

Unmodeled searches for long-lasting gravitational-wave
signals with LIGO and studies of underground seismic noise
for future gravitational-wave detectors

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Tanner Prestegard

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Dr. Vuk Mandic, Advisor

July, 2016

© Tanner Prestegard 2016
ALL RIGHTS RESERVED

Acknowledgments

The journey to obtain my doctoral degree has been the most difficult and most rewarding experience of my life. At times, it has been exhausting, exhilarating, challenging, and fulfilling. It would not have been possible without the support and guidance of my friends, family, and colleagues.

First, I would like to thank my advisor, Vuk Mandic. Thanks to his support, I have been able to work on one of the most exciting frontiers of physics with people from all over the world. Vuk has always respected and encouraged me and pushed me to step outside of my comfort zone. He has set a great example of both a scientist and person, and I count myself lucky to have been able to learn from him.

I would also like to thank the entire LIGO stochastic group for the many things I have learned while working with them. I am especially grateful to Eric Thrane for everything he has taught me about being a researcher. I have learned a lot about data analysis and programming from him, as well how to lead a project and collaborate with others effectively. Nelson Christensen has also been an invaluable person to work with; I am very thankful for his advice and insight on several of my projects.

I have been lucky to work with some excellent people in the LIGO group here at UMN, including Gwynne Crowder, Shivaraj Kandhasamy, Pat Meyers, and (honorary member) Michael Coughlin. We have shared some great times collaborating on projects, sharing ideas, and traveling to conferences. Thanks especially to Pat and Michael for some very insightful discussions and enjoyable diversions along the way.

For collaborative work on the STAMP all-sky search, I thank Marie-Anne Bizouard and Samuel Franco. It was a long journey and an uphill battle, but we did it! I also thank Samuel for the many good times we had, both in Minnesota and in France.

Planning and installing the seismic array at Homestake has been one of the most

enjoyable parts of my graduate experience. I thank Daniel Bowden, Gary Pavlis, and Victor Tsai for being great colleagues and friends throughout the many weeks we spent at Homestake. This project would not have been possible (or enjoyable) without assistance from many of the SURF employees, especially Jaret Heise, Tom Regan, and Jamey Tollefson.

Thanks to Scott Bergman, who randomly sent me an article on the LHC, which triggered a change of major to physics and set me on this path. I am also grateful to Dan Dahlberg for a lot of advice and encouragement, especially during my undergraduate years.

A big thanks to many of the graduate students with whom I've shared this experience: Barry Costanzi, Terry Bretz-Sullivan, Jared Turkewitz, Mark Pepin, Dominick Rocco, Andy Galkiewicz, Tobias Gulden, Andy Julin, and Peter Martin. All the hours we spent working on homework together are now fond memories, and Fridays won't be the same without the Holy Burrito Crew.

I am deeply grateful to my family for their constant encouragement, love, and support. My parents have always emphasized the importance of education, but also forced me to put my books down once in a while and go outside. I appreciate their willingness to let me find my own path, while still providing some guidance along the way; I am a better, more well-rounded person because of it. Finally, I thank my wife, Katie, for helping me to keep things in perspective and for being by my side throughout this entire experience. I am very grateful for her unwavering support and encouragement over the last six years.

Dedication

For my parents, who have supported and encouraged me throughout my entire life.

Abstract

The Laser Interferometer Gravitational-wave Observatory (LIGO) has recently reported the first two direct detections of gravitational waves, confirming yet another prediction of general relativity and providing an arena for testing gravity in the strong-field, high-velocity regime. These detections herald the beginning of the era of gravitational-wave astronomy, in which frequent gravitational-wave observations will provide a means for probing currently inaccessible regions of the universe and periods in its history.

In this dissertation, I focus on detecting and characterizing a stochastic gravitational-wave background and long-lasting gravitational-wave transients, which will be important components of this era for the information they can yield about the beginnings of the universe and the objects contained within it. I implement a method for estimating the parameters of a model of a stochastic gravitational-wave background and apply it to a model based on core-collapse supernovae. Using the expected sensitivities of Advanced LIGO and the Einstein Telescope, I estimate the detectability of such a background and compare the results to simulations of core-collapse events. I also develop an unmodeled all-sky search for long-lasting gravitational-wave transients and apply it to LIGO S5 and S6 data, setting the first upper limits on signals lasting between 10–500 s in the 40–1000 Hz band.

Finally, I describe the implementation of a 3D array of seismometers in and around the Homestake Mine, with the goal of mitigating seismic and Newtonian noise for future generations of gravitational-wave detectors. A seismic radiometer algorithm is developed and applied to simulated and real data; I demonstrate its ability to separate various components of the seismic field and map their directional dependence. In order to improve models of surface seismic waves in the radiometer algorithm, I use measurements of Rayleigh waves from the Homestake array and perform parameter estimation to fit a biexponential model of the Rayleigh wave eigenfunctions.

Table of Contents

Acknowledgments	i
Dedication	iii
Abstract	iv
List of Tables	x
List of Figures	xii
1 Gravitational waves and cosmology	1
1.1 Introduction	1
1.2 General relativity	2
1.2.1 Equivalence principles	3
1.2.2 Formalism	5
1.2.3 The Einstein field equations	11
1.2.4 Tests of general relativity	13
1.3 Gravitational waves	15
1.3.1 Linearized gravity and gauge conditions	16
1.3.2 Generation of gravitational waves	20
1.3.3 Sources of gravitational waves	23
1.3.4 Detection of gravitational waves	27
1.3.5 GW150914	31
1.4 Cosmology discussion	33
1.4.1 The standard cosmological model	34

1.4.2	Λ CDM model parameters	38
1.4.3	Cosmological timeline	39
2	Laser Interferometer Gravitational-wave Observatory	49
2.1	Detector description	51
2.1.1	Detector layout	51
2.1.2	Noise sources	54
2.2	Data analysis	59
2.2.1	Compact binary coalescences	60
2.2.2	Gravitational-wave bursts	62
2.2.3	Continuous waves	63
2.2.4	Stochastic backgrounds	64
2.3	Advanced LIGO	66
2.3.1	Upgrades	66
2.3.2	First observing run	69
2.3.3	Future plans	69
3	Searches for a stochastic gravitational-wave background	71
3.1	Sources of stochastic gravitational-wave backgrounds	72
3.1.1	Cosmological models	74
3.1.2	Astrophysical models	76
3.2	LIGO stochastic search	78
3.2.1	Effect of a plane gravitational wave on an interferometer	78
3.2.2	Cross-correlation estimator	80
3.2.3	Search pipeline	83
3.3	Parameter estimation for SGWB models	83
3.3.1	Formalism	83
3.3.2	Nested sampling	86
3.4	Star formation rates	87
3.4.1	Observations of star-forming galaxies	88
3.4.2	Gamma-ray burst rate	88
3.4.3	SFR models	89
3.4.4	Chemical evolution	90

3.4.5	High-mass, high-redshift SFR	93
3.4.6	Discussion	97
3.5	Parameter estimation for an SGWB from stellar core-collapse	97
3.5.1	Core-collapse spectrum	98
3.5.2	Parameter estimation	107
3.5.3	Discussion	117
4	Searches for long-lasting gravitational-wave transients	119
4.1	Astrophysical sources	120
4.1.1	Core-collapse supernovae	120
4.1.2	Neutron stars	122
4.1.3	Dynamically formed binary black holes	122
4.2	Analysis framework	123
4.2.1	A plane gravitational wave from a point source	123
4.2.2	Estimator for GW power	124
4.2.3	Frequency-time maps	128
4.2.4	Pattern recognition	128
4.3	Data quality	131
4.3.1	Formalism	133
4.3.2	Algorithm design	134
4.3.3	Background study	136
4.3.4	Safety for GW signals	140
4.3.5	Comparison with LIGO data quality flags	143
4.3.6	Conclusions	145
4.4	Seed-based clustering	146
4.4.1	Algorithm description	146
4.4.2	Examples	148
4.5	All-sky searches with STAMP	150
4.5.1	Clustering	153
4.5.2	Pre-processing of data	161
4.5.3	Data quality	162

5	An all-sky search for long-duration gravitational-wave transients in LIGO S5 and S6 data	164
5.1	Background studies	165
5.1.1	Data selection	165
5.1.2	Data quality	166
5.1.3	Background estimation	167
5.2	Search sensitivity	170
5.2.1	Waveforms	170
5.2.2	Sensitivity study	175
5.3	Loudest triggers	178
5.4	Upper limits	181
5.4.1	False alarm density and visible volume calculations	182
5.4.2	Bayesian formalism for FAD calculation	184
5.4.3	Efficiency calculation	186
5.4.4	Amplitude calibration uncertainties	186
5.4.5	Final upper limits	187
5.5	Discussion	187
5.6	Conclusions	191
6	Seismic noise and gravitational-wave detectors	192
6.1	Seismic waves	193
6.1.1	Body waves	193
6.1.2	Surface waves	194
6.2	Seismic noise in GW interferometers	197
6.3	Three-dimensional seismic array at Homestake	199
6.3.1	Array description	199
6.3.2	Station configuration	200
6.3.3	Data acquisition	210
6.3.4	Underground GPS timing	212
6.3.5	Data quality monitoring	215
6.3.6	Power spectra	219
6.4	Seismic radiometer	223

6.4.1	Formalism	225
6.4.2	Simulations	229
6.4.3	Analysis of real data	231
6.5	Parameter estimation for Rayleigh wave eigenfunctions	235
6.5.1	Rayleigh wave measurements	239
6.5.2	Parameter estimation	244
6.5.3	Application to seismic radiometer	248
6.5.4	Discussion	249
6.6	Summary	251
7	Conclusions	252
	References	255

List of Tables

1.1	Summary of different “fluids” used in the standard cosmological model. .	38
1.2	Current best values for measured and derived parameters of the Λ CDM cosmological model.	39
1.3	Predicted light element abundances from Big Bang nucleosynthesis, compared to observations. Abundances are given in terms of mass fraction relative to hydrogen.	45
2.1	Overview of the six LIGO science runs, including dates of operation and strain sensitivity achieved.	51
3.1	Parameters of three models of the star formation rate.	90
3.2	Best-fit parameters for the high-mass, high-redshift mode of star formation, determined by minimization in each bin.	96
3.3	Comparison of upper limits on ξ and A' for the combined model and the individual low- and high-frequency models.	115
3.4	Injected parameters of a simulated SGWB from stellar core-collapse, along with maximum likelihood parameters and 95% confidence intervals. . . .	117
5.1	Parameters of ADI waveforms used to test the sensitivity of the STAMP all-sky search.	171
5.2	Parameters of sinusoidal waveforms used to test the STAMP all-sky search sensitivity.	174
5.3	Parameters of sine-Gaussian waveforms used to test the sensitivity of the STAMP all-sky search.	174
5.4	Parameters of white noise burst waveforms used to test the sensitivity of the STAMP all-sky search.	175

5.5	Most significant triggers from the STAMP all-sky S5 and S6 zero-lag analyses.	181
5.6	Visible volume-based upper limits on ADI waveforms from the STAMP all-sky search.	187
5.7	Signal strength where the search achieves 50% efficiency for each waveform and corresponding energy estimate (S5 dataset).	189
5.8	Signal strength where the search achieves 50% efficiency for each waveform and corresponding energy estimate (S6 dataset).	190
6.1	Descriptions of the 24 seismic stations in the Homestake 3D array. . . .	202
6.2	Locations of the seismic stations in the Homestake 3D array.	204
6.3	Description of Q330 and Antelope seismic and state-of-health channels. .	211
6.4	Normal operating ranges for seismic station state-of-health channels. . .	219
6.5	Total recovered power across all directions as a function of frequency for P-waves, both S-wave polarizations, and three parametrizations of the Rayleigh wave model.	239
6.6	Details of seismic events used in the Rayleigh wave eigenfunction parameter estimation.	242
6.7	Rayleigh wave eigenfunction parameters estimated with <code>MultiNest</code>	246

List of Figures

1.1	Effect of a passing gravitational wave on a ring of freely-falling particles.	21
1.2	A simulated GW waveform from the inspiral, merger, and ringdown of a binary black hole system.	25
1.3	Measurements of the orbital period of the Hulse-Taylor binary system over a period of thirty years.	28
1.4	Detector strain time-series, reconstructed waveforms for GW150914, and a time-frequency plot of the strain data.	32
2.1	Aerial view of the two LIGO observatories at Hanford, WA and Livingston, LA.	50
2.2	Overview of the optical and sensing system of the LIGO interferometers.	52
2.3	Primary sources of displacement noise and sensing noise in the LIGO H1 detector.	55
2.4	Advanced LIGO sensitivity curve for different signal recycling cavity tunings and laser powers.	67
2.5	Optical configuration of the Advanced LIGO detectors.	68
2.6	Strain sensitivity for LIGO during O1 and SNR versus redshift plot for the merger of two $30 M_{\odot}$ black holes.	70
3.1	Landscape plot of SGWB energy density spectrum for various cosmological and astrophysical models.	73
3.2	Accessible regions of parameter space for a power law model of the stochastic gravitational-wave background.	86
3.3	Plot of three SFR models versus redshift.	91
3.4	2D maps of minimal χ^2 in the ν, z_m plane and a, b plane.	95

3.5	Star formation rate versus redshift for Model 4 and a high-mass mode of star formation.	96
3.6	$\Omega_{\text{GW}}(f)$ for various parameter choices for a model of an SGWB produced by stellar core-collapse.	100
3.7	$f \tilde{h}(f) $ from the <code>s27f_{heat}1.00</code> simulation, exhibiting a low-frequency peak at ≈ 60 Hz.	101
3.8	Raw $f \tilde{h}(f) $ spectrum and the resulting fits for several simulations of stellar core-collapse.	103
3.9	$\Omega_{\text{GW}}(f)$ for various parameter choices for a low-frequency model of an SGWB produced by stellar core-collapse.	106
3.10	Projected sensitivities of the stochastic search for Advanced LIGO and the Einstein Telescope.	107
3.11	95% sensitivity contours for model parameters of an SGWB from stellar core-collapse.	109
3.12	95% sensitivity contours for low-frequency peak parameters for SGWB from stellar core-collapse.	111
3.13	Confidence intervals on 2D posteriors for the combined model featuring low- and high-frequency components.	114
3.14	Projected 2D posteriors for a simulated SGWB from stellar core-collapse, recovered with the Einstein Telescope.	116
3.15	Plot of injected and maximum likelihood SGWB spectra compared to Einstein Telescope sensitivity.	117
4.1	Distribution of STAMP SNR statistic.	127
4.2	STAMP ft -maps of SNR, \hat{Y} , and $\hat{\sigma}_Y$ for simulated Gaussian noise and Gaussian noise plus a simulated GW signal from a black hole accretion disk instability.	129
4.3	STAMP ft -maps of P and \hat{P}^{adj} corresponding to both detectors for two cases: simulated Gaussian noise and Gaussian noise plus a simulated GW signal from an accretion disk instability.	130
4.4	Ft -maps of an example glitch in H1.	132
4.5	Ft -maps of SNR and SNR_{Ξ} for a glitch, well-behaved noise, and a simulated GW signal.	135

4.6	Histograms of SNR_{Ξ} ft -map pixels for LIGO data, Gaussian noise, glitches, and a simulated GW signal.	137
4.7	Scatter plots of injected ADI signals, time-shifted LIGO S5 data, and Gaussian noise in the plane of the glitch identification parameters $R_I(t)$ and \mathcal{F}_+	138
4.8	Plot of p -value versus SNR_{Γ} using a density-based clustering algorithm.	139
4.9	P -value versus SNR_{Γ} for three frequency bands: 100–250 Hz, 375–525 Hz, and 40–100 Hz.	141
4.10	Results of a safety study for simulated ADI signals.	142
4.11	ft -maps of SNR and SNR_{Ξ} for a ≈ 1 s accretion disk fragmentation waveform at 1 Mpc injected into Gaussian noise.	143
4.12	ft -maps of an ADI signal injected into LIGO time-shifted data containing a glitch.	144
4.13	P -value versus SNR_{Γ} for $1 \text{ s} \times 1 \text{ Hz}$ pixels of time-shifted LIGO S5 data with no flags applied, with the SNR_{Ξ} -based flag applied, and with the LIGO data quality flags applied (Categories 1–3).	145
4.14	Flow chart describing the burstegard clustering algorithm.	149
4.15	Recovery of a simulated ADI signal with the burstegard clustering algorithm.	151
4.16	Recovery of a simulated EBBH signal with the burstegard clustering algorithm.	152
4.17	SNR ft -map (including a signal) recovered with the wrong sky direction.	154
4.18	Recovery of a simulated GW signal using brute force iteration over sky positions compared to the fastring method.	155
4.19	Step-by-step example of zebragard clustering method.	157
4.20	Comparison of zebragard sensitivity using different numbers of sky directions.	159
4.21	Comparison of zebragard sensitivity using different methods to choose the sky directions.	160
4.22	Comparison of fastring , zebragard , and $\text{abs}(\text{SNR})$ -based searches.	160
4.23	Glitch cut safety for ADI and EBBH signals with an all-sky search	163

5.1	False alarm rate versus SNR_Γ for 100 time-shifts of LIGO data from the S5 and S6 science runs.	169
5.2	SNR ft -maps of four ADI waveforms used to test the sensitivity of the STAMP all-sky search.	172
5.3	SNR ft -maps of six sinusoidal signals used to test the sensitivity of the STAMP all-sky search.	173
5.4	SNR ft -maps of two sine-Gaussian signals used to test the STAMP all-sky search sensitivity.	175
5.5	SNR ft -maps of three white noise burst signals used to test the sensitivity of the STAMP all-sky search.	176
5.6	Detection efficiency curves with the STAMP all-sky search for all waveforms in S5 data.	178
5.7	Detection efficiency curves with the STAMP all-sky search for all waveforms in S6 data.	179
5.8	Distribution of FAR versus SNR_Γ for S5 and S6 zero-lag analyses.	180
5.9	Efficiency-based upper limits on all 15 waveforms used in the STAMP all-sky search.	188
6.1	Plot of Homestake array layout.	201
6.2	Photo of a Quanterra Q330 datalogger.	203
6.3	Photo of a Quanterra PB14 Packet Baler.	203
6.4	Seismometer vault enclosure at DEAD surface station.	205
6.5	Sealed seismometer vault at TPK surface station.	206
6.6	Full RRDG station, including sealed seismometer vault, electronics doghouse, solar panels, GPS antenna, and radio antenna.	207
6.7	Seismometer enclosure for underground stations.	208
6.8	Underground station data acquisition system and electronics.	209
6.9	Real-time data flow for Homestake stations using telemetry.	212
6.10	Flow chart of underground GPS timing distribution system for the Homestake 3D seismic array.	214
6.11	Example of TTL to RS-232 conversion at the 4850 level.	215
6.12	Front and rear views of a GPS transceiver.	216
6.13	Screenshot of main page of Homestake data quality monitor.	217

6.14	Screenshot of a daily page from the Homestake data quality monitor. . .	218
6.15	Plots of all 10 channels used in the Homestake data quality monitor. . .	220
6.16	Median amplitude spectral densities for all Homestake stations.	222
6.17	Histograms of amplitude spectral density in each frequency bin for the RRDG and A4850 stations.	224
6.18	Directional recovery of a simulated P-wave using the seismic radiometer analysis.	230
6.19	Directional recovery of a simulated S-wave using the seismic radiometer analysis.	232
6.20	Directional recovery of a simulated Rayleigh wave using the seismic ra- diometer algorithm.	233
6.21	Analysis of background noise using the P-wave component of the seismic radiometer algorithm.	234
6.22	S-wave analysis of background noise using the seismic radiometer algo- rithm (0.1 Hz and 0.2 Hz).	235
6.23	S-wave analysis of background noise using the seismic radiometer algo- rithm (0.5 Hz and 1 Hz).	236
6.24	Rayleigh wave analysis of background noise using the seismic radiometer algorithm (0.1 Hz and 0.2 Hz).	237
6.25	Rayleigh wave analysis of background noise using the seismic radiometer algorithm (0.5 Hz and 1 Hz).	238
6.26	Time-series plot of all seismic channels for the 300 station during an earth- quake.	240
6.27	Raw and filtered time-series traces and velocity heatmap corresponding to a seismic event.	243
6.28	Rayleigh wave amplitude measurements and eigenfunction models for 0.1– 0.6 Hz.	247
6.29	Rayleigh wave amplitude measurements and eigenfunction models for 0.7– 1.0 Hz.	248
6.30	Analysis of background noise with the biexponential Rayleigh wave model in the seismic radiometer algorithm.	250

Chapter 1

Gravitational waves and cosmology

1.1 Introduction

Over the past 30 years, cosmology has come into its own as an experimental science, thanks to increasingly more precise electromagnetic observations of the universe. Detection and study of the Cosmic Microwave Background (CMB), coupled with observations of Type Ia supernovae and large scale structure have facilitated the development of a “standard model” of cosmology, the Λ CDM model. Coupled with inflation, this model explains the abundance of light nuclei, the anisotropy of the CMB, and many other cosmological and astrophysical phenomena.

We are just beginning a new phase in our study of the universe; not only can we “see” the universe electromagnetically, but we can also “hear” the universe through observations of gravitational waves (GWs). Predicted by Einstein’s theory of general relativity, these ripples in spacetime carry information about the earliest moments of the universe and the most exotic objects contained in it. The quest for the direct detection of gravitational waves began in the 1960s with Joseph Weber’s bar detectors, and has recently been achieved by the large-scale interferometric detectors of the Laser Interferometer Gravitational-Wave Observatory (LIGO) [1]; further detections of GWs will establish the new era of gravitational-wave astronomy. This dissertation will describe analyses of LIGO data with the goal of detecting long-lasting gravitational waves and estimating the parameters of their sources, as well as improving future generations of gravitational-wave detectors.

The remainder of this chapter will provide an introduction to general relativity, gravitational waves, and cosmology. In Sec. 1.2, general relativity will be introduced, culminating in the Einstein field equations. Sec. 1.3 will provide an overview of gravitational waves. Finally, Sec. 1.4 will describe the Λ CDM model of cosmology and where gravitational waves fit into this model.

1.2 General relativity

In 1905, Einstein first proposed his theory of special relativity, which unified classical Newtonian mechanics with Maxwell’s electrodynamics. This theory contended that the laws of physics should be invariant across all inertial reference frames, and that the speed of light should be the same for all observers and is the “universal speed limit.” One of the results of special relativity is that space and time are not separate entities as we commonly experience them, but are joined together in the construct of *spacetime*.

Special relativity is only valid under “special” circumstances; namely, when dealing with a flat spacetime (represented by the Minkowski metric $\eta_{\mu\nu}$).

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1.1)$$

In general, the metric is used to calculate the spacetime interval ds^2 between two events; in special relativity, it is given by:

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu = -dt^2 + dx^2 + dy^2 + dz^2. \quad (1.2)$$

Here and in most of this section, we set $c = 1$.

At small enough scales, spacetime can be shown to be locally flat, and special relativity is applicable. However, the gravitational redshift, one of Einstein’s first steps toward general relativity, implied that sufficiently extended regions of spacetime must be curved [2]. Along with this, Newton’s laws of gravitation (Eqs. 1.3 and 1.4) are clearly

not compatible with special relativity.

$$\frac{d^2\mathbf{x}}{dt^2} = -\frac{\partial\Phi}{\partial\mathbf{x}} \quad (1.3)$$

$$\nabla^2\Phi = 4\pi G\rho \quad (1.4)$$

Here, Φ is the gravitational potential, G is the gravitational constant, and ρ is the mass density. The equations of motion given in Eq. 1.3 are not Lorentz-invariant, nor is the gravitational potential given in Eq. 1.4. Eq. 1.4 also implies that changes in the density ρ propagate to the potential Φ instantaneously, which violates the idea of a maximum speed limit, one of the primary principles of relativity.

Einstein was able to resolve these issues in his theory of general relativity by realizing that gravity is a result of curved spacetime caused by the presence of mass or energy. Or, as John Wheeler said [3]: “Spacetime tells matter how to move, matter tells spacetime how to curve.”

1.2.1 Equivalence principles

There are a few useful principles which provide qualitative insight into the workings of general relativity and led Einstein along the path to its full development.

The weak equivalence principle

The weak equivalence principle states that the inertial mass m_i and the gravitational mass m_g are equal.

$$m_i = m_g \quad (1.5)$$

In Newtonian mechanics, a general force can be described by $F = m_i a$, while the gravitational force is given by $F_g = m_g \nabla\Phi$. This implies that the acceleration of an object under the influence of gravity is independent of its mass.

At first glance, this does not seem especially surprising, but if we instead think of gravitational mass as “gravitational charge,” it becomes more intriguing. For example, a particle with electric charge q and mass m in an electric field E will experience an acceleration $a = (q/m)E$, which is proportional to the ratio of its electric charge and inertial mass. In the case of gravity, it just so happens that the ratio of gravitational

charge to inertial mass is one.

The weak equivalence principle implies that a gravitational field cannot be differentiated from a uniformly accelerated frame in small enough regions of spacetime. Imagine a physicist in a small box, who has no contact with outside world, and is observing the motion of freely-falling particles under the influence of only gravitational forces. Because all of the particles accelerate at the same rate (independent of mass), this physicist cannot tell whether he or she is in a gravitational field or in an accelerating spaceship. Again, this assumes that the experiments are restricted to “small enough” regions of spacetime. This is certainly not the case for charged particles in an electromagnetic field, since there exist particles with many different combinations of electric charge and mass.

The Einstein equivalence principle

Einstein expanded on the weak equivalence principle with the Einstein equivalence principle, which states that the outcome of any local non-gravitational experiment is independent of the velocity of its freely-falling reference frame and its location in space and time. Although inertial frames cannot be extensive (due to gravity), we can define locally inertial frames that track the motion of freely-falling particles in small enough regions of spacetime [4]. This implies that the laws of physics should reduce to those of special relativity in these regions. As we will see in Sec. 1.2.2, this will allow the laws of special relativity to be generalized to curved spacetime if they are formulated in a generally covariant way.

One of the most famous predictions of the Einstein equivalence principle is that of the gravitational redshift. This means that a photon emitted in a gravitational field will be reduced in frequency if observed in a weaker region of a gravitational field than it was emitted in. This can be seen through an analogy with two rockets, separated by a distance d , and both moving with some acceleration a . The trailing rocket emits a photon with wavelength λ_0 , which the leading rocket receives after some time $\Delta t = d/c$ (assuming that the rockets are moving at a velocity much less than the speed of light). When the leading rocket receives the photon, it will have picked up an additional velocity $\Delta v = a\Delta t$. Due to the Doppler effect, the leading rocket will perceive the photon to be redshifted by an amount

$$\frac{\Delta\lambda}{\lambda_0} = \frac{\Delta v}{c} = \frac{ad}{c^2}. \quad (1.6)$$

The Einstein equivalence principle states that the same effect should occur in a uniform gravitational field, since we cannot distinguish between a gravitational field and a uniformly accelerated frame of reference. For example, imagine a photon of wavelength λ_0 emitted by an observer on the ground, which is perceived by an observer standing on top of a tower of height d . This observer should also see a redshifted photon, just as the observer in the leading rocket did. The amount of redshift in this situation can be calculated with Eq. 1.6, replacing a with the gravitational acceleration a_g .

There is also a strong equivalence principle, which generalizes the Einstein equivalence principle from local “non-gravitational” experiments to any type of local experiment, but it will not be discussed further here.

1.2.2 Formalism

Much of the formalism of general relativity is based around the mathematics of differential geometry and tensor calculus. The cornerstone of general relativity, the Einstein field equations, are based around elementary concepts in these fields; the real genius of Einstein was his realization that these ideas could be applied to space and time, and that spacetime curvature is generated by mass and energy.

One of the fundamental principles involved in Einstein’s formulation of general relativity was the idea of general covariance, which states that the laws of physics should not depend on the choice of coordinate system. This concept seems intuitive, considering that coordinates are not actually a part of nature, but are constructs used by humans to describe nature. As an example, consider some physical law $A^\mu = B^\mu + C^\mu$ defined in one coordinate system, where the index μ runs over coordinates. If each side of the equation is a tensor (i.e., it transforms in a specific way that will be discussed later), then general covariance states that this law should have the same general form in any other coordinate system (denoted by primes): $A'^\mu = B'^\mu + C'^\mu$. This is an extremely powerful statement and allows the generalization of physical laws formulated in one coordinate system to any other coordinate system, but it is only true if each side of the equation transforms as a tensor.

We begin our study of general relativity with a metric tensor $g_{\mu\nu}$ which, in general, encapsulates the geometric structure of spacetime and defines an inner product for tangent vectors. Unlike special relativity, we make no assumptions about the curvature (or

lack thereof) of spacetime.

Affine connection

To proceed further, it is instructive to consider a freely-falling coordinate system ξ^μ in which a particle is at rest. The equations of motion in this frame are given by

$$\frac{\partial^2 \xi^\lambda}{\partial \tau^2} = 0, \quad (1.7)$$

where τ is the proper time. Because the coordinate system is locally inertial, we use the Minkowski metric to calculate the spacetime interval: $ds^2 = \eta_{\mu\nu} \xi^\mu \xi^\nu$.

But how would we describe this particle's motion in another coordinate system with coordinates x^μ ? The equations of motion in the new coordinate system can be derived by considering that $\frac{\partial^2 \xi^\lambda}{\partial \tau^2} = \frac{\partial}{\partial \tau} \left(\frac{\partial \xi^\lambda}{\partial x^\nu} \frac{\partial x^\nu}{\partial \tau} \right)$. This results in the following equations of motion:

$$\frac{\partial^2 x^\lambda}{\partial \tau^2} + \frac{\partial x^\lambda}{\partial \xi^\alpha} \frac{\partial^2 \xi^\alpha}{\partial x^\mu \partial x^\nu} \frac{\partial x^\mu}{\partial \tau} \frac{\partial x^\nu}{\partial \tau} = 0. \quad (1.8)$$

This equation, which defines the paths that the world lines of freely-falling particles follow, is known as the geodesic equation. The shortest path between two points is not always a straight line; in a curved spacetime, it is a geodesic.

Of particular importance is the new term which has appeared in Eq. 1.8, proportional to two first derivatives of the coordinates x^μ . The coefficient multiplying these two first derivatives is known as the affine connection of the transformation:

$$\Gamma^\lambda_{\mu\nu} = \frac{\partial x^\lambda}{\partial \xi^\alpha} \frac{\partial^2 \xi^\alpha}{\partial x^\mu \partial x^\nu}. \quad (1.9)$$

In general, the affine connection gives a way of comparing tensors which lie in different tangent spaces by parallel transporting them so that they are in the same tangent space. This is not something we previously had to worry about in the flat, Euclidean geometry of special relativity, but is crucial when working in a curved spacetime. In the 4D Lorentzian spacetime of general relativity, the affine connection coefficients are also known as Christoffel symbols. The label of ‘‘symbols’’ here emphasizes the fact that these objects do not transform like tensors under general coordinate transformations.

Rather, the Christoffel symbols must be recalculated when the coordinate system is redefined.

Typically, we begin with a metric $g_{\mu\nu}$ and from there, calculate the Christoffel symbols and other important aspects of the curvature. To do this, we consider the spacetime interval in each coordinate system. For the coordinate system in which the particle is freely-falling, the spacetime interval is given by

$$ds^2 = \eta_{\mu\nu} d\xi^\mu d\xi^\nu. \quad (1.10)$$

In the other coordinate system, the spacetime interval is

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu. \quad (1.11)$$

This gives a relation between the general metric $g_{\mu\nu}$ and the metric in the freely-falling coordinate system:

$$g_{\mu\nu} = \eta_{\alpha\beta} \frac{\partial \xi^\alpha}{\partial x^\mu} \frac{\partial \xi^\beta}{\partial x^\nu} \quad (1.12)$$

Taking a derivative of $g_{\mu\nu}$ with respect to x^λ and doing a few algebraic tricks allows us to obtain the Christoffel symbols as a function of the metric.

$$\Gamma^\lambda_{\mu\nu} = \frac{1}{2} g^{\sigma\lambda} \left(\frac{\partial g_{\mu\sigma}}{\partial x^\nu} + \frac{\partial g_{\nu\sigma}}{\partial x^\mu} - \frac{\partial g_{\mu\nu}}{\partial x^\sigma} \right) \quad (1.13)$$

Covariant derivative

Going back to the principle of general covariance, it is evident that it will be very useful to define physical laws in a coordinate-invariant way. We first consider a general transformation of a vector (or a tensor of rank 1) from an unprimed coordinate system to a primed one:

$$V'^\mu = \frac{\partial x'^\mu}{\partial x^\nu} V^\nu. \quad (1.14)$$

For a tensor of arbitrary rank, each index transforms independently in the same way as the vector's single index did. To see this, we consider the transformation of a mixed tensor of rank 3:

$$T'^{\mu}_{\nu}{}^{\lambda} = \frac{\partial x'^\mu}{\partial x^\eta} \frac{\partial x^\rho}{\partial x'^\nu} \frac{\partial x'^\lambda}{\partial x^\sigma} T^{\eta}_{\rho}{}^{\sigma}. \quad (1.15)$$

One problem that we will run into right away, however, is that the partial derivative operator is not generally covariant. This is apparent if we define a new tensor $B_\mu{}^\nu$ of rank 2 as the partial derivative of a vector:

$$B_\mu{}^\nu = \frac{\partial}{\partial x^\mu} A^\nu. \quad (1.16)$$

Then, after performing a general coordinate transformation, we obtain

$$\begin{aligned} B_\mu{}^\nu &= \frac{\partial}{\partial x'^\mu} A'^\nu \\ &= \frac{\partial}{\partial x'^\mu} \left(\frac{\partial x'^\nu}{\partial x^\sigma} A^\sigma \right) \\ &= \frac{\partial x'^\nu}{\partial x^\sigma} \frac{\partial x^\rho}{\partial x'^\mu} \frac{\partial A^\sigma}{\partial x^\rho} + \frac{\partial^2 x'^\nu}{\partial x^\sigma \partial x^\rho} \frac{\partial x^\rho}{\partial x'^\mu} A^\sigma. \end{aligned} \quad (1.17)$$

The first term in Eq. 1.17 is what we would expect for a tensor transformation, but the presence of the second term indicates that $B_\mu{}^\nu$ is not a tensor. Thus, the partial derivative operator acting on a tensor generates an object that is not a tensor and is not generally covariant. Intuitively, this makes sense: the partial derivative operator acts essentially as a projection operator in the tangent direction, but the direction of the tangent planes are also functions of the coordinates. This hints that the affine connection may be helpful here, as it connects tangent spaces on a manifold.

We would like to define a new operator which acts similarly to the partial derivative, but transforms like a tensor and is generally covariant. With a bit of insight and some algebra, along with the assistance of the affine connection, we can obtain the *covariant derivative*, shown here acting on a vector V^ν :

$$\nabla_\mu V^\nu = \partial_\mu V^\nu + \Gamma^\nu{}_{\mu\sigma} V^\sigma. \quad (1.18)$$

Here, $\partial_\mu = \frac{\partial}{\partial x^\mu}$. Although neither the first nor second terms in Eq. 1.18 transform independently as tensors, the non-tensorial components in the transformation cancel each other out exactly; as a result, the sum of the terms does transform as a tensor. This covariant derivative allows us to specify a derivative in a coordinate-invariant way because it transforms tensors into tensors. We also note here that the covariant derivative reduces to the ordinary partial derivative in a flat spacetime, since the Christoffel symbols

are trivially zero in that case.

It is also useful to introduce the comma and semicolon notation for partial and covariant derivatives. A comma indicates a partial derivative,

$$V^\nu_{,\mu} = \partial_\mu V^\nu \quad (1.19)$$

while a semicolon indicates a covariant derivative.

$$V^\nu_{;\mu} = \nabla_\mu V^\nu = V^\nu_{,\mu} + \Gamma^\nu_{\mu\sigma} V^\sigma \quad (1.20)$$

This brings to mind the “comma-goes-to-semicolon” rule, which states that to convert an equation from its special relativistic version in a local inertial frame to a generally covariant form valid in any coordinate system, we just replace the partial derivatives with covariant derivatives, or equivalently, commas with semicolons. From this point on, we will primarily use the comma and semicolon notation.

Curvature

Consider parallel transporting a vector around a closed loop. In Euclidean space, choose any loop you like: in the end, the vector will point in the same direction as it did initially. But now imagine doing the same thing around a closed loop on a sphere. As an example, begin at the equator, move 90° clockwise, then up to the north pole, then back down to the starting point along the initial line of longitude. The vector will now point 90° clockwise of its initial orientation. Thus, it is evident that the assumption of invariance under parallel transport does not hold in a curved geometry.

The Riemann curvature tensor (sometimes known as the Riemann-Christoffel curvature tensor, or just the curvature tensor) $R^\rho_{\sigma\mu\nu}$ describes the amount of deviation under parallel transport from that of Euclidean space. It can be constructed from the commutator of two covariant derivatives acting on a vector.

$$\begin{aligned} [\nabla_\mu, \nabla_\nu] V^\rho &= \nabla_\mu \nabla_\nu V^\rho - \nabla_\nu \nabla_\mu V^\rho \\ &= R^\rho_{\sigma\mu\nu} V^\sigma \end{aligned} \quad (1.21)$$

This commutator effectively measures the difference between parallel transporting a

vector halfway around a loop in one direction versus halfway around the loop in the opposite direction. After a bit of algebra, we obtain the following expression for the Riemann curvature tensor:

$$R^\rho{}_{\sigma\mu\nu} = \Gamma^\rho{}_{\sigma\nu,\mu} - \Gamma^\rho{}_{\sigma\mu,\nu} + \Gamma^\rho{}_{\sigma\mu}\Gamma^\alpha{}_{\alpha\nu} - \Gamma^\rho{}_{\sigma\nu}\Gamma^\alpha{}_{\alpha\mu}. \quad (1.22)$$

Here, we have given the definition of the Riemann curvature tensor with the first index raised; this tensor is commonly found in this form or with all lower indices. In Euclidean space, the commutator of two covariant derivatives is zero (since the Christoffel symbols and their derivatives are zero); the Riemann curvature tensor is trivially zero only in this special case.

There are several useful properties of the Riemann curvature tensor, including symmetry and antisymmetry under interchange of certain indices and pairs of indices. However, we will not cover these properties other than to state the Bianchi identities, which can be derived from them:

$$R_{\rho\sigma\mu\nu;\lambda} + R_{\sigma\lambda\mu\nu;\rho} + R_{\lambda\rho\mu\nu;\sigma} = 0. \quad (1.23)$$

The Bianchi identities will be useful later when deriving the Einstein field equations.

The Ricci tensor $R_{\mu\nu}$ can be formed from the Riemann curvature tensor by contracting the first and third indices:

$$R_{\mu\nu} = g^{\alpha\beta} R_{\alpha\mu\beta\nu} \quad (1.24)$$

It can be shown that all other contractions of indices of the Riemann curvature tensor either give something proportional to the Ricci tensor, as defined here, or zero. This occurs due to the properties of the Christoffel symbols which form the affine connection in general relativity and is not the case for a curvature tensor formed from arbitrary connection coefficients. Thus, for our general relativistic purposes, this is the only meaningful rank 2 tensor that we can derive from the Riemann curvature tensor.

Finally, the trace of the Ricci tensor is called the Ricci scalar, or scalar curvature:

$$R = R^\mu{}_\mu = g^{\mu\nu} R_{\mu\nu}. \quad (1.25)$$

1.2.3 The Einstein field equations

Now that we have developed some background in curvature, the real problem is to determine how one might relate the curvature of spacetime to mass and energy. Put another way, we would like to find some tensor generalization of Eqs. 1.3 and 1.4 to curved spacetime. One straightforward modification would be to change the mass density ρ in Eq. 1.4 to its tensor version, $T_{\mu\nu}$. The stress-energy tensor is most easily defined by considering its components [2]:

$$T_{00} = \text{mass-energy density} \quad (1.26)$$

$$T_{j0} = \text{density of } j\text{th component of momentum} \quad (1.27)$$

$$T_{jk} = j, k \text{ component of stress or} \quad (1.28)$$

$k\text{th component of flux of } j\text{th component of momentum.}$

First, we consider the geodesic equation (Eq. 1.8) in the Newtonian limit: low velocity and weak gravitational fields. In this limit, the following simplifications can be made:

$$\begin{aligned} \frac{dt}{d\tau} &\approx 1 \\ v &\ll 1 \\ g_{\mu\nu} &\approx \eta_{\mu\nu}. \end{aligned} \quad (1.29)$$

Using these relations, Eq. 1.8 reduces to

$$\begin{aligned} \frac{d^2 x_i}{dt^2} &= -\Gamma_{i00} \\ &= -\frac{1}{2} h_{00,i} \end{aligned} \quad (1.30)$$

Equating this with Eq. 1.3 implies that $h_{00} = -2\Phi$, and thus, $g_{00} = -(1 + 2\Phi)$ in the Newtonian limit. From this comparison, we should expect that the generally covariant form of our equation will include second derivatives of the metric instead of the potential Φ . We know that the Riemann curvature tensor is formed of second derivatives of the metric since it contains derivatives of the Christoffel symbols, which themselves contain first derivatives of the metric. However, it has too many indices compared to $T_{\mu\nu}$;

instead, we might consider the Ricci tensor, since it also contains second derivatives of the metric and only has two indices.

At this point, we have something like

$$R_{\mu\nu} = \alpha T_{\mu\nu}, \quad (1.31)$$

where α is a constant of proportionality. This seems reasonable (and was actually considered by Einstein at one point), but there is a problem with it: energy conservation [4]. This is enforced by requiring the divergence of the stress-energy tensor to be zero:

$$T_{\mu\nu}{}^{;\mu} = 0. \quad (1.32)$$

This indicates that the divergence of $R_{\mu\nu}$ should also be zero (see Eq. 1.31). However, this is not the case; if we contract two of the indices in the Bianchi identities (Eq. 1.23), we obtain

$$R_{\mu\nu}{}^{;\mu} = \frac{1}{2}R_{;\nu}. \quad (1.33)$$

From this, we can see that Eq. 1.31 is not the equation we are looking for. However, Eq. 1.33 clearly suggests a new idea for the left side of Eq. 1.31 which is divergence-free: $R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}$. Thus, we might consider

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \alpha T_{\mu\nu}. \quad (1.34)$$

Does this equation fulfill our requirements? First, it relates curvature to energy and momentum in a tensor equation and it also holds under conservation of energy. We required that the left-hand side of the equation should contain second derivatives of the metric, and it does. Finally, we can show that it will reduce to Eqs. 1.3 and 1.4 in the Newtonian limit if the constant of proportionality α is equal to $8\pi G$.

This gives the final form of the canonical equations of general relativity, the Einstein field equations.

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi GT_{\mu\nu} \quad (1.35)$$

The left-hand side of Eq. 1.35 is often grouped into one object called the Einstein tensor;

using this, the Einstein field equations may also be written as

$$G_{\mu\nu} = 8\pi GT_{\mu\nu}. \quad (1.36)$$

Given some distribution of mass and/or energy, the goal is typically to solve the field equations for a metric. In certain cases, it is possible to solve the system of equations exactly; in other cases, post-Newtonian approximations are made. These approximations use expansions in small parameters which express the deviation of the theory from the Newtonian limit.

Einstein also considered including a cosmological constant term proportional to the metric in order to allow for a static universe. With the inclusion of this constant, the field equations take on the form

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi GT_{\mu\nu}, \quad (1.37)$$

where Λ is the cosmological constant. At the time, it was thought that the universe was static, although Hubble's observations of the expanding universe would soon disprove that theory. However, a universe with a cosmological constant may also be expanding, so including this term was certainly not "the biggest blunder of Einstein's life," as he put it.

We also note that the Einstein field equations may be rigorously derived from the Einstein-Hilbert action (rather than the informal derivation given here) [4].

1.2.4 Tests of general relativity

In the field of physics, it may be that experimental results generate new theoretical developments or that theoretical predictions drive new experimental work; in the case of general relativity, it has primarily been the latter. It has taken many years for measurement technology to become accurate enough for precision testing of Einstein's theory of general relativity. During this time, some challengers to general relativity have arisen, but none have proved as compatible with observations (or as elegant) as Einstein's theory. Below, I will describe some well-known experimental tests of general relativity.

Anomalous perihelion precession of Mercury

In 1859, Urbain Le Verrier discovered that the perihelion position of Mercury’s orbit was precessing at a rate slightly different than that predicted by Newtonian mechanics [5]. The difference between prediction and observation was a mere 43 arcseconds per century, but this was well above the experimental uncertainty of the observations. Several other hypotheses were proposed, including unobserved dust between the Sun and Mercury, and the presence of an undiscovered moon or planet; however, none of these objects were ever found. Einstein explained this discrepancy in 1915 using his new theory of general relativity to the level of $\approx 0.1\%$ [2, 6].

The gravitational redshift

The gravitational redshift was described previously in Sec. 1.2.1. It was first experimentally observed in 1959 by Pound and Rebka [7]. They used two samples of ^{57}Fe : a moving “source,” from which the emitted gamma rays would be slightly Doppler shifted, and a stationary “receiver” either below or above the moving source (both configurations were used). The experiment was designed such that when the gravitational redshift was exactly matched by the Doppler shift due to the source motion, the emitted gamma rays would have the required frequency for absorption by the receiver. They observed a net fractional frequency shift of $(-5.13 \pm 0.51) \times 10^{-15}$, compared to general relativity’s prediction of -4.92×10^{-15} [7].

Gravitational lensing

Gravitational lensing is the bending of light by a massive object. The existence of this phenomenon becomes apparent when we consider that in a curved spacetime, light should follow geodesics, rather than the Euclidean equivalent (a straight line). Thus, as the presence of mass and energy causes spacetime to warp, the paths of light rays should be bent, as well. This was first observed by Eddington during a solar eclipse in 1919 [8]. Gravitational lensing is also predicted by Newtonian gravity, but it is too weak by a factor of two compared to observations.

Frame dragging

Frame dragging is an effect where a rotating massive object distorts nearby spacetime and causes inertial frames to be “dragged” along with the object’s rotation. The Gravity Probe B experiment tested this phenomenon in 2004, using gyroscopes orbiting Earth in a satellite [9]. In 2011, the results of the analysis were published and shown to be compatible with general relativity within the relevant experimental and theoretical uncertainties.

1.3 Gravitational waves

The recent direct detection of gravitational waves [10] can be heralded as another triumph of general relativity. Gravitational waves can be thought of as small ripples in the curvature of spacetime; they may be treated as small deviations from a flat metric and are generally studied using perturbation theory. GWs were first predicted by Einstein in 1916, but it is unlikely that he gave them much consideration; due to the incredibly small strain amplitudes which were expected, GWs seemed far beyond the precision of our measurement technologies until the last few decades. However, the fact that gravitational waves interact very weakly is also a benefit; any gravitational waves that we may detect will reach us essentially untarnished. Because of this, gravitational waves carry significant information about the astrophysical objects or cosmological mechanisms which produced them.

As mentioned in Sec. 1.1, most of our knowledge of the universe up to this point is based on electromagnetic observations, in spite of the fact that light couples relatively strongly to matter. Further observations of weakly interacting gravitational waves are expected to open a new window through which we may study the universe and the objects contained in it.

In this section, I will discuss the prediction of gravitational waves from general relativity and the mechanics of how they are generated. I will also give a brief overview of some of the most interesting sources of gravitational waves and some strategies for detecting them.

1.3.1 Linearized gravity and gauge conditions

Gravitational waves appear most clearly in general relativity when we consider a “linearized gravity” framework, in which GWs are treated as perturbations around a flat spacetime described by the Minkowski metric. This describes a scenario where the observer is far from the actual source of gravitational waves and spacetime is essentially flat, apart from the GW perturbations. The full metric in linearized gravity is given by

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.38)$$

where $h_{\mu\nu}$ is the GW perturbation. Under the requirement that the GW perturbation is small (i.e., $|h_{\mu\nu}| \ll 1$), terms beyond first-order in h can be dropped.

First, it is useful to calculate the Christoffel symbols, as they can be used to determine the Riemann curvature tensor, the Ricci tensor, and the Ricci scalar. The Christoffel symbols are given by Eq. 1.9 in general; using the assumptions of linearized gravity, we obtain

$$\Gamma^\alpha{}_{\beta\gamma} = \frac{1}{2} (h^\alpha{}_{\beta,\gamma} + h^\alpha{}_{\gamma,\beta} - h_{\beta\gamma}{}^{,\alpha}). \quad (1.39)$$

The terms proportional to Γ^2 in the Ricci tensor can be dropped, since they are second-order in h . This gives

$$\begin{aligned} R_{\mu\nu} &\simeq \Gamma^\alpha{}_{\mu\nu,\alpha} - \Gamma^\alpha{}_{\mu\alpha,\nu} \\ &\simeq \frac{1}{2} (h^\alpha{}_{\nu,\mu\alpha} + h_{\mu\alpha}{}^{,\alpha}{}_{\nu} - \square h_{\mu\nu} - h_{,\mu\nu}), \end{aligned} \quad (1.40)$$

where $\square = \partial_\mu \partial^\mu$ is the d’Alembertian, and $h = h^\alpha{}_\alpha$ is the trace of $h_{\mu\nu}$. Finally, the Ricci scalar is calculated from the Ricci tensor using Eq. 1.25:

$$R = h_{\alpha\beta}{}^{,\alpha\beta} - \square h \quad (1.41)$$

This yields the Einstein field equations in linearized gravity:

$$h^\alpha{}_{\nu,\mu\alpha} + h_{\mu\alpha}{}^{,\alpha}{}_{\nu} - \square h_{\mu\nu} - h_{,\mu\nu} + \eta_{\mu\nu} h_{\alpha\beta}{}^{,\alpha\beta} - \eta_{\mu\nu} \square h = 16\pi T_{\mu\nu} \quad (1.42)$$

At this point, it is common to change variables to the trace-reversed perturbation

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h, \quad (1.43)$$

which has a trace $\bar{h} = -h$. We note here that $\bar{h}_{\mu\nu}$ contains the same information as $h_{\mu\nu}$. If we recalculate the metric in terms of $\bar{h}_{\mu\nu}$, this becomes clear.

$$g_{\mu\nu} = \left(1 - \frac{1}{2}\bar{h}\right)\eta_{\mu\nu} + \bar{h}_{\mu\nu} \quad (1.44)$$

Using the trace-reversed perturbation, the Einstein field equations can be written as

$$\bar{h}_{\nu\alpha}{}^{,\alpha}{}_{,\mu} + \bar{h}_{\mu\alpha}{}^{,\alpha}{}_{,\nu} - \square\bar{h}_{\mu\nu} - \eta_{\mu\nu}\bar{h}_{\alpha\beta}{}^{,\alpha\beta} = 16\pi T_{\mu\nu}. \quad (1.45)$$

Harmonic gauge

At this point, it is useful to fix the gauge by choosing a coordinate system. For the problem at hand, the most common choice is the harmonic gauge (also called the de Donder gauge or the Lorenz gauge), which imposes the requirement that

$$\square x^\mu = 0 \quad (1.46)$$

This can be equivalently formulated as

$$g^{\alpha\beta}\Gamma^\mu{}_{\alpha\beta} = 0 \quad (1.47)$$

or, in linearized gravity,

$$\bar{h}{}^{\mu\nu}{}_{,\nu} = 0. \quad (1.48)$$

The formulation of this gauge given in Eq. 1.48 is analogous to the Lorenz gauge condition for the four-potential in electromagnetism ($A^\mu{}_{,\mu} = 0$).

In the harmonic gauge, only the third term in Eq. 1.45 is non-zero. The full Einstein field equations become

$$\square\bar{h}_{\mu\nu} = -16\pi GT_{\mu\nu}, \quad (1.49)$$

while the vacuum equations are

$$\square \bar{h}_{\mu\nu} = 0. \quad (1.50)$$

Eqs. 1.49 and 1.50 are wave equations; this is the theoretical prediction of wave solutions to Einstein's field equations, or gravitational waves.

This simplest solutions to the vacuum equation are plane waves:

$$\bar{h}_{\mu\nu} = \text{Re} \left[A_{\mu\nu} e^{ik_\alpha x^\alpha} \right] \quad (1.51)$$

where $A_{\mu\nu}$ is the wave amplitude tensor and $k_\alpha = (\omega, k_x, k_y, k_z)$ is the wave vector. For non-trivial values of $\bar{h}_{\mu\nu}$, the field equations (Eq. 1.50) require that the wave vector is a null vector; i.e., $k_\alpha k^\alpha = 0$, or $k^0 = (k_x^2 + k_y^2 + k_z^2)^{1/2}$. This can also be interpreted as the requirement that gravitational waves should propagate at the speed of light.

Applying the harmonic gauge condition (Eq. 1.48) to the plane wave gives the requirement that

$$A_{\mu\alpha} k^\alpha = 0, \quad (1.52)$$

which is the requirement that the wave amplitude tensor is orthogonal to the wave vector.

Transverse-traceless gauge

The wave amplitude tensor $A_{\mu\nu}$ has 16 components in a 4D spacetime, but several of the requirements that we have imposed have restricted the independence of these components. Because the stress-energy tensor is symmetric in general, $\bar{h}_{\mu\nu}$ should be symmetric, and so should $A_{\mu\nu}$ (see Eq. 1.49). This leaves us with 10 independent components. The harmonic gauge condition applied to our GW (see Eq. 1.52) applies four more constraints, leaving us with 6 independent components in $A_{\mu\nu}$.

However, the field equations and the harmonic gauge condition in linearized gravity are still invariant under coordinate transformations like

$$x'^\mu = x^\mu + \epsilon^\mu \quad (1.53)$$

$$h'_{\mu\nu} = h_{\mu\nu} - \partial_\mu \epsilon_\nu - \partial_\nu \epsilon_\mu. \quad (1.54)$$

This freedom allows us to set four more gauge conditions, which will reduce the independent degrees of freedom of $A_{\mu\nu}$ to two. These conditions are typically chosen to be

$$A_{0\mu} = 0 \tag{1.55}$$

$$A^\mu{}_\mu = 0, \tag{1.56}$$

commonly known as the *transverse-traceless* gauge. Note that Eq. 1.55 only represents three independent constraints, since there is a redundancy with one of the constraints from the harmonic gauge condition given in Eq. 1.52. We also note that in the transverse-traceless gauge, $\bar{h}_{\mu\nu} = h_{\mu\nu}$.

If we choose a coordinate system which has a GW propagating in the z -direction, our wave amplitude tensor will look like

$$A_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx}^{TT} & A_{xy}^{TT} & 0 \\ 0 & A_{xy}^{TT} & A_{yy}^{TT} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{1.57}$$

where the TT superscripts indicate that we are working in the transverse-traceless gauge. The orthogonality conditions from Eq. 1.52 require that the z components of $A_{\mu\nu}$ are zero, while one of our new gauge conditions (Eq. 1.55) sets the t components to be zero.

This leaves us with the three independent components shown in Eq. 1.57, however, the traceless condition (Eq. 1.56) requires that A_{yy} is the opposite of A_{xx} . This results in a wave amplitude tensor with two independent components, which represent the two independent polarizations of gravitational waves, commonly referred to as “plus” and “cross” polarizations. Thus, the components of $A_{\mu\nu}^{TT}$ may be renamed as

$$h_+ = A_{xx}^{TT} \tag{1.58}$$

$$h_\times = A_{xy}^{TT}. \tag{1.59}$$

Any generic wave amplitude tensor $A_{\mu\nu}^{TT}$ can be represented as a sum of the plus and

cross polarizations:

$$A_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} = h_+ \mathbf{e}_{\mu\nu}^+ + h_\times \mathbf{e}_{\mu\nu}^\times. \quad (1.60)$$

Here $\mathbf{e}_{\mu\nu}^+$ and $\mathbf{e}_{\mu\nu}^\times$ are the unit polarization tensors, given by

$$\mathbf{e}_{\mu\nu}^+ = \mathbf{e}_x \otimes \mathbf{e}_x - \mathbf{e}_y \otimes \mathbf{e}_y \quad (1.61)$$

$$\mathbf{e}_{\mu\nu}^\times = \mathbf{e}_x \otimes \mathbf{e}_y + \mathbf{e}_y \otimes \mathbf{e}_x \quad (1.62)$$

in this particular case. Note that these expressions are not correct for an arbitrary coordinate system; however, we always have the freedom to choose the z -axis of our coordinate system to lie along the direction of GW propagation.

Qualitatively, a gravitational wave that is propagating perpendicularly to the plane of a ring of test particles will cause the ring to deform into an ellipse first along one axis and then the other, oscillating between these two configurations as a function of time. This unique behavior is illustrated in Fig. 1.1; this signature of gravitational waves lends itself to a interferometer-based strategy for GW detection (discussed further in Sec. 1.3.4).

Finally, it is important to note that linearized gravity, however useful, is only an approximate tool; while it may be quite accurate in some cases, it provides at best a rough qualitative understanding in the strong-field and high-velocity regimes of gravity. In these cases, the quantitative behavior of the system can be best determined by complex numerical calculations using post-Newtonian approaches based on perturbation theory or numerical relativity.

1.3.2 Generation of gravitational waves

When considering the generation of gravitational waves from the motion of massive objects, it is typical to use an analogy with electromagnetism, taking an object's mass to be its "gravitational charge." Beginning with lowest order, it seems clear that there




Figure 1.1: Effect of a passing gravitational wave (propagating into the page) with plus polarization (top) and cross polarization (bottom) on a ring of freely-falling particles. The effect is exaggerated for illustrative purposes.

can be no monopole gravitational radiation. In electromagnetism, monopole radiation would require a charge to be created or destroyed; in gravitation, it would require mass or energy to be created or destroyed. Obviously, neither of these are allowed.

The predominant form of radiation in electromagnetism is from the electric dipole, where the dipole moment \mathbf{p} for a collection of charges is given by

$$\mathbf{p} = \sum_i q_i \mathbf{x}_i, \quad (1.63)$$

where the index i runs over individual particles with charge q_i and position \mathbf{x}_i . The power radiated by an electric dipole is proportional to the second time derivative of the electric dipole moment $\ddot{\mathbf{p}}$.

For gravitational radiation, the mass dipole moment is

$$\mathbf{d} = \sum_i m_i \mathbf{x}_i. \quad (1.64)$$

If we consider the first time derivative of this expression ($\dot{\mathbf{d}}$), we see that it is proportional to the sum of the momenta of the particles. Due to conservation of momentum, $\dot{\mathbf{d}} = 0$, and thus, its second time derivative must also be zero. We could also consider an

analogue to magnetic dipole radiation, however, we would find that conservation of angular momentum precludes the production of gravitational radiation in this case. Thus, there cannot be any gravitational radiation from a mass dipole.

However, there are no rules preventing quadrupole gravitational radiation. Similar to dipole radiation, quadrupole radiation should be proportional to accelerations of the quadrupole moment $I_{\mu\nu}$. We also expect that, like electromagnetism, the observed wave amplitude should be inversely proportional to the distance from the source (R). Filling in factors of G and c , we obtain

$$h_{\mu\nu} = \frac{2G}{Rc^4} \ddot{I}_{\mu\nu}. \quad (1.65)$$

Here, $I_{\mu\nu}$ is the “reduced” or “trace-free” quadrupole moment, given by

$$I_{\mu\nu} = \int dV \left(x_\mu x_\nu - \frac{1}{3} \delta_{\mu\nu} r^2 \right) \rho, \quad (1.66)$$

where ρ is the mass density. Eq. 1.65 can also be rigorously derived from Eq. 1.49 by using the corresponding Green function for the d’Alembertian (see Refs. [2, 4]).

Following Ref. [11], we consider a simple case of a binary system comprising two point masses, each with mass M . The masses are separated by a distance $2r_0$ and are orbiting each other in the x - y plane with frequency f . In this case, the relevant components of the reduced quadrupole moment tensor are

$$I_{xx} = 2Mr_0^2 \left[\cos^2(2\pi ft) - \frac{1}{3} \right] \quad (1.67)$$

$$I_{yy} = 2Mr_0^2 \left[\sin^2(2\pi ft) - \frac{1}{3} \right] \quad (1.68)$$

$$I_{xy} = I_{yx} = 2Mr_0^2 \sin(2\pi ft) \cos(2\pi ft). \quad (1.69)$$

The zz component is a constant (arising only from the $\delta_{\mu\nu}$ term), thus disappearing when we take time derivatives. Cross-terms involving z and x or y are zero.

Using Eq. 1.65, we calculate the two polarization components:

$$h_+ = h_{xx} = -h_{yy} = \frac{32\pi^2 G}{Rc^4} Mr_0^2 f^2 \cos(4\pi ft) \quad (1.70)$$

$$h_\times = h_{xy} = h_{yx} = -\frac{32\pi^2 G}{Rc^4} Mr_0^2 f^2 \sin(4\pi ft) \quad (1.71)$$

It is interesting to note that the gravitational wave frequency is actually equal to twice the orbital frequency f of the binary system.

1.3.3 Sources of gravitational waves

While the motion of everyday objects here on Earth should generate gravitational waves, the masses involved are too small for these GWs to be detected (at present). This becomes clear with a few examples.

Using Eq. 1.70, we can see that the maximum amplitude of a GW is

$$\begin{aligned} |h| &= \frac{32\pi^2 G}{Rc^4} Mr_0^2 f^2 \\ &= 2.60 \times 10^{-42} \left(\frac{M}{1 \text{ kg}} \right) \left(\frac{1 \text{ m}}{R} \right) \left(\frac{r_0}{1 \text{ m}} \right)^2 \left(\frac{f}{1 \text{ Hz}} \right)^2. \end{aligned} \quad (1.72)$$

For an astrophysical system, the orbital frequency can be derived from the masses and their separation by using Kepler's third law:

$$f^2 = \frac{GM}{16\pi^2 r_0^3}. \quad (1.73)$$

Thus, for an astrophysical system, the GW amplitude is given by

$$|h| = 1.1 \times 10^{-54} \left(\frac{M}{1 \text{ kg}} \right)^2 \left(\frac{1 \text{ m}}{R} \right) \left(\frac{1 \text{ m}}{r_0} \right) \quad (1.74)$$

For example, consider a binary system composed of two neutron stars, each with mass $1.4 M_\odot$, separated by 10 km. If this system is at a distance of 10 Mpc from Earth, GWs emitted by this binary pair would reach Earth with an amplitude of $|h| \approx 2.7 \times 10^{-21}$.

Now, imagine setting up a gravitational-wave generator in a terrestrial laboratory, made of two spheres with masses of 1000 kg each, separated by a distance of 10 m, and

rotating at a frequency of 100 Hz. This already seems like a difficult experiment to construct, but perhaps not beyond the limit of human technology. It would be nice if our detector could be in the same laboratory as the source, but to be in the wave zone, our detector must be at least one wavelength away (in this case, 3000 km). Plugging our numbers into Eq. 1.72, we get a GW amplitude of $|h| \approx 9 \times 10^{-40}$.

From this comparison, it is apparent that it will be much easier to detect astrophysical GWs than those generated here on Earth. In the rest of this section, I will discuss some astrophysical sources of gravitational waves. Some of these sources will be discussed in greater detail in later chapters of this dissertation, in relation to analyses focused on their detection.

Compact binary coalescences

Perhaps the most promising class of astrophysical GW sources is the coalescence of compact binary systems. These systems may be composed of two neutron stars (NS-NS), two black holes (BH-BH), or one neutron star and one black hole (NS-BH), and are qualitatively similar to the point mass toy model used previously to estimate GW amplitudes. One important fact about these systems was omitted in the toy model: because gravitational waves carry energy (and momentum), the binary system must lose orbital energy as it emits gravitational waves. This leads to shrinkage of the binary system's orbit, resulting in a higher orbital frequency. After enough time, the objects approach each other, inspiral, and merge to form a single black hole. Following the merger, there is expected to be a ringdown phase, in which the resulting black hole emits gravitational waves due to decaying quasinormal modes [12–14].

These binary systems constantly produce GWs during their orbit, but the frequency is too low to be detected by current interferometric GW detectors until the final inspiral, merger, and ringdown. As these phases only last tenths of a second or less, these signals are typically considered to be relatively short. A simulated GW waveform from a black hole binary coalescence is shown in Fig. 1.2.

There are also some proposed scenarios where two black holes dynamically “capture” each other to form a binary system with a highly eccentric orbit [16,17]. The eventual fate of these binary systems is the same as in other cases where the binary system is formed before the stars collapse to compact objects since the orbit is expected to circularize




Figure 1.2: A simulated GW waveform from the inspiral, merger, and ringdown of a binary black hole system. The black holes each have mass $30 M_{\odot}$ and no spin. The full waveform is shown (top), along with a zoomed-in view of the final milliseconds, highlighting the final plunge, merger, and ringdown. The simulation is based on an effective-one-body model in numerical relativity [15].

before the final inspiral; however, in the dynamical capture scenario, the spins of the two objects may differ significantly, leading to modulation of the GW signal.

Isolated neutron stars

Isolated neutron stars are also expected to be proficient emitters of gravitational waves. A rotating, spherically symmetric star would not produce gravitational waves since its mass quadrupole moment would not vary with time; however, neutron stars may become deformed due to a number of phenomena, including strong magnetic fields, seismic activity, or accretion [18, 19]. There are also several classes of instabilities which may occur in neutron stars and lead to normal mode oscillations, including r -modes [20, 21] and f -modes [22]. These deformations and oscillations may create a quadrupolar asymmetry, which then leads to GW emission as the neutron star rotates.

In general, the GW signals produced by isolated neutron stars are thought to be relatively long. Depending on the cause of the deformation or the type of instability, GW signals from isolated neutron stars may last for hours or days, or even be essentially continuous, lasting for weeks or longer [19, 23]. The frequency of these signals is expected to contain a characteristic modulation due to the Doppler shift from the motion of the Earth.

Stellar collapse

Although the violent dynamics of core-collapse supernovae are still a subject of intense research, it is generally expected that deviations from spherical symmetry occur during these explosions, and should lead to production of gravitational waves. There are a number of phenomena which may lead to emission of GWs during the core-collapse and ensuing bounce, including asymmetric neutrino emission, neutrino-driven convection, and standing accretion shock instability [18, 24]. If a protoneutron star remnant is formed after the core bounce, it may emit GWs due to convection, rotational instabilities, magnetic stresses, or other mechanisms associated with isolated neutron stars. Many of the GW emission scenarios occurring during the core-collapse and bounce stages are expected to produce GW signals lasting a few seconds or less [18], while those associated with the resulting isolated neutron stars are expected to produce longer signals, as previously discussed.

Stochastic backgrounds

A stochastic background of gravitational waves is a persistent, universal background which may be produced by an incoherent superposition of GWs from weak, individually undetectable astrophysical sources, and/or by cosmological mechanisms. Such a background can be thought of as the gravitational analogue of the electromagnetic CMB. An astrophysical background may be produced by a superposition of individual emitters of GWs, including compact binary coalescences, magnetars, core-collapse supernovae, and others [25]. A cosmological background may be produced by a number of phenomena, including a pre-Big Bang cosmology [26–28], cosmic strings [29,30], or the amplification of vacuum fluctuations during inflation [31,32]. Stochastic gravitational-wave backgrounds will be covered in more detail in Chapter 3.

1.3.4 Detection of gravitational waves

The first strong experimental evidence for the existence of gravitational waves was obtained from observations of the Hulse-Taylor binary system (PSR B1913+16) [33–35]. By observing the pattern of radio pulses from the binary system, Taylor and Weisberg were able to determine many of the physical characteristics of the system, including the component masses, orbit inclination, and orbital shape and period. Measurements of the periastron passage over a period of six years indicated that the orbital period of the binary was decreasing [34]. These measurements were found to be in close agreement with the predicted energy loss due to gravitational-wave emission from general relativity and provided the first strong experimental evidence for the existence of gravitational waves. Fig. 1.3 shows observations of PSR B1913+16’s orbital period over a period of 30 years [35].

The first direct detection of gravitational waves was made recently, on September 14, 2015, by the Advanced LIGO detectors. More information on this GW detection is given in Sec. 1.3.5. As we accumulate more direct GW observations, we may learn much more about the dynamics of gravity in the strong-field and high-velocity regimes, as well as the history of the universe. The rest of this subsection will discuss different experimental methods of detecting gravitational waves.




Figure 1.3: Measurements of the orbital period of the Hulse-Taylor binary system over a period of thirty years. The observational data show remarkable agreement with the general relativistic prediction (solid line) which includes energy loss due to gravitational radiation. Figure reproduced from Ref. [35], with the permission of J. Weisberg; copyright 2010 AAS, reproduced with permission.

Resonant mass detectors

Joseph Weber built and operated the first resonant mass (or “bar”) detector for gravitational waves in the early 1960s. A basic resonant mass detector is a large metal cylinder with an accelerometer attached to one end; a passing gravitational wave should cause a tidal force across the bar, which the accelerometer will pick up. However, reasonable sensitivity to GWs is only achieved within a few Hz of resonant mode frequencies, which are determined by the inherent properties of the bar.

In the late 1960s, Weber reported several observations of gravitational waves with his resonant mass detector [36–38]. This led to the construction of several new resonant mass detectors; however, none of the new experiments were able to replicate Weber’s finding, nor were independent reanalyses of his data.

Further generations of resonant mass detectors were constructed in the 1980s and

1990s. These include the ALLEGRO [39], NIOBE [40], NAUTILUS [41], and AURIGA [42] detectors. Many of these detectors were cryogenically cooled to a few degrees K in order to reduce thermal noise. There is also a spherical detector in Brazil, which is cooled to temperatures as low as 20 mK [43]. The best strain sensitivity achieved by resonant bar detectors was around 10^{-21} , which is comparable to that of a GW interferometer; however, the narrow frequency band over which this sensitivity applies is a significant detriment.

Pulsar timing arrays

Pulsar timing arrays focus on using observations from many millisecond pulsars over long periods of time to detect gravitational waves. Gravitational waves passing between a pulsar and the Earth should cause changes in the arrival times of the pulses at Earth (called timing residuals). These experiments typically target gravitational waves with very low frequency, e.g., 10^{-8} – 10^{-9} Hz, including stochastic gravitational-wave backgrounds and mergers of supermassive black holes in galactic centers [19]. At these frequencies, the timing residuals induced by GWs should be on the order of tens of nanoseconds; however, millisecond pulsars are typically very stable and can reach this level of reliability. Currently operational pulsar timing arrays include the International Pulsar Timing Array [44], the Parkes Pulsar Timing Array [45], the European Pulsar Timing Array [46], and the North American Nanohertz Observatory for Gravitational Waves [47]. Some of these recent experiments have attained strain sensitivities as low as 10^{-15} in the nHz band [19].

Interferometry

Interferometric gravitational-wave detectors were conceived of in the 1970s and first built in the 1990s. In an interferometer, a laser sends light into a beamsplitter, which splits the light into two perpendicular paths, or “arms.” At the end of each arm, the light strikes a mirror and reflects, traveling back toward the beamsplitter, where the light from the two arms will recombine. If the arms are of equal length, constructive interference will occur; however, if a gravitational wave is present, its quadrupolar nature should cause one interferometer arm to effectively stretch and the other to shrink, causing some

amount of destructive interference. Thus, the light intensity at the beamsplitter output provides information about the amplitude of an incident gravitational wave.

While this idea of the arms stretching or shrinking is useful for understanding how an interferometer works, it is not the whole story. In reality, a GW will also cause the light waves in the interferometer to stretch or shrink (similar to the cosmological redshift due to expansion of the metric). However, this leads to an accumulated phase difference between the light in the two arms of the interferometer, which in the end still causes interference at the output port. The more time the light spends in the interferometer arms, the greater the phase it will accumulate; toward this end, interferometric GW detectors feature Fabry-Pérot cavities for effectively increasing the interferometer arm length by causing the laser light to take multiple trips down the arms. Some interferometers also feature power recycling mirrors (to re-use light which exits the beamsplitter toward the laser), as well as signal recycling mirrors at the output port to increase the storage time of the interferometer, and thus, the accumulated phase.

The first-generation network of GW interferometers included the Laser Interferometer Gravitational-wave Observatory (LIGO) [1,48], Virgo [49], GEO 600 [50], and TAMA 300 [51]. The initial version of LIGO operated three interferometers in total: two with 4-km arms and one with 2-km arms. These interferometers were mostly sensitive in a frequency band from 40–2000 Hz. Of the first-generation interferometers, LIGO achieved the best sensitivity during its S6 science run, with a strain sensitivity of $\approx 2 \times 10^{-22}$ in the 100–200 Hz band [52].

The second generation of GW interferometers is just beginning observations, with Advanced LIGO [52] collecting its first observational data from September 2015 to January 2016. Compared to initial LIGO, it features many significant upgrades, including improved seismic isolation and increased laser power. When it reaches design sensitivity, it is expected to be up to 10 times more sensitive than initial LIGO. Advanced LIGO is also capable of probing lower frequencies, with considerable sensitivity down to 10 Hz (compared to 40 Hz for initial LIGO). Advanced Virgo [53] is still in commissioning, but is expected to come online in 2017. The underground detector KAGRA [54] (formerly called LCGT) is being built in Japan by the collaboration which operated the TAMA 300 detector and features cryogenic cooling to reduce thermal noise. Along with these detectors, plans are in place to construct a LIGO detector in India in the next ≈ 6 years,

which will significantly improve source localization capabilities [55].

Future generations of GW interferometers are mostly in the proposal and design stages at present. The Einstein Telescope [56,57] is an underground, triple-interferometer proposal being coordinated by the European Gravitational Observatory. The Evolved Laser Interferometer Space Antenna (eLISA) project has recently launched a pathfinder mission to demonstrate their technology for a space-based interferometer [58,59]. Other proposed space-based interferometers include the Big Bang Observer (BBO) [60,61], the Advanced Laser Interferometer Antenna (ALIA) [61], and the DECi-hertz Interferometer Gravitational wave Observatory (DECIGO) [62]. The space-based detectors will be able to avoid the problem of seismic noise (the primary low-frequency noise source for GW interferometers) altogether and are mostly focused on detecting extremely low frequency gravitational waves and gravitational-wave backgrounds.

Chapter 2 of this dissertation will focus on GW interferometry with LIGO in more detail.

1.3.5 GW150914

On September 14, 2015 at 09:50:45 UTC, the Advanced LIGO detectors made the first direct detection of a gravitational-wave signal, called GW150914 [10]. The signal was produced by the inspiral, merger, and ringdown of a black hole binary system with masses of $36_{-4}^{+5} M_{\odot}$ and $29_{-4}^{+4} M_{\odot}$, and a final black hole mass of $62_{-4}^{+4} M_{\odot}$ (all masses in source frame). A total energy of $3.0_{-0.5}^{+0.5} M_{\odot} c^2$ was emitted in gravitational waves, making this one of the most energetic astrophysical events ever observed. The source was measured to be at a luminosity distance of 410_{-180}^{+160} Mpc, corresponding to a redshift of $z = 0.09_{-0.04}^{+0.03}$. The signal has a characteristic “up-chirp” waveform, increasing in frequency from 35–250 Hz and reaching a maximum strain amplitude of 1.0×10^{-21} . Plots of the data from both LIGO detectors and the reconstructed waveforms are shown in Fig. 1.4.

GW150914 was initially detected by a pipeline designed to search for generic transients. Using 16 days of coincident data collected between September 12 to October 20, its significance was estimated at 1 in 22,500 years, equivalent to 4.6σ . Subsequent follow-up by a matched-filtering pipeline using binary black hole templates estimated its significance at 1 in 203,000 years, or $> 5.1\sigma$. A second, independent matched-filtering

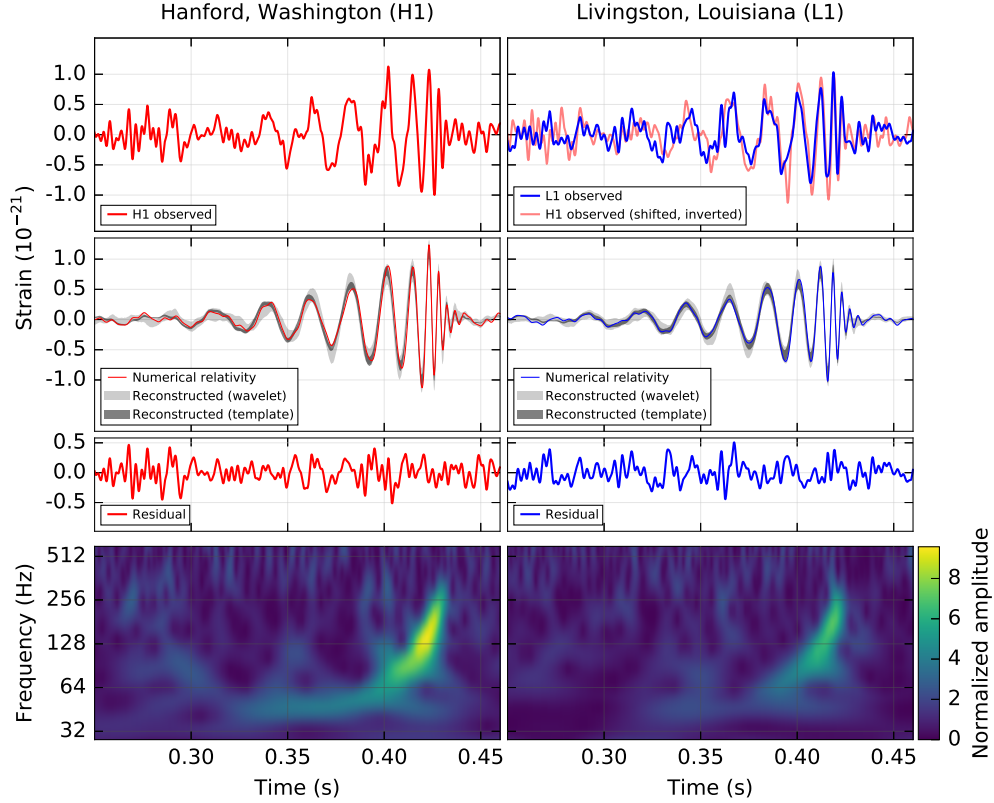


Figure 1.4: GW150914, as detected by the H1 (left column) and L1 (right column) Advanced LIGO detectors. Times shown are relative to 09:50:45 UTC on September 14, 2015. Top row: detector strain time-series. Middle row: reconstructed waveforms, using numerical relativity, black hole binary template matching, and a sine-Gaussian wavelet reconstruction. Bottom row: time-frequency plots of the strain data, clearly illustrating the “chirp-up” signature of the waveform. Figure reproduced from Ref. [10] under the Creative Commons Attribution 3.0 License.

pipeline also analyzed the data and detected GW150914 with similar significance.

Several companion papers have been published focusing on specific results which can be inferred from the detection of GW150914. More details on the generic transients search and the matched filtering searches can be found in Refs. [63] and [64], respectively. Information on the estimation of parameters and analyses of potential

biases is in Ref. [65]. The overall BBH merger rate was estimated to be between 2–400 $\text{Mpc}^{-3}\text{yr}^{-1}$ [66]. Quantitative tests of general relativity using GW150914 were performed in Ref. [67]; no evidence for violations of GR were found. The astrophysical implications from the detection of GW150914 and its estimated parameters were explored in Ref. [68], including binary black hole formation mechanisms. Estimates of the energy density in a stochastic gravitational-wave background from the inspiral and merger of binary black hole systems were calculated in Ref. [69]. A follow-up search for high-energy neutrinos around the time of GW150914 was performed with IceCube and ANTARES, but no deviations from the expected background were found [70].

The detection of GW150914 definitively proves the existence of gravitational waves and that black hole binaries can form and merge within a Hubble time [10].

1.4 Cosmology discussion

The idea of a “Big Bang” was first conceived in the 1920s and built on the theory of general relativity. The basic idea is simple: an expanding universe should be smaller, denser, and hotter in the past than it is today. If we carry this idea to its extreme and look far enough back in time, the universe may have been as small as a single point, otherwise known as a singularity. Due to the extremely large energies and small scales associated with this time, the laws of physics as we know them would break down; thus, it is difficult to make any speculations about the universe before the singularity.

In 1929, Hubble’s observations led to the conclusion that we live in an expanding universe. Based on further observations and the eventual detection of the CMB in 1965, the Big Bang theory has become more and more entrenched in modern cosmology. This simple idea, coupled with basic thermodynamics and nuclear physics, has made a number of predictions of cosmological phenomena which match observations to a remarkable level of accuracy.

In this section, I will provide an overview of the standard model of cosmology, including a cosmological timeline and a discussion of how gravitational waves fit into this model.

1.4.1 The standard cosmological model

The Λ CDM model is generally based on a hot Big Bang cosmology, and includes a cosmological constant (indicated by Λ) and cold dark matter (CDM). In this model, the expansion of the universe is parametrized by a scale factor $a(t)$, and the universe is assumed to be isotropic and homogeneous when viewed on a large enough scale (this assumption is also referred to as the cosmological principle).

Friedmann-Lemaître-Robertson-Walker metric

Following Ref. [71], the spacetime interval for an isotropic, homogeneous universe can be written as

$$ds^2 = dt^2 - a^2(t) \left(\frac{dr^2}{1 - kr^2} + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \right). \quad (1.75)$$

This is known as the Friedmann-Lemaître-Robertson-Walker (FLRW) metric. Here, t , r , θ , and ϕ are comoving coordinates, and k represents the curvature (+1 for positive curvature, 0 for no curvature (flat), or -1 for negative curvature). We have taken $c = 1$ in this equation and throughout much of this section.

From here, we consider the Einstein field equations, including a cosmological constant (see Eq. 1.37). The first step is to calculate the non-zero components of the Ricci tensor and the Ricci scalar from the metric:

$$R_{00} = -3 \frac{\ddot{a}}{a} \quad (1.76)$$

$$R_{ij} = - \left(\frac{\ddot{a}}{a} + 2 \frac{\dot{a}^2}{a^2} + 2 \frac{k}{a^2} \right) g_{ij} \quad (1.77)$$

$$R = -6 \left(\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2} \right). \quad (1.78)$$

Stress-energy tensor

Now, we must consider the stress-energy tensor $T_{\mu\nu}$. Because the FLRW metric is symmetric and diagonal, the Einstein field equations require $T_{\mu\nu}$ to be symmetric and diagonal, as well. The assumption of isotropy also requires the spatial components of $T_{\mu\nu}$ to be equal. The simplest model which satisfies these requirements is that of a

perfect fluid, defined by its rest energy density ρ and a pressure p . The corresponding stress-energy tensor is

$$T^\mu{}_\nu = \text{diag}(\rho, -p, -p, -p). \quad (1.79)$$

For simplicity, we assume the equation of state of the fluid to be

$$p = w\rho, \quad (1.80)$$

where w is a time-independent constant for a given fluid. With this equation of state, conservation of energy gives a relation between the energy density and the scale parameter:

$$\rho \propto a^{-3(1+w)}. \quad (1.81)$$

Keep in mind that both ρ and a are functions of time. There are several different options for w :

- $w = 0$: non-relativistic matter or “dust.” This makes sense because matter has negligible pressure compared to its rest mass. This type of fluid will have $\rho \propto a^{-3}$; thus, as the universe expands, the energy density of matter will change proportionally to the volume.
- $w = 1/3$: radiation. This has an energy density of $\rho \propto a^{-4}$. The extra power of a^{-1} comes from the cosmological redshifting of photons as the universe expands.
- $w = -1$: vacuum energy, or the cosmological constant, is often assumed to have a negative pressure which drives the expansion of the universe. This has $\rho = \text{constant}$.
- $w = -1/2$: curvature. There is not an easily intuitive explanation for this case, but it gives $\rho \propto a^{-2}$.

Solving the field equations

Now that we have a metric and have derived the important curvature objects from it, and we have a description of the stress-energy tensor in an isotropic, homogeneous universe, we can work with the Einstein field equations. The 00 component of the field equations

gives

$$\frac{\dot{a}^2}{a^2} + \frac{k}{a^2} = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3}, \quad (1.82)$$

while the trace of the spatial terms gives

$$2\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2} = -8\pi Gp + \Lambda. \quad (1.83)$$

Taking the difference of these two equations yields information about the acceleration of the universe's expansion:

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}. \quad (1.84)$$

It is typical to define the Hubble parameter, which encapsulates the expansion rate of the universe, as

$$H \equiv \frac{\dot{a}}{a}. \quad (1.85)$$

The Hubble *constant* is defined as the present-day value of the Hubble parameter, and is denoted as H_0 . The scale factor today is usually taken to be $a_0 = 1$.

Ignoring the cosmological constant for a moment, we may rearrange Eq. 1.82 as

$$k = \frac{8\pi G}{3}\rho - H_0^2. \quad (1.86)$$

We define the critical density ρ_c as

$$\rho_c = \frac{3H_0^2}{8\pi G} \quad (1.87)$$

so that if ρ is greater than, equal to, or less than ρ_c , the curvature will be positive, non-existent (flat geometry), or negative, respectively. ρ_c is also referred to as the energy density required to “close” the universe.

Density parameters

If we divide Eq. 1.82 by H_0^2 , a unitless equation is obtained:

$$\frac{H^2}{H_0^2} = \frac{\rho}{\rho_c} - \frac{k}{H_0^2 a^2} + \frac{\Lambda}{3H_0^2}. \quad (1.88)$$

If we consider the present-day form of this equation, we can take $H = H_0$ and $a = a_0 = 1$. We can also separate the energy density ρ into matter and radiation ($\rho = \rho_m + \rho_r$).

$$1 = \frac{\rho_m}{\rho_c} + \frac{\rho_r}{\rho_c} - \frac{k}{H_0^2} + \frac{\Lambda}{3H_0^2} \quad (1.89)$$

It is most common to define the unitless density parameters

$$\Omega_m = \frac{\rho_m}{\rho_c} \quad (1.90)$$

$$\Omega_r = \frac{\rho_r}{\rho_c} \quad (1.91)$$

$$\Omega_k = -\frac{k}{H_0^2} \quad (1.92)$$

$$\Omega_\Lambda = \frac{\Lambda}{3H_0^2} \quad (1.93)$$

so that Eq. 1.89 becomes

$$1 = \Omega_m + \Omega_r + \Omega_k + \Omega_\Lambda. \quad (1.94)$$

This requires $\Omega = \Omega_m + \Omega_r + \Omega_\Lambda = 1$ for the universe to be flat ($k = \Omega_k = 0$).

Discussion

From this model, it is possible to predict the rate of expansion of the universe for different types of “fluid.” If we assume that one type of “fluid” (matter, radiation, vacuum energy, curvature) dominates at a particular time in the history of the universe, it is simple to solve Eq. 1.82 and determine the time dependence of the scale factor [72].

In general, the universe is thought to have been initially radiation-dominated after the Big Bang [73]. The energy density of radiation should fall off faster than that of matter as the universe expands (since ρ goes as a^{-4} for radiation and a^{-3} for matter), so at some point, the energy density of radiation should drop below that of matter. This transition occurred a few tens of thousands of years after the Big Bang, and matter remained the dominant form of energy in the universe until approximately 4 billion years ago. Since then, vacuum energy, or “dark energy” has been the dominant form of energy.

Fluid	w	$\rho(\mathbf{a})$	$\mathbf{a}(\mathbf{t})$
Matter	0	a^{-3}	$t^{2/3}$
Radiation	$\frac{1}{3}$	a^{-4}	$t^{1/2}$
Vacuum energy	-1	constant	e^{Ht}
Curvature	$-\frac{1}{2}$	a^{-2}	t

Table 1.1: Summary of different “fluids” used in the standard cosmological model. The parameter w in the second column indicates the relation between pressure p and energy density ρ in the fluid’s equation of state (see Eq. 1.80). The third column shows the dependence of the energy density on the scale factor a , and the fourth column shows the time dependence of the scale factor, assuming that each particular fluid is dominant.

1.4.2 Λ CDM model parameters

The Λ CDM model is very simple in that it only requires six parameters to be measured; all other parameters of interest may be derived from the measurements. The six parameters which are measured depend on the experiment, but are typically the following:

- Physical baryon density $\Omega_b h^2$: this is the unitless baryon density parameter multiplied by the square of the reduced Hubble constant $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$.
- Physical dark matter density $\Omega_c h^2$: analogous to the physical baryon density, but for cold dark matter.
- Vacuum energy density Ω_Λ .
- Scalar spectral index (n_s): inflation predicts a power-law form for primordial density fluctuations with this parameter as the spectral index.
- Reionization optical depth (τ): the optical depth at which reionization of hydrogen occurred (more discussion of this in Sec. 1.4.3).
- Curvature fluctuation amplitude (Δ_R^2): amplitude of primordial curvature fluctuations.

In Table 1.2, we list the current best values of these parameters from the 2015 Planck results [74], along with some derived parameters of interest.

Parameter	Symbol	Value
Physical baryon density	$\Omega_b h^2$	0.02230 ± 0.00014
Physical dark matter density	$\Omega_c h^2$	0.1188 ± 0.0010
Vacuum energy density	Ω_Λ	0.6911 ± 0.0062
Matter density	Ω_m	0.3089 ± 0.0062
Age of the universe	t_0	$(13.799 \pm 0.021) \times 10^9 \text{ yr}$
Reionization optical depth	τ	0.066 ± 0.012
Reionization redshift	z_{re}	$8.8^{+1.2}_{-1.1}$
Scalar spectral index	n_s	0.9667 ± 0.0040
Curvature fluctuation amplitude	Δ_R^2	$(2.14 \pm 0.05) \times 10^{-9}$

Table 1.2: Current best values for measured and derived parameters of the Λ CDM cosmological model. All values taken from Ref. [74].

1.4.3 Cosmological timeline

This subsection will highlight some of the important events in the cosmological timeline. Many of these occur due to “freeze-out” of certain nuclear or chemical processes as the universe expands, cools, and decreases in matter and radiation density. At certain temperatures, some important processes fall out of thermodynamic equilibrium, leading to significant changes in the overall composition of the universe.

Here, I will discuss some of these transitions, along with other important events in the history of the universe.

Grand unification

At the extremely high energy scales of the early universe, it is thought that all four fundamental forces were of equal strength and unified into a single force. Before a Planck time (5.4×10^{-44} s), it is expected that quantum gravitational interactions dominated; as a result, general relativity and cosmology as we know them break down in these early time periods and so far, most theories which attempt to describe this time are only speculative. After a Planck time, the gravitational force is thought to have separated, while the remaining three forces were still unified in the electroweak force.

At about 10^{-36} seconds, the strong force also separated, leaving only the electromagnetic and weak forces unified [71]. The energy scale of this transition is about 10^{14} GeV. Finally, the electromagnetic and weak forces separated at about 10^{-12} seconds and 1 TeV, due to spontaneous symmetry breaking.

Inflation

Non-inflationary cosmology (as described in Sec. 1.4.1) is very successful, but there are a few important questions that it does not answer:

- Horizons: why does the CMB appear to be so isotropic (differences on the order of 10^{-5})? In a matter or radiation-dominated universe without inflation, certain parts of the universe would have never been in causal contact with each other, so there is no reason that they should be in thermal equilibrium.
- Flatness: why does the universe appear to be so flat ($k = \Omega_k = 0$ within experimental uncertainties)?
- Relics: a very hot early universe should have produced many heavy magnetic monopoles, so why have we never observed any of these?

The most widely accepted theory which answers these questions is the theory of *inflation*, first proposed by Alan Guth in 1981 [75]. In general terms, inflation was a very rapid expansion of spacetime, during which the scale factor increased by a factor of $\approx 10^{25}$, occurring approximately 10^{-36} seconds after the Big Bang. In most theories, inflation is thought to be driven by the evolution of a scalar field, referred to as the inflaton, in a potential. At the end of inflation, it is thought that oscillations of the inflaton field around the bottom of its potential may have led to the production of matter; this is commonly referred to as *reheating* [71].

Inflation's answers to the three problems mentioned previously are as follows:

- Horizons: in an inflationary universe, the universe we observe today would have been small enough to have been in causal contact and thermodynamic equilibrium before inflation; thus, the high degree of isotropy that we observe today is expected.
- Flatness: because Ω_k is proportional to a^{-2} (see Eq. 1.92), it should be reduced by a factor of $\approx 10^{-50}$ compared to its value before inflation. The universe may

not really be flat, but we only see a small part of it, which may appear locally flat to us.

- Relics: based on expected production of magnetic monopoles during the GUT transition, inflation is thought to have decreased the overall density of these relics to the point where it is unlikely that we would observe any [72].

There is not currently one “best” theory of inflation which is commonly accepted, however, the qualitative concept of inflation is still regarded as the best (and most simple) explanation for many of our observations of the universe. Guth’s original theory of inflation considered a single scalar field which was initially trapped in a false minimum [75] and eventually tunneled through to the true vacuum state. However, it was found that this theory was problematic, as different parts of the universe may inflate at different times, creating “bubbles” which would not be large enough to account for the size of the currently observable universe.

Linde [76] and Albrecht and Steinhardt [77] independently proposed “slow-roll” inflation in 1982, which was based on the inflaton rolling down a potential hill with a very slight slope toward the true vacuum state. There also exist many other theories, including chaotic inflation [78], in which multiple scalar fields exist, with random, position-dependent values, and eternal inflation [79], in which inflation continues on forever in certain parts of the universe.

It is expected that quantum mechanical fluctuations in the inflaton field should be enlarged into classical density perturbations during inflation [80]. These inhomogeneities are thought to form the basic seeds for the eventual formation of large-scale structures in the universe. Along with these scalar perturbations, inflation is expected to generate perturbations in the metric tensor, or gravitational waves. Gravitational waves at a particular comoving scale should be “frozen in” until the Hubble parameter has become comparable to their wavelength; after this point, they undergo redshifting due to the expansion of the universe [80]. This may lead to a GW spectrum which is flat within the Hubble length and goes as λ at larger wavelengths.

Baryogenesis

After the end of inflation, the universe was still hot and dense enough for quarks to roam freely. This state is typically described as a *quark-gluon plasma*. As the universe expanded and cooled, quarks were eventually bound into baryons and mesons.

Theories of baryogenesis in the early universe seek to understand why there are more baryons than antibaryons today. Clearly, we observe matter almost exclusively, and while we cannot rule out the possibility that other galaxies composed entirely of antimatter may exist, it seems unlikely.


Based on thermodynamic calculations and estimates for the amount of the matter in the universe today, it is thought that there were approximately $10^9 + 1$ protons for every 10^9 antiprotons during the early universe [80]. Although it is possible that the universe came into being with a baryon asymmetry, it is also possible that this asymmetry was generated by certain processes in the early universe which had a preference for matter over antimatter. The latter theory is generally preferred by cosmologists, as the notion of overall symmetry is generally important throughout physics. Sakharov [81] defined three conditions which a process must satisfy to produce an excess of baryons versus antibaryons:

1. Baryon number (B) violation.
2. Charge-parity (CP) violation.
3. Non-equilibrium conditions.

Interactions which violate CP symmetry have been observed in neutral kaon decays. Along with this, the expansion of the universe would eventually drive certain processes away from equilibrium, satisfying condition 3. However, no processes have ever been observed which do not conserve baryon number. One possibility lies in the framework of a grand unified theory; in these theories, baryons and leptons may potentially be treated different states of a single entity. Although overall conservation of baryon and lepton number ($B+L$) may be required, it may be possible to violate each one individually [80]. This is currently an open problem in cosmology and particle physics, and an area of active research.

Nucleosynthesis

Nucleosynthesis describes the process of the formation of light elements in the early universe. Around 1 second after the Big Bang, the temperature of the universe cooled to about 10^{10} K, which is approximately the energy scale of nuclear processes which convert between neutrons and protons.



Because the proton is lighter than the neutron ($\Delta m = 1.3$ MeV), Eqs. 1.95 and 1.96 require 1.8 MeV and 0.8 MeV to go from left to right, respectively. When the temperature of the universe is around this energy level, there is a preference for these processes to proceed from right to left.

These processes freeze out when the expansion rate of the universe becomes larger than the reaction rate. Because the temperature distribution of photons follows a black-body spectrum, the temperature of freeze-out is also highly dependent on the ratio of baryons to photons η . For these nuclear reactions, freeze-out occurs at a temperature of 0.8 MeV [71]. The equilibrium ratio of neutrons to protons at some temperature (in the non-relativistic limit) can be estimated as

$$\frac{N_n}{N_p} = e^{\Delta mc^2/kT}, \quad (1.97)$$

which gives a ratio of $\approx 1/6$ at a temperature of 0.8 MeV.

After this freeze-out occurs, it takes about two minutes for the temperature to decrease enough to allow deuterium to form.



It is important to note that, during this time, a non-negligible fraction of the neutrons decay (the neutron half-life is $\tau_n = 881$ s), changing the ratio of neutrons and protons to $\approx 1/7$. Formation of deuterium allows for the creation of heavier elements, including:

- Tritium



- Helium-3



- Helium-4



- Beryllium



- Lithium



Lack of stable isotopes at atomic numbers of 5 and 8 prevents significant formation of heavier elements.

Assuming all neutrons go into forming ${}^4\text{He}$ gives a helium mass fraction of 0.25 relative to hydrogen. Further calculations are needed to estimate other abundances; some of these are given in Table 1.3 and are compared to observations.

The predicted abundances given in Table 1.3 depend both on the baryon-to-photon ratio η and the effective number of neutrino species N_ν ; because these parameters are not narrowly restricted, the predictions given are only rough estimates. For the most part, there is good agreement between the predictions and observations, with the sole exception being ${}^7\text{Li}$. The most promising explanation for the dearth of this isotope involves deep convection zones in stellar atmospheres, which may be capable of trapping and burning ${}^7\text{Li}$ [71, 72].

Element	Predicted abundance	Observed abundance
^2H	$\approx 3 \times 10^{-5}$	$2.78_{-0.38}^{+0.44} \times 10^{-5}$ [72]
^3He	$\approx 1 \times 10^{-5}$	$(1.04 \pm 0.06) \times 10^{-5}$ [82]
^4He	≈ 0.25	0.2477 ± 0.0029 [83]
^7Li	$\approx 3 \times 10^{-10}$	$(2.37 \pm 0.05) \times 10^{-10}$ [84]

Table 1.3: Predicted light element abundances from Big Bang nucleosynthesis, compared to observations. Abundances are given in terms of mass fraction relative to hydrogen.

This model of nucleosynthesis allows the placement of upper limits on gravitational-wave energy density in the early universe. Light element abundances are effectively determined by the freeze-out temperature of Eqs. 1.95 and 1.96, which is determined by the point at which the expansion rate of the universe (represented by the Hubble parameter) exceeds the reaction rate. Since the Hubble parameter is proportional to the overall energy density of the universe, the freeze-out temperature is a function of the overall energy density. Thus, significant energy density in a background of gravitational waves would have an effect on the freeze-out temperature, and as the freeze-out temperature increases, so does the ratio of neutrons to protons (up to a ratio of 1), which should lead to the production of more light elements. As a result, observations of light element abundances make it possible to put a limit on the total GW energy density [85, 86]:

$$\Omega_{\text{BBN}} < 1.1 \times 10^{-5} (N_\nu - 3). \quad (1.107)$$

Analysis of WMAP data has constrained $N_\nu < 4.4$ [87], setting $\Omega_{\text{BBN}} < 1.54 \times 10^{-5}$. Searches for a stochastic gravitational-wave background using LIGO data [88, 89] have surpassed this bound in the LIGO frequency band, with a current best upper limit of $\Omega_{\text{GW}} < 5.6 \times 10^{-6}$ in the 41.5–169.25 Hz band.

Recombination and the CMB

Initially, the universe was too hot for electrons and protons to be bound as neutral hydrogen; instead, it was filled with a plasma of charged particles and hot radiation. Because the ratio of baryons to photons was small ($\eta \approx 10^{-9}$), an appreciable density of neutral hydrogen could not form until the universe cooled well below its binding energy

of 13.6 eV.

It is difficult to assign an exact time to this “recombination” of protons and electrons into neutral hydrogen because it did not occur instantaneously. However, since this transition defines the beginning of a new phase in the universe, we should require that there is a negligible proportion of charged particles compared to neutral hydrogen. This is thought to have occurred approximately 380,000 years after the Big Bang, at a temperature of ≈ 3500 K, or 0.3 eV. This temperature can be derived from assuming a blackbody spectrum for the photon energy distribution, with knowledge of η and the desired ionization fraction of hydrogen. Due to the cosmological redshift, the CMB today has a temperature of 2.7 K, or 2.3×10^{-4} eV.

Soon after recombination, the rate of Compton scattering of photons became comparable to the expansion rate of the universe, leading to freeze-out of this reaction. This allowed photons to decouple from matter and “free-stream” through the universe without interacting. The background of thermal radiation formed by these photons constitutes the Cosmic Microwave Background. The prediction of the CMB and its features is one of the most significant achievements of the standard Big Bang cosmological model.

The CMB was first observed in 1965 by Penzias and Wilson [90]. In recent years, it has been studied extensively by space-based experiments, including the Cosmic Background Explorer [91], the Wilkinson Microwave Anisotropy Probe [92], and the Planck satellite [93, 94]. These experiments have revealed the remarkable homogeneity of the CMB (to the level of 10^{-5} in temperature); however, the slight anisotropies of the CMB provide even more information about the universe.

Primary anisotropies in the CMB occur due to density fluctuations in the primordial plasma. Areas of higher density form gravitational potential wells, which compress the plasma in these regions. This compression is resisted by radiation pressure, and the interaction between these two forces sets up oscillations in the plasma. The thermal radiation is generally hotter in compressed regions and cooler in rarefied regions. These hot and cold spots were “frozen-in” when decoupling occurred and can still be seen in the CMB today.

Typically, the CMB anisotropy is characterized in terms of a multipole expansion in spherical harmonics $Y_{lm}(\theta, \phi)$. There are several peaks in this spectrum, the angular size of which is determined by the wavelength of the acoustic oscillation at the time of

last scattering. The first peak corresponds to a mode that only had time to compress before last scattering, the second peak corresponds to a mode that had time to compress and expand once, and so on for the other peaks. The even peaks tend to be damped compared to the odd peaks; this is because the even peaks correspond to rarefaction at the time of last scattering, and the presence of baryons in gravitational potential wells caused increased compression and decreased rarefaction (this is called “baryon drag”). Measurements of these peaks are in excellent agreement with predictions; see Fig. 50 from Ref. [95] for an example.

Secondary anisotropies in the CMB mostly occur due to effects after decoupling. The Sunyaev-Zel’dovich effect can distort the CMB through inverse Compton scattering off of high-energy electrons. The Sachs-Wolfe effect occurs when photons are gravitationally redshifted at the surface of last scattering (non-integrated effect) or by gravitational potential wells after last scattering (integrated effect). Gravitational lensing of CMB photons can occur due to large scale structure in the universe; this effect tends to blur the acoustic peaks and can also modify the polarization.

The CMB contains polarization anisotropies at the level of 10^{-6} due to quadrupolar anisotropies in the radiation intensity. In theory, this anisotropy can be generated by scalar, vector, or tensor (gravitational-wave) perturbations; however, vector contributions to the polarization are thought to be negligible. A quadrupole anisotropy generated by scalar perturbations is expected to produce a curl-free polarization pattern (called an E-mode), while gravitational waves should produce a divergence-free pattern (called a B-mode). The E-mode polarization was first detected in 2002 by the Degree Angular Scale Interferometer (DASI) [96]; further studies [97–100] have confirmed the existence of this polarization and established its conclusive agreement with the standard cosmological model.

Observations of B-mode polarization in the CMB may provide evidence for a gravitational-wave background in the early universe. The B-mode polarization has been observed due to gravitational lensing by the South Pole Telescope [101] and at degree angular scales by BICEP2 [102], although the BICEP2 result may possibly be contaminated by dust.

First stars and reionization

Somewhere between a few hundred million and a billion years after the Big Bang, the first stars begin to form. It is likely that these were very massive and metal-poor stars, often referred to as Population III stars. This generation of stars produced the first significant amounts of heavier elements, including carbon, oxygen, and iron.

The UV photons produced by these stars were energetic enough to reionize neutral hydrogen. However, due to the expansion of the universe, the interaction rate between these photons and atoms was significantly reduced compared to the early universe. Thus, the universe remains mostly transparent today.

Chapter 2

Laser Interferometer Gravitational-wave Observatory

The Laser Interferometer Gravitational-wave Observatory (LIGO) is a large-scale scientific effort focused on the direct detection of gravitational waves. From 2002 to 2007, the LIGO Scientific Collaboration (LSC) operated three kilometer-scale Michelson interferometers: one 4-km interferometer (H1) and one 2-km interferometer (H2) in Hanford, WA, and another 4-km interferometer (L1) in Livingston, LA; this will be referred to as the “initial” LIGO configuration. After a two-year period of upgrades to the “enhanced” LIGO configuration and the decommissioning of H2, H1 and L1 collected data from 2009 to 2010, achieving their best strain sensitivity ($\approx 2 \times 10^{-22}$ in a 100 Hz frequency band) during this time [52]. Photographs of the two LIGO observatories are shown in Fig. 2.1.

From 2010 to 2015, the LIGO detectors collected no science data, instead undergoing a series of upgrades to become “Advanced LIGO.” The Advanced LIGO detectors recently acquired their first observational data from September 2015 to January 2016. During this observational run, the strain sensitivity was a factor of 2 or 3 better than the enhanced LIGO sensitivity. However, this chapter will focus primarily on the initial LIGO detectors, since the analyses presented in this dissertation utilized data collected in the initial LIGO configuration. Henceforth, the initial LIGO detectors will simply be referred to as the LIGO detectors.

In Sec. 2.1, I provide a description of the initial LIGO detectors and some of their



Figure 2.1: Aerial view of the LIGO observatories at Hanford, WA (top) and Livingston, LA (bottom). The laser and main optics are contained in the large buildings at the vertex of the two beam tubes. Figure reproduced from Ref. [1]. Copyright IOP Publishing. Reproduced with permission. All rights reserved.

most significant noise sources. Sec. 2.2 gives an overview of data analysis techniques and observational results from the LSC. Finally, Sec. 2.3 briefly introduces Advanced LIGO and discusses its first observational run.

2.1 Detector description

The LIGO detectors were modified power-recycled Michelson interferometers, designed to be capable of detecting GWs over a frequency band of 40–7000 Hz, and to be sensitive to GW strain amplitudes as low as 10^{-21} [48]. They collected data during six science runs which occurred between 2002 and 2010. Tab. 2.1 summarizes these science runs and the strain sensitivities achieved during them.

Science Run	Dates	Strain Sensitivity
S1 [103]	23 August 2002 - 9 September 2002	$2 \times 10^{-20} \text{ Hz}^{-1/2}$
S2 [104]	14 February 2003 - 14 April 2003	$3 \times 10^{-21} \text{ Hz}^{-1/2}$
S3 [105]	31 October 2003 - 9 January 2004	$2 \times 10^{-21} \text{ Hz}^{-1/2}$
S4 [106]	22 February 2005 - 23 March 2005	$5 \times 10^{-22} \text{ Hz}^{-1/2}$
S5 [1]	5 November 2005 - 30 September 2007	$3 \times 10^{-22} \text{ Hz}^{-1/2}$
S6 [107]	7 July 2009 - 20 October 2010	$2 \times 10^{-22} \text{ Hz}^{-1/2}$

Table 2.1: Overview of the six LIGO science runs, including dates of operation and strain sensitivity achieved. Strain sensitivities are estimated for the most sensitive 100 Hz frequency band ($\approx 100\text{--}200$ Hz).

This section will give an overview of the LIGO detectors, including the different subsystems and main sources of noise in the detectors.

2.1.1 Detector layout

The design goal of the initial LIGO detectors was to reach a strain sensitivity of 10^{-21} , corresponding to a length change of $\approx 10^{-18}$ m. In order to achieve this level of sensitivity, the LIGO detectors feature several modifications compared to a standard Michelson interferometer. A diagram of the LIGO detectors is shown in Fig. 2.2.

Each arm in a LIGO interferometer features a Fabry-Pérot cavity, which is a type of resonant optical cavity. A Fabry-Pérot cavity is made of specially designed mirrors and

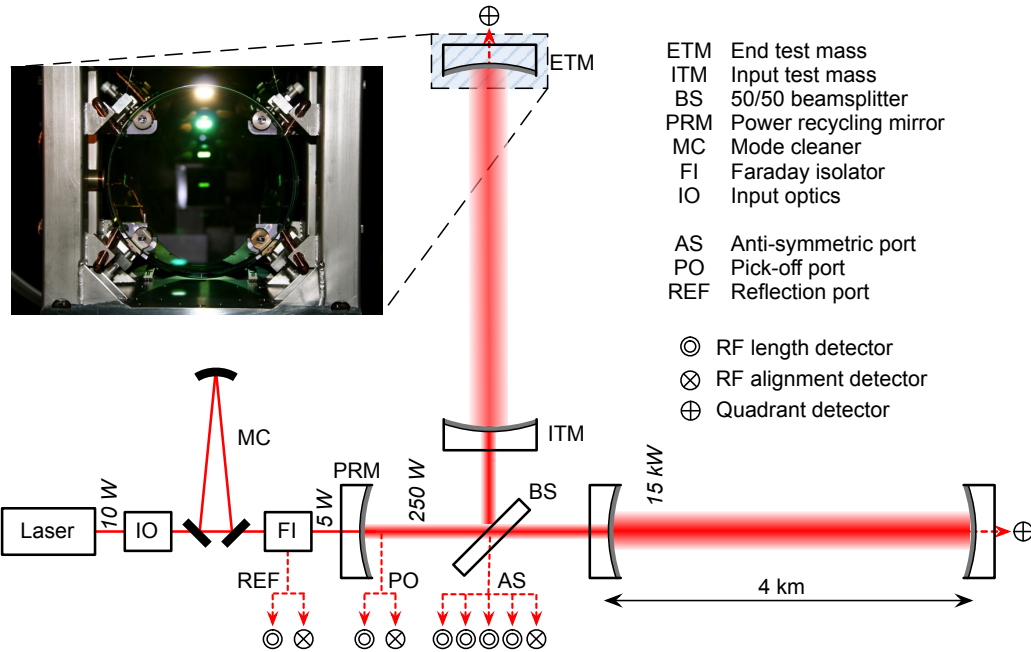


Figure 2.2: Overview of the optical and sensing system of the LIGO interferometers. The Fabry-Pérot cavities are between the input test masses and the end test masses. The power recycling mirror is placed between the beamsplitter and the laser. The laser powers given here are generic and are only meant to give a relative description of the laser power in different parts of the detector. The inset image shows one of the suspended input test masses. Figure and information reproduced from Ref. [1]. Copyright IOP Publishing. Reproduced with permission. All rights reserved.

mirror coatings which allow the laser light to resonate in the cavity. This causes the laser light to circulate multiple times in an arm, effectively increasing the interferometer arm length and the accumulated phase of the light (if a GW is present). For the initial LIGO detectors, a GW signal at 100 Hz would traverse the arms approximately 100 times [1].

The LIGO detectors have a intentional asymmetry of 355 mm in length of their two arms, called a Schnupp asymmetry [19]. This asymmetry is induced in the inner Michelson interferometer formed by the beamsplitter and the input mirrors (see Fig. 2.2); the two Fabry-Pérot cavities are identical in length. This is done so that the operating point of the interferometer is on a dark fringe due to destructive interference; thus, changes in the arm length will increase the power at the output port.

Another modification from a standard Michelson interferometer is the implementation of power recycling [108, 109]. Because the output of the GW channel is designed to be on a dark fringe when there is no GW power in the interferometer, much of the light incident on the beamsplitter returns toward the laser. Power recycling involves the insertion of another partially transmitting mirror between the beamsplitter and the laser so that much of this light is returned to the arms of the interferometer, thus “recycling” some of the laser power. This technique significantly increases the light power stored in the interferometer and reduces shot noise at the output.

The interferometer’s laser emits 10 W of power from an Nd:YAG source [1] at a wavelength of 1064 nm. It is stabilized in power by implementing a feedback loop between a photodetector and the power amplifier and is stabilized in frequency by locking to a reference cavity using the Pound-Drever-Hall (PDH) technique [110]. All spatial modes other than the fundamental Gaussian are filtered out by a pre-mode cleaner and a mode cleaner, both of which operate in vacuum. After the pre-mode cleaner, the laser beam passes into the LIGO vacuum system, which houses the primary optical components of the detector [1]. A vacuum of $\approx 10^{-8}$ – 10^{-9} Torr is maintained in this system, which includes the 4 km-long, 1.2 m-diameter beam tubes. The vacuum system also includes approximately 200 baffles for trapping scattered light, which can significantly contaminate the detector readout.

The LIGO detectors employ a suspension system that aims to mitigate local environmental and seismic effects. The mirrors are each suspended as a pendulum by a loop of steel wire, which provides $1/f^2$ damping above the pendulum’s resonant frequency. The pendula have a resonant frequency of 0.76 Hz [19], which allows the mirrors to be treated as freely-falling test masses in the GW band (40 Hz and higher). This provides an isolation factor of 2×10^4 at 100 Hz, which is the most sensitive region of the frequency band for the LIGO detectors [1]. Seismic noise is also reduced by the use of isolation stacks formed of alternating layers of springs and metal slabs, which provide f^{-8} attenuation above 10 Hz [111]. The structures from which the mirror pendula hang are mounted on these isolation stacks.

The input and end test masses are 10.7 kg cylinders of silica, 25 cm in diameter and 10 cm thick. The silica mass substrates are polished in order to minimize optical losses;

surface deviations from sphericity are $\lesssim 0.5$ nm. They also feature low-absorption coatings made of multiple dielectric layers, which aim to reduce thermal noise and distortion of the mirror surface due to thermal lensing. However, this effect still occurs, especially in the input test masses where the laser beam passes through the substrates. To counteract this effect, a CO₂ laser subsystem is used to heat the masses and actively control the thermal lens [1].

There are several degrees of freedom in the LIGO interferometers which must be precisely controlled to achieve high enough sensitivity to detect gravitational waves. First, the length of each optical cavity in the interferometer must be kept at an integer multiple of the laser’s wavelength in order to maintain resonance. These lengths are sensed using a modified version of the PDH reflection scheme [112]. Error signals are generated using heterodyne detection of radio-frequency sidebands at the REF, PO, and AS ports shown in Fig. 2.2. The REF port gives a signal proportional to deviations in the laser frequency relative to the average arm length, the PO port gives error signals for the recycling cavity length and the Michelson phase, and the AS port gives an error signal proportional to differential arm length changes [1]. The error signal from the REF port is used as a feedback signal for frequency stabilization of the laser, while the REF, PO, and AS error signals are used as feedback controls to drive coil-magnet actuators which control the arm lengths. The pitch and yaw of each mirror are also controlled using a hierarchy of feedback loops. These lengths and alignments must be carefully maintained in order to keep the interferometer “locked” so that resonance is achieved and science data can be collected.

2.1.2 Noise sources

The LIGO detectors are subject to a variety of instrumental and environmental sources of noise. These noise sources are typically divided into two classes: displacement noise, which causes changes in the cavity lengths, and sensing noise, which limits our ability to measure the cavity lengths. A plot of the major displacement and sensing noise sources in the LIGO H1 detector is shown in Fig. 2.3.

The remainder of this subsection will discuss some of the primary noise sources in the LIGO detectors.

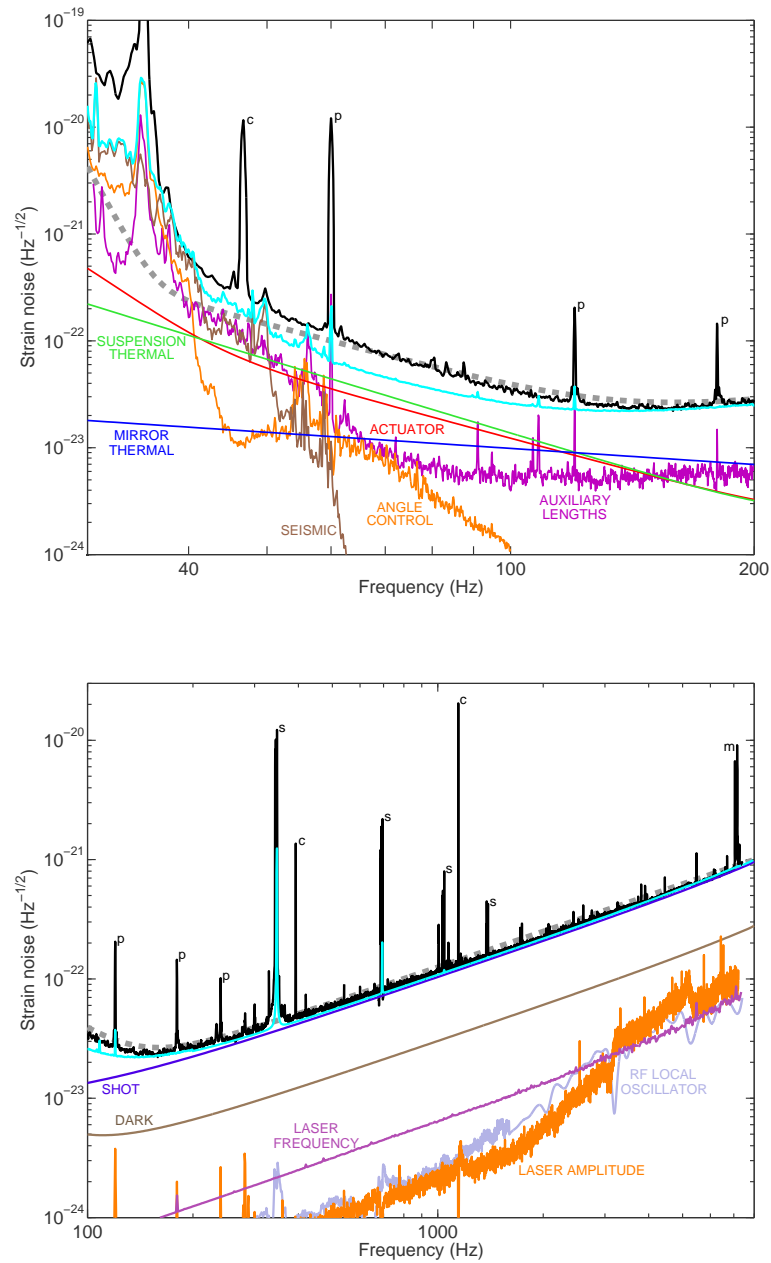


Figure 2.3: Primary sources of displacement noise (top) and sensing noise (bottom) in the LIGO H1 detector. The black curves show the measured strain noise, the dashed gray curve shows the design goal, and the cyan curve shows the root-sum-square of all known noise sources. Noise peaks correspond to calibration lines (c), power line harmonics (p), suspension wire vibrational modes (s), and mirror vibrational modes (m). Figure and information reproduced from Ref. [1]. Copyright IOP Publishing. Reproduced with permission. All rights reserved.

Seismic noise

Seismic noise is the dominant noise source for the LIGO detectors below 45 Hz; its contribution to the detector noise spectrum sharply increases with decreasing frequency and sets the lower limit of 40 Hz for the GW band of the detectors. It is generated by large-scale motions of the earth from seismic waves, earthquakes, ocean tides, and human activity. Seismic noise causes motions of the test masses and is managed through the use of pendula and isolation stacks. A related source of noise, called “gravity-gradient noise,” or “Newtonian noise,” is generated by variations in the local gravitational field due to seismic or atmospheric density perturbations. Newtonian noise was not a concern for the initial LIGO detectors since there were other, more significant noise sources in the relevant frequency band; however, it may become more problematic in future GW detectors where these noise sources are reduced. This will be discussed further in Chap. 6.

Shot noise

Shot noise is a type of sensing noise in the LIGO detectors, and is the dominant noise source above ≈ 100 Hz. Although the laser power is carefully regulated, statistical fluctuations in the number of photons emitted arise due to the particle nature of light. This leads to variations in the readout power, which we interpret as length changes in the interferometer arms.

For a GW which induces a strain $h(t) = h e^{-2\pi i f_{\text{GW}} t}$ in an interferometer, the resulting phase change for a single beam of light is [11, 108]

$$\delta\phi = \frac{hc}{f_{\text{GW}}\lambda} \sin(2\pi f_{\text{GW}}\tau_s) e^{-\pi i f_{\text{GW}}\tau_s}, \quad (2.1)$$

where λ is the wavelength of the laser and τ_s is the storage time of the light. In a real Fabry-Pérot cavity, the output light is a combination of beams which have taken different numbers of round trips through the cavity. Ref. [108] gives the result for this case:

$$\delta\phi = \frac{h\pi Gc}{\lambda} \frac{\tau_s}{[1 + (2\pi f_{\text{GW}}\tau_s)^2]^{1/2}}. \quad (2.2)$$

Here, G is the cavity gain. From here, we use basic uncertainty propagation and Poisson

statistics to derive the expected strain noise in an interferometer due to shot noise:

$$\sigma_{h,\text{shot}} = \left[\frac{\hbar\lambda [1 + (2\pi f_{\text{GW}}\tau_s)^2] \Delta\nu}{4\pi P_{\text{BS}} c \tau_s^2} \right]^{1/2}. \quad (2.3)$$

Here, $\Delta\nu$ is the measurement bandwidth and P_{BS} is the laser power at the beamsplitter. We note that $\sigma_{h,\text{shot}}$ increases with frequency and decreases with increasing laser power.

Radiation pressure noise

Radiation pressure noise occurs in GW interferometers because photons transfer momentum to the test masses when they are reflected; the resulting motion of the test masses is a source of displacement noise. In the initial LIGO detectors, radiation pressure was well below seismic noise and was not considered a significant noise source; however, this may change for Advanced LIGO, since the laser power will be significantly increased.

In some time interval Δt , the average number of photons passing through the beamsplitter is

$$N_\gamma = \frac{P_{\text{BS}}\lambda\Delta t}{4\pi\hbar c}, \quad (2.4)$$

with fluctuations $\sqrt{N_\gamma}$ by Poisson statistics. The momentum imparted to a single test mass by these fluctuations is

$$p_{\text{mass}} = \left(\frac{4\pi\hbar P_{\text{BS}}\Delta t}{c\lambda} \right)^{1/2}. \quad (2.5)$$

From here, it is trivial to calculate the strain spectrum from the force, using $F = dp/dt = 4\pi^2 f^2 m x$. We add an extra factor of 2, since each cavity contains two mirrors, and an extra factor of N accounts for the number of trips taken by the light in the cavity [11]. The end result is

$$\sigma_{h,\text{rad}} = \frac{N}{m f^2 L} \sqrt{\frac{2\hbar P_{\text{BS}}}{\pi^3 c \lambda}}, \quad (2.6)$$

where m is the mirror mass. We note that $\sigma_{h,\text{rad}}$ decreases with frequency and increases with increasing laser power; this is the opposite behavior of that observed for shot noise. Because of the counterplay between these two quantum noise sources, there is an optimal

laser power $P_{\text{BS,opt}}$ which minimizes the total quantum noise at a particular frequency:

$$P_{\text{BS,opt}}(f) = \frac{\pi m L \lambda f^2}{4N\tau_s} \left[\frac{1 + (2\pi f \tau_s)^2}{2} \right]^{1/2}. \quad (2.7)$$

Using this power results in the minimum achievable quantum noise for an interferometer, called the “standard quantum limit:”

$$h_{\text{SQL}}(f) = \frac{1}{\pi f} \left[\frac{\hbar N}{m L c \tau_s} \sqrt{\frac{1 + (2\pi f \tau_s)^2}{2}} \right]^{1/2}. \quad (2.8)$$

This limit is only achievable at a single frequency at a time, since $P_{\text{BS,opt}}$ is a function of frequency.

This quantum limit can be circumvented by the use of squeezed states of light, which contain quantum-correlated photons [113]. Through the use of an optical parametric amplifier, it is possible to decrease phase noise at the expense of amplitude noise, and vice versa. Squeezing has been tested in both the GEO 600 and LIGO detectors and has demonstrated a factor of 13 dB improvement below shot noise [113]. This technique may be utilized to improve the sensitivity of future generations of GW interferometers.

Thermal noise

Thermal noise in GW interferometers arises due to mechanical losses and dissipation in various parts of the detector. The most significant thermal noise in LIGO occurs due to resonant modes in the test masses, in the test mass suspensions, and in the optical coatings of the test masses. Since most of the energy is concentrated near the resonant frequencies of these modes, all of these components are designed such that their resonances are as far as possible from the most sensitive frequency band of the detector. They are also engineered to have extremely low mechanical dissipation in order to reduce thermal energy deposition.

In the mirror pendula, thermal noise arises due to flexing near the end of the suspension wire [1], which is made as thin as possible in order to reduce this effect. While the silica test masses themselves have extremely low mechanical losses, the dielectric coatings on them have relatively high dissipation. This effect is amplified by the fact

that these coatings are on the faces of the test masses which reflect the laser beam.

Overall, thermal noise is a significant noise source in the 40–100 Hz band and near resonant frequencies; in the lower and higher parts of the LIGO frequency band, seismic noise and quantum noise are more problematic.

Other noise sources

Several other sources of noise have been identified in the LIGO detectors, although not all of them are thoroughly or quantitatively understood. These include laser frequency and amplitude noise, vibrational modes of the mirror pendula (called “violin modes”), harmonics of the 60 Hz power lines, charge build-up on the test masses, noise in the actuation coils, and a variety of noise sources associated with the servo system which controls many degrees of freedom in the interferometers [1]. Along with these sources, there are also local environmental effects which can couple to the GW channel of the detector. For example, airplanes passing within the vicinity of the detector, lightning strikes, and stray magnetic fields can inject significant bursts of non-stationary noise into the detector. To monitor these effects near the LIGO observatories, the detectors are outfitted with a set of physical environmental monitors, including microphones, accelerometers, seismometers, and magnetometers. The LIGO Detector Characterization (DetChar) group uses the information from these monitors to identify and flag stretches of data which may be contaminated by these effects [1].

2.2 Data analysis

Gravitational-wave signals are expected to encompass a wide range of characteristics, owing to the variety of astrophysical and cosmological mechanisms which may produce them. The duration of these signals may be anywhere from milliseconds to essentially persistent, and they can be monochromatic, narrowband, or broadband in frequency. Some types of astrophysical systems, like compact binaries, are well-modeled in numerical relativity, and the expected GW signals from these sources are relatively well-understood. However, other cases, like core-collapse supernovae, are more chaotic and difficult to model; along with this, there may also be GW production mechanisms which are presently unknown.

Depending on the morphology and the availability of models for a particular GW signal, different data analysis strategies may be optimal. For this reason, data analysis within the LSC is divided into four main groups, each specializing in a different class of GW signal. Here, I will discuss the potential sources and analysis strategies of each group, as well as some of the published observational results.

Many of the recent analysis results include data from the Virgo detector [114], which is a 3 km-long interferometer near Pisa, Italy. The LSC and Virgo collaborations have a joint agreement which includes data sharing; the availability of data from three detectors is immensely useful for improving overall sensitivity to GWs and localizing source sky positions. The first-generation Virgo detector collected data in four science runs (referred to as Virgo Science Runs (VSR) 1, 2, 3, and 4) between 2007 and 2011, before shutting down operations for upgrades to the Advanced Virgo configuration.

2.2.1 Compact binary coalescences

The compact binary coalescence (CBC) search focuses on transient GW signals, lasting $\mathcal{O}(1\text{ s})$ or less, produced by the coalescence of neutron star-neutron star, neutron star-black hole, and black hole-black hole binary systems. These binary systems continually produce gravitational waves over their lifetime, but only enter the sensitive band of GW interferometers during the final seconds of their inspiral, merger, and ringdown, when they increase sharply in both amplitude and frequency. More discussion of compact binary coalescences can be found in Sec. 1.3.3.

The main goal of the CBC data analysis group is to detect gravitational waves produced by the coalescence of two compact objects, which has recently been accomplished [10]. They also hope to use future detections to estimate rates of compact binary coalescences and to create a catalog of events for developing a better understanding of the overall population of these binary systems. These signals can be used to test general relativity (especially in the strong-field regime), post-Newtonian theory, and other theories of gravity. Measurements of orbital eccentricity and black hole spins can also provide significant insight into binary formation scenarios. Finally, CBC signals can be used to gain a better understanding of the neutron star equation-of-state and gamma-ray bursts produced during coalescences involving a neutron star [115].

CBC signals are the most well-understood source of GWs and there exist many

numerical models of this type of waveform. A technique called “matched filtering” can be used to search for known signal waveforms buried in well-behaved (nearly Gaussian) noise [116]. This technique is qualitatively similar to doing a cross-correlation between a model waveform (or “template”) and the data. However, the model waveform is not known exactly; it depends on several parameters, including the masses and spins of the two compact objects, the sky direction and inclination of the source, the distance to the source, the polarization angle, and the ending time and phase of the inspiral. As a result, a template bank of waveforms using different parameters is generated and used to analyze the data. Due to computational limitations, the bank cannot be constructed to fit every possible set of parameters exactly; instead, the template bank has parameter spacings such that the loss of signal-to-noise ratio (SNR) due to template mismatch should be $< 5\%$ [117].

Triggers are generated by filtering the time-series data from a single detector with matched filters from the template bank designed to detect the expected signal. Each trigger has an associated SNR, coalescence time, and other parameters associated with the template. This is done with data from multiple detectors, and coincidence is required between detectors in order to eliminate spurious triggers occurring due to instrumental artifacts.

The CBC data analysis group also includes a strong focus on parameter estimation. This analysis uses GW candidate triggers generated by the matched filtering search and performs a full Bayesian analysis using Markov chain Monte Carlo (MCMC) techniques to estimate and set limits on the properties of the GW source. This has been done using both simulated signals and blind hardware injections [118], as well as for GW150914 [65].

Recently, two CBC analysis pipelines detected GW150914 with the highest significance [10] of any LIGO analyses; more information on the analysis details can be found in Ref. [64]. Results from this search were used to test general relativity [67] and to set limits on the merger rates of binary black hole systems [66].

Past CBC analyses have also made significant contributions. A search for low mass binary coalescences ($2\text{--}25 M_{\odot}$) with LIGO S6 and Virgo VSR 2 and 3 data set upper limits on merger rates for NS-NS, NS-BH, and BH-BH binaries [119]. There have also been efforts to analyze the data at extremely low latency (on the order of minutes) so that any significant triggers can be immediately sent to electromagnetic observatories

for follow-up [120].

2.2.2 Gravitational-wave bursts

The LIGO burst analyses are geared towards detecting short-duration GW signals whose waveform is not known. Many astrophysical mechanisms which may produce GWs are not fully understood; as a result, the amplitude and frequency evolution of the GWs are relatively uncertain. Potential sources of GW bursts include core-collapse supernovae, black hole mergers, and seismic activity in neutron stars [1]. Until now, the burst group has focused on GW signals lasting a few seconds or less; however, new developments may enable them to extend the parameter space of their searches in the future.

To search for GW bursts whose precise waveform is not known, excess power methods are used. This technique identifies triggers by finding periods of data which are not consistent with regular detector noise; this requires careful usage of data quality flags and consistency checks to avoid triggering on data artifacts. Most burst searches use a time-frequency representation of the strain data and some form of clustering to localize triggers in both time and frequency. Some burst analyses are incoherent, meaning that triggers are first identified in individual detectors, and then a time coincidence between multiple detectors is required for possible GW candidate events (in order to remove triggers due to data artifacts). Other analyses are coherent, using either cross-correlation or a coherent network analysis [115].

Even though precise signal waveforms are not known, burst searches still use simulated GW signals to estimate their sensitivity to GWs based on their morphologies, frequency content, and amplitude. Many of these are *ad hoc* waveforms, meaning they are not generated by simulating a specific astrophysical event, but are chosen to encapsulate various morphologies which real GWs could have. Sine-Gaussian bursts have been used to test sensitivity to monochromatic signals, while Gaussian bursts and white noise bursts are used to test broadband sensitivity. Black hole ringdown waveforms were also used in a recent burst search [107].

A variety of astrophysical events are expected to produce both gravitational waves and other energetic phenomena, including gamma-ray bursts and high-energy neutrinos. Triggered burst analyses use observations of these phenomena to identify times and sky directions which may contain GW signals. Observing an astrophysical event through

both GWs and other phenomena would greatly enhance confidence in both the event itself and in our astrophysical understanding of the process which generated it. The burst and CBC groups analyzed LIGO S6 and Virgo VSR 2 and 3 data which coincided with 154 gamma-ray bursts and placed lower bounds on the distance to the progenitor, with some assumptions about the energy emitted [121]. There are also many cases where electromagnetic observatories follow up on potential GW candidates; initial GW triggers are generated using a low-latency pipeline which identifies significant triggers within minutes [122, 123]. GW triggers have been followed up on by the *Swift* experiment [124] and the IceCube neutrino observatory [125]. A follow-up search for neutrino triggers was also done for GW150914 by ANTARES and IceCube, although no significant neutrino emission was detected [70].

There are also several all-sky burst searches, which make no assumptions about the arrival time or sky direction of GWs. These searches tend to be more computationally intensive due to the large amount of data they must process and the number of sky directions they must consider. An all-sky burst analysis pipeline was actually the first to detect GW150914 [63], although its significance was lower compared to CBC matched-filtering analyses. Past all-sky searches have set limits on the rate of generic strong GW bursts [107], and have also specifically targeted intermediate mass black-hole binaries ($100\text{--}450 M_{\odot}$) [126].

2.2.3 Continuous waves

The continuous waves search aims to detect long-lasting, narrowband gravitational waves which may be emitted by rapidly rotating neutron stars. These neutron stars are expected to have a time-varying mass quadrupole moment due to elastic or magnetic deformations, instabilities like r -modes, and free precession [127]. The waveform of this signal type is expected to be nearly sinusoidal and relatively weak; as a result, it is expected that LIGO may only be able to detect these sources within our galaxy. However, there are many observed pulsars within this range which are expected to be emitters of continuous GWs; the ellipticity of these sources will determine the signal strength and thus, their detectability.

The waveform of a GW signal from a rotating neutron star is a function of six parameters: two for the source position, two for the source orientation, and two for

the spin frequency and its derivative. For searches focused on a known pulsar, some of these parameters can be estimated from electromagnetic observations; with this reduced parameter space, it is possible to perform a more sensitive coherent search. However, for all-sky searches without a specific target, no such simplification is possible. For short datasets, it is still possible to perform a coherent search, but for datasets longer than a few months, semi-coherent approximations are used due to computational limitations. These semi-coherent searches are sensitive only to power and neglect the phase of the signal [1].

The coherent searches are primarily designed for signal detection and parameter estimation. They use a Bayesian method in the time-domain to form a joint likelihood for all detectors involved in the searches. The incoherent/semi-coherent methods use some form of strain power summing in each frequency bin; if any bins are found to contain high SNR, they are often followed up on by coherent search methods.

Targeted searches for continuous GWs have focused on several known objects. Searches directed towards the Crab pulsar and the Vela pulsar [128–130] have beaten the so-called “spindown limit,” which is an indirect upper limit on the strain based on the assumption that the slowing rotation of the pulsar is entirely due to GW emission. These searches have also targeted other known pulsars [131], supernova remnants [132], the galactic center [133], and the low-mass x-ray binary, Scorpius X-1 [134].

One all-sky search used LIGO S5 data to search for any periodic gravitational waves in the 50–1100 Hz range and set strain limits below 10^{-24} in a 200-Hz band [135]. Another search used an \mathcal{F} -statistic to search Virgo VSR1 data from 100–1000 Hz and certain regions of the frequency derivative space which were previously not explored, setting strain upper limits in the most sensitive region of the frequency band of $\approx 5 \times 10^{-24}$ [136]. Finally, another all-sky search targeted continuous waves from spinning neutron stars in binary systems, and was sensitive to binary orbits with eccentricities as high as 0.9 [137].

2.2.4 Stochastic backgrounds

The stochastic search is geared towards detecting a random background of gravitational waves, which may be produced by cosmological processes or the superposition of weak,

individually unresolvable astrophysical sources. This stochastic gravitational-wave background (SGWB) is generally expected to be weak enough that it will be below the intrinsic noise of a single GW interferometer (at current sensitivities). Thus, stochastic analyses are based around the cross-correlation of data from multiple detectors, in order to enhance GW signals and suppress uncorrelated noise sources. Detector characterization work is very important for the stochastic group, in order to identify and remove frequency bins containing correlated noise, which affects the search sensitivity and may masquerade as a signal. The sensitivity of stochastic searches is also dependent on the total integration time, so these analyses generally use as much data as possible.

Stochastic analyses characterize different GW backgrounds based on their frequency spectrum of gravitational-wave energy density, $\Omega_{\text{GW}}(f)$. There are many astrophysical models which predict SGWB spectra for different sources, although most models predict something similar to a power law in the frequency band of LIGO. These analyses divide the data into segments and estimate the power in narrow frequency bins for each segment. All of the data segments are eventually combined using a weighted average to give an overall estimate of the GW power in each frequency bin. A typical segment size is 60 s, with a frequency resolution of 0.25 Hz.

The all-sky stochastic search is focused on detecting an isotropic background of gravitational waves coming from all directions on the sky; this is the gravitational analogue of the electromagnetic Cosmic Microwave Background. This search has been performed several times, most recently analyzing LIGO S6 and Virgo VSR 2 and 3 data. It has set a limit on the total isotropic GW energy density of $\Omega_{\text{GW}} < 5.6 \times 10^{-6}$ in the 41.5–169.25 Hz frequency band, which has surpassed the limit from Big Bang nucleosynthesis in this band [88, 89].

The directional stochastic searches are similar to the all-sky search, but instead of integrating over the entire sky isotropically, they use different mappings to look for localized sources of unmodeled continuous GWs. The radiometer search [138] is designed to look for point sources, and the spherical harmonics search looks for extended sources with an arbitrary angular distribution [139, 140]. These searches have analyzed and set limits on GW strain power across the entire sky. From the resulting sky maps, they were also able to set limits on GWs coming from certain known astrophysical objects, including Scorpius X-1 and the galactic center.

The stochastic radiometer search has also been adapted into a search for “long-lasting GW transients,” called the Stochastic Transient Analysis Multi-detector Pipeline. This search constructs frequency-time maps of cross-correlation SNR and parses these maps with image processing techniques to find groups of pixels which contain significant power and may correspond to GWs. A targeted search using GRB triggers in S5 data has been performed [141], as well as an all-sky search in S5 and S6 data [142]; this all-sky search will be discussed in greater detail in Chap. 5. These searches have focused on transients lasting a few seconds to a few hundred seconds so far, but may eventually be adapted to even longer transient signals, lasting hours to weeks.

Finally, the stochastic group has used the detection of GW150914 to estimate the SGWB energy density due to binary black hole coalescences [69]. More details on the stochastic search and its analysis methods will follow in Chap. 3.

2.3 Advanced LIGO

The Advanced LIGO detectors began collecting their first observational data in September 2015, after a five-year period of upgrades and commissioning. Although the work presented in this dissertation does not use Advanced LIGO data, this is a significant development in the fields of cosmology, astronomy, gravitational waves, and general relativity, so it is worth discussing briefly. In this section, I will describe some of the upgrades to the Advanced LIGO detectors and their expected sensitivity, discuss some results from the first observing run, and mention developments expected in the future.

2.3.1 Upgrades

Significant upgrades have been made to the LIGO detectors in order to increase their strain sensitivity by a factor of ≈ 10 (relative to the initial LIGO detectors) and to improve their sensitivity in the 10–40 Hz range, which was previously dominated by seismic noise. One important change for the Advanced LIGO configuration is the introduction of a signal recycling mirror between the beamsplitter and the detector readout. This cavity increases the storage time of the interferometer and helps increase the sensitivity across a broad range of frequencies [52]. The cavity length can also be adjusted, which allows for some tuning of the interferometer’s frequency response. It can even be tuned for

narrowband operation with improved sensitivity at a single frequency (that of a known pulsar, for example), although the noise levels at other frequencies would be increased in this configuration. Fig. 2.4 shows the Advanced LIGO sensitivity curve for several different tunings and laser powers.




Figure 2.4: Estimated Advanced LIGO sensitivity curves for different tunings of the signal recycling cavity and different laser powers: the case of no signal recycling cavity (solid blue), optimized for neutron star-neutron star inspirals (solid red), optimized for $30\text{-}30 M_{\odot}$ binary black hole inspirals (solid black), no detuning and low laser power (dash-dot gray), no detuning and high laser power (dash-dot yellow), and a narrowband tuning at 1 kHz (dash-dot orange).

The Advanced LIGO configuration can be seen in Fig. 2.5, with the signal recycling cavity defined by the SRM, SR2, SR3, and ITM mirrors.

The laser power in Advanced LIGO has been significantly increased from the 4.5 W used in initial LIGO, in order to reduce shot noise at high frequencies. Once Advanced LIGO reaches design sensitivity, the input laser power may be as high as 125 W. Although




Figure 2.5: Optical configuration of the Advanced LIGO detectors. ITM: input test mass. ETM: end test mass. ERM: end reaction mass. PRM: power recycling mirror. SRM: signal recycling mirror. CP: compensation plate. BS: beamsplitter. FI: Faraday isolator. PD: photodetector. Figure and information reproduced from Ref. [52]. Copyright IOP Publishing. Reproduced with permission. All rights reserved.

this increase in power also leads to increased radiation pressure noise, the weight of the test masses have been significantly increased (from 10.7 kg to 40 kg) to counteract this effect. A potential increase in test mass thermal noise is also to be expected from higher laser power; however, this effect is mitigated by the usage of larger beam spots on the mirrors. The arm cavity length measurement is essentially an average from all photons in the beam, so a wider beam allows for averaging over more of the test mass surface [52]. To suppress seismic noise, the test masses in Advanced LIGO are hung from quadruple pendulum suspensions, with each pendulum providing $1/f^2$ suppression above its resonant frequency.

Although the increased sensitivity of the Advanced LIGO detectors is obviously

beneficial for detecting gravitational waves, it also reveals new noise sources which were less significant for initial LIGO. The detector sensitivity is approaching a regime where certain non-local phenomena, including seismic activity and Schumann resonances, may cause correlated noise between the detectors. Research is ongoing in both of these areas.

2.3.2 First observing run

Advanced LIGO's first observing run (O1) began on September 18, 2015 at 15:00 UTC and ended on January 12, 2016 at 16:00 UTC. The primary scientific result from this period was the first direct detection of gravitational waves [10]. This event, called GW150914, was observed on September 14, 2015 at 09:50:45 UTC, before the actual beginning of O1. However, the detectors were locked and in the same configuration as they would be for the observing run; the only difference was that the system for alerting electromagnetic telescopes for follow-up of significant triggers was not in place. More information about GW150914 can be found in Sec. 1.3.5. All currently published results using O1 data utilize only the first 16 days of coincident data collected around GW150914. The remainder of the data is still being studied by many of the data analysis groups and will be published in the coming months.

The sensitivity of the Advanced LIGO detectors during O1 is shown in Fig. 2.6. Above 200 Hz, it is approximately a factor of two better than the best sensitivity achieved by the enhanced LIGO detectors during S6; below 200 Hz, the sensitivity increase is even more dramatic. GW150914 would likely not have been detected during S6, since its estimated SNR given the S6 sensitivity is ≈ 6 [143]. The design sensitivity of Advanced LIGO is also shown; future upgrades may even be able to improve the sensitivity further.

2.3.3 Future plans

Advanced LIGO's second observing run (O2) is expected to begin in the fall of 2016 and last for six months, while plans for the third observing run (O3) place its start in late 2017, with a duration of nine months [144]. Advanced Virgo is expected to begin observations in early 2017 and may join Advanced LIGO for part of O2.

Advanced LIGO expects to achieve design sensitivity by 2019 and plans to collect data for extended periods of time after this. It may be possible to increase the sensitivity




Figure 2.6: Left: LIGO strain sensitivity during O1 compared to S6. H1 strain is shown in dark red, while L1 is in light red. Right: single-detector SNR as a function of redshift for the merger of two $30 M_{\odot}$ black holes (assuming optimal orientation). Reprinted figure with permission from Ref. [143], copyright 2016 by the American Physical Society.

of the detectors by as much as a factor of two beyond the design sensitivity [143, 145] by using squeezed light, increased mirror masses, and new materials to reduce thermal noise, but this remains to be seen.

Plans to place the former H2 detector in India have recently been approved by the Indian government; however, the earliest this detector is expected to come online is 2022, and it will take a few more years after that to achieve design sensitivity [144]. The Japanese GW interferometer, KAGRA, is expected to begin observations in 2018; the addition of KAGRA and LIGO-India to the worldwide GW detector network will significantly improve localization of GW sources.

Chapter 3

Searches for a stochastic gravitational-wave background

A stochastic gravitational-wave background (SGWB) is expected to result from a superposition of many weak and unresolved sources of gravitational waves. These may be cosmological processes in the early universe or individual astrophysical emitters of GWs. Detection of an SGWB holds the promise of insight into the earliest moments of the universe (cosmological backgrounds) and formation rates and ensemble properties of GW source populations (astrophysical backgrounds).

The layout of this chapter is as follows. Sec. 3.1 will describe some potential sources of a stochastic gravitational-wave background. The LIGO stochastic search group has developed a cross-correlation-based algorithm for analyzing long periods of LIGO data with the goal of detecting an SGWB; this search method will be presented in Sec. 3.2. In Sec. 3.3, I discuss a technique for estimating SGWB model parameters using the results of searches for an isotropic SGWB, and show how it can be applied to models with high-dimensional parameter spaces.

The star formation rate is a key component of astrophysical SGWB models; in Sec. 3.4, I discuss an effort to estimate the star formation rate at high redshift using a hierarchical model of cosmic chemical evolution [146]. In particular, I use this chemical evolution model to estimate the parameters of a high-mass mode of star formation occurring at high redshift. Finally, in Sec. 3.5, I apply our parameter estimation

method to an SGWB produced by stellar core-collapse and estimate its detectability by Advanced LIGO and the Einstein Telescope.

3.1 Sources of stochastic gravitational-wave backgrounds

In general, a gravitational-wave background may be characterized in terms of how its energy density is distributed in frequency. Keeping in line with traditional cosmological definitions, we use a normalized energy density

$$\Omega_{\text{GW}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}, \quad (3.1)$$

where $d\rho_{\text{GW}}$ is the GW energy density contained in the frequency range between f and $f + df$, and ρ_c is the critical density, defined in Eq. 1.87.

Fig. 3.1 shows $\Omega_{\text{GW}}(f)$ for a variety of cosmological and astrophysical SGWB models (solid lines). This includes SGWBs produced by binary neutron stars and binary black holes [147], cosmic strings [29], axion inflation [148], slow-roll inflation [149], and a stiff equation of state [150] (i.e., some energy with $w > 1/3$; see Eq. 1.81 and the following text). Limits from Big Bang nucleosynthesis [85], CMB observations [151], pulsar timing [152], and analyses of LIGO and Virgo data [88, 89] are shown, as well. The CMB limits come from large-scale observations of temperature and polarization fluctuations, as well as constraints on the effective number of neutrino species, while the BBN limit comes from observations of the light element abundances (see Eq. 1.107). Two of the most recent analyses of initial LIGO data have surpassed the BBN limit in the LIGO band [88, 89].

Note that the spectra of these models depend on their various parameters; the spectra shown here are for a specific choice of model parameters intended to be generally representative of the most realistic parametrization of each model. All of these models can be well-approximated by a power law in the few Hz–few kHz region where the LIGO detectors are most sensitive.

Any type of SGWB can be completely characterized by an appropriate choice of $\Omega_{\text{GW}}(f)$ under the assumptions that the background is isotropic, unpolarized, stationary, and Gaussian [153]. Isotropy is a reasonable requirement for a cosmological SGWB, based on observations of the CMB, but it is possible that an astrophysical SGWB could




Figure 3.1: Landscape plot of SGWB energy density spectrum for various cosmological and astrophysical models. Each curve represents a specific choice of parameters for that particular model. Bounds from LIGO analyses are shown as thick black lines in their respective frequency bands. Other bounds are from Big Bang nucleosynthesis (BBN) constraints, CMB observations, and pulsar timing analyses. Figure reproduced from Ref. [89] with permission. Copyright 2016 by the American Physical Society.

be distributed anisotropically across the sky. Regarding polarization, there is no obvious reason why one GW polarization should contain more energy than the other, although there are some models of axion inflation which predict a preference for one polarization due to parity breaking [154]. Based on the expected age of potential SGWBs (Gyr), it seems unlikely that they would vary on the timescale of LIGO observation periods; thus, an SGWB should be effectively stationary. Finally, the assumption of Gaussianity is safe assuming that the SGWB is produced by a large number of independent sources of GWs; by the central limit theorem, the sum of these independent random variables should be Gaussian-distributed.

In this section, I will describe some potential sources of a stochastic gravitational-wave background.

3.1.1 Cosmological models

Numerous models predict a cosmological SGWB produced by inflation-related phenomena, cosmic strings, and other mechanisms in the early universe. Many of these models can be considered speculative, since there are large theoretical uncertainties in the production mechanisms, and in some cases, physics beyond the Standard Model is required [85].

Inflationary models

A stochastic gravitational-wave background may be generated during inflation by a variety of mechanisms, although most theories are related to initial fluctuations which are greatly amplified during the rapid expansion of the universe. These may be zero-point vacuum fluctuations, as well as fluctuations in the inflaton field or in the gravitational field itself. Refs. [31] and [155] describe this amplification during standard slow-roll inflation. There are also models where the inflaton is a pseudoscalar axion [148]; this may extend the duration of inflation, leading to significant non-Gaussianity and increased GW amplitude. Another model predicts significant large-scale inhomogeneities produced by parametric resonance during preheating, a process which causes exponentially increasing excitations of a quantum field and may lead to significantly enhanced particle production [85]. The particle production process and the resulting inhomogeneities are expected to generate gravitational waves [156–159]. For most of these models, the details of the spectra are highly dependent on the form of the inflaton potential.

Cosmic strings

Cosmic strings are one-dimensional topological defects which are hypothesized to have been generated during a symmetry-breaking phase transition in the early universe. They may also be created by the expansion of fundamental strings due to inflation. Since formation, the network of cosmic strings may have evolved by stretching, interacting, oscillating, and evaporating. These interactions may have led to the formation of loops;

kinks and cusps on the loops near the pinch-off point produce strong bursts of gravitational waves. The superposition of these bursts across the cosmic string network could produce an SGWB. The resulting spectrum has two main features: a low-frequency peak and a flat spectrum at higher frequencies; the specific frequency ranges which correspond to these components depends on the average string tension and loop size. This behavior can be seen in the cosmic string spectrum shown in Fig. 3.1.

The overall GW energy density radiated by cosmic strings is primarily dependent on the linear mass density of the string, μ . It can also be parametrized in terms of a deficit angle, $\delta = 4\pi G\mu$, which represents the deviation from Euclidean geometry caused by the presence of a string. If the symmetry breaking which led to cosmic string formation occurred at a mass scale m , then μ is proportional to m^2 [160]. No evidence for GUT scale strings with $\delta \approx 10^{-5}$ has been obtained from observations of large scale structure and CMB anisotropy, but smaller strings could exist and produce gravitational waves [160].

Pre-Big Bang cosmology

Pre-Big Bang cosmology describes a scenario based on string cosmology through which the initial singularity in the standard cosmological model may be avoided [161]. In this theory, the universe begins in the string perturbative vacuum [28], which is in a state of accelerated expansion driven by dilatons. The expansion also leads to an increase in curvature; eventually, higher-order corrections to the string action reduce or stop the growth of curvature. The point of maximum curvature can be associated with the Big Bang and the transition to the radiation-dominated epoch of the standard cosmological model [27].

The model can be parametrized in terms of the string coupling parameter g_s at the transition point between the dilaton-driven and stringy phases and the total redshift acquired during the expansion z_s [26]. The GW energy density spectrum generally has two characteristic features: an initial range where the amplitude increases with frequency, and then an essentially flat region of the spectrum, with some slight oscillatory behavior. However, these features depend some on the choice of model parameters.

Phase transitions

The universe has undergone a number of phase transitions throughout its history, including GUT scale symmetry breaking, electroweak symmetry breaking, and inflation. These transitions are often dramatic events in the history of the universe and provide an environment which is conducive to the production of gravitational waves [85].

During first-order phase transitions, the universe is in a false vacuum state; transition to the true vacuum state may proceed through quantum tunneling in small volumes or “bubbles” [162]. These bubbles quickly expand and eventually reach relativistic speeds due to the pressure difference between their interior (true vacuum) and their exterior (false vacuum). This process transfers the energy from the false vacuum into the kinetic energy of the expanding bubble walls. Eventually, bubbles collide and release significant amounts of energy, generating gravitational waves in the process. The GW spectrum from this process depends primarily on two parameters: the ratio of the rate of change of bubble nucleation to the expansion rate of the universe, and the temperature at which the phase transition occurs.

3.1.2 Astrophysical models

An astrophysical SGWB may be produced by the superposition of many faint and unresolved individual sources of GWs. Such a background can be produced by a significant population of essentially any type of GW source, including supernovae, neutron stars, and compact binary systems. Astrophysical SGWBs are generally expected to contain more energy than cosmological backgrounds (except, perhaps, at relatively low frequencies and very high frequencies); as a result, they may act as foregrounds to the cosmological SGWBs in the frequency band of ground-based GW detectors. In order to make an eventual detection of a cosmological SGWB, it may be necessary to carefully analyze and subtract the astrophysical SGWB(s).

According to Ref. [25], the SGWB spectrum $\Omega_{\text{GW}}(f)$ is primarily dependent on the GW energy spectrum of an individual source dE_{GW}/df . It is also necessary to account for cosmological effects and the rate of occurrence of a particular source, through the comoving star formation rate $R_*(z)$ and the mass fraction of stars converted into progenitors, λ .

The complete formula is given as [25]:

$$\Omega_{\text{GW}}(f) = \frac{8\pi G}{3H_0^3 c^2} \lambda f \int_0^{z_{\text{sup}}} dz \frac{R_*(z)}{(1+z)\sqrt{\Omega_\Lambda + \Omega_m(1+z)^3}} \frac{dE_{\text{GW}}(f_e)}{df_e}. \quad (3.2)$$

Here, f is the frequency in the detector frame, related to the emitted frequency (or frequency in the source frame) f_e by $f = f_e/(1+z)$, and Ω_Λ and Ω_m are the current dark energy and matter densities, respectively, normalized by the critical density ρ_c . The upper limit of the integral (z_{sup}) depends on the maximal redshift (z_{max}) considered in the the star formation rate model and the maximal emission frequency of the source (f_{max}):

$$z_{\text{sup}} = \begin{cases} z_{\text{max}} & \text{if } f < \frac{f_{\text{max}}}{1+z_{\text{max}}} \\ \frac{f_{\text{max}}}{f} - 1 & \text{otherwise.} \end{cases} \quad (3.3)$$

With this framework and some astrophysical knowledge of a particular GW source, a model can be formed for an astrophysical SGWB produced by an ensemble of these sources. Below, I discuss a few astrophysical SGWB models.

Magnetars

Neutron stars with very strong magnetic fields may experience a magnetically-induced quadrupolar deformation and emit gravitational waves as they rotate [163]. The gravitational-wave emission depends on the neutron star mass and amount of deformation, along with the magnetic field strength and the initial rotational period of the neutron star. The energy spectrum from a single magnetar is given by [25]:

$$\frac{dE_{\text{GW}}}{df} = I\pi^2 f^3 \left[\frac{5c^2 R^6}{192\pi^2 G I^2} \left(\frac{B_p}{\epsilon_B} \right)^2 + f^2 \right]^{-1}, \quad (3.4)$$

where I is the neutron star moment of inertia, R is its radius, and B_p/ϵ_B is the ratio of the magnetic field strength at the poles to the to the ellipticity. The initial rotational period of the neutron star fixes the maximum frequency of GW emission and sets the upper limit of the redshift integral (see Eq. 3.3). The full GW spectrum is computed by inserting the energy spectrum into Eq. 3.2.

Binary coalescences

An SGWB may be formed from the superposition of weak and unresolved coalescences of compact binaries. Under the assumption of a circular orbit and considering only quadrupole radiation, the GW energy spectrum depends only on the chirp mass [147]:

$$\frac{dE_{\text{GW}}}{df} = \frac{(G\pi)^{2/3}}{3} (M_c^z)^{5/3} f^{-1/3}. \quad (3.5)$$

Here $M_c^z = (1+z)M_c$ is the observed redshifted chirp mass and M_c is the physical chirp mass.

One refinement is needed to calculate the full GW energy density spectrum for this model. Since there may be a significant delay between the formation of a binary system and its merger, the rate of coalescences at redshift z is dependent on the star formation rate at an earlier redshift. This is accounted for by introducing a time delay term when computing the observed rate of binary coalescences per volume from the star formation rate (see Ref. [147] for more details).

3.2 LIGO stochastic search

The LIGO stochastic search strategy is primarily based around the cross-correlation of data from multiple detectors to enhance GW correlations and suppress uncorrelated noise. A thorough understanding of this search strategy is essential for the parameter estimation work presented later in this chapter. In this section, I describe this cross-correlation search framework and how it is applied to LIGO data.

3.2.1 Effect of a plane gravitational wave on an interferometer

We begin with a plane wave expansion for gravitational metric perturbations in the transverse-traceless gauge at time t and position \vec{x} [153]:

$$h_{ab}(t, \vec{x}) = \sum_A \int_{-\infty}^{\infty} df \int_{S^2} d\hat{\Omega} h_A(f, \hat{\Omega}) e^{2\pi i f(t - \hat{\Omega} \cdot \vec{x}/c)} e_{ab}^A(\hat{\Omega}). \quad (3.6)$$

The sky direction $\hat{\Omega}$ is defined by an azimuthal angle ϕ and a polar angle θ , with the origin at the center of the Earth. The index A runs over the two independent GW polarizations, $+$ and \times , and the indices ab run over spatial coordinates; thus, $e_{ab}^A(\hat{\Omega})$ represents the GW polarization tensor for polarization A . Finally, $\tilde{h}_A(f, \hat{\Omega})$ is a complex Fourier amplitude for polarization A which satisfies the requirement that $\tilde{h}_A(-f, \hat{\Omega}) = \tilde{h}_A^*(f, \hat{\Omega})$ [153].

It is useful to define a coordinate system around the sky direction $\hat{\Omega}$. This is equivalent to rotating the original coordinate system CCW around the \hat{z} -axis by ϕ and then CCW around the new \hat{y} -axis by θ so that the final z -axis points in the direction of the source, $\hat{\Omega}$. In order to make this coordinate system as general as possible, we also include a possible CCW rotation of the source along the sky direction $\hat{\Omega}$ by an angle ψ , which is commonly referred to as the source polarization angle. This results in

$$\begin{aligned} \hat{m} = & (\cos \theta \cos \phi \cos \psi - \sin \phi \sin \psi) \hat{x} + (\cos \theta \sin \phi \cos \psi + \cos \phi \sin \psi) \hat{y} \\ & - \sin \theta \cos \psi \hat{z} \end{aligned} \quad (3.7)$$

$$\begin{aligned} \hat{n} = & (-\sin \phi \cos \psi - \cos \theta \cos \phi \sin \psi) \hat{x} + (\cos \phi \cos \psi - \cos \theta \sin \phi \sin \psi) \hat{y} \\ & \sin \theta \sin \psi \hat{z} \end{aligned} \quad (3.8)$$

$$\hat{\Omega} = \sin \theta \cos \phi \hat{x} + \sin \theta \sin \phi \hat{y} + \cos \theta \hat{z}, \quad (3.9)$$

such that $\{\hat{m}, \hat{n}, \hat{\Omega}\}$ form a right-handed coordinate system.

General polarization tensors for a gravitational wave coming from any direction on the sky $\hat{\Omega}$ can be written as:

$$\begin{aligned} e_{ab}^+(\hat{\Omega}) &= \hat{m} \otimes \hat{m} - \hat{n} \otimes \hat{n} \\ e_{ab}^\times(\hat{\Omega}) &= \hat{m} \otimes \hat{n} + \hat{n} \otimes \hat{m}. \end{aligned} \quad (3.10)$$

This may be compared to Eqs. 1.61 and 1.62, which give the polarization tensors for the specific case of a GW propagating along the z -axis ($\phi = \theta = \psi = 0$).

Following [85], the detector tensor for interferometer I can be defined as

$$d_I^{ab}(t) \equiv \frac{1}{2} [\hat{u}_I(t) \otimes \hat{u}_I(t) - \hat{v}_I(t) \otimes \hat{v}_I(t)]. \quad (3.11)$$

Here, $\hat{u}_I(t)$ and $\hat{v}_I(t)$ correspond to the directions of the two arms of the interferometer; the time dependence comes about due to the motion and rotation of the Earth. This

formulation does not require the arms to be orthogonal, although this is the case for all currently operating GW interferometers.

The antenna pattern functions define the directional sensitivity of an interferometer to each GW polarization as a function of sky direction and source orientation [85]. These functions are given by the product of the detector tensor with the GW polarization tensors.

$$\begin{aligned} F_I^+(\phi, \theta, \psi) &= e_{ab}^+ d_I^{ab} \\ &= \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi \end{aligned} \quad (3.12)$$

$$\begin{aligned} F_I^\times(\phi, \theta, \psi) &= e_{ab}^\times d_I^{ab} \\ &= \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \sin 2\psi + \cos \theta \sin 2\phi \cos 2\psi \end{aligned} \quad (3.13)$$

The GW strain in interferometer I can be defined as the product of the GW plane wave (Eq. 3.6) and the detector tensor (Eq. 3.11); after performing this operation and taking a Fourier transform, we obtain

$$\tilde{h}_I(f) = \sum_A \int_{S^2} d\hat{\Omega} \tilde{h}_A(f) e^{-2\pi i f \hat{\Omega} \cdot \vec{x}/c} F_I^A(\hat{\Omega}) \quad (3.14)$$

We note that Eq. 3.6 and the resulting formalism describes an isotropic search which integrates over the entire sky without preference; however, this may be adapted to perform directional searches. The stochastic radiometer search [138, 164] applies a delta function in sky position to search for point sources of continuous gravitational waves, while the spherical harmonics search [139, 140] integrates over extended regions of the sky using a spherical harmonic decomposition.

3.2.2 Cross-correlation estimator

Using the formalism developed thus far, a cross-correlation estimator for GW power can be constructed. The strain time-series in interferometer I may be written as

$$s_I(t) = h_I(t) + n_I(t), \quad (3.15)$$

where $h_I(t)$ is the strain due to GWs and $n_I(t)$ is the strain due to noise sources. For some observation time T (i.e., the amount of data being analyzed), a cross-correlation estimator S for two detectors I and J can be defined in terms of the Fourier transforms of each detector's time-series:

$$S = \frac{T}{2} \int_{-\infty}^{\infty} df \tilde{s}_I^*(f) \tilde{s}_J(f) \tilde{Q}(f). \quad (3.16)$$

Here, \tilde{s}_I is the Fourier transform of the strain time-series from detector I , \tilde{Q} is a to-be-defined filter function, and the star denotes complex conjugation.

Assuming that there is no correlated noise between the detectors ($\langle \tilde{n}_I^*(f) \tilde{n}_J(f) \rangle = 0$), the expectation value of S depends only on $\langle \tilde{h}_I^*(f) \tilde{h}_J(f) \rangle$:

$$\mu = \langle S \rangle = \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df' \delta_T(f - f') \langle \tilde{h}_I^*(f) \tilde{h}_J(f') \rangle \tilde{Q}(f'). \quad (3.17)$$

Here, $\delta_T(f - f') = \sin(\pi f T) / (\pi f)$ is a finite-time approximation to the Dirac delta function. The expectation value $\langle \tilde{h}_I^*(f) \tilde{h}_J(f) \rangle$ can be calculated based on the four assumptions we have made about the SGWB (see Ref. [153] for more details); the end result for μ is

$$\mu = \frac{3H_0^2}{20\pi^2} T \int_{-\infty}^{\infty} df |f|^{-3} \Omega_{\text{GW}}(|f|) \gamma(|f|) \tilde{Q}(f). \quad (3.18)$$

Here, $\gamma(f)$ is the overlap reduction function, which quantifies the sensitivity reduction which may occur if the two detectors are not colocated or coaligned [165]:

$$\gamma(f) \equiv \frac{5}{8\pi} \sum_A \int_{S^2} d\hat{\Omega} e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c} F_I^A(\hat{\Omega}) F_J^A(\hat{\Omega}). \quad (3.19)$$

Here $\Delta \vec{x}_{IJ} \equiv \vec{x}_I - \vec{x}_J$ is the vector pointing from detector J to detector I .

In the small-signal limit (compared to the intrinsic detector noise), the variance of S is calculated as $\sigma^2 \approx \langle S^2 \rangle$ [153]:

$$\sigma^2 \approx \frac{T}{4} \int_{-\infty}^{\infty} P_I(|f|) P_J(|f|) |\tilde{Q}(f)|^2, \quad (3.20)$$

where $P_I(|f|)$ is the estimator for the one-sided noise power spectrum in detector I . This quantity is normally calculated for time segment t by averaging P_I from some number

of adjacent time segments.

The final quantity of interest is the signal-to-noise ratio, or SNR, defined as $\text{SNR} = \mu/\sigma$. With the current definitions of μ and σ , the square of the SNR should be proportional to

$$\text{SNR}^2 \propto \frac{\left(\tilde{Q}(f), \frac{\gamma(|f|)\Omega_{\text{GW}}(f)}{|f|^3 P_I(|f|) P_J(|f|)}\right)^2}{\left(\tilde{Q}(f), \tilde{Q}(f)\right)}, \quad (3.21)$$

where the parentheses represent the following inner product:

$$(A(f), B(f)) = \int_{-\infty}^{\infty} df P_I(|f|) P_J(|f|) A(f) B(f). \quad (3.22)$$

It is trivial to maximize this ratio by choosing \tilde{Q}_{IJ} to be equal to the right-hand term in the numerator of Eq. 3.21 (up to a constant); thus

$$\tilde{Q}(f) = \mathcal{N} \frac{\gamma(|f|)\Omega_{\text{GW}}(|f|)}{|f|^3 P_I(|f|) P_J(|f|)}. \quad (3.23)$$

Here, \mathcal{N} is a normalization constant, generally chosen so that $\langle S \rangle = \Omega_\alpha T$. Using this filter, the SNR is given by

$$\text{SNR} = \frac{3H_0^2}{10\pi^2} \sqrt{T} \left[\int_{-\infty}^{\infty} df \frac{\gamma^2(|f|)\Omega_{\text{GW}}^2(|f|)}{f^6 P_I(|f|) P_J(|f|)} \right]^{1/2}. \quad (3.24)$$

The SNR is proportional to the square root of the observation time T ; thus, accumulating more data increases the sensitivity of this estimator to an SGWB.

One interesting thing to note is that the optimal filter contains the energy density spectrum $\Omega_{\text{GW}}(f)$, which is not known *a priori*. However, most SGWB models are well approximated by a power law in the LIGO band:

$$\Omega_{\text{GW}}(f) = \Omega_\alpha (f/f_{\text{ref}})^\alpha, \quad (3.25)$$

where Ω_α is a constant and the reference frequency f_{ref} is normally taken to be 100 Hz. The analysis can be repeated for different values of α ; often, stochastic analyses use $\alpha = 0$ for the flat spectrum expected of many cosmological models, and $\alpha = 2/3$ or $\alpha = 3$ for astrophysical models (see Fig. 3.1).

3.2.3 Search pipeline

To analyze data from a LIGO science run, stochastic analyses typically divide the data into segments of a particular duration (often 60 s). A Fourier transform is applied to each segment and coarse graining may or may not be applied to achieve the desired frequency resolution. Typical frequency resolutions of these analyses are 1/4 Hz or 1/32 Hz, although this may change for future analyses.

Next, the cross-correlation estimator S and its variance σ are calculated in each time segment. A weighted average is used to combine the estimates for S and σ from each segment into an overall estimate for the entire science run. The end result is an estimator for Ω_{GW} in each frequency bin. Any frequency bins where the magnitude of the estimator exceeds its variance ($\text{SNR} > 1$) will warrant further study and may be evidence of an SGWB. The overall energy density is also calculated by assuming a power law with exponent α (see Eq. 3.25) and integrating over the entire frequency band of the search.

3.3 Parameter estimation for SGWB models

Even without the detection of an SGWB, the information obtained from stochastic searches can be used to constrain theoretical models, improving our understanding of the universe and its history. This section will describe a Bayesian method which uses stochastic search results to estimate and constrain SGWB model parameters.

This method can use the results of LIGO searches for an isotropic SGWB, as well as expected sensitivity curves of future GW detectors. Using search results will reveal which regions of the model parameter space are not consistent with the estimates for Ω_{GW} from that search, while the use of projected detector sensitivities will illustrate which regions of the model parameter space we can expect to probe with new generations of detectors.

3.3.1 Formalism

Bayesian inference allows us to estimate a set of parameters $\vec{\theta}$ for a particular model given our data D and prior knowledge about the parameters. Bayes' theorem is given

as

$$P(\vec{\theta}|D) = \frac{P(D|\vec{\theta})P(\vec{\theta})}{P(D)}, \quad (3.26)$$

where $P(\vec{\theta}|D)$ is the posterior probability distribution of the parameters, $P(\vec{\theta})$ is the prior probability distribution of the parameters, $P(D|\vec{\theta})$ is the likelihood, and $P(D)$ is the Bayesian evidence, which may be discarded as a normalization factor in parameter estimation. The prior encapsulates our previous knowledge of the parameters (if any), while the posterior represents our new knowledge after incorporating the data. The likelihood function can be thought of as the device which allows us to quantitatively use the data to move from the prior to the posterior and is the key feature in parameter estimation.

By adjusting the integration limits of Eqs. 3.18 and 3.20, the stochastic cross-correlation estimator and its variance can be calculated in a number of small frequency bins. Assuming a flat GW energy density spectrum in each frequency bin, S_i is an estimator for Ω_0 in the frequency bin f_i , and σ_i^2 is the associated variance. Since the goal is to constrain the parameters of an SGWB model, we compare the measured value S_i to the predicted GW energy spectrum of the model $\Omega_M(f_i)$. Because the estimator is formed by averaging over many time segments over the course of a science run, it is expected to be nearly Gaussian; thus, the following likelihood function can be used to do this comparison [166]:

$$\mathcal{L}(S_i, \sigma_i | \vec{\theta}) \propto \exp \left[-\frac{1}{2} \sum_i \frac{(S_i - \Omega_M(f_i; \vec{\theta}))^2}{\sigma_i^2} \right]. \quad (3.27)$$

Here, the index i runs over frequency bins f_i , and $\vec{\theta}$ represents the set of model parameters.

The goal of parameter estimation is to determine the posterior distribution of the parameters $\vec{\theta}$; using the posteriors, confidence intervals for the parameters $\vec{\theta}$ may be calculated. If a prior distribution for any of the parameters $\vec{\theta}$ is known (potentially from previous searches), it can be used to compute the posterior. In cases where a prior distribution is not defined, uniform priors are used. In some cases, these priors may be defined by theoretical or empirical limits; in other cases where our knowledge is more speculative, the limits are chosen so as to encapsulate all of the significant posterior

mass.

As an example, I apply this framework to a power law model (see Eq. 3.25) with two parameters: the amplitude Ω_α and the spectral index α . Because this model has only two parameters, it is possible to perform this analysis by “brute force” calculation: the likelihood function is simply evaluated at each point on a grid in the parameter space of the model and used to calculate the evidence (or marginalized likelihood) and the posterior distributions. From the posteriors, confidence intervals on the model parameters may be estimated.

Using a 200×200 grid in the two-dimensional parameter space of the model, I calculate the GW energy spectrum $\Omega_{\text{GW}}(f)$ at each point and compare it to either LIGO stochastic search results or projected detector sensitivities using Eq. 3.27. For this example, I use the results of the S5 isotropic search as well as the expected design sensitivities of the Advanced LIGO and Einstein Telescope detectors.

I use uniform priors with ranges $\Omega_\alpha \in [10^{-15}, 10^{-5}]$ and $\alpha \in [-10, 10]$, and a reference frequency of $f_{\text{ref}} = 100$ Hz. The resulting posteriors are shown in Fig. 3.2; the regions of parameter space to the left of the curves are consistent with the data or projected sensitivity at 95% confidence. Thus, the S5 result allows us to eliminate the portion of the parameter space to the right of the solid curve. For $\alpha = 0$, this corresponds to an upper limit on $\Omega_0 < 7.8 \times 10^{-6}$, compared to the limit from the full S5 analysis of $\Omega_0 < 6.9 \times 10^{-6}$ [88]. This difference likely arises from the fact that a larger parameter space is used in this study.

For the projected sensitivities of Advanced LIGO and the Einstein Telescope, the regions of parameter space to the right of the respective curves would be incompatible with a null result at 95% confidence; thus, if these detectors do not observe an SGWB, we can safely eliminate these regions of the parameter space.

However, this brute force approach can be quite computationally expensive for models with many parameters due to the large dimensionality of the parameter space. This simple model required only 40000 likelihood evaluations using a 200×200 grid, but a model with even four parameters would require 1.6 billion likelihood evaluations to achieve a similar grid density. Because of this, it is necessary to pursue other techniques for evaluating the likelihood function of more complex models.




Figure 3.2: Two-dimensional parameter space of a simple power law model of an SGWB. Regions to the right of the solid curve are not consistent with the results of the S5 isotropic search and can be eliminated at 95% confidence. Given the projected sensitivities of Advanced LIGO and Einstein Telescope, we will be able to eliminate regions of parameter space to the right of the respective curves, even without a detection of an SGWB.

3.3.2 Nested sampling

For models with high-dimensional parameter spaces, we use nested sampling to efficiently evaluate the multidimensional likelihood function. Nested sampling is a computational technique for model comparison in Bayesian statistics which was developed by John Skilling [167–169]. It is primarily intended to efficiently calculate the evidence; however, it also produces estimates of the posterior distributions, which can be used for parameter estimation.

The nested sampling procedure begins with some set of n live points in the model

parameter space which are randomly sampled from the prior distribution. The likelihood function is evaluated at each point, and the point with the lowest likelihood L_{\min} is replaced by another point drawn from the prior. There exist many different methods for selecting this new point, but one procedure which is often used is to select another of the live points and then take a random step in each parameter to generate a new point. This procedure is repeated over and over in order to sample the regions of parameter space with the highest likelihood. The analysis is eventually terminated when the evidence has been determined to some user-specified precision.

For SGWB parameter estimation, I have utilized a nested sampling algorithm called `MultiNest` [170–173]. This algorithm features several improvements over a basic implementation of nested sampling. Typically, the most time-consuming aspect of nested sampling is replacing the point with $L = L_{\min}$ by a new point with higher likelihood; if the points are drawn randomly from a uniform distribution, it would take increasingly longer as the set of live points increases in likelihood (i.e., as L_{\min} increases). For an N -dimensional parameter space, `MultiNest` generates an approximate likelihood contour with $L = L_{\min}$ in the form of an N -dimensional ellipsoid in order to efficiently find points in parameter space with $L > L_{\min}$.

`MultiNest` is also capable of sampling multimodal distributions and separating the modes. To do this, it groups the live points into clusters, generating a separate ellipsoidal bound for each cluster which is sufficiently separated from other clusters. The algorithm can then evaluate the local evidence and separately constrain the parameters for each mode.

A kernel density estimation function from the `CosmoMC` package [174] is utilized for estimating posterior distributions based on the samples from `MultiNest`.

3.4 Star formation rates

A key component of an astrophysical SGWB model is the star formation rate (SFR) since most, if not all astrophysical SGWBs are expected to be the product of some stellar process or the dynamics of a stellar remnant. Here, I describe a study aimed at estimating new models of the SFR for use with our SGWB parameter estimation method.

3.4.1 Observations of star-forming galaxies

In the mid-2000s, Hopkins and Beacom [175, 176] compiled a large number of luminosity density observations and used those to estimate the star formation rate up to $z \approx 6$. More recently, Behroozi et al. [177] have used Markov chain Monte Carlo methods on observations up to $z = 8$ in an attempt to extract the SFR and reduce systematic uncertainties. The resulting SFR is generally in agreement with that of Hopkins and Beacom, although it peaks at a slightly lower redshift.

Recently, luminosity density measurements have been extended out to $z \approx 11$ [178–186], with a steep drop-off after $z > 8$. It is trivial to extract the star formation rate under the assumption that it is proportional to the measured luminosity [187, 188]. However, there are well-known uncertainties related to measurements of the luminosity density at high redshift, including correcting for dust attenuation in the far-UV portion of the spectrum and the fact that surveys may miss faint galaxies which are products of early star formation [146].

3.4.2 Gamma-ray burst rate

Another method for estimating the SFR, especially in low-luminosity galaxies, is based around connecting the rate of gamma-ray bursts (GRBs) at high redshift to the SFR [189–196]. The rate of GRBs ($d\dot{N}/dz$) may be estimated from the SFR by

$$\frac{d\dot{N}}{dz} = F(z) \frac{\mathcal{E}(z) R_*(z) dV/dz}{\langle f_{\text{beam}} \rangle (1+z)}, \quad (3.28)$$

where $F(z)$ is the fraction of bursts that can be seen at a given redshift, $\mathcal{E}(z)$ is the fraction of stars that produce GRBs, $\langle f_{\text{beam}} \rangle$ is the fraction of GRBs which are observable (due to beaming), dV/dz is the comoving volume element, and $R_*(z)$ is the SFR [197–199]. Note that $F(z)$ depends on the experimental capability of observing the initial burst of gamma rays and extracting a redshift from the afterglow [197].

Assuming that $F(z)$ and $\langle f_{\text{beam}} \rangle$ are roughly constant, the GRB-producing fraction of stars can be modeled as $\mathcal{E}(z) = \mathcal{E}_0(1+z)^\eta$. Ref. [199] estimated $\eta \approx 1.2$ in their analysis, implying that a larger fraction of stars produced GRBs in the past; they hypothesize that the GRB rate may be related to cosmic metallicity.

This functional form and the relative abundance of accurate GRB and SFR measurements at low redshift can be used to calculate the unknown quantities. The expected number of GRB observations between redshifts $z=1-4$ can be calculated as:

$$\begin{aligned} N_{1-4}^{\text{exp}} &= \Delta t \frac{\Delta\Omega}{4\pi} \int_1^4 dz F(z) \mathcal{E}(z) \frac{R_*(z)}{\langle f_{\text{beam}} \rangle} \frac{dV/dz}{1+z} \\ &= \mathcal{A} \int_1^4 dz R_*(z) (1+z)^\eta \frac{dV/dz}{1+z}. \end{aligned} \quad (3.29)$$

To calculate the unknown constant \mathcal{A} , one may insert the actual number of GRB observations between redshifts of $z=1-4$, invert the equation, and solve for \mathcal{A} . Then, this can be extended to higher redshifts: in a given redshift range z_1-z_2 , the average SFR is proportional to the observed GRB rate, $N_{z_1-z_2}^{\text{obs}}$.

$$\langle R_* \rangle_{z_1-z_2} = N_{z_1-z_2}^{\text{obs}} \left[\mathcal{A} \int_{z_1}^{z_2} dz (1+z)^\eta \frac{dV/dz}{1+z} \right]^{-1}, \quad (3.30)$$

This allows for the extrapolation of the SFR to higher redshifts based on GRB observations. For this study, we use observations from Ref. [199], which extend out to $z \approx 9$.

3.4.3 SFR models

Using recent experimental results which estimate the star formation rate from these methods, we have studied three different models for the star formation rate. For each model, we use the well-measured data at low redshift [177] and add several data points at high redshift from the corresponding references:

Model 1: The SFR estimated from GRB observations by the *Swift* experiment [199].

Model 2: The SFR from Model 1, scaled down by a factor of ≈ 2 at high redshift. This is based on models from Refs. [200, 201] which argue that the GRB rate saturates beyond $z = 4$ due to low metallicity.

Model 3: The SFR based on measurements of the luminosity function in high-redshift galaxies [177].

We then use a least squares method to fit the data¹ from the above references to a

¹Fits performed by V. Mandic.

functional form for the SFR developed by Springel and Hernquist [202]:

$$R_*(z) = \nu \frac{a \exp(b(z - z_m))}{a - b + b \exp(a(z - z_m))} \quad (3.31)$$

The ν and z_m parameters give the amplitude (in $M_\odot/\text{yr}/\text{Mpc}^3$) and redshift of the maximal SFR, while $b - a$ and b are related to the logarithm of the SFR's slope at high and low redshift, respectively. The resulting model parameters are shown in Tab. 3.1 and the SFRs are plotted in Fig. 3.3.

Although Model 3 has a higher peak (both in amplitude and redshift) than the GRB-based star formation rates, the magnitude of $b - a$ is significantly larger, leading to a steeper downward slope with increasing redshift. As a result, Models 1 and 2 have significantly increased star formation at high redshift compared to Model 3.

Model	ν [$M_\odot/\text{yr}/\text{Mpc}^3$]	z_m	a	b
1	0.145	1.86	2.80	2.62
2	0.146	1.72	2.80	2.46
3	0.178	2.00	2.37	1.80

Table 3.1: Parameters of three models of the star formation rate (see text for more details). Fit uncertainties on all parameters are on the order of a few percent (see Ref. [146]).

3.4.4 Chemical evolution

Although our three SFR models are comparable at low redshift, they predict considerable differences in the amount of star formation at high redshift. As a result, each model should lead to significant differences in observable quantities, especially the stellar metallicity and ionization fraction at different redshifts. To expand the domain of this study and further analyze our SFR models, we use a hierarchical model of cosmic chemical evolution, developed in Refs. [203–205]. We apply this chemical evolution model to each of our SFR models in order to estimate several observable quantities and compare them to observations.

In this chemical evolution model, baryons are divided into three reservoirs: stars (and their remnants), the interstellar medium (ISM) or gas, and the intergalactic medium




Figure 3.3: Plot of three SFR models versus redshift. The GRB-based SFRs (Models 1 and 2) indicate significantly more star formation at high redshift compared to the luminosity function-based SFR (Model 3).

(IGM). The mass evolution of each reservoir is governed by set of differential equations based on baryon accretion from the IGM to the ISM, baryon ejection from the ISM to the IGM, the star formation process, which transfers baryons from the ISM into stars, and mass ejection from stars to the ISM due to stellar winds and explosions. Along with the mass evolution, the chemical evolution of the ISM and IGM are tracked as a function of redshift. The chemical fluxes are derived from the mass fluxes with several assumptions [205]:

1. The chemical composition of the accretion flow from the IGM to the ISM is the composition of the IGM at redshift z .
2. Outflows from the ISM to the IGM have the same chemical composition as the ISM at redshift z .

3. The chemical composition of stars forming at redshift z is the same as that of the ISM at redshift z .
4. Matter ejected into the ISM by stars is given by stellar yields.

Along with these assumptions, the chemical evolution model depends on the SFR, the initial mass function, and stellar data, along with the minimum dark matter halo mass in collapsed structures and the efficiency at which supernovae may power outflows from these structures [205]. With this information, the model can track mass flows, individual element abundances, and the ionization fraction over time. Other useful information, including the optical depth and redshift of reionization may be extracted, as well.

Each of the three models for the star formation rate was studied in detail² using this chemical evolution model [146]. Model 1, which uses GRB-based data, results in an excessive optical depth of $\tau = 0.150$ compared to the WMAP value of $\tau = 0.089 \pm 0.014$ [206], although the redshift of reionization for this model, $z_{\text{re}} = 12.47$, is close to the 1σ limit of the WMAP value ($z_{\text{re}} = 11.1 \pm 1.1$). These values are based on an assumption that the fraction of ionizing photons which escape from a galaxy is $f_{\text{esc}} = 0.2$; tests which used reduced escape fractions resulted in more reasonable optical depths, but values for the redshift of reionization which were too low by 3–4 σ .

These results motivate the use of Model 2, which is essentially a scaled-down version of Model 1. The chemical evolution results for this model include a significantly better fit to the optical depth, with $\tau = 0.087$ (using $f_{\text{esc}} = 0.2$); however, the redshift of reionization remains quite low at $z_{\text{re}} = 8.62$. Both Models 1 and 2 are relatively good fits to overall metallicity observations and the rate of Type II supernovae [146].

Model 3 results in an optical depth of $\tau \approx 0.063$, about 2σ below the WMAP value, and $z_{\text{re}} = 7.02$, which is about 4σ below the WMAP value [206]. However, the metallicity and rate of Type II supernovae are both in good agreement with observations (see Ref. [146] for more details). To supplement the low-redshift star formation of Model 3, we consider the addition of a high-mass mode of star formation at high redshift. This scenario would lead to increased optical depth and an earlier redshift of reionization, while maintaining agreement with metallicity and SNII rate observations. This type of

²These studies were performed by E. Vangioni and K. A. Olive.

star formation may comprise Population III stars: a population of massive and metal-poor stars, hypothesized to have been the first generation of stars which formed at high redshift.

We note here that more recent results from the Planck experiment [74] indicate a lower optical depth and a later redshift of reionization than those used in this study; this is discussed in further detail at the end of this section.

3.4.5 High-mass, high-redshift SFR

Because there are essentially no observational constraints on the star formation rate at $z > 11$, we have considerable freedom to choose the parameters of this high-mass, high-redshift SFR model. In order to estimate the high-mass, high-redshift SFR parameters, I have scanned the parameter space of the model, calculated relevant quantities using the hierarchical cosmic chemical evolution model, and minimized a χ^2 function based on the following constraints:

- Optical depth: I use the WMAP value for the optical depth of $\tau = 0.089 \pm 0.014$ [206]. I also use the redshift of reionization, $z_{\text{re}} = 11.1 \pm 1.1$, defined as the point at which the ionization fraction reaches 50%.
- Metallicity: the overall metallicity (relative to solar metallicity) at redshifts of $z = 4$ and $z = 0$. The observed values are taken to be $Z(4) = 0.03 \pm 0.01$ and $Z(0) = 0.5 \pm 0.5$ [146].
- Individual element abundances: measurements of carbon and oxygen abundances are used to break degeneracies in the overall metallicity. I use high-redshift measurements of these quantities from the SAGA database [207, 208].
- SFR at high redshift: I use five SFR measurements between $z \approx 6$ –10 from Refs. [182, 185] as a constraint.

The individual χ^2 values are defined by taking the difference between the computed value of a parameter and its observational value, divided by the observational uncertainty; the overall χ^2 value is the sum of the individual values. Note that these constraints are evaluated for the combined SFR including Model 3 and the new high-mass mode.

For this high-mass mode, I use a Salpeter initial mass function with a restricted mass range from 36–100 M_{\odot} . I use the Springel and Hernquist functional form to model the SFR (see Eq. 3.31); to estimate the model parameters, I have divided the 2D parameter space in ν, z_m into a 25×25 grid and minimized the overall χ^2 value in each bin using a Nelder-Mead simplex method [209]. In each bin, ν and z_m are not strictly fixed, but are allowed to vary within the bin, while the a and b parameters are allowed to vary between 0 and 30. The overall parameter space for ν and z_m is restricted to ranges of 10^{-5} – $1 M_{\odot}$ and 5–20, respectively. These limits have proved sufficient to fully explore the regions of interest within the parameter space. I repeat this procedure in the 2D a, b parameter space, allowing ν and z_m to vary within the previously stated limits.

Fig. 3.4 shows the minimal χ^2 value in each bin for the ν, z_m parameter space (top panel) and the a, b parameter space (bottom panel). The 68% and 95% frequentist confidence intervals are calculated around the minimal χ^2 value using $\Delta\chi^2$ values of 2.30 and 5.99, respectively, for a χ^2 distribution with two degrees of freedom.

The best fit value is indicated by a star in the upper panel of Fig. 3.4. In the lower panel, the χ^2 function is nearly constant along the diagonal where $a \gtrsim b$. There is also a secondary region of low χ^2 values around $z_m \lesssim 6$, which corresponds to some improvements in terms of the oxygen and carbon abundances. However, this region comprises cases where the a and b parameters are close to 1, which causes the high-mass mode of star formation to be very broad and extend to high redshift. The horizontal feature around $\nu \approx 3 \times 10^{-4} M_{\odot}$ corresponds to the limit where $a = b$ and the SFR becomes a constant equal to ν .

The best fit parameters for the high-mass SFR (with associated uncertainties) are given in Tab. 3.2; however, there are no useful uncertainties for the a and b parameters due to the degeneracy of the parameter space. We note that the functional form of the SFR breaks down when $a \leq b$, effectively eliminating that region of the parameter space.

We define a new model of the star formation rate, Model 4, as the sum of Model 3 and the high-mass SFR model; a plot of this star formation rate is shown in Fig. 3.5 in solid blue, while the high-mass mode is shown in dashed black.

The high-mass mode is peaked around a redshift of $z = 12$, which should ensure that the correct epoch of reionization is achieved. This makes sense because the low-mass portion of Model 4 is already a good fit to the low-redshift SFR on its own.




Figure 3.4: 2D maps of minimal χ^2 in the ν, z_m plane (top) and a, b plane (bottom). 68% (solid black) and 95% (dashed gray) confidence intervals are shown. The location of the point of minimal χ^2 is marked in the top panel with a star. In the bottom panel, the χ^2 value is nearly constant along the diagonal $a \gtrsim b$. Reproduced from Fig. 5 in Ref. [146], with permission.

Model	ν [$M_{\odot}/\text{yr}/\text{Mpc}^3$]	z_m	a	b
High mass	$0.00218^{+0.38}_{-0.03}$	$11.87^{+0.07}_{-0.99}$	13.81	13.36

Table 3.2: Best-fit parameters for the high-mass, high-redshift mode of star formation. Uncertainties on ν and z_m come from the 68% frequentist confidence intervals. There are no useful uncertainties on the best-fit values of a and b due to the degeneracy of the parameter space along $a \gtrsim b$.




Figure 3.5: Star formation rate versus redshift for Model 4 (solid blue) and a high-mass mode of star formation (dashed black). Model 4 is the sum of Model 3 and the high-mass mode.

Overall, the addition of the high-mass mode to Model 3 leads to excellent agreement with the WMAP data in terms of optical depth and the redshift of reionization. The low-redshift portion of this model was already a good fit to the metallicity data and the results of the chemical evolution model indicate that the agreement with metallicity observations is not strongly affected by the addition of the high-mass mode. However,

there is some improvement with regards to oxygen and carbon abundances in metal-poor stars (likely due to the use of these abundances in the χ^2 function). The interested reader should refer to Ref. [146] for more details.

3.4.6 Discussion

The “modified” star formation rates (Model 2 and Model 4) have been used in Ref. [210] to study a stochastic gravitational-wave background due to stellar core-collapse to black holes. However, they found that the detectability of such a background by Advanced LIGO or the Einstein Telescope was not significantly different for either SFR model.

The recent detection of GW150914 [10], a GW signal produced by the coalescence of two heavy black holes, each of mass $M \approx 30 M_\odot$, suggests that the stochastic gravitational-wave background due to binary black hole coalescence may be stronger than previously suspected. Ref. [69] has used the star formation rate from Model 2 to estimate limits on the energy density of such a background and to project when it might be accessible by interferometric GW detectors.

Recent results from the Planck collaboration indicate a lower optical depth than previously expected, with slightly larger uncertainties: $\tau = 0.066 \pm 0.016$ [74]. This also leads to a later redshift of reionization: $z_{\text{re}} = 8.7^{+1.7}_{-1.4}$. As detailed in Ref. [211], Model 3 is compatible with these results at the 1σ level; as a result, the necessity of adding a high-mass, high-redshift component of the star formation rate (i.e., Model 4) is significantly reduced. At the very least, use of the Planck results in this study would likely lead to different best-fit parameters for the high-mass mode of star formation.

Overall, Model 2 is still a good fit to these new data; however, Model 1 is even less compatible with these results than the previous measurements.

3.5 Parameter estimation for an SGWB from stellar core-collapse

The violent and asymmetric process of stellar core-collapse is one of the most promising potential sources of gravitational waves. There are many mechanisms which may produce gravitational waves during the collapse, including rotation during the collapse and core bounce [212], neutrino-driven convection and standing-accretion-shock instability

(SASI) [213–215], asymmetric neutrino emission [214–216], aspherical outflows [217–221], and magnetic stresses [220, 221].

Core-collapse simulations have become more advanced in recent years, and three-dimensional simulations are becoming more common. These simulations are extremely computationally intensive, and as a result, it is not currently possible to provide the full general-relativistic treatment for all relevant processes during the collapse. However, using simplified treatments of certain parts of the core-collapse process, the gravitational-wave signature has been extracted from both 2D and 3D simulations [222–225].

From these simulations, many studies have estimated the stochastic gravitational-wave background produced by these events due to both Pop. I/II and Pop. III stars [226–237]. We propose to take a more systematic look at this model: in this section, I use the parameter estimation techniques detailed in Sec. 3.3 and some of the star formation rates from Sec. 3.4 to study the stochastic gravitational-wave background produced by stellar core-collapse.

3.5.1 Core-collapse spectrum

Functional form

Since the dynamics of the core-collapse process are incredibly complex and not fully understood, we use an empirical form to describe the frequency spectrum of GWs from a single core-collapse event. This functional form was determined by Refs. [226, 233] to be

$$f|\tilde{h}(f)| = \frac{G}{\pi c^4 D} E_\nu \langle q \rangle \left(1 + \frac{f}{a}\right)^3 e^{-f/b}, \quad (3.32)$$

where D is the distance from the star, E_ν is the total energy emitted in neutrinos, $\langle q \rangle$ is the average value of the luminosity-weighted neutrino anisotropy, which encapsulates the angular variance in neutrino emission, and a and b are empirical parameters without a physical interpretation, both with units of Hz.

The GW energy spectrum emitted by a single source can be computed from $f|\tilde{h}(f)|$ [238] as

$$\frac{dE}{df}(f) = \frac{\pi^2 c^3 D^2}{G} \left(f|\tilde{h}(f)|\right)^2. \quad (3.33)$$

For a core-collapse event, this gives

$$\frac{dE}{df}(f) = \frac{G}{c^5} E_\nu^2 \langle q \rangle^2 \left(1 + \frac{f}{a}\right)^6 e^{-2f/b}. \quad (3.34)$$

Using this formula in Eq. 3.2, the spectrum of GW energy density for an SGWB produced by stellar core-collapse is:

$$\Omega_{\text{GW}}(f) = \frac{8\pi G}{3H_0^3 c^2} f \xi \int dz \frac{R_*(z)}{(1+z)\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}} \left(1 + \frac{f(1+z)}{a}\right)^6 e^{-2f(1+z)/b}. \quad (3.35)$$

Because Ω_{GW} is degenerate with regard to E_ν , $\langle q \rangle$, and λ_{CC} , and the values of these parameters are not well-known, we have combined them into a single scaling parameter, ξ , which has units of m^2/s :

$$\xi = \frac{G\lambda_{\text{CC}}}{c^5} E_\nu^2 \langle q \rangle^2. \quad (3.36)$$

Here, λ_{CC} is the mass fraction of stars which undergo core-collapse, which is also not known *a priori*.

Overall, this model has three free parameters: a , b , and ξ . In general, the GW energy density increases as a decreases; at frequencies above a few Hz, it essentially scales proportionally to a^{-6} . As b increases, the rate at which the exponential term in Eq. 3.35 suppresses the GW energy density is decreased; this leads to an increase in the overall energy density and pushes the peak of the spectrum to higher frequencies. Finally, the GW energy density scales proportionally to ξ . Fig. 3.6 shows the GW energy density for various parameter choices in order to illustrate their effects.

Fitting to simulations

Here, we have attempted to fit the functional form for $f|\tilde{h}(f)|$ given in Eq. 3.32 to the GW waveforms extracted from simulations of core-collapse supernovae³. This is done in order to estimate realistic ranges on the a and b parameters for use in parameter estimation.

The general procedure is as follows: from a given simulation's waveform time-series, compute the FFT and multiply by frequency to obtain $f|\tilde{h}(f)|$. Since we are primarily

³Most of the fits described here were performed by K. Crocker [239].




Figure 3.6: $\Omega_{\text{GW}}(f)$ for various parameter choices for a model of an SGWB produced by stellar core-collapse. The units of a and b are Hz, while the units of ξ are m^2/s . See the text for more details on the effects of varying the model parameters.

interested in the lower frequency part of the signal, we do not consider frequencies above 400 Hz for our fit; this is because detector sensitivities decrease above ≈ 100 Hz, the stochastic search sensitivity goes as f^{-3} , and the overlap reduction function (see Eq. 3.19) decreases in overall amplitude with increasing frequency. We also resample the waveform FFT using a logarithmic frequency spacing so that the low-frequency part of the signal (which is of most interest to us) contributes comparatively to the high-frequency part when performing the fit. We then fit our functional form (Eq. 3.32) to the resampled waveform FFT using a least-squares approach.

We also note that several of the strain spectra from these simulations exhibit a low frequency peak which is not well-described by the functional form we have chosen. As a result, we ignore this feature in these fits and attempt to model it with a separate functional form, described in the following section. An example of this peak can be seen

in Fig. 3.7, which shows the `s27fheat1.00` simulation from Ref. [223].




Figure 3.7: Raw $f|\tilde{h}(f)|$ (black) from the `s27fheat1.00` simulation from Ref. [223], assuming a plus-polarized signal and a polar observer. The least-squares fit to the data (excluding the low frequency peak at ≈ 60 Hz) is shown in blue.

Most of the simulations used in this study were obtained from `stellarcollapse.org` [240]; however, a few were obtained directly from their authors [224, 225].

The first simulation that we use comes from Ref. [223], which studies the hydrodynamics of the post-bounce phase of the core-collapse process in a $27 M_{\odot}$ star. The simulation is carried out in three dimensions and uses general relativity, along with an energy-averaged neutrino leakage model which agrees well with full hydrodynamical simulations of neutrino leakage. The full simulation process is repeated four times with different neutrino heating rates and simulation durations. The top-left panel of Fig. 3.8 shows the strain from this simulation for a cross-polarized signal observed along the polar axis of the star.

The second simulation comes from Ref. [241] and focuses on studying the angular momentum dependence of a $12 M_{\odot}$ progenitor. This is a 2D simulation which assumes axisymmetry and uses general-relativistic hydrodynamics. The simulation is performed several times with varying values for the core angular momentum. The top-right panel

of Fig. 3.8 shows the strain from this simulation for a plus-polarized signal observed along the equatorial plane of the star.

Another simulation from Ref. [225] studies the collapse of four stellar progenitors with masses between $12\text{--}25 M_{\odot}$ in two dimensions, assuming axisymmetry. It uses multifrequency neutrino transport with relativistic corrections and Newtonian hydrodynamics, although some effects of strong-field gravity are accounted for using an effective potential. All models used in the simulation are non-rotating. The strain from this simulation is shown in the bottom-left panel of Fig. 3.8 for a plus-polarized signal.

Finally, we consider a waveform from Ref. [224], which uses a two-dimensional model and multi-group, three-flavor neutrino transport. These simulations were carried out for six different progenitor masses between $8.1\text{--}27 M_{\odot}$ and use the neutrino hydrodynamics codes `Vertex-CoCoNuT` [242] and `Vertex-Prometheus` [243, 244]. A plot of the strain from one of these simulations is shown in the bottom-right panel of Fig. 3.8 for a plus-polarized signal and an equatorial observer.

The simulated GW waveforms shown in Fig. 3.8 are only a small selection of those included in our study; in total, more than 70 waveforms were studied in order to estimate the ranges of the a and b parameters. From these studies, we have determined that the following parameter ranges allow for reasonable fits of the GW strain spectra produced by these simulations:

- $a \in [5 \text{ Hz}, 150 \text{ Hz}]$
- $b \in [10 \text{ Hz}, 400 \text{ Hz}]$

The ξ parameter is related to the overall SGWB energy density and is not tied to the shape of the strain spectrum from an individual source; thus, we have not estimated its range from this study.

Low-frequency peak

The low-frequency peak which we have observed in the simulations is of particular interest because it tends to occur around $10\text{--}100$ Hz, which is in the most sensitive region of the frequency band of Advanced LIGO and the Einstein Telescope. The LIGO stochastic search will be especially sensitive to this feature, since the sensitivity of the stochastic estimator for GW power goes as f^{-3} (see Eq. 3.18).

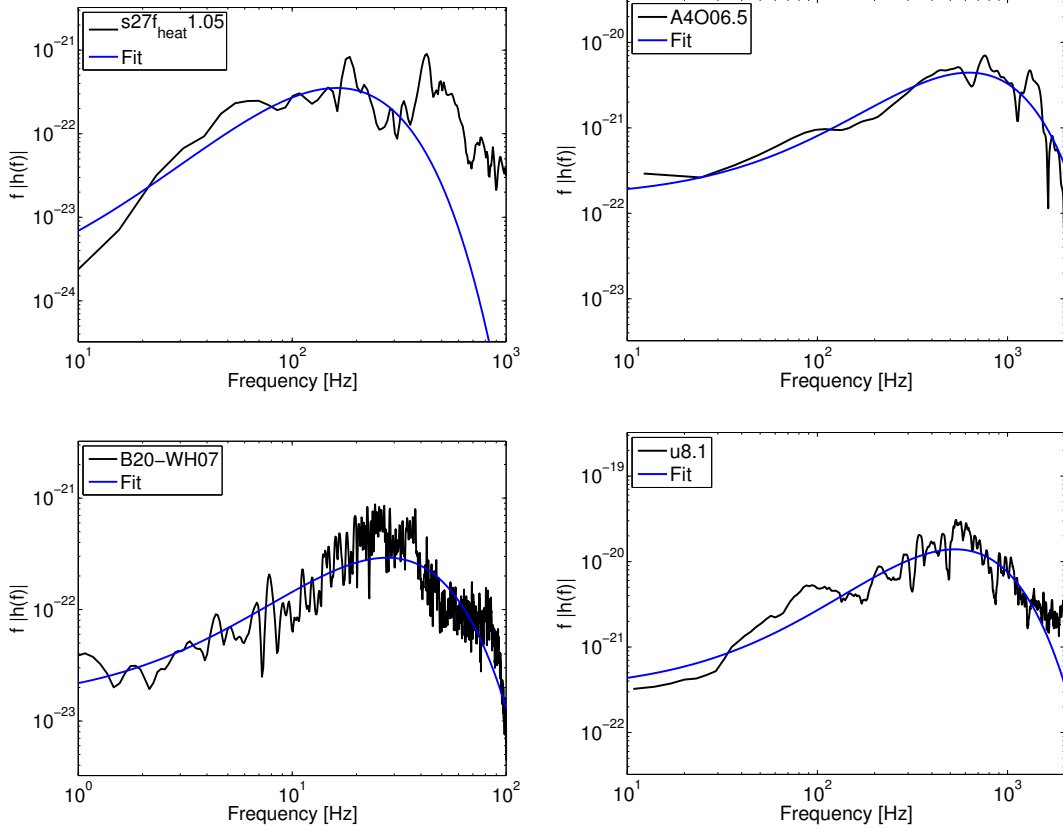


Figure 3.8: Raw $f|\tilde{h}(f)|$ strain spectra (black) and the resulting fits (blue) of our functional form (Eq. 3.32) for four different simulations of core-collapse supernovae. Top-left: $s27f_{\text{heat}}1.05$ simulation from Ref. [223]. Top-right: $A4O06.5$ simulation from Ref. [241]. Bottom-left: $B20\text{-}WH07$ simulation from Ref. [225]. Bottom-right: $u8.1$ simulation from Ref. [224].

Several papers have made note of this low-frequency peak in the GW spectrum and attempted to identify its source. Ref. [224] indicates that a quasi-periodic signal around 100 Hz arises due to prompt post-shock convection during the first tens of milliseconds following core bounce. However, they rule out convection as the direct source of this signal because the entropy and lepton gradients which drive the convection quickly disperse, while the GW emission lasts for tens of milliseconds; instead, they speculate that the convection leads to the development of acoustic waves which produce a GW signal due to the resulting mass motions. Finally, they note that this signal always peaks

near 100 Hz in their simulations; there appears to be no significant dependence on the properties of the progenitor.

Ref. [223] found that early GW emission due to prompt convection sets in at ≈ 10 ms after core bounce and lasts for about 30 ms. They also note that the GW signal produced by prompt convection is particularly sensitive to the perturbations which drive the convection; this is compatible with Ref. [224]’s statement that the signal from prompt convection may depend on the width of the layer in which the convection occurs.

Finally, Ref. [245] found that the speed of core rotation during the collapse process may impact the frequency of the signal from prompt convection. Simulations in which the core was rapidly rotating led to the emission of spiral waves; the acoustic waves which produce the low-frequency GW signal were Doppler shifted by emission on top of the spiral waves, leading to a low-frequency peak closer to 200 Hz. They also note that accurate neutrino transport is a necessary component of these simulations in order to have a reliable prediction for the GW signal; this agrees with similar results from Ref. [224], which found significantly reduced GW emission at low frequencies when simplified neutrino rates were used.

To summarize, there is significant evidence that the low-frequency peak is a realistic feature of the GW signal produced during core-collapse; it is present in many different simulations and there is a general agreement that the signal is related to prompt post-bounce convection. Although it seems to vary for the different simulations, the peak in the GW frequency spectrum tends to occur between ≈ 60 –120 Hz, or even as high as 200 Hz for the rapidly-rotating models of Ref. [245]. We also note that not all of the simulations included in our study exhibited this peak; however, it seems that the more modern simulations, which feature general relativistic hydrodynamics and neutrino transport, generally include this feature.

To model this low-frequency peak in the strain spectrum, I use a Gaussian function:

$$f|\tilde{h}(f)| = \frac{A}{D} e^{-\frac{(f-\mu)^2}{2\sigma^2}}. \quad (3.37)$$

Here, μ and σ are the mean and standard deviation of the Gaussian, respectively, A is the amplitude (in meters), and D is the source distance. This results in a GW energy

spectrum of

$$\frac{dE}{df}(f) = \frac{\pi^2 c^3 A^2}{G} e^{-\frac{(f-\mu)^2}{\sigma^2}} \quad (3.38)$$

and a GW energy density of

$$\Omega_{\text{GW}}(f) = \frac{8\pi^3 c}{3H_0^3} f A'^2 \int dz \frac{R_*(z)}{(1+z)\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}} e^{-\frac{(f(1+z)-\mu)^2}{\sigma^2}}. \quad (3.39)$$

Here, I have combined the two scaling parameters A and λ_{CC} into $A' = A\lambda_{\text{CC}}^{1/2}$, with units of $\text{m}^2/\text{kg}^{1/2}$.

This model of the low-frequency peak has three free parameters: μ , σ , and A' . The μ parameter is the frequency at which the peak in the strain spectrum occurs and it controls the effective upper bound of the resulting GW energy density spectrum. This is because as the model is integrated over sources at different redshifts, the spectrum may acquire power at lower frequencies (due to redshifting), but not at higher frequencies.

Referring to Eq. 3.37, σ defines the width of the Gaussian peak and has no effect on the peak amplitude of the strain spectrum; however, when it is integrated over redshift to calculate Ω_{GW} , a wider peak results in an increased GW energy density since we are effectively increasing the amount of GW energy emitted during a core-collapse by making the strain spectrum more broad. A larger value of σ also allows for some amount of GW energy density at frequencies above μ , while low values of σ result in a sharp cutoff at $f = \mu$. Finally, $\Omega_{\text{GW}}(f)$ goes as A'^2 , so this is simply a scaling parameter. In Fig. 3.9, I show the resulting GW energy density for this model of the low-frequency peak, using various parameter choices in order to illustrate their effects.

Combined case

I also consider a combined case which includes both the low frequency peak and the higher frequency functional form for the core-collapse spectrum. This results in an energy spectrum of

$$\frac{dE}{df}(f) = \frac{\xi}{\lambda_{\text{CC}}} \left[\left(1 + \frac{f}{a}\right)^3 e^{-f/b} + \pi A' \sqrt{\frac{c^3}{\xi G}} e^{-\frac{(f-\mu)^2}{2\sigma^2}} \right]^2. \quad (3.40)$$

Using this spectrum to calculate the gravitational-wave energy density results in




Figure 3.9: $\Omega_{\text{GW}}(f)$ for various parameter choices for a low-frequency model of an SGWB produced by stellar core-collapse. μ and σ are in units of Hz, while A' is in units of $\text{m}/\text{kg}^{1/2}$. See the text for more details on the effects of varying the model parameters.

$$\Omega_{\text{GW}}(f) = \frac{8\pi G}{3H_0^3 c^2} f \xi \int dz \left\{ \frac{R_*(z)}{(1+z)\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}} \times \left[\left(1 + \frac{f}{a}\right)^3 e^{-f/b} + \pi A' \sqrt{\frac{c^3}{\xi G}} e^{-\frac{(f-\mu)^2}{2\sigma^2}} \right]^2 \right\}. \quad (3.41)$$

The combined model has 6 free parameters: a , b , ξ , μ , σ , and A' . Refer to the previous two sections, especially Figs. 3.6 and 3.9 and the surrounding text, for descriptions of the parameters and their effect on Ω_{GW} .

3.5.2 Parameter estimation

Here, I utilize the parameter estimation framework described in Sec. 3.3 to systematically study the SGWB model described in this section thus far. The goal of this study is to project the sensitivity of Advanced LIGO and the Einstein Telescope to this model. For Advanced LIGO, I use the estimated stochastic search sensitivity with the H1 and L1 detectors operating at design sensitivity for one year. For the Einstein Telescope, the expected sensitivity assumes colocated detectors operating for one year. A plot of the detector sensitivity curves is shown in Fig. 3.10; the SGWB sensitivity with the Einstein Telescope is expected to be about a factor of 100 better than Advanced LIGO above ≈ 50 Hz and even better at lower frequencies.




Figure 3.10: Projected sensitivities of the LIGO stochastic search for Advanced LIGO (solid) and the Einstein Telescope (dashed). In both cases, these projections are for the design sensitivity of the detectors. The Einstein Telescope sensitivity assumes colocated detectors; the sharp peaks in this curve are due to expected instrumental noise lines.

To perform this study, I scan the parameter space of each model; for each choice

of parameters, I calculate the model spectrum $\Omega_M(f)$ and compare it to the detector sensitivity using Eq. 3.27. Since the goal is to project the sensitivity of future detectors to this model, I set $S_i = 0$ and simply compare Ω_M to the projected sensitivity, σ_i .

Uniform priors are used for all parameters; these are liberally defined such that all of regions of significant posterior mass are included.

High frequency

The priors for a and b are based on fitting the higher frequency functional form to core-collapse simulations as previously described; these limits are $5 \text{ Hz} < a < 150 \text{ Hz}$ and $10 \text{ Hz} < b < 400 \text{ Hz}$. Limits on ξ are designed so that we are able to explore the most interesting regions of parameter space. If ξ is too small, the SGWB will not be detectable; too large and it will be unrealistically loud and clearly would have been detected already. These limits are different for each detector; for Advanced LIGO, I use a range of $\xi = 10^6\text{--}10^{13} \text{ m}^2/\text{s}$, and for ET, I use $\xi = 10^2\text{--}10^9 \text{ m}^2/\text{s}$.

This formulation of the SGWB model has only three free parameters, and since the log-likelihood scales proportionally to ξ , this is effectively a two-dimensional problem. As a result, brute force calculation is the simplest solution to determine the posterior distributions. I scan a 200×200 grid in the two-dimensional parameter space in a and b (using a fixed value of ξ) and calculate the likelihood at each point. From here, it is trivial to calculate the likelihood for different values of ξ at the fixed values of a and b .

Using this method, I have estimated the posterior distributions on the model parameters and calculated expected sensitivity contours in the parameter space of the model at 95% confidence. Fig. 3.11 shows these contours for both Advanced LIGO and the Einstein Telescope in the ξ - a plane and the ξ - b plane, after marginalizing over the third parameter in each case.

In all panels, the area below the curves indicates the regions of parameter space which are compatible with a null result; thus, if no SGWB is detected, it will be possible to eliminate the regions of parameter space above the curves.

The posterior in the ξ - a plane tends to decrease with a at constant ξ because $\Omega_{\text{GW}}(f)$ essentially scales inversely to a at frequencies above 1 Hz. This can be inferred from the functional form of the spectrum given in Eq. 3.32. The contours in the ξ - b plane are decreasing with b since as b increases, the GW energy spectrum is pushed to higher

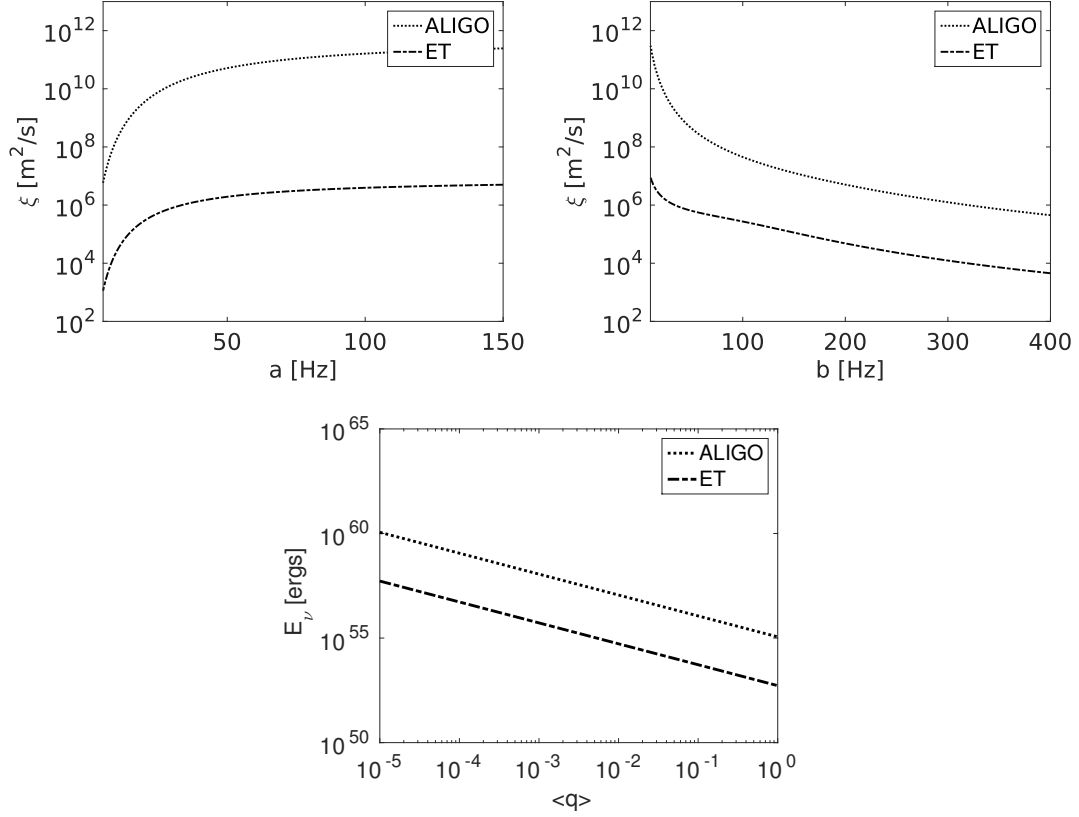


Figure 3.11: 95% sensitivity contours for model parameters of an SGWB from stellar core-collapse. Contours are shown for Advanced LIGO (dotted) and for the Einstein Telescope (dot-dash). The bottom panel translates the sensitivity on ξ into limits on E_ν and $\langle q \rangle$, assuming $\lambda_{\text{CC}} = 0.01 M_\odot^{-1}$. More details are provided in the text.

frequencies, away from the sensitive range of the detectors.

Marginalizing over both a and b gives the posterior distribution on ξ . We calculate an expected sensitivity on this parameter by determining the value of ξ for which 95% of the posterior mass is contained. This results in an expected sensitivity of $1.8 \times 10^{11} \text{ m}^2/\text{s}$ for Advanced LIGO and $3.9 \times 10^6 \text{ m}^2/\text{s}$ for ET.

In the bottom panel of Fig. 3.11, I have used these expected sensitivities on ξ to invert Eq. 3.36 and plot the corresponding $E_\nu - \langle q \rangle$ curves. For this calculation, I have fixed $\lambda_{\text{CC}} = 0.01 M_\odot^{-1}$, based on expected masses required for core collapse and the use of a modified Salpeter initial mass function. Taking $\langle q \rangle = 0.0045$, as Refs. [226, 233] do for

collapse to a neutron star, I obtain neutrino energies of 2.5×10^{57} ergs and 1.2×10^{55} ergs for ALIGO and ET, respectively. This gives an estimate of the average energy required to be emitted in neutrinos during a core-collapse for this background to be detectable. For both detectors, the resulting energy value is quite high; the best energy sensitivity comes from ET and is about two orders of magnitude above expected neutrino energies from supernovae ($\approx 3 \times 10^{53}$ ergs [226, 233]).

Low frequency peak

The priors for the low frequency peak are based on the previously described work with the individual simulations. For most of the simulations, the low frequency peak seemed to occur between 60–120 Hz; however, there are some cases described in the literature where this peak may occur as high as 200 Hz due to rapid rotation during the collapse process [245]. Thus, I have used a liberally defined prior on μ to be uniform in 30–200 Hz.

Most of the waveforms that we studied from Ref. [223] had relatively narrow low-frequency peaks, with typical widths between 40–60 Hz. The low-frequency peaks in the waveforms from Ref. [241] tended to be wider, with some as broad as 140 Hz. As a result, I have used a uniform prior on σ between 10–80 Hz.

Because A' is a scaling parameter like ξ , the prior limits must be carefully defined using the same strategy as described for the limits on ξ in the previous section. This results in ranges of $A' = 10^{-15}$ – 10^{-12} m/kg^{1/2} for Advanced LIGO and $A' = 10^{-17}$ – 10^{-13} m/kg^{1/2} for ET.

I have performed a scan of the parameter space using an identical method to that described in the previous section. Fig. 3.12 shows 95% sensitivities in two-dimensional parameter spaces where I have marginalized over the third model parameter. Regions of parameter space above each of the curves should be detectable by Advanced LIGO (dotted) or the Einstein Telescope (dash-dot); if no detection is made, those parts of the parameter space can be eliminated at 95% confidence.

In the A' – μ plane, the Advanced LIGO curve reaches a minimum near $\mu = 80$ Hz, since this is where the stochastic search achieves its best sensitivity with this detector (see Fig. 3.10). However, the Einstein Telescope is most sensitive near 1 Hz, which is not included in our parameter space for μ . As a result, this contour decreases with decreasing μ .




Figure 3.12: 95% sensitivity contours for low-frequency peak parameters. Contours are shown for Advanced LIGO (dotted) and for the Einstein Telescope (dot-dash). The bottom panel translates the expected sensitivity on A' into limits on A and λ_{CC} . More details are provided in the text.

For both detectors, the contours in the A' - σ parameter space decrease as σ increases. This is because σ acts similarly to a scaling parameter for $\Omega_{GW}(f)$, as previously discussed.

I have also marginalized over the μ and σ parameters to estimate an expected sensitivity on A' at 95% confidence: I obtained values of $A' < 1.4 \times 10^{-14}$ m/kg^{1/2} for Advanced LIGO and $A' < 1.4 \times 10^{-15}$ m/kg^{1/2} for ET. In the bottom panel of Fig. 3.12, I have used these expected sensitivities on A' to calculate contours in the A - λ_{CC} parameter space. Taking λ_{CC} to be $0.01 M_{\odot}^{-1}$ gives sensitivities of $A < 201.7$ m for Advanced LIGO and $A < 19.9$ m for the Einstein Telescope.

I have studied the `s27fheat1.00` simulation of Ref. [223] and the `A1005.5` simulation of Ref. [241] in order to make a comparison to our expected sensitivities on A . For each simulation, I compute the strain spectrum $f|\tilde{h}(f)|$ and fit the functional form from Eq. 3.32 to the simulation (not including the frequencies where the low-frequency peak is prevalent). Next, I subtract this fit from the data and then fit the Gaussian functional form given in Eq. 3.37 to the resulting low-frequency peak.

The best fits for this functional form corresponded to values of $A = 0.18\text{ m}$ for the `s27fheat1.00` simulation and $A = 0.85\text{ m}$ for the `A1005.5` simulation. These values are approximately 3 and 2 orders of magnitude below the expected Advanced LIGO sensitivity to A , respectively, and 2 and 1 orders of magnitude below the expected Einstein Telescope sensitivity. Increasing the assumed value for λ_{CC} from $0.01\text{ M}_{\odot}^{-1}$ to 1 M_{\odot}^{-1} would improve the resulting sensitivities by a factor of 10; however, it seems unrealistic to expect one core-collapse per solar mass of star-forming material. Even in this extreme case, the expected sensitivities would still be larger than the values of A given by the fits to the simulations. Thus, we do not expect Advanced LIGO or the Einstein Telescope to be sensitive to these types of signals unless our assumptions about the rate of core-collapse supernovae are relatively inaccurate.

We note that these expected sensitivities do not make any projections regarding the detectability of gravitational-wave bursts produced by core-collapse supernovae. Rather, we are projecting the sensitivity of the detectors to an SGWB composed of unresolvable bursts for which the “average” individual event is similar to the simulated waveforms we have considered.

Combined model

I have also performed a parameter estimation study for the combined model, which is the sum of the low-frequency peak and the higher frequency model, as shown in Eq. 3.41. In this case, the model has six parameters, and a brute force exploration of the model parameter space is computationally infeasible. In order to study this model, I employ the `MultiNest` algorithm, as detailed in Sec. 3.3.2. The same priors as given in the previous two sections are used for all parameters.

Because the ranges on the ξ and A' parameters span several orders of magnitude, I have used logarithmic priors for these parameters. After calculating the posterior

distributions, it is simple to convert from the posterior on the logarithmic parameter to the posterior on the parameter itself.

$$\begin{aligned} P(x) &= P(\log_{10}(x)) \frac{\partial \log_{10}(x)}{\partial x} \\ &= P(\log_{10}(x)) \frac{1}{x \ln(10)} \end{aligned} \tag{3.42}$$

Fig. 3.13 shows 95% confidence intervals in the b - a and A' - ξ planes for both Advanced LIGO and the Einstein Telescope. In the b - a plane, the region to the right of the curves is enclosed by the contours; thus, these detectors should be sensitive to parts of the parameter space to the left of their respective contours. This behavior is apparent because a acts as an inverse scaling parameter in our model and b acts similar to a scaling parameter.

In the A' - ξ plane, the region below the curves is enclosed by the contours, so these detectors should be sensitive to areas of the parameter space above the contours. The shape of the contours is immediately apparent since both of the parameters illustrated here are scaling parameters.

The 2D posterior in the σ - μ plane spans most of the parameter space and does not yield significant insights into these parameters; as a result, this posterior is not included here.

In Tab. 3.3, I compare the expected sensitivities on the ξ and A' parameters for the combined model to those from the individual low- and high-frequency models. Counterintuitively, the upper limits for the combined model tend to be lower; this likely occurs because the combined model includes both low- and high-frequency contributions, increasing the overall GW energy density and improving the detectability of such an SGWB. This can be seen most significantly for the case of the ξ parameter for Advanced LIGO: the combined model upper limits are four orders of magnitude below that of the individual high-frequency model. This occurs because small values of b result in a spectrum which is peaked at lower frequencies where Advanced LIGO has little sensitivity, requiring extremely large values of ξ to be detectable. When the low-frequency peak is added in, the SGWB is much more detectable in these cases and does not require such extreme ξ values.




Figure 3.13: Posterior distributions on combined model parameters in the b - a plane and the A' - ξ plane. Confidence intervals are shown for Advanced LIGO (dotted) and the Einstein Telescope (dot-dash).

Combined model with injection

In order to demonstrate the capabilities of our parameter estimation method, I have analyzed a simulated SGWB based on this model of stellar core-collapse, including the

Model	ξ [m^2/s]		A' [$\text{m}/\text{kg}^{1/2}$]	
	ALIGO	ET	ALIGO	ET
Low frequency	—	—	1.4×10^{-14}	1.4×10^{-15}
High frequency	1.8×10^{11}	3.9×10^6	—	—
Combined	3.0×10^7	2.3×10^5	1.1×10^{-14}	1.2×10^{-15}

Table 3.3: Comparison of upper limits on ξ and A' for the combined model and the individual low- and high-frequency models. The limits for the combined model are lower compared to the individual models due to the overall increase in SGWB energy density for the combined model.

empirical addition of the low-frequency peak in the form of a Gaussian function. For this simulation, I have assumed 1 year of observation time and Einstein Telescope noise levels. The parameters of this injected signal have been chosen so as to produce distinct low- and high-frequency content and to be above the sensitivity of the Einstein Telescope. These parameter choices are summarized in Tab. 3.4. A plot of $\Omega_{\text{GW}}(f)$ for this model can be seen in Fig. 3.15.

I use the same priors on all six parameters as described in the previous subsection and again use `MultiNest` to study the high-dimensional parameter space of this model. In Fig. 3.14, I show the resulting 2D posteriors in the b - a , σ - μ , and A' - ξ planes, after marginalizing over four of the parameters in each case. The color scale shows the relative height of the posterior throughout the parameter space, while the solid white line indicates the 95% confidence contours. The white star indicates the values of the parameters which were chosen for the injection. In all cases, the posterior distribution and resulting confidence contours are in good agreement with the injected values of the parameters.

However, there is a more significant deviation from the injection parameters for the case of the a and ξ parameters. The posterior distribution is more strongly peaked at $a \approx 53$ Hz than at the injected value of 40 Hz, and at $\xi \approx 3 \times 10^5$ m^2/s than at the injected value of 1×10^5 m^2/s . This is likely due to a slight degeneracy between these two parameters: as previously discussed, a acts as an inverse scaling parameter at frequencies above a few Hz.




Figure 3.14: Projected 2D posteriors for a simulated SGWB from stellar core-collapse, recovered with the Einstein telescope. The color scale indicates the relative height of the posterior throughout the parameter space, the solid white line indicates the contours which enclose 95% of the posterior mass, and the white star indicates the injected parameter values.

I have also extracted the maximum likelihood point sampled by `MultiNest` and plotted the resulting SGWB spectrum, compared to the injected SGWB spectrum and the Einstein Telescope sensitivity, in Fig. 3.15. There are no significant differences between the injected and maximum likelihood spectra in the sensitive band of the detector.

For each parameter, I have marginalized over the other five parameters in order to calculate 95% confidence intervals. These intervals are given in Tab. 3.4 and can be compared to the injected values and the maximum likelihood values. In all cases, the injected value of the parameter is within the resulting confidence interval; however, in




Figure 3.15: Plot of $\Omega_{\text{GW}}(f)$ for injected and maximum likelihood parameters for an SGWB from stellar core-collapse, compared to the expected Einstein Telescope sensitivity. More details are provided in the text.

the case of a , it is somewhat close to the lower limit of the interval. This is likely due to the slight degeneracy between a and ξ , as previously discussed.

Type	a [Hz]	b [Hz]	ξ [m^2/s]	μ [Hz]	σ [Hz]	A' [$\text{m}/\text{kg}^{1/2}$]
INJ	40	120	1×10^5	70	20	3×10^{-14}
ML	54.2	125.4	3.8×10^5	70.0	19.9	3.0×10^{-14}
CI	36.6–65.7	114.3–134.5	$(0.2\text{--}8.3) \times 10^5$	65.3–73.3	18.3–21.4	$(2.1\text{--}3.7) \times 10^{-14}$

Table 3.4: Injected (INJ) parameters of a simulated SGWB from stellar core-collapse, along with maximum likelihood (ML) parameters and 95% confidence intervals (CI) from the resulting parameter estimation. Maximum likelihood parameters are taken from the point in parameter space found by `MultiNest` to have the highest likelihood.

3.5.3 Discussion

In this section, I have studied an SGWB produced by stellar core-collapse. I used an empirical form developed by Refs. [226, 233] to model the GW spectrum produced by

such a background. Simulations of the core-collapse process were utilized to estimate the potential ranges of the parameters for this model. Several of these simulations exhibited a low-frequency peak near 100 Hz which was not captured by the functional form of our model; as a result, I modeled this peak separately using a Gaussian function. I studied both of these models using the parameter estimation techniques detailed in Sec. 3.3 and projected the sensitivity of Advanced LIGO and the Einstein Telescope to them. I also compared the expected sensitivities on the low-frequency peak to GW waveforms produced from individual simulations of the core-collapse process and found that the resulting limits were 1–3 orders of magnitude above the peak amplitudes derived from fits to the simulations.

Finally, I combined the low- and high-frequency parts of the model into a single, six-parameter model, performed parameter estimation using `MultiNest`, and compared the resulting limits to those of the individual low- and high-frequency components. I also added an injection of a simulated SGWB based on this model and demonstrated that it was possible to accurately recover the parameters of the signal using the Einstein Telescope sensitivity.

Chapter 4

Searches for long-lasting gravitational-wave transients

In this chapter, we focus on the detection of long-lasting (or long-duration) transient gravitational waves. The defining characteristic of this class of GW signals is a duration lasting between a few seconds and a few weeks. Signals lasting longer than this may be better classified as continuous or persistent, while shorter signals may be classified as short-duration transients. This distinction is important because it drives the choice of algorithm used for analyzing the data and attempting to detect these signals.

While matched filtering [116, 119, 246, 247] is an effective technique in some searches for short-duration GW transients, it requires that the astrophysical process behind the emission of GWs is relatively well-understood. This is so that waveforms of simulated GW signals can be produced to form a template bank for comparison with the data. However, in the case of long-lasting GW transients, knowledge of the astrophysical sources and emission mechanisms is more speculative than for many other types of GWs.

The Stochastic Transient Analysis Multi-detector Pipeline (STAMP) is an adapted form of the stochastic radiometer search [138, 140, 164], and is focused on the detection of point sources of long-lasting GW transients. Similarly to the radiometer search, STAMP is based on the cross-correlation of data from multiple GW interferometers. However, STAMP searches are based around the construction of a spectrogram (or frequency-time map) and the use of pattern recognition algorithms to identify possible GW signals. Also,

STAMP analyses generally use a finer time resolution and coarser frequency binning compared to the stochastic radiometer, due to the difference in targeted signal duration and morphology.

The rest of this chapter will describe various aspects of STAMP and its objectives. In Sec. 4.1, I will discuss some possible sources of long GW transients. Sec. 4.2 will provide some of the details of the STAMP analysis framework and the cross-correlation estimator for GW power around which the pipeline is based. In Sec. 4.3, I develop and apply a flag for identifying and removing probable data artifacts in order to improve GW search sensitivity. I will construct a new clustering algorithm and describe its use for detecting GWs with STAMP in Sec. 4.4. Finally, Sec. 4.5 will discuss several adaptations I have made to STAMP for use in all-sky searches.

4.1 Astrophysical sources

There exist several astrophysical mechanisms which are expected to produce long-lasting gravitational-wave transients. Often, these transients are associated with stellar core-collapse or instabilities in a stellar remnant, like a neutron star or black hole. A case where the inspiral of a dynamically formed system of black holes may produce long-lasting GW transients will also be discussed.

In general, the detection of long-lasting GW transients holds the promise of revealing the astrophysical properties of the objects which produce them. Eventual detections will also allow us to better understand the populations of these sources, including the formation rates and scenarios for black holes and neutron stars. There is also significant evidence which relates long-lasting gamma-ray bursts and core-collapse supernovae [248]; detection of a long-transient GW coincident with a long GRB would provide great insight into the GRB mechanism, which is currently not well-understood.

4.1.1 Core-collapse supernovae

Several processes take place during a core-collapse supernova which may contribute to the overall GW signature of these events. These processes include rotation during the collapse and core bounce [212], neutrino-driven convection and standing-accretion-shock instability (SASI) [213–215], asymmetric neutrino emission [214–216], aspherical

outflows [217–221], and magnetic stresses [220, 221]. Depending on the stellar properties and the core-collapse mechanism, different processes may dominate over others [18, 249].

If an explosion occurs during the supernova process, a protoneutron star (PNS) is left behind and will cool on a Kelvin-Helmholtz timescale [250]. In certain scenarios, fallback accretion [251] or a late-time hadron-quark phase transition [252] may occur, causing the PNS to collapse and form a black hole. If the PNS survives, significant convection may develop due to thermal and lepton gradients, potentially lasting for tens of seconds and making it a source of long-lasting GW transients [216, 253–255]. Current models estimate a GW strain of $h \approx 3 \times 10^{-23}$ at a distance of 10 kpc for this process [255–257].

Several types of rotational instabilities may develop in PNSs in order to redistribute energy and momentum to reach a minimal energy state [18]. It is thought that most PNSs are born with a core in solid-body rotation and an outer region which is differentially rotating [258]; magneto-rotational instabilities [259] and hydrodynamic shear instabilities [260] will act to drive the system toward uniform rotation. These instabilities, especially the latter, may lead to significant deformation of the PNS and the emission of GWs [261]. There are also cases where a secular gravitational radiation-reaction instability may develop, depending on the ratio of rotational kinetic energy to gravitational energy. This may lead to a non-axisymmetric “bar-like” deformation and cause significant emission of GWs lasting for ≈ 10 –100 s [262, 263].

Particularly massive stars may collapse directly to a black hole either without an initial explosion or due to fallback accretion after a weak explosion. Gravitational waves may be emitted due to asymmetric mass motions within the black hole’s accretion disk, or due to asymmetric clumping of the disk. One model of a gravitational instability in the outer accretion disk leads to disk fragmentation due to helium photodisintegration cooling [264]. The inspiral of the fragments produces a GW signal lasting $\mathcal{O}(10$ –100 s) with strain amplitudes as high as $h \approx 10^{-22}$ at 100 Hz for a source at 100 Mpc [264]. Another model proposes a scenario where the central black hole and the accretion disk are linked by strong magnetic fields; spin-down of the black hole creates an instability in the disk which drives strong magnetoturbulence, resulting in a time-varying quadrupole moment and the emission of GWs [265–267]. This “accretion disk instability” model will be commonly referred to as the “ADI model” throughout the remainder of the text.

4.1.2 Neutron stars

Neutron stars may undergo a variety of phenomena which induce a time-varying quadrupole moment and the emission of gravitational waves. Pulsing neutron stars (pulsars) sometimes experience “glitches,” during which the rotational frequency increases over ≈ 2 minutes, followed by a relaxation period lasting on the order of weeks [268–270]. The underlying mechanism and physical process behind pulsar glitches is relatively unknown, although it is a subject of current research. Ref. [271] assumes that the emitted GW energy is comparable to the change in rotational energy and estimates that a glitch in a pulsar at 1 kpc could produce a GW signal lasting ≈ 10 s, with a strain of $h \approx 8 \times 10^{-24}$ at 3.8 kHz, although this is well below the Advanced LIGO detector noise at this frequency.

Seismic activity in the crust of highly-magnetized neutron stars is thought to be responsible for flares from soft-gamma repeaters and anomalous X-ray pulsars. There is expected to be an $\mathcal{O}(100 \text{ s})$ -long “ringdown” period of quasiperiodic oscillations after one of these flares, which is a potential source of long-lasting GW transients [272].

There are also several classes of quasinormal mode oscillations which may occur in neutron stars, including *r*-modes. *R*-modes are quasitoroidal oscillations which have the Coriolis force as their restoring force; these modes are unstable with respect to gravitational radiation-reaction and are expected to grow until reaching some saturation amplitude. *R*-modes are generally expected to be a source of very long-lasting GW signals (essentially continuous); however, in cases where the saturation amplitude is sufficiently high, the timescale of GW emission may be much shorter, producing a long-lasting GW transient signal [273].

4.1.3 Dynamically formed binary black holes

Dynamically formed binary systems comprising two black holes may be a potential source of long-lasting GW transients. These binary systems are expected to have highly eccentric orbits; although most binaries are expected to circularize during the inspiral process, if there is any residual eccentricity, these waveforms may not be fit sufficiently well by circular waveform template banks [274]. As the two black holes pass the point of closest approach in their orbit, small bursts of GWs are emitted at low frequencies; this has been observed in simulations of these systems [275] and provides a long-lasting,

albeit non-continuous signal. A specially-designed excess-power based search may be able to leverage this feature in order to detect such a signal once interferometric GW detectors become more sensitive at the low frequencies where these bursts occur.

4.2 Analysis framework

The STAMP analysis strategy for detecting long-duration GW transients is based around an excess-power statistic using the cross-correlation of data from two GW interferometers. We first cross-correlate the data to form this excess-power statistic (discussed further in Sec. 4.2.2) and then project it in the form of a frequency-time map (henceforth, ft -map; also known as a spectrogram). Then, pattern recognition algorithms are utilized to pick out groups of pixels in the ft -maps which are significant and may correspond to GW signals. ft -maps of other relevant statistics are also formulated for data quality and other analysis purposes.

4.2.1 A plane gravitational wave from a point source

The stochastic analysis framework described in Sec. 3.2 may be adapted by taking the isotropic plane wave description of Eq. 3.6 and using a delta function over sky position to consider a point source of GWs:

$$h_{ab}(t, \vec{x}) = \sum_A \int_{-\infty}^{\infty} df e_{ab}^A(\hat{\Omega}) \tilde{h}_A(f) e^{2\pi i f(t - \hat{\Omega} \cdot \vec{x}/c)}. \quad (4.1)$$

Following the remainder of Sec. 3.2.1, an expression for the point source GW time-series strain in detector I may be obtained by contracting Eq. 4.1 with the detector tensor (Eq. 3.11):

$$h_I(t) = \sum_A \int_{-\infty}^{\infty} df \tilde{h}_A(f) e^{2\pi i f(t + \hat{\Omega} \cdot \vec{x}/c)} F_I^A. \quad (4.2)$$

Using a discrete Fourier transform, the GW strain frequency spectrum in detector I is given by:

$$\tilde{h}_I(f) = \sum_A \tilde{h}_A(f) e^{-2\pi i f \hat{\Omega} \cdot \vec{x}/c} F_I^A(\hat{\Omega}). \quad (4.3)$$

A cross-correlation of these spectra from two detectors will result in a term with two Fourier amplitudes $\tilde{h}_A(f, \hat{\Omega})$; thus, we define the one-sided power spectrum of GW strain as

$$H_{AA'}(t; f) = 2 \left\langle \tilde{h}_A^*(t; f) \tilde{h}_{A'}(t; f) \right\rangle. \quad (4.4)$$

Assuming no correlations between the two GW polarizations, the sum over polarizations gives the total GW power:

$$H(t; f) \equiv \text{Tr}(H_{AA'}(t; f)). \quad (4.5)$$

This gives a single spectrum which is invariant under change of polarization bases.

4.2.2 Estimator for GW power

Taking the strain time-series $s_I(t)$ to be composed of gravitational-wave strain and noise strain (as in Eq. 3.15), we can define an estimator for GW power \hat{Y} as

$$\hat{Y}(t; f; \hat{\Omega}) \equiv 2 \text{Re} \left[\tilde{Q}_{IJ}(t; f; \hat{\Omega}) \tilde{s}_I^*(t; f) \tilde{s}_J(t; f) \right]. \quad (4.6)$$

Here, \tilde{s}_I is the Fourier transform of the strain time-series in detector I , t is the start time of the segment for which the Fourier transform is calculated, f is the frequency bin, $\hat{\Omega}$ is the sky direction, and \tilde{Q}_{IJ} is some to-be-determined filter function. For an appropriately defined \tilde{Q}_{IJ} , the estimator \hat{Y} should be real; thus, we consider only the real part of the estimator.

Assuming that the noise components in different detectors are uncorrelated (and are not correlated with the GW strain), the expectation value of this estimator is

$$\left\langle \hat{Y}(t; f; \hat{\Omega}) \right\rangle = 2 \text{Re} \left[\tilde{Q}_{IJ}(t; f; \hat{\Omega}) \left\langle \tilde{h}_I^*(t; f) \tilde{h}_J(t; f) \right\rangle \right]. \quad (4.7)$$

Using Eqs. 4.3 and 4.4 gives

$$\left\langle \hat{Y}(t; f; \hat{\Omega}) \right\rangle = 2 \text{Re} \left[\tilde{Q}_{IJ}(t; f; \hat{\Omega}) \sum_{AA'} \frac{1}{2} H_{AA'}(t; f) e^{-2\pi i f \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c} F_I^A(t; \hat{\Omega}) F_J^{A'}(t; \hat{\Omega}) \right], \quad (4.8)$$

where $\Delta\vec{x}_{IJ} \equiv \vec{x}_I - \vec{x}_J$ is the vector pointing from detector J to detector I . For unpolarized sources, we take $H_{AA'}(t; f) = \frac{1}{2}H(t; f)\delta_{AA'}$ [18], which gives

$$\langle \hat{Y}(t; f; \hat{\Omega}) \rangle = 2 \operatorname{Re} \left[\tilde{Q}_{IJ}(t; f; \hat{\Omega}) H(t; f) e^{-2\pi i f \hat{\Omega} \cdot \Delta\vec{x}_{IJ}/c} \sum_A F_I^A(t; \hat{\Omega}) F_J^A(t; \hat{\Omega}) \right]. \quad (4.9)$$

Now, we may choose the filter function \tilde{Q}_{IJ} to be

$$\tilde{Q}_{IJ} = \frac{e^{2\pi i f \hat{\Omega} \cdot \Delta\vec{x}_{IJ}/c}}{\epsilon_{IJ}(t; \hat{\Omega})}, \quad (4.10)$$

where $\epsilon_{IJ}(t; \hat{\Omega})$ is the unpolarized pair efficiency, given by

$$\epsilon_{IJ}(t; \hat{\Omega}) \equiv \frac{1}{2} \sum_A F_I^A(t; \hat{\Omega}) F_J^A(t; \hat{\Omega}). \quad (4.11)$$

With this definition for \tilde{Q}_{IJ} , we obtain $\langle \hat{Y} \rangle = H(t; f)$; thus, \hat{Y} is an estimator for GW power.

The variance of \hat{Y} is defined as

$$\sigma_Y^2 \equiv \langle \hat{Y}^2(t; f; \hat{\Omega}) \rangle - \langle \hat{Y}(t; f; \hat{\Omega}) \rangle^2, \quad (4.12)$$

and can be estimated as (see [18, 153])

$$\hat{\sigma}_Y^2 = \frac{1}{2} \left(N_I(t; f) N_J(t; f) \left| \tilde{Q}_{IJ}(t; f; \hat{\Omega}) \right|^2 \right) \quad (4.13)$$

in the small-signal limit, where $N_I(t; f)$ is the one-sided noise power spectrum for detector I .

The power spectral density (or ‘‘auto-power’’) in time segment t and frequency bin f is defined as

$$\hat{P}_I(t; f) \equiv 2 |\tilde{s}_I(t; f)|^2. \quad (4.14)$$

However, this includes contributions from both signal and noise. In practice, we estimate $N_I(t; f)$ using an average of the power spectra in time segments adjacent to t ; henceforth, this estimator will be denoted as $\hat{P}_I^{\text{adj}}(t; f)$. Using Eq. 4.14, we can calculate \hat{P}_I^{adj} in

time segment t_0 and frequency bin f :

$$\hat{P}_I^{\text{adj}}(t_0; f) = \frac{1}{2n} \left[\left(\sum_{t=t_0-n\delta t}^{t=t_0+n\delta t} \hat{P}_I(t; f) \right) - \hat{P}_I(t_0, f) \right]. \quad (4.15)$$

Here n is a positive integer, usually taken to be between 4 and 8, and δt is the time resolution of the ft -map. This estimator is defined so that for many GW signal morphologies, \hat{P}_I^{adj} will not contain contributions from the signal. However, in some cases, especially monochromatic signals, \hat{P}_I^{adj} will contain signal contributions; this will increase the noise estimate and reduce the overall SNR, making STAMP less sensitive to these types of signals. We note that future work with STAMP may include new definitions of \hat{P}_I^{adj} , possibly using the median rather than the mean, or including nearby frequency bins in the estimate rather than just adjacent time segments in the same frequency bin.

Using Eqs. 4.12 and 4.15, the estimator for $\hat{\sigma}_Y^2$ is calculated as

$$\hat{\sigma}_Y^2 = \frac{1}{2} \left| \tilde{Q}_{IJ}(t; f; \hat{\Omega}) \right|^2 \hat{P}_I^{\text{adj}}(f) \hat{P}_J^{\text{adj}}(f). \quad (4.16)$$

For analysis purposes, we typically look at the signal-to-noise ratio (SNR) of \hat{Y} , defined as

$$\text{SNR}(t; f; \hat{\Omega}) \equiv \frac{\hat{Y}(t; f; \hat{\Omega})}{\hat{\sigma}_Y(t; f; \hat{\Omega})}. \quad (4.17)$$

The theoretical distribution of the SNR statistic was calculated in [276]; it is given by

$$f_{\text{SNR}}(x) = \frac{N^{2N}}{2^{2N-1} \xi^{4N+2} \Gamma^2(N)} \int_0^\infty |s| e^{-|sx|} K_0 \left(\frac{Ns}{\xi^2} \right) s^{2N-1} ds, \quad (4.18)$$

where N is the number of adjacent segments used to estimate \hat{P}^{adj} , K_0 is a modified Bessel function of the second kind, and ξ is a factor which accounts for the effect of windowing on the variance of the data. In Fig. 4.1, we show the distribution of SNR for real time-shifted LIGO data from the S5 science run, and for Gaussian noise.

After constructing an ft -map of SNR, pattern recognition algorithms are used to




Figure 4.1: Distribution of STAMP SNR statistic in real time-shifted LIGO data (blue), Gaussian noise colored with the initial LIGO sensitivity curve (red), and the theoretical distribution (black). The ft -map pixel size is $1 \text{ s} \times 1 \text{ Hz}$. We have used $N = 17$, and ξ is calculated empirically as 0.847. The real data matches Gaussian noise relatively well, but some deviations are apparent beyond $\text{SNR} \approx 5$. Deviations from the theoretical model are likely due to signal processing effects, including high-pass filtering and zero-padding.

identify significant groupings of pixels. The resulting clusters are ranked using a multi-pixel SNR statistic, SNR_Γ :

$$\text{SNR}_\Gamma \equiv \frac{\sum_{(t,f) \in \Gamma} \hat{Y}(t; f; \hat{\Omega}) \hat{\sigma}_Y^{-2}(t; f; \hat{\Omega})}{\left(\sum_{(t,f) \in \Gamma} \hat{\sigma}_Y^{-2}(t; f; \hat{\Omega}) \right)^{1/2}}. \quad (4.19)$$

Here, the summations run over pixels in the cluster, represented by Γ . This can be shown to be a minimum-variance estimator for the GW power in the cluster Γ , assuming no correlations between pixels in \hat{Y} [18].

4.2.3 Frequency-time maps

In this subsection, we show some STAMP ft -maps for illustrative purposes. Fig. 4.2 shows SNR, \hat{Y} , and $\hat{\sigma}_Y$ ft -maps for Gaussian noise and for Gaussian noise plus a simulated ADI signal. The signal is visible as a bright, downward-curving track in the SNR and \hat{Y} ft -maps. In the $\hat{\sigma}_Y$ ft -map, it appears as several closely-spaced tracks due to the use of adjacent segments in estimating \hat{P}^{adj} , which is used to calculate $\hat{\sigma}_Y$. The frequency variation in \hat{Y} and $\hat{\sigma}_Y$ is due to the inherent noise spectrum of the LIGO detector, which has its minimum between ≈ 100 – 200 Hz.

Fig. 4.3 shows the auto-power and adjacent auto-power ft -maps for both detectors. Again, the simulated signal is visible as a single bright track in the \hat{P}_{H1} and \hat{P}_{L1} ft -maps, while it appears as several closely-spaced tracks in the $\hat{P}_{\text{H1}}^{\text{adj}}$ and $\hat{P}_{\text{L1}}^{\text{adj}}$ ft -maps due to the use of adjacent segments.

4.2.4 Pattern recognition

One of the most important and interesting aspects of a STAMP search is the choice of pattern recognition algorithm. This often depends on the morphology of the GW signal of interest, as well as the computational resources available for the search.

The first pattern recognition algorithm developed for use with STAMP was the box search. This algorithm is based around constructing rectangles in frequency-time space and calculating the SNR_Γ statistic (see Eq. 4.19) using the pixels in each rectangle. Typically, the boxes are tiled to attain coverage of the entire ft -map. This algorithm is useful for broadband signals and for relatively monochromatic signals.

The Radon algorithm has been used to search for line-like tracks in STAMP ft -maps. It uses a Radon transform to integrate SNR along lines with different intercepts and different slopes [278]. A linear ft -map track containing significant power will be visible as a point in Radon space. The Hough algorithm extends the idea of the Radon transform to polynomial tracks of arbitrarily high order [279].

The Locust algorithm is a gradient-driven wandering algorithm which uses local maxima to chain pixels together into traces. It is generally sensitive to signals of arbitrary shape; however, its dependence on the presence of local maxima and its relatively high computational demands present some significant drawbacks.




Figure 4.2: STAMP ft -maps for simulated Gaussian noise (left column) and Gaussian noise plus a simulated GW signal from an accretion disk instability [265–267, 277] (right column). The simulated Gaussian noise is colored with the initial LIGO noise curve. Top row: SNR ft -maps. Middle row: \hat{Y} ft -maps. Bottom row: $\hat{\sigma}_Y$ ft -maps. On the right, the $\hat{\sigma}_Y$ estimate in the adjacent time segments is contaminated by the GW signal.




Figure 4.3: STAMP ft -maps for simulated Gaussian noise (left column) and Gaussian noise plus a simulated GW signal from an accretion disk instability [265–267,277] (right column). The simulated Gaussian noise is multiplied by the initial LIGO noise curve. Top row: \hat{P}_{H1} ft -maps. Second row: $\hat{P}_{H1}^{\text{adj}}$ ft -maps. Third row: \hat{P}_{L1} ft -maps. Bottom row: $\hat{P}_{L1}^{\text{adj}}$ ft -maps. The use of adjacent segments in the calculation of \hat{P}^{adj} is apparent.

The `BurstCluster` algorithm was first developed for use in the flare pipeline of the LIGO burst group [280] and was adapted to STAMP by S. Kandhasamy. It is a density-based clustering algorithm that connects different regions of the ft -map based on the relative density of pixels which exceed a given threshold on SNR (seeds). This algorithm is most effective at identifying “blob-like” signals which exhibit high densities of above-threshold pixels. Narrowband, “track-like” signals do not usually have high enough seed pixel density to be identified as significant by `BurstCluster`.

In Sec. 4.4, I will discuss the development of a new seed-based clustering algorithm for use in STAMP analyses.

4.3 Data quality

The LIGO detectors are often affected by noise transients due to environmental effects or instrumental problems, including mechanical vibrations, electromagnetic disturbances, power shorts, misalignment of optical cavities, and more. These generate data artifacts, or “glitches,” which pose a problem for analyses of gravitational-wave data. Glitches may diminish the sensitivity of a GW search by raising the threshold required for a trigger to be statistically significant. They may also be erroneously identified as GW candidates, since they often contain significant power. As a result, it is important to identify and remove glitches from the data in order to improve search sensitivity and increase our confidence in GW candidates which do not resemble glitches.

An example glitch is shown in Fig. 4.4; the glitch is clearly visible as a single bright column in the $\hat{P}_{\text{H1}} ft$ -map and as a sequence of alternating bright columns in the $\hat{P}_{\text{H1}}^{\text{adj}} ft$ -map. However, since glitches are caused by local effects, there is no significant structure in the L1 auto-power ft -maps. This still leads to significant cross-power, as the glitch can be seen in the SNR ft -map as a single column of black and white pixels at $t \approx 22$ s.

Several algorithms have been devised to identify periods of non-stationary noise in LIGO data in order to remove glitches [281–288]. These periods are marked with LIGO data quality flags so that they can be discarded [287–290]. However, analyses of LIGO data with STAMP have indicated that many glitches remain in the data even with the use of these flags.

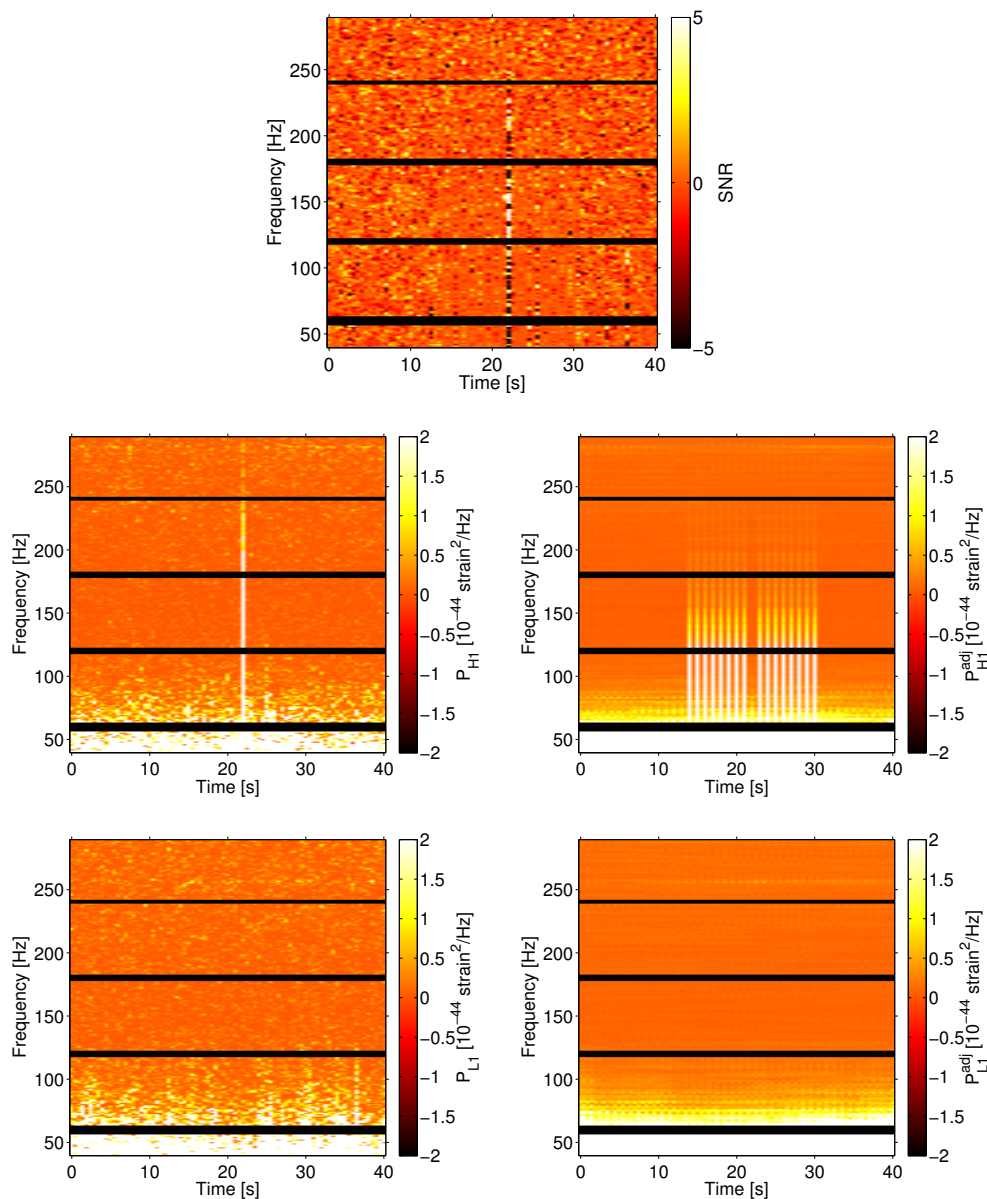


Figure 4.4: Ft -maps of a glitch in H1 at $t \approx 22$ s. Top: SNR ft -map. Middle row: \hat{P}_{H1} ft -map (left) and $\hat{P}_{\text{H1}}^{\text{adj}}$ ft -map (right). Bottom row: \hat{P}_{L1} ft -map (left) and $\hat{P}_{\text{L1}}^{\text{adj}}$ ft -map (right). The glitch is visible as a column of black and white pixels in the SNR ft -map and a single bright column in the \hat{P}_{H1} ft -map. It appears as a sequence of alternating bright columns in the $\hat{P}_{\text{H1}}^{\text{adj}}$ ft -map due to the use of adjacent segments in estimating the noise auto-power. No structure is visible in the L1 ft -maps since the glitch is local to the H1 interferometer. Horizontal black lines are frequencies which have been removed due to the presence of contamination from power line harmonics.

To combat this issue, we have developed a glitch-identification algorithm which acts essentially as a consistency check between two GW interferometers [291]. This algorithm compares the auto-power in each detector and checks whether it is consistent with our expectations for a real GW signal; if not, it is marked as a potential glitch. This work expands on [284], which demonstrates how physical environmental monitoring channels can be used to identify long-lasting noise transients; here we are primarily interested in identifying short noise transients in the presence of long-lasting GW signals. Along with this, our algorithm only utilizes GW strain channels to identify noise transients, rather than auxiliary environmental channels. Note that this algorithm checks the consistency of raw data segments before any clustering procedure is applied, as opposed to checking the consistency of GW triggers generated by an analysis. This section will describe this glitch-identification algorithm in detail.

4.3.1 Formalism

The goal of this glitch identification algorithm is to discriminate between significant power deposited in STAMP ft -maps due to GW signals and due to glitches. As discussed in Sec. 4.2.2, we use the average auto-power in adjacent segments (\hat{P}_I^{adj}) as an estimate of the auto-power due to detector noise alone. Thus, the difference between \hat{P}_I and \hat{P}_I^{adj} will give an estimate for the auto-power due to GWs in detector I :

$$\hat{P}_I^{\text{GW}}(t; f) = \frac{\hat{P}_I(t; f) - \hat{P}_I^{\text{adj}}(t; f)}{\epsilon_{II}}. \quad (4.20)$$

Here, ϵ_{II} represents the antenna factor for detector I ; it accounts for the direction-dependent sensitivity of the interferometer (see 4.11). Using this normalization, we can make a comparison between the GW auto-power in two GW detectors. The quantity $\hat{\Xi}$ is defined as an estimator for the GW auto-power difference between detectors I and J :

$$\hat{\Xi}(t; f) \equiv \frac{\hat{P}_I(t; f) - \hat{P}_I^{\text{adj}}(t; f)}{\epsilon_{II}} - \frac{\hat{P}_J(t; f) - \hat{P}_J^{\text{adj}}(t; f)}{\epsilon_{JJ}}. \quad (4.21)$$

By construction, $\langle \hat{\Xi}(t; f) \rangle = 0$ for well-behaved noise and for GW signals.

We would like to normalize $\hat{\Xi}(t; f)$ to a unitless quantity with near-unity variance.

The variance of $\hat{\Xi}$ can be calculated as

$$\sigma_{\Xi}^2(t; f) \equiv \left\langle \left(\frac{\hat{P}_I(t; f) - \hat{P}_I^{\text{adj}}(t; f)}{\epsilon_{II}} - \frac{\hat{P}_J(t; f) - \hat{P}_J^{\text{adj}}(t; f)}{\epsilon_{JJ}} \right)^2 \right\rangle - \left\langle \frac{\hat{P}_I(t; f) - \hat{P}_I^{\text{adj}}(t; f)}{\epsilon_{II}} - \frac{\hat{P}_J(t; f) - \hat{P}_J^{\text{adj}}(t; f)}{\epsilon_{JJ}} \right\rangle^2. \quad (4.22)$$

This results in terms with four-point correlations like $\langle \hat{P}_I \hat{P}_J^{\text{adj}} \rangle = \langle s_I^* s_I s_J^{*, \text{adj}} s_J^{\text{adj}} \rangle$. By calculating the four-point correlations, we obtain an estimator for the variance:

$$\sigma_{\Xi}^2(t; f) \equiv \frac{\hat{P}_I^2(t; f) - \hat{P}_I^{2, \text{adj}}(t; f)}{\epsilon_{II}^2} + \frac{\hat{P}_J^2(t; f) - \hat{P}_J^{2, \text{adj}}(t; f)}{\epsilon_{JJ}^2} - \frac{2\epsilon_{IJ}^2}{\epsilon_{II}\epsilon_{JJ}} |\hat{Y}(t; f)|^2. \quad (4.23)$$

This allows us to construct a unitless signal-to-noise ratio for the $\hat{\Xi}$ statistic, SNR_{Ξ} :

$$\text{SNR}_{\Xi} \equiv \frac{\hat{\Xi}}{\hat{\sigma}_{\Xi}}. \quad (4.24)$$

Considering Eqs. 4.21, 4.22, and 4.24 qualitatively, the behavior of SNR_{Ξ} should be different for real GW signals ($\hat{Y}(t; f) > 0$) and glitches ($\hat{P}_I(t; f) > \hat{P}_J(t; f)$ or vice versa). Loud glitches in detectors I and J should cause $\text{SNR}_{\Xi} \approx \pm 1$, respectively, surrounded by columns of $\text{SNR}_{\Xi} \mp 1$, respectively (adjacent segments are affected by our noise estimation technique). Real GW signals, however, should have $\text{SNR}_{\Xi} \approx 0$, surrounded by $\text{SNR}_{\Xi} \approx 0$, although there should be slightly larger fluctuations in the time segments adjacent to the signal. This behavior is qualitatively illustrated in Fig. 4.5, which shows SNR and SNR_{Ξ} ft -maps for a glitch, well-behaved noise, and a simulated GW signal.

4.3.2 Algorithm design

To identify glitches, it is clear that we want to look for ft -map columns which contain an excess of pixels with $|\text{SNR}_{\Xi}| \approx 1$. We define $\mathcal{F}_+(t)$ and $\mathcal{F}_-(t)$ as the fraction of pixels in time segment t with $0.95 < \text{SNR}_{\Xi} < 1.05$ and $-1.05 < \text{SNR}_{\Xi} < -0.95$, respectively. Based on the definition of SNR_{Ξ} (Eq. 4.24), \mathcal{F}_+ helps to identify glitches in detector I , while \mathcal{F}_- identifies glitches in detector J . In Fig. 4.6, I show the distribution of SNR_{Ξ} for time-shifted LIGO data, Gaussian noise, glitches, and simulated GW signals; this




Figure 4.5: Ft -maps of SNR (left column) and SNR_Ξ (right column) for time-shifted LIGO data. Top row: a likely glitch; the adjacent columns of $\text{SNR}_\Xi \approx +1$ occur due to our method of calculating \hat{P}^{adj} . Middle row: nearly stationary noise. Bottom row: stationary noise plus a relatively strong simulated GW signal. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

helps to illustrate the choice of parameters for the \mathcal{F} thresholds.

However, using only \mathcal{F}_+ and \mathcal{F}_- to define this glitch-identification flag would result in the loss of significant amounts of non-glitchy data. As seen in the top right panel of Fig. 4.5, the use of adjacent segments for our estimates for the noise auto-power, \hat{P}^{adj} , causes $|\text{SNR}_{\Xi}| \approx 1$ in ft -map columns adjacent to the glitch. In order to discriminate between actual glitches and columns adjacent to glitches, I have implemented an ‘‘auto-power stationarity ratio’’:

$$R_I(t) = \frac{1}{N_f} \sum_f \frac{\hat{P}_I(t; f)}{\hat{P}_I^{\text{adj}}(t; f)}. \quad (4.25)$$

Here, I denotes the detector; this statistic is essentially a signal-to-noise ratio for the auto-power statistic, averaged over all frequency bins in each time segment.

Using the \mathcal{F} and R statistics, the glitch identification algorithm flags ft -map columns as *glitch-like* if either of the following criteria are satisfied:

$$\mathcal{F}_+(t) > 0.027 \text{ and } R_I(t) > 2 \text{ and } R_J(t) < 2 \quad (4.26)$$

$$\mathcal{F}_-(t) > 0.027 \text{ and } R_J(t) > 2 \text{ and } R_I(t) < 2 \quad (4.27)$$

The final requirement in each case ensures that when there is significant power in *both* detectors (i.e., $R_I(t) > 2$ and $R_J(t) > 2$), the time segments will not be flagged. Although real GW signals should not fail the \mathcal{F} portion of the flag, this requirement helps to ensure that signals will not be flagged since they will deposit power in both detectors. The parameter values given were primarily chosen to optimize the efficiency of this algorithm at flagging glitches, while still assuring the safety of GW signals. Fig. 4.7 shows scatter plots of our glitch identification parameters $R_I(t)$ and \mathcal{F}_+ for simulated ADI signals (left) and time-shifted data and Gaussian noise (right). The upper right-hand quadrant contains cases where both components of the glitch flag are satisfied and the corresponding time segments would be flagged as glitch-like (neglecting the safety feature mentioned previously).

4.3.3 Background study

To determine the effectiveness of this flag at identifying glitches, I have analyzed time-shifted LIGO S5 data and Gaussian noise, both with and without the flag applied.

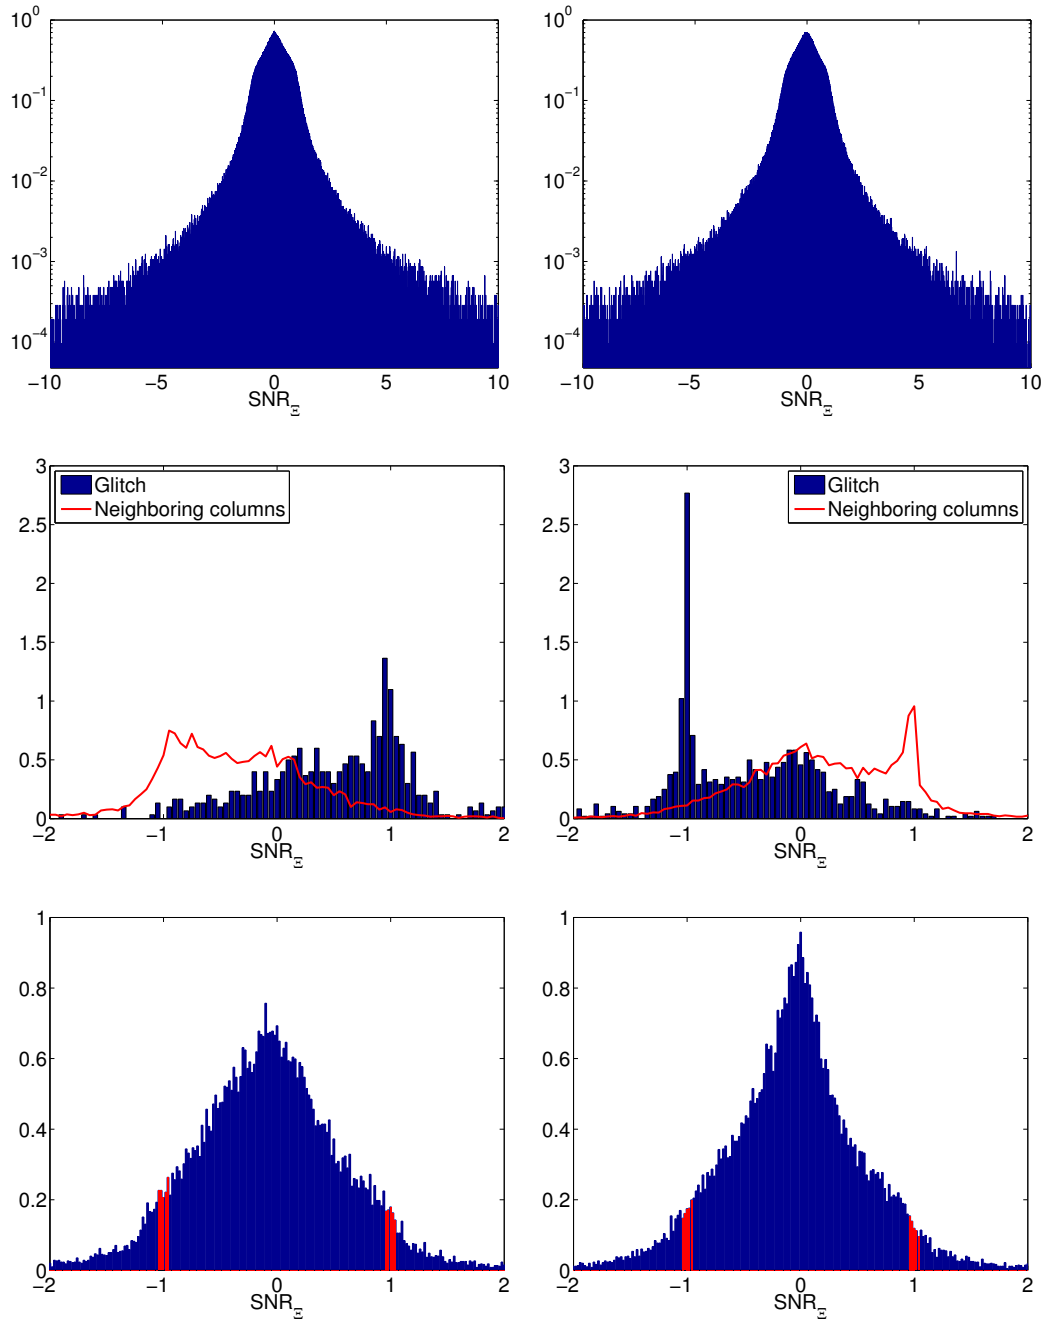


Figure 4.6: Histograms of SNR_Ξ ft -map pixels. Top row: 2200 s of clean time-shifted LIGO S5 data (left) and Gaussian noise (right). Middle row: 2 s of data containing glitches in H1 (left) and L1 (right). Bottom row: 40 s-long ADI injections at 30 Mpc (left) and at 5 Mpc (right). Red bars indicate $0.95 < \text{SNR}_\Xi < 1.05$. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.




Figure 4.7: Scatter plots of injected ADI signals (left), time-shifted LIGO S5 data (right red \times 's), and Gaussian noise (right green \circ 's) in the plane of our glitch identification parameters $R_I(t)$ and \mathcal{F}_+ . $R_J(t)$ and \mathcal{F}_- are not shown here due to overall qualitative similarities with the displayed parameters. Injection distances range from 5 – 40 Mpc with smaller distances corresponding to redder markers. The glitch identification thresholds for each parameter are represented by solid black lines with the upper right quadrant corresponding to the glitch-likely case. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

The `BurstCluster` algorithm [280] was used to identify clusters of pixels containing significant power in $120\text{ s} \times 150\text{ Hz}$ ft -maps (100 – 250 Hz) with $4\text{ s} \times 0.25\text{ Hz}$ pixels. In Fig. 4.8, I show plots of p -value (false alarm probability) versus SNR_Γ for time-shifted data and Gaussian noise with and without the glitch-likely flag. The results indicate a significantly better agreement between time-shifted data and Gaussian noise when the glitch-likely flag is applied. The largest SNR_Γ due to a background noise trigger is reduced by two orders of magnitude with the use of the glitch-likely flag; this significantly increases the ability of the pipeline to identify real GW signals.

Having demonstrated our glitch-identification algorithm for the case of $4\text{ s} \times 0.25\text{ Hz}$ pixels in the 100–250 Hz band, we look at a few other special cases. An exhaustive exploration of the domain of utility of our algorithm is beyond the scope of this discussion (and is probably not computationally feasible); thus, we examine cases which we expect to be of primary interest for future STAMP analyses. In the top-left panel of Fig. 4.9,




Figure 4.8: Plot of p -value versus SNR_Γ using a density-based clustering algorithm to identify groupings of pixels with significant power for time-shifted LIGO data (TS) and Gaussian Monte Carlo noise (MC), with and without the use of our glitch identification algorithm. The dashed cyan curve (MC with cut) overlaps directly with the green curve (MC) because the glitch identification algorithm does not have any effect on the MC data. For real time-shifted data, the glitch identification algorithm improves the sensitivity at $p = 0.1\%$ by several orders of magnitude. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

I show p -value versus SNR_Γ for $1\text{ s} \times 1\text{ Hz}$ pixels in the 100–250 Hz range. This demonstrates that this glitch-identification flag can effectively clean even shorter segments of data so as to achieve relatively good agreement with Gaussian noise.

The top-right panel of Fig. 4.9 shows p -value versus SNR_Γ for $4\text{ s} \times 0.25\text{ Hz}$ in the 375–525 Hz band. Again, with the use of this algorithm, we achieve relatively good agreement between cleaned time-shifted data and Gaussian noise. This is to be expected due to the relative stationarity of the detector noise (primarily shot noise) at high frequencies.

The bottom-left panel in Fig. 4.9 shows p -value versus SNR_Γ for $4\text{ s} \times 0.25\text{ Hz}$ pixels in the 40–100 Hz frequency band. While the agreement between time-shifted data and Gaussian noise is improved with the use of the glitch-identification algorithm, significant disagreement remains due to non-stationary noise, which is more common at lower frequencies in the LIGO detectors. This is illustrated in the bottom-right panel of Fig. 4.9, which shows an SNR_Ξ ft -map from a period of low-frequency noise. The

noise in this case appears to be due to quasi-continuous broadband noise rather than infrequent glitches. This could potentially be mitigated by the use of additional vetoes utilizing physical environmental monitors, including microphones and seismometers.

Not only should this flag be effective at identifying glitches, it should also rarely flag well-behaved noise as glitch-like. Using simulated Gaussian noise colored with the initial LIGO noise curve, I have estimated the false glitch rate of this flag to be less than $1 \times 10^{-3} \text{ day}^{-1}$ in the frequency range of 100–250 Hz [291].

4.3.4 Safety for GW signals

The final aspect which must be considered for this glitch-identification algorithm is its *safeness*: the probability that it falsely identifies a time segment containing a GW signal as glitch-like. I have tested the safeness of this flag by adding simulated GW signals (based on the ADI model) to Gaussian noise. The signals are simulated in 312 different sky directions, with 20 different noise realizations for each direction. This is done for sources at distances between 1–50 Mpc. For each injection, I recorded the fraction of time segments satisfying the $R(t)$ criterion in either detector, the $F_+(t)$ or $F_-(t)$ criteria, and the relevant combination of both flags which corresponds to a glitch-like segment.

The results of this study can be seen in Fig. 4.10. For realistic and marginally-detectable signals ($d > 20$ Mpc), the fraction of flagged segments is negligible. When the signal is very loud ($d \leq 5$ Mpc), the fraction of flagged segments becomes as high as $\approx 3\%$; however, this is not a matter of significant concern since these signals are unrealistically loud and it is not expected that their sources will be present at such close distances.

I also considered a shorter ($t \approx 1$ s) GW signal produced by accretion disk fragmentation [264, 277]. In this model, an accretion disk forms clumps due to helium photodisintegration; as the clumps inspiral into the central black hole, a chirp-like GW signal is produced. For typical pixel sizes used in current STAMP analyses, this signal will be fully contained in only 1 or 2 ft -map columns, making it morphologically similar to a glitch. Ft -maps of SNR and SNR_{Ξ} for a simulated accretion disk fragmentation signal can be seen in Fig. 4.11.

As with the ADI signals, I have varied the sky location and distance, using multiple realizations of noise. For a very loud signal at 1 Mpc, the false dismissal probability is




Figure 4.9: Top-left: p -value versus SNR_Γ for the 100–250 Hz band using $1\text{ s} \times 1\text{ Hz}$ pixels; the relatively good agreement between Gaussian Monte Carlo (MC) and time-shifted data (TS) suggests that even short segments of cross-correlated data can be cleaned with our glitch identification flag. Top-right: plot of p -value versus SNR_Γ for the 375–525 Hz range with $4\text{ s} \times 0.25\text{ Hz}$ pixels. This frequency band exhibits good agreement between Monte Carlo and time-shifted data due to the nearly stationary noise associated with higher frequencies. Bottom-left: p -value versus SNR_Γ for the 40–100 Hz band using $4\text{ s} \times 0.25\text{ Hz}$ pixels. While the glitch-identification flag dramatically improves the agreement between Monte Carlo and time-shifted data, significant disagreement remains, possibly due to non-stationary noise associated with this band. Bottom-right: an ft -map of SNR_Ξ for time-shifted LIGO S5 data demonstrating the non-stationary noise sometimes associated with low frequencies. Note that the MC with cut curve (dashed cyan) overlaps with the MC curve (solid green) in the top row and the bottom-left panel. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.




Figure 4.10: Safety study for simulated ADI signals, using 312 sky directions, 20 noise realizations per sky direction, and source distances from 1 – 50 Mpc. The y -axis shows the fraction of time segments containing a GW signal which were flagged by the two individual components of the glitch flag, and by the full flag. The fraction of segments flagged as glitch-like decreases at distances $\lesssim 5$ Mpc because both detectors exceed the threshold on $R(t)$, triggering the safety feature and preventing the signal from being flagged as glitch-like. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

21%; however, at $d > 2.4$ Mpc, it drops below 1%. Thus, the results are similar to the ADI case: unrealistically loud signals may be frequently flagged, while more realistic signals have a very low false dismissal probability.

Finally, I have demonstrated that this glitch-identification flag can be used to more accurately reconstruct a GW signal by removing one or more glitchy time segments. Fig. 4.12 shows a 700 s-long ADI signal which has been injected into a period of time-shifted data containing a glitch (visible as a vertical column at $t \approx 490$ s). I have used the `BurstCluster` algorithm to recover the signal both without (bottom-left) and with (bottom-right) the glitch identification algorithm applied. This demonstrates that it is possible to improve the reconstruction of our signal by removing glitches, and that it is possible to recover a long-lasting GW transient which is disrupted by a glitch.




Figure 4.11: Ft -maps of SNR (left) and SNR_Ξ (right) for a ≈ 1 s accretion disk fragmentation waveform at 1 Mpc injected into Gaussian noise. The signal is contained in only two ft -map columns and spans a large range of frequencies, making it morphologically similar to a glitch. The pixel size is $1 \text{ s} \times 1 \text{ Hz}$. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

4.3.5 Comparison with LIGO data quality flags

The LSC’s detector characterization group has devised numerous data quality flags which can be used to determine when the interferometer’s strain channel is contaminated or corrupted by environmental or subsystem noise [287–290]. It is important to understand how this glitch-identification algorithm compares to these flags and whether it provides complementary information.

During the S5 science run, LIGO data quality flags were divided into four categories, in order of severity. Category 1 flags indicate periods of data which are severely contaminated by known and identified processes. Data marked with Category 2 flags may be analyzed, but various vetoes may be applied in post-processing of GW candidate triggers. Category 3 and 4 flags are advisory flags which exert caution in the case of a detection candidate. Detailed descriptions of these flags can be found in Refs. [281, 288].

I have analyzed time-shifted LIGO data, using several combinations of LIGO data quality flags to remove contaminated data, and compared it to the case where no data quality flags are used, but the glitch-identification algorithm is applied. Fig. 4.13 shows p -value versus SNR_Γ for all of these cases.




Figure 4.12: Ft -maps of an ADI signal injected into LIGO time-shifted data containing a glitch (at $t \approx 490$ s). Top-left: SNR ft -map. Top-right: SNR_Ξ ft -map. Bottom row: recovery without (left) and with (right) the use of the glitch-identification algorithm. The glitch-identification algorithm excludes the glitch from the reconstructed track, improving the recovery of the signal. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

The results indicate that our SNR_Ξ -based algorithm removes a large number of glitches which are not already identified by the LIGO data quality flags. Clearly, the two types of flags are complementary: the LIGO data quality flags identify and characterize specific environmental and instrumental disturbances, while our glitch-identification flag finds inconsistencies in auto-power between detectors.




Figure 4.13: P -value versus SNR_Γ for $1\text{ s} \times 1\text{ Hz}$ pixels of time-shifted (TS) LIGO S5 data with no flags applied (solid blue), with the SNR_Ξ -based flag applied (dashed red), and with the LIGO data quality flags applied in succession (CAT0 representing no flags applied, CAT1 representing the application of Category 1 flags, CAT2 representing the application of Category 1 and 2 flags, etc.). The data are parsed in $120\text{ s} \times 150\text{ Hz}$ ft -maps. Figure reproduced from Ref. [291], copyright IOP Publishing, reproduced with permission, all rights reserved.

4.3.6 Conclusions

In this section, we have developed and tested a glitch-identification flag based on checking auto-power consistency between detectors. It has proven to be effective at identifying glitches and cleaning time-shifted LIGO data so that it approaches Gaussian noise, with minimal losses in data and livetime. This flag is safe for realistic signals and, although it flags some time segments containing very loud GW signals, it does not do so at a high rate.

The glitch-identification algorithm as presented uses the entire frequency band to calculate the R and \mathcal{F} statistics. For analyses using large frequency bands, glitches may be “diluted” due to the sheer number of pixels in each time segment, and not exceed the thresholds set in Eqs. 4.26 and 4.27. Thus, it may be useful to calculate the glitch statistics in multiple frequency bands in order to identify glitches which are relatively localized in frequency. It may also be necessary to tune the glitch-likely thresholds for

different analyses, depending on the frequency range used, the GW signals of interest, and the overall characteristics of the data.

Finally, we note that this glitch-identification algorithm has been used successfully in two STAMP-based analyses of LIGO data: a search for GWs coincident with long gamma-ray bursts [141], and an all-sky search for long-lasting GW transients [142]. Although this algorithm is not inherently safe for cases where the source direction is not known (since the antenna factors ϵ_{II} and ϵ_{JJ} will be incorrect), it has proven effective and relatively safe for narrowband GW signals in an all-sky search (more details in Chap. 5).

4.4 Seed-based clustering

As mentioned previously, the choice of pattern recognition algorithm is an important decision for performing an analysis with STAMP. However, the GW signal morphology is not always known *a priori*, which makes it difficult to choose the most effective algorithm. Along with this, many of these algorithms have significant computational demands, which must be accounted for when analyzing large datasets which may span hundreds of days or more. In this section, I will develop a new seed-based clustering algorithm which can be used in STAMP searches to identify arbitrarily-shaped groups of pixels containing significant power and potentially corresponding to a GW signal. This algorithm has been used in both of the STAMP analyses published by the LSC to date [141, 142].

4.4.1 Algorithm description

The *burstegard* clustering algorithm [292] is designed to find GW signals of arbitrary shape in STAMP *ft*-maps; typically, this algorithm is applied to the SNR *ft*-map, but it could be extended to *ft*-maps of other statistics in the future. The algorithm depends on five user-defined parameters:

- *Pixel threshold*: only pixels which exceed this threshold on SNR will be involved in the clustering procedure.
- *Minimum cluster size*: a cluster must contain at least this many pixels to be saved for further processing. This is useful for preventing clustering of small detector

noise fluctuations.

- *Clustering radius*: pixels separated by a distance larger than this parameter will not be clustered together. Distance is defined in units of bins in the ft -map.
- *Time metric*: size of the clustering ellipse in the time direction.
- *Frequency metric*: size of the clustering ellipse in the frequency direction.

The time and frequency metrics, along with the clustering radius, define an ellipse for clustering to a particular pixel.

A general outline of the `burstegard` clustering procedure is as follows:

1. Identify all above-threshold pixels in the ft -map.
2. Starting in the bottom-left corner of the ft -map and moving from left-to-right through the rows, we designate each pixel as the seed for a cluster. Any above-threshold pixels that fall within the designated clustering ellipse (centered at the current seed pixel) are added to the cluster.
3. Once all nearby above-threshold pixels are added to the cluster, we choose another pixel from the cluster to be the seed and repeat this process. Pixels which are already identified as being part of the cluster are excluded at this point in order to eliminate redundancy and improve the computational speed.
4. When all pixels in the cluster have been used as the seed, the cluster cannot be connected to anything else and is designated as complete. All pixels associated with this cluster are excluded from further clustering.
5. If the total number of pixels in this cluster is greater than the minimum cluster size designated by the user, the cluster is saved for further analysis.
6. Move to the next remaining above-threshold pixel in the ft -map and repeat this process to begin a new cluster. This continues until all above-threshold pixels in the ft -map have been used as a seed.
7. Calculate the cluster ranking statistic SNR_Γ for each saved cluster. The user may retain all of these clusters for further analysis, or only the cluster with the maximal SNR_Γ value.

A flow chart of the process is shown in Fig. 4.14 for illustrative purposes.

The `burstegard` algorithm has been developed in the C++ programming language in order to maximize computational speed and efficiency. Although the clustering procedures are different and therefore, not directly comparable, `burstegard` is ≈ 100 times faster than the Matlab-based `BurstCluster` algorithm.

4.4.2 Examples

To demonstrate the performance of the `burstegard` clustering algorithm, I have injected two different simulated GW waveforms into Gaussian noise which is colored with the initial LIGO noise spectrum. In each case, I have applied `burstegard` to the SNR ft -map in an attempt to recover the signal. This procedure has been repeated with different clustering parameters in order to demonstrate their effect on the signal recovery. The standard clustering parameters used are as follows:

1. *Pixel threshold*: SNR = 1
2. *Minimum cluster size*: 80 pixels
3. *Clustering radius*: 2 pixels
4. *Time metric*: 1
5. *Frequency metric*: 1

The first injected signal is from an accretion disk instability [265–267,277]; the signal manifests itself as a ≈ 240 s “chirp-down” track in the ft -map of cross-correlation SNR. Fig. 4.15 shows the full SNR ft -map, along with signal recoveries using two different parameter sets, and plots showing all identified clusters. The middle row shows the recovery using the standard clustering parameters, while the bottom row uses a reduced pixel threshold of SNR = 0.5. In the first case, there are fewer noise clusters identified, but much of the signal is not recovered since the lower frequency tail is at a lower SNR and does not exceed the required threshold. Using a reduced pixel threshold allows the recovery of the entire signal, although many noise pixels are also included in the final cluster. The overall significance of the signal is reduced due the inclusion of noise pixels, as the SNR_Γ is 103.6 in the first case and 85.9 in the second case, even though more

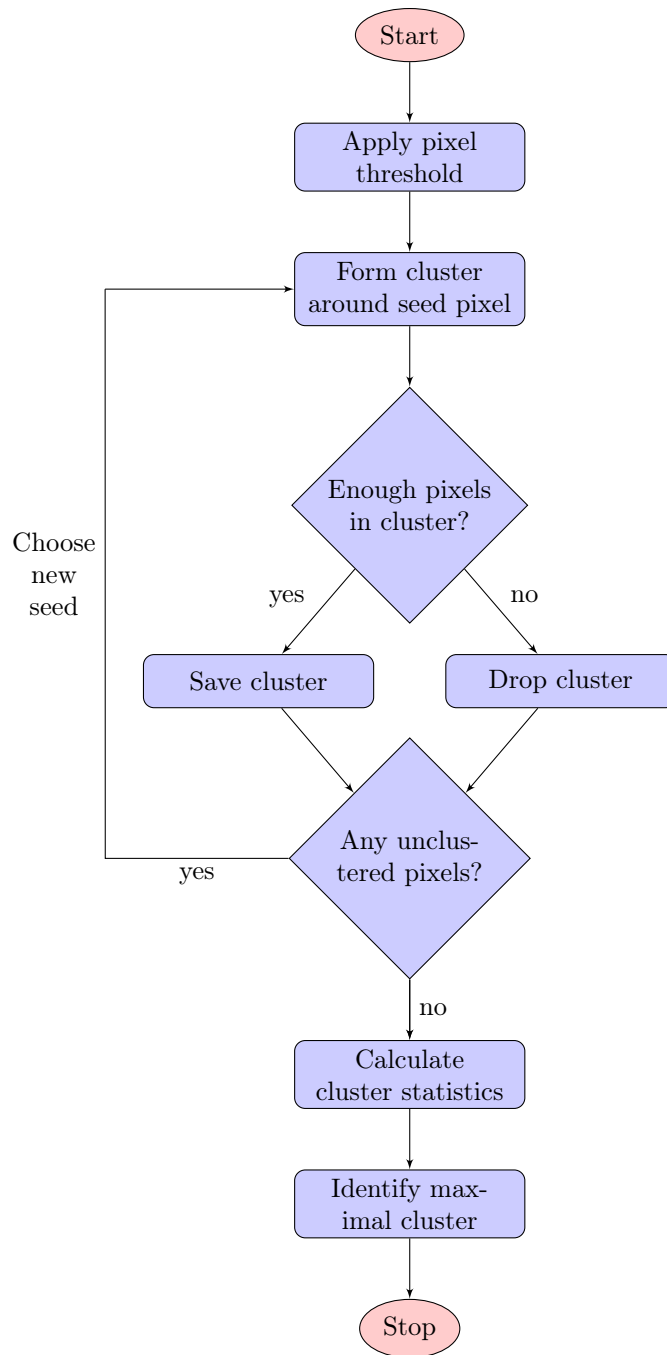


Figure 4.14: Flow chart describing the **burstegard** clustering algorithm. Forming a cluster around a seed pixel involves finding nearby above-threshold pixels using a clustering ellipse, and recursively clustering around those pixels until no further connections can be made. Cluster statistics are calculated for saved clusters only.

of the signal is technically recovered. There are also significantly more noise clusters identified when using a reduced pixel threshold (49, compared to two in the standard case).

I also inject and recover a simulated signal from the inspiral of an eccentric binary black hole (EBBH) system [275]. Fig. 4.16 shows the full SNR ft -map (the injection is visible as a short, “chirp-up” track), along with signal recoveries and all identified clusters. The middle row shows the recovery with the standard clustering parameters; only one cluster is identified, with an SNR_F of 126.1, however, the signal is not fully recovered. In the bottom row, I have increased the clustering radius from two to three pixels; this allows the clustering algorithm to jump over small gaps in the signal due to noise fluctuations for a more complete recovery, however, more noise pixels are included in the cluster, as well. The overall result is again a decrease in sensitivity, since the maximal cluster in this case has an SNR_F of 116.2. Along with this, there are also more noise clusters identified in this case.

These examples demonstrate some of the strengths and weaknesses of the **burstegard** clustering algorithm and the effects of modifying the clustering parameters.

4.5 All-sky searches with STAMP

Traditionally, gravitational-wave searches fall into one of two categories: targeted or all-sky. Targeted searches focus on detecting gravitational waves which may be emitted along with other astrophysical phenomena, like gamma-ray bursts or high-energy neutrinos. Detections of these phenomena by other observatories provide a sky direction and time window in which to search for GWs. This typically reduces some of the uncertainty in the signal parameter space since the sky direction is relatively well known, and also lowers the computational requirements of the search by reducing the amount of data which must be analyzed. The STAMP analysis pipeline, as described thus far, is designed for targeted searches: the optimal filter \tilde{Q}_{IJ} requires knowledge of the signal’s sky location in order to use the correct time delay and antenna factors.

All-sky searches instead focus on analyzing as much data as possible and detecting GWs which may come from any direction. These are typically excess-power searches,

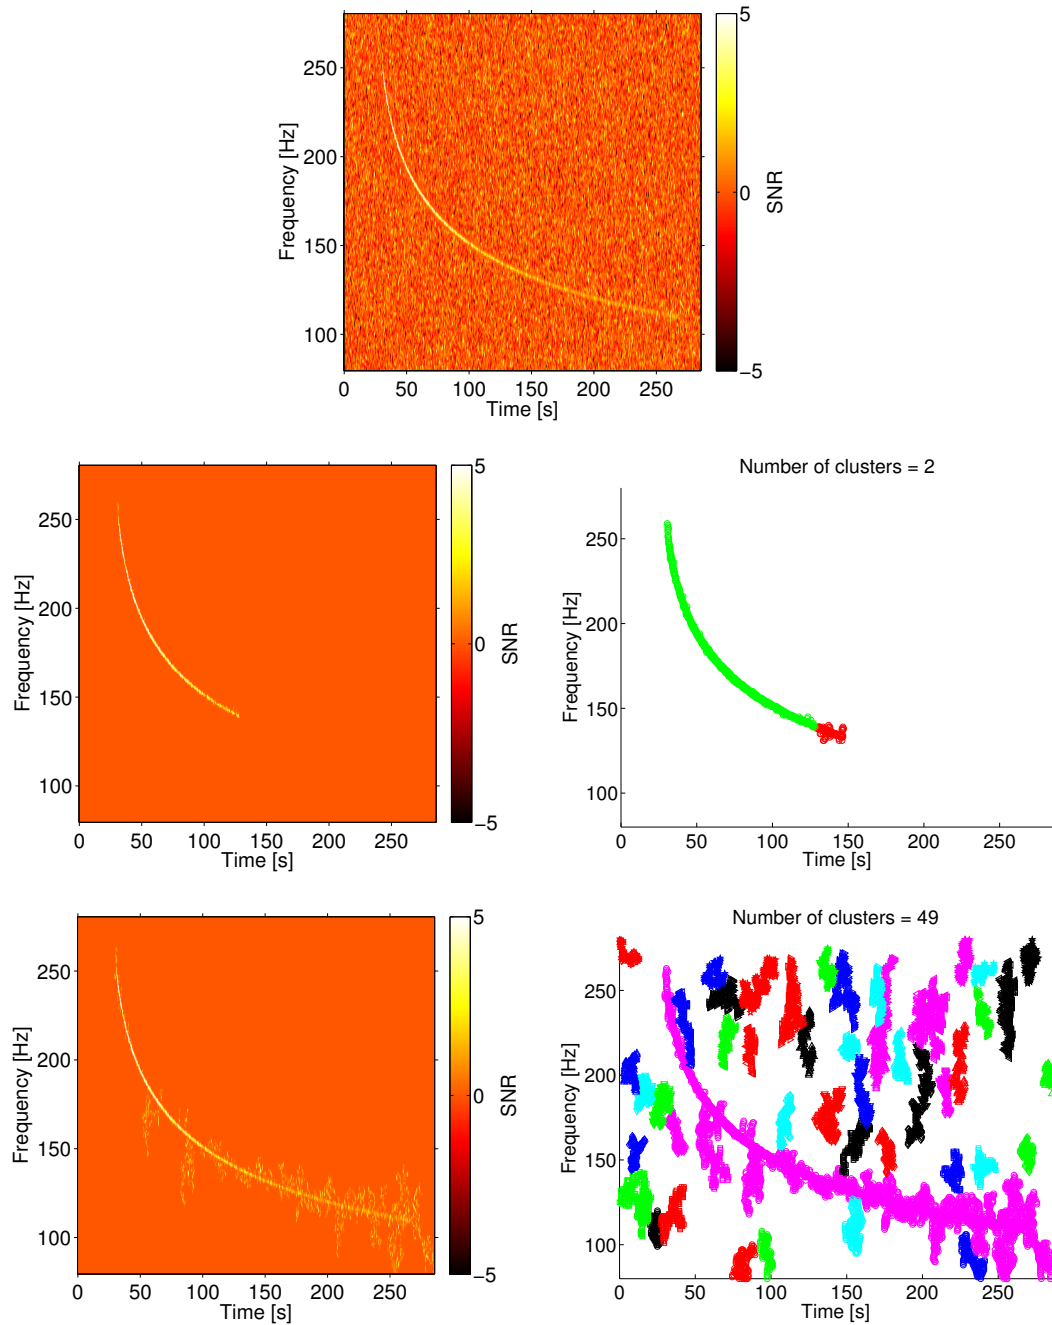


Figure 4.15: Recovery of a simulated ADI signal with the `burststegard` clustering algorithm. Top: SNR ft -map. The signal is visible as a bright, “chirp-down” track. Middle row: recovered cluster with the highest SNR_Γ (left) and all clusters identified (right) using the standard clustering parameters. Bottom row: same as middle row, but using a reduced pixel threshold (0.5 instead of 1).

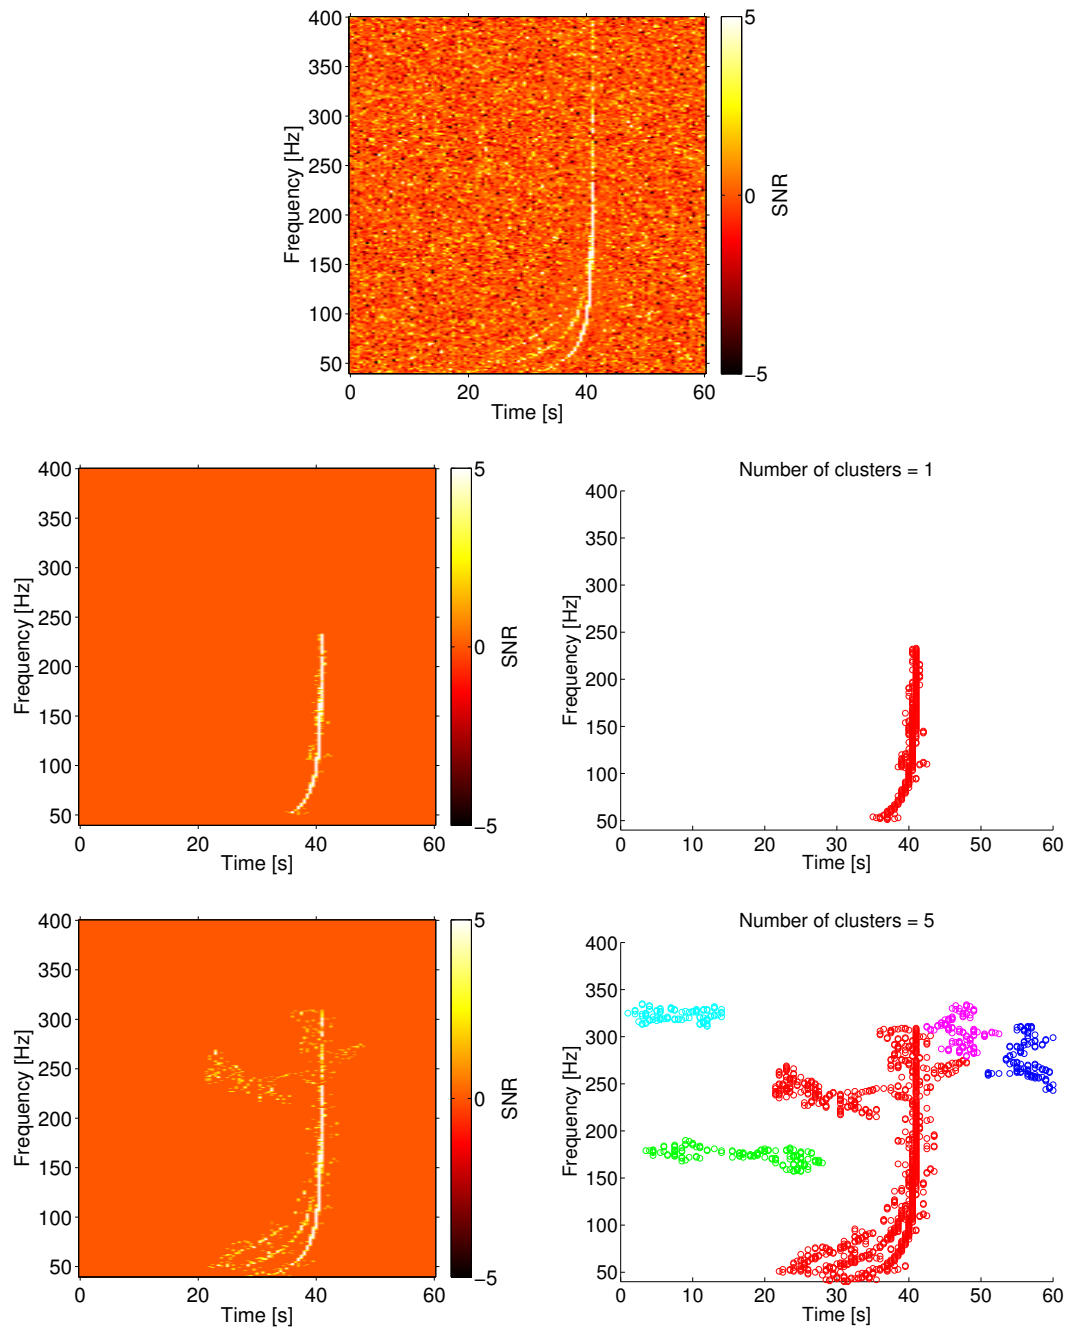


Figure 4.16: Recovery of a simulated EBBH signal with the `burstegard` clustering algorithm. Top: SNR ft -map. The signal is visible as a bright, “chirp-up” track. Middle row: recovered cluster with the highest SNR_Γ (left) and all clusters identified (right) using the standard clustering parameters. Bottom row: same as middle row, but using an increased clustering radius (3 pixels instead of 2).

which are generally sensitive to any process which induces significant power in the detectors; thus, all-sky searches are useful for detecting GWs from unknown or unmodeled mechanisms. These searches usually sacrifice some degree of sensitivity (compared to their targeted counterparts) for improved flexibility and computational efficiency, since all-sky searches typically analyze considerably larger datasets.

In this section, I will describe several adaptations which I have developed and implemented which facilitate an all-sky search with STAMP.

4.5.1 Clustering

The primary difficulty associated with all-sky searches is that the direction $\hat{\Omega}$ of a potential GW signal is not known. For a targeted search pipeline like STAMP, this means that the “wrong” sky direction may be chosen for the analysis, and as a result, the corresponding filter \tilde{Q}_{IJ} is not the optimal filter and the signal may actually induce negative power. An example of this behavior is shown in Fig. 4.17; the signal is visible as a track with alternating light and dark bands. The width and frequency of these bands is related to the difference in the time delay between detectors for the true signal direction and the direction chosen for the filter.

This situation is problematic because the clustering algorithms utilized in STAMP are designed to only identify groups of pixels with positive SNR. There are two naive solutions to this problem. First, we could repeat the STAMP analysis for a given chunk of data many times, using a different sky direction for each iteration. Because the filter \tilde{Q}_{IJ} is an optimal filter, the case where we analyze the signal using the true sky direction to construct the filter should result in a fully positive signal with maximal SNR_{Γ} , assuming that our clustering algorithm is well-suited to the task. Thus, we could divide the sky into tiles and repeat the analysis for each tile, keeping only the cluster with the highest value of SNR_{Γ} in all tiles. However, analyzing the entire sky with 1° resolution requires 65160 tiles, which requires significant computational resources.

The second obvious solution is to pick a random sky direction, take the absolute value of the SNR map, and run the normal clustering algorithm. However, this doubles the number of above-threshold noise pixels and reduces the effectiveness of the clustering algorithm significantly. The rest of this subsection will explore two alternative solutions to this problem.




Figure 4.17: SNR ft -map of a simulated ADI signal recovered with the wrong filter. The signal was injected at $(ra,dec) = (6h,30^\circ)$ and recovered with a filter designed for a sky position of $(ra,dec) = (15h,20^\circ)$.

The fastring algorithm

Because STAMP uses a two-detector baseline, many sky directions appear nearly identical, since the filter function \tilde{Q}_{IJ} is primarily dependent on the signal time delay between detectors (it also has some dependence on the antenna pattern functions F_I^A , but this does not affect the resulting signal morphology in STAMP ft -maps since this is only a scaling factor). The time delay is a function of source sky position and it can be shown geometrically that there are “rings” of constant time delay on the sky for a two-detector network.

In order to eliminate redundant sky positions and speed up an all-sky analysis, we have developed the **fastring** algorithm. This routine divides the sky into tiles, calculates the time delay for each tile, converts it to a phase delay (dependent on the maximum frequency used in the analysis and the separation between the detectors), and groups sky positions by similar phase delays. The tile size is determined by requiring that the maximum phase error is less than $\pi/4$; we enforce this requirement at $f = f_{\max}$ since this gives the largest phase error for a corresponding error in time delay.

Since each sky position in a group has essentially the same phase delay and filter

function (up to a scaling factor), we can simply analyze one sky direction from each group. This allows us to reduce the number of sky positions analyzed per ft -map from 65160 to a few hundred, although this number depends on the maximum frequency used in the analysis and on the criteria for defining phase delays as similar.

Fig. 4.18 shows the recovered SNR_Γ as a function of sky position for a simulated ADI signal at $(\text{ra}, \text{dec}) = (6\text{h}, 30^\circ)$ using both brute force iteration and the `fastring` routine. The recovered SNR_Γ using `fastring` is within 20% of that from brute force iteration, but uses only 8 sky positions instead of 316. Note that the white patch in the `fastring` recovery contains three rings of nearly equal SNR_Γ .




Figure 4.18: Recovery of a simulated ADI signal using brute force iteration (left) and the `fastring` method (right) to define sky positions. The resulting SNR_Γ values are very comparable in each ring, but the brute force method requires analysis of 316 sky positions, while `fastring` uses only 8.

Searches using `fastring` must still analyze a few hundred sky positions per ft -map to cover the entire sky; depending on the amount of data which must be analyzed for a given search, this may still be computationally unrealistic.

The zebra algorithm

Another solution to the all-sky problem relies on taking advantage of the signal morphology shown in Fig. 4.17. When the wrong sky position is used to recover a signal, it always results in a banded pattern of alternating positive and negative SNR. In order to

take advantage of this structure, I have developed the **zebragard** clustering algorithm, which is capable of recovering signals even when the wrong sky position is used for the analysis. This algorithm uses the **burstegard** clustering strategy in several steps:

1. Cluster only above-threshold positive pixels using a reduced requirement on the minimum number of pixels to form a cluster (≈ 20).
2. Re-cluster the resulting positive clusters using a slightly larger clustering radius and a larger requirement on cluster size. This allows for connection of nearby positive clusters which may be slightly disconnected due to noise fluctuations and removes small noise clusters which may be identified during step 1.
3. Multiply the ft -map by -1 and repeat steps 1 and 2. This effectively clusters pixels which were initially negative.
4. Take a positive cluster and try to connect it to all negative clusters using a large clustering radius. Repeat this process for all possible connections between positive clusters and negative clusters.
5. Combine any connected clusters into a single cluster. A higher requirement on minimum cluster size is enforced at the end of this step.
6. Take the absolute value of the ft -map and calculate the cluster ranking statistic, SNR_Γ .

Note that in step 4, we do not try to connect positive clusters to other positive clusters or negative clusters to other negative clusters. Also, although the **burstegard** clustering algorithm is applied multiple times during the full **zebragard** procedure, the number of pixels involved in each clustering step is considerably reduced (compared to the full ft -map); as a result, the algorithm runs relatively quickly.

Fig. 4.19, shows a step-by-step example of the **zebragard** clustering strategy, following the steps detailed above. Using this algorithm, it is possible to fully reconstruct the signal, even though the wrong sky direction was used in the filter.

The example shown in Fig. 4.17 is a case where although the signal is partially negative, most of the power is still apparent in the real part of the cross-correlation. Depending on how the sky direction of the source differs from that used in the filter,

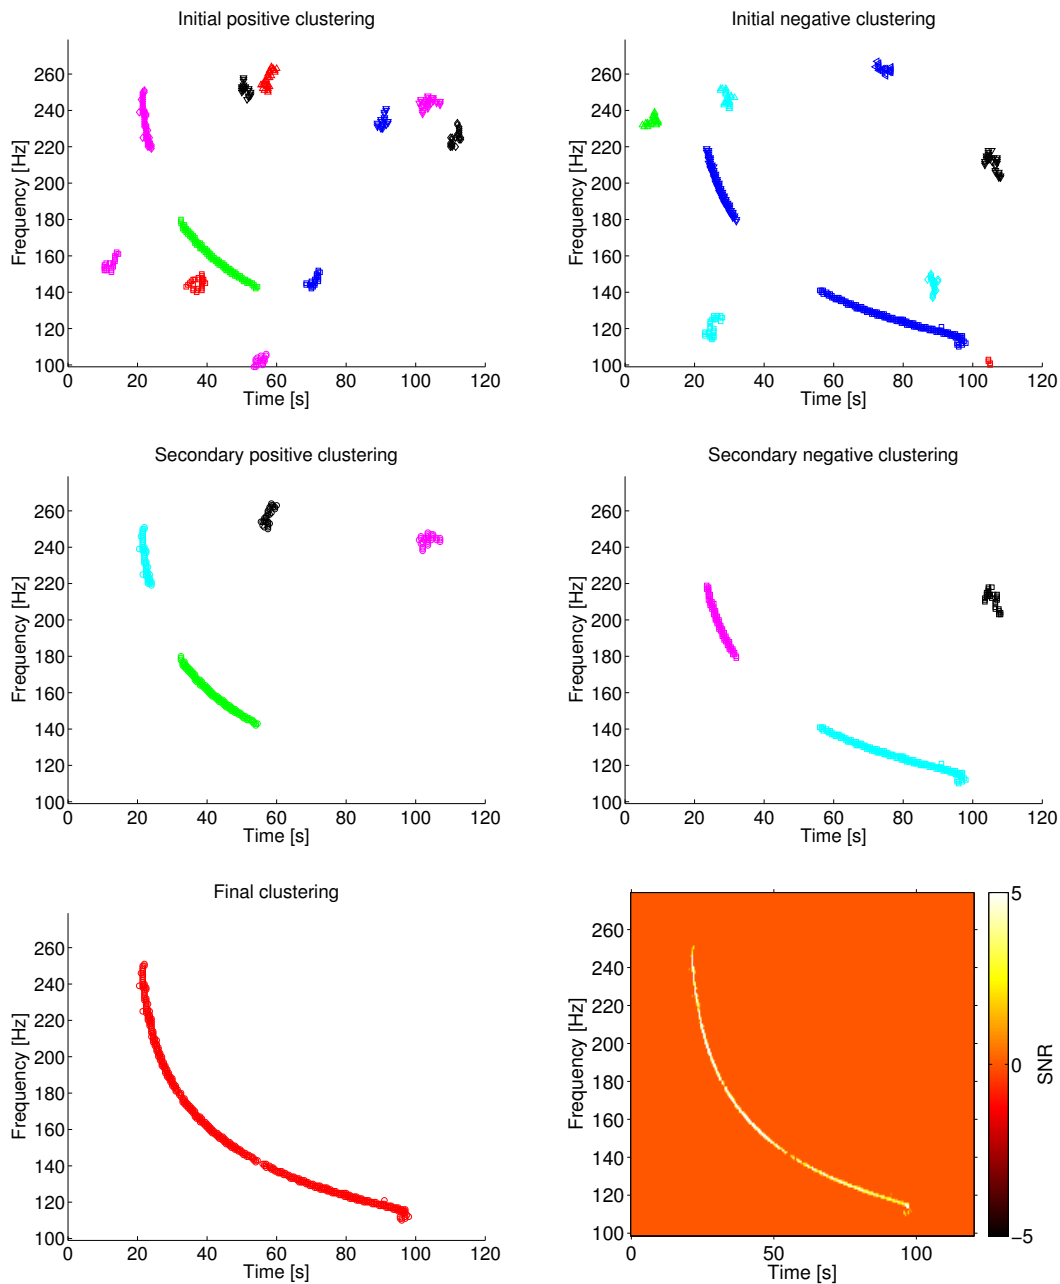


Figure 4.19: Step-by-step example of zebra-gard clustering method. Top row: initial clustering of positive pixels (left) and negative pixels (right). Middle row: secondary clustering of positive pixels (left) and negative pixels (right). This step expands the cluster radius and requires more pixels in each cluster. Bottom row: final results after connecting positive and negative clusters. The cluster plot (left) and the SNR ft -map (right, with absolute value applied) of the maximal resulting cluster are both shown.

significant power may be deposited in the imaginary part of the cross-correlation and will be lost when the real part is taken to calculate \hat{Y} (see Eq. 4.6). In these cases, the **zebragard** clustering strategy may fail since the overall signal power may be low or there may be large gaps between the bands. Because of this, it is often useful to analyze each period of data with **zebragard** using multiple sky positions in order to ensure a well-recovered signal (hopefully fewer than that required for **fastring**).

To test this, I have applied **zebragard** to several months of time-shifted LIGO data and to data containing simulated ADI signals at various source distances. The results of this study are shown in Fig. 4.20, in terms of false alarm probability versus distance for a fixed signal recovery rate of 50%, using anywhere from 1 to 10 sky directions with **zebragard** (sky directions chosen randomly for each *ft*-map). For cases where more than one sky direction is used to analyze an *ft*-map, **zebragard** returns only the cluster with the largest SNR_F found across all sky directions.

Effectively, this plot shows the probability that a noise trigger has higher SNR_F than a recovered signal as a function of distance. At constant false alarm probability, curves which are further to the right indicate higher sensitivity, meaning that we can detect the signal when its source is further away. In general, increasing the number of sky directions increases the sensitivity, but the relative gain diminishes beyond approximately 5 sky directions. The computational time increases linearly with the number of sky directions used in the analysis; thus, in order to balance sensitivity and computational efficiency, I recommend using ≈ 5 sky directions for analyses with **zebragard**.

I have also considered whether there is an optimal way of choosing the sky directions for analyzing the data. To answer this question, I have repeated the same study as detailed above for testing the number of sky directions; however, I instead vary the methods for choosing the sky directions. The methods tested are as follows:

- Random: a different set of sky directions is generated and used for analyzing each *ft*-map.
- Time delay: the sky directions are chosen so that the set of possible time delays is spanned (≈ -10 – 10 ms for the LIGO Hanford and Livingston detectors).
- Time delay (high sensitivity): same as the previous method, except we only use directions where the antenna pattern functions indicate high sensitivity.

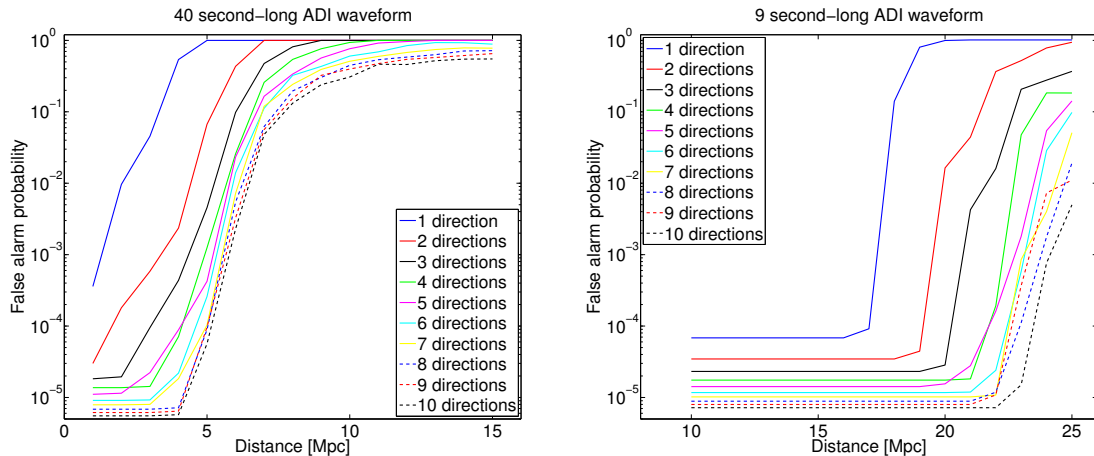


Figure 4.20: ZebraGard sensitivity as a function of number of sky directions used for a 40 second-long ADI waveform (left) and a 9 second-long ADI waveform (right). See text for more details.

- Consistent random: one set of sky directions is randomly generated and used to analyze all ft -maps.

For each method, 5 sky directions are used in total.

Fig. 4.21 shows the results of this study; again, curves which are further to the right at constant false alarm probability indicate higher sensitivity. Overall, the methods prove to be quite similar; although there are indications of some differences using the longer waveform, these are not apparent for the shorter waveform. Thus, it seems that for an all-sky search with potentially unknown GW waveform morphologies, there is no significant difference between methods for choosing the sky directions when using zebraGard.

Finally, I compare the zebraGard clustering strategy to the fastring method and to a strategy using normal burstegard clustering on the absolute value of the SNR ft -map. A plot of false alarm probability versus distance for a 40 second-long ADI waveform is shown in Fig. 4.22. Overall, the fastring-based method is the most sensitive, but it is only $\approx 10\%$ more sensitive than zebraGard with 5 sky directions, and takes about 50 times longer to run. ZebraGard is also more sensitive than the absolute value-based method and comparable in computational speed.

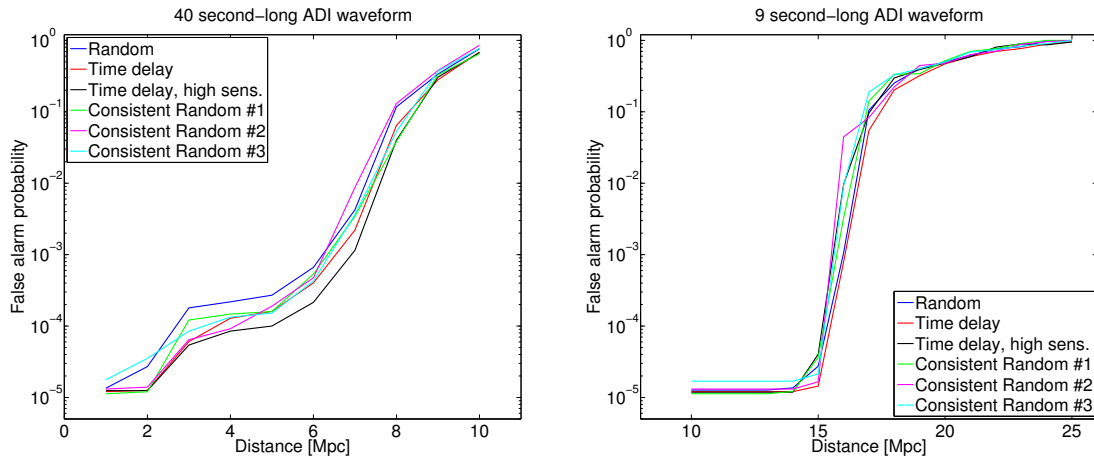


Figure 4.21: ZebraGard sensitivity as a function of the method for choosing the sky directions. Five sky directions were used for each method. See text for more details.




Figure 4.22: False alarm probability versus distance for analyses of a 40 second-long ADI waveform using `fastring`, `zebraGard`, and a method based on `burstegard` clustering of the absolute value of the SNR ft -map.

Overall, `zebraGard` with 5 sky directions appears to provide the best balance of sensitivity and computational efficiency.

4.5.2 Pre-processing of data

The STAMP analysis strategy has traditionally been separated into two phases: pre-processing and analysis. First, the auto-power ft -maps for each detector and the cross-correlation $\tilde{s}_I^*(f)\tilde{s}_J(f)$ are calculated and saved into intermediate data files. Second, these pre-processed data are loaded, the \hat{Y} , $\hat{\sigma}_Y$, and SNR ft -maps are calculated, and the clustering is performed. This strategy allows the analysis to be performed multiple times for different sky positions or choices of other analysis parameters without having to do the pre-processing each time.

The time slide (or time shift) method is commonly used in gravitational-wave data analysis to characterize the inherent detector noise [293]. In this method, one detector’s time-series is shifted relative to the other by a set amount (typically at least larger than the light travel time between detectors, and usually longer than the duration of the signal being targeted) in order to remove gravitational-wave correlations from the data. Thus, any triggers identified in the time-shifted data should be produced by detector noise alone. This is typically repeated for multiple time shifts in order to study more realizations of the detector noise. When the unshifted (or “zero-lag”) data are analyzed, the significance of any triggers can be estimated in terms of a false alarm rate (FAR): how often the detector noise alone would be expected to produce a trigger with SNR_T equal to or larger than a given value.

However, the previously discussed strategy for pre-processing would involve saving an instance of pre-processed files for each time shift analyzed, since the cross-correlation is calculated during the pre-processing. This is a problem for all-sky searches due to the sheer quantity of data which is analyzed: hundreds of days of data and hundreds or thousands of different time shifts. The amount of storage needed quickly adds up to petabytes or more, and the pre-processing time could be several months.

As a result, I have modified the pre-processing strategy outlined above. Simply put, only the FFTs are calculated in the pre-processing and the cross-correlation is now done as the first step in the analysis phase; this requires only one set of pre-processed files for an entire GW search since the time shifting procedure is now part of the analysis phase. Extensive tests have been performed in order to confirm that the end result is identical to the old pre-processing method.

4.5.3 Data quality

The STAMP glitch identification flag discussed in Sec. 4.3 was designed and tested for a targeted search where the sky direction of the GW source is known. For a search using `zebragard` where random sky directions are used, it is probable that the ϵ_{II} and ϵ_{JJ} factors will not match up with that of the true sky direction of the source, leading to an incorrect estimation of the GW auto-power difference (see Eq. 4.21). In this way, GW signals could be falsely labeled as glitches by this algorithm.

I have tested the safety of the glitch flag for simulated GW signals recovered with the wrong sky direction using a method similar to that described in Sec. 4.3.4. Because the glitch statistics are based on the fraction of “glitchy” pixels in each ft -map column, it tends to prove safe for narrowband signals which do not make up a significant portion of each column, assuming the overall frequency range of the ft -map is comparatively large. However, this is not the case for more broadband signals like the EBBH waveforms. Future searches which include waveforms similar to the EBBH signals may require a modified form of the glitch identification algorithm in order to ensure safety for these signals.

Fig. 4.23, shows the results of applying the glitch flag to an ADI waveform (left) and an EBBH waveform (right) for an all-sky search where the source direction is not known. The ADI waveform is narrowband and is not flagged as glitch-like for any signal strengths used in this study. However, the EBBH signal is broadband and makes up a large fraction of pixels in an ft -map column (see Fig. 4.16). As a result, it may be identified as glitch-like when the sky direction is not known; up to 14% of the ft -map columns containing an EBBH signal were flagged as glitch-like in this study, which may seriously impact the sensitivity of an all-sky search for these signals.




Figure 4.23: Fraction of ft -map columns containing signal which were identified as glitch-like for an ADI waveform (left) and an EBBH waveform (right) when the signal sky direction is not known. No portion of the ADI signal is flagged as glitch-like due to its narrowband morphology (all data series overlap in this plot); however, the EBBH waveform has up to 14% of its columns flagged as glitch-like, especially at closer distances when the signal is louder. Different colors indicate different parametrizations of the glitch flag.

Chapter 5

An all-sky search for long-duration gravitational-wave transients in LIGO S5 and S6 data

Thus far, searches for unmodeled bursts of GWs [107, 294, 295] have targeted source objects including core-collapse supernovae [249], neutron star to black hole collapse [296], cosmic string cusps [297], binary black hole mergers [298–300], starquakes in magnetars [301], pulsar glitches [302], and signals associated with gamma-ray bursts [121]. These searches have been primarily focused on signals with a duration of ≈ 1 second or less.

At the other end of the spectrum, searches for persistent stochastic gravitational-wave backgrounds have also been conducted, including isotropic [88] and anisotropic backgrounds [140]. However, this leaves the intermediate duration region of the unmodeled GW parameter space relatively unexplored. Multiple proposed astrophysical scenarios predict the emission of GW transients lasting from a few seconds to hundreds of seconds, or even longer (see Sec. 4.1). The first search for unmodeled long-duration GW transients used the STAMP analysis pipeline to analyze data from the LIGO S5 science run which was associated with GRB triggers from the *Swift* telescope.

In this chapter, I describe and apply a STAMP-based search for long-lasting GW transients coming from all sky directions and all times during the LIGO S5 and S6

science runs. The all-sky adaptations to STAMP (described in Sec. 4.5) are integral to this search, including the `zebragard` clustering algorithm and the new pre-processing strategy. Sec. 5.1 describes efforts to analyze the background distribution of triggers due to detector noise, including the data analyzed, data quality efforts, and the overall analysis strategy. In Sec. 5.2, 15 simulated GW signals are analyzed in order to estimate the sensitivity of the search. The results of analyzing the unshifted (zero-lag) data and the loudest events are presented in Sec. 5.3. Finally, I discuss and apply a procedure for setting upper limits in Sec. 5.4 and wrap up with a brief discussion of the results in Sec. 5.5 and a summary of the search in Sec. 5.6.

5.1 Background studies

A crucial part of any GW search is to understand the background of accidental triggers due to detector noise; this is important for preventing false identification of noise triggers as GW candidates and for improving the overall sensitivity of the search. To estimate the false alarm rate (FAR), i.e., the rate of accidental triggers due to detector noise, we introduce a non-physical time-shift between the H1 and L1 strain data before computing the cross-correlation. Analysis of the time-shifted data proceeds identically to that of unshifted data (more details in Sec. 5.1.3); assuming the number of hypothetical GW signals is small, the data should not contain a correlated GW signal, so any triggers will be generated by the detector noise. This process is repeated for multiple time-shifts in order to gain a more accurate estimate of the FAR from detector noise.

In this section, I describe the data analyzed in this search, efforts to clean the data in order to remove large data artifacts, details of the analysis, and the results of analyzing the time-shifted data.

5.1.1 Data selection

During the fifth LIGO science run (S5, November 5, 2005 to September 30, 2007), the 4 km and 2 km detectors in Hanford, WA (H1 and H2) and the 4 km detector in Livingston, LA (L1) recorded data for nearly two years. The Virgo detector (V1) in Pisa, Italy joined the LIGO detectors during this period and began its first science run on May 18, 2007.

After a two-year period of upgrades to H1 and L1 and the decommissioning of H2, the sixth LIGO and second and third Virgo science runs were organized jointly from July 7, 2009 to October 10, 2010. During the entire period spanned by S5 and S6, the H1 and L1 detectors achieved the best strain sensitivity, reaching $\approx 2 \times 10^{-23} \text{ Hz}^{-1/2}$ around 150 Hz in 2010 [303, 304]. Because of its reduced arm length, H2’s average sensitivity was more than a factor of two lower than H1 in the 40–400 Hz range. V1 sensitivity varied over time, but was always lower than the sensitivity of H1 and L1 by a factor of ≈ 1.5 –5 at frequencies higher than 60 Hz. Along with this, the total amount of coincident data for the H1-L1 pair was more than twice that of the H1-V1 and L1-V1 pairs added together. As a result, Virgo data has not been included in this analysis since it would have increased the overall sensitivity of the search by only a few percent or less at the cost of analyzing significantly more data. Thus, this search focuses only on S5 and S6 data from the H1-L1 pair.

5.1.2 Data quality

LIGO data quality flags are used to remove portions of the S5 and S6 data which have been contaminated by instrumental problems or environmental conditions that decrease the sensitivity of the detectors and increase the rate of data artifacts. Some of the phenomena which these data quality flags cover include hardware injections of simulated signals, excitation of suspension thermal modes, power fluctuations in the photon calibration laser, and severe glitching in the sideband frequencies used in the PDH locking scheme. See the discussion in Sec. 4.3.5 and Refs. [287–290] for a further description of the LIGO data quality flags. Overall, 5.8% and 2.2% of H1-L1 coincident data are discarded for S5 and S6, respectively, due to our choices for Category 1 data quality flags¹.

Category 2 data quality flags are also used to veto possible GW candidate triggers which are generated by a clearly identified noise source. For this purpose, we considered all Category 2 flags which have been used by past modeled and unmodeled GW transient searches [290]; however, we found that the most useful flags for our analysis were related to malfunctions of the longitudinal control of the Fabry-Pérot cavities and to increases in seismic noise.

¹Data quality flag selection performed by M. A. Bizouard.

We also apply two data cleaning techniques concurrently with the data processing. First, we remove frequency bins which are known to be contaminated by instrumental and environmental effects. This includes the resonant frequencies of the mirror suspensions (341–349 Hz and harmonics), 60 Hz power line harmonics, and sinusoidal signals injected for calibration purposes (399–405 Hz). In total, we have removed 47 1 Hz-wide frequency bins from the S5 data and 64 1 Hz-wide frequency bins from the S6 data. Second, we apply the STAMP glitch identification flag discussed in Secs. 4.3 and 4.5.3, in order to remove data artifacts. Both of these methods are applied before any triggers are generated by the clustering algorithm.

However, preliminary analyses of the data revealed the presence of glitches in the H1 and L1 data which randomly coincide during the time-shifting procedure. The STAMP glitch flag does not remove these glitches due to its built-in safety mechanism, which prevents it from flagging segments where there is significant power in both detectors. These glitches are typically much less than 1 s in duration, and as a result, nearly all of their power is concentrated in a single 1 s time segment. To suppress these glitches, I have developed a discriminant variable, `SNRfrac`, which measures the fraction of a trigger’s SNR_T deposited in each time segment (or ft -map column). Using studies of the detector noise (see Sec. 5.1.3) and a set of simulated GW signals (see Sec. 5.2.1), I found that eliminating any triggers with $\text{SNRfrac} \geq 0.45$ optimized the search sensitivity for all simulated waveforms studied in this search. This cut is applied in post-processing, after the initial analysis which generates triggers using a clustering algorithm.

5.1.3 Background estimation

Analysis details

The Category 1 data quality flags that we use divide each science run into analysis segments; we further divide each analysis segment into 500 s-long intervals which overlap by 50% and span the entire dataset. As an example, there are 18837 analysis segments in the S5 dataset with an average duration of about 1500 seconds each; this corresponds to five 500-second, 50% overlapping ft -maps per analysis segment on average.

For each interval, we calculate the standard STAMP ft -maps, including SNR , \hat{Y} , and $\hat{\sigma}_Y$, with $1 \text{ s} \times 1 \text{ Hz}$ pixel resolution. We apply the `zebragard` clustering algorithm

with 5 randomly chosen sky directions in order to generate triggers; the sky directions are consistent across intervals within an analysis segment.

In terms of frequency content, we restrict the analysis to the 40–1000 Hz band. The lower limit is set due to the presence of seismic noise, which is the dominant low-frequency noise source for current terrestrial GW detectors and rises sharply with decreasing frequency. The upper limit is set to include the most likely regions of frequency space for long-lasting GW transients, while keeping the computational requirements of the search at an acceptable level. We note that the frequency range of our analysis includes the most sensitive part of the LIGO frequency band, namely 100–200 Hz.

In order to generate more realizations of detector noise, we utilize the time-shift method, briefly described in Sec. 4.5.2. For a given time-shift i , the H1 data from interval n are cross-correlated with L1 data from interval $n + i$. Since there may be a time gap between two consecutive analysis segments, the actual time-shift applied in this process is always at least $500i$ seconds. The time-shifting method is also circular: if $n + i > N$, where N is the total number of 500 s intervals required to span the dataset, then the H1 data from interval n are cross-correlated with L1 data from interval $n + i - N$. It is important to note that the minimum time-shift duration is much longer than the light travel time between the two detectors and also longer than the signal models we consider (see Sec. 5.2.1) in order to prevent accidental correlations.

Using this method, I have analyzed 100 time-shifts to estimate the background due to detector noise during S5, amounting to a total analyzed livetime of 84.1 years. I have also studied 100 time-shifts of S6 data, with a total analyzed livetime of 38.7 years.

Background results

The cumulative rates of background triggers for the S5 and S6 datasets can be seen in Fig. 5.1. After applying the `SNRfrac` and Category 2 data quality flags in post-processing, the background distribution for time-shifted data is comparable to that resulting from a study of simulated Gaussian detector noise, indicating that the resulting data are relatively well-cleaned.

Looking at the “loudest” remaining triggers in the cleaned distributions from Fig. 5.1 (darkest gray), we would expect S5 detector noise to generate triggers with $\text{SNR}_\Gamma \gtrsim 35$ at a rate of 4×10^{-10} Hz, or once for every 79 years of data; for S6, we expect triggers

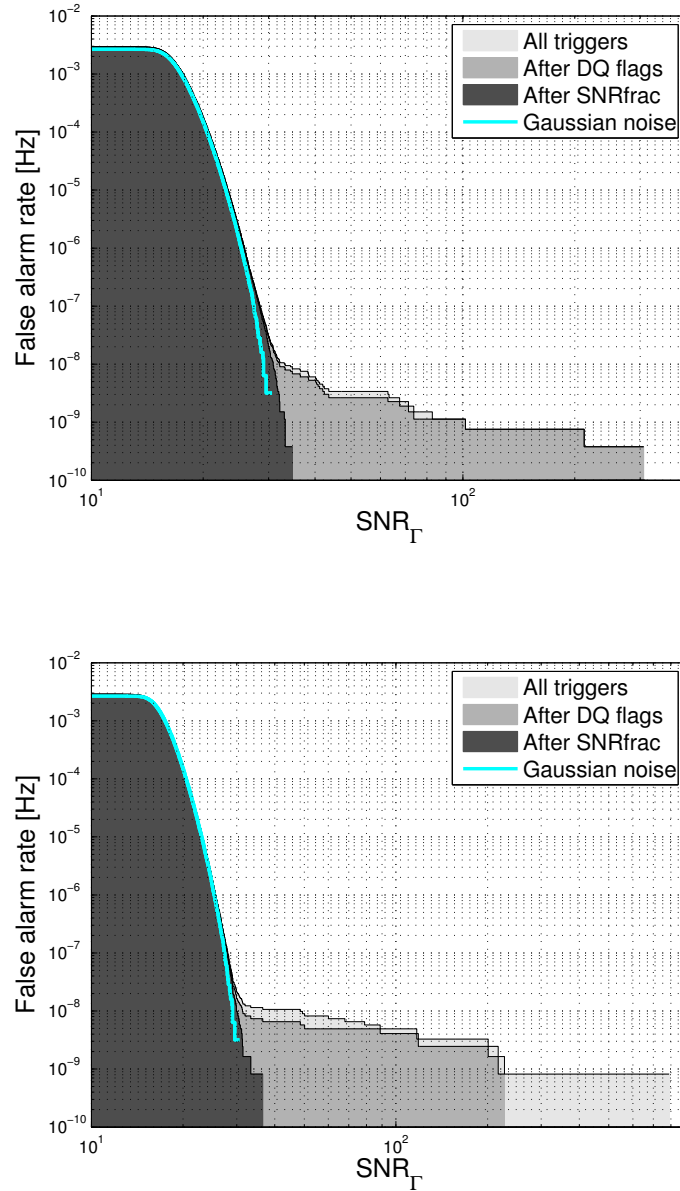


Figure 5.1: The false alarm rate is shown as a function of the trigger ranking statistic SNR_T for 100 time-shifts of S5 data (top) and S6 data (bottom). Distributions are shown before and after post-processing cuts are applied. Also shown is the FAR distribution generated by a Monte Carlo simulation assuming Gaussian detector noise. Reprinted figure with permission from Ref. [142], copyright 2016 by the American Physical Society.

with $\text{SNR}_\Gamma \gtrsim 37$ to occur once for every 35 years of data. This study indicates that the data cleaning methods are working as expected and can be applied to the zero-lag data; any resulting high- SNR_Γ triggers from that study should only be produced by GWs.

5.2 Search sensitivity

In this section, I describe the application of our search algorithm to time-shifted LIGO data plus several simulated GW signals, in order to test the effectiveness of our search at identifying and recovering GWs.

5.2.1 Waveforms

To assess the sensitivity of our search to realistic GW signals, we use 15 different simulated GW waveforms, classified into four different waveform “families.” Four of these waveforms are based on an astrophysical model of a black hole accretion disk instability [265–267, 277]. The other 11 waveforms are not based on a model of an astrophysical GW source, but are chosen to encapsulate various characteristics that long-transient GW signals may possess, including their duration, frequency content, and bandwidth. These *ad hoc* waveforms can be divided into three families: sinusoids with time-dependent frequency content, sine-Gaussians, and band-limited bursts of white noise. I have developed simulations which produce the strain time-series for each of these signal types and used them to generate the waveforms used in this search.

All simulated waveforms have 1 s-long Hann-like tapers applied to the beginning and end of the waveforms in order to prevent data artifacts which may occur when the signals have high intensity. In this section, I will provide a brief description of each GW signal model and the sensitivity of our search algorithm to them.

Accretion disk instabilities

In our sensitivity study, we include four variations on the ADI model (described further in Sec. 4.1.1). The parameters for the four waveforms are described in Table 5.1; we note that although this set of waveforms does not span the entire parameter space of the ADI model, it does encapsulate most of the possible variations in the signal morphology

in terms of durations, frequency ranges and derivatives, and amplitudes². While these waveforms may not be precise representations of realistic signals, they capture the important features of many proposed models and produce long-lived spectrogram tracks. A summary of these waveforms is shown in Tab. 5.1; Fig. 5.2 illustrates the different signal morphologies with SNR ft -maps.

Waveform	$M[M_{\odot}]$	a^*	ϵ	Duration [s]	Frequency [Hz]
ADI-A	5	0.30	0.050	39	135–166
ADI-B	10	0.95	0.200	9	110–209
ADI-C	10	0.95	0.040	236	130–251
ADI-E	8	0.99	0.065	76	111–234

Table 5.1: Parameters of ADI waveforms [265–267] used to test the sensitivity of the search. Here, M is the mass of the central black hole, a^* is the dimensionless Kerr spin parameter of the black hole, and ϵ is the fraction of the disk mass that forms clumps. Frequency refers to the ending and starting frequencies of the GW signal, respectively. All waveforms have an accretion disk mass of $1.5 M_{\odot}$.

Sinusoids

The sinusoidal waveforms are characterized by a sine function with time-dependent frequency content. The waveforms are described by:

$$h_{+}(t) = \frac{1 + \cos^2 \iota}{2} \cos 2\psi \cos \phi(t) - \cos \iota \sin 2\psi \sin \phi(t), \quad (5.1)$$

$$h_{\times}(t) = \frac{1 + \cos^2 \iota}{2} \sin 2\psi \cos \phi(t) + \cos \iota \cos 2\psi \sin \phi(t), \quad (5.2)$$

where ι is the inclination angle of the source, ψ is the source polarization, and $\phi(t)$ is a phase time-series, given by

$$\phi(t) = 2\pi \left(f_0 t + \frac{1}{2} \left(\frac{df}{dt} \right) t^2 + \frac{1}{6} \left(\frac{d^2 f}{dt^2} \right) t^3 \right). \quad (5.3)$$

Two of the sinusoidal waveforms are completely monochromatic, two have a linear frequency evolution with time, and two have a quadratic frequency evolution with time.

²Although amplitude depends on the distance to the source, it is also affected by the choice of model parameters, including mass and disk clumpiness.




Figure 5.2: SNR ft -maps of the four ADI waveforms used to test the STAMP all-sky search sensitivity. Top-left: ADI-A. Top-right: ADI-B. Bottom-left: ADI-C. Bottom-right: ADI-E. See Tab. 5.1 for more details.

This family of waveforms is summarized in Tab. 5.2 and displayed in Fig. 5.3. For the monochromatic waveforms, some spectral leakage is apparent in neighboring frequency bins. The central frequency bin of these signals is reduced in amplitude due to the use of adjacent segments in the noise estimation, which includes contributions from the signal in this case.




Figure 5.3: SNR ft -maps of six sinusoidal signals used to test the sensitivity of the STAMP all-sky search. Top row: monochromatic waveforms MONO-A (left) and MONO-B (right). Middle row: linear waveforms LINE-A (left) and LINE-B (right). Bottom row: quadratic waveforms QUAD-A (left) and QUAD-B (right). See Tab. 5.2 for more details.

Waveform	Duration [s]	f_0 [Hz]	$\frac{df}{dt}$ [Hz/s]	$\frac{d^2f}{dt^2}$ [Hz/s ²]
MONO-A	150	90	0.0	0.00
MONO-B	250	505	0.0	0.00
LINE-A	250	50	0.6	0.00
LINE-B	100	900	-2.0	0.00
QUAD-A	30	50	0.0	0.33
QUAD-B	70	500	0.0	0.04

Table 5.2: Parameters of sinusoidal waveforms used to test the sensitivity of the search. Here, f_0 is the initial frequency of the signal, df/dt is the frequency derivative, and d^2f/dt^2 is the second derivative of the frequency.

Sine-Gaussians

The sine-Gaussian waveforms are essentially monochromatic signals multiplied by a Gaussian envelope:

$$e^{-t^2/\tau^2}. \quad (5.4)$$

Here, τ is the decay time, which defines the width of the Gaussian envelope. This set of waveforms is summarized in Tab. 5.3 and illustrated in Fig. 5.4.

Waveform	Duration [s]	f_0 [Hz]	τ [s]
SG-A	150	90	30
SG-B	250	505	50

Table 5.3: Parameters of sine-Gaussian waveforms used to test the sensitivity of the search. Here, τ is the decay time of the Gaussian envelope.

Band-limited white noise bursts

I have generated random white noise and used a 6th-order Butterworth band-pass filter to restrict the noise to a given frequency band. Each polarization component of the simulated waveforms is generated independently so that the two components are uncorrelated. This family of waveforms is summarized in Tab. 5.4 and illustrated in Fig. 5.5.




Figure 5.4: SNR ft -maps of two sine-Gaussian signals used to test the sensitivity of the STAMP all-sky search. Left: SG-A waveform. Right: SG-B waveform. See Tab. 5.3 for more details.

Waveform	Duration [s]	Frequency band [Hz]
WNB-A	20	50–400
WNB-B	60	300–350
WNB-C	100	700–750

Table 5.4: Parameters of band-limited white noise burst waveforms used to test the sensitivity of the search.

5.2.2 Sensitivity study

Using the waveforms described in the previous section, we performed a sensitivity study to determine the overall detection efficiency of this search as a function of waveform amplitude. The procedure for this study is as follows:

1. For each waveform, randomly generate 1500 times between the beginning and end of the dataset, such that the injected waveform would be fully included in at least one 500 second-long analysis window. Consecutive injections are required to be at least 1000 seconds apart.
2. For each of these injection times, generate a simulated signal with random sky position, source inclination, and polarization angle.
3. Add the simulated signal to time-shifted data and perform the analysis as described




Figure 5.5: SNR ft -maps of three white noise burst waveforms used to test the sensitivity of the STAMP all-sky search. Top row: WNB-A (left) and WNB-B (right). Bottom: WNB-C. See Tab. 5.4 for more details.

in Sec. 5.1.3.

4. The simulated signal is considered “recovered” if the following requirements are met:

- The trigger was found between the known start and end time of the simulated signal.
- The trigger was found within the known frequency band of the signal.
- The SNR_Γ of the trigger exceeded a threshold determined by the loudest

trigger found in the zero-lag dataset (without injections).

5. Repeat this procedure for 16 different amplitudes for each waveform in order to fully characterize the search’s detection efficiency as a function of signal strength.

This procedure is performed independently for each of the S5 and S6 datasets³, and the results are combined using a method which will be discussed in Sec. 5.4.

Figs. 5.6 and 5.7 show the efficiency (ratio of recovered signals to the total number of injected signals) for the S5 and S6 datasets, respectively. The efficiency is plotted as a function of either the distance to the source (for the ADI waveforms) or the root-sum-squared strain amplitude (h_{rss}) at Earth (for the *ad hoc* waveforms), defined as

$$h_{\text{rss}} = \sqrt{\int (|h_+(t)|^2 + |h_\times(t)|^2) dt}. \quad (5.5)$$

In general, the pipeline is more sensitive to signals which have a higher efficiency at a given distance or h_{rss} .

Within each of the four waveform families, the pipeline efficiency has a significant frequency dependence which traces the strain sensitivity of the LIGO detectors (see Fig. 2.3). For example, the search is more sensitive to LINE-A than LINE-B, primarily due to the fact that LINE-A is in the most sensitive portion of the LIGO band, while LINE-B is at much higher frequencies where the detectors are less sensitive. Overall, the search is more sensitive to all of the signal models in the S6 dataset (compared to S5) due to the improvement in strain sensitivity for the enhanced LIGO configuration used during S6.

In general, it may be more difficult to detect longer signals; if all waveforms are fixed at a particular amplitude, the power will be more spread out for longer signals and may fall below the detector noise (depending on the fixed amplitude chosen). Finally, the rate of change of frequency for a waveform also influences its detectability for this search; since STAMP uses adjacent time segments for estimating $\hat{\sigma}_Y$, signals whose frequency changes slowly with time will deposit significant power in $\hat{\sigma}_Y$, resulting in reduced SNR. This is particularly problematic for the monochromatic and sine-Gaussian waveforms.

³Parts of the sensitivity study were performed by M. A. Bizouard and S. Franco.

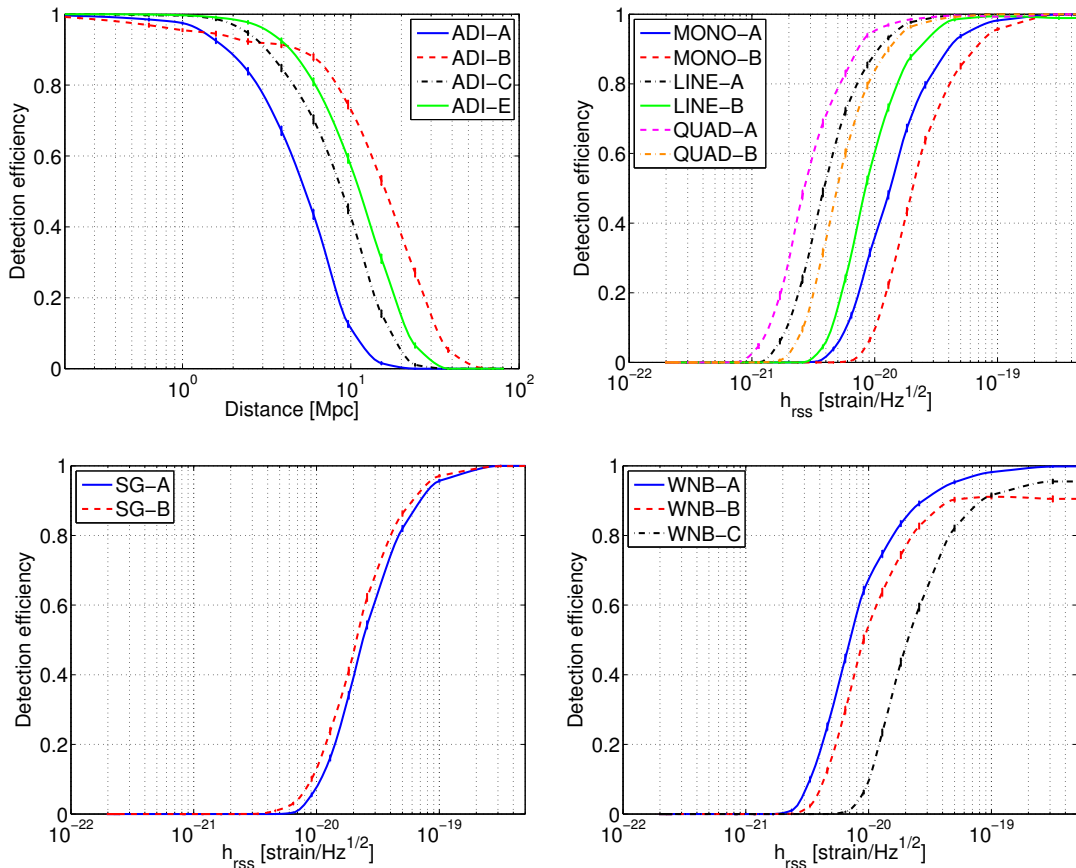


Figure 5.6: Detection efficiency for different waveforms in S5 data as a function of distance to the source (ADI waveforms) or the signal strength h_{rss} (all others). We use an `SNRfrac` threshold of 0.45 and a recovery threshold of $\text{SNR}_\Gamma = 29.65$ as determined by the loudest unshifted S5 trigger (see Tab. 5.5). Error bars are computed using binomial statistics. Top-left: ADI waveforms. Top-right: sinusoidal waveforms. Bottom-left: sine-Gaussian waveforms. Bottom-right: white noise burst waveforms. Reprinted figure with permission from Ref. [142], copyright 2016 by the American Physical Society.

5.3 Loudest triggers

After studying the distribution of background triggers from time-shifted data and optimizing the `SNRfrac` veto using both background and simulated GW signals, the final step in the analysis is to apply the search algorithm to the zero-lag data in order to

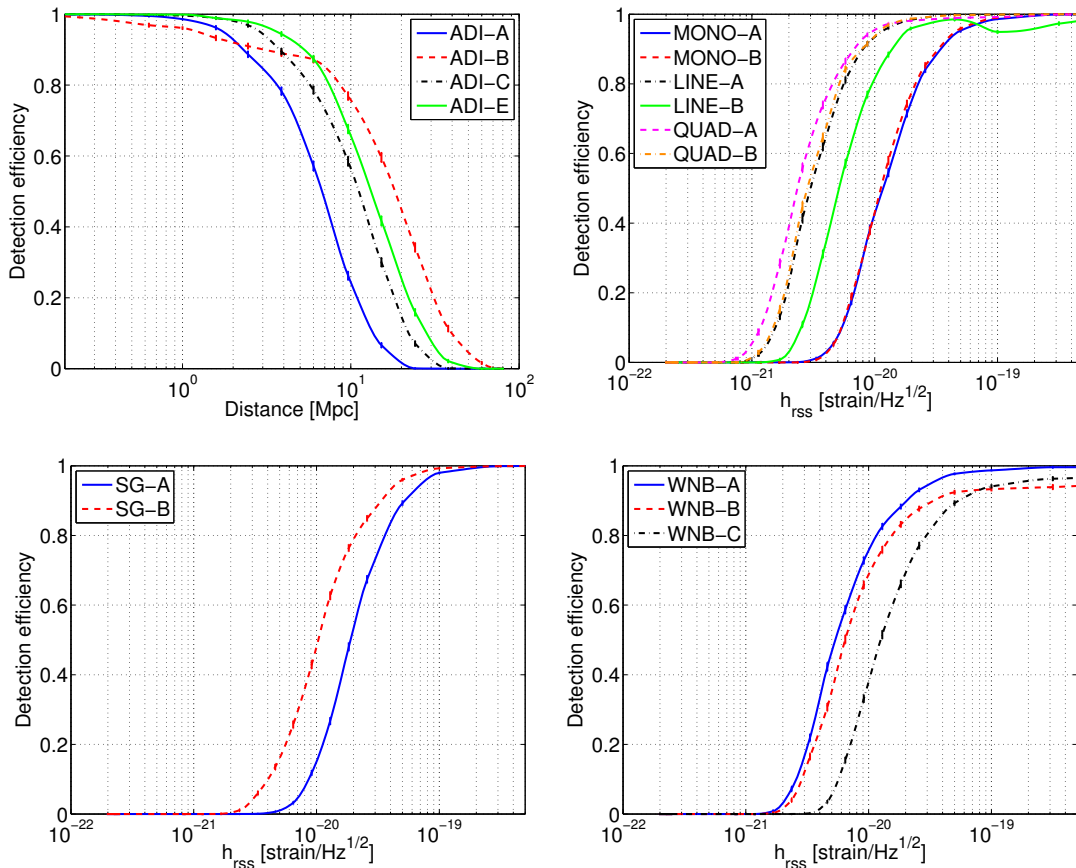


Figure 5.7: Detection efficiency for different waveforms in S6 data as a function of distance to the source (ADI waveforms) or the signal strength h_{rss} (all others). We use an SNR_{frac} threshold of 0.45 and a recovery threshold of $\text{SNR}_{\Gamma} = 27.13$ as determined by the loudest unshifted S6 trigger (see Tab. 5.5). Error bars are computed using binomial statistics and are on the order of a few percent. Top-left: ADI waveforms. Top-right: sinusoidal waveforms. Bottom-left: sine-Gaussian waveforms. Bottom-right: white noise burst waveforms. Reprinted figure with permission from Ref. [142], copyright 2016 by the American Physical Society.

search for GW candidates. The analysis proceeds identically to that of the time-shifted data, as described in Sec. 5.1.3 (without the time-shifting, of course). I have run this analysis and studied the resulting distributions of FAR versus SNR_{Γ} for both S5 and S6; these distributions are compared to the corresponding background distributions in

Fig. 5.8. A slight deficit of triggers is present in the S6 zero-lag, but it remains within one standard deviation of what is expected from the background.




Figure 5.8: Distribution of FAR versus SNR_Γ for zero-lag triggers from S5 (black circles) and S6 (red triangles). The distributions are compared to the estimated background distributions for S5 (solid black) and S6 (dashed red). We observe a slight deficit of triggers in the S6 zero-lag that remains within one standard deviation of the expectation based on the background. Also shown is the FAR distribution generated by a Monte Carlo simulation assuming Gaussian detector noise (solid cyan). Reprinted figure with permission from Ref. [142], copyright 2016 by the American Physical Society.

Here, I consider the significance of the “loudest” triggers from the S5 and S6 zero-lag datasets, including calculation of both the FAR and FAP (false alarm probability) of each trigger. The FAR is defined as

$$\text{FAR}(\text{SNR}_\Gamma^*) = \frac{N(\text{SNR}_\Gamma^*)}{T_{\text{bkg}}}, \quad (5.6)$$

where $N(\text{SNR}_\Gamma^*)$ is the number of background triggers with SNR_Γ larger than a given threshold SNR_Γ^* , and T_{bkg} is the total background livetime from all time-slides. The FAP is the probability of observing at least N background triggers with SNR_Γ higher

than the threshold SNR_Γ^* :

$$\text{FAP}(N) = 1 - \sum_{n=0}^{n=N-1} \frac{\mu_{\text{bkg}}^n}{n!} e^{-\mu_{\text{bkg}}}. \quad (5.7)$$

Here, μ_{bkg} is the number of background triggers expected from a Poisson process, defined as

$$\mu_{\text{bkg}} = T_{\text{obs}} \times \text{FAR}(\text{SNR}_\Gamma^*), \quad (5.8)$$

where T_{obs} is the zero-lag livetime. For the loudest triggers in each dataset, we take $N = 1$ to estimate the FAP.

The most significant triggers from the S5 and S6 zero-lag analyses occurred with false alarm probabilities of 54% and 92%, respectively. They also have respective false alarm rates of 1.00 yr^{-1} and 6.94 yr^{-1} . This shows that triggers of this significance are frequently generated by detector noise alone, and thus, these triggers cannot be considered GW candidates. More details on these triggers are shown in Tab. 5.5.

Dataset	SNR_Γ	FAR [yr^{-1}]	FAP	GPS time	Dur. [s]	Freq. [Hz]
S5	29.65	1.00	0.54	851136555.0	23.5	129–201
S6	27.13	6.94	0.92	958158359.5	23.5	537–645

Table 5.5: The most significant triggers from the S5 and S6 datasets. GPS times given correspond to trigger start times.

Additional follow-up on these triggers indicated that no Category 2 data quality flags in H1 nor L1 were active at the time of these triggers⁴. Examination of the ft -maps, the whitened time-series around the time of the triggers, and the monitoring records indicate that these triggers were due to a small excess of noise in H1 and/or L1, and were not associated with a well-identified environmental or instrumental noise source.

5.4 Upper limits

Since no GW candidates were identified, I proceed to place upper limits on the rate of long-duration GW transients at Earth, assuming an isotropic and uniform distribution

⁴This follow-up work was performed by M. A. Bizouard.

of sources. I use two implementations of the loudest event statistic [305] to set upper limits at 90% confidence. The first method uses the false alarm density formalism, which accounts for both the background noise and the sensitivity of the search to simulated GW signals. This method is only applicable for signals which are distance-calibrated; in our case, this is true only for the ADI signals. The second method uses the search pipeline's efficiency to set loudest event statistic rate upper limits. In this section, I describe both methods for calculating upper limits, associated uncertainties, and the resulting upper limits from this search.

5.4.1 False alarm density and visible volume calculations

In order to constrain the rate and source density of the GW signals studied in this search, I estimate the volume of the sky in which the search algorithm is sensitive to these signals. For this, I use the visible volume [306, 307]

$$V_{\text{vis}}(\text{SNR}_{\Gamma}) = \sum_i 4\pi r_i^2 \left(\frac{dN_{\text{inj}}}{dr}(r_i) \right)^{-1}. \quad (5.9)$$

Here, the index i runs over detected injections (determined by the detection threshold, SNR_{Γ}), r_i is the distance to the i th detection, and dN_{inj}/dr is the radial density of injections.

The false alarm density (FAD) statistic is based on the visible volume and is useful for comparing the results of searches over different datasets or even using different detector networks [306–308]. It provides an estimate of the number of background triggers expected given the visible volume and background livetime of the search. The classical FAD is defined in terms of the FAR divided by the visible volume:

$$\text{FAD}_c(\text{SNR}_{\Gamma}) = \frac{\text{FAR}(\text{SNR}_{\Gamma})}{V_{\text{vis}}(\text{SNR}_{\Gamma})}. \quad (5.10)$$

In this way, the FAD accounts for the network sensitivity to GW sources as well as the detector noise level and the accumulated livetime.

Following Ref. [308], we define a new FAD which produces a monotonic ranking of triggers:

$$\text{FAD}(\text{SNR}_{\Gamma,i}) = \min(\text{FAD}_c(\text{SNR}_{\Gamma,i}), \text{FAD}_c(\text{SNR}_{\Gamma,i-1})), \quad (5.11)$$

where the index i runs over triggers in increasing order of SNR_Γ . The resulting FAD statistic may be used to combine results from searches over different datasets or with different detector networks by calculating the time-volume productivity of the combined search:

$$\nu(\text{FAD}) = \sum_k V_{\text{vis},k}(\text{FAD}) \times T_{\text{obs},k}. \quad (5.12)$$

Here, the index k runs over datasets or detector networks. The uncertainty on ν can be calculated as

$$\sigma_\nu(\text{FAD}) = \sqrt{\sum_k T_{\text{obs},k}^2 \sigma_{V_{\text{vis},k}}^2(\text{FAD})} \quad (5.13)$$

This combined time-volume product is eventually used to calculate final upper limits.

First, I have calculated V_{vis} and the FAD as a function of SNR_Γ using the background distributions and the waveform sensitivity studies. Note that these quantities are calculated individually for each signal model since the detection efficiency is different for each waveform. Then, these quantities are used to estimate the FAD of the zero-lag triggers; a detection threshold is set for the combined search based on the SNR_Γ of the “loudest event,” or in this case, the zero-lag trigger with the lowest overall FAD from both datasets. The 90% confidence upper limits are calculated (following [305]) as

$$R_{90\%,\text{VT}} = \frac{2.3}{\nu(\text{FAD}^*)}, \quad (5.14)$$

where ν is given by Eq. 5.12 and FAD^* is the FAD of the loudest zero-lag trigger. The factor of 2.3 in the numerator is the 90% upper limit on the Poisson rate for zero observed events. The subscript VT indicates that these upper limits are in terms of number of observations per volume per time.

This form for the upper limits does not account for the calibration uncertainty of the detectors, which translates into an uncertainty on the visible volume. Statistical uncertainties are also present from the sensitivity study, but are only on the order of a few percent compared to the systematic uncertainties. These uncertainties are marginalized over using a Bayesian formalism detailed in the next subsection.

5.4.2 Bayesian formalism for FAD calculation

The observed astrophysical rate of triggers is given by

$$R = \frac{N}{\nu(\text{FAD}^*)}, \quad (5.15)$$

where N is the number of triggers observed by the search, ν is the time-volume product described in Eq. 5.12, and FAD^* is the false alarm density of the most significant zero-lag trigger. Although this search found no GW candidates, N is Poisson distributed, with some underlying mean and variance μ . To properly handle the uncertainty associated with the time-volume product, I have utilized a Bayesian formalism to account for uncertainties in both quantities in the calculation of rate upper limits. Here, the goal is to constrain the rate using the mean expected number of triggers, μ .

Given an observation of N triggers with a time-volume product of $\bar{\nu}$, a posterior distribution for μ and ν can be calculated by using Bayes' theorem.

$$P(\mu, \nu | N, \bar{\nu}) = \frac{P(N, \bar{\nu} | \mu, \nu) P(\mu, \nu)}{P(N, \bar{\nu})}, \quad (5.16)$$

where $P(\mu, \nu | N, \bar{\nu})$ is the posterior distribution, $P(N, \bar{\nu} | \mu, \nu)$ is the likelihood, $P(\mu, \nu)$ are the priors, and $P(N, \bar{\nu})$ is the evidence. In this case, the evidence may be discarded, since it is simply a normalization factor.

The priors on μ and ν are uniform:

$$\mu \in [0, \mu_{\max}] \quad (5.17a)$$

$$\nu \in [0, \nu_{\max}]. \quad (5.17b)$$

Here μ_{\max} and ν_{\max} are defined to be large enough that all of the significant posterior mass is enclosed. Note that the choices of these limits vary for each signal model, since they are dependent on the search sensitivity.

The likelihood function can be framed as the product of two separate distributions: a Poisson distribution with mean μ and a Gaussian distribution with mean $\bar{\nu}$ and sigma σ_ν (see Eq. 5.13). Assuming that μ and $\bar{\nu}$ are uncorrelated random variables, the joint probability density function is just the product of the two individual probability density

functions:

$$P(N, \bar{\nu} | \mu, \nu) = \frac{e^{-\mu} \mu^N}{N!} \frac{e^{-(\nu - \bar{\nu})^2 / (2\sigma_\nu^2)}}{\sqrt{2\pi}\sigma_\nu}. \quad (5.18)$$

Thus far, this formalism has been laid out to calculate the posterior distribution as a function of μ and ν . However, we are primarily interested in setting limits on the rate, given as the ratio of μ and ν . The posterior distribution can be transformed into a function of two new variables by multiplying the original posterior by the Jacobian determinant of the transformation:

$$P(R, \nu | N, \bar{\nu}) = P(\mu, \nu | N, \bar{\nu}) \times |J|. \quad (5.19)$$

Choosing the new variables to be R and ν (for simplicity) gives a Jacobian determinant of

$$|J| = \begin{vmatrix} \frac{\partial \mu}{\partial R} & \frac{\partial \mu}{\partial \nu} \\ \frac{\partial \nu}{\partial R} & \frac{\partial \nu}{\partial \nu} \end{vmatrix} = \nu. \quad (5.20)$$

This gives a posterior distribution of

$$P(R, \nu | N, \bar{\nu}) = \nu \frac{e^{-\mu} \mu^N}{N!} \frac{e^{-(\nu - \bar{\nu})^2 / (2\sigma_\nu^2)}}{\sqrt{2\pi}\sigma_\nu}. \quad (5.21)$$

Finally, marginalizing over ν gives the posterior distribution of R :

$$P(R | N, \bar{\nu}) = \int_0^{\nu_{\max}} P(R, \nu | N, \bar{\nu}) d\nu. \quad (5.22)$$

Using this distribution, the 90% limit on R is given by $R_{90\%, \text{VT}}$, such that 90% of the posterior mass is enclosed. This can be determined by numerically solving the following equation:

$$0.9 = \int_0^{R_{90\%, \text{VT}}} P(R | N, \bar{\nu}) dR. \quad (5.23)$$

The resulting upper limits are conservative since the associated uncertainties have been included using this method.

5.4.3 Efficiency calculation

For signal models without a physical distance calibration, I use the loudest event statistic to calculate rate upper limits at 90% confidence based on the search pipeline’s efficiency:

$$R_{90\%,T} = \frac{2.3}{\sum_k \epsilon_k(\text{SNR}_{\Gamma,k}^*) \times T_{\text{obs},k}}. \quad (5.24)$$

Here, the index k runs over searches (S5 and S6 in our case), and ϵ_k is the efficiency of search k (for a given waveform) using a detection threshold $\text{SNR}_{\Gamma,k}^*$, which corresponds to the loudest zero-lag trigger from search k . The resulting upper limits are in terms of number of observations per time, as indicated by the subscript T. These upper limits are a function of the signal strength (in terms of either source distance or h_{rss}), because the efficiency depends on signal strength.

Again, amplitude calibration uncertainty is the primary source of uncertainty for these upper limits. In this case, I account for this uncertainty by adjusting the signal amplitudes used in the sensitivity study (and shown in the efficiency curves in Figs. 5.6 and 5.7) upward by a multiplicative factor corresponding to the respective 1σ amplitude calibration uncertainty.

5.4.4 Amplitude calibration uncertainties

During the S5 science run, the amplitude calibration uncertainty was measured to be 10.4% and 14.4% for the H1 and L1 detectors, respectively, in the 40–2000 Hz frequency band [309]. Assuming that the individual detector uncertainties are uncorrelated, the quadrature sum gives a total calibration uncertainty of 17.8% on the amplitude, and thus, an uncertainty of 53.4% on the visible volume. For S6, the amplitude calibration uncertainty was measured at 4.0% for H1 and 7.0% for L1 in the 40–2000 Hz band [310], resulting in a total calibration uncertainty of 8.1% on the amplitude and 24.2% on the visible volume. These uncertainties are included in both of the upper limit calculations using the methods described in the previous subsections.

5.4.5 Final upper limits

Here, I present final upper limits using the formalisms developed in the previous subsections. The upper limits given are conservative and include both statistical and systematic uncertainties as described previously. We again note that the dominant source of uncertainty for both the FAD-based and efficiency-based upper limits is the amplitude calibration uncertainty of the detectors.

The VT upper limits are only calculated for the ADI waveforms since they are the only distance-calibrated waveforms used in this search. These upper limits are shown in Tab. 5.6. For all waveforms, I have calculated loudest event statistic upper limits in terms of the pipeline efficiency using Eq. 5.24. These upper limits are a function of signal strength and are shown in Fig. 5.9.

Waveform	V_{vis} [Mpc ³]		$R_{90\%,\text{VT}}$ [Mpc ⁻³ yr ⁻¹]
	S5	S6	
ADI-A	1.8×10^3	3.6×10^3	9.4×10^{-4}
ADI-B	5.7×10^4	9.1×10^4	3.4×10^{-5}
ADI-C	7.8×10^3	1.6×10^4	2.2×10^{-4}
ADI-E	1.6×10^4	3.2×10^4	1.1×10^{-4}

Table 5.6: Rate upper limits on ADI waveforms calculated with Eq. 5.14. The visible volume of each search is shown to illustrate relative search sensitivities and does not include uncertainties (53% and 24% for all waveforms in S5 and S6, respectively). However, 1σ uncertainties on the visible volume are included in the resulting upper limits using the Bayesian method described in Sec. 5.4.2.

5.5 Discussion

Given the absence of detection candidates in the search, I have calculated and presented upper limits on the event rate for different GW signal families. Specifically, Figure 5.9 shows that, along with signal morphology, both the frequency and the duration of a signal influence the search sensitivity. The h_{rss} values for a search efficiency of 50% obtained with S5 and S6 data are reported in Tabs. 5.7 and 5.8; for the ADI waveforms, the 50% efficiency distance is also given. Although these limits cannot be precisely compared to




Figure 5.9: Loudest event statistic upper limits for the 15 simulated GW signals used to test the sensitivity of the search (calculated with Eq. 5.24). 1σ calibration uncertainties are included by increasing the signal amplitudes by the corresponding multiplicative factor. Top-left: ADI waveforms. Top-right: sinusoidal waveforms. Bottom-left: sine-Gaussian waveforms. Bottom-right: white noise burst waveforms. Reprinted figure with permission from Ref. [142], copyright 2016 by the American Physical Society.

the results of unmodeled short transient GW searches [107] because different waveforms were used by the two searches, it is clear that in order for long transient GW signals to be observed, it is necessary for the source to be more energetic: the total energy radiated is spread over hundreds of seconds instead of a few hundred milliseconds.

Waveform	$h_{\text{rss}}^{50\%}$ [$\text{Hz}^{-1/2}$]	distance ^{50%} [Mpc]	E_{GW} [$M_{\odot}c^2$]
ADI-A	1.8×10^{-21}	5.4	1.5×10^{-7}
ADI-B	1.9×10^{-21}	16.3	1.9×10^{-7}
ADI-C	3.6×10^{-21}	8.9	9.9×10^{-7}
ADI-E	2.3×10^{-21}	11.5	3.3×10^{-7}
MONO-A	1.3×10^{-20}	...	3.1×10^{-6}
MONO-B	2.1×10^{-20}	...	2.3×10^{-4}
LINE-A	3.9×10^{-21}	...	4.9×10^{-7}
LINE-B	8.5×10^{-21}	...	9.6×10^{-5}
QUAD-A	2.7×10^{-21}	...	2.4×10^{-7}
QUAD-B	5.1×10^{-21}	...	1.6×10^{-5}
SG-A	2.4×10^{-20}	...	1.0×10^{-5}
SG-B	2.1×10^{-20}	...	2.5×10^{-4}
WNB-A	7.2×10^{-21}	...	5.5×10^{-6}
WNB-B	9.1×10^{-21}	...	1.8×10^{-5}
WNB-C	2.1×10^{-20}	...	5.0×10^{-4}

Table 5.7: Values of h_{rss} and distance where the search achieves 50% efficiency for each of the simulated GW signals studied in the S5 dataset. E_{GW} is an estimate of the energy released (given by Eq. 5.25) by a source located at the detection distance for the ADI waveforms, or at 10 kpc for the *ad hoc* waveforms.

Energy estimates

We can estimate the amount of energy emitted by a source located at a distance r where the search efficiency drops below 50% ($h_{\text{rss}}^{50\%}$) using the quadrupolar radiation approximation to estimate the energy radiated by a pair of rotating point masses:

$$E_{\text{GW}} \approx (h_{\text{rss}}^{50\%})^2 r^2 \pi^2 f_{\text{GW}}^2 \frac{c^3}{G}. \quad (5.25)$$

Considering the mean frequency of each GW waveform, we obtain an estimate of the energy that would have been released by a source that would be detected by this search. For the ADI waveforms, the corresponding energy is between $9 \times 10^{-8} M_{\odot}c^2$ and $6 \times 10^{-7} M_{\odot}c^2$ (refer to Tabs. 5.7 and 5.8). For the *ad hoc* waveforms, we must fix a fiducial distance at which we might expect to observe a signal. For instance, considering

Waveform	$h_{\text{rss}}^{50\%}$ [$\text{Hz}^{-1/2}$]	distance ^{50%} [Mpc]	E_{GW} [$M_{\odot}c^2$]
ADI-A	1.4×10^{-21}	6.8	9.4×10^{-8}
ADI-B	1.7×10^{-21}	18.6	1.5×10^{-7}
ADI-C	2.9×10^{-21}	11.3	6.3×10^{-7}
ADI-E	2.0×10^{-21}	13.4	2.4×10^{-7}
MONO-A	1.2×10^{-20}	...	2.4×10^{-6}
MONO-B	1.1×10^{-20}	...	7.0×10^{-5}
LINE-A	3.1×10^{-21}	...	3.1×10^{-7}
LINE-B	5.2×10^{-21}	...	3.7×10^{-5}
QUAD-A	2.4×10^{-21}	...	1.9×10^{-7}
QUAD-B	2.9×10^{-21}	...	5.3×10^{-6}
SG-A	1.9×10^{-20}	...	6.2×10^{-6}
SG-B	1.0×10^{-20}	...	5.9×10^{-5}
WNB-A	5.5×10^{-21}	...	3.2×10^{-6}
WNB-B	6.4×10^{-21}	...	9.2×10^{-6}
WNB-C	1.3×10^{-20}	...	1.7×10^{-4}

Table 5.8: Values of h_{rss} and distance where the search achieves 50% efficiency for each of the simulated GW signals studied in the S6 dataset. E_{GW} is an estimate of the energy released (given by Eq. 5.25) by a source located at the detection distance for the ADI waveforms, or at 10 kpc for the *ad hoc* waveforms.

a Galactic source at 10 kpc, the emitted energy would be in the range 2×10^{-7} – $2 \times 10^{-4} M_{\odot}c^2$. This is still 2–4 orders of magnitude larger than the amount of energy estimated in Ref. [254] ($4 \times 10^{-9} M_{\odot}c^2$) for protoneutron star convection over timescales of ≈ 30 s.

Future work

We note that the search for long-duration transient signals is also closely related to the effort by LIGO and Virgo to observe a stochastic background of GWs. One or more long-lived transient GW events, with a duration of days or longer, could produce an apparent signal in either the isotropic [88, 89, 311] or directional [140] stochastic GW searches. It was for this reason that this long-duration transient search pipeline was originally developed [18]. A dedicated search for long-duration transient GW signals which last for days or longer will be a necessary component in the effort to understand

the origin of apparent stochastic background signals which may be observed by LIGO and Virgo in the future [312].

After 2010, the LIGO and Virgo interferometers went through a series of upgrades [52, 53]. LIGO has just concluded its first observational campaign with its advanced configuration and will be joined by Virgo in 2017 [313]. The strain sensitivity of the advanced detectors is expected to eventually reach a factor of 10 better than the first-generation detectors. This development alone should increase the distance reach of our search by a factor of 10, the energy sensitivity by a factor of 100, and the volume of space which we can probe by a factor of 1000.

Improvements are also being made to this search pipeline; a technique for generating triggers called “seedless clustering” has been shown to increase the sensitivity of the search by 50% or more in terms of distance [314–319]. These improvements to the search pipeline, coupled with the increased sensitivity of LIGO and Virgo, will drastically improve the prospects for detecting long-duration transient GWs.

5.6 Conclusions

In this chapter, I have presented an all-sky search for long-lasting GW transients with durations between a few seconds and a few hundred seconds. I performed the search on data from the LIGO H1 and L1 detectors collected during the S5 and S6 science runs, using a cross-correlation pipeline to analyze the data and identify potential GW candidate triggers. To reject high-SNR_T triggers due to detector noise, I developed a discriminant cut based on the trigger morphology; along with this, LIGO data quality flags were used to veto well-identified instrumental or environmental noise sources to remove significant outliers. No GW candidates were identified in this search, and as a result, I set upper limits on several types of simulated GW signals. These are the first upper limits from an unmodeled all-sky search for long-transient GWs (see Tab. 5.6 and Fig. 5.9).

Chapter 6

Seismic noise and gravitational-wave detectors

Below ≈ 10 Hz, seismic noise is currently the dominant noise source for terrestrial gravitational-wave detectors (see Fig. 2.3). Seismic noise affects interferometric GW detectors in two ways: first, ground motion due to seismic noise causes the interferometer support structures and mirrors to move, effectively changing the length of the optical cavities and inducing strain. Second, density fluctuations in the rock and in the atmosphere due to seismic waves cause changes in the local gravitational field, which couple to the mirrors and induce length changes in the arms of the interferometer; this is commonly referred to as Newtonian noise or gravity-gradient noise.

For current gravitational-wave detectors, actual ground motion is the primary source of seismic noise. Newtonian noise has never actually been observed in GW detectors, but is predicted to be a limiting noise source for future GW detectors, which are expected to have improved sensitivity at lower frequencies [320].

In this chapter, I will describe an experiment which aims to estimate the seismic and Newtonian noise contributions at different depths underground using a three-dimensional array of seismometers. I begin with an introduction to seismic waves in Sec. 6.1 and a discussion of seismic noise in interferometric GW detectors in Sec. 6.2. In Sec. 6.3, I describe the implementation of a three-dimensional seismic array in the Homestake Mine in Lead, SD, along with associated details regarding station operation and data

collection. In Sec. 6.4, I introduce a radiometer method for identifying directional sources of seismic waves and separating them into their components. I apply this radiometer method to simulations of seismic waves and to real data. Finally, in Sec. 6.5, I use measurements of Rayleigh wave amplitudes to estimate parameters for a new model of Rayleigh wave eigenfunctions. Sec. 6.6 summarizes the results of this chapter.

6.1 Seismic waves

Seismic waves are elastic waves which carry energy through the earth and are commonly produced by an impulse, such as an earthquake or an explosion. These waves span a range of frequencies, from the lowest free oscillations of the Earth at $\approx 3 \times 10^{-4}$ Hz up to a few Hz for seismic waves generated by earthquakes [321]. Anthropogenic activities may create seismic waves at frequencies as high as 100 Hz, although these frequencies are quickly damped due to the length scale of anisotropies in the Earth's crust. The overall background of seismic noise generally peaks between 0.1–0.2 Hz, which is commonly referred to as the microseismic peak. This phenomenon arises due to the transfer of energy between resonant standing waves on the ocean and the crust, which may be driven by atmospheric pressure and wind [321].

Seismic waves may be generally divided into two classes: body waves and surface waves. In this chapter, I will describe the main types of seismic waves and their properties.

6.1.1 Body waves

Body waves propagate through the interior of the earth and generally travel more quickly than surface waves. There are two fundamental ways in which strain may be induced in a body: volume change without a change of shape, and shape change without a change in volume [321]. As a result, there are two classes of body waves: longitudinal (compressional) waves and transverse (shear) waves. The longitudinal waves cause compression and rarefaction along the direction of wave propagation, while the transverse waves produce shear and rotation perpendicular to the direction of propagation.

Longitudinal waves are also referred to as P-waves (for *primary*) because they travel more quickly than shear waves and are the first to arrive after an earthquake. In an

isotropic medium, the P-wave speed may be described in terms of the Lamé parameters λ and μ :

$$v_p = \sqrt{\frac{\lambda + 2\mu}{\rho}}. \quad (6.1)$$

Here, the first Lamé parameter λ is related to the bulk modulus K of the material by

$$\lambda = K - \frac{2}{3}\mu, \quad (6.2)$$

and the second Lamé constant μ is the shear modulus of the medium, while ρ is its mass density. P-wave velocities vary from 2–5 km/s in sedimentary rocks and may be as high as 6 km/s in the upper crust [321]. Since P-waves are the first to arrive after a seismic event, they are the most accurately measured type of seismic wave and are routinely used to locate the epicenter of an earthquake.

Transverse waves are often called S-waves, which refers both to the fact that they are the second type of seismic wave to arrive after an earthquake, and to the shearing motion they produce perpendicular to the direction of wave propagation. The S-wave velocity is given by

$$v_s = \sqrt{\frac{\mu}{\rho}}. \quad (6.3)$$

Since the bulk and shear moduli always take non-negative values, it follows that $v_p > v_s$; this is the primary motivation for the naming convention of the two classes of seismic body waves. On average, the S-wave velocity is approximately 60% of the P-wave velocity, although this depends on the properties of the medium. S-waves may provide significant insights into material rigidity and fluid content, since a perfect fluid has no shear modulus [321].

In anisotropic media, P-waves may not be entirely longitudinal and S-waves may not be entirely transverse. In this case, these waves may be referred to as “quasi-P” and “quasi-S,” respectively.

6.1.2 Surface waves

Surface waves are seismic waves which travel relatively closely to the surface of the Earth (compared to body waves). When a seismic event occurs within approximately a

wavelength or so of the surface, a considerable fraction of the energy may be trapped in a waveguide formed by the surface layers, resulting from the interference, reflection, and refraction of waves interacting with the surface layer [321]. The existence of surface waves can be derived from the equations of motion in a continuous medium:

$$\rho \frac{\partial^2 u_i}{\partial t^2} = \frac{\partial \tau_{ij}}{\partial x_j}, \quad (6.4)$$

where τ_{ij} is the stress tensor, u_i is the displacement in the i th direction, and x_i is the i th coordinate. We assume a simple relation between stress and strain in a material, with one component due to volumetric strain and one due to shear strain:

$$\tau_{ij} = \lambda \epsilon_{kk} \delta_{ij} + 2\mu \epsilon_{ij}. \quad (6.5)$$

Here, summation is implied over repeated indices; ϵ_{ij} is the strain tensor, given by

$$\epsilon_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right). \quad (6.6)$$

This is also known as the strain-to-displacement relationship.

From these equations of motion, it can be shown that there are two possible types of surface waves: Love waves, which comprise only a horizontally-polarized transverse component, and Rayleigh waves, which contain a longitudinal component and a transverse component in the vertical direction. The three displacement components for a Love wave are given by

$$\begin{aligned} u_1 &= 0 \\ u_2 &= l_1(z, k, \omega) \cos(kx_1 - \omega t) \\ u_3 &= 0, \end{aligned} \quad (6.7)$$

where the coordinate system is aligned such that x_1 is the direction of wave propagation and $z = x_3$ is the vertical direction. Here, k is the wavenumber, ω is the angular frequency, and l_1 is the eigenfunction solution obtained by solving the equations of motion with this functional form. For a homogeneous half-space, where the shear modulus is constant, Love waves cannot be generated; rather, they require a surface layer of lower

velocity. If the shear modulus is a function only of depth (i.e., $\mu = \mu(z)$), the result is a first-order ordinary differential equation for l_1 :

$$-\rho\omega^2 l_1 = -\mu(z)k^2 l_1 + \frac{\partial}{\partial z} \left[\mu(z) \frac{\partial l_1}{\partial z} \right]. \quad (6.8)$$

This equation can be numerically solved for a specific functional form of $\mu(z)$.

Rayleigh waves may be described by the following three displacement components:

$$\begin{aligned} u_1 &= r_1(z, k, \omega) \cos(kx_1 - \omega t) \\ u_2 &= 0 \\ u_3 &= r_2(z, k, \omega) \sin(kx_1 - \omega t), \end{aligned} \quad (6.9)$$

where r_1 and r_2 are the horizontal and vertical eigenfunctions, respectively.

Under the assumption that the material properties are only functions of depth (i.e., $\lambda = \lambda(z)$ and $\mu = \mu(z)$), the system of ODEs resulting from the use of this functional form in the equations of motion can be solved. For the case of a homogeneous half-space, the solutions are analytic:

$$\begin{aligned} r_1 &= C_1 e^{-a_1 kz} + C_2 e^{-a_2 kz} \\ r_2 &= C_3 e^{-a_3 kz} + C_4 e^{-a_4 kz}. \end{aligned} \quad (6.10)$$

The C and a constants are parameters which depend on the material properties of the crust in a given region.

One important feature of surface waves is their velocity dispersion: the surface wave velocity is frequency-dependent. In general, lower frequency surface waves, which have a longer wavelength and are more sensitive to deeper material, travel more quickly than higher frequency waves.

Surface wave velocities are generally less than those of body waves, partially due to the presence of lower density rock and sediments in the upper layers of the crust. Rayleigh waves are the slowest of the surface waves, while Love waves may travel as quickly as $0.9 v_s$ [321].

6.2 Seismic noise in GW interferometers

Seismic noise is a significant obstacle which is currently preventing terrestrial GW interferometers from achieving better sensitivity at low frequencies. The current generation of GW interferometers is sensitive down to about 10 Hz, and the next generation may reach as low as 1 Hz. As seismic noise due to ground motion is mitigated, Newtonian noise is expected to provide a significant contribution to the low-frequency detector noise.

There is significant motivation for reducing seismic noise in GW interferometers in order to achieve better low-frequency sensitivity. First, many models of a stochastic gravitational-wave background predict considerable power at low frequency. This is especially true for cosmological SGWBs, which should be produced in the very early universe and would be significantly redshifted today. Along with this, some astrophysical SGWBs may have significant power around or even below 1 Hz. The LIGO stochastic search is also more sensitive at low frequencies, since the SNR described in Eq. 3.24 is proportional to f^{-3} . Increased sensitivity at lower frequencies may also lead to significantly longer and more detailed observations of compact binary coalescences (currently only the final chirp is in the LIGO band).

Several different strategies may be employed to mitigate seismic noise. In the LIGO detectors, the test masses are suspended as pendula, which provides $1/f^2$ damping above the resonant frequency. The interferometer support structures are mounted on isolation stacks, which use active and passive vibration isolation in the form of alternating layers of springs and metal slabs. More details on the initial LIGO configuration are available in Sec. 2.1.1. The Advanced LIGO configuration features quadruple pendula for the mirror suspensions, with each pendula providing $1/f^2$ damping, and an active isolation system which uses seismometers and geophones to drive hydraulic actuators in order to counteract ground motion [322].

However, the pendulum suspension is only practical for vibration isolation above ≈ 1 Hz, since the resonant frequency of the pendulum is inherently tied to its length. As an example, a pendulum with a resonant frequency of 0.1 Hz would be almost 25 m long. Other methods of decreasing the resonant frequency of the pendulum, including magnetic damping, may be explored for future GW detectors.

The strategies mentioned so far are most effective for dealing with ground motion due

to seismic noise, rather than Newtonian noise. One option for dealing with Newtonian noise is the construction of moats around each test mass. Hughes and Thorne [323] studied this technique and suggested that to be effective, these moats should be ≈ 9 – 15 m deep and ≈ 20 – 35 m in radius in order to suppress Rayleigh waves around 10 Hz. To suppress lower frequency seismic waves with longer wavelengths, the moats should be wider and deeper. Hughes and Thorne also note that this strategy would be most effective at suppressing shear waves, since they cannot travel through non-rigid media; compressional waves could still be transmitted through the moats to some degree.

It may also be possible to mitigate Newtonian noise through an active noise cancellation technique called Wiener filtering [320, 324]. A Wiener filter uses correlations between multiple sensors to estimate the data in another sensor at a different location; the effectiveness of the filter depends on the data correlations, the frequency range of interest, and optimization of the filter parameters [324]. For a GW interferometer, the goal would be to use an array of seismometers to estimate the seismic field at each test mass, convert to strain, and subtract that contribution from the strain data. Using data from the DUGL array, Wiener filtering has achieved effectively an order of magnitude of seismic noise suppression in the 0.05–1 Hz band [324], with even better results expected from the new array at Homestake (see Sec. 6.3). We note that this suppression is achieved with a seismometer as the target sensor; for use in an actual GW detector, the transfer function and the couplings between the seismometers and the GW detector would be different, so the given results are not directly transferable.

Because surface waves are the dominant type of seismic wave and are exponentially attenuated with increasing depth, it is expected that an underground GW interferometer would be subject to less overall seismic noise compared to a detector on the surface. Going to depth should also suppress Newtonian noise from ground-atmosphere density perturbations, which is expected to be significant [320]. With this goal in mind, the KAGRA collaboration is currently constructing a 3 km GW interferometer in the Kamioka mine, which is expected to begin operation in 2018 [54]. However, a space-based detector is the clearest strategy for completely eliminating seismic noise. This idea is being pursued by the eLISA project, which has recently launched a pathfinder mission to demonstrate their technology [58, 59].

To test underground conditions for future generations of gravitational-wave detectors,

we have planned and installed a three-dimensional array of seismometers in and around the Homestake Mine in Lead, SD. This array will be described in the following section.

6.3 Three-dimensional seismic array at Homestake

The Homestake Mine in Lead, SD was one of the largest and deepest gold mines in North America and has a rich history of scientific involvement. In the 1960s, scientific experiments first came to the mine, including Ray Davis' famous observation of solar neutrinos, which led to the solar neutrino problem and the eventual discovery of neutrino oscillations. The mine officially closed in 2002, but reopened in 2007 as the Sanford Underground Research Facility (SURF) thanks to donations from T. Denny Sanford and funding from the NSF. SURF currently features several dark matter and neutrino experiments (primarily at the 4850-ft level), which benefit from the cosmic ray shielding provided by thousands of feet of rock.

Since seismic surface waves are typically the largest source of seismic noise, going underground may also provide significant benefit for gravitational-wave detectors, as these waves are damped exponentially with depth. Being underground also reduces the potential Newtonian noise due to atmospheric perturbations and interactions at the ground-atmosphere interface. To investigate the effect of depth on seismic noise, an array of seismometers was installed in the Homestake Mine and operated from 2008 to 2010 [325]; this experiment was called the Deep Underground Gravity Laboratory (DUGL). Three stations were initially installed and three others were added soon after; there were originally nine planned stations, but three were not used due to poor operating conditions. In this section, I describe the second generation of this experiment, which was developed in collaboration with geophysicists from Caltech and Indiana University, and features a significantly expanded array in terms of aperture, depth, and number of stations, as well as more robust and accurate data acquisition systems and improved timing accuracy.

6.3.1 Array description

The current version of the seismic array at Homestake features 24 seismic stations in total, with 15 underground stations and 9 surface stations. Six of the surface stations

are on SURF property near the mine, while three others are on private land and are further separated from the other stations in terms of horizontal distance. All of the stations were installed between September 2014 and May 2015 and the array is expected to collect data until December 2016, resulting in about 1.5 years of data for the entire array. However, some surface stations may be removed in the fall of 2016 due to weather concerns.

Station information

In Tab. 6.1, I have listed each station by its ID and provided the station type and a short description of its location. Note that 4 stations were originally planned for the 4100 level, but only three were implemented due to difficulties installing power and networking capabilities at one of the proposed sites. Due to the chronology of the planning and installation process, the three stations on that level are named A4100, C4100, and D4100; there is no B4100 station.

The coordinates of each station are given in Tab. 6.2; uncertainties for these coordinates are on the order of ≈ 2 m. The array layout is shown in Fig. 6.1. Underground station locations were obtained from maps of the mine drifts based on work by mine surveyors, while surface station coordinates come from GPS data.

6.3.2 Station configuration

Each seismic station includes a seismometer, a data acquisition system, and auxiliary electronics related to power and timing. Nearly all of the equipment used in the experiment has been rented from the Portable Array Seismic Studies of the Continental Lithosphere (PASSCAL) instrument center, which is part of the Incorporated Research Institutions for Seismology (IRIS).

Most stations use a Streckheisen STS-2 high-sensitivity broadband seismometer. The following stations have potential moisture concerns and, as a result, use a more water-resistant Guralp CMG-3T broadband seismometer: DEAD, ROSS, YATES, and 300. Overall, the STS-2 seismometers are more robust and less “finicky” than the Guralps, based on anecdotal evidence.

At each station, the seismometer is directly connected to a Quanterra Q330 datalogger (shown in Fig. 6.2), which powers the seismometer and digitizes the analog data. The




Figure 6.1: Homestake array layout, relative to YATES station (at origin). A zoomed-in view with the underground stations labeled is shown on top (DEAD, SHL, and TPK not pictured). All stations are shown in the lower plot, but only the surface stations are labeled.

Station ID	Station Type	Station Description
DEAD	Surface	Private landowner's yard in Deadwood
LHS	Surface	At Lead-Deadwood High School
ORO	Surface	Near Oro Hondo shaft
ROSS	Surface	Near Ross warehouse
RRDG	Surface	Rhyolite Ridge
SHL	Surface	Private landowner's yard on Strawberry Hill
TPK	Surface	Private landowner's yard near Terry Peak
WTP	Surface	Near Yates water treatment plant
YATES	Surface	Near Yates administration building
300	Underground	300 level, near side entrance
800	Underground	800 level, near Ross shaft
1700	Underground	1700 level, near Ellison shaft
A2000	Underground	2000 level, near ramp at Y intersection
B2000	Underground	2000 level, near Ross lunch room
C2000	Underground	2000 level, 7 ledge, A limb
D2000	Underground	2000 level, top of ramp
E2000	Underground	2000 level, north 9 ledge
A4100	Underground	4100 level, powder magazine west of Oro Hondo
C4100	Underground	4100 level, Bill Roggenthen's room
D4100	Underground	4100 level, Yates station old locomotive barn
A4850	Underground	4850 level, 17 ledge
B4850	Underground	4850 level, roll-up door
C4850	Underground	4850 level, 4 Winze Wye
D4850	Underground	4850 level, Davis incline

Table 6.1: Station IDs and descriptions for the 24 seismic stations in the second-generation Homestake 3D array.

Q330 is capable of remote telemetry using an ethernet connection or a radio antenna; in this way, digitized data packets may be sent to a remote computer for further processing and storage. Q330 datalogging parameters may be programmed using Willard,¹ a GUI interface provided by Quanterra. The Q330 can also save data to a local hard drive through a serial connection; in the field, we use Quanterra PB14 Packet Balers, which are rugged and water-resistant enclosures featuring a serial interface, a 20 GB hard drive, and some limited networking capability. A photograph of a Baler is shown in Fig. 6.3.

In order to synchronize the data from the entire array to microsecond accuracy, each

¹Willard is designed to run in Windows, but may be run in Linux using `Wine`.



Figure 6.2: Photo of a Quanterra Q330 datalogger. The seismometer is connected through one of the SENSOR ports on the left. A GPS antenna may be connected to the GPS ANT port or an external GPS source may be connected to the EXT GPS port. The Q330 is connected to a Baler via the SERIAL port and/or to a local network through the QNET port. The Q330 is powered either via the POWER port or through the QNET port. All connections (other than GPS ANT) are military-type; more details can be found in the Q330 manual [326].



Figure 6.3: Photo of a Quanterra PB14 Packet Baler. The Baler is connected to a Q330 via the QNET + SERIAL port and receives its power through this connection. The ATTN button is used to “wake” the Baler from its low-power state; when the Baler has completed its boot cycle, the Q330 it is connected to will automatically begin transferring data to it.

Station ID	Easting [m]	Northing [m]	Elevation [m]	Install Date
DEAD	-187	3386	-127.0	05/18/2015
LHS	-1849	-572	59.0	05/19/2015
ORO	-49	-967	-82.0	03/30/2015
ROSS	-474	-796	3.0	11/19/2014
RRDG	-1120	794	52.0	05/20/2015
SHL	3385	-3869	147.0	05/17/2015
TPK	-3687	-1322	115.0	05/19/2015
WTP	730	200	-70.0	01/22/2015
YATES	0	0	0.0	11/19/2014
300	-421	-651	-119.9	01/13/2015
800	-582	-542	-275.0	01/15/2015
1700	-549	-64	-551.2	03/25/2015
A2000	-859	-135	-641.7	03/24/2015
B2000	-712	-344	-641.7	03/25/2015
C2000	-971	-81	-641.9	03/25/2015
D2000	-1389	101	-642.0	03/25/2015
E2000	-1609	442	-641.9	03/25/2015
A4100	-147	-813	-1282.5	01/21/2015
C4100	65	-110	-1282.7	01/21/2015
D4100	46	-955	-1282.6	01/21/2015
A4850	-857	-1313	-1509.8	04/02/2015
B4850	-516	-663	-1510.1	11/19/2014
C4850	-94	-657	-1510.4	11/19/2014
D4850	78	90	-1509.8	11/19/2014

Table 6.2: Locations of the seismic stations in the Homestake 3D array. We use the UTM coordinate system in zone 13T; all locations are relative to the YATES surface station. The YATES station has absolute coordinates of (Easting, Northing, Elevation) = (599504, 4911750, 1625.0) m. Surface station locations are obtained from GPS data, while underground coordinates are based on maps of the mine drifts. Uncertainties are on the order of ≈ 2 m.

station receives a GPS timing signal. The GPS timing and power configurations differ for surface stations and underground stations and will be described further in the following sections. In order to ensure uniformity across stations, the seismometers are carefully oriented with respect to geodetic north. For the surface stations, this alignment was accomplished with a compass (accounting for magnetic declination); underground, we have used an survey-grade gyrocompass called an OCTANS [327].

Surface stations

For the surface stations, we have placed the seismometers in holes which are approximately $4' \times 4'$ wide and 3' deep in order to reach bedrock. Each seismometer is situated on a leveled concrete cylinder, which is about 6" tall and 1' in diameter. The seismometer is surrounded by a "vault" featuring an inner layer of foam board insulation and an outer layer of timber in order to provide thermal and acoustic insulation; the edges of the inner layer are sealed with insulating foam. A photograph of this arrangement at the DEAD station (not including the inner insulating layer) can be seen in Fig. 6.4.



Figure 6.4: Seismometer vault enclosure at DEAD surface station. The timber enclosure and the concrete cylinder are shown, but the layer of foam board insulation is not pictured.

After installation of the seismometer, a sheet of plywood is screwed down to the top of the timber vault and the hole is filled back in. Additional dirt is piled on top of the vault in order to improve thermal and acoustic insulation. A sealed vault at the TPK station can be seen in Fig. 6.5. The seismometer cable is inside of the blue conduit and runs from the vault to the Q330 (not pictured). The conduit provides some level of protection from the weather and from curious animals.



Figure 6.5: Seismometer vault at TPK surface station, sealed with plywood. The seismometer cable is inside of the blue conduit.

The Q330, Baler, and other electronics are contained in a mostly sealed enclosure on the surface. For five of the surface stations, large doghouses were used as enclosures since they provided some measure of ventilation, protection from the elements, and water drainage capabilities (see Fig. 6.6). The other four surface stations used large military-grade totes to hold the electronics.

All of the surface stations used GPS antennas connected to the GPS ANT port of the Q330 to provide a timing signal. The GPS antennas were attached to the top of a $\approx 7'$ tall wooden post in order to provide an unobstructed view of the sky. The GPS antenna is visible in Fig. 6.6 as the small hemispherical object on the very top of the wooden post (just above the green radio antenna). Overall, timing for the surface stations was very consistent, generally within $\pm 5 \mu\text{s}$.

Each surface station was powered by two solar panels which were approximately $3' \times 1'$ and mounted to the wooden post. The panels were secured with guy lines and oriented to maximize the surface area exposed to the sun over the course of a day. These panels were connected to a power distribution box which connected to the Q330 and to a 55 A-h battery, which was used to store power for overnight operation and during cloudy periods. Only a few of the surface stations (primarily WTP and DEAD) had difficulty



Figure 6.6: Full RRDG station, including the sealed seismometer vault (under dirt pile), electronics doghouse (beige igloo near center), solar panels (middle right, attached to center of wooden post), GPS antenna (on top of wooden post), and radio antenna (green cylinder just below GPS antenna).

due to lack of sun exposure.

The five surface stations located on the mine property (ORO, ROSS, RRDG, WTP, YATES) included a radio antenna (mounted on the wooden post; shown as a horizontal green cylinder in Fig. 6.6), which was connected to the QNET port of the Q330 and transmitted data to a receiver antenna mounted on the SURF administration building. This receiver was connected to a central computer in the administration building server room which handled the incoming data packets and stored them in a database (see Fig. 6.9). The LHS station also featured a radio antenna, which transmitted data to a receiver inside Lead-Deadwood High School and then to a computer, which processed the data from that station. The three remaining surface stations were on private property and several kilometers from the mine, so telemetry was not possible. As a result, the data were stored locally at each station using Balers and required manual retrieval.

Underground stations

For the underground stations, the seismometers were situated on leveled granite tiles which were securely attached to concrete slabs. We used existing concrete at stations for which it was available and poured new concrete slabs at other stations. To provide thermal and acoustic isolation, each seismometer was surrounded by two “huts” constructed of 1" thick foam insulation board (shown in Fig. 6.7).




Figure 6.7: Seismometer enclosure used for underground stations (C2000 station shown here). The seismometer is on a leveled granite tile which is grouted to a concrete slab. The two foam huts are sealed with insulating foam sealant. Each hut also includes a top cover which is not in place in this picture.

The data acquisition system and other electronics at each station were usually set on a table 5–10 feet from the seismometer in order to minimize electronic noise. They were also covered with a plastic tub to reduce potential exposure to moisture and dust. Fig. 6.8 shows an example of the underground electronics configuration at the B4850 station.

The underground stations all had access to AC power; however, the electronics were



Figure 6.8: Data acquisition system and electronics for the underground sites (B4850 station shown here). The Q330 and Baler are shown on a wooden table along with a power box (on top of Q330, to left of Baler) and a GPS transceiver (small black object on top of Baler). The battery is in the purple tub on the ground.

directly connected to small (≈ 12 A-h) batteries. The batteries were kept in a high-charge state using low-power “battery tenders” which were connected to AC power. This configuration allowed for continued data collection for approximately a day in case of dropouts in AC power due to mine construction work or any other unforeseen issues. The batteries were typically placed in a plastic tub separate from the electronics in case of corrosion.

A GPS timing signal was available at each underground station through a network of optical cables, splitters, and transceivers which was connected to a GPS antenna mounted on the roof of the SURF administration building. Each station featured a GPS transceiver (visible in Fig. 6.8 as the small black box on top of the Baler; also shown in Fig. 6.12) which connected directly to the EXT GPS port of the Q330. More details about the GPS transceivers and the timing network in the mine are provided in Sec. 6.3.4.

All of the underground stations were directly connected to one of the mine’s networks through an ethernet cable and transmitted data to the central computer. Each station also featured a Baler for storing data locally in case of network connection issues.

For some of the underground stations (300, 800, A2000, B2000, A4100, and D4100), we were able to use pre-existing sites from the original DUGL experiment.

6.3.3 Data acquisition

The central computer in the SURF administration building runs the Antelope real time system for acquiring and storing seismic data [328]. The Q330s at each station send data packets over the network which are collected by a `q3302orb` process. The data packets from each station are then stored in a ring buffer which is managed by an `orbserver` process. An `orb2orb` process may be used to send the data in the ring buffer to another ring buffer operating on a separate machine. For this experiment, the `orb2orb` process at SURF sends the data to another `orbserver` operating at UMN. The data are then pulled from UMN to Caltech and IU using `orb2orb` processes at those locations. There is also an instance of Antelope operating on a computer at Lead-Deadwood High School in order to collect the LHS station data. Due to restrictions on the school district’s network, these processes bypass the SURF computer and use `orb2orb` to send the data directly to UMN’s `orbserver`. An `orb2db` process is used at both SURF and UMN to write the data in the ring buffer to a local database. A chart describing the real-time data flow of the experiment is shown in Fig. 6.9.

The `orb2db` processes are configured to write the data into “dayvolumes,” which means that a single file is saved for each station and each channel every day. Along with the seismic channels, there are many state-of-health channels for the seismometer, Q330, and Antelope processes, which are described further in Tab. 6.3. Each dayvolume is saved in `miniSEED` format and the following file naming convention is used for the dayvolumes: `{station_name}.{experiment_net_code}..{channel_name}.{year}.{day_of_year}`, where the net code for this experiment is `X6`².

At UMN, we convert the `miniSEED` dayvolumes into gravitational-wave data frames (`.gwf` files) following the specification in Ref. [329].

²The two periods in a row arise due to a null parameter value.

Channel code	Description
HHE	East seismic channel (100 Hz)
HHN	North seismic channel (100 Hz)
HHZ	Vertical seismic channel (100 Hz)
LHE	East seismic channel (1 Hz)
LHN	North seismic channel (1 Hz)
LHZ	Vertical seismic channel (1 Hz)
LCQ	Q330 clock quality (percent)
LCE	Q330 clock phase error (μ s)
LCC	GPS clock quality (percent)
LPL	Q330 clock phase lock loop status
LCL	Time since GPS lock was lost (seconds)
QBD	Total number of Q330 reboots in last 24 hours
QBP	Logical port buffer percent full from real-time status
QDG	Data gaps (s)
QDL	Current data latency (s)
QDR	Current total input and output data rate (bits/s)
QEF	Overall communications efficiency (percent)
QG1	Total number of data gaps in last hour
QGD	Total number of data gaps in last 24 hours
QID	Total number of Q330 IP address changes in last 24 hours
QLD	Total number of comm link cycles in last 24 hours
QPD	Total number of POCs received in last 24 hours
QRD	Total number of bytes read in last 24 hours
QRT	Current run time (s)
QTH	Current throttle setting (bits/s)
QTP	Ratio of seconds read to real-time clock
QWD	Total number of bytes written in last 24 hours
VCO	Voltage controlled oscillator value
VEA	Antenna current (A)
VEC	Main system current (A)
VEP	Main system voltage (V)
VKI	Main system temperature ($^{\circ}$ C)
VM1	Mass position for channel 1
VM2	Mass position for channel 2
VM3	Mass position for channel 3
VPB	Packet buffer
VTW	Main system opto inputs

Table 6.3: Description of Q330 and Antelope seismic and state-of-health channels. Channel sampling rates are indicated by the first letter of the channel code, where H indicates 100 Hz, L indicates 1 Hz, V indicates 0.1 Hz, and Q indicates 0.05 Hz.

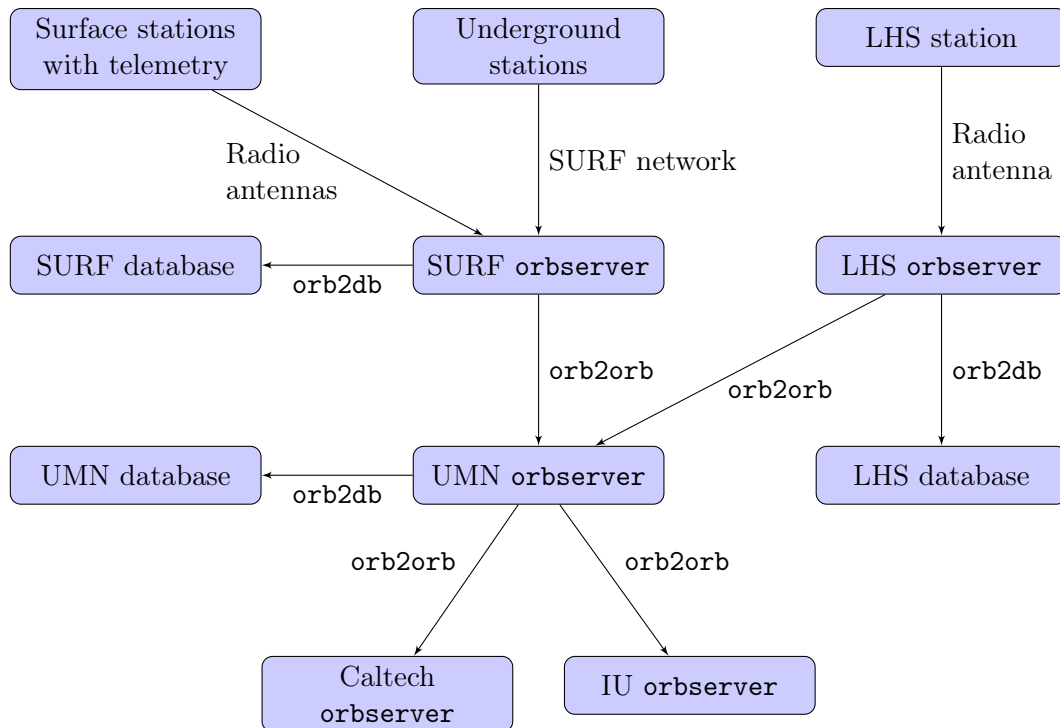


Figure 6.9: Real-time data flow for Homestake seismic stations using telemetry. Steps between stations and `orbservers` also include an instance of `q3302orb` in order to collect data packets and put them in the ring buffer managed by the `orbservers` process. Due to network restrictions, the `orb2orb` processes at SURF and LHS operate to push data to UMN, while the `orb2orb` processes at Caltech and IU pull data from UMN. We also note that data from the three remote surface stations (DEAD, SHL, and TPK) is periodically extracted and manually added to the databases.

6.3.4 Underground GPS timing

In this section, I describe a GPS timing distribution system which allows the transportation of the GPS timing signal into the mine and to all of the seismic stations.

Timing distribution system

The timing system begins with a GPS antenna mounted on the roof of the SURF administration building and connected to a Q330 in the server room. This “master” Q330 is an independent datalogger which is not connected to a seismometer; its primary purpose is to convert the high-frequency GPS signal (received by the antenna) into its separate

1PPS (1 pulse-per-second) and NMEA metadata components, which can be used as an external timing signal by the other Q330s. The output from this Q330's EXT GPS port is fed into a GPS transceiver which converts the electric signal into an optical signal (more details on the electronics are provided in the next subsection).

A network of optical splitters and GPS transceivers is used to send the GPS signal down both the ROSS and YATES shafts and down the drifts on the 300, 800, 1700, 2000, 4100, and 4850 mine levels. At each station, a GPS transceiver is used to convert the signal from optical to electrical. The electric signal is then sent into the Q330's EXT GPS port. A diagram of the timing distribution system in the mine is shown in Fig. 6.10.

We note here that there is also a necessary conversion of the electrical GPS signal between the TTL and RS-232 logic standards which occurs between the master Q330 and the first transceiver in the network, as well as between the transceiver and Q330 at each underground station. The master Q330 outputs the external GPS signal using the RS-232 standard, which treats signals between -15 V and -3 V as a digital 1 and signals between 3 V and 15 V as a digital 0. However, the GPS transceivers receive and transmit TTL logic levels, which accepts $2\text{--}5\text{ V}$ as digital 1 input and $0\text{--}0.8\text{ V}$ as digital 0 input. Outputs are usually acceptable in the range of $2.7\text{--}5\text{ V}$ for digital 1 and $0\text{--}0.5\text{ V}$ for digital 0.

As a result, we use powered converters at several places in the GPS timing chain. We convert the GPS signal from the master Q330 from RS-232 to TTL before the first GPS transceiver in the chain. The signal remains in TTL form throughout the rest of the timing chain, but must be converted to RS-232 at each station before being sent to the Q330's EXT GPS port. An example of this process is shown in Fig. 6.11.

Electronics

The GPS transceivers and optical splitters were custom-made for this application by Liteway, Inc. [330]. The splitters are passive (unpowered), so the two output signals are halved in amplitude compared to the input. As a result, the signal cannot pass through two consecutive splitters without dropping below required TTL logic levels; thus, we include a GPS transceiver between consecutive splitters in order to boost the signal. Several transceivers are part of the timing distribution system specifically for

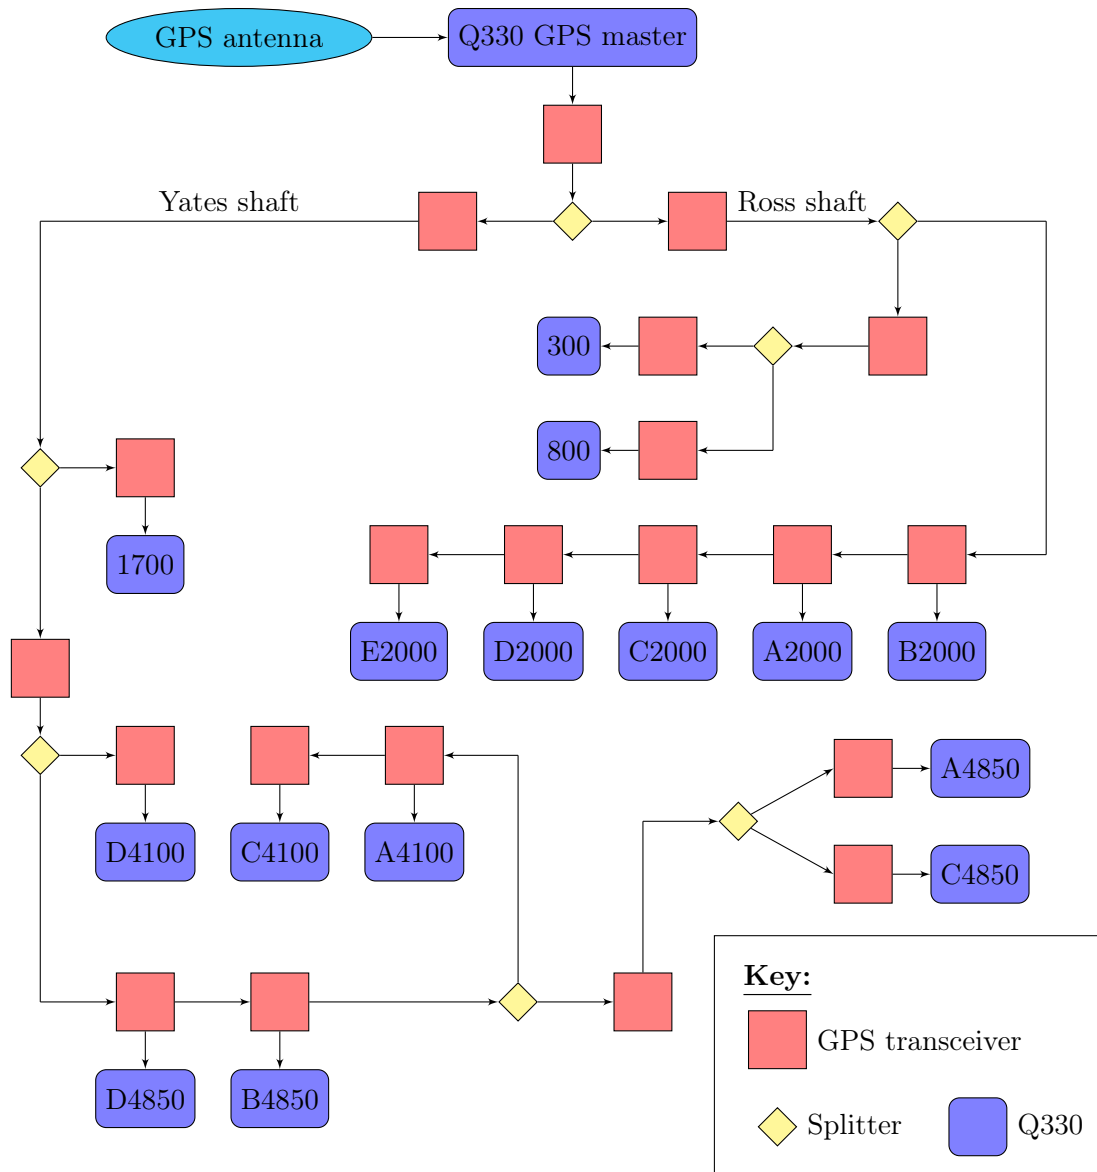


Figure 6.10: Flow chart of underground GPS timing distribution system for the Homestake 3D seismic array. The A4100 and C4100 stations originally received the timing signal through the Ross shaft, but due to a broken fiber, the network was reconfigured so that they receive timing from the 4850 level. RS-232 to TTL converters are not shown in this diagram, but are present between the Q330 GPS master and its transceiver, and between all transceivers and Q330s.




Figure 6.11: Example of TTL to RS-232 conversion at the 4850 level. The small green rectangles represent the converters.

this purpose, including the two transceivers which are in the chain immediately before the signal travels down the Yates and Ross shafts (see Fig. 6.10).

The GPS transceivers (model number GPSX-1001) are designed to transmit, receive, and repeat standard GPS NMEA and 1PPS signals using TTL logic. They take an electrical or optical signal as input and may output either or both types of signals. The functionality of the transceivers may be controlled by setting the dip switches appropriately. Fig. 6.12 shows the front and back of a GPS transceiver.

The converters are from B & B Electronics and are capable of converting RS-232 to TTL or vice versa [331]. Finally, we use custom-made cables between the Q330 and the converter, and between the converter and the GPS transceiver.

6.3.5 Data quality monitoring

I have developed a web interface for low-latency monitoring of the incoming data from the underground stations and surface stations with telemetry. This is done so that any issues with the array can be identified and responded to quickly in order to prevent significant data loss.

The `miniSEED` database of dayvolumes is continually updated by the Antelope `orb2db` process as the `orbserver` ring buffer accumulates more data. I have developed a Python-based package which loads in all available `miniSEED` data for the day in question and produces a time-series plot of each channel. A `cron` job is set to run this code on an hourly basis to produce plots of ten channels for each station: three seismic channels (LHE, LHN, LHZ), the clock quality (LCQ) and timing error (LCE) channels, Q330



Figure 6.12: Front and rear views of a GPS transceiver. The 1PPS and NMEA lights blink when receiving the corresponding signal. The optical signal is input through the Rx port and is output from the Tx port. The electrical 1 PPS and NMEA signals are output (or input, depending on dip switch settings) through the BNC jack and the terminal block, respectively.

voltage (VEP), temperature (VKI), and the three mass positions (VM1, VM2, and VM3).

The plots are organized into folders by date and within the folders by station and channel name. Several PHP scripts are used to operate the webpage. The home page (`index.php`) obtains a list of dates for which plots are available (using the plot folder names) and generates a month-by-month calendar of dates for which plots have been made (see Fig. 6.13). Each date for which there are plots available is a link to a daily page; dates for which plots are missing (either due to station malfunctions or network issues) have a red background, the tint of which is determined by the number of missing plots (more missing plots corresponds to a redder background).

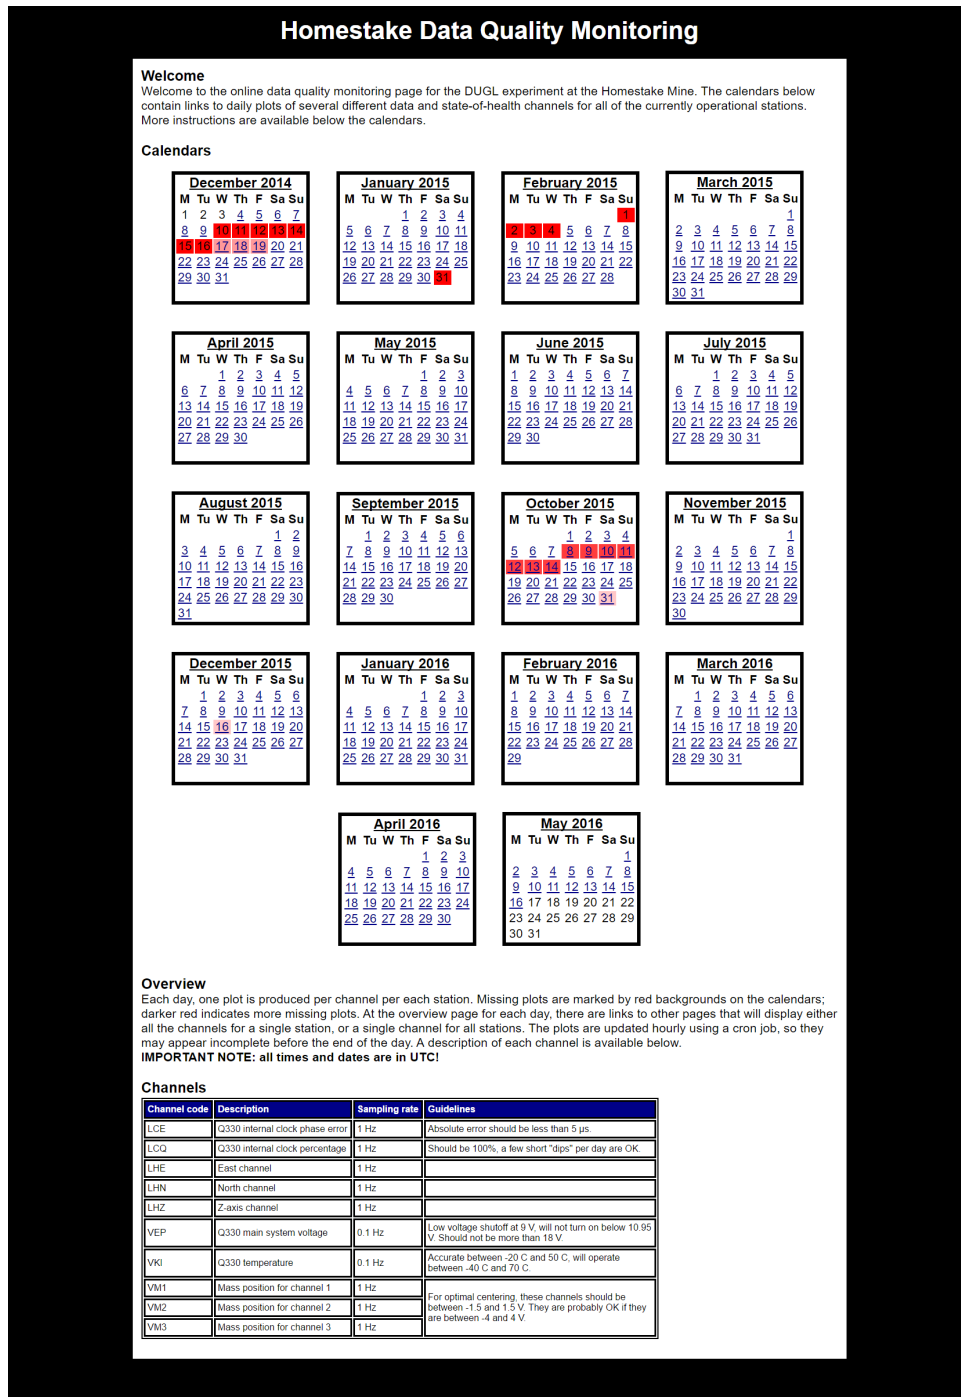


Figure 6.13: Main page of the Homestake data quality monitor. Calendars with links are visible for months with available data. A table with guidelines indicating the normal operating range for each channel is shown at the bottom.

Each daily page is generated by another PHP script called `day_overview.php`. This script takes in the day's date as an argument from the URL, looks in the corresponding folder, and generates links to all available plots for that date. The plots may be organized by station (i.e., all channels for one station for that date) or by channel (a single channel for all stations for that date). A screenshot of this configuration is shown in Fig. 6.14.

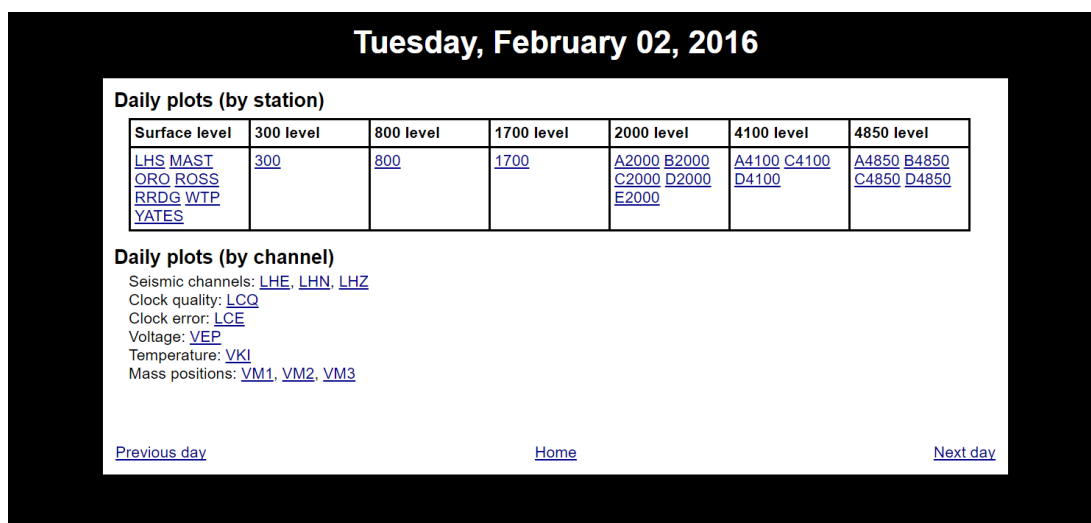


Figure 6.14: Screenshot of a daily page from the Homestake data quality monitor. The station links are organized by depth and each link leads to a page which shows plots of all channels for a particular station. The channel links are organized by channel type and each link leads to a page which shows plots of a particular channel for all stations.

The links from the daily pages lead to another PHP script, `show_plots.php`, which takes in arguments for the date, the type of page (station or channel), and the label (station name or channel name, respectively). This code generates the webpage for viewing the plots (either all channels for one station or one channel for all stations), using the arguments passed through the URL.

The channel data are plotted in blue when within the normal operating range and in red when the operating range thresholds are exceeded. This is done so that a user who is checking the plots for any data quality issues can quickly identify channels which may be out of range. Normal operating ranges are only defined for the state-of-health channels and are shown in Tab. 6.4. These are generally based on information in the Q330 and seismometer manuals, along with our strict requirements for timing accuracy.

Channel	Operating range
Clock quality	100%
Clock error	$-5 \mu\text{s} < \Delta t < 5 \mu\text{s}$
Voltage	$11 \text{ V} < V < 18 \text{ V}$
Temperature	$-20^\circ\text{C} < T < 50^\circ\text{C}$
Mass position	$-4 \text{ V} < \Delta V_m < 4 \text{ V}$

Table 6.4: Normal operating ranges defined for the seismic stations in the Homestake 3D array. These state-of-health channels are plotted in red by the data quality monitor when these ranges are exceeded (see Fig. 6.15).

Fig. 6.15 shows an example of the daily plots which were generated for the LHS station on February 1, 2016. The timing error and timing quality channels briefly deviate from their normal operating ranges (plotted in red), but in this case, the discrepancies are not significant and are not a cause for concern.

This entire package, along with detailed instructions, is available in the UMN LIGO SVN repository at https://support.physics.umn.edu/svn/ligo/Homestake_DQ.

6.3.6 Power spectra

For an overview of the data collected from the Homestake 3D seismic array, I have calculated the amplitude spectral density (ASD) of seismic noise over long periods. This has been done for different stations and different seismic channels (east, north, vertical) and compared to expected limits on the seismic background noise from the Peterson low- and high-noise models [332].

In Fig. 6.16, I have plotted the amplitude spectral density of ground displacement using one year of data (from June 1, 2015–May 31, 2016), split into 400 second intervals. All spectra shown use the vertical seismic channel. Note that the DEAD, SHL, and TPK stations only used ≈ 3 months of data; no more data are currently available from these stations since they do not have telemetry. The YATES station is also missing several weeks of data during this time period due to power and communication issues. The curves shown are the median amplitudes in each frequency bin; this is determined based on the amplitudes resulting from the ASDs calculated from all 400 second intervals which span the available data for each station. All available data are utilized, including data which may contain noise transients.

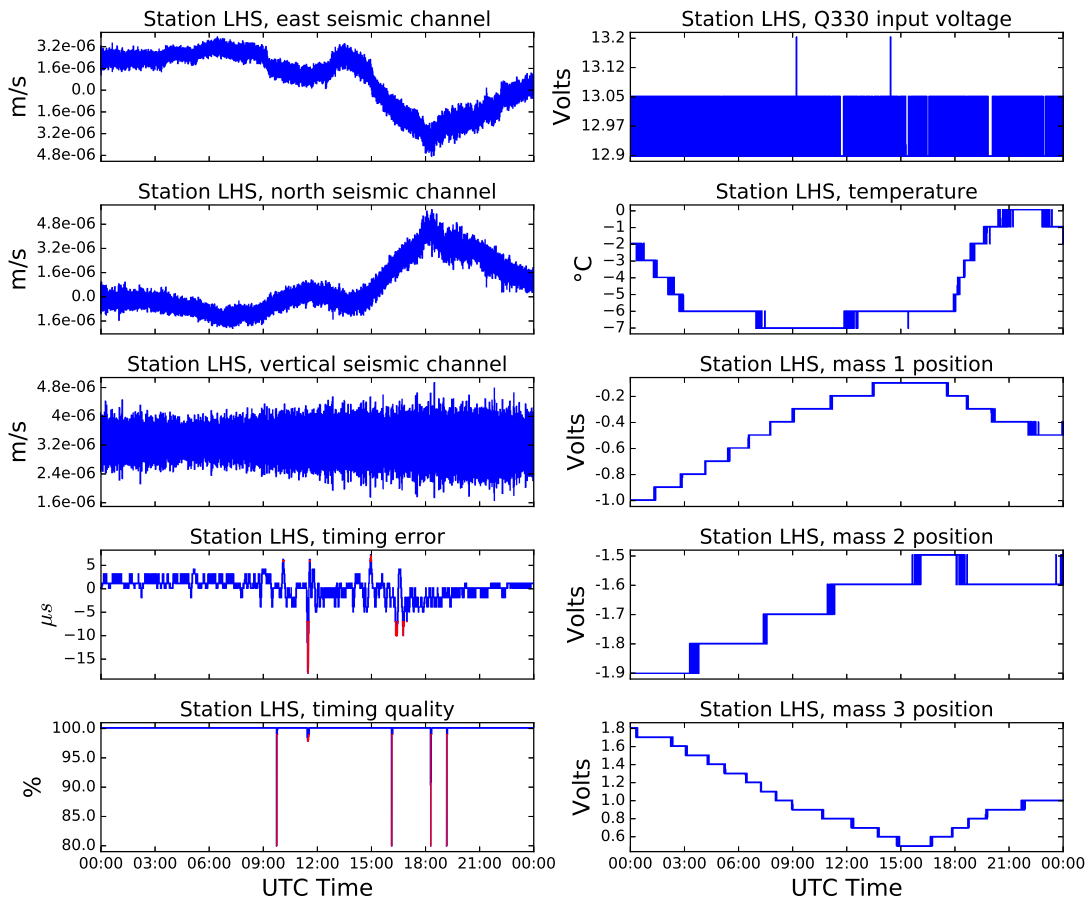


Figure 6.15: Plots of all 10 channels used in the Homestake data quality monitor for the LHS station for a single day. Most channels are plotted in all blue, indicating that they are within the acceptable operating range, but the timing error and timing quality have some brief spots of red where the normal ranges are exceeded. In this case, the deviation is not extreme and is not of significant concern.

The top-left panel shows the ASD for several different stations to provide a direct comparison of the noise level at different depths. All of the stations are in close agreement in the middle range of frequencies, which corresponds to the microseismic peak. At higher frequencies, there is significantly less noise with depth; above 0.5 Hz, the C4100 and D4850 stations are nearly an order of magnitude quieter than other stations. At the lowest frequencies, there is fairly good agreement between the stations, although a slight trend of decreasing noise with depth is apparent; this may be due to larger temperature

variations closer to the surface or may be a coincidence since this can also depend on the quality of the concrete pad at each station. In some cases, the spectra fall relatively close to the low-noise Peterson model, especially at low frequencies and between 0.4–0.9 Hz.

For the surface stations (top-right panel) there is a wide range of variability; this is due to differences in the local environment in terms of thermal insulation and proximity to human activity. Differences in the microseismic peak (0.1–0.2 Hz) are due to differences in the amount of data used in this analysis; the microseism experiences seasonal variations and appears differently for stations which do not include the full year of data (DEAD, SHL, TPK, and YATES).

The middle-left panel shows spectra for the 300, 800, and 1700 stations. The noise levels are reduced with depth at higher and lower frequencies; the higher level of low-frequency noise at the 800 station is likely due to its proximity to one of the mine shafts.

In the middle-right panel, I show spectra for the 2000-level stations. There is generally good agreement between the stations across all frequencies. B2000 experiences increased noise, especially at high frequencies, likely due to its location near a lunch room and a mine shaft.

Spectra for the 4100-level stations are shown in the bottom-left panel. The C4100 station appears to have the least noise at low frequencies, while the D4100 station experiences the most high-frequency noise. These variations are not well-understood based on station locations and expected proximity to human activity.

Finally, spectra for the 4850 level are shown in the bottom-right panel. Here, the B4850 station experiences significantly increased high-frequency noise due to its location near a large fan and ongoing construction. Above 1 Hz, the other stations are in fairly good agreement, although each station seems to have its own individual noise peaks in the spectrum. This is likely due to the unique environment surrounding each station: C4850 is in a storage room and very close to a rail line, and D4850 is very close to other scientific experiments, ventilation equipment, and human activity. The A4850 station seems to have more overall high-frequency noise than these two stations, which is not well-understood since it is one of the most isolated stations in the entire array.

In Fig. 6.17, I show ASD histograms for the RRDG surface station (left column) and the A4850 underground station (right column) as examples of a relatively good surface station and our deepest and most isolated underground station. Here, I have taken all

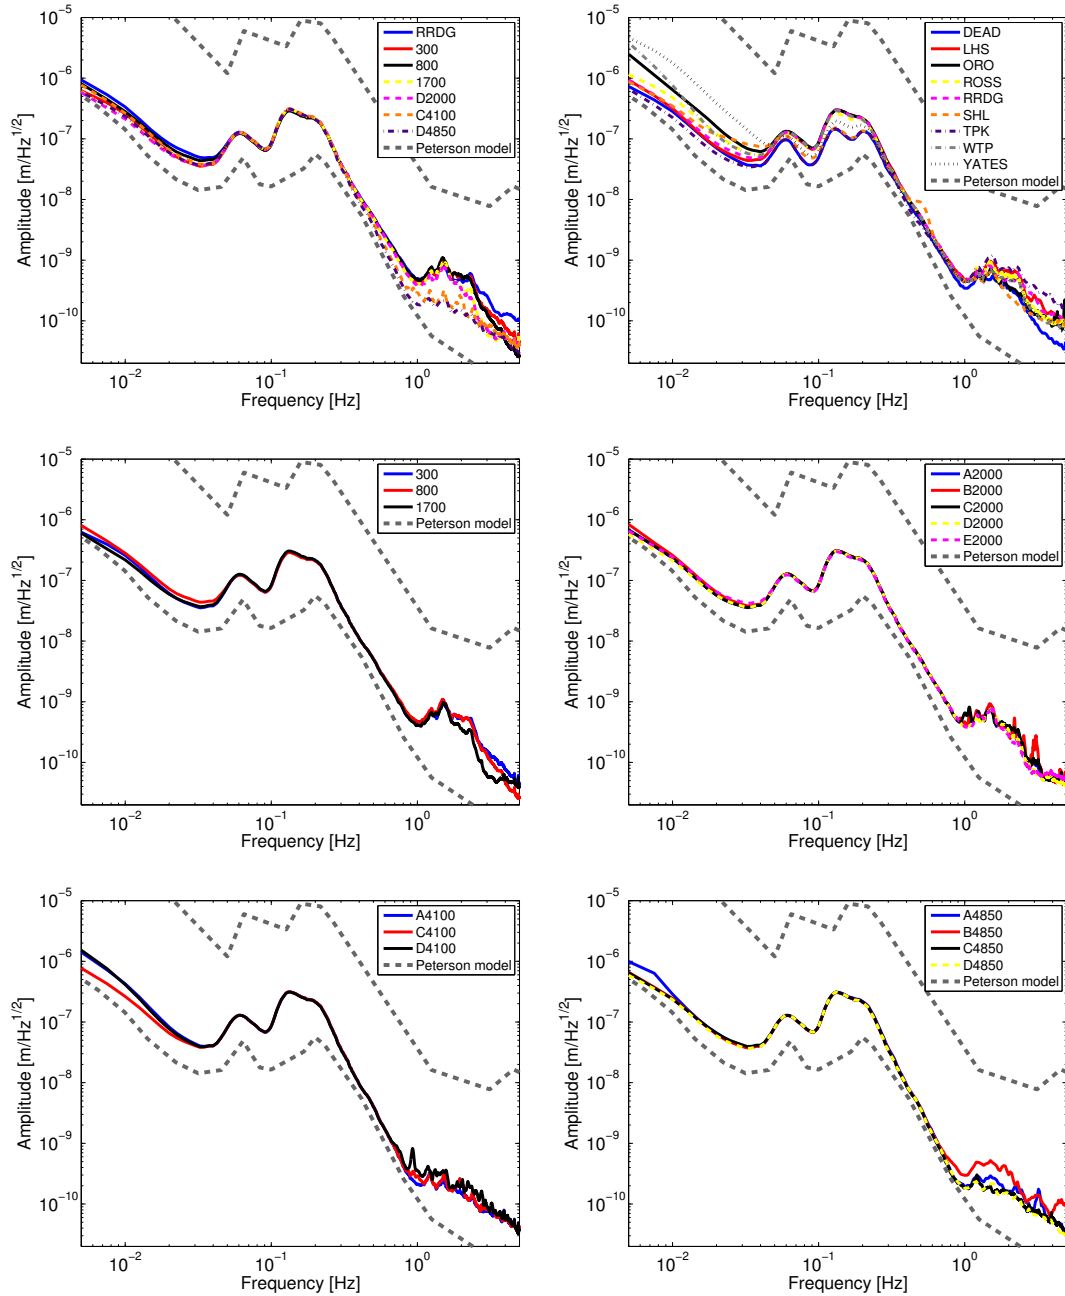


Figure 6.16: Median amplitude spectral densities for all Homestake stations; 1 year of data is used except for DEAD, SHL, and TPK, for which only ≈ 3 months of data are currently available, and YATES, which is missing data due to power and communication issues. More details are provided in the text.

of the ASDs calculated from the 400-second data intervals over 1 year and calculated a histogram in each frequency bin to understand the overall variability of the noise background. The white curve represents the median ASD (identical to those shown in Fig. 6.16), the black curves represent the 95% confidence intervals in each frequency bin, and the color scale shows the overall distribution. The Peterson low- and high-noise models are shown in dashed gray.

The histograms display about two orders of magnitude of variation across all frequencies for both the RRDG station and the A4850 station. The A4850 station measures less noise in general and appears to have less overall variation than RRDG. There also appears to be significantly more high-frequency noise in the RRDG station; this is likely due to anthropogenically generated surface waves which are suppressed with depth.

Both stations stay within the low- and high-noise Peterson models most of the time; in the 0.3–0.9 Hz range, the A4850 station is actually below the low-noise model a significant fraction of the time. There is also a considerable difference between the vertical channel and the horizontal channels at low frequencies: at 0.01 Hz and below for both stations, the vertical channel is almost an order of magnitude lower than the horizontals.

6.4 Seismic radiometer

The stochastic radiometer technique, which is used to search for directional sources of a stochastic gravitational-wave background and unmodeled sources of continuous gravitational waves [138, 164], can be adapted to analyze seismic data. By applying this technique to seismic waves, we hope to identify significant components of the seismic field and ascertain any potential directional dependence in the seismic environment around Homestake.

In this section, I will discuss some of the models used in this technique, developed in Ref. [333], demonstrate their application to simulated data and real data, and discuss preliminary results and plans for further development.

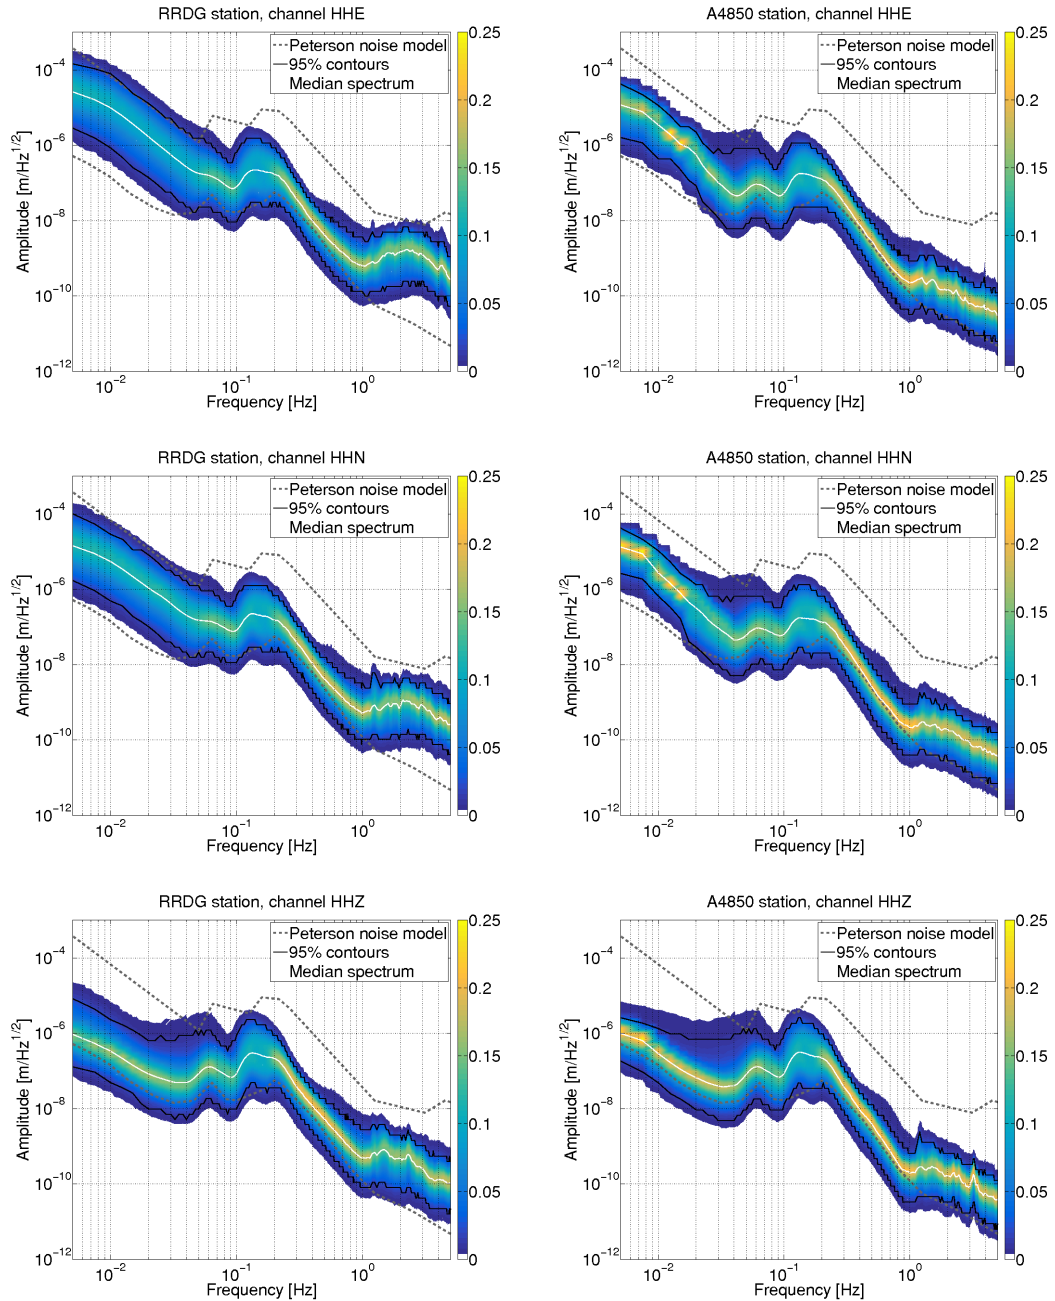


Figure 6.17: Histograms of amplitude spectral density in each frequency bin for the RRDG (left column) and A4850 (right column) stations. The plots are divided into rows by channel: east (top), north (middle), and vertical (bottom). Median ASDs (solid white), 95% confidence intervals for each frequency bin (solid black), and the Peterson low- and high-noise models (dashed gray) are shown. See the text for more details.

6.4.1 Formalism

Body waves

Beginning with body waves, the seismic field at location \vec{x} and time t due to P-waves is given by the following plane-wave expansion:

$$\vec{p}(\vec{x}, t) = \int df d\hat{\Omega} P(f, \hat{\Omega}) \hat{\Omega} e^{2\pi i f(t - \hat{\Omega} \cdot \vec{x} / v_p)}. \quad (6.11)$$

Here, $\hat{\Omega}$ is the direction of wave propagation, f is the frequency, v_p is the P-wave velocity, and $P(f, \hat{\Omega})$ is the frequency- and direction-dependent P-wave amplitude.

The corresponding seismic field due to S-waves can be written as

$$\vec{s}(\vec{x}, t) = \sum_A \int df d\hat{\Omega} S_A(f, \hat{\Omega}) \hat{e}_A(\hat{\Omega}) e^{2\pi i f(t - \hat{\Omega} \cdot \vec{x} / v_s)}. \quad (6.12)$$

Here, the index A runs over the two independent S-wave polarizations (typically parametrized as horizontal and vertical), $\hat{e}_A(\hat{\Omega})$ is the unit vector in the direction of polarization A , v_s is the S-wave velocity, and $S_a(f, \hat{\Omega})$ is the S-wave amplitude.

With this decomposition of the seismic field, two-point correlations between the P-wave and S-wave amplitudes can be determined:

$$\langle P^*(f, \hat{\Omega}) P(f', \hat{\Omega}') \rangle = \delta(f - f') \delta^2(\hat{\Omega}, \hat{\Omega}') H_P(f, \hat{\Omega}) \quad (6.13)$$

$$\langle S_A^*(f, \hat{\Omega}) S_{A'}(f', \hat{\Omega}') \rangle = \delta_{AA'} \delta(f - f') \delta^2(\hat{\Omega}, \hat{\Omega}') H_{S,A}(f, \hat{\Omega}) \quad (6.14)$$

$$\langle S_A^*(f, \hat{\Omega}) P(f', \hat{\Omega}') \rangle = 0. \quad (6.15)$$

Here, the δ s indicate Kronecker or Dirac delta functions, $H_P(f, \hat{\Omega})$ indicates the power spectral density in P-waves, and $H_{S,A}(f, \hat{\Omega})$ indicates the power spectral density in S-waves of polarization A . These correlations represent assumptions that waves at different frequencies, from different directions, and with different polarizations are uncorrelated. These assumptions may not hold in all cases since reflections can change the polarization of a wave (a reflected P-wave can contain longitudinal and transverse components) and due to the existence of quasi-P and quasi-S waves in anisotropic media.

The measured ground displacement in a single seismometer due to body waves can

be calculated as the sum of the P and S components. For seismometer i at a location \vec{x}_i , the measurement in channel α is

$$d_{i\alpha}(\vec{x}_i, t) = (\vec{p}(\vec{x}_i, t) + \vec{s}(\vec{x}_i, t)) \cdot \hat{\alpha}, \quad (6.16)$$

where $\hat{\alpha}$ is the unit vector in the direction aligned with seismometer channel α . We assume that the seismic noise floor is well above the instrument noise and do not consider it here.

Cross-correlating data from channel α in seismometer i at location \vec{x}_i and channel β from seismometer j at location \vec{x}_j gives

$$\begin{aligned} \langle Y_{i\alpha j\beta} \rangle &= \int_{-T/2}^{T/2} dt d_{i\alpha}(\vec{x}_i, t) d_{j\beta}(\vec{x}_j, t) \\ &= \int_{-T/2}^{T/2} dt \int df d\hat{\Omega} \left(\sum_A H_{S,A}(f, \hat{\Omega}) (\hat{e}_A(\hat{\Omega}) \cdot \hat{\alpha}) (\hat{e}_A(\hat{\Omega}) \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_s} \right. \\ &\quad \left. + H_P(f, \hat{\Omega}) (\hat{\Omega} \cdot \hat{\alpha}) (\hat{\Omega} \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_p} \right) \end{aligned} \quad (6.17)$$

Here, $\Delta \vec{x} = \vec{x}_i - \vec{x}_j$. The time integral in this expression is trivial, and if the analysis is performed in a small frequency bin Δf , the frequency integral is also trivial. Using the one-sided frequency spectrum also gives a factor of 2; the result of these operations is

$$\begin{aligned} \langle Y_{i\alpha j\beta} \rangle &= 2T \Delta f \int d\hat{\Omega} \left(\sum_A H_{S,A}(\hat{\Omega}) (\hat{e}_A(\hat{\Omega}) \cdot \hat{\alpha}) (\hat{e}_A(\hat{\Omega}) \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_s} \right. \\ &\quad \left. + H_P(\hat{\Omega}) (\hat{\Omega} \cdot \hat{\alpha}) (\hat{\Omega} \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_p} \right). \end{aligned} \quad (6.18)$$

The amplitudes $H_{S,A}$ and H_P are written here as functions of only $\hat{\Omega}$ since the analysis is performed in a single frequency bin. These amplitudes can be parametrized in terms

of a basis to study the directional distribution of the components of the seismic field:

$$H_P(\hat{\Omega}) = \sum_n P_n Q_n(\hat{\Omega}) \quad (6.19)$$

$$H_{S,A}(\hat{\Omega}) = \sum_n S_{A,n} Q_n(\hat{\Omega}). \quad (6.20)$$

Here, $S_{A,n}$ corresponds to the power in basis element n due to S-waves of polarization A and P_n corresponds to the power in basis element n due to P-waves. For this analysis, the pixel basis ($Q_{\hat{\Omega}_0}(\hat{\Omega}) = \delta(\hat{\Omega} - \hat{\Omega}_0)$) is used; however, another useful basis could be the spherical harmonics basis, $Q_{lm}(\hat{\Omega}) = Y_{lm}(\hat{\Omega})$.

Using any basis, the following gamma functions may be defined, which capture the direction dependence of the cross-correlation.

$$\gamma_{P,n,i\alpha j\beta} = \int d\hat{\Omega} Q_n(\hat{\Omega})(\hat{\Omega} \cdot \hat{\alpha})(\hat{\Omega} \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_p} \quad (6.21)$$

$$\gamma_{SA,n,i\alpha j\beta} = \int d\hat{\Omega} Q_n(\hat{\Omega})(\hat{e}_A(\hat{\Omega}) \cdot \hat{\alpha})(\hat{e}_A(\hat{\Omega}) \cdot \hat{\beta}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_s} \quad (6.22)$$

With these gamma functions, the full cross-correlation can be written as

$$\begin{aligned} \langle Y_{i\alpha j\beta} \rangle &= 2T \Delta f (S_{H,a} \gamma_{S_{H,a},i\alpha j\beta} + S_{V,b} \gamma_{S_{V,b},i\alpha j\beta} + P_c \gamma_{P,c,i\alpha j\beta}) \\ &= S_d \gamma_d \end{aligned} \quad (6.23)$$

Note that summation is implied for repeated indices. P-waves and S-waves may be combined as written in Eq. 6.23 or divided into two separate analyses.

Using seismic data from the Homestake 3D array and the locations of the seismometers, all of the quantities in Eq. 6.23 other than the coefficients $S_{H,a}$, $S_{V,b}$, and P_c can be calculated. The goal of our analysis is to estimate these coefficients. In order to do this, we define a likelihood function

$$\mathcal{L} \propto \exp \left[(Y_i^* - \gamma_{id}^* S_d) \Sigma^{-1} (Y_i - \gamma_{id} S_d) \right], \quad (6.24)$$

where the index i runs over all pairs of seismometers and seismic channels and d runs over seismic field components. The covariance matrix Σ becomes proportional to the identity matrix under the assumption that all of the seismometers have a similar noise

floor which is stationary in time. In order to maximize this likelihood, the exponent should be zero. Under this requirement, we obtain the following expression for the coefficients:

$$\vec{S} = (\gamma^\dagger \gamma)^{-1} \gamma^\dagger \vec{Y}. \quad (6.25)$$

In practice, I have implemented a least-squares method to invert the matrix $\gamma^\dagger \gamma$ and determine \vec{S} .

The velocities v_p and v_s are not known *a priori*; we use $v_p = 5600$ m/s and $v_s = 3300$ m/s based on studies of the Homestake environment in Ref. [334].

Rayleigh waves

Extending this method to the case of Rayleigh waves, the resulting displacement field may be described by

$$\vec{r}(\vec{x}, t) = e^{-z/\alpha} \left(\cos(\omega t - \vec{k} \cdot \vec{x}) \hat{k} + \epsilon \sin(\omega t - \vec{k} \cdot \vec{x}) \hat{z} \right), \quad (6.26)$$

where z is the vertical coordinate, α is an unknown decay constant, ω is the angular frequency, \vec{k} is the wave vector, and ϵ is a unitless parameter which encapsulates the relative amplitude of the vertical component to the horizontal component (typical values are 0.5–1.5). The parameters α and ϵ must be determined in advance and are not estimated by the seismic radiometer algorithm. Note that \vec{x} describes the location where the seismic field is being measured, rather than the x coordinate.

The formulation of Rayleigh waves described here differs from that shown in Eq. 6.10, where a biexponential functional form is used for the amplitude eigenfunctions. The two cases may be compatible in situations where the exponential coefficients in Eq. 6.10 are similar, but this cannot be assumed in general. I will first test the more simplistic formulation described here, and in Sec. 6.5, I will implement and test the biexponential functional form.

With this field, we calculate the following plane wave expansion for Rayleigh waves:

$$\vec{r}(\vec{x}, t) = \int df d\hat{\Omega} e^{-z/\alpha} R(f, \hat{\Omega}) e^{2\pi i f(t - \hat{\Omega} \cdot \vec{x}/v_r)} \left(\hat{\Omega} + \epsilon e^{i\pi/2} \hat{z} \right). \quad (6.27)$$

Here, v_r is the Rayleigh wave velocity.

The two-point correlation for Rayleigh waves is defined as

$$\langle R^*(f, \hat{\Omega})R(f', \hat{\Omega}') \rangle = \delta(f - f')\delta^2(\hat{\Omega}, \hat{\Omega}')H_R(f, \hat{\Omega}), \quad (6.28)$$

where $H_R(\hat{\Omega})$ is the direction-dependent power in Rayleigh waves. Analogous to the body waves case, the cross-correlation estimator for Rayleigh waves is

$$\begin{aligned} \langle Y_{i\alpha j\beta} \rangle &= 2T\Delta f \int d\hat{\Omega} e^{(z_i+z_j)/\alpha} H_R(\hat{\Omega}) e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / v_r} \\ &\times \left[(\hat{\Omega} \cdot \hat{\alpha} + \epsilon e^{i\pi/2} \hat{z} \cdot \alpha) (\hat{\Omega} \cdot \hat{\beta} + \epsilon e^{-i\pi/2} \hat{z} \cdot \beta) \right]. \end{aligned} \quad (6.29)$$

Expanding this in some basis $Q_n(\hat{\Omega})$

$$H_R(\hat{\Omega}) = \sum_n R_n Q_n(\hat{\Omega}) \quad (6.30)$$

gives the following gamma functions:

$$\gamma_{R,n,i\alpha j\beta} = \int d\hat{\Omega} Q_n(\hat{\Omega}) e^{2\pi i f \hat{\Omega} \cdot \vec{x} / v_r} \left[(\hat{\Omega} \cdot \hat{\alpha} + \epsilon e^{i\pi/2} \hat{z} \cdot \alpha) (\hat{\Omega} \cdot \hat{\beta} + \epsilon e^{-i\pi/2} \hat{z} \cdot \beta) \right]. \quad (6.31)$$

Finally, the cross-correlation estimator can be written as a sum over the gamma functions

$$\langle Y_{i\alpha j\beta} \rangle = 2T\Delta f e^{-(z_i+z_j)/\alpha} R_n \gamma_{R,n,i\alpha j\beta} \quad (6.32)$$

From here, we can apply the same likelihood treatment as for the body waves in order to determine the basis element coefficients R_n . We use a Rayleigh wave velocity of $v_r = 2500$ m/s for this model [334].

6.4.2 Simulations

Based on the models described in the previous subsections, I have developed simulations of P-waves, S-waves, and Rayleigh waves. In this subsection, I will simulate seismic waves of different types, frequencies, and with different locations, and demonstrate their recovery with the seismic radiometer algorithm. In all cases, I use a virtual array of 24 seismometers which correspond to the locations of the actual Homestake stations. The cross-correlation is performed for all combinations of seismometers and channels.

First, I have simulated a P-wave coming from $(\phi, \theta) = (300^\circ, -30^\circ)$ with an amplitude of 10^{-6} m, including Gaussian white noise with an amplitude of 10^{-4} m. The simulated data have been analyzed using the pixel basis and the P-wave likelihood estimator. Fig. 6.18 shows the directional distribution of P_n . The regions with the most recovered power are indicated in red and match up well with the injected signal location; the dotted black contour indicates the region which contains 95% of the total power (43 square degrees). In general, the direction can be more tightly constrained for seismic waves whose wavelength is less than the aperture of our array (see Fig. 6.1); based on expected P-wave velocities, this corresponds to frequencies of a few Hz or more. However, I have focused here on frequencies which may be expected to contain more overall seismic noise. Summing up the power in all directions, I obtain an estimate for H_P of $0.99 \times 10^{-12} \text{ m}^2$, which is in good agreement with the injected amplitude (squared).




Figure 6.18: Recovery of a simulated P-wave using the seismic radiometer analysis. The location of the injected seismic wave is denoted by a white star. The dotted black contour indicates the directions which contain 95% of the total power.

Next, I repeat this procedure using a simulated S-wave with horizontal polarization

and the S-wave likelihood estimator. The injected signal is coming from a direction of $(\phi, \theta) = (80^\circ, 20^\circ)$ with an amplitude of 10^{-6} m and I again add Gaussian white noise with an amplitude of 10^{-4} m. The recovery across all directions is shown in Fig. 6.19 for both horizontal and vertical polarizations. I estimate that $H_{S,H} = 1.07 \times 10^{-12}$ m² and $H_{S,V} = 0.07 \times 10^{-12}$ m². This is in good accordance with expectations, although the total power ($H_{S,H} + H_{S,V}$) is slightly high. The dotted black contour indicates the region which contains 95% of the total power; in this case, it corresponds to an area of 48 square degrees. Compared to the P-wave case, the recovered power is not quite as accurate and the direction is less constrained; this likely occurs because S-waves have twice as many degrees of freedom as P-waves.

Finally, I inject a Rayleigh wave with an amplitude of 10^{-6} m, a direction of $(\phi, \theta) = (210^\circ, 0^\circ)$, and parameters $\alpha = 1000$ m and $\epsilon = 1$, adding Gaussian white noise with an amplitude of 10^{-4} m. The known values of α and ϵ are used to perform the recovery in this case; for an analysis of real data, it is necessary to make an estimate of realistic values of these parameters. Fig. 6.20 shows the distribution of recovered power across the sky; the white star denotes the direction of the injected signal and the dotted black contour shows the directions which contain 95% of the total power, corresponding to 38 square degrees in this case. Summing over all sky positions, I obtain an estimate for H_R of 1.10×10^{-12} m², which is in fairly good agreement with the injected amplitude.

6.4.3 Analysis of real data

Next, I test the seismic radiometer algorithm using 1 hour of real data (beginning at GPS time 1121335217) from the Homestake 3D array. The data chosen contain no explicit seismic signal so that we can characterize the general background of seismic noise.

The noise data were analyzed with the P-wave, S-wave, and Rayleigh wave estimators at 0.1, 0.2, 0.5, and 1 Hz. For the Rayleigh wave analysis, I use $\epsilon = 1$ based on analyses from Ref. [335], which indicate that this a reasonable choice for this parameter, at least near frequencies of 0.1 Hz, and I test three values for α : 100, 1000, and 2000 m. In principle, the value of α should be frequency-dependent, but here I use the simplest implementation of this model (same α for all frequencies) due to the fact that we only have a general idea of what values α might take. Testing three different values of α should provide some insight into whether a particular value is more appropriate for a

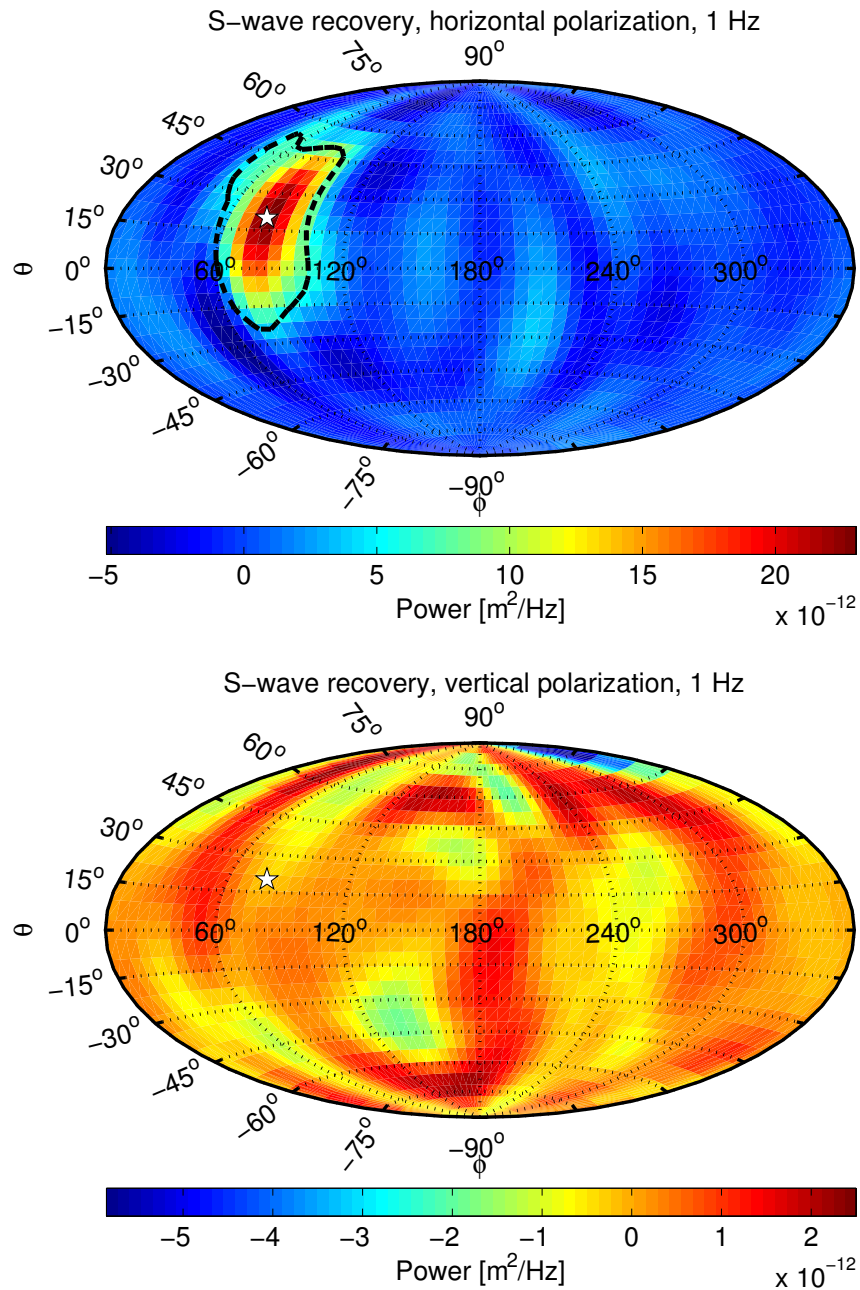


Figure 6.19: Recovery of a simulated S-wave using the seismic radiometer analysis. The top panel shows power in the horizontal polarization and the bottom panel shows power in the vertical polarization. The injected signal is horizontally polarized and its direction is denoted by a white star in the top panel. The dotted black contour in the top panel indicates the directions which contain 95% of the total power.




Figure 6.20: Recovery of a simulated Rayleigh wave using the seismic radiometer analysis. The location of the injected signal is indicated by a white star and the dotted black contour indicates the directions which contain 95% of the total power.

given frequency.

Fig. 6.21 shows the recovery for P-waves at 0.1 Hz, 0.2 Hz, 0.5 Hz, and 1 Hz. The total power recovered across all directions is given in Tab. 6.5. Overall, most of the power comes from 0.1 Hz and 0.2 Hz, likely due to the microseismic peak. There also seems to be a preference for $\theta = 0$ at all frequencies other than 0.5 Hz.

In Fig. 6.22, I show the recovery for horizontally and vertically polarized S-waves at 0.1 Hz and 0.2 Hz; Fig. 6.23 shows the recovery at 0.5 Hz and 1 Hz. The total power in each frequency and polarization is given in Tab. 6.5. At 0.1 Hz and 0.2 Hz, the horizontal S-wave power is almost equal to the P-wave power, while the vertical S-wave power is significantly less; at 0.5 Hz and 1 Hz, the power is considerably decreased and there is more variability in the overall power in the two polarizations, implying that the noise level at these frequencies is less significant. Overall, the horizontally polarized S-waves have little θ dependence, which indicates that they may be picking up on surface waves.

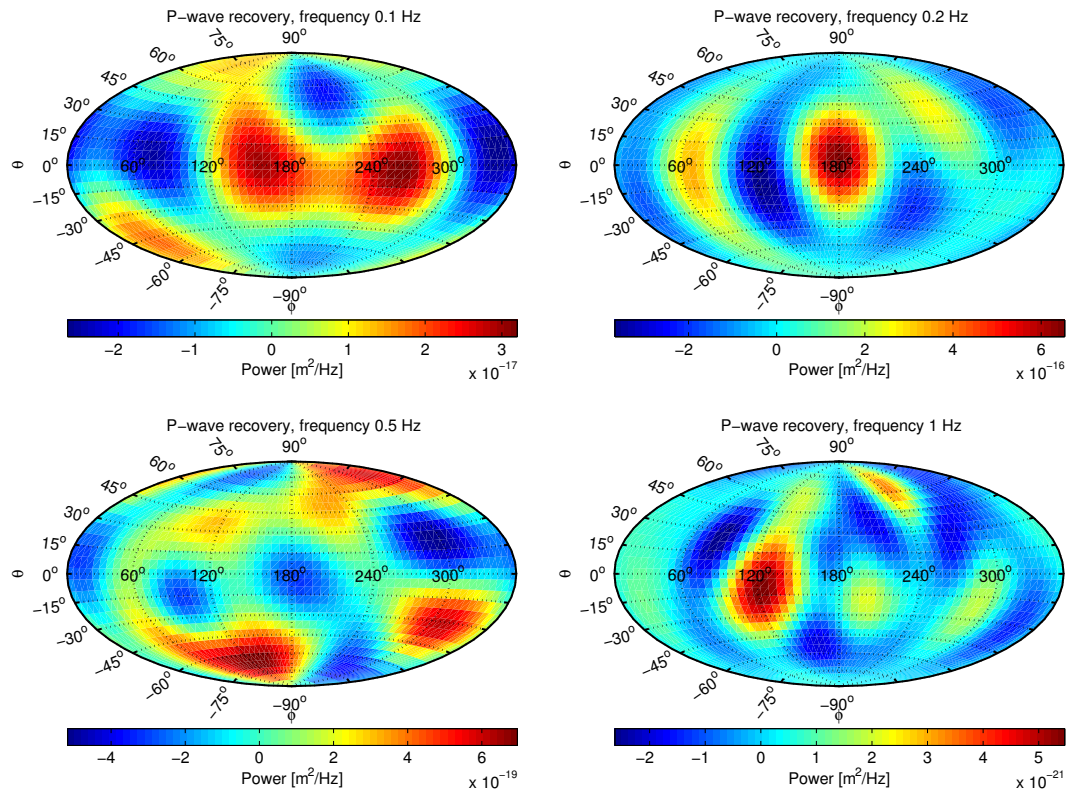


Figure 6.21: P-wave analysis of 1 hour of background noise with the seismic radiometer algorithm. Recoveries are performed at 0.1 Hz (top-left), 0.2 Hz (top-right), 0.5 Hz (bottom-left), and 1 Hz (bottom-right).

The vertical polarization seems to be more spread across different directions, which may occur due to reflections off of the surface and other interfaces in the crust.

Fig. 6.24 shows the Rayleigh wave recovery at 0.1 Hz and 0.2 Hz with $\alpha = 100$ m, 1000 m, and 2000 m; Fig. 6.25 shows the same analyses at 0.5 Hz and 1 Hz. Again, the most power is recovered at the lower frequencies (see Tab. 6.5 for more details) and the analysis with $\alpha = 1000$ m seems to recover the most power across all frequencies, while the analysis with $\alpha = 100$ m consistently recovers the least power. However, the overall structure seems to be fairly consistent for different values of α at 0.1 Hz and 0.2 Hz; at the higher frequencies, the case with $\alpha = 100$ m produces relatively different recoveries compared to the other two cases. At the lower frequencies, especially 0.2 Hz, there is little θ dependence (as expected) since Rayleigh waves are surface waves and should not

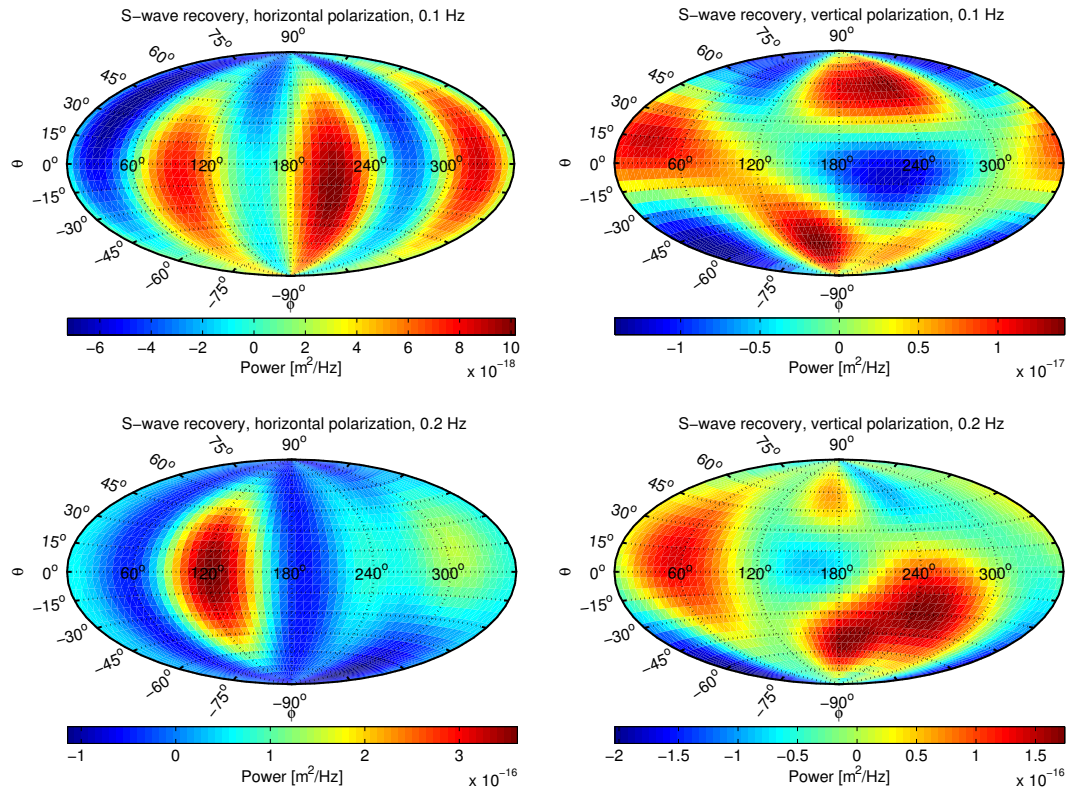


Figure 6.22: S-wave analysis of 1 hour of background noise with the seismic radiometer algorithm. Recoveries are performed at 0.1 Hz (top row) and 0.2 Hz (bottom row) for the horizontal (left column) and vertical (right column) polarizations.

be coming from depth.

6.5 Parameter estimation for Rayleigh wave eigenfunctions

One issue with the current Rayleigh wave model that I have implemented is that the parameters ϵ and α are not known *a priori*. This model also features several simplifying assumptions, including that the depth suppression is not frequency dependent and is the same for the both the horizontal and vertical components, and that the vertical and horizontal amplitudes are related by a constant coefficient which does not depend on depth or frequency (see Eq. 6.26). In explicit terms, this model takes the Rayleigh wave eigenfunctions (as described by Eq. 6.9) to be $r_1(z, k, \omega) = r_1(z) = e^{-z/\alpha}$ and

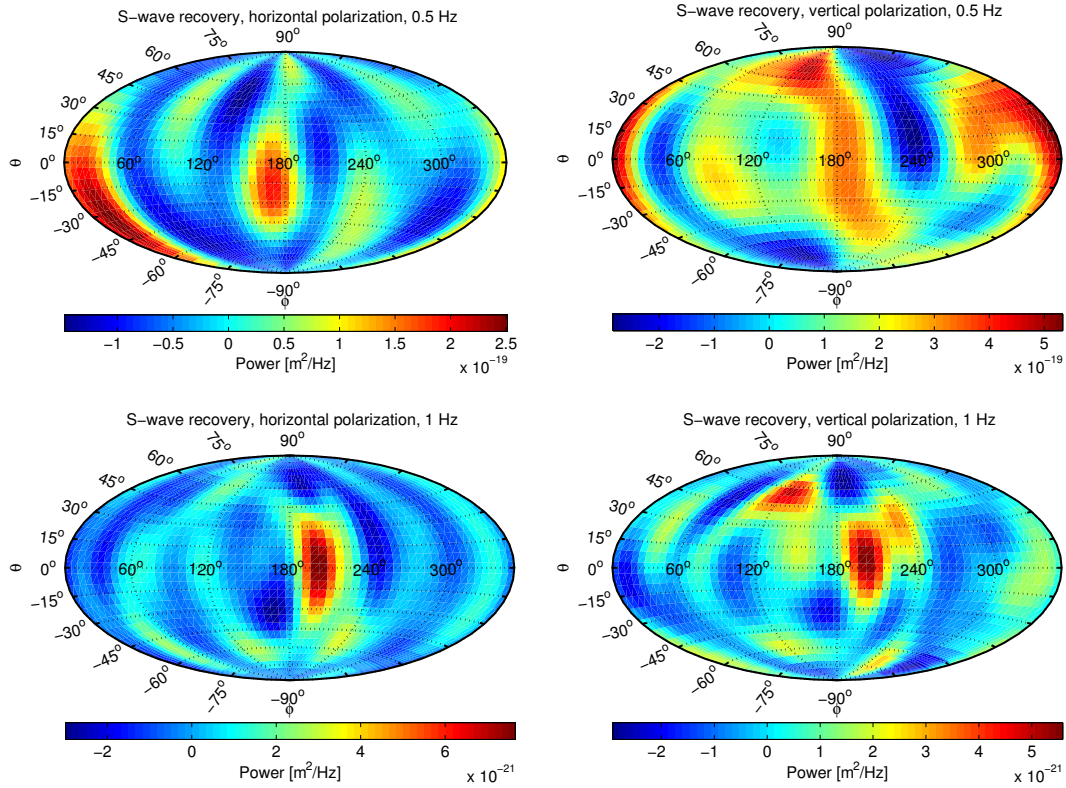


Figure 6.23: S-wave analysis of 1 hour of background noise with the seismic radiometer algorithm. Recoveries are performed at 0.5 Hz (top row) and 1 Hz (bottom row) for the horizontal (left column) and vertical (right column) polarizations.

$r_2(z, k, \omega) = r_2(z) = \epsilon e^{-z/\alpha}$, where α and ϵ depend on the material constants.

In this section, I aim to increase the complexity of our Rayleigh wave model by using the biexponential form for the eigenfunctions given in Eq. 6.10. This functional form mitigates some of the assumptions described previously and has proven to be a reasonable approximation to the eigenfunctions even in somewhat heterogeneous media [336]. I will use observations of Rayleigh waves in the Homestake 3D array to estimate the parameters of these eigenfunctions at several frequencies.

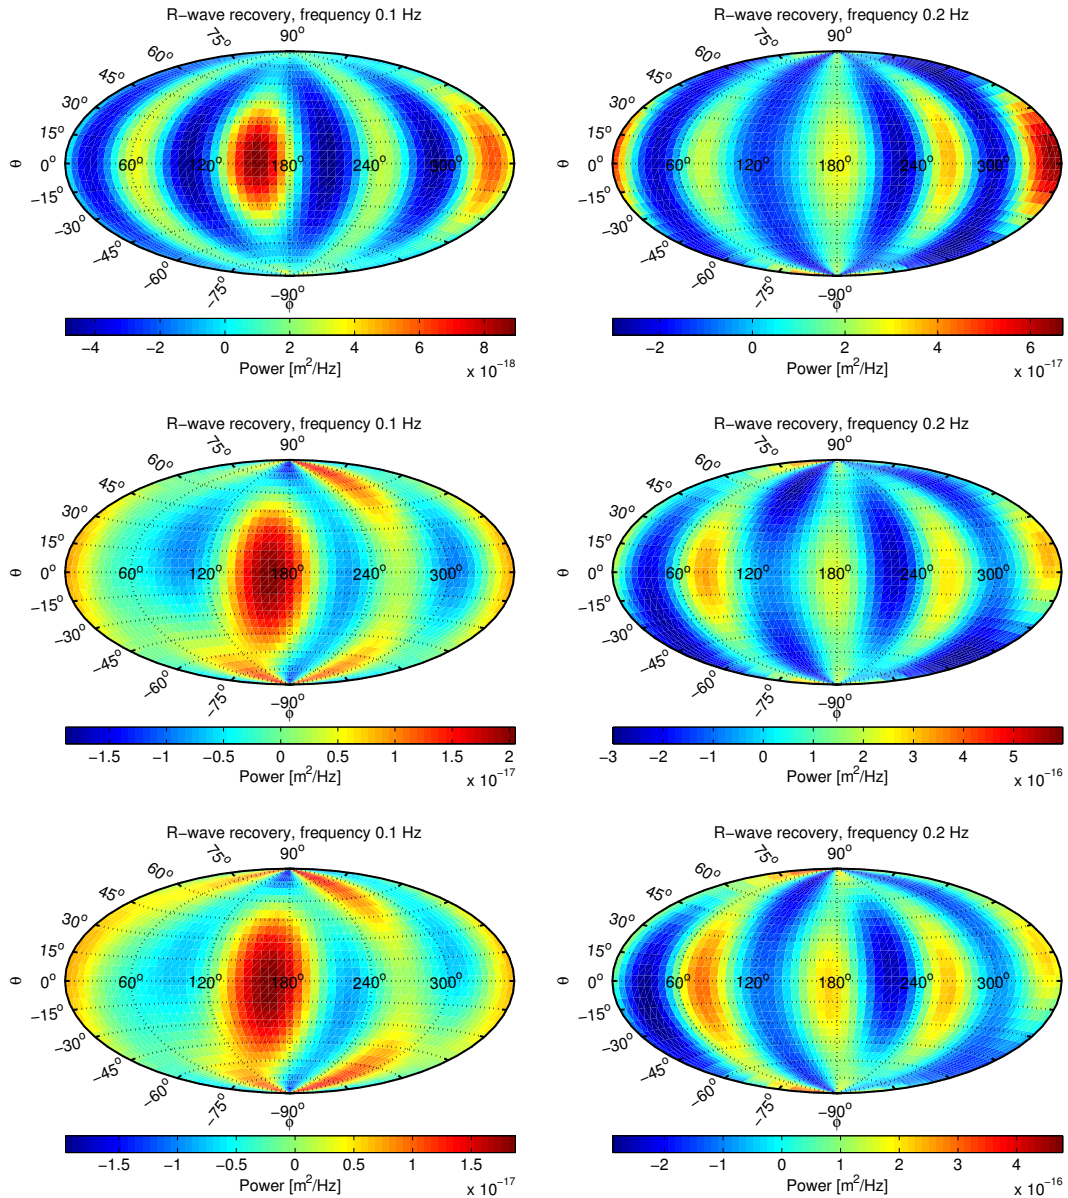


Figure 6.24: Rayleigh wave analysis of 1 hour of background noise with the seismic radiometer algorithm. Recoveries are performed at 0.1 Hz (left column) and 0.2 Hz (right column) using three different values of α : 100 m (top row), 1000 m (middle row), and 2000 m (bottom row).

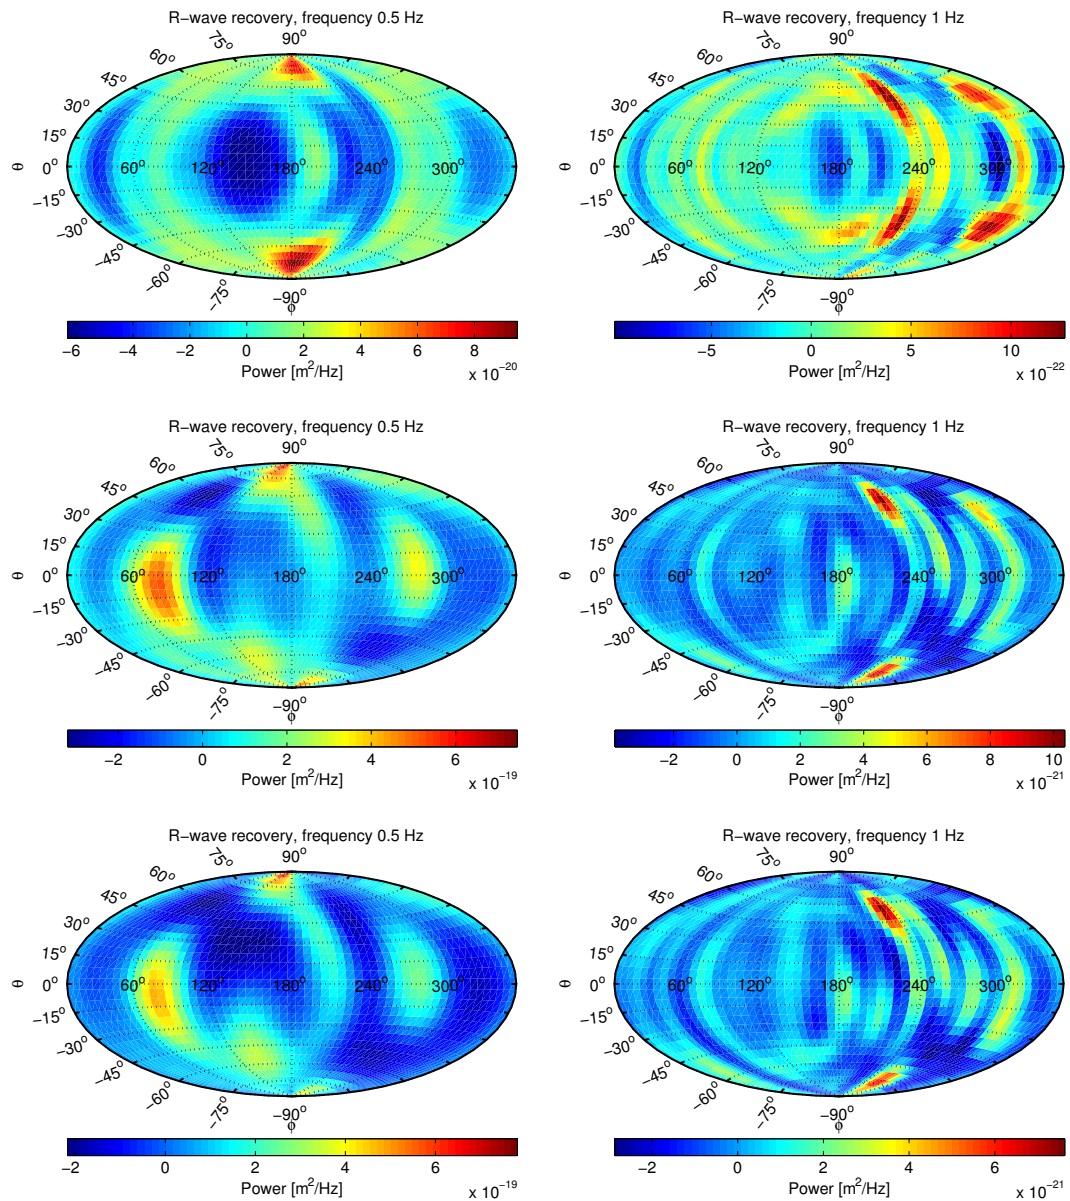


Figure 6.25: Rayleigh wave analysis of 1 hour of background noise with the seismic radiometer algorithm. Recoveries are performed at 0.5 Hz (left column) and 1 Hz (right column) using three different values of α : 100 m (top row), 1000 m (middle row), and 2000 m (bottom row).

Recovery type	Recovered power (m ²)			
	0.1 Hz	0.2 Hz	0.5 Hz	1 Hz
P-wave	4.8×10^{-19}	1.3×10^{-17}	3.8×10^{-20}	1.3×10^{-22}
S-wave horizontal	4.8×10^{-19}	1.6×10^{-17}	3.4×10^{-21}	1.5×10^{-22}
S-wave vertical	2.9×10^{-20}	3.4×10^{-19}	4.7×10^{-20}	1.0×10^{-22}
R-wave ($\alpha = 100$ m)	7.8×10^{-20}	1.1×10^{-18}	1.5×10^{-21}	1.3×10^{-23}
R-wave ($\alpha = 1000$ m)	4.0×10^{-19}	8.2×10^{-18}	3.8×10^{-20}	1.2×10^{-22}
R-wave ($\alpha = 2000$ m)	3.4×10^{-19}	7.7×10^{-18}	3.7×10^{-20}	1.0×10^{-22}

Table 6.5: Total recovered power across all directions as a function of frequency for P-waves, both S-wave polarizations, and three parametrizations of the Rayleigh wave model. See the text for more details.

6.5.1 Rayleigh wave measurements

In simple terms, the Rayleigh wave eigenfunctions encapsulate how the Rayleigh wave amplitudes change with depth and frequency, as well as how the horizontal amplitudes relate to the vertical amplitudes. I aim to understand this behavior by measuring Rayleigh wave amplitudes generated by seismic events, using the data from the Homestake 3D seismic array.

A time-series plot of a seismic event usually has at least two identifiable features: the arrival of the first P-waves from the event, and the surface waves, which arrive slightly later and are larger in amplitude. The S-wave feature may be visible between the P-waves and surface waves, but is sometimes difficult to separate from the P-waves. In Fig. 6.26, I show an example time-series from the 300 station for a 5.8 magnitude earthquake which occurred off the coast of Oregon. The P-wave arrival is visible at about 250 s as a large spike in the east and vertical channels. In the north channel, the S-wave arrival is apparent just before 500 s, and is followed closely by the surface waves at about 550 s. The velocity dispersion of the surface waves is clearly seen in the shape of the time-series: the low-frequency surface waves arrive first and higher frequencies arrive later, leading to an increasing frequency modulation of the surface wave portion of the time-series.

The surface wave portion of the signal comprises both Love waves and Rayleigh waves. These components can be separated by a rotation about the vertical axis from the




Figure 6.26: Time-series plot of all seismic channels (from top to bottom: east, north, vertical, as indicated by the channel names on the right) for the 300 station during an earthquake. The y-axis is in units of m/s and indicates ground velocity. The P-wave arrival is apparent as a large spike at approximately 250 s. The S-wave arrival is primarily noticeable in the north channel just before 500 s. The surface waves begin around 550 s and last for several minutes; the velocity dispersion is also readily apparent in the time-series.

standard coordinate system defined by the seismometer channels (east, north, vertical) into a coordinate system where one axis is aligned along the path of the seismic wave. In this coordinate system, the Rayleigh wave has a compressional component in the radial direction and a vertical shear component. The Love waves only have a horizontal shear component, so these waves should be present only in the direction transverse to the wave's path. This requires that the location of the seismic event is known so that we may calculate the geodesic path from the event's epicenter to our array.

From here, I extract the horizontal and vertical Rayleigh wave amplitudes at different frequencies and at different depths in order to determine the eigenfunctions. The full procedure is given below for completeness.

1. Load the data corresponding to a given seismic event and manually determine the location of Rayleigh waves in the time-series.
2. Apply \sin^2 tapers to the beginning and end of the time-series to prevent artifacts due to sharp edges when performing a Fourier transform.
3. Rotate from the east-north-vertical coordinate system to radial-transverse-vertical, such that the radial axis is aligned with the direction of wave propagation.
4. Filter the time-series in the radial and vertical channels at a given frequency f_0 by performing a Fourier transform, multiplying by a Gaussian function, and performing an inverse Fourier transform. I used the following functional form for the Gaussian: $\exp[-\alpha(f - f_0)^2/f_0^2]$; I tested several values and chose $\alpha = 100$ based on the resulting filter's capability to pick out well-defined wave packets within a narrow frequency range.
5. Calculate an amplitude envelope for the radial and vertical time-series using a Hilbert transform³ and find the maximum amplitude in each envelope.
6. Repeat this procedure for all stations in the array.
7. Calculate the average vertical amplitude for this event across all surface stations and normalize both the horizontal and vertical channels by this value.
8. Repeat this procedure at several frequencies for the same seismic event.
9. Repeat this procedure for several seismic events.

Note that the eigenfunctions given in Eq. 6.10 are in terms of ground displacement, while our data contain measurements of ground velocity. However, since I have normalized the amplitudes relative to the surface verticals, this should cancel out any additional coefficients which arise from the time derivative required to compute ground

³The Hilbert transform calculates the analytic representation of a signal by removing the negative frequencies. This is done by using a transfer function which is 0 for $f < 0$, 1 at $f = 0$ and 2 for $f > 0$. The result is a complex-valued signal whose magnitude is a smooth function which traces the extrema of the original signal. Examples of this amplitude envelope can be seen in blue in the top panel of Fig. 6.27.

velocity from ground displacement. As a result, it will still be possible to extract the displacement eigenfunctions based on measurements of ground velocity amplitudes.

In Fig. 6.27, I show several time-series traces corresponding to a particular seismic event and the resulting heatmap of wave velocity versus frequency. A raw time-series trace is shown along with several filtered time-series and amplitude envelopes (solid blue). The bottom panel shows a heatmap in terms of calculated wave velocity (based on the geodesic distance to the event divided by the wave travel time) versus frequency. The color scale indicates the amplitude (in terms of ground velocity) as shown in the time-series traces. The velocity dispersion is clear in the heatmap, as evidenced by the higher velocities (i.e., earlier arrival time) at lower frequencies compared to higher frequencies.

For this analysis, I have used 14 seismic events between May 20–28, 2015, many of which are generated by blasting in mines within the region. These events were identified by Gary Pavlis using data from the Homestake 3D array, as well as several regional seismic stations. Details of these events are shown in Tab. 6.6.

Event ID	Date	UTC Time	Latitude	Longitude
1	05/20/2015	21:41:32	43.8030	-105.4968
2	05/20/2015	23:03:02	44.3627	-106.0725
3	05/22/2015	16:40:53	43.8744	-105.4298
4	05/22/2015	20:23:44	43.9408	-105.4559
5	05/23/2015	21:23:46	43.8150	-105.5652
6	05/23/2015	23:33:32	43.5233	-104.5268
7	05/25/2015	18:27:49	44.1526	-105.2807
8	05/25/2015	19:13:30	44.2315	-105.3453
9	05/25/2015	20:00:26	43.9373	-105.3073
10	05/25/2015	20:13:27	43.8249	-105.4232
11	05/26/2015	20:03:36	43.6702	-105.4549
12	05/27/2015	20:35:50	43.9628	-105.5024
13	05/27/2015	22:56:36	43.8224	-105.5756
14	05/28/2015	00:28:00	43.9059	-105.4762

Table 6.6: Details of seismic events used to estimate parameters of Rayleigh wave eigenfunctions.

Not all of the stations have good data corresponding to each event; in total, I have extracted approximately 300 amplitude measurements at each frequency from the 14





Figure 6.27: Raw and filtered time-series traces and velocity heatmap corresponding to a seismic event. The solid blue envelopes are produced by the Hilbert transform of the time-series. The heatmap shows velocity versus frequency and the color scale indicates the amplitude of ground velocity corresponding to the envelopes.

events and 24 stations.

6.5.2 Parameter estimation

Method

Using the measurements extracted from seismic events, I performed parameter estimation with `MultiNest` to fit the biexponential model described by Eq. 6.10. In this model, the amplitudes are relative, so we have the freedom to normalize one of the coefficients; thus, we set the parameter C_1 to be 1. This leaves eight unknown parameters in total: the coefficients C_2 , C_3 , and C_4 , the exponents a_1 , a_2 , a_3 , and a_4 , and the phase velocity c_p , which is needed to calculate the wavevector k in the exponential terms and is not known *a priori*.

Uniform priors between 0 and 1 were used for all four a parameters. For the C parameters, uniform priors were used as well, but with different ranges for each parameter. C_2 was allowed to vary between -3-3; however, there is some degeneracy between C_3 and C_4 . Specifically, C_3 and C_4 can be interchanged with no loss of generality as long as a_3 and a_4 are interchanged as well. Preliminary studies indicated that in general, there was a preference for the two parameters to have different signs such that their sum was ≈ 1 ; as a result, I have fixed the prior on C_3 to be between 0-3 and the prior on C_4 to be between -3-0. Finally, the prior on c_p was set to be uniform between 1000-4000 m/s, based on general knowledge of expected Rayleigh wave velocities.

I have used a likelihood function which is based on a χ^2 comparison of the model (for a given set of parameters) to the data.

$$\mathcal{L} = \exp \left\{ - \sum_i \left[(M_H(C_2, a_1, a_2, c_p; z_i) - H_i)^2 + (M_V(C_3, C_4, a_3, a_4, c_p; z_i) - V_i)^2 \right] \right\} \quad (6.33)$$

Here, the index i runs over measurements; z_i is the depth corresponding to measurement i , and H_i and V_i are the normalized horizontal and vertical amplitudes from measurement i . M_H is the normalized horizontal amplitude at depth z_i , generated from the eigenfunction model with parameters C_2 , a_1 , a_2 , and c_p (see Eq. 6.10, line 1). M_V is the normalized vertical amplitude at depth z_i , generated from the eigenfunction model with parameters C_3 , C_4 , a_3 , a_4 , and c_p (see Eq. 6.10, line 2). Equal weighting was used for all

measurements. Note that an independent parameter estimation analysis was performed for each frequency since the parameters are expected to be a function of frequency.

Results

I have performed this analysis at frequencies from 0.1 Hz up to 1 Hz in intervals of 0.1 Hz. For each frequency, I extracted the maximum likelihood point sampled by `MultiNest` and used the resulting parameters to construct the eigenfunctions. The eigenfunction parameters for all frequencies are summarized in Tab. 6.7.

Overall, the C_2 parameter is close to zero at all frequencies, which makes the r_1 eigenfunction act essentially as a single exponential; in general, this also allows large variability in the a_2 parameter to have little effect on the resulting amplitudes. This is understandable given the measurements and the depth of our array; with a deeper array, it might be possible to see the eigenfunction amplitudes cross zero and increase again, which would necessitate the biexponential form.

Overall, the phase velocities mostly follow a decreasing trend with frequency up to about 0.7 Hz, as we would expect based on the velocity dispersion shown in Fig. 6.27. At higher frequencies, the eigenfunctions do not provide a particularly good fit to the measurements and the resulting phase velocities do not fit well with expectations. This may be due to some degeneracy between the a parameters and the phase velocity: in the future, it may be more effective to use four parameters which represent the product of the a parameters with c_p^{-1} , rather than treating the a parameters and c_p as five independent parameters.

To illustrate the results, I have plotted the mean amplitude measurement at each depth along with its standard error, compared to the eigenfunctions generated using the estimated parameters. Fig. 6.28 shows this comparison at frequencies between 0.1–0.6 Hz, while Fig. 6.29 shows the higher frequencies between 0.7–1.0 Hz. The horizontal data at 0.1 Hz are quite messy and the resulting eigenfunction, although mostly within the error bars, does not trace the behavior of the mean measurements very well. This is likely due to the lack of significant seismic content below ≈ 0.2 Hz (see Fig. 6.27, bottom panel).

There is very good agreement between the measurements and the models at 0.2 and 0.3 Hz. As the frequency increases, the horizontal models deviate more from the

Frequency [Hz]	C_2	C_3	C_4	a_1	a_2	a_3	a_4	c_p [m/s]
0.1	0.12	1.77	-0.78	0.99	0.21	0.12	0.32	1799
0.2	-0.07	2.36	-1.36	0.59	0.43	0.36	0.78	2599
0.3	-0.11	2.16	-1.14	0.63	0.94	0.30	0.79	2025
0.4	-0.05	2.46	-1.47	0.45	0.86	0.29	0.69	1795
0.5	0.07	2.31	-1.32	0.55	0.04	0.30	0.87	1781
0.6	0.001	2.92	-1.94	0.38	0.21	0.39	0.84	1997
0.7	0.04	2.09	-1.11	0.28	0.81	0.25	0.77	1620
0.8	-0.03	2.00	-1.00	0.37	0.76	0.28	0.86	2087
0.9	0.04	2.11	-1.12	0.11	0.63	0.16	0.46	1107
1.0	-0.05	1.95	-0.95	0.08	0.06	0.19	0.99	1208

Table 6.7: Parameters of Rayleigh wave eigenfunctions resulting from parameter estimation with `MultiNest`.

measurements and generally predict lower amplitudes at depth than observed. At higher frequencies, the horizontal measurements seem to plateau at the deepest levels; one hypothesis for this behavior is that the Rayleigh wave amplitude drops below that of the background noise at these frequencies. This may also occur due to the lack of seismic content above ≈ 0.7 Hz, as is evident in the example heatmap shown in Fig. 6.27.

Overall, the eigenfunctions provide a good fit to the data (based on a reduced χ^2 calculation), other than the horizontals at 1 Hz. A few of the vertical fits also have a somewhat high reduced χ^2 , even though they trace the mean measurements quite well. This may indicate that the errors have been underestimated.

Several features of the resulting eigenfunctions match well with expectations. First, the depth suppression occurs on longer length scales for lower frequencies; i.e., lower frequencies are suppressed less at depth compared to higher frequencies. Second, the overall shape of the eigenfunctions agrees with that of Ref. [336]: the horizontals begin dropping off immediately with depth, while the verticals actually increase slightly at first and then begin to decrease. Finally, the ratio of horizontal to vertical amplitudes at the surface is between 0.9–1.1 for all frequencies other than 0.1 Hz [335].

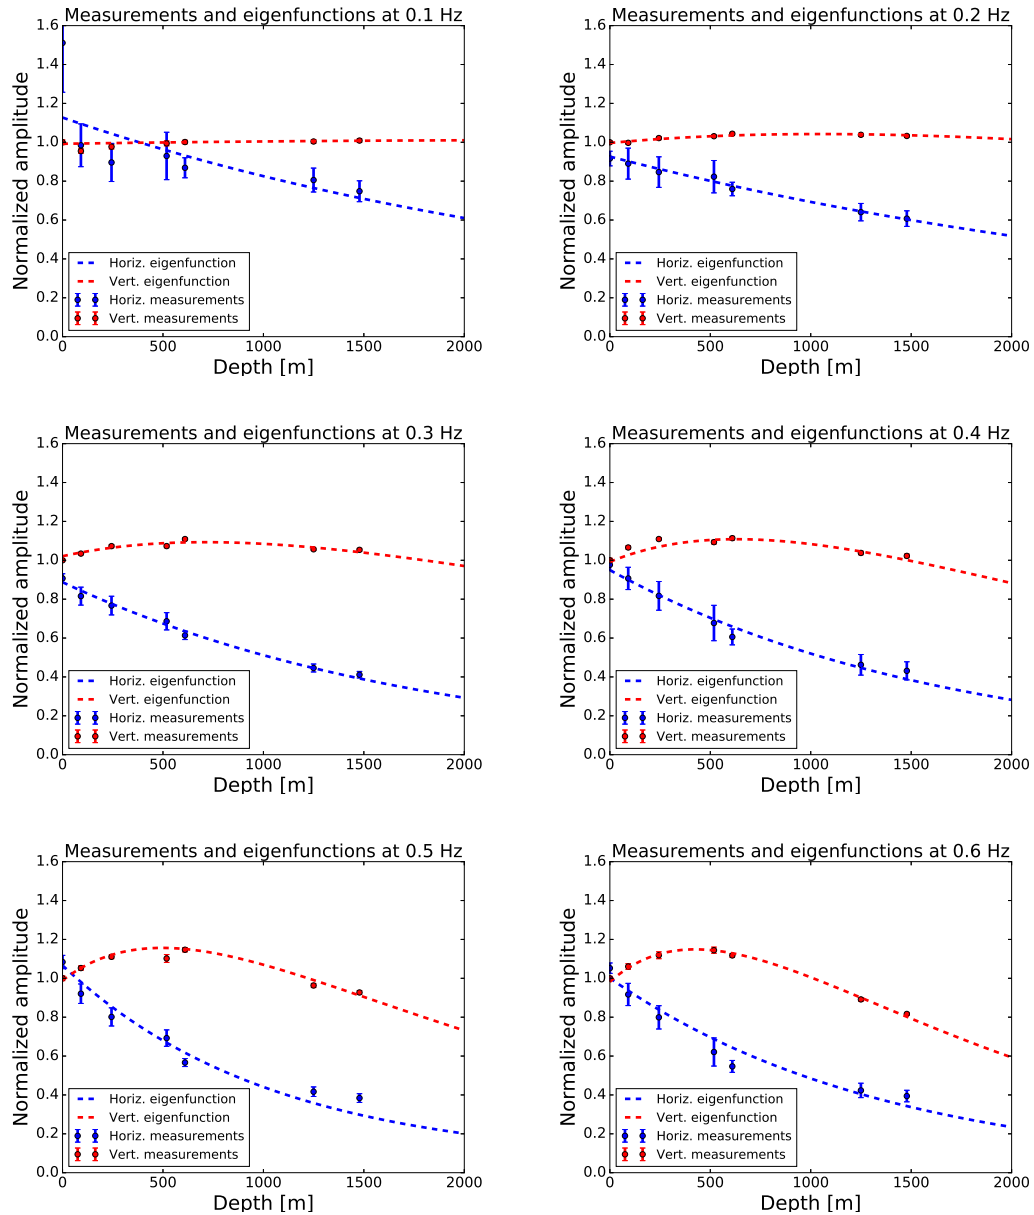


Figure 6.28: Rayleigh wave amplitude measurements (solid circles) and eigenfunction models (dashed lines) for 0.1–0.6 Hz. Measurements shown are mean measurements at each depth and the error bars show the standard error on the mean. In some cases, error bars are too small to be seen.

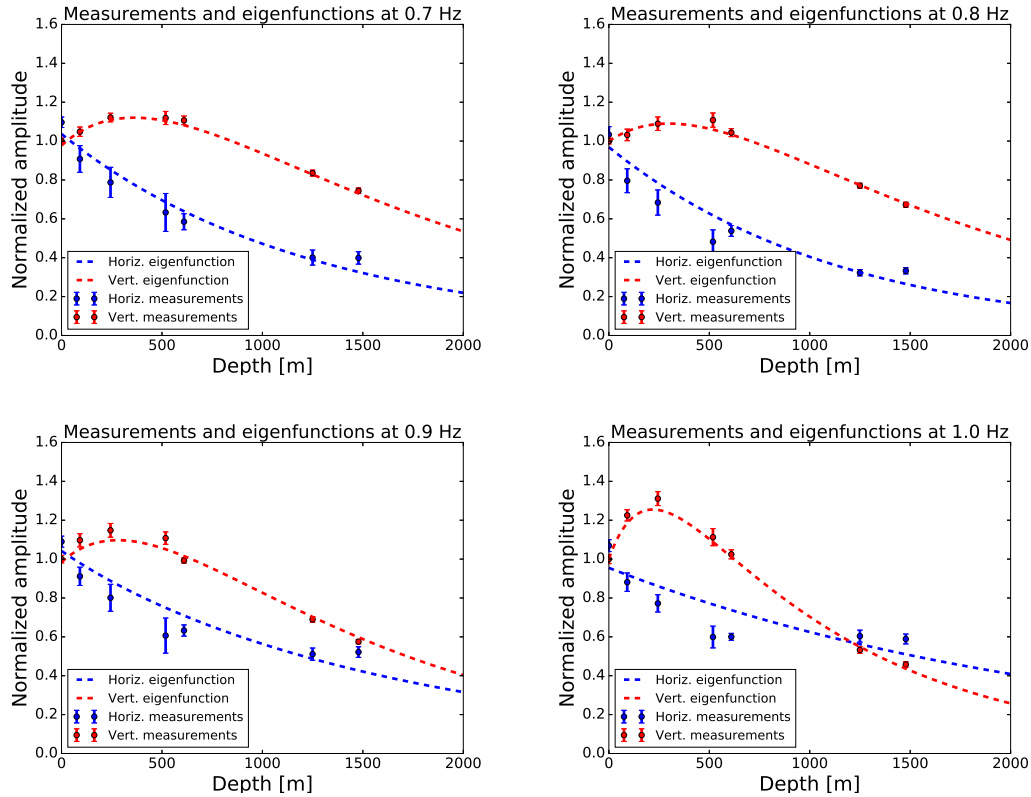


Figure 6.29: Rayleigh wave amplitude measurements (solid circles) and eigenfunction models (dashed lines) for 0.7–1.0 Hz. Measurements shown are mean measurements at each depth and the error bars show the standard error on the mean. In some cases, error bars are too small to be seen.

6.5.3 Application to seismic radiometer

Using these eigenfunctions, the Rayleigh wave displacement field looks like

$$\vec{r}(\vec{x}, t) = \left(e^{-a_1 k z} + C_2 e^{-a_2 k z} \right) \cos(\omega t - \vec{k} \cdot \vec{x}) \hat{k} + \left(C_3 e^{-a_3 k z} + C_4 e^{-a_4 k z} \right) \sin(\omega t - \vec{k} \cdot \vec{x}) \hat{z}, \quad (6.34)$$

where the wavevector k is given by $2\pi f/c_p$. Following the previous calculation for Rayleigh waves in Sec. 6.4, we obtain the following gamma functions:

$$\gamma_{R,n,i\alpha j\beta} = \int d\hat{\Omega} Q_n(\hat{\Omega}) e^{2\pi i f \hat{\Omega} \cdot \vec{x}/c_p} \left[(r_1 \hat{\Omega} \cdot \hat{\alpha} + r_2 e^{i\pi/2} \hat{z} \cdot \hat{\alpha}) (r_1 \hat{\Omega} \cdot \hat{\beta} + r_2 e^{-i\pi/2} \hat{z} \cdot \hat{\beta}) \right], \quad (6.35)$$

where we have abbreviated the eigenfunctions as r_1 and r_2 ; see Eq. 6.10. Finally, the cross-correlation estimator is given as a sum over the new gamma functions:

$$\langle Y_{i\alpha j\beta} \rangle = 2T\Delta f R_n \gamma_{R,n,i\alpha j\beta}. \quad (6.36)$$

I have implemented the model between 0.1–1 Hz in the seismic radiometer analysis; in Fig. 6.30, I show the Rayleigh wave recoveries of 1 hour of background noise at 0.1 Hz, 0.2 Hz, 0.5 Hz, and 1 Hz. Summing up the total power in all directions, I obtained (in order of increasing frequency): $1.6 \times 10^{-19} \text{ m}^2$, $6.1 \times 10^{-18} \text{ m}^2$, $3.7 \times 10^{-22} \text{ m}^2$, and $1.2 \times 10^{-29} \text{ m}^2$. These numbers are comparable to the initial Rayleigh wave analysis at 0.1 Hz and 0.2 Hz (see Tab. 6.5); however, they are significantly decreased at the higher frequencies. Again, the recovery at 0.2 Hz contains the most overall power, likely due to the microseismic peak.

The overall distribution of power is in fairly good agreement with the old model, other than at 1 Hz. At 0.1 Hz and 0.2 Hz, we see little θ dependence, as expected; however, this is not the case at the higher frequencies and is likely due to the lack of power at those frequencies.

Overall, it is not obvious that the radiometer recoveries have been improved by the implementation of the new Rayleigh wave model; however, it is a more robust and technically correct model. The previous model cannot account for the frequency dependence of the eigenfunctions or the initial increase in the vertical eigenfunctions, both of which are clearly displayed by our measurements of Rayleigh wave amplitudes. Further work should be done to compare these results to other analyses which map the directional distribution of power and/or track the microseismic peak in order to obtain better criteria for performing this comparison.

6.5.4 Discussion

I have shown that we can measure Rayleigh wave amplitudes and obtain relatively good fits using a biexponential functional form for the eigenfunctions. Future work may improve on the results shown here in several ways. First, a more sophisticated algorithm for extracting Rayleigh wave amplitudes could improve the quality of the resulting measurements. This could take the form of some sort of “clustering” across

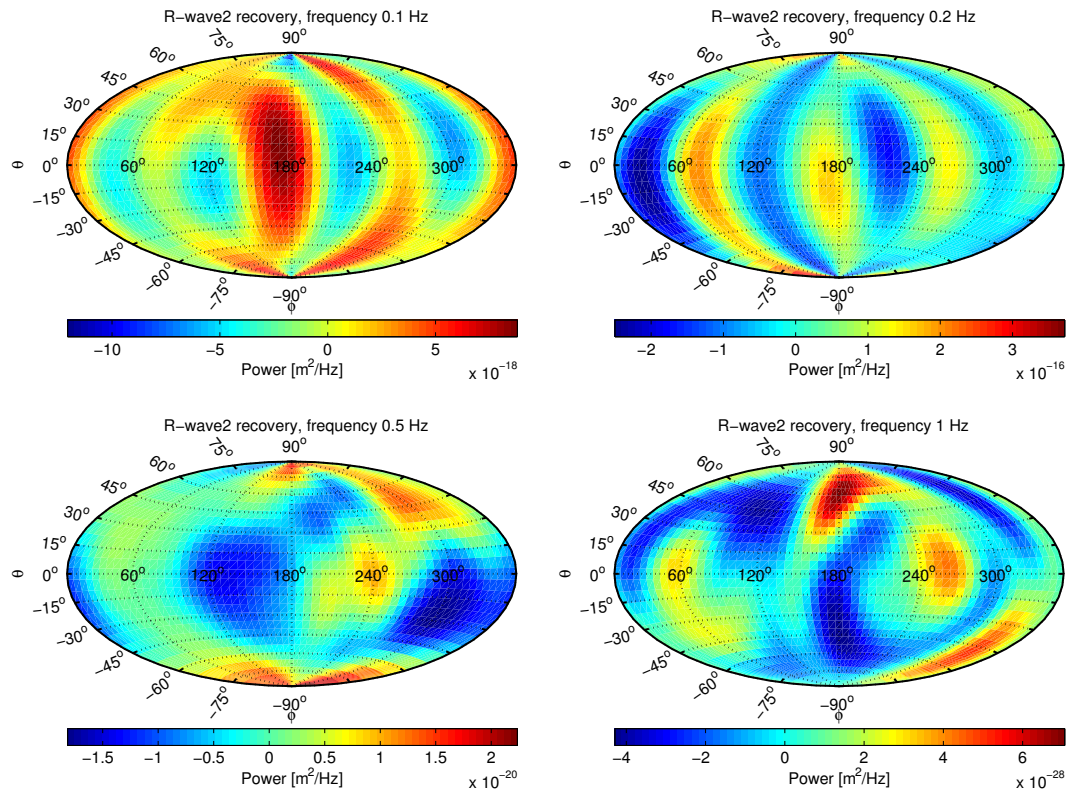


Figure 6.30: Analysis of 1 hour of background noise with the biexponential Rayleigh wave model in the seismic radiometer algorithm. The analysis is performed at 0.1 Hz (top-left), 0.2 Hz (top-right), 0.5 Hz (bottom-left) and 1 Hz (bottom-right).

different frequencies and/or stations in order to ensure that the same peak is used for each measurement. Along with this, a more complete event catalog would allow for a significant increase in the number of data points and likely result in more robust measurements, as well as decreasing the effect of any spurious measurements which might be present.

For this study, I have used nearby mine blasts as the source of Rayleigh waves; however, it would also be useful to pursue a similar study using earthquakes. These data are easily accessible through IRIS and will likely provide very clean and consistent data across stations. However, these distant events often carry little information above ≈ 0.2 Hz, since these frequencies are damped due to the heterogeneity of the material in the surface layers. As a result, an earthquake analysis may prove challenging, since our

array is better designed to probe the eigenfunctions of higher-frequency seismic waves, whose structure changes appreciably over length scales comparable to our array depth.

6.6 Summary

In this chapter, I have focused on an effort to study seismic noise using a 3D array of seismometers in the Homestake Mine with the goal of improving the sensitivity of future interferometric GW detectors. I discussed the effect of seismic and Newtonian noise on these detectors, as well as some methods for mitigating these noise sources. I gave an overview of the recently installed 3D seismic array in the Homestake Mine and described a web interface which I developed to monitor the status and quality of data from the seismic stations in the array. To study the seismic data and decompose it into different components, I introduced a seismic radiometer method which may be used to identify directional sources of seismic waves or the seismic noise background. I used this algorithm to analyze simulations of body and surface waves and a period of real data from the array. Finally, I extracted measurements of Rayleigh wave amplitudes from data collected at Homestake and performed parameter estimation to implement a more complex model of Rayleigh waves, which was then added to the seismic radiometer algorithm.

Future studies are expected to continue this work and eventually estimate strain spectra due to seismic and Newtonian noise, following the theory outlined in Ref. [320].

Chapter 7

Conclusions

Within the last few months, LIGO has announced the first two direct detections of gravitational waves [10, 337]. Advanced LIGO is expected to continually improve in sensitivity during its upcoming observational runs, and additional detectors, including Advanced Virgo and KAGRA, should come online in the next few years. The field of gravitational-wave astronomy is just beginning and should yield an abundance of new insights into the universe.

The central goal of this dissertation is the detection and characterization of long-lasting GW signals, in terms of both a persistent background of gravitational waves and long-lasting transient GW signals. These signals are expected to be integral components of gravitational-wave astronomy for the important information they may yield about the beginnings of the universe and the various objects contained within it.

Chap. 3 focused on efforts to detect and model a stochastic gravitational-wave background. I gave an overview of possible sources of such a background and described the LIGO search strategy geared towards its detection. Using a Bayesian parameter estimation method, I projected the sensitivity of Advanced LIGO and the Einstein Telescope to an SGWB produced by core-collapse supernovae. I compared the results of this study to GW waveforms produced by core-collapse simulations and found the resulting limits to be 1–2 orders of magnitude higher than the expected strain from the simulations. Using a model of cosmic chemical evolution model [146], I also fit a model of the star formation rate at high redshift, since it is an important component of astrophysical SGWB models.

In Chap. 4, I described the STAMP search pipeline, a new search algorithm designed

to detect GW signals lasting between tens and hundreds of seconds. This unmodeled search pipeline uses an excess power method and image processing techniques to identify significant groupings of energy and localize them in time and frequency. I have developed a new data quality method which significantly improves the sensitivity of STAMP searches by identifying and removing instrumental artifacts from the data. I have also developed two new clustering algorithms which parse the data and are able to identify groupings of pixels of arbitrary shape which may correspond to GW signals. Finally, I have implemented features in STAMP which make it capable of detecting long-lasting transient GWs coming from any sky direction at any time.

In Chap. 5, I described the application of this STAMP all-sky search to LIGO S5 and S6 data, targeting signals with durations between 10–500 s. In total, I analyzed over 400 days of data in the 40-1000 Hz frequency band, with minimal assumptions about the signal waveform, polarization, source direction, or time of occurrence. To project the sensitivity of this search to real signals, 15 different simulated GW waveforms were analyzed, including signals produced by a black hole accretion disk instability, and several *ad hoc* waveforms. Using the results of the analysis, I set the first upper limits from an all-sky search for unmodeled long-duration transient gravitational waves. An analysis of the first Advanced LIGO data with this pipeline is currently underway, but is outside the scope of this dissertation.

Finally, in Chap. 6, I considered the effect of seismic noise and Newtonian noise on terrestrial interferometric GW detectors. Because these noise sources primarily affect the lower end of LIGO’s frequency band, mitigating them will be particularly important for detecting an SGWB, as well as increasing the length of time a binary coalescence signal is in the sensitive band of the detectors. I described the planning and installation of a three-dimensional array of seismometers in and around the Homestake Mine in Lead, SD, with the goal of characterizing the seismic noise background in order to test underground conditions for future generations of GW detectors. For this array, we have implemented a novel timing network for transporting a GPS signal underground, as well as a low-latency data quality interface for monitoring the overall state of most of the seismic stations in the network. I also implemented and tested a seismic radiometer algorithm for decomposing the seismic field into different components and studying its directional distribution. Finally, I extracted measurements of Rayleigh wave amplitudes

from the Homestake data and used these measurements to perform parameter estimation in order to fit a model of Rayleigh wave eigenfunctions. Future studies will expand on this work in order to estimate strain spectra from seismic and Newtonian noise, which can be used for planning the next generation of gravitational-wave detectors.

References

- [1] B. P. Abbott *et al.*, *LIGO: the Laser Interferometer Gravitational-Wave Observatory*, Rep. Prog. Phys. **72**, 076901 (2009).
- [2] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*, 1st ed. (W. H. Freeman and Company, 1973).
- [3] J. A. Wheeler, *A Journey Into Gravity and Spacetime*, 1st ed. (W. H. Freeman and Company, 1990).
- [4] S. M. Carroll, *Spacetime and Geometry: An Introduction to General Relativity*, 1st ed. (Addison Wesley, 2005).
- [5] U. Le Verrier, Cr. Hebd. Acad. Sci. **49**, 379 (1859).
- [6] A. Einstein, Preuss. Akad. Wiss. Berlin, Sitzber. **1915 (part 2)**, 831 (1915).
- [7] R. V. Pound and G. A. Rebka, Jr., Phys. Rev. Lett. **4**, 337 (1960).
- [8] F. W. Dyson, A. S. Eddington, and C. Davidson, Philos. T. R. Soc. **220A**, 291 (1920).
- [9] C. W. F. Everitt *et al.*, Phys. Rev. Lett. **106**, 221101 (2011).
- [10] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Observation of gravitational waves from a binary black hole merger*, Phys. Rev. Lett. **116**, 061102 (2016).
- [11] P. R. Saulson, *Fundamentals of interferometric gravitational wave detectors* (World Scientific, 1994).

- [12] K. A. Postnov and L. R. Yungelson, *Liv. Rev. Rel.* **9**, 6 (2006).
- [13] J. A. Faber and F. A. Rasio, *Liv. Rev. Rel.* **15**, 8 (2012).
- [14] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **89**, 102006 (2014).
- [15] A. Taracchini *et al.*, *Phys. Rev. D* **89**, 061502 (2014).
- [16] S. Sigurdsson and L. Hernquist, *Nature* **364**, 423 (1993).
- [17] S. F. Portegies Zwart and S. L. W. McMillan, *Astrophys. J.* **528**, L17 (2000).
- [18] E. Thrane, S. Kandhasamy, C. D. Ott, *et al.*, *Phys. Rev. D* **83**, 083004 (2011).
- [19] K. Riles, *Prog. Part. Nucl. Phys.* **68**, 1 (2013).
- [20] L. Lindblom, B. J. Owen, and S. M. Morsink, *Phys. Rev. Lett.* **80**, 4843 (1998).
- [21] N. Andersson, *Astrophys. J.* **502**, 708 (1998).
- [22] E. Gaertig, K. Glampedakis, K. D. Kokkotas, and B. Zink, *Phys. Rev. Lett.* **107**, 101102 (2011).
- [23] A. Mytidis, M. Coughlin, and B. Whiting, *Astrophys. J.* **810**, 27 (2015).
- [24] C. D. Ott, *Class. Quantum Grav.* **26**, 6 (2009).
- [25] T. Regimbau and V. Mandic, *Class. Quantum Grav.* **25**, 184018 (2008).
- [26] R. Brustein, M. Gasperini, M. Giovannini, and G. Veneziano, *Phys. Lett. B* **361**, 45 (1995).
- [27] A. Buonanno, M. Maggiore, and C. Ungarelli, *Phys. Rev. D* **55**, 3330 (1997).
- [28] V. Mandic and A. Buonanno, *Phys. Rev. D* **73**, 063008 (2006).
- [29] X. Siemens, V. Mandic, and J. D. Creighton, *Phys. Rev. Lett.* **98**, 111101 (2007).
- [30] T. W. B. Kibble, *J. Phys. A: Math. Gen.* **9**, 1387 (1976).
- [31] L. P. Grishchuk, *Sov. Phys. JETP* **40**, 409 (1975).

- [32] L. P. Grishchuk, *Class. Quantum Grav.* **10**, 2449 (1993).
- [33] R. A. Hulse and J. H. Taylor, *Astrophys. J.* **195**, L51 (1975).
- [34] J. H. Taylor and J. M. Weisberg, *Astrophys. J.* **253**, 908 (1982).
- [35] J. M. Weisberg, D. J. Nice, and J. H. Taylor, *Timing measurements of the relativistic binary pulsar PSR B1913+16*, *Astrophys. J.* **722**, 1030 (2010).
- [36] J. Weber, *Phys. Rev. Lett.* **17**, 1228 (1967).
- [37] J. Weber, *Phys. Rev. Lett.* **20**, 1307 (1968).
- [38] J. Weber, *Phys. Rev. Lett.* **22**, 1320 (1969).
- [39] E. Mauceli *et al.*, *Phys. Rev. D* **54**, 1264 (1996).
- [40] D. G. Blair *et al.*, *Phys. Rev. Lett.* **74**, 1908 (1995).
- [41] P. Astone *et al.*, *Phys. Lett. B* **385**, 421 (1996).
- [42] M. Cerdonino *et al.*, *Class. Quantum Grav.* **14**, 1491 (1997).
- [43] O. D. Aguiar *et al.*, *Class. Quantum Grav.* **25**, 114042 (2008).
- [44] G. Hobbs *et al.*, *Class. Quantum Grav.* **27**, 084013 (2010).
- [45] D. J. Reardon *et al.*, *Mon. Not. R. Astron. Soc.* **455**, 1751 (2016).
- [46] L. Lentati *et al.*, *Mon. Not. R. Astron. Soc.* **453**, 2576 (2015).
- [47] F. Jenet *et al.*, arXiv: 0909.1058 (2009).
- [48] A. Abramovici *et al.*, *Science* **256**, 325 (1992).
- [49] F. Acernese for the Virgo Collaboration, *Class. Quantum Grav.* **23**, S63 (2006).
- [50] S. Hild for the LIGO Scientific Collaboration, *Class. Quantum Grav.* **23**, S643 (2006).
- [51] K. Kawabe and the TAMA collaboration, *Class. Quantum Grav.* **14**, 1477 (1997).

- [52] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Class. Quantum Grav.* **32**, 074001 (2015).
- [53] F. Acernese *et al.* (Virgo Collaboration), *Class. Quantum Grav.* **32**, 024001 (2014).
- [54] K. Kuroda, *Class. Quantum Grav.* **27**, 084004 (2010).
- [55] B. Iyer *et al.*, LIGO DCC #M1100296 (2011).
- [56] M. Punturo *et al.*, *Class. Quantum Grav.* **27**, 084007 (2010).
- [57] B. S. Sathyaprakash *et al.*, *Class. Quantum Grav.* **30**, 079501 (2013).
- [58] K. Danzmann and the LISA study team, *Class. Quantum Grav.* **13**, A247 (1996).
- [59] P. Amaro-Seoane *et al.*, arXiv: 1201.3621 (2012).
- [60] G. M. Harry, P. Fritschel, D. A. Shaddock, W. Folkner, and E. S. Phinney, *Class. Quantum Grav.* **23**, 4887 (2006).
- [61] J. Crowder and N. J. Cornish, *Phys. Rev. D* **72**, 083005 (2005).
- [62] S. Sato *et al.*, *J. Phys.: Conf. Ser.* **154**, 012040 (2009).
- [63] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **93**, 122004 (2016).
- [64] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **93**, 122003 (2016).
- [65] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241102 (2016).
- [66] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), arXiv: 1602.03842 (2016).
- [67] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 221101 (2016).
- [68] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **818**, L22 (2016).

- [69] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 131102 (2016).
- [70] S. Adrián-Martínez *et al.* (ANTARES Collaboration, IceCube Collaboration, LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **93**, 122010 (2016).
- [71] E. W. Kolb and M. S. Turner, *The Early Universe*, 1st ed. (Westview Press, 1994).
- [72] S. Weinberg, *Cosmology*, 1st ed. (Oxford University Press, 2008).
- [73] B. Ryden, *Introduction to Cosmology*, 1st ed. (Addison-Wesley, 2002).
- [74] P. A. R. Ade *et al.* (Planck Collaboration), arXiv: 1502.01589 (2015).
- [75] A. H. Guth, *Phys. Rev. D* **23**, 347 (1981).
- [76] A. D. Linde, *Phys. Lett. B.* **108**, 389 (1982).
- [77] A. Albrecht and P. J. Steinhardt, *Phys. Rev. Lett.* **48**, 1220 (1982).
- [78] A. D. Linde, *Mod. Phys. Lett. A* **1**, 81 (1986).
- [79] P. J. Steinhardt, *The Very Early Universe: Proceedings of the Nuffield Workshop, Cambridge, 21 June - 9 July, 1982*, pp. 251–266, 1983.
- [80] J. A. Peacock, *Cosmological Physics*, 1st ed. (Cambridge University Press, 1999).
- [81] A. D. Sakharov, *JETP Lett.* **5**, 24 (1967).
- [82] T. M. Bania, R. T. Rood, and D. S. Balsler, *Nature* **45**, 54 (2002).
- [83] M. Peimbert, V. Luridiana, and A. Peimbert, *Astrophys. J.* **666**, 636 (2007).
- [84] J. Meléndez and I. Ramírez, *Astrophys. J.* **615**, L33 (2004).
- [85] M. Maggiore, *Phys. Rep.* **331**, 283 (2000).
- [86] B. Allen, *The stochastic gravity wave background: sources and detection*, in *Relativistic gravitation and gravitational radiation. Proceedings, School of Physics, Les Houches, France, September 26 - October 6, 1995*, 1996, gr-qc/9604033.

- [87] R. H. Cyburt, B. D. Fields, K. A. Olive, and E. Skillman, *Astropart. Phys.* **23**, 313 (2005).
- [88] B. P. Abbott *et al.*, *Nature* **460**, 990 (2009).
- [89] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Improved upper limits on the stochastic gravitational-wave background from 2009-2010 LIGO and Virgo data*, *Phys. Rev. Lett.* **113**, 231101 (2014).
- [90] A. A. Penzias and R. W. Wilson, *Astrophys. J.* **142**, 419 (1965).
- [91] N. W. Boggess *et al.*, *Astrophys. J.* **397**, 420 (1992).
- [92] C. L. Bennett *et al.*, *Astrophys. J. Suppl. S.* **208**, 20 (2013).
- [93] J. M. Lamarre *et al.*, *New Astron. Rev.* **47**, 1017 (2003).
- [94] R. Adam *et al.* (Planck Collaboration), arXiv: 1502.01582 (2015).
- [95] N. Aghanim *et al.* (Planck Collaboration), arXiv: 1507.02704 (2015).
- [96] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel, *Nature* **420**, 772 (2002).
- [97] C. Pryke *et al.*, *Astrophys. J.* **692**, 1247 (2009).
- [98] H. C. Chiang *et al.*, *Astrophys. J.* **711**, 1123 (2010).
- [99] D. Barkats *et al.*, *Astrophys. J.* **783**, 67 (2014).
- [100] A. T. Crites *et al.*, *Astrophys. J.* **805**, 36 (2015).
- [101] D. Hanson *et al.* (SPTpol Collaboration), *Phys. Rev. Lett.* **111**, 141301 (2013).
- [102] P. A. R. Ade *et al.* (BICEP2 Collaboration), *Phys. Rev. Lett.* **112**, 241101 (2014).
- [103] B. P. Abbott *et al.*, *Nucl. Instr. Meth. A* **517**, 154 (2004).
- [104] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 062001 (2005).

- [105] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. Lett.* **95**, 221101 (2005).
- [106] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Class. Quantum Grav.* **24**, 5343 (2007).
- [107] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **85**, 122007 (2012).
- [108] B. J. Meers, *Phys. Rev. D* **38**, 2317 (1988).
- [109] S. Sato *et al.*, *Appl. Optics* **39**, 4616 (2000).
- [110] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, *Appl. Phys. B* **31**, 97 (1983).
- [111] J. Giaime, P. Saha, D. Shoemaker, and L. Sievers, *Rev. Sci. Instrum.* **67**, 208 (1996).
- [112] M. W. Regehr, F. J. Raab, and S. E. Whitcomb, *Opt. Lett.* **20**, 1507 (1995).
- [113] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Nature Phys.* **7**, 962 (2011).
- [114] F. Acernese *et al.*, *Class. Quantum Grav.* **25**, 184001 (2008).
- [115] The LIGO Scientific Collaboration and the Virgo Collaboration, LIGO DCC #T0900389 (2009).
- [116] C. W. Helmstrom, *Statistical Theory of Signal Detection*, 2nd ed. (Pergamon Press, 1968).
- [117] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [118] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **88**, 062001 (2013).
- [119] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **85**, 082002 (2012).

- [120] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astron. Astrophys.* **541**, A155 (2012).
- [121] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **760**, 12 (2012).
- [122] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astron. Astrophys.* **539**, A124 (2012).
- [123] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J. Suppl. S.* **211**, 7 (2014).
- [124] P. A. Evans *et al.* (Swift Collaboration, LIGO Scientific Collaboration, Virgo Collaboration), *Astrophys. J. Suppl. S.* **203**, 28 (2012).
- [125] M. G. Aartsen *et al.* (IceCube Collaboration, LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **90**, 102002 (2014).
- [126] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **89**, 122003 (2014).
- [127] The LIGO Scientific Collaboration and the Virgo Collaboration, LIGO DCC #T1500055 (2015).
- [128] B. P. Abbott *et al.*, *Astrophys. J.* **683**, L45 (2008).
- [129] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **737**, 93 (2011).
- [130] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **91**, 022004 (2015).
- [131] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **785**, 119 (2014).
- [132] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **813**, 39 (2015).

- [133] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **88**, 102002 (2013).
- [134] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **91**, 062008 (2015).
- [135] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. Lett.* **102**, 111102 (2009).
- [136] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Class. Quantum Grav.* **31**, 165014 (2014).
- [137] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **90**, 062010 (2014).
- [138] S. W. Ballmer, *Class. Quantum Grav.* **23**, S179 (2006).
- [139] E. Thrane *et al.*, *Phys. Rev. D* **80**, 122002 (2009).
- [140] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **107**, 271102 (2011).
- [141] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **88**, 122004 (2013).
- [142] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **93**, 042005 (2016).
- [143] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *GW150914: The Advanced LIGO detectors in the era of first discoveries*, *Phys. Rev. Lett.* **116**, 131103 (2016).
- [144] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Living Rev. Relativ.* **19**, 1 (2016).
- [145] J. Miller, L. Barsotti, S. Vitale, P. Fritschel, and M. Evans, *Phys. Rev. D* **91**, 062005 (2015).

- [146] E. Vangioni, K. A. Olive, T. Prestegard, J. Silk, P. Petitjean, and V. Mandic, *The impact of star formation and gamma ray burst rates at high redshift on cosmic chemical evolution and reionization*, Mon. Not. R. Astron. Soc. **447**, 2575 (2015).
- [147] C. Wu, V. Mandic, and T. Regimbau, Phys. Rev. D **84**, 104024 (2012).
- [148] N. Barnaby, E. Pajer, and M. Peloso, Phys. Rev. D **85**, 023525 (2012).
- [149] M. S. Turner, Phys. Rev. D **55**, R435(R) (1997).
- [150] L. A. Boyle and A. Buonanno, Phys. Rev. D **78**, 043531 (2008).
- [151] I. Sendra and T. L. Smith, Phys. Rev. D **85**, 123002 (2012).
- [152] Z. Arzoumanian *et al.* (The NANOGrav Collaboration), Astrophys. J. **821**, 13 (2016).
- [153] B. Allen and J. D. Romano, Phys. Rev. D **59**, 102001 (1999).
- [154] A. Lue, L. Wang, and M. Kamionkowski, Phys. Rev. Lett. **83**, 1506 (1999).
- [155] A. A. Starobinskii, Sov. Phys. JETP Lett. **30**, 719 (1979).
- [156] S. Khlebnikov and I. Tkachev, Phys. Rev. D **56**, 653 (1997).
- [157] L. Kofman, A. Linde, and A. A. Starobinsky, Phys. Rev. Lett. **73**, 3195 (1994).
- [158] L. Kofman, A. Linde, and A. A. Starobinsky, Phys. Rev. Lett. **76**, 1011 (1996).
- [159] L. Kofman, A. Linde, and A. A. Starobinsky, Phys. Rev. D **56**, 3258 (1997).
- [160] C. J. Hogan, AIP Conf. Proc. **456**, 79 (1998).
- [161] M. Gasperini and G. Veneziano, Astropart. Phys. **1**, 317 (1993).
- [162] A. Kosowsky, M. S. Turner, and R. Watkins, Phys. Rev. Lett. **69**, 14 (1992).
- [163] C. Wu, V. Mandic, and T. Regimbau, Phys. Rev. D **87**, 042002 (2013).
- [164] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007).

- [165] E. E. Flanagan, *Phys. Rev. D* **48**, 2389 (1993).
- [166] V. Mandic, E. Thrane, S. Giampanis, and T. Regimbau, *Phys. Rev. Lett.* **109**, 171102 (2012).
- [167] J. Skilling, *AIP Conf. Proc.* **735**, 395 (2004).
- [168] J. Skilling, *Bayesian Anal.* **1**, 833 (2006).
- [169] D. S. Sivia with J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, 2006).
- [170] F. Feroz and M. P. Hobson, *Mon. Not. R. Astron. Soc.* **384**, 449 (2008).
- [171] F. Feroz, J. R. Gair, M. P. Hobson, and E. K. Porter, *Class. Quantum Grav.* **26**, 215003 (2009).
- [172] F. Feroz, M. P. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [173] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettit, arXiv: 1306.2144 (2014).
- [174] A. Lewis and S. Bridle, *Phys. Rev. D* **66**, 103511 (2002).
- [175] A. M. Hopkins, *Astrophys. J.* **615**, 209 (2004).
- [176] A. M. Hopkins and J. F. Beacom, *Astrophys. J.* **651**, 142 (2006).
- [177] P. S. Behroozi, R. H. Wechsler, and C. Conroy, *Astrophys. J.* **770**, 57 (2013).
- [178] R. J. Bouwens, G. D. Illingworth, M. Franx, and H. Ford, *Astrophys. J.* **670**, 928 (2007).
- [179] R. J. Bouwens, G. D. Illingworth, M. Franx, and H. Ford, *Astrophys. J.* **686**, 230 (2008).
- [180] R. J. Bouwens, G. D. Illingworth, P. A. Oesch, *et al.*, *Astrophys. J.* **737**, 90 (2011).
- [181] R. J. Bouwens, L. Bradley, A. Zitrin, *et al.*, *Astrophys. J.* **795**, 126 (2014).
- [182] R. J. Bouwens, G. D. Illingworth, P. A. Oesch, *et al.*, *Astrophys. J.* **803**, 34 (2015).

- [183] P. A. Oesch, R. J. Bouwens, G. D. Illingworth, *et al.*, *Astrophys. J.* **759**, 135 (2012).
- [184] P. A. Oesch, R. J. Bouwens, G. D. Illingworth, *et al.*, *Astrophys. J.* **773**, 75 (2013).
- [185] P. A. Oesch, R. J. Bouwens, G. D. Illingworth, *et al.*, *Astrophys. J.* **786**, 108 (2014).
- [186] P. A. Oesch, R. J. Bouwens, G. D. Illingworth, *et al.*, *Astrophys. J.* **808**, 104 (2015).
- [187] P. Madau, L. Pozzetti, and M. Dickinson, *Astrophys. J.* **498**, 106 (1998).
- [188] R. C. Kennicutt, Jr., *Annu. Rev. Astron. Astr.* **36**, 189 (1998).
- [189] T. Totani, *Astrophys. J.* **486**, L71 (1997).
- [190] S. Mao and H. J. Mo, *Astron. Astrophys.* **339**, L1 (1998).
- [191] R. A. M. J. Wijers, J. S. Bloom, J. S. Bagla, and P. Natarajan, *Mon. Not. R. Astron. Soc.* **294**, L13 (1998).
- [192] C. Porciani and P. Madau, *Astrophys. J.* **548**, 522 (2001).
- [193] V. Bromm and A. Loeb, **575**, 111 (2002).
- [194] F. Daigne, E. M. Rossi, and R. Mochkovitch, *Mon. Not. R. Astron. Soc.* **372**, 1024 (2006).
- [195] R. Chary, E. Berger, and L. Cowie, *Astrophys. J.* **672**, 272 (2007).
- [196] E. E. O. Ishida, R. S. de Souza, and A. Ferrara, *Mon. Not. R. Astron. Soc.* **418**, 500 (2011).
- [197] M. D. Kistler, J. Yüksel, J. F. Beacom, and K. Z. Stanek, *Astrophys. J.* **673**, L119 (2008).
- [198] M. D. Kistler, J. Yüksel, J. F. Beacom, A. M. Hopkins, and J. S. B. Wyithe, *Astrophys. J.* **705**, L104 (2009).

- [199] M. D. Kistler, J. Yüksel, and A. M. Hopkins, arXiv: 1305.1630 (2013).
- [200] M. Trenti, R. Perna, and S. Tacchella, *Astrophys. J.* **773**, L22 (2013).
- [201] P. S. Behroozi and J. Silk, *Astrophys. J.* **799**, 32 (2014).
- [202] V. Springel and L. Hernquist, *Mon. Not. R. Astron. Soc.* **339**, 312 (2003).
- [203] F. Daigne, K. A. Olive, E. Vangioni-Flam, J. Silk, and J. Audouze, *Astrophys. J.* **617**, 693 (2004).
- [204] F. Daigne, K. A. Olive, J. Silk, F. Stoehr, and E. Vangioni, *Astrophys. J.* **647**, 773 (2006).
- [205] E. Rollinde, E. Vangioni, D. Maurin, K. A. Olive, F. Daigne, J. Silk, and F. H. Vincent, *Mon. Not. R. Astron. Soc.* **398**, 1782 (2009).
- [206] G. Hinshaw *et al.*, *Astrophys. J. Suppl. Ser.* **208**, 19 (2013).
- [207] T. Suda *et al.*, *Publ. Astron. Soc. Jpn.* **60**, 1159 (2008).
- [208] T. Suda, S. Yamada, Y. Katsuta, Y. Komiya, C. Ishizuka, W. Aoki, and M. Y. Fujimoto, *Mon. Not. R. Astron. Soc.* **412**, 843 (2011).
- [209] J. A. Nelder and R. Mead, *Comput. J.* **7**, 308 (1965).
- [210] K. Crocker, V. Mandic, T. Regimbau, K. Belczynski, W. Gladysz, K. A. Olive, T. Prestegard, and E. Vangioni, *Phys. Rev. D* **92**, 063005 (2015).
- [211] I. Dvorkin, E. Vangioni, J. Silk, J.-P. Uzan, and K. A. Olive, arXiv: 1604.04288 (2016).
- [212] H. Dimmelmeier, C. D. Ott, A. Marek, and H.-T. Janka, *Phys. Rev. D* **78**, 064056 (2008).
- [213] J. W. Murphy, C. D. Ott, and A. Burrows, *Astrophys. J.* **707**, 1173 (2009).
- [214] A. Marek, H.-T. Janka, and E. Müller, *Astron. Astrophys.* **496**, 475 (2009).
- [215] K. Kotake, W. Iwakami, N. Ohnishi, and S. Yamada, *Astrophys. J.* **697**, L133 (2009).

- [216] C. D. Ott, *Class. Quantum Grav.* **26**, 204015 (2009).
- [217] A. Burrows and J. Hayes, *Phys. Rev. Lett.* **76**, 352 (1996).
- [218] C. L. Fryer, D. E. Holz, and S. A. Hughes, *Astrophys. J.* **609**, 288 (2004).
- [219] M. Shibata, Y. T. Liu, S. L. Shapiro, and B. C. Stephens, *Phys. Rev. D* **74**, 104026 (2006).
- [220] T. Takiwaki and K. Kotake, *Astrophys. J.* **743**, 30 (2011).
- [221] M. Obergaulinger, M. A. Aloy, H. Dimmelmeier, and E. Müller, *Astron. Astrophys.* **457**, 209 (2006).
- [222] C. D. Ott *et al.*, *Phys. Rev. Lett.* **106**, 161103 (2011).
- [223] C. D. Ott *et al.*, *Astrophys. J.* **768**, 115 (2013).
- [224] B. Müller, H.-T. Janka, and A. Marek, *Astrophys. J.* **766**, 43 (2013).
- [225] K. N. Yakunin *et al.*, *Phys. Rev. D* **92**, 084040 (2015).
- [226] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni, *Phys. Rev. D* **73**, 104024 (2006).
- [227] V. Ferrari, S. Matarrese, and R. Schneider, *Mon. Not. R. Astron. Soc.* **303**, 247 (1999).
- [228] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Phys. Rev. D* **61**, 124015 (2000).
- [229] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Phys. Rev. D* **330**, 651 (2002).
- [230] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Class. Quantum Grav.* **19**, 1335 (2002).
- [231] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Mon. Not. R. Astron. Soc.* **348**, 1373 (2004).

- [232] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Class. Quantum Grav.* **21**, 545 (2004).
- [233] A. Buonanno, G. Sigl, G. Raffelt, H.-T. Janka, and E. Müller, *Phys. Rev. D* **72**, 084001 (2005).
- [234] D. Coward, R. Burman, and D. Blair, *Mon. Not. R. Astron. Soc.* **329**, 411 (2002).
- [235] X.-J. Zhu, E. Howell, and D. Blair, *Mon. Not. R. Astron. Soc.* **409**, L132 (2010).
- [236] S. Marassi, R. Schneider, and V. Ferrari, *Mon. Not. R. Astron. Soc.* **398**, 293 (2009).
- [237] F. Pacucci, A. Ferrara, and S. Marassi, *Mon. Not. R. Astron. Soc.* **449**, 1076 (2015).
- [238] E. S. Phinney, arXiv: astro-ph/0108028 (2001).
- [239] K. Crocker *et al.*, (2016), in preparation.
- [240] C. D. Ott, *stellarcollapse.org*, <https://stellarcollapse.org> (2016).
- [241] E. Abdikamalov, S. Gossan, A. M. DeMaio, and C. D. Ott, *Phys. Rev. D* **90**, 044001 (2014).
- [242] B. Müller, H.-T. Janka, and H. Dimmelmeier, *Astrophys. J. Suppl. S.* **189**, 104 (2010).
- [243] R. Buras, M. Rampp, H.-T. Janka, and K. Kifonidis, *Astron. Astrophys.* **447**, 1049 (2006).
- [244] M. Rampp and H.-T. Janka, *Astron. Astrophys.* **396**, 361 (2002).
- [245] T. Kuroda, T. Takiwaki, and K. Kotake, *Phys. Rev. D* **89**, 044011 (2014).
- [246] J. Abadie *et al.*, *Phys. Rev. D* **85**, 082002 (2012).
- [247] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **87**, 022002 (2013).

- [248] S. E. Woosley and J. S. Bloom, *Annu. Rev. Astron. Astr.* **44**, 507 (2006).
- [249] C. D. Ott, *Class. Quantum Grav.* **26**, 063001 (2009).
- [250] A. Burrows and J. M. Lattimer, *Astrophys. J.* **307**, 178 (1986).
- [251] W. Zhang, S. E. Woosley, and A. Heger, *Astrophys. J.* **679**, 639 (2008).
- [252] K. Sumiyoshi, C. Ishizuka, A. Ohnishi, S. Yamada, and H. Suzuki, *Astrophys. J.* **690**, L43 (2009).
- [253] L. Dessart, A. Burrows, E. Livne, and C. D. Ott, *Astrophys. J.* **645**, 534 (2006).
- [254] E. Müller, M. Rampp, R. Buras, H.-T. Janka, and D. H. Shoemaker, *Astrophys. J.* **603**, 221 (2004).
- [255] W. Keil, H.-T. Janka, and E. Müller, *Astrophys. J.* **473**, L111 (1996).
- [256] J. A. Miralles, J. A. Pons, and V. A. Urpin, *Astrophys. J.* **543**, 1001 (2000).
- [257] J. A. Miralles, J. A. Pons, and V. A. Urpin, *Astron. Astrophys.* **420**, 245 (2004).
- [258] C. D. Ott, A. Burrows, T. A. Thompson, E. Livne, and R. Walder, *Astrophys. J. Suppl. Ser.* **164**, 130 (2006).
- [259] S. Akiyama, J. C. Wheeler, D. L. Meier, and I. Lichtenstadt, *Astrophys. J.* **584**, 954 (2003).
- [260] A. L. Watts, N. Andersson, and D. I. Jones, *Astrophys. J.* **618**, L37 (2005).
- [261] S. Scheidegger, R. Käppeli, S. C. Whitehouse, T. Fischer, and M. Liebendörfer, *Astron. Astrophys.* **514**, A51 (2010).
- [262] A. Corsi and P. Mészáros, *Class. Quantum Grav.* **26**, 204016 (2009).
- [263] D. Lai and S. L. Shapiro, *Astrophys. J.* **442**, 259 (1995).
- [264] A. L. Piro and E. Pfahl, *Astrophys. J.* **658**, 1173 (2007).
- [265] M. H. P. M. van Putten, *Astrophys. J.* **575**, L71 (2002).

- [266] M. H. P. M. van Putten, *Phys. Rev. Lett.* **87**, 091101 (2001).
- [267] M. H. P. M. van Putten, *Astrophys. J.* **684**, L91 (2008).
- [268] A. Krawczyk, A. G. Lyne, J. A. Gil, and B. C. Joshi, *Mon. Not. R. Astron. Soc.* **340**, 1087 (2003).
- [269] P. D. Morley and R. García-Pelayo, *Europhys. Lett.* **23**, 185 (1993).
- [270] P. M. McCulloch, P. A. Hamilton, D. McConnell, and E. A. King, *Nature (London)* **346**, 822 (1990).
- [271] N. Andersson, G. L. Comer, and D. Langlois, *Phys. Rev. D* **66**, 104002 (2002).
- [272] K. D. Kokkotas, E. Gaertig, and A. Colaiuda, *J. Phys. Conf. Ser.* **222**, 012031 (2010).
- [273] B. Owen and L. Lindblom, *Class. Quantum Grav.* **19**, 1247 (2002).
- [274] B. Vaishnav, I. Hinder, D. Shoemaker, and F. Herrmann, *Class. Quantum Grav.* **26**, 204008 (2009).
- [275] P. Csizmadia, G. Debreczeni, I. Rácz, and M. Vasúth, *Class. Quantum Grav.* **29**, 245002 (2012).
- [276] M. Coughlin, N. Christensen, J. Gair, S. Kandhasamy, and E. Thrane, *Class. Quantum Grav.* **31**, 165012 (2014).
- [277] L. Santamaría and C. D. Ott, LIGO DCC #T1100093 (2011).
- [278] S. R. Deans, *The Radon transform and some of its applications*, 2nd ed. (Dover Publications, 2007).
- [279] P. Raffai, Z. Frei, Z. Márka, and S. Márka, *Class. Quantum Grav.* **24**, S457 (2007).
- [280] R. Khan and S. Chatterji, *Class. Quantum Grav.* **26**, 155009 (2009).
- [281] L. Blackburn *et al.*, *Class. Quantum Grav.* **25**, 184004 (2008).
- [282] J. R. Smith *et al.*, *Class. Quantum Grav.* **28**, 235005 (2011).

- [283] P. Ajith *et al.*, *Class. Quantum Grav.* **23**, 5825 (2006).
- [284] M. W. Coughlin for the LIGO Scientific Collaboration and the Virgo Collaboration, *Class. Quantum Grav.* **28**, 235008 (2011).
- [285] T. Isogai for the LIGO Scientific Collaboration and the Virgo Collaboration, *J. Phys. Conf. Ser.* **243**, 012005 (2010).
- [286] T. Ballinger for the LIGO Scientific Collaboration and the Virgo Collaboration, *Class. Quantum Grav.* **26**, 204003 (2009).
- [287] N. Christensen for the LIGO Scientific Collaboration and the Virgo Collaboration, *Class. Quantum Grav.* **27**, 194010 (2010).
- [288] J. Slutsky *et al.*, *Class. Quantum Grav.* **27**, 165023 (2010).
- [289] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Class. Quantum Grav.* **29**, 155002 (2012).
- [290] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Class. Quantum Grav.* **32**, 115012 (2015).
- [291] T. Prestegard, E. Thrane, *et al.*, *Class. Quantum Grav.* **29**, 095018 (2012).
- [292] T. Prestegard and E. Thrane, LIGO DCC #L1200204 (2012).
- [293] M. Wąs *et al.*, *Class. Quantum Grav.* **27**, 194014 (2010).
- [294] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **81**, 102001 (2012).
- [295] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **112**, 131101 (2014).
- [296] L. Baiotti, I. Hawke, and L. Rezzolla, *Class. Quantum Grav.* **24**, S187 (2007).
- [297] T. Damour and A. Vilenkin, *Phys. Rev. D* **64**, 064008 (2001).
- [298] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. C. Miller, and M. Walczak, *Astrophys. J.* **789**, 120 (2014).

- [299] M. Shibata and K. Taniguchi, *Living Rev. Relativ.* **14**, 6 (2011).
- [300] Z. B. Etienne, J. A. Faber, Y. T. Liu, S. L. Shapiro, K. Taniguchi, and T. W. Baumgarte, *Phys. Rev. D* **77**, 084002 (2008).
- [301] S. Mereghetti, *Astron. Astrophys. Rev.* **15**, 225 (2008).
- [302] N. Andersson and G. L. Comer, *Phys. Rev. Lett.* **87**, 241101 (2001).
- [303] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), LIGO DCC #T0900499 (2010).
- [304] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), arXiv: 1203.2674 (2012).
- [305] R. Biswas *et al.*, *Class. Quantum Grav.* **26**, 175009 (2009).
- [306] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **83**, 122005 (2011).
- [307] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **85**, 102004 (2012).
- [308] S. Klimentenko, C. Pankow, and G. Vedovato, LIGO DCC #T1300869 (2014).
- [309] J. Abadie *et al.* (LIGO Scientific Collaboration), *Nucl. Instrum. Meth. A* **624**, 223 (2010).
- [310] I. Bartos *et al.*, LIGO DCC #T1100071 (2011).
- [311] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **91**, 022003 (2015).
- [312] E. Thrane, V. Mandic, and N. Christensen, *Phys. Rev. D* **91**, 104021 (2015).
- [313] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Living Rev. Relativ.* **19**, 1 (2016).
- [314] E. Thrane and M. Coughlin, *Phys. Rev. D* **88**, 083010 (2013).

- [315] E. Thrane and M. Coughlin, Phys. Rev. D **89**, 063012 (2014).
- [316] M. Coughlin, E. Thrane, and N. Christensen, Phys. Rev. D **90**, 083005 (2014).
- [317] M. Coughlin, P. Meyers, E. Thrane, J. Luo, and N. Christensen, Phys. Rev. D **91**, 063004 (2015).
- [318] M. Coughlin, P. Meyers, S. Kandhasamy, E. Thrane, and N. Christensen, Phys. Rev. D **92**, 043007 (2015).
- [319] E. Thrane and M. Coughlin, Phys. Rev. Lett. **115**, 181102 (2015).
- [320] J. Harms, Living Rev. Relativ. **18**, 3 (2015).
- [321] H. Doyle, *Seismology*, 1st ed. (Wiley, 1995).
- [322] F. Matichard *et al.*, Class. Quantum Grav. **32**, 185003 (2015).
- [323] S. A. Hughes and K. S. Thorne, Phys. Rev. D **58**, 122002 (1998).
- [324] M. Coughlin, J. Harms, N. Christensen, V. Dergachev, R. DeSalvo, S. Kandhasamy, and V. Mandic, Class. Quantum Grav. **31**, 215003 (2014).
- [325] J. Harms *et al.*, Class. Quantum Grav. **27**, 225011 (2010).
- [326] Kinematics, *Q330*, <http://www.kinematics.com/p-78-Q330.aspx> (2015).
- [327] iXBlue, *Octans*, <https://www.ixblue.com/products/octans> (2016).
- [328] Boulder Real Time Technologies, Inc., *Boulder Real Time Technologies*, <https://www.brtd.com> (2016).
- [329] The LIGO Scientific Collaboration and the Virgo Collaboration, LIGO DCC #T970130 (2009).
- [330] Liteway, Inc., *LuxLink Fiber Optic Communications Systems*, <http://www.luxlink.com> (2015).
- [331] B & B Electronics, *TTL Converters*, <http://www.bb-elec.com/Products/Serial-Connectivity/Serial-Converters/TTL-Converters.aspx> (2016).

- [332] J. Peterson, U.S. Department of Interior Geological Survey Report No. OFR 93-322, 1993 (unpublished).
- [333] V. Mandic, E. Thrane, and V. Tsai, Directional analysis of seismic data, 2015, internal document.
- [334] W. M. Roggenthen and C. D. Kock, American Rock Mechanics Association Report No. ARMA-2013-493, 2013 (unpublished).
- [335] F.-C. Lin, V. C. Tsai, and B. Schmandt, *Geophys. J. Int.* **198**, 656 (2014).
- [336] M. M. Haney and V. C. Tsai, *Geophysics* **80**, 167 (2015).
- [337] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016).