

**GRAVITATIONAL WAVE ASTROPHYSICS:
INSTRUMENTATION, DETECTOR CHARACTERIZATION,
AND A SEARCH FOR GRAVITATIONAL
SIGNALS FROM GAMMA-RAY BURSTS**

A Dissertation Presented

by

DANIEL HOAK

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2015

Physics

© Copyright by Daniel Hoak 2015

All Rights Reserved

**GRAVITATIONAL WAVE ASTROPHYSICS:
INSTRUMENTATION, DETECTOR CHARACTERIZATION,
AND A SEARCH FOR GRAVITATIONAL
SIGNALS FROM GAMMA-RAY BURSTS**

A Dissertation Presented

by

DANIEL HOAK

Approved as to style and content by:

Laura Cadonati, Chair

Jennie H. Traschen, Member

Stephane Y. Willocq, Member

Min S. Yun, Member

Rory Miskimen, Department Chair
Physics

ACKNOWLEDGEMENTS

*Hereafter, in a better world than this,
I shall desire more love and knowledge of you.*

Science is not done in a (metaphorical) vacuum, and this thesis is no exception. I have been incredibly fortunate to work with a large cast in my time at UMass and within LIGO. The many individuals I mention here are grouped by our immediate connection, but their influence and friendship cannot be described with brief categories. While the ideas and results presented here are due to many, I have tried to make the derivations, explanations, and measurements reflect my own understanding. Any mistakes are, of course, mine alone.

I owe my start in LIGO to Mike Zucker and Joe Giaime, who took me on as an operator shortly before the S5 science run. My early education in gravitational waves came through watching an exceptional team of scientists work on the initial LIGO detector, especially Valera Frolov, Brian O'Reilly, Rana Adhikari, Matt Evans, Stefan Ballmer, and Sam Waldman. The grad students I knew in that time, especially Kate Dooley, Anamaria Effler, Jeff Kissel, Adam Mullavey, Nic Smith and Rob Ward, were my inspiration to make the jump back into academia.

The most intense battles of grad school were fought alongside my fellow students at UMass: Ben Gamari, Jaime Hopkins, Sereres Johnston, Jess McIver, and Lee Walsh. I would not have succeeded without their spirit and skill. I am grateful to Jane, Ingrid, Mary Ann and the rest of the UMass administration for always being better organized than I was.

I am immensely proud to be a member of the UMass LIGO team, especially Satyanarayan Mohapatra, Sebastian Fischetti, James Clark, Jess McIver, Alex Lombardi, Joshua Kerrigan and James McInerney. Satya, the best-dressed man in LIGO, deserves a shout-out for blazing the path at UMass and giving me my first lessons in Python. James Clark was the postdoc who helped me navigate data analysis and amplitude spectral densities for the first time. My grad school sibling Jess McIver has paralleled my career in coursework and research, and whose work informed much of my own.

The data analysis component of this thesis was only possible thanks to the support and advice of LIGO-Virgo GRB enthusiasts and the Burst group. I especially wish to thank the developers of X-pipeline - Patrick Sutton, Michal Was, Gareth Jones, and others - for giving me the opportunity to conduct a search, and to learn coherent data analysis on a really kickass pipeline. Tom Adams was my partner in the GEO-GRBs search, and our patient and careful reviewers were Jonathan Leong, Nicolas Leroy, and Florent Robinet. The results presented here would not have been possible without them.

Detector characterization has a special place in my heart for its challenge, importance, and room for creativity. This is as much due to the subject as it is due to the tone set by the chairs of the DetChar group: Keith Riles, Gabriela Gonzalez, and Josh Smith. Many individuals in DetChar have made essential contributions to code infrastructure and my own understanding of the field, without whom my work would have been impossible; thanks to Duncan MacLeod, Laura Nutall, and Andrew Lundgren, among many others.

I have been fortunate to make two extended trips to the LIGO Hanford Observatory during grad school. These visits would not have been possible without the support and good wishes of Dave Reitze, Peter Fritschel, and David Shoemaker. At LHO I have worked and learned from the best team in the gravitational wave business: Daniel Sigg, Keita Kawabe, Mike Landry, Robert Schofield, Dick Gustafson, Rick Savage, and the many operators and staff. Visitors during my time at the site have pushed me to work harder and better than I realized I could, especially John Miller, Alberto Stochino, Lisa Barsotti, Gabriele Vajente, Rich Abbott, Aidan Brooks, and Jamie Rollins. Many thanks to Koji Arai for entrusting his exquisitely designed and built OMC to my care, for his patience with my beginner's questions, and for teaching me the meaning of 'orz'.

An exceptional team of young scientists brought the H1 instrument from hodgepodge to wholeness. The late nights spent with this crew have been exhausting and rewarding, an experience that is unlikely to be matched in this life. Many thanks to Sheila Dwyer, Evan Hall, Kiwamu Izumi, Sudarshan Karki, Jeff Kissel, Elli King, and Alexa Staley for the beers, laughs, bugs, features, locklosses, loop instabilities, partially loaded filter modules, sunsets, sunrises, DAQ restarts, Guardian reloads, undamped resonances, memorable aLOGs, and many Mpc of fun.

No LIGO thesis that includes work at the Hanford observatory would be complete without acknowledging Daniel Sigg's whiskey cabinet. But neither would it be complete without mentioning his humor, encouragement, support, patience and wisdom. Thank you, Daniel.

It is no understatement to say that my graduate education would not have happened without my advisor, Laura Cadonati. The group that Laura built at UMass was the reason I pursued a PhD with LIGO, and her support and patience through six full years of graduate school has been energizing and humbling. Grazie mille, Laura!

To my friends outside of physics, who have made it all worthwhile: here's hoping that life after grad school can become Shakespeare's "better world than this."

To my family: Mom and Jim, Dad and Linda, Madeline, and Ahma: you have kept me grounded, supported, educated, fed, clothed, and lifted, and none of this would have been possible without you. Thank you.

Finally, to Jessiaha, who will always hit, and whom I will always love: thank you. How about we chuck it all and move to Tuscany?

ABSTRACT

GRAVITATIONAL WAVE ASTROPHYSICS: INSTRUMENTATION, DETECTOR CHARACTERIZATION, AND A SEARCH FOR GRAVITATIONAL SIGNALS FROM GAMMA-RAY BURSTS

SEPTEMBER 2015

DANIEL HOAK

B.A., BOSTON UNIVERSITY

M.S., UNIVERSITY OF MASSACHUSETTS, AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Laura Cadonati

In the coming years, the second generation of interferometric gravitational wave detectors are widely expected to observe the gravitational radiation emitted by compact, energetic events in the nearby universe. The field of gravitational wave astrophysics has grown into a large international endeavor with a global network of kilometer-scale observatories. The work presented in this thesis spans the field, from optical metrology, to instrument commissioning, to detector characterization and data analysis. The principal results are a method for the precise characterization of optical cavities, the commissioning of the advanced LIGO Output Mode Cleaner at the Hanford observatory, and a search for gravitational waves associated with gamma-ray bursts.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iv
ABSTRACT	vi
LIST OF TABLES	ix
LIST OF FIGURES	x
 CHAPTER	
1. INTRODUCTION	1
2. GRAVITATIONAL WAVES	3
2.1 General Relativity	3
2.2 Generation of Gravitational Waves	8
2.3 Gravitational Waves from Binary Systems	11
2.4 E_{GW} and Strain Amplitude	14
2.5 Astrophysical Sources of Gravitational Waves	17
3. GRAVITATIONAL WAVE DETECTORS	19
3.1 Optical Gain of the Initial LIGO Detectors	21
3.2 Limits to Detection	25
3.3 Detector Antenna Patterns	27
3.4 Results from the Initial Detector Era and Current Astrophysical Limits	29
3.5 Can Interferometers Detect Metric Perturbations?	32
4. OPTICAL RESONATORS	33
4.1 Optical Cavities	33
4.2 Pound-Drever-Hall Locking	38
4.3 Precision Measurements of Cavity Length and Linewidth	42
5. ADVANCED LIGO	50
5.1 The Advanced LIGO Interferometers	50
5.2 Cavities & Control	52
5.3 The Output Mode Cleaner	55
5.4 OMC Cavity Length Noise	60
5.5 The OMC as a Spectrum Analyzer	64
5.6 Measurement of the Interferometer Contrast Defect	67

6. DETECTOR CHARACTERIZATION	72
6.1 Locking and Unlocking	72
6.2 Stationarity of the Noise & Effects on the Searches	74
6.3 Instrumental Correlations of Glitches	78
6.4 Glitches in aLIGO: OMC Scattering	82
7. A SEARCH FOR GRAVITATIONAL WAVES ASSOCIATED WITH GAMMA-RAY BURSTS	86
7.1 Gamma-ray Burst Physics	87
7.2 GRB Satellite Missions	90
7.3 Gravitational Waves from GRBs and Previous Searches	90
7.4 Formalism of Coherent Searches	92
7.5 Sky Localization Uncertainty	99
7.6 A Search With Data From GEO 600	102
7.7 Sky Localization for <i>Fermi</i> -GBM Events	109
8. CONCLUSION	113
 APPENDICES	
A. NOTES ON AMPLITUDE SPECTRAL DENSITY	116
B. DETECTION STATISTIC FORMALISM	124
C. COHERENT CONSISTENCY CHECKS	131
D. EPIGRAPHS	135
 BIBLIOGRAPHY	 136

LIST OF TABLES

Table	Page
3.1 Sensing parameters of the initial LIGO detectors.	24
4.1 Best fit parameters and 67% confidence level statistical uncertainties for the Input Mode cleaner data. The results presented here are for an audio modulation frequency of 1 kHz and an RF modulation of 9.1 MHz.	48
4.2 Parameters of the H1 aLIGO Input Mode Cleaner and Y-arm optical cavities, measured using the double-modulation technique.	49
5.1 Modulation depths at H1 for the RF sidebands, measured using the OMC. For the calculation of the contrast defect in the next section, we assume uncertainties in $\Gamma_{9.1}$ and $\Gamma_{45.5}$ of 0.005.	67
5.2 Peak heights for OMC cavity sweeps in full lock, with DARM offset off (on). Heights are given in milliamps of DCPD photocurrent. Peaks without a measurement were either too small to fit, or degenerate with another mode. We conservatively estimate the uncertainty of the carrier sums to be 0.05 mA.	69

LIST OF FIGURES

Figure	Page
2.1 The effect of the two polarizations of gravitational waves on a ring of test masses. The direction of propagation of the wave is perpendicular to the page.	8
2.2 Cumulative excess phase in the orbit of PSR B1913+16, compared to the prediction for damping from gravitational radiation. Negative values represent a progressively earlier periastron, indicative of a decreasing orbital period. The rate of orbital period decay is $-2.423(1) \times 10^{-12}$ seconds/second. Figure from [143].	13
3.1 Optical layout of a simple Michelson interferometer. The detection photodiode is placed at the antisymmetric port of the interferometer.	20
3.2 Shot noise limit from Eq. 3.13 for the initial LIGO detectors, and a representative sensitivity curve for the H1 detector during the S6 science run.	26
3.3 F_+ antenna pattern. The z -axis is vertical.	28
3.4 F_\times antenna pattern.	28
3.5 Limits on the rate of compact binary coalescences set by the most recent LIGO-Virgo science runs. The light grey regions are the limits set by the S5 science run, and the dark grey regions are the limits set by the combined results of the S5 and S6 science runs. The blue regions are the range of predicted rates based on population studies, as detailed in [12]. Figure from [18].	29
3.6 Upper limits on h_{rss} from continuous wave sources. The blue points are limits derived from the first year of data from the S5 science run, and red points are derived from the full two-year dataset. Figure from [2].	30
3.7 Upper limits on the rate of generic gravitational wave sources emitting $M_\odot c^2$ of energy isotropically in short-duration, narrowband gravitational wave bursts. The limits degrade as the frequency of the hypothesized source increases due to the decreasing sensitivity of the detectors and from the factor of f_0^2 in Eq. 2.46. Figure from [15].	31
3.8 Limits on stochastic gravitational wave energy as a fraction of ρ_{critical} , the energy density required to close the universe: $\Omega_{GW} = \frac{1}{\rho_{\text{critical}}} \frac{d\rho_{GW}}{d \ln f}$. Figure from [9].	31
4.1 Fields around an optical cavity. Figure by J. Adamopoulos.	34
4.2 Power and phase of the transmitted and reflected beams from an optical cavity. The cavity finesse is 14, and the detuning phase $\phi_f = 4\pi fL/c = 2\pi f/f_{FSR}$	36
4.3 Different optical cavity topologies encountered in this thesis. A two-mirror cavity (a), such as the LIGO arm cavities; a three-mirror triangular cavity (b), such as the Input Mode Cleaner; and (c), a four-mirror ‘bowtie’ cavity, similar to the Output Mode Cleaner.	38
4.4 Demodulated signals for PDH in reflection. The sideband frequency is 15% of the cavity FSR. The cavity finesse is 156, and the detuning phase $\phi_f = 2\pi f/f_{FSR}$	41

4.5	Zoom of the linear region of the PDH error signal.	41
4.6	I- and Q-phase signals on transmission, as a function of the RF sideband detuning phase $\phi_f = 2\pi f_{RF}/f_{FSR}$, with the carrier frequency held slightly off-resonance. The cavity finesse is 156. Note the zero crossings in the Q-phase signal.	44
4.7	Double demodulation: a diagram of the carrier, RF sidebands and audio sidebands for the cavity length and linewidth measurement. Figure by A. Staley, from [125].	45
4.8	Experimental setup for the double-modulation measurement. Figure by K. Izumi, from [125].	46
4.9	Experimental results for the scan of the 16 m input mode cleaner; the x-axis shows the frequency of the RF modulation. The audio modulation frequency for this dataset was 1 kHz. The best-fit parameters are given in Table 4.2. The feature around 9.095 MHz is most likely due to a mechanical resonance in the EOM.	47
4.10	Zoom of the zero-crossing at the FSR frequency, for the IMC measurement at 9.1 MHz and 1 kHz. The best-fit FSR value has been subtracted from the x-coordinate for clarity. Note the precise null value for all four I,Q signals at the FSR.	47
5.1	Schematic of the advanced LIGO interferometers. Figure from [7]. The Fabry-Perot arm cavities are formed between the <i>end test mass</i> (ETM) and <i>inner test mass</i> (ITM) mirrors. Power-recycling and signal-recycling cavities are formed between the ITMs and the <i>power recycling mirror</i> (PRM) and the <i>signal recycling mirror</i> (SRM), respectively. Additional large mirrors are used to align and shape the beam. All of the detector components in this figure after the phase modulation (ϕ_m) are housed in large vacuum chambers and are isolated from ground motion by multi-stage suspensions.	51
5.2	Coupled optical cavity topologies used in the advanced LIGO detectors. Power recycling (a) forms a coupled cavity between the Power Recycling Mirror (PRM) and the Fabry-Perot arm cavities. Intensity and frequency fluctuations from the input beam are suppressed by the coupled cavity pole, with $\omega_{cc} \simeq 1$ Hz. Signal recycling (b) forms a resonant cavity between the Signal Recycling Mirror (SRM) and the arms. Acoustic-frequency fluctuations generated by gravitational wave interactions with the arms (the ‘signal’) are resonant in this cavity, and the response of the interferometer is broadened.	52
5.3	Photo of the author in the ETMY vacuum chamber, working on alignment of the green laser. I am kneeling, Keita Kawabe is standing.	53
5.4	Locking sequence of the advanced LIGO detectors. Figure from [126]. The top figure shows the power buildup of the 1064 nm light in the arms. In states I and II, the arms are locked using the 532 nm lasers. In states III and IV, the arms are moved off-resonance and the cavities of the DRMI are locked. In states V and VI the arms are brought back to resonance, this time with much larger power buildup due to the power recycling cavity. In state VII the detector is fully resonant.	54
5.5	Recent performance of the H1 advanced LIGO detector, compared to previous best sensitivity. At low frequencies, the improved seismic isolation and suspensions have reduced the control noise that couples to the gravitational wave channel. At 100 Hz the sensitivity improvement is due to the increased arm cavity finesse. At high frequencies the signal recycling cavity has broadened the response of the detector.	56

5.6	Sweep of the differential arm (DARM) offset with the full interferometer locked. Here, 2.3 W of power is incident on the input mode cleaner, and about 80 W is incident on the beamsplitter (due to the gain of the power recycling cavity). The power measured by the antisymmetric port DCPDs is a quadratic function of the DARM offset.	57
5.7	Design of the advanced LIGO Output Mode Cleaner, from [30].	57
5.8	Arrangement of optics in the advanced LIGO HAM6 vacuum chamber. The optical table is approximately two meters wide. The beam from the signal recycling mirror enters from the left. Pick-off beams for monitoring the state of the detector leave the vacuum chamber at the top and are collected on an in-air optics bench. The rectangles, mostly at the bottom of the figure, are 10 kg ballast masses to balance the seismic isolation platform. Figure from [35].	59
5.9	Photo of the author working in the HAM6 chamber in August, 2014. The OMC suspension cage is the tall structure on my left.	59
5.10	When an optical cavity is locked at the resonance, small cavity length fluctuations (δL) couple quadratically to the transmitted power (δP). At the half-resonance point, length noise maximally couples to the transmitted power. The same relationship is true for fluctuations in the alignment of the input beam. When the cavity is well-aligned, power fluctuations in the transmitted light from input beam jitter are suppressed.	60
5.11	Relative intensity noise observed for a half-fringe lock of the OMC. Above 100 Hz the signal is limited by input intensity noise.	61
5.12	Two measurements of OMC cavity length noise. Above 100 Hz the half-fringe method is limited by input intensity noise. Between 1 and 100 Hz the PDH method is limited by noise of the PZT actuators. There is agreement at low frequency, and also around 1 kHz, where resonances of the seismic platform disturb the cavity. The body modes of the OMC silica breadboard are visible at 10 kHz in the PDH measurement.	62
5.13	Estimate of DARM noise due to OMC length fluctuations (loop-suppressed), compared to the recent sensitivity of the H1 interferometer. Note the y-axis is in units of displacement sensitivity; divide by the arm length to get the equivalent strain sensitivity.	64
5.14	Scan of the OMC with a single-bounce beam off an ITM. The scan data are plotted in black and fits to the peaks are plotted as colored Lorentzian functions (purple for sidebands, blue for higher-order modes, green for an extended sample of higher-order modes). The gold trace is the sum of the fitted peaks. As the voltage on the PZT decreases, the cavity length decreases, and higher-order modes with increasing optical frequency are resonant.	65
5.15	Scan of the OMC with the full interferometer locked with zero DARM offset. Higher-order modes of the carrier and sidebands are labeled. Capital letters are used for the 45.5 MHz sidebands and lower-case letters are used for the 9.1 MHz sidebands.	68
5.16	Scan of the OMC with the full interferometer locked with the DARM offset on, $x_0 = 15.8$ pm. Note the height of the CR0 peaks compared to Fig. 5.15.	69
5.17	Fit to mode optical frequency, for OMC cavity sweep in full lock with DARM offset on. The optical frequency for each mode was determined using the sideband frequencies and the OMC higher-order mode spacing. The peaks were fit to Lorentzian functions of the PZT voltage, and the set of voltages were fit to the mode frequencies with a 4th-order polynomial.	71

6.1	Accounting of the sources of downtime for the L1 detector during the two-year S5 science run. Above-normal ground motion, in the form of earthquakes, wind, trains, construction and logging activities near the site, and the microseismic peak due to ocean waves, was responsible for 100 days of lost science data. The H1 detector did somewhat better in S5, with 76.9% uptime. In the S6 science run, between July 2009 and October 2010, the instruments had slightly worse duty factor due to preparations for the advanced LIGO installation.	73
6.2	Duration of science segments collected by the LIGO interferometers in their sixth science run, July 2009 - October 2010. The L1 instrument recorded 3,276 segments, H1 recorded 2,395.	74
6.3	Example of a lock loss recovery for the H1 advanced LIGO detector. Figure by Duncan MacLeod.	75
6.4	Distribution of sensitivity per frequency bin for the H1 detector during the S6D epoch (6.28 million seconds of data, or about 72 days). The data were divided into 50% overlapping segments with 6-second duration. The FFTs of these segments were collected and used to calculate percentiles of the noise for each frequency bin. For 0.5% of the time, the noise was greater than the black trace.	76
6.5	Median-normalized sensitivity percentiles per frequency bin. The expected ratio for Rayleigh-distribution noise for each percentile is given by the dashed lines. There is a clear deviation from Rayleigh-distributed noise, especially at low frequency.	77
6.6	Cumulative histogram of time-slide events for the Coherent WaveBurst search in S6D [15]. The long tail of events, especially at low frequencies, extends to an astrophysically unlikely SNR, and statistically blinds the search to plausible signals. Vetoing the events based on correlations to instrumental and environmental noise improved the results, but at the expense of duty cycle. Typically around 10-20% of the data are vetoed. The error bars are \sqrt{N} . This plot is my own work, and should not be interpreted as a result of the CWB group.	77
6.7	Correlation between glitches and ETMX actuation current during highly elevated, narrowband microseismic noise on Oct 22 2009 at the H1 detector. The value of the coil current at the time of a glitch in the gravitational wave channel is given by the red squares. The grey traces are the coil current around the time of the glitch. The seismic noise on this day was sharply peaked around 0.125 Hz.	79
6.8	Location of Barkhausen-type glitches in a phase-space plot of the ETMX coil current and its time derivative. Note the correlation between glitch times and the phase and amplitude of the coil current. The data here are the same as in Fig. 6.7.	80
6.9	Glitch probability for two months of data in H1 S6B as a function of coil current and its time derivative. Nearly all of the time the interferometer state is located in the very center of the plot, where the glitch probability is around 10^{-3} . During times of very large, low-frequency seismic disturbances, the coil current amplitude is very large, and the interferometer state falls in points on the extremes of the plot. Barkhausen glitches tend to occur in the upper right and lower left quadrants.	80
6.10	Performance of a veto constructed by rejecting times when the coil current was in a probable configuration to generate Barkhausen noise. Statistics are presented per science segment.	81
6.11	Cutoff of Bessel functions of the first kind as the order n grows larger than Γ	84

6.12	Top panel: spectrogram of scattering noise in the calibrated gravitational wave channel. Bottom panel: velocity of the longitudinal degree of freedom of the OMC suspension in $\mu\text{m}/\text{sec}$. Scattering arches are correlated with peaks in the velocity of the OMC, indicating that the path length between a scattering source and the interferometer is being modulated by large motions of the OMC suspension. (The motion of the OMC is measured at the middle stage of the suspension, and the bottom stage, with the optical components, is moving faster.)	84
6.13	Photo of the author working to check beam alignment into the OMC black glass shroud, designed to enclose the output mode cleaner in a non-reflective covering to prevent scattered light. The black welder's glass can be seen around the OMC suspension cage in the upper left of the image.	85
7.1	Comparison of results from previous searches for gravitational wave associated with GRBs and the observed redshift distribution for GRB progenitors. On the left, we plot lower limits on the distance to 150 GRBs from a search for unmodeled gravitational wave signals, assuming $E_{GW} = 10^{-2} M_{\odot} c^2$ in a narrowband signal around 150 Hz. On the right are lower limits to CBC events associated with 26 short GRBs. The dashed curves are extrapolations for the advanced detectors, assuming a 10x improvement in sensitivity and a 5x increase in the number of GRBs analyzed. The red curves are the redshift distributions for <i>Swift</i> GRBs (left: all GRBs, right: short-hard GRBs). Figure from [17].	92
7.2	Vector space of detector signals for the 3-detector case. The green plane is the space spanned by the antenna response vectors. Figure from [127].	93
7.3	Time-frequency map for the 'loghbayesiancirc' detection statistic for a simulated gravitation wave event. The waveform is an NSBH inspiral. The color of each pixel indicates the value of the detection statistic. Adjacent time-frequency pixels that lie in the 99th percentile or larger are used to calculate the total event significance. For this event the sum of the significant pixels is 126.	96
7.4	Energy and incoherent energy in the circular energy plane for an example analysis with the H1-L1 detector network. The x- and y-axes have arbitrary units of non-normalized signal energy. The starred event is the signal shown in Fig. 7.3. The many hundreds of small crosses indicate background events. Background events and signal usually have roughly equal energy and incoherent energy in the non-null direction.	96
7.5	Null energy vs. null incoherent energy for an example analysis with the H1-L1 detector network, in units of non-normalized signal energy. The starred event is the signal shown in Fig. 7.3. The purple dashed line indicates the threshold for the coherent cut. Events lying below this line are inconsistent with a gravitational wave signal from the sky location of the GRB and are rejected by the search.	97
7.6	Foreground and background event rates for an example GRB analysis. The traces are cumulative histograms that count the number of events over a particular significance threshold. The significance of each event is the value of the detection statistic given in Eq. 7.10. Traces for the on-source (foreground) and off-source (background) event distributions are shown, before and after the coherent cuts are applied. The on-source has an outlier, with a significance of about 120. This is the NSBH signal shown in Fig. 7.3.	98

7.7	Event significance as a function of phase shift error between two detectors, for three sets of narrowband signals with different central frequencies. When the error in the time shift between detectors corresponds to a multiple of 180° , the coherent energy is essentially nulled, and the signal will not be detected by the search. This plot was generated using a search grid of a single point and scattering the simulated waveforms widely across the sky.	100
7.8	Example linear and circular grids for a search for gravitational wave signals up to 1792 Hz. The localization for the <i>Fermi</i> GBM event GRB 080906B is shown. The linear search grid contains 41 sky positions, arranged in the direction of the gradient of the time shift between the H2 and G1 detectors. The circular grid contains 1324 sky positions and would require several days to analyze on a massively parallel computing cluster.	102
7.9	Demonstrative noise spectra for the detectors used in The GEO-GRBs search. GEO's best sensitivity is around 500 Hz in the early epoch, and above 1 kHz in the late epoch when quantum-noise squeezing was applied to reduce the shot noise [14].	103
7.10	Distribution of p-values for the loudest event in each of the 129 GRBs analyzed in the search.	105
7.11	Illustration of possible p-value distributions using 500 pseudo-experiments, each with 129 GRBs. The results from the search, in red, fall within the distribution of results from the pseudo-experiments, in grey.	106
7.12	Distribution of upper limits on strain amplitude for gravitational wave signals from the 129 GRBs analyzed in the GEO-GRBs search. The simulated gravitational wave signals are $Q = 9$ sine-Gaussians with central frequencies 500 Hz (red) and 1 kHz (black). The upper limits have been marginalized over detector calibration uncertainties and GRB sky localization uncertainty. Figure from [11].	106
7.13	Distribution of lower limits on distance to gravitational wave signals from the GRBs in the search. A standard gravitational wave emission energy of $10^{-2} M_\odot$ is assumed, for sine-Gaussian waveforms with central frequencies 500 Hz (red) and 1 kHz (black). Figure from [11].	107
7.14	Upper limits on gravitational wave strain amplitude for sine-Gaussian waveforms at 500 Hz and 1 kHz, for <i>Swift</i> and <i>Fermi</i> events. The upper limits are largely the same; this indicates that the sky localization of the <i>Fermi</i> events (and the use of the linear search grid) did not affect the sensitivity of the search.	108
7.15	Ratio of upper limits on gravitational wave strain amplitude to the least sensitive detector, as a function of the ratio of detector sensitivities.	109
7.16	Sky location reconstruction with the linear grid. Figure from [11]. The true GRB location is marked by the star. The red dot is the sky localization from the <i>Fermi</i> GBM, and the black circle illustrates a typical 95% containment region around the GBM localization. A GRB search with X-PIPELINE could further localize a detected signal along the linear search grid, with a localization uncertainty illustrated by the grey band.	110
7.17	Accuracy of sky localization for various waveforms, for the H1-V1 detector network in S6 noise. The first four waveforms are $Q=9$ sine-Gaussians with the central frequencies given in the legend. High-frequency signals can be localized to roughly 1° in the direction parallel to the linear grid. Waveforms from CBCs are longer-duration and can be localized to less than 1°	111

8.1	Anticipated number of BNS detections for proposed advanced LIGO observing runs. The average detector sensitivity to BNS mergers is expressed in megaparsecs (x-axis). The expected number of detected signals is calculated using estimates of the rate density for BNS mergers, for observing runs of different sensitivity and duration. The sensitive volume is given by the cube of the value on the x-axis, and the duration of the observing run is given in the legend. Data points are the expected number of events for the most likely rate derived from population synthesis models and extrapolation from the known population of BNS systems. Error bars indicate the range of possible rates consistent with observation, which is quite large. The dashed line indicates the threshold at which the probability for at least one detection crosses 50%.	114
A.1	Distributions of ASD values for three frequency bins in L1 data from S6B. The deviation from the expected Rayleigh distribution is more dramatic for the low frequency data where glitches are more common.	120
A.2	Distribution of sensitivity per frequency bin for the L1 detector during the S6B epoch. The plot contains data from 952 science segments between Sept 21 2009 and Jan 8 2010, for a total of 3.42 million seconds of data, or about 40 days.	121
A.3	Median-normalized sensitivity percentiles per frequency bin, for L1 in S6B. The expected ratio for each percentile for Rayleigh-distributed noise is given by the dashed lines. There is a broadband deviation from Gaussian noise due to excessive seismic disturbances during this four-month epoch.	121
A.4	Median ASDs for the H1 detector for the epochs of the S6 science run. Note the steady improvement of the noise due to commissioning efforts.	122
A.5	Comparison of the 95th percentile noise for the H1 detector for the epochs of the S6 science run, to the median noise of the S6D epoch.	122
A.6	Median ASDs for the L1 detector for the epochs of the S6 science run.	123
A.7	Comparison of the 95th percentile noise for the L1 detector for the epochs of the S6 science run, to the median noise of the S6D epoch.	123

CHAPTER 1

INTRODUCTION

*A space whose every cubit
Seems to cry out*

There is currently a worldwide effort to observe gravitational waves from compact, energetic events in the nearby universe. These detections will mark the beginning of the field of gravitational wave astrophysics, and directly probe regions of strong-field gravity for the first time. The first generation of interferometric detectors collected data between 2005 and 2010, after which the main detector components were de-commissioned to make way for second-generation instruments. These advanced detectors are now preparing for their first data-taking operations.

The full promise of the advanced detectors will only be realized if certain persistent challenges are met. First, commissioning the detectors themselves is an ongoing research and development project, and reaching the final sensitivity goals will require significant effort in the coming years. Second, the rate of transient noise events in the detectors is a limiting factor for searches for short-duration unmodeled signals. Reducing the rate of these transients would substantially improve the sensitivity to short-duration signals. Finally, information from electromagnetic observations can be used to design searches with improved sensitivity to gravitational waves. These so-called externally triggered searches are some of the most promising avenues for gravitational wave detection.

The original research presented in this thesis touches on all of these topics. Briefly, the highlights are the following:

- I have worked with colleagues from the LIGO Hanford observatory to develop a method for the precise characterization of optical cavities. We have demonstrated this method using two of the optical cavities of the H1 instrument. The results have been published in *Optics Express* [125].
- Alongside a large number of talented scientists, I have worked to prepare the H1 instrument for the first aLIGO observing run. During the spring of 2012 I assisted with the installation of the Arm-Length Stabi-

lization system (ALS) which is used in the lock acquisition process. A paper describing the performance of ALS was published in *Classical and Quantum Gravity* [126]. Between May 2014 and July 2015 I worked at the LIGO Hanford observatory to commission the detector, with a focus on the operation of the Output Mode Cleaner (OMC). A paper describing the OMC is in preparation.

- Throughout my graduate studies I have worked on techniques to identify and mitigate the various noise couplings that generate an excess of loud transient signals in the detectors. In particular I have developed a detector characterization pipeline that correlates loud transient signals with pairs of instrumental parameters.
- Using data from the previous observing runs of the LIGO, Virgo, and GEO 600 detectors, I collaborated with data analysis colleagues to complete a search for gravitational waves associated with gamma-ray bursts (GRBs). The results of the search have been published in *Physical Review D* [11]. A novel component of this search was a study of the analysis pipeline's ability to localize a detected signal on the sky, within the uncertainties provided by gamma-ray satellite experiments.

This thesis is organized in the following way. In Chapter 2, I review the formalism of gravitational waves, derive the waveform emitted by a simple binary system, discuss the emission energy for transient gravitational wave sources, and present the current evidence and expected rates for a variety of sources. Chapter 3 describes the basic design of modern interferometric detectors, derives the optical response of the initial detectors, calculates the limit from a fundamental noise source, and briefly presents the results of the initial detector era. Chapter 4 introduces the formalism of optical cavities and Pound-Drever-Hall locking, and presents the methods and results of a technique for precise measurements of optical cavity properties. In Chapter 5 I describe the commissioning of the advanced detector at the LIGO Hanford observatory, with a focus on the Output Mode Cleaner. In Chapter 6 I discuss the data quality of the initial LIGO instruments, and present results from an analysis pipeline designed to identify the instrumental noise couplings that generate the background of transient signals. In Chapter 7 I describe the methods and results of a search for gravitational waves associated with GRBs at high frequencies (> 500 Hz) using data from GEO 600 and other detectors. Finally, in Chapter 8 I summarize the results of this thesis and discuss implications for the future of the field.

This thesis has been given LIGO document number P1500136. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. The work presented here was supported by the NSF under Grant No. 0955773. Any opinions, findings, conclusions or recommendations expressed in this thesis are, of course, my own, and do not reflect the views of the National Science Foundation.

CHAPTER 2

GRAVITATIONAL WAVES

*The heavens themselves, the planets, and this center
Observe degree, priority, and place*

Gravitational waves are one of the earliest predictions of the theory of general relativity. Despite numerous precision tests of GR in the past century [146], gravitational waves have never been directly detected. Their indirect detection, in the binary system PSR B1913+16 [143], has motivated a large international effort to directly detect gravitational waves and inaugurate the promising field of gravitational-wave astrophysics.

In this chapter, I derive gravitational waves as solutions to Einstein's equations in free space, utilizing the standard weak-field approximation. Considering the source term for general relativity, I show that gravitational waves are generated by accelerating, asymmetric mass distributions, and illustrate this with the example of a binary system. I discuss the energy released by plausible gravitational wave sources and the strain induced at the Earth from astrophysical sources. Finally I give a summary of likely rates for gravitational wave sources and the sensitivity required for their detection.

The derivations in this chapter follow the text by Carroll [42], while the discussion of the energy carried by gravitational waves follows a note by Sutton [129].

2.1 General Relativity

For a given point in a four-dimensional coordinate system x_μ , we define the infinitesimal length element ds as a function of the *metric tensor*, $g_{\mu\nu}$, at the point x :

$$ds^2 = g_{\mu\nu}(x)dx^\mu dx^\nu. \quad (2.1)$$

The force of gravity is manifested through the shape of the metric, under the assumption that freely-falling test particles will follow the shortest path through spacetime. These paths are referred to as *geodesics* of the metric,

and it can be shown that a unique geodesic exists for any particular point in spacetime in the presence of a particular metric. The metric for flat spacetime is referred to as the *Minkowski metric*, and is given by

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.2)$$

For small perturbations to flat spacetime, we typically write the metric tensor $g_{\mu\nu}$ as the Minkowski metric plus a small *metric perturbation* denoted by $h_{\mu\nu}$:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad |h_{\mu\nu}| \ll 1. \quad (2.3)$$

This form allows us to expand the equations of gravity in terms of the metric perturbation h , and ignore higher-order terms. This is referred to as the *weak-field approximation* and is valid as long as the observer is far from a gravitational source.

Einstein's equation for general relativity relates the *Einstein tensor* $G_{\mu\nu}$ to the stress-energy tensor $T_{\mu\nu}$ in the following way:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (2.4)$$

where G is Newton's constant, c is the speed of light, and Λ is the *cosmological constant*, a term with positive energy density that models the observational evidence for the accelerating expansion of the universe. Since Λ is important only on very large distance scales we will neglect it for the remainder of this thesis.

The Einstein tensor $G_{\mu\nu}$ is a function of the Ricci tensor $R_{\mu\nu}$ and the curvature scalar $R = g^{\mu\nu} R_{\mu\nu}$,

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R. \quad (2.5)$$

The Ricci tensor is a contracted form of the Riemann curvature tensor, which describes the parallel transport of a 4-vector around a closed path. To explicitly write the Ricci tensor in a given coordinate system we introduce the Christoffel symbols Γ_{bc}^a :

$$R_{\mu\nu} = R_{\mu\sigma\nu}^{\sigma} = \partial_{\sigma} \Gamma_{\nu\mu}^{\sigma} - \partial_{\nu} \Gamma_{\sigma\mu}^{\sigma} + \Gamma_{\sigma\alpha}^{\sigma} \Gamma_{\nu\mu}^{\alpha} - \Gamma_{\nu\alpha}^{\sigma} \Gamma_{\sigma\mu}^{\alpha}. \quad (2.6)$$

The Christoffel symbols are known as *affine connections*, and arise from the definition of the derivative on tangent spaces to smooth manifolds. They are defined using the local spacetime metric $g_{\mu\nu}$:

$$\Gamma_{\mu\nu}^{\sigma} = \frac{1}{2}g^{\sigma\lambda} (\partial_{\nu}g_{\nu\lambda} + \partial_{\nu}g_{\mu\lambda} - \partial_{\lambda}g_{\mu\nu}) . \quad (2.7)$$

The Christoffel symbols have one derivative of the metric, and the Ricci tensor and curvature scalar have two. Hence, the Einstein tensor is a second-order operator of the metric $g_{\mu\nu}$, and the Einstein equations are nonlinear second-order equations relating the spacetime metric to the stress-energy tensor.

To see how metric perturbations of spacetime can act as freely propagating waves, we turn to the weak-field approximation. If we ignore higher-order terms in h , the inverse of the metric is given by:

$$g^{\mu\nu} = \eta^{\mu\nu} - h^{\mu\nu} . \quad (2.8)$$

If we ignore higher-order terms again in the construction of the Christoffel symbols, we have:

$$\Gamma_{\mu\nu}^{\sigma} = \frac{1}{2}\eta^{\sigma\lambda} (\partial_{\nu}h_{\nu\lambda} + \partial_{\nu}h_{\mu\lambda} - \partial_{\lambda}h_{\mu\nu}) . \quad (2.9)$$

Turning to the Ricci tensor, we note that the Γ^2 terms will give rise to second-order terms in the derivative of h , i.e. $(\partial_{\mu}h_{\nu\lambda})^2$. We can assume these terms will be vanishingly small, and so we only include terms which are first-order in the derivatives of Γ . We are left with:

$$\begin{aligned} R_{\mu\nu} &= \partial_{\sigma}\Gamma_{\nu\mu}^{\sigma} - \partial_{\nu}\Gamma_{\sigma\mu}^{\sigma} \\ &= \frac{1}{2} (\partial_{\sigma}\partial_{\nu}h_{\mu}^{\sigma} + \partial_{\sigma}\partial_{\mu}h_{\nu}^{\sigma} - \partial_{\mu}\partial_{\nu}h_{\sigma}^{\sigma} - \partial_{\sigma}\partial^{\sigma}h_{\mu\nu}) . \end{aligned} \quad (2.10)$$

The trace of the metric perturbation is $h_{\sigma}^{\sigma} = h$, and the term $\partial_{\sigma}\partial^{\sigma}h_{\mu\nu}$ is the D'Alembertian, the Lorentz-invariant extension of the Laplace operator to four-dimensional spacetime, denoted by $\square h_{\mu\nu}$. To find the curvature scalar R we contract the Ricci tensor:

$$\begin{aligned} R &= \eta^{\mu\nu}R_{\mu\nu} = \frac{1}{2} (\partial_{\sigma}\partial_{\nu}h^{\sigma\nu} + \partial_{\sigma}\partial_{\mu}h^{\sigma\mu} - \partial_{\mu}\partial^{\mu}h - \square h_{\mu}^{\mu}) \\ &= \partial_{\mu}\partial_{\nu}h^{\mu\nu} - \square h . \end{aligned} \quad (2.11)$$

Putting it all together, the Einstein tensor in the weak-field approximation is:

$$G_{\mu\nu} = \frac{1}{2} (-\square h_{\mu\nu} + \partial_{\sigma}\partial_{\nu}h_{\mu}^{\sigma} + \partial_{\sigma}\partial_{\mu}h_{\nu}^{\sigma} - \eta_{\mu\nu}\partial_{\mu}\partial_{\nu}h^{\mu\nu} + \eta_{\mu\nu}\square h - \partial_{\mu}\partial_{\nu}h) . \quad (2.12)$$

Now, in our expression for $g_{\mu\nu}$, we have left the door open to coordinate transforms. In particular, there may be more than one coordinate system in which the metric may be written as the Minkowski metric plus a small metric perturbation. So far, the metric perturbation is invariant under infinitesimal transformations to the coordinates:

$$x^\mu \rightarrow \tilde{x}^\mu = x^\mu + \epsilon^\mu . \quad (2.13)$$

This amounts to a gauge invariance of the system, and we are free to fix the gauge in a convenient manner. The typical choice is the *harmonic gauge*, defined by the requirement $\square x^\mu = 0$. It can be shown that this is equivalent to the requirement $g^{\mu\nu}\Gamma_{\mu\nu}^\rho = 0$, and in the weak-field approximation this leads to the following relationship:

$$\partial_\sigma h_\nu^\sigma = \frac{1}{2}\partial_\nu h . \quad (2.14)$$

To the harmonic gauge condition, it is convenient to add the definition of the *trace-reversed* perturbation $\tilde{h}_{\mu\nu}$:

$$\tilde{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h . \quad (2.15)$$

The harmonic gauge condition can now be written as

$$\partial_\sigma \tilde{h}_\nu^\sigma = 0 . \quad (2.16)$$

If we use this expression in our weak-field expression for the Einstein tensor, we have

$$G_{\mu\nu} = -\frac{1}{2}\square\tilde{h}_{\mu\nu} . \quad (2.17)$$

Or, bringing back the energy-momentum tensor,

$$\square\tilde{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} . \quad (2.18)$$

In vacuum, the result $\square\tilde{h}_{\mu\nu} = 0$ implies the existence of wavelike solutions for the metric perturbation. Plane-wave solutions will have the form:

$$\tilde{h}_{\mu\nu} = A_{\mu\nu}e^{ik_\alpha x^\alpha} , \quad (2.19)$$

where k_α is the 4-dimensional wave vector and $A_{\mu\nu}$ is the amplitude of the wave. Inserting this solution into the vacuum result for Einstein's equation, we see that the wave vector k_μ must satisfy the relation $k_\mu k^\mu = 0$.

Furthermore, by asserting the harmonic gauge condition $k^\mu A_{\mu\nu} = 0$, we can eliminate two final gauge freedoms by asserting that $A^\mu{}_\mu = 0$ and $A^{0\mu} = 0$. The first equation ($k_\mu k^\mu = 0$) states that the wave vector must be a null vector, which is general relativity's way of saying that the plane wave solutions must travel at the speed of light. The second equation states that the amplitude of the waves is transverse to the direction of motion, and the third states that the wave amplitude is traceless; this choice of gauge is referred to as the *transverse-traceless gauge*. (The fourth equation, $A^{0\mu} = 0$, amounts to a choice of frame, which sounds like a violation of special relativity, but amounts to an assertion that the wave be orthogonal to an observer with a timelike velocity u , $A_{\mu\nu}u^\nu = 0$. See [42] for a rigorous derivation.)

If we choose the direction of propagation of the wave to be in the z-direction,

$$k^\alpha = (\omega, 0, 0, \omega/c), \quad (2.20)$$

then the requirement $k_3 A^{3\nu} = 0$ implies that $A_{3\nu}$ is equal to zero. If $A_{\mu\nu}$ is to be traceless, we are left with only two free parameters, which we label in the following way:

$$A_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_+ & A_\times & 0 \\ 0 & A_\times & -A_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2.21)$$

There are two independent polarizations, referred to as the *plus polarization* and the *cross polarization*. The plane waves themselves can be written as the real part of our solution to Einstein's equation in vacuum. For a wave traveling parallel to the z-axis the solutions take the following form:

$$\begin{aligned} h_+(t, z) &= A_+ \cos(w(t - z/c) + \phi_0) \\ h_\times(t, z) &= A_\times \cos(w(t - z/c) + \phi_0). \end{aligned} \quad (2.22)$$

To see how a gravitational wave might interact with a set of freely-falling test masses, we can write down an expression for the spacetime metric in the presence of a wave with components h_+ and h_\times :

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + h_+ & h_\times & 0 \\ 0 & h_\times & 1 - h_+ & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.23)$$




Figure 2.1. The effect of the two polarizations of gravitational waves on a ring of test masses. The direction of propagation of the wave is perpendicular to the page.

The time-varying spacetime interval separating any pair of masses is given by Eq. 2.1:

$$ds^2 = -dt^2 + (1 + h_+)dx^2 + (1 - h_+)dy^2 + dz^2 + 2h_\times dx dy. \quad (2.24)$$

The motion induced by a passing gravitational wave is illustrated in Fig. 2.1. From this result we see why the effect of gravitational waves on matter is often referred to as a *strain*, since the change in the spacetime interval between two points is proportional to their initial separation (dx , dy , etc).

2.2 Generation of Gravitational Waves

Now, let us consider how a metric perturbation $h_{\mu\nu}$ might be generated from the source term, $T_{\mu\nu}$. We first note that since energy and momentum are conserved, $\partial^\nu T_{\mu\nu} = 0$, i.e. the stress-energy tensor must be divergenceless. (It is interesting to note that the same must be true for the other side of Einstein's equation, i.e. $\partial^\nu G_{\mu\nu} = 0$. An examination of Eq. 2.12 shows that the terms neatly cancel under this operation!)

To solve Eq. 2.18 we employ the Green's function for the D'Alembertian. The Green's function is the solution of the differential equation for a delta function source, $\square G(x^\sigma - x'^\sigma) = \delta^{(4)}(x^\sigma - x'^\sigma)$. The general solution for an arbitrary source $T_{\mu\nu}(x'^\sigma)$ is given by the integral:

$$h_{\mu\nu}(x^\sigma) = -\frac{16\pi G}{c^4} \int G(x^\sigma - x'^\sigma) T_{\mu\nu}(x'^\sigma) d^4 x'. \quad (2.25)$$

The form of the Green's function for the D'Alembertian is:

$$G = -\frac{\delta(|\mathbf{x} - \mathbf{x}'|/c - (t - t'))}{4\pi|\mathbf{x} - \mathbf{x}'|} \theta(t - t'), \quad (2.26)$$

where the function $\theta(t - t') = 1$ for $t > t'$ and zero elsewhere, since we are concerned with the radiation field traveling outward from the source.

Using this expression we solve for $h_{\mu\nu}$:

$$h_{\mu\nu}(x^\sigma) = \frac{4G}{c^4} \int \frac{\delta(|\mathbf{x} - \mathbf{x}'|/c - (t - t'))}{|\mathbf{x} - \mathbf{x}'|} \theta(t - t') T_{\mu\nu}(x'^\sigma) dt' d^3 x'. \quad (2.27)$$

Using the delta function we can solve the integral over time, giving us $t' = t - |\mathbf{x} - \mathbf{x}'|/c$. This leaves:

$$h_{\mu\nu}(t, \mathbf{x}) = \frac{4G}{c^4} \int \frac{T_{\mu\nu}(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}, \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d^3 x'. \quad (2.28)$$

We expand the denominator in the usual way, $|\mathbf{x} - \mathbf{x}'| = d - \mathbf{x}' \cdot \hat{\mathbf{n}} + \mathcal{O}\left(\frac{r^2}{d}\right)$, under the assumption that we are calculating the metric perturbation in the far-field regime, $r \ll d$, where r is the radius of the source and d is the distance to the source. Keeping only the leading term and dropping the prime:

$$h_{\mu\nu}(t, d) = \frac{4G}{c^4} \int \frac{T_{\mu\nu}(t - d/c, \mathbf{x})}{d} d^3 x. \quad (2.29)$$

Calculating the stress-energy tensor for an arbitrary system is quite complicated. Often, the derivation is performed using the multipole expansion of $T_{\mu\nu}$. In this approach, it can be shown that (i) the monopole and dipole terms are zero, and (ii) that the quadrupole terms of the momentum can be written as time derivatives of the mass monopole term. Here, we will take a more mathematical approach (again, following [42]), and use the divergencelessness of $T_{\mu\nu}$ to get the same result.

First, we Fourier transform the previous equation:

$$\tilde{h}_{\mu\nu}(\omega, d) = \int e^{i\omega t} h_{\mu\nu}(t, d) dt = \frac{4G}{dc^4} \int e^{i\omega t} T_{\mu\nu}(t - d/c, \mathbf{x}) d^3 x dt. \quad (2.30)$$

Changing variables from t to the retarded time $t_r = t - d/c$ and Fourier transforming the right hand side, we have:

$$\begin{aligned} \tilde{h}_{\mu\nu}(\omega, d) &= \frac{4G}{dc^4} \int e^{i\omega(t_r - d/c)} T_{\mu\nu}(t_r, \mathbf{x}) d^3 x dt_r \\ &= \frac{4G}{dc^4} \int e^{-i\omega d/c} \tilde{T}_{\mu\nu}(\omega, \mathbf{x}) d^3 x. \end{aligned} \quad (2.31)$$

Now we invoke the divergencelessness of $T_{\mu\nu}$. This is expressed as $\partial^\mu T_{\mu\nu} = 0$, and note that in Fourier space, derivatives over the time coordinate can be expressed as $\partial^0 \tilde{T}_{0\nu} = i\omega \tilde{T}_{0\nu}$. In this way we can relate the time

components to the spatial components, by $i\omega\tilde{T}_{0\nu} = -\partial^k\tilde{T}_{k\nu}$. (Latin indices indicate spatial dimensions of the tensor.) From this we see that we only need calculate the spatial components of $\tilde{T}_{\mu\nu}$. The same logic must apply to $\tilde{h}_{\mu\nu}$, where $i\omega\tilde{h}_{0\nu} = -\partial^k\tilde{h}_{k\nu}$. We employ this trick while integrating $\tilde{T}_{\mu\nu}$ by parts:

$$\int \tilde{T}_{ij}(\omega, \mathbf{x}) d^3x = \int \partial^k \left(x_i \tilde{T}_{kj}(\omega, \mathbf{x}) \right) d^3x - \int x_i \left(\partial^k \tilde{T}_{kj}(\omega, \mathbf{x}) \right) d^3x. \quad (2.32)$$

In the first term the integral over $\partial^k(x_k\tilde{T}_{k\nu})$ is zero since we assume that $T_{\mu\nu}$ goes to zero at infinity, and we can relate the second term to the time component:

$$\int d^3x \tilde{T}_{ij}(\omega, \mathbf{x}) = i\omega \int d^3x x_i \tilde{T}_{0j}(\omega, \mathbf{x}). \quad (2.33)$$

Exploiting the symmetry of \tilde{T}_{ij} , we apply this formula twice, picking up a factor of 1/2 in the process:

$$\begin{aligned} \tilde{h}_{ij}(\omega, d) &= \frac{4G}{dc^4} e^{-i\omega d/c} \int \tilde{T}_{ij}(\omega, \mathbf{x}) d^3x \\ &= \frac{4G}{dc^4} e^{-i\omega d/c} i\omega \int x_i \tilde{T}_{0j}(\omega, \mathbf{x}) d^3x \\ &= \frac{4G}{dc^4} e^{-i\omega d/c} i\omega \int \frac{1}{2} \left(x_i \tilde{T}_{0j}(\omega, \mathbf{x}) + x_j \tilde{T}_{0i}(\omega, \mathbf{x}) \right) d^3x \\ &= \frac{2G}{dc^4} e^{-i\omega d/c} i\omega \int \left[\partial^l \left(x_i x_j \tilde{T}_{0l}(\omega, \mathbf{x}) \right) - x_i x_j \left(\partial^l \tilde{T}_{0l}(\omega, \mathbf{x}) \right) \right] d^3x \\ &= -\frac{2G}{dc^4} e^{-i\omega d/c} i\omega \int x_i x_j \left(\partial^l \tilde{T}_{0l}(\omega, \mathbf{x}) \right) d^3x \\ &= -\frac{2G}{dc^4} e^{-i\omega d/c} \omega^2 \int x_i x_j \tilde{T}_{00}(\omega, \mathbf{x}) d^3x. \end{aligned} \quad (2.34)$$

The term $\int x_i x_j \tilde{T}_{00}(\omega, \mathbf{x}) d^3x$ is the quadrupole moment of the mass-energy distribution of the source, written as \mathcal{M}_{ij} . Fourier transforming back to the time domain, we note that the factor of $-\omega^2$ can be accounted for by taking ∂_t^2 of the time domain function. In the end, we have an expression for the metric perturbation as a function of the retarded time:

$$h_{ij}(t_r) = \frac{2G}{dc^4} \ddot{\mathcal{M}}_{ij}(t_r), \quad \ddot{\mathcal{M}}_{ij}(t_r) = \frac{d^2}{dt^2} \int x_i x_j T_{00}(t_r, \mathbf{x}) d^3x. \quad (2.35)$$

Gravitational radiation is thus generated by the second time derivative of the quadrupole moment of the source energy density. (Remember that we have assumed a nonrelativistic source, where the energy density is larger than the momentum density. This will be a valid assumption for all but the most exotic astrophysical objects.) To phrase this in a more tangible way, gravitational waves are generated by accelerating, asymmetric mass distributions.

What kinds of sources might generate a detectable strain signal $h_{\mu\nu}$? Note that in Eq. 2.35, $\ddot{\mathcal{M}}_{ij}$ has units of energy. A quick manipulation of the constants can give us a sense of the strain generated by potential sources on the Earth:

$$h_{\mu\nu}(t) \simeq 10^{-47} \left(\frac{\ddot{\mathcal{M}}_{\mu\nu}}{1 \text{ Joule}} \right) \left(\frac{1 \text{ km}}{\text{d}} \right), \quad (2.36)$$

...from sources within the galaxy:

$$h_{\mu\nu}(t) \simeq 10^{-22} \left(\frac{\ddot{\mathcal{M}}_{\mu\nu}}{10^{-6} M_{\odot} c^2} \right) \left(\frac{1 \text{ kpc}}{\text{d}} \right), \quad (2.37)$$

...and for astrophysical sources:

$$h_{\mu\nu}(t) \simeq 10^{-23} \left(\frac{\ddot{\mathcal{M}}_{\mu\nu}}{0.01 M_{\odot} c^2} \right) \left(\frac{10 \text{ Mpc}}{\text{d}} \right), \quad (2.38)$$

Strain amplitude of $\mathcal{O}(10^{-23})$ is within the realm of measurement for modern gravitational wave observatories; this will be motivated in the following chapter. Signals from terrestrial sources, with energy scales which we might control in a laboratory, are evidently far below the reach of detectors. Sources from the galaxy and beyond would appear to be good candidates, although the energy scales involved are, for lack of a better word, astronomical.

We should note here that $\ddot{\mathcal{M}}_{ij}$ is the kinetic energy contained in the quadrupolar mass moment of the source energy distribution, it is not the energy carried away by gravitational radiation. That expression will be derived two sections forward.

2.3 Gravitational Waves from Binary Systems

As an example, consider a binary system of two identical stars with mass m in a circular orbit with radius r and orbital angular frequency ω . For simplicity, we consider motion in the x - y plane, and the observer located some distance d along the z axis. We assume the orbit is stable, so ω and r are constants. Recall from Newtonian mechanics that the orbital frequency is given by $\omega = \sqrt{Gm/4r^3}$.

To compute the energy density $T_{00}(t, \mathbf{x})$, we model the stars as point masses, with the first located at $x = r \cos(\omega t)$, $y = r \sin(\omega t)$, and similarly for the second, with a phase shift of π relative to the first. Integrating over $x_i x_j d^3x$ returns factors of $\cos^2 \omega t$ and $\sin^2 \omega t$, which through the familiar trigonometric identities result in terms with frequency 2ω .

Thus, for an observer viewing the system face-on, the metric perturbation will be:

$$\begin{aligned} h_+(t) &= -\frac{8G}{c^4} mr^2 \omega^2 \cos(2\omega t), \\ h_\times(t) &= -\frac{8G}{c^4} mr^2 \omega^2 \sin(2\omega t). \end{aligned} \quad (2.39)$$

When viewed along the axis of rotation, gravitational waves from a rotational system are said to be *circularly polarized*, since the two polarizations have a constant relative phase shift and equal amplitudes. As the observer moves off-axis, the waves will become increasingly elliptically polarized, as the amplitude of the cross-polarization decreases with respect to the plus-polarization. When viewed along the plane of the orbit, binary systems emit linearly polarized gravitational waves. The dependence on the inclination ι of the system relative to the observer is given by:

$$\begin{aligned} h_+(t) &\rightarrow \frac{1}{2}(1 + \cos^2 \iota) h_+(t), \\ h_\times(t) &\rightarrow \cos \iota h_\times(t), \end{aligned} \quad (2.40)$$

where $\iota = 0$ corresponds to an observer located along the axis of rotation and $h_{+,\times}$ are given by Eq. 2.39.

All binary systems will emit energy in the form of gravitational radiation, and from Eq. 2.39 we see that the frequency of this radiation will be twice the binary's orbital frequency. This is a general feature of gravitational radiation from rotating sources.

Energy loss to gravitational waves will cause the orbit to slowly decay. Our best observational evidence for gravitational waves comes from just such a system: the pulsar PSR B1913+16, also known as the Hulse-Taylor pulsar. This is a binary neutron star (BNS) system located 6400 pc from Earth, with masses $1.441M_\odot$ (for the pulsar) and $1.387M_\odot$ (for the companion), in a highly elliptical ($e = 0.617$) orbit with period of about 7.75 hours [143]. Since very shortly after its discovery in 1974, the system has been observed to have a decaying orbit exactly in agreement with the predicted energy loss due to emission of gravitational radiation (see Fig. 2.2). The power emitted by the system in gravitational radiation is 7.35×10^{24} watts, or 1.9% of the power radiated by the Sun.

Over the next three hundred million years, the Hulse-Taylor pulsar will lose energy via gravitational radiation, and the orbit will circularize and steadily contract. In the final moments before merger, the system will emit a significant fraction of a solar mass of energy in gravitational waves. The orbital period will quickly decay from $T = 1$ second to 0.01 seconds in the span of a few hours. As the orbital separation decreases, the expression for the orbital frequency which we derived from Newton's law of gravitation will no longer be valid. To calculate the orbital mechanics as the neutron stars acquire velocities that approach the relativistic limit, various




Figure 2.2. Cumulative excess phase in the orbit of PSR B1913+16, compared to the prediction for damping from gravitational radiation. Negative values represent a progressively earlier periastron, indicative of a decreasing orbital period. The rate of orbital period decay is $-2.423(1) \times 10^{-12}$ seconds/second. Figure from [143].

approximations to the full general relativistic calculation can be made. These approximation techniques include terms of the multipole expansion of the stress-energy tensor higher than the quadrupole term, and also apply relativistic corrections of order $(v/c)^2$ to the quadrupolar terms we have calculated above. These *post-Newtonian* corrections account for various relativistic effects such as frame-dragging.

Even then, the validity of the post-Newtonian approximation is limited to orbits where the size of the binary companions is small compared to the orbital radius. For neutron stars, with radius $\mathcal{O}(10\text{km})$, the approximation breaks down when the orbital frequency is above 1kHz and the non-point-like nature of the neutron stars must be considered. In particular, the neutron stars are expected to experience tidal deformations in the final orbits of the system before merger [71]. The magnitude of the tidal deformations and their effect on the bulk motion of the matter in the system is dependent on the neutron star equation of state, which has not been constrained by observation [84]. Modeling the dynamics of the nuclear matter in BNS mergers using fully relativistic hydrodynamic simulations which account for magnetic field interactions is an active and challenging area of research. See for example Ref. [110].

Binary black holes systems (BBH), on the other hand, can be accurately modeled as point-like particles using numerical solutions to Einstein's equations [45]. Many dozens of BBH waveforms have been generated by groups using supercomputers to model systems with a variety of masses and spins.

For all binary systems, the inspiral stage of the binary evolution will end when the orbit becomes a plunge into the gravitational well. The start of the plunge is typically approximated by the smallest radius of a stable orbit in the Schwarzschild metric. This is referred to as the *innermost stable circular orbit* (ISCO). The orbital frequency at this point is given by:

$$f_{ISCO} = \frac{1}{6\pi\sqrt{6}} \frac{c^3}{GM}, \quad (2.41)$$

where M is the total mass of the system. For a system with two neutron stars, $M = 2.8M_{\odot}$, and $f_{ISCO} \simeq 1570\text{Hz}$. The frequency of the gravitational wave emission at the point of the plunge will be $2f_{ISCO}$, but the evolution of the system from $f = \frac{f_{ISCO}}{2}$ to $f = f_{ISCO}$ is so rapid that f_{ISCO} is generally taken to be the upper limit on gravitational wave emission from the inspiral of a binary system with mass M .

2.4 E_{GW} and Strain Amplitude

Some complications arise if we wish to calculate the energy flux from gravitational radiation starting from Eq. 2.18. One issue is our assumption of the weak-field limit in the derivation of a linearized $G_{\mu\nu}$, namely that the gravitational wave is propagating through flat space with no energy density. An added challenge is the relationship between the energy-momentum tensor and $h_{\mu\nu}$: classically, the energy flux from a propagating wave is quadratic in the perturbed field, but we have derived the weak-field approximation using only the first order terms of h . The reader who is troubled by these philosophical and technical inconsistencies is invited to consult the literature. For now, we will quote the standard solution for the energy flux and move forward from there.

If we expand the Ricci tensor to second order (quadratic in h) and equate to the stress-energy tensor (written here in lowercase to acknowledge our ambivalence with the calculation¹), we have:

$$t_{00} = \frac{c^5}{32\pi G} \langle \partial_0 h_{ij}(t) \partial_0 h_{ij}(t) \rangle, \quad (2.42)$$

where the angle brackets indicate the average over many cycles. Noting that $\partial_0 = \partial_t/c$, and taking the time derivative over the nonzero elements in Eq. 2.23, the energy flux from a passing gravitational wave is:

$$F_{GW} = \frac{c^3}{16\pi G} \langle \dot{h}_+^2(t) + \dot{h}_\times^2(t) \rangle. \quad (2.43)$$

¹In defense of this approach the results are in perfect agreement with the Hulse-Taylor pulsar.

We wish to relate the energy emitted by the source in gravitational radiation to the strain measured by the observer. To continue, we need to assume some form for $h_{+, \times}(t)$. A natural choice would be the results from our example of the binary system; however, as will be motivated in Chapter 7, we will instead use *sine-Gaussians*, a family of generic waveforms that lends itself to models for transient signals. While sine-Gaussians do not directly model an astrophysical source, they are a good approximation for a number of rotating systems and are widely used in gravitational wave data analysis.

Sine-Gaussian waveforms are given by sine and cosine functions of fixed frequency, modulated by a Gaussian envelope function:

$$\begin{bmatrix} h_+(t) \\ h_{\times}(t) \end{bmatrix} = \frac{h_{r_{ss}}}{(2\pi\sigma_t^2)^{1/4}} \begin{bmatrix} \frac{(1+\cos^2\iota)}{2} \\ \cos\iota \end{bmatrix} \begin{bmatrix} \sin 2\pi f_0(t-t_0) \\ \cos 2\pi f_0(t-t_0) \end{bmatrix} e^{-(t-t_0)^2/4\sigma_t^2}, \quad (2.44)$$

where t_0 is the central time of the waveform, f_0 is the central frequency, σ_t is the duration, and we define the quality factor $Q = 2\pi f_0 \sigma_t$ such that the bandwidth of the waveform is given by $\sigma_f = f_0/2Q$. The inclination of the source is again given by ι , with $\iota = 0$ indicating an optimally-oriented, face-on rotating system with respect to the observer.

Eq. 2.44 is defined such that the root sum of the squares is normalized, i.e.:

$$\sqrt{\int_{-\infty}^{\infty} |h_+(t)|^2 + |h_{\times}(t)|^2 dt} = h_{r_{ss}}, \quad (2.45)$$

and we call the amplitude $h_{r_{ss}}$ the *root sum square* strain amplitude. $h_{r_{ss}}$ is a useful quantity for the calculation of signal strength in gravitational wave data analysis, as we shall see in Chapter 7. Note that $h_{r_{ss}}$ has units of amplitude spectral density, $\text{Hz}^{-1/2}$. These units are explained in Appendix A.

Now we are in a position to integrate Eq. 2.43 and calculate the energy carried by a gravitational wave. To compute the time average of the derivative of the strain signals, we approximate the sine-Gaussian waveforms as narrowband signals in frequency; this gives us an overall factor of $(2\pi f_0)^2$. (Riles [111] has shown that this approximation is accurate for $Q \gtrsim 6$.) We assume that the system is optimally oriented (again, a choice motivated in Chapter 7), and integrating over the whole sky gives a factor of $2\pi r^2$ (2π instead of 4π because we have fixed the inclination angle). We have:

$$\begin{aligned} E_{GW} &= \int F_{GW} d\Omega dt = \frac{c^3}{16\pi G} (2\pi f_0)^2 2\pi r^2 h_{r_{ss}}^2 \left[\int_{\iota=0} d\cos\iota \frac{(1+\cos^2\iota)^2}{4} + \cos^2\iota \right] \\ &= \frac{2}{5} \frac{\pi^2 c^3}{G} f_0^2 r^2 h_{r_{ss}}^2. \end{aligned} \quad (2.46)$$

The prefactor of $2/5$ is due to the rotational, face-on nature of the system; if we had assumed isotropic emission from the source the prefactor would be unity. (For the same h_{rss} measured by an observer, isotropic emission leads to larger E_{GW} .) A quick calculation of the constants shows we are in agreement with the estimates from two sections prior: for a source 20 Mpc away, emitting gravitational waves at 200Hz, a strain of 10^{-22} indicates $E_{GW} = 0.003M_{\odot}c^2$. Note that we have made no statement about the duration of the signal. Our approximation of $Q \gtrsim 6$ is equivalent to an assumption that the signal is at least several cycles long, and h_{rss} is the average strain amplitude over this time. For $Q = 10$ and $f_0 = 200$ Hz, $\sigma_t = 8$ msec.

If we integrate F_{GW} over the whole sky and use our expression for $h_{ij}(t)$ from the transverse-traceless gauge, Eq. 2.35, we get the following approximation for P_{GW} :

$$P_{GW} = \frac{c^3}{16\pi G} 4\pi r^2 \left(\frac{2G}{rc^4} \right)^2 \langle \ddot{\mathcal{M}}^2 \rangle = \frac{G}{c^5} (\epsilon MR^2 \omega^3)^2, \quad (2.47)$$

where we have expressed the quadrupolar mass moment $\mathcal{M} = \epsilon MR^2$, with $\epsilon \leq 1$ the measure of the nonuniformity of the mass distribution, and we have assumed the motion of the quadrupolar moment to be periodic with angular frequency ω . If we say that $\omega = v/R$, then we can simplify the expression:

$$P_{GW} = \frac{G}{c^5} \left(\frac{M}{R} v^3 \right)^2. \quad (2.48)$$

Hence the best gravitational wave sources will be massive, compact, and rapidly moving.

In section 36.4 of Misner, Thorne, and Wheeler [94], an argument is made for calculating the maximum possible emission of gravitational waves from a compact source. Taking Eq. 2.48 to the extreme, we say that the upper limit on the velocity is the speed of light, and the lower limit for R is the Schwarzschild radius for the total mass of the system, $R_S = 2GM/c^2$. After some dimensional gymnastics, this gives $P_{GW} \simeq c^5/4G = 5 \times 10^4 M_{\odot} c^2 s^{-1}$. Of course we have used relativistic parameters in a nonrelativistic model – this result is perhaps only useful for its comic effect!

Using more reasonable numbers, a system of two neutron stars ($M = 2.8 M_{\odot}$) with an orbital radius $R = 50$ km ($= 6R_S$) and a frequency $\omega = \sqrt{1.4M_{\odot}G/4R^3} = 2\pi \times 100$ Hz (for $v \sim 0.1c$) emit $P_{GW} = 0.002 M_{\odot} c^2 s^{-1}$. This gives some sense of the energy released into gravitational waves in the final moments of the merger of two neutron stars; it is tremendous. The gravitational wave energy flux at the Earth from such a system in the Virgo cluster, 20 Mpc away, is ~ 1 mW/m³ – about the same as the electromagnetic energy flux from a full moon.

2.5 Astrophysical Sources of Gravitational Waves

To conclude this chapter, we give a brief review of astrophysical sources for that could release sufficient E_{GW} to be detectable.

As implied by our example, binary systems are reliable sources of gravitational waves. Binaries composed of main-sequence stars have masses too small and separations too large to generate $h_{\mu\nu}$ of significant amplitude, although there are several known white dwarf binaries that could be detected by a space-based gravitational wave experiment. (In fact the expected signal for a space-based experiment is so reliable that these systems are referred to as ‘verification binaries’.)

Compact binaries, where one or both objects are neutron stars or black holes, are efficient radiators of gravitational waves. The mergers of compact binaries are known as *compact binary coalescences* (CBCs). As of this writing there are ten known binary neutron star (BNS) systems in the galaxy, including the Hulse-Taylor pulsar. None of these systems release sufficient energy to be detectable, but the population is used to constrain the rate of BNS mergers in the galaxy, using assumptions of pulsar luminosity distributions and the typical lifetime of the system before merger. Similarly, while there are no known neutron star-black hole (NSBH) binaries or binary black hole (BBH) systems, their populations can be estimated from the galactic star forming rate and observations of high-mass X-ray binaries. BNS systems are expected to merge about once every 10,000 years in a Milky-Way-like galaxy, and NSBH and BBH rates are about two orders of magnitude lower [12]. Due to the extrapolation from a small population and the poor constraints on theoretical parameters, these rate estimates are subject to large uncertainties, typically two orders of magnitude in either direction.

CBCs are transient signals which emit large E_{GW} for only a few seconds. We can imagine other astrophysical systems that might generate rapidly varying quadrupolar moments for a short amount of time. Among the best candidates are supernovae, which are known to have significant asymmetry and liberate a significant fraction of a solar mass of gravitational binding energy. Type Ia supernovae, generated by the thermonuclear collapse of a white dwarf in a binary system, are not energetic enough to be detectable. But core-collapse supernovae (CCSN) of Type II or Ib,c can result in the formation of a neutron star or black hole. The mechanism for core-collapse supernovae is not constrained by observation or theory and the gravitational wave emission from these systems is highly speculative. If lower estimates of E_{GW} from supernovae are correct ($10^{-6} M_{\odot} c^2$), only galactic CCSN will be detectable [28].

Isolated neutron stars can have large angular momenta, with spin frequencies approaching 1 kHz. Any deformations from spherical symmetry will generate gravitational waves, although the deformation would have to be quite large for the signal to be detectable. Neutron stars that have some oblateness which is not aligned with the spin axis will precess and emit gravitational waves. This oblateness can be generated by strong magnetic fields (again, not aligned with the spin axis) or anisotropies in the star’s interior. Gravitational signals from nonax-

isymmetric neutron stars will be long-duration and fixed frequency, up to Doppler shifts due to relative motion between the source and the observer. These signals are known as *continuous wave* (CW) signals in gravitational wave data analysis. Direct measurements of neutron star physics via gravitational waves, either from observations of pulsars or from BNS mergers, are considered to be the best chance to constrain the equation of state for dense nuclear matter [48, 109].

Neutron stars with large magnetic fields are known as *magnetars* [92]. Magnetars can occasionally experience dramatic realignments of the magnetic field lines which can be associated with significant shifts in the star's mass distribution due to crustal cracking and starquakes. The magnetic field realignment generates streams of high-energy particles and photons, which are detected at Earth as bursts of gamma-rays; magnetars that generate bursts of this kind are called *soft gamma repeaters* or SGRs. The gravitational wave emission associated with SGR events, sometimes called *flares*, is not well constrained by theory, but there are known SGRs within several kiloparsecs of Earth.

Finally, a diffuse population of random, unresolved sources can lead to a *stochastic background* of gravitational wave energy [26]. This is typically expressed as Ω_{GW} , or the ratio of the gravitational wave energy density to the critical density required to close the universe. A stochastic gravitational background could arise from the galactic population of binary systems (at low frequency, $\sim\mu\text{Hz}$), the cosmological population of BNS or BBH mergers (at high frequency, $\sim 100\text{ Hz}$), the cosmological population of supermassive black hole binaries (at *very* low frequency, $\sim\text{nano Hz}$), the gravitational analogue to the cosmic microwave background (which has a flat frequency content in basic models of inflation), or from more exotic sources such as cosmic strings or axions. Gravitational waves that were excited in the very early universe can leave their imprint on the polarization of the CMB; this indirect measurement via experiments like BICEP2 probes gravitational radiation with frequencies around an inverse Hubble time (H_0^{-1}), or 10^{-18} Hz .

CHAPTER 3

GRAVITATIONAL WAVE DETECTORS

Not from the stars do I my judgment pluck

Gravitational wave detectors are transducers for gravitational wave strain: they convert a metric perturbation from a passing gravitational wave into a voltage signal appropriate for digital analysis. The first practical gravitational wave detectors were *bar detectors*, pioneered by Joseph Weber [142]. Bar detectors are composed of large masses of metal, typically aluminum and typically cylindrical in shape. The surface of the bar is instrumented with low-noise piezoelectric transducers; a passing gravitational wave excites the solid body modes of the mass and generates oscillations that are measured by the piezo crystals. As a result bar detectors are very narrow-band instruments (a few Hz or less), only sensitive to gravitational waves very near the frequency of the bar's body modes (usually, around several hundred Hz, depending on the size of the bar). Weber's methods revolutionized the until-then theoretical field of gravitational wave astronomy and motivated a generation of experiments. While Weber's work is regarded as pioneering, his later results – which included claims of detecting gravitational waves associated with SN1987A – are now presented as a cautionary tale in experimental physics.

Modern gravitational wave detectors measure strain as the relative displacement between two test masses. Early work on a simplified design was carried out in the 1970s by R. Forward and others [62]. The lasting model for modern experiments (along with a careful analysis of what continue to be the dominant limits to sensitivity) was developed by Rainer Weiss [144] with important contributions from R. Drever. In the most basic sense, these detectors are kilometer-scale versions of the Michelson-Morely experiment to measure the speed of light in different inertial reference frames, although a variety of optical enhancements are used to boost the signal. Terrestrial gravitational wave detectors are more broadband than bar detectors, with sensitivity across the full acoustic frequency range, between 10 Hz and a few kiloHertz.

There are currently a handful of kilometer-scale implementations of the experiment that Weiss envisioned. The LIGO observatories, in Livingston, LA, and Hanford, WA, have four-kilometer-long arms [20]; Virgo, in Cascina, Italy, is three kilometers [21]; and GEO 600, in Hannover, Germany, is 600 meters long [69]. A Japanese




Figure 3.1. Optical layout of a simple Michelson interferometer. The detection photodiode is placed at the antisymmetric port of the interferometer.

collaboration is currently building the KAGRA detector in an underground facility in the Kamioka mine in Japan with 3 km arms [83]. Finally, a third LIGO site is planned for India, built in cooperation with the Indian government. Within the gravitational wave community, the detectors are known by their shorthand names: the LIGO instruments are H1 and L1, Virgo is V1, GEO is G1. A third LIGO instrument, H2, was formerly installed at the Hanford site, and had 2 km arms. H2 ceased operations in 2010. All of the detectors are housed within elaborate vacuum systems that limit acoustic interference and phase noise caused by refraction from density fluctuations air. These vacuum systems are some of the largest in the world.

In this chapter, I describe the formalism of gravitational wave signal extraction for a simple Michelson interferometer, and derive the response function for the first generation of detectors. This response function is used to calculate the limit to detector sensitivity for shot noise, a fundamental noise source. I discuss the antenna response of the detectors, and conclude with a brief discussion of astrophysical results from the initial detector era.

I note in passing that the search for gravitational waves associated with gamma-ray bursts presented in Chapter 7 uses data from the initial detectors. The derivations in this chapter are intended to provide the reader with a sense for how the detectors used in that search achieved their strain sensitivity.

3.1 Optical Gain of the Initial LIGO Detectors

Let us consider a Michelson interferometer on the surface of the Earth with two arms of length L oriented along the x and y axes (see Fig. 3.1). The gravitational wave is incident from above. For simplicity, this wave will have $h_{\times} = 0$. Using Eq. 2.24, we can calculate the change in path length for the electromagnetic wave of a laser propagating down the x -arm. For all electromagnetic waves the spacetime interval between events is always light-like, $ds^2 = 0$ (recall Eq. 2.1). Using this, we can write the following integral:

$$L_x = \int_0^L c dt = \int_0^L \sqrt{1 + h_+(t)} dx. \quad (3.1)$$

Since $|h_+(t)| \ll 1$, we approximate the term in the square root, and have the following:

$$L_x \simeq L + \frac{1}{2} L h_+(t). \quad (3.2)$$

For the y -arm, the expression is the same, but the term with the metric perturbation is negative.

Consider now the laser field that propagates up and down the arms and arrives at the *antisymmetric port* or AS port, the side of the beamsplitter opposite that of the input laser. The input beam is described by an electric field E_{BS} incident on the beamsplitter, and the field at the AS port is given by:

$$E_{AS} = E_{BS} (t_{bs} e^{ikL_x} r_x (-r_{bs}) e^{ikL_x} + r_{bs} e^{ikL_y} r_y t_{bs} e^{ikL_y}), \quad (3.3)$$

where the reflectivity and transmissivity of the beamsplitter is given by r_{bs} and t_{bs} , the amplitude reflectivities of the end mirrors are r_x and r_y , and the length of the arms is L_x and L_y . Note that we have followed the convention that the reflection off the back surface of the beamsplitter acquires an overall minus sign; this is required for energy conservation.

We will assume that the end mirrors are perfectly reflecting ($r_x = r_y = 1$) and the beamsplitter is perfectly 50-50 and lossless, so $r_{bs} = t_{bs} = 1/\sqrt{2}$. Using our expression for the arm length in terms of gravitational wave strain (Eq. 3.2):

$$E_{AS} = -\frac{E_{BS}}{2} e^{2ikL} (e^{ikLh_+(t)} - e^{-ikLh_+(t)}) = -iE_{BS} e^{2ikL} \sin[kLh_+(t)], \quad (3.4)$$

The power measured by a photodetector at the antisymmetric port is:

$$P_{AS} = |E_{AS}^* E_{AS}| = P_{BS} \sin^2[kLh_+(t)], \quad (3.5)$$

where $P_{BS} = |E_{BS}|^2$ is the power in Watts incident on the beamsplitter. Note that due to our phase convention, a strain of zero leads to destructive interference at the antisymmetric port between the fields from the two arms, and none of the input power is transmitted to the AS port. This is commonly referred to as the *dark fringe*; the terminology *fringe* dates to the original Michelson-Morely experiment. To take advantage of the benefits of nullity [114], gravitational wave interferometers are usually operated with the antisymmetric port on the dark fringe, and nearly all of the incident power is reflected back from the beamsplitter towards in the input beam. Operating on the dark fringe provides technical benefits, since our sensitivity to fluctuations in the input field (E_{BS}) or the laser frequency (k) will be zero when the strain is zero.

Unfortunately, there is a trade-off. Operating on the dark fringe limits our sensitivity to the strain. This can be seen by differentiating the expression above:

$$\frac{\delta P_{AS}}{\delta h} = 2P_{BS} kL \sin(kLh) \cos(kLh). \quad (3.6)$$

This expression is often referred to as the *optical gain* of the interferometer because it relates Watts of power at the anti-symmetric port to the strain signal experienced by the arms. It provides some basic guidance for interferometer design: we want to maximize the arm length L , maximize the power incident on the beamsplitter P_0 , and minimize the laser wavelength $\lambda = \frac{2\pi}{k}$. Since the strain h is expected to be very small, the cosine term is very close to unity and can be neglected. But, note that due to the sine term, the optical gain scales linearly with h . For $h \simeq 0$ we will have no signal!

Extracting a measurable signal while maintaining the noise-cancellation benefits of the dark fringe is one of the primary complications of interferometer design. The typical solution is to employ a *local oscillator field* (E_{LO}) at the antisymmetric port in addition to the field that carries the signal from the arms. The power at the AS port is then given by $P_{AS} = |E_{LO} + E_{AS}|^2$, where E_{AS} is from Eq. 3.4. Cross-terms between the LO field and the AS field will result in a signal in P_{AS} that is linear in δh even when the phase shift between the arms is zero. In the first generation of interferometric detectors this local oscillator was a radio frequency modulation of the carrier light. Such a detection scheme is referred to as *heterodyne detection*, because the local oscillator field has a different frequency than the carrier field that contains the signal. Second-generation interferometers use *homodyne detection*, in which the local oscillator field is drawn from the same carrier field that propagates down the arms. This field is generated using a static differential arm offset, which permits a small amount of carrier light to exist at the AS port. Details on the homodyne readout scheme for advanced LIGO are presented in Chapter 5.

Modern gravitational wave interferometers are significantly more complicated than our simple Michelson example. In particular, the arms of the interferometer are formed by *Fabry-Perot optical cavities* which amplify

the phase shift acquired by the electric field due to gravitational wave strain. Another Fabry-Perot cavity is formed using an additional mirror placed at the symmetric port of the Michelson; this *recycling mirror* amplifies the power incident on the beamsplitter. The advanced LIGO detectors have yet another cavity at the output of the interferometer called *signal recycling cavity*, which broadens the frequency response of the detector to gravitational waves. The physics of Fabry-Perot cavities will be derived in the following chapter and a description of the advanced LIGO detectors will be given in Chapter 5. For now we will provide the response function for the initial LIGO detectors. Our intention is to illustrate the physics of the first generation gravitational wave detectors, which collected data for the analysis presented in Chapter 7.

The response of a power-recycled Fabry-Perot Michelson interferometer (PRFPMI) is given by:

$$P_{AS}(f) = P_{BS} \sin^2[g_\phi(f) kLh], \quad (3.7)$$

where the phase shift from the arms is amplified by a frequency-dependent phase gain $g_\phi(f)$. The phase gain is expressed using \mathcal{F} , the *finesse* of the arm cavities, and the cavity pole frequency f_{pole} , above which the cavity response to length fluctuations is reduced:

$$g_\phi(f) = \frac{2\mathcal{F}}{\pi} \left(1 + i \frac{f}{f_{\text{pole}}} \right)^{-1}. \quad (3.8)$$

Cavity finesse is a measure of the light storage time in the cavity, and the cavity pole frequency defines the cavity linewidth. These quantities will be explained in greater detail in the following chapter. We should note here that the cavity pole is a function of the cavity finesse (see Eq. 4.9), and increasing the cavity finesse improves the detector sensitivity at low frequencies but reduces the frequency range where the detector is most sensitive. This frequency range is referred to as the detector *bandwidth*¹.

To derive the full optical gain for a power-recycled interferometer, we differentiate Eq. 3.7, drop the cosine term, and use the undifferentiated expression to relate the sine term to P_{AS} and P_{BS} :

$$\begin{aligned} \delta P_{AS} &= 2P_{BS} \sin[g_\phi(f) kLh] \cos[g_\phi(f) kLh] g_\phi(f) kL \delta h \\ &= 2P_{BS} \sqrt{\frac{P_{AS}}{P_{BS}}} g_\phi(f) kL \delta h \\ &= \frac{4\mathcal{F}kL}{\pi} \sqrt{P_{AS}P_{BS}} \left(1 + i \frac{f}{f_{\text{pole}}} \right)^{-1} \delta h. \end{aligned} \quad (3.9)$$

¹Increasing the finesse does not sacrifice sensitivity at high frequency, and in the ideal case there is no reason not to increase the finesse as much as possible. In reality there are significant control difficulties for high-finesse cavities, and the sensitivity at low frequencies is limited by a variety of noise sources. See Appendix C of [121] for a discussion.

Parameter	iLIGO Design
Arm length, L	3994.5 m
Input power, P_{IN}	20 W
Laser wavelength, λ	1064 nm
Carrier recycling gain, g_{cr}^2	50
Arm cavity finesse, \mathcal{F}	220
Arm cavity pole, f_c	85 Hz

Table 3.1. Sensing parameters of the initial LIGO detectors.

Finally, we can relate the power on the beamsplitter, P_{BS} , to the input power to the interferometer using the gain of the power recycling cavity: $P_{BS} = g_{cr}^2 P_{IN}$. Using these expressions, we calculate the optical gain to be:

$$\frac{\delta P_{AS}}{\delta h} = 8g_{cr} \frac{\mathcal{F}L}{\lambda} \sqrt{P_{IN}P_{AS}} \left(1 + i \frac{f}{f_{pole}}\right)^{-1}. \quad (3.10)$$

This expression defines the signal at the antisymmetric port due to a small strain perturbation in the arms.

It is appropriate here to point out an important assumption that we have made, namely that the light travel time down the arm is small compared to any time evolution of the metric perturbation $h(t)$. An equivalent statement is that the arm length is small compared to the wavelength of the gravitational radiation. For gravitational waves of high frequency (small wavelength), the detector response will be degraded, and for signals in $h(t)$ with periods equal to the light travel time in the arms the response will be zero. For the 4 km LIGO arm cavities this null point is approximately 37.5 kHz, and well above the practical limit to sensitivity due to the cavity pole of the Fabry-Perot arms.

Parameters for the initial LIGO detectors are given in Table 3.1. In homodyne readout, P_{AS} will depend on a differential arm offset that is large compared to the amplitude of gravitational wave signals but small compared to the wavelength of the laser. The size of this offset determines the static value for P_{AS} and the response to small metric perturbations. In practice this offset is adjusted on a daily timescale to respond to changes in detector operations, and small changes to the overall detector response are dynamically calculated and accounted for.

Prior to the S6 science run in 2009-10, the LIGO detectors used heterodyne readout for sensing the antisymmetric port power. In this case P_{AS} is given by the power in radio-frequency sidebands that are impressed onto the laser before the input to the interferometer. This leads to a more complicated expression for Eq. 3.7 that includes terms for the power in the RF sidebands and their transmissivity through the power-recycled Michelson. Since a focus of this thesis is the instrumentation used for homodyne detection, we have neglected any discussion of heterodyne readout.

3.2 Limits to Detection

How small a strain δh can we hope to measure with a terrestrial interferometer? Since the Michelson interferometer converts a change in path length into a change in intensity at the antisymmetric port, the question changes from “How small a strain can we measure?” to “How small a change in laser intensity can we detect?” Modern photodetectors can easily measure changes in intensity in the parts per million; even Michelson and Morely were able to measure relative phase shifts of about $\lambda/20$, with visible light. The limiting factors for measuring relative length changes between the arms are *noise sources* in the detector itself. These noise sources can be indistinguishable from time-varying metric perturbations and determine the limit of the detector sensitivity.

The noises that can impact a detector are legion and include forces that act on the mirrors (seismic motions from the ground, residual gas in the vacuum system), thermoelastic noise that moves the mirrors at microscopic scales (thermally excited modes in the mirror suspensions, thermal noise in the mirror coatings), noise in the laser field itself (small variations in intensity or wavelength), and noise from the analog and digital electronics that sense and control the position and orientation of the mirrors.

In this section we outline the sensitivity limits due to the quantum nature of light. The signal that is measured by a photodiode at the antisymmetric port is subject to the Poisson nature of photon counting statistics on the photodiode itself. The statistical variation in the number of photons arriving at the photodiode limits the size of true intensity fluctuations which are detectable. This is referred to as *shot noise*. For most of the LIGO detection band, shot noise is the limiting noise source. The remainder of this section discusses shot noise.

The mean rate of arrival of photons at the antisymmetric port photodiode is given by $N = P_{AS}/\hbar\omega$, where ω is the laser frequency. From Poisson statistics, the power measured at the AS port will have standard deviation given by:

$$\sigma = \hbar\omega\sqrt{N} = \hbar\omega\sqrt{\frac{P_{AS}}{\hbar\omega}} = \sqrt{\hbar\omega P_{AS}}. \quad (3.11)$$

Let us now shift to the frequency domain. For Gaussian noise with variance σ^2 , the value of the *two-sided power spectral density* (PSD) is a constant, given by $G_n = \sigma^2$, in units of Hz^{-1} . Noise of this kind is called *white noise*. Since the output of a Michelson interferometer is proportional to the amplitude of the strain of the gravitational wave, we typically visualize the frequency-domain noise in terms of the *amplitude spectral density*, or ASD, in units of $\text{Hz}^{-1/2}$. The one-sided ASD for shot noise is given by:

$$S_n = \sqrt{2G_n} = \sqrt{2}\sigma = \sqrt{2\hbar\omega P_{AS}}. \quad (3.12)$$

This expression gives us the variation in P_{AS} due to photon counting statistics. (Since we often deal in terms of induced photocurrent rather than incident power on a photodiode, the expression $S_n = \sqrt{2qI}$, where q is the




Figure 3.2. Shot noise limit from Eq. 3.13 for the initial LIGO detectors, and a representative sensitivity curve for the H1 detector during the S6 science run.

fundamental electric charge and I the photocurrent, is also useful.) For some notes on amplitude spectral density, see Appendix A.

Using our expression for the shot-noise variation in P_{AS} , we can calculate the equivalent strain signal this will generate in our interferometer, using our expression for the initial LIGO optical gain (Eq. 3.10):

$$\delta h_{shot} = \frac{S_n}{\delta P_{AS}/\delta h} = \frac{1}{\sqrt{\epsilon}} \frac{1}{4g_{cr}FL} \sqrt{\frac{\pi \hbar c \lambda}{P_{IN}}} \left(1 + i \frac{f}{f_{pole}} \right). \quad (3.13)$$

The prefactor of $\epsilon^{-1/2}$ is an overall power efficiency which accounts for losses between the input to the interferometer and the antisymmetric port. These include losses in auxiliary optics such as mode-matching telescopes and Faraday isolators, and also small percentages of the light which are picked off from the main beam and used for sensing and control. For the H1 detector in the S6 science run, $\epsilon = 0.42$.

In Fig. 3.2 we compare the sensitivity limit predicted by Eq. 3.13 and the sensitivity realized during the last LIGO-Virgo science run. The projection for the shot noise uses the parameters from Table 3.1. At low frequencies, the detector is limited by noise sources such as seismic vibrations and thermal noise in the mirror coatings. Above roughly 300 Hz the sensitivity is limited only by shot noise. In this region, an increase in the detector's input power P_{IN} would directly translate into a sensitivity improvement of $\sqrt{P_{IN}}$. Increasing the input power is, however, technically challenging for many reasons [54, 117].

Note that the shot noise limit to sensitivity does not depend on P_{AS} . While the variance in the rate of arrival of photons at the detector increases like $\sqrt{P_{AS}}$, so does the optical gain of the interferometer, and the two effects cancel.

Due to the quantum nature of light, shot noise is a fundamental lower limit on interferometer sensitivity. In a fully quantum-mechanical treatment [43, 44], shot noise is realized as the introduction of *vacuum fluctuations* that enter the interferometer through the antisymmetric port. A trick that has been shown to be effective in reducing shot noise is to constrain the amplitude or phase uncertainty of the vacuum fluctuations by injecting another laser field into the antisymmetric port of the interferometer. This laser, phase-locked to the input laser, *squeezes* the uncertainty of the vacuum into one of the two degrees of freedom of the laser (amplitude, or phase). This technique has been implemented on the GEO 600 detector in Germany and the LIGO H1 detector [14, 3] with good results, and will likely be included in the design of future experiments.

3.3 Detector Antenna Patterns

In general, gravitational wave detectors are more omni-directional than a typical electromagnetic dipole antenna. Their sensitivity as a function of sky location depends on the polarization of the gravitational wave. For a interferometer with two orthogonal arms, the detector antenna functions are given by:

$$\begin{aligned} F_+(\theta, \phi) &= \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi, \\ F_\times(\theta, \phi) &= \cos \theta \sin 2\phi, \end{aligned} \tag{3.14}$$

where $F_+(\theta, \phi)$ is the detector's sensitivity to +-polarized waves originating from the sky location defined by (θ, ϕ) , and likewise for F_\times . Here, the arms of the detector lie along the positive x and y axes, ϕ is the azimuthal angle ($0 \leq \phi \leq 2\pi$) which is zero along the x axis, and θ is the polar angle ($0 \leq \theta \leq \pi$). These equations are plotted in Fig. 3.3 and Fig. 3.4. The interferometer is maximally sensitive in the $\pm z$ direction, and there is no sensitivity in the plane of the detector to \times -polarized gravitational waves, due to the $\cos \theta$ term in h_\times . There are four *null points* in the x - y plane where the detector has no sensitivity to gravitational waves of any polarization.

The time-series data from a gravitational wave detector due to a source at a sky location (θ, ϕ) emitting gravitational waves given by $(h_+(t), h_\times(t))$ can be written as:

$$d(t, \theta, \phi) = h_+(t)F_+(\theta, \phi) + h_\times(t)F_\times(\theta, \phi) + n(t), \tag{3.15}$$

where $n(t)$ is a noise term intrinsic to the detector. This expression will be important when we describe methods of coherent data analysis in Chapter 7.




Figure 3.3. F_+ antenna pattern. The z -axis is vertical.




Figure 3.4. F_x antenna pattern.




Figure 3.5. Limits on the rate of compact binary coalescences set by the most recent LIGO-Virgo science runs. The light grey regions are the limits set by the S5 science run, and the dark grey regions are the limits set by the combined results of the S5 and S6 science runs. The blue regions are the range of predicted rates based on population studies, as detailed in [12]. Figure from [18].

3.4 Results from the Initial Detector Era and Current Astrophysical Limits

In the previous decade, the first generation of interferometric detectors completed multiyear observing runs. A variety of astrophysical searches were performed using this data, and no detection candidates were identified. In this section we provide a brief survey of the most general results.

The most promising signals for terrestrial detectors are compact-binary coalescences (CBCs). Searches for CBCs perform a matched-filter analysis of the data using the known waveform of the binary inspiral, for a wide parameter space of binary masses, spins, and orientations. Results from all-sky searches for BNS, NSBH, and BBH signals were most recently published in Ref. [18], and the upper limits are illustrated in Fig. 3.5. The sensitivity of the searches to date is still an order of magnitude above the predicted rates for CBCs from population synthesis models.

A variety of different methods have been used to search for continuous-wave signals from isolated pulsars. Analyses of this kind search for monochromatic signals in the detectors which evolve following the expected sidereal Doppler shift from the Earth’s rotation relative to a single point on the sky, as well as the frequency derivative from the spin-down of the signal. Continuous-wave (CW) searches can be directed, using ephemeris from radio observations to analyze known pulsars, or undirected, searching for previously undiscovered neutron stars. Undirected CW searches test an extremely large parameter space and are typically run on volunteered distributed computing resources. Results from the latest undirected CW search using LIGO data are shown in




Figure 3.6. Upper limits on $h_{r_{SS}}$ from continuous wave sources. The blue points are limits derived from the first year of data from the S5 science run, and red points are derived from the full two-year dataset. Figure from [2].

Fig. 3.6, from [2]. To generate this plot some assumption has been made about the ellipticity of the source. A related plot in the same paper sets lower limits on the distance to CW sources with some assumptions on ellipticity, and the spin down is assumed to be entirely due to gravitational radiation. The best limits are around 3 kpc for sources with frequency around 150 Hz.

Searches for unmodeled, short duration gravitational wave signals, or *bursts*, present the best chance for entirely unexpected discoveries. Limits from searches of this kind are shown in Fig. 3.7, from [15]. Here, we plot upper limits on the rate of a hypothetical population of sources emitting $M_{\odot}c^2$ in short-duration gravitational wave bursts. The limits extend below $10^{-6} \text{ Mpc}^{-3} \text{ year}^{-1}$, indicating that there is no population of extremely loud events in the nearby universe. A related figure from the same paper plots the limits on the energy released to gravitational waves by short-duration, narrowband sources, within a distance of 10 kpc (essentially covering the Milky Way galaxy). The limits extend below $10^{-6} M_{\odot}c^2$, indicating that there is no galactic source of quiet transient signals. Note that the limits on the rate of hypothetical sources that are isotropically distributed in the universe will improve by a factor of 10^3 in the advanced detector era. Advanced LIGO could detect sources that occur once per cubic gigaparsec per year, emitting $M_{\odot}c^2$ around 200 Hz.

Finally, the initial detectors placed upper limits on the energy density of gravitational waves, Ω_{GW} . This is shown in Fig. 3.8, as a function of frequency and in terms of the fraction of the critical energy density to close the universe. The initial detectors improved upon the upper limit on Ω_{GW} that is derived from Big Bang Nucleosynthesis (BBN), and the advanced detectors could probe exotic models for inflation and cosmic strings. See [5, 9] for the latest results.




Figure 3.7. Upper limits on the rate of generic gravitational wave sources emitting $M_{\odot}c^2$ of energy isotropically in short-duration, narrowband gravitational wave bursts. The limits degrade as the frequency of the hypothesized source increases due to the decreasing sensitivity of the detectors and from the factor of f_0^2 in Eq. 2.46. Figure from [15].




Figure 3.8. Limits on stochastic gravitational wave energy as a fraction of ρ_{critical} , the energy density required to close the universe: $\Omega_{GW} = \frac{1}{\rho_{\text{critical}}} \frac{d\rho_{GW}}{d \ln f}$. Figure from [9].

There are a variety of proposed techniques to detect gravitational waves besides the terrestrial Michelson interferometers. LISA, a proposed space-based gravitational wave antenna with freely-falling test masses separated by one million kilometers [34], would be sensitive to gravitational waves in the millihertz frequency range. There are known galactic binary systems that radiate gravitational waves in this band with sufficient amplitude to be detectable by LISA. Other source of low-frequency gravitational waves are supermassive black-hole binaries at the center of galaxies, and gravitational waves emitted from the inflationary period immediately following the Big Bang. Even lower-frequency gravitational waves could be detected using the signals from millisecond pulsars, treating the Earth and the pulsar as opposite ends of a long interferometer arm in space [53]. An array of stable millisecond pulsars would be sensitive to a low-frequency stochastic background of gravitational waves due to supermassive black hole binaries.

3.5 Can Interferometers Detect Metric Perturbations?

Gravitational wave scientists are often confronted with the question: “In a Michelson interferometer, shouldn’t the laser wavelength be stretched along with the length of the arm, and the gravitational wave signal will cancel?”

The answer is, of course, no. Interferometers do not use the wavelength of the laser as a ruler to measure the arm length. Rather, the phase evolution of the laser field as it propagates down the arm is used as a clock to measure the time-of-flight between the beamsplitter and the end test masses. The Fabry-Perot arm cavities are used to amplify this phase shift through many bounces between the input test mass and end test mass. The laser field in the arm cavities samples the time-of-flight down the arm many times, and the reflected field from the arm cavities carries the information of the built-up phase shift. The relative phase shift between the arms is measured by the interference pattern at the antisymmetric port.

As mentioned previously, if the period of the gravitational wave matches the time-of-flight down the arm, the end test mass will be in the same position relative to the beamsplitter when the laser field began the trip down the arm, and the phase shift will be zero. (Also the phase buildup of a Fabry-Perot cavity is reduced at frequencies above the cavity pole; this is why the LIGO interferometers are less sensitive at high frequency.) But, in the low frequency limit, where the response of the arm cavities is flat and the time variation of the gravitational wave is slow compared to the time of light in the arms, the interferometer will act as a transducer for strain.

CHAPTER 4

OPTICAL RESONATORS

*Thy glass will show thee how thy beauties wear;
Thy dial how thy precious minutes waste;
The vacant leaves thy mind's imprint will bear,
And of this book, this learning mayst thou taste.*

An *optical resonator*, also called an optical cavity or a Fabry-Perot cavity, is an arrangement of partially-transmissive mirrors that permit standing electromagnetic waves to form between them. This condition is known as *resonance*. For an incident electromagnetic wave with wavelength λ , an optical cavity of length L will be resonant when $L = N\lambda/2$, where N is some integer. When this condition is met, the small amount of light that leaks through the first mirror of the cavity will build into a standing wave between the mirrors. The power in the standing wave will grow until the losses of the cavity – due to transmission through the mirrors, absorption, or scattering – are equal to the input power through the first mirror. This is a magical effect: place one mirror in front of a laser and the light is reflected, but place two mirrors the correct distance apart and they become transmissive!

Optical cavities are one of the fundamental technologies used in modern gravitational wave interferometers, and a detailed understanding of their response is crucial for understanding the operation of the detectors. In this section I will describe simple optical cavities and their formalism, and derive the Pound-Drever-Hall locking technique for controlling optical cavities on resonance. In the second half of the chapter these foundations are used to outline a method for the precise characterization of optical cavity length and linewidth. This method has been used to measure the properties for two of the optical cavities of the advanced LIGO H1 instrument to high precision.

4.1 Optical Cavities

In this section we will derive the basic formalism for an optical cavity of two mirrors. The description is easily generalized to cavities with a larger number of optical components.




Figure 4.1. Fields around an optical cavity. Figure by J. Adamopoulos.

Consider a pair of mirrors, each constructed of a transparent dielectric substrate with negligible thickness, with a concave surface of spherical radius ρ . The mirrors are separated by a length L ; this is called the *one-way cavity length*. The curved surface has an optical coating that is partially reflective with a power reflectivity given by $R = r^2$, where r is the field reflectivity. The transmissivity is given by $T = 1 - R = t^2$. In a complete treatment of optical cavities, the laser field is written as an electromagnetic field with a Gaussian profile, which is focused by the concave surface of the mirrors. While Gaussian beam propagation methods are crucial for a proper understanding of the interactions between lasers and mirrors, we will neglect this for now, and model the beams as plane waves. For an introduction to Gaussian beams, the reader is encouraged to study Kogelnik and Li [81].

Our incident electric field is given by:

$$E_{in}(x, t) = E_0 e^{i\omega x/c - i\omega t}, \quad (4.1)$$

where $\omega = ck$ is the laser angular frequency and $\lambda = 2\pi/k$ is the laser wavelength. For the remainder of this section we will neglect the time-dependence of the laser field. We label the first mirror the beam encounters the *input mirror*, and the second as the *end mirror*; these mirrors have amplitude reflectivities r_i and r_e , and corresponding transmissivities. For simplicity we assume the mirrors are lossless, i.e. $r_i^2 + t_i^2 = 1$.

We define the circulating electric field as the field that is incident on the end mirror:

$$E_{circ} = E_0 t_i e^{i\omega L/c} + E_{circ} r_i r_e e^{2i\omega L/c}. \quad (4.2)$$

Rearranging, we have:

$$E_{circ} = E_0 \frac{t_i e^{i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}}. \quad (4.3)$$

The reflected field from the cavity is given by the prompt reflection plus the backward-going transmission of E_{circ} :

$$E_{refl} = -r_i E_0 + t_i r_e e^{i\omega L/c} E_{circ} = E_0 \left[-r_i + \frac{t_i^2 r_e e^{2i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}} \right] = E_0 \frac{(r_e e^{2i\omega L/c} - r_i)}{1 - r_i r_e e^{2i\omega L/c}}, \quad (4.4)$$

where we have used the convention from the previous chapter, in which reflections off the back surface of mirrors acquire a phase shift of π (again, this is necessary for energy conservation). The transmitted field is given by:

$$E_{trans} = t_e E_{circ} = E_0 \frac{t_e t_i e^{i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}}. \quad (4.5)$$

Finally, note that the circulating field in the cavity can be expressed in terms of the cavity field gain g_c :

$$g_c = \frac{E_{circ}}{E_0} = \frac{t_i e^{i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}}. \quad (4.6)$$

The square of the cavity field gain, g_c^2 , determines the power buildup inside the cavity. On resonance ($\omega L/c = 1$) this can be large, especially for cavities with mirror reflectivities very close to unity. These transmission and reflection formulae are the basic rules for modeling the behavior of electromagnetic fields around optical cavities.

The reflectivity and transmissivity of optical cavities are functions of the mirror reflectivities and transmissivities, the laser frequency, and the cavity length. If the reflectivities of the input and end mirrors are the same, the cavity is said to be *critically coupled*, and on resonance all of the incident laser power will be transmitted, except for losses within the cavity. The two cavities which are mentioned in greatest detail in this thesis, the advanced LIGO input mode cleaner (IMC) and output mode cleaner (OMC), are both critically coupled cavities. Optical cavities with $r_i < r_e$ are called *overcoupled*, while cavities with $r_i > r_e$ are *undercoupled*. We will forgo a discussion of the subtleties between these types of cavities, but note that the LIGO arm cavities are overcoupled. In this configuration the reflected light from the cavity is dominated by the field resonating in the cavity. This field has a phase shift of π relative to the promptly reflected field (see Fig. 4.2), so the sign of the reflected field changes when the cavity is on resonance. This sign flip will be revisited in our discussion of coupled cavities in Chapter 5.

A cavity is said to be *resonant* when the length L is equal an integer number of half-wavelengths of the laser. This condition can be achieved by either actuating on the mirror positions to control the cavity length, or by controlling the laser frequency. When a cavity/laser pair is held at the resonant condition, the cavity (or laser) is said to be *locked*.

The phase shift acquired by the laser as it traverses the length of the cavity is given by $\phi = \omega L/c = 4\pi f L/c$, where f is the laser frequency. At resonance, ϕ must equal an integer multiple of 2π . If ϕ is not equal to $N \times 2\pi$ the cavity is said to be *detuned* from resonance. In Fig. 4.2 we plot the power reflected and transmitted




Figure 4.2. Power and phase of the transmitted and reflected beams from an optical cavity. The cavity finesse is 14, and the detuning phase $\phi_f = 4\pi fL/c = 2\pi f/f_{FSR}$.

by a critically-coupled cavity, as a function of the detuning phase. Also we show the phase of the reflected and transmitted field. Note the sign flip of the reflected field on resonance (red trace).

The spacing between successive resonances is called the *free spectral range* of the cavity, or *FSR*:

$$f_{FSR} = \frac{c}{2L}. \quad (4.7)$$

The circulating power in the cavity as a function of laser frequency (or cavity length) has the form of a Lorentzian function of L and ω . The cavity *linewidth* is the full width at half maximum (FWHM) of the Lorentzian that describes the cavity. In terms of frequency this is known as the *cavity pole*. The cavity *finesse*, denoted by \mathcal{F} , is a dimensionless measure of the narrowness of the cavity linewidth. Finesse is defined as the FSR divided by the FWHM and has the functional form:

$$\mathcal{F} = \frac{FSR}{FWHM} = \frac{\pi}{2 \arcsin\left(\frac{1-r_i r_e}{2\sqrt{r_i r_e}}\right)} \simeq \frac{\pi\sqrt{r_i r_e}}{1 - r_i r_e}. \quad (4.8)$$

In terms of the finesse, the cavity pole is given by:

$$f_{\text{pole}} = \frac{FSR}{2\mathcal{F}} = \frac{c}{4\mathcal{F}L}, \quad (4.9)$$

and the cavity storage time is given by:

$$\tau_s = \frac{1}{2\pi f_{\text{pole}}} = \frac{2\mathcal{F}L}{\pi c}. \quad (4.10)$$

The FSR and pole frequency determine the dynamics of the optical cavity. For a cavity very near resonance, the transmission of power fluctuations on the input beam will be suppressed above the pole frequency. The same is true for frequency fluctuations. The cavity response time is finite, and only fluctuations that happen slowly compared to the cavity storage time will be perfectly transmitted. For this reason, and because optical cavities store light, optical cavities are sometimes called “capacitors for light”. This is something of a misnomer: a capacitor in series will suppress DC signals, while an optical cavity will suppress AC signals above the cavity pole.

Light reflected from a detuned cavity will acquire a phase shift proportional to the amount of detuning. If we imagine the detuning is due to fluctuations in the cavity length L at frequency f , the phase gain is given by:

$$g_\phi(f) = \frac{2\mathcal{F}}{\pi} \left(1 + i \frac{f}{f_{\text{pole}}} \right)^{-1}. \quad (4.11)$$

Note that for $f \ll f_{\text{pole}}$ the phase gain is given by $g_\phi \simeq \frac{c}{L} \tau_s$, or one-half the number of cavity round trips. Recall how LIGO uses Fabry-Perot cavities in the arms to build up the phase shift due to a passing gravitational wave; it is exactly due to this effect, in which the phase shift acquired by the laser in a cavity that is slightly off resonance builds up through multiple passes of the cavity.

In order to maintain a cavity in a resonant state, we use servo control, with an *error signal* derived from some measure of how detuned the cavity is from resonance. This error signal is filtered to generate an appropriate *control signal*, which is applied through appropriate actuators to the laser frequency (or the cavity length, or both) to maintain the resonant condition.

One method for generating an error signal is to employ lock-in detection. In this technique an oscillatory signal of known frequency is applied to the length of the cavity, typically by actuating on one of the mirrors. This dither signal generates a fixed-frequency modulation of the transmitted light. Low-frequency modulations of the transmitted light, generated by slight offsets of the cavity from the resonant condition, beat against the dither signal and generate sidebands. Demodulating the transmitted light around the dither frequency recovers these low-frequency sidebands and provides an error signal that is sensitive to the derivative of the cavity resonance curve. The LIGO output mode cleaner is locked using dither sensing for both the length and alignment degrees of freedom.

There are several drawbacks to the dither technique for cavity locking. The bandwidth of the error signal is typically limited to around 10 Hz or less (due to sensing noise when measuring the dither signal amplitude on the transmitted light) and robust control out to 100 Hz is very challenging. Dithering requires actuation on




Figure 4.3. Different optical cavity topologies encountered in this thesis. A two-mirror cavity (a), such as the LIGO arm cavities; a three-mirror triangular cavity (b), such as the Input Mode Cleaner; and (c), a four-mirror ‘bowtie’ cavity, similar to the Output Mode Cleaner.

the mirrors at high frequency, and larger mirrors require large actuation forces. In general dither sensing is very useful for small optical cavities with fixed mirrors and fast actuators but is impractical for large suspended cavities such as the LIGO arms.

We note from Fig. 4.2 that the phase of the light reflected and transmitted by the cavity is a sensitive measure of how detuned the cavity is from resonance. If we could sense this phase, locking an optical cavity would be easy. For optical-frequency electromagnetic fields, where the phase oscillates at 10^{14} Hz, measuring the phase directly is infeasible. But, by using a local oscillator field, we can construct a signal that is sensitive to the phase. This is described in the next section.

4.2 Pound-Drever-Hall Locking

The Pound-Drever-Hall (PDH) technique [55] utilizes a phase-modulated reference field to generate an error signal proportional to the phase of the reflected light. The phase modulation is performed at high frequency, much higher than the cavity pole frequency. For cavities that fit inside a typical lab this implies radio-frequency modulation. Modulation at these high frequencies can be generated using an electro-optical modulator (EOM), a crystal with a birefringence that varies as a function of applied electric field. By varying the electric field across the EOM crystal with a frequency f_{RF} , the incident electric field acquires an oscillating phase shift. This is written as:

$$E_{inc} = E_0 e^{i\omega_0 t} e^{i\Gamma \sin(\omega_{RF} t)}, \quad (4.12)$$

where Γ is the *modulation depth*, the amplitude of the phase modulation. We can use the Jacobi-Anger identity to expand the phase modulation in terms of Bessel functions:

$$e^{i\Gamma \sin(\omega_{RF} t)} = \sum_{n=-\infty}^{\infty} J_n(\Gamma) e^{in\omega_{RF} t}. \quad (4.13)$$

Keeping only the first-order terms we rewrite the incident electric field as:

$$E_{inc} = E_0 \left[J_0(\Gamma)e^{i\omega_0 t} + J_1(\Gamma)e^{i(\omega_0 + \omega_{RF})t} - J_1(\Gamma)e^{i(\omega_0 - \omega_{RF})t} \right]. \quad (4.14)$$

The first term is the *carrier* field, with frequency ω_0 , and the second and third terms are the radio-frequency *sidebands* with frequency $\omega_0 \pm \omega_{RF}$. The amplitude of the sidebands is usually a small fraction of the input field amplitude.

Upon reflection or transmission of the cavity, the amplitude of the field at each frequency will depend on the reflection and transmission formulae that we derived in the previous section. For convenience, we can express these formulae as:

$$\begin{aligned} E_{refl} &= r(\omega)E_0e^{i\omega t}, \quad r(\omega) = \frac{(r_e - r_i)e^{2i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}} \\ E_{trans} &= t(\omega)E_0e^{i\omega t}, \quad t(\omega) = \frac{t_e t_i e^{i\omega L/c}}{1 - r_i r_e e^{2i\omega L/c}}, \end{aligned} \quad (4.15)$$

where $r(\omega)$ and $t(\omega)$ are referred to as the reflection and transmission coefficients. The reflected field with the sidebands is given by:

$$E_{refl} = E_0 \left[r(\omega_0)J_0(\Gamma)e^{i\omega_0 t} + r(\omega_0 + \omega_{RF})J_1(\Gamma)e^{i(\omega_0 + \omega_{RF})t} - r(\omega_0 - \omega_{RF})J_1(\Gamma)e^{i(\omega_0 - \omega_{RF})t} \right]. \quad (4.16)$$

For brevity we shall refer to $r(\omega_0 \pm \omega_{RF})$ as $r(\pm\omega_{RF})$, but the relative nature of the sideband frequency to the carrier is always understood.

The power measured by a photodiode in reflection of the cavity is given by $P_{refl} = E_{refl}^* E_{refl}$. Since the Bessel functions are real for real inputs this will be given by:

$$\begin{aligned} P_{refl} &= |E_0|^2 \left[r^*(\omega_0)r(\omega_0)J_0^2(\Gamma) + r^*(\omega_{RF})r(\omega_{RF})J_1^2(\Gamma) + r^*(-\omega_{RF})r(-\omega_{RF})J_1^2(\Gamma) + \right. \\ &\quad + r^*(\omega_0)r(\omega_{RF})J_0(\Gamma)J_1(\Gamma)e^{i\omega_{RF}t} - r^*(\omega_0)r(-\omega_{RF})J_0(\Gamma)J_1(\Gamma)e^{-i\omega_{RF}t} \\ &\quad \left. + r(\omega_0)r^*(\omega_{RF})J_0(\Gamma)J_1(\Gamma)e^{-i\omega_{RF}t} - r(\omega_0)r^*(-\omega_{RF})J_0(\Gamma)J_1(\Gamma)e^{i\omega_{RF}t} + \mathcal{O}(2\omega_{RF}) + \dots \right]. \end{aligned} \quad (4.17)$$

The first line of this expression is a DC term with no frequency content. The second and third lines are oscillatory signals with frequency $\pm\omega_{RF}$ – these are the beat notes between the RF sidebands and the carrier. Beat notes at larger multiples of ω_{RF} will be suppressed like $J_1(\Gamma)$ and are neglected. The amplitude of the beat notes is determined by the reflection coefficients $r(\omega_0)$ and $r(\pm\omega_{RF})$, and we can measure their amplitude using a photodiode with RF sensitivity.

Keeping only terms with frequency ω_{RF} and rearranging, we have:

$$P_{refl}(\omega_{RF}) = P_0 J_0(\Gamma) J_1(\Gamma) \left[(r^*(\omega_0)r(\omega_{RF}) - r(\omega_0)r^*(-\omega_{RF})) e^{i\omega_{RF}t} + (r(\omega_0)r^*(\omega_{RF}) - r^*(\omega_0)r(-\omega_{RF})) e^{-i\omega_{RF}t} \right]. \quad (4.18)$$

Note the symmetry between the terms. If we use the shorthand expressions

$$\begin{aligned} A &= r^*(\omega_0)r(\omega_{RF}) \\ B &= r^*(\omega_0)r(-\omega_{RF}) \end{aligned} \quad (4.19)$$

and expand the exponentials, we can collect cosine and sine terms with frequency ω_{RF} and have:

$$P_{refl}(\omega_{RF}) = P_0 J_0(\Gamma) J_1(\Gamma) \left[[(A + A^*) - (B + B^*)] \cos \omega_{RF}t + i[(A - A^*) + (B - B^*)] \sin \omega_{RF}t \right]. \quad (4.20)$$

Recall that for a complex number $z = a + ib$, we can write $z + z^* = 2a = 2\text{Re}[z]$, and $z - z^* = 2ib = 2i\text{Im}[z]$.

Using this, we have:

$$P_{refl}(\omega_{RF}) = 2P_0 J_0(\Gamma) J_1(\Gamma) \left(\text{Re}[A - B] \cos \omega_{RF}t - \text{Im}[A + B] \sin \omega_{RF}t \right). \quad (4.21)$$

We make one final manipulation: since $\text{Re}[z] = \text{Re}[z^*]$ and $\text{Im}[z] = -\text{Im}[z^*]$ we use the conjugate of A to make the sign of the two expressions the same. Restoring the original expressions for A^* and B , we have:

$$P_{refl}(\omega_{RF}) = 2P_0 J_0(\Gamma) J_1(\Gamma) \left[\text{Re}[r(\omega_0)r^*(\omega_{RF}) - r^*(\omega_0)r(-\omega_{RF})] \cos \omega_{RF}t + \text{Im}[r(\omega_0)r^*(\omega_{RF}) - r^*(\omega_0)r(-\omega_{RF})] \sin \omega_{RF}t \right]. \quad (4.22)$$

The term that oscillates with $\sin \omega_{RF}t$ is called the *in-phase*, or *I-phase* component, because it is in phase with the original modulation from the EOM. The term that goes like $\cos \omega_{RF}t$ is called the *quadrature phase* or *Q-phase* component. Demodulation of P_{refl} at ω_{RF} will recover the coefficients of these two oscillatory signals.

A comparison of the I-phase and Q-phase signals is in Fig. 4.4, and a zoom of the region around the resonance is given in Fig. 4.5. The I-phase term has a nearly linear response around the cavity resonance in the region where




Figure 4.4. Demodulated signals for PDH in reflection. The sideband frequency is 15% of the cavity FSR. The cavity finesse is 156, and the detuning phase $\phi_f = 2\pi f/f_{FSR}$.

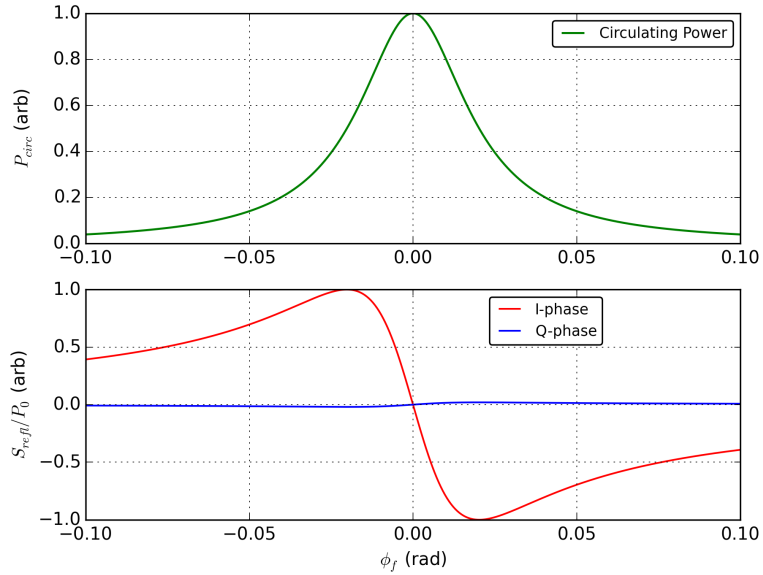


Figure 4.5. Zoom of the linear region of the PDH error signal.

the cavity buildup is greater than about 70% of the total. By normalizing the I-phase signal by the circulating power the linear region can be expanded [57]. The I-phase signal is referred to as the *PDH error signal*, and within the linear region it provides a precise measure of the cavity detuning from resonance. The cavity can be locked by feeding back the PDH error signal to an actuator that changes either the cavity length or the laser frequency. PDH locking has become a standard technique in metrology and optical engineering. Among many other applications it is used to provide exquisitely precise frequency references for spectroscopy.

4.3 Precision Measurements of Cavity Length and Linewidth

The LIGO interferometers use the PDH technique to control most of the optical cavities in the detector. The input beam to the interferometer is passed through a *input mode cleaner* (IMC), a short, high-finesse optical cavity that rejects higher order modes of the laser field and provides an initial stage of frequency stabilization. The RF sidebands used for the PDH error signals are generated before the IMC. These sidebands must be precisely resonant in the IMC along with the carrier so the error signals they generate are not polluted by length noise in the IMC. In this section, we present a method for precision measurements of an optical cavity's length and linewidth. Measurements of this kind are important step for characterizing the subsystems of a gravitational wave detector, for assessing the mirror reflectivity and absorption, and for checking the width of the linear region of PDH error signal. In particular the method was used to accurately measure the IMC length for H1, and verify

There are a variety of techniques for measuring the cavity pole frequency and the mirror reflectivity and absorption [76, 96]. The method presented here is unique in that it can be performed easily in-situ using the transmitted light from an optical cavity and with the typical RF modulation techniques that are used in modern optical experiments. As part of validating the method, we performed a careful analysis of errors, and demonstrated the technique with measurements of two of the optical cavities in the H1 instrument. We found that our measurement of the cavity pole frequency was limited by uncertainties in our tuned RF electronics, and a simpler experimental design could potentially measure the cavity pole with precision better than 1000ppm. Our cavity length measurement was precise enough to distinguish different fringes of the input mode cleaner (i.e. different multiples of $\lambda/2$).

To outline the technique, consider the transmission coefficient given in Eq. 4.15. In the previous section, we derived the PDH error signal using in-phase component of the reflected light, for an RF sideband that was not resonant in the cavity. Here, we use the I-phase and Q-phase signals of the transmitted light, for an RF sideband that is resonant in the cavity, i.e. that is a multiple of the cavity FSR. The expression for the power of the RF beat notes in transmission is given by:

$$P_{trans}(\omega_{RF}) = 2P_{in} J_0(\Gamma)J_1(\Gamma) \left[\text{Re}[t(\omega_0)t^*(\omega_{RF}) - t^*(\omega_0)t(-\omega_{RF})] \cos \omega_{RF}t \right. \\ \left. + \text{Im}[t(\omega_0)t^*(\omega_{RF}) - t^*(\omega_0)t(-\omega_{RF})] \sin \omega_{RF}t \right]. \quad (4.23)$$

Recall that the transmission coefficient is given by Eq. 4.15. If we demodulate P_{trans} at ω_{RF} , we will get in-phase and quad-phase signals in exactly the same way as in the PDH technique. In terms of $t(\omega)$, these in-phase and quad-phase signals are given by:

$$\begin{aligned} \text{In phase : } & \text{Re} \left[t(\omega_0)t^*(\omega_{RF}) - t^*(\omega_0)t(-\omega_{RF}) \right] \\ \text{Quad phase : } & \text{Im} \left[t(\omega_0)t^*(\omega_{RF}) - t^*(\omega_0)t(-\omega_{RF}) \right]. \end{aligned} \quad (4.24)$$

Consider three possible cases:

- **The carrier is resonant** ($\omega_0 = N * \omega_{FSR}$). In this case $t(\omega_0)$ is purely real and $t^*(\omega_{RF}) = t(-\omega_{RF})$. The terms cancel, and both the I- and Q-phase signals are zero.
- **The carrier is slightly off-resonance** ($\omega_0 = N * \omega_{FSR} + \delta\omega$), **but the RF sidebands are resonant** ($\omega_{RF} = N * \omega_{FSR}$). Here, $t(\pm\omega_{RF}) = t(\omega_0)$. Again the terms cancel and the demodulated signals vanish.
- **The carrier is slightly off-resonance and the RF sidebands are slightly off-resonance.** In this case there is a phase mismatch between the upper and lower sidebands and the I and Q signals in transmission will be nonzero.

For the third case, the signals we expect to observe are shown in Fig. 4.6. In this figure, we model the response of a 16.5 meter cavity with finesse 156. The carrier laser wavelength is 1064×10^{-9} meters and the cavity is offset from resonance by 10 picometers. The FSR for the cavity is $f_{FSR} \simeq 9.1$ MHz and the RF sidebands are scanned between 8.95 MHz and 9.25 MHz. (Note that the demodulation of P_{trans} must be done at f_{RF} , which is changing with the scan.)

We note two things from Fig. 4.6. First, when $\phi_f = 0$ ($f_{RF} = f_{FSR}$, corresponding to the second case above), the I and Q signals are zero. This is how we will measure the cavity FSR and, hence, the cavity length.

Second, the Q-phase signal has secondary zero crossings. We will call the frequencies for these points by f_{\pm} . These secondary zero crossings carry information about the cavity linewidth. To see this, we need to rewrite the expression for the PDH signals in transmission.




Figure 4.6. I- and Q-phase signals on transmission, as a function of the RF sideband detuning phase $\phi_f = 2\pi f_{RF}/f_{FSR}$, with the carrier frequency held slightly off-resonance. The cavity finesse is 156. Note the zero crossings in the Q-phase signal.

Starting with Eq. 4.23, let us call R is the round-trip reflectivity of the cavity, $R = r_i r_e$. If we use the following approximations,

$$R = r_e r_i, \quad t(\phi) = \frac{g_{tr}(1-R)}{1-Re^{i\phi}}, \quad J_0(\Gamma) \approx 1, \quad J_1(\Gamma) \approx \Gamma/2, \quad e^{i\phi_c} \approx 1 + i\phi_c, \quad (4.25)$$

then we get the following for the transmitted power:

$$P_{trans} = \phi_c g_{tr}^2 \Gamma P_{in} R e^{i\phi_f} \frac{1 - e^{i\phi_f}}{(e^{i\phi_f} - R)^2}, \quad (4.26)$$

where $\phi_c = 2\pi\delta f/f_{FSR}$ is the carrier detuning phase from resonance. We can expand this expression to find explicit terms for the I and Q-phase signals. By multiplying top and bottom by the complex conjugate of the denominator, $(e^{-i\phi_f} - R)^2$, we get:

$$\begin{aligned} P_{trans} &= \phi_c g_{tr}^2 \Gamma P_{in} R e^{i\phi_f} (1 - e^{i\phi_f}) \frac{(e^{-i\phi_f} - R)^2}{(1 - 2R \cos \phi_f + R^2)^2} \\ &= \phi_c g_{tr}^2 \Gamma P_{in} \frac{-2R}{(1 - 2R \cos \phi_f + R^2)^2} \sin \frac{\phi_f}{2} \\ &\times \left[\sin \frac{\phi_f}{2} (1 + 2R - (1 + 2 \cos \phi_f) R^2) + i \cos \frac{\phi_f}{2} (1 - 2R - (1 - 2 \cos \phi_f) R^2) \right]. \quad (4.27) \end{aligned}$$

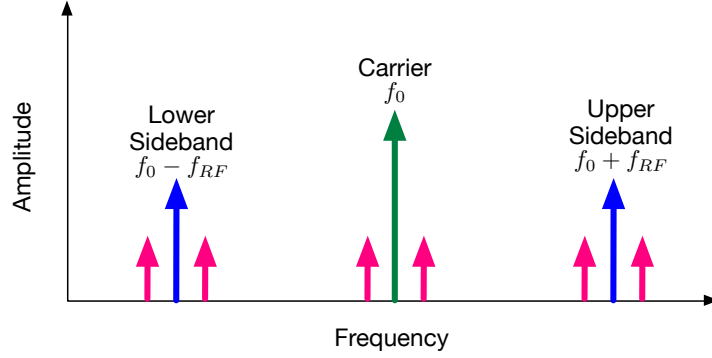


Figure 4.7. Double demodulation: a diagram of the carrier, RF sidebands and audio sidebands for the cavity length and linewidth measurement. Figure by A. Staley, from [125].

The Q-phase signal is the imaginary part. We can solve for the zero crossings of the Q-phase signal by setting the coefficient equal to zero,

$$1 - 2R - (1 - 2 \cos \phi_f)R^2 = 0. \quad (4.28)$$

The two solutions have the form $f_{\pm} = f_{\text{pole}}/\sqrt{R}$. Since $f_{\text{pole}} = f_{FSR}/2\mathcal{F}$, then $f_{\pm} = \frac{f_{FSR}}{2} \frac{1-R}{2R}$. We precisely measured f_{FSR} using the central zero crossing, and we can easily invert f_{\pm} to calculate R and from there, f_{pole} . Thus we have a complete analytical expression for measuring the cavity pole frequency using zero crossings of the Q-phase signal in transmission.

In practice, measurements of zero crossings are subject to various systematic uncertainties. In particular, measurements using demodulated signals are sensitive to relative phase shifts between the detected RF power and the local oscillator. Also there can be residual amplitude modulation (RAM) in the laser at the same frequency as the RF sidebands. These and other noise sources such as electronics offsets and timing uncertainties can bias measurements of the FSR and f_{\pm} .

To attain the best possible precision, we add another modulation frequency at a fraction of the cavity pole frequency. This audio modulation is denoted by f_a . It can be applied either at the laser frequency or in the cavity length. In our measurement, we dithered the laser frequency. The demodulated RF signals are demodulated again at f_a , which eliminates our sensitivity to DC electronics offsets and suppresses the effect of RAM. This *double-modulation* technique was first demonstrated by a Japanese group [33, 25, 24].

The sideband picture for the full experiment is illustrated in Fig. 4.7. The complete analytical expression for the set of four demodulation products (RF $\{I,Q\} \times$ AF $\{I,Q\}$) is given in Ref. [125], although the results are directly calculable from the transmission coefficient.

Our procedure for the measurement is the following: lock the optical cavity to the carrier frequency ω_0 and introduce an RF sideband, with frequency close to the cavity FSR. Introduce an audio sideband by dithering the

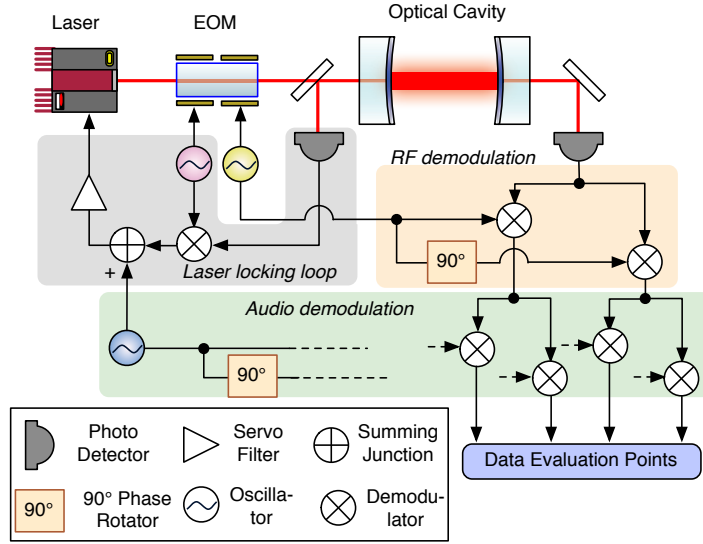


Figure 4.8. Experimental setup for the double-modulation measurement. Figure by K. Izumi, from [125].

laser, with a frequency much less than the cavity pole. Sweep the RF sidebands (f_{RF}) across a range slightly larger than the cavity linewidth. Demodulate the transmitted signal at the RF frequency, and then again at the AF frequency. The zero crossings of the various demodulation products will carry information about the cavity length and pole frequency.

In Fig. 4.8 we sketch the setup of the double-modulation measurement. We used this technique to characterize two of the optical cavities in the advanced LIGO detector, the 16 meter input mode cleaner and one of the 4 km arm cavities. Due to the length of the arm cavity measuring and demodulating the transmitted signals was not practical, so the measurement of the cavity FSR was made in reflection with $f_{RF} = 24.5$ MHz, or approximately $666 \times f_{FSR}$. While this prevented a measurement of the cavity pole (estimated to be 40 Hz), the large RF frequency provided an extremely precise measurement of cavity length.

For the measurement of the IMC cavity pole, the results were complicated by the resonant circuits used to generate and detect the RF sidebands. These resonant circuits can be modeled by Lorentzian functions of frequency, with central frequency f_{EOM} and width f_{EOM}/Q , for a quality factor Q defined by the circuit. The resonant elements lead to a frequency-dependent amplitude response and phase shift across the span of the RF demodulation. We estimated that the resonant circuit for the EOM was the dominant source of phase change, although there are also resonant circuits in the detection photodiode and the demodulation circuit. The presence of a frequency-dependent phase rotation significantly complicates the analytical solution for f_{\pm} , since any errors in the phase estimation will rotate the I-phase signal (maximal at f_{\pm}) into the Q-phase signal (zero at f_{\pm}) and bias the result.

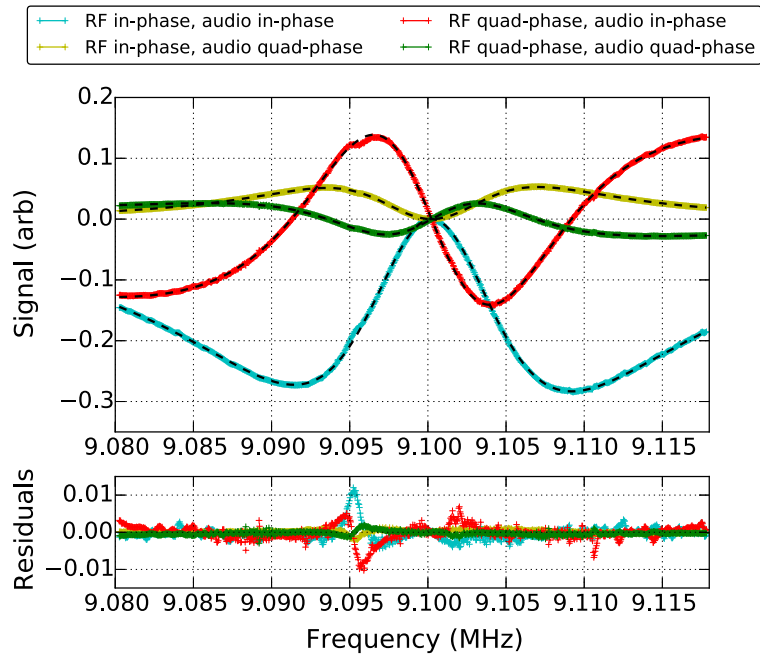


Figure 4.9. Experimental results for the scan of the 16 m input mode cleaner; the x-axis shows the frequency of the RF modulation. The audio modulation frequency for this dataset was 1 kHz. The best-fit parameters are given in Table 4.2. The feature around 9.095 MHz is most likely due to a mechanical resonance in the EOM.

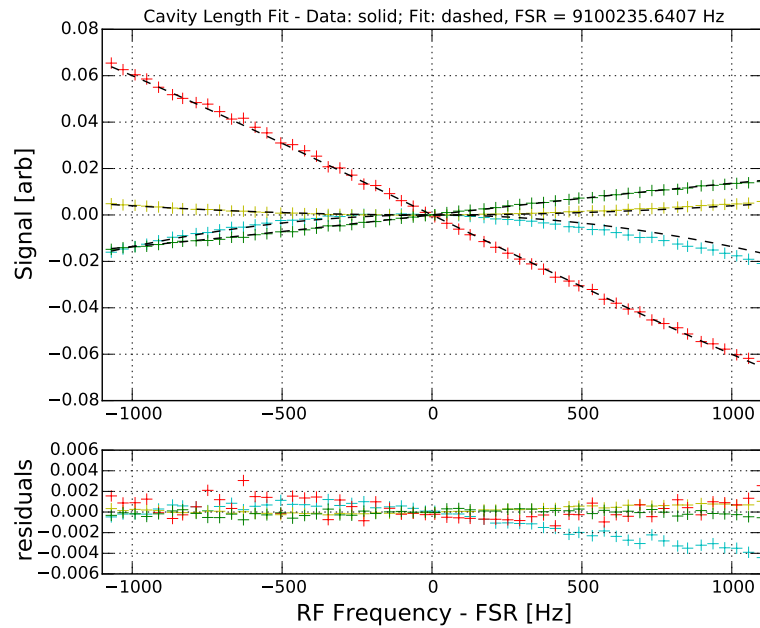


Figure 4.10. Zoom of the zero-crossing at the FSR frequency, for the IMC measurement at 9.1 MHz and 1 kHz. The best-fit FSR value has been subtracted from the x-coordinate for clarity. Note the precise null value for all four I,Q signals at the FSR.

Parameter	Best-Fit Value	Stat. Uncertainty
f_{FSR}	9 100 234 Hz	2 Hz
T (=1-R)	6068 ppm	3 ppm
f_{EOM}	9 113 000 Hz	120 Hz
Q_{EOM}	61.1	0.2
θ_0	54.0°	0.2°
ϕ_0	10.23°	0.08°

Table 4.1. Best fit parameters and 67% confidence level statistical uncertainties for the Input Mode cleaner data. The results presented here are for an audio modulation frequency of 1 kHz and an RF modulation of 9.1 MHz.

To account for this frequency-dependent response we performed a nonlinear least-squares fit to the full set of four demodulated I,Q-signal pairs, allowing for all demodulation phases to be simultaneously constrained. Our fit had seven free parameters: an overall amplitude coefficient (A), the free spectral range of the cavity (f_{FSR}), the round-trip reflectivity of the mirrors (R), the parameters of the RF resonant circuit (f_{EOM} and Q), and two static demodulation phase shifts (θ_0 and ϕ_0) for RF phase and audio phase respectively.

Data were collected using two RF modulation frequencies (9.1 MHz and 45.5 MHz) and three audio modulations (1 kHz, 303 Hz, and 103 Hz). Overall, the seven fit parameters were found to be consistent across the six data sets. Differences in the RF phase estimation were used to estimate our systematic errors. For the measurement of the cavity pole, we found that a linear frequency dependence in the RF demodulation phase was our dominant source of systematic error. Unfortunately the experimental setup was constrained by the available electronics, which were designed for gravitational wave detection and not for a measurement of this kind. By using electronics with flat frequency response, the systematic errors could likely be reduced to below the statistical errors.

The results for one of the six datasets collected for the IMC are shown in Table 4.1. The data and fit are compared in Fig. 4.9, and a zoom of the region around the FSR is shown in Fig. 4.10. Final results for the three physical parameters we sought to measure – IMC cavity length and cavity pole, and Y-arm cavity length – are given in Table 4.2. This table includes estimates of our systematic uncertainties. As part of Ref. [125] we derived analytic expressions for the dominant systematic effects that limit the precision of the technique. For our measurements of the IMC and the Y-arm, the potential bias from these effects was estimated using the parameters and noise characterization of our analysis. The use of double-modulation reduces most of the systematic effects, although the precision of the IMC cavity pole measurement was limited by the uncertainty in the resonant RF circuit parameters as described above.

For the Y-arm measurement, the precision was remarkable: a 4 km distance was measured to a precision of 70 parts per billion, taking the statistical and systematic error in quadrature. For this cavity we expect the ultimate

Parameter	Measured Value	Stat. Uncertainty	Sys. Uncertainty
Y-arm Length	3994.4692 m	0.2 mm	0.2 mm
IMC Length	16.471701 m	3 μ m	1 μ m
IMC f_{pole}	8806 Hz	10 Hz	52 Hz

Table 4.2. Parameters of the H1 aLIGO Input Mode Cleaner and Y-arm optical cavities, measured using the double-modulation technique.

precision could be as low as 1 ppb; in our case the precision was limited by a manufacturing error of the green coatings that reduced the cavity finesse.

For the aLIGO IMC, the design power transmissivity of the input and end mirrors is 6000 ± 200 ppm. From our combined measurements, the best-fit value for the round-trip reflectivity is 6068 ± 3 ppm, neglecting systematic uncertainties. The characterization of the IMC using this technique was an important confirmation that the mirrors were not absorbing or scattering more light than intended. This is an important safety check before a high-power beam can be delivered to the interferometer.

CHAPTER 5

ADVANCED LIGO

*We'll make an instrument of this;
Omit nothing may give us aid.*

The first generation LIGO instruments completed a series of observing runs in 2010. Following this data-taking, nearly all of the components were disassembled and removed to make way for the second-generation instruments. These are referred to as the *advanced LIGO* interferometers (“aLIGO” for short), and are designed to achieve a strain sensitivity ten times better than the initial detectors. Essentially the only components that were not modified as part of the upgrade were the components of the vacuum system, the facilities, and the humans involved in the project, along with their experience and knowledge from the initial detector era. A complete description of the aLIGO instruments is given in [7].

In this section I provide some details of the design, installation, and commissioning of the advanced LIGO interferometers, including a brief description of the lock acquisition process. A large focus is given to the Output Mode Cleaner (OMC) for the H1 instrument, which was installed in the spring and summer of 2014 and first used for sensing the gravitational strain channel in February 2015. I present the details of a measurement of the OMC cavity length noise and an estimation of its contribution to the H1 noise floor. Finally I discuss some results of using the OMC as an optical spectrum analyzer, in particular a measurement of the contrast defect of the H1 instrument.

5.1 The Advanced LIGO Interferometers

The aLIGO detectors employ several technological improvements compared to the first-generation detectors. The layout of the main optical elements is shown in Fig. 5.1. The major changes from initial LIGO are as follows:

- **Active seismic isolation** is applied in-vacuum, between the ground and the suspension point for the optics. Except for the external isolator systems used in initial LIGO, all previous seismic isolation was passive.

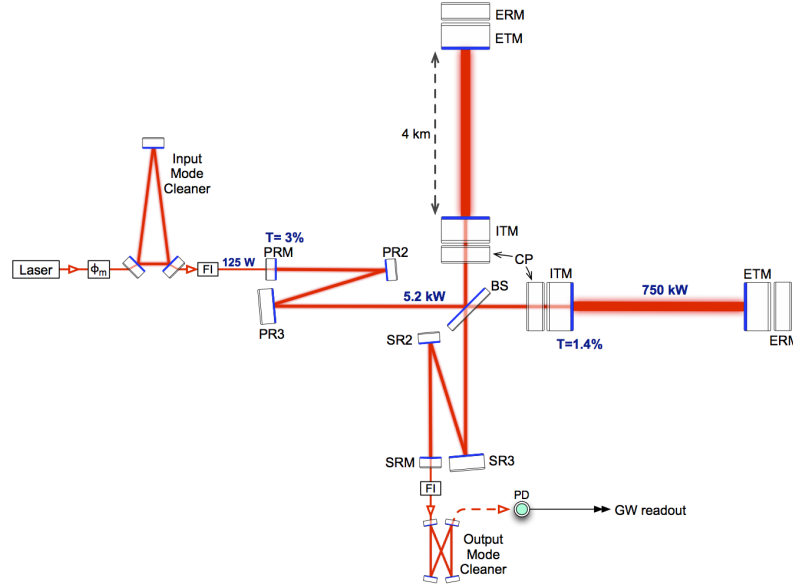


Figure 5.1. Schematic of the advanced LIGO interferometers. Figure from [7]. The Fabry-Perot arm cavities are formed between the *end test mass* (ETM) and *inner test mass* (ITM) mirrors. Power-recycling and signal-recycling cavities are formed between the ITMs and the *power recycling mirror* (PRM) and the *signal recycling mirror* (SRM), respectively. Additional large mirrors are used to align and shape the beam. All of the detector components in this figure after the phase modulation (ϕ_m) are housed in large vacuum chambers and are isolated from ground motion by multi-stage suspensions.

The new in-vacuum seismic tables have three stages of isolation (two for the small vacuum chambers) with each level supported by blade springs that provide $1/f^2$ isolation above the resonant frequencies, which are typically 1 Hz. Inertial seismometers and capacitive position sensors measure the position and acceleration of each stage. These signals are used as error signals to damp the resonant frequencies of the blade springs and provide additional isolation using magnetic actuators. The seismic isolation platforms for the large vacuum chambers are the size of a small car, in both volume and mass.

- **Multi-stage suspended optics** add several more orders of magnitude of isolation at frequencies above 1 Hz. For the test masses whose relative position is measured to detect the gravitational wave strain, the final stage of suspension is monolithic, i.e. the suspension wires and the fused-silica mirrors form a single welded body. This greatly decreases the mechanical loss within the final stage of the suspensions and limits the interaction of thermally-driven fluctuations with the longitudinal degree of freedom.
- **Higher-finesse arm cavities** and a 200 W laser reduce the sensing noise limit due to shot noise. In addition, **larger test masses** of 40 kg and 34 cm diameter reduce the motion imparted to the mirrors from variations in the number of photons; this radiation pressure noise is a limiting noise source at low frequencies. The




Figure 5.2. Coupled optical cavity topologies used in the advanced LIGO detectors. Power recycling (a) forms a coupled cavity between the Power Recycling Mirror (PRM) and the Fabry-Perot arm cavities. Intensity and frequency fluctuations from the input beam are suppressed by the coupled cavity pole, with $\omega_{cc} \simeq 1$ Hz. Signal recycling (b) forms a resonant cavity between the Signal Recycling Mirror (SRM) and the arms. Acoustic-frequency fluctuations generated by gravitational wave interactions with the arms (the ‘signal’) are resonant in this cavity, and the response of the interferometer is broadened.

larger-diameter mirrors also enable larger beam sizes on the optics. By sampling a larger area of the mirror surface, the displacement noise due to Brownian motion of the mirror coatings is made smaller.

- **Signal recycling**, also known as **signal extraction**, is a technique applied to Michelson interferometers with Fabry-Perot arm cavities to broaden the response of the detector [95, 91]. An additional mirror at the antisymmetric port, the *signal recycling mirror*, forms a coupled cavity with the arms and changes the effective finesse experienced by the acoustic sidebands generated by gravitational waves. This increases the optical gain of the interferometer at high frequencies. See Fig. 5.2.

5.2 Cavities & Control

The primary optics of the advanced LIGO detectors form four optical cavities: the power recycling cavity, the signal recycling cavity, the X-arm, and the Y-arm. Because the circulating power in the arm cavities is strongly coupled by the addition of the recycling cavities, the length degrees of freedom for the arms, given by L_x and L_y , are often expressed in a different basis that considers length changes that are common or different between the arms:

$$L_+ = \frac{L_x + L_y}{2} \quad , \quad L_- = \frac{L_x - L_y}{2} . \quad (5.1)$$




Figure 5.3. Photo of the author in the ETMY vacuum chamber, working on alignment of the green laser. I am kneeling, Keita Kawabe is standing.

Here, L_+ is called the *common arm* degree of freedom or CARM, and L_- is the *differential arm* degree of freedom, or DARM. Similarly, the lengths of the power recycling, signal recycling, and small Michelson degrees of freedom are referred to as PRCL, SRCL, and MICH.

In the resonant, low-noise state of the detectors, CARM, PRCL, SRCL, and MICH are sensed and controlled using PDH signals. These PDH signals are generated with RF sidebands at 9.1 MHz and 45.5 MHz. By design, the 9.1 MHz sideband is resonant in the PRC and the 45.5 MHz sideband is resonant in the SRC. There is a significant amount of cross-coupling in the length degrees of freedom, and the overall control scheme for the interferometer is quite complex. An analytical description of the control scheme for the initial LIGO detectors is given by Fritschel et al. [65], and preliminary studies of the advanced LIGO design were performed at the Caltech 40 m lab by Rob Ward [139] and others.

For advanced LIGO, the task of bringing four optical cavities into simultaneous resonance is a daunting technical challenge. The process by which the detectors transition from an uncontrolled, nonresonant state to a fully resonant, low-noise state appropriate for data collection is called *lock acquisition*. In initial LIGO, a so-called stochastic locking process was used in which the mirrors swung freely until by chance they were aligned in a resonant condition, at which time fast actuation was applied to stop their motion and control the cavity degrees of freedom.

Stochastic locking was deemed unsuitable for advanced LIGO for several reasons. Among these were the larger test masses and higher finesse of the arms, and the additional complexity introduced by the signal recycling cavity. Advanced LIGO utilizes a three-step acquisition process in which the arms are locked independently of the corner station optics (power recycling cavity, signal recycling cavity, and short-armed Michelson) using




Figure 5.4. Locking sequence of the advanced LIGO detectors. Figure from [126]. The top figure shows the power buildup of the 1064 nm light in the arms. In states I and II, the arms are locked using the 532 nm lasers. In states III and IV, the arms are moved off-resonance and the cavities of the DRMI are locked. In states V and VI the arms are brought back to resonance, this time with much larger power buildup due to the power recycling cavity. In state VII the detector is fully resonant.

532 nm green lasers that are injected from the endstations. This *Arm-Length Stabilization system*, or ALS, provides independent control over the arm cavities during the lock acquisition procedure. The performance of the ALS system and the procedure for bringing the interferometers to resonance is described in [126, 124, 87].

A timeline of lock acquisition with the ALS system is illustrated in Fig. 5.4. The following is a brief outline of the steps in the locking procedure:

1. The 4 km arm cavities are aligned and locked. The PDH signal for sensing the length of the cavity is generated using green lasers (532 nm wavelength) that are injected from the end stations. The green lasers are phase-locked to the main 1064 nm laser, which ensures that the arm lengths are maintained at a precise distance from the resonance condition for the 1064 nm laser. The PDH signal from each arm is used to lock the green laser frequency to the arm length, and the green light transmitted by the arms is beat against a frequency-doubled pickoff from the main laser. This beat note signal is used as an error signal to control the arm lengths at a precise offset from the resonance point of the main laser.
2. Once the arms are controlled relative to the main laser, they are brought off-resonance by several kHz to prevent interference with the corner station cavities. (During this process the green laser frequency remains locked to the arm length, and the 532 nm light is fully resonant in the arm cavities.) At this

stage the optics in the corner station are aligned and the cavities that form the Dual-Recycled Michelson Interferometer (DRMI) are locked. Note that if the arms are resonant for the 1064 nm main laser, the laser field experiences a sign flip upon reflection from the arm cavity. This sign flip would spoil control of the DRMI cavities and is the reason the arms are held off-resonance during DRMI acquisition. Following DRMI acquisition, the sensing of the corner station cavities is transitioned from the traditional PDH signals to a set of sideband harmonic frequencies which are insensitive to the state of the arm cavities. This ‘3F’ technique was first demonstrated by a Japanese group [29].

3. With the DRMI controlled by the 3F signals, the offset between the arm cavity length and the main laser resonance point is reduced. As the power buildup in the arms increases, control of the arm cavities is switched from the green PDH signal to various combinations of the transmitted light through the arms, which is properly filtered and normalized to provide a linear error signal. This procedure is called the *CARM offset reduction*. Once the 1064 nm light is fully resonant in the arms, control of the arm length is handed off to a PDH signal derived from the 9.1 MHz sideband. At this point, the lock acquisition process is complete.

With the arm cavities in their fully resonant state the detector can act as a transducer for gravitational waves. Due to the use of signal recycling, the sensitivity of the detector as a function of frequency to a gravitational signal differs from the expression we derived in Chapter 3. The signal recycling mirror forms a *coupled optical cavity* with the Fabry-Perot arms. This increases the frequency of the cavity pole in the interferometer response function (Eq. 3.10) and slightly reduces the response at low frequencies. The net effect is to improve the shot-noise limited sensitivity of the detector. The details of shot noise limited sensing for signal-recycled interferometers have been worked out by Buonanno and Chen [41], and a good introduction to the topic is given in Chapter 3 of [139].

In Fig. 5.5 we plot a recent noise curve from the H1 advanced LIGO interferometer. Included in the plot are a curve representing the aLIGO design sensitivity (for 125W input power, and higher SRM reflectivity) and the best performance of the H1 detector from the initial LIGO era. The sensitivity improvement from signal recycling is clearly visible. Both detectors have about the same power incident on the beamsplitter. The increased arm cavity finesse in the aLIGO instruments has improved the sensitivity around 100 Hz, and the response of the aLIGO instrument at high frequencies has improved due to signal recycling.

5.3 The Output Mode Cleaner

Gravitational wave signals will generate fluctuations in the differential arm length degree of freedom, L_{-} . In advanced LIGO, this length is measured using homodyne detection, often referred to as *DC readout*. Homodyne detection provides an improvement in the shot-noise-limited sensitivity and decouples the gravitational wave




Figure 5.5. Recent performance of the H1 advanced LIGO detector, compared to previous best sensitivity. At low frequencies, the improved seismic isolation and suspensions have reduced the control noise that couples to the gravitational wave channel. At 100 Hz the sensitivity improvement is due to the increased arm cavity finesse. At high frequencies the signal recycling cavity has broadened the response of the detector.

detection channel from many technical noise sources in the interferometer, such as oscillator noise from the RF sidebands. In homodyne detection, a small static differential arm offset (called the ‘DARM offset’) is applied to the arms to deliver a constant amount of carrier light at the antisymmetric port. A gravitational signal will generate small intensity fluctuations on top of this static carrier field. The intensity of the carrier light is detected with low-noise direct current photodiodes (DCPDs). Since any power fluctuation on the DCPDs will be interpreted as a signal, any laser fields sensitive to other interferometric degrees of freedom must be removed prior to detection. The *Output Mode Cleaner* (OMC) is used to filter the light at the antisymmetric port and reject optical fields that do not directly couple to the gravitational wave signal.

A plot of the carrier power detected at the AS port as a function of the DARM offset is shown in Fig. 5.6. For offsets of tens of picometers, the change in power at the AS port is well-modeled by a quadratic. From Chapter 3, recall Eq. 3.5:

$$P_{AS} = |E_{AS}^* E_{AS}| = P_{BS} \sin^2[kLh] \simeq P_{BS} (kx)^2. \quad (5.2)$$

Here we have written $Lh = x$ as the static DARM offset of $\mathcal{O}(10^{-12})$ meters. A gravitational wave strain will induce small fluctuations around this static offset. Typically the interferometer is operated with a DARM offset between 10 and 20 pm, and the residual DARM motion is of order 10^{-13} meters, at frequencies around 1 Hz. This residual motion is small enough that the response of P_{AS} is well approximated by a linear function, and gravitational wave perturbations of 10^{-19} m around 100 Hz are readily detectable. A detailed expression will be derived in the next section (Eq. 5.6).

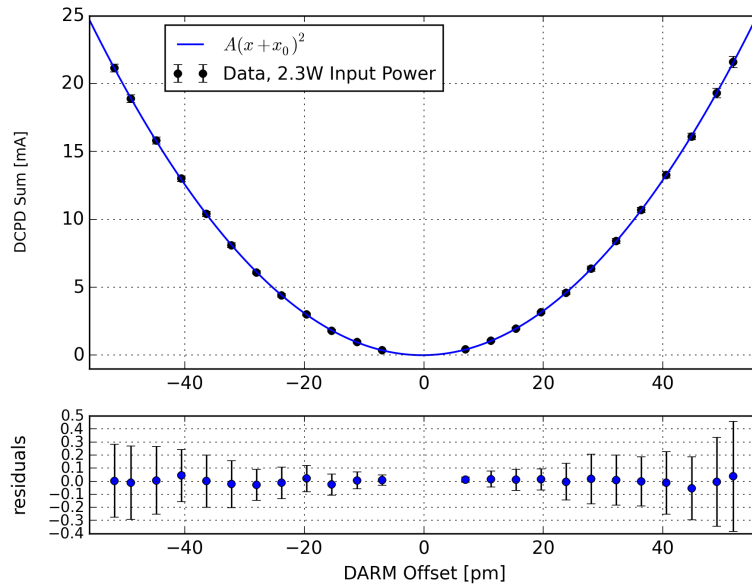


Figure 5.6. Sweep of the differential arm (DARM) offset with the full interferometer locked. Here, 2.3 W of power is incident on the input mode cleaner, and about 80 W is incident on the beamsplitter (due to the gain of the power recycling cavity). The power measured by the antisymmetric port DCPDs is a quadratic function of the DARM offset.




Figure 5.7. Design of the advanced LIGO Output Mode Cleaner, from [30].

A diagram of the OMC is shown in Fig. 5.7. The OMC is a four-mirror optical cavity in a ‘bowtie’ configuration with a cavity round-trip length of $2L = 1.132$ meters and finesse of 390. The cavity mirrors, alignment optics, and photodiodes are bonded to a fused-silica breadboard 45 cm long, 15 cm wide, and 4.175 cm thick. The beam size ($1/e^2$ radius) inside the cavity is about 0.5 mm. Two of the cavity mirrors are instrumented with PZT actuators which are used to control the cavity length and lock the OMC on the carrier field. The transmitted light intensity is measured using two low-noise DCPDs (arranged after a beamsplitter), and the position and angle of the input beam is measured using a small pick-off beam split between two quadrant photodiodes (QPDs). The advanced LIGO OMC was designed and built by Koji Arai, Zach Korth, and others.

To reduce path length variations between the OMC and the main optics of the interferometer, the OMC breadboard is suspended from steel wires in a two-stage suspension. The suspension resonance frequencies are around 1 Hz, and the coupling of ground motion to the OMC breadboard is suppressed by approximately a factor of 10^4 at 10 Hz. The OMC suspension is mounted on a seismic isolation platform inside of a vacuum chamber, along with a variety of steering mirrors and sensors. A schematic of the vacuum chamber is shown in Fig. 5.8, and Fig 5.9 is a photograph of the author working in the chamber.

The alignment of the beam incident on the OMC has four degrees of freedom, corresponding to the position and angle of the beam in the horizontal and vertical directions. We control the alignment of the input beam using three small suspended mirrors with electromagnetic actuators, called ‘tip-tilt’ mirrors. The position and angle of the input beam is measured using either the centering information from the OMC QPDs (assuming some prior information about the ‘good’ beam spot position on the QPD), or by constructing a lock-in amplifier using acoustic frequency dithering of the tip-tilt mirrors¹. In both cases, the error signals (from the QPD sensors, or the demodulated signals from the tip-tilt dither) are decomposed into a the cavity axis position and angle basis. This decomposition is described in [31].

Misalignments in the position or angle of the input beam will change the amount of transmitted light measured by the DCPDs. If the alignment of the input beam into the OMC is poor, then small fluctuations in the input beam alignment will linearly couple to the power transmitted by the cavity. If the input beam is well aligned to the cavity axis, the coupling between alignment fluctuations and the transmitted light will be quadratic. Low-noise operations require a well-aligned input beam to the OMC, in order to suppress the coupling of input beam jitter to the gravitational wave channel.

¹The OMC that was used during the S6 science run was aligned using a different method, the *beacon dither*, which maximized the signal of the ETM solid-body mode at 8 kHz in the transmitted light through the cavity. [122]




Figure 5.8. Arrangement of optics in the advanced LIGO HAM6 vacuum chamber. The optical table is approximately two meters wide. The beam from the signal recycling mirror enters from the left. Pick-off beams for monitoring the state of the detector leave the vacuum chamber at the top and are collected on an in-air optics bench. The rectangles, mostly at the bottom of the figure, are 10 kg ballast masses to balance the seismic isolation platform. Figure from [35].




Figure 5.9. Photo of the author working in the HAM6 chamber in August, 2014. The OMC suspension cage is the tall structure on my left.




Figure 5.10. When an optical cavity is locked at the resonance, small cavity length fluctuations (δL) couple quadratically to the transmitted power (δP). At the half-resonance point, length noise maximally couples to the transmitted power. The same relationship is true for fluctuations in the alignment of the input beam. When the cavity is well-aligned, power fluctuations in the transmitted light from input beam jitter are suppressed.

5.4 OMC Cavity Length Noise

Fluctuations in the OMC cavity length will also lead to variations in the transmitted power. The OMC cavity length is controlled using a lock-in amplifier signal generated by a 4.1 kHz dither, applied using one of the PZT actuators. The other PZT actuator is used to adjust the cavity length and maintain the resonance condition, using the error signal provided by the length dither. In principle, OMC cavity length noise cannot be distinguished from a gravitational wave signal, and could pollute the noise floor if it is not adequately suppressed.

There are a few methods to characterize the OMC length noise *in situ*. One is calibration of the demodulated signal from the length dither; this is used as the error signal for the cavity length control and is intended to be an accurate measure of the displacement from the fringe. Typically this approach is only accurate for low frequencies (< 10 Hz), since the bandwidth of the lock-in dither signal is limited by the sensing noise of the OMC DCPDs.

A second method uses the 45MHz sidebands present at the AS port of the interferometer to PDH lock the OMC. This is somewhat invasive and requires an RF photodiode to sense the power of the reflected light from the cavity. However, the PDH signal should be immune to many noise sources, for example intensity noise or scattering from optical elements upstream of the OMC. The H1 OMC was PDH locked during interferometer downtime in January, 2015, and the length noise was characterized to high frequency.

Another method is the so-called ‘half-fringe’ or ‘offset-locking’ technique. This is illustrated in Fig. 5.10. When the OMC is locked at the resonance, length noise couples quadratically to the transmitted power and

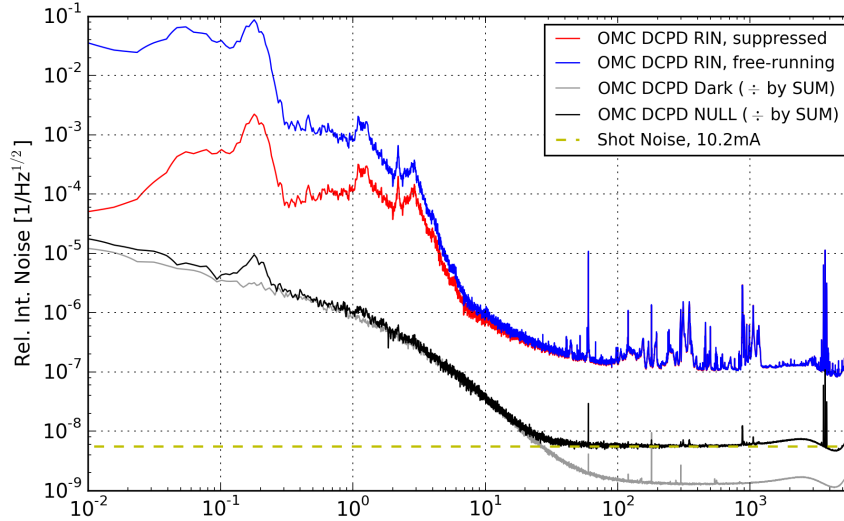


Figure 5.11. Relative intensity noise observed for a half-fringe lock of the OMC. Above 100 Hz the signal is limited by input intensity noise.

generates only small amplitude signals in the DCPDs. If the OMC is locked halfway up the resonance, length noise will couple linearly and generate a larger signal.

In Fig. 5.11 we show results from a half-fringe lock of the H1 OMC with the full interferometer locked at low input power (2.3 W). *Relative Intensity Noise* (RIN) is the variation in the intensity of light measured by the DCPDs as a function of frequency, normalized by the average intensity. In the plot, the red trace is the RIN measured during the half-fringe lock. During this measurement the OMC was locked, and length fluctuations were corrected using the dither lock-in signal and PZT actuation, which suppressed RIN noise from length fluctuations at low frequencies where the control servo has significant gain. In the blue trace we have compensated the observed RIN for the loop suppression. This is our estimate of the RIN due to ‘free-running’ length noise.

The plot includes estimates for the noise floor of the measurement. The gray trace is the measured electronics noise for the DCPDs and the dashed gold line is the shot-noise limit for the DCPD photocurrent at the time of the measurement (10.2mA). The black trace is the DCPD ‘null’ channel, the residual signal after the signal from one DCPD is subtracted from the other. Note that the shot noise limit agrees well with the floor of the null channel above 60 Hz. In this region, the residual noise between the DCPDs was incoherent and matched the expectation from shot noise. Our signal is significantly higher than the shot noise limit, implying we were measuring real intensity noise. The flat floor of the red curve above 100 Hz is suspicious, and the measurement was most likely limited by input intensity noise in this band. There was some coherence above 100 Hz between the DCPDs and intensity sensors upstream of the OMC, which indicates the RIN in this region was not entirely due to cavity length fluctuations.




Figure 5.12. Two measurements of OMC cavity length noise. Above 100 Hz the half-fringe method is limited by input intensity noise. Between 1 and 100 Hz the PDH method is limited by noise of the PZT actuators. There is agreement at low frequency, and also around 1 kHz, where resonances of the seismic platform disturb the cavity. The body modes of the OMC silica breadboard are visible at 10 kHz in the PDH measurement.

Using the known resonance curve for the OMC (Fig. 5.10), we can calculate the $\delta P/\delta L$ slope at the half-fringe locking point; this is 2.72×10^{-9} meters/RIN. For our measurement in Fig. 5.11, we calculate the length noise assuming the RIN on transmission is entirely due to cavity length fluctuations. This is shown in Fig. 5.12, where we also plot results from the PDH lock of the OMC. Note that the PDH method, which is not sensitive to intensity noise, is lower than our measurement above 100 Hz. Importantly, the two measurements agree on the forest of lines around 1 kHz. These lines are due to a resonance of the support struts in the seismic isolation table that supports the OMC. Table vibrations at these frequencies propagate to the OMC glass breadboard and warp the cavity, generating cavity length noise.

Our goal is to predict the OMC length noise coupling into the gravitational wave channel. We assume that the free-running cavity length noise is the same whether the cavity is locked on the half-fringe (when we made the length noise measurement) or on the full fringe (the configuration when the detector is in low-noise), and we correct for suppression of the length noise due to the OMC control loop. On the full fringe, the coupling of length noise to transmitted power is quadratic, and we can estimate the coupling of length noise into the DCPD power measurement using the formalism for optical cavities. This coupling is given by:

$$\delta P_{AS}^{OMC}(f) = \left(\frac{4\mathcal{F}}{\lambda}\right)^2 \text{ASD}\{L^2(t)\}, \quad (5.3)$$

where $\mathcal{F} = 390$ is the finesse of the OMC, $\lambda = 1064 \text{ nm}$, and $\text{ASD}\{L^2(t)\}$ is the amplitude spectral density of the square of the timeseries of the cavity length noise. We measured $L(t)$ while locked on the half fringe - it is the time series corresponding to the red trace in Fig 5.11. Note that it is important to square the time-series before calculating the ASD; in this way we correctly account for the quadratic coupling of length fluctuations around the cavity resonance.

We convert P_{AS} into δx (the variation in the differential arm length) using the quadratic form of Eq. 5.2 and our knowledge of the DARM offset. Let us rewrite Eq. 5.2, and say that the relationship between the power at the antisymmetric port and the DARM displacement x is given by [22]:

$$P_{AS} = P_{BS} (kx)^2 = \frac{P_0}{P_{ref}} \left(\frac{x}{x_f} \right)^2. \quad (5.4)$$

Here, P_0 is the input power to the interferometer, P_{ref} is a normalization constant, and x_f is a reference offset from zero displacement. When the DARM displacement is equal to x_f , $P_{AS} = P_0/P_{ref}$. (It is convenient to normalize P_{AS} to the input power, since it is easier to measure than P_{BS} .) If we assume the static DARM offset is some value x_0 and we wish to measure small fluctuations δx , then we can expand the term in the numerator:

$$P_{AS} = \frac{P_0}{P_{ref}} \left(\frac{x_0^2 + 2x_0\delta x + \mathcal{O}(\delta x^2)}{x_f^2} \right). \quad (5.5)$$

Solving for δx , and dropping the DC term:

$$\delta x(f) = \delta P_{AS}(f) \frac{P_{ref}}{P_0} \frac{x_f^2}{2x_0}. \quad (5.6)$$

This is how we measure the differential length fluctuations using homodyne readout. From here, we convert to strain using the total arm length: $\delta h = \delta x/L$. During typical low-noise operations at H1, $x_f = 14 \text{ pm}$, $P_{ref} = 1560$, and $P_0 = 24.1 \text{ W}$. In this state, for $x_0 = 16 \text{ pm}$, P_{AS} is about 20 mW.

Before we can compare this result to the detector sensitivity in Fig. 5.5, we must correct δx for the detector response, in particular the coupled-cavity pole of the arms and the signal recycling cavity. With this, our expression becomes:

$$\delta x^{OMC}(f) = \delta P_{AS}^{OMC}(f) \frac{P_{ref}}{P_0} \frac{x_f^2}{2x_0} \left(1 + i \frac{f}{f_{pole}} \right), \quad (5.7)$$

with $f_{pole} = 389 \text{ Hz}$ and $\delta P_{AS}^{OMC}(f)$ calculated from Eq. 5.3. The result of this calculation is shown in Fig. 5.13, where we compare the noise in the gravitational wave channel due to OMC length fluctuations to the observed displacement sensitivity. At all frequencies the noise from the OMC is lower than the current noise spectrum, and we conclude that the OMC is not a limiting noise source. Furthermore, recall that above 100 Hz the measurement




Figure 5.13. Estimate of DARM noise due to OMC length fluctuations (loop-suppressed), compared to the recent sensitivity of the H1 interferometer. Note the y-axis is in units of displacement sensitivity; divide by the arm length to get the equivalent strain sensitivity.

of OMC length noise was likely to be an overestimate due to input intensity noise. The only region for worry is the forest of lines around 1 kHz. These peaks are due to resonances in the seismic isolation table, as mentioned above. There are currently efforts to design dampers to mitigate these resonances.

5.5 The OMC as a Spectrum Analyzer

We can use a high-finesse cavity like the OMC as an optical spectrum analyzer to measure the frequency and spatial mode content of the laser field incident to the cavity. The resonant spatial modes in an optical cavity are the familiar Hermite-Gauss or Laguerre-Gauss transverse electromagnetic modes (TEM). They are characterized by two indices labeling the mode order in the transverse directions and are written TEM_{mn} with $m, n \geq 0$. Laser fields composed of higher order spatial modes acquire excess round-trip phase as they circulate in the cavity. This phase is referred to as the *Gouy phase*, ζ , and it increases the effective frequency of the laser field. Thus, higher order modes are resonant for cavity lengths slightly offset from the resonance condition for the carrier. The frequency spacing for higher order spatial modes is calculated using the Gouy phase,

$$f_{\text{HOM}} = \frac{\zeta}{\pi} f_{\text{FSR}}, \quad (5.8)$$

and the optical frequency for a mode of order m, n is

$$f_{mn} = f_0 + (m + n + 1)f_{\text{HOM}}, \quad (5.9)$$




Figure 5.14. Scan of the OMC with a single-bounce beam off an ITM. The scan data are plotted in black and fits to the peaks are plotted as colored Lorentzian functions (purple for sidebands, blue for higher-order modes, green for an extended sample of higher-order modes). The gold trace is the sum of the fitted peaks. As the voltage on the PZT decreases, the cavity length decreases, and higher-order modes with increasing optical frequency are resonant.

where f_0 is the carrier frequency. The round-trip Gouy phase of an optical cavity is calculated from the mirror curvatures (for spherical mirrors this is quite simple, using the radii of curvature for the mirror surfaces). The calculation of the Gouy phase for cavities with more than two mirrors is somewhat more complicated, but straightforward [32]. For the aLIGO OMC, $\zeta = 39.5^\circ$, $f_{\text{FSR}} = 261$ MHz, and $f_{\text{HOM}} = 58$ MHz. Even for an input beam with pure TEM_{00} mode content, higher order modes can be generated by imperfect overlap of the input beam with the resonant beam in the cavity, due to misalignments of the input beam or a mismatch in beam size or shape.

To perform a spectral analysis with the OMC, we sweep the cavity length using the PZT actuator and measure the transmitted power using the DCPDs. As the cavity length changes, modes of various orders and input fields of various frequencies will resonate in the cavity. For an input beam with RF sidebands of frequency ω_{RF} , we can distinguish pairs of resonances at $\omega_0 \pm \omega_{RF}$. A sweep of the OMC cavity across slightly more than one cavity FSR is shown in Fig. 5.14. For this sweep, most of the optics in the interferometer were misaligned and none of the cavities were locked. The input beam to the OMC was a ‘straight shot’ beam from the input mode cleaner. Essentially it was a measurement of the quality of the laser field that is injected into the interferometer, and the quality of the mode-matching and alignment into the OMC.

There are several notable features in the cavity scan shown in Fig. 5.14:

- The fundamental (00-mode) **carrier peaks** are the tallest feature, with 35mA of photocurrent. Most of the input beam power was at the carrier frequency and well-aligned to the cavity. The peaks are approximately located at 15 V and 56 V on the PZT voltage. These peaks are resonances of the same input field, but the cavity round-trip length has been shifted by $\lambda/2$. The PZTs used on the OMC are roughly calibrated at 14 nm/V and change the cavity round-trip length by twice the PZT displacement: $14 \text{ nm/V} \times 2 \times \sim 40 \text{ V} \simeq 1064 \text{ nm}$. There is some nonlinearity in the PZT response, this will be discussed at the end of the next section. (Note, from Chapter 4, the cavity resonance condition is $L = N\lambda/2$ for the cavity one-way length L . The round-trip cavity length ($2L$) changes by λ between resonances.)
- The **9.1 MHz sidebands** are the small peaks on the shoulders of the carrier peaks, approximately one volt higher or lower. These peaks are fit with Lorentzians, the fits are shown in purple.
- The **45.5 MHz sidebands** are offset from the carrier by about six volts and their height is about 0.07mA in the DCPD signal. Their Lorentzian fits are also shown with a purple trace.
- The **carrier TEM₂₀ mode** is almost halfway between the two carrier peaks, at 36 V. Sometimes referred to as the *doughnut mode*, the 20/02 mode is symmetric around the cavity axis and is sensitive to a mismatch in beam size between the input beam and the ideal cavity beam. We calculate *mode matching* of an input beam into a cavity using the relative power in the 00 and 20 modes, $1 - (\text{TEM}_{20}/\text{TEM}_{00})$. For this data the mode matching is 0.91. The height of the 20/02 mode was later reduced using the thermal compensation system (TCS), which uses a CO₂ laser to generate a thermal lens in the ITM mirror substrates and tune the effective mirror curvature.
- The **sideband TEM₂₀ modes** are offset from the carrier TEM₂₀ mode. These peaks (and others) are fit with green traces.
- **Higher-order carrier modes** are visible up to $n + m = 8$. A single FSR only contains up to the 4th-order mode, since $5 \times f_{\text{HOM}} > f_{\text{FSR}}$, but 5th-order modes and higher wrap around the FSR due to the repetitive resonance structure. These modes are fit with blue curves.

Peaks in the cavity scan are fit using a three-parameter Lorentzian function:

$$P(V) = \frac{A}{1 + \left(\frac{V-V_0}{V_c}\right)^2}, \quad (5.10)$$

where A , V_0 , and V_c are free parameters. The peak amplitudes A are used to derive information about the relative power in each mode in the input beam. The peak locations V_0 are used to calibrate the PZT voltage in terms of

Modulation frequency	Γ
9.1 MHz	0.191
45.5 MHz	0.284

Table 5.1. Modulation depths at H1 for the RF sidebands, measured using the OMC. For the calculation of the contrast defect in the next section, we assume uncertainties in $\Gamma_{9.1}$ and $\Gamma_{45.5}$ of 0.005.

optical frequency using prior knowledge of f_{HOM} and the sideband frequencies ω_{RF} . In principle the peak width V_c carries information about the finesse of the cavity, although in practice this measurement is limited by the response of the DCPD electronics.

One useful measurement that can be performed with a cavity scan is a characterization of the sideband modulation depth. From Chapter 4, recall the expression for the laser field after the phase modulation by the EOM, using the Jacobi-Anger identity:

$$E_{inc} = E_0 \left[J_0(\Gamma)e^{i\omega_0 t} + J_1(\Gamma)e^{i(\omega_0 + \omega_{RF})t} - J_1(\Gamma)e^{i(\omega_0 - \omega_{RF})t} \right]. \quad (5.11)$$

Here, Γ is the modulation depth, measured in radians. If we were to measure the power at frequency $\omega_0 \pm \omega_{RF}$, we would measure $P_0 J_0(\Gamma) J_1(\Gamma)$. Similarly, the power of the carrier would be $P_0 J_0^2(\Gamma)$. With these two expressions, we can use the ratio of power in the carrier to power in the sideband and provide an independent measure of Γ :

$$\frac{P_{carrier}}{P_{sb}} = \frac{P_0 J_0^2(\Gamma)}{P_0 J_0(\Gamma) J_1(\Gamma)} = \frac{J_0(\Gamma)}{J_1(\Gamma)}. \quad (5.12)$$

We can numerically invert this result to find Γ . Measuring the sideband modulation depth using an optical spectrum analyzer downstream of the EOM such as the OMC is an important check of the RF modulation electronics.

Using the peak heights for the carrier and the sidebands we can follow Eq. 5.12 and calculate the modulation depths for the RF sidebands. The results for this cavity sweep are given in Table 5.1.

5.6 Measurement of the Interferometer Contrast Defect

In the full lock, the higher-order mode content at the antisymmetric port is considerably richer than with a straight-shot beam from the input. Partially, this is due to the sidebands and their interactions with the optics of the DRMI cavities. There is also carrier light present due to asymmetries between the carrier fields resonating in the arms. One source of asymmetry, the DARM offset, is intentionally applied and should only contribute power into the TEM₀₀ mode. But there are other sources of asymmetry, such as slight differences in the mirror curvature, imperfections in the mirror coatings or substrates, interactions between the carrier light and the DRMI optics,




Figure 5.15. Scan of the OMC with the full interferometer locked with zero DARM offset. Higher-order modes of the carrier and sidebands are labeled. Capital letters are used for the 45.5 MHz sidebands and lower-case letters are used for the 9.1 MHz sidebands.

and most importantly any imperfections or misbalance in the beamsplitter reflectivity. These imperfections add up to what is called the *contrast defect* between the arms. The contrast defect results in carrier power at the dark port, even when no DARM offset is applied. Since these modes do not couple to gravitational waves, they can act as a source of noise.

To measure the contrast defect, we performed sweeps of the OMC cavity with the full interferometer locked, with the DARM offset on and off. The results of the sweeps and the fit to the complicated mode structure are shown in Figs. 5.15 and 5.16.

First, note the height of the 45.5 MHz sideband peaks. This sideband frequency is resonant in the signal recycling cavity and is the dominant source of light at the antisymmetric port, about 40 mA of total photocurrent. Next, note the large order to which the carrier peaks extend: there is substantial power up to CR9, and other peaks are present but not labeled. Finally, note the difference in height of the CR0 peak between the data with the DARM offset off (Fig. 5.15) and the DARM offset on (Fig. 5.16). Without a DARM offset, there is almost no TEM₀₀ carrier field at the antisymmetric port. This gives us confidence that the contrast defect is small.

To calculate the contrast defect we match every peak to an optical mode and fit the peak using the Lorentzian function given in Eq. 5.10. The amplitudes from the fits for the identifiable peaks are shown in Table 5.2.

We calculate the contrast defect in two ways:

Mode	Carrier	+45.5 MHz	-45.5 MHz	+9.1 MHz	-9.1 MHz
0	0.040 (15.940)	19.450 (19.803)	18.95 (19.370)	0.041 (0.047)	
1	0.231 (0.123)	1.421 (0.221)	0.650 (0.580)	0.294 (0.131)	
2	0.327 (0.638)	0.794 (0.768)	1.31 (1.093)	0.116 (0.110)	0.201 (0.284)
3	0.501 (0.647)		0.123 (0.056)	0.087 (0.065)	0.180 (0.472)
4	0.247 (0.060)		0.044 (0.052)	0.123 (0.151)	0.632 (0.622)
5	0.437 (0.450)			0.055 (0.085)	0.074 (0.125)
6	0.211 (0.134)			0.013 (0.022)	0.150 (0.291)
7	0.185 (0.145)			0.025 (0.017)	0.033 (0.038)
8	0.178 (0.146)			0.021 (0.014)	0.077 (0.133)
9	0.435 (0.275)				
Sum	2.79 (18.56)	21.67 (20.81)	21.07 (21.15)	1.14 (0.64)	1.35 (1.97)

Table 5.2. Peak heights for OMC cavity sweeps in full lock, with DARM offset off (on). Heights are given in milliamps of DCPD photocurrent. Peaks without a measurement were either too small to fit, or degenerate with another mode. We conservatively estimate the uncertainty of the carrier sums to be 0.05 mA.




Figure 5.16. Scan of the OMC with the full interferometer locked with the DARM offset on, $x_0 = 15.8$ pm. Note the height of the CR0 peaks compared to Fig. 5.15.

1. We assume that all the power in the carrier TEM_{*nm*} modes with $n + m > 1$ with the DARM offset **on** is due to the contrast defect. This sum (2.62 ± 0.03 mA) is likely an overestimate since some of the carrier TEM₀₀ power will be lost to higher order modes due to misalignments. The mode-matching for the carrier field is around 90% for a single-bounce beam (as measured in the previous section), so the overestimate may be as much as 10%.
2. We assume that all the carrier power with the DARM offset **off** (2.79 ± 0.03 mA) is due to the contrast defect. In this case we are susceptible to overestimation from residual offsets due to the RF electronics or electrostatic actuation on the ETMs.

In each case, we calculate the contrast defect as the ratio of carrier power at the antisymmetric port to the total carrier power available before the destructive interference of the Michelson interferometer. We must account for losses between the beamsplitter and the OMC to properly measure the available power.

The total power available at the antisymmetric port is equal to the product of the input power (at the time of the measurement), the efficiency of the input optics (measured to be 88%), the power after the phase modulation for the sidebands (using our measurement of Γ from the previous section), the carrier gain in the power recycling cavity, and the transmission of the SRM:

$$\begin{aligned}
P_{AS}^{total} &= P_{in} \times J_0^2(\Gamma_{9.1}) \times J_0^2(\Gamma_{45.5}) \times \epsilon_{IO} \times g_{cr}^2 \times T_{SRM} \\
&= (2.8 \pm 0.1 \text{ W}) \times (0.982 \pm 0.001) \times (0.960 \pm 0.001) \\
&\quad \times (0.88 \pm 0.02) \times (33 \pm 2) \times (0.37 \pm 0.001) = 28.3 \pm 2.1 \text{ W}. \tag{5.13}
\end{aligned}$$

Losses between the SRM and the OMC are due to losses in the output Faraday isolator, the transmission of the tip-tilt mirrors, and the responsivity of the OMC DCPDs:

$$\begin{aligned}
P_{AS}^{loss} &= T_{OFI} \times T_{OM1} \times T_{OM3} \times PD_{QE} \\
&= (0.95 \pm 0.02) \times (0.95 \pm 0.001) \times (0.99 \pm 0.001) \times (0.75 \pm 0.02 \text{ A/W}) \\
&= 0.67 \pm 0.017 \text{ A/W}. \tag{5.14}
\end{aligned}$$

The contrast defect is then given by:

$$\begin{aligned}
\frac{\sum P_{carrier}}{P_{AS}^{total} \times P_{AS}^{loss}} &= \frac{2.62 \pm 0.05 \text{ mA}}{28.3 \text{ W} \times 0.67 \text{ A/W}} = \mathbf{137.8 \pm 11.1 \text{ ppm}} \quad (\text{first case}) \\
&= \frac{2.79 \pm 0.05 \text{ mA}}{28.3 \text{ W} \times 0.67 \text{ A/W}} = \mathbf{146.8 \pm 11.8 \text{ ppm}} \quad (\text{second case}). \tag{5.15}
\end{aligned}$$




Figure 5.17. Fit to mode optical frequency, for OMC cavity sweep in full lock with DARM offset on. The optical frequency for each mode was determined using the sideband frequencies and the OMC higher-order mode spacing. The peaks were fit to Lorentzian functions of the PZT voltage, and the set of voltages were fit to the mode frequencies with a 4th-order polynomial.

The two methods are in good agreement. The dominant source of uncertainty is from the carrier recycling gain, g_{cr}^2 . Surprisingly, the first case, which we thought would overestimate the contrast defect, returned a smaller result. There are a variety of potential systematic uncertainties for this measurement, especially the alignment control of the interferometer, which was not in a complete state at the time the data were collected. In any case we can be confident that the contrast defect is quite small. Note that if we have not accounted for all of the losses between the beamsplitter and the OMC, our result is biased and the true contrast defect will be larger.

Finally, one last result: we can use the mode content from the output of the full interferometer to calibrate the OMC PZT actuation in terms of optical frequency. We fit the PZT voltage to optical frequency using the sideband frequencies and the higher-order mode spacing of the OMC, with a 4th order polynomial. The fit parameters are set using thirteen data points: the four peaks of the 45.5 MHz sidebands, and the carrier peaks from CR0 to CR9. The result of the fit was compared to all of the identified peaks in the sweep; this is shown in Fig. 5.17, for data with the DARM offset on (Fig. 5.16). The response of the PZT is not quite linear. The residuals of the fit are extremely small; this gives us confidence that we have correctly identified the optical modes.

CHAPTER 6

DETECTOR CHARACTERIZATION

*Be not afeard; the isle is full of noises,
Sounds and sweet airs, that give delight and hurt not*

Gravitational wave observatories are complex instruments, and despite over a decade of experience their daily operation are an ongoing research and development project. Observatories operate in two modes: *commissioning*, when detector scientists make invasive changes to understand and improve the sensitivity to gravitational signals, and *observing*, when the detectors are maintained in a resonant state as much as possible, with a minimum of changes to the detector configuration. Typically, the instruments can be held at resonance for a few hours until and environmental disturbance or control instability causes them to lose lock.

Calibrated strain data from the LIGO detectors is stored in a time-series format with a sample rate of 16,384 Hz. The strain output from the detectors is not a pristine, Gaussian-noise dataset. Numerous noise artefacts are present in the data and can masquerade as gravitational wave signals. Excess noises range from constant narrowband features such as lines, which are a problem for searches for continuous-wave signals, to short-duration bursts of broadband excess noise. The latter case, referred to as *glitches*, are presumed to arise from time-varying couplings of environmental or instrumental noise into the gravitational wave channel. Tremendous effort has been applied to studying, categorizing, and searching for the cause of glitches [8]. Glitches form a background that limits the sensitivity of astrophysical searches for short-duration signals.

In this chapter, I will broadly describe the typical data quality of the LIGO detectors, and quantify the non-Gaussian component of the noise. I will present some results from a novel technique for identifying the instrumental cause of glitches, and conclude with an example of glitches in the advanced LIGO detectors.

6.1 Locking and Unlocking

During an observing run, the detectors are maintained in their low-noise state for as much time as possible by round-the-clock operations staff. The duty cycle or *uptime* of a detector is limited by a variety of factors. During




Figure 6.1. Accounting of the sources of downtime for the L1 detector during the two-year S5 science run. Above-normal ground motion, in the form of earthquakes, wind, trains, construction and logging activities near the site, and the microseismic peak due to ocean waves, was responsible for 100 days of lost science data. The H1 detector did somewhat better in S5, with 76.9% uptime. In the S6 science run, between July 2009 and October 2010, the instruments had slightly worse duty factor due to preparations for the advanced LIGO installation.

the two-year S5 science run a careful accounting of the causes for downtime was made at the L1 detector, and the results are shown in Fig. 6.1. The downtime is organized into the basic categories that plague gravitational wave observatories: environmental disturbances (wind, earthquakes, local seismic noise and microseism from ocean waves); equipment failures; regular maintenance and commissioning activities; and the irreducible time necessary to recover from a lock loss by running the ‘Up’ and ‘Down’ scripts.

The uptime achieved by the LIGO detectors has historically been a major focus of commissioning efforts. Uptime between the LIGO sites is rarely correlated. If each site manages to collect data with 70% duty cycle in randomly-distributed segments, the network duty cycle is less than 50%. This results in a factor of two reduction in the anticipated event rate for an observing run and can make the difference between detections and upper limits.

The durations of the *science segments* that represent undisturbed periods of data acquisition are highly variable. A histogram of segment duration for the LIGO detectors in the S6 science run is shown in Fig 6.2. The median duration for L1 segments was 2,400 seconds, the H1 median was 4,700 seconds.

Early duty cycle results from the advanced LIGO detectors are encouraging. The advanced seismic isolation has dramatically reduced the coupling of seismic disturbances into aLIGO interferometers. This is especially the case in the anthropogenic seismic band of 1-10 Hz, generated by human activity such as cars, trucks, trains, and machinery, and also by logging activities near the Livingston site.




Figure 6.2. Duration of science segments collected by the LIGO interferometers in their sixth science run, July 2009 - October 2010. The L1 instrument recorded 3,276 segments, H1 recorded 2,395.

Despite the increased complexity of the advanced detectors, the recovery from lock loss can be a straightforward, rapid process. In Fig. 6.3 we plot the recovery from a lock loss in the H1 detector. The ‘Guardian’ is the automation software which controls the state of the instrument and manages the complex transitions between the uncontrolled, unlocked state and the resonant low-noise state. Various steps in the path from unlocked to resonant are given on the left-hand side of the plot. In this example, the instrument was locked for the first sixteen minutes of the hour, lost lock, and recovered the full low-noise state (‘LSC_FF’) in less than twenty minutes. The process was entirely automated with no human intervention.

The first two weeks of June 2015 were set aside for an engineering run to quantify the duty cycle of the instruments and provide data for tests of the analysis pipelines. The H1 instrument achieved a duty cycle of 63%, and L1 42%. Many of the reasons for downtime during the engineering run were relatively well-understood and are expected to be remedied before the first observing run in September 2015. These engineering run results can be taken as a lower bound on future performance.

6.2 Stationarity of the Noise & Effects on the Searches

In the ideal case, the noise in the detector is described by a stationary Gaussian process. In reality, the noise is nonstationary with many loud bursts of excess noise that are extremely improbable for noise drawn from a constant Gaussian distribution. We can characterize the Gaussianity of the noise by studying the amplitude spectral density (ASD) of the strain signal from the detectors. Recall that the power spectral density (PSD) of a time series $h(t)$ is given by:




Figure 6.3. Example of a lock loss recovery for the H1 advanced LIGO detector. Figure by Duncan MacLeod.

$$S(f) = \frac{1}{N} \sum_{i=0}^{N-1} |\tilde{H}_i(f)|^2, \quad (6.1)$$

where $\tilde{H}_i(f)$ are samples of the Fourier transform of $h(t)$, taken over N discrete steps. If $h(t)$ is drawn from a Gaussian distribution, then its Fourier transform $\tilde{H}_i(f)$ will have real and imaginary components that are independent Gaussian random variables with zero mean and variance σ^2 . The ASD is the square root of the PSD, and for Gaussian noise the set of samples from each frequency bin in the ASD will follow a *Rayleigh distribution*. Details of ASD estimation and the statistical distribution of the ASD for stationary Gaussian noise are given in Appendix A.

We can compare the observed ASD of the detector to a Rayleigh distribution and quantify the excursions from Gaussianity. This was done in detail for the LIGO detectors in the S6 science run, between July 2009 and October 2010. S6 was divided into four epochs, based on commissioning breaks that resulted in significant changes to the sensitivity. For the study, the science-quality data from each epoch was Fourier transformed and converted to an ASD. An example result is shown in Fig. 6.4 for the H1 detector. Here, we plot the percentile of the noise ASD for each frequency bin, with 0.167 Hz resolution.

In Fig. 6.5, we normalize each percentile by the median, and compare the result to the expectation for a Rayleigh distribution. Visualizing the data in this way gives us a picture of what fraction of time the noise behaves in a non-Gaussianity manner, as a function of frequency. Notice that the median-normalized percentiles largely match the dashed lines which give the expectation for Gaussian noise. At low frequencies, a deviation from Gaussianity is clear after roughly the 95th percentile. This observation agrees with our experience from the detectors. The rate of loud glitches is about 0.5 Hz and their median duration is about 100 msec, thus there is




Figure 6.4. Distribution of sensitivity per frequency bin for the H1 detector during the S6D epoch (6.28 million seconds of data, or about 72 days). The data were divided into 50% overlapping segments with 6-second duration. The FFTs of these segments were collected and used to calculate percentiles of the noise for each frequency bin. For 0.5% of the time, the noise was greater than the black trace.

some excess noise present about 5% of the time. Most glitches occur at low frequency. Broadband glitches up to 1 kHz are much less frequent, thus there are only small deviations from the dashed lines at high frequency.

The population of glitches in the detector acts as background noise for searches looking for transient gravitational waves. Searches for compact binary coalescences (CBCs) are somewhat immune to the problem, since the CBC signal morphology is well-defined, and searches can perform χ^2 tests to check that event candidates agree with the signal model. Searches for unmodeled, short-duration bursts have significant problems with glitches, since they make no assumptions about the signal morphology. In Chapter 7 and Appendix C we discuss coherent analysis techniques for glitch rejection. Since the noise in the detectors is non-Gaussian and glitch mechanisms are largely unexplained, an analytical model for the background noise in transient searches is not possible. Instead, the background is measured empirically, by repeating the analyses with many unphysical *time-slides*, or time offsets between the detector data. In this way the searches estimate the rate of random noise correlations between the detectors.

An example of the results of the background estimation for an unmodeled burst search is shown in Fig. 6.6. There is a long tail of high-SNR events, much higher than would be expected from astrophysical sources. In order to claim a detection, an event candidate would need to have an SNR that was much higher than events which are found to occur from chance noise fluctuations in the detectors. The events in the tail are so frequent and so loud that the SNR threshold for detection all but excludes plausible astrophysical signals.




Figure 6.5. Median-normalized sensitivity percentiles per frequency bin. The expected ratio for Rayleigh-distribution noise for each percentile is given by the dashed lines. There is a clear deviation from Rayleigh-distributed noise, especially at low frequency.




Figure 6.6. Cumulative histogram of time-slide events for the Coherent WaveBurst search in S6D [15]. The long tail of events, especially at low frequencies, extends to an astrophysically unlikely SNR, and statistically blinds the search to plausible signals. Vetoing the events based on correlations to instrumental and environmental noise improved the results, but at the expense of duty cycle. Typically around 10-20% of the data are vetoed. The error bars are \sqrt{N} . This plot is my own work, and should not be interpreted as a result of the CWB group.

Ideally, the glitch background would be eliminated through studies of the noise couplings and modifications to the detectors. Alternatively, we can monitor the noise sources or the coupling parameters and advise the search groups to ignore times when we expect excess noise to be present in the gravitational wave channel. This second approach is known as *vetoing* the data. The identification of glitch coupling mechanisms is not easy. The coupling of any individual noise source to the gravitational wave channel can be modulated by time-varying parameters, and the parameter space of the detectors is described by many hundreds of channels. In the next section we describe a technique to identify instrumental couplings that generate glitches.

6.3 Instrumental Correlations of Glitches

A number of algorithms have been developed [23, 75, 120, 136] to study correlations between glitches in the gravitational wave channel and auxiliary instrumental channels – for example, the alignment error signals, length control signals for the DRMI optical cavities, seismic noise, environmental data from magnetometers and microphones, and so on. These techniques have performed very well, and glitch backgrounds such as those shown in Fig. 6.6 can generally be vetoed such that the analyses are sensitive to astrophysically plausible SNRs. But these algorithms rarely give a clear answer for the instrumental cause of glitches. In a sense they answer the question of ‘What caused the glitch in the gravitational wave channel?’ with the question ‘What cause the glitch in the auxiliary channel?’ Sometimes the cause is readily apparent, for example when a photodiode is glitching due to an electronics problem, but often the glitch in the auxiliary channel looks like a louder, filtered version of the glitch in the gravitational wave channel. This does not always provides useful information.

The detector parameter space is described by hundreds of measurements: the length of optical cavities, the position of the beam on mirrors, the alignment of the beam, and so on. We can visualize the interferometer state at an instant in time as a high-dimensional vector in the parameter space of the instrumental channels which are recorded. The veto algorithm presented here calculates the probability that the interferometer will experience a glitch at each point in this high-dimensional space. There are two assumptions in the method. First, that the duration of glitches is short (~ 100 msec) and the noise coupling to the gravitational wave channel is changing on this timescale. Second, we assume that the coupling is modulated by two or more parameters which are measured by auxiliary channels (or parameters derived from them), and that studying a single channel is not sufficient to describe the glitch mechanism. These assumptions are motivated by previous studies which correlated glitches to the value of single channels over short and long timescales, with varying degrees of success.

For now we consider a two-parameter glitch coupling model. To illustrate the method, we will introduce a source of glitches known as Barkhausen noise [145]. Barkhausen noise is generated by the discrete transitions of the microscopic magnetic domains in a ferromagnetic material as an external magnetic field is varied. In




Figure 6.7. Correlation between glitches and ETMX actuation current during highly elevated, narrowband microseismic noise on Oct 22 2009 at the H1 detector. The value of the coil current at the time of a glitch in the gravitational wave channel is given by the red squares. The grey traces are the coil current around the time of the glitch. The seismic noise on this day was sharply peaked around 0.125 Hz.

initial LIGO, the final-stage actuators on the test masses were electromagnetic coils, and Barkhausen noise was observed during times when the amplitude of the drive to the actuators was very large.

In Fig. 6.7 we show the correlation between glitches and the coil current sent to the electromagnetic actuators on ETMX. The signature of Barkhausen noise is clear. The bursts of noise arrive just after the time of most rapid change in the magnetic field of the voice-coil actuators. These are the times when the largest number of magnetic domains are flipping to align with the changing external field from the actuators.

Note that if we only examined the value of the coil current at the time of the glitches, the pattern would be hidden from us. Furthermore, if we only examined the behavior of the channel on long timescales (say, the average value of the channel during a science segment) we would miss the correlation. Earlier vetoes for Barkhausen noise looked for times when the RMS coil current was elevated. These vetoes worked, but the amount of time vetoed was very large.

The timing of the bursts of excess noise due to the Barkhausen effect can be anticipated using a model with two parameters, the amplitude of the coil current and the phase. A plot of the current and its derivative extracts this information; this is analogous to a phase-space diagram from classical mechanics. In Fig. 6.8 we plot the location of glitches in the parameter space of the coil current (x-axis) and the time-derivative of the current (y-axis). Each pixel in the plot represents a discrete value of the coil current and its time derivative.




Figure 6.8. Location of Barkhausen-type glitches in a phase-space plot of the ETMX coil current and its time derivative. Note the correlation between glitch times and the phase and amplitude of the coil current. The data here are the same as in Fig. 6.7.




Figure 6.9. Glitch probability for two months of data in H1 S6B as a function of coil current and its time derivative. Nearly all of the time the interferometer state is located in the very center of the plot, where the glitch probability is around 10^{-3} . During times of very large, low-frequency seismic disturbances, the coil current amplitude is very large, and the interferometer state falls in points on the extremes of the plot. Barkhausen glitches tend to occur in the upper right and lower left quadrants.




Figure 6.10. Performance of a veto constructed by rejecting times when the coil current was in a probable configuration to generate Barkhausen noise. Statistics are presented per science segment.

By dividing the glitch count in each pixel of Fig. 6.8 by the amount of time the coil current and its time derivative had the value corresponding to that pixel, we can build a glitch probability map for the two-dimensional parameter space of the coil current and its time derivative. This is shown in Fig. 6.9. Most of the time, the interferometer state is located at the very center of this plot where the glitch probability is low. Glitches are correlated with excursions from the center, and in particular there is a phase dependence shown by the two lobes of low glitch probability. From here the definition of a veto is clear: we flag any times the coil current and its time derivative cross a pixel with a large glitch probability.

The performance of a veto derived in this way is shown in Fig. 6.10. In the top panel we plot the veto efficiency (the fraction of glitches vetoed) and the veto deadtime (the fraction of science time lost to the veto) per science segment. In the middle panel, we plot the *use percentage*, the fraction of veto segments that contain a glitch. In the bottom panel we show the veto efficiency-over-deadtime, a common measure of veto effectiveness¹. Vetoes with a large efficiency-over-deadtime indicate that we have understood a glitch coupling mechanism. When you build a good veto, it means you have learned something about the detector.

The periods of best performance for this veto occur during periods of elevated microseismic noise (ground motion in the 0.1-1 Hz band) due to storms in the northern Pacific in the fall of 2009. During these times, the control loops had to apply large drives to the mirrors to maintain the interferometer lock, and the coil current

¹A little algebra will reveal that the efficiency-over-deadtime statistic is equivalent to the ratio of the glitch rate during the veto segments to the average glitch rate.

was very large. Outside of these times the amplitude of the coil current drive was small and very little time was vetoed. Note that the veto use percentage is quite high. More than 60% of the time, a veto segment contained a glitch.

In the advanced LIGO detectors, great care has been taken to remove all ferromagnetic materials from the test masses, and the final stage of the suspensions use an electrostatic drive rather than magnetic actuation. This invites other problems, such as varying drive gain as the charge on the test mass changes due to cosmic rays, UV emission from ion pumps in the vacuum system, and so on. But there has been no evidence of Barkhausen noise.

The veto algorithm presented here has been run on a variety of instrumental channels in data from the S6 science run, and some pairs of channels (or pairs of channels and time derivatives) have demonstrated impressive veto performance. An example is the combination of a measure of the spot location on the OMC QPDs (recall that the LIGO detectors used homodyne readout in S6) and a measure of the differential misalignment of the arm cavities. The glitch mechanism implied by this combination of channels is not immediately clear. For the advanced detector era we plan to implement a multidimensional model for glitch couplings that can identify noise sources which are modulated by more than three instrument parameters. An added feature, still in development, is a glitch clustering algorithm to classify glitches based on their morphology.

6.4 Glitches in aLIGO: OMC Scattering

Early data collected by the aLIGO instruments has shown that the glitch background is roughly the same as that observed with the initial detectors. One class of glitches which has been studied, and which is germane to this thesis, are bursts of noise due to scattered light from the OMC. Scattered light is a significant problem for gravitational wave detectors, since even small changes in the amplitude of the carrier field at the antisymmetric port will generate a signal. A scattering source that reflects light back into the beam path with a varying path length will generate an electromagnetic field at the DCPDs with a phase shift relative to the carrier field. In this section we provide a brief derivation of the signal from scattered light (following [64]) and give some initial results from the H1 aLIGO instrument.

Imagine a scattering source located somewhere between the beamsplitter and the DCPDs. This source will reflect a small amount of the electric field incident on the antisymmetric port back into the interferometer, with a path length given by $x(t)$. When viewed from the AS port, the interferometer is almost perfectly reflective, so all of this light will be reflected back towards the AS port. Let us assume that the scattering source reflectivity r_s is small, so higher-order reflections can be neglected. The field at the DCPDs in the presence of a scattering source is given by:

$$E_{AS}^{scatter} = \sqrt{1 - r_s^2} E_{AS} e^{i\omega_0 t} + r_s E_{AS} e^{i\omega_0 t + ikx(t)}, \quad (6.2)$$

where E_{AS} is the field leaving the interferometer. The path length from the interferometer to the scattering source and back, for a source moving at frequency ω_s with amplitude A , is changing like:

$$x(t) = 2A \sin(\omega_s t). \quad (6.3)$$

Therefore the field detected by the OMC is being phase-modulated like:

$$E_{AS}^{scatter} = \sqrt{1 - r_s^2} E_{AS} e^{i\omega_0 t} + r_s E_{AS} e^{i\omega_0 t} e^{i \frac{4\pi A}{\lambda} \sin(\omega_s t)}. \quad (6.4)$$

We can expand the second term using the Jacobi-Anger identity (Eq. 4.13, with $\Gamma = 4\pi A/\lambda$). This gives us the following expression for the field incident on the OMC:

$$E_{AS}^{scatter} = \sqrt{1 - r_s^2} E_{AS} e^{i\omega_0 t} + r_s E_{AS} e^{i\omega_0 t} \left[\sum_{n=-\infty}^{\infty} J_n(\Gamma) e^{in\omega_s t} \right]. \quad (6.5)$$

If we assume the scattering source does not change the TEM_{nm} mode content of the light, this field will be transmitted by the OMC and be detected by the DCPDs. Dropping higher-order terms in r_s , and assuming that $1 - r_s^2$ is close to unity, the excess signal is given by:

$$E_{AS}^{scatter} = |E_{AS}|^2 \left(1 + r_s \left[\sum_{n=-\infty}^{\infty} J_n(\Gamma) e^{in\omega_s t} \right] \right). \quad (6.6)$$

In the frequency domain, the Fourier transform of the second term is a series of delta functions at frequencies $\omega_0 \pm n\omega_s$ with amplitudes $r_s J_n(\Gamma)$. In practice the carrier field $E_{AS} e^{i\omega_0 t}$ has broadband frequency content, and this broadens the delta function peaks.

The distinguishing feature of scattered light noise is the presence of a sharp cutoff frequency, sometimes called the *scattering shelf*. The shelf is due to the cutoff in the amplitude of the scattering harmonics, since $J_n(\Gamma)$ drops rapidly after $n > \Gamma$ (see Fig. 6.11). The cutoff frequency can be related to the properties of the scattering source by:

$$f_{shelf} = n_{shelf} \omega_s = \Gamma \omega_s = \frac{4\pi A}{\lambda} \omega_s = \frac{2v_s}{\lambda}, \quad (6.7)$$

where v_s is the velocity of the scattering source, $v_s = A\omega_s$, from the time derivative of $x(t)/2$.

A spectrogram of scattering noise in the H1 detector is shown in Fig. 6.12. The maximum frequency of the scattering noise, out to 100 Hz or more, implies a scattering source velocity of tens of microns per second relative to the OMC if we assume a single reflection. In this case it was motion of the OMC suspension itself that was modulating the path length between the scattering source and the interferometer. During a test run of the




Figure 6.11. Cutoff of Bessel functions of the first kind as the order n grows larger than Γ .

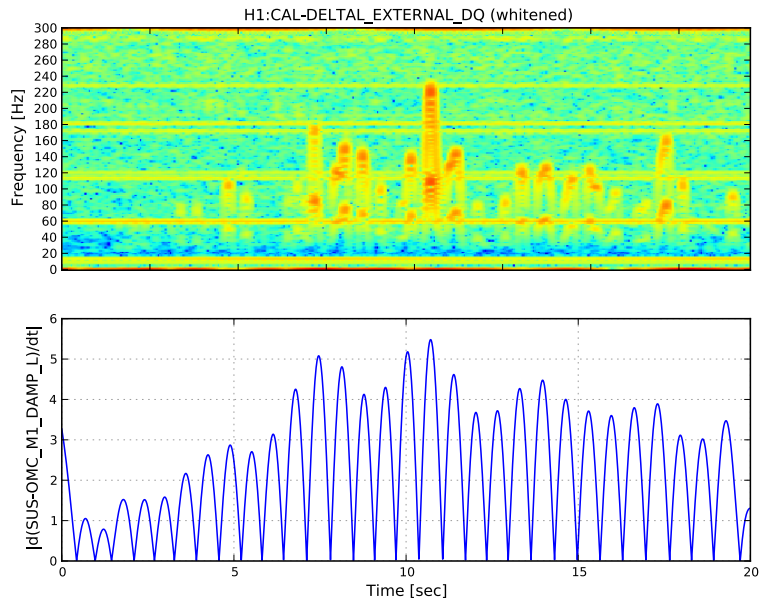


Figure 6.12. Top panel: spectrogram of scattering noise in the calibrated gravitational wave channel. Bottom panel: velocity of the longitudinal degree of freedom of the OMC suspension in $\mu\text{m}/\text{sec}$. Scattering arches are correlated with peaks in the velocity of the OMC, indicating that the path length between a scattering source and the interferometer is being modulated by large motions of the OMC suspension. (The motion of the OMC is measured at the middle stage of the suspension, and the bottom stage, with the optical components, is moving faster.)



Figure 6.13. Photo of the author working to check beam alignment into the OMC black glass shroud, designed to enclose the output mode cleaner in a non-reflective covering to prevent scattered light. The black welder's glass can be seen around the OMC suspension cage in the upper left of the image.

instruments in early June, noise variations due to scattered light were a significant source of background events in the searches for unmodeled burst signals. Later that month, a black glass shroud was installed around the OMC suspension cage to mitigate the effects of scattered light. The installation of the shroud is shown in Fig. 6.13. In the future scattered light noise is expected to be a frequency obstacle to improved sensitivity at low frequencies.

CHAPTER 7

A SEARCH FOR GRAVITATIONAL WAVES ASSOCIATED WITH GAMMA-RAY BURSTS

*Yet from those flames
No light, but rather darkness visible.*

Gamma-ray bursts (GRBs) are intense flashes of high-energy photons which are observed approximately once per day and are isotropically distributed in the sky [93]. Since their public discovery in 1973 [79], astronomers have found that most GRBs are extra-galactic in origin, and there is substantial evidence that they emit gamma-rays in tightly beamed relativistic jets [61]. GRBs are grouped into two broad classes based on their spectral hardness and the duration of the initial gamma-ray flash [82]. The progenitors of *long-soft* GRBs are generally accepted as core-collapse supernovae (CCSN) in massive, rapidly rotating stars [148]. The progenitors of *short-hard* GRBs have yet to be constrained by observation, but are widely thought to be associated with compact binary coalescences (CBCs), such as the mergers of binary neutron star (BNS) or neutron star-black hole (NSBH) systems [97]. Both the merger and core-collapse supernovae scenarios result in the formation of a stellar-mass black hole or magnetar with an accretion disk. CBCs are expected to be bright sources of gravitational waves, while the gravitational wave emission by supernovae is more speculative [100].

In this chapter I present the methods and results of a search for gravitational waves associated with 129 GRBs using data from the LIGO, Virgo, and GEO detectors [11]. The analysis was primarily carried out by myself and Tom Adams, a graduate student at Cardiff University in the United Kingdom. This was the first search for gravitational wave transients to use data from GEO 600, and the first to use data from a squeezed-light interferometer [14]. Due to the sensitivity of the GEO detector, the search focused on a higher frequency band than previous searches. This motivated novel methods for analyzing GRBs with sky localizations from the *Fermi* GBM instrument, which typically have uncertainty regions of dozens of square degrees. The use of a linear search grid for an analysis with two detectors was found to significantly reduce the computational cost of the search. As a part of this work, we characterized the ability of the analysis pipeline to localize a detected gravitational signal. In the search itself, no gravitational wave events were detected. We placed upper limits on

the gravitational wave emission from the GRBs, for generic narrowband waveforms at 500 Hz and 1 kHz, and placed lower limits on the distance to the GRBs assuming a gravitational wave emission energy of $10^{-2} M_{\odot} c^2$.

7.1 Gamma-ray Burst Physics

The prevailing model of GRBs is the so-called *fireball model* [63]. A compact central engine generates a relativistic explosion of material composed of hadrons and electrons, with a bulk Lorentz factor $\Gamma \gtrsim 10^2$. Variations in the density of matter in the relativistic outflow are called *shock fronts*. As the material moves away from the central engine, electrons in the shock front plasma interact with magnetic fields and emit photons through synchrotron radiation. Since the photon energy in the rest frame of the shock front is in the x-ray band, photon-photon interactions such as pair production are suppressed, and the fireball remains optically thin. In the frame of the observer, the photons are blueshifted (Lorentz-boosted) in the direction of the relativistic outflow, and reach energies beyond 1 GeV [104]. Further interactions between the shock fronts and the circumstellar environment will generate a long-lived afterglow. The hadronic content of the fireball produces pions, which decay to electrons, muons, and neutrinos. These high-energy neutrinos are a target for detectors such as IceCube [1]. The fireball model is in agreement with many of the observed properties of gamma-ray bursts, such as the millisecond structure in their light curves, the non-thermal radiation spectrum of the gamma-ray flash, and the presence of long-lived afterglows in the x-ray and optical bands. But, the model is agnostic regarding the properties of the central engine, other than it being compact and capable of releasing tremendous energy on a short timescale.

An additional property of GRB progenitors, which strictly speaking is not required as part of the fireball model, is that the central engine is rapidly rotating. This assumption is motivated by the luminosity of the gamma-ray emission, which implies a tightly beamed jet rather than isotropic emission. (If the gamma-rays are emitted isotropically, the energy released would be equivalent to tens of solar masses, which would strain the bounds of theoretical astrophysics.) Beamed emission is known to be generated by collimated jets that are produced by magnetized, rotating systems, such as the accretion disk around a black hole. Thus, independent of the specific physics of the central engine, the sources of GRBs are likely to be compact, rapidly rotating objects. This makes GRBs good candidates for short-duration gravitational wave signals.

Observational evidence for beamed emission comes in the form of *jet breaks*, or the characteristic dimming of the GRB afterglow as the fireball slows and becomes nonrelativistic. As the Lorentz factor of the outflow approaches unity, the electromagnetic radiation is no longer beamed in the direction of the outflow, and the emission from any point in the fireball becomes isotropic in the frame of the observer. If the initial explosion is isotropic, the intensity measured by a distant observer will remain the same. But if the explosion is beamed,

the collimated emission from the jet will broaden as the jet becomes nonrelativistic, and the intensity of photons emitted in the original direction of the jet will decrease. This results in a sharp discontinuity of the GRB light curve, the jet break. The timing of the jet break after the GRB prompt emission can be used to determine the initial Lorentz factor of the jet, and from there the opening angle of the prompt gamma-ray emission (see Eq. 6 of [37], and references therein). The typical timescale for GRB jet breaks is $\mathcal{O}(1)$ day after the event, implying an opening angle of less than 10° .

Over the past two decades, astronomers have constrained the physics of GRB central engines through a variety of observations. The cosmological nature of GRBs was first established in 1997 by the *BeppoSax* satellite, which localized GRB 970228 to a galaxy at redshift $z=0.695$, a distance of 2.5 Gpc [51]. Soon after this discovery, *BeppoSax* astronomers associated long GRBs with core-collapse supernovae for the first time, with GRB 980425, also known as SN 1998bw [149]. Further observations have established a strong connection between long GRBs (duration > 2 sec) and type Ib,c core-collapse supernova. These supernovae originate from massive, rapidly rotating stars which have shed their outer envelope of hydrogen. Long GRBs tend to occur in young spiral galaxies with active star forming regions [39], consistent with the hypothesis that they are generated by CCSN from short-lived massive stars.

Short GRBs are a challenge to localize due to their short duration in the gamma-ray band ($\simeq 2$ sec) and their faint x-ray and optical afterglows. The first optical afterglow of a short GRB was detected by the *Swift* satellite, from GRB 050509B [68]. This and subsequent observations have not established a clear progenitor for short GRBs, but several properties have been established (see [37] for a review):

- Short GRBs originate from non-star-forming regions in late-type elliptical or irregular galaxies [60]. This implies a progenitor population of old stars.
- Short GRBs tend to be offset from their host galaxies by several kiloparsecs. Galaxy offsets are predicted for binary star systems in which one or both of the stars have undergone a supernova, giving the system an angular momentum ‘kick’ out of the galaxy.
- Short GRBs have smaller redshifts than long GRBs, with a median of $z = 0.2$ compared to $z = 2.0$ for long GRBs, after correcting for luminosity and detector sensitivity. Small redshifts are evidence for progenitors that require a long period of evolution before the GRB event.
- No short GRB has been associated with a supernovae, out to deep observing limits.
- One short GRB has been associated with an infrared afterglow consistent with a ‘kilonovae’, or a reddish re-brightening due to the radioactive decay of neutron-rich matter [131].

All of these properties are consistent with the model of BNS or NSBH mergers as the progenitors of short GRBs. Numerical simulations of BNS mergers have provided further evidence, in particular that BNS mergers can generate large magnetic fields capable of powering a relativistic jet along the axis of rotation [108, 110].

The beamed nature of GRBs has significant implications on their rate density. Call the rate of short GRBs R_{GRB} , and the rate of observed short GRBs R_{obs} . The opening angle of the relativistic jet, θ_j , is inversely proportional to the Lorentz factor Γ of the jet, and determines the fraction of GRBs whose gamma-ray emission is collimated in the direction of the Earth:

$$R_{obs} = (1 - \cos \theta_j) R_{GRB}. \quad (7.1)$$

As mentioned above, astronomers measure θ_j by observing the jet breaks in GRB afterglows. Jet breaks have been observed in many long GRBs and a handful of short GRBs, and both populations have median inferred opening angles of $\theta_j \sim 6^\circ$ (although the measurements for short GRBs are more uncertain). This implies that only one in a hundred short GRB events are aligned such that the gamma-ray flash is directed towards Earth.

For short GRBs, the observed beaming angles imply an overall rate that is in tantalizing agreement with the projected rate for CBCs discussed at the end of Chapter 2. Consider the following calculation, starting with the observed rate of short GRBs, $R_{obs} \simeq 3 \times 10^{-3} \text{ Mpc}^{-3} \text{ Myr}^{-1}$ [118]. Some fraction of short GRBs could be due to soft gamma repeater flares (SGRs) or other events in the nearby universe; call f_γ the fraction of GRBs due to a BNS or NSBH. We calculate the rate of CBCs giving rise to short GRBs as:

$$R_{CBC} = \frac{R_{GRB}}{f_\gamma} = \frac{R_{obs}}{f_\gamma(1 - \cos \theta_j)}. \quad (7.2)$$

For $\theta_j = 6^\circ$ and $f_\gamma = 1/2$, R_{CBC} is equal to the best guess for the rate of BNS mergers, $R_{merger} \sim 1 \text{ Mpc}^{-3} \text{ Myr}^{-1}$ [12]. While we should stress the large uncertainty in estimates of θ_j , f_γ , and the BNS rate, the fact that R_{merger} and R_{CBC} are in agreement for plausible values of θ_j and f_γ is encouraging. See [47] for a detailed discussion.

Whatever their precise mechanism, the luminosity of GRBs, combined with their extra-galactic distances, implies such an immense energy budget that GRBs are almost certainly associated with the birth of a black hole. Black holes are the most efficient converters of mass into energy, and they are essentially the only way a large fraction of a solar mass rest energy can be converted into electromagnetic radiation in a timescale of seconds.

7.2 GRB Satellite Missions

Photons in the gamma-ray portion of the spectrum do not penetrate the Earth's atmosphere, so GRBs are only observed by satellite experiments. (Some terrestrial detectors have been built to observe the Cherenkov radiation from a GRB-induced electromagnetic shower in the atmosphere, but their sensitivity does not match the spaceborne experiments.) There are several gamma-ray detectors onboard satellites in Earth orbit or elsewhere in the solar system. Most GRBs are detected by dedicated experiments on the *Swift* and *Fermi* satellites.

Swift, launched in 2005, is a dedicated GRB mission that uses a coded-mask gamma-ray detector to accurately localize the burst's position on the sky [67]. *Swift* takes its name from its rapid slew rate. The satellite can point its onboard X-ray and optical telescopes at the GRB location within minutes (sometimes seconds) of the event. The *Swift* Burst Alert Telescope (BAT) can localize a GRB to within several arcminutes on the sky, and x-ray or optical followup can localize the afterglow to within several arcseconds.

Fermi, launched in 2009, is a high-energy astrophysics mission with an observing window ranging from the soft X-ray to the hard gamma-ray end of the spectrum [90]. The Gamma-ray Burst Monitor (GBM) is an array of scintillation detectors onboard *Fermi* that detects GRBs with roughly half-sky coverage. The localization ability of the *Fermi*-GBM is much rougher than *Swift*; the 95% containment radius for *Fermi* triggers is typically 10-12 degrees [49, 74, 119].

After *Swift* or *Fermi* detect an event, the sky location, gamma-ray flux, energy spectrum and other data are transmitted to the ground and disseminated to astronomers via the Gamma-ray burst Coordinates Network (GCN). The information is typically available within minutes of the arrival of the gamma-ray burst at Earth.

7.3 Gravitational Waves from GRBs and Previous Searches

In this section we provide a non-exhaustive survey of theoretical predictions for gravitational waves from GRBs. The progenitors of GRBs are expected to have unique gravitational wave signatures [50], and most of the models of GRB central engines predict gravitational wave emission at high frequencies (> 500 Hz), where the search presented in this chapter was sensitive.

Short GRBs are believed to be associated with the inspiral and merger of a neutron star either with another neutron star or with a black hole. The inspiral phase of these mergers is expected to be a bright source of gravitational radiation [133], although most of the gravitational energy flux from the inspiral occurs at frequencies below 500 Hz. Numerical simulations of the merger phase have shown that substantial gravitational wave emission can occur at frequencies greater than 1 kHz [78, 115]. BNS mergers may result in the formation of a *hyper-massive neutron star* (HMNS), which can produce strong gravitational wave emission as it collapses to a black hole [99, 77].

The progenitors of long GRBs are believed to be core-collapse supernovae (CCSN) in rapidly rotating massive stars. Simulations of CCSN indicate several methods for gravitational wave emission at frequencies of several hundred Hz to 1 kHz, but the amplitude of the emission is highly uncertain [100]. The most optimistic emission models arise from the pulsations of a proto-neutron star core, which may release $10^{-7} M_{\odot} c^2$ in gravitational waves in a narrow frequency band around 1 kHz [101, 102].

Both types of GRBs are expected to result in the formation of a black hole with a rapidly-rotating accretion disk. Instabilities in the accretion disk can emit significant energy in gravitational waves, perhaps $10^{-2} M_{\odot} c^2$. The so-called bar mode instability, in the $l = 2, m = 2$ non-axisymmetric mode, is an optimistic model for gravitational wave emission. Typical frequencies are between 500 Hz and 2 kHz [66, 80]. An accretion disk that cools rapidly enough to become self-gravitating may fragment into one or more smaller bodies and generate an inspiral-like signal that persists to higher frequencies [105]. The micro structure of GRB light curves may be explained by precession of the accretion disk and the jet due to the Lense-Thirring effect [112] which can produce gravitational waves up to 1 kHz. Instead of an accretion disk, a torus may form around the black hole and convert the spin energy of the black hole into gravitational waves in the 1-2 kHz band [134, 135]. Numerical simulations have produced similar signals [116, 78]. Finally, the infall of matter from a rapidly-rotating accretion disk could produce long-duration gravitational waves in the 700 Hz–2.4 kHz band [106].

As an alternative to the merger scenario, a fraction of short GRBs could be produced by giant flares from a local population of *soft gamma repeaters* (SGRs) which are expected to produce some gravitational wave energy ($\lesssim 10^{-8} M_{\odot} c^2$) in the 1 kHz range [56, 130, 73, 89].

Although most GRB progenitors will be at distances too large for the resulting gravitational signals to be detectable by gravitational observatories, it is possible that a few GRBs could be located nearby. Recent studies indicate the existence of a local population of under-luminous long GRB with an observed rate density approximately 10^3 greater than the high-luminosity population [123, 86, 85, 46, 137, 72]. GRBs of this kind could be due to quenched jets or the breakout of relativistic shocks from a CCSN [103].

The LIGO and Virgo collaborations have performed several searches for gravitational waves coincident with GRBs. Most recently, data from the sixth LIGO science run (S6) and the second and third Virgo science runs (VSR2 and 3) were examined for generic gravitational wave bursts and for signals from CBCs associated with GRBs [17]. No signals were detected. In separate analyses, the non-detection of gravitational wave counterparts for GRB 051103 [16] and GRB 070201 [19], two short GRBs with error boxes overlapping the M81 galaxy at 3.6 Mpc and the Andromeda galaxy at 770 kpc respectively, ruled out the progenitor objects being CBCs in M81

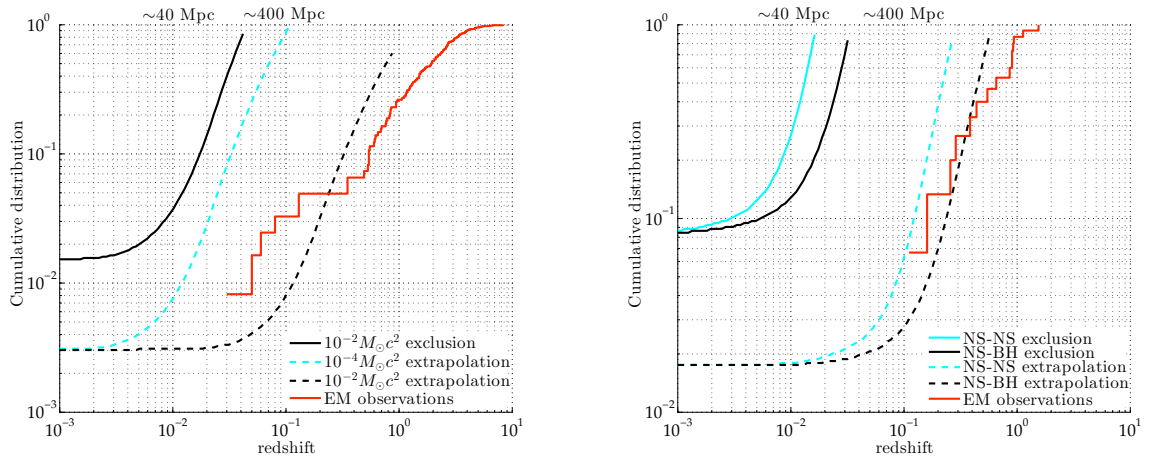


Figure 7.1. Comparison of results from previous searches for gravitational wave associated with GRBs and the observed redshift distribution for GRB progenitors. On the left, we plot lower limits on the distance to 150 GRBs from a search for unmodeled gravitational wave signals, assuming $E_{GW} = 10^{-2} M_{\odot} c^2$ in a narrowband signal around 150 Hz. On the right are lower limits to CBC events associated with 26 short GRBs. The dashed curves are extrapolations for the advanced detectors, assuming a 10x improvement in sensitivity and a 5x increase in the number of GRBs analyzed. The red curves are the redshift distributions for *Swift* GRBs (left: all GRBs, right: short-hard GRBs). Figure from [17].

or M31 with high confidence¹. A summary of recent lower limits on the distance GRB progenitors from LIGO-Virgo GRB searches is given in Fig. 7.1. In this figure, the exclusion distances set by the searches are compared to the observed distribution of GRB redshifts. A projection for five years of data from the advanced detectors indicates that a coincident detection with a short GRB is possible, if short GRBs are generated by BNS or NSBH mergers.

7.4 Formalism of Coherent Searches

Searches for gravitational waves from GRBs coherently combine the data from widely separated detectors to increase the significance of a signal and reduce the effect of noise. In this section, we present an outline of coherent analysis methods, especially in the context of searches for gravitational waves associated with GRBs. The algorithm that implements these methods is known as X-PIPELINE [127, 141] and it has been used in searches for unmodeled gravitational wave signals associated with GRBs since 2005 [13, 17, 6]. The coherent analysis performed by X-PIPELINE assumes that the gravitational signal is short duration, with a peak time known to within $\mathcal{O}(10)$ minutes and a sky location known to within $\sim 20^\circ$. Searches of this kind are referred to as *externally triggered* searches, since the event time and sky location are constrained by independent observations.

¹These non-detections support the hypothesis that some fraction of short GRBs are due to less energetic events in the nearby universe, rather than BNS or NSBH signals in distant galaxies.




Figure 7.2. Vector space of detector signals for the 3-detector case. The green plane is the space spanned by the antenna response vectors. Figure from [127].

As shown in Chapter 3, the sensitivity of a Michelson interferometer to gravitational waves is not isotropic, and the sensitivity of the detector to an event at a given sky location depends on the polarization of the gravitational waves. Taken by itself, the data from a single detector cannot distinguish the location of a transient gravitational wave source on the sky. Furthermore, the large background of glitches in gravitational wave detectors makes it unlikely that a single instrument could distinguish a short-duration event from noise.

These problems can be overcome by analyzing the data from a *network* of two or more detectors simultaneously. Recall our expression for the time-series data from a single detector due to a source at (θ, ϕ) from Chapter 3:

$$d(t, \theta, \phi) = h_+(t)F_+(\theta, \phi) + h_\times(t)F_\times(\theta, \phi) + n(t). \quad (7.3)$$

For a network of D gravitational wave detectors, we rewrite the elements of the equation as matrices:

$$\mathbf{d} = \mathbf{F}\mathbf{h} + \mathbf{n}, \quad (7.4)$$

where the matrices are given by:

$$\begin{bmatrix} d_1(t) \\ d_2(t) \\ \vdots \\ d_D(t) \end{bmatrix} = \begin{bmatrix} F_1^+ & F_1^\times \\ F_2^+ & F_2^\times \\ \vdots & \vdots \\ F_D^+ & F_D^\times \end{bmatrix} \begin{bmatrix} h_+(t) \\ h_\times(t) \end{bmatrix} + \begin{bmatrix} n_1(t) \\ n_2(t) \\ \vdots \\ n_D(t) \end{bmatrix}. \quad (7.5)$$

The detector antenna factors F_i^+ and F_i^\times are functions of the i th detector and the sky position of the GRB; recall Eq. 3.14 and Figs. 3.3 and 3.4.

We exploit the fixed dimensionality of the signal vector \mathbf{h} and the D -dimensionality of the network to construct linear combinations of the data vector \mathbf{d} which are sensitive to gravitational waves from the GRB sky location. In this way we construct a *detection statistic* that is particularly sensitive to gravitational waves from the GRB location. Furthermore, we can construct a linear combination of \mathbf{d} for which a gravitational wave signal will cancel. This *null vector* gives us a powerful method for distinguishing signal from noise. This is illustrated in Fig. 7.2, for a network of three detectors.

Note that since the detector noise is uncorrelated between detectors, the vector \mathbf{n} has an identity covariance matrix, which is invariant under change of orthonormal basis. In other words, no linear combination of \mathbf{n} should be special.

In the D -dimensional space, a signal \mathbf{h} is restricted to the hyperplane defined by the column space of the matrix of antenna factors, \mathbf{F} . This is the plane spanned by the \mathbf{F}^+ and \mathbf{F}^\times vectors. The unit vectors of the hyperplane are $\mathbf{e}^+ = \mathbf{F}^+ / |\mathbf{F}^+|$ and $\mathbf{e}^\times = \mathbf{F}^\times / |\mathbf{F}^\times|$. The *null vector* orthogonal to the $(\mathbf{e}^+, \mathbf{e}^\times)$ hyperplane is constructed using the vector cross product, $\mathbf{e}^n = \mathbf{e}^+ \times \mathbf{e}^\times$. We can measure the content of the data vector \mathbf{d} in each of these directions by taking the dot product, $(\mathbf{e}^{(+,\times,n)} \cdot \mathbf{d})$.

We will measure our signal strength using the *coherent energy*, defined for a projection \mathbf{e}^α as:

$$E_\alpha = |\mathbf{e}^\alpha \cdot \mathbf{d}|^2 = \sum_{i,j=1}^D e_i^{\alpha*} e_j^\alpha d_i d_j^*. \quad (7.6)$$

We will also make use the *incoherent energy*:

$$I_\alpha = \sum_{i=1}^D |e_i^\alpha|^2 |d_i|^2. \quad (7.7)$$

For a true gravitational wave signal, the coherent energy for the projections \mathbf{e}^+ and \mathbf{e}^\times will be much larger than the coherent energy in the null direction, \mathbf{e}^n , where the contribution of a signal vanishes. For a glitch that occurs in only one of the detectors, the coherent energy in the null direction \mathbf{e}^n will be roughly equal to the energy in the $(\mathbf{e}^+, \mathbf{e}^\times)$ plane. Thus we can distinguish signal from glitches by making a cut on the ratio between the coherent energy in the $\mathbf{e}^{+,\times}$ direction and the null direction. In practice, using the ratio of the coherent null energy to the incoherent null energy is more robust. Examples of how this cut is applied are given in Appendix C.

Using the unit vectors for the two polarizations, we can construct other directions, for example those that align with left or right circularly polarized gravitational waves:

$$\begin{aligned} \mathbf{e}^\circlearrowleft &= \mathbf{e}^+ + i\mathbf{e}^\times, \\ \mathbf{e}^\circlearrowright &= \mathbf{e}^+ - i\mathbf{e}^\times. \end{aligned} \quad (7.8)$$

For GRBs, we can assume that we are viewing a rotating system along the axis of angular momentum. The gravitational radiation from such a system will be circularly polarized, $h_{\times} = \pm ih_{+}$. This is especially useful if the detector network is composed of only two misaligned detectors, i.e. detectors with different values for F_i^{+} and F_i^{\times} . In this case the plane formed by F^{+} and F^{\times} spans the space of possible data vectors and there is no null direction. In order to distinguish signal from noise, we need to make an assumption about the signal morphology. The assumption of circular polarization is well-motivated by GRB physics and allows us to define the ‘circular null’ direction as orthogonal to the direction that maximizes the coherent energy:

$$E_{n\circ} = \min(|e^{\circ} \cdot \mathbf{d}|^2, |e^{\circ} \cdot \mathbf{d}|^2). \quad (7.9)$$

The detection statistic used to rank events is a Bayesian log-likelihood ratio of the probability of circularly polarized gravitational signal to the probability of pure noise with a prior on signal amplitude. This is called the ‘loghbayesiancirc’ statistic (where ‘h’ indicates the prior on the signal amplitude). The statistic is given by, for example in the right-circularly-polarized direction:

$$2L(\mathbf{d} | \circ, \sigma_h) = \left[\frac{|e^{\circ} \cdot \mathbf{d}|^2}{1 + 1/(\sigma_h |\mathbf{f}^{\circ}|)^2} \right] - \log(1 + \sigma_h^2 |\mathbf{f}^{\circ}|^2), \quad (7.10)$$

where \mathbf{f}° is the noise-weighted unit vector in the right-circular-polarization direction, and σ_h is an amplitude that corresponds to our prior expectations for a plausible signal. The final statistic is calculated by marginalizing over σ_h . A similar expression is calculated using \mathbf{f}^{\times} . The noise weighting accounts for differences in sensitivities between the detectors as a function of frequency. A derivation of the statistic is given in Appendix B.

In practice, the detection statistic is calculated using the Fourier transforms of the time-series data from each detector. The time-frequency maps that are generated by the Fourier transform are projected onto the left and right circular polarizations; the direction with maximum coherent energy is taken to be the signal direction, and the direction with minimum coherent energy is the null direction. An example time-frequency map for the loghbayesiancirc statistic for a simulated gravitational wave signal is given in Fig. 7.3. The signal was added to live data from the H1 and L1 detectors as a “blind injection”, intended to test the effectiveness of the search pipelines. Note the chirping shape of the waveform, as the amplitude and frequency increase with time. At high frequencies (above 250 Hz) the sensitivity of the detectors degrades and the signal is no longer distinguishable from noise. The simulated waveform corresponds to an NSBH merger about 25 Mpc away. The distance was chosen to provide a signal with reasonable SNR in the initial detectors, and is quite optimistic when compared with expectations from astrophysical population studies.

Signal events are distinguished from noise using the ratio of the coherent energy to the incoherent energy, in the null direction. This is illustrated in Fig. 7.5. In this figure, events that are due to true gravitational signals




Figure 7.3. Time-frequency map for the ‘loghbayesiancirc’ detection statistic for a simulated gravitation wave event. The waveform is an NSBH inspiral. The color of each pixel indicates the value of the detection statistic. Adjacent time-frequency pixels that lie in the 99th percentile or larger are used to calculate the total event significance. For this event the sum of the significant pixels is 126.




Figure 7.4. Energy and incoherent energy in the circular energy plane for an example analysis with the H1-L1 detector network. The x- and y-axes have arbitrary units of non-normalized signal energy. The starred event is the signal shown in Fig. 7.3. The many hundreds of small crosses indicate background events. Background events and signal usually have roughly equal energy and incoherent energy in the non-null direction.




Figure 7.5. Null energy vs. null incoherent energy for an example analysis with the H1-L1 detector network, in units of non-normalized signal energy. The starred event is the signal shown in Fig. 7.3. The purple dashed line indicates the threshold for the coherent cut. Events lying below this line are inconsistent with a gravitational wave signal from the sky location of the GRB and are rejected by the search.

have very little coherent energy in the null direction, and lie above the diagonal. Loud background events, due to glitches in one detector or another, have roughly the same coherent and incoherent null energy and lie on the diagonal. In this way we can define *coherent cuts* which distinguish between signal and noise. The coherent cut threshold is shown by the dashed purple line. Events lying below this line are inconsistent with a gravitational signal from the GRB sky location. The cut threshold is tuned independently for each GRB analysis to maximize the sensitivity to simulated gravitational signals; this is necessary to account for different sky locations, detector sensitivities, and background distributions, which can vary substantially from one GRB to another. A detailed description of the coherent cut method used by X-PIPELINE is given in Appendix C.

Data from the detectors surrounding the time of a GRB are divided into an *off-source window*, which is used to characterize the background of transient signals around the time of the GRB, and an *on-source window*, which is searched for gravitational wave events. To allow for possible gravitational precursors from, for example, the CCSN associated with long GRBs, the on-source window is $[-600,+60]$ sec around the start time of the gamma-ray signal observed by a satellite experiment. For very long-lasting GRBs, the on-source window is extended to include the entire T_{90} time of the event, defined as the time interval over which 90% of the total background-subtracted photon counts are observed. The standard off-source window is ± 1.5 hours around the time of the GRB trigger, excluding the on-source window.




Figure 7.6. Foreground and background event rates for an example GRB analysis. The traces are cumulative histograms that count the number of events over a particular significance threshold. The significance of each event is the value of the detection statistic given in Eq. 7.10. Traces for the on-source (foreground) and off-source (background) event distributions are shown, before and after the coherent cuts are applied. The on-source has an outlier, with a significance of about 120. This is the NSBH signal shown in Fig. 7.3.

To estimate the rate of background events, the analysis of the off-source data is repeated many hundreds of times with unphysical *time-slides* of >3 sec applied to the data from one of the detectors. This technique is used to estimate the probability that random noise fluctuations in the detectors may appear to be a true gravitational wave signal in the on-source window. A typical search with X-PIPELINE will perform $\mathcal{O}(10^3)$ time-slides on the data in the off-source window, enough to quantify the rate of background events to a *false-alarm rate* (FAR) below 10^{-6} Hz. Repeating the analysis for hundreds of time-slides is the most computationally intensive portion of the search.

To determine if a gravitational wave is present in the data, the loudest on-source event is compared to the distribution of off-source events. The *false-alarm probability* (FAP), or *p-value* of this event is defined as the fraction of off-source events with equal or greater significance. This is an empirical measure of the probability of obtaining such an event in the on-source under the null hypothesis. Events with $p < 0.01$ are followed up with detailed investigations to determine if the events can be associated with noise artefacts in the detectors.

Examples of the on-source and off-source event distributions for an X-PIPELINE analysis are shown in Fig. 7.6. This analysis contained the event shown in Fig. 7.3 in its on-source region. The blue trace is the distribution of background events before the coherent cuts are applied. Many of the background events are determined to be inconsistent with the signal model and are cut, leaving behind the red trace. The purple trace is

the distribution of on-source events following the coherent cuts. The significance of the loudest on-source event is 126, and no event in the background (red trace) has a significance above 90. Since the foreground event has a significance larger than all of the estimated background, we place an upper limit on its p-value, using the total area under the off-source curve, after vetoing: $p \lesssim 10^{-4}$.

To claim evidence for a detection, the analysis must estimate a sufficient number of background trials to estimate the p-value to less than 3σ , or $1/370$. To claim a detection at the 5σ threshold we would need to analyze $\mathcal{O}(10^6)$ background trials. Due to the computational cost of the analysis, the standard search performs a background analysis sufficient for p-value estimation out to 3σ . GRBs with a significant outlying event are followed up with a full background estimation out to the 5σ level.

For a search that includes many GRBs analyzed together we must account for the *trials factor*, sometimes called the *look-elsewhere effect*. This is a multiplicative factor, equal to the number of GRBs in the search, that is applied to the individual p-values returned by the collection of individual, independent searches. To quickly motivate the trials factor, note that if p-values are randomly drawn from a uniform distribution, and 100 GRBs are analyzed in a search, the probability to measure a $p = 0.01$ in at least one of the individual analyses is quite high. This we must weight the significance of any one event by the probability to observe an event of that significance is a set of N independent searches.

If no single GRB analysis contains a significant event, the set of all GRBs in the search can be tested for an excess of sub-threshold events. This is referred to as the *weighted binomial test*, described in the appendix of [17]. The weighted binomial test calculates the probability to observe the collection of p-values returned by a population of independent analyses, while de-weighting the significance of GRBs whose analysis had relatively poor sensitivity to gravitational waves. Since the test collects the results from all of the GRBs analyzed in the search and returns a single p-value describing the significance of the foreground events, it usually presented as the final, one-line summary of the significance of the search results.

7.5 Sky Localization Uncertainty

The sky localizations of GRBs detected by the Fermi GBM can have uncertainty regions covering hundreds of square degrees depending on the gamma-ray flux and energy spectrum. In a coherent search for gravitational waves associated with GRBs, performing the analysis using an incorrect sky location can reduce the significance of a signal in two ways.

First, the analysis will incorrectly estimate the sensitivity of each detector to gravitational waves from the sky location of the GRB. This can result in a loss of coherent signal energy when the time-frequency maps from each detector are combined. Over most of the sky, the antenna factors vary slowly, and the loss of signal is of order a




Figure 7.7. Event significance as a function of phase shift error between two detectors, for three sets of narrow-band signals with different central frequencies. When the error in the time shift between detectors corresponds to a multiple of 180° , the coherent energy is essentially nulled, and the signal will not be detected by the search. This plot was generated using a search grid of a single point and scattering the simulated waveforms widely across the sky.

few percent. The effect of this small variation on the analysis has been checked with a variety of empirical tests, and we conclude that this effect is not large enough to alter the results of our search.

Second, and more significantly, an error in the sky location will lead to an incorrect time-shift of the data when synchronizing the arrival time of a gravitational signal across detectors. For pairs of ground-based detectors the difference in arrival time for a gravitational wave signal is $\mathcal{O}(10)$ milliseconds, and an error in the sky location of a few degrees could introduce incorrect synchronizations of a millisecond or more. For waveforms with frequency content above 1 kHz this results in the misalignment of the gravitational signal by several periods, and when the data vectors are combined the coherent signal energy will be diminished. In the worst case, the waveform will be shifted by a half-period between the detectors and the signal will cancel entirely in the coherent summation. An illustration of signal loss due to a mismatch between the searched sky location and the event location is shown in Fig. 7.7.

The standard solution in gravitational wave searches is to repeat the analysis over a discrete grid of sky positions covering the uncertainty region. The grid step is chosen such that the timing synchronization error between any position in the GRB error box and the nearest grid point is less than 25% of the period for the highest-frequency gravitational wave signals included in the search. The time delay between the detectors is a function of sky position, but for simplicity the grid step size is held constant across the search area, at the smallest

step size required within the uncertainty region. For uncertainty regions with radii of $\mathcal{O}(10)$ degrees the variation in the magnitude of the time-of-arrival correction is not large enough to warrant a variable grid spacing.

For two detectors separated by a distance d , the difference in time of arrival of the gravitational wave is:

$$t = \frac{d \cos \theta}{c}, \quad (7.11)$$

where θ is the angle between the inter-detector baseline and the line-of-sight to the GRB and c is the speed of light. Call α the value for 25% of a full period for the highest frequency included in the search. We require that the maximum time delay error dt is less than α for every point within the uncertainty region. Thus the spacing $d\theta$ between grid points is:

$$|d\theta| \leq \frac{2c}{d \sin \theta} \alpha. \quad (7.12)$$

Previous GRB searches with LIGO and Virgo data have used regular grids of concentric circles around the best estimate of the source location, covering at least 95% of the sky location probability distribution. For the *Fermi* GBM, the 68% containment radius is typically $2^\circ - 3^\circ$ due to statistical effects, but the localizations have additional systematic errors of several degrees. As a result, the 95% containment region can cover hundreds of square degrees, and a search for gravitational signals with frequencies larger than $\mathcal{O}(100)$ Hz will require many hundreds of search points. Since the search is essentially repeated for every point in the grid, the computational costs can be unrealistic, even for computing clusters with thousands of CPU cores.

For a two-detector network, the problem can be simplified. Within the GRB error box, the time shift between the detectors will only change across one dimension. To minimize the computational cost of the search, we are free to cover the search region with a linear search grid, arranged parallel to the maximum gradient of change in the time-of-arrival between detectors. In the case of a 2-detector network, this linear pattern is sufficient to capture the dominant source of coherent energy variability as the likelihood is calculated across the GRB uncertainty region. A comparison of the circular and linear search grids for the Fermi event GRB 080906B is shown in Fig. 7.8.

We verified the effectiveness of the linear search grid by analyzing a handful of GRBs using both the linear and circular tiling. The GRBs in our test sample were chosen to represent a variety of grid densities, detector antenna responses, and sizes of the 95% containment region. The results for the two methods were nearly identical. For simulated gravitational signals, the minimum amplitude for detection was the same for each method within a few percent. Furthermore, the analysis using the linear grid was completed in a fraction of the time required for the circular grid, and typically required $\mathcal{O}(10^3)$ or fewer CPU hours. Using computing clusters with thousands of CPU cores, it was possible to analyze some GRBs localized by the GBM in less than three hours.




Figure 7.8. Example linear and circular grids for a search for gravitational wave signals up to 1792 Hz. The localization for the *Fermi* GBM event GRB 080906B is shown. The linear search grid contains 41 sky positions, arranged in the direction of the gradient of the time shift between the H2 and G1 detectors. The circular grid contains 1324 sky positions and would require several days to analyze on a massively parallel computing cluster.

7.6 A Search With Data From GEO 600

We have applied the linear grid technique in a search for gravitational waves associated with 152 GRBs detected between June 2006 and September 2011. These GRB events occurred during times when the GEO 600 detector in Hannover, Germany was operating in coincidence with one other kilometer-scale detector. Previous searches for gravitational waves associated with GRBs focused on the best sensitivity region of the kilometer-scale detectors, typically around 100-200 Hz. GEO 600 reaches its best sensitivity above 500 Hz, and for searches using GEO data it is sensible to extend the search bandwidth above 1 kHz to explore the entire sensitive band (see Fig. 7.9). The density of search grid points required by this frequency range would have made a circular search pattern computationally infeasible; this was the motivation behind the development of the linear search grid.

The times and localizations of the GRB events were obtained from the Gamma-ray burst Coordinates Network (GCN) [36], supplemented by the *Swift* and *Fermi* on-line catalogues², as well as the published *Fermi* four-year catalogue [138]. Most of the GRBs in our sample were detected by *Swift* and *Fermi*; a few of the GRBs were detected by other space borne experiments such as *INTEGRAL* [147], *AGILE* [58], or *MAXI* [88].

We analyzed GRBs which were observed when GEO 600 plus one other observatory was taking science-quality data. For the LIGO and Virgo interferometers, *science mode* is a rigorous definition, and identifies

²<http://heasarc.gsfc.nasa.gov/W3Browse/fermi/fermigbrst.html>, http://swift.gsfc.nasa.gov/archive/grb_table




Figure 7.9. Demonstrative noise spectra for the detectors used in The GEO-GRBs search. GEO’s best sensitivity is around 500 Hz in the early epoch, and above 1 kHz in the late epoch when quantum-noise squeezing was applied to reduce the shot noise [14].

times when the detector configuration is stable and the interferometer is operating in a resonant, low-noise state. GEO 600 has no strictly-defined science mode, and collects data on an opportunistic basis between commissioning activities.

In our search, we made no distinction between short GRBs and long GRBs, and the analysis was performed without regard to the observed GRB fluence or redshift (if known). As part of our data quality checks, we excluded segments of the data which were flagged as having poor noise quality, and we only analyzed GRBs with sufficient good-quality data around the time of the event. We also discarded GRBs that were determined to have exceptionally high rates of background events or exceptionally poor sensitivity to gravitational waves from the sky location of the GRB. This, for example, can result from environmental or instrumental noise at the time of the GRB (identified using methods described in [38, 8, 10, 120, 75]), or a GRB sky location that includes one of the sensitivity null points of the detectors.

In addition to these sources of noise, the sensitivity of the GEO 600 detector can change by 20% or more at frequencies >1 kHz, depending on whether squeezed light states are being injected. A change in sensitivity of this magnitude may bias the background estimation if it occurs partway through the off-source window around a GRB. In this analysis, no GRB off-source (or on-source) window includes time when GEO 600 changed from a squeezing to a non-squeezing state, or vice versa. As final check on data quality, we inspected the detector noise

characteristics at the time of each GRB by hand, and verified that the sensitivity of the detectors in the off-source window agreed with the on-source window. In this way we were confident that the background estimation was not biased compared with the foreground.

In the epoch of our search (Feb 4th 2006 to Nov 3rd 2011), there were 152 GRBs with sufficient science data to analyze. For 130 GRBs had good sensitivity to potential gravitational wave signals based on the results of the background estimation. For each of these GRBs we calculated the p-value for the loudest event in the on-source. Three GRBs in our sample had on-source events with $p < 0.01$:

- GRB 060502A, a *Swift* BAT detection with $T_{90} = 28.4$ sec and an observed redshift of $z = 1.51$ [52], was analyzed using data from the L1 and G1 detectors. There were three significant events in the on-source window. An examination of the data quality around the time of the GRB revealed non-stationary noise in the L1 detector associated with ground motion due to a magnitude 5.0 earthquake in Costa Rica. All three on-source events occurred during a segment of time that was identified *a priori* as likely to experience an increased rate of transient signals due to elevated seismic noise. Since 35% of the on-source window for this GRB was flagged as having elevated ground motion, we vetoed the three events and do not include this GRB in the cumulative results.
- GRB 080816A, a *Fermi* GBM detection with $T_{90} = 4.6$ sec [138], was analyzed using data from the H2 and G1 detectors. There was one significant on-source event, with $p = 0.001$. A signal processing algorithm revealed multiple instrumental channels in the H2 detector with excess noise at the time of the event. The nature of the instrumental noise is not understood, but the signal in the gravitational wave channel is unlikely to be of astrophysical origin. No redshift observations are available for this GRB.
- GRB 090712A, a *Swift* BAT detection with $T_{90} = 145$ sec [138], was analyzed using data from the G1 and V1 detectors. There was one significant on-source event, with $p = 0.003$. While we found no plausible instrumental or environmental cause for the signal, the observed p-value for this GRB is not significant in a data set containing 129 GRBs. No redshift observations are available for this GRB.

The distribution of p-values for the most significant event in the on-source window for the 129 GRBs is shown in Fig. 7.10. To test the sample of GRBs for a population of sub-threshold signals, we used the weighted binomial test to check that the distribution of p-values was compatible with the uniform distribution expected from the null hypothesis. The test yielded a background probability of 0.193, indicating that the distribution was consistent with no gravitational events being present. An illustration of the statistical significance of our observed p-value distribution is shown in Fig. 7.11, where we compare our results to a collection of pseudo-experiments. For 129 events we select a random number between zero and 1, following a uniform distribution. This is repeated




Figure 7.10. Distribution of p-values for the loudest event in each of the 129 GRBs analyzed in the search.

for 500 trials and plotted in the gray points. The observed p-value distribution is completely contained by the distribution of p-values from the pseudo-experiments.

As part of the analysis, we measured the sensitivity of the search to simulated gravitational wave signals, as a function of amplitude. We simulated gravitational waves from GRBs using circularly polarized sine-Gaussians with quality factor $Q = 9$ (recall Eq. 2.44). To account for errors in the detector calibration, timing, and the GRB sky location, we jittered the simulated waveforms in amplitude, phase, and central time before adding them to the detection data. The magnitude of the jitter was Gaussian-distributed with variance proportional to the calibration uncertainties of each detector. For GRBs detected by the *Fermi* GBM, the sky positions of the simulated waveforms were distributed according to the systematic uncertainties of the GBM detector [49]. This sky position jittering was performed across the entire uncertainty region, and was not restricted to the axis of the linear search grid. For each GRB analysis, several hundred simulated signals were generated for a range of amplitudes, and the overall detection efficiency was measured using the entire population of jittered waveforms. In this way we estimated the search sensitivity and marginalized over the uncertainty in detector calibration and the GRB sky localization.

For each GRB we calculated the total amplitude in gravitational wave-induced strain that would result in a detection for 90% of the simulated signals. In terms of gravitational wave strain amplitude h_{rSS} , the median 90% upper limit for our GRB sample was $2.8 \times 10^{-21} \text{ Hz}^{-1/2}$ for circularly-polarized sine-Gaussian signals at 500 Hz and $3.4 \times 10^{-21} \text{ Hz}^{-1/2}$ at 1 kHz. The search sensitivity to signals below 300 Hz was limited by the sensitivity




Figure 7.11. Illustration of possible p-value distributions using 500 pseudo-experiments, each with 129 GRBs. The results from the search, in red, fall within the distribution of results from the pseudo-experiments, in grey.




Figure 7.12. Distribution of upper limits on strain amplitude for gravitational wave signals from the 129 GRBs analyzed in the GEO-GRBs search. The simulated gravitational wave signals are $Q = 9$ sine-Gaussians with central frequencies 500 Hz (red) and 1 kHz (black). The upper limits have been marginalized over detector calibration uncertainties and GRB sky localization uncertainty. Figure from [11].




Figure 7.13. Distribution of lower limits on distance to gravitational wave signals from the GRBs in the search. A standard gravitational wave emission energy of $10^{-2} M_{\odot}$ is assumed, for sine-Gaussian waveforms with central frequencies 500 Hz (red) and 1 kHz (black). Figure from [11].

of the GEO 600 detector and by the choice of coherent cut thresholds, which were tuned to minimize the effect of nonstationary noise in GEO at low frequencies. In general low-frequency signals were not detectable by the search, including the inspiral waveform from a BNS or NSBH merger. Our upper limits on gravitational wave strain at 500 Hz and 1 kHz for the 129 GRBs are shown in Fig. 7.12.

If we assume a fixed gravitational emission energy, we can use the upper limits on h_{rss} to calculate lower limits on the distance to the GRB, using Eq. 2.46:

$$D_{\text{excl}} = \sqrt{\frac{5}{2}} \sqrt{\frac{G}{\pi^2 c^3}} \frac{\sqrt{E_{GW}}}{f_0 h_{rss}}. \quad (7.13)$$

Here, E_{GW} is the energy released by the GRB central engine in gravitational waves, G is Newton's constant, f_0 is the central frequency of the rigid rotator model, and h_{rss} is the strain amplitude upper limit observed by the search. Recall that the factor of $\sqrt{5/2}$ arises from the beamed gravitational emission from a rotating system, viewed on-axis. The lower limit on the distance to a GRB, using an assumption of the energy emitted in gravitational waves, is referred to as the GRB *exclusion distance*.

The distribution of exclusion distances for waveforms with central frequencies of 500 Hz and 1 kHz for the 129 GRBs is shown in Fig. 7.13. The median exclusion distance for the 500 Hz and 1 kHz waveforms are 0.8 Mpc and 0.3 Mpc respectively, where we have assumed the GRB central engine releases $E_{GW} = 10^{-2} M_{\odot} c^2$ total energy in gravitational waves. The distance limits scale with the square root of the assumed emission energy,

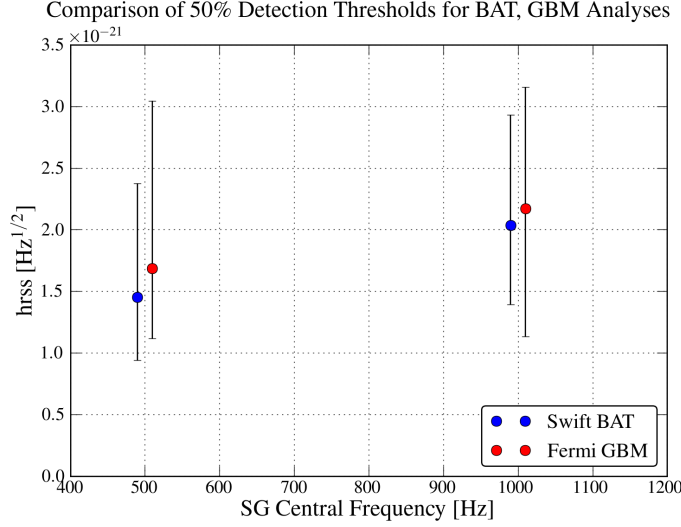


Figure 7.14. Upper limits on gravitational wave strain amplitude for sine-Gaussian waveforms at 500 Hz and 1 kHz, for *Swift* and *Fermi* events. The upper limits are largely the same; this indicates that the sky localization of the *Fermi* events (and the use of the linear search grid) did not affect the sensitivity of the search.

$E_{GW}^{1/2}$, and for a pessimistic assumption of $E_{GW} = 10^{-8} M_{\odot} c^2$ the median exclusion distance becomes 0.8 kpc at 500 Hz. For comparison, the GRB in our sample with the smallest observed redshift is GRB 080905A, with $z = 0.1218$ [113] or $D \simeq 590$ Mpc [150].

As a check to verify that the linear grids did not degrade the sensitivity, we examined the upper limits from the sample of *Swift* and *Fermi* GRBs separately. The results are shown in Fig. 7.14. There is no significant difference between the populations.

Finally, we verified that the search sensitivity was limited by the sensitivities of the detectors the GRB sky locations. Excursions from this hypothesis might indicate problems with the coherent cut thresholds or particularly bad distributions of background noise. The results of the check are shown in Fig. 7.15, which plots the ratio of the upper limits from each of the 129 GRB analyses to the sensitivity of the least-sensitive detector. The results are plotted as a function of the ratio of the sensitivities of the two detectors in the search. When the two detectors are well-matched, the upper limits are a factor of 6-8 above the detector noise floor, corresponding to a detection threshold SNR of about 8-12. This detection threshold SNR is in good agreement with past experience from gravitational wave searches, indicating that the sensitivity achieved by the search was in agreement with reasonable expectations, given the sensitivity of the detectors. As the detector sensitivities diverge, the upper limits of the search approach the level of the noise floor in the least sensitive detector.




Figure 7.15. Ratio of upper limits on gravitational wave strain amplitude to the least sensitive detector, as a function of the ratio of detector sensitivities.

7.7 Sky Localization for *Fermi*-GBM Events

One of the primary goals of gravitational wave searches is the prompt localization of the source sky location for follow-ups by electromagnetic (EM) astronomers. Currently, very few GRBs detected by the *Fermi* GBM are examined for optical counterparts, due to the telescope resources that are required to search an uncertainty region of hundreds of square degrees. The detection of a gravitational wave signal associated with a GRB will be of tremendous interest to the astronomical community, and any improvement of the GBM localization can increase the chances that astronomers will detect an optical, radio, or X-ray counterpart.

If a gravitational wave signal is detected in the on-source window, X-PIPELINE can localize the source to within a few degrees along the axis of the linear grid, depending on the frequency content of the signal and its duration. The reconstructed sky location is the point on the search grid that maximizes the coherent energy of the signal. For a search using data from two widely-separated detectors, the localization using the gravitational wave signal is limited to an annulus on the sky encircling the line connecting the two detectors [4]. In principle, this localization cannot be improved in the direction perpendicular to the search grid. A hypothetical localization by the linear grid for a typical *Fermi* GBM event is illustrated in Fig. 7.16.

For searches using data from three or more widely-separated detectors, a gravitational wave signal can be localized in both dimensions, and in this case the computational cost of the full circular tiling may be worthwhile. Efforts to characterize X-PIPELINE’s ability to localize signals using three or more detectors are ongoing.

We characterized X-PIPELINE’s localization errors using a linear search grid for a handful of GRBs. A typical result is given in Fig. 7.17, for an analysis with the H1 and V1 detectors. We plot the observed errors




Figure 7.16. Sky location reconstruction with the linear grid. Figure from [11]. The true GRB location is marked by the star. The red dot is the sky localization from the *Fermi* GBM, and the black circle illustrates a typical 95% containment region around the GBM localization. A GRB search with X-PIPELINE could further localize a detected signal along the linear search grid, with a localization uncertainty illustrated by the grey band.

for 600 simulated short-duration gravitational wave signals with central frequencies of 150 Hz and 1.5 kHz in a cumulative histogram. High-frequency signals have better localizations, due to larger sensitivity to time delays as the sky position changes across the grid. For signals at 1.5 kHz, 95% of the simulations were reconstructed to less than 2.5° along the axis of the linear grid. This provides an empirical measurement of the 2σ uncertainty in X-PIPELINE's localization.

The usefulness of improved GRB localizations using coincident gravitational wave detections is limited, but important. One of the primary science goals of the advanced LIGO and Virgo instruments is an independent measurement of the Hubble constant, using BNS mergers to directly measure the source luminosity distance (D_L) and redshift (z). Recall that the Hubble constant is given by:

$$H_0 = c \frac{z}{D_L} \quad (7.14)$$

for $z \ll 1$. The waveform of a BNS merger can be parametrized in terms of the *chirp mass*, defined as a function of the masses of the binary system m_1 and m_2 :

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (7.15)$$




Figure 7.17. Accuracy of sky localization for various waveforms, for the H1-V1 detector network in S6 noise. The first four waveforms are $Q=9$ sine-Gaussians with the central frequencies given in the legend. High-frequency signals can be localized to roughly 1° in the direction parallel to the linear grid. Waveforms from CBCs are longer-duration and can be localized to less than 1° .

To first order, the phase evolution of the waveform and the absolute amplitude are determined by the chirp mass. For this reason, BNS inspirals are sometimes called *standard sirens*: the absolute amplitude of the signal can be constrained by the signal itself. The luminosity distance is then directly calculated from the observed amplitude.

There are two challenges to this approach, which make a direct calculation of the Hubble constant difficult without some assumptions about the BNS system. First, the phase evolution (and thus the chirp mass) of the waveform is degenerate with the redshift of the source, and more highly redshifted systems will appear to have higher mass. Second, the observed amplitude of the waveform is inversely proportional to the inclination of the binary system, and systems that are tilted with respect to the line of sight will appear dimmer. A variety of solutions to these challenges have been proposed to enable a measurement of the Hubble constant from BNS detections alone, for example the source redshift can be constrained by galaxy distributions from catalogs [107] or the chirp mass can be assumed to follow the distribution of known neutron star masses [132]. Each of these solutions relies on assumptions of the source population which add uncertainty to the overall measurement.

These assumptions can be avoided if a population of BNS inspirals is detected in coincidence with short GRBs. In this case, the mass/redshift degeneracy is broken using observations of the GRB host galaxy, and the distance/inclination degeneracy is broken by assuming a BNS associated with a GRB is viewed nearly on-axis ($\iota \simeq 0$). Nissanke et al. have shown that a measurement of H_0 to better than 10% accuracy can be made with as few as twenty coincident detections [98].

The crucial step in this method is the localization of the source host galaxy, which requires an observation of the GRB afterglow. A localization with X-PIPELINE could significantly improve the chances for a detection and ultimately a measurement of H_0 . Targeted searches like X-PIPELINE can have sensitivity up to two times better than the equivalent all-sky search [141], and while the analysis has greater latency than realtime all-sky searches, the sensitivity improvement from a targeted search could make the difference between a missed signal and a detection candidate. In principle, a gravitational wave signal detected by X-PIPELINE could reduce the localization uncertainty from the *Fermi* GBM within a few hours of the GRB detection, quickly enough that electromagnetic astronomers could detect an optical afterglow.

CHAPTER 8

CONCLUSION

*If we do meet again, why, we shall smile;
If not, why then this parting was well made.*

This thesis has endeavored to cover the broad field of experimental gravitational-wave astrophysics, from instrumentation, to detector characterization, and finally data analysis.

In the coming years, the advanced LIGO and Virgo instruments will begin to take data at unprecedented sensitivity. In Fig. 8.1 we visualize the expected number of detections for a two-detector network with a given average sensitivity to BNS mergers, assuming the rate density predictions in Ref. [12]. The initial detector observing runs, labeled as ‘S5+S6’, were not expected to make any detections, even for the most optimistic merger rates. The second observing run, expected to start in September 2016 and last six months, could provide the first detections if the ‘best guess’ rate estimates are true. As we pointed out in Chapter 7, these ‘best guess’ rates are in agreement with the rate of short GRBs, which are presumed to be due to BNS and NSBH mergers.

Despite this optimistic future, we must keep in mind that advanced LIGO is a detection machine. Even with optimistic rates, the SNR of BNS events in the aLIGO detectors will generally be between ten and thirty, and SNRs greater than forty are unlikely. These signals will be too faint to constrain the sort of fundamental questions these detectors were built to answer. For example, measurements of the neutron star equation of state require high-SNR signals, either to quantify the tidal disruption of the stars in the final orbits before the plunge, or detect ringing of the body modes of the hyper-massive neutron star that results from the merger.

Another example: constraining the mechanism for core-collapse supernovae requires a CCSN detection. Models of gravitational wave emission from CCSN are not optimistic, and advanced LIGO is likely not sensitive enough to detect a CCSN signal unless the event is within the Milky Way. To increase the probability of a detection to better than one per century the instruments must be upgraded to detect signals beyond the Local Group.

Finally, gravitational waves from BNS mergers could provide an independent measure of the Hubble expansion, if the distance-inclination and mass-redshift degeneracies can be broken (see the discussion in Chapter 7).




Figure 8.1. Anticipated number of BNS detections for proposed advanced LIGO observing runs. The average detector sensitivity to BNS mergers is expressed in megaparsecs (x-axis). The expected number of detected signals is calculated using estimates of the rate density for BNS mergers, for observing runs of different sensitivity and duration. The sensitive volume is given by the cube of the value on the x-axis, and the duration of the observing run is given in the legend. Data points are the expected number of events for the most likely rate derived from population synthesis models and extrapolation from the known population of BNS systems. Error bars indicate the range of possible rates consistent with observation, which is quite large. The dashed line indicates the threshold at which the probability for at least one detection crosses 50%.

Both of these can be accomplished if the gravitational wave signal is associated with a gamma-ray burst. The beaming angle for short GRBs implies the rate of coincident detections with the advanced gravitational observatories will be very small, perhaps one in five years. Without a significant upgrade in sensitivity it will be challenging to collect a large population of coincident detections.

The prospects for Advanced LIGO are bright, but the true discoveries will be made by the next generation – of detectors, and of grad students.

APPENDIX A

NOTES ON AMPLITUDE SPECTRAL DENSITY

In this appendix we discuss the calculation of the amplitude spectral density for a discrete time-series signal, and present the general expression for the signal-to-noise ratio in terms of the ASD. We conclude with a discussion of the statistical distribution of the ASD for ideal Gaussian noise and present some extended results from a study of the noise during the LIGO S6 science run.

A.1 Amplitude Spectral Density

Let us say we have a discrete time-series signal h with sample rate F_s (measured in Hz) and duration T (measured in seconds). The total number of samples is $N = F_s \cdot T$. The Fourier transform of h will have frequency resolution $dF = 1/T$ (measured in Hz) and a Nyquist frequency of $F_s/2$ (again, in Hz). Remember the rules of Fourier transforms: the result will have values for both positive and negative frequencies. Thus, our DFT of h will have $F_s/(2 * dF) = F_s \cdot T/2$ samples from $f = 0$ to $f = F_s/2$, and again from $f = 0$ to $f = -F_s/2$. So the total number of samples in the Fourier transform will be $F_s \cdot T = N$. This is equal to the number of samples in the time series; all the information in the original data is preserved.

We can write this down in a more explicit analytical form. Recall that the Fourier transform and inverse Fourier transform for continuous, infinite signals are given by:

$$\tilde{H}(f) = \int_{-\infty}^{\infty} h(t)e^{-2\pi ift} dt \quad h(t) = \int_{-\infty}^{\infty} \tilde{H}(f)e^{2\pi ift} df. \quad (\text{A.1})$$

We convert the continuous time-series $h(t)$ to a discrete time-series h_n :

$$h(t) \Rightarrow h_n, \quad n = 0 \dots N - 1, \quad N = F_s \cdot T. \quad (\text{A.2})$$

Frequency-domain data are written like so:

$$\tilde{H}(f) \Rightarrow \tilde{H}_k, \quad k = 0 \dots N - 1. \quad (\text{A.3})$$

The Discrete Fourier Transform (DFT) is then defined as:

$$\tilde{H}_k = \sum_{n=0}^{N-1} h_n e^{-2\pi i k n / N} \quad h_n = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{H}_k e^{2\pi i k n / N}. \quad (\text{A.4})$$

We normalize the Fourier components by the sampling frequency, in Hz. This will give us a set of numbers which have units [signal]/Hz, where [signal] stands for whatever units the initial timeseries was measured in – meters, volts, or (as in the case of strain data from LIGO) unitless. Also we can use only the magnitude of the complex numbers and drop the phase information. This is done like so:

$$\tilde{h}_k = |\tilde{H}_k| / F_s. \quad (\text{A.5})$$

We now convert the normalized Fourier components \tilde{h}_k into the *power spectral density* (PSD). In the strictest mathematical sense, the PSD of a time series $x(t)$ is the Fourier transform of the expectation value of the auto-correlation of $x(t)$:

$$G(f) = \int_{-\infty}^{\infty} R_x(\tau) e^{-2\pi f \tau} d\tau, \quad (\text{A.6})$$

where

$$R_x(\tau) = \langle x(t) * x(t - \tau) \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T x(t) x(t - \tau) dt, \quad (\text{A.7})$$

and the angle brackets indicate the expectation value, or average value. This is a statement of the *Weiner-Khinchin theorem*. Consider, for a moment, what this relationship means: if the PSD of a signal is a delta function, the autocorrelation function will be constant in time, and vice versa. Signals which are correlated in time will have structure in the frequency domain.

For our purposes, we can use a simpler definition for the PSD, which can be derived from the Weiner-Khinchin definition by taking $\tau = 0$ and using Parseval's theorem. We say:

$$G(f) = \lim_{T \rightarrow \infty} \frac{1}{T} \langle |\tilde{h}(f)|^2 \rangle, \quad (\text{A.8})$$

where $\tilde{h}(f)$ are the normalized Fourier components, and $\langle |\tilde{h}(f)|^2 \rangle$ denotes the expectation value for the square of each of our Fourier components $\tilde{h}(f)$. Here, the PSD is defined as an average of the magnitude squared of the Fourier coefficients.

Note that since $\tilde{h}(f)$ has units [signal]/Hz, and $1/T$ has units of Hz, the PSD has units of [signal]²/Hz. *The power spectral density is the mean square amplitude per hertz of bandwidth, as a function of frequency.* Put another way, the PSD is a measurement of the power of the signal in each frequency bin, normalized by the

bandwidth of the frequency bin – it is the the expected energy spectral density per unit time. The bandwidth of each frequency bin is determined by the duration of the FFT, $BW = dF = 1/T$. The quantity $1/T$ tells you how narrowly you can measure signal frequency; longer duration FFTs will be able to discern narrower frequency characteristics. (But: longer FFTs will be less sensitive to short-duration signals; there is a trade-off!)

The PSD has positive and negative frequency information; we call this the *two-sided power spectral density*. It's more natural to restrict ourselves to the positive frequencies. For a real input $h(t)$, the negative frequency values will have the same magnitude as their positive frequency counterparts: $\tilde{h}(f) = -\tilde{h}(-f)$. We can ignore the negative frequency bins and express the same information if we multiply the positive frequencies by a factor of two. The *one-sided* PSD is typically written as $S(f)$:

$$S(f) = 2G(f) , f > 0 . \quad (\text{A.9})$$

From here, we calculate the *one-sided amplitude spectral density* (ASD) of the data, which is just the square root of the PSD. Note the units: $1/\sqrt{\text{Hz}}$. The ASD is a *density*: it represents the root mean square (rms) amplitude of the signal in each frequency bin, normalized by the square-root of the width of the bin (the bandwidth). To get the true rms amplitude for each frequency component, we multiply the ASD by $\sqrt{1/T}$, the width of each bin, which has units $\text{sqrt}[\text{Hz}]$. The result is known as the *amplitude spectrum*.

An important quantity in data analysis is the *signal-to-noise ratio*, or SNR, usually written as ρ . To calculate the SNR, we take the ratio of a signal PSD ($G(f)$) with the PSD of the noise ($G_n(f)$), and integrate over all frequencies. In integral form (for continuous signals) this looks like:

$$\rho^2 = \int_{-\infty}^{\infty} \frac{G(f)}{G_n(f)} df = 2 \int_0^{\infty} \frac{G(f)}{G_n(f)} df = 2 \int_0^{\infty} \frac{\tilde{h}^*(f)\tilde{h}(f)}{G_n(f)} df . \quad (\text{A.10})$$

Note that since ρ^2 is the ratio of PSDs, then the SNR itself is given by the ratio of ASDs. This leads to a helpful rule for narrowband signals: the SNR for a narrowband signal (such as a low-Q sine-Gaussian waveform as was used in the GRB search in Chapter 7) is approximately $\sqrt{2}$ times the ratio of the ASD of the signal to the ASD of the noise. For a narrowband signal, the ASD at the central frequency is given by h_{rss} , the quantity we derived in Chapter 2.

A.2 Statistics of Stationary Gaussian Noise

Ideally, the sensitivity of gravitational wave detectors is limited by a collection of stationary Gaussian processes with different frequency characteristics. The combination of these Gaussian noise sources produces the

familiar sensitivity curve when we calculate the ASD of the strain sensitivity of the detectors, for example Fig. 5.5.

If the noise is truly Gaussian and stationary, we can make quantitative predictions for the statistical behavior of each frequency bin in the PSD and ASD. From Eq. A.8 we see that the PSD is the magnitude of the Fourier components of the time-series data. The Fourier components are complex numbers, $\tilde{h}(f) = x + iy$, and if the noise is Gaussian the real and imaginary part of $\tilde{h}(f)$ will be independent, identically distributed Gaussian random variables with zero mean and variance given by σ^2 . In general the variance will be a function of frequency. The PSD is the quadrature sum of the real and imaginary parts of the Fourier terms:

$$S(f) = x(f)^2 + y(f)^2, \quad x = \text{Re}[\tilde{h}(f)], \quad y = \text{Im}[\tilde{h}(f)]. \quad (\text{A.11})$$

A random variable constructed from the quadrature sum of two i.i.d. Gaussian random variables with variance σ^2 follows a Gamma distribution with a shape parameter $k = 1$ and scale parameter $\theta = 2\sigma^2$. (This is also known as an exponential distribution with a rate parameter $\lambda = 1/\theta$.) An important property of Gamma distributions of this type, $\Gamma(k = 1, \theta = 2\sigma^2)$, is that their mean and standard deviation are both equal to $\sqrt{2}\sigma$. This is used to construct the so-called *Rayleigh statistic* for a PSD, which is the ratio of the standard deviation in each frequency bin to the mean value of that bin,

$$R = \sigma/\mu. \quad (\text{A.12})$$

For Gaussian noise the Rayleigh statistic is equal to one. Deviations from unity can be used to study the noise characteristics [128]. $R > 1$ indicates variance in the noise larger than expected from Gaussian processes alone (i.e., the presence of glitches), and $R < 1$ indicates coherent variation in the noise (for example, a mechanical resonance with constant amplitude).

There is a minor but common misconception that each frequency bin of the PSD of Gaussian data should follow a Poisson distribution. In a strict (pedantic?) sense this is not correct, since a Poisson distribution is defined only for discrete values. The correct description for the PSD of Gaussian noise is the continuous Gamma distribution. Importantly, the Rayleigh statistic can be calculated for both discrete Poisson distributions and for the $\Gamma(k = 1, \theta = 2\sigma^2)$ distribution, and so noise studies that assume the PSD follows a Poisson distribution are nonetheless correct.

The ASD of Gaussian noise is the square root of a Gamma-distributed random variable:

$$\sqrt{S(f)} = \sqrt{x(f)^2 + y(f)^2}. \quad (\text{A.13})$$

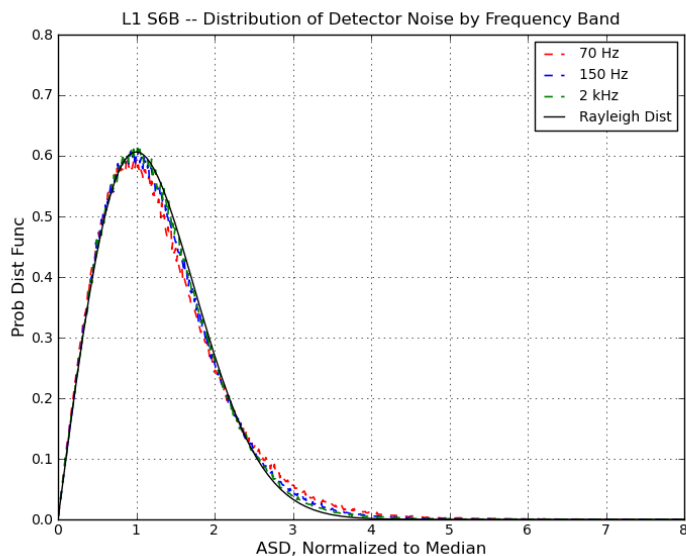


Figure A.1. Distributions of ASD values for three frequency bins in L1 data from S6B. The deviation from the expected Rayleigh distribution is more dramatic for the low frequency data where glitches are more common.

The value of the ASD for each frequency bin of Gaussian noise will follow a Rayleigh distribution. Deviations from the expected Rayleigh behavior for Gaussian noise can be studied to characterize the properties of the noise, for example in the results presented in Chapter 6. In Fig. A.1 we plot a direct comparison between the Rayleigh distribution and data for three frequency bins of the noise ASD from the L1 detector in S6B.

A subtle but important point to consider when calculating the ASD for Gaussian noise is the inherent bias in the method used to estimate the ASD. This bias is important for searches which rely on the precise estimation of the ASD of the noise to calculate the SNR for detected signals, using Eq. A.10, see for example Appendix B of [27]. The details of ASD estimation and the biases inherent in the usual methods have been described by Evan Hall [70].

A.3 Studies of Noise from the S6 Science Run

In this section we include some interesting plots from a study of the long-term noise quality during the S6 science run. Plots for all of the S6 epochs and details of the study are available at the following address: <https://wiki.ligo.org/DetChar/S6RayleighDistributions>.

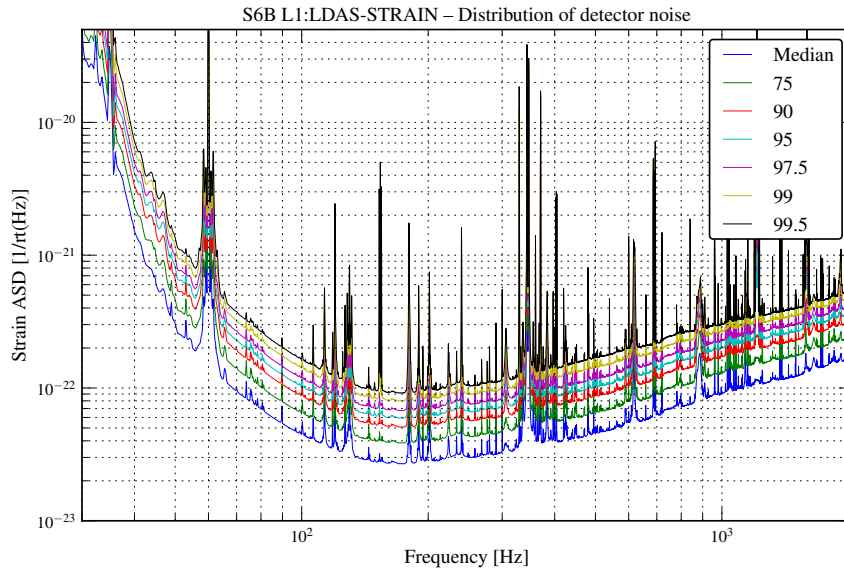


Figure A.2. Distribution of sensitivity per frequency bin for the L1 detector during the S6B epoch. The plot contains data from 952 science segments between Sept 21 2009 and Jan 8 2010, for a total of 3.42 million seconds of data, or about 40 days.

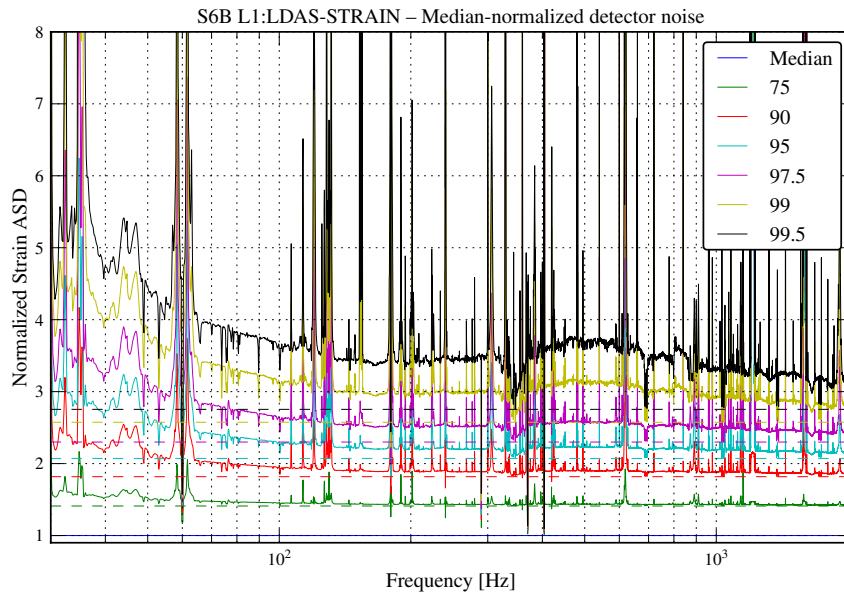


Figure A.3. Median-normalized sensitivity percentiles per frequency bin, for L1 in S6B. The expected ratio for each percentile for Rayleigh-distributed noise is given by the dashed lines. There is a broadband deviation from Gaussian noise due to excessive seismic disturbances during this four-month epoch.




Figure A.4. Median ASDs for the H1 detector for the epochs of the S6 science run. Note the steady improvement of the noise due to commissioning efforts.




Figure A.5. Comparison of the 95th percentile noise for the H1 detector for the epochs of the S6 science run, to the median noise of the S6D epoch.




Figure A.6. Median ASDs for the L1 detector for the epochs of the S6 science run.




Figure A.7. Comparison of the 95th percentile noise for the L1 detector for the epochs of the S6 science run, to the median noise of the S6D epoch.

APPENDIX B

DETECTION STATISTIC FORMALISM

In this appendix we will work through the derivation of the logbayesiancirc detection statistic used by the X-PIPELINE GRB search to rank events. The derivation follows Ref. [140] with some added details.

B.1 Bayesian Formalism

Consider the time-series output of a gravitational wave detector, $d(t)$, which may be pure noise, $d(t) = n(t)$ or contain signal and noise, $d(t) = h(t) + n(t)$. For now, we will assume that the noise $n(t)$ is stationary zero-mean Gaussian noise with a two-sided power spectral density $G_n(f)$. The signal, $h(t)$, is a particular waveform with a specific amplitude, peak time, and other parameters.

We wish to know the probability that $d(t)$ contains a signal of type $h(t)$; we write this probability as $p(h|d)$ (pronounced “the probability of h, given d”). Using Bayes’ Theorem, we can rewrite $p(h|d)$ into an expression that lends itself to measurement (see [40] for an explicit derivation of the following steps):

$$p(h|d) = \frac{p(h)p(d|h)}{p(d)}. \quad (\text{B.1})$$

Here, the quantity $p(d)$ is called the *a priori* probability of observing a specific detector output $d(t)$, and $p(h)$ is the *a priori* probability to observe a signal $h(t)$. $p(h|d)$ is the *a posteriori* probability that $h(t)$ was observed, given a particular detector output $d(t)$. These quantities are informally referred to as the “prior” and the “posterior” probabilities. The expression $p(d|h)$ is called the *likelihood* of observing a particular detector output $d(t)$ in the presence of the signal $h(t)$.

We can write the probability to observe a particular detector output in the following way:

$$p(d) = p(h)p(d|h) + p(0)p(d|0), \quad (\text{B.2})$$

where $p(0)$ is the prior probability of observing no signal, and $p(d|0)$ is the posterior probability of detector output $d(t)$ in the absence of signal. (This is a mathematical expression for a simple truth: there either is a signal,

or there isn't.) Finally, we define a *likelihood ratio* in the following way:

$$\Lambda(d|h) = \frac{p(d|h)}{p(d|0)}. \quad (\text{B.3})$$

$\Lambda(d, h)$ is the ratio of two hypotheses: the denominator is the probability of the *null hypothesis*, or the probability that $d(t)$ is pure noise, and the numerator is the probability of the *signal hypothesis* for a particular signal $h(t)$. The ratio of the two tells us how much more likely the signal hypothesis is than the null hypothesis. We can guess that Λ needs to be a function of the typical noise background $n(t)$. This will become explicit in the next section.

Bringing all of this together, we rewrite the probability that we have observed a signal $h(t)$ as:

$$\begin{aligned} p(h|d) &= \frac{p(h)p(d|h)}{p(d)} = \frac{p(h)p(d|h)}{p(h)p(d|h) + p(0)p(d|0)} \\ &= \frac{\Lambda}{\Lambda + p(0)/p(h)}. \end{aligned} \quad (\text{B.4})$$

The ratio of priors in the denominator, $p(0)/p(h)$, is a constant for a given signal $h(t)$ and does not depend on the detector output. Therefore, the probability that we have observed a signal of type $h(t)$ is a monotonically increasing function of the likelihood ratio $\Lambda(d, h)$, and approaches unity in the limit of Λ large compared to the ratio of priors.

The likelihood ratio is thus an acceptable *detection statistic* for the signal $h(t)$. In the sense of the Neyman-Pearson lemma, it is an *optimal* statistic: for a given false-alarm probability α , we can always find an appropriate threshold Λ^* , above which the probability of the null hypothesis is less than α . For most experiments the probability of detecting a signal is much, much less than the probability of the null hypothesis. In this case Λ must be very large before $p(h|d)$ is significant. Since the range of interesting Λ is so large, what is typically discussed and plotted is the *log-likelihood ratio*, $\ln \Lambda$.

An important consequence of the formalism that follows is that we need not compute the full likelihood ratio every time we want to search the data. From the form of the likelihood ratio we can calculate a sufficient statistic for the data which monotonically increases with Λ .

B.2 Calculation of the Likelihoods

In order for the likelihood ratio to be a useful detection statistic, we need to calculate it for a given discrete time series $d(t)$. So, we need to write an analytical expression for $p(d|h)$ and $p(d|0)$.

In the presence of Gaussian noise (possibly colored, or changing with frequency) with zero mean and covariance matrix Σ_n , the probability of observing a data vector $X = \{x_i\}$ of length N is given by:

$$P(X) = \frac{1}{\sqrt{(2\pi)^N |\Sigma_n|}} \exp \left(-\frac{1}{2} \sum_{i,j=0}^{N-1} x_i \Sigma_{n,ij}^{-1} x_j \right). \quad (\text{B.5})$$

Here, Σ_n^{-1} is the inverse of the noise covariance matrix and $|\Sigma_n|$ is the determinant. Now we apply a trick: take the limit of an large number of samples $N \gg 1$, such that the discrete time series X can be considered a continuous function $x(t)$. In this limit, we can express the covariance matrix Σ_n as the autocorrelation function of the noise $R_n(\tau)$, where $\tau = t - t'$, and write the sum in the exponent as an integral:

$$\sum_{i,j=0}^{N-1} x_i \Sigma_{n,ij}^{-1} x_j = \int \int_{-\infty}^{\infty} x(t) R_n^{-1}(t - t') x(t') dt dt'. \quad (\text{B.6})$$

See [59] for a rigorous derivation. For our next trick, we use the Weiner-Khinchin theorem to write the integral of the autocorrelation function in the time domain as an integral of the Fourier transform of $G_n(f)$ (the PSD of the noise) in the frequency domain:

$$R_n(t - t') = \int_{-\infty}^{\infty} G_n(f) e^{2\pi i f(t-t')} df. \quad (\text{B.7})$$

Finally, by integrating over t and t' and exploiting the orthogonality of the Fourier basis, we have the following:

$$\int \int \int_{-\infty}^{\infty} \frac{x(t) x(t')}{G_n(f)} e^{2\pi i f(t'-t)} dt dt' df = \int_{-\infty}^{\infty} \frac{\tilde{x}^*(f) \tilde{x}(f)}{G_n(f)} df. \quad (\text{B.8})$$

This expression is familiar - it's the same as our expression for ρ^2 from Eq. A.10.

From here, we are motivated to define an *inner product* for the vector space of data functions $a(t)$ and $b(t)$ in the presence of noise with has a one-sided PSD given by $S_n(f) = 2G_n(f)$:

$$\langle a|b \rangle = \int_0^{\infty} \frac{\tilde{a}(f) \tilde{b}^*(f) + \tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} df, \quad (\text{B.9})$$

where $\tilde{a}^*(f)$ is the complex conjugate of the Fourier transform of $a(t)$. For a given $a(t)$ and $b(t)$, this inner product is a mapping from the infinite-dimensional vector space of signals onto \mathbb{R} .

Armed with this new definition, our probability to observe a detector output $d(t)$ becomes an exponential function of the inner product. For example, the probability to observe $d(t)$ under the null hypothesis is:

$$p(d|0) \propto \exp\left(-\frac{1}{2}\langle d|d\rangle\right). \quad (\text{B.10})$$

where we have neglected constants that do not depend on d or h . Finally, we can write the likelihood ratio Λ in terms of the inner product:

$$\begin{aligned} \Lambda(d|h) &= \frac{p(d|h)}{p(d|0)} = \frac{p(d-h|0)}{p(d|0)} = \exp\left(-\frac{1}{2}\langle d-h|d-h\rangle + \frac{1}{2}\langle d|d\rangle\right) \\ &= \exp\left(-\frac{1}{2}[\langle d|d\rangle - 2\langle d|h\rangle + \langle h|h\rangle] + \frac{1}{2}\langle d|d\rangle\right) \\ &= \exp\left(\langle d|h\rangle - \frac{1}{2}\langle h|h\rangle\right). \end{aligned} \quad (\text{B.11})$$

Note that the likelihood, $\Lambda(d|h)$, is a function of a signal waveform $h(t)$ – and the log-likelihood ratio is just the SNR minus a constant!

You might ask how we can calculate the likelihood ratio if we do not know the signal waveform ahead of time. The answer is we are free to use any basis of waveforms $h(t)$ to calculate Λ as a function of the waveform parameters, and then sum the likelihood over the subspace of waveforms that contain the complete signal. For the GRB search with X-PIPELINE we calculate the likelihood in the Fourier basis of sine waves, by decomposing the signal into time-frequency pixels using the DFT. Adjacent pixels with significant likelihood are summed to construct a final value for Λ .

B.3 The Detection Statistic

In this section we present the full expression for the logbayesian detection statistic used in the GEO-GRBs search (Eq. 7.10) and work through an example calculation. The derivation follows section 5.3 of [140].

For GRB searches with X-PIPELINE, we construct a detection statistic from the likelihood ratio given by Eq. B.11, with the inner product defined by Eq. B.9. Signals are assumed to be circularly polarized, and we only test waveforms h which have right- or left-circular polarization. For the remainder of this section we will restrict ourselves to right-circularly-polarized signals; the expressions for left-circular-polarization are the same. We write our signal as $\mathbf{h} = Ae^{\odot}$, where e^{\odot} is the unit vector in the right-circularly-polarized direction.

We make a further assumption regarding the amplitude of the signals, motivated by our expectation that detectable signals will be distributed around some characteristic amplitude. Signals that are very strong (i.e., very close by) are unlikely, since we assume our sources are distributed uniformly (and sparsely) in the nearby

universe. Signals that are very weak will fall below the sensitivity of the detector and are indistinguishable from noise. This assumption is expressed as a prior $p(A)$, and we construct the detection statistic by marginalizing over A :

$$\Lambda(d|h) = \int \Lambda(d, h) p(A) dA, \quad (\text{B.12})$$

where $p(A)$ is a Gaussian prior with characteristic amplitude σ_h :

$$p(A) = \frac{1}{\sqrt{2\pi\sigma_h^2}} \exp \left[-\frac{1}{2} \left(\frac{A}{\sigma_h} \right)^2 \right]. \quad (\text{B.13})$$

Using Eq. B.11 and with some rearrangement, we have:

$$\Lambda(d|h, \sigma_h) = \int \frac{dA}{\sqrt{2\pi\sigma_h^2}} \exp \left[\langle d|h \rangle - \frac{1}{2} \langle h|h \rangle - \frac{1}{2} \left(\frac{A}{\sigma_h} \right)^2 \right]. \quad (\text{B.14})$$

We can perform this integral using the usual tricks for Gaussian integration. Recall the general form of a Gaussian integral is solved using:

$$\int_{-\infty}^{\infty} e^{-ax^2+bx+c} dx = \sqrt{\frac{\pi}{a}} e^{\frac{b^2}{4a}+c}. \quad (\text{B.15})$$

Expressing the signal \mathbf{h} as $A\mathbf{e}^\odot$ and the data as a vector, we rewrite the integral as:

$$\Lambda(\mathbf{d}|\odot, \sigma_h) = \int \frac{dA}{\sqrt{2\pi\sigma_h^2}} \exp \left[-\frac{A^2}{2} \left(\langle \mathbf{e}^\odot|\mathbf{e}^\odot \rangle + \frac{1}{\sigma_h^2} \right) + A \langle \mathbf{d}|\mathbf{e}^\odot \rangle \right]. \quad (\text{B.16})$$

The result is:

$$\Lambda(\mathbf{d}|\odot, \sigma_h) = \frac{1}{\sqrt{1 + \sigma_h^2 \langle \mathbf{e}^\odot|\mathbf{e}^\odot \rangle}} \exp \left[\frac{1}{2} \frac{\langle \mathbf{d}|\mathbf{e}^\odot \rangle^2}{\langle \mathbf{e}^\odot|\mathbf{e}^\odot \rangle + 1/\sigma_h^2} \right]. \quad (\text{B.17})$$

Recall that our expression for the inner product, Eq. B.9, carried a factor of $1/S_n(f)$ as part of the calculation of the SNR. In practice, this is included in the calculation of the unit vectors, and we call $\mathbf{f}^\odot = \mathbf{e}^\odot / \sqrt{S_n(f)}$ the *noise-weighted* unit vector in the right-circularly-polarized direction. With this shorthand we can simplify the inner products, and we get the expression:

$$\Lambda(\mathbf{d}|\odot, \sigma_h) = \frac{1}{\sqrt{1 + \sigma_h^2 |\mathbf{f}^\odot|^2}} \exp \left[\frac{1}{2} \frac{|\mathbf{e}^\odot \cdot \mathbf{d}|^2}{1 + 1/(\sigma_h |\mathbf{f}^\odot|^2)} \right], \quad (\text{B.18})$$

where $|\mathbf{f}^\odot|^2$ is of order $1/S_n(f)$. Finally, by taking the log, we have an expression for the log-likelihood that \mathbf{d} is consistent with a right-circularly polarized gravitational wave with amplitude σ_h :

$$2L(\mathbf{d}|\odot, \sigma_h) = \left[\frac{|\mathbf{e}^\odot \cdot \mathbf{d}|^2}{1 + 1/(\sigma_h |\mathbf{f}^\odot|^2)} \right] - \log(1 + \sigma_h^2 |\mathbf{f}^\odot|^2). \quad (\text{B.19})$$

The overall factor of two is a nonstandard choice to simplify the expression.

B.4 The Detection Statistic in Practice

Now, we turn to an example calculation of the detection statistic. The waveform that will maximize L is of the form $\mathbf{h} = Ae^{\odot}$, a right-circularly polarized gravitational wave with an amplitude A . For now we drop the noise term, and say that $d = h$. As part of the data conditioning, the time series from each detector is *whitened* to filter out steady features in the noise such as lines, and is normalized to the detector noise floor to produce a unit-variance time series. For data containing a gravitational wave with amplitude A we write this as $\mathbf{d} = Ae^{\odot}/\sqrt{S_n(f)} = A\mathbf{f}^{\odot}$.

The log-likelihood given this data vector is:

$$\begin{aligned} 2L(\mathbf{d} | \odot, \sigma_h) &= \left[\frac{|e^{\odot} \cdot A\mathbf{f}^{\odot}|^2}{1 + 1/(\sigma_h|\mathbf{f}^{\odot}|)^2} \right] - \log(1 + \sigma_h^2|\mathbf{f}^{\odot}|^2) \\ &= \frac{A^2|\mathbf{f}^{\odot}|^2}{1 + 1/(\sigma_h|\mathbf{f}^{\odot}|)^2} - \log(1 + \sigma_h^2|\mathbf{f}^{\odot}|^2). \end{aligned} \quad (\text{B.20})$$

The likelihood ratio $\Lambda(\mathbf{d} | \odot, \sigma_h)$ is then:

$$\Lambda(\mathbf{d} | \odot, \sigma_h) = \sqrt{\frac{1}{1 + \sigma_h^2|\mathbf{f}^{\odot}|^2}} \exp \left[\frac{1}{2} \frac{A^2|\mathbf{f}^{\odot}|^2}{1 + 1/(\sigma_h|\mathbf{f}^{\odot}|)^2} \right]. \quad (\text{B.21})$$

Working out the denominator in the exponent,

$$1 + \frac{1}{(\sigma_h|\mathbf{f}^{\odot}|)^2} = \frac{(\sigma_h|\mathbf{f}^{\odot}|)^2}{(\sigma_h|\mathbf{f}^{\odot}|)^2} + \frac{1}{(\sigma_h|\mathbf{f}^{\odot}|)^2} = \frac{1 + (\sigma_h|\mathbf{f}^{\odot}|)^2}{(\sigma_h|\mathbf{f}^{\odot}|)^2}. \quad (\text{B.22})$$

So,

$$\Lambda(\mathbf{d} | \odot, \sigma_h) = \sqrt{\frac{1}{1 + \sigma_h^2|\mathbf{f}^{\odot}|^2}} \exp \left[\frac{1}{2} \frac{A^2|\mathbf{f}^{\odot}|^2 (\sigma_h|\mathbf{f}^{\odot}|)^2}{1 + (\sigma_h|\mathbf{f}^{\odot}|)^2} \right]. \quad (\text{B.23})$$

The noise-weighted antenna response vector \mathbf{f}^{\odot} is of order $1/\sqrt{S_n(f)}$, where $S_n(f)$ is the one-sided PSD of the noise. For the initial LIGO detectors the amplitude spectral density of the noise was a minimum of about $\sqrt{S_n(f)} \simeq 2 \times 10^{-23} \text{ Hz}^{-1/2}$ in a 100 Hz band around $f = 150 \text{ Hz}$. If we assume our signal is a narrowband waveform around this frequency, we can make a back-of-the-envelope estimate for the value of the likelihood under various assumptions of signal strength and the amplitude prior. The expression for Λ is:

$$\Lambda(\mathbf{d} | \odot, \sigma_h) = \sqrt{\frac{1}{1 + \sigma_h^2/S_n(f)}} \exp \left[\frac{A^2\sigma_h^2/S_n^2(f)}{1 + \sigma_h^2/S_n(f)} \right]. \quad (\text{B.24})$$

The term in the square root amounts to a normalization constant, and for the purposes of illustration we will neglect it. Let us focus on the term in the exponent, and calculate its value for a variety of signal amplitudes.

For initial detector searches, the amplitude width σ_h ranges from 10^{-23} to 10^{-21} $\text{Hz}^{-1/2}$. For a loud signal, for example with $A \sim 10^{-21}$ $\text{Hz}^{-1/2}$, we can calculate the likelihood ratio for each value of the amplitude prior. Here we go:

$$\begin{aligned}\Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-23}) &\propto \exp \left[\frac{10^{-42} \times 10^{-46} / (16 \times 10^{-92})}{1 + 10^{-46} / (4 \times 10^{-46})} \right] = \exp \left[\frac{10^4 / 16}{1 + 1/4} \right] \simeq \exp \left[\frac{10^4}{20} \right] \\ \Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-22}) &\propto \exp \left[\frac{10^{-42} \times 10^{-44} / (16 \times 10^{-92})}{1 + 10^{-44} / (4 \times 10^{-46})} \right] = \exp \left[\frac{10^6 / 16}{1 + 100/4} \right] \simeq \exp \left[\frac{10^6}{400} \right] \\ \Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-21}) &\propto \exp \left[\frac{10^{-42} \times 10^{-42} / (16 \times 10^{-92})}{1 + 10^{-42} / (4 \times 10^{-46})} \right] = \exp \left[\frac{10^8 / 16}{1 + 10^4/4} \right] \simeq \exp \left[\frac{10^8}{4 \times 10^4} \right].\end{aligned}$$

For a weak signal, for example with $A \sim 10^{-23}$ $\text{Hz}^{-1/2}$, we have:

$$\begin{aligned}\Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-23}) &\propto \exp \left[\frac{10^{-46} \times 10^{-46} / (16 \times 10^{-92})}{1 + 10^{-46} / (4 \times 10^{-46})} \right] = \exp \left[\frac{1/16}{1 + 1/4} \right] \simeq \exp \left[\frac{1}{20} \right] \\ \Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-22}) &\propto \exp \left[\frac{10^{-46} \times 10^{-44} / (16 \times 10^{-92})}{1 + 10^{-44} / (4 \times 10^{-46})} \right] = \exp \left[\frac{10^2/16}{1 + 100/4} \right] \simeq \exp \left[\frac{10^2}{400} \right] \\ \Lambda(\mathbf{d} | \odot, \sigma_h = 10^{-21}) &\propto \exp \left[\frac{10^{-46} \times 10^{-42} / (16 \times 10^{-92})}{1 + 10^{-42} / (4 \times 10^{-46})} \right] = \exp \left[\frac{10^4/16}{1 + 10^4/4} \right] \simeq \exp \left[\frac{10^4}{4 \times 10^4} \right].\end{aligned}$$

Note that the value of the detection statistic for a given signal amplitude scales like A^2 (or, like ρ^2 , where ρ is the SNR). The final calculation of the logbayesian detection statistic is performed by marginalizing over σ_h .

In this calculation, we have only measured the strength of the signal with respect to the detector noise floor $S_n(f)$. Importantly, we have not quantified the consistency of the signal to a coherent circularly-polarized signal across the detector network. The method by which we quantify the signal coherence is described in the next section.

APPENDIX C

COHERENT CONSISTENCY CHECKS

In this appendix we illustrate the method we use to distinguish gravitational signals from loud noise transients (glitches). The methods of this section rely on the assumption that the noise of the detectors in the network is stationary and Gaussian, with occasional large-amplitude transients at random, uncorrelated times. The example given here is an expansion on section 5.7.2 of [140].

For a given projection with unit vector \mathbf{e} , we calculate the the *coherent energy* as:

$$E = |\mathbf{e} \cdot \mathbf{d}|^2 = \sum_{i,j=1}^D e_i^* e_j d_i d_j^*. \quad (\text{C.1})$$

The *incoherent* component of the energy is the sum of the only the diagonal (autocorrelation) terms:

$$I = \sum_{i=1}^D |e_i|^2 |d_i|^2. \quad (\text{C.2})$$

For a circular signal, the circular energy E_{\circ} is taken to be the maximum of either the left- or right-circular energy; the incoherent part is equal in both cases. The null projection \mathbf{e}^n is used to find the null energy.

Consider a detector network comprised of two instruments. Our data vector is composed of the data streams from each detector, and is equal to whatever gravitational signal is present, plus noise (which may be Gaussian, or due to a loud glitch):

$$\mathbf{d} = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = \mathbf{s} + \mathbf{n} = \begin{bmatrix} s_1 + n_1 \\ s_2 + n_2 \end{bmatrix}. \quad (\text{C.3})$$

We consider two examples: one where a gravitational wave signal is present in the data and the amplitude is large compared to the Gaussian background noise ($|\mathbf{s}| \gg |\mathbf{n}|$, $s_1 = s_2 = A$), and one where no signal is present and one detector has a large glitch ($\mathbf{s} = 0$, $n_1 = A \gg n_2$). We will calculate the coherent and incoherent energies for the circular and null projections in both cases to see how the coherent consistency checks are constructed.

The general expression for the coherent energy in the 2-detector case is:

$$E = |\mathbf{e} \cdot \mathbf{d}|^2 = |\mathbf{e} \cdot \mathbf{s} + \mathbf{e} \cdot \mathbf{n}|^2. \quad (\text{C.4})$$

In the circular polarization projection, this becomes:

$$\begin{aligned} E_{\circ} &= \min(E_{\circ}, E_{\circ}) = \min(|\mathbf{e}^{\circ} \cdot \mathbf{e}|^2, |\mathbf{e}^{\circ} \cdot \mathbf{e}|^2) \\ &= |\mathbf{e}^{\circ} \cdot \mathbf{s} + \mathbf{e}^{\circ} \cdot \mathbf{n}|^2 = (|e_1^{\circ}|s_1 + |e_2^{\circ}|s_2 + |e_1^{\circ}|n_1 + |e_2^{\circ}|n_2)^2. \end{aligned} \quad (\text{C.5})$$

For the incoherent energy:

$$I_{\circ} = |e_1^{\circ}|^2(s_1 + n_1)^2 + |e_2^{\circ}|^2(s_2 + n_2)^2. \quad (\text{C.6})$$

In the null direction, any signal component will vanish ($\mathbf{e}^n \cdot \mathbf{s} = 0$), so the coherent null energy expression is only noise:

$$E_n = |\mathbf{e}^n \cdot \mathbf{s} + \mathbf{e}^n \cdot \mathbf{n}|^2 = (|e_1^n|n_1 + |e_2^n|n_2)^2, \quad (\text{C.7})$$

and the incoherent null energy is the same,

$$I_n = |e_1^{\circ}|^2(s_1 + n_1)^2 + |e_2^{\circ}|^2(s_2 + n_2)^2. \quad (\text{C.8})$$

Following through the arithmetic and collecting only the leading order terms in amplitude A , we have the following for case of a gravitational wave signal:

$$\begin{aligned} E_{\circ} &\propto A^2, \\ I_{\circ} &\propto A^2, \\ E_n &\propto |\mathbf{n}|^2, \\ I_n &\propto A^2. \end{aligned} \quad (\text{C.9})$$

For the case of a glitch where $\mathbf{s} = 0$ and $n_1 = A$,

$$\begin{aligned} E_{\circ} &\propto A^2, \\ I_{\circ} &\propto A^2, \\ E_n &\propto A^2, \\ I_n &\propto A^2. \end{aligned} \quad (\text{C.10})$$

We see that in the signal case, the signal energy will cancel in the null direction, and $I_n > E_n$, whereas in the case of noise the incoherent and coherent null energies will be roughly equal. In this way we can apply a ratio test on the null energies to distinguish signal from noise. Cuts of this sort are very effective for rejecting loud glitches.

For low-amplitude glitches, we consider two more expressions. First, the difference of the coherent and incoherent null energies, in which the terms which are quadratic in noise components will cancel:

$$\begin{aligned}
E_n - I_n &= (|e_1^n|n_1 + |e_2^n|n_2)^2 - |e_1^n|^2(s_1 + n_1)^2 + |e_2^n|^2(s_2 + n_2)^2 \\
&= 2|e_1^n|n_1|e_2^n|n_2 - |e_1^n|^2s_1^2 - 2|e_1^n|^2s_1n_1 - |e_1^n|^2s_1^2 - 2|e_2^n|^2s_2n_2 \\
&\propto n_1n_2 - s_1^2 - s_2^2 - s_1n_1 - s_2n_2.
\end{aligned} \tag{C.11}$$

We call this the *cross-correlation* term, since it removes the autocorrelation terms from the coherent energy.

We also need the sum of the coherent and incoherent null energies:

$$E_n + I_n \propto n_1^2 + n_2^2 + n_1n_2 + s_1^2 + s_2^2 + s_1n_1 + s_2n_2. \tag{C.12}$$

In the case of a loud glitch in one detector with no signal, the magnitude of the cross-correlation term, $|E_n - I_n|$, will be proportional to the glitch amplitude A . In the presence of a true gravitational wave, the signal is present in both detectors, so $|E_n - I_n| \propto A^2$. In all cases, the magnitude of the summation expression is proportional to A^2 .

We therefore use the ratio of the two to set a threshold r on the acceptance of events:

$$\frac{|E_n - I_n|}{(E_n + I_n)^\alpha} > r. \tag{C.13}$$

Note the exponent α in the denominator; if we choose $\alpha = 1$, the denominator will dominate in the noise case, and glitches will be rejected with probability increasing with amplitude:

$$\begin{aligned}
\frac{|E_n - I_n|}{(E_n + I_n)} &\propto \frac{A^2}{A^2} = \text{const} \quad (\text{for signal}) \\
\frac{|E_n - I_n|}{(E_n + I_n)} &\propto \frac{A}{A^2} = \frac{1}{A} \quad (\text{for glitches}).
\end{aligned} \tag{C.14}$$

If we choose $\alpha = 1/2$, the numerator will dominate in the signal case. Signals will fall above the cutoff with a probability increasing with amplitude, and glitches will be rejected at a constant fraction:

$$\begin{aligned} \frac{|E_n - I_n|}{(E_n + I_n)^{1/2}} &\propto \frac{A^2}{A} = A && \text{(for signal)} \\ \frac{|E_n - I_n|}{(E_n + I_n)^{1/2}} &\propto \frac{A}{A} = \text{const} && \text{(for glitches)}. \end{aligned} \tag{C.15}$$

In either case, lowering the threshold r will accept more events, whether they be signal or noise. In practice, we used a fixed $\alpha = 0.8$ and tune for the value of r which optimizes the background rejection while maintaining the lowest possible threshold for detection.

One final observation: if both detectors have loud glitches at the same time, each with significant amplitude in the null projection, the difference $|E_n - I_n|$ can be of the order A^2 , due to the crossterm $n_1 n_2$; in this case the consistency check can fail. Coincident noise transients between detectors are a source of indistinguishable background for the search and limit the sensitivity. The rate of coincident glitches is estimated empirically, by repeating the analysis many times with unphysical time offsets in the data from one of the detectors.

APPENDIX D

EPIGRAPHS

The chapter epigraphs in this thesis are from the following sources:

Acknowledgements: Shakespeare, *As You Like It*, Act I Scene 2.

Chapter 1 (Introduction): Shakespeare, *The Tempest*, Act II Scene 1.

Chapter 2 (Gravitational Waves): Shakespeare, *Troilus and Cressida*, Act I Scene 3.

Chapter 3 (Detectors): Shakespeare, Sonnet 14.

Chapter 4 (Optics): Shakespeare, Sonnet 77.

Chapter 5 (Advanced LIGO): Shakespeare, *The Winter's Tale*, Act IV Scene 4.

Chapter 6 (Detector Characterization): Shakespeare, *The Tempest*, Act III Scene 2.

Chapter 7 (GRB Search): Milton, *Paradise Lost*, Book 1.

Chapter 8 (Conclusion): Shakespeare, *The Tragedy of Julius Caesar*, Act V Scene 1.

BIBLIOGRAPHY

- [1] Aartsen, M. G., et al. Search for prompt neutrino emission from gamma-ray bursts with icecube. *The Astrophysical Journal Letters* 805, 1 (2015), L5.
- [2] Aasi, J., et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys. Rev. D* 87 (Feb 2013), 042001.
- [3] Aasi, J., et al. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics* 7 (2013), 613–619.
- [4] Aasi, J., et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories, 2013. arxiv:1304.0670.
- [5] Aasi, J., et al. Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009~2010 LIGO and Virgo Data. *Phys. Rev. Lett.* 113 (Dec 2014), 231101.
- [6] Aasi, J., et al. Search for Gravitational Waves Associated with Gamma-ray Bursts Detected by the Interplanetary Network. *Phys. Rev. Lett.* 113 (2014), 011102.
- [7] Aasi, J., et al. Advanced ligo. *Classical and Quantum Gravity* 32, 7 (2015), 074001.
- [8] Aasi, J., et al. Characterization of the LIGO detectors during their sixth science run. *Classical and Quantum Gravity* 32, 11 (2015), 115012.
- [9] Aasi, J., et al. Searching for stochastic gravitational waves using data from the two colocated LIGO Hanford detectors. *Phys. Rev. D* 91 (Jan 2015), 022003.
- [10] Aasi, J., et al. The characterization of Virgo data and its impact on gravitational-wave searches. *Classical and Quantum Gravity* 29, 15 (2012), 155002.
- [11] Aasi, J., et al. Methods and results of a search for gravitational waves associated with gamma-ray bursts using the GEO 600, LIGO, and Virgo detectors. *Phys. Rev. D* 89 (Jun 2014), 122004.
- [12] Abadie, J., et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity* 27, 17 (2010), 173001.
- [13] Abadie, J., et al. Search for gravitational-wave inspiral signals associated with short gamma-ray bursts during ligo's fifth and virgo's first science run. *The Astrophysical Journal* 715, 2 (2010), 1453.
- [14] Abadie, J., et al. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nature Physics* 7 (2011), 962–965.
- [15] Abadie, J., et al. All-sky search for gravitational-wave bursts in the second joint LIGO-Virgo run. *Phys. Rev. D* 85 (Jun 2012), 122007.
- [16] Abadie, J., et al. Implications for the Origin of GRB 051103 from LIGO Observations. *The Astrophysical Journal* 755, 1 (2012), 2.
- [17] Abadie, J., et al. Search for Gravitational Waves Associated with Gamma-Ray Bursts during LIGO Science Run 6 and Virgo Science Runs 2 and 3. *The Astrophysical Journal* 760, 1 (2012), 12.
- [18] Abadie, J., et al. Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys. Rev. D* 85 (Apr 2012), 082002.
- [19] Abbott, B. P., et al. Implications for the Origin of GRB 070201 from LIGO Observations. *The Astrophysical Journal* 681, 2 (2008), 1419.

- [20] Abbott, B. P., et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics* 72, 7 (2009), 076901.
- [21] Accadia, T., et al. Virgo: a laser interferometer to detect gravitational waves. *Journal of Instrumentation* 7, 03 (2012), P03012.
- [22] Adhikari, R., Evans, M., Fricke, T., Frolov, V., Kawabe, K., Smith, N., , and Waldman, S. J. DC readout Normalization for Enhanced LIGO, 2009. LIGO T0900023.
- [23] Ajith, P., et al. Instrumental vetoes for transient gravitational-wave triggers using noise-coupling models: The bilinear-coupling veto. *Phys. Rev. D* 89 (2014), 122001.
- [24] Aketagawa, Masato, Kimura, Shohei, Yashiki, Takuya, Iwata, Hiroshi, Banh, Tuan Quoc, and Hirata, Kenji. Measurement of a free spectral range of a Fabry-Perot cavity using frequency modulation and null method under off-resonance conditions. *Meas. Sci. and Tech.* 22, 2 (2011), 025302.
- [25] Aketagawa, Masato, Yashiki, Takuya, Kimura, Shohei, and Banh, TuanQuoc. Free spectral range measurement of Fabry-Perot cavity using frequency modulation. *Int. Jour. of PEM* 11, 6 (2010), 851–856.
- [26] Allen, Bruce. 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).
- [27] Allen, B., et al. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D* 85 (2012), 122006.
- [28] Andersson, Nils et al. The transient gravitational-wave sky. *Classical and Quantum Gravity* 30, 19 (2013), 193002.
- [29] Arai, K., et al. New signal extraction scheme with harmonic demodulation for power-recycled fabryperotmichelson interferometers. *Physics Letters A* 273, 12 (2000), 15 – 24.
- [30] Arai, K., et al. Output Mode Cleaner Design, 2013. LIGO T1000276.
- [31] Arai, K., et al. Matrix calculations for the angular control of H1 OMC, 2014. LIGO T1400585.
- [32] Arai, K. On the accumulated round-trip Gouy phase shift for a general optical cavity, 2013. LIGO T1300189.
- [33] Araya, A., et al. Absolute-length determination of a long-baseline Fabry-Perot cavity by means of resonating modulation sidebands. *Appl. Opt.* 38, 13 (May 1999), 2848–2856.
- [34] Babak, S., Gair, J. R., Petiteau, A., and Sesana, A. Fundamental physics and cosmology with LISA. *Classical and Quantum Gravity* 28, 11 (2011), 114001.
- [35] Barostti, L., et al. HAM6 Overview, 2013. LIGO D1000342-v14.
- [36] Barthelmy, S D. GCN and VEvent: A status report. *Astronomische Nachrichten* 329, 3 (Mar. 2008), 340.
- [37] Berger, Edo. Short-Duration Gamma-Ray Bursts. *Annual Review of Astronomy and Astrophysics* 52, 1 (2014).
- [38] Blackburn, L, et al. The LSC glitch group: monitoring noise transients during the fifth LIGO science run. *Classical and Quantum Gravity* 25, 18 (2008), 184004.
- [39] Bloom, J. S., Kulkarni, S. R., and Djorgovski, S. G. The Observed Offset Distribution of Gamma-Ray Bursts from Their Host Galaxies: A Robust Clue to the Nature of the Progenitors. *The Astronomical Journal* 123, 3 (2002), 1111.
- [40] Brown, D. *Searching for Gravitational Radiation from Binary Black Hole MACHOs in the Galactic Halo.* PhD thesis, University of Wisconsin, Milwaukee, 2004. arxiv:0705.1514.

- [41] Buonanno, Alessandra, and Chen, Yanbei. Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors. *Phys. Rev. D* 64 (Jul 2001), 042006.
- [42] Carroll, Sean. *Spacetime and Geometry: An Introduction to General Relativity*. Addison-Wesley, 2003.
- [43] Caves, Carlton M. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Phys. Rev. Lett.* 45 (Jul 1980), 75–79.
- [44] Caves, Carlton M. Quantum-mechanical noise in an interferometer. *Phys. Rev. D* 23 (Apr 1981).
- [45] Centrella, J., et al. Black-hole binaries, gravitational waves, and numerical relativity. *Rev. Mod. Phys.* 82 (2010), 3069–3119.
- [46] Chapman, R., et al. How common are long gamma-ray bursts in the local Universe? *MNRAS* 382 (2007), L21–L25.
- [47] Clark, J., et al. Prospects for joint gravitational wave and short gamma-ray burst observations. arXiv:1409.8149.
- [48] Clark, J., et al. Prospects for high frequency burst searches following binary neutron star coalescence with advanced gravitational wave detectors. *Phys. Rev. D* 90 (2014), 062004.
- [49] Connaughton, V. GCN circular 11574, 2011.
- [50] Corsi, A. Gravitational waves and gamma-ray bursts. In *Death of Massive Stars: Supernovae and Gamma-Ray Bursts*, P. Roming, N. Kawai, and E. Pian, Eds., vol. 279. Proceedings International Astronomical Union Symposium, 2012. arXiv:1204.4110.
- [51] Costa, E., et al. Discovery of an X-ray afterglow associated with the gamma-ray burst of 28 February 1997. *Nature* 387, 6635 (1997), 783785.
- [52] Cucchiara, A. GCN GRB Observation Report 5052, 2006.
- [53] Demorest, P. B., et al. Limits on the Stochastic Gravitational Wave Background from the North American Nanohertz Observatory for Gravitational Waves. *The Astrophysical Journal* 762, 2 (2013), 94.
- [54] Dooley, K. L., et al. Angular control of optical cavities in a radiation-pressure-dominated regime: the Enhanced LIGO case. *J. Opt. Soc. Am. A* 30, 12 (Dec 2013), 2618–2626.
- [55] Drever, R.W.P., Hall, J.L., Kowalski, F.V., Hough, J., Ford, G.M., Munley, A.J., and Ward, H. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B* 31, 2 (1983), 97–105.
- [56] Duncan, R. C., and Thompson, C. Formation of very strongly magnetized neutron stars - Implications for gamma-ray bursts. *The Astrophysical Journal Letters* 392 (June 1992), L9–L13.
- [57] Evans, M. *Lock Acquisition in Resonant Optical Interferometers*. PhD thesis, California Institute of Technology, 2001.
- [58] Feroci, M, et al. SuperAGILE: The hard X-ray imager for the AGILE space mission. *Nuclear Instruments and Methods in Physics Research Section A* 581, 3 (2007), 728–754.
- [59] Finn, Lee S. Detection, measurement, and gravitational radiation. *Phys. Rev. D* 46 (Dec 1992), 5236–5249. arxiv:9209010.
- [60] Fong, W., Berger, E., and Fox, D. B. Hubble space telescope observations of short gamma-ray burst host galaxies: Morphologies, offsets, and local environments. *The Astrophysical Journal* 708, 1 (2010), 9.
- [61] Fong, W. et al. A Jet Break in the X-Ray Light Curve of Short GRB111020A: Implications for Energetics and Rates. *The Astrophysical Journal* 756, 2 (2012), 189.
- [62] Forward, Robert L. Wideband laser-interferometer gravitational-radiation experiment [sic]. *Phys. Rev. D* 17 (1978), 379–390.

- [63] Fox, D B, and Meszaros, P. GRB fireball physics: prompt and early emission. *New Journal of Physics* 8, 9 (2006), 199.
- [64] Fricke, T. *Homodyne Detection for Laser-Interferometric Gravitational Wave Detectors*. PhD thesis, Louisiana State University, 2011.
- [65] Fritschel, P., et al. Readout and control of a power-recycled interferometric gravitational-wave antenna. *Appl. Opt.* 40, 28 (Oct 2001), 4988–4998.
- [66] Fryer, Chris L., Holz, Daniel E., and Hughes, Scott A. Gravitational Wave Emission from Core Collapse of Massive Stars. *The Astrophysical Journal* 565, 1 (2002), 430.
- [67] Gehrels, N., et al. The Swift Gamma-Ray Burst Mission. *The Astrophysical Journal* 611, 2 (2004), 1005.
- [68] Gehrels, N., et al. A short gamma-ray burst apparently associated with an elliptical galaxy at redshift $z=0.225$. *Nature* 437 (2005), 851–854.
- [69] Grote, H. The geo 600 status. *Classical and Quantum Gravity* 27, 8 (2010), 084003.
- [70] Hall, E., et al. More notes on median spectrum estimation, 2015. LIGO T1500300.
- [71] Hinderer, T., et al. Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral. *Phys. Rev. D* 81 (2010), 123016.
- [72] Howell, E. J., and Coward, D. M. A redshift-observation time relation for gamma-ray bursts: evidence of a distinct subluminal population. *MNRAS* 428 (2013), 167–181.
- [73] Hurley, K., et al. A new analysis of the short-duration, hard-spectrum GRB 051103, a possible extragalactic soft gamma repeater giant flare. *MNRAS* 403 (2010), 342–352.
- [74] Hurley, K., et al. The Interplanetary Network Supplement to the Fermi GBM Catalog of Cosmic Gamma-Ray Bursts. *The Astrophysical Journal Supplement Series* 207, 2 (2013), 39.
- [75] Isogai, T., for the LVC. The Used Percentage Veto for LIGO S6 and Virgo VSR2 Gravitational Wave Searches. *J. Phys. Conf. Ser.* 243 (2010), 012005.
- [76] Isogai, T., Miller, J., Kwee, P., Barsotti, L., and Evans, M. Loss in long-storage-time optical cavities. *Opt. Express* 21, 24 (Dec 2013), 30114–30125.
- [77] Kenta, H., et al. Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform. *Phys. Rev. D* 88 (Aug 2013), 044026.
- [78] Kiuchi, K., Sekiguchi, Y., Shibata, M., and Taniguchi, K. Long-term general relativistic simulation of binary neutron stars collapsing to a black hole. *Phys. Rev. D* 80 (Sep 2009), 064037.
- [79] Klebesadel, R. W., Strong, I. B., and Olson, R. A. Observations of Gamma-Ray Bursts of Cosmic Origin. *The Astrophysical Journal* 182 (1973), L85.
- [80] Kobayashi, S., and Mészáros, P. Gravitational radiation from gamma-ray burst progenitors. *Astrophysical Journal* 589, 2 (2003), 861.
- [81] Kogelnik, H., and Li, T. Laser beams and resonators. *Proceedings of the IEEE* 54, 10 (Oct 1966), 1312–1329.
- [82] Kouveliotou, C., et al. Identification of two classes of gamma-ray bursts. *The Astrophysical Journal* 413 (Aug. 1993), L101.
- [83] Kuroda, K, and the LCGT Collaboration. Status of LCGT. *Classical and Quantum Gravity* 27, 8 (2010), 084004.
- [84] Lattimer, James M., and Prakash, Madappa. Neutron star observations: Prognosis for equation of state constraints. *Physics Reports* 442, 16 (2007), 109 – 165.

- [85] Le, T., and Dermer, C. D. On the Redshift Distribution of Gamma-Ray Bursts in the Swift Era. *The Astrophysical Journal* 661, 1 (2007), 394.
- [86] Liang, E., et al. Low-luminosity gamma-ray bursts as a unique population: Luminosity function, local rate, and beaming factor. *The Astrophysical Journal* 662, 2 (2007), 1111.
- [87] Martynov, D. *Lock Acquisition and Sensitivity Analysis of Advanced LIGO Interferometers*. PhD thesis, California Institute of Technology, 2015.
- [88] Matsuoka, Masaru, et al. The MAXI Mission on the ISS: Science and Instruments for Monitoring All-Sky X-Ray Images. *Publications of the Astronomical Society of Japan* 61, 5 (2009), 999–1010.
- [89] Mazets, E. P., et al. A Giant Flare from a Soft Gamma Repeater in the Andromeda Galaxy (M31). *The Astrophysical Journal* 680, 1 (2008), 545.
- [90] Meegan, C., et al. The Fermi Gamma-ray Burst Monitor. *The Astrophysical Journal* 702, 1 (2009), 791.
- [91] Meers, Brian J. Recycling in laser-interferometric gravitational-wave detectors. *Phys. Rev. D* 38 (Oct 1988), 2317–2326.
- [92] Mereghetti, Sandro. The strongest cosmic magnets: soft gamma-ray repeaters and anomalous x-ray pulsars. *The Astronomy and Astrophysics Review* 15, 4 (2008), 225–287.
- [93] Mészáros, P. Gamma-ray bursts. *Reports on Progress in Physics* 69, 8 (2006), 2259.
- [94] Misner, Charles, Thorne, Kip, and Wheeler, John. *Gravitation*. W. H. Freeman, 1973.
- [95] Mizuno, J. *Comparison of optical configurations for laser-interferometric gravitational-wave detectors*. PhD thesis, Leibniz Universität Hannover, 1995.
- [96] Mueller, Chris L., et al. In situ characterization of the thermal state of resonant optical interferometers via tracking of their higher-order mode resonances. *Classical and Quantum Gravity* 32, 13 (2015), 135018.
- [97] Nakar, E. Short-hard gamma-ray bursts. *Physics Reports* 442 (2007), 166 – 236.
- [98] Nissanke, S., et al. Determining the Hubble constant from gravitational wave observations of merging compact binaries, 2013. arXiv:1307.2638.
- [99] Oechslin, R., and Janka, H.-T. Gravitational Waves from Relativistic Neutron-Star Mergers with Microphysical Equations of State. *Phys. Rev. Lett.* 99 (Sep 2007), 121102.
- [100] Ott, C. D. The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity* 26, 6 (Mar. 2009), 3001.
- [101] Ott, C. D., Burrows, A., Dessart, L., and Livne, E. A New Mechanism for Gravitational-Wave Emission in Core-Collapse Supernovae. *Phys. Rev. Lett.* 96 (May 2006), 201102.
- [102] Ott, Christian D. Probing the core-collapse supernova mechanism with gravitational waves. *Classical and Quantum Gravity* 26, 20 (2009), 204015.
- [103] Piran, T. et al. The long, the short and the weak: the origin of gamma-ray bursts. *Phil. Trans. R. Soc. A* 371 (2013), 1992.
- [104] Piran, Tsvi. The physics of gamma-ray bursts. *Rev. Mod. Phys.* 76 (2005), 1143–1210.
- [105] Piro, A. L., and Pfahl, E. Fragmentation of Collapsar Disks and the Production of Gravitational Waves. *The Astrophysical Journal* 658, 2 (2007), 1173.
- [106] Piro, A. L., and Thrane, E. Gravitational Waves from Fallback Accretion onto Neutron Stars. *The Astrophysical Journal* 761, 1 (2012), 63.

- [107] Pozzo, W Del. Measuring the Hubble constant using gravitational waves. *Journal of Physics: Conference Series* 484, 1 (2014), 012030.
- [108] Price, D. J., and Rosswog, S. Producing Ultrastrong Magnetic Fields in Neutron Star Mergers. *Science* 312, 5774 (2006), 719–722.
- [109] Read, Jocelyn S., et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys. Rev. D* 79 (2009), 124033.
- [110] Rezzolla, L., et al. The missing link: Merging neutron stars naturally produce jet-like structures and can power short Gamma-Ray Bursts. *The Astrophysical Journal Letters* 732, L6 (2011).
- [111] Riles, Keith. Notes on Relating Sine-Gaussian Waveforms to Astrophysical Source Strengths, 2004. LIGO-T040055.
- [112] Romero, G. E., Reynoso, M. M., and Christiansen, H. R. Gravitational radiation from precessing accretion disks in gamma-ray bursts. *Astronomy & Astrophysics* 524 (2010), A4.
- [113] Rowlinson, A., et al. Discovery of the afterglow and host galaxy of the low-redshift short GRB 080905A. *Monthly Notices of the Royal Astronomical Society* 408, 1 (2010), 383–391.
- [114] Saulson, Peter R. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Pub Co Inc, 1994.
- [115] Sekiguchi, Y., Kiuchi, K., Kyutoku, K., and Shibata, M. Gravitational Waves and Neutrino Emission from the Merger of Binary Neutron Stars. *Phys. Rev. Lett.* 107 (Jul 2011), 051102.
- [116] Shibata, M., and Taniguchi, K. Merger of black hole and neutron star in general relativity: Tidal disruption, torus mass, and gravitational waves. *Phys. Rev. D* 77 (Apr 2008), 084015.
- [117] Sidles, John A., and Sigg, Daniel. Optical torques in suspended Fabry-Perot interferometers. *Physics Letters A* 354, 3 (2006), 167 – 172.
- [118] Siellez, K., Boer, M., and Gendre, B. Simultaneous event detection rates by electromagnetic and gravitational wave detectors in the advanced era of LIGO and Virgo. *Mon. Not. R. Astron. Soc.* 437, 1 (2014), 649–655.
- [119] Singer, L. P., et al. Discovery and Redshift of an Optical Afterglow in 71 sq deg: iPTF13bxi and GRB130702A. *The Astrophysical Journal Letters* 776, 2 (2013), L34.
- [120] Smith, J. R., et al. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Classical and Quantum Gravity* 28, 23 (2011), 235005.
- [121] Smith, N. *Techniques for Improving the Readout Sensitivity of Gravitational Wave Antennae*. PhD thesis, Massachusetts Institute of Technology, 2012.
- [122] Smith-Lefebvre, N., et al. Optimal alignment sensing of a readout mode cleaner cavity. *Opt. Lett.* 36, 22 (Nov 2011), 4365–4367.
- [123] Soderberg, A. M., et al. Relativistic ejecta from X-ray flash XRF 060218 and the rate of cosmic explosions. *Nature* 442 (2006), 1014–1017.
- [124] Staley, A. *Achieving Resonance in the Advanced-LIGO Gravitational Wave Detector: with a focus on the Arm Length Stabilization Technique*. PhD thesis, Columbia University, 2015.
- [125] Staley, A., and Hoak, D., et al. High precision optical cavity length and width measurements using double modulation. *Opt. Express* 23, 15 (Jul 2015), 19417–19431.
- [126] Staley, A., and Martynov, D., et al. Achieving resonance in the Advanced LIGO gravitational-wave interferometer. *Classical and Quantum Gravity* 31, 24 (2014), 245010.

- [127] Sutton, P. J., et al. X-Pipeline: an analysis package for autonomous gravitational-wave burst searches. *New Journal of Physics* 12, 5 (2010), 053034.
- [128] Sutton, P. RayleighMonitor Overview, 2004. LIGO G040422.
- [129] Sutton, Patrick J. A Rule of Thumb for the Detectability of Gravitational-Wave Bursts, 2013. arxiv:1304.0210.
- [130] Tanvir, N. R., et al. An origin in the local Universe for some short gamma-ray bursts. *Nature* 438 (2005), 991–993.
- [131] Tanvir, N. R., et al. A ‘kilonova’ associated with the short-duration γ -ray burst GRB130603B. *Nature* 500, 7 (Aug. 2013), 547–549.
- [132] Taylor, Stephen R., Gair, Jonathan R., and Mandel, Ilya. Cosmology using advanced gravitational-wave detectors alone. *Phys. Rev. D* 85 (2012), 023535.
- [133] Thorne, K. S. . In *Three hundred years of gravitation*, S. W. Hawking and W. Israel, Eds. Cambridge University Press, 1987, pp. 330–458.
- [134] van Putten, M. H. P. M. Proposed Source of Gravitational Radiation from a Torus around a Black Hole. *Phys. Rev. Lett.* 87 (Aug 2001), 091101.
- [135] van Putten, M. H. P. M., Levinson, A., Lee, H. K., Regimbau, T., Punturo, M., and Harry, G. M. Gravitational radiation from gamma-ray burst-supernovae as observational opportunities for LIGO and VIRGO. *Phys. Rev. D* 69 (Feb 2004), 044007.
- [136] Vaulin, R., et al. Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data. *Phys. Rev. D* 88 (Sep 2013), 062003.
- [137] Virgili, F. J., Liang, E., and Zhang, B. Low-luminosity gamma-ray bursts as a distinct GRB population: a firmer case from multiple criteria constraints. *MNRAS* 392 (2009), 91–103.
- [138] von Kienlin, Andreas, et al. The Second Fermi GBM Gamma-Ray Burst Catalog: The First Four Years. *The Astrophysical Journal Supplement Series* 211, 1 (2014), 13.
- [139] Ward, R. *Length Sensing and Control of a Prototype Advanced Interferometric Gravitational Wave Detector*. PhD thesis, California Institute of Technology, 2010.
- [140] Was, M. *Searching for gravitational waves associated with gamma-ray bursts in 2009-2010 LIGO-Virgo data*. PhD thesis, Universite Paris-Sud 11, 2012.
- [141] Was, M., Sutton, P. J., Jones, G., and Leonor, I. Performance of an externally triggered gravitational-wave burst search. *Phys. Rev. D* 86 (Jul 2012), 022003.
- [142] Weber, J. Observation of the Thermal Fluctuations of a Gravitational-Wave Detector. *Phys. Rev. Lett.* 17 (1966), 1228–1230.
- [143] Weisberg, J. M., Nice, D. J., and Taylor, J. H. Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16. *The Astrophysical Journal* 722, 2 (2010), 1030.
- [144] Weiss, Rainer. Electromagnetically Coupled Broadband Gravitational Antenna. *Tech. Rep. (MIT)* (Apr 1972). LIGO-P720002.
- [145] Weiss, Rainer. Collection of Reports on Barkhausen Noise, 2009. LIGO T0900061.
- [146] Will, Clifford M. The Confrontation between General Relativity and Experiment. *Living Reviews in Relativity* 9, 3 (2006).
- [147] Winkler, C, et al. The INTEGRAL mission. *Astronomy and Astrophysics* 411, 1 (2003), L1–L6.

- [148] Woosley, S. E., and Bloom, J. S. The Gamma-Ray Burst - Supernova Connection. *Annual Review of Astronomy and Astrophysics* 44 (2006), 507–556.
- [149] Woosley, S. E., Eastman, Ronald G., and Schmidt, Brian P. Gamma-Ray Bursts and Type Ic Supernova SN 1998bw. *The Astrophysical Journal* 516, 2 (1999), 788.
- [150] Wright, E. L. A Cosmology Calculator for the World Wide Web. *Publications of the Astronomical Society of the Pacific* 118, 850 (2006), pp. 1711–1715.