

Searching for Continuous Gravitational Waves from Unknown
Isolated Neutron Stars in Advanced LIGO Data

by
Orion Eli Sauter

A dissertation in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2018

Doctoral Committee:

Professor Keith Riles, Chair
Associate Professor Robert D. Deegan
Professor David W. Gerdes
Assistant Professor Emanuel Gull
Professor James T. Liu

Orion Eli Sauter

osauter@umich.edu

ORCID iD: 0000-0003-2293-1554

© 2018 Orion Eli Sauter

Acknowledgements

First and foremost, my deepest thanks to Doctors David Ebb and Shannon MacDonald, along with their entire staff of residents, technicians, and nurses at the Massachusetts General Hospital Pediatric Oncology Ward, and Francis Burr Proton Center. Through their efforts, I am alive today, and have the distinct privilege of telling my colleagues I have lain at the end of a particle accelerator, and had protons fired into my brain at $0.6c$.

Recovering from that experience was a difficult period, and I appreciate Professor Finn Larsen's patience. I was not the student either of us expected, but I hope he is as satisfied with the outcome as I am.

In spite of my difficulties, Professor Keith Riles welcomed me into his group without judgment. He has been the ideal advisor on this work – Every time I was ready to throw my hands up in defeat, he gave me the nudge (or more often, shove) in the right direction that I needed.

Vladimir Dergachev was an invaluable resource in learning to use the PowerFlux analysis pipeline he developed. Thanks to him, I was able to contribute to the continuous wave search efforts.

My parents, Steve and Sally, always fostered an interest in science from the very beginning – There was never anything else I wanted to do. With their support through everything, I have reached that goal.

Finally, my love and thanks to my wife Marika, who gives me strength every day.

We've each had our share of hardships, but I know together there is nothing we can't face. I look forward to a lifetime of adventures together!

Table of Contents

Acknowledgements	ii
List of Tables	viii
List of Figures.	x
Abstract	xiv
Chapter I Introduction.	1
Chapter II Gravitational Waves	4
2.1 General Relativity	4
2.2 Theoretical Derivation of Gravitational Waves	6
2.3 Source Requirements	9
2.4 Expected Sources	11
2.4.1 Compact Binary Coalescences	12
2.4.2 Bursts	12
2.4.3 Stochastic Background	13
2.4.4 Continuous Waves	13
Chapter III LIGO	15
3.1 LIGO Design Principles	16
3.2 Initial LIGO	18
3.3 Advanced LIGO Upgrades	19
3.4 Detector Noise	20

Chapter IV	Detecting Continuous Waves	24
4.1	Neutron Stars as GW Sources	24
4.2	Inherent Difficulty of Detection	26
4.3	Search Inputs	26
4.4	Search Classes	27
4.5	PowerFlux	28
4.5.1	StackSlide	28
4.5.2	Power Sum	29
4.5.3	Software Injections	29
4.5.4	Line Detection	30
4.5.5	Skymarks	31
4.6	Stages of Analysis	32
4.6.1	Semi-coherent	32
4.6.2	Fully Coherent	33
4.6.3	Loosely Coherent	33
4.7	Vetos	35
4.7.1	S Parameter	36
4.7.2	Line Response Veto	36
4.8	Pipeline Performance	38
Chapter V	S6 Mock Data Challenge	40
5.1	Hardware Injections	41
5.2	Software Injections	41
5.3	Pipelines	42
5.3.1	Time-domain \mathcal{F} -statistic	42
5.3.2	Frequency/Sky Hough	43
5.3.3	Einstein@Home	43

5.4	Challenge Parameters	43
5.5	Results	45
Chapter VI Line Cleaning in O1 Searches		49
6.1	Low-Frequency Lines	50
6.1.1	Combs	50
6.1.2	Non-Stationary Lines	51
6.1.3	Excessive Outliers	52
6.2	Cleaning	52
6.2.1	SFTClean	56
6.2.2	PEM Correlations	56
6.2.3	Comparison	57
6.2.4	Fraction Cleaned	61
Chapter VII Exploiting Longer Coherence Times		67
7.1	Losses and Gains	68
7.1.1	Line Avoidance	68
7.1.2	Signal Smear	68
7.1.3	Antenna Pattern	69
7.2	Choosing Coherence Length	70
7.3	Comparison Search	76
7.3.1	Wandering Lines	77
7.4	Conclusions	82
Chapter VIII LALBarycenter Approximations		84
8.1	Capabilities of LALBarycenter	85
8.1.1	LALBarycenterEarth	85
8.1.2	LALBarycenter	86
8.2	TEMPO2	86

8.3	Principles of Fits	87
8.3.1	Fitting Methods	87
8.4	Earth Fit	89
8.5	Emission Time Fit	99
8.5.1	Fitting Procedure	100
8.5.2	Definitions of Variables	103
8.5.3	Direction-independent terms	104
8.5.4	Difference-independent terms	105
8.5.5	Time-independent terms	105
8.6	Application to Loosely Coherent Searches	106
Chapter IX Conclusions		108
Bibliography		109

List of Tables

3.1	Summary of upgrades to LIGO systems	21
4.1	Types of Continuous Wave Search	28
4.2	Parameters for line detection	31
4.3	Parameters used for each stage of the O1 analysis	32
4.4	Maximum phase changes (in degrees) due to template mismatch for 1800s SFTs [31, Tab. 1]	35
5.1	S6 MDC Parameters	40
5.2	Pipeline search parameters	42
6.1	Summary of physical environment monitors [81]	57
6.2	Average ratio of SNRs, $\frac{\text{Run 2}}{\text{Run 1}}$	66
6.3	Number of templates with SNR above threshold. Cleaned runs show sharp decline in high-SNR points.	66
7.1	Maximum frequency to prevent Doppler shift for sources in the ecliptic plane from occupying multiple bins.	69
8.1	Summary of EarthState and EmissionTime structures	85
8.2	Fitting methods used in Optimx	88
8.3	Orbital periods used in Earth position fit	98

8.4	Parameters used in fits	103
8.5	Term significance analysis. The max error column shows errors when the specified terms are omitted.	106

List of Figures

2.1	Model of a star system made by the University of Michigan Physics Demo Lab	7
2.2	Plus (top) and cross (bottom) linear polarizations for wave traveling into the page	11
3.1	Diagram of Initial LIGO detectors	16
3.2	Advanced LIGO optical configuration	20
3.3	Noise budget for initial and advanced detectors	21
3.4	Sample strain spectrum at LHO from start of O1 showing significant effects from mass suspensions (violin modes), marked in red.	22
4.1	Example of bin realignment used in StackSlide	29
4.2	PowerFlux upper limit validation	30
4.3	Accounting for signal phase in PowerFlux searches	34
4.4	Max SNR of O1 templates over full sky	37
4.5	Weight histogram for outlier near 77.5 Hz line. Note concentration of weight in a few frequency bins.	38
4.6	O1 Upper Limits for each pipeline in low-frequency search	39
5.1	Photon calibrator beam path	41
5.2	Critical frequency second derivatives	44

5.3	Overall detection efficiency	47
5.4	Detection efficiency inside/outside bands containing known lines [75, Fig. 2].	47
5.5	Detection efficiency for each frequency band [75, Fig. 3].	48
5.6	Difference between injected parameters, and parameters associated with highest SNR point in band [75, Fig. 8].	48
6.1	H1 low-frequency normalized spectrum from 3600s O1 SFTs	50
6.2	H1 low-frequency normalized spectrum from 3600s O2 SFTs	51
6.3	Example of non-stationary line in O1 at 75.5 Hz in raw frequency-time spectrogram	53
6.4	75.5 Hz line in residual spectrogram after subtraction	54
6.5	Spectrogram with 75.5 Hz line removed after cleaning.	55
6.6	Coherence between detector strain channel, and LHO Y end-station vacuum equipment area magnetometer channel. Data are from 60-second SFTs, averaged over one hour, and sampled once per day.	58
6.7	Median SNR across spindowns	59
6.8	Max SNR for each frequency/spindown in test run	60
6.9	Upper limits in search band. Most lines are absent from cleaned runs.	62
6.10	Distribution of time outliers spent in cleaned regions. Error bars estimated as $\Delta N = \sqrt{N}$	63
6.11	Relation between time outliers spent in cleaned regions and spindown. Error bars estimated as $\Delta y = \sigma_y/N$	64
6.12	Relation between time outliers spent in cleaned regions and SNR. Error bars estimated as $\Delta y = \sigma_y/N$	65
7.1	Median binned weight over signal polarizations (ψ) for SFTs of various coherence times. Sources refer to equatorial coordinates.	71

7.2	Mean signal concentration for test data.	73
7.3	Standard deviation of signal bins over SFTs.	74
7.4	Total available H1 observation time for each T_{coh} during the O1 run.	75
7.5	Average T_{coh} chosen for sources on the equator for various wave orientations during the O1 run	75
7.6	Errors in recovered parameters with respect to injected values	77
7.7	Comparison of values for SNR and UL between injection runs	78
7.8	Validation of upper limits. Ideally, 95% of injections would lie above red line.	79
7.9	Comparison of values for SNR and UL between runs with no injections	80
7.10	Wandering line at 99.975 Hz after background subtraction	80
7.11	Hypothetical wandering line to study background subtraction	81
7.12	Wandering line after background subtraction	82
7.13	Power sums along stationary frequency. Frequency shifts represent difference from f_0	83
8.1	x-position of Earth in ICRF [37]	90
8.2	Fit residual after including solar year	91
8.3	Fit residual after including half solar year	92
8.4	Fit residual after including period of Jupiter	93
8.5	Fit residual after including period of Saturn	94
8.6	Fit residual after including half period of Jupiter	95
8.7	Fit residual after including period of Moon	96
8.8	Fit residual after including period of Uranus	97
8.9	Emission time-difference function. For a search over a small patch of sky, emission times for each point in the patch can be calculated from $T + \Delta$	99

8.10	Emission time $T(u, p)$ for the start of O1 over the sky.	100
8.11	Emission time-difference $\Delta(u, p)$ for $\Delta\alpha$ at the start of O1.	101
8.12	Emission time-difference $\Delta(u, p)$ for $\Delta\delta$ at the start of O1.	101
8.13	Example skymaps of selected fit terms at the start of O1	102

Abstract

With the advent of the advanced gravitational wave detector era, the Laser Interferometer Gravitational-wave Observatory (LIGO) has made several detections of gravitational waves from coalescing binary black holes and binary neutron stars, but other sources are also predicted to exist. Among these are continuous waves from isolated neutron stars. This type of signal is expected to be significantly weaker than the sources observed so far, but to last for years. In this work we discuss the efforts and obstacles involved in searching for continuous gravitational waves. In particular, we give details of the PowerFlux analysis pipeline and evaluate its performance compared to other pipelines. We also discuss searches in the first observing run of Advanced LIGO, including methods for mitigating the effects of the many spectral noise lines present at low frequencies. Finally, we discuss approximations to the barycentering routines used by LIGO to account for the Earth's motion, approximations that provide justification for more computationally efficient loosely coherent searches.

Chapter I

Introduction

On September 14, 2015 a new window was opened on the universe. The Laser Interferometer Gravitational-wave Observatory (LIGO) made the first-ever direct detection of gravitational waves [10]. Such waves had been predicted by Einstein 100 years ago, but until this date only indirect evidence had been observed. Most notably, after Hulse and Taylor found the first known binary pulsar system [42], changes in its orbit measured over decades precisely matched the predictions from General Relativity of energy lost due to gravitational radiation.

Just as electromagnetic waves can be produced by time-dependent charge dipoles, gravitational waves rely on time-dependent mass quadrupoles. Several objects are predicted to have such quadrupoles strong enough to be accessible to LIGO. The detections so far have come from binary black hole, and binary neutron star mergers, which fall into the category of compact binary coalescences (CBCs). Other potential sources include bursts from supernovae, a stochastic background radiation left over from the Big Bang, and the focus of this work, continuous waves (CWs) from isolated pulsars. The nature of gravitational waves and their detection is described in Chapters II and III, respectively.

Different methods for detecting CWs in LIGO data have been developed. This

work focuses on improving searches carried out with the PowerFlux program, developed by Vladimir Dergachev [30]. The program was created during his graduate work at the University of Michigan, but it has evolved greatly over the years. PowerFlux is a multi-stage quick-look pipeline, designed for all-sky pulsar searches. A description of the PowerFlux algorithms is given in Chapter IV.

After completing analysis of the 6th Science run (S6), PowerFlux participated in a Mock Data Challenge, in which pipelines were compared in their ability to detect artificial signals added to the S6 data after collection [75]. The parameters and results of this challenge, focusing on stages 3 and 4, are given in Chapter V.

The first observing run of Advanced LIGO (O1) was plagued by many spectral noise lines at low frequency, making detections difficult in these regions. Tools have been developed to clean noisy data, reducing false-positives [8]. We investigate the use of cleaning with PowerFlux searches, including validations of upper-limits, in Chapter VI.

Since LIGO is sensitive to waves coming from many different directions, we rely on changes in the signal frequency from Doppler shifts to resolve a source in space. However, Doppler shift depends on source frequency, and at low frequencies causes shifts too small to be well resolved by the standard 30-minute Fourier transforms previously used in the analysis. Therefore we also investigate the use of longer coherence times, which have finer frequency resolution, but pose other challenges. The results are discussed in Chapter VII.

Like other pipelines, PowerFlux makes use of the LIGO Analysis Library (LAL) for common astronomical calculations. However, some of LAL's models are more precise than necessary for portions of our analysis. In particular LALBarycenter, which tracks the position of the detectors relative to the Solar System Barycenter (SSB), is used where a faster, more approximate tool would suffice. In

Chapter VIII, we examine the original model for LALBarycenter, TEMPO2 [34], and suggest a more appropriate calculation.

Some results presented in this work are preliminary, and have not been reviewed yet by the LIGO Scientific Collaboration.

Chapter II

Gravitational Waves

2.1 General Relativity

General Relativity is based on the principle that there is no special frame of reference – fundamental laws should work regardless of position and motion in space. A classic thought-experiment is to imagine observing a ball fall to the floor in a lab at 9.8 m/s^2 . It could be that the lab is on Earth, but based on the ball's motion alone, it's equally possible that the lab is accelerating through space at 9.8 m/s^2 in the opposite direction.

This suggests that being in a gravitational field is equivalent to observing from an accelerating reference frame. Einstein explained this equivalence by describing gravity as a warping of spacetime. The precise shape of spacetime is given by the metric tensor,

$$g_{\mu\nu} = \begin{pmatrix} g_{00} & g_{01} & g_{02} & g_{03} \\ g_{10} & g_{11} & g_{12} & g_{13} \\ g_{20} & g_{21} & g_{22} & g_{23} \\ g_{30} & g_{31} & g_{32} & g_{33} \end{pmatrix}. \quad (2.1)$$

This tensor is required to be symmetric ($g_{\mu\nu} = g_{\nu\mu}$) and non-singular. It describes

the relationship between coordinates. This is analogous to the way unit vectors for non-Cartesian coordinate systems change direction depending on position. In plane-polar coordinates,

$$\hat{\mathbf{r}} = \cos(\theta)\hat{\mathbf{x}} + \sin(\theta)\hat{\mathbf{y}}. \quad (2.2)$$

Similarly, in the presence of a massive object, warping of space changes the path that particles follow, leading to effects like gravitational lensing. The metric tensor gives the information needed to measure distance, and determine how particles move in space.

The line element used to find distances is given by

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu. \quad (2.3)$$

The motion of particles is determined by the geodesic equation,

$$\frac{d}{d\lambda}(g_{\mu\nu}\dot{x}^\nu) - \frac{1}{2}\frac{\partial g_{\sigma\rho}}{\partial x^\mu}\dot{x}^\sigma\dot{x}^\rho = 0, \quad (2.4)$$

where λ is a parameterization of the curve, $\dot{x}^\mu = \frac{dx^\mu}{d\lambda}$, and Einstein summation notation is used (repeated indices are summed). This equation tells us that the acceleration of a particle is intimately connected to the shape of the space around it. John A. Wheeler summarized this equivalence principle (EP) as, “Spacetime tells matter how to move; matter tells spacetime how to curve” [83].

The curvature of a metric space is described by the Christoffel symbols,

$$\Gamma^\gamma_{\alpha\beta} = \frac{1}{2}g^{\gamma\mu}(\partial_\beta g_{\mu\alpha} + \partial_\alpha g_{\mu\beta} - \partial_\mu g_{\alpha\beta}), \quad (2.5)$$

and the Riemann tensor,

$$R^\rho{}_{\sigma\mu\nu} = \partial_\mu \Gamma^\rho{}_{\nu\sigma} - \partial_\nu \Gamma^\rho{}_{\mu\sigma} + \Gamma^\rho{}_{\mu\lambda} \Gamma^\lambda{}_{\nu\sigma} - \Gamma^\rho{}_{\nu\lambda} \Gamma^\lambda{}_{\mu\sigma}. \quad (2.6)$$

Using these, we can write Einstein's equation

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = -\frac{8\pi G_N}{c^4} T_{\mu\nu}, \quad (2.7)$$

where $T_{\mu\nu}$ is the stress-energy tensor

$$T_{\mu\nu} = \left(\begin{array}{c|c} \text{energy density} & \text{energy flux} \\ \hline \text{momentum density} & \text{stress tensor} \end{array} \right). \quad (2.8)$$

Using this, we can determine the properties of space around a given mass/energy distribution. A common analogy for spacetime warping is to imagine a bowling ball on a trampoline (see Fig. 2.1 for a miniature version). The weight of the ball pulls down the surface, drawing in smaller objects. This might lead one to wonder whether the trampoline could be made to vibrate, sending out ripples in the surface far from a massive object.

2.2 Theoretical Derivation of Gravitational Waves

We begin by assuming a small perturbation, $h_{\mu\nu}$, to flat spacetime,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (2.9)$$



Figure 2.1: Model of a star system made by the University of Michigan Physics Demo Lab [72]. An object placed on the surface would fall inward toward the weight in the same way the curvature created by our Sun allows the Earth to orbit.

where $\eta_{\mu\nu}$ is the Minkowski metric,

$$\eta_{\mu\nu} \equiv \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.10)$$

The curvature of the perturbed metric is given by the Christoffel symbols, rewritten as

$$\Gamma^{\gamma}_{\alpha\beta} = \frac{1}{2} (\eta^{\gamma\mu} - h^{\gamma\mu}) (h_{\mu\alpha,\beta} + h_{\mu\beta,\alpha} - h_{\alpha\beta,\mu}). \quad (2.11)$$

For small $h_{\mu\nu}$ though, we can drop higher orders:

$$\Gamma^\gamma_{\alpha\beta} = \frac{1}{2}\eta^{\gamma\mu} (h_{\mu\alpha,\beta} + h_{\mu\beta,\alpha} - h_{\alpha\beta,\mu}). \quad (2.12)$$

From those, we get the Riemann tensor

$$R^\alpha_{\mu\beta\nu} = \frac{1}{2}\eta^{\alpha\delta} (\partial_\mu\partial_\beta h_{\delta\nu} + \partial_\delta\partial_\nu h_{\mu\beta} - \partial_\delta\partial_\beta h_{\mu\nu} - \partial_\mu\partial_\nu h_{\delta\beta}), \quad (2.13)$$

and hence the Ricci tensor

$$R_{\mu\nu} = R^\alpha_{\mu\alpha\nu} = \frac{1}{2} (\partial_\alpha\partial_\nu h^\alpha_\mu + \partial_\mu\partial_\alpha h^\alpha_\nu - \square h_{\mu\nu} - \partial_\mu\partial_\nu h), \quad (2.14)$$

and Ricci scalar

$$R = \eta^{\gamma\lambda} R_{\gamma\lambda} = \partial_\gamma\partial_\lambda h^{\gamma\lambda} - \square h, \quad (2.15)$$

where h is the trace of $h_{\mu\nu}$ and $\square \equiv \partial_\mu\partial^\mu$.

We introduce the trace-reversed strain, $\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$, and choose the Lorentz gauge, $\partial^\mu\bar{h}_{\mu\nu} = 0$. Applying these constraints to Eq. (2.14) and Eq. (2.15) gives

$R_{\mu\nu} = -\frac{1}{2}\square h_{\mu\nu}$ and $R = -\frac{1}{2}\square h$, respectively. We now plug into Eq. (2.7) to find

$$\square h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\square h = \square\bar{h}_{\mu\nu} = -\frac{16\pi G_N}{c^4}T_{\mu\nu}. \quad (2.16)$$

When $T_{\mu\nu} = 0$, this takes the form of the wave equation, suggesting the existence of freely propagating gravitational waves:

$$\left[\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2 \right] \bar{h}_{\mu\nu} = 0. \quad (2.17)$$

2.3 Source Requirements

The question now becomes, what sort of mass configurations will produce gravitational waves? To determine potential sources, we first need to express the stress-energy tensor, $T_{\mu\nu}$, in terms of the mass configuration. Following Cheng's derivation [26], energy conservation requires $D^\mu T_{\mu\nu} = 0$; since $h_{\mu\nu}$ is small, however, this requirement can be approximated with the regular derivative, $\partial^\mu T_{\mu\nu} = 0$. Differentiating the conservation equation with respect to time gives

$$\frac{\partial^2 T^{00}}{c^2 \partial t^2} = -\frac{\partial^2 T^{0j}}{c \partial t \partial x^j} \quad (2.18)$$

where Roman indices are summed only over spatial dimensions. Applying the conservation equation $\partial_0 T^{0j} = -\partial_i T^{ij}$ once again gives

$$\frac{\partial^2 T^{00}}{c^2 \partial t^2} = \frac{\partial^2 T^{ij}}{\partial x^i \partial x^j}. \quad (2.19)$$

We now multiply both sides by $x^k x^l$ and integrate over the source volume:

$$\frac{\partial^2}{c^2 \partial t^2} \int T^{00} x^k x^l d^3 \mathbf{x} = \int \frac{\partial^2 T^{ij}}{\partial x^i \partial x^j} x^k x^l d^3 \mathbf{x} \quad (2.20)$$

$$= \oint \left[x^k x^l \frac{\partial T^{kj}}{\partial x^j} - 2x^k T^{kj} \right] dS + 2 \iiint T^{kl} d^3 x, \quad (2.21)$$

where Eq. (2.21) is obtained through two integrations by parts. The terms in the first integral extend only over the surface of the source, and may be neglected.

Thus, we have

$$\frac{\partial^2}{c^2 \partial t^2} \int T^{00} x^k x^l d^3 \mathbf{x} = 2 \iiint T^{kl} d^3 \mathbf{x}. \quad (2.22)$$

In the Newtonian limit, $T_{00} = \rho c^2$ for mass-density ρ . Substituting into the previous equation, we recognize the form of the quadrupole moment tensor,

$I_{ij} = \int \rho(\mathbf{x}) x_i x_j d^3 \mathbf{x}$, resulting in*

$$\frac{\partial^2 I_{ij}}{\partial t^2} = 2 \int T_{ij} d^3 \mathbf{x}. \quad (2.23)$$

Inside the source, we can write the solution to Eq. (2.16) as

$$\bar{h}_{\mu\nu}(\mathbf{x}, t) = \frac{4G_N}{c^4} \int \frac{T_{\mu\nu}(\mathbf{x}', t - |\mathbf{x} - \mathbf{x}'|/c)}{|\mathbf{x} - \mathbf{x}'|} d^3 \mathbf{x}', \quad (2.24)$$

where the integration is over the volume of mass. In the limit of long-wavelength and long-range, we can approximate the vector differences in Eq. (2.24) as a constant distance r :

$$\bar{h}_{\mu\nu}(\mathbf{x}, t) = \frac{4G_N}{rc^4} \int T_{\mu\nu}(\mathbf{x}', t - r/c) d^3 \mathbf{x}'. \quad (2.25)$$

At this point, we make one more gauge transformation to make the perturbation transverse, $k^\mu h_{\mu\nu} = 0$ for wave vector k^μ , and traceless, $h^\mu{}_\mu = \bar{h}^\mu{}_\mu = 0$. In this gauge, $h_{\mu\nu} = \bar{h}_{\mu\nu}$, so substituting in Eq. (2.23) we have

$$h_{ij}(\mathbf{x}, t) = \frac{2G_N}{rc^4} \frac{\partial^2}{\partial t^2} I_{ij} \left(t - \frac{r}{c} \right). \quad (2.26)$$

This suggests any mass configuration with a time-varying quadrupole moment tensor will produce gravitational waves.

In the transverse-traceless gauge, the polarization of the wave may be written

*Note that we have lowered the indices on both sides. The equation remains valid, since it is linear in the tensors.

(for a wave traveling along the z -axis)

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \epsilon_+ & \epsilon_\times & 0 \\ 0 & \epsilon_\times & -\epsilon_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (2.27)$$

where ϵ_+ and ϵ_\times are the two linear polarization amplitudes for gravitational waves.

An illustration of these polarization modes can be seen in Fig. 2.2.




Figure 2.2: Plus (top) and cross (bottom) linear polarizations for wave traveling into the page. Blue curves show the movement of a ring of mass as the wave passes.

2.4 Expected Sources

In Eq. (2.26), the factor $\frac{2G_N}{c^4} \approx 1.65 \times 10^{-44}$ Newtons⁻¹ indicates that extreme sources are needed to produce a detectable signal. In spite of this high bar, there are a few astronomical objects that are predicted to produce waves strong enough to

detect.

2.4.1 Compact Binary Coalescences

Compact Binary Coalescence (CBC) signals are produced by pairs of black holes (BH–BH), neutron stars (NS–NS), or a combination (BH–NS) that spiral around each other before colliding [21, 27]. As the bodies get closer, their orbital velocities increase, which produces a characteristic “chirp” in the frequency of the emitted waves.

Starting with GW150914, LIGO has made several detections in the advanced detector era, all from this category [10, 9, 12, 13, 14, 15]. In doing so, LIGO has discovered definitive proof of the existence of gravitational waves, and of black holes, neither of which had been directly measured before. With the discovery of binary neutron star merger GW170817 and its electromagnetic counterpart, LIGO was also able to set limits on the speed of gravity, and provide an explanation for observed short gamma-ray bursts. Together, the detections serve as a test of General Relativity. The results so far all include the predicted behavior within their errors.

2.4.2 Bursts

During a supernova, a star goes through rapid density changes as it first collapses, then explodes. If there is any anisotropy during this process, there is the potential to release a burst of powerful gravitational waves. This type of source was once believed to be the most promising for detection [76]. Ott *et al.* predict strains of order 3×10^{-23} , which would be detectable by LIGO in the event of a galactic supernova [59].

2.4.3 Stochastic Background

Just as light from the rest of the universe impinges on the Earth from all directions, it should be possible to detect a sum total of all gravitational radiation. More specifically, we also expect gravitational waves from the Big Bang, analogous to the Cosmic Microwave Background. Such a background could be detected through correlation between the two detector outputs. With a long enough observation time, even a weak signal could be found [16].

2.4.4 Continuous Waves

While CBC and burst signals overcome the inherent weakness of gravity via their large source masses and rapid changes, continuous wave signals offer the possibility of long integration times, which reduce the effect of noise. Since continuous waves are expected to maintain coherence for long periods, adding up the detected strain at specific frequencies will tend to enhance the signal, while the incoherent noise cancels itself.

In 1968, pulsars were identified as rotating neutron stars [70], which have the potential to generate continuous gravitational waves. If a neutron star has an axial asymmetry, *e.g.* a bump on the surface, the star's rotation will give the time-varying quadrupole moment tensor needed to produce waves. For small ellipticity, $\epsilon \equiv \frac{I_{xx} - I_{yy}}{I_{zz}}$, we can write the amplitude of Eq. (2.26) as

$$h_0 \approx \frac{16\pi^2 G}{Rc^4} \epsilon I_{zz} f^2 \quad (2.28)$$

where f is the rotational frequency of the pulsar. There are a number of mechanisms predicted to create ellipticities that would emit gravitational waves [64, 49].

In general, ellipticity scales with the square of the average magnetic field inside

the star [40], but there is uncertainty in the strength of this field. Younger, fast-spinning pulsars are expected to have the strongest fields, and therefore are the most promising sources of continuous gravitational waves.

For non-isolated neutron stars that are accreting matter, observations have shown a mismatch in the angular momentum transferred to the star, and its rotational frequency [25]. One possibility is that the accreted matter is producing gravitational waves, which carry the excess energy [74]. We can set a torque-balance limit on the strength of the gravitational waves by assuming all the missing energy is accounted for by GW emission.

Pulsars have been observed to change frequency suddenly, or “glitch”, gradually returning to their original frequency band. This is believed to be due to transfers in angular momentum between the star’s core and its crust [39]. Along with the initial burst of waves from the glitch, a continuous wave signal is possible as the star recovers.

Based on estimates of the neutron star equation of state, Krastev *et al.* put the maximal value of ϵ at 2.4×10^{-6} , and the resulting h_0 at $[0.4 - 1.5] \times 10^{-24}$ [47]. This falls above the upper limits for emission strength reported in the O1 search, 4×10^{-25} [11], suggesting the detector is capable of picking up waves from a spinning neutron star if a suitable source can be found. This type of signal is the focus of this work, and further details of search efforts will be given in Chapter IV.

Chapter III

LIGO

The previous chapter discussed the sources of gravitational waves, so now we turn to the issue of detecting them. A number of methods have been proposed, including Weber’s pioneering work with resonant masses [77, 78] and pulsar timing suggested by Hulse, Taylor, and Weisberg [42, 79]. This work focuses on the efforts of the Laser Interferometer Gravitational-wave Observatory (LIGO). The idea of using lasers to sense gravitational strain was first proposed by Gertsenshtein and Pustovoit in 1963 [38]. Rainer Weiss expanded on the idea in 1972 [80], developing the models LIGO would be based on. Conceptually, the detector is similar to an “L”-shaped laser interferometer. As a gravitational wave passes, one arm is compressed, while the other extends. This causes the photons in the shorter arm to arrive at the photodetector sooner, producing an interference pattern. In practice, the detectors are more complicated, employing resonant cavities, and noise-reduction techniques [3].

The LIGO detectors sense GWs from all directions, with varying sensitivity. To localize signals, multiple observations are needed, either from separate observatories, or separate times due to the movement of the Earth. The latter option is not as helpful for short-lived events, so LIGO uses multiple detectors: in Hanford, WA




Figure 3.1: Diagram of Initial LIGO detectors [29, Fig. 5]

(LHO), and in Livingston, LA (LLO). LIGO was joined by Virgo and the European Gravitational Observatory in August, 2017 adding a third detector to the network.

3.1 LIGO Design Principles

The LIGO detectors are based on a Michelson interferometer with Fabry-Perot resonant cavities (Fig. 3.1). In such a setup, a laser, which produces coherent light, is sent into a beam splitter and then through two perpendicular arms. At the end of each arm, a mirror reflects the light back into the beam splitter, and combines the two beams at a photodetector. With the addition of input mirrors in the arms, the light can be made to circulate in the arms many times before exiting.

For a mathematical description of the interferometer, we follow Saulson's

derivation [65]. Suppose the electric field of the input light is

$$E_{\text{in}} = E_0 e^{i(2\pi ft - kx)}. \quad (3.1)$$

An ideal 50-50 beam splitter has a reflection coefficient of $R = 1/\sqrt{2}$ and transmission coefficient $T = i/\sqrt{2}$, so the beams traveling along the x - and y -axes have electric fields

$$\begin{aligned} E_x &= (iE_0/\sqrt{2})e^{i(2\pi ft - k_x x)} \\ E_y &= (E_0/\sqrt{2})e^{i(2\pi ft - k_y y)} \end{aligned} \quad (3.2)$$

After reflecting at the end mirror, each electric field gets a factor of -1 . When recombined at the beam splitter, the electric field at the output is

$$\begin{aligned} E_{\text{out}} &= (-iE_0/2)e^{i(2\pi ft - 2k_x L_x)} + (-iE_0/2)e^{i(2\pi ft - 2k_y L_y)} \\ &= (-iE_0/2)e^{2\pi i ft} (e^{-2ik_x L_x} + e^{-2ik_y L_y}) \\ &= (-iE_0/2)e^{i(2\pi ft - k_x L_x - k_y L_y)} (\exp[i(k_x L_x - k_y L_y)] + \exp[-i(k_x L_x - k_y L_y)]) \\ &= -iE_0 e^{i(2\pi ft - k_x L_x - k_y L_y)} \cos(k_x L_x - k_y L_y). \end{aligned} \quad (3.3)$$

Since the power $P \propto |E|^2$, we find the time-averaged power delivered to the photodetector is

$$P_{\text{out}} = P_{\text{in}} \cos^2(k_x L_x - k_y L_y). \quad (3.4)$$

This implies that the power measured at the photodetector depends on the difference in phase accumulated along each of the arms. This derivation is a simplified description of the interferometer's function. Its response to gravitational waves is discussed in the next section.

3.2 Initial LIGO

The LIGO detectors' arms are each 4 km in length, with light-storage times of about 1 ms, due to circulation in resonant cavities: $\frac{(1 \text{ ms}) \times c}{8 \text{ km}} \approx 37$ round trips. The differential arm-length (DARM) is calibrated to create a nearly dark fringe at the photodetector. A linearly polarized gravitational wave compresses space in one direction, while expanding it in the perpendicular direction, leading to a change in DARM. The photodetector registers this difference as a change in intensity, allowing us to detect the wave.

One might object that the spatial stretching also changes the wavelength of the light beams, canceling the change in travel distance, and resulting in no phase change. This neglects the fact, however, that the speed of light remains constant, and the beam in one arm will complete the round trip faster than the other [66]. As an example, consider a GW with strain $h(t)$, traveling along the z -axis, and aligned in polarization with the detector. The metric will be

$$ds^2 = -c^2 dt^2 + [1 + h(t)] dx^2 + [1 - h(t)] dy^2 + dz^2. \quad (3.5)$$

For light, $ds^2 = 0$, so the proper time for a round-trip along the x-axis will be

$$\tau_x = \frac{2}{c} \int_0^L \sqrt{1 + h(t)} dx. \quad (3.6)$$

Since $h(t) \ll 1$, $\sqrt{1 + h(t)} \approx 1 + \frac{1}{2}h(t)$ and

$$\tau_x \approx \frac{L}{c} [2 + h(t)]. \quad (3.7)$$

Similarly, for the y-axis,

$$\tau_y \approx \frac{L}{c} [2 - h(t)]. \quad (3.8)$$

The phase of the signal along each axis is

$$\phi_r = \omega\tau_r - 2\pi \frac{r(ds/dr)}{\lambda(ds/dr)}. \quad (3.9)$$

As stated above, the spatial component is unchanged, due to cancelation of the distance and wavelength changes. Therefore, the total phase difference is

$$\Delta\phi = \frac{2}{c}h(t)\omega L = \frac{4\pi L}{\lambda}h(t). \quad (3.10)$$

3.3 Advanced LIGO Upgrades

The initial upgrades to the detector that constituted the first stages of Advanced LIGO were completed in 2015. A diagram of the optical configuration is shown in Fig. 3.2, and some of the physical parameters that changed between Initial and Advanced LIGO are shown in Table 3.1.

The increased laser power gives LIGO access to higher frequencies by reducing shot noise:

$$(\Delta I)^2 \propto I \quad (3.11)$$

for laser intensity I . The extra power, though, increases radiation pressure noise, so to compensate, the test masses were made heavier. The increased number of pendulum stages allow for better seismic isolation, reducing the low-frequency noise. Improved mirror finesse results in longer light storage times, increasing sensitivity overall.

In total, the Advanced LIGO upgrades will give a 10-fold strain sensitivity


Table 3.1: Summary of upgrades to LIGO systems. Advanced LIGO storage time is calculated from cavity finesse: $\tau = \frac{2L}{c} \frac{F}{\pi}$

	Initial [29]	Advanced [3]
Laser Power (W)	6	125
Test Mass Mass (kg)	12	40
Test Mass Diameter (cm)	28	34
Input Mirror Transmission	3%	1.4%
Pendulum Stages	1	4
Light Storage Time (ms)	0.88	3.8


modifications, like charging noise from excess electrons on the masses.

Environmental noise results from sources outside the detector, such as seismic noise.

See Fig. 3.3 for a comparison of the ideal expected noise budgets between the initial and advanced detectors.



(a) Initial LIGO noise budget from Weiss and Shoemaker [82]



(b) Advanced LIGO noise budget from LSC [3]

Figure 3.3: Noise budget for initial and advanced detectors

Most of these sources contribute to a “noise floor” that varies slowly in frequency, but some create sharp spectral artifacts with strain far above the

surrounding points. These “lines” can create spurious outliers during analysis. There are a number of sources of these lines, some of which will be detailed in Chapter VI. Here we discuss the example of “violin modes”, the resonant frequencies of the mirror suspensions [52].

Violin modes are vibrationally resonant frequencies of the silica fibers that suspend the test masses. The modes can be excited by coupling with thermal noise. Most LIGO searches are carried out in the frequency range 50-2000 Hz, which contains the resonant frequency of the test mass suspensions, and several harmonics. An example of the strain-frequency distribution using a spectral average from many 7200-second duration short Fourier transforms (SFTs) for LHO at the start of O1 is shown in Fig. 3.4. Notice the large sets of spikes near 500, 1000, 1450, and 1950 Hz. These correspond to the fundamental, 2nd, 3rd, and 4th harmonics of the suspension resonance. Noise lines like these can cause problems for continuous gravitational wave searches, since sources are also very narrowband.




Figure 3.4: Sample strain spectrum at LHO from start of O1 showing significant effects from mass suspensions (violin modes), marked in red.

Efforts to reduce the noise in the detector are ongoing, and improvements are still

being made to bring the detectors up to design specifications. With every decrease in the noise floor, our reach out toward potential sources extends a little further.

Chapter IV

Detecting Continuous Waves

The sources observed by LIGO so far have been collisions between pairs of black holes and neutron stars, referred to as compact binary coalescences (CBCs). These signals are some of the loudest we expect to see, with strain magnitude of order 10^{-21} . Continuous wave signals on the other hand are expected to be significantly weaker, with current upper limits of order 10^{-25} [11]. Promising sources of this type of gravitational wave are spinning neutron stars with axial asymmetries.

4.1 Neutron Stars as GW Sources

Neutron stars are compact objects that result from supernova collapse of stars not massive enough to create a black hole. Some neutron stars spin rapidly; charged particles near the surface are accelerated by the powerful magnetic field surrounding the star, and emit light in the star's magnetosphere. Such stars are called pulsars, after the pulsed signals received by Earth due to the misalignment of their magnetic and rotational poles [46].

The possibility of neutron stars was first postulated by Landau, based on balancing gravitational forces with outward nuclear pressure derived from quantum

mechanics [48]. Baade and Zwicky suggested that such an object could result from supernovae [20]. When a star undergoes supernova, the neutron star remnant retains much of the angular momentum and charge, but the significantly smaller radius ($\mathcal{O}(35\,000)$ times smaller) creates a far more powerful, rapidly rotating magnetic field. This field will accelerate nearby charged particles as the star rotates. Accelerating charges emit light, and if the field axis has some misalignment from the spin axis, the light traveling toward Earth will be broken into regular short bursts, characteristic of a pulsar. However, most pulsars have been observed to slow their rotation gradually, indicating a loss of rotational energy. Some loss can be accounted for in the light emitted, but this is not enough to match observations [61].

If this energy deficit is to be made up with gravitational wave emission, the neutron star must have some non-axisymmetry. Neutron star structure is broadly divided into three regions: the surface, a crust, and a core. The surface is relatively low-density, at 10^6 g/cm³, with a strong magnetic field. The actual strength of the field depends on the rotational properties of the star:

$$B \propto \sqrt{P\dot{P}}, \tag{4.1}$$

where P and \dot{P} are the rotation period and its time-derivative. Young pulsars recently formed from supernovae tend to have a high spindown, corresponding to a large B-field, $\mathcal{O}(10^{12}$ Gauss). Millisecond pulsars have high rotation rate, low spindown, and low B-field, $\mathcal{O}(10^8$ Gauss). The crust of a neutron star is thought to be of similar density to the surface with a thickness of 1–2 km, but made up almost entirely of neutrons. The core contains 99% of the total mass, in a homogeneous neutron-proton fluid. Here, the density rises to 10^{14} g/cm³ [24].

One possible GW emission from pulsars is from a defect in the surface that creates a time-varying quadrupole moment. An analysis of the neutron star

equation of state by Ushomirsky *et al.* suggests that temperature gradients within the crust could create a quadrupole moment that would emit gravitational waves [73]. Temperature governs the distribution of electrons within the crust, so a laterally-varying temperature leads to asymmetric distribution of mass. As described in Chapter II, ellipticities in a rotating object lead to GW emission.

4.2 Inherent Difficulty of Detection

By analyzing known pulsars, estimates can be made of expected strain. Zimmerman calculates the strains from the Crab and Vela pulsars as $\mathcal{O}(10^{-26})$, based on misalignments between the spin axis and symmetry axis [85]. Even at Advanced LIGO design sensitivity the noise floor is not expected to go below $\mathcal{O}(10^{-24} \text{ Hz}^{-1/2})$ [54]. How then do we expect to detect such signals? The key is integration: continuous wave signals are expected to be coherent, maintaining a consistent phase evolution over long periods. Noise sources, on the other hand, are generally incoherent and/or transient. By adding up a signal over several weeks or months, we increase our sensitivity to gravitational waves, while decreasing the influence of uncorrelated noise.

4.3 Search Inputs

During operation, the detectors produce a time-series of strain measurements. Since continuous wave signals are long lasting, and change frequency only gradually, we typically divide this data into segments of length T_{coh} , called the coherence time, and perform a discrete Fourier transform to extract the frequency information. The

number of samples in a Fourier transform is inversely related to the resolution,

$$\Delta f = \frac{1}{T_{\text{coh}}}. \quad (4.2)$$

For the data used with PowerFlux, the typical T_{coh} is 30 minutes, with time segments overlapping by 50%. We apply a Hann window [57] to each time-series before performing the Fourier transform. These are referred to as Short Fourier Transforms (SFTs). The SFTs covering a particular run, such as O1, are divided into frequency chunks 29 Hz in width, overlapping by 4 Hz, and SFTs for each frequency band are joined. This allows us to load a single data set for any given search.

4.4 Search Classes

There are three broad classes of CW search performed by the LSC: targeted, directed, and all-sky [62, 64]. For known sources, like the Crab Pulsar, electromagnetic observations give us precise sky positions and rotation frequency, related to GW frequency by

$$f_{\text{GW}} = 2f_{\text{rot}}. \quad (4.3)$$

With this information, we're able to run *targeted* searches over a narrow band of frequencies and sky positions. In some cases, we know the location of an object that could be a pulsar, such as a supernova remnant, but we have yet to make any electromagnetic observations. We can do a *directed* search in the area of the object over a broad frequency spectrum. The final type of CW search makes no assumption of sky direction or frequency. These *all-sky* searches are computationally expensive, but offer the possibility of detecting previously unknown, or unexpected sources. These search classes are summarized in Table 4.1.

Table 4.1: Types of Continuous Wave Search

Class	Position Known?	Frequency Known?
Targeted	✓	✓
Directed	✓	×
All-Sky	×	×

4.5 PowerFlux

There are several methods used by the CW group to search for continuous gravitational waves. The pipeline used in this work is PowerFlux [5], similar to StackSlide [55], but more powerful.

4.5.1 StackSlide

The StackSlide pipeline was designed to sum the powers in sequential SFTs to estimate signal strength. SFT frequency bins are laid out in columns, such that each row is constant in frequency, and each column is constant in time. For each column i , the total bin-shift due to source spindown and Doppler shift is calculated, and the column is shifted that many bins up or down:

$$\Delta b_i = f_0(\hat{\mathbf{e}} \cdot \mathbf{v}(t_i))T_{\text{coh}} + \dot{f}T_{\text{coh}}t_i, \quad (4.4)$$

where T_{coh} is the coherence time of the SFT, f_0 is the frequency of the source at t_0 , $\hat{\mathbf{e}}$ is the sky direction vector, $\mathbf{v}(t_i)$ is the detector velocity vector (in units of c), \dot{f} is the frequency derivative, and t_i is the median time for SFT i . The total power for a template can then be summed from a single row. See Fig. 4.1 for an example.

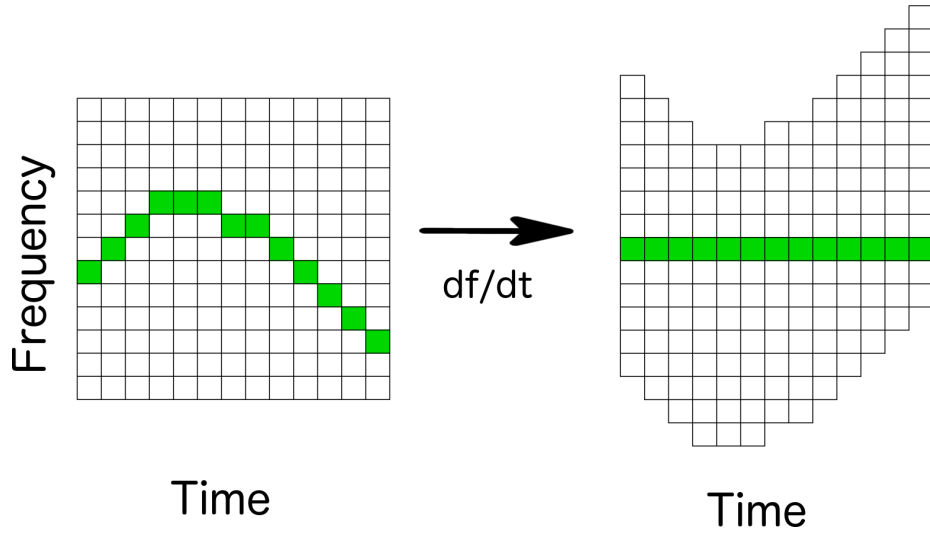


Figure 4.1: Example of bin realignment used in StackSlide. Doppler shift and spin-down are calculated for each time, and the total \dot{f} used to adjust the frequencies.

4.5.2 Power Sum

The primary difference between StackSlide and PowerFlux is that PowerFlux weights SFT data based on detector sensitivity to a signal template at each time and frequency [7]. The signal estimator is defined as

$$R = \frac{2}{T_{\text{coh}}} \left(\sum_i \frac{[(F_\psi^i)^2]^2}{(\tilde{P}_i)^2} \right)^{-1} \sum_i \frac{(F_\psi^i)^2 P_i}{(\tilde{P}_i)^2}, \quad (4.5)$$

where F_ψ^i is the antenna pattern for polarization ψ , P_i is the power in SFT bin i , and \tilde{P}_i is an estimate of the power spectral density of noise in the bin. This weighting scheme results in SFTs with high noise and poor antenna response being devalued.

4.5.3 Software Injections

Because we have yet to detect a continuous wave signal, the best way to test our ability to detect them is with fake data, where a generated signal has been added to

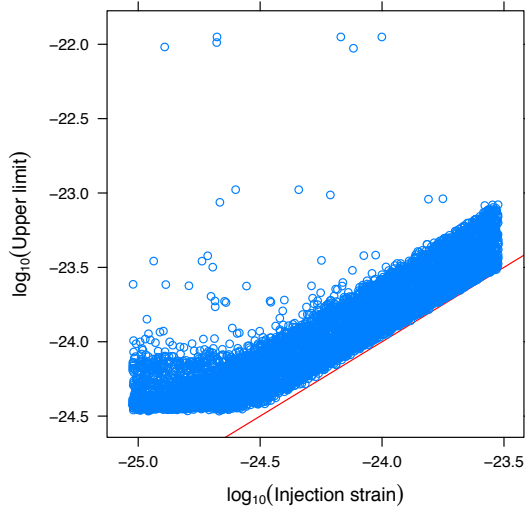


Figure 4.2: PowerFlux upper limit validation. Each point represents a separate injection in the 50-200Hz frequency range. Each established upper limit (vertical axis) is compared against the injected strain value (horizontal axis, red line) [11, Fig. 1].

the detector’s strain. PowerFlux includes tools for injecting, then running a search for, fake signals. This allows us to test signals of varying parameters to quantify our detection capabilities. In particular, PowerFlux sets 95% upper limit confidence intervals on the power of actual signals. Injections are regularly used to verify these limits. An example of the UL validation for the O1 All-sky search is shown in Fig. 4.2. Upper limits are tuned such that $h_{UL} > h_0$ for 95% of injections.

4.5.4 Line Detection

PowerFlux includes a test for sharp, isolated strain in a single bin used to remove lines [5]. The procedure is as follows:

1. Sort bins from a narrow search band in descending magnitude, $z[i]$
2. Set median $M = z[N_{\text{bins}}/2]$
3. Set $Q_{\text{lines}} = z[N_{\text{lines}}]$

Table 4.2: Parameters for line detection

Parameter	Value
N_{bins}	501
N_{lines}	5
N_{little}	30
N_{clust}	5

4. Set $Q_{\text{most}} = z[N_{\text{little}}]$
5. For each bin (Pass 1):
 - (a) If $z[i] > Q_{\text{most}}$ mark as **HIGH**
 - (b) If $z[i] > Q_{\text{lines}}$ mark as **CANDIDATE**
 - (c) If $z[i] > 2Q_{\text{most}} - M$ mark as **VERY HIGH**
6. For each bin (Pass 2), if there are at least N_{clust} **HIGH** bins in a row, mark all as **CLUSTERED**
7. Any bins marked **CANDIDATE** and **VERY HIGH**, but not **CLUSTERED** are vetoed.

The parameter values currently in-use are listed in Table 4.2

4.5.5 Skymarks

Since saving every point searched would use an excessive amount of disk space, PowerFlux includes the ability to divide the sky into regions, with only the highest-SNR point saved for each. These regional designations are referred to as skymarks. In the initial search stage, the sky is divided into bands according to the z-component of sky direction, e_3 . Another form of skymark is the vetos (see Section 4.7), which remove points showing characteristics of noise lines, leaving the other bands clean.

Table 4.3: Parameters used for each stage of the O1 analysis. For all stages > 0 , the loosely coherent pipeline was used. [11, Tab. 1]

Stage	Instrument sum	Phase coherence (rad)	Spindown step (Hz/s)	Sky refinement	Frequency refinement	SNR increase (%)
0	semi-coherent	NA	1×10^{-10}	1	1/2	NA
1	incoherent	$\pi/2$	1×10^{-10}	1/4	1/8	20
2	coherent	$\pi/2$	5×10^{-11}	1/4	1/8	10
3	coherent	$\pi/4$	2.5×10^{-11}	1/8	1/16	10
4	coherent	$\pi/8$	5×10^{-12}	1/16	1/32	7

4.6 Stages of Analysis

PowerFlux is a multi-stage pipeline, which marks outliers for followup in subsequent stages. Each stage refines the search region, and imposes stricter requirements on coherence. One of the key parameters defining the stages of analysis is the degree of coherence required. Increasing coherence emphasizes errors in template parameters, meaning that more searches must be carried out for the same space. A correct template, however, will gain significant SNR with more coherence. Due to computational costs, fully coherent searches are only performed in the later stages of followup, when few candidates remain. The parameters for the various stages in the O1 search are shown in Table 4.3.

4.6.1 Semi-coherent

In the initial PowerFlux stage, phase changes from one SFT to the next are disregarded. Coherence is maintained within each SFT, but the phase with the maximum SNR is selected in each case. By ignoring phase change between SFTs, semi-coherent searches can account for mismatch from the true signal template. This

approach is illustrated in Fig. 4.3a. The power estimate for a semi-coherent sum is

$$P = \sum_{l=1}^N |a_l|^2 \quad (4.6)$$

for N SFTs with signal amplitude a_l .

4.6.2 Fully Coherent

The most stringent requirement for coherence is to maintain phase over the full observation period. A downside is that any mismatch in template parameters could lead to a significant penalty in SNR. Many searches are required to cover the dense parameter space, so to reduce computational costs, full coherence is only used for candidates that have passed followup requirements. The power estimate for a fully coherent sum is

$$P = \sum_{k,l=1}^N a_k^* a_l e^{i(\phi_k - \phi_l)}, \quad (4.7)$$

where ϕ is the phase of the signal.

4.6.3 Loosely Coherent

The loosely coherent method [31, 32, 33] was introduced into PowerFlux during the S5 search, and has been used in the searches that followed [4, 6, 11]. The concept is illustrated in Fig. 4.3. In an initial PowerFlux search, change in signal phase across SFTs is disregarded (Fig. 4.3a), but a true signal would maintain coherence over the span of the search (Fig. 4.3b). However, if the template's frequency does not precisely match the signal's, the phase can drift in a systematic fashion (Fig. 4.3c). The loosely coherent method is designed to account for this.

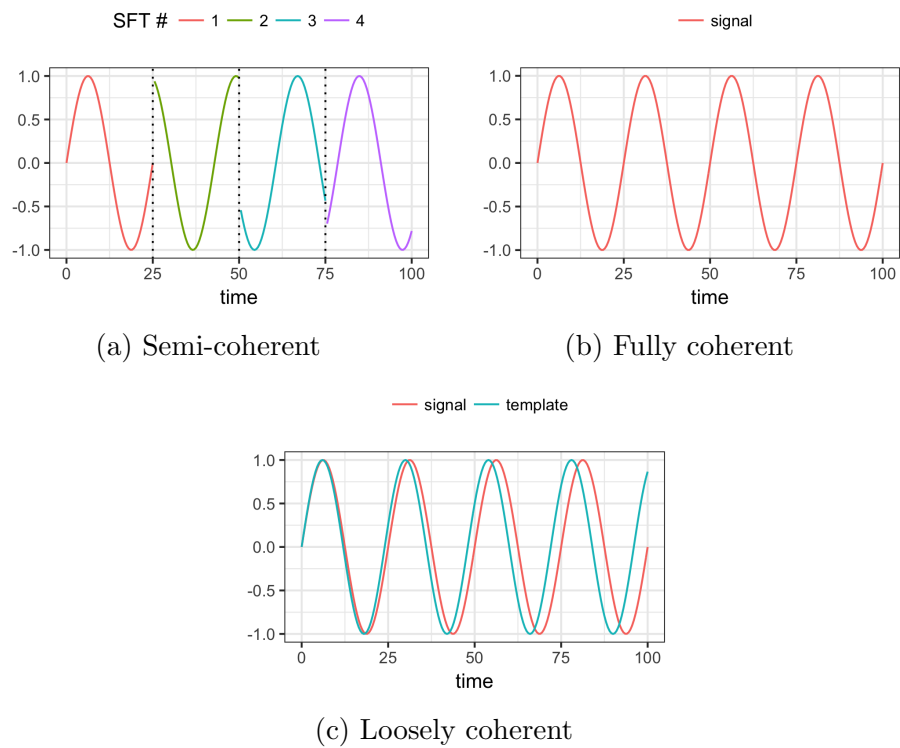


Figure 4.3: Accounting for signal phase in PowerFlux searches. Curves show true/assumed wave amplitude in time.

Table 4.4: Maximum phase changes (in degrees) due to template mismatch for 1800s SFTs [31, Tab. 1]

Phase shift cause	100 Hz	500 Hz	1000 Hz	2000 Hz
Frequency mismatch of $\Delta f = 0.1/\Delta t$	36	36	36	36
Sky position mismatch of $\Delta r = 1^\circ$	1.1	6	11	23
Spindown mismatch of $\Delta \dot{f} = 10^{-12} \text{ Hz s}^{-1}$ for $T = 1 \text{ yr}$	20	20	20	20
Source modulation for $\rho = 1$ and $r = 0.1 \text{ AU}$	0.1	0.6	1	2

Mathematical Statement

In a typical CW search, the detector output is broken into N segments and Fourier transformed. In the ideal case, we want to find a signal that resides in a single bin, $\{a_k\}_{k=1}^N$ in the k th transform, with phase $\{\phi_k\}_{k=1}^N$. For a fully coherent signal, we would have $\phi_{l+1} - \phi_l = 0$. However, if the template does not match the signal exactly, the phase would change from one SFT to the next in a systematic way. The power estimate for the loosely coherent sum is [31, Eq. 10]

$$P = \sum_{k,l=1}^N a_k^* a_l \left(\frac{\sin \delta}{\delta} \right)^{|k-l|}, \quad (4.8)$$

where δ is the phase error due to template mismatch. Some example values for δ are given in Table 4.4.

4.7 Vetos

To reduce the effect of noise lines, a number of methods have been developed to identify and/or eliminate spurious outliers. These vetos create special skymark designations that prevent anomalous templates from being selected for followup. If a template matches one of these vetos, it will not be included in the usual skybands, and may be discarded or examined separately.

4.7.1 S Parameter

The S parameter, introduced in the S4 analysis [7], is a measure of how stationary a template is in detector frequency. It is defined as

$$S = \dot{f} + [(\boldsymbol{\Omega} \times \mathbf{v}_{\text{Earth}}/c) \cdot \hat{\mathbf{n}}] \hat{f}_0, \quad (4.9)$$

where $\boldsymbol{\Omega}$ is the Earth’s angular velocity, $\mathbf{v}_{\text{Earth}}$ is the average velocity of the Earth for the current SFT, and \hat{f}_0 and \dot{f} are the frequency and frequency derivative of the potential signal. See Fig. 4.4 for an example of the S Parameter applied to O1 data.

Templates with $|S| < S_{\text{large}}$ are considered at-risk for contamination, and for the S4 run, were not included in upper limits. By default,

$$S_{\text{large}} = \frac{N_{\text{det}}}{T_{\text{obs}} T_{\text{coh}}}, \quad (4.10)$$

where N_{det} is the minimum number of occupied bins for detection, and T_{obs} is the total observing time of the run. However, in practice the limit needs to be increased to remove outliers due to pervasive combs.

4.7.2 Line Response Veto

Similar to the S Parameter Veto, the Line Response Veto introduced in the S5 and S6 analyses [4, 6] looks for templates that remain stationary in detector frequency. However, this veto sets a limit on the fraction of weight accumulated in a given binwidth over the run. This gives more precision than the S parameter approximation, by vetoing only points that are specifically inflated by a narrow frequency source.

The procedure is as follows:




Figure 4.4: Max SNR of O1 templates over full sky. Frequency change in detector frame depends on both source spindown and frequency (via Doppler modulation). Note enhancement to the right of a chosen S Parameter line. These templates are stationary in detector frequency, and potentially pick up noise lines.

1. For each SFT, calculate the weight $w_i = T_i w_d$, where T_i is the background strain calculated for SFT i , and w_d is the overall weight for detector dataset d
2. For each SFT, calculate the total bin change Δb_i over the integration. This is the same as Eq. (4.4)
3. Calculate the offset $b_0 = (\sum_i w_i \Delta b_i) / (\sum_i w_i)$
4. Calculate adjusted bin change, $\Delta b'_i = \Delta b_i - b_0$
5. Calculate fraction of weight with $|\Delta b'_i| < b_{\text{tol}}$

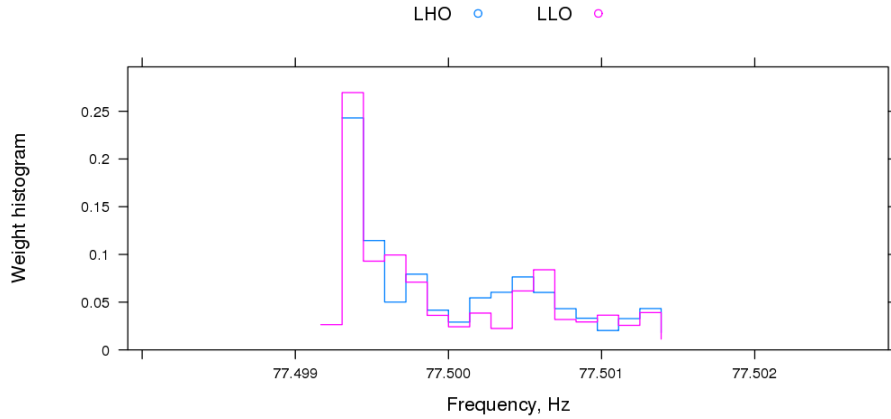


Figure 4.5: Weight histogram for outlier near 77.5 Hz line. Note concentration of weight in a few frequency bins.

For recent searches, the veto threshold was at least 5% weight in any 3 bins. An example weight histogram is shown in Fig. 4.5.

4.8 Pipeline Performance

A comparison of the upper limits set by the various CW pipelines is shown in Fig. 4.6. PowerFlux offers upper limits for both circularly polarized, and linearly polarized waves, which tend to bracket the results given by other pipelines. By distinguishing these two, PowerFlux is able to set the strictest upper limits in most cases. The next chapter addresses PowerFlux’s performance compared to other continuous wave pipelines in detection capability. In general, PowerFlux is significantly more computationally efficient than other search techniques, while still maintaining good detection efficiency.




Figure 4.6: O1 Upper Limits for each pipeline in low-frequency search [11, Fig. 21]. PowerFlux’s best/worst-case polarizations tend to bracket other pipeline results.

Chapter V

S6 Mock Data Challenge

As we have yet to make the first detection of continuous gravitational waves, we have no real data with which to compare the sensitivity of different detection algorithms. Therefore, to fill this void, the LIGO collaboration has added fake signals to its data sets to model the detector’s response to various theoretical sources.

A number of such “signal injections” were added to data from the 6th Science run of LIGO (S6). The Mock Data Challenge (MDC) was designed to compare the CW pipelines in their ability to detect these signals. The challenge was divided into 4 stages, summarized in Table 5.1, though this work concerns only stages 3 and 4, the results of which were published in [75].

Table 5.1: S6 MDC Parameters

Stage	No. of Inj.	Frequency Rng. (Hz)	Visibility
1	21	200–240	Open
2	184	100–500	Open
3	1561	40–2000	Open
4	1550	40–2000	Blind

5.1 Hardware Injections

Some of the signal injections performed by LIGO use the detectors' Photon Calibrators (PCAL) to apply radiation pressure to the mirrors (Fig. 5.1). During S6, eleven such CW signals were injected, including one blind injection for which parameters were not immediately released [51]. Hardware injections give a more accurate test of the sensitivity of our analysis pipelines, but they run the risk of contaminating true signals. For high-statistics studies, injections are inserted into data copies after collection, using software tools.




Figure 5.1: Photon calibrator beam path [44, Fig. 2]. Laser applies varying pressure to the test mass, simulating a gravitational wave.

5.2 Software Injections

Since the MDC required a large sample size, it made more sense to use software injections to avoid over-contamination of the data. Sets of data were created with generated signals added on top of the real detector output for the S6 run. From these, 30-minute SFTs were created for analysis by the various CW pipelines.

Table 5.2: Pipeline search parameters [75]. In cases with multiple values, initial stage and low frequency parameters are shown.

Pipeline	T_{coh} (h)	δf (Hz)	$\delta \dot{f}$ (Hz/s)
PowerFlux	0.25	2.78×10^{-4}	2×10^{-10}
Time-domain \mathcal{F} -statistic	48	5.79×10^{-6}	3.3×10^{-11}
Frequency Hough	0.284	9.76×10^{-4}	2.4×10^{-11}
Sky Hough	0.5	5.55×10^{-4}	1.37×10^{-11}
Einstein@Home	60	3.61×10^{-6}	1.16×10^{-10}

5.3 Pipelines

The pipelines participating in the challenge were PowerFlux, Time-domain \mathcal{F} -statistic, Frequency Hough, Sky Hough, and Einstein@home. With the exception of PowerFlux, detailed previously, these are briefly described below, with some differences in search parameters summarized in Table 5.2.

5.3.1 Time-domain \mathcal{F} -statistic

For signal parameters \mathcal{A} , and Doppler modulation λ , the signal model is

$$h(t; \mathcal{A}, \lambda) = \mathcal{A}^\mu h_\mu(t; \lambda), \quad (5.1)$$

and we can define the maximum likelihood as

$$\mathcal{F} \equiv \max_{\mathcal{A}} \log \mathcal{L}(x; \mathcal{A}, \lambda) = \frac{1}{2} x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (5.2)$$

where $\mathcal{M}^{\mu\nu}$ is the Fisher matrix for the parameters \mathcal{A} , and x_μ is the detected strain.

The detection statistic is given by

$$2\mathcal{F} = \mathcal{A}_{\text{ML}}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}_{\text{ML}}^\nu, \quad (5.3)$$

where $\mathcal{A}_{\text{ML}}^\mu \equiv \mathcal{M}^{\mu\nu} x_\nu$. The data are split into chunks of 2 sidereal days each, and $2\mathcal{F}$ is calculated. Candidates above a certain threshold are checked for coincidence within each 2-day period. [2, 18]

5.3.2 Frequency/Sky Hough

The Hough pipelines digitize SFTs by using a power threshold to set each bin to either 0 or 1. For each template, a path is plotted through this peakgram, and the bins are summed and compared to Gaussian noise. Sky Hough groups candidates by frequency and sky location, and checks for coincidence. Coincident candidates are clustered and passed on for followup. Frequency Hough maps points from the peakgram onto a set of templates with high resolution in frequency. Candidates for followup are once again chosen from groups of frequencies and sky locations, but with wider frequency bands than Sky Hough. [1, 5, 17, 19]

5.3.3 Einstein@Home

The Einstein@Home pipeline uses distributed processing on volunteers' computers to perform searches of LIGO data. Before analysis, SFTs are cleaned of known noise lines, replacing the removed frequency bins with Gaussian noise. The data are split into segments of a few days each, and $2\mathcal{F}$ values calculated. These values are averaged over all segments, and ranked according to the logBSGL statistic. [45, 69]

5.4 Challenge Parameters

For each injected signal, a window for each parameter is given to be searched: 0.1 Hz in frequency, $\max(2 \times 10^{-9} \text{ Hz/s}, \pm 3 \times \text{true } \dot{f})$ in spindown, and a sky region with radius $30 \text{ degrees} \times \min(200 \text{ Hz/injected frequency}, 1)$. For 25% of the injections, a

braking index n is applied, giving second and third frequency derivatives [60],

$$\ddot{f} = \frac{n\dot{f}^2}{f}, \quad \dddot{f} = \frac{n\dot{f}}{f} \left(2\ddot{f} - \frac{\dot{f}^2}{f} \right). \quad (5.4)$$

None of the pipelines explicitly search for second or higher order frequency derivatives, so the presence of a braking index can degrade detection ability. The critical \ddot{f} that would cause a signal to smear over multiple frequency bins is defined as

$$\ddot{f}_{\text{crit}} = \frac{\delta f}{T_{\text{obs}}^2}. \quad (5.5)$$

Comparisons between the injected braking indices and the \ddot{f}_{crit} for each pipeline is shown in Fig. 5.2.




Figure 5.2: Critical frequency second derivatives [75, Fig. 5]. Few injections exceed the limit for any pipeline.

Each pipeline has requirements that define detection, which were applied to the injected signals. For PowerFlux, we demanded that outliers survive the first followup search stage, which requires coincident signals in both detectors, and increased SNR from the combined data. The detection efficiency (no. detected/no.

injected) is measured as a function of sensitivity depth,

$$d = \frac{\sqrt{S_h(f)}}{h_0}, \quad (5.6)$$

where h_0 is the injected strain amplitude, f is the injected frequency, and S_h is the harmonic sum of the power spectral density over the detectors at that frequency. This measure quantifies the difficulty of detection by accounting for the level of noise that may cover an injected signal.

Approximately half the injections were classed as “open,” and the others “blind.” The open injections provided both a search region, and the precise signal parameters that were injected. This allowed for self-assessment and tuning during the first phase of the MDC, to optimize the pipelines’ sensitivity in the injection bands. The second phase used blind injections, meaning that only the search regions were known a priori. True injection parameters were revealed only after the search groups submitted their final results. The blind injections were designed to emulate a real all-sky search, where signal parameters are unknown. The pipelines had similar performance between the open and blind injections, so the efficiency plots shown below do not distinguish between the classes.

5.5 Results

There are two measures used for the success of each pipeline: the detection efficiency, and the accuracy of recovered parameters. Comparisons of the former are shown in Fig. 5.3, and those for the latter in Fig. 5.6. Due to the prevalence of detector artifacts, particularly at low frequencies, the injections were separated into those intersecting known lines (Fig. 5.4a) and those in quiet bands (Fig. 5.4b). Larger Doppler shifts also affect detection capability, so the results are separated by

frequency band (Fig. 5.5).

The results show the value in using multiple strategies to analyze data. Einstein@Home is extremely sensitive in the absence of noise lines, but the cleaning procedure used severely degrades its efficiency for signals that overlap with noise. PowerFlux tends to fall in the middle as far as detection efficiency, but has some of the best parameter recovery. It is also computationally efficient, using the widest spindown step of the pipelines as well as the shortest coherence time. PowerFlux finished the MDC analysis earliest, and the actual O1 search finished 9 months before other pipelines.

PowerFlux achieves optimum sensitivity only after extensive tuning, which was not performed for the MDC data. Instead, the parameters from the previous S6 search were used, which did not include frequencies < 400 Hz. After tuning for O1, PowerFlux proved to be the most sensitive of the CW searches, reporting the lowest upper limits [11].

Given that no real CW signals have been yet detected, it's important to maintain a wide variety of search techniques, increasing our chances of discovery. At the same time, testing against simulated signals offers the opportunity to optimize pipelines, and identify the strengths and weaknesses present in each. The MDC was an important test case heading into the advanced detector era, to evaluate both our data, and the status of our search methods.

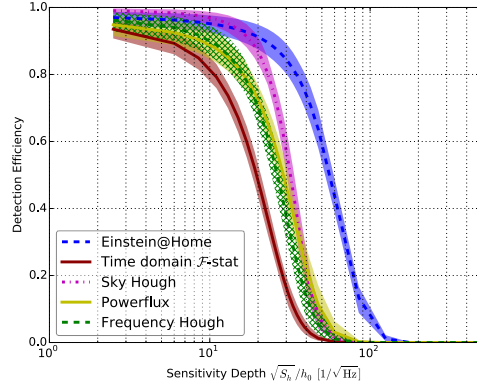


Figure 5.3: Overall detection efficiency [75, Fig. 1]. Error regions represent binomial distributions.

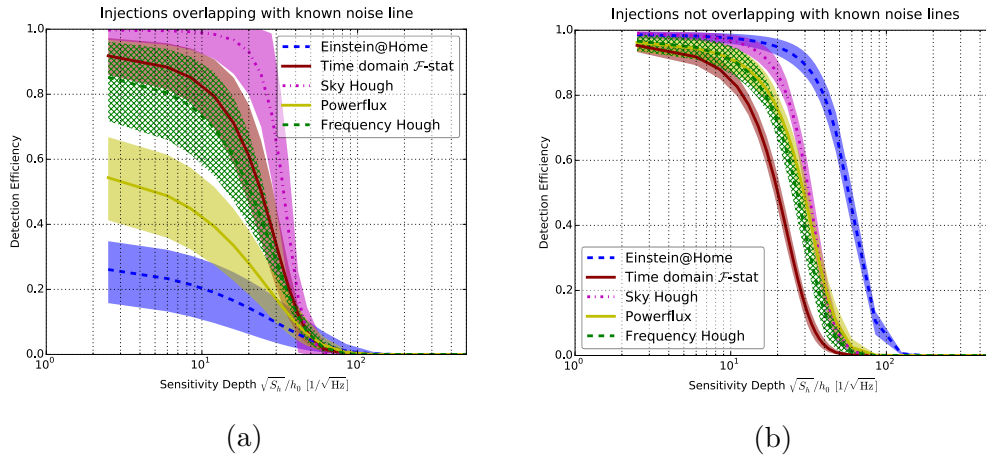
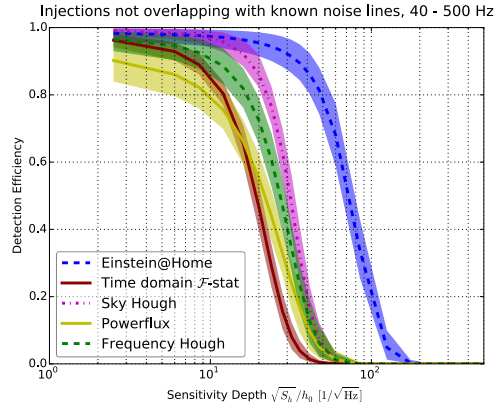
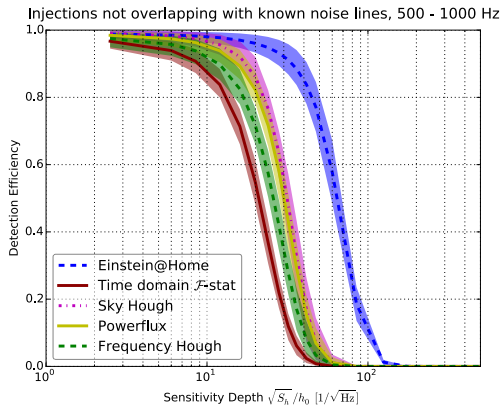


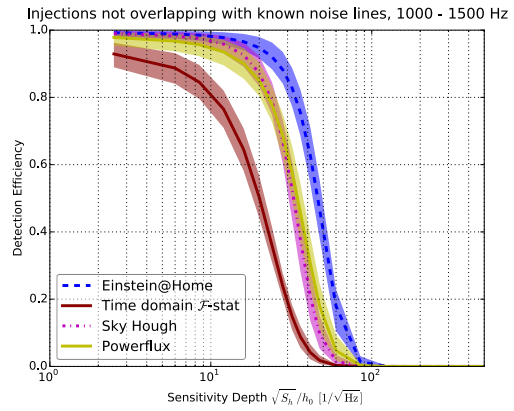
Figure 5.4: Detection efficiency inside/outside bands containing known lines [75, Fig. 2].



(a)

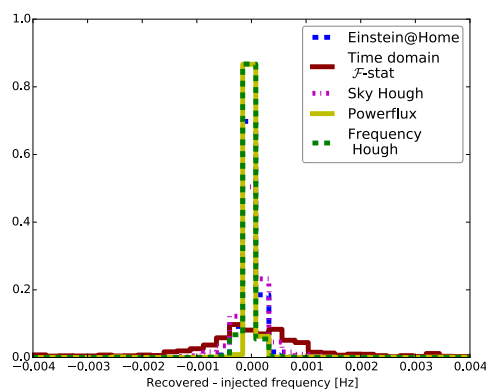


(b)




(c)

Figure 5.5: Detection efficiency for each frequency band [75, Fig. 3].



(a)



(b)

Figure 5.6: Difference between injected parameters, and parameters associated with highest SNR point in band [75, Fig. 8].

Chapter VI

Line Cleaning in O1 Searches

The first observing run of Advanced LIGO began on September 18, 2015 and ran through January 12, 2016. While a significant improvement over Initial LIGO, the resulting data contain instrumental artifacts. Many of the artifacts in the low-frequency region (10-150 Hz) are sharp “lines” that each contribute significant power to a few detector frequency bins. Examples of lines include the resonant frequencies of the test mass suspension (“violin modes”), and excitations from power systems at 60 Hz. While most CW signals will show a time-varying detector frequency, some templates remain stationary for extended periods. If the template overlaps with one of these lines, its inferred power may be erroneously inflated by the artifact, giving a false-positive. Extraneous outliers are expensive to follow-up, so the collaboration makes efforts to mitigate these noise sources, both by tracking and eliminating their cause before further data collection, and by cleaning the data already in hand.

6.1 Low-Frequency Lines

6.1.1 Combs

Many of the lines in the low-frequency range of O1 have been identified as combs – series of lines with equal spacing, due to harmonics of the noise source. One example, a comb spaced at 1 Hz intervals and offset from zero by 0.5 Hz, is shown in Fig. 6.1.




Figure 6.1: H1 low-frequency normalized spectrum from 3600s O1 SFTs. Dots mark members of half-Hertz comb. Normalization sets the median power in each frequency bin to 1.

This set of lines, dubbed the “half-Hertz comb,” was the result of blinking LEDs on the GPS timing cards at the detector end stations. The cards indicated their GPS synchronization by cycling between on for one second, then off for one second. This precisely-timed minute current draw coupled to the GW strain channel, creating the noise lines. Harmonics of the square-wave appear in the data as a series of evenly-spaced peaks. The problem was resolved by reprogramming the cards to

keep the LEDs on after synchronization rather than blinking, which significantly reduced the strength of the comb [28]. The improved O2 spectrum is shown in Fig. 6.2.




Figure 6.2: H1 low-frequency normalized spectrum from 3600s O2 SFTs. The half-Hertz comb (red dots) is almost non-existent.

6.1.2 Non-Stationary Lines

Most lines have a relatively stable frequency and power, but sometimes a line's power can fluctuate, which leads to problems with the background subtraction used by PowerFlux. The algorithm is described in [7]. SFT powers are laid out in a grid, with each row indexed by time, and each column indexed by frequency. The following procedure is then carried out

1. For each row, subtract the median $\log(\text{power})$ from each element in the row
2. For each column, subtract the median $\log(\text{power})$ from each element in the column
3. Repeat from (1) until changes are sufficiently small

In the case of a non-stationary line, however, this algorithm can subtract too much or too little in different grid regions. See Figures 6.3 and 6.4, for example. The line gets stronger partway through the run, which results in a median that does not represent the true noise behavior in that band. The algorithm assumes a consistent statistic for the noise, and cannot remove this line.

6.1.3 Excessive Outliers

The presence of these instrumental lines can lead to extraneous outliers. As discussed regarding the S parameter veto (Section 4.7.1), some signal templates can spend long periods of time in a single detector frequency bin. If that bin contains a line, the template can yield an inflated SNR that causes it to be marked for followup. Because of the limited computational resources available, we have a strong incentive to reduce the number of outliers that are unlikely to be useful.

6.2 Cleaning

In the case of known instrumental lines, the contaminated frequency regions can be “cleaned” by replacing the affected regions with Gaussian noise. The mean and standard deviation of the noise generated are determined from the noise floor at either end of the cleaned region, and interpolated between. This technique has been used extensively by the Einstein@Home search group since the S4 search [8], and we explored its use in the low-frequency PowerFlux searches. The non-stationary line at 75.5 Hz described above is among the lines removed by the cleaning process (Fig. 6.5).




Figure 6.3: Example of non-stationary line in O1 at 75.5 Hz in raw frequency-time spectrogram. Horizontal gaps represent times detector was offline.




Figure 6.4: 75.5 Hz line in residual spectrogram after subtraction. Excess power remains, due to non-stationarity of the line.




Figure 6.5: Spectrogram with 75.5 Hz line removed after cleaning.

6.2.1 SFTClean

The *SFTClean* tool was introduced in the Einstein@Home S4 search because of the large number of narrow spectral lines from sources such as harmonics of the power mains and violin modes of the mirror suspensions [8]. The tool takes a list of line regions and a complex-valued SFT as input and removes the data in the specified frequency bands from the SFT. The lost data are replaced with generated Gaussian noise, with mean corresponding to an assumed noise floor interpolated between the end points of each removed region.

The analysis carried out here used the Version 1 line lists for O1, released January 31, 2016 [28].

6.2.2 PEM Correlations

When choosing frequency bands for cleaning, it's important to ensure the excess power removed could not be due to a true signal. The LIGO detectors are equipped with a wide array of physical environment monitors (PEMs), shown in Table 6.1. Any strain line that correlates with a line in one of these monitor channels would be highly suspect, and acceptable to remove.

A number of lines remained in the O1 data after cleaning, so we attempted to match these to PEM channels so they could be added to the next version of the line list. An example is shown in Figure 6.6. This plot shows coherence between the H1 strain channel, and the magnetic field detected in an area of the Y end-station. The coherence is calculated by the tool *gwpy* [53] according to

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)} \quad (6.1)$$

for cross-spectral density S_{xy} and auto-spectral densities S_{xx} and S_{yy} . There is

Table 6.1: Summary of physical environment monitors [81]

Type	Sensor	Type	Sensor
Seismic noise	3-axis seismometer	Cosmic ray muons	Scintillation detector
Seismic noise	2-axis tiltmeter	Power line fluctuations	Line monitor
Seismic noise	1-axis accelerometer	Residual gas	Gas analyzer
Acoustic noise	Electret microphone	Contamination	Crystal monitor head
Magnetic field	3-axis magnetometer	Contamination	RGA head
Radio interference	Multichannel receiver	Contamination	Control

intermittent coherence at 29.8 Hz during the latter half of O1, suggesting some terrestrial interference in the strain channel, though further investigation would be needed before adding this frequency to the H1 line list.

6.2.3 Comparison

As a test of the cleaning procedure, we ran all-sky searches of O1 data over the frequency range 50–135 Hz. The first set of searches used standard PowerFlux parameters. The second did not perform background subtraction on the SFTs before searching. The third used SFTs that had been cleaned of known lines, and background subtracted. Each PowerFlux run covers 501 frequency bins, or about 69.6 mHz for the 7200-second SFTs used here. SNR values are based on the estimated noise in each frequency band.

As discussed in Section 6.1.2, non-stationary lines can result in the background subtraction removing too much or too little from bins. The effect of this can be seen in Fig. 6.7 – For small assumed $|f|$, the SNR of the original run (with background subtraction) is artificially lowered as a result of too much power being removed from certain bins. Background subtraction can be disabled, but doing so inflates the SNR because of the lines that are no longer removed. Removing the known lines via cleaning avoids most of the SNR increase from detector artifacts, including those that the background subtraction misses. There remains some SNR increase,




Figure 6.6: Coherence between detector strain channel, and LHO Y end-station vacuum equipment area magnetometer channel. Data are from 60-second SFTs, averaged over one hour, and sampled once per day.





Figure 6.7: Median SNR across spindowns. Dotted line marks beginning of S parameter veto for $S_{\text{large}} = -2.5 \times 10^{-11}$. Templates to the right of this line may be stationary in detector frequency, and pick up noise lines. Also note excessive background subtraction in original run.

however, since only well-vetted lines are included in the list for removal, so some artifacts are still present.


The stationarity of a template also depends on frequency, due to Doppler shift. Fig. 6.8 shows the changes in SNR for runs with/without background subtraction and cleaning over the full range of templates tested. As above, the enhanced SNR region to the right of the S parameter line is greatly diminished with cleaning.

In addition to determining maximum SNR, PowerFlux also sets upper limits on the strength of signals. Until a CW detection is made, these upper limits should represent the weakest signal the pipeline could detect. Typically, noise lines appear



(a) Original run

(b) No background subtraction



(c) Cleaned

Figure 6.8: Max SNR for each frequency/spindown in test run. Dotted line marks S parameter veto for $S_{\text{large}} = -2.5 \times 10^{-11}$.

in the upper limits, since a stationary signal at the same frequency would need to exceed the noise to be detected. Sample upper limits reported by PowerFlux are shown in Fig. 6.9.

The vast majority of lines do not appear in the cleaned upper limits, but we must question whether this is an accurate representation of the detectors' capabilities. If a signal lay on top of one of these lines, it would be removed as well. In light of this, we investigated the extent to which certain templates would be affected by cleaning.

6.2.4 Fraction Cleaned

By its nature, cleaning introduces artificial data into the analysis. It's important to keep track of how much data has been replaced for any affected outlier. For example, it may happen that a line is wider than recorded in the line list, and as a result is not completely removed. The artificial data that replaces the removed region is based on the end points, so in this case we would be inserting a block of elevated signal. The pipeline finds the template with the highest SNR, but if a significant amount of that SNR is coming from artificial data, we should veto that outlier just as we did for noise lines in Section 4.7.

By calculating the Doppler shift and spindown for each time in O1, we can track a source's frequency in the detector frame. We mark each time for which the frequency intersects with a cleaned region. We may wish to set an upper limit on the amount of time a signal is in a certain frequency band. We can also use this tool to measure how much of a signal in an uncleaned search comes from a known contaminated region.

There are some lines wide enough to cover a full search band when using 2-hour SFTs. As above, the search bands are 69.6 mHz wide. In the H1 line list, there are 20 lines wider than this, and in L1 there are 25. Among these is the 60 Hz line

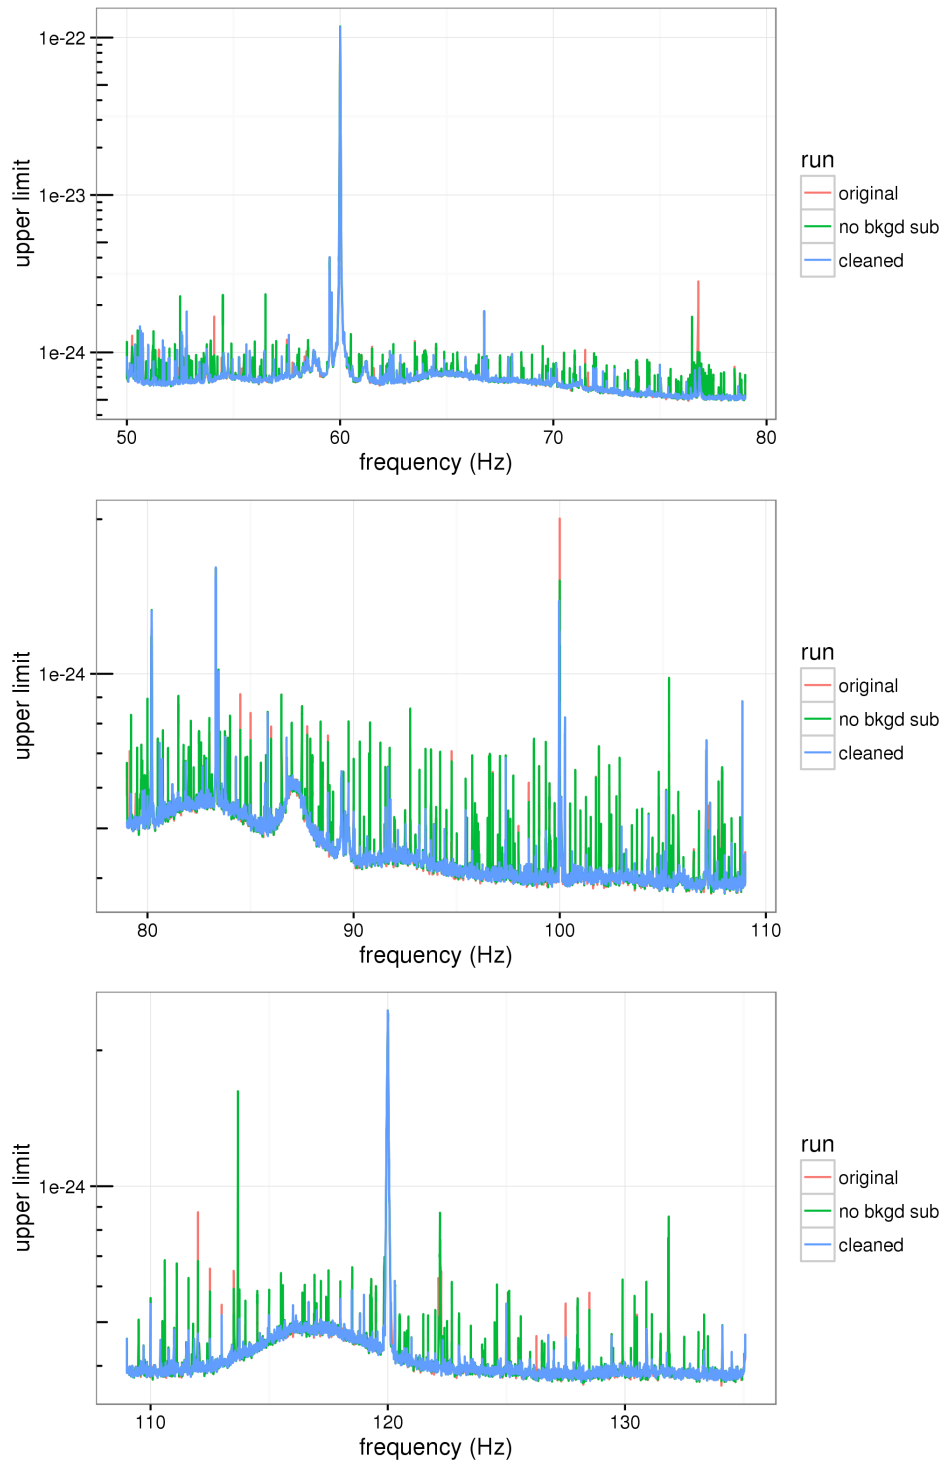


Figure 6.9: Upper limits in search band. Most lines are absent from cleaned runs.

resulting from the U.S. power mains, which is within the frequency band searched. This line, including shoulders, has a width of 120 mHz, enough to cover almost 2 search bands.




Figure 6.10: Distribution of time outliers spent in cleaned regions. Error bars estimated as $\Delta N = \sqrt{N}$.

Fig. 6.10 shows a histogram of the fraction of SFTs for which outliers were in a cleaned region. The runs with and without background subtraction yield nearly




Figure 6.11: Relation between time outliers spent in cleaned regions and spindown. Error bars estimated as $\Delta y = \sigma_y/N$.

equal distributions, but the cleaned data histogram shows less time in the regions marked as lines. This shows that cleaning does indeed help to disfavor outlier templates that get much of their power from an instrumental line. The spike at 100% cleaned is due to the issue discussed above.

Fig. 6.11 shows a similar plot to Fig. 6.7, where low $|f|$ results in Doppler




Figure 6.12: Relation between time outliers spent in cleaned regions and SNR. Error bars estimated as $\Delta y = \sigma_y/N$.

stationarity. The cleaned searches tend to choose templates that spend less time in marked regions, in spite of their stationarity.

Fig. 6.12 shows a significant drop in SNR for the cleaned plots, mainly for templates with around 50% cleaning. Also of note is the relative SNR of the runs with/without background subtraction. In most cases, the templates without

Table 6.2: Average ratio of SNRs, $\frac{\text{Run 2}}{\text{Run 1}}$

Run 1	Run 2	SNR ratio
original	no bkgd sub	0.9826544
original	cleaned	0.7637940
no bkgd sub	cleaned	0.7499136

Table 6.3: Number of templates with SNR above threshold. Cleaned runs show sharp decline in high-SNR points.

Run	SNR > 5	SNR > 7	SNR > 9
original	58958	18825	3982
no bkgd sub	58950	19352	4327
cleaned	58591	15708	1461

background subtraction have higher SNR, as we might expect, but near 50% cleaning, the relation inverts. This is likely due to the background subtraction issue described in Section 6.1.2. The overall change in SNR is shown in Table 6.2, and the distribution of SNR is summarized in Table 6.3. The SNR of the cleaned runs is significantly diminished, while the runs with background subtraction disabled trend toward lower SNR, but retain some exceptionally high cases.

In summary, cleaning shows promise for reducing outliers from known noise lines, though care must be taken to keep track of affected templates. Further testing and validation will be required to use cleaning in production of upper limits should future LIGO data runs exhibit the level of line contamination seen in the O1 run.

Chapter VII

Exploiting Longer Coherence Times

Broadband CW analyses are typically carried out by taking the detectors' output strain over time, dividing it into segments, and performing a Fourier transform on each. The duration of each segment is referred to as the coherence time, T_{coh} . This quantity determines the frequency resolution,

$$\delta f = \frac{1}{T_{\text{coh}}}, \quad (7.1)$$

and the time resolution,

$$\delta t = T_{\text{coh}}. \quad (7.2)$$

As discussed in Chapter VI, the data contain many narrow-frequency lines. By increasing the coherence time, we make the frequency bins narrower, allowing the evolving or modulated signal to move off of a line more quickly; this choice, however, comes at the cost of worse time resolution, which can cause a signal to be smeared across multiple bins in a single SFT. In this chapter, we examine the advantages and limitations of using long coherence times in a search, and compare

the results of searches using different coherence times.

7.1 Losses and Gains

7.1.1 Line Avoidance

Due to coherence with GPS timing, most of the lines in the low-frequency O1 data are narrow and bin-centered. For Hann-windowed SFTs, this means the line will occupy 3 bins, regardless of the coherence time used. The total frequency band contaminated will therefore be

$$\Delta f_{\text{line}} = \frac{3}{T_{\text{coh}}}. \quad (7.3)$$

This means that when we use tools like the line response veto or cleaning, the amount of data lost is significantly smaller than for the typical 30 minute coherence time. As a signal moves in detector frequency due to Doppler shift and spindown, it will also change bins more rapidly with longer coherence time.

7.1.2 Signal Smear

The rapid bin changes introduce a problem, though: If the signal occupies multiple bins during a single SFT, the power will be spread among them, diminishing the strength for the nominal bin. The total number of bins a signal will enter during a single SFT is, to a good approximation,

$$\Delta b = |\dot{f}_{\text{det}}|(T_{\text{coh}})^2, \quad (7.4)$$

where \dot{f}_{det} is the total frequency rate of change in the detector frame, from both source spindown and Doppler modulation. Since Doppler modulation depends on frequency, the requirement of a signal remaining in one bin imposes a limit on the

Table 7.1: Maximum frequency to prevent Doppler shift for sources in the ecliptic plane from occupying multiple bins.

T_{coh} (hours)	Max f (Hz)
0.5	16 000
1	4 000
2	1 000
4	250
8	60

maximum frequency searched for each coherence time (see Table 7.1).

7.1.3 Antenna Pattern

The sensitivity of the LIGO detectors depends in part on the direction of the source. As the detectors change orientation during the Earth’s daily rotation, this sensitivity will vary. With short coherence times, we can downweight SFTs with poor sensitivity, but long coherence requires averaging the antenna pattern over the integration time. Using the antenna pattern expressions in Jarankowski, *et al.* [43], we can see the effect of this degradation. For a set of times, we calculate the antenna pattern $F^\psi(t)$ for polarization ψ . Normally, two parameters are required to specify polarization, but by decomposing the signal into plus and cross polarizations, we can shift ι into the strain equations, and use only ψ in the antenna response functions:

$$h(t, \psi, \iota, \alpha, \delta, d) = F^+(t, \psi, \alpha, \delta, d)h^+(t, \psi, \iota) + F^\times(t, \psi, \alpha, \delta, d)h^\times(t, \psi, \iota) \quad (7.5)$$

We assign a weight to the pattern proportional to the power, for circular and linear polarizations,

$$W^\circ(t) = (F^+(t))^4 + (F^\times(t))^4, \quad (7.6)$$

$$W^+(t) = (F^+(t))^4, \quad (7.7)$$

$$W^\times(t) = (F^\times(t))^4. \quad (7.8)$$

We then average this pattern over the coherence time to get a value for SFT n ,

$$\tilde{F}_n^\psi = \frac{1}{T_{\text{coh}}} \int_{nT_{\text{coh}}}^{(n+1)T_{\text{coh}}} F^\psi(t) dt, \quad (7.9)$$

and calculate the weight for each SFT,


$$\tilde{W}_n^\circ = (\tilde{F}_n^+)^4 + (\tilde{F}_n^\times)^4, \quad (7.10)$$

and similarly for $\tilde{W}_n^+(t)$ and $\tilde{W}_n^\times(t)$. The variation in these functions over an example day is shown in Fig. 7.1 for sources aligned with the equatorial pole, the equatorial plane, and halfway in between.


7.2 Choosing Coherence Length

Before carrying out a full analysis using longer coherence time, it is helpful to consider the various ways this parameter affects the data. Ideally, we would like to formulate a figure of merit that characterizes the expected performance of a given T_{coh} value.

To prevent signal smearing, we should only use long coherence times when $|\dot{f}_{\text{det}}|$ is small. Since Doppler shift increases with frequency, this suggests that coherence time should be inversely proportional to signal frequency, in cases where Doppler shift dominates over source spindown. In the ideal case, the signal frequency in the detector frame would change by less than one bin during a single SFT, but still move several bins over the course of the run to better avoid artifacts. The frequency



(a) \tilde{W}_n°



(b) \tilde{W}_n^+

Figure 7.1: Median binned weight over signal polarizations (ψ) for SFTs of various coherence times. Sources refer to equatorial coordinates.

bin offsets visited due to Doppler shift are

$$b(t) = \lfloor (\hat{\mathbf{e}} \cdot \mathbf{v}(t)/c) T_{\text{coh}} f_0 \rfloor. \quad (7.11)$$

For each SFT we define t_{count}^i , the total time spent in each b^i . From this, we can define a signal concentration for each SFT_{*n*},

$$t_{\text{conc}}^n = \frac{\max(t_{\text{count}}^i)}{\sum_i t_{\text{count}}^i}. \quad (7.12)$$

This quantity is shown in Fig. 7.2. Choosing the bin with the longest time, we assign a particular bin b_{SFT}^n to each SFT_{*n*}. To avoid narrow lines in the detector, we would like the signal bin to change over the course of the run. The standard deviation of b_{SFT}^n is shown in Fig. 7.3. Detector observing periods are broken up over the run, which limits the time that longer SFTs can cover. The total live-time t_{live} for O1 is shown in Fig. 7.4. We considered a function to measure the quality of a particular T_{coh} for given signals,

$$Q(T_{\text{coh}}, f_0, \hat{\mathbf{e}}) = \text{mean}_n \left(\tilde{W}_n t_{\text{conc}}^n \right) t_{\text{live}} \sigma_b, \quad (7.13)$$

where σ_b is the standard deviation of b_{SFT}^n . This contains the key metrics that interest us: It devalues cases where antenna pattern is poor, or where the signal occupies many bins in a single SFT, while also favoring cases that sample many frequencies across SFTs. We can calculate this quantity for a source on the equator at the start of O1 for several coherence times, and pick the one with the largest result. The mean coherence time chosen over a range of wave orientations is shown in Fig. 7.5. As expected, the chosen T_{coh} decreases as frequency (and Doppler modulation) increases. Unfortunately, tests of this measure, though generally

conforming to expectations, showed it lacks consistency. Nearby templates could receive starkly different coherence times, and the measure favors significantly longer coherence times than typically used. In spite of this, the parameters going into Q (Figs. 7.1 to 7.4) provide valuable insight into the choice of T_{coh} .




Figure 7.2: Mean signal concentration for test data.

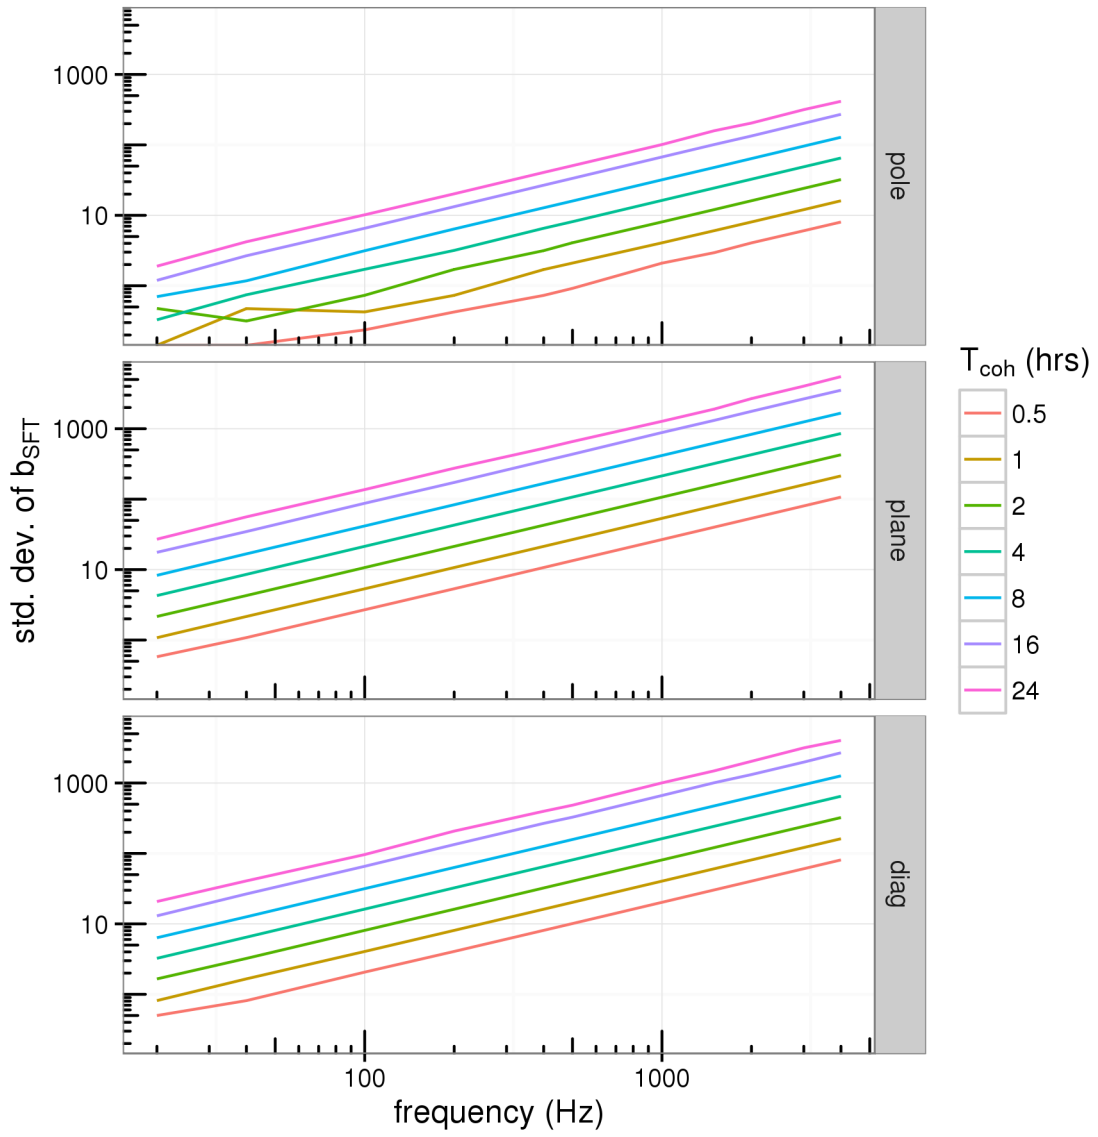


Figure 7.3: Standard deviation of signal bins over SFTs.




Figure 7.4: Total available H1 observation time for each T_{coh} during the O1 run.




Figure 7.5: Average T_{coh} chosen for sources on the equator for various wave orientations during the O1 run. As expected, T_{coh} decreases with frequency. There is little difference in results between linear and circular polarizations.

7.3 Comparison Search

As a test of the performance of long coherence time searches, we ran searches in the O1 data for 5000 injected signals using 2- and 4-hour SFTs. The injections were distributed over the whole sky, with frequency in the range 90-100 Hz, and spindowns between -1×10^{-8} and 1×10^{-9} Hz/s. Searches were nearly centered on the injections, with a spindown range of $\dot{f} \pm 2 \times 10^{-10}$ and a sky search radius of 0.15 radians. PowerFlux searches usually cover 501 frequency bins, but 4-hour SFTs have twice the resolution, so we used 2 searches to cover the same band as the 2-hour SFTs, a total of 69.6 mHz. For each injection, we choose the outlier from the 4-hour runs that was nearest to the injected frequency.

Fig. 7.6 shows the errors between injected and recovered signal frequency, spindown, and sky position. In each case, the 4-hour SFTs yielded values closer to the correct values. Fig. 7.7 shows the respective SNR and UL values for each point in the run for the two coherence times. The 4-hour searches had increased SNR, and in many cases increased UL, but the minimum UL is lower (Fig. 7.8b). Comparing the upper limits to the injected strain shows that the 4-hour searches are exceeding the 95% confidence level demanded for the upper limits (Fig. 7.8). For this set of injections, 97.5% of the 2-hour searches reported upper limits above the injected strain, but 99.1% of 4-hour searches were above. Normally, PowerFlux parameters are tuned to give a conservative estimate of upper limits, but as this is an exploratory study, that has not yet been done.

We also wish to know what advantages are to be gained from long coherence time when there is no signal present. We repeated the searches described, using the same sky disks and spindown ranges, but with the injections disabled. The SNR and UL for these runs are shown in Fig. 7.9. Without injections, the 4-hour runs

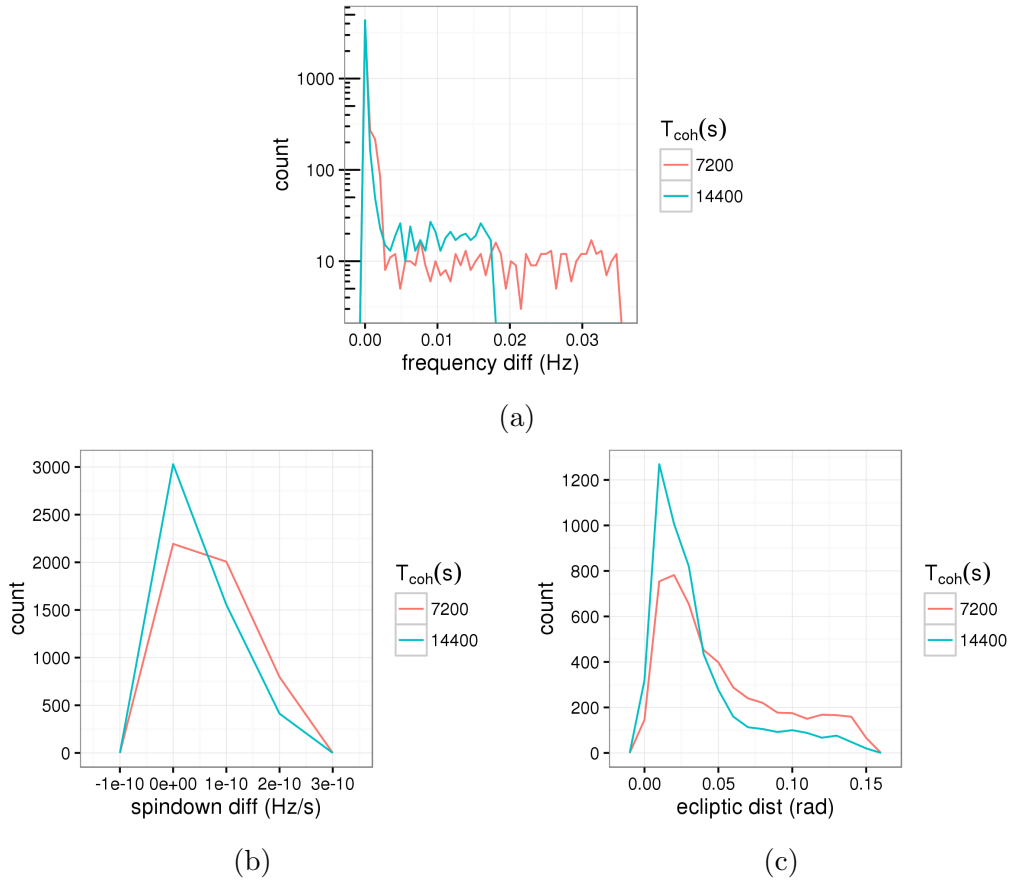


Figure 7.6: Errors in recovered parameters with respect to injected values. In each case, the 4-hour runs achieve better matching to the true parameters.

achieve better upper limits in spite of being overestimates. The SNRs though, show a collection of 2-hour points with abnormally high values. These points are due to instrumental artifacts, discussed in the next section.

7.3.1 Wandering Lines

A number of the cases with high SNR disparity had signal templates that overlapped with a wandering line near 99.975 Hz. This line shifted in frequency at the beginning of the run, causing the background subtraction to miss part of the line. In both spectrograms, an area of over-subtraction is visible (Fig. 7.10), but the




Figure 7.7: Comparison of values for SNR and UL between injection runs. Red line marks equality. 4-hour runs show increased SNR and UL in the presence of injections.


maximum SNR for the 2-hour SFTs was 40.12, while for the 4-hour SFTs it was 18.30, indicating different behavior in the background subtraction between the two coherence times.

To test this explanation, we created a model of a wandering line, with frequency


$$f(t) = \begin{cases} f_0 & t < T/2 \\ f_0 + \left(\frac{t-T/2}{T/2}\right)^2 \Delta f & t \geq T/2 \end{cases} \quad (7.14)$$

for base frequency f_0 , total time T , and frequency deviation Δf . The line remains fixed for the first half of the run, then moves by Δf quadratically over the second half. We do not expect a Doppler-modulated source to show this type of behavior. We can construct a time-series for the line with

$$h(t) = \sin(2\pi t f(t)). \quad (7.15)$$



(a) Full range



(b) Zoomed: Blue lines show minimum upper limit reported for each run.

Figure 7.8: Validation of upper limits. Ideally, 95% of injections would lie above red line.




Figure 7.9: Comparison of values for SNR and UL between runs with no injections. Red line marks equality. 4-hour runs generally show decreased SNR and UL, indicating less sensitivity to noise.

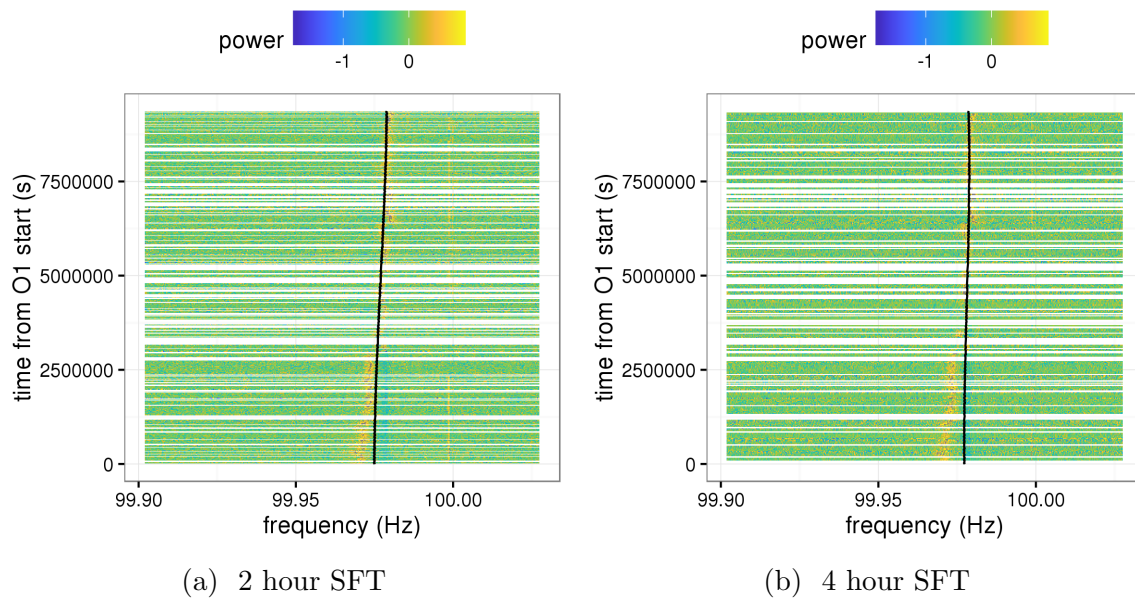


Figure 7.10: Wandering line at 99.975 Hz after background subtraction. Curve shows template path selected by each search. Note over-subtraction (blue) to the right of the curve, and under-subtraction (yellow) to the left.

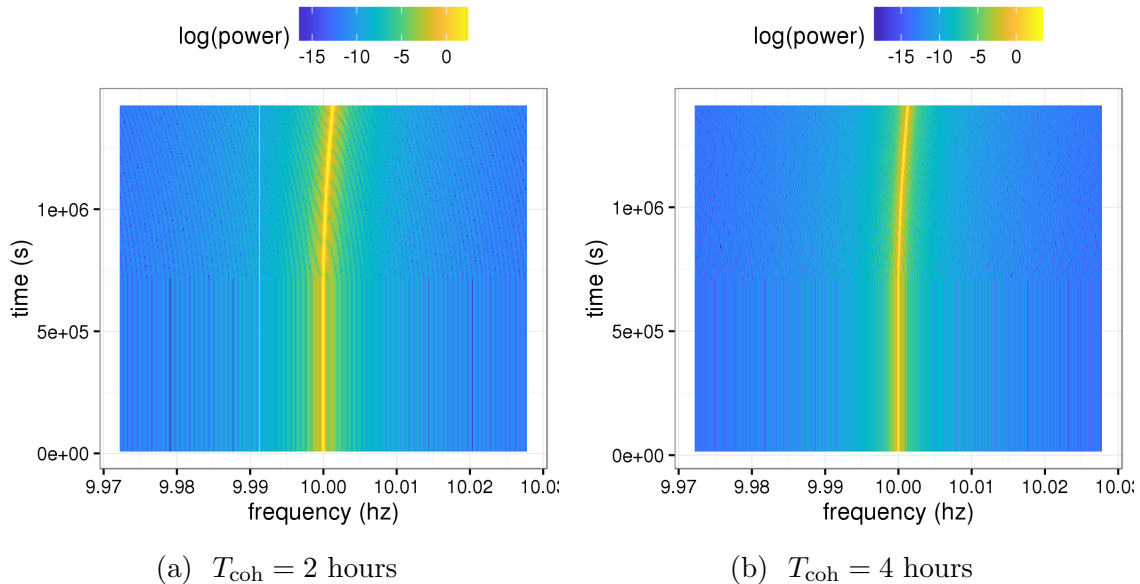


Figure 7.11: Hypothetical wandering line to study background subtraction. The line has fixed-frequency for the first half of the timespan, then shifts quadratically by a small amount. Some aliasing is present due to low sampling rate.

For the test data, we calculate this value for $0 < t < 200$ hours, sampled at $f_s = 50$ Hz, with $\Delta f = 0.28$ mHz, or 2 frequency bins for the 2-hour SFTs. The time series is then divided into segments of T_{coh} with 50% overlap, and a Hann window applied [57]:

$$h[i] = \left(\frac{1}{2} - \frac{1}{2} \cos \left(\frac{2\pi i}{N} \right) \right) h(t_i), \quad (7.16)$$

where t_i are the time samples within each segment, and the total number of samples is $N = T_{\text{coh}} f_s$. We then apply a discrete Fourier transform to each segment to create the spectrograms shown in Fig. 7.11a for $T_{\text{coh}} = 2$ hours, and Fig. 7.11b for $T_{\text{coh}} = 4$ hours. We now perform background subtraction, using the algorithm described in Section 6.1.2. The results are shown in Fig. 7.12.

The frequency shift in the line causes part to be missed by background subtraction, while also subtracting too much in the region the line left. We can sum the total power for a template that remains stationary in detector frequency




Figure 7.12: Wandering line after background subtraction. As in the real data, the change in frequency causes the background subtraction to miss part of the line, as well as over-subtracting the region it left.

(Fig. 7.13). We consider the effect of changing the base frequency of the template, in essence sweeping the spectrogram horizontally.

In the band shown, the power found from the 4-hour SFTs is less than that from the 2-hour SFTs in 98% of the fixed-frequency trajectories. This suggests that the increased frequency resolution of the 4-hour SFTs makes them less sensitive to instrumental lines. These results support the decreased SNR and upper limits seen above when no injections were present (Fig. 7.9).

7.4 Conclusions

The use of coherence times beyond 2 hours has the potential to give greater sensitivity to signals, while decreasing the effects of instrumental lines. The increased frequency resolution that brings these benefits also carries costs. For signals with large spindown or Doppler shift, long coherence times can smear the




Figure 7.13: Power sums along stationary frequency. Frequency shifts represent difference from f_0 .

power over several bins. The increased requirements on detector live-time can also decrease the amount of usable data. Considering the limitations, long coherence searches can prove useful for low-frequency bands, where Doppler shifts are smaller, and the advantages gained by line avoidance outweigh the costs.

Chapter VIII

LALBarycenter Approximations

The LIGO detectors are better described as antennas than telescopes, because they pick up signals from all directions with varying sensitivity. This means that to localize signal sources, we must precisely determine each detector's position in space at any given time. The LIGO Analysis Library (LAL) Suite provides two functions for this purpose: `LALBarycenterEarth`, and `LALBarycenter` [50]. `LALBarycenter` also considers the effects other Solar System bodies have on a GW traveling to Earth, such as Shapiro and Roemer delays, to determine the signal's emission time from a distant body. This allows us to transform from the detector reference frame to the source frame, and measure the parameters associated with the source.

Calling these routines for every time and sky-point in a search can become expensive. Ideally, we would be able to account for the small differences in emission times in a region of the sky. The Loosely Coherent method for `PowerFlux` was developed to reuse signal templates and allow for a continuous drift in signal phase [31, 32, 33]. In this chapter, we discuss the origins and capabilities of `LALBarycenter`, and present a model for approximating emission times for points near a reference. This model serves as a proof-of-concept for the template reuse employed by loose coherence.

Table 8.1: Summary of EarthState and EmissionTime structures

EarthState	EmissionTime
Earth position and velocity	Detector position and velocity
Einstein delay	Roemer delay
Greenwich Mean Sidereal Time	Shapiro delay
Sun-Earth vector	Travel time from source to SSB

Much of the work in this chapter has been submitted for publication in Physical Review D [67].

8.1 Capabilities of LALBarycenter

LALBarycenter is based on the widely-used radio astronomy timing package TEMPO2 [34, 41]. LALBarycenter generates two structures of importance to our calculations: EarthState and EmissionTime. The elements of these structures are summarized in Table 8.1.

8.1.1 LALBarycenterEarth

LALBarycenterEarth takes as input the time of measurement and ephemeris data giving details of Sun/Earth position and velocity at reference times. Since this routine is only concerned with the Earth’s center of mass, sky direction is not needed. For the input time, positions and velocities for the Earth and Sun are interpolated from the given ephemeris data. The Sun-Earth vector is calculated simply by taking the difference in positions. Calculation of the Einstein delay uses a truncated form of the equation in TEMPO2. Greenwich Mean Sidereal Time (GMST) is calculated with the equation in [68, p. 50].

8.1.2 LALBarycenter

LALBarycenter takes as input: observation time, source direction and distance, detector site, and EarthState from LALBarycenterEarth. Using the inputs, the routine calculates the Earth’s rotational state, as well as Shapiro and Roemer delays. All the information is added together to find the signal’s emission time. This procedure allows us to translate measurements made at the detectors into the source frame.

8.2 TEMPO2

LALBarycenter adapts several features developed for TEMPO2 in calculating emission times. The time difference between the detector and the SSB is summarized in [35, Eq. 8]:

$$\Delta_{\odot} = \Delta_A + \Delta_{R\odot} + \Delta_p + \Delta_{D\odot} + \Delta_{E\odot} + \Delta_{S\odot} \quad (8.1)$$

The terms used are:

- Δ_A Atmospheric delays
- $\Delta_{R\odot}$ Roemer delay describing light travel-time from observatory to SSB
- Δ_p Parallax delay due to curved wavefront
- $\Delta_{D\odot}$ Dispersion due to free electrons in solar wind
- $\Delta_{E\odot}$ Einstein delay, including both gravitational redshift and special relativistic time dilation
- $\Delta_{S\odot}$ Shapiro delay from curved spacetime around Solar System bodies

Gravitational waves are believed to travel at light speed regardless of medium and frequency, so no dispersion is present, and the terms Δ_A and $\Delta_{D\odot}$ are neglected

in the LALBarycenter implementation.

8.3 Principles of Fits

8.3.1 Fitting Methods

The purpose of fitting a model to data is to make the model match the data as closely as possible, but there are several ways to measure the degree of matching. In fitting emission times, we are also fitting signal phase for an assumed frequency. We decided that the maximum tolerable error was a 30-degree phase change for a 2 kHz signal, which is equivalent to a 42 μs error in emission time, in order to lose no more than $\sim 15\%$ SNR in all-sky CW searches reaching as high as 2 kHz. Since we do not want to exceed this error for any single point, our fitting routine would ideally determine the set of coefficients that minimizes the maximum error for any point:

$$\min_i \left[\max_j |f_i(x_{j,1}, \dots, x_{j,k}) - y_j| \right] \quad (8.2)$$

where f_i is the fitting function for input training points $x_{j,k}$, with corresponding coefficients $\{a_k\}_i$, and y_j are the output training points.

There is no generalized method to find the coefficient set that satisfies this equation; however, there are methods for approximately minimizing functions. In the statistical analysis program R [63], a variety of such methods are implemented in the package Optimx [56]. The methods implemented in Optimx are listed in Table 8.2.

Unfortunately, these methods are computationally expensive, so we also used the

Table 8.2: Fitting methods used in Optimx

Name	Reference
BFGS	Nocedal and Wright (1999)
CG	Fletcher and Reeves (1964)
Nelder-Mead	Nelder and Mead (1965)
L-BFGS-B	Byrd et. al. (1995)
nlm	Dennis and Schnabel (1983) and Schnabel et al. (1985)
nlminb	http://netlib.bell-labs.com/cm/cs/cstr/153.pdf
spg	Varadhan and Gilbert (2009)
ucminf	H. B. Nielsen (2000)
newuoa	M. J. D. Powell (2009)
bobyqa	M. J. D. Powell (2009)
nmkb	Varadhan and Borchers (2011)
hjk	Varadhan and Borchers (2011)

more typical least-squares fitting, which minimizes the sum of the squared errors:

$$\min_i \left[\sum_j (f_i(x_{j,1}, \dots, x_{j,k}) - y_j)^2 \right] \quad (8.3)$$

The model can be written as a matrix equation,

$$\mathbf{y} = \mathbf{X}\mathbf{a} + \boldsymbol{\epsilon}, \quad (8.4)$$

where X_{jk} is the k th fitting term of the j th input vector (and \mathbf{X} is generally not square), \mathbf{a} is a vector of coefficients, and $\boldsymbol{\epsilon}$ is a vector of errors. We can write the sum of the squared errors as

$$\boldsymbol{\epsilon}^T \boldsymbol{\epsilon} = (\mathbf{y} - \mathbf{X}\mathbf{a})^T (\mathbf{y} - \mathbf{X}\mathbf{a}) \quad (8.5)$$

$$= \mathbf{y}^T \mathbf{y} - 2(\mathbf{X}\mathbf{a})^T \mathbf{y} + (\mathbf{X}\mathbf{a})^T \mathbf{X}\mathbf{a} \quad (8.6)$$

$$= \mathbf{y}^T \mathbf{y} - 2(\mathbf{X}\mathbf{a})^T \mathbf{y} + \mathbf{a}^T \mathbf{X}^T \mathbf{X}\mathbf{a} \quad (8.7)$$

Note that $\mathbf{y}^T \mathbf{X} \mathbf{a} = (\mathbf{X} \mathbf{a})^T \mathbf{y}$, since both are 1x1 scalars. We now take the derivative with respect to \mathbf{a} , and set it to zero:

$$0 = 2\mathbf{X}^T \mathbf{X} \mathbf{a} - 2\mathbf{X}^T \mathbf{y} \quad (8.8)$$

The requirement of minimizing the squared errors gives

$$\mathbf{a} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}. \quad (8.9)$$


8.4 Earth Fit

Since the bodies that determine the orbit of the Earth affect, by definition, the shape of spacetime in the vicinity of the detectors, we began by making fits to the position of the Earth given by LALBarycenterEarth. We hoped these fits would inform the choice of parameters for the emission time fits.


The Earth's x-position for a ~ 95 year period is shown in Fig. 8.1a. The frequency spectrum of this data is shown in Fig. 8.1b, with various solar system periods labeled. We performed the fits by examining the residual after including each period in the fit, and choosing the next period to add. These residuals are shown in Figs. 8.2a to 8.8a. The model functions are of the form

$$x(t) = \sum_{i=1}^N A_i \cos \left(2\pi \frac{t - \phi_i}{T_i} \right), \quad (8.10)$$

where A_i and ϕ_i are fitting parameters, and T_i is the period of the body under consideration. Periods used are listed in Table 8.3.




(a)




(b)

Figure 8.1: x-position of Earth in ICRF [37]




(a)




(b)

Figure 8.2: Fit residual after including solar year




(a)




(b)

Figure 8.3: Fit residual after including half solar year




(a)




(b)

Figure 8.4: Fit residual after including period of Jupiter




(a)




(b)

Figure 8.5: Fit residual after including period of Saturn




(a)




(b)

Figure 8.6: Fit residual after including half period of Jupiter




(a)




(b)

Figure 8.7: Fit residual after including period of Moon



(a)



(b)

Figure 8.8: Fit residual after including period of Uranus

Table 8.3: Orbital periods used in Earth position fit

Body	Period (days)
Earth	365.256363004
Half Earth	182.628181502
Jupiter	4332.589
Saturn	10759.22
Moon	27.3217
Uranus	30689.0222
Half Jupiter	2166.2945




Figure 8.9: Emission time-difference function. For a search over a small patch of sky, emission times for each point in the patch can be calculated from $T + \Delta$.

8.5 Emission Time Fit

The fits for emission time are designed to use an exact value for the center of a patch of sky, then calculate the difference in emission time for surrounding points in the patch. Mathematically stated, the emission time T is a function of detection time t , source location u , and source parameters p (for a moving source).

Calculating $T(t, u, p)$ directly for every point is expensive, so instead we're looking for a less expensive function Δ such that

$$T(t, u, p) \approx T(t, u_0, p_0) + \Delta(t, u, p), \quad (8.11)$$

where (u_0, p_0) is a reference point that can be reused for all the points in a patch.

The principle is illustrated in Fig. 8.9.

We chose to use points shifted in right ascension (α) and declination (δ) to generate the differences in emission time input to the fits. For right ascension, this corresponds to a rotation around the equatorial pole, but the shift in declination is not a proper rotation; instead, it is a flow toward either the North or South Pole. This means we must impose the requirement $\frac{\pi}{2} - |\delta| > \Delta\delta$ to avoid crossing the poles. Without this limit, some fit terms would be discontinuous (see *e.g.*




Figure 8.10: Emission time $T(u, p)$ for the start of O1 over the sky.

Figs. 8.13e to 8.13f). Fig. 8.10 shows the emission times for a set of points covering the sky. Fig. 8.11 and Fig. 8.12 show the emission time difference between those points, and ones shifted by 0.01 radians in α and δ , respectively.

The fits were performed in R, using the tool `lm` (Linear Model), which is an implementation of the least-squares method described above. Terms were selected by comparing residuals to maps of various combinations of right ascension, declination, and locations of solar system bodies. Some selected terms are shown in Fig. 8.13. Note, for example, the similarities of Fig. 8.11 and Fig. 8.12 to Fig. 8.13p and Fig. 8.13q.

8.5.1 Fitting Procedure

The fits use as input data points drawn from the search space of O1. The data use a set of times \mathcal{T} , and a set of skypoints G_t , with minimum spherical separation ϵ . We add to G_t a set of $N_\odot \times N_\odot$ points with spacing ϵ_\odot covering the Sun, to account for Shapiro and Roemer delays. The points in G_t are shifted in right ascension and declination by a set ΔG . The parameters used in this study are listed in Table 8.4.




Figure 8.11: Emission time-difference $\Delta(u, p)$ for $\Delta\alpha$ at the start of O1.




Figure 8.12: Emission time-difference $\Delta(u, p)$ for $\Delta\delta$ at the start of O1.




Figure 8.13: Example skymaps of selected fit terms at the start of O1

Table 8.4: Parameters used in fits

Parameter	Value
\mathcal{T}	Every hour between t_{\min} and t_{\max}
t_{\min}	From start to end of O1, spaced every 200 000 seconds
t_{\max}	$t_{\min} + 250\,000$
ϵ	0.1040524 radians
ΔG	All combinations of $\Delta\alpha$ and $\Delta\delta$
$G_{t,r}$	Random subset of G_t with 7.5×10^5 points
$\Delta\alpha$	$\{-0.01, -0.00667, -0.00333, 0, 0.00333, 0.00667, 0.01\}$
$\Delta\delta$	$\{-0.01, -0.00667, -0.00333, 0, 0.00333, 0.00667, 0.01\}$
N_{\odot}	5
ϵ_{\odot}	0.001 radians

For clarity, we divide the terms into several categories:

- Direction-independent terms depending on GPS time and shift in sky position
- Direction differential-independent terms depending on source sky position and GPS time
- Time-independent terms depending on source sky position and shift in position

8.5.2 Definitions of Variables

The sky position variables are defined as

$$\begin{aligned}
 e_1 &= \cos(\delta) \cos(\alpha) \\
 e_2 &= \cos(\delta) \sin(\alpha) \\
 e_3 &= \sin(\delta)
 \end{aligned}
 \tag{8.12}$$

with $\alpha \in [-\pi, \pi]$ and $\delta \in [-\frac{\pi}{2} + 0.01, \frac{\pi}{2} - 0.01]$. The adjustment by 0.01 radians prevents flow over the poles, which would lead to ambiguous right ascension, as well as discontinuities in some terms. The change in e_i for a shift in right ascension $\Delta\alpha$,

and in declination $\Delta\delta$ can be approximated via Taylor expansion:

$$\begin{aligned}
\Delta e_1 &= \left(-\frac{1}{2}\Delta\alpha^2 \cos\alpha \cos\delta - \frac{1}{2}\Delta\delta^2 \cos\alpha \cos\delta + \frac{1}{4}\Delta\alpha^2\Delta\delta^2 \cos\alpha \cos\delta - \Delta\alpha \cos\delta \sin\alpha \right. \\
&\quad \left. + \frac{1}{2}\Delta\alpha\Delta\delta^2 \cos\delta \sin\alpha - \Delta\delta \cos\alpha \sin\delta + \frac{1}{2}\Delta\alpha^2\Delta\delta \cos\alpha \sin\delta + \Delta\alpha\Delta\delta \sin\alpha \sin\delta\right) \\
\Delta e_2 &= \left(\Delta\alpha \cos\alpha \cos\delta - \frac{1}{2}\Delta\alpha\Delta\delta^2 \cos\alpha \cos\delta - \frac{1}{2}\Delta\alpha^2 \cos\delta \sin\alpha - \frac{1}{2}\Delta\delta^2 \cos\delta \sin\alpha \right. \\
&\quad \left. + \frac{1}{4}\Delta\alpha^2\Delta\delta^2 \cos\delta \sin\alpha - \Delta\alpha\Delta\delta \cos\alpha \sin\delta - \Delta\delta \sin\alpha \sin\delta + \frac{1}{2}\Delta\alpha^2\Delta\delta \sin\alpha \sin\delta\right) \\
\Delta e_3 &= \left(\Delta\delta \cos\delta - \frac{1}{2}\Delta\delta^2 \sin\delta\right)
\end{aligned} \tag{8.13}$$

LALBarycenter provides vectors with information on the state of the Sun and Earth that are useful:

$$\begin{aligned}
\mathbf{S} &\quad \text{Vector pointing from Sun to Earth} \\
\mathbf{v} &\quad \text{Detector velocity vector} \\
\Delta t &\quad \text{Time since reference point} \\
\Omega_{\oplus} &\quad 2\pi/\text{sidereal day}.
\end{aligned} \tag{8.14}$$

We also define an array of the sin/cos of the reference point's right ascension and declination:

$$\mathbf{z} = \{\sin\alpha, \sin\delta, \cos\alpha, \cos\delta\} \tag{8.15}$$

and the second-order terms, excepting \sin^2 terms due to identity:

$$\begin{aligned}
\mathbf{z}' &= \{\cos^2\alpha, \cos^2\delta, \sin\alpha \sin\delta, \sin\alpha \cos\alpha, \\
&\quad \sin\alpha \cos\delta, \sin\delta \cos\alpha, \sin\delta \cos\delta, \cos\alpha \cos\delta\}
\end{aligned} \tag{8.16}$$

8.5.3 Direction-independent terms

The following terms are constant in sky-direction, and can be precomputed for every GPS time, and direction difference.

$$\sum_i a_{1,i} \Delta e_i \quad (8.17)$$

(Figs. 8.13p to 8.13u)

8.5.4 Difference-independent terms

The following terms are constant in direction-difference.

$$\sum_i a_{2,i} \sin(\Omega_{\oplus} \Delta t) z'_i + b_{5,i} \sin(\Omega_{\oplus} \Delta t) z'_i \quad (8.18)$$

$$a_{3,1} \Delta t \cos \delta + a_{3,2} \Delta t^2 \cos \delta + a_{3,3} \Delta t \cos^2 \delta \quad (8.19)$$

$$\sum_i a_{4,i} \Delta t^2 z'_i \quad (8.20)$$

$$\sum_i a_{5,i} \mathcal{S}_i e_i \quad (8.21)$$

(Figs. 8.13m to 8.13o)

8.5.5 Time-independent terms

The following terms vary only in sky-direction.

$$\sum_i a_{6,i} z_i \quad (8.22)$$

(Figs. 8.13a to 8.13d)

$$\sum_i a_{7,i} z'_i \quad (8.23)$$

(Figs. 8.13e to 8.13l)

Each of the terms in Eqs. (8.17) to (8.23) is multiplied by

Table 8.5: Term significance analysis. The max error column shows errors when the specified terms are omitted.

Term Group	Equation	Max Fit Error (s)
2nd Order Sinusoids	8.23	3.4319967642
1st Order Sinusoids	8.22	0.4335595581
Δt	8.19	0.0235629364
Δt^2	8.20	0.0010017764
Sun Direction	8.21	0.0002486640
Sidereal Rotation	8.18	0.0001820357
Direction-difference	8.17	0.0001655650

$\Delta\alpha, \Delta\delta, \Delta\alpha^2, \Delta\delta^2, \Delta\alpha\Delta\delta$. In addition, we include direction-time differential terms

$$\sum_i a_{8,i} \Delta t \Delta e_i \quad (8.24)$$

without $\Delta\alpha/\Delta\delta$ factors. Each term goes to zero when the rotation angle goes to zero. Note that Sun-Earth and detector velocity vectors are those for the saved points. In each term, any parts greater than order 3 in $\Delta\alpha$ and $\Delta\delta$ are removed. The effect of removing any terms is shown in Table 8.5.

8.6 Application to Loosely Coherent Searches

The phase change due to mismatches in template and source parameters shown in Table 4.4 can cancel and lead to multiple templates with the same relative phase, *e.g.* the template's frequency depends in part on Doppler shift, which is determined by sky position. The loosely coherent method employs bundles of sky positions to reuse templates. That is, the difference in phase between one sky position and a neighboring point is taken to be a smoothly varying function. The approximations to LALBarycenter described in this chapter show exactly that.

We also note that the fit depends on the difference in sky-direction, implying

that error in source position can lead to significant differences in emission time. Such errors could therefore severely degrade the sensitivity of fully coherent searches, requiring finely spaced templates. Computational resource limitations make this infeasible, so the advantages of loosely coherent searches are appealing.

In principle, the fits presented in this chapter could be used as a substitute for those in LALBarycenter, but careful implementation would be required for any sizable gain in efficiency. Instead, they serve better as both a proof-of-concept, and justification for using an efficient implementation of the loosely coherent method in followup searches [33].

Chapter IX

Conclusions

With the existence of gravitational waves definitively proven, the field has seen a surge of interest. This work described the principles behind these waves and their detection with LIGO. As part of the S6 Mock Data Challenge, we evaluated the performance of PowerFlux against other pipelines [75]. We investigated lines in the first O1 all-sky search [11], and tested the effects of cleaning in low frequency regions. Longer coherence times can also allow for noise reduction through finer frequency resolution. We examined the tradeoffs present in increasing coherence time to 4 hours. Finally, we developed approximations to the barycentering routines, which provide justification for the loosely coherent search technique. Continuous wave signals are among the most difficult to detect, but the methods described here show promise, and the algorithms are always being improved. The first continuous wave signal may already reside in our data, simply waiting to be recognized.

Bibliography

- [1] J. Aasi *et al.* Application of a Hough search for continuous gravitational waves on data from the fifth LIGO science run. *Classical and Quantum Gravity*, 31(8):085014, 2014.
- [2] J. Aasi *et al.* Implementation of an \mathcal{F} -statistic all-sky search for continuous gravitational waves in Virgo VSR1 data. *Classical and Quantum Gravity*, 31(16):165014, 2014.
- [3] J. Aasi *et al.* Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015.
- [4] J. Abadie *et al.* All-sky search for periodic gravitational waves in the full S5 LIGO data. *Physical Review D*, 85(2):022001, 2012.
- [5] B. Abbott *et al.* All-sky search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 77:022001, Jan 2008.
- [6] B. Abbott *et al.* Comprehensive all-sky search for periodic gravitational waves in the sixth science run LIGO data. *Physical Review D*, 94(4):042002, 2016.
- [7] B. P. Abbott *et al.* All-sky search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 77(2), 2008.
- [8] B. P. Abbott *et al.* Einstein@Home search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 79, 2009.
- [9] B. P. Abbott *et al.* GW151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.*, 116:241103, Jun 2016.
- [10] B. P. Abbott *et al.* Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116(6), 2016.
- [11] B. P. Abbott *et al.* All-sky search for periodic gravitational waves in the O1 LIGO data. *Phys. Rev. D*, 96:062002, Sep 2017.

- [12] B. P. Abbott *et al.* GW170104: Observation of a 50-solar-mass binary black hole coalescence at redshift 0.2. *Phys. Rev. Lett.*, 118:221101, Jun 2017.
- [13] B. P. Abbott *et al.* GW170608: Observation of a 19 solar-mass binary black hole coalescence. *The Astrophysical Journal Letters*, 851(2):L35, 2017.
- [14] B. P. Abbott *et al.* GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Phys. Rev. Lett.*, 119:141101, Oct 2017.
- [15] B. P. Abbott *et al.* GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.
- [16] B. Allen. The stochastic gravity-wave background: sources and detection. In J.-A. Marck and J.-P. Lasota, editors, *Proceedings of the Les Houches School on Astrophysical Sources of Gravitational Waves*. Cambridge University Press, 1996.
- [17] F. Antonucci *et al.* Detection of periodic gravitational wave sources by Hough transform in the f versus \dot{f} plane. *Classical and Quantum Gravity*, 25(18):184015, 2008.
- [18] P. Astone *et al.* Data analysis of gravitational-wave signals from spinning neutron stars. V. a narrow-band all-sky search. *Phys. Rev. D*, 82:022005, Jul 2010.
- [19] P. Astone *et al.* Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform. *Phys. Rev. D*, 90:042002, Aug 2014.
- [20] W. Baade and F. Zwicky. Cosmic rays from super-novae. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 20, pages 259–263, 1934.
- [21] H. A. Bethe and G. E. Brown. Evolution of binary compact objects that merge. *The Astrophysical Journal*, 506:780–789, 1998.
- [22] C. Biwer *et al.* Validating gravitational-wave detections: The Advanced LIGO hardware injection system. *Phys. Rev. D*, 95:062002, Mar 2017.
- [23] C. Blair *et al.* First demonstration of electrostatic damping of parametric instability at Advanced LIGO. *Phys. Rev. Lett.*, 118:151102, Apr 2017.
- [24] M. Bruggen and S. Rosswog. *Introduction to High-Energy Astrophysics*. Cambridge University Press, 2007.

- [25] D. Chakrabarty, E. H. Morgan, M. P. Muno, *et al.* Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature*, 424:42 EP–, 07 2003.
- [26] T.-P. Cheng. *Relativity, Gravitation and Cosmology: A Basic Introduction*. Oxford University Press, 2nd edition, 2010.
- [27] J. P. A. Clark, E. P. J. van den Heuvel, and W. Sutantyo. Formation of neutron star binaries and their gravitational radiation. *Astronomy and Astrophysics*, 72:120–128, 1979.
- [28] P. B. Covas, A. Effler, E. Goetz, *et al.* Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97:082002, Apr 2018.
- [29] D. C. Coyne. The laser interferometer gravitational-wave observatory (LIGO) project. In *1996 IEEE Aerospace Applications Conference. Proceedings*, volume 4, pages 31–61 vol.4, Feb 1996.
- [30] V. Dergachev. *An All-Sky Search for Continuous Gravitational Waves*. PhD thesis, University of Michigan, 2009.
- [31] V. Dergachev. On blind searches for noise dominated signals: a loosely coherent approach. *Classical and Quantum Gravity*, 27(20):205017, 2010.
- [32] V. Dergachev. Loosely coherent searches for sets of well-modeled signals. *Phys. Rev. D*, 85(062003), 2012.
- [33] V. Dergachev. Efficient loosely coherent searches for medium scale coherence lengths. *arXiv:1807.02351*, 2018.
- [34] R. T. Edwards, G. B. Hobbs, and R. N. Manchester. www.atnf.csiro.au/research/pulsar/tempo.
- [35] R. T. Edwards, G. B. Hobbs, and R. N. Manchester. TEMPO2, a new pulsar timing package – II. the timing model and precision estimates. *Monthly Notices of the Royal Astronomical Society*, 372(4):1549–1574, 2006.
- [36] A. Einstein and N. Rosen. On gravitational waves. *Journal of the Franklin Institute*, 223(1):43–54, January 1937.
- [37] A. L. Fey *et al.* The second realization of the international celestial reference frame by very long baseline interferometry. *The Astronomical Journal*, 150(58), August 2015.

- [38] M. E. Gertsenshtein and V. I. Pustovoit. On the detection of low frequency gravitational waves. *Journal of Experimental and Theoretical Physics (U.S.S.R.)*, 16(2):433–435, 1963.
- [39] B. Haskell and A. Melatos. Models of pulsar glitches. *International Journal of Modern Physics D*, 24(03):1530008, 2015.
- [40] B. Haskell, L. Samuelsson, K. Glampedakis, and N. Andersson. Modelling magnetically deformed neutron stars. *Monthly Notices of the Royal Astronomical Society*, 385(1):531–542, 2008.
- [41] G. Hobbs, R. Edwards, and R. Manchester. TEMPO2: a new pulsar timing package. *Chinese Journal of Astronomy and Astrophysics*, 6(S2):189, 2006.
- [42] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51–L53, 1975.
- [43] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58:063001, Aug 1998.
- [44] S. Karki, D. Tuyenbayev, S. Kandhasamy, *et al.* The Advanced LIGO photon calibrators. *Review of Scientific Instruments*, 87(11):114503, 2016.
- [45] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi. Search for continuous gravitational waves: Improving robustness versus instrumental artifacts. *Phys. Rev. D*, 89:064023, Mar 2014.
- [46] M. M. Komesaroff. Possible mechanism for the pulsar radio emission. *Nature*, 225:612, 02 1970.
- [47] P. G. Krastev, B.-A. Li, and A. Worley. Nuclear limits on gravitational waves from elliptically deformed pulsars. *Physics Letters B*, 668(1):1 – 5, 2008.
- [48] L. D. Landau. On the theory of stars. In *Collected papers of L. D. Landau*. Gordon and Breach, 1965.
- [49] P. D. Lasky. Gravitational waves from neutron stars: A review. *Publications of the Astronomical Society of Australia*, 32:e034, 2015.
- [50] LIGO Scientific Collaboration. LIGO Analysis Library Suite.
- [51] LIGO Scientific Collaboration. S6 continuous wave hardware injections.
- [52] N. A. Lockerbie, L. Carbone, B. Shapiro, *et al.* First results from the ‘violin-mode’ tests on an Advanced LIGO suspension at MIT. *Classical and Quantum Gravity*, 28(24):245001, 2011.

- [53] D. Macleod, A. L. Urban, S. Coughlin, *et al.* gwpy, July 2018.
- [54] D. V. Martynov *et al.* Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93:112004, Jun 2016.
- [55] G. Mendell and M. Landry. StackSlide and Hough search SNR and statistics. Technical report, LIGO Hanford Observatory, 2005.
- [56] J. C. Nash and R. Varadhan. Unifying optimization algorithms to aid software system users: optimx for R. *Journal of Statistical Software*, 43(9):1–14, 2011.
- [57] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck. *Discrete-Time Signal Processing*. Prentice-Hall, 2nd edition, 1999.
- [58] J. P. Ostriker and J. E. Gunn. On the nature of pulsars. I. theory. *Astrophysical Journal*, 157:1395, 1969.
- [59] C. D. Ott, A. Burrows, L. Dessart, and E. Livne. A new mechanism for gravitational-wave emission in core-collapse supernovae. *Phys. Rev. Lett.*, 96:201102, May 2006.
- [60] C. Palomba. Pulsars ellipticity revised. *Astronomy and Astrophysics*, 354:163–168, Feb. 2000.
- [61] D. Prialnik. *An Introduction to the Theory of Stellar Structure and Evolution*. Cambridge University Press, 2008.
- [62] R. Prix. *Neutron Stars and Pulsars*, volume 357 of *Astrophysics and Space Science Library*, chapter Gravitational Waves from Spinning Neutron Stars. Springer, Berlin, Heidelberg, 2009.
- [63] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015.
- [64] K. Riles. Recent searches for continuous gravitational waves. *Modern Physics Letters A*, 32(39):1730035, 2017.
- [65] P. R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing Company, 1994.
- [66] P. R. Saulson. If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves? *American Journal of Physics*, 65(6):501–505, 1997.
- [67] O. Sauter, V. Dergachev, and K. Riles. Efficient estimation of barycentered relative time delays for distant gravitational wave sources. Submitted to *Phys. Rev. D*, *arXiv:1712.06118v2*.

- [68] P. K. Seidelmann, editor. *Explanatory Supplement To The Astronomical Ephemeris*. The Nautical Almanac Office, U.S. Naval Observatory, 1992.
- [69] A. Singh, M. A. Papa, H.-B. Eggenstein, *et al.* Results of an all-sky high-frequency Einstein@Home search for continuous gravitational waves in LIGO’s fifth science run. *Phys. Rev. D*, 94:064061, Sep 2016.
- [70] F. G. Smith. *Pulsars*. Cambridge University Press, 1977.
- [71] United States Naval Observatory Nautical Almanac Office, Great Britain Nautical Almanac Office, Science and Engineering Research Council (Great Britain), *et al.* *The Astronomical Almanac for the Year 2000*. U.S. Government Printing Office, 2015.
- [72] University of Michigan Physics Demo Lab. 8C10.u2 – Accretion Disk.
- [73] G. Ushomirsky, C. Cutler, and L. Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *Monthly Notices of the Royal Astronomical Society*, 319(3):902–932, 2000.
- [74] R. V. Wagoner. Gravitational radiation from accreting neutron stars. *ApJ*, 278:345–348, Mar. 1984.
- [75] S. Walsh, M. Pitkin, M. Oliver, *et al.* Comparison of methods for the detection of gravitational waves from unknown neutron stars. *Phys. Rev. D*, 94, 2016.
- [76] J. Weber. Observation of the thermal fluctuations of a gravitational-wave detector. *Phys. Rev. Lett.*, 17:1228–1230, Dec 1966.
- [77] J. Weber. Evidence for discovery of gravitational radiation. *Phys. Rev. Lett.*, 22:1320–1324, Jun 1969.
- [78] J. Weber. Anisotropy and polarization in the gravitational-radiation experiments. *Phys. Rev. Lett.*, 25:180–184, Jul 1970.
- [79] J. M. Weisberg and J. H. Taylor. The relativistic binary pulsar b1913+16: Thirty years of observations and analysis. In F. A. Rasio and I. H. Stairs, editors, *Binary Radio Pulsars*, volume 328. ASP Conference Series, 2005.
- [80] R. Weiss. Electromagnetically coupled broadband gravitational antenna. Technical report, Massachusetts Institute of Technology, 1972.
- [81] R. Weiss. Physical environment monitor system. Technical Report T-960029, LIGO, 1994.
- [82] R. Weiss and D. Shoemaker. Paths to improving detector performance. Technical Report G010206-00-M, LIGO, 2001.

- [83] J. A. Wheeler and K. W. Ford. *Geons, black holes, and quantum foam : a life in physics*. Norton, 1998.
- [84] H. Wickham. The split-apply-combine strategy for data analysis. *Journal of Statistical Software*, 40(1):1–29, 2011.
- [85] M. Zimmermann. Revised estimate of gravitational radiation from Crab and Vela pulsars. *Nature*, 271:524, 02 1978.