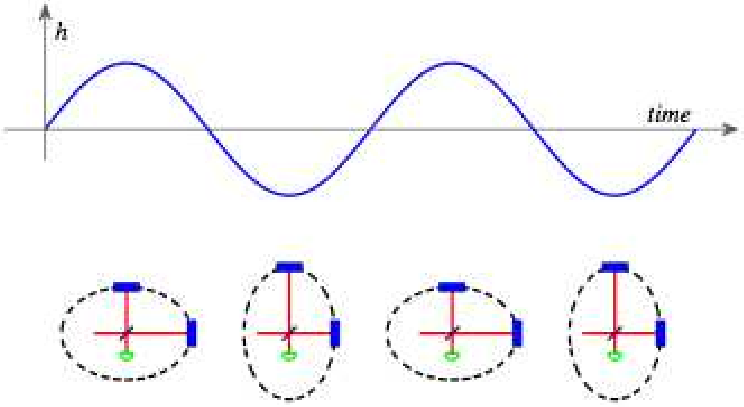

Figure 2.3: Cartoon showing the effect of a passing gravitational wave on a Michelson interferometer. The wave is assumed to have the plus polarization and is propagating perpendicular to the diagram. Figure courtesy [40].

2.1.1 Michelson Interferometers

The LIGO interferometers are, at the most basic level, Michelson interferometers, which are a common optical configuration, designed to create an interference pattern by splitting coherent light into two arms and reflecting them off a couple of mirrors. A cartoon showing the effect of a passing gravitational wave on a Michelson is shown in Figure 2.3. The essential idea is to maintain an interferometer in a "locked" position, with the difference in length of its arms fixed by measuring it from the light bounced back from it and feeding that back to the mirrors. The error signal that needs to be fed back would then be the gravitational wave signal, since it measures the deviation from the resonant working point. An instrument working on such a principle is known as a null instrument.

If the two arms are termed the x arm and the y arm, then the phase change of the light due to round-trip travel in each arm is given by

$$\phi_x = \frac{2\pi}{\lambda} 2L_x \tag{2.1}$$

and

$$\phi_y = \frac{2\pi}{\lambda} 2L_y, \tag{2.2}$$

where λ is the wavelength of the light and L_x and L_y are the arm lengths. At the anti-symmetric (AS) of the interferometer (which is kept dark), the power of light received is a function of the difference in arm lengths and is given by

$$P_{\mathrm{AS}} = P_0 \sin^2(\phi_x - \phi_y) = P_0 \sin^2\left(\frac{2\pi}{\lambda}(2\Delta L)\right), \tag{2.3}$$

where P_{AS} is the power seen at the AS port, P_0 is the input laser power and $\Delta L = L_x - L_y$. The incident gravitational wave strain can be related to the length changes in the arm by

$$h = \frac{\Delta L}{L}. \tag{2.4}$$

Thus the power seen at the AS port in terms of the gravitational wave strain is

$$P_{\mathrm{AS}} = P_0 \sin^2\left(\frac{2\pi h L}{\lambda}\right). \tag{2.5}$$

Using typical numbers for the LIGO detectors like $\lambda = 1064$nm, $L = 4000$m and $h < 10^{-21}$ in equation 2.5 $\frac{2\pi hL}{\lambda} \approx 10^{-8}$. Thus the $\sin^2\left(\frac{2\pi hL}{\lambda}\right)$ term in equation 2.5 can be written as $\left(\frac{2\pi hL}{\lambda}\right)^2 \approx 10^{-16}$. Thus the power change at the AS port due to a GW is extremely small. It is not possible to distinguish the sign of h using equation 2.5 and thus using the power as a readout scheme is not the optimal technique.

2.1.2 Fabry-Perot Cavities

The frequency band that first generation detectors like LIGO could possibly detect gravitational waves in was constrained by ground seismic noise. This meant that the detectors would only be sensitive to frequencies greater than about 30 Hz. Astrophysical sources were expected to have maximum frequencies in the range of a few kHz. These frequency constraints set the optimal length of the interferometers to a few hundred km or so. However building an interferometer that is a hundred km long is impractical and thus instead of having a simple Michelson configuration, two more mirrors were added to each of the arms, which would form a Fabry-Perot cavity with the inner mirrors. These cavities increase the light recycling time by a factor of about 75, increasing the effective distance from 4 km to around 300 km or so. This also increases the power recycling in the cavities. The gravitational wave sensitivity is proportional to the phase change caused by the incoming GW and the power recycling in the interferometer.

Another method of increasing the GW signal in the interferometer is to add a power recycling mirror (PRM) before the beam splitter (BS), which would reflect the light coming out of the interferometer back into it. The power recycling factor, which is the number that the power gets multiplied by due to the addition of the PRM depends sensitively on the optical losses in the as-built interferometer; it is 60 for H1, 45 for L1 and 70 for H2 [40]. Thus with the addition of the PRM and the end test masses (ETMs), with an input power of 4.5 W for H1 and L1, the power stored in the arm cavities is 20 kW and 15 kW respectively [40]. About 10 kW are stored in H2's arm cavities with an input of about 2 W [40]. Thus the initial LIGO interferometers are known as power-recycled Fabry-Perot Michelson interferometers and their configurations are shown schematically in Figure 2.4.

The LIGO input optics add two pairs of RF sidebands to the laser before it is input into the interferometer (see Section 2.2.1). The first pair of RF sidebands have a phase modulation at 62.5 MHz and the second pair at 25 MHz [40]. The LIGO interferometers are built with a Schnupp asymmetry, which makes the two arm lengths unequal by precisely an integral multiple of the carrier light's wavelength. At the operating point, i.e. in the resonant lock mode, the carrier light is resonant in the arm and recycling cavities and on a Michelson fringe. The RF sidebands resonate differently. One pair of the RF sidebands is resonant in the recylcing Michelson cavity, but not in the arm cavities. The other pair is not resonant in the interferometer at all and is reflected back from the recycling mirror. In this configuration, the RF phase modulation sidebands, which are reflected from the cavity input mirror, serve as a local oscillator to mix with the carrier field. The

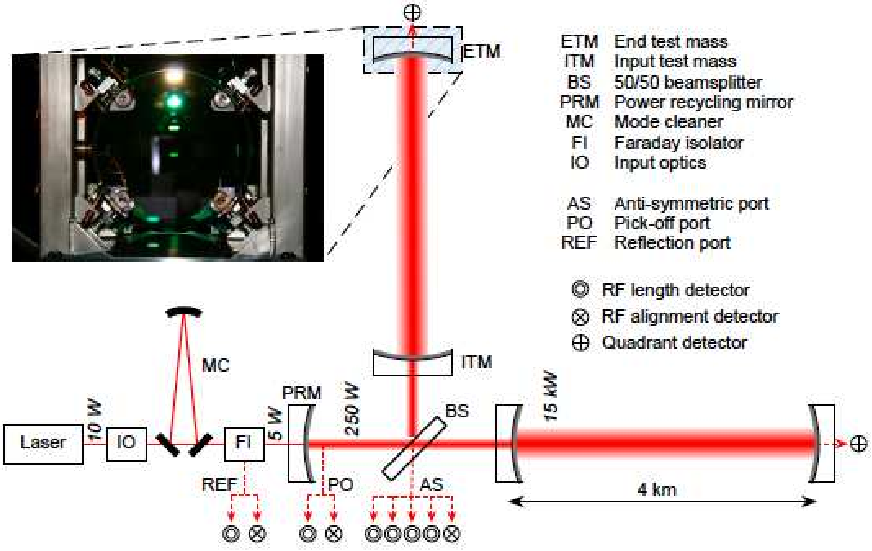

Figure 2.4: Optical and sensing configuration of the LIGO 4 km interferometers. The IO block includes laser frequency and amplitude stabilization, and electro-optic phase modulators. The power recycling cavity is formed between the PRM and the two ITMs, and contains the BS. The inset photo shows an input test mass mirror in its pendulum suspension. The near face has a highly reflective coating for the infrared laser light, but transmits visible light. Through it one can see mirror actuators arranged in a square pattern near the mirror perimeter. Figure and caption courtesy [40].

GW signal can then be read out as a phase shift that the carrier light experiences in reflection, which can be read out as the RF amplitude modulation. This RF amplitude modulation is linear in amplitude for small deviations. For more details, please see references [40, 41, 42].

2.2 LIGO subsystems

2.2.1 Pre-Stabilized Laser

The pre-stabilized laser (PSL) is the subsystem which contains the laser, the input optics (IO), the input mode cleaner (MC) and the Faraday isolator (FI) as shown in Figure 2.4. The laser used in initial LIGO was a diode pumped Nd:YAG laser with a power output of around 10 W at a single frequency of 1064 nm [40]. The IO block includes a couple of EOMs that are used to add in the RF sidebands that are used for the Schnupp transmission locking and PDH reflection locking; a reference cavity, which serves as part of a frequency stabilization servo (FSS); an intensity stabilization servo (ISS); and a halfwave plate, which is used to adjust the input power into the interferometer. The mode cleaner is used to optically filter the spatial modes to pass only the spatial 00 mode, which resonates in the arm cavities. It is also used to aid in stabilizing the laser frequency so that the light will resonate in the long arms.

2.2.2 Suspension System

The suspension system (SUS) is the system that is responsible for supporting and stabilizing the suspended optics of the interferometer. An example of this can be seen in the inset of Figure 2.4. It shows an optic with 4 optical sensor and electromagnetic actuators (OSEMs) on the face and one side OSEM. These OSEMs act on little magnets that are glued on each optic. The OSEMs also sense the position of the optic by using a shadow optical sensor. Thus it acts as a means of measuring the position of each optic and also as a method of actuation.

The 5 OSEMs are divided in 4 different basis vectors (which are not independent) of pitch, yaw, position and side. The side OSEM is the only one that actuates on and reads the side degree of freedom. The pitch motion is controlled by using the top two OSEMs as one unit and the the bottom two as another unit, but with the opposite sign. The position motion can be achieved by using both the top and bottom units with the same sign. The yaw motion divides up the left two and right two OSEMs into different but opposite units.

Under normal operating conditions, the shadow sensor readouts from the OSEMs are used to calculate the position, pitch, yaw and side parameters for each optic and this is then fed back into the OSEM actuators with a simple low-pass filter to keep the optic steady from moving, near the suspension resonant frequencies (around 1 Hz). The optics usually sway due to seismic or inherent interferometric noise and they need to be kept near their working position at all times in order to maintain lock. These servo loops accomplish this task under normal circumstances. Under certain conditions like earthquakes or the failure of the high-voltage line that controls the OSEMs, the control loops can malfunction or go out of range. The optic can then sway more than desired and in order to avoid this, watchdog sensors are used to stop the control system from actuating the optics beyond a certain point.

Most of the large optics are also equipped with an optical lever, which is a laser that is bounced off their face and into a sensor. Changes in the angle of incidence and reflection off the optic would have a linear effect on the position at which the laser beam hits the sensor. This readout can also be fed back into the suspension control system. The optical levers are much better than the OSEMs at high frequency changes in the optic's angles and also have a larger working range.

One of the most common chores that need to be done after an earthquake, which trips some watchdogs and switches off the SUS damping systems, is to reset the watchdog controls and get the optic back to a controlled state. This usually involves a series of operations like checking if the optics are within nominal parameters, if it is safe to switch on the watchdogs and if optical lever damping is required. If optical lever damping is required, then its control loop needs to be switched on and then one needs to wait for the optic to return to within the nominal working range of the OSEMs and then switch the OSEM loops on and turn the optical lever loops off. All of these tasks are easily accomplished by a cleverly written script, which mimics an operator's action. After spending many hours manually damping each of the 10 or so optics, I wrote a simple python script which did the damping automatically. This script has now become a standard for damping all the optics for both interferometers at the Hanford site.

2.2.3 Length Sensing and Control

The length sensing and control (LSC) system is responsible for maintaining the interferometer in a resonant lock by changing the position degrees of freedom. It uses as input, the in-phase (I) and quadrature-phase (Q) readouts taken at three different points in the interferometers. These points are the AS port, the pick-off (PO) after the PRM and the reflected beam (REF) as shown in Figure 2.4.

The LSC system controls four length degrees of freedom referred to as the differential arm (DARM), Michelson (MICH), power recycling cavity (PRC) and the common mode arm (CARM). The DARM_ERR or the error signal in the DARM loop is also the gravitational wave channel. The DARM signal is derived from the AS port and this loop tries to minimize the carrier light coming out of the AS port by keeping the differential motion of the arms to a minimum. The MICH loop serves to keep the Michelson cavity constant, i.e. the distance between the PRM and two input test masses (ITMs) constant. The MICH loop gets its readout from the Q-phase of the PO photodiode. The PRC loop keeps the recycling cavity resonant to the sidebands and the carrier by driving the PRM and it also gets its readout from the PO photodiode, but it uses the I-phase. The CARM signal is used to control the common-mode motion of the arms, i.e. the motion that is correlated with each other, unlike DARM, which is anti-correlated. The CARM loop signal is derived from the REF photodiodes and instead of feeding it to the optics like the other three signals, it is fed to the voltage controlled oscillator (VCO) of the laser, which controls the frequency of the laser. The CARM signal is fed back to the arms during lock acquisition, but is switched to the VCO during the operation in observing mode.

2.2.4 Angular Sensing and Control

The angular sensing and control (ASC) system consists of a series of wavefront sensors (WFSs) and quadrant photo diodes (QPDs), which keep the angular degrees of freedom of the interferometer constant. The two QPDs are at the two ends of the interferometer after the ETMs and sense using the transmission light. The WFS are located along with the major photodiodes in REF, PO and AS ports. The angular sensing and control system is extremely important for high power operation and has been improved in the recent enhanced LIGO upgrade.

2.2.5 Thermal Compensation System

The major LIGO optics have reflectivities in the range of 99.997%. The reflective coating on their surface is specified to abosrb no more than about 1 ppm (parts per million) of laser power, and the substrate is rated to absorb about 3–4 ppm. The ETMs and ITMs were designed to absorb some of the kWs of power circulating in the arm cavities and distort into the correct shape for optimal operation. However after the first few runs, it was discovered that the mirrors were absorbing a lot more heat than was expected.

The thermal compensation system (TCS) was installed to correct this problem. It is designed to apply heating to the test masses by shining a CO_2 laser operating at 10.6 μm wavelength on to the inner faces of the ITMs of the interferometers. This laser can be illuminated via two different masks, one with a central heating pattern and another with an annulus pattern. The central heating pattern can be used to rectify insufficient heating of the optic and the annulus pattern is used to correct for excessive heating of the optic. A Bessel mask further downstream from the masks clips the higher order maxima of the Airy diffraction pattern, leaving only the central lobe of the Airy disk. The incident power is varied using a polarizer. For more details please see [43, 44].

2.3 Astrowatch

After the major science run S5, which started in November 2005 and ended in October 2007 was followed by an upgrade of the two 4 km detectors H1 and L1 as shown in Figure 2.1. These upgrades were called the enhanced LIGO upgrade. Enhanced LIGO includes an increase in laser power from a maximum of 10 W to 35 W and the implementation of a DC readout scheme. These upgrades took a little more than a year of commissioning work and during that time, both H1 and L1 were completely inoperative. It was decided to run H2, the 2 km interferometer during this time as an Astrowatch instrument, which would try to collect as much data as possible, so as to not miss any opportunistic events like galactic supernovae or close GRBs or SGRs etc.

Due to an acute shortage of skilled manpower, most of which was being used for the enchanced LIGO upgrade, the LSC decided to let graduate students operate and run H2 on a voluntary basis. I was one such volunteer along with Berit Behnke (Albert Einstein Institute), Jericho Cain (Univ of Mississippi), Nicolas Fotopoulus (Univ of Wisconsin Milwaukee), Evan Goetz (Univ of Michigan), Shivaraj Kandasamy (Univ of Minnesota), Satyaprakash Mohapatra (Univ of Amherst), Adam Mullavey (Aust. Natl. Univ), Philip Roberts (Andrews Univ), Jacob Slutsky (Louisiana State Univ), Szymon Steplewski (Univ of Washington Pullman), Matthew West (Syracuse Univ) and Junyi Zhang (Univ of Michigan).

Astrowatch was a best-effort operation and thus shifts were organized for after hours and on weekends to keep the interferometer running when there were no people actively working on H1 commissioning. At first this meant that shifts would run from 10 PM local time to about 6 AM. However it was observed later on that the interferometer was very stable from around midnight or so and rarely needed intervention from an operator. In addition, anthropogenic noise was reduced during the night hours and this meant that the interferometer could acquire a resonant lock on its own most of the time. Thus due to a lack of manpower, it was decided to change the shifts to 6 PM to 2 AM, so that all the available manpower would be used when the interferometer required the most intervention.

2.3.1 Astrowatch highlights

Data that has been certified to be science quality is referred to as “science mode” data. An operator or an experienced astrowatcher is required to manually certify data as science data. It requires a series of sanity checks on if there is any work going on in the Large Vacuum Equipment Area (LVEA), which houses the interferometer, if the data servers are online, if there are no unknown injections taking place, etc. After the data has been certified, it can be used by analysis groups after calibration.

Since science mode requires a manual input, if the detector reached a stage where the data that is being collected is as good as science data, but is not certified, it is called “Up” time. Most of Up time is clustered around the nights and early mornings, when no one is around. However some Up time occurs when the operator on shift is unable to put the detector in science mode due to someone working the LVEA or if there is some other maintenance going on. These Up times are often salvageable and as good as science mode data. Working with Jacob Slutsky, we certified some of the Up time as usable for an LSC search. For details, please see Section 2.3.3.

Figure 2.5 shows a pie-chart of the various states of the interferometer during Astrowatch. Besides Science and Up time, the other states are off, in which the input optics were off and no coherent light was entering the interferometer; mode cleaner locked, in which the input optics were on and locked optimally and common mode, in which the interferometer is locked, but not at full power.

A figure of merit used to assess the quality of the data is the inspiral range, which is defined as the distance at which a neutron star–neutron star inspiral whose orientation has been averaged over all possible orientations, can been detected with an SNR of 8. It hovers in the range of 14 Mpc for the two 4 km detectors and about half that for H2. It is a good figure of merit because such an inspiral would sample most of the sensitive LIGO frequency band and thus it is an average of

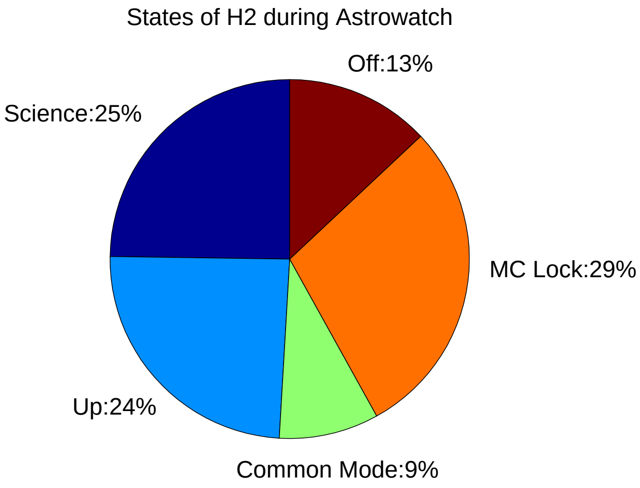

Figure 2.5: Pie chart of various states of H2 during astrowatch.

the entire noise curve. Figure 2.6 shows the inspiral range of H2 for both science mode and Up mode data during Astrowatch.

2.3.2 Acoustic Tuning

H2 required maintenance from time to time during Astrowatch. This was divided in various different tasks, which included things to check daily, scripts to run once in a while etc. One of the tasks that I worked on with Robert Schofield of U Oregon was to identify paths through which acoustic noise can enter the detector's GW readout. As shown in Figure 2.4, there are multiple photodiodes looking at the AS port. There are four major AS photo diodes (ASPDs), which have a low noise performance and are used when the detector is locked. They are called ASPD1, ASPD2, ASPD3 and ASPD4. For acquiring lock, another photodiode called ASPD5 is used, which does not operate as well as the other ASPDs, but has a greater tolerance to large fluctuations of light, which are likely to happen during lock acquisition.

Acoustic noise can manifest itself in the GW channel through any defects/deadspots on any of the ASPDs. There are a series of mirrors and beam splitters that guide

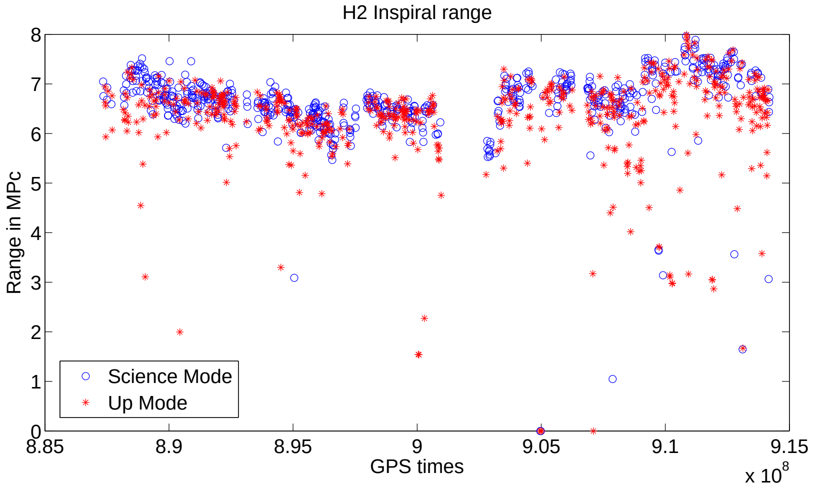

Figure 2.6: Inspiral range of H2 during Astrowatch.

the output of the AS port to each of the ASPDs. If the laser beam is incident near a defect on an ASPD, it can raster across it through the shaking of the post holding one of the steering mirrors or beam splitters. The post is likely to be shaken by ambient acoustic noise, since the AS tables are in air. These tables are shielded acoustically, but in spite of this, some noise manages to creep in once in a while. Figure 2.7 shows a measurement done on the GW channel, while running on all four ASPDs and then each one of them one by one.

A way to identify some of the sources of this noise is to run the detector on one ASPD at a time, which uses multiple speakers to output a large amplitude comb, 25 Hz apart of acoustic signal in the AS port table room. This noise would show up as a series of lines in the GW channel, if it is run using one ASPD at a time. The beam that hits an ASPD is then adjusted to minimize the comb of lines in the GW channel. Figures 2.9 and 2.8 show the results of the acoustic tuning.

2.3.3 SGR search data quality

During the course of Astrowatch between August 21 2008 and September 5 2008, over a hundred SGR flares were observed. The SGR named SGR-0501+4516 was located at 0.8 ± 0.4 kpc from the Galactic center, was an order of magnitude closer

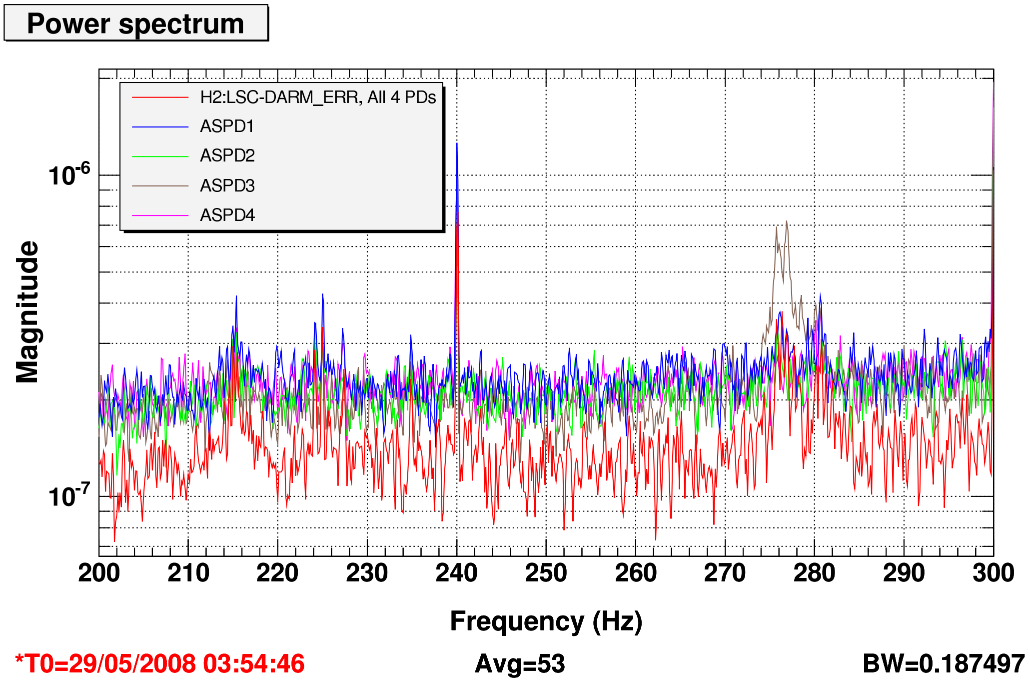

Figure 2.7: Figure showing the GW channel DARM_ERR, when running on all four ASPDs and then each one of them one by one. ASPD3 is the only one to show a peak around 275 Hz and thus some acoustic tuning is required.

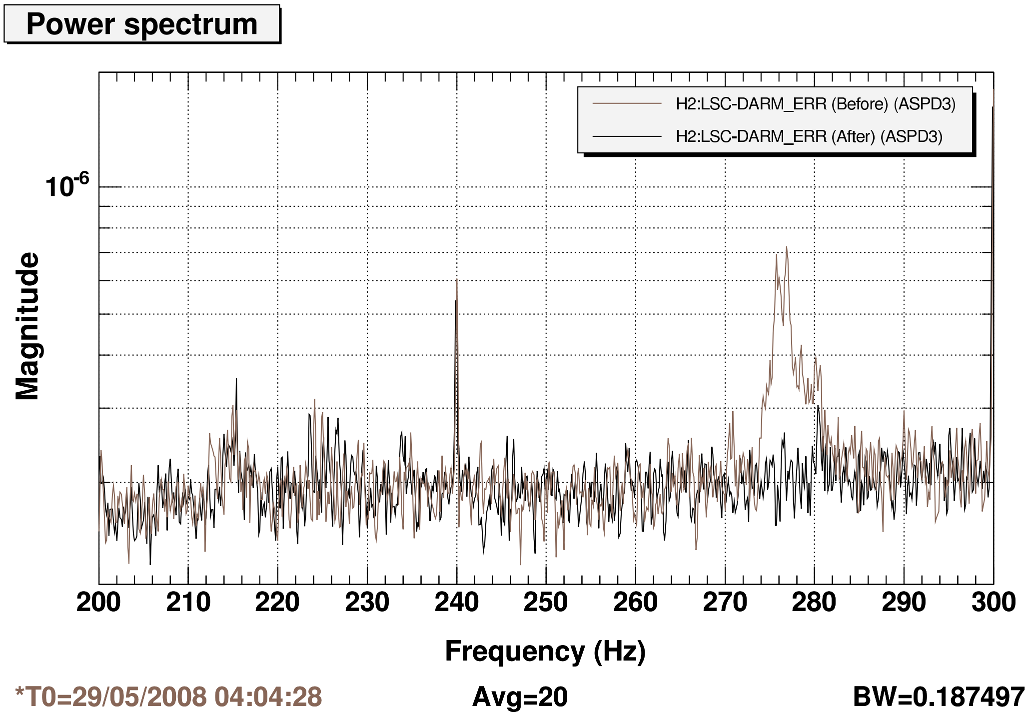

Figure 2.8: The GW channel running on ASPD3 alone, before and after the acoustic tuning was performed.

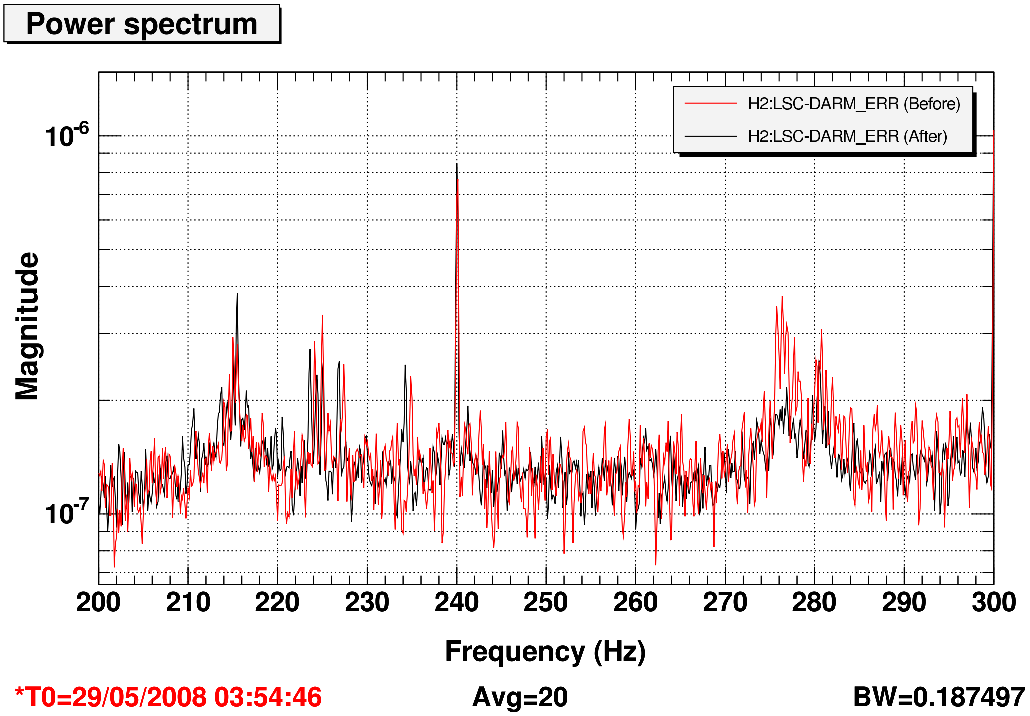

Figure 2.9: The GW channel running on all fours ASPDs, before and after the acoustic tuning was performed on ASPD3.

than any SGR burst that had occured during LIGO's previous observations [31]. Thus in spite of not possessing data from the 4 km interferometers, which are more sensitive than H2 by a factor of 2 or so, a search for these SGR flares was feasible.

Such a search is underway right now and it is being led by Peter Kalmus of Caltech. In collaboration with Jacob Slutsky, we made a first order data quality cut on the Astrowatch data. One of the biggest priorities was analyzing the Up times for any usable segments. We trawled through the Hanford e-logs and the automated figures of merit and came up with first level cuts on the data, allowing us to identify significant periods of Up time data that were usable for science observations. For more details, please see [45].

After using the Up time segments that were deemed usable, the number of flares that could be analyzed went up from the 37 that occurred during science time to 68. This sort of data quality work was unique to Astrowatch, since normally there would not be any significant amount of Up time and its quality would not be under debate.

Chapter 3

Sources of continuous gravitational waves

3.1 Introduction

Neutron stars are remarkable objects which bring together all the four fundamental forces of gravity, weak-force, strong-force and electromagnetism. They are also at the heart of general relativity, as they strongly curve spacetime around them. Thus they are unique laboratories of physics. Understanding them can give insight into the confluence of a large portion of known physics, as well as charter the unknown. When dealing with neutron stars, one has to deal with large magnetic fields, nuclear densities (even supernuclear densities), relativistic velocities, extreme curvature of spacetime and particle physics.

Neutron stars are believed to be born in some of the largest explosions since the Big Bang, called supernovae, which occur when ultra massive stars undergo collapse because thermal pressure from nuclear fusion in the stellar core can no longer sustain the gravitational pull of the star on itself. These supernovae will produce a rapidly changing non-axisymmetric mass quadrupole moment and thus create gravitational waves. Neutron stars retain most of the angular momentum of the collapsing star and are extremely compact. Two neutron stars rotating around each other would slowly inspiral due to the emission of gravitational radiation, which could be detectable in the final moments before and during the eventual merger.

Not only are neutron stars one of the most promising sources of gravitational waves, but gravitational waves could be a significant probe into their nature. Gravitational waves can help constrain the equation of state (relation between pressure and energy density) of a neutron star, and help us understand their composition as well. This chapter summarizes some of the current knowledge on the emission of gravitational radiation from neutron stars. It also describes important properties of pulsars, which are neutron stars that are observed to produce pulses in the electromagnetic spectrum. This leads into the discussion of how searches for gravitational waves can be conducted in Chapter 4. The object described in Chapter 6 is a suspected neutron star.

3.2 Neutron Stars

A neutron star is a compact object of a radius around 10 km and a mass of about 1.4 solar mass. They are believed to be composed mostly of neutrons with central densities ranging from $0.8 - 1.6$ neutrons$\cdot$fm^{-3}, which is many times greater than the nuclear density found in ordinary atomic matter (0.16 neutrons$\cdot$fm^{-3}). In addition to neutrons, the core of a neutron star also contains degenerate Fermi gases of protons and electrons in beta equilibrium. The Fermi energies are sufficient to likely also support more exotic particle matter, including particles like strangeness-bearing baryons [46, 47], condensed mesons [49, 50] or deconfined quarks [51]. Fermions in such conditions of high density are expected to exhibit superfluidity and superconductivity as well. Observationally, some neutron stars are known to have very strong magnetic fields, which vary implicitly as a function of their age and composition and can range from 10^8 Gauss (milli-second pulsars) to 10^{15} Gauss (young pulsars/magnetars).

Some neutron stars emit electromagnetic radiation along their magnetic axis with an angular spread which ranges from a few degrees to as much as 30 degrees [47]. If this beam of radiation sweeps by the Earth as it rotates, it appears to pulse at the rotation frequency of the star and the object is known as a pulsar. Figure

3.1 illustrates this lighthouse like effect. Pulsars are known to have very stable rotation periods, although some of them are known to glitch, i.e., experience abrupt changes in the frequency or time derivative of the frequency of the observed pulses. The electromagnetic radiation frequency emitted (pulsed or otherwise) by some pulsars spans an enormous range of frequencies from radio to X-rays. The first pulsar was discovered by Hewish and Bell in 1967 with a period of 1.337 seconds [52], now known as PSR 1919+21 (PSR stands for pulsating source of radio). Several hundreds of pulsars have been discovered so far and almost all of them lie within our galaxy, with a few discovered in the relatively nearby Magellanic Clouds. Most of these pulsars were detected by radio observations from various observatories around the planet. Some pulsars exist in binary systems with low mass companions like white dwarf stars in the form of Low-mass X-ray binaries (LMXBs). A double pulsar system is also known to exist (PSR J0737-3039A and PSR J0737-3039B). Globular clusters are rich in old spun-up milli-second pulsars (see next section).

Neutron stars are one of the most promising sources of gravitational waves. There are several possible mechanisms for the emission of gravitational waves and some of these are discussed in this chapter. This chapter will also give a brief overview of pulsar detection, their properties and composition.

3.3 Pulsars

As mentioned above, pulsar is an observational term referring to the class of objects that are observed by the detection of an electromagnetic pulse at periodic intervals of great precision. Some of these pulsars have broadband emissions ranging from radio to X-rays, but most of them are only seen in radio. The pulsations are extremely stable and in some cases they are bound by the accuracy of terrestrial clocks. An example of this is the precision to which the period of PSR 1937+21 is known, 1.55780644887275±0.00000000000003 ms, as of 29 November 1982. That is an accuracy of 1 part in 10^{14}. Such precise timing coupled with their compactness,

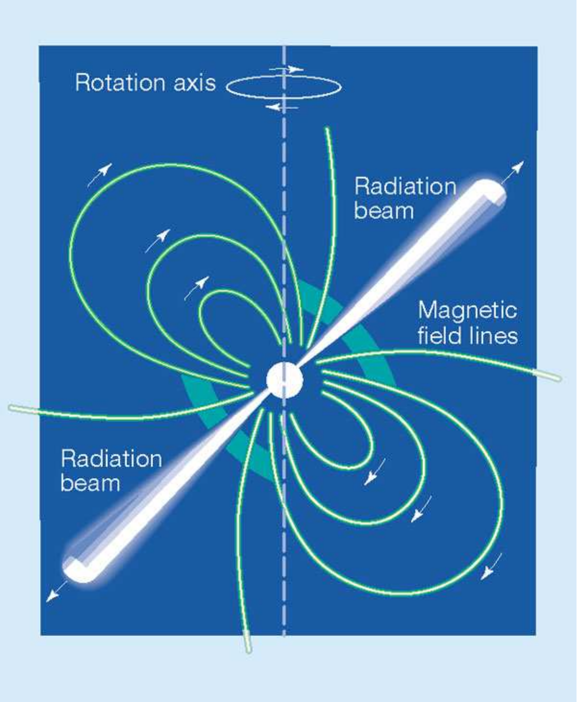

Figure 3.1: A cartoon of a pulsar. The lighthouse effect is shown. If the Earth is within the light emission cone, which is around the magnetic field lines, the neutron star will be seen as a pulsar. Figure courtesy [53].

gives us the opportunity to use these objects as tests of General Relativity.

There are some significant biases in the detection of pulsars. One of these is the beaming nature of each pulsar. Pulsars often have their magnetic axis at an angle to the rotation axis. The pulsar emits electromagnetic radiation in a tight cone of about 10 degrees around the magnetic axis [47]. Thus each pulsar in effect beams. If the Earth happens to fall within the path of this beam, the pulsar may be detected. Pulsar emissions are usually too weak to be detected extragalactically. Due to this effect, more pulsars are known close to us than far from us. Distance to pulsars are determined by various techniques like parallax measurements in the case of close pulsars, association with known supernova remnants, absorption spectral features due to intervening, but known hydrogen clouds, etc. Taking these biases into account, over 10^5 active pulsar-like objects are suspected to exist in our galaxy [55]. However, since active pulsars last only a few 10's of millions of years and not all neutron stars pulse, the number of neutron stars in the galaxy is expected to be orders of magnitude larger [47].

The detection of pulsars usually involves taking a fast Fourier transform (FFT) of radio observatory data. Any periodic signal that is sufficiently loud to rise above the noise will be detected as a pulsar. However only a few pulsars are strong enough to overcome the spread of their signal into various frequencies due to the Doppler shift caused by the motion of the Earth. To detect such pulsars, one of the techniques used is known as barycentric resampling. In this technique, a pulsar's location, frequency and other parameters are guessed and then a correction applied to the data to remove the effects of the Doppler shift of the Earth's motion. The data is then processed through a FFT algorithm. If there is any periodicity in the signal, it will show up in the FFT data, given that a sufficient amount of data is used to ensure that the signal will rise above the noise. In the case of pulsar detection, the signal is often strong enough to require only a few hours of data to attain a sufficient signal to noise ratio (SNR). This is not the case in gravitational wave signal data analysis, in which the signal is very weak and integration of days to years of data is required. Other techniques used to detect pulsars are time domain folding and heterodyning.

Figure 3.2 shows a scatter plot of known pulsars, plotted as a function of their periods and the first derivative of the period. Three distinct classes of pulsars can be seen. The first group with an average period of around 1 s and magnetic fields (Bs) ranging from 10^{11} to 10^{13} Gauss are young pulsars which are no more than a few million years old. The second group consists of pulsars which have a large magnetic field of greater than 10^{14} Gauss. These pulsars are known as Anomalous X-ray pulsars (AXPs) or Magnetars. The last set of pulsars are known as milli-second pulsars (MSP), which have weaker fields in the range of 10^8 to 10^9 Gauss and periods of a few milli-seconds. MSPs are are believed to be old "recycled" pulsars which have been spun up by accretion, thereby increasing their spin angular momentum. They are believed to have had a companion star from

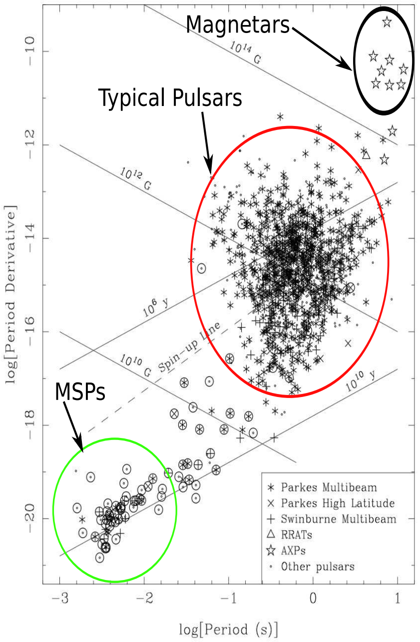

Figure 3.2: A scatter plot of the periods and first period derivatives of pulsars observed by various observatories (logarithmic scales on both axes). Magnetars, typical young pulsars and milli-second pulsars are shown. Figure courtesy [109]. Lines of constant age and constant magnetic field are shown (see equations 3.7 and 3.16).

which they accreted matter coherently and were spun up. The impact of the matter on to the star leads to a rise in temperature, from which X-rays are produced and emitted. Sometimes the companions of these MSPs go supernova themselves and are kicked out of orbit, leaving the MSP isolated. An isolated MSP can also occur when the NS evaporates its companion, for example the Black Widow system [48].

3.3.1 Pulsar emission and ages

Pulsars emit prodigious amounts of energy, most of which is drawn from their rotational energy. The Crab pulsar, which resides in the Crab nebula is a relatively young pulsar, which spins with a period of around 30 milli-seconds. It puts out about 5×10^{38} ergs/s of power, in wideband electromagnetic emission [56]. However the pulsar glow itself is a tiny fraction and most of the energy goes in lighting up the nebula and accelerating particles away from the star in the form of a pulsar wind. Since most of the energy output is powered by the rotational energy of the pulsar, we can equate the loss in energy to an effective spindown of the pulsar as

$$-\frac{dE}{dt} = -\frac{d}{dt}\left(\frac{1}{2}I\Omega^2\right) = -I\Omega\dot{\Omega} - \frac{1}{2}\Omega^2\dot{I} = \dot{E}, \tag{3.1}$$

where $\frac{dE}{dt}$ is the power emitted, I is the moment of inertia of the pulsar and Ω is the angular velocity. The moment of inertia of a pulsar does not change significantly with time, since its mass and radius are roughly constant. Thus we can ignore the $\dot{I}$ term in equation 3.1 and we can relate the energy loss to the spindown as

$$I = -\frac{\dot{E}}{\Omega\dot{\Omega}} \tag{3.2}$$

in the convention in which the spindown $\dot{\Omega}$ is negative and $\dot{E}$ is positive. If we assume that the energy loss is due to magnetic dipole emission [57], then the energy output of such a dipole is given by

$$-\frac{dE}{dt} = \frac{2}{3}R^6 B^2 \Omega^4 \sin^2\alpha \tag{3.3}$$

where R is the radius of the pulsar, B the magnetic field at the surface, and α the angle between the magnetic axis and the rotation axis. Equating equations 3.1 and 3.3, we get

$$B = \sqrt{\frac{-3I\dot{\Omega}}{2R^6 \sin^2\alpha\Omega^3}}. \tag{3.4}$$

This magnetic field is real, when $\dot{\Omega}$ is negative definite, which is the case if the energy radiated is produced by the spindown of a pulsar. Astronomers are more used to using the period P and $\dot{P}$, which are related to the angular velocity Ω by

$$P = \frac{2\pi}{\Omega} \tag{3.5}$$

and

$$\dot{P} = -\frac{2\pi\dot{\Omega}}{\Omega^2} \tag{3.6}$$

which turns equation 3.4 to

$$B = \sqrt{\frac{3IP\dot{P}}{8\pi^2 R^6 \sin^2\alpha}}. \tag{3.7}$$

For typical R's of around 10 km, I of 10^{38} kg·m^2 and periods corresponding to the Crab nebula, B is in the range of 10^{12} Gauss and the rotational energy is about 10^{49} ergs. These numbers give an indication that pulsars are compact objects with extremely large rotational energies and surface magnetic fields, which favors the hypothesis that they are neutron stars, which are theorized to have similar properties.

Another form of radiation that could carry away the rotational energy of a pulsar is gravitational radiation, which is emitted when a non-axisymmetric mass quadrupolar moment is generated by the pulsar. The loss of rotational energy given by the emitted GW flux is

$$L_{\mathrm{grav}} = \frac{1}{5}\langle \dddot{I}_{\mu\nu} \dddot{I}^{\mu\nu} \rangle, \tag{3.8}$$

where $\langle\rangle$ represents a time average and $I_{\mu\nu}$ is the moment of inertia tensor. The second time derivative of the moment of inertia tensor for a rotating neutron star as shown in Figure 3.3 is given by [64]

$$\ddot{I}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -16\pi^2 f^2 (I_{\mathrm{xx}} - I_{\mathrm{yy}}) \cos(4\pi f t) & -32\pi^2 f^2 (I_{\mathrm{xx}} - I_{\mathrm{yy}}) \sin(4\pi f t) & 0 \\ 0 & -32\pi^2 f^2 (I_{\mathrm{xx}} - I_{\mathrm{yy}}) \sin(4\pi f t) & 16\pi^2 f^2 (I_{\mathrm{xx}} - I_{\mathrm{yy}}) \cos(4\pi f t) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{3.9}$$

Using equations 3.8 and 3.9, the GW flux can be written as

$$L_{\mathrm{grav}} = \frac{dE}{dt} = \frac{32}{5}\Omega^6 I_{\mathrm{zz}}^2 \left(\frac{I_{\mathrm{xx}} - I_{\mathrm{yy}}}{I_{\mathrm{zz}}}\right)^2. \tag{3.10}$$

The various mechanisms which can generate this kind of asymmetry are discussed in the following sections.

If the observed spindown of a pulsar were purely due to magnetic dipole radiation ($\dot{\Omega}_{\mathrm{mag}}$), then using equations 3.2 and 3.3 we find

$$\dot{\Omega}_{\mathrm{mag}} \propto \Omega^3. \tag{3.11}$$

And if the emission were purely due a gravitational quadrupole ($\dot{\Omega}_{\mathrm{grav}}$), then using equations 3.2 and 3.10 we find

$$\dot{\Omega}_{\mathrm{grav}} \propto \Omega^5. \tag{3.12}$$

For multipole radiation in general, the spindown will follow a power law:

$$\dot{\Omega} = -K\Omega^n, \tag{3.13}$$

where K contains geometric terms like the moment of inertia I and source terms like the magnetic field (B) and the ellipticity. n is called the braking index of the pulsar and $n = 2l + 1$, and $l = 1$ for a dipole and $l = 2$ for a quadrupole. The braking index of a pulsar is indicative of its mechanism of energy loss. If the pulsar were losing all of its energy due to magnetic dipolar radiation, the value of n would be 3. The value of n would be 5 if the energy loss were due to gravitational

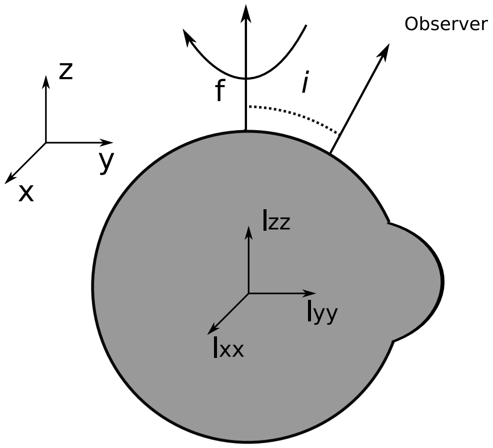

Figure 3.3: An illustration of a rotating neutron star with a non-axisymmetric bump on it, emitting gravitational waves. The neutron star is rotating around the z axis and the observer is at an inclination angle ι off the z axis.

radiation. Differentiating both sides of equation 3.13, we get

$$\ddot{\Omega} = -Kn\Omega^{n-1}\dot{\Omega}. \tag{3.14}$$

Multiplying both sides with Ω and using equation 3.13, we get

$$n = \frac{\Omega\ddot{\Omega}}{\dot{\Omega}^2}. \tag{3.15}$$

For some pulsars $\ddot{\Omega}$ is known well enough to allow for a direct calculation of n. For example, the Crab pulsar has $n \simeq 2.5$ and Vela pulsar (PSR J0835-4510) has $n \simeq 1.7$. These are less than the 3 expected for magnetic dipole emission, suggesting that this may not be the only mechanism for emission. Indeed, both pulsars are recent supernova remnants (SNRs) with strong winds powering the observed nebula. Moreover, equation 3.13 assumes that the moment of inertia I is constant. For $\dot{I}$ to make a significant contribution, it needs to change on a scale comparable to the spindown timescale: $\frac{I}{\dot{I}} \approx \frac{P}{\dot{P}}$. This is only likely to happen if the star is nearly at the breakup spin frequency (Kepler frequency), which is only the case for a few MSPs. Pulsars have mostly been observed to have $n \leq 3$. For gravitational wave searches in which the frequency evolution of a suspected neutron star is unknown, it is prudent to be conservative while setting up the search parameter space and thus values of the braking index n ranging from 2 to 7 are used.

After shuffling terms and integrating both sides of equation 3.13, we get an expression for the "spindown age" of the pulsar,

$$t = \frac{1}{(n-1)K}\left(\frac{1}{\Omega^{n-1}} - \frac{1}{\Omega_0^{n-1}}\right). \tag{3.16}$$

Here we have assumed that $t = 0$ at birth and Ω_0 is the initial angular velocity. Substituting $K = -\frac{\dot{\Omega}}{\Omega^n}$ we get

$$t = -\frac{1}{n-1}\frac{\Omega}{\dot{\Omega}}\left(1 - \left(\frac{\Omega}{\Omega_0}\right)^{n-1}\right). \tag{3.17}$$

After a long time, when $\Omega << \Omega_0$, we can ignore the $\frac{\Omega}{\Omega_0}$ term and a characteristic age can be defined for a pulsar as

$$\tau_c = -\frac{1}{n-1}\frac{\Omega}{\dot{\Omega}} = \frac{1}{n-1}\frac{P}{\dot{P}}. \tag{3.18}$$

The characteristic age of a pulsar calculated in such a manner for the Crab pulsar (1200 years) compares well with its historical age of about 1100 years. These age estimates are useful for setting up parameter spaces for gravitational wave searches, since the expected gravitational wave parameters (like spindown) are usually a function of the expected age of the source.

3.4 Composition of neutron stars

Conventional neutron stars are believed to contain a liquid nuclear-matter core, enclosed in a thin solid crust. The study of the composition of neutron stars was motivated by the glitchiness of pulsars. Many pulsars such as the Vela pulsar are known to undergo a rapid change in frequency and frequency derivatives in a matter of minutes [65]. Some theories suggest that these glitches are associated with starquakes [66]. These starquakes could occur in the solid crust of a conventional neutron star or in the solid interiors of more exotic compact stars like quark stars or hybrid stars. Other theories involve some sort of superfluid vortex pinning that slips through the crust on occasion.

A simple model of a neutron star is that of a static, spherically symmetric, relativistic star. Solutions of Einstein's equations under these conditions are called the Oppenheimer-Volkoff (OV) equations. The only quantities that show up in these equations are the mass (M), energy density (ρ) and pressure (p), all of which are functions of the radial coordinate r. The relationship between the pressure and energy density describes the state of the matter forming the fluid star. These relationships are called the equations of state (EOS) and are written as either

$$p = p(\rho) \text{ or } \rho = \rho(p). \tag{3.19}$$

The total gravitational mass M (which includes the baryon mass and the gravitational binding energy and all other forms of energy in it) is the first OV equation as shown in [47]

$$M(r) = 4\pi \int_0^r \rho(r')r'^2 dr', \tag{3.20}$$

where r is the radial coordinate. $M(R)$ is the total gravitational mass of the star, if R is the radius of the star, i.e. the point where the pressure $p(R)$ vanishes. This is because 0 pressure cannot support any matter against the attractive gravitational force and thus R would be the effective radius of the star. The second equation relates the pressure to the radial coordinate r as shown in [47]

$$\frac{dp}{dr} = -\frac{[p(r)+\rho(r)][M(r)+4\pi r^3 p(r)]}{r[r-2M(r)]}. \tag{3.21}$$

The OV equations are derived from Einstein's equations applied to the interior of a spherical, static and relativistic star. All of these equations are in units of $G = c = 1$, where G is the gravitational constant and c is the speed of light.

The OV equations can be rewritten in a form, such that the relativistic corrections are separated from the Newtonian solution

$$dM(r) = 4\pi r^2 \rho(r) dr \tag{3.22}$$

and

$$4\pi r^2 dp(r) = -\frac{M(r)dM(r)}{r^2} \times \left(1+\frac{p(r)}{\rho(r)}\right) \times \left(1+\frac{4\pi r^3 p(r)}{M(r)}\right) \times \left(1-\frac{2M(r)}{r}\right)^{-1}. \tag{3.23}$$

Equation 3.23 can be interpreted as balancing the force acting on an infinitesimal shell of matter due to the pressure and the gravitational attraction. The left-hand side of equation 3.23 if the force acting outward on the shell due to the pressure. The first term of the right-hand side is the Newtonian gravitational pull and the other terms are the corrections due to general relativity. Solving equations 3.22 and 3.23 requires an equation of state. One can solve for quantities like the radius of the star and the pressure and energy densities etc. by integrating both sides of equations 3.22 and 3.23.

According to current models, a typical neutron star would consist of a thin "atmosphere" and an "envelope" consisting of light elements such as Hydrogen or Helium or Carbon, whose dynamic is fully controlled by the magnetic field. Below the envelope lies a thin solid "crust" which is composed of ordinary atomic nuclei crushed into a lattice with a sea of electrons. These atomic nuclei are most likely Iron nuclei, since it is the most stable nucleus due to its high binding energy per nucleon [69]. Going deeper into the star, free neutrons leak out of the nuclei and form a superfluid. This part is known as the core of the neutron star and the exact properties (the EOS) of this superfluid is not known. Near the very center of the star could lie a transition between the neutron superfluid and other more exotic matter like strange quarks or pions and kaons in addition to neutrons (hybrid stars) [47]. The neutron star structure is illustrated in figure 3.4 [70].

3.5 Mechanisms of Emission

Neutron stars can produce gravitational waves through several possible mechanisms. A pair of neutron stars can inspiral into one another and produce an inspiral-merger-ringdown signature, which would be detectable by LIGO for a few minutes while the signal lies within its band. Neutron stars can also produce bursts of gravitational radiation associated with the starquakes that are responsible for their frequency glitches or through some unknown mechanism. However the emission type that is of interest here is the continuous emission of gravitational radiation by a non-zero mass quadrupolar moment, powered by its rotational energy. Some of the mechanisms of emission include non-zero ellipticity caused by "mountains" on the pulsar surface (most of the quadrupole moment comes from the corresponding depression of the surface), fluid modes like the r-modes and free precession due to the mis-alignment of the moment of inertia and rotation axes. Mass accretion from a companion star can drive any of the emission mechanisms mentioned above.

3.5.1 Non-axisymmetric mass tensor

A promising mechanism for emission of continuous gravitational waves from a neutron star is the existence of a non-axisymmetric distortion of its solid crust. These non-axisymmetric distortions are akin to mountains on the Earth's surface and their effects can usually be represented by a single unitless number called the ellipticity, defined as

$$\epsilon = \frac{I_{\mathrm{xx}} - I_{\mathrm{yy}}}{I_{\mathrm{zz}}} \tag{3.24}$$

where I_{xx}, I_{yy} and I_{zz} are the moments of inertia along the x, y and z axes respectively as seen in Figure 3.3. They form the diagonal elements of the moment of inertia tensor if the object is spinning along the z axis. This non-axisymmetric ellipticity produces a non-zero mass quadrupole moment, which accelerates due to the spin. This in turn produces gravitational radiation, with loss of rotational energy given by the emitted GW flux:

$$L_{\mathrm{grav}} = \frac{dE}{dt} = \frac{32}{5}\Omega^6 I^2 \epsilon^2 \tag{3.25}$$

where, I is the moment of inertia along the spinning axis, i.e. $I = I_{\mathrm{zz}}$. The neutron star emits gravitational radiation if the mountain produces a net mass quadrupolar moment and this emission is at twice the rotation frequency of the neutron star. The frequency is twice the rotation frequency because the ellipticity manifests itself in such a way that the object looks symmetric under a 180 degree rotation. The amplitude of the gravitational wave received by a detector is a function of the inclination angle ι that the source makes with the detector.

The gravitational strain $h_{\mu\nu}$ emitted by such a distortion can be shown to be [71]

$$h_{\mu\nu} = \frac{2}{d}\ddot{I}_{\mu\nu}, \tag{3.26}$$

where d is the distance to the source. Using equations 3.9, 3.26 and 3.24, we get

$$h_{\mu\nu} = \frac{32\pi^2 f^2 \epsilon I}{d} \times \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\cos(4\pi f t)(1+\cos^2(\iota)) & -2\sin(4\pi f t)\cos(\iota) & 0 \\ 0 & -2\sin(4\pi f t)\cos(\iota) & \cos(4\pi f t)(1+\cos^2(\iota)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{3.27}$$

The gravitational strain that would be received by a detector if the source were optimally oriented and emitting at a frequency f, would be

$$h_0 = \frac{4\pi^2 G}{c^4}\frac{I f^2}{d}\epsilon. \tag{3.28}$$

Uncertainties in the moment of inertia I and the distance d lead to uncertainties in h_0 as large as factors of 2 or 3.

The maximum ellipticity that can be maintained on a neutron star is a function of its composition. A completely liquid star would be unable to sustain any ϵ due to elasticity, since its shear modulus would be 0. However even a completely liquid neutron star could maintain an effective magnetic ellipticity when a neutron star with a large toroidal magnetic field is distorted into a prolate shape [67]. This form of magnetic ellipticity is discussed later in this section. A nice summary of the possible values of ϵ that can be sustained by various compositions of neutron stars is given in [77]. I summarize the results here.

The ellipticity ϵ can be defined in terms of the corresponding mass quadrupole moment Q_{22} and the moment of inertia I as

$$\epsilon = \sqrt{\frac{8\pi}{15}\frac{Q_{22}}{I}}. \tag{3.29}$$

As [77] notes, the moment of inertia of a neutron star was estimated by [68] as

$$I = 9.2 \times 10^{37}\mathrm{kg}\cdot\mathrm{m}^2 \left(\frac{M}{1.4M_\odot}\right)\left(\frac{R}{10\mathrm{km}}\right)^2 \times \left(1 + 0.7\left(\frac{M}{1.4M_\odot}\right)\left(\frac{10\mathrm{km}}{R}\right)\right), \tag{3.30}$$

which is accurate to a few percent for various equations of state. The maximum quadrupole moment is also a function of the maximum breaking strain σ_{max} of the crust. For a conventional neutron star with a solid crust and a liquid interior it is calculated by [77] as

$$Q_{22,\mathrm{max}} = 2.4 \times 10^{31}\mathrm{kg}\cdot\mathrm{m}^2 \left(\frac{\sigma_{\mathrm{max}}}{10^{-2}}\right)\left(\frac{R}{10\mathrm{km}}\right)^{6.26}\left(\frac{1.4M_\odot}{M}\right)^{1.2}. \tag{3.31}$$

The maximum ellipticity of a conventional neutron star with a solid crust and a liquid interior can then be calculated using equations 3.29, 3.30 and 3.31, as shown in equation 4 of [77] to give

$$\epsilon_{\max} = 3.4\times10^{-7}\left(\frac{\sigma_{\max}}{10^{-2}}\right)\left(\frac{1.4M_\odot}{M}\right)^{2.2}\left(\frac{R}{10\mathrm{km}}\right)^{4.26}\times\left(1+0.7\left(\frac{M}{1.4M_\odot}\right)\left(\frac{10\mathrm{km}}{R}\right)\right)^{-1}. \tag{3.32}$$

The fiducial value of $\sigma_{\max} = 10^{-2}$ is obtained from the maximum strains seen in iron (or other terrestrial alloys). The other quantities in the equation are not subject to much change and thus the maximum breaking strain that the material can withstand determines the ellipticity that can be sustained. Thus for a neutron star with conventionally expected solid iron crust, the maximum ϵ would be of the order of 10^{-7}. However recent results [73] show that the breaking strain for pure single crystals to be as high as 0.1, which would push the maximum ϵ to 10^{-6}.

Neutron stars containing a substantial component of strange quarks in the solid state were first proposed in [78]. For these solid strange stars, the shear modulus could be 3 orders of magnitude higher than that for a typical solid iron crust neutron star with a liquid interior. This leads to a maximum ellipticity of the order of 10^{-4}.

At neutron star densities, pure neutron matter is not in its lowest energy state and some protons and electrons are present such that the chemical potentials of these particles are given by

$$\mu_n = \mu_p + \mu_e. \tag{3.33}$$

This is the equation for Beta-equilibrium, where μ_n, μ_p and μ_e are the chemical potentials of neutrons, protons and electrons. In general charge-neutral matter consists of a majority population of neutrons with a few protons and electrons present in Beta-equilibrium. However this high neutron-to-proton ratio results in a high isospin asymmetry. It is not energetically favorable to produce protons because of the necessity to produce associated electrons. It is possible for these stars to produce negative quark matter which can lead to the production of positive baryonic matter, restoring some of the isospin asymmetry. This transition can happen in phases and mixed phases and hybrid stars are also possible. References [77] and [47] talk about how this phase transition can begin with the formation of droplets of quark matter, which eventually merge to form rods and then slabs and then the reverse happens to baryonic matter which then becomes the minority. Reference [46] gives an estimate of the shear modulus of one such hybrid star consisting of a body centered cubic lattice of quark matter within a bulk of baryonic matter as a function of the charge density, the diameter of the quark matter droplets and the spacing between these droplets. The maximum ellipticity is then computed in [77] to be of the order of 5×10^{-6}.

As mentioned above, another mechanism for the production of an ellipticity, which would emit gravitational waves is via a "frozen in" magnetic field. According to [67], such magnetic distortions tend to dominate over the natural oblateness of the neutron star of the magnetic field

$$B_t > 3.4 \times 10^{12} \text{ Gauss} \times \left(\frac{\nu_s}{300\text{Hz}}\right), \tag{3.34}$$

if ν_s is the frequency of rotation of the neutron star. Such a neutron star is unstable with respect to the wobble angle between its angular momentum axis and the magnetic axis. The wobble angle evolves on a dissipation timescale such that the angular momentum axis and the magnetic axis are eventually orthogonal, making it optimal for GW emission. Reference [67] discusses the possibilities of such a mechanism taking effect for various types of neutron stars, like young neutron stars, accreting neutron stars and milli-second pulsars. For milli-second pulsars that are about 10 million years old, which is an optimistic estimate for Calvera (see Chapter 6), [67] a maximum magnetic ellipticity ϵ_B of around 6.0×10^{-6} is possible.

Thus the maximum ellipticities that can be sustained on neutron stars are highly uncertain and can range from as high as 10^{-4} to as low as 10^{-7}.

3.5.2 Non-axisymmetric oscillations

Non-axisymmetric oscillations of neutron stars are a promising source of gravitational radiation. The gravitational radiation from these oscillations can be steady if there is some kind of dynamical or secular instability to stabilize the oscillation. The scenarios in which such oscillations are most likely to occur are in the case of recently formed neutron stars and accreting neutron stars. For a review of these oscillations, please see [58].

While a rotating neutron star may possess a whole range of pulsation modes, the ones that are of interest for GW emission include acoustic modes called f-modes and coriolis restored modes called r-modes. These r-modes are of particular interest, as they have been theorized to be long lasting and of high enough amplitude to be detectable by LIGO [137].

r-modes belong to the family of modes driven to instability by the Chandrasekhar-Friedman-Schutz (CFS) instabilities [59, 60], which are similar to the Kelvin-Helmholtz(KH) instability. The KH instability is caused when a velocity shear is present within a continuous fluid. An example is the giant eye of Jupiter and some cloud formations on the Earth. Imagine a neutron star rotating at an angular velocity of Ω. If in the frame of a fluid element rotating with the star (i.e. in the co-rotating frame), an r-mode instability were moving at $\frac{-\Omega}{3}$, then to a non-rotating frame of reference, it will be travelling at $\frac{2\Omega}{3}$ in the direction of the rotating neutron star. The gravitational waves emitted by this star would then carry away positive angular momentum and impart a negative angular momentum back-reaction to the star. But this reverse momentum, when seen in the co-rotating frame, helps the instability to grow. This sort of positive feedback continues till viscosity takes over, under the most interesting circumstances, only indirectly via nonlinear hydrodynamical mode coupling which may result in the generationa of a stable amplitude of gravitational waves.

While the generation of GWs through an ellipticity mechanism as mentioned in the section above is due to a mass quadrupole moment with $l = m = 2$, GWs from r-modes are generated through a $l = m = 2$ current quadrupole [61]. This affects the parameter estimation and the effectiveness of the detection statistics used for conducting a search for r-mode emission. As shown in [62] the traditional data analysis approaches are effective for r-modes with relatively minute corrections needed. An r-mode is an Eulerian velocity perturbation and using the approximations of Newtonian gravity and slow-rotation of the neutron star it can be shown [62]

$$\delta v_j = \alpha \Omega R \left(\frac{r}{R}\right)^l Y_j^{B,l,l} e^{i\omega t}, \tag{3.35}$$

where, r is the radial coordinate, R is the radius of the neutron star, Ω is the angular velocity of the star, ω is the angular frequency of the GW and α is a dimensionless amplitude and is an equivalent to h_0 for a traditional mass quadrupole for r-modes. $Y_j^{B,l,l}$ is a magnetic-parity vector spherical harmonic, which are functions of the scalar spherical harmonics $Y^{l,m}$. α can be related to the intrinsic strain amplitude h_0 by [62]

$$h_0 = \sqrt{\frac{8\pi}{5}} d^{-1} \omega^3 \alpha M R^3 \tilde{J}, \tag{3.36}$$

where d is the distance to the source, M is the mass of the neutron star and $\tilde{J}$ is a dimensionless functional of the neutron star equation of state. As is noted in [62], the EOS-related uncertainties, which would determine the uncertainties in equation 3.36 are dominated by uncertainties in M and R, which are usually in the range of a factor of $2-3$. The gravitational wave frequency is obtained in terms of the angular frequency of the wave, which is related to the angular velocity of the star as [62]

$$f = \frac{\omega}{2\pi} \tag{3.37}$$

and

$$\omega = -\frac{(l+2)(l-1)}{(l+1)}\Omega. \tag{3.38}$$

For the $l = 2$ as mentioned above, $\omega = -\frac{4}{3}\Omega$ for gravitational wave emission. Searches done for GWs emitted by r-modes from known pulsars would have have to search $\frac{4}{3}^{\text{rd}}$ the observed electromagnetic frequency.

While it is believed that r-modes are most likely to play a role in young and accreting neutron stars, there is at least one reference in the literature [63] which mentions r-modes as a possible spindown mechanism for milli-second pulsars. Reference [63] shows low-amplitude r-modes along with the emission of thermal X-rays, would allow for a spin down due to a combination of gravitational radiation and standard magnetic torque. This result is of particular interest to the search done for GWs from Calvera in Chapters 6 and 7.

3.5.3 Torque-free Precession

A simple mechanism of emission of gravitational radiation from a neutron star is torque-free precession. If the principal moment of inertia axes were not aligned with the angular momentum axes, then the neutron star would wobble like a top and a GW back reaction would damp out these wobbles [64]. This GW emission would occur at the rotation frequency. However, it would be much weaker than the other forms of emission discussed in this chapter. As shown in [72], even for optimistic values of parameters like the angle of wobble (between the principal axis and the angular momentum axis) of 0.1 radians, for a neutron star at a distance of 1 kpc and spinning at 500 Hz, the expected gravitational wave strain amplitude would be of the order of 10^{-27}.

While the calculations in [64] were done for a rigid body, this is not strictly true for neutron stars, since they are theorized to have liquid interiors. A more realistic calculation was done in [74], which also showed that this sort of emission can be relatively long lived and last up to 10^5 years. However the amplitude is still too small to be detectable by initial LIGO. Moreover this mechanism of emission would be improbable for a MSP type neutron star, since any wobble would have damped out a long time ago, given its age.

3.5.4 Accretion

Some pulsars are known to have companions from which they accrete matter. These companions can range from supergiants, which overflow their so called Roche lobe, to white dwarves. Accretion is the process via which old pulsars are "recycled" and spun-up to become milli-second pulsars. Accretion leads to the formation of an accretion disk around the neutron star, which emits X-rays.

In the context of GW generation, accretion is a natural way of generating and maintaining crustal deformations. Accretion is not an isotropic process and there are many “hot spots” on the neutron star surface at which the accreted material falls. These hot spots can lead to the formation of crustal asymmetries, which would be damped by GW emission. This can lead to a situation in which the accretion spin-up is balanced by the GW spin-down, which could be in equilibrium, leading to a continuous GW signal. This was first suggested in [75]. Rapidly accreting stars are of particular interest for this sort of emission mechanism, for example ones in low-mass X-ray binaries. One of the most promising of such sources of GW is Sco-X1, which as the brightest X-ray source in the sky (besides the Sun) and according to [75], it could emit GWs that would be detectable with a gravitational strain amplitude of about 10^{-26}. A search for GWs from Sco-X1 using data from the second science run of LIGO can be found in reference [95].

According to [76], the magnetic field of a neutron star can help localize the accreting matter to a small area of the surface. The magnetic field lines can be compressed into a narrow belt at the magnetic equator, which then confines the accreted material on to the poles of the neutron star. Since the magnetic axis is not generally aligned to the angular momentum axis, this accumulation of material takes place asymmetrically, thus creating a mass quadrupole and consequently GWs are emitted. The dissipation of these mountains is slow because the matter is highly conductive and thus it crosses the field lines slowly. This leads to a scenario in which the pile up is matched by steady GW emission.

The next chapter is a discussion of LIGO searches for CWs from NSs emitted via the mechanisms discussed in this chapter.

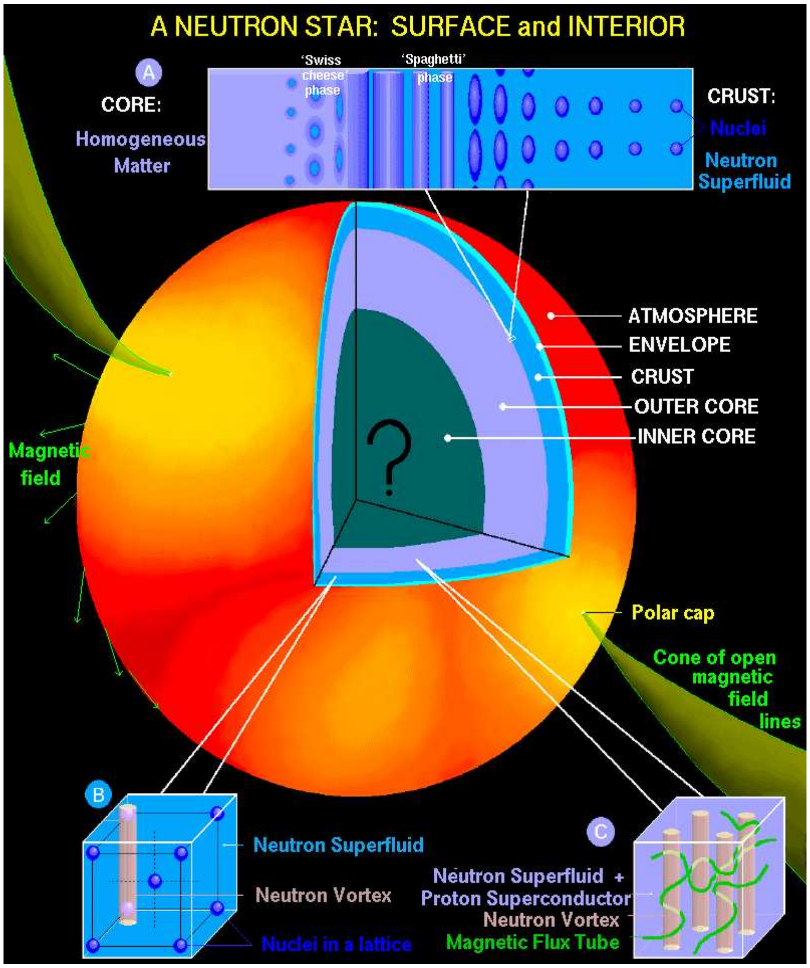

Figure 3.4: The composition of neutron stars. A neutron star consists of an atmosphere and an envelope of light elements. A core consisting of a neutron superfluid forms the bulk of the neutron star. The inner core could consist of exotic matter like strange quarks or pions and kaons in addition to neutrons (hybrid stars). Figure courtesy Danny Page [70].

Chapter 4

Finding weak quasi-periodic signals in noisy broadband data

4.1 Introduction

Continuous GW emitted by neutron stars are expected to be exceedingly weak relative to the noise spectral density(S_h) of ground based detectors like LIGO. While it is possible to detect loud sources of GWs like neutron star - neutron star inspiral-merger-ringdowns using a few seconds of data, one has to dig CW signals out of the noise by integrating several days to years of data. An optimal method for detecting signals buried deep within noise is *matched filtering*. Matched filtering is also called coherent integration in the field of CW data analysis. For a search done using the optimal matched filtering technique using an integration time of T, the signal scales as T and the noise scales as $\sqrt{T}$, and thus the sensitivity or SNR of the search scales as $\sqrt{T}$.

However for matched filtering, one has to compare the data to an assumed form of the signal, known as a template. The significance of each template must then be calculated. For a small number of templates, when certain search parameters are well constrained, a fully coherent integration is the preferred search method. However, the number of templates for a search over unknown parameters can increase dramatically. This increases the chances that a random fluctuation in the noise will seem like a real signal and thus we have to raise the threshold of detection.

This leads to a loss of sensitivity. Moreover these wide-parameter searches can get computationally prohibitive and thus less efficient, but computationally more feasible methods are more practical. These are known as semi-coherent methods, whose sensitivity scales as $\sqrt[4]{T}$. This chapter summarizes both the coherent and semi-coherent methods and their applications.

4.2 Summary of searches for continuous gravitational waves with LIGO data

4.2.1 Known pulsar searches

There are over 1500 known pulsars [109], among which an order of 200 have spin frequencies above 25 Hz, which would be detectable by LIGO. And about half of them are isolated (i.e., not part of a binary). From radio observations of these pulsars, their sky locations are known to within a few arc-seconds and their frequency to within a few mHz. For these pulsars, it is possible to confine a GW search to a few templates and in some cases to even a single template. This removes the computational constraint and one is able to integrate all the available data using a matched filtering algorithm and achieve the best possible sensitivity.

For sources for which the sky location and frequency information is accurate enough to require only a single search template, the integration is done in the time domain and the parameters of the source are calculated as probabilities within a Bayesian framework. It is possible to even incorporate ephemeris data from radio instruments to correct for frequency shifts due to pulsar glitches. These corrections assume that the gravitational wave frequencies follow twice the radio pulse frequency. Some of the searches done by the LIGO Scientific Collaboration using the time domain technique are in [79], [80] and [81].

The Crab pulsar is a young pulsar (1000 years) with a rotation frequency of around 30 Hz and a relatively large frequency spindown, making it one of the most important potential sources of continuous gravitational waves. Searches for GWs from the Crab pulsar were done in [80] and [81] using LIGO data. In [80], two different methods are used to search for the Crab pulsar. The first one uses the time domain technique within a Bayesian framework as discussed above, for a single template search. The second method uses a frequency domain "wide" parameter search, allowing for a slight discrepancy between gravitational wave signal frequency and twice the radio pulse frequency. This frequency domain search is done within a frequentist framework. While these wide parameter searches only encompass a small frequency band around twice the radio pulse frequency, the number of searched templates get prohibitively large and thus are benefitted by any speed-up of the frequency domain searches. Barycentric resampling helps increase the computational efficiency, and thus for fixed computational resources, the sensitivity of these kinds of searches. An ongoing effort in the LIGO scientific collaboration is to conduct wide parameter searches on all the known isolated pulsars including the ones in [81], as well as some other pulsars that could not be searched for with a single template. This search would not have been possible to conduct given limited computational resources, without the use of the resampling algorithm.

4.2.2 All-sky searches

It is estimated that there are about 10^8 neutron stars in the galaxy [47]. A large fraction of them are not observable as pulsars as they either do not emit electromagnetic radiation or their beams do not sweep across the Earth. These neutron stars are very difficult to detect electromagnetically, but they could be detectable in GWs, which are not strongly beamed. The difficulty in searching for sources of this kind is the lack of any information about them, including their sky position, frequency or frequency derivatives. For an observation over a time T, the resolution in the sky position improves like T^2. Consequently, the number of sky position templates needed to recover a signal with minimal (fixed) loss of SNR due to mismatch between the signal and template parameters scales like T^2. The number of frequency templates scales as T and the first frequency derivative as T^2 and so on. Thus even for a simple search involving just the first frequency derivative, the number of templates to search over scales as T^5. Since the sensitivity scales as $\sqrt{T}$, at best it scales with computational power required C as $C^{\frac{1}{10}}$. Thus limited computational power severely limits the sensitivity of these kinds of searches. A discussion on the relationship between the search parameters, number of templates and computational power required as a function of integration time can be found in Chapters 6 and 7.

Given the prohibitive scaling of the number of templates required for a coherent search, semi-coherent methods such as the PowerFlux method [82, 83], the Hough transform [84] and StackSlide have been developed by the LSC for all-sky searches. While these methods are less sensitive than the coherent methods, they are much more efficient and for all-sky searches, give the best scientific result for a given amount of computational time.

The optimal method [87] for all-sky CW searches is one that combines the sensitivity of a coherent search with the efficiency of a semi-coherent step. It involves splitting the time of integration into sub units which are coherently integrated and then combined semi-coherently. The basic idea behind these kinds of searches is to conduct a coherent wide parameter search over a smaller integration time, allowing for a coarser parameter spacing and consequently requiring a smaller number of templates. This is then followed by a semi-coherent combination of a series of these coherent integrations. The thresholds are set low for the coherent step, to allow a weak signal to rise above the noise, followed by a higher threshold for the semi-coherent integration. Several schemes of implementing these kinds of *hierarchical* methods have been studied in [85],[86] and [87]. Hierarchical methods have been used in LSC analyses for all-sky CW sources, as described in [88, 89]. These searches were implemented on a public distributed computing project called Einstein@Home [91], based on a distributed computing platform known as Berkeley Open Infrastructure for Network Computing (BOINC) [90]. BOINC has been previously used on other projects like SETI@Home. Einstein@Home has attracted more than 100,000 participants, delivering more than 80 Teraflops of computing power.

The sensitivity of hierarchical methods lies somewhere in between $\sqrt{T}$ and $\sqrt[4]{T}$. The longer the coherent step, the more sensitive the search gets. The resampling algorithm would make the coherent step significantly faster (by about a factor of 10, see Chapter 5) and thus for given computational resources allow for an even larger coherent step. This would increase the sensitivity of Einstein@Home searches for the same amount of computational resources by almost 25%. Implementation of the resampling algorithm in Einstein@Home is in progress.

4.2.3 Searches for suspected neutron stars

Some potential CW sources lie in between the known pulsars with well known parameters and the completely unknown neutron stars, whose locations and frequencies are completely unknown. These sources include objects which are suspected to be neutron stars, which could be pulsating at a rate, which could be detectable by LIGO. Some examples include the central compact object CasA in a supernova remnant in the constellation Cassiopeia, and Calvera, a compact X-ray source at high galactic latitude, which is suspected of being a close and rapidly-spinning neutron star, like a milli-second pulsar (see Chapter 6). For all these sources, their sky locations are known to within a few arc seconds thus requiring only a single sky location template for each object.

For each of these objects, the frequency and frequency derivatives are unknown. A search for gravitational waves from these objects would have to encompass all the frequencies and frequency derivatives, at which, given the source distance and apparent age, the source could be detectable. For a given tolerance of the maximum loss of SNR, the resolution required for the search would be determined by the integration time of the search. A description of how the template space was selected for the LSC's CasA search can be found in [136]. See Chapter 6 for details on how the parameter space was selected for the Calvera search.

The resampling algorithm is very efficient for these kinds of searches, since it can resample the data for a single sky location, and then calculate multiple frequency and frequency derivative templates very cheaply. The planned CasA search [136] has a coherent integration time of about 12 days. The search, which is currently ongoing uses another method to calculate the so called $\mathcal{F}$-statistic (see Section 5.1). The resampling algorithm can be used to calculate the $\mathcal{F}$-statistic much more efficiently. The speed up factor is about a factor of 50. Since the search involves frequency derivatives up to to the 2^{nd} order, the maximum coherent integration time possible scales as T^6. Thus for the same amount of computational resources, a coherent integration time of about 23 days becomes possible. This is a sensitivity increase of about 40% for the same amount of computational resources.

4.2.4 Binary systems

Neutrons stars in binary systems like LMXB systems and other accreting systems (accreting X-ray pulsars) are another promising source of gravitational waves (see Chapter 3). For these sources, the computational problem is even more acute than for searches for isolated neutron stars. This is because there are additional parameters that one has to search over. These include binary doppler parameters like the orbital period of the neutron star around its companion, the semi-major axis of this orbit, its orbital phase etc. These parameters are sometimes partly determined by electromagnetic observations, but even in those cases, the resolution required to search over these additional parameters scales as a power law of T. A few algorithms are being developed to do a short coherent integration and then incoherently combine the frequencies over which the signal would have been spread due to the Doppler effect of the binary system's motion. These include the TwoSpect method [92], Quadratic search [93] and the Sideband search [94].

One of the more promising candidates among the LMXBs is Sco X-1, which is a neutron star orbiting a low-mass companion with a period of 18.9 hours. While the sky location of Sco X-1 and its orbital period are well determined from electromagnetic observations (X-ray mostly), the semi-major axis and orbital phase have large uncertainties. The rotation frequency and any associated frequency derivatives are also undetermined. A search for GWs from this system was done using the third LIGO science run(S3) data [95]. A search for Sco X-1 is being developed which will use the Sideband search for known binary systems.

4.2.5 Searches directed at a region of the sky

Some regions of the sky are more likely to contain neutron stars than others. About 10^6 neutron stars are estimated to be present in the galactic center of the Milky Way [96]. A large population of old milli-second neutron stars exist in globular clusters like 47 Tuc [97]. Since most of the neutron stars are unknown, a wide parameter search over a large frequency band, associated frequency derivatives and multiple sky locations is needed. The number of sky locations required even for a square-degree in the sky is prohibitive and thus targeting some of the most compact globular clusters or the innermost core of the galactic center is a more practical strategy. Efforts are ongoing in the LSC to study the feasibility of these searches. They will ultimately use the barycentric resampling code described here in conjunction with semi-coherent methods to accomplish these searches. The barycentric resampling code will help increase the sensitivity of these searches by increasing the maximum coherent integration time possible, given limited computational resources.

4.3 Coherent techniques

In this section, coherent techniques are described in moderate detail. A frequentist statistic calculated using coherent integration, called the $\mathcal{F}$-statistic is described in Section 5.1.

If a gravitational wave incident on a detector has a strain amplitude of $h(t)$ and the intrinsic noise of the detector is $n(t)$, then the data collected by the detector $x(t)$ would be

$$x(t) = n(t) + h(t; \mathcal{D}, \mathcal{A}) \tag{4.1}$$

where $\mathcal{D}$ are collectively called the "Doppler" parameters and $\mathcal{A}$ collectively called the "Amplitude" parameters. Splitting the parameters into $\mathcal{D}$ and $\mathcal{A}$ is convenient, since $\mathcal{D}$ only affects the frequency-dependent search parameters like frequency, frequency derivatives, sky location, binary parameters etc., while $\mathcal{A}$ includes parameters that only affect the amplitude, like the inclination angle of the source, initial phase, intrinsic amplitude etc.

A classical Wiener filter of matched-filtering theory [98] is defined as

$$(x||y) = 4\mathcal{R}\left(\int_0^\infty \frac{\tilde{x}(f)\tilde{y}^*(f)}{S_n(f)}df\right), \tag{4.2}$$

where $\tilde{x}(f)$ and $\tilde{y}^*(f)$ are the Fourier transform and complex conjugate of the Fourier transform of $x(t)$ and $y(t)$ respectively. $\mathcal{R}(\bullet)$ is the real part of $\bullet$. The factor of 4 is needed because the integratiion is done only for one of four quadrants. $S_n(f)$ is the single-sided power spectral density, which is the Fourier transform of the auto-correlation function of the noise $n(t)$, defined as

$$S_n(f) = 2\int_{-\infty}^{\infty} \mathcal{E}[n(0)n(t)] \exp^{-i2\pi ft} dt, \tag{4.3}$$

where $\mathcal{E}[\bullet]$ is the expectation value of $\bullet$. This expectation value is calculated by average over multiple instantiations of random noise. The factor of 2 is needed because the integration only involves the positive frequencies. In practice, such an integration from $-\infty$ to ∞ is not possible and a good approximation of $S_n(f)$ is given by

$$S_n(f) \approx \frac{2}{T_{\text{span}}}\mathcal{E}[|\tilde{n}(f)|^2], \tag{4.4}$$

where T_{span} is the time span of the analysis and $\tilde{n}(f)$ is the Fourier transform of the noise $n(t)$, defined as

$$\tilde{n}(f) \approx \int_0^{T_{\text{span}}} n(t)e^{2\pi ift}dt \tag{4.5}$$

Using the inner product defined in Equation 4.2, the probability $P\left(x|\mathcal{A},\mathcal{D},S_n\right)$ of measuring the strain $x(t)$ in the presence of Gaussian noise $n(t)$ (information of which is encoded in S_n) and a signal $h(t)$ as described in Equation 4.1 is given by

$$P\left(x|\mathcal{A},\mathcal{D},S_n\right) \propto e^{-\frac{1}{2}(x-h||x-h)}. \tag{4.6}$$

The signal-to-noise ratio (SNR), ρ for CW searches as shown in [110, 99, 100], is given by

$$\rho = \sqrt{h||h} \propto \frac{h_0}{\sqrt{S_n}} \sqrt{T_{\mathrm{span}} \mathcal{N}}, \tag{4.7}$$

where $\mathcal{N}$ is the number of detectors that are equally sensitive. The computational cost scales approximately as $\mathcal{N}$. Thus incorporating data from multiple detectors is computationally the cheapest way of increasing SNR. The coherent techniques can be split into two categories: one for multiple template searches, the frequentist framework; and another for single template searches, the Bayesian framework. The difference between these frameworks lies in the definition of probability. But in the presence of Gaussian noise and large number of data points, they both give similar results.

4.3.1 Frequentist framework

The frequentist definition of probability involves the relative frequency of occurrence of an event in the limit of a large number of trials. If the number of trials is n_t and the number of times an event x occurs is n_x, then the frequentist probability is given by

$$P(x) \approx \frac{n_x}{n_t}. \tag{4.8}$$

If the number of trials goes to ∞, then the probability converges to

$$P(x) = \lim_{n_t \to \infty} \frac{n_x}{n_t}. \tag{4.9}$$

In the case of gravitational wave data analysis, we usually define a frequentist statistic $\Psi\left(x(t), h(t)\right)$, which is calculated for each form of the signal $h(t; \mathcal{A}, \mathcal{D})$. The question that is of interest to us, is whether there is a non-zero $h(t)$ present in the signal $x(t)$. This question can be answered by testing the hypothesis H_0, that there is no signal in the data, i.e. h=0; or H_1, that there is a signal of strength $h(t)$ present in the data. A couple of quantities that are used to understand the effectiveness of the statistic Ψ are the *False Alarm* probability (ξ) and the

Efficiency (κ) (or its complimentary *False Dismissal* probability (ϱ)). κ and ξ are set by the parameters and tolerances of a particular search. Given these set values of $\kappa_{\rm set}$ and $\xi_{\rm set}$, one can compute a threshold value Ψ^*, such that

$$\xi(\Psi^*) = \int_{\Psi^*}^{\infty} P(\Psi|H_0) d\Psi = \xi_{\rm set} \tag{4.10}$$

and

$$\varrho(\Psi^*) = \int_{-\infty}^{\Psi^*} P(\Psi|H_1) d\Psi = 1 - \kappa_{\rm set} \tag{4.11}$$

because

$$\kappa = 1 - \varrho, \tag{4.12}$$

where $P(\Psi|H_0)$ and $P(\Psi|H_1)$ are the probability distributions of Ψ given the two hypothesis H_0 and H_1 respectively. These probabilty distributions are normalized,

$$\int_{-\infty}^{\infty} P(\Psi) d\Psi = 1. \tag{4.13}$$

Equation 4.10 gives the probability that Ψ crosses Ψ^* despite the absence of any signal. Equation 4.11 gives the probability that Ψ does not cross Ψ^* despite the presence of a signal $h(t)$ as hypothesized in H_1. Equation 4.12 gives us the probability of detecting a signal of signal strength $h(t)$ as hypothesized in H_1.

An optimal detection strategy is to select a statistic $\Psi(x; h)$ such that the efficiency κ is maximized, while minimizing the false alarm rate ξ. According to the Neyman-Pearson lemma [101], this optimal statistic is given by

$$\Psi(x; h) = \frac{P(x|h)}{P(x|0)}, \tag{4.14}$$

which is called the likelihood ratio. In the case of Gaussian noise, using equation 4.6, the log-likelihood ratio is

$$\ln \Psi(x; h) = (x||h) - \frac{1}{2}(h||h). \tag{4.15}$$

A search over multiple templates is done by calculating this $\Psi(x; h)$ for each template $h(t)$. But as we saw in equation 4.1, each $h(t)$ is a function of Doppler parameters $\mathcal{D}$ and amplitude parameters $\mathcal{A}$. Since the amplitude parameters $\mathcal{A}$ are not known, one needs to choose a statistic for which the log-likelihood function is analytically maximized over $\mathcal{A}$. Such a statistic is known as the $\mathcal{F}$-statistic, which is described in [110] and summarized in Section 5.1. The resampling algorithm, described in the next chapter is a computationally efficient technique of calculating this $\mathcal{F}$-statistic.

In the presence of a signal, when the template is chosen to perfectly match this signal

$$\mathcal{E}[2\mathcal{F}] = 4 + \rho^2. \tag{4.16}$$

Also shown in [110], when applied to Gaussian noise, $2\mathcal{F}$ is a random variable, which is drawn from a non-central χ^2-distribution with 4 degrees of freedom and a non-centrality parameter ρ^2. For the noise-only case, the non-centrality parameter reduces to 0 and probability distribution of $2\mathcal{F}$ becomes

$$P(2\mathcal{F}) = \frac{1}{2}\mathcal{F}e^{-\mathcal{F}}. \tag{4.17}$$

Using equations 4.10, 4.11, 4.12 and 4.17, the false alarm probability and efficiency can be calculated. In Chapter 6, I discuss how this is done for an actual search.

4.3.2 Bayesian framework

The probability of two events A and B occurring is given by

$$P(A \cap B) = P(A) \times P(B|A) = P(B) \times P(A|B). \tag{4.18}$$

Rearranging this, we get

$$P(A|B) = \frac{P(A) \times P(B|A)}{P(B)}. \tag{4.19}$$

This is known as Bayes theorem. $P(A|B)$ is known as the posterior probability, $P(A)$ is known as the prior probability, as it reflects our prior knowledge (or lack thereof) of A. $P(B|A)$ is known as the likelihood and $P(B)$ is just a normalizing factor. A better form of equation 4.19 in terms of a hypothesis H and data Data is

$$p(H|\mathrm{Data}) = \frac{p(H) \times p(\mathrm{Data}|H)}{\mathrm{Norm}}. \tag{4.20}$$

The question that we are interested in asking is: what is the probability that our hypothesis (H) is correct given the data? H is usually a function of both $\mathcal{A}$ and $\mathcal{D}$. The term $p(\mathrm{Data}|H)$ is the probability distribution function (pdf) of the data as observed, given the set of parameters in H. $p(H)$ embodies the state of the prior knowledge of the parameters. For some parameters $p(H)$ is sharply peaked, take for example, the sky location of a known pulsar. On the other hand $p(H)$ is uniform in certain parameters like the initial phase of the incoming wave (Φ_0), which reflects the complete lack of knowledge of such a parameter. Norm is just a normalization constant and it includes the $p(\mathrm{Data})$ which is just 1.

The posterior pdf, $p(H|\mathrm{Data})$ contains all the information about the analysis, but it is very difficult to interpret. It is often useful to reduce the dimensionality of the posterior pdf by marginalizing (integrating) over less interesting (nuisance) parameters. The marginal distribution, so derived is the weighted average of the parameter of interest, given all possible combinations of all other parameters. If one were only interested in the gravitational wave strain amplitude h_0, then its marginal pdf can be given by

$$p(h_0|\mathrm{Data}) \propto \int\int\int p(\mathrm{Data}|\mathcal{A}^*)p(\mathcal{A}^*)\ d\phi_0\ d\psi\ d\cos\iota, \tag{4.21}$$

where $\mathcal{A}^*$ are all the amplitude parameters except h_0, which are ϕ_0, the initial phase of the gravitational wave, ψ, the polarization angle and ι, the inclination angle of the source. This marginalization is often computationally expensive and it is not feasible to marginalize over a lot of the Doppler parameters in $\mathcal{D}$. Thus, this method is optimal for known $\mathcal{D}$ and unknown $\mathcal{A}$. If there are uncertainties in some parameters of $\mathcal{D}$, like the frequency derivatives of the source, this method becomes very expensive.

For more details on Bayesian CW data analysis, please see [102]. In the LSC, Bayesian statistics are used when the data is analyzed in the time domain and frequentist statistics are used when the data is analyzed in the frequency domain. This is mostly due to historical reasons and either of the two statistics can be used for both analysis techniques.

4.3.3 Template bank scalings

As mentioned before, CW data analysis is a computationally intensive problem. And this problem is exacerbated in the cases where none of the search parameters are constrained. The number of templates required for a particular search is set by a parameter called the mismatch parameter m. This mismatch paramter is the relative loss in the detection statistic Ψ (or in particular $\mathcal{F}$), incurred due to a Doppler-offset $d\mathcal{D} = \mathcal{D} - \mathcal{D}_0$, where $\mathcal{D}_0$ is the set of correct Doppler parameters as

$$m(\mathcal{D}_0, d\mathcal{D}) = \frac{|\mathcal{E}[\mathcal{F}(\mathcal{D}_0)] - \mathcal{E}[\mathcal{F}(\mathcal{D})]|}{\mathcal{E}[\mathcal{F}(\mathcal{D}_0)]}. \tag{4.22}$$

This mismatch can be interpreted as a distance measure over a multi-dimensional manifold of parameters. A corresponding local metric g_{ij} can then be introduced on the parameter space [104, 105] as

$$m(\mathcal{D}_0, d\mathcal{D}) = g_{ij}(\mathcal{D}_0) d\mathcal{D}^i d\mathcal{D}^j + \mathcal{O}(d\mathcal{D}^3), \tag{4.23}$$

where $\mathcal{E}[\mathcal{F}(\mathcal{D}_0)]$ is a local maximum of $\mathcal{F}$ if there is a signal at $\mathcal{D}_0$. This metric is laid out in detail in [100]. The main parameters that one searches over in $\mathcal{D}$, are frequency f, frequency derivatives f^k (where k represents the k^{th} frequency derivative and $k > 0$), sky positions (right ascension, α and declination, δ). These parameters are often correlated, for example a shift in f^k can also be seen as a little shift in α and δ. The metric required for an all sky search can get very complicated as seen in [100]. For targeted searches, the metric can be simplified, since α and

δ are known. If the number of templates required are to be calculated for f and f^k, then the metric components $g_{ff}, g_{ff^k}, g_{f^j f^k}$ are of interest. The resolution is determined by the amount of shift in a parameter like spindown, that is required to shift the signal by one frequency bin. For the k^{th} derivative, the resolution required scales as $\frac{1}{T^{k+1}}$. A careful calculation of these metric components, like in [100], gives

$$g_{ff} = \frac{(\pi T)^2}{3}, \tag{4.24}$$

$$g_{ff^k} = \frac{(2\pi)^2 T^{k+2} f}{k!(k+2)(k+3)}, \tag{4.25}$$

and

$$g_{f^j f^k} = \frac{(2\pi f)^2 T^{j+k+2}}{j!k!(j+2)(k+2)(j+k+3)}. \tag{4.26}$$

Because of the fact that the resampling algorithm uses a fast Fourier transform method to calculate the $\mathcal{F}$-statistic for multiple frequency templates, the spacing between frequency templates is fixed. Thus the cross diagonal terms (g_{ff^k}) cannot be used and the frequency parameter is projected out. In this projected metric, the template spacings are given by

$$df = 2\sqrt{\frac{m}{g_{ff}}}, \tag{4.27}$$

$$d\dot{f} = 2\sqrt{\frac{m}{\gamma_{\dot{f}\dot{f}}}}, \tag{4.28}$$

$$d\ddot{f} = 2\sqrt{\frac{m}{\gamma_{\ddot{f}\ddot{f}}}}. \tag{4.29}$$

Where,

$$\gamma_{\mu\nu} = \left(g_{\mu\nu} - \frac{g_{f\mu} g_{f\nu}}{g_{ff}} \right). \tag{4.30}$$

This leads to

$$df = \frac{2\sqrt{3m}}{\pi} \frac{1}{T}, \tag{4.31}$$

$$d\dot{f} = \frac{12\sqrt{5m}}{\pi} \frac{1}{T^2} \tag{4.32}$$

and

$$d\ddot{f} = \frac{20\sqrt{7m}}{\pi}\frac{1}{T^3} \tag{4.33}$$

The template scalings for sky position parameters are explained in [100] in detail. We define a unit vector $\vec{n}$, which points to a source located at a right ascension α and declination δ as

$$\vec{n} = (\cos\delta\cos\alpha, \cos\delta\sin\alpha, \sin\delta). \tag{4.34}$$

As mentioned before, sky position parameters are correlated with spindown parameters. If one ignores the antenna-pattern functions and for obesevation times much less than a year, the metric components scale as follows

$$g_{n^i n^j} \approx \left(\frac{fv}{c}\right)^2 T^2 e_i e_j \tag{4.35}$$

and

$$g_{f^k n^i} \approx \frac{T^{k+1} f T v e_i}{c} \tag{4.36}$$

where, v is the orbital velocity of the Earth and e_i is the unit vector as defined in equation 4.34. For longer integration times, the scaling with T gets even larger. But the scaling is at least of order T^2. A simple targeted search involving just the first two frequency derivatives, would have its template bank scale by a factor of T^6. Multiply by the scalings for unknown sky locations for all-sky searches and the computational task becomes quite daunting.

4.4 Semi-coherent techniques

As discussed in the section above, the number of templates required for a fully coherent search grows very quickly with the integration time. For all-sky searches, semi-coherent methods are better suited, which trade off the optimality of the coherent searches for computational efficiency. These search methods do not keep track of the phase information of the signal, instead they add up the power in the

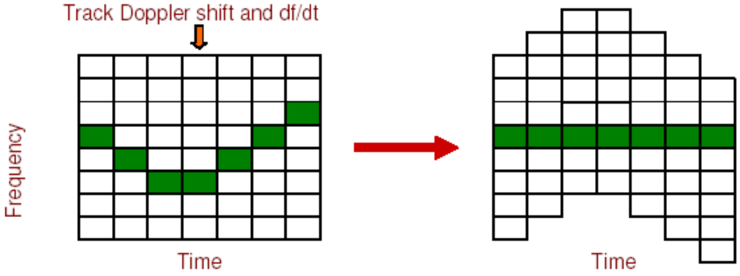

Figure 4.1: An illustration of the semi-coherent techniques. The signal is shown in the shaded green boxes. If the Doppler parameters are matched for a given sky location, excess power can be detected by sliding the appropriate frequency bins and then stacking and adding them together. The method of adding is different in each technique, but the essential idea remains the same.

appropriate frequency bins. There are three different semi-coherent methods that are used by the LSC, which are the StackSlide method, the PowerFlux method and the Hough transform. All of these methods are described in detail in [103].

In all these methods the obervation time T_{obs} is divided into N coherent observational periods called stacks, each of length T_{coh}. A Fourier transform is then done over these smaller stacks. The simplest of the three methods is the StackSlide method, which averages normalized power from each of these stacks after adjusting the frequency bins for the slowly-varying Doppler shift computed for each sky location bin. The PowerFlux method is similar to the StackSlide method, but before averaging, the power is weighted by the noise floor and the detector's antenna pattern for each stack. The weighing function changes as a function of time, sky location and other template parameters. The Hough transform sums a series of binary numbers (zeros or ones) for each stack. A frequency bin in a given stack contributes zero if it fails to cross a threshold set a priori and one if the threshold is crossed. This method is robust to large fluctuations in the detector noise as a function of time. A schematic of semi-coherent techniques can be seen in Figure 4.1.

The sensitivity of semi-coherent techniques scales as

$$\mathcal{E}[h_0] \propto N^{\frac{1}{4}} \sqrt{\frac{S_n}{T_{\mathrm{coh}}}} \tag{4.37}$$

where, h_0 is the strain amplitude. Notice that if $T_{\mathrm{coh}} = T_{\mathrm{obs}}$, then $N = 1$ and thus this reduces to the scaling for a fully coherent search. While semi-coherent searches take a hit in sensitivity, the computational benefits are significant. This can be seen by looking at the scaling of the templates. The frequency templates scale only with the coherent stack time T_{coh} as

$$df \approx \frac{1}{T_{\mathrm{coh}}}. \tag{4.38}$$

As seen in equation 4.31, for coherent searches, df scales as the observation time, which could be as long as years. The typical coherent stack time for semi-coherent searches is of the order of 1800 seconds. Thus the number of frequency templates for a year long search is about 17000 times less for a semi-coherent search. Similarly the scaling for $\dot{f}$ is

$$d\dot{f} \approx \frac{1}{T_{\mathrm{obs}} T_{\mathrm{coh}}}, \tag{4.39}$$

which, when compared to equation 4.32, give us another factor of 17000 for a year long search. These are huge gains in computational efficiency, which make all-sky searches more feasible if these techniques are used.

While all the methods described above take a fast Fourier transform of each stack, it is also possible to use the $\mathcal{F}$-statistic instead [106]. A modified version of the StackSlide method, which sums the $\mathcal{F}$-statistic calculated for each of the coherent stacks is shown to be an optimal method by [107]. For given computational power, the larger the coherent stack time, the greater the sensitivity of the search. The resampling algorithm speeds up the computation of the $\mathcal{F}$-statistic for each coherent stack, allowing for an increase in the coherent time for the same available computational power. Thus the resampling algorithm will increase the efficiency and sensitivity of semi-coherent searches in the future.

Chapter 5

Implementation of Barycentric Resampling

This chapter outlines the implementation of barycentric resampling. It follows reference [110] to define the $\mathcal{F}$-statistic and follows reference [108] in discussing the implementation. The theory is introduced first, followed by a discussion of the pre-processing required to deal with the non-stationarity of the detector and computational constraints. The pre-processing is easier to understand in the time domain and is thus explained that way first. However the actual implementation is in the frequency domain and this is explained in the context of the time domain implementation. The barycentering and antenna patterns follow next and the discussion is completed with a discussion of the practical difficulties like gaps in real data, discreteness of the data and interpolation losses.

5.1 Preliminaries

In this section we closely follow the method of Jaranowski, Krolak, and Schutz [110] to provide the background on the signal and the detection statistic. Power-recycled Fabri-Perot Michelson interferometers such as those used by the Laser Interferometer Gravitational Wave Observatory (LIGO) are sensitive to the strain caused by gravitational waves passing through it. The strain measured at a detector can be written as [110]

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t), \tag{5.1}$$

where t is the time in the detector frame, and h_+ and $h_\times$ are the "plus" and "cross" polarizations of gravitational wave. $F_+(t)$ and $F_\times(t)$ are the beam-pattern functions of the interferometer and are given by

$$F_+(t) = \sin\zeta[a(t)\cos 2\psi + b(t)\sin 2\psi], \tag{5.2}$$

and

$$F_\times(t) = \sin\zeta[b(t)\cos 2\psi - a(t)\sin 2\psi], \tag{5.3}$$

where ψ is the polarization angle of the wave and ζ is the angle between detector arms (which in the case of LIGO is 90°). The functions $a(t)$ and $b(t)$ both depend on time and location of source and detector, but are independent of the polarization angle ψ.

In the detector frame the phase of a gravitational wave produced by an isolated neutron star can be written as [110]

$$\Psi(t) = \Phi_0 + 2\pi \sum_{k=0}^{s} f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c}\mathbf{n}_0 \cdot \mathbf{r}_d(t) \sum_{k=0}^{s} f_0^{(k)} \frac{t^k}{k!}, \tag{5.4}$$

where Φ_0 is the phase at the start time of the observation, $f_0^{(k)}$ is the k^{th} derivative of the frequency, c is the speed of light, α and δ are the right ascension and declination of the source, $\mathbf{n}_0 = \mathbf{n}_0(\alpha, \delta)$ is the unit vector of the source in the Solar System barycenter (SSB) reference frame, $\mathbf{r}_d$ is the position vector of the detector in the same frame, and s is the order of the expansion. Neglecting changes in the proper motion of the star, the third term in Equation 5.4 is a correction to the phase due to the detector motion relative to the neutron star.

We can define $\Phi(t) = \Psi(t) - \Phi_0(t)$, as well as defining

$$\Phi_s(t) = 2\pi \sum_{k=1}^{s} f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c}\mathbf{n}_0 \cdot \mathbf{r}_d(t) \sum_{k=1}^{s} f_0^{(k)} \frac{t^k}{k!} \tag{5.5}$$

and

$$t_m = \frac{\mathbf{n}_0 \cdot \mathbf{r}_d(t)}{c}. \tag{5.6}$$

Equations 5.5 and 5.6 let us write

$$\Phi(t) = 2\pi f[t + t_m(t;\alpha,\delta)] + \Phi_s(t; f^{(k)}, \alpha, \delta), \tag{5.7}$$

which has the modulation due to the detector's motion around the SSB clearly separated from the modulation due to the gravitational wave's intrinsic frequency, although not the derivatives of the frequency.

An almost optimal statistic for the detection of continuous gravitational wave signals is called the $\mathcal{F}$-statistic [110, 111]. It is the logarithm of the likelihood function maximized over the extrinsic and unknown signal parameters. The $\mathcal{F}$-statistic is given by

$$\mathcal{F} = \frac{4}{S_h(f)T_0}\frac{B|F_a|^2 + A|F_b|^2 - 2CR(F_aF_b^*)}{D}. \tag{5.8}$$

where $S_h(f)$ is the one-sided spectral density of the detector's noise at frequency f and T_0 is the observation time. A, B, C, and D are given by

$$A = (a\|a); B = (b\|b); C = (a\|b); D = A \cdot B - C^2 \tag{5.9}$$

with

$$(p\|q) = \frac{2}{T_0}\int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} p(t)q(t)dt. \tag{5.10}$$

F_a and F_b are integrals defined as

$$F_a(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t)x(t)e^{-i\Phi_s(t)}dt \tag{5.11}$$

and

$$F_b(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} b(t)x(t)e^{-i\Phi_s(t)}dt, \tag{5.12}$$

where $x(t)$ is the time series data output by an interferometer. We define a new time variable called t_b as follows:

$$t_b = t + t_m. \tag{5.13}$$

Taking a derivative with respect to t on both sides of Equation 5.13, we get

$$\frac{dt_b}{dt} = 1 + \frac{dt_m}{dt} \tag{5.14}$$

From Equations 5.6 and 5.14, we get

$$\frac{dt_m}{dt} = \frac{\mathbf{n}_0 \cdot \mathbf{v}_d(t)}{c} \tag{5.15}$$

where $\mathbf{v}_d(t)$ is the velocity of the detector in the SSB frame and thus $\frac{\mathbf{n}_0 \cdot \mathbf{v}_d(t)}{c}$ is the Doppler shift of the source with respect to the detector. For a detector located on Earth, the maximum Doppler shift experienced is of the order of 10^{-4}. Using this fact and Equation 5.14 we get $\delta t_b \approx \delta t$. This is the key step in the resampling procedure, which is shown schematically in Figure 5.1.

We can thus rewrite the Equations for F_a and F_b as

$$F_a(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t_b) x(t_b) e^{-2\pi i f t_b} e^{-i\Phi_s(t_b)} dt_b, \tag{5.16}$$

and

$$F_b(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} b(t_b) x(t_b) e^{-2\pi i f t_b} e^{-i\Phi_s(t_b)} dt_b \tag{5.17}$$

which are just the Fourier transforms of the resampled data and the detector response, multiplied by a phase $e^{-i\Phi_s(t_b)}$ [110]. Equations 5.16 and 5.17 can be efficiently evaluated using FFTs. Details of the resampling procedure can be found in Section 5.2.1.

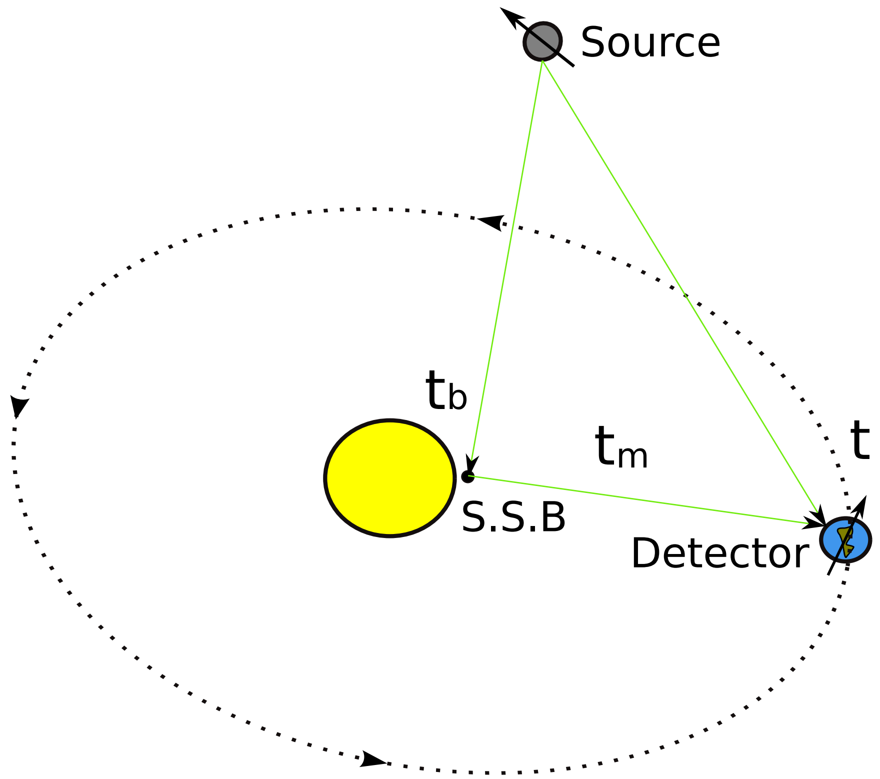

Figure 5.1: Cartoon showing the time differences used in resampling. S.S.B is the solar system barycenter and it is outside the sun because of the mass and large lever arm of Jupiter. The time elapsed on the detector is t and on the solar system barycenter is t_b.

5.2 Implementation of barycentric resampling

Gravitational wave detectors collect data at the rate of about 16-20 kHz for spans of time on the order of a year. This means that typical searches for gravitational waves will involve on the order of a terabyte (TB) of data. Computers currently have memories of a few gigabytes (GB), making it necessary to break up the data into pieces that can fit in the memory of a single computer. To analyze the full data set hundreds to thousands of these computers can then be used together in the form of a Beowulf cluster, or tens to hundreds of thousands with distributed computing systems such as Einstein@Home [88].

A flow-chart showing the resampling algorithm, as it is imlpemented here is shown in Figure 5.2. This flow-chart is for the analysis as it is done in the frequency domain. However conceptually the time domain analysis is easier to understand and thus we present it first, followed by the equivalent, but more complicated frequency domain analysis.

5.2.1 Time Domain Analysis

The $\mathcal{F}$-statistic can be calculated from a time series directly by following the steps outlined in Section 5.1. However, due to the large amounts of data involved, it is impractical to do this for the entire data set. One way to address this problem is to divide the data into band-limited time series, making it possible to analyze one small sub-band at a time. Time series spanning different frequency bands are then analyzed in parallel on a Beowulf cluster or a distributed computing system. In this section we provide details on how this is accomplished in the time domain, and address some of the difficulties that arise.

Heterodyning, low-pass filtering, and downsampling

Let the output of the instrument be the time series $x(t)$, and its Fourier transform be

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt. \tag{5.18}$$

If we consider the Fourier transform of the complex time series $x_h(t) = x(t)e^{-2\pi i f_h t}$,

$$\begin{aligned} \tilde{x}_h(f) &= \int_{-\infty}^{\infty} x(t) e^{-2\pi i f_h t} e^{-2\pi i f t} dt \\ &= \int_{-\infty}^{\infty} x(t) \cdot e^{-2\pi i (f+f_h) t} dt \\ &= \tilde{x}(f+f_h), \end{aligned} \tag{5.19}$$

it is obvious that multiplying the time series $x(t)$ by $e^{-2\pi i f_h t}$ has shifted all the frequencies in the time series $x(t)$ by f_h. This procedure is referred to as complex heterodyning.

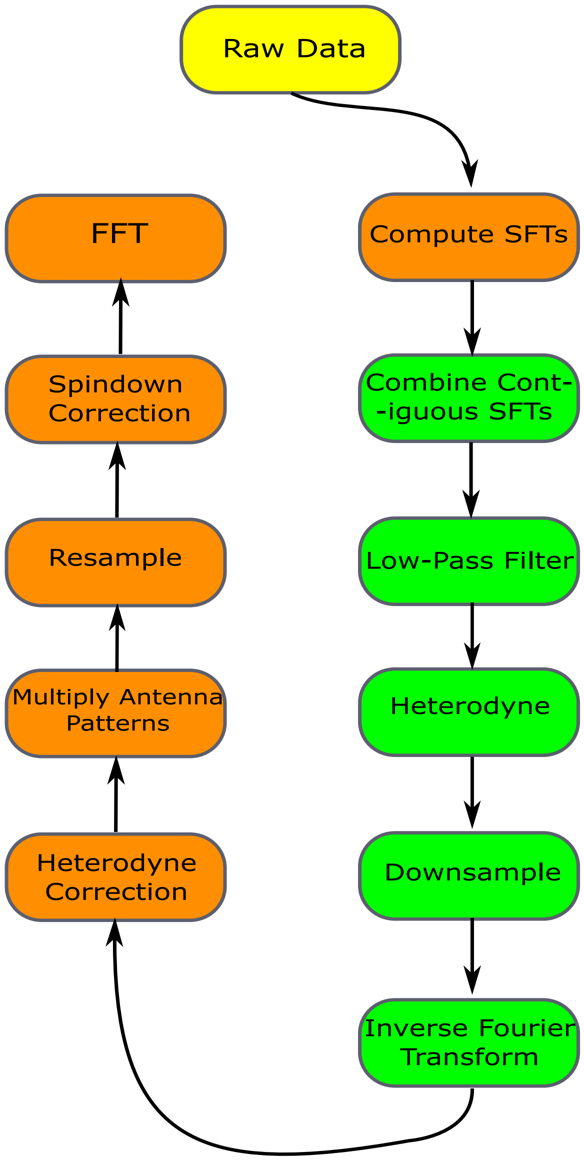

Figure 5.2: A flow-chart showing the resampling algorithm. The green refers to an operation done in the frequency domain and the orange refers to an operation done in the time domain.

If just a small frequency band B of data around f_h is of interest, low-pass filtering followed by downsampling can be used to reduce the bandwidth of the data appropriately. Specifically, if we wish to downsample by a factor D, the new Nyquist frequency of our time series will be given by

$$f_{\mathrm{Nyq,new}} = \frac{f_{\mathrm{Nyq,old}}}{D} = \frac{B}{2}. \tag{5.20}$$

A simple but effective downsampling technique involves picking every D^{th} point in the time series. To avoid aliasing effects however, prior to downsampling a low pass filter must be applied to the data with a sharp fall-off around the new Nyquist frequency. The heterodyned, band-limited, downsampled complex time series will have a sampling time $\Delta t = 1/B$. For example, suppose we are only interested in analyzing data between 990 Hz and 1 kHz. By multiplying the data with the phase factor $e^{-2\pi(995)it}$, data at 995 Hz moves to 0 Hz (DC), 990 Hz moves to -5 Hz, and 1 kHz to +5 Hz (we have taken t to be measured in seconds). To avoid aliasing problems when we downsample, we low-pass filter the data at 5Hz, the new Nyquist frequency. We can then downsample by picking one point out of every 100. The resulting complex time series will be sampled at 10 Hz and contain all the information in the original time series between 990 Hz and 1 kHz.

Barycentric resampling and heterodyne correction

In this section we explain how to use the low bandwidth heterodyned complex time series to compute the $\mathcal{F}$-statistic given by Eq. (5.8).

In the following we will work only with F_a. The procedure for F_b is completely analogous. It is easiest to begin with the integral definition for F_a in Eq. (5.11) with the phase explicitly written out, namely,

$$F_a(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t)x(t)e^{-2\pi i f(t+t_m)}e^{-i\Phi_s(t)}dt, \tag{5.21}$$

and a similar expression holds for F_b. The heterodyned version of F_a is

$$F_a(f+f_h) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t)x(t)e^{-2\pi i(f+f_h)(t+t_m)}e^{-i\Phi_s(t)}dt. \tag{5.22}$$

If we already have a complex heterodyned time series $x_h(t)$ (heterodyned in the detector frame), we can use it to absorb some (but not all) of the heterodyne exponent in Eq. (5.22) as follows:

$$x(t)e^{-2\pi i(f+f_h)(t+t_m)} = x_h(t)e^{-2\pi i f_h t_m}e^{-2\pi i f(t+t_m)}. \tag{5.23}$$

This means that rather than Eq. (5.22), we should evaluate

$$F_a(f+f_h) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t)z(t)e^{-2\pi i f(t+t_m)}e^{-i\Phi_s(t)}dt, \tag{5.24}$$

where

$$z(t) = x_h(t)e^{-2\pi i f_h t_m}. \tag{5.25}$$

At this point we have an expression which looks like Eqs. (5.11) and (5.12), and we can write the integral over t instead as an integral over t_b:

$$F_a(f+f_h) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} a(t_b)z(t_b)e^{-2\pi i f t_b}e^{-i\Phi_s(t_b)}dt_b, \tag{5.26}$$

with a similar expression for F_b.

The discrete version of Eq. (5.26) for a time series with N points reads

$$F_a(f+f_h) = \sum_{k=1}^{N} a(t_b^k)z(t_b^k)e^{-2\pi i f t_b^k}e^{-i\Phi_s(t_b^k)}dt_b, \tag{5.27}$$

and a similar expression holds for F_b:

$$F_b(f+f_h) = \sum_{k=1}^{N} b(t_b^k)z(t_b^k)e^{-2\pi i f t_b^k}e^{-i\Phi_s(t_b^k)}dt_b, \tag{5.28}$$

where t_b^k is the k^{th} datum in the time series as measured in the barycentric frame and $dt_b = t_b^{k+1} - t_b^k$. The relationship between t_b and t can be written as

$$t_b^k = t^k + t_m(t^k; \alpha, \delta). \tag{5.29}$$

This relationship between t^k and t_b^k can be used to calculate $z(t_b^k)$ from the time series $z(t^k)$. In practice, one starts out with $z(t^k)$, i.e. data sampled regularly in the detector frame. Then we calculate $T^k(t_b^k)$, which are detector times corresponding to regularly spaced samples in the barycentric frame. These $T^k(t_b^k)$ are irregularly sampled in the detector frame, but since we have $z(t^k)$, we can calculate $z(T^k(t_b^k))$ by using interpolation. The interpolated time series $z(T^k(t_b^k))$ is the $z(t_b^k)$ of Eqs. (5.27) and (5.28). A similar procedure may be used to calculate the $a(t_b^k)$ from $a(t^k)$, and the $b(t_b^k)$ from $b(t^k)$. The factor of $e^{i\Phi_s(t_b^k)}$ in Equations 5.27 and 5.28 is calculated using Equation 5.5. In this case, instead of calculating $\Phi_s(t^k)$, we calculate $\Phi_s(T^k(t_b^k))$, which is equivalent to calculating $\Phi_s(t_b^k)$. While in theory one has to calculate the quantity $\mathbf{n}_0 \cdot \mathbf{r}_d(t)$ in Equation 5.5, in practice this information is already encoded in $T^k(t_b^k)$ as

$$\mathbf{n}_0 \cdot \mathbf{r}_d(t) = t_m \cdot c = (t_b^k - T^k(t_b^k)) \cdot c\,. \tag{5.30}$$

With all the parts of Equations 5.27 and 5.28 in hand, we can compute $F_a(f + f_h)$ and $F_b(f + f_h)$.

In summary, as shown in Figure 5.3, the procedure is the following:

1. Start with a heterodyned, band-limited, downsampled $x_h(t^k)$ with t^k regularly spaced in time, in the frame of reference of the detector.

2. Correct the $x_h(t^k)$ for the heterodyning done in the detector frame by multiplying with $e^{-2\pi i f_h t_m}$ to produce the $z(t^k)$.

3. The $z(t^k)$ correspond to data irregularly spaced in the barycentric frame. Calculate $T^k(t_b^k)$, which are times in the detector frame corresponding to regularly sampled solar system barycenter times.

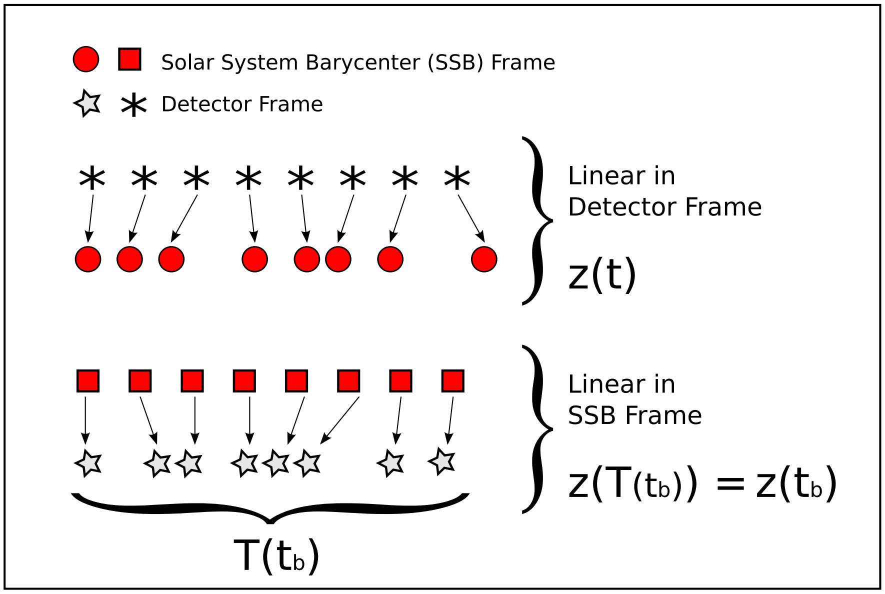

Figure 5.3: Graphical description of the resampling procedure

4. Using interpolation, calculate $z(T^k(t_b^k))$ from $z(t^k)$, which is the $z(t_b^k)$ used in Eqs. (5.27) and (5.28).

5. Similarly, from $a(t^k)$ and $b(t^k)$ calculate $a(t_b^k)$ and $b(t_b^k)$ respectively.

6. Using FFTs, evaluate Eqs. (5.27) and (5.28) to calculate $F_a(f + f_h)$ and $F_b(f + f_h)$.

7. Use Eq. (5.8) to calculate the $\mathcal{F}$-statistic.

5.2.2 Frequency Domain Analysis

In the previous section we describe a practical way of calculating the $\mathcal{F}$-statistic from time series data. However, in practice the calculation is done in the frequency domain for a couple of reasons. One is that much of the code written in the LIGO Scientific Collaboration's (LSC) Continuous Waves working group is tailored to an analysis performed in the frequency domain and hence there exist many data processing and validation tools to process the data that are useful to this code. Another reason is that gravitational wave detectors are subject to many sources of noise, some of which change daily or even hourly, such as wind, microseism, earthquakes, anthropogenic noise, etc. These change the noise floor of any analysis as a function of time. Working in the frequency domain is a natural way to deal with this problem.

We begin a frequency domain analysis by taking short time-baseline Fourier transforms of the time domain data, called short Fourier transforms (SFTs). When we calculate the $\mathcal{F}$-statistic, we divide by the noise in the instrument at that frequency, as shown in Eq. (5.8). However, Eq. (5.8) assumes the noise is stationary. To account for the non-stationarity of the noise we need to weight by the noise over time, which is done on a per SFT basis. This normalization process is described in the next section.

The computational cost of estimating the noise per SFT scales with the number of SFTs and thus for a fixed observation time scales inversely with the time-baseline. A compromise is needed between the demands of computational time and relative stationarity of the detector for a given time-baseline. In LIGO, SFTs are usually 1800 seconds long, since the detector is reasonably stationary for that time.

Dealing with non-stationary and colored data

To deal with non-stationarities, variations in the noise floor from SFT to SFT, and colored data, we can normalize our SFT data to absorb the $1/S_h(f)$ term in the definition of the $\mathcal{F}$-statistic in Eq. (5.8). If $X_{\alpha,k}$ is the k^{th} frequency bin of the α^{th} SFT, then we can redefine a normalized data point $\hat{X}_{\alpha,k}$ as

$$X_{\alpha,k} \longrightarrow \hat{X}_{\alpha,k} = \frac{X_{\alpha,k}}{\sqrt{S_{\alpha,k}}} \,, \tag{5.31}$$

where $S_{\alpha,k}$ is an estimate of the one-sided power spectral density for the k^{th} frequency bin of the α^{th} SFT. Estimators used for this purpose should be robust in the presence of spectral features in the data, such as a running median.

Merging SFTs into long time-baseline Fourier transforms

There are many practical difficulties that arise when dealing with SFTs. Often contiguous chunks of data have to be divided up into multiple SFTs and it is necessary to coherently combine them into one long time-baseline SFT. This is done using the Dirichlet kernel, which is the equivalent of a sinc interpolation (ideal interpolation) done in the time domain. In order to keep the computational cost down, the Dirichlet kernel is truncated at a finite number of points (usually around 16). This introduces a slight interpolation error, which cannot be avoided without sacrificing a large amount of computational power.

Suppose we divide the data $x(t)$ of length T_0 into M short chunks of length $T_{\rm SFT}$ each with N points, so that $T_0 = MT_{\rm SFT}$. The discrete Fourier transform (DFT) of the data is

$$X_b = \sum_{l=0}^{NM-1} x_l e^{-2\pi i l b/NM}, \tag{5.32}$$

where $x_l = x(l\Delta t)$, Δt is the sampling time, and b is a long time-baseline frequency index. We can write the Fourier transform in terms of two sums:

$$X_b = \sum_{\alpha=0}^{M-1} \sum_{j=0}^{N-1} x_{\alpha,j} e^{-2\pi i b(j+N\alpha)/NM}\ , \tag{5.33}$$

where $x_{\alpha,j} = x((j+N\alpha)\Delta t)$. We can express the $x_{\alpha,j}$ in terms of an inverse DFT of a short chunk of data,

$$x_{\alpha,j} = \frac{1}{N} \sum_{k=0}^{N-1} X_{\alpha,k}\, e^{2\pi i j k/N}, \tag{5.34}$$

where the $X_{\alpha,k}$ are the starting SFT data,

$$X_{\alpha,k} = \sum_{j=0}^{N-1} x_{\alpha,j}\, e^{-2\pi i j k/N}. \tag{5.35}$$

Replacing $x_{\alpha,j}$ with Eq. (5.34) in Eq. (5.33) gives

$$\begin{aligned} X_b &= \sum_{\alpha=0}^{M-1}\sum_{j=0}^{N-1}\left(\frac{1}{N}\sum_{k=0}^{N-1} X_{\alpha,k}\, e^{2\pi ijk/N}\right) e^{-2\pi ib(j+N\alpha)/NM} \\ &= \frac{1}{N}\sum_{\alpha=0}^{M-1} e^{-2\pi ib\alpha/M}\sum_{k=0}^{N-1} X_{\alpha,k}\sum_{j=0}^{N-1} e^{-2\pi ij(b/M-k)/N}\ . \end{aligned} \tag{5.36}$$

The last sum in this expression can be evaluated analytically. In particular,

$$\sum_{j=0}^{N-1} z^{cj} = \frac{1-z^{Nc}}{1-z^{c}}. \tag{5.37}$$

We take $z = e$, $c = -iy/N$, with $y = 2\pi(b/M - k)$, so that the sum is given by

$$\sum_{j=0}^{N-1} e^{-iyj/N} = \frac{1-e^{-iy}}{1-e^{-iy/N}}\ . \tag{5.38}$$

In the large N limit the exponent of the denominator will be small so that

$$\begin{aligned} \frac{1-e^{-iy}}{1-e^{-iy/N}} &\approx \frac{1-e^{-iy}}{1-(1-iy/N)} = \frac{iN}{y}(e^{-iy}-1) \\ &= N(\frac{\sin y}{y} - i\frac{1-\cos y}{y}). \end{aligned} \tag{5.39}$$

This means we can write Eq. (5.36) as

$$X_b = \sum_{\alpha=0}^{M-1} e^{-2\pi ib\alpha/M}\sum_{k=0}^{N-1} X_{\alpha,k}P_{b,k}, \tag{5.40}$$

with the Dirichlet kernel

$$P_{b,k} = \frac{\sin y}{y} - i\frac{1-\cos y}{y}, \tag{5.41}$$

and $y = 2\pi(b/M - k)$. The function $P_{b,k}$ is very strongly peaked around $y = 0$, which is near a value of the frequency index $k^* = \text{floor}(b/M)$. This means one only needs to evaluate the sum over k for a few terms Δk around k^*. With this in mind we write

$$X_b \approx \sum_{\alpha=0}^{M-1} e^{-2\pi i b\alpha/M} \sum_{k=k^*-\Delta k}^{k^*+\Delta k} X_{\alpha,k} P_{\alpha,k}. \tag{5.42}$$

To produce a heterodyned time series a sub-band of the X_b may be selected and inverse Fourier transformed.

Normalized long time-baseline Fourier transforms

With the normalized SFT data $\hat{X}_{\alpha,k}$ from Eq. (5.31) we can construct a normalized version of the long time-baseline Fourier transform

$$\hat{X}_b \approx \sum_{\alpha=0}^{M-1} e^{-2\pi i b\alpha/M} \sum_{k=k^*-\Delta k}^{k^*+\Delta k} \hat{X}_{\alpha,k} P_{\alpha,k}, \tag{5.43}$$

and take a sub-band of $\hat{X}_b$, inverse Fourier transform it, and produce the heterodyned time series, and correct it to produce $\hat{z}(t_b^k)$. In terms of this time series, we can write

$$\hat{F}_a(f+f_h) = \sum_{k=1}^{N} \hat{z}(t_b^k) a(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \tag{5.44}$$

and

$$\hat{F}_b(f+f_h) = \sum_{k=1}^{N} \hat{z}(t_b^k) b(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \ , \tag{5.45}$$

and thus

$$\mathcal{F} = \frac{4}{T_0} \frac{B|\hat{F}_a|^2 + A|\hat{F}_b|^2 - 2C\Re(\hat{F}_a \hat{F}_b^*)}{D}. \tag{5.46}$$

Heterodyning

As shown before in Eqs. (5.18) and (5.19), heterodyning is a procedure by which the frequency of interest can be shifted arbitrarily. When one applies the kind of correction in Eq. (5.18), we effectively move all the frequencies by a set amount. By doing so, we convert the time series from a real time series to a complex time series, with the same amount of information content.

Heterodyning in the frequency domain can be done in two ways, one in which the time series produced after inverse Fourier transforming is real and another in which it is complex. A cosine transform used to heterodyne would produce a real time series, but this method is not used in an implementation of the technique (see Section 5.3). A complex heterodyned time series is produced by inverse Fourier transforming a relabeled band of the frequencies. Since in Eq. (5.18), all frequencies are shifted by a fixed amount, the equivalent procedure in the frequency domain is just relabeling the heterodyne frequency f_h as DC and subsequently all the other frequencies relative to this new DC.

Taking the example from Section 5.2.1, we can just internally change the labels of the 995 Hz frequency bin to DC and 1000 Hz to 5 Hz. Once this relabeling is done, the original data will have all shifted by 995 Hz, with the 10 Hz from -5 Hz to +5 Hz containing all the relevant information. If one were using the whole band without downsampling or filtering, then this relabeling would have to wrap around the Nyquist frequency edge, but since the whole purpose of heterodyning is to downsample, it is never necessary to do so.

Downsampling and low-pass filtering

Following the time domain algorithm, after heterodyning the data, it needs to be downsampled and low-pass filtered. The downsampling and low-pass filtering is achieved by simply throwing out the data that is not in the band of interest. The heterodyning is done in such a way as to keep the center of the band of interest at DC. A Tukey window applied to the band of interest, keeping a little bit of data on both edges to facilitate the rise of the window from 0 to 1, is a good choice of a low-pass filter. Once an inverse Fourier transform is performed on this smaller subset in the frequency domain, it generates the same heterodyned, downsampled, and low-pass filtered time series as the time domain algorithm.

Gaps in the data

Data collected by an interferometer will have gaps due to periods of downtime. These gaps need to be dealt with in a manner that preserves the phase coherence of the segments around the gaps. The gaps increase the analysis time without contributing any power to the $\mathcal{F}$-statistic, and thus act like a zero padding.

The data is divided up into a series of contiguous chunks and gaps. For each contiguous chunk, the SFTs in that chunk are normalized, patched up and then a heterodyned, downsampled and low-pass-filtered time series is calculated from it. Heterodyning done by relabeling is equivalent to multiplying with $e^{-2\pi i f_h (t-t_c)}$, where t_c is the start time of the data chunk being heterodyned and f_h is the heterodyne frequency. If we have multiple chunks that are separately being heterodyned, then t_c is different for each chunk. In the time domain analysis, we assumed that the heterodyne reference time is the same as the start time of the analysis. In order to achieve the same kind of heterodyning, one needs to multiply each newly created time series with a correcting phase factor, namely

$$e^{-2\pi i f_h (t_c - t_s)}, \tag{5.47}$$

where t_s is the start time of the overall analysis.

A Tukey window can then be applied to each of these time series to smoothly bring the data to zero at the edges, which correspond to the gaps. The gaps are then filled with zeros, as no data was collected during those times. This procedure is repeated for all the gaps and contiguous chunks. At the end, a time series is produced, which is contiguous and spans the time of the analysis. By ensuring that the timestamps of the first datum of each contiguous chunk correspond with the start time of that chunk, we ensure that the phase coherence is maintained throughout.

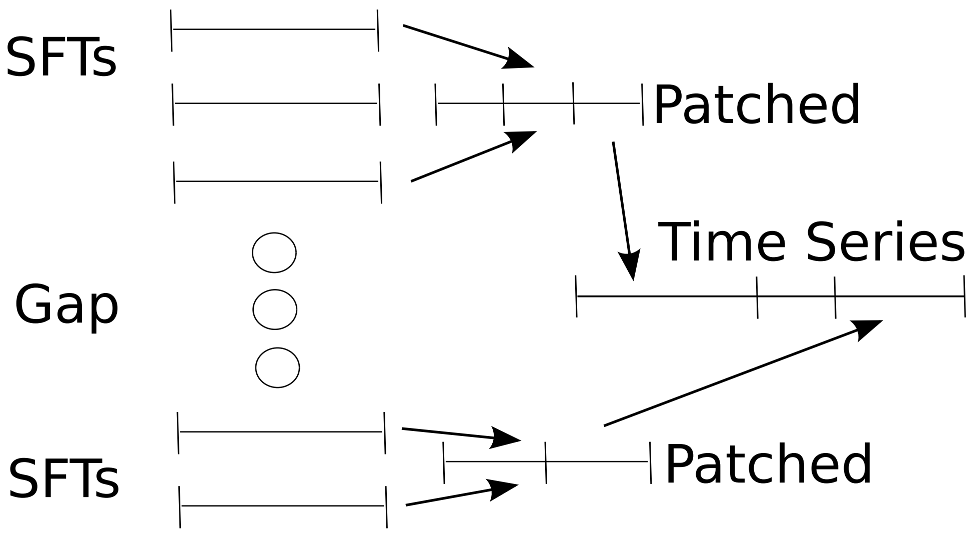

Figure 5.4: Pictorial description of data pre-processing

Summary

To summarize, as shown in Figure 5.4, a simple algorithm to produce a time series equivalent to the one used for the time domain analysis is as follows:

1. Divide the data into time chunks and Fourier transform them to create SFTs.
2. Normalize these SFTs and assign them weights.
3. Identify contiguous sets of SFTs.
4. Combine each contiguous chunk of SFTs into one long time-baseline Fourier transform (FT).
5. Create a downsampled, heterodyned, and low-pass-filtered time series by inverse Fourier transforming the desired frequencies from the FT.
6. Stitch all these time domain chunks together, filling gaps with zeros.

5.3 Results

5.3.1 Speed

The scheme previously used to compute the $\mathcal{F}$-statistic, involved the use of the Dirichlet kernel to combine a series of SFTs [115] [112], which were calculated for 30 minutes of data taken at 16 kHz. The 30 minute window was set by the maximum Doppler shift due to the motion of the Earth. A C code called ComputeFStatistic_v2 [113] was written in the LIGO Analysis Library (LAL) to calculate the $\mathcal{F}$-statistic using this algorithm. The code which implements our method is also written in C and is called ComputeFStatistic_resamp [114]. Henceforth we will refer to the previous implementation as the LAL implementation and our implementation as Resampling.

The $\mathcal{F}$-statistic is calculated for a series of templates looping over various parameters such as sky location, α and δ, spin-downs f^k, and various frequencies f. We can ignore the way the two implementations deal with loops over α, δ, and f^k, since they both loop over them in the same manner. The speed of computation for a loop over frequencies f is worth comparing, however. Figure 5.5 gives a pictorial description of the computational costs associated with both the implementations.

Assume that we have N data points (take for example 10^6 seconds of data at 100 Hz, i.e. 10^8 data points). For a simple case, N is also the number of frequency templates that need to be computed. In case of a specified mismatch parameter, the number of frequency templates scale as N. For a given sky location and spindown parameters, if the number of operations required for computation in the previous implementation is $N_{\mathrm{LAL\_ops}}$ (in Figure 5.5, it is the sum of the operations required by the Dirichlet Kernel and Numerical operations), then the computational cost for computing all the frequency templates is given by

$$N_{\mathrm{Tot}}^{\mathrm{LAL}} = N_{\mathrm{LAL_{ops}}} \cdot N_{\mathrm{SFTs}} \cdot N \ , \tag{5.48}$$

where N_{SFTs} is the number of SFTs used in the analysis. $N_{\mathrm{LAL\_ops}}$ is approximately of the order of 30.

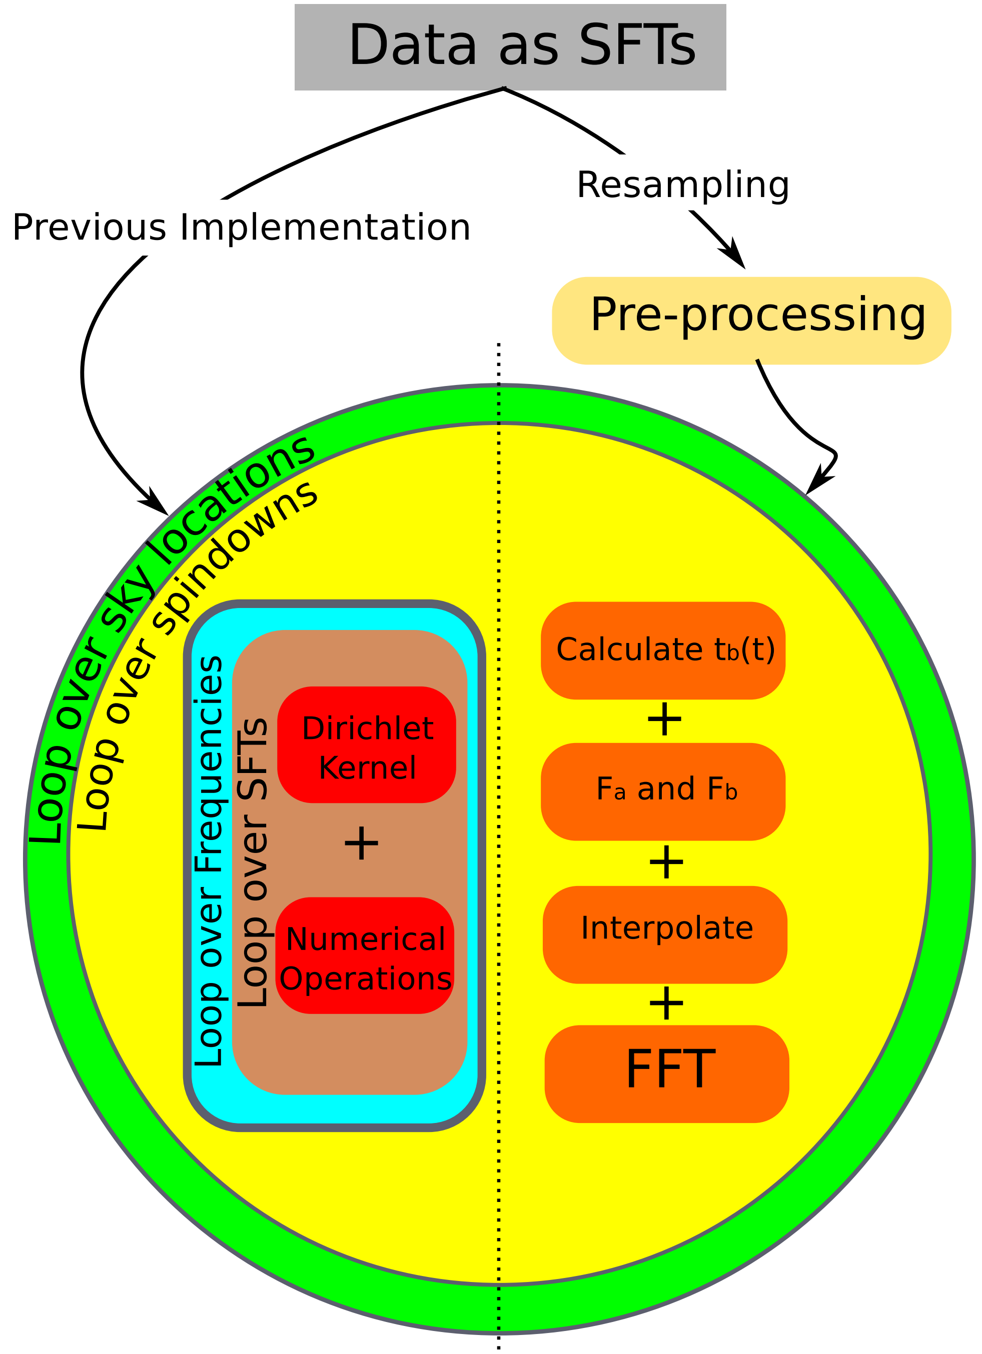

Figure 5.5: Pictorial representations of the computational costs associated with the previous implementation of the $\mathcal{F}$-statistic algorithm and the resampling algorithm. Notice the reduction in the number of loops required by resampling.

Compare this to the resampling method, which consists of 4 major steps:

1. Calculating $t_b(t)$, given a sky location and time.
2. Calculating the integrands of F_a and F_b.
3. Interpolating and calculating the beam patterns.
4. Taking the Fourier transform.

Each of these steps involves order 10 operations, but all of these steps are sequential, therefore they only add, resulting in a total number of operations per data point, $N_{\text{Resamp\_ops}}$, of approximately 30 operations. The last step is the Fourier transform, which is of order $N \log N$, therefore the total number of steps is:

$$N_{\text{Tot}}^{\text{Resamp}} = (N_{\text{Resamp\_ops}} + \log N) \cdot N \ . \tag{5.49}$$

Therefore the ratio of operations between the two methods is

$$\frac{N_{\text{Tot}}^{\text{LAL}}}{N_{\text{Tot}}^{\text{Resamp}}} = \frac{N_{\text{LAL\_ops}} \cdot N_{\text{SFTs}}}{N_{\text{Resamp\_ops}} + \log N} . \tag{5.50}$$

To first order, we have

$$\frac{N_{\text{Tot}}^{\text{LAL}}}{N_{\text{Tot}}^{\text{Resamp}}} \approx \frac{N_{\text{SFTs}}}{\log N} \ . \tag{5.51}$$

Therefore for large observation times, this method of calculating the $\mathcal{F}$-Statistic is faster and, in the case of a targeted search, it allows for a large parameter space in $F^{(k)}$'s.

The speed-up in practice is reduced by a few practical issues as seen in section 5.4. However, Resampling is still considerably more efficient than the LAL implementation. For Einstein@Home, because of the relatively small coherent integration time, the speed-up is around 10. But for targeted searches that span multiple months or years, the improvement can be as high as a factor of 2000. Thus, while some targeted searches which integrate over a couple of years were impossible to do previously, they are now possible.

5.3.2 Validations

The probability density distribution of the $\mathcal{F}$-statistic for Gaussian noise of zero mean and unity standard deviation is a χ^2 distribution with four degrees of freedom. In the presence of a signal, the distribution is a χ^2 of four degrees of freedom with a non-centrality parameter given by the $\mathcal{F}$-statistic in the absence of noise for the particular signal.

Resampling uses various approximate methods in the calculation of the $\mathcal{F}$-statistic, and this can lead to disagreements between the theoretical $\mathcal{F}$-statistic probability density function and the output of the code. These changes are of the order of a few percent and are within acceptable limits. The validity of the code can be tested by using a Monte Carlo simulation of about a million different injections of the same signal in different instances of noise. The noise is generated as a Gaussian noise of zero mean and unity standard deviation, and the signal is added into this noise. For each individual injection the signal is chosen with a given set of amplitude parameters and a fixed sky location and spindowns, and the search is conducted over these exact chosen parameters in order to avoid any mismatches. These Monte Carlos are then repeated with another set of parameters, which are themselves chosen randomly. While it is not an exhaustive test, randomly chosen parameters ensure that we are not biased in the validation test. The plot in Figure 5.6 is produced by performing one such Monte Carlo simulation. In this case, both the LAL implementation and Resampling were run on the same set of data. The $\mathcal{F}$-statistic was picked out at the appropriate frequency and this was repeated about a million times. A histogram of these $\mathcal{F}$-statistic values was then plotted. As one can see, there is very good agreement in between the expected distribution of the $\mathcal{F}$-statistic and the two implementations.

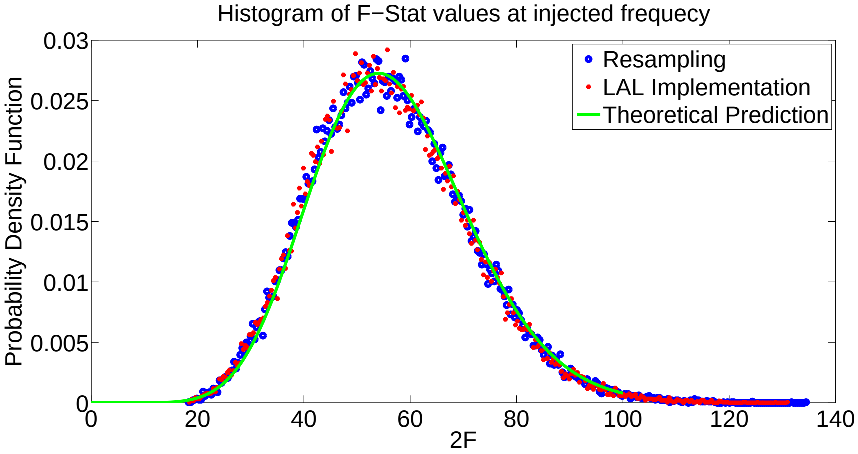

Figure 5.6: Histogram of results of Monte Carlo simulation with signals injected in different instances of noise.

5.4 Practical Considerations

5.4.1 Discreteness

In the implementation of the algorithm explained above, one major obstacle is the fact that the data collected by any physical instrument is discrete and thus must be handled appropriately. Take, for example, the heterodyne frequency used in the calculation. This frequency cannot be chosen arbitrarily, as only certain frequencies are sampled and thus there are only certain permitted choices.

Most major FFT computation algorithms output the frequency series in a specific format, which split the data into two parts. The first bin output by these algorithms is the DC followed by the first positive frequency bin up to positive Nyquist and then follows this up with the negative frequencies starting at the negative Nyquist frequency. This order of placing frequency bins speeds up computation and is necessary for the internal workings of these algorithms. Thus when an inverse FFT is performed on the frequency domain data in the form of SFTs, a simple reshuffling needs to be done. The frequency selected to be the first bin will become the new DC and thus the data will have been heterodyned by that said frequency. In order to ensure that the same frequency bin is chosen as DC, one needs an odd number of bins per SFT. If the number of bins are even, then upon increasing the amount of data it can shift this number to an odd number as the increase is always done by changing the number of SFTs. But if the number of bins per SFT is odd, then it will remain odd for any number of SFTs. This ensures that there is no mismatch in choosing the appropriate bin as the heterodyne frequency.

5.4.2 Interpolation Issue

The resampling algorithm is always used on discrete data in practice. In order to calculate the resampled time series, as shown in the sections above, one needs to interpolate. In theory, for a Nyquist limited time series, one can find the exact value of the function at any time. But one needs to use a sinc interpolation technique which would require a computation involving all the data points in the time series. This is computationally very expensive and is impractical for GW data analysis. Thus a computationally cheaper, but lossy interpolation technique needs to be used.

The reason why interpolation leads to losses, can be seen in the extreme case of a time series with a signal at Nyquist (a triangle wave), which is interpolated using a linear interpolater to compute the points exactly in the middle of the currently sampled points. This would lead to a complete loss of the signal. However this is an extreme example of the lossy nature of interpolation and in most instances interpolation is not quite that lossy. Figure 5.7 is the interpolation response of a linear interpolater when calculated at different offsets.

The resampling algorithm does not use a constant offset, since the point to interpolate depends on the relationship between the barycentric and detector times, which are changing as a function of time, sky location and other parameters. This means that the interpolation response is a convolution of all a set of constance offset interpolation responses. The actual response is also a function of the interpolation technique used.

In order to isolate the effects of interpolation, a toy-model of resampling was constructed, which consisted of a reciever rotating around the Sun in a perfectly

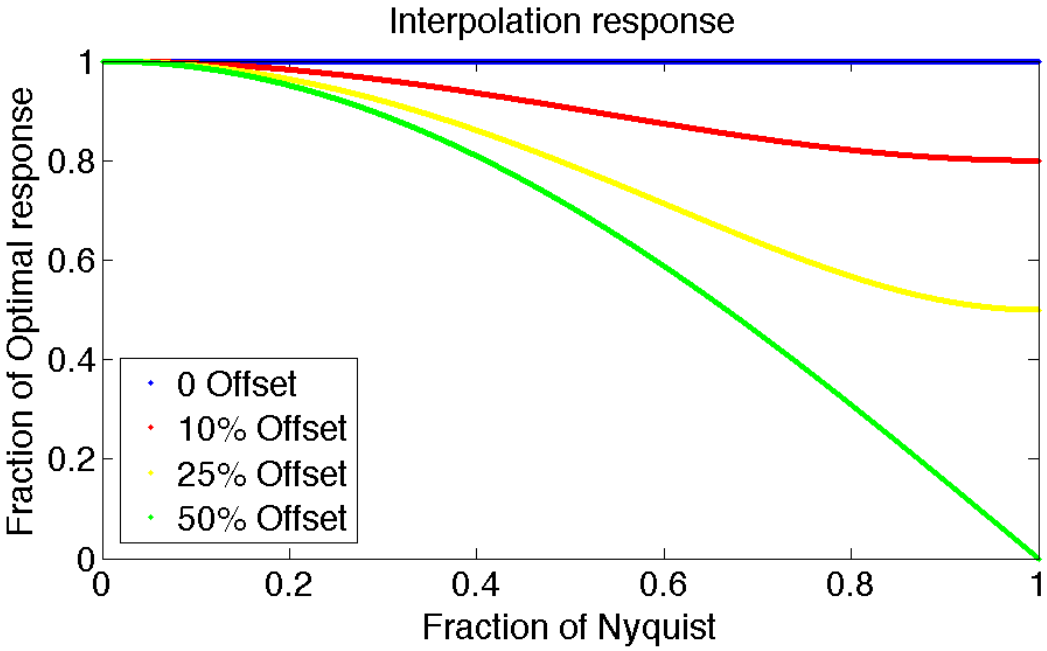

Figure 5.7: Linear interpolation response. The figure shows the loss of signal (y axis) as a function of the frequency of the signal (x axis). The four curves are the offset at which the interpolation is done. 0% refers to no interpolation and thus a 100% response. 50% is the worst case and the 0% response at Nyquist is the triangle wave case explained above.

circular orbit and a source that was optimally oriented and emitting a mono-amplitude sinusoid. This helped remove the effects of sky location, antenna patterns and spindowns from the computation. The signal was then generated at many different frequencies ranging from near 0Hz to the Nyquist frequency. The response was then normalized to what was expected. A similar analysis was done by using noise, in which the comparison was done with the injected amplitude, since the resampling procedure ideally, should not change the amplitude of the noise. The results for a linear interpolater and a cubic spline interpolater are shown in Figures 5.8 and 5.9.

Some amount of loss, while using interpolation is inevitable, since it is an approximation of the exact solution. However, the amount of loss that can be

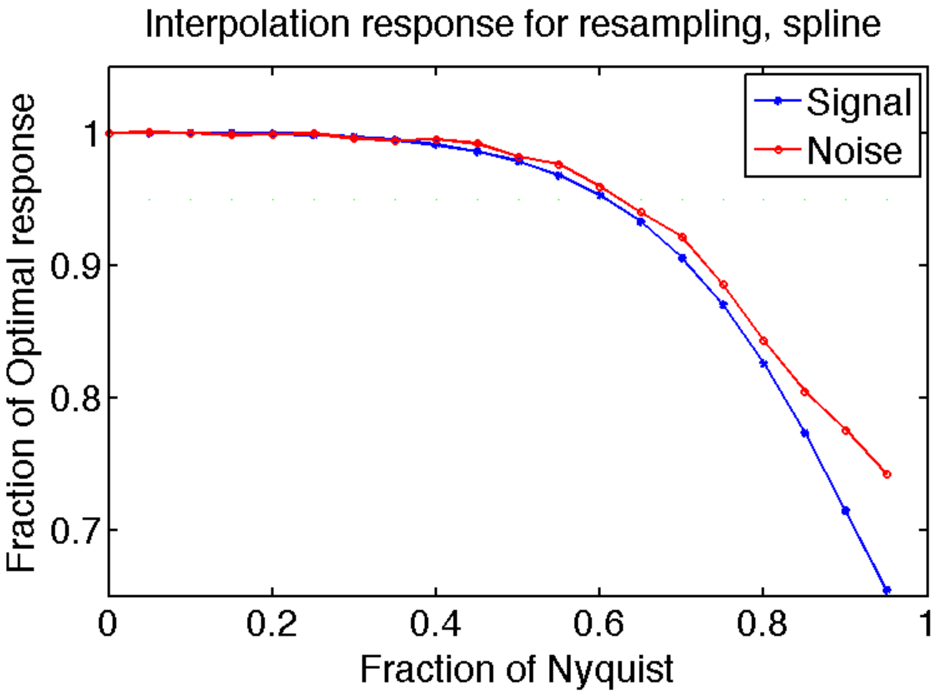

Figure 5.8: Interpolation response to signal and noise using a cubic spline interpolater.

tolerated is an external parameter, which is set usually to about 5% by the LSC's CW group. A quick glance at the zoomed figures of 5.10 and 5.11 shows that in order to maintain a maximum tolerance of 5%, only about 10% of the band can be used with linear interpolation and about 60% can be used with cubic spline interpolation. Linear interpolation requires two computations for each interpolation, while cubic spline requires three. Thus using cubic spline is much more economical, since for about 1.5 times the computational cost, the band that is usable goes up by a factor of 5. Cubic spline interpolation turns out to be the most economical of all the interpolations that were tried, most of which are not mentioned here.

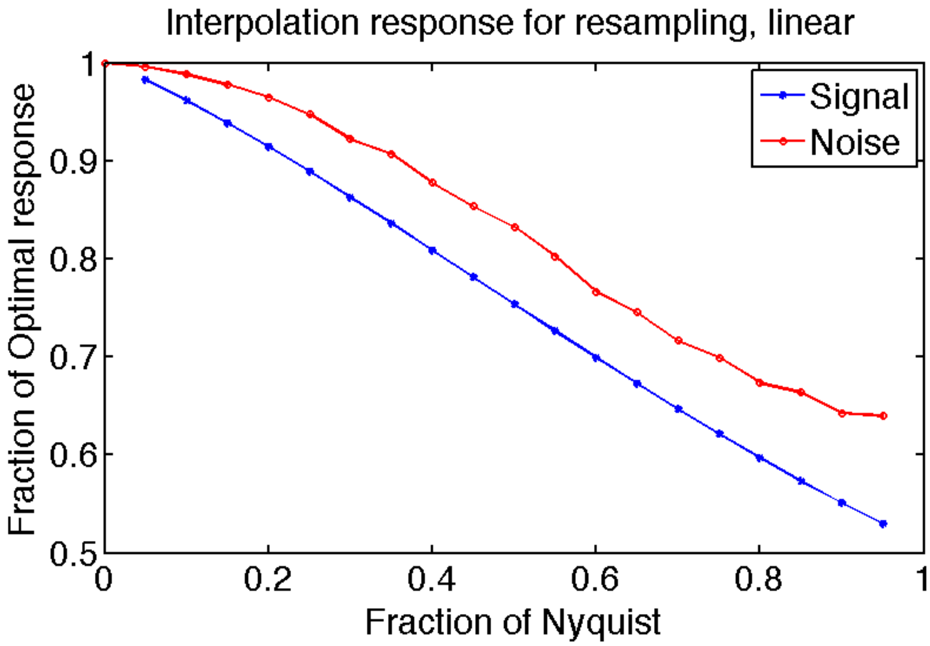

Figure 5.9: Interpolation response to signal and noise using a linear interpolater.

5.5 Summary and conclusions

In this chapter, I describe an efficient implementation of the barycentric resampling technique, which deals with the non-stationarity of the detector and calculates the $\mathcal{F}$-statistic. Although the calculation of the $\mathcal{F}$-statistic has been targeted, this technique can be used for many other kinds of searches. The major contribution of this technique is to remove the Doppler shift of the Earth's motion in a gravitational wave signal. Thus, once this Doppler shift is removed, both frequentist and Bayesian techniques can be applied to the data. In the process of implementing this algorithm, a series of practical issues are dealt with, including constraints of modern computer memory, discreteness of the data taken, losses due to interpolation, and gaps in real data.

The computational savings due to this technique can be used in various ways. One such use is to increase the coherent integration time for all-sky searches like the Einstein@Home searches. Currently Einstein@Home [88] uses a 40 hour long

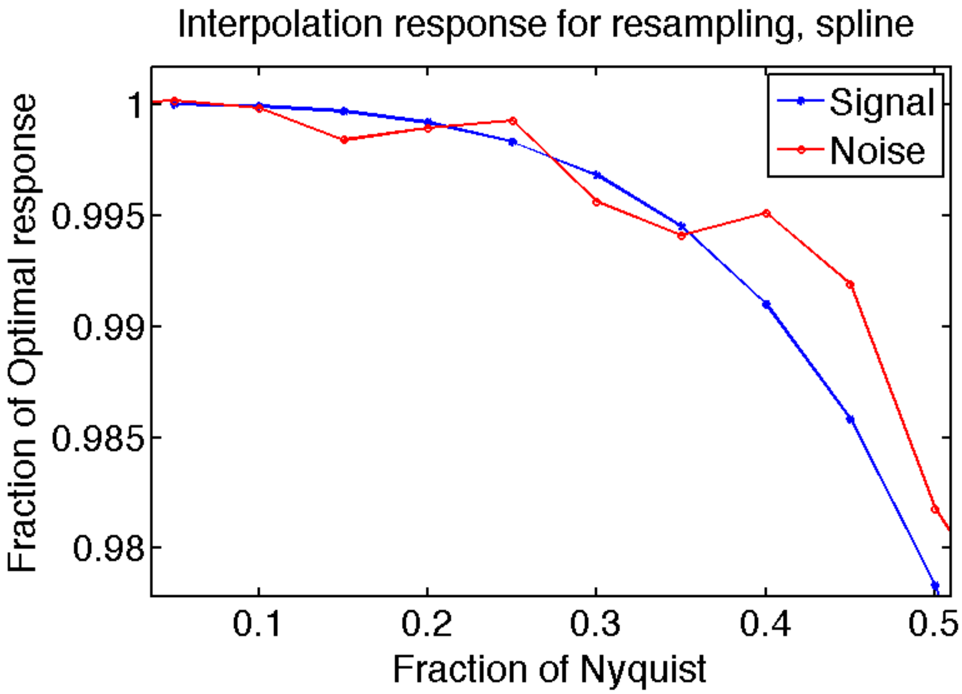

Figure 5.10: Interpolation response to signal and noise using a cubic spline interpolater. Zoomed to show the random fluctuations in the response due to noise, but a smooth response of the signal.

coherent integration time. The resampling code will be about 10 times faster for such integration times, and for the same computational power and keeping the same scaling for the search, we can coherently integrate 64 hours instead, which corresponds to a sensitivity increase of about 25%.

The resampling technique is most effective for long integration times, which are feasible for targeted searches like the search for gravitational waves from the Crab pulsar [80]. The computational savings can be used to search over wider parameter spaces like more spindown parameters or to search over binary systems.

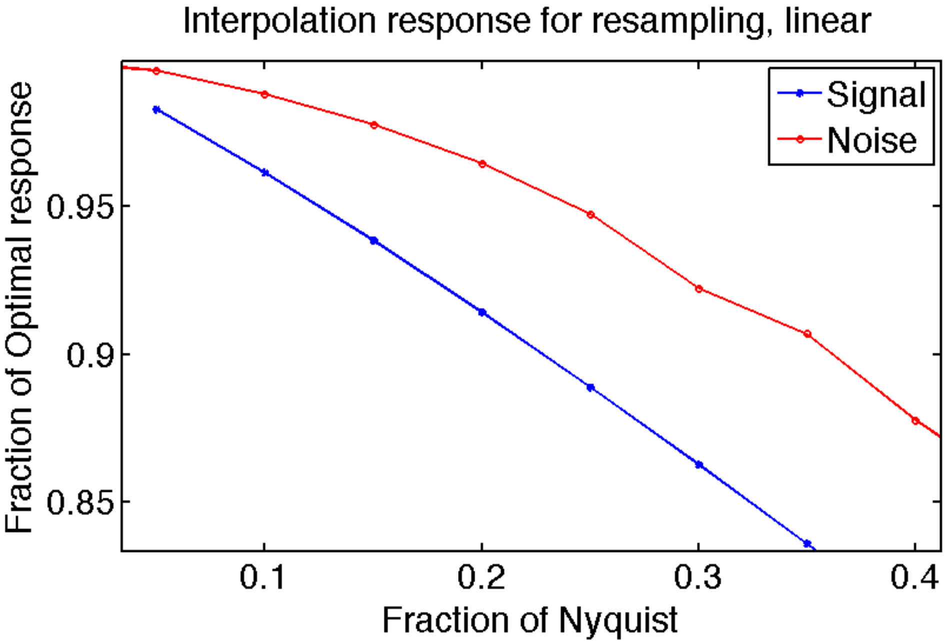

Figure 5.11: Interpolation response to signal and noise using a linear interpolater. Zoomed to show the random fluctuations in the response due to noise, but a smooth response of the signal.

Chapter 6

Introduction to Calvera

6.1 Overview

In this chapter, I introduce an object dubbed *Calvera* [116] as a promising potential source of continuous gravitational waves. This source was found in the *ROSAT* All-Sky Survey Bright Source Catalog. Follow up investigations with *Swift*, Gemini-North and *Chandra* X-ray Observatory confirmed its compact object status and gave some observational data to constrain its nature [116]. I review the the results and the conclusions of these results as done in [116].

6.1.1 Observable properties and classification of neutron stars

The quantities that can be measured for an astronomical object are its sky location, the frequency range of its emission spectrum and the incident flux as a function of frequency. If one were to then suspect thermal emission from the object, then the incident flux can be fit to a thermal spectrum and an effective temperature (T_{eff}) can be computed. This effective temperature can then be related to the luminosity (L) and an effective black-body radius (R_{bb}) of the object. Assuming that the object emitting the thermal black-body spectrum is at a temperature of T_{eff} and a sphere of radius R_{bb}, the luminosity is given by

$$L = 4\pi R_{\text{bb}}^2 \sigma T_{\text{eff}}^4, \tag{6.1}$$

where, σT_{eff}^4 comes from the Stefan-Boltzmann's law and σ is the Stefan-Boltmann constant. This can then be related to the incident flux F by

$$F = \frac{L}{4\pi d^2} = \frac{R_{\text{bb}}^2 \sigma T_{\text{eff}}^4}{d^2}, \tag{6.2}$$

where, d is the distance to the source.

The distance to some objects can be calculated using different methods like parallax measurements, measurements of distances to associated nebulae etc. If one has a distance measurement, then equation 6.2 can be used to calculate the luminosity and subsequently the black-body radius. Figure 6.1 shows a plot of the X-ray luminosity and black-body radii that were computed for different types of neutron stars. Figure 6.2 shows the same objects in a plot of R_{bb} vs kT_{eff}, where k is the Boltzmann constant. These neutron stars can be divided into four different categories from Figure 6.2. These four categories are

1. Isolated neutron stars (INSs) like the so called Magnificent Seven [121], which are thought to be young and hot and expected to be thermally powered due to core and crustal cooling.

2. Magnetars, which are neutron stars with very powerful magnetic fields. These fields can sometimes be as high as 10^{14} Gauss. They are powered by a decaying magnetic field. Thus their spectra is not thermal and they are expected to be brighter in X-rays than thermally powered INSs.

3. Central Compact Objects (CCOs), which are X-ray sources in the centers of supernova remnants (SNRs). These objects have smaller blackbody radii (R_{bb}) than INSs and Magnetars. This is because a substantial fraction of their radiant energy goes into powering the nebula, making them appear dimmer than they would otherwise.

4. Milli-second Radio Pulsar (MSPs), which are old neutron stars that have been spun up to their fast milli-second scale rotation periods. All the MSPs plotted in Figure 6.1 are found in the globular cluster 47 Tuc.

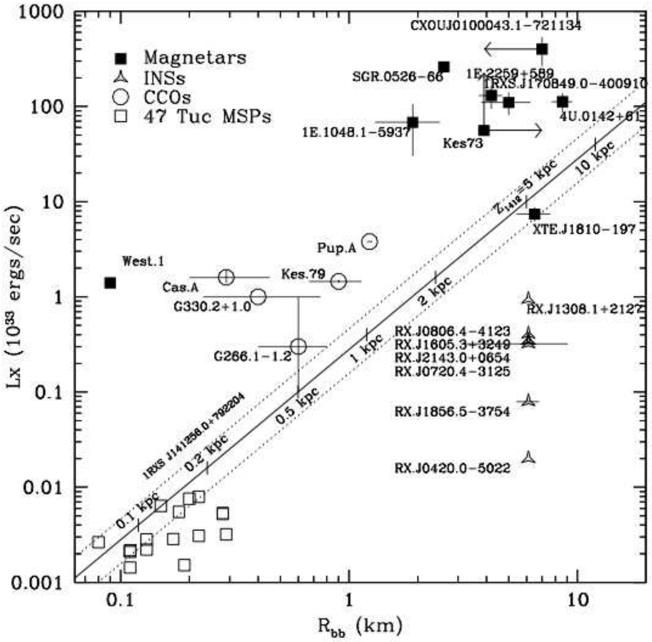

Figure 6.1: Figure reproduced with permission from [116]. Plot of blackbody radius R_{bb} and the X-Ray Luminosity L_X. Best fits for Calvera lie on the black line and the error bars are the dotted lines. All four possibilities are plotted on the same plot, so comparisons can be made with Calvera.

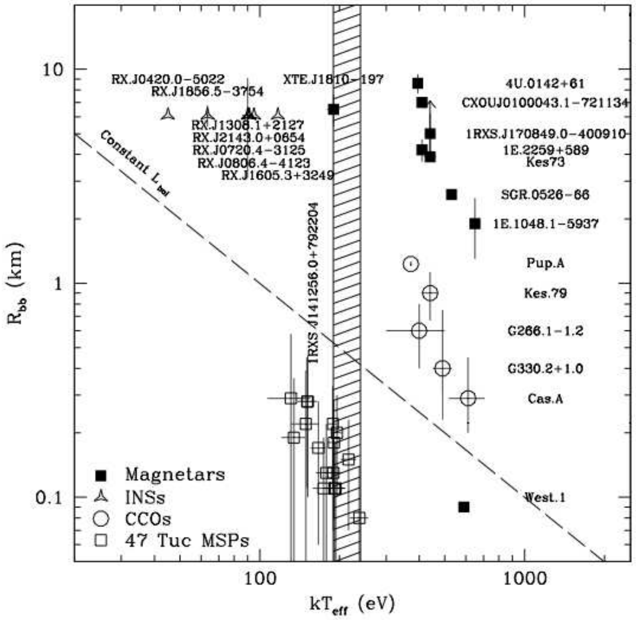

Figure 6.2: Figure reproduced with permission from [116]. Plot of effective temperature kT_{eff} vs blackbody radius R_{bb}. Best fits for Calvera lie in the hatched region. All four possibilities are plotted on the same plot, so comparisons can be made with Calvera.

6.2 Calvera

6.2.1 Observations

Calvera was first identified by Rutledge et al. [116] as a bright X-ray source from the *ROSAT* All-Sky Bright Source Catalog, as 1RXS J141256.0+792204. It was suspected to have a very large X-ray to optical flux ratio, making it an intriguing target for a follow up with other X-ray and optical instruments. Follow up observations conducted by using the *Swift* satellite, Gemini-North and the *Chandra* X-ray Observatory confirmed this suspicion. The X-ray flux in the $0.1-2.4$ keV band, $F_{\mathrm{X}}(0.1-2.4\mathrm{keV})$ was measured to be 2.5×10^{-13} erg cm^{-2} s^{-1}. No corresponding optical flux was detected. This set a lower limit on the flux ratio $\frac{F_{\mathrm{X}}(0.1-2.4\mathrm{keV})}{F_{\mathrm{optical}}} > 8700$, where F_{optical} is the optical flux. It was found at a right ascension of $14^{\mathrm{h}}12^{\mathrm{m}}55^{\mathrm{s}}.885$ and a declination of $79\,^{\circ}22'04''.10$ with an uncertainty of $0.57''$. An important fact to note about Calvera is it high galactic declination. A high delination is usually indicative of close proximity to Earth, since few distant galactic objects are found at such declinations (i.e., out of the plane of the galaxy).

Rutledge et al. obtained about 1900 seconds of *Swift* data using both its X-ray Telescope (XRT) [119] and the Ultraviolet/Optical Telescope (UVOT) [120] simultaneously looking at Calvera. Using the data obtained from the XRT, they fit the X-ray spectrum to a thermal blackbody model to an effective temperature of $kT_{\mathrm{eff}} = 215 \pm 25$eV, where k is the Boltzmann constant. The error bars are using a 90% confidence interval. The corresponding UVOT observation did not find an optical counterpart down to about 21 magnitudes. Observations done with *Chandra* and Gemini-North strengthened this hypothesis.

With its high effective temperature of 2.5 million Kelvin, Calvera is about two orders of magnitude hotter than the hottest uncollapsed stars ($\approx 6\times10^{4}$ K). We have no information of Calvera's luminosity or its distance. If one were to assume that it is close, which seems likely given its high galactic latitude, then upon using equation 6.2, we find that its R_{bb} would be small (≈ 0.1 km) compared to a main sequence star. This would make it a compact object like a neutron star.

The X-ray to Optical flux ratios of most known X-ray source classes are much smaller than that measured for Calvera. For instance stars typically have such ratios $\leq 10^{-3}$, Active Galactic Nuclei have ratios $\approx 0.1 - 10$, and white dwarfs and X-ray binaries have ratios typically $10 - 100$. The measured ratio for Calvera was ≥ 8700. This is consistent with its high effective temperature. All the objects which possess a large X-ray to Optical flux ratio and for whom their distances have been measured have blackbody radii smaller than about 10 km and thus are compact objects. Examples include INSs like 1RXS J185635.1-375433 [117], RXJ1605.3+3249 [118], etc. No counter-example exists of a non-compact object that has a high effective temperature. Thus Rutledge et al. assume that Calvera is a compact object and classify it among the various types of known compact objects.

In Figure 6.1, given Calvera's T_{eff} and for various distances d, it lies along the 45 degree black line and the uncertainty in T_{eff} shows up as the dotted lines. In Figure 6.2, Calvera lies in the hashed region.

6.2.2 Classifying Calvera

Rutledge et al. use the X-ray luminosities (L_X), the thermal blackbody radii (R_{bb}) and the effective temperatures (kT_{eff}) to distinguish these various classes and to classify Calvera. For the MSPs, CCOs and INSs, they use the thermal bolometric luminosities. For magnetar luminosities, they use only the 2-10 keV band. This is because magnetars are known to have a non-thermal spectrum and thus only the band spanning the data is used to compare the two objects.

Isolated Neutron Star

The methodology used to mark Calvera as an interesting object has been previously used to identify the so-called Magnificent Seven [121], a set of observationally homogenous X-ray dim isolated neutron stars. The spectrum of these objects is not very well understood, but they all display thermal spectra in the X-ray band, i.e. their spectrum in the X-ray band can be fit well to the emission by an isotropically radiating thermal blackbody of a certain blackbody radius R_{bb} at an effective temperature kT_{eff}. The blackbody radius R_{bb} is given by

$$R_{bb} = \frac{d}{T_{\text{eff}}^2}\sqrt{\frac{(\text{B.C})F_X}{\sigma}}, \tag{6.3}$$

where d is the distance to the source, B.C is the bolometric correction, F_X is the X-ray flux in the 0.1-2.4 keV band and σ is the Stefan-Boltzmann constant. Some of the Magnificent Seven are close enough to Earth for a parallax measurement to be possible, for example the distance of 1RXS J185535.1-375433 [122] was measured to be 167^{+18}_{-12} pc.

If we interpret Calvera as an INS, extrapolating from the Magnificent Seven and using Calvera's effective temperature, then from Figure 6.1, it would lie at a distance of 8.4 kpc from the Earth and have a z of 5.1 kpc from the midplane of the galactic disk, with a Galactocentric distance of 14.0 kpc. This places it in the Galactic halo. There are no known isolated neutron stars in the Galactic halo, although there is the possibility that some radio pulsars populate the halo [123]. Due to the high kT_{eff}(225 ± 25 eV) of Calvera, under this interpretation its X-ray luminosity L_X would be 1.0×10^{34} erg s^{-1}, which is an order of magnitude greater than the next most luminous isolated neutron star.

Since in this interpretation, the source lies in the galactic halo, it is unlikely to be powered by accretion from the interstellar medium, which is very sparse in the halo. If it is powered by the remnant heat of a supernova, then one of the following scenarios have to be true. If it was created in the Galactic plane (as seems likely), then it would receive a kick during the supernova explosion and attain a kick velocity whose z-component could be v_z. Under this scenario, it would travel a distance $z = v_z\tau$ in the time τ from the galactic plane. If we assume standard cooling curves [124], it would require a cooling time $\tau_c < 10^6$ years. In order for $\tau <= \tau_c$, the kick velocity $v_z > 5100$ km s^{-1}, which is an order of magnitude greater than the average kick velocity (380 km s^{-1}) observed [125]. On the other hand, if we assume the average kick velocity for Calvera, then it would require

13 million years to get to its calculated position in this interpretation, which is disallowed by standard cooling curves.

Thus there is a conundrum here, which strongly disfavors this interpretation for Calvera.

Magnetar

A magnetar is a neutron star with an extremely powerful magnetic field which can range from 10^{12} to 10^{14} Gauss. Their electromagnetic emissions are usually powered by the decay of this magnetic field. Work done recently has shown that magnetar X-ray spectra can be described as a soft blackbody with a power law dominating at higher energies. If one were to assume that these magnetars are standard-candles (as is suspected in [126, 127]) and assume the corresponding standard-candle luminosity as calculated in [127], it puts Calvera at a distance of 66 kpc from Earth and about 40 kpc above the Galactic disk. This is an even more extreme scenario than the one for X-ray dim isolated neutron stars, making this an even more unfavorable hypothesis.

On the other hand, if one were to ignore current evidence and assume that the X-ray luminosity of magnetars were a free parameter, then we can place Calvera within 0.1 kpc of the Galactic disk, since all known Magnetars are located within 0.1 kpc of the Galactic disk [127, 128]. This would place Calvera at a distance of around 0.17 kpc from the Earth and its X-ray luminosity would then be $L_X < 8.7 \times 10^{29}$ erg s^{-1}, which would be about 10^5 times less than the standard candle luminosity [127], or about 10^3 times less than the faintest known Magnetar [128]. This result also makes the hypothesis of Calvera being a Magnetar very unlikely.

Central Compact Object

A central compact object (CCO) is a point-like X-ray source associated with a supernova remnant, which has similar spectral properties [129, 130]. This category includes the strong X-ray source Cas A. Since the defining characteristic of this category is the presence of a supernova remnant, Calvera is ruled out by the fact that there is no supernova remnant observed in the optical band within 2 degrees of it. Rutledge et al. explore the possibility that it could be the first unhosted CCO, which would have similar spectral properties to the known CCOs. Upon analysis, it is found that if it were an average CCO, the luminosity of Calvera would be 10 times lower than the faintest known CCO.

Millisecond Radio Pulsar

The final class of objects that Rutledge et al. compare Calvera to are millisecond radio pulsars (MSPs). These are old spun-up neutron stars, which spin with periods of a few milliseconds. These are usually found in globular clusters, like 47 Tuc [131]. Comparing Calvera to the homogenous survey of these MSPs in 47 Tuc [131], it is found that Calvera has a similar effective temperature (T_{eff}) and for distances ranging from 80 pc to 260 pc, it has a similar thermal bolometric luminosity and blackbody radius (R_{bb}). This makes MSPs a favored hypothesis. If it were a MSP, it would not be the first to have been discovered from an X-ray sample. The first and only other object discovered in such a manner can be found in [132].

6.2.3 Properties of Calvera

As discussed above, Calvera is most likely to be a millisecond radio pulsar type object. It is suspected to be at a distance of 80-260 Pc. This would make it one of the closest MSP like objects to us. It is also expected to be old, since most millisecond pulsars are older than a few million years [109]. However given its alleged proximity to us, Calvera is a very promising source of continuous gravitational waves. It is also likely to be spinning fast, since MSPs typically have spin frequencies in the range of 50-700 Hz, which places it in the LIGO band. The sky location of Calvera is known to sub arcsecond precision, which is sufficient resolution for a targeted search to require only one sky location template. Thus Calvera is a good candidate for a broadband CW search and its properties make such a search possible. A search done for GWs from Calvera is described in chapter 7.

Chapter 7

Search for GWs from Calvera

7.1 Introduction

This chapter documents a search done using the resampling algorithm for a suspected milli-second pulsar type object. This object dubbed *Calvera* [116] is suspected to be as close as $80 - 260$ parsecs from Earth. I present the selection of the parameter space of the search. A description of the search pipeline is then followed by the interpretation of the results. An overview of a followup scheme is mentioned last.

7.1.1 Possibility of a search

As discussed in the previous chapter, there is a good chance that Calvera is a millisecond radio pulsar that is very close to Earth (80-260 pc), possibly with a spin period of a few milliseconds (which has not been observed as such, presumably because its radio beam does not sweep past the Earth). MSPs have typical spin frequencies ranging from 10 Hz to a few 100 Hz, putting it right in the LIGO band, which spans from 40 Hz to about 2000 Hz.

The gravitational wave luminosity is bounded by the time derivative of the total rotational kinetic energy

$$\left(\frac{dE}{dt}\right)_{\rm gw} = \frac{32G}{5c^5} I_{\rm zz}^2 \epsilon^2 \left(\pi f\right)^6 \le - \left(\frac{dE}{dt}\right)_{\rm rot} = -\frac{d}{dt}\left(\frac{\pi^2 I_{\rm zz} f^2}{2}\right), \tag{7.1}$$

where ϵ is the equatorial ellipticity of the star, I_{zz} is the principle moment of inertia of the star and f is the gravitational wave frequency, G is the gravitational constant and c, the speed of light. Solving Equation 7.1 for the ellipticity gives us

$$\epsilon \leq \sqrt{\frac{5c^5}{32\pi^4 G I_{\mathrm{zz}}} \frac{-\dot{f}}{f^5}}. \tag{7.2}$$

The spinning down star can be modelled with a frequency derivative, which is dependent on a power law of the frequency of rotation. i.e. $\dot{f} \propto f^n$, where n is called the braking index of the star [133, 135] (see Section 3.3.1). A characteristic age τ_c can also be defined by using only the frequency, frequency derivatives and a braking index as follows

$$\tau_c = \frac{1}{n-1}\left(\frac{f}{-\dot{f}}\right). \tag{7.3}$$

This characteristic age τ_c is a very good approximation for the actual age of the pulsar.

Using equation 7.3 in 7.2, we get an ellipticity independent of $\dot{f}$ and as a function of the age τ,

$$\epsilon_{\mathrm{age}} \leq \sqrt{\frac{5c^5}{32\pi^4 G I_{\mathrm{zz}} \left(n-1\right) \tau_c f^4}}. \tag{7.4}$$

In order to convert these ellipticities to measurable gravitational wave strain amplitude h_0, we use the following equation,

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{\mathrm{zz}} \epsilon f^2}{D}, \tag{7.5}$$

where D is the distance to the source [110]. Equations 7.5 and 7.4 give us an age based GW strain amplitude,

$$h_{\mathrm{age}} \leq \frac{1}{D} \sqrt{\frac{5 G I_{\mathrm{zz}}}{2c^3 \tau_c \left(n-1\right)}}, \tag{7.6}$$

Notice that the age based strain amplitude h_{age} is independent of the frequency of the gravitational wave and inversely proportional to the distance of the source. It is also inversely proportional to $\sqrt{\tau_c}$, the age of the source.

Using S5 data with the two 4 km detectors H1 and L1, the minimum gravitational wave strain that can be detected at the frequency of 150 Hz or so, where the detectors are most sensitive is around 10^{-25}. So a source that is 1000 years old and at a distance of 1 kpc would be detectable. So would a source at 100 pc even if it were 10^5 years old. Millisecond pulsars are usually quite old (10 million years or so) [109] and thus usually undetectable to LIGO. But Calvera is suspected to be anamoulously close to us (80-260 pc). Thus it would be detectable if it were around 10^7 years old. This is a decent assumption for the age of a typical MSP [109].

Plugging in the numbers for Calvera into equation 7.6, including the most optimistic distance of 80 pc; age of 10^7 years; braking index n of 7 and the moment of intertia of a typical neutron star, with a radius of 10 km and a mass of $1.4M_\odot$, 10^{38} kg m^2; the age based spindown limit h_{age} comes out to about 2.83×10^{-25}. This is a limit that can be beaten for a band of about 300 Hz centered around 200 Hz by using all of LIGO's S5 data and integrating coherently over 2 years of calendar time (1 year of live time) with data from 2 detectors. Thus a search, while not being exhaustive, is possible.

7.1.2 Parameter Selection

The number of templates in a coherent search can increase prohibitively with the integration time. For a search which has a single sky location, multiple frequencies, the first two frequency derivatives and spanning a time T, the number of templates scales as T^3. Since for a coherent search, the sensitivity of a search scales as $\sqrt{T}$, the sensitivity of the search scales as $C^{\frac{1}{6}}$ for given computational power C. Thus it is sensible to cap the number of templates searched, such that the search can be completed in a given time frame.

For the Calvera search, the available computational power was set to 1 week of operation on 300 Intel Core2 Duo processors. The parameters of the search were set using an iterative algorithm, in which an integration time and an age τ is set and then the computational cost is calculated. If it is above the allocated computational cost, then the age is lowered and the computational cost recalculated. If changing the age is not good enough or crosses certain bounds, then the integration time is lowered.

Frequency and spindown templates

The frequency and spindown parameters are calculated using a procedure very similar to the one used in [136]. Frequency derivative ranges are calculated based on braking indices ranging from 2 to 7. This range covers all known pulsars, except for the Vela pulsar, which is too young (Calvera is assumed to be much older) and interacts with its wind nebula (not observed for Calvera), including the static dipole (n=3), quadrupole (n=5) and the saturated r-mode (n=7) [137]. The spindown upper limit (h_{age}) is calculated using a braking index of 5, which is the braking index if the emission from the star is dominated by gravitational waves from a constant mass quadrupole. This is the upper limit we use to set the parameters of the search, as it is the absolute maximum of emission of gravitational waves. It is pointless to search in regions of the parameter space where the noise strength is above this GW spindown upper limit.

The search parameter space must still search over the whole range of braking indices and this is reflected in the choice of frequency derivatives. The frequency range can be chosen from Figure 7.1 to be from 80 Hz to 360 Hz. The spindown parameters derived from equation 7.3 and for braking indices ranging from 2 to 7 are as follows [136]

$$\frac{f}{6\tau_c} \leq -\dot{f} \leq \frac{f}{\tau_c}, \tag{7.7}$$

where, we take $\tau_c = 1 \times 10^7$ years and

$$\frac{2\dot{f}^2}{f} \leq \ddot{f} \leq \frac{7\dot{f}^2}{f}. \tag{7.8}$$

In practice, the frequency parameter space is divided up in to small chunks of

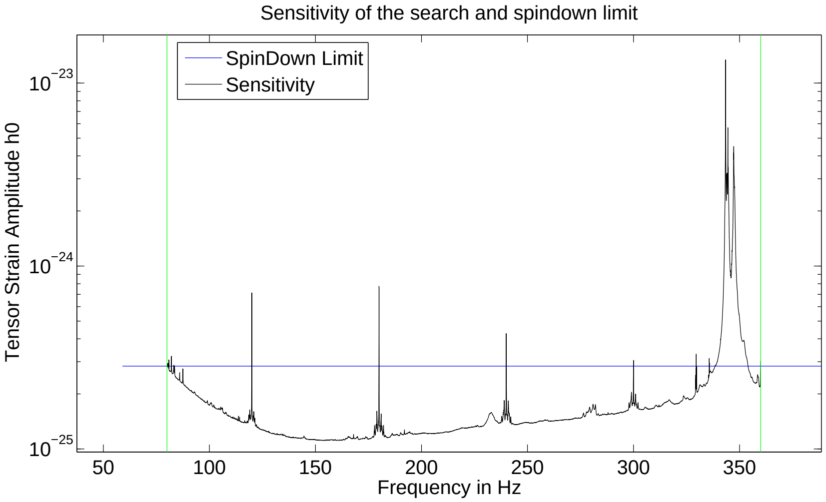

Figure 7.1: Choice of Parameter space, comparing the spin down upper limit from a star which is at 80 pc and 10 million years old with the LIGO science curve, integrated with 1.4 years of H1 data and 1.1 years of L1 data.

about 0.025 Hz and the upper and lower bounds of the frequency derivatives are calculated for each band using equations 7.7 and 7.8. This sort of partitioning is done to make the distribution of the computational processes among the various nodes easier. The number of templates required with a mismatch parameter of 0.15 as a function of frequency band can be seen in Figure 7.2. The choice of the value of the mismatch parameter is in accordance with previous LIGO searches.

Sky position templates

Given the computational limitations on the search for Calvera, it is best if only one sky position template be required to span the parameter space. This turns out to be the case and the reasoning is explained below.

The uncertainty radius in the position of Calvera is $0.57''$ [116]. Other results like reference [134] reduce this uncertainty radius down to about $0.33''$. The metric for sky position was shown in Chapter 4. If θ is the angle between the Earth's velocity vector and the vector point from the Earth to the source, then the frequency

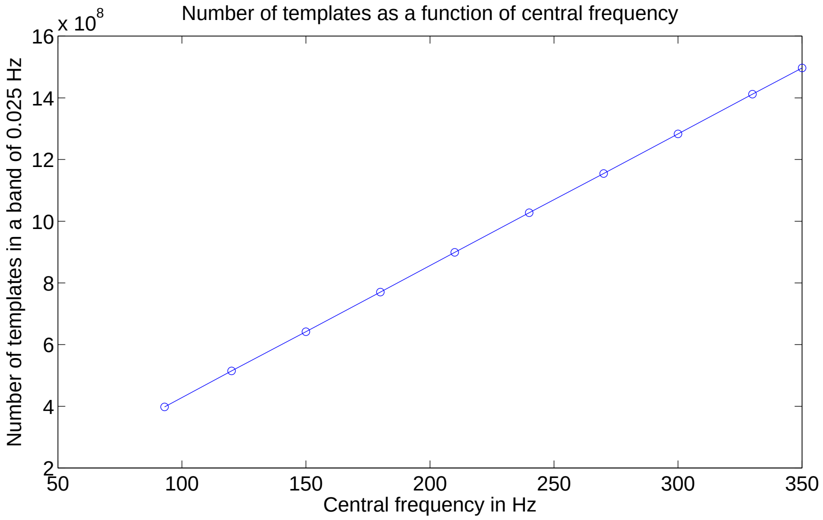

Figure 7.2: Number of templates (distinct values of $f, \dot{f}, \ddot{f}$ and one sky pixel in right ascension and declination) required as a function of frequency band (0.025 Hz) to span the parameter space.

difference caused by a deviation from this perfect angle is given by

$$\Delta f = \frac{|v|\sin(\theta)\Delta\theta}{c}\left(f_0 + \dot{f}T_{\rm obs} + ...\right), \tag{7.9}$$

where $|v|$ is the magnitude of the velocity of the Earth and f_0 and $\dot{f}$ are the frequency and spindowns respectively. $T_{\rm obs}$ is the observation time and the coherent integration time, since the search for Calvera is a fully coherent search. We can safely ignore the $\dot{f}T_{\rm obs}$ term, as $f_0 >> \dot{f}T_{\rm obs}$. If we then approximate Δf to about $\frac{1}{T_{\rm obs}}$ and take a time average of the velocity of the Earth over the span of the search, then the minimum spacing required in the angle θ is

$$\Delta\theta_{\rm min} = \frac{2c}{v_{\rm max}}\frac{1}{f_{\rm max}T_{\rm obs}}. \tag{7.10}$$

Using the numbers for the Calvera search, $f_{\rm max}$ is 360 Hz; $T_{\rm obs}$ is about 6.0×10^7 seconds and given its high declination, the $v_{\rm max}$ is $3.5 \times 10^{-5}c$ and thus the minimum angular resolution required is about $0.54''$. This does not cover the complete uncertainty radius (however it does cover [134]). It is helpful to notice that equation 7.10 is an approximation of the actual angular resolution required. Given that the search is already pushing the computational power available to its limits, a slight compromise in angular resolution is acceptable (loss of SNR is less than 5%) and only one sky location is sufficient. Moreover this is only the case for the higher frequencies.

7.1.3 Search pipeline

The search for Calvera was done on a computing cluster known as ATLAS at the Albert Einstein Institute (AEI), in Hannover, Germany. It is a beowulf cluster (a loosely networked cluster of many inexpensive linux processors), consisting of 7000 processors, split into quad-core machines, with each machine possessing a total of 8 GB of memory. This meant that each processor could use up a maximum of 2 GB of memory. Bands larger than about 0.025 Hz cannot be processed without running into this memory constraint.

The search was thus split into 11200 jobs spanning the frequency parameter space of 280 Hz from 80 Hz to 360 Hz in equal bands of 0.025 Hz. In each job approximately 6×10^7 seconds of data from H1 and L1 was filtered through approximately 1×10^9 templates as shown in Figure 7.2, producing about 1×10^9 values of $2\mathcal{F}$ for that frequency band. These 11200 jobs were submitted as parallel jobs by using the batch queueing system called Condor [138]. Each job is given separate command line arguments, which tell the job where it can find the SFTs it needs to load, the output directory, the parameters that it needs to search over, etc. These arguments are specified in a condor dag ("directed acyclic graph" job dependency) file, which is then submitted to the cluster. The total computing time required for completing all the jobs was about 6 cpu-years. Since on average, about a 1000 machines were used, about 2 days were sufficient to complete the search.

The search pipeline consists of the following scripts and files -

1. **Search configuration file** - A python file, consisting of the search param-

eters for the entire search. It includes items like the right ascension and declination of the source, the frequencies to search over, the braking indices to search over, the location of input SFTs, the output location, condor parameters and all the other parameters of the search that are not specific to any given job.

2. **DAG generation file** - A python script, that generates a condor dag file with the input of the configuration file. This script also generates a condor submit file, which specifies the output parameters and other internal condor parameters. This script splits the jobs into frequency bands and specifies the spindowns and input SFT locations specific to each job.
3. **Resampling $\mathcal{F}$-statistic code [114]** - Code written in C. This is the program that calculates the $\mathcal{F}$-statistic using the resampling algorithm. For more details, see Chapter 5.
4. **FStat Output** - The output of the resampling code. This contains usually the top 100 (or whatever number is specified) candidates with the highest $\mathcal{F}$-statistic value. For each candidate, its frequency, sky location and spindowns are specified. An Fstat Output file is produced by each job.
5. **Histogram Output** - Another output of the resampling code. This contains the histogram of the $\mathcal{F}$-statistic output for each job.
6. **Kolmogorov-Smirnov test statistic calculator** - A python script which takes in the histogram output files as input and generates a Kolmogorov-Smirnov(KS) test statistic for each band. This is then used as a veto against wandering lines. It is explained in detail in the vetoes section.
7. **Upper Limit calculation script** - A matlab script that takes in the FStat Output file and uses the loudest $\mathcal{F}$-statistic value and the noise of the detector to calculate an upper limit.

The search pipeline consists of the use of all these scripts in the order that they have been presented. The scripts are located on the ATLAS cluster in the `/home/ppatel/CalveraSearch/Scripts` directory.

An overview of all the important search parameters can be found in Table 7.1.

Detectors	H1 & L1 (Coherently combined)
Science Run	S5
Start Date	Nov 7th 2005
End Date	Oct 01 2007
Total Observation Time	6.0×10^{7} seconds
H1 SFTs	23787
L1 SFTs	19197
H1 Live Time	4.28×10^{7} seconds
L1 Live Time	3.46×10^{7} seconds
Minimum Frequency (f)	80.0 Hz
Maximum Frequency (f)	360.0 Hz
Minimum First Spindown ($\dot{f}$)	-1.1×10^{-12} Hz s^{-1}
Maximum First Spindown ($\dot{f}$)	-4.3×10^{-14} Hz s^{-1}
Minimum Second Spindown ($\ddot{f}$)	9.9×10^{-30} Hz s^{-2}
Maximum Second Spindown ($\ddot{f}$)	1.1×10^{-25} Hz s^{-2}
Search Frequency Band Size	$\frac{1}{40}$ Hz
Mismatch Parameter	15%
Total Frequency Bands	11200
Total Sky Position Templates	1
Total Templates	1×10^{13}

Table 7.1: Calvera Search Parameters.

7.2 Probabilities and False Alarm Rates

While the Calvera search uses the $\mathcal{F}$-statistic, it is more convenient to use $2\mathcal{F}$, since it has a distribution function that is well known and easy to manipulate and compute. For Gaussian noise of $\mu = 0$ and a $\sigma = 1$, $2\mathcal{F}$ is distributed as a χ^2 distribution of 4 degrees of freedom [110]. If we refer to $2\mathcal{F}$ as x from now onwards, the probability distribution function in the absence of signal is

$$p(x) = \frac{x}{4}e^{-x/2}. \tag{7.11}$$

In the absence of signal, the expectation value of x would be 4. In the presence of signal, this expectation value becomes $4+\rho^2$, where ρ is the optimal signal to noise ratio for that given signal strength [110]. It can be shown that in the presence of signal, x becomes a χ^2 distribution with a non-centrality parameter ρ^2 [110] -

$$p(x) = \frac{1}{2}e^{-(x+\rho^2)/2}\sqrt{\frac{x}{\rho^2}}I_1\left(\sqrt{\rho^2 x}\right), \tag{7.12}$$

where I_1 are the modified Bessel functions of the first kind of order 1.

The Calvera search consists of over 10^{13} computations of the $\mathcal{F}$-statistic for different templates. The result of each such computation is drawn from equation 7.11 under the assumption that no signal is present. Since there are so many instances of the $\mathcal{F}$-statistic calculation, there is a possibility of noise rising to a level loud enough to be confused with signal. We can calculate such a threshold at which the probability of such an occurrence is lower than the desired false alarm probability. The probability that a single calculation will be lower than a certain threshold x_t is given by

$$P(x <= x_t) = \int_0^{x_t} \frac{x}{4}e^{-x/2} = 1 - e^{-x_t/2}\left(\frac{x_t}{2}+1\right). \tag{7.13}$$

For N independent templates, the probability that the loudest $\mathcal{F}$-statistic value will be lower than the threshold of x_t would be given by

$$P_N(x <= x_t) = \left(1 - e^{-x_t/2}\left(\frac{x_t}{2} + 1\right)\right)^N \tag{7.14}$$

For a false alarm rate of 1%, i.e., $P_N = 1 - 0.01 = 0.99$ and a single template search, the threshold x_t would be 13.3. For the Calvera search with $N = 10^{13}$, the threshold is about 76.5. This means that if any template in the whole search has a value of $2\mathcal{F}$ larger than 76.5, the probability of that being due to the noise is less than 1%.

Another important statistic is the expected loudest event for a multi-template search. This can be computed by calculating the probability that x is exactly equal to some maximum x_m. It is given by picking one template to be exactly x_m and impose that all other templates have $x < x_m$. Such an arrangement can happen in N ways if there are N templates in total. Therefore the probability density function is

$$p(x = x_m) = N\frac{x_m}{4}e^{-x_m/2}\left[1 - e^{-x_m/2}\left(\frac{x_m}{2} + 1\right)\right]^{N-1}. \tag{7.15}$$

For a given N, the expectation value of the loudest event would be given by a simple integration of the probability density function as

$$E[x_m] = \int_0^\infty p(x = x_m)x_m dx_m. \tag{7.16}$$

An estimate of the error in this expectation value can be computed from the second moment as follows

$$\sigma_{x_m} = \sqrt{E[x_m^2] - (E[x_m])^2}, \tag{7.17}$$

where,

$$E[x_m^2] = \int_0^\infty p(x = x_m)x_m^2 dx_m. \tag{7.18}$$

As expected, the expectation value for a single template search would be 4 and for the Calvera search, it would be 67.1 ± 2.1. This serves as a sanity check for the search. All of these values were calculated before doing the actual search or looking at the results. Only templates whose $2\mathcal{F}$ values were above the nominal threshold of 76.5 with a false alarm rate less than 1% would be followed up.

7.3 Vetoes

One of the major sources of noise for CW searches are spectral lines. Some of these lines which are sometimes orders of magnitude larger than the noise floor can also wander, i.e. they change their frequency as a function of time. These spectral features, also known as lines in the data can give anamoulously large $\mathcal{F}$-statistic values, which can possibly be confused for a signal.

The resampling algorithm uses a spectral running median estimator [139] to suppress sharp spectral features. This running median estimator is an effective way of cutting out line features that are extremely narrow. In the resampling code, we run the running median estimator on 50 bins of 1800s long SFTs, which corresponds to a bandwidth of 27.8mHz. Any lines that are wider than 27.8mHz would not be vetoed by this method, since they would appear as an elevation in noise rather than as a spectral feature.

7.3.1 Known Lines

There are many spectral lines in LIGO data that are known to be caused by certain known physical phenomenon. Examples of this include the 60Hz line and its harmonics at 120Hz,180Hz and so on, the violin modes of the suspension upon which the mirrors hang etc. These lines are “known” because they have been studied for a long time now and causal connections between them have been found or at the very least, an extremely strong correlation has been noted.

One method for finding these lines has been to look for correlations between the gravitational wave channel of the detector and other auxiliary channels, like for example, the input power or magnetometers channels for the 60Hz harmonics. A list of such lines has been compiled by the LSC CW group with the help of

Robert Schofield, Nelson Christensen and Keith Thorne. Table 7.2 notes all the lines that are within the band of interest for the Calvera search ($80-360$ Hz). The table contains the central frequency of the line, along with the lower and upper frequencies that are affected by it, the reason for identifying them as lines and the detector from which the data was taken. Among the physical causes that caused these persistant lines, besides the ones that were mentioned before, include voltage supplies and optical levers.

Any anomalously large $\mathcal{F}$-statistic value in the presence of these lines can be safely vetoed as an instrumental artifact. Marginal cases can exist, where a large $\mathcal{F}$-statistic value is seen near a very weak line, which could warrant further investigation. A more quantitative veto for these artifacts is presented in the next section.

7.3.2 Kolmogorov-Smirnov Test Statistic

As mentioned previously, in the presence of normalized Gaussian noise, $2\mathcal{F}$ is distributed as a χ^2 variable with 4 degrees of freedom. For the Calvera search, the band size of 0.025Hz is small enough to ensure that in the absences of any instrumental artifacts, the noise spectrum would look flat, while restricting the search to that particular band. This data is then normalized as shown in Chapter 5. Thus one would expect that in the absence of any lines, the histogram of the $\mathcal{F}$-statistic values would follow a χ^2 distribution on band by band basis.

In the presence of a large instrumental line, like a 60Hz harmonic, the input spectrum would no longer look flat and would possibly have a shoulder. This would distort the output of the resampling code and alter the distribution from the expected χ^2 of 4 degrees of freedom. Thus deviation from the expected distribution can be used as a veto for instrumental lines. Wandering lines would also distort the distribution and can be vetoed relatively easily. Figures 7.4 and 7.3 show the differences between the distributions of two bands, one with a known line present

Frequency(Hz)	**Lower**	**Upper**	**Cause**	**Detector**
85.80	85.79	85.81	LVEA Rack Magnetometer	H1
89.90	89.84	89.96	Auxilliary Power Supply	H1
93.05	93.04	98.06	Beam Splitter Mic and Magnetometer	H1
119.99	119.74	120.24	60Hz harmonic	H1
139.95	139.94	139.96	+15V Supply Ripple	H1
180.0	179.95	180.05	60Hz harmonic	H1
329.59	329.50	329.87	Beam Splitter Violin Modes	H1
343-344	343.27	344.8	Test Mass Violin Modes	H1
347.5	347.1	347.7	Test Mass Violin Modes	H1
93.29	93.27	93.3	Optical Lever A	L1
96.70	96.696	96.72	Optical Lever B	L1
119.97	119.73	120.01	60Hz harmonic	L1
139.938	139.92	139.958	Optical Lever A	L1
145.062	145.047	145.078	Optical Lever B	L1
180.05	179.97	180.08	60Hz harmonic	L1
186.587	186.565	186.61	Optical Lever A	L1
193.416	193.395	193.437	Optical Lever B	L1
233.231	233.185	233.277	Optical Lever A	L1
241.777	241.713	241.842	Optical Lever B	L1
343-344	342.94	344.4	Test Mass Violin Modes	L1
346-347	346.6	347.02	Test Mass Violin Modes	L1

Table 7.2: Known spectral lines that were persistent throughout S5.

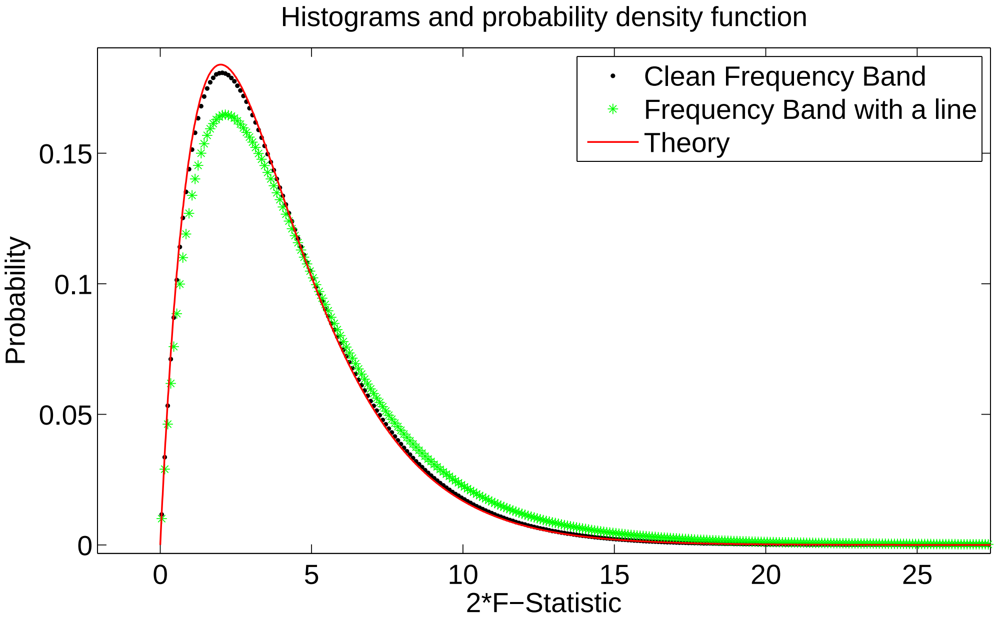

Figure 7.3: The probability distribution functions of two bands, one containing a known line and another with no known lines in it. The theoretical curve is the χ^2 probability distribution function with 4 degrees of freedom. The clean frequency band is from 93.0-93.025 Hz and the frequency band with a line is from 119.975-120.0 Hz. About 10^9 templates were used to calculate the histograms.

and another without any known lines. The data in both figures is from real data (both H1 and L1 detectors).

An obvious question about this approach is the danger that we might veto a real signal with this procedure. A real signal coming from the right sky location and with spindowns matching will only affect a few templates around the "right" or closest template. In the presence of a spurious signal, say at the wrong sky location, one would expect that more templates are rung up. However these spurious signals are noise themselves, since we are only interested in looking for signals coming from a single sky location in this search. Thus vetoing other signals like hardware injections at the wrong sky location is a positive feature of this method. Monte Carlo simulations were conducted to see if a real signal is vetoed by this method

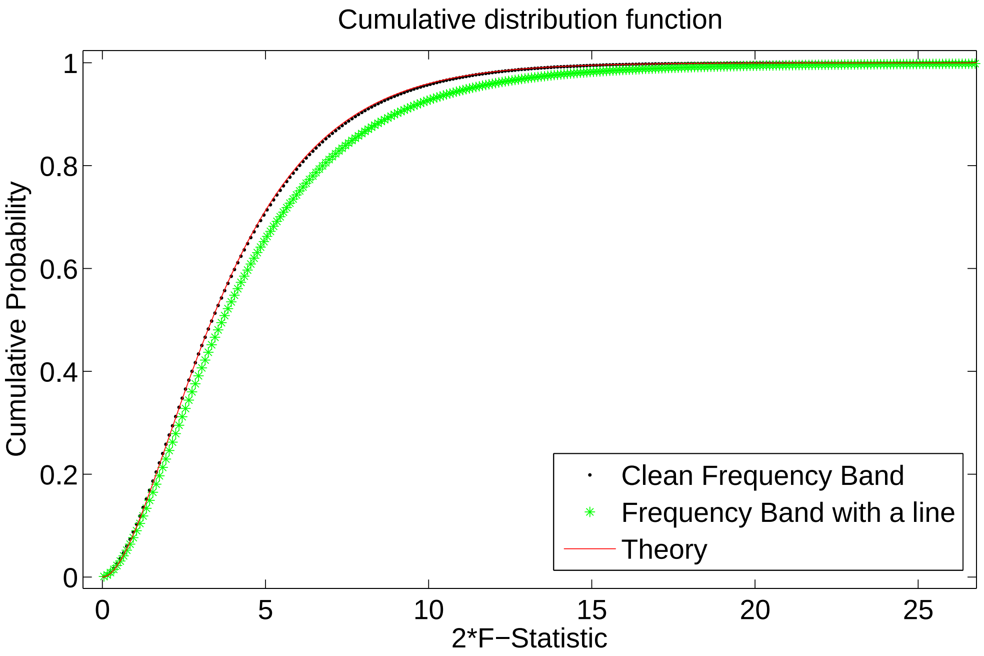

Figure 7.4: The cumulative distribution functions of two bands, one containing a known line and another with no known lines in it. The theoretical curve is the cumulative distribution function of a χ^2 probability distribution function with 4 degrees of freedom. The clean frequency band is from 93.0-93.025 Hz and the frequency band with a line is from 119.975-120.0 Hz. About 10^9 templates were used to calculate the histograms.

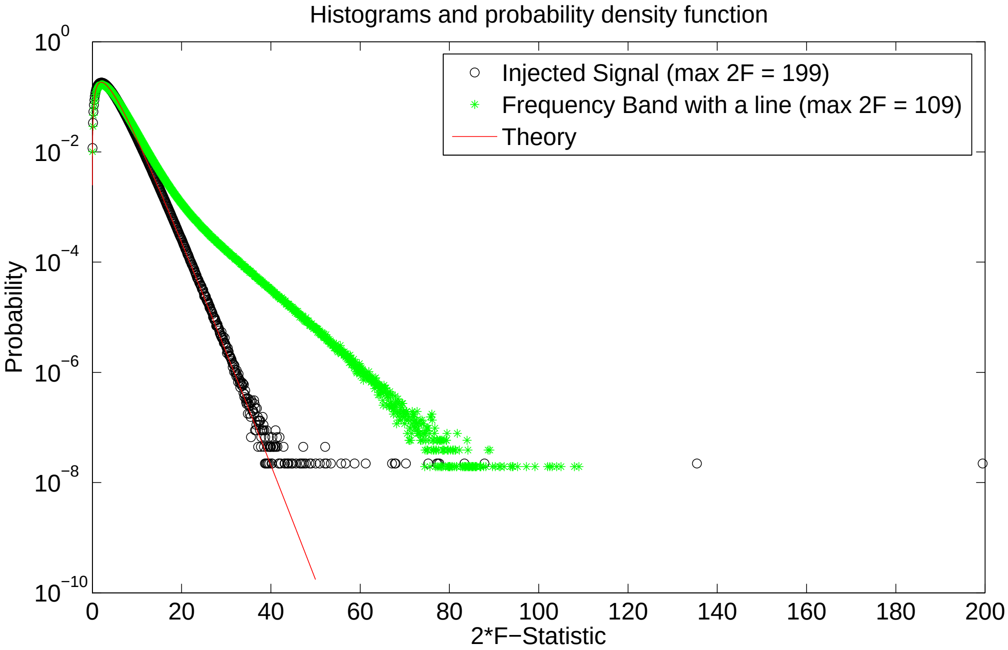

Figure 7.5: A comparison between the effects of an injected signal in a band of real data with no known instrumental artifacts (93.0-93.025 Hz) and a band with a line on the distribution of $2\mathcal{F}$. The y axis is plotted on a log scale and the signal amplitude was chosen to be louder than the loudest event due to the presence of the line.

or not and signal ranging from SNRs of 1 to 100 were used. None of these signals were vetoed. A comparison between the effect of a signal on the distribution of $2\mathcal{F}$ and the effect of a line, can be seen in Figures 7.5 and 7.6.

A convenient method of distinguishing between two distributions is the Kolmogorov-Smirnov (KS) test statistic. This statistic is the maximum vertical distance between two cumulative distribution functions. One can calculate the probability that two distributions are different if the total number of samples and the statistic itself is taken into account. However for the purposes of vetoing, such a probability calculation is unnecessary and just the KS statistic is enough. Monte-Carlo simulations with Gaussian noise with no line present have shown that the KS statistic

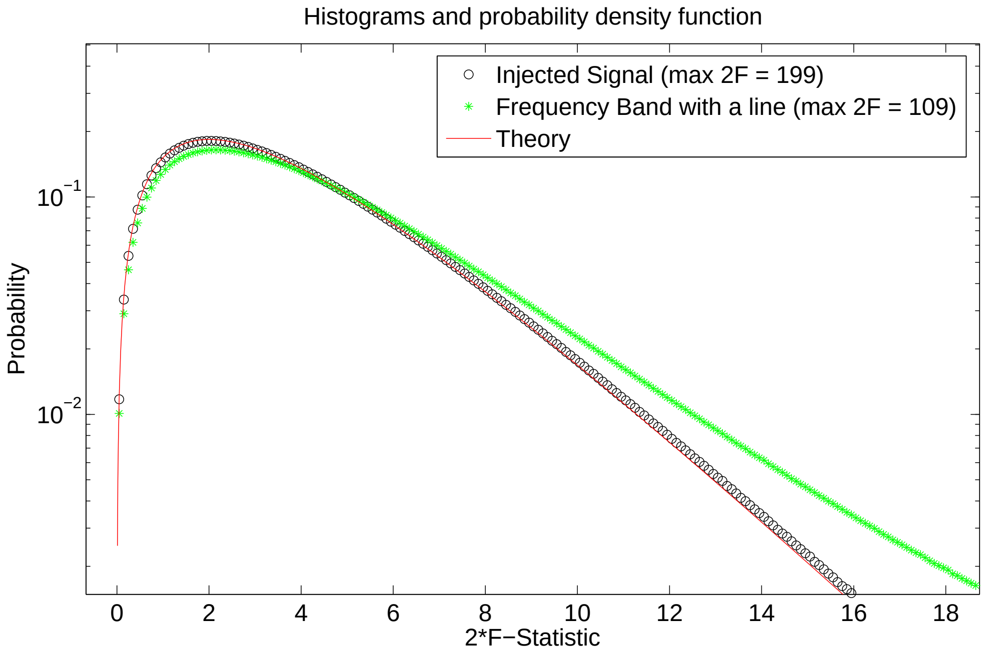

Figure 7.6: A zoom of Figure 7.5.

is less than about 7×10^{-3}. This can be used as a veto threshold for the search. This is also confirmed by Figure 7.7, which is a plot of the histogram of all the KS statistic outputs from each frequency band and a best fit.

7.4 Search Results

The raw search results are shown in Figure 7.8, which shows the maximum value of $2\mathcal{F}$ over all of the 10^9 or so templates in each frequency band. Also shown in the Figure are the expected loudest $2\mathcal{F}$ for each frequency band and also the loudest $2\mathcal{F}$ expected for the entire search (assuming only Gaussian noise). The 1% false alarm rate threshold is also shown.

The KS test statistic, which is calculated for each band from its histogram output, is shown in Figure 7.9. A couple of zoomed in plots show how the KS test statistic picks out the well known lines like the 60Hz harmonics (Figure 7.10) and the violin modes (Figure 7.11).

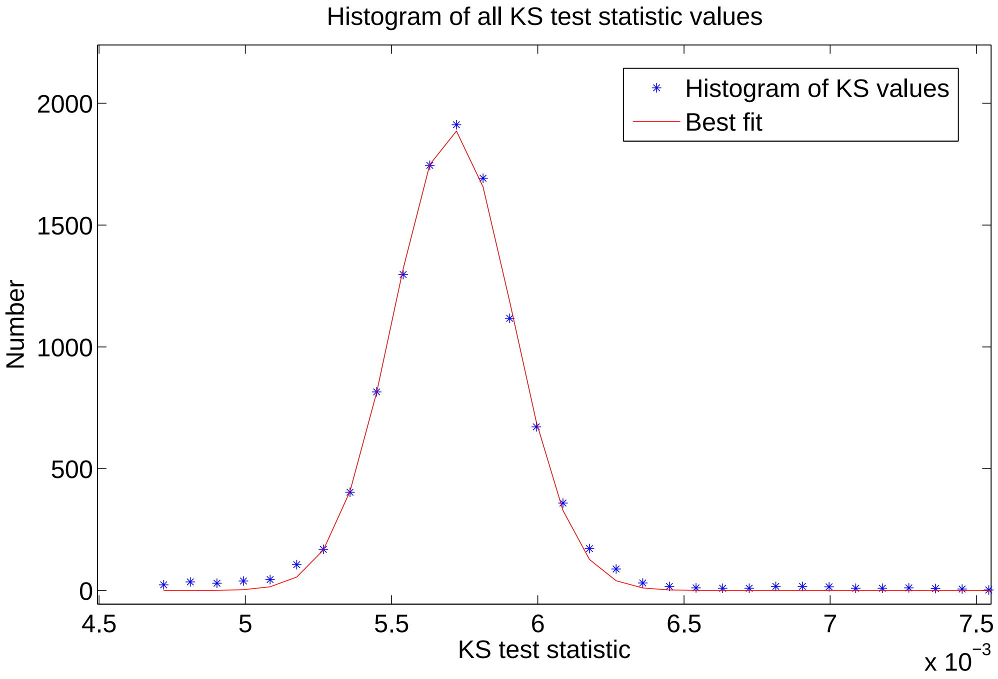

Figure 7.7: A histogram of the KS test stastitic. A best fit to a Gaussian is also shown in red. The mean of the best fit is at 5.7×10^{-3} and the sigma is 2.5×10^{-4}. Thus the veto threshold of 7×10^{-3} is much greater than 3σ.

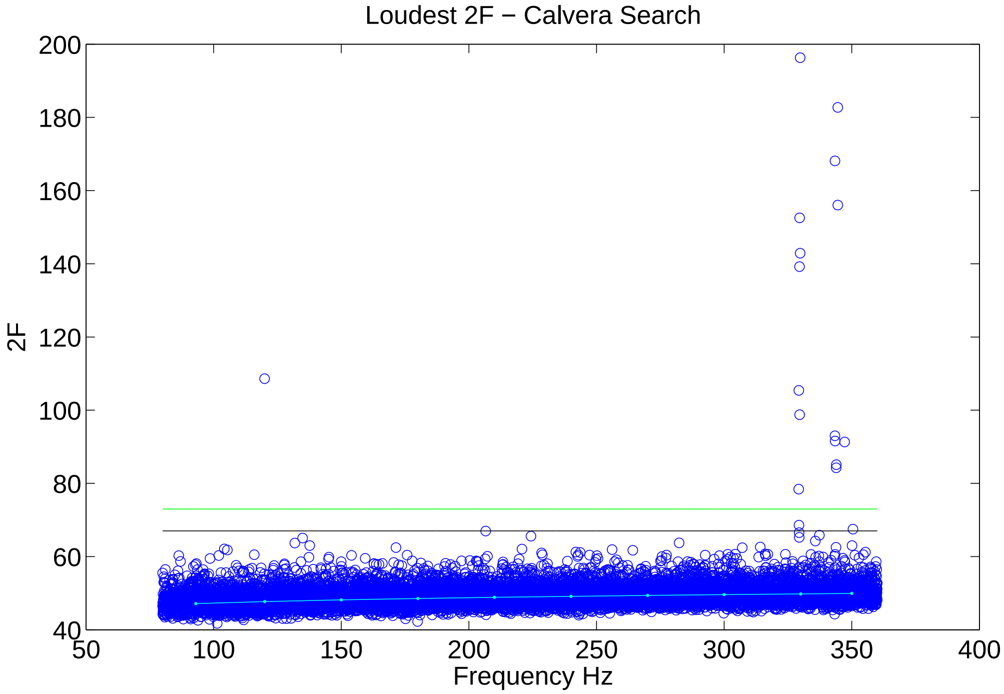

Figure 7.8: The raw results of the search. The maximum $2\mathcal{F}$ is shown for each frequency band. The black line is the expected loudest $2\mathcal{F}$ for the whole search. The green line is the 1% false alarm rate threshold. The cyan line is the expected loudest $2\mathcal{F}$ for each band.

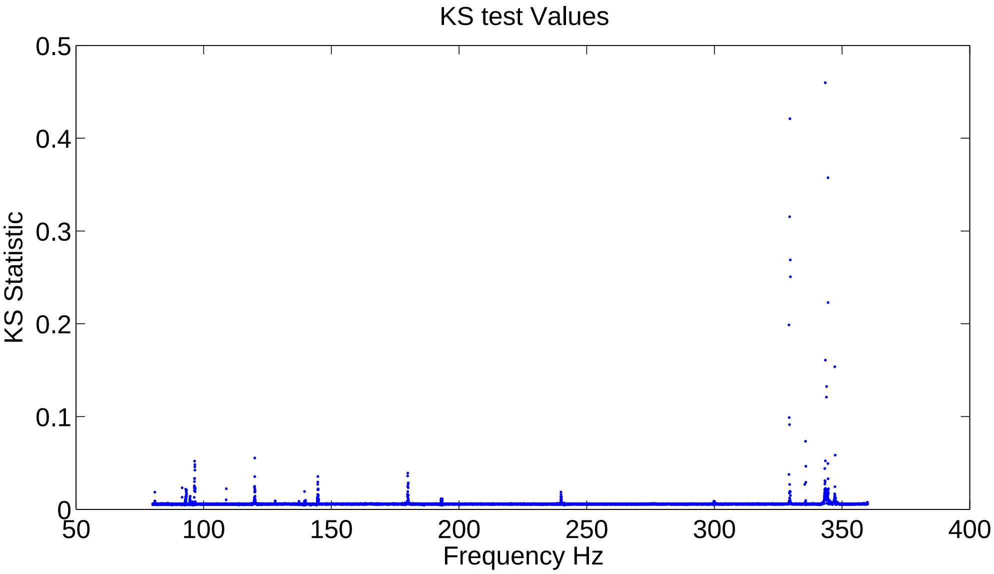

Figure 7.9: The KS test statistic for each frequency band.

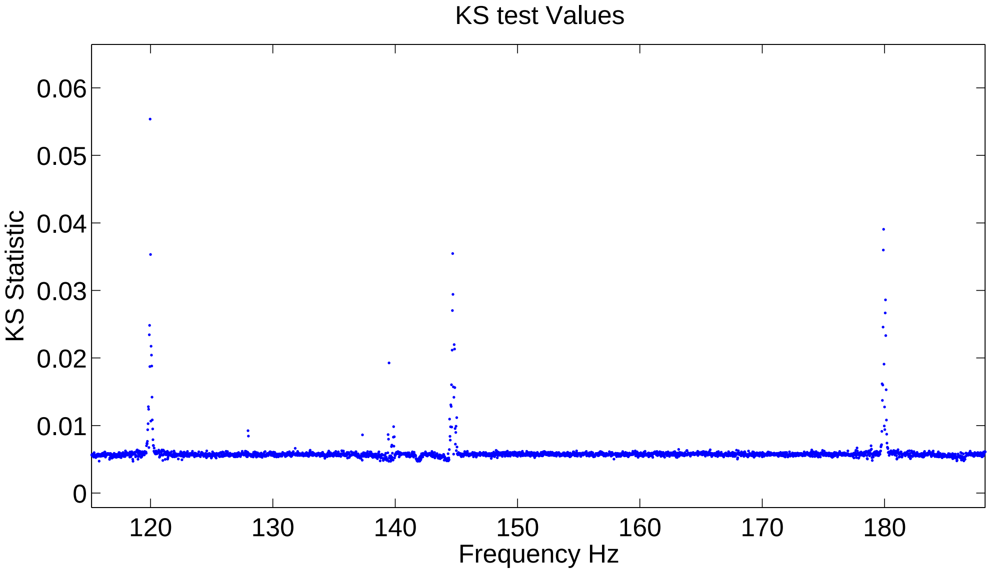

Figure 7.10: A zoom of Figure 7.9.

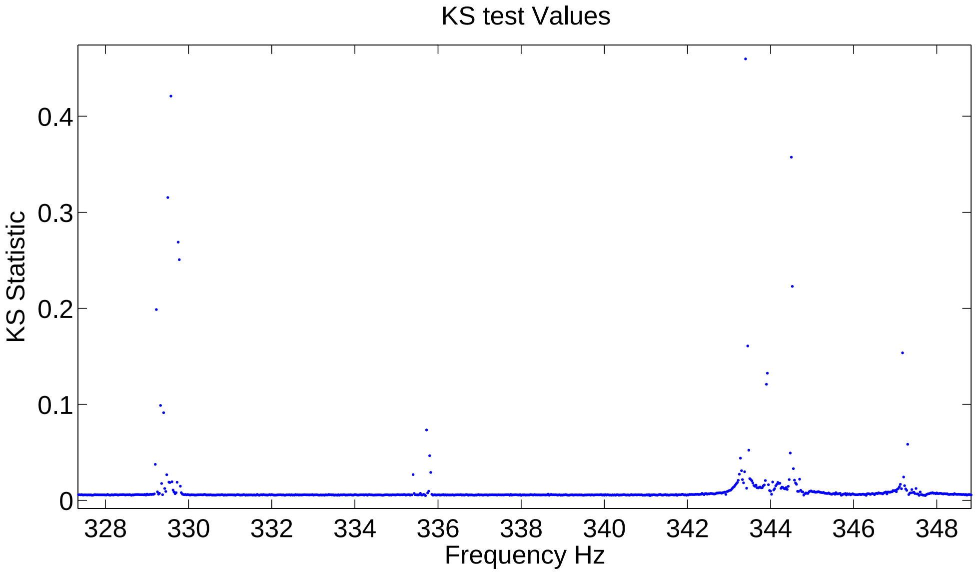

Figure 7.11: A zoom of Figure 7.9 in the vicinity of the suspension violin modes.

As discussed in the previous section, the KS test statistic can be used as a veto and simulations have shown that a threshold of approximately $7 \times 10^{-3}$is sufficient to veto most lines. The vetoed frequency bands are shown in Figure 7.12. These bands were investigated further. Most of the frequency bands vetoed belong to the known lines shown in table 7.2. Some of the frequency bands belong to non-persistent lines, which are not shown in table 7.2. All the $2\mathcal{F}$ values that are above the 1% false alarm rate are due to either the 60Hz harmonics or the violin modes of the detector and are vetoed by the KS statistic.

After the $2\mathcal{F}$ values of the vetoed frequency bands are removed, the relevant results are shown in Figure 7.13. None of the templates have crossed the 1% false alarm rate threshold, thus, without any further investigation, this can be considered a null result. However a follow up scheme in the event of a potential signal is discussed in the section below. The loudest event has a $2\mathcal{F}$ of 67.47, while the expected loudest event was 67.1 ± 2.1. The data also follows the expected loudest events on a per band basis, which is frequency dependent because of the difference in the number of templates. All of this suggests that the results of the

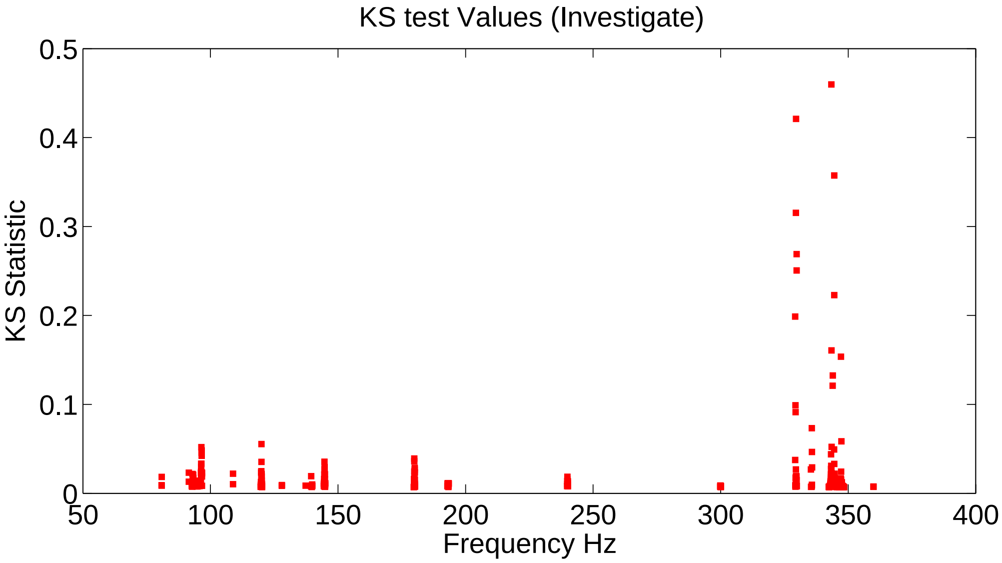

Figure 7.12: Frequency bands vetoed using the KS test statistic

search are consistent with the absence of signal in Gaussian noise. The noise is very close to being Gaussian, after the removal of the bands which have KS test statistic above the veto threshold.

7.5 Follow up

While the search yielded no results worth following up, set by the pre-determined criteria that only events crossing the 1% false alarm rate threshold would be followed up, a set of follow up rules were laid down in case of such an eventuality. This section details some of these rules and also talks about following up triggers from other CW searches.

7.5.1 Criteria for Follow up

The first step in determining a follow up strategy is to set thresholds and other parameters which determine the "interest" in the output from a given template (i.e., the value of $2\mathcal{F}$ obtained from filtering the data through a given template) before looking at the results of the search. This is crucial, since we do not want to

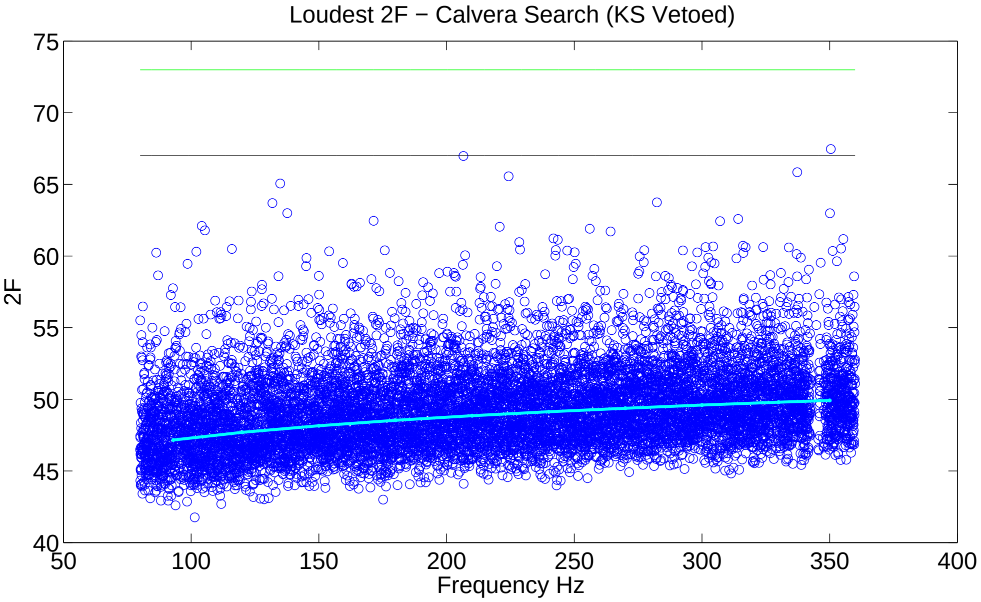

Figure 7.13: $2\mathcal{F}$ for frequency bands not vetoed by the KS test statistic. The black line is the expected loudest $2\mathcal{F}$for the whole search. The green line is the 1% false alarm rate threshold. The cyan line is the expected loudest $2\mathcal{F}$ for each band.

be biased by the results themselves. A good criterion for determining if a template needs to be followed up or not, is the probability of a noise fluctuation being equal to or greater than the output of the template (the false alarm rate). For CW searches like the Calvera search, it is relatively straightforward, since the noise is very close to Gaussian and thus the false alarm rates can be estimated to relatively great accuracy using theoretical predictions.

In the case of the Calvera search, the pre-determined threshold was set at a $2\mathcal{F}$ of 76.5, which corresponded to a false alarm rate of 1%. All the events that crossed this threshold were investigated. Since after vetoing known lines and bands failing the KS test, none of the templates crossed this threshold, there was no need to follow any of these events up. Some erroneous SFTs were included in the search to begin with and this was detected with the help of the follow up procedure. This is discussed at the end of this section.

A follow up scheme for CW searches would include the following steps -

1. **Increasing the sensitivity of the search** - A real signal's significance would increase with the increase in sensitivity of the search, while a noise fluctuation would be unaffected by such an increase in sensitivity. This can be done by increasing the integration time or by incorporating a more sensitive technique. One such follow up scheme is under development at Caltech [140], which would use triggers from a semi-coherent search like the Powerflux method [82, 83] and use the resampling algorithm to coherently integrate the same data used in the semi-coherent technique. This technique cannot be used for the Calvera search, since the Calvera search already uses the most efficient technique and all of the collected data.

2. **Adding an extra detector** - One of the most computationally cost effective ways of improving the sensitivity of a search is to include another detector in the search. When one is pushing the limit on the available computational power, it is sometimes necessary to leave out a less sensitive detector out of the search, for example, in the case of the Calvera search, H2 was left out.

However, while following up triggers, the parameter space to search over is in a narrow range around the suspected trigger and thus it becomes easier to include another detector into the analysis. If the sensitivity of the search does not increase in relation to the appropriate addition of a new interferometer, then it casts doubts about the validity of the trigger as a signal. On the other hand, an increase in the search statistic by an appropriate amount would lead to an increased confidence in the assumption that the trigger is a real signal.

3. **Splitting the data in time** - For a search like Calvera, in which all the available data has been used, it is sometimes useful to split the data into multiple parts and to conduct a narrow search around the trigger of interest. While the signal strength is reduced by decreasing the integration time, it is also offset by the reduced parameter space that needs to be searched over. The reduced parameter space decreases the chances of a statistical noise fluctuation looking like a signal. If one were to split the total integrated data into two parts, assuming that the signal was on continuously and that the sensitivity of the detectors was the same for both parts, the power ought to reduce by a factor of about 1.4. If the power were however only concentrated for a short period of time, this reduces the confidence of it being a continuous wave source. Further splitting of the data should also follow this pattern of having the signal strength distributed evenly throughout the data.

4. **Splitting the data by detector** - This is very similar to the splitting the data in time. Here again we reduce the parameter space around the trigger of interest and split the data by detectors. If there are two detectors like in the Calvera search, which have similar sensitivities, then the power ought to reduce by a factor of 1.4. If however this does not happen and all the power is seen only in one detector, then the trigger is most likely due to an instrument artifact.

A Follow Up Example

The first time the Calvera search was conducted, all the SFTs present on the ATLAS cluster were used for the analysis. It was erroneously assumed that only science mode data from S5 was used to generate the SFTs. Unfortunately some of the SFTs after the official end of S5 on Oct 1st 2007 00:00 Universal Time, were included along with all the S5 SFTs. Thus the search was run with 9 SFTs that were not officially science mode and should not have been included in the search to begin with.

I was however unaware of this and it only showed up because the results of the first search showed two loud events at 91 Hz and 91.1 Hz. The $2\mathcal{F}$ at 91 Hz was about 210 and at 91.1 Hz it was 355. These were well above the 1% false alarm rate cutoff and thus warranted further investigation. The bands containing both these frequencies failed the KS test. This suggested that a line might be present around those frequency bands. However there were no known persistant or non-persistant lines around the band of interest. Thus these two bands warranted further investigation. I will only discuss the follow up of the 91.1 Hz event. The 91 Hz event was followed up in a very similar fashion.

The first step in the follow up was to divide the data by detector. A search run on all the data, but by using only H1 around the 91.1 Hz frequency band gave a maximum $2\mathcal{F}$ of 37, which was consistent with noise. The same search done with the L1 detector gave a maximum $2\mathcal{F}$ of 858 or so. This $2\mathcal{F}$ was higher than the $2\mathcal{F}$ found using both the detectors because the template space used was more dense and the noise of H1 was not added in. This meant L1 was the only detector responsible for the event.

The second step was to discover the physical cause of the L1 instrumental artifact. In order to isolate the time of the instrumental artifact, all of S5 was divided into 10 parts of equal length and the search was conducted on each of these parts seperately. The only time span at which the $2\mathcal{F}$ exceeded what was expected from noise was for the last part. Upon further investigation, it was narrowd down to 9 SFTs after the end of the S5 run on Oct 1st 2007.

During the first week of October, a graduate student was conducting experiments with the photon calibrator at the L1 detector. The photon calibrator consisted of two separate injections at 91 and 91.1 Hz and during the time that the photon calibrator was operational, these lines were active. During this time, some of the data was accidently marked as science mode data and this was then converted to SFTs, which were then stored along with the rest of the S5 SFTs on the ATLAS cluster.

This example shows how a follow up would have proceeded in the event of the suspicion of a real signal. If all of these tests were passed, then further data or astrophysical observation could have been used to increase our confidence of detection.

7.6 Upper Limit Calculation

The search for GWs from Calvera was a discovery search and thus the parameter space was set to maximize the detection potential. However there was no detection and thus upper limits were set for the gravitational strain expected from an object like Calvera if it were to meet the optimistic assumptions of its position, age and nature.

7.6.1 Theoretical Predictions

The $\mathcal{F}$-statistic is by design maximized over nuisance variables like the inclination angle, polarization angle and initial phase of the GWs. However these variables change the amplitude of the signal and thus the probability that it would rise above the noise and be detected. In order to compute an analytical upper limit, we need to compute the optimal signal to noise ratio and then use equation 7.12 to compute the probability that the $\mathcal{F}$-statistic value will be lower than the maximum $\mathcal{F}$-statistic found in the search. An upper limit would be set for each frequency band separately.

The optimal SNR ρ^2 can be computed through the following equation [110]

$$\rho^2(f) \approx A_2(\delta, \psi, \iota) T_0 \sin^2 \chi \frac{h_0^2}{S_h(f)}, \tag{7.19}$$

where T_0 is the integration time, f is the frequency of the frequency band, δ is the declination of the source, ψ is the polarization angle and ι is the inclination angle of the source. S_h is the power spectral density and h_0 is the gravitational strain. χ is the angle between the arms of the interferometer. A_2 can be calculated by [110]

$$A_2(\delta, \psi, \iota) = F_2(\iota) e_1(\delta) \cos 4\psi + G_2(\iota) e_2(\delta), \tag{7.20}$$

where,

$$F_2(\iota) = \frac{\sin^4 \iota}{4} \tag{7.21}$$

and

$$G_2(\iota) = \frac{1}{4}(1 + 6\cos^2 \iota + \cos^4 \iota). \tag{7.22}$$

e_1 and e_2 are different for each interferometer and depend on the latitude and longitude of its location -

$$e_1(\delta) = 4 j_1 \cos^4 \delta \tag{7.23}$$

and

$$e_2(\delta) = 4 j_2 - j_3 \cos 2\delta + j_1 \cos^2 2\delta. \tag{7.24}$$

j_1, j_2 and j_3 are functions of λ (latitude) and γ (longitude) -

$$j_1(\lambda, \gamma) = \frac{1}{256}(4 - 20\cos^2 \lambda + 35 \sin^2 2\gamma \cos^4 \lambda), \tag{7.25}$$

$$j_2(\lambda, \gamma) = \frac{1}{1024}(68 - 20\cos^2 \lambda - 13 \sin^2 2\gamma \cos^4 \lambda), \tag{7.26}$$

and

$$j_3(\lambda, \gamma) = \frac{1}{128}(28 - 44\cos^2 \lambda + 5 \sin^2 2\gamma \cos^4 \lambda). \tag{7.27}$$

The upper limit for each frequency band is calculated as follows -

1. Loop over all frequency band.
2. Calculate $S_h(f)$.
3. Compute the loudest $2\mathcal{F}$ value. This value comes from the result of the search.
4. Make a matrix of ψ and $\cos\iota$ values.
5. Pick a h_0 and compute ρ^2 using equation 7.19 for each ψ and $\cos\iota$.
6. Using the loudest $2\mathcal{F}$ value as threshold, integrate equation 7.12 with a noncentral parameter of ρ^2 to get a probability that $2\mathcal{F}$ would be greater than the loudest value.
7. Adjust h_0 according to probability output and probability desired. A simple bisection method works just fine to hone in the h_0 required for the desired probability. For the Calvera search, probability required was 95%.

This procedure works well when the noise is assumed to be Gaussian. For the Calvera search, the noise is very close to Gaussian for the non-vetoed bands. The upper limits on h_0 calculated analytically as described above can be seen in Figure 7.14, along with the injections which were done to verify them.The injections are consistently about 10% higher than the analytic upper limits. This 10% is due to the inherent lossiness of the resampling technique and due to the fact that exactly matched templates were not used to span the paramter space of the search. A 15% mismatch tolerance was used instead. Thus the upper limits quoted for the search are 10% greater than the analytic upperlimits and can be seen in Figure 7.15. The age-based limit expected for Calvera, if it were 1×10^7 years old and at 80 pc from Earth, would be 2.83×10^{-25} and the lowest upper limit is at 152 Hz is 1.14×10^{-25}.

7.6.2 Injections

The upper limits shown in Figure 7.14 are analytical upper limits and do not take into account some of the approximations that went into the calculation of the $\mathcal{F}$-

statistic, like the inherent losses due to interpolation, the use of a limited number of templates with a specific mismatch parameter. These approximations can be of the order of a few percent and its very difficult to predict the exact number. These systematic errors are discussed in Sec 7.6.3

In previous searches like the Crab search [80], upper limits were set by using large scale Monte Carlo simulations. These Monte Carlo simulations involved picking a h_0 and then injecting signals with different polarizations and inclination angles into the data. The search was then run on a restricted parameter space and the number of times the loudest $2\mathcal{F}$ value crosses the loudest $2\mathcal{F}$ value that was recorded during the search is tallied. This is then converted to a percentage calculation and the h_0 for which the probability crossed 95% is set as the upper limit.

Monte Carlo simulations were not possible for the Calvera search due to the prohibitive computational cost. However Monte Carlo simulations were done at 5 Hz intervals for the entire band with about 10000 injections each. The results gave upper limits that were about 10% greater than the ones calculated analytically. This is consistent with the 2% expected losses due to interpolation and about 8% expected losses due to the mismatch parameter of 15%. Thus all the upper limits in 7.15 and 7.16 are quoted 10% higher than the analytically computed upper limits.

Using equation 7.5, the upper limits in Figure 7.15 can be converted to ellipticities. The upper limits in the ellipticities are shown in Figure 7.16

7.6.3 Systematic Errors

The upper limits presented above for this search are subject to several sources of systematic errors. Most of these systematics are very small and, do not change the results significantly and are unavoidable. These kinds of errors include errors in the ephemeris data that is used to calculate the difference in photon arrival times from a given source between the detector and the solar system barycenter, timing errors that exist in the data recording computers, errors that exist due to

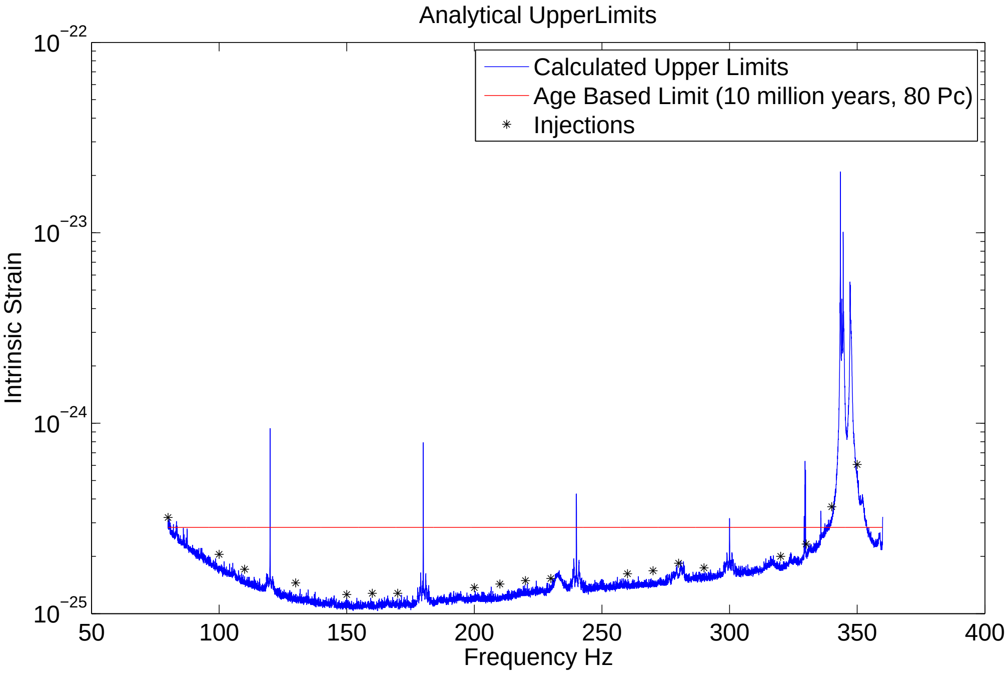

Figure 7.14: Analytically computed upper limits for the Calvera search. The upper limits computed through injections, carried out to verify certain bands are shown in black.

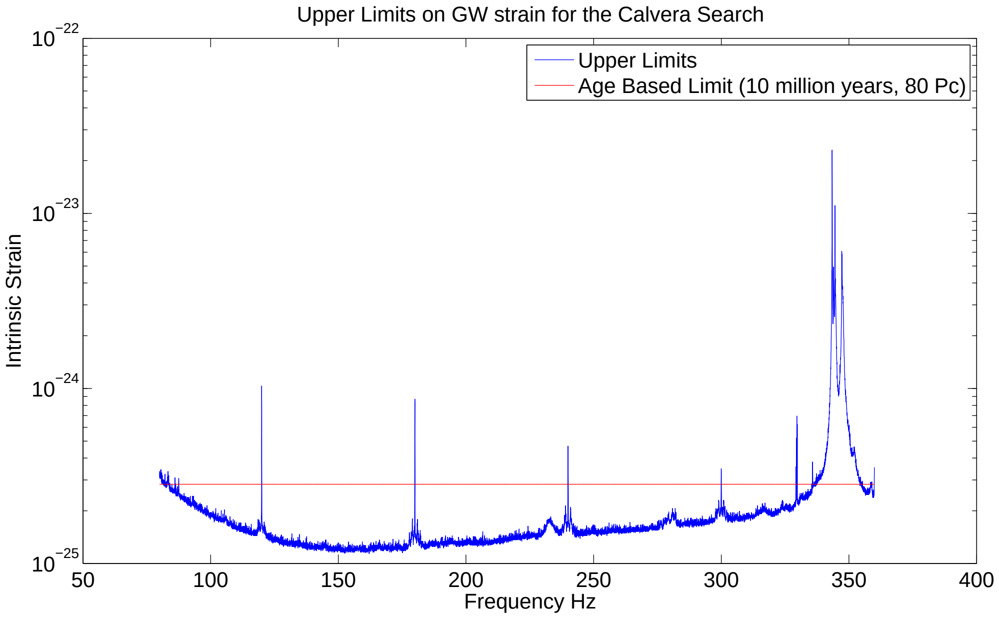

Figure 7.15: Upper limits on GW strain for the Calvera search.

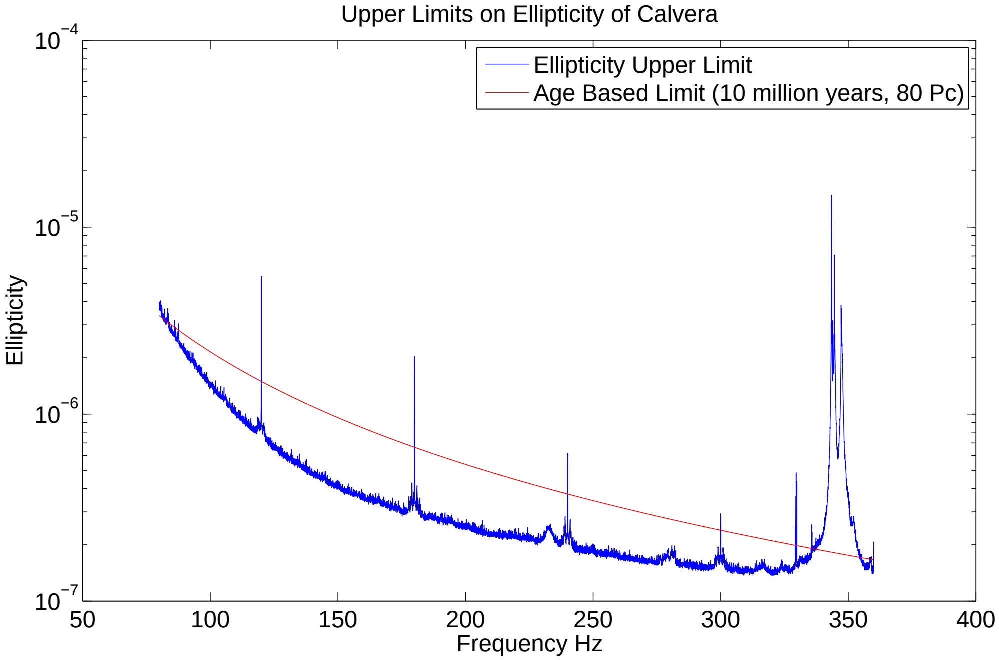

Figure 7.16: Upper limits on ellipticity of Calvera.

the inaccurate calibration of the detector. I discuss some of these systematics in this section.

Losses in the resampling code

The code that is used to compute the $\mathcal{F}$-statistic using the barycentric resampling technique (see Chapter 5) makes some assumptions that can lead to a slight loss in SNR from the optimal that can be estimated theoretically. These approximations include assumptions that the antenna patterns of the detector do not change over the course of about half an hour. These assumptions reduce the computational cost significantly without affecting the SNR very much at all. Other losses include interpolation losses that were discussed in Chapter 5.

The interpolation error is impossible to predict exactly since the interpolation pattern changes with sky position. Thus it is hard to precisely correct for interpolation losses. The solution is to compute the $\mathcal{F}$-statistic for a larger frequency band than desired and then to only use the innermost frequency bins. Innermost frequency bins are the frequencies which are closer to DC (both positive and negative), since the effects of interpolation increase as the frequency gets closer to the Nyquist frequency. As mentioned in Chapter 5, cubic spline interpolation is used in the resampling code and a frequency band that is approximately twice the desired frequency band is used. In spite of using only the middle half of the computed frequency band, losses of SNR or a few percent can be expected. On average we can expect about 2% loss in SNR. The effect of this loss and the loss associated with metric mismatch discussed below, is estimated and included in the quoted upper limits using the procedure described in Sec 7.6.1 above.

Metric mismatch

As discussed in Chapter 4, a metric mismatch parameter is used to define the grid spacing in search parameter spaces. The mismatch is the maximum SNR loss that can be tolerated and thus templates are placed overlapping in such a manner that the maximum loss is as specified by the mismatch. For exact equations please see

Chapter 4.

A 15% mismatch was used to lay down the Calvera search parameter space. While allowing for the second frequency derivative, only single templates were needed to cover the second frequency derivative parameter space. Since only one template was used to cover the sky position parameter space as well, all the templates were divided into the nearly orthogonal frequency and frequency derivative space.

Monte Carlo simulations done on these kinds of parameter spaces such as in [80, 136] show that on average a randomly chosen template lying within the parameter space covered manifests itself as an SNR loss of the total mismatch divided by the number of orthogonal parameters that are being covered. Similar Monte Carlo simulations performed by the author (not as extensive as in [80, 136]), show this to be the case for the Calvera search and an SNR loss of about 8% is seen. estimated. The effect of this loss and the loss associated with interpolation in the resampling code, is estimated and included in the quoted upper limits using the procedure described in Sec 7.6.1 above.

Calibration

The calibration error uncertainties arise from the inability to measure the absolute calibration of the detector response to differential strain caused by GWs on the LIGO detectors. The methods used to compute these uncertainties are discussed in [141]. In [141], the calibration errors of H1 were estimated to be around 10.1% and for L1 to be about 14.1%. These estimates were used to adjust the analytically calculated upper limits presented in Sec 7.6.1.

Chapter 8

Conclusions

This thesis presents the results of a search for continuous, quasi-periodic gravitational waves from Calvera, an X-ray source that is suspected of being as close as 80 pc from the Earth and may be rapidly spinning. This kind of deep and wide-parameter search is made possible because of the development of barycentric resampling to greatly speed up the computation of the optimal detection statistic (the $\mathcal{F}$-statistic). No evidence for GWs was found. Upper limits on the gravitational strain that could have been detected from the search were set in a band from 80-360 Hz. Upper limits were also set on the ellipticity of Calvera.

This thesis summarized the work I did in the LSC CW group. In Chapter 2 we briefly reviewed how the initial LIGO interferometers worked. We summarized how continuous GWs are emitted from neutron stars and how the GWs can help us understand the underlying physics of these extreme objects. The resampling algorithm and its implementation was introduced. The resampling algorithm fits into various analysis techniques used by the CW group and will help improve them in the near future. An object dubbed Calvera was then introduced and its potential as a GW source was discussed. A search was then conducted to look for a gravitational signal from Calvera. No significant results were found and the data was shown to be remarkably consistent with Gaussian noise. Thus upper limits were set on the gravitational wave strain and the ellipticities expected from Calvera if certain optimistic assumptions were met.

The upper limits on the Calvera search are the most sensitive broadband GW searches done using LIGO data. The lowest h_0 was in the range of 1×10^{-25} and at the highest frequencies examined in this search, the ellipticity of the neutron star was constrained to be less than 2×10^{-7}. Without the technical innovations described in this thesis, with the same amount of computational power, the upper limit on h_0 would have been 5 times higher. This would place the upper limits above the age based upper limit of GW signal expected from Calvera, making the search uninteresting.

The ellipticity upper limits lie within a realistic range of possible ellipticities as seen in Chapter 3. However this does not constrain any of the theoretical neutron star models, since it is possible that the ellipticities were simply much lower than the maximum sustainable. Or the neutron star may not be spinning at a frequency accessible in this search. Moreover, some or all of the optimistic assumptions about Calvera could be false. The Calvera search was meant to be a discovery search and upper limits were only calculated because no events were found.

With advanced LIGO commissioning taking place from now till about 2014, the detector sensitivity will go up by about an order of magnitude. Advanced LIGO will also push the lowest searchable frequency down to about 10 Hz, bringing many pulsars in the LIGO band. The searches for continuous GWs will remain computationally bound. The implementation of the resampling algorithm as discussed in this thesis will go a long way in improving searches in the future.

Bibliography

[1] J. M Weisberg, J. H Taylor, FA Rasio, IH Stairs - ASP Conf. Ser. Vol. 328 (2005)

[2] J. B. Hartle, GRAVITY, An introduction to Einstein's general relativity

[3] B. Allen, χ^2 time-frequency discriminator for gravitational wave detection, Phys. Rev. D. 71, 062001 (2005)

[4] A. Rodriguez Master's thesis, Louisiana State University

[5] B. Abbott *et al.*, Analysis for LIGO data for gravitaional waves from binary neutron stars, Phys. Rev. D. 69, 122001 (2004)

[6] B. Abbott *et al.*, Search for gravitational waves from the galactic and extra-galactic binary neutron stars, Phys. Rev. D. 72, 082001 (2005)

[7] B. Abbott *et al.*, Search for gravitational waves from primordial black hole binary coalescences in the galactic halo, Phys. Rev. D. 72, 082002 (2005)

[8] B. Abbott *et al.*, Search for gravitational waves from binary black hole inspirals in LIGO data, Phys. Rev. D. 73, 062001 (2006)

[9] B. Abbott *et al.*, TAMA Collaboration, Joint LIGO and TAMA300 search for gravitational waves from inspiralling neutron star binaries, Phys. Rev. D. 73, 102002 (2006)

[10] B. Abbott *et al.*, Search for Gravitational-Wave Bursts from Soft Gamma Repeaters, Phys. Rev. D. 78, 062003 (2008)

[11] B. Abbott *et al.*, Search for gravitational waves from binary inspirals in S3 and S4 LIGO data, Phys. Rev. D. 77, 062002 (2008)

[12] B. Abbott *et al.*, Search for gravitational waves from low mass binary coalescences in the first year of LIGOs S5 data, Phys. Rev. D. 79, 122001 (2009)

[13] B. Abbott *et al.*, Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data, Phys. Rev. D. 80, 062001 (2009)

[14] B. Abbott *et al.*, Search for gravitational waves from low mass compact binary coalescence in 186 days of LIGOs fifth science run, Phys. Rev. D. 80, 047101 (2009)

[15] B. Abbott *et al.*, Search for gravitational-wave bursts associated with t gamma-ray bursts using data from LIGO science run 5 and VIRGO science run 1, ApJ. 715, 1453 (2010)

[16] C. D. Ott, Probing the core-collapse supernova mechanism with gravitational waves, Class. Quant. Grav. 26:063001 (2009)

[17] B. Abbott *et al.*, First upper limits from LIGO on gravitational wave bursts, Phys. Rev. D. 69, 102001 (2004)

[18] B. Abbott *et al.*, Search for gravitational waves associated with the gamma ray burst GRB030329 using the LIGO detectors, Phys. Rev. D. 72, 042002 (2005)

[19] B. Abbott *et al.*, Upper limits on gravitational wave bursts in LIGOs second science run, Phys. Rev. D. 72, 062001 (2005)

[20] B. Abbott *et al.*, Upper limits from the LIGO and TAMA detectors on the rate of gravitational-wave bursts, Phys. Rev. D. 72, 122004 (2005)

[21] B. Abbott *et al.*, Search for gravitational-wave bursts in LIGO's third science run, Class. Quant. Grav. 23, S29-S39 (2006)

[22] B. Abbott *et al.*, A Joint Search for Gravitational Wave Bursts with AURIGA and LIGO, Class. Quant. Grav. 25, 095004 (2008)

[23] B. Abbott *et al.*, Search for gravitational wave radiation associated with the pulsating tail of the SGR 1806-20 hyperflare of December 27, 2004 using LIGO, Phys. Rev. D. 76, 062003 (2007)

[24] B. Abbott *et al.*, Search for gravitational-wave bursts in LIGO data from the fourth science run, Class. Quant. Grav. 24, 5343-5369 (2007)

[25] B. Abbott *et al.*, Search for Gravitational Waves Associated with 39 Gamma-Ray Bursts Using data from the Second, Third, and Fourth LIGO Runs, Phys. Rev. D. 77, 062004 (2008)

[26] B. Abbott *et al.*, First joint search for gravitational-wave bursts in LIGO and GEO600 data, Class. Quant. Grav. 25, 245008 (2008)

[27] B. Abbott *et al.*, Implications for the Origin of GRB 070201 from LIGO Observations, ApJ. 681, 1419 (2008)

[28] B. Abbott *et al.*, Search for Gravitational Wave Bursts from Soft Gamma Repeaters, Phys. Rev. Lett. 101, 211102 (2008)

[29] B. Abbott *et al.*, First LIGO search for gravitational wave bursts from cosmic (super)strings, Phys. Rev. D. 80, 062002 (2009)

[30] B. Abbott *et al.*, Search for High Frequency Gravitational Wave Bursts in the First Calendar Year of LIGO's Fifth Science Run, Phys. Rev. D. 80, 102002 (2009)

[31] B. Abbott *et al.*, Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm, ApJ. 701, L68-L74 (2009)

[32] B. Abbott *et al.*, All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run, Phys. Rev. D. 80, 102001 (2009)

[33] B. Abbott *et al.*, All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run, Phys. Rev. D. 81, 102001 (2010)

[34] B. Abbott *et al.*, Analysis of first LIGO science data for stochastic gravitational waves, Phys. Rev. D. 69, 122004 (2004)

[35] B. Abbott *et al.*, Upper limits on a stochastic background of gravitational waves, Phys. Rev. Lett. 95, 221101 (2005)

[36] B. Abbott *et al.*, Searching for Stochastic Background of Gravitational Waves with LIGO, ApJ. 659, 918 (2007)

[37] B. Abbott *et al.*, Upper limit map of a background of gravitational waves, Phys. Rev. D. 76, 082003 (2007)

[38] B. Abbott *et al.*, First Cross-Correlation Analysis of Interferometric and Resonant-Bar Gravitational-Wave Data for Stochastic Backgrounds, Phys. Rev. D. 76, 022001 (2007)

[39] B. Abbott *et al.*, An upper limit on the stochastic gravitational-wave background of cosmological origin, Nature 460, 990 (2009)

[40] B. Abbott *et al.*, LIGO: The Laser Interferometer Gravitational-Wave Observatory, Rep. Prog. Phys. 72, 076901 (2009)

[41] R. Drever *et al.*, Laser phase and frequency stabilization using an optical resonator, Appl. Physics B 31 (1983)

[42] E. Black, An introduction to PoundDreverHall laser frequency stabilization, Am. J. Phys. 69, 79 (2001)

[43] S. Balmer, Upper Limit Map of a Stochastic Background of Gravitational Waves, PhD thesis, Massachusetts Institute of Technology (2006)

[44] J. Betzweiser, Analysis of spatial mode sensitivity of a gravitational wave interferometer and a targeted search for gravitational radiation from the Crab pulsar, PhD thesis, Massachusetts Institute of Technology (2007)

[45] J. Slutsky, Quantifying the impact of data quality on searches for gravitational waves from binary coalescing systems with LIGO, PhD thesis, Louisiana State University (2010)

[46] N. K. Glendenning, Strangeness in compact stars and signal of deconfinement, ApJ 23, 12 (1997)

[47] N. K. Glendenning, Compact Stars, Nuclear Physics, Particle Physics and General Relativity (1997)

[48] Fruchter, A.S. *et al.*, A millisecond pulsar in an eclipsing binary. Nature 333: 237-239.

[49] D. B. Kaplan and A. E. Nelson, Strange goings on in dense nucleonic matter, Phys. Lett. B 175,57 (1986)

[50] D. B. Kaplan and A. E. Nelson, Strange condensate realignment in relativistic heavy ion collisions, Phys. Lett. B 179,409 E (1986)

[51] J. C. Collins and M. J. Perry, Superdense Matter: Neutrons or Asymptotically Free Quarks?, Phys. Rev. Lett. 34, 1353 (1975)

[52] A. Hewish *et al.*, Observation of a Rapidly Pulsating Radio Source, Nature 217, 709 (1968)

[53] J. H. Seiradakis, Pulsar astronomy: Older than they look, Nature 406, 139-140 (200)

[54] L. A. Rawley *et al.* ApJ 326, 947 (1988)

[55] A. G. Lyne in *High Energy Phenomena Around Collapsed Stars* (1987)

[56] V. Trimble, Astron. J. 73, 535 (1968)

[57] F. Pacini, Rotating Neutron Stars, Pulsars and Supernova Remnants, Nature 219, 145 (1968)

[58] N. Andersson, Gravitational waves from instabilities in relativistic stars, CQG, 20, R105 (2003)

[59] S. Chandrasekhar, Solutions of two problems in theory of gravitational radiation, Phys. Rev. Lett. 24, 611 (1970)

[60] J. L. Friedman and B. F. Schutz, ApJ 222, 281 (1978)

[61] L. Lindblom *et al.*, Gravitational Radiation Instability in Hot Young Neutron Stars, Phys. Rev. Lett. 80, 4843 (1998)

[62] B. J. Owen, arXiv:1006.1994 (2010)

[63] A. Reisenegger and A.Bonačić, Millisecond Pulsars with r-Modes as Steady Gravitational Radiators, Phys. Rev. Lett, 91, 20 (2008)

[64] M. Zimmermann and E. Szadenits, Phys. Rev. D., 20 (1979)

[65] G. Baym and D. Pines, Ann. Phys. 66, 816 (1971)

[66] M. A. Ruderman, Nature 223, 597 (1969)

[67] C. Cutler, Gravitational waves from neutron stars with large toroidal B fields, Phys. Rev. D. 66, 084025 (2002)

[68] M. Bejger and P. Haensel, Astron. Astrophysics 396, 917 (2002)

[69] V. S. Beskin, Radio pulsars, PHYS-USP, 1999, 42 (11), 10711098.

[70] D. Page, http://www.astroscu.unam.mx/neutrones/NS-Picture/NStar/NStar.html, on August 16 2010.

[71] C. Misner, K. Thorne and J. Wheeler, *Gravitation* (1973)

[72] R. Prix for the Ligo Scientific Collaboration, LIGO-P060039-00-Z

[73] C. J. Horowitz and K. Kadau, Breaking Strain of Neutron Star Crust and Gravitational Waves, Phys. Rev. Lett. Volume 102, Issue 19

[74] C. Cutler and D. I. Jones, Gravitational wave damping of neutron star wobble, Phys. Rev. D. 63(2), 024002 (2001)

[75] L. Bildsten, Gravitational radiation and rotation of accreting neutron stars, ApJ 501 L89 (1998)

[76] A. Melatos and D. J. B. Payne, Gravitational Radiation from an Accreting Millisecond Pulsar with a Magnetically Confined Mountain, ApJ 623 1044 (2005)

[77] B. J. Owen, arXiv:astro-ph/0503399v2

[78] A. R. Bodmer, Collapsed Nuclei, Phys. Rev. D. 4, 1601 (1971)

[79] B. Abbott *et al.*, Limits on gravitational wave emission from selected pulsars using LIGO data, Phys. Rev. Lett. 94 181103 (2005)

[80] B. Abbott *et al.*, Beating the spin-down limit on gravitational wave emission from the Crab pulsar, ApJ Lett. 683, 45 (2008)

[81] B. Abbott *et al.*, Searches for gravitational waves from known pulsars with S5 LIGO data, ApJ 713, 671 (2010)

[82] B. Abbott *et al.*, All-sky search for periodic gravitational waves in LIGO S4 data, Phys. Rev. D. 77 022001 (2008)

[83] B. Abbott *et al.*, All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data, Phys. Rev. Lett. 102 , 122001 (2009)

[84] B. Abbott *et al.*, First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform, Phys. Rev. D 72, 102004 (2005)

[85] P. R. Brady and T. Creighton, Searching for periodic sources with LIGO. II. Hierarchical searches, Phys. Rev. D. 61(8), 082001 (2000)

[86] S. Frasca *et al.*, Class. Quant. Grav. 22, 1013 (2005)

[87] C. Cutler *et al.*, Improved stack-slide searches for gravitational-wave pulsars, Phys. Rev. D. 72(4), 042004 (2005)

[88] B. Abbott *et al.*, The Einstein@Home search for periodic gravitational waves in LIGO S4 data, Phys. Rev. D. 79, 022001 (2009)

[89] B. Abbott *et al.*, Einstein@Home search for periodic gravitational waves in early S5 LIGO data, Phys. Rev. D. 80, 042003 (2009)

[90] http://boinc.berkeley.edu/

[91] http://einstein.phys.uwm.edu/

[92] TwoSpect search algorithm, being developed by E.Goetz and K.Riles.

[93] Quadratic search algorithm, being developed by S.van der Puten and H.J.Bulten

[94] Sideband search algorithm, being developed by C.Messenger, L.Sammut, B.Owen and A.Melatos

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

[96] M. .P. Muno *et al.*, A deep chandra catalog of x-ray point sources toward the galatic center, ApJ 589:225-241 (2003)

[97] R. N. Manchester *et al.*, Discovery of ten millisecond pulsars in the globular cluster 47 Tuc, Nature 352,219-221 (1991)

[98] L. Wainstein and V. Zubakov, *Extraction of Signals from Noise*, Prentice-Hall, Englewood Cliffs (1962)

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

[100] R. Prix, Search for continuous gravitational waves: Metric of the multidetector $\mathcal{F}$-statistic, Phys. Rev. D. 75, 023004 (2006)

[101] J. Neyman and E. Pearson, Phil. Trans. R. Soc. Lond. A (1933)

[102] R. J. Dupuis and G. Woan, Bayesian estimation of pulsar parameters from gravitational wave data, Phys. Rev. D. 72, 102002 (2005)

[103] B. Abbott *et al.*, Upper Limits on Gravitational Wave Emission from 78 Radio Pulsars, Phys. Rev. D. 76 042001 (2007)

[104] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, Phys. Rev. D. 53, 6749 (1996)

[105] R. Balasubramanian *et al.*, Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters, Phys. Rev. D. 53, 3033 (1996)

[106] B. Krishnan *et al.*, Phys. Rev. D. 70, 802001 (2004)

[107] H. J. Pletsch and B. Allen, Exploiting Large-Scale Correlations to Detect Continuous Gravitational Waves, Phys. Rev. Lett. 103:181102 (2009)

[108] P. Patel *et al.*, Implementation of barycentric resampling for continuous wave searches in gravitational wave data, Phys. Rev. D. 81, 084032 (2010)

[109] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The ATNF Pulsar Catalogue. http://www.atnf.csiro.au/research/pulsar/psrcat/.

[110] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev.* **D58** 063001, 1998.

[111] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, *Class. Quant. Grav.* **26**, 204013 (2009)

[112] P.R. Williams and B.F. Schutz, An efficient Matched Filtering Algorithm for the Detection of Continuous Gravitational Wave Signals, *arXiv*:gr-qc, **9912029** v1

[113] ComputeFStatistic_v2, LIGO Analysis Library(LAL), https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.htm

[114] ComputeFStatistic_resamp, LAL, https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.htm

[115] R. Prix, LIGO Document **T0900149-v1**, https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument? docid=1665

[116] R. E. Rutledge *et al.*, Discovery of an Isolated Compact Object at High Galactic Latitude, *arXiv*:0705.1011v2 [astro-ph] (2007)

[117] F. M. Walter *et al.*, Discovery of a nearby isolated neutron star, Nature, 379, 233 (1996)

[118] C. Motch *et al.*, The isolated neutron star candidate RX J1605.3+3249, Astronomy & Astrophysics, 351, 177 (1999)

[119] D. N. Burrows *et al.*, X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy (2003), Proceedings of the SPIE, Volume 4851, 1320

[120] P. W. Roming *et al.*, Ultraviolet/Optical Telescope of the Swift MIDEX mission, Proc. SPIE Vol. 4140, p. 76-86 (2000)

[121] F. Haberl *et al.*, Evidence for precession of the isolated neutron star RX J0720.4-3125, Astronomy & Astrophysics, 451, L17 (2006)

[122] D. L. Kaplan *et al.*, The Parallax and Proper Motion of RX J1856.5-3754 Revisited, ApJ, 571, 447 (2002)

[123] R. N. Manchester *et al.*, The Australia Telescope National Facility Pulsar Catalogue, AJ, 129, 1993 (2005)

[124] D. Page and Reddy, annu. Rev. Nucl. and Part. Sci, in press

[125] C. -A. Faucher-Giguère and Kaspi, Birth and Evolution of Isolated Radio Pulsars, ApJ, 643, 332 (2006)

[126] P. M. Woods and C. Thompson, Soft gamma repeaters and anomalous X-ray pulsars: magnetar candidates (Compact stellar X-ray sources), 547-586 (2006)

[127] M. Durant and M. H. van Kerkwijk, Distances to Anomalous X-Ray Pulsars Using Red Clump Stars, ApJ, 650, 1070 (2006)

[128] M. P. Muno *et al.*, A Neutron Star with a Massive Progenitor in Westerlund 1, ApJ, 636, L41 (2006)

[129] G. G. Pavlov *et al.*, The Compact Central Object in the RX J0852.0-4622 Supernova Remnant, ApJ, 559, L131 (2001)

[130] S. Park *et al.*, Discovery of a Candidate Central Compact Object in the Galactic Nonthermal SNR G330.2+1.0, ApJ, 653, L37 (2006)

[131] S. Bogdanov *et al.*, Chandra X-Ray Observations of 19 Millisecond Pulsars in the Globular Cluster 47 Tucanae, ApJ, 646, 1104 (2006)

[132] A. Zepka *et al.*, Discovery of Three Radio Pulsars from an X-Ray–selected Sample, ApJ, 456, 305 (1996)

[133] J. P. Ostriker and J. E. Gunn, On the Nature of Pulsars. I. Theory, ApJ, 157 1395 (1969)

[134] A. S. H. Shevchuk *et al.*, Chandra Observations of 1RXS J141256.0+792204 (Calvera), ApJ, 705 391 (2009)

[135] C. Palomba, Pulsars ellipticity revised, Astronomy and Astrophysics, 354, 163 (2000)

[136] K. Wette *et al.*, Searching for gravitational waves from Cassiopeia A with LIGO, Class. Quantum Grav. 25, 235011 (2008)

[137] B. J. Owen *et al.*, Gravitational waves from hot young rapidly rotating neutron stars, Phys. Rev. D. 58 084020 (1998)

[138] `http://www.cs.wisc.edu/condor/`

[139] S. D. Mohanty, LIGO Document T030168-00-D

[140] J. Betzweiser and V. Dergachev, Private communication

[141] J. Abadie *et al.*, Calibration of the LIGO Gravitational Wave Detectors in the Fifth Science Run, arXiv:1007.3973v1 [gr-qc] (2010)