

2017

The Effects of Instrumental Noise on Searches for Generic Transient Gravitational Waves in Advanced LIGO

Marissa Walker

Louisiana State University and Agricultural and Mechanical College

Follow this and additional works at: https://digitalcommons.lsu.edu/gradschool_dissertations



Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Walker, Marissa, "The Effects of Instrumental Noise on Searches for Generic Transient Gravitational Waves in Advanced LIGO" (2017). *LSU Doctoral Dissertations*. 4349.

https://digitalcommons.lsu.edu/gradschool_dissertations/4349

This Dissertation is brought to you for free and open access by the Graduate School at LSU Digital Commons. It has been accepted for inclusion in LSU Doctoral Dissertations by an authorized graduate school editor of LSU Digital Commons. For more information, please contact gradetd@lsu.edu.

THE EFFECTS OF INSTRUMENTAL NOISE ON SEARCHES FOR GENERIC TRANSIENT
GRAVITATIONAL WAVES IN ADVANCED LIGO

A Dissertation

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Physics and Astronomy

by

Marissa Walker

B.S., Abilene Christian University, 2010

May 2017

ACKNOWLEDGMENTS

There are so many people I want to thank for helping me get through the past few years - people who have not just helped me survive graduate school but have truly made this a wonderful experience. I'm grateful to have so many excellent and admirable people in my life. If I took the time to give everyone the thanks they deserve here, my dissertation would double in length.

First of all, thanks to Gabriela González for being a wonderful advisor. Your passion for research is contagious, and your love of teaching is always evident when you patiently explain new concepts to your students. I'm so grateful I had the opportunity to learn from such a wise and kind mentor.

Thanks to all of the LSU LIGO group. Duncan Macleod, for graciously answering a million questions, teaching me about python, providing me with encouragement and caffeine, and making us all be more sociable. Thanks to the other wonderful LSU postdocs Keiko, Marie, and Adam, who have been good role models and friends, and my fellow LSU students - including Thomas, Jon, Nutsinee, Chris, and many others - who learned about LIGO along with me and made our dingy basement office a really pleasant place to be throughout the years. Thank you all for your help in understanding the intricacies of gravitational wave physics, your support during my anxious preparations for papers or presentations, and your incredible ability to make me smile and laugh even when I was most overwhelmed. And Shivaraj, an honorary member of Team LSU, thank you for being a great source of encouragement and chocolate.

Thanks to the LIGO detector characterization team. Josh Smith, Andy Lundgren, and David Shoemaker, for leading the detector characterization group and making it a positive, welcoming, and productive environment. Jess McIver, for introducing me to the world of subsystem characterization and offering great advice on research, writing, and life. Special thanks to Brennan Hughey, TJ Massinger, and Laura Nuttall, for being amazing data quality experts and for developing many of the vetoes used to improve the searches. I've really enjoyed working with you and learning from you, and the work on burst data quality presented here would not be possible without you. Thanks

to Chris Biwer and all my fellow fellows at Hanford who made the three months I spent there so wonderful and memorable.

Thanks to the instrumental experts in LIGO who shared their knowledge with me, especially to Stuart Aston and Jeff Kissel for answering many questions about suspensions.

Thank you to Marco Drago, Sergey Klimenko, Gabriele Vedovato, Shubhanshu Tiwari, and the rest of the cWB analysis team, for answering all my questions about burst searches, helping me learn to run cWB, and producing extra plots and trigger sets for me to use.

Outside of work, many others have supported me in meaningful ways. Courtney, you are a magnificent person, you made Baton Rouge feel like home to me, and I cannot imagine my life without you. I owe so much to you and my other fantastic roommates here - Meagan, Ana, and JP. Thank you all for walking by my side through all of the joys and challenges of grad school. Minter family and college roommates, especially Sarah, Kelsey, Mark, and Noemí, who have continued to send me love and support from far away. Thanks for being my forever friends.

Thanks to my other dear friends here in Baton Rouge who have enriched my life here. Lucy, for being an amazingly hospitable person who has brought so much joy to my life. Kristen and Brad, for the best coffee discussions and lovely hospitality. Tia and all the Landry family, for being excellent neighbors and beautiful friends throughout my time here. And to all of my great friends from South who have made me feel welcomed and loved.

Last but certainly not least - Mom, Dad, Martin, Tsunami, Marcella, Ben, and Simon. You're the best family anyone could ask for, and I could not have done this without your love and support. You all mean more to me than I could ever express in words.

PREFACE

The research presented in this thesis was carried out within the LIGO Scientific Collaboration (LSC). Although much of this work stemmed from discussion and collaboration with LSC members, this dissertation has not been reviewed by the collaboration. Opinions expressed here are my own and not necessarily those of the LSC.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iii
PREFACE	iv
LIST OF TABLES	vii
LIST OF FIGURES	viii
ABSTRACT	xii
CHAPTER	
1 GRAVITATIONAL WAVES	1
1.1 Gravitational Waves	2
1.1.1 Linearized Gravity	2
1.1.2 Astronomical sources of gravitational waves	5
1.2 Gravitational wave detectors	8
1.3 Advanced LIGO	9
1.3.1 Noise sources	11
2 SEARCHING FOR UNMODELED GRAVITATIONAL WAVES	14
2.1 Unmodeled sources of gravitational waves	14
2.2 Generic searches for unmodeled sources	15
2.2.1 Coherent WaveBurst search method	15
2.2.2 Estimating false alarm rates	18
2.2.3 Effects of non-stationary noise on the background	19
2.3 Detector Characterization	20
2.3.1 Finding instrumental glitches	20
2.3.2 Analyzing glitch morphology	21
2.3.3 Analyzing correlations between glitches	27
3 DATA QUALITY IN O1: TRANSIENTS	30
3.1 Comparison between cWB and Omicron	31
3.2 O1 Data Quality	35
3.2.1 Category 1	39
3.2.2 Category 2	44
3.2.3 Category 3	50
3.2.4 Summary of veto impact on Omicron triggers	51
4 SEARCH RESULTS IN O1: GRAVITATIONAL WAVES	54
4.1 Remaining glitches	54
4.1.1 Blip glitch investigations	55
4.1.2 cWB search bins	59
4.2 O1 search results	61
4.2.1 Data quality checks	62

4.2.2	GW150914 and search background	63
4.3	A new window to the universe	66
5	TRANSIENTS IN SUSPENSION SYSTEMS	69
5.1	Advanced LIGO Suspensions	69
5.2	Motion transients in suspensions	73
5.2.1	Typical suspension behavior	74
5.2.2	Suspension transients in O1	75
5.2.3	Correlations between suspensions stages	80
5.3	Propagation of motion transients	81
5.3.1	Correlations with gravitational wave strain data	86
5.4	Conclusions	90
6	CONCLUSIONS	91
	REFERENCES	93
	VITA	99

LIST OF TABLES

3.1	Summary of vetoes applied to the single detector Omicron triggers with SNR greater than 10. Deadtime and efficiency are given as a percentage of the remaining time and triggers left after the previous category.	53
4.1	A few of the most significant channels correlating with blip times from cWB.	57

LIST OF FIGURES

1.1	An illustration of the effects of the two gravitational wave polarizations on a ring of test masses.	6
1.2	An illustration of the effects of a gravitational wave on a Michelson interferometer. As the length of each arm changes, the interference of the returning beams changes because of the varying phase difference caused by differences in the optical path length.	9
1.3	Basic optical path of the Advanced LIGO interferometers.	10
1.4	Estimated noise levels predicted to limit the final Advanced LIGO design sensitivity. .	12
1.5	Sensitivity of the Advanced LIGO detectors during the first observing run.	12
2.1	GW150914 as seen by each detector in different time-frequency resolutions.	16
2.2	The coherent WaveBurst background distribution from O1, compared with a background of Gaussian noise.	19
2.3	Single detector Omicron triggers in the 10 minutes surrounding GW150914.	22
2.4	Omega scan images of GW150914.	23
2.5	Effects of glitches near 500 hz and harmonics on the Livingston strain Omicron triggers.	24
2.6	Omega scan of one of the short duration glitches near 500 Hz and harmonics which affected the instrumental background in early engineering runs	26
2.7	A comparison of the frequency distribution of the background events before and after violin mode damping was modified on June 6.	26
2.8	Example glitch caused by the beating of different radio frequency controls.	27
3.1	Receiver Operator Characteristic (ROC) curve showing correlation between Livingston (L1) and Hanford (H1) single detector Omicron triggers and cWB background, before vetoes. The efficiency is shown as the percentage of cWB background triggers which are coincident with a single detector Omicron trigger. The asterisks mark the efficiency and false alarm for time windows of 0.1 and 0.5 seconds on either side of the cWB trigger time, as reference points.	33
3.2	Properties of cWB background triggers not coincident with an Omicron trigger from either detector.	34

3.3	Comparison of cWB properties with single detector Omicron triggers coincident within 0.1 seconds.	36
3.4	ROCs of cWB and Omicron triggers, showing the fraction of Omicron triggers coincident with cWB background events.....	37
3.5	The impact of Category 1 data quality vetoes on single interferometer Omicron triggers from Livingston.	41
3.6	Spectrum comparison between a typical time during O1 and the hours during the Beckhoff system error.	42
3.7	The impact of Category 1 data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford during O1	43
3.8	Omega plots showing H1 DAC saturations.	45
3.9	The impact of the ETMY saturation flags and ETMY saturation with Omicron SNR greater than 200 on cWB background triggers from September 12 to October 20.	46
3.10	The impact of Category 2 data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford.	47
3.11	Elevated seismic motion corresponding to increased glitch rate in GW strain.	48
3.12	These flags were used for the CBC search but not implemented for the generic searches.	49
3.13	ROC Curve showing the efficiency of the most significant auxiliary channel in the hveto analysis from September 12 to 26, at three different SNR thresholds placed on the auxiliary triggers.....	51
3.14	The impact of data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford during O1, shown by the distribution of transients in Signal-to-Noise Ratio (SNR).....	52
3.15	The impact of data quality vetoes from each category on single interferometer Omicron triggers from Livingston and Hanford, showing the SNR and frequency of the triggers vetoed.	52
4.1	A typical blip glitch time-frequency morphology.	55
4.2	Several blips from each detector, at different <i>non-coincident</i> times during O1.	56

4.3	The potential effects of applying a veto based on the auxiliary channel with the highest significance in correlations to blip glitches.	58
4.4	A nonstationary line in the strain data creates excess cWB background triggers around its frequencies as the feature moves between 520-530 Hz.....	60
4.5	Coherent WaveBurst background before and after signal-based cuts to separate blips into a separate bin, and the constrained background after the application of data quality vetoes.	61
4.6	The effects of vetoes on the ρ distribution before and after signal constraints, during the analysis period of GW150914, September 12 to October 20, 2015.	62
4.7	The cWB background distribution over time during the analysis of GW150914.	64
4.8	GW150914 as seen by the Livingston and Hanford detectors.	65
4.9	The cWB search background and results from September 12 to October 20, 2015.	67
5.1	Quadruple suspensions design.	70
5.2	Bode diagram for the modeled longitudinal displacement from the top stage of the quadruple suspension to the two subsequent stages.	71
5.3	OSEM schematic	72
5.4	OSEM coherence	73
5.5	Typical amplitude spectral density (ASD) for the ITMY longitudinal motion monitors from O1.	75
5.6	An example of non-stationary noise in the suspension monitors.	76
5.7	Omicron trigger distributions from the ITMY suspension monitors from one week of data during O1.	77
5.8	ETMY high frequency triggers in O1	78
5.9	Time series around loud high frequency glitch in ETMX.....	79
5.10	A filtered time series in ITMY displacement and locking channels, showing a correlation between actuation from the control signals and the resulting displacement of the suspension.....	80

5.11	ROC curves showing correlations between Omicron triggers in different stages of the ITMY suspension.....	81
5.12	Modelled transient response of a simple pendulum to a sine-Gaussian injection.	83
5.13	Amplitude distribution of injections at the top stage of the suspension, for all injections and for those coincident with lower stages.	84
5.14	Ratio of amplitudes of injection Omicron triggers at lower stages to the top stage, plotted against the peak frequency estimated by Omicron for the top stage motion, compared with the modeled transfer function.	84
5.15	Time series from one injection at 4.1 Hz, after application of a bandpass filter with a window of 1 Hz around the injection frequency.	85
5.16	Amplitude ratios as calculated by the Omega algorithm, and frequency estimated using the maxima and minima of bandpassed time series.	87
5.17	Receiver Operator Characteristic (ROC) curves showing the correlation between noise transients in the GW strain and ITMY suspension data from November 1 to 8, 2015.....	88
5.18	Livingston ITMY triggers from a week in the first observing run, multiplied by the transfer function to the lowest stage and divided by the arm length to convert the displacement into equivalent strain amplitude.	89

ABSTRACT

One hundred years after Albert Einstein predicted the existence of gravitational waves as a result of his theory of general relativity, the Laser Interferometer Gravitational-Wave Observatory (LIGO), made the first direct detection of a gravitational-wave signal from a binary black hole merger, GW150914. GW150914 was found not only by search methods specifically developed to find the distinctive waveform produced by coalescing binaries, but also by generic searches designed to find any arbitrary short-duration signal in the LIGO data. The impact of noise on the searches must be carefully investigated in order to reduce the search background and enable confident gravitational-wave detections. In this dissertation, I will present my work on characterizing transient noise sources in the detectors and implementing data quality vetoes to reduce their effects on the generic transient gravitational-wave searches.

Chapters 3 and 4 describe my work on the data quality of the searches for generic transient gravitational waves. I worked on the development of data quality vetoes during the first observing run and the decisions about which vetoes to implement in the transient searches. I also analyzed the transient noise sources that the vetoes were unable to eliminate, using statistical methods to search for potential instrumental causes. Since the development of data quality vetoes requires a thorough understanding of every component of the detectors, I have also conducted a detailed investigation into the transients in the suspension systems used to isolate the LIGO optics from seismic motion. Chapter 5 presents the details of this work.

The first gravitational wave detection was only the beginning an exciting era of gravitational-wave astronomy that will give us a new way of understanding the universe. Even in the first observing run, a second binary black hole merger was observed. The methods used in this dissertation to investigate and reduce background noise will continue to play an important role in making these detections possible. As the detectors improve in the future and continue to take data, more signals will be detected, bringing us a wealth of new information about black holes and other types of sources.

CHAPTER 1

GRAVITATIONAL WAVES

In 1915, Albert Einstein published his General Theory of Relativity, which describes gravity as the interaction between matter, energy, space, and time.[1] Extending gravitational theory beyond Newton's law of gravitation to be reconciled with special relativity, this formulation of the theory of gravity in terms of geometry and relativity revolutionized our understanding of the universe. While in a weak gravitational field the theory agrees with the Newtonian picture, general relativity introduced several important theoretical insights and predictions beyond classical physics. It provided a satisfactory explanation for the anomalous precession of Mercury's orbit, which disagreed with Newtonian calculations. Einstein also used the theory of relativity to accurately predict the bending of starlight around the sun, which was famously first observed by Arthur Eddington and his colleagues in 1919, and further confirmed by later measurements.

In 1916, Einstein predicted the existence of gravitational waves as a result of the theory of relativity. [2] Analogous to the acceleration of charge that produces electromagnetic waves, accelerated mass can dynamically change the curvature of spacetime in such a way that produces ripples in the gravitational field, which travel outward at the speed of light. As these waves travel, they stretch and squeeze space itself, transverse to the direction of propagation. While several of Einstein's predictions were confirmed within a few decades after his initial formulation of the theory of relativity, gravitational waves remained a mystery for much longer. In fact, the existence of gravitational waves as a physical reality remained controversial among physicists until a conference in 1957 on the role of gravitation in physics.[3] Indirect detection of the effects of gravitational radiation on a binary pulsar provided evidence for the theory by 1982,[4] but direct observation took over thirty more years of careful planning and experimentation. Finally, on September 14, 2015, a gravitational wave signal from the merger of two black holes was detected by the Laser Interferometer Gravitational-Wave Observatory (LIGO). [5]

This dissertation will present only a small but exciting piece of the 100-year history of gravitational waves, relating to the tuning of the detectors, data analysis, and validation of the ground-breaking gravitational wave detection through careful instrumental investigations. To begin, this chapter will briefly discuss the prediction of gravitational waves from the theory of general relativity, an overview of some astrophysical sources, and an introduction to the instrumental methods used by ground-based interferometric gravitational wave detectors, specifically LIGO.

1.1 Gravitational Waves

Gravitational radiation is a key prediction of the theory of general relativity, though conceptually the idea of gravitational waves is a natural consequence of applying the principles of special relativity to classical gravity. In fact, even without using the thorough formalization of general relativity, one can use relativistic corrections to the Newtonian gravitational potential to make rough approximations on the nature of gravitational waves. [6] In Newton's law of universal gravitation, the force of gravity acts instantaneously between masses, regardless of the distance between them. If one object moves or its mass changes, the gravitational force on the other mass is altered with no time delay. This action violates the laws of special relativity, which dictate that nothing can travel faster than the speed of light. To communicate any information about movement of one mass to the other, there must be some kind of radiation that travels at a finite speed: a gravitational wave.

1.1.1 Linearized Gravity

In general relativity, the relationship between matter and spacetime is expressed by the Einstein field equations:

$$R_{\alpha\beta} - \frac{1}{2}g_{\alpha\beta}R = \frac{8\pi G}{c^4}T_{\alpha\beta} \quad (1.1)$$

The lefthand side of the differential equation contains information about the geometrical nature of spacetime, described by the spacetime metric $g_{\mu\nu}$, from which is derived the Riemann tensor, $R_{\alpha\beta}$, and the Ricci scalar, R . These tensors define the curvature of space and time. On the righthand side is $T_{\alpha\beta}$, the stress-energy tensor, which represents the mass and energy of the system.

The stress-energy tensor is the relativistic generalization of the Newtonian stress tensor, which uses only the spacial dimensions. Essentially, each component $T^{\alpha\beta}$ gives the α component of the momentum flux of the across a surface in the direction of x^β . Within the stress-energy tensor, one can thus find the mass density, energy, pressure, and momentum of the system in all dimensions. This equivalency between the geometrical components of spacetime and the mass and energy of a system demonstrates the connection between spacetime and matter. The matter and energy warp the surrounding spacetime, and the curved spacetime dictates the motion of the matter.

As simple as Equation 1.1 may appear, the Einstein field equations are actually ten coupled, nonlinear differential equations, and exact solutions can only be found using some special assumptions. Manipulating the equations to predict the intricacies of gravitational waveforms involves advanced computational techniques. So for the purposes of demonstrating the prediction of gravitational radiation from from the theory of general relativity, we will only use the weak-field approximation, so that the spacetime metric can be described as flat (Minkowski) with only a small added perturbation. The Minkowski metric describing flat spacetime is written here in Cartesian coordinates (t,x,y,z) as $\eta_{\alpha\beta}$:

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

Flat spacetime can also be represented equivalently using the spacetime interval:

$$ds^2 = -cdt^2 + dx^2 + dy^2 + dz^2 \tag{1.2}$$

Although the Einstein equations tell us that the curvature will be influenced by mass and energy, we can approximate the metric far away from large mass as the flat spacetime with only a small perturbation, $h_{\alpha\beta}$, such that the magnitude $h \ll 1$:

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta} \quad (1.3)$$

The Riemann curvature tensor can be greatly simplified using approximations to first order in h , and by taking advantage of gauge freedom, the gravitational wave equation results from the Einstein field equations. Away from the matter at the source of the gravitational waves, the stress-energy tensor becomes zero and the wave equations in the weak field far from the source becomes:

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\alpha\beta} = 0 \quad (1.4)$$

This wave equation shows that gravitational waves are a direct consequence of the theory of general relativity.

In the transverse and traceless gauge, solutions for a gravitational wave traveling in the z direction are of the form:

$$h = Ae^{2\pi ft - \mathbf{k} \cdot \mathbf{x}} \quad (1.5)$$

where the amplitude of the wave is given by A , the frequency by f , and the wave number \mathbf{k} . For a wave traveling in the z direction, the transverse traceless gauge provides a simple form of $h_{\alpha\beta}$:

$$h_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_x & 0 \\ 0 & h_x & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

The wave is *transverse* because the strain is only in the x and y dimensions; the z-component of the strain is zero. This tensor is traceless because the sum of the diagonal components is zero. As can be seen from the opposite signs of h_{11} and h_{22} in Equation 1.1.1, the strain has the effect of contracting one dimension while simultaneously lengthening another. The magnitudes h_+ and h_x demonstrate the fact that the gravitational waves can then be represented as a linear combination of two polarization states. The diagonal terms are called the “plus” polarization h_+ because of the strain effects in the x and y directions. The off-diagonal terms are the “cross” polarization h_x , which is rotated by 45 degrees like an X. Figure 1.1 illustrates the effects of these two polarizations on a ring of test masses. For the plus polarization, as the wave passes through the ring, the distance between masses is first increased in the y-dimension while being decreased in the x-dimension. Half of a cycle later, the x-dimension is stretched while the y-dimension is squeezed.

1.1.2 Astronomical sources of gravitational waves

The strength of a gravitational wave is described by its strain, h , the proportional change in separation L between free masses. The production of gravitational waves can be understood by using the Einstein field equations near to the source, instead of the weak field limit. The dominant term in the amplitude of the strain is proportional to the second derivative of the gravitational quadrupolar moment of the source, $\ddot{I}_{\mu\nu}$, and inversely proportional to the observer’s distance from the source, R [8]:

$$h = \frac{\Delta L}{L} \simeq \frac{2G}{Rc^4} \ddot{I}_{\mu\nu} \quad (1.6)$$




Figure 1.1: An illustration of the effects of the two gravitational wave polarizations on a ring of test masses. [7]

Masses and gravitation near Earth are not nearly great enough to produce detectable gravitational waves, but accelerations and violent interactions of massive objects far out in the universe could make such a splash in space-time that the ripples will reach the earth with enough strength to be detected. Even so, estimates of the strength of the waves predict that they should only change the separation between masses on earth by a tiny scale factor. For example, using rough estimates of the quadrupolar moment and distance of a neutron star binary near coalescence in the Virgo cluster, the above equation yields a strain of just 10^{-21} .

Some sources of gravitational waves are predicted to emit continuous waves over a long period of time.[9, 10] For example, non-axisymmetric rapidly rotating neutron stars and orbiting binaries should produce gravitational radiation at twice their orbital frequency, which can be modeled by a sinusoid of fairly steady frequency over a long time. For many pulsars, this frequency is known precisely from astronomical observations. Different models for calculating the upper-limits on the expected strength of these continuous waves yield strains of 10^{-24} or smaller. Less well-modeled long-duration signals come from sources such as the stochastic background of primordial

gravitational radiation, analogous to the cosmic microwave background, or from the incoherent background made up of the superposition of waves from many sources.

Other gravitational waves are emitted from violent astrophysical events over a short period of time. The best modeled transient sources are compact binary systems in the final stages of spiraling in towards one another and merging. The huge acceleration of these highly massive and dense systems provides a large amount of energy released as gravitational radiation. As the orbiting bodies approach their final coalescence, the orbital frequency and the amplitude of the gravitational waves increases rapidly, emitting a characteristic gravitational-wave signature. The inspiral phase of the waveform can be modeled with post newtonian approximations, while numerical relativity is needed to model the merger.[11] The first direct detection of gravitational waves, which will be discussed in greater detail in Chapter 3, came from the merger of two stellar mass black holes, with a signature well-matched to the predictions from relativistic models. Similar signals would also be produced from the coalescence of neutron stars or larger black holes, but with important differences in the evolution of waveforms.

Finally, gravitational wave sources could come in short-duration bursts from sources that are not well modeled, or from sources that deviate unexpectedly from the model. For example, core-collapse supernovae should produce gravitational radiation, but many uncertainties in the models make the waveform difficult to predict. For a supernova at a distance of 10 kiloparsecs from earth, emitting an energy of $10^{-7} M_{\odot} c^2$ in gravitational radiation, the strain amplitude would be roughly 6×10^{-21} . [9]

The various possibilities for astrophysical gravitational wave sources creates an exciting potential for gravitational wave detectors to make interesting astronomical observations. Detected waves could offer insight on models of core collapse mechanisms,[12] constrain predictions of rates of binary mergers,[13] and uncover new information about black holes, which cannot be seen electromagnetically. The first gravitational wave detection alone was an advance in astrophysics; its implications discussed in [14] include proving the existence of intermediate mass black holes and

black hole mergers, as well as showing consistency with some models of event rates, and severe constraints on others. [15, 16]

1.2 Gravitational wave detectors

In 1974, Joseph Taylor and Russell Hulse discovered a binary pulsar system,[17] and in the following years the period of the orbit was observed to decay by the precise amount that would be predicted by the energy loss due to gravitational radiation.[4] These measurements provided strong experimental evidence for the existence of gravitational waves and increased interest in developing experiments to directly observe gravitational waves, which had already begun with the work of Joseph Weber. [18] The first gravitational wave detectors were resonant aluminum cylinders, pioneered by Weber. These bar detectors were designed to respond to gravitational waves at their resonant frequencies. Although the mass detectors improved over time and with the addition of cryogenic detectors, the sensitivity never reached the sensitivity necessary to detect any gravitational waves.

The idea for building gravitational wave detectors using interferometry was first proposed in the early 1970s. Interferometric detectors build upon the basic Michelson-Morley interferometer: a laser beam is sent through a beam splitter, reflects off of two end mirrors, and recombines again at the beam splitter. The returning light, sensed by a photodetector at the output, carries with it information about the phase difference between the returning beams. As a quadrupolar gravitational wave with the appropriate polarization impinges upon the interferometer, the length of one arm increases as the other decreases, as illustrated in Figure 1.2. This variation in the path length difference will affect the relative phase between the returning beams, which in turn determines the intensity of light sensed at the output.[19]

The initial set of gravitational wave interferometers included the two Laser Interferometer Gravitational-wave Observatory (LIGO) detectors in the USA,[21, 22] the Virgo instrument in Italy,[23] and GEO 600 in Germany.[24] These detectors collected data over several years in the 2000s, gradually increasing sensitivity with various improvements through several periods of observation. Beginning in 2010, the LIGO and Virgo detectors began to be decommissioned to make a




Figure 1.2: An illustration of the effects of a gravitational wave on a Michelson interferometer. As the length of each arm changes, the interference of the returning beams changes because of the varying phase difference caused by differences in the optical path length. [20]

series of major upgrades, including additional optical cavities, higher laser power, and stronger seismic isolation systems.[25, 26] Although the Advanced LIGO detectors have not yet reached their full design sensitivity, they were ready to begin searching for gravitational waves in September 2015, with a sensitivity over three times better than the initial detectors.

1.3 Advanced LIGO

The historic first direct gravitational wave detection was made by the Laser Interferometer Gravitational-Wave Observatory (LIGO), a project comprising two kilometer-scale interferometers in the USA.

To actually reach the sensitivity required to detect the tiny signals expected, the instrument is much more complex than a simple Michelson interferometer. Each component of the system has been designed to maximize the sensitivity to gravitational waves, and to minimize the effects of terrestrial noise sources. Figure 1.3 shows the basic optical path of the laser beam.

Instead of single end mirrors, LIGO has four-kilometer arms with Fabry-Pérot cavities to increase the effective optical path length. Instead of only making one trip through the beam

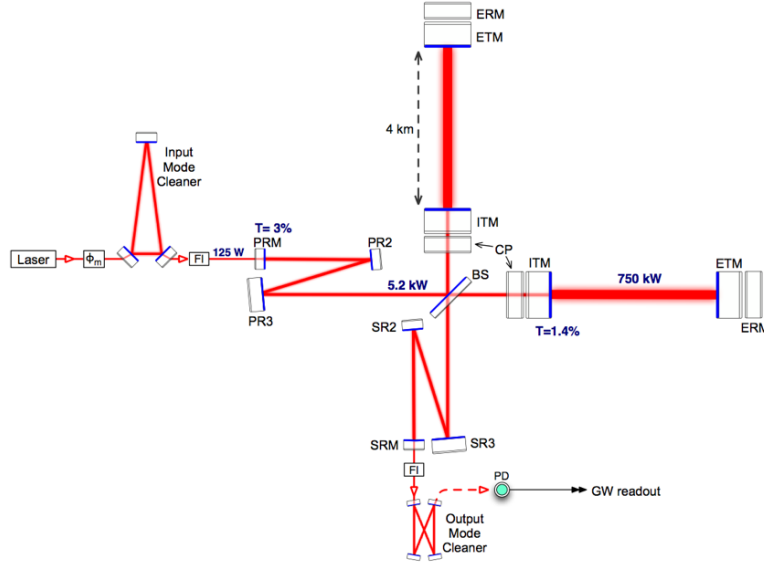


Figure 1.3: Basic optical path of the Advanced LIGO interferometers.[25]

splitter and down the arms, additional partially reflective mirrors are placed between the beam splitter and end mirrors, creating the resonant cavities in the arms. Light makes multiple trips in the resonant cavities, which increases the optical path length. Since the strain length change is proportional to the total length of the arms, the accumulated phase shift caused by a gravitational wave signal is amplified, and the sensitivity of the detector is increased by a factor of nearly 300 from a simple Michelson configuration.

Several additional optical cavities further improve the sensitivity of the detector. The Input Mode Cleaner (IMC), filters the laser light before it goes through the beam splitter, ensuring that the beam profile is the Gaussian 00 mode. The Power Recycling Cavity (PRC) acts as a complex mirror between the laser and the beam splitter, sending any reflected light back into the interferometer, further building up the resonant power in the arms. The Signal Recycling Cavity (SRC) at the output amplifies the gravitational wave signal before readout.

1.3.1 Noise sources

There are several noise sources that affect LIGO's sensitivity in different frequency regions. For each predicted source of noise, efforts have been made to improve the instrumentation of Advanced LIGO to push the limits of sensitivity. Figure 1.4 shows the estimated noise curve that can be eventually achieved with Advanced LIGO,[27] while Figure 1.5 is the actual sensitivity reached during the first observing run.[28]

The main limiting factor at low frequency is the seismic noise, varying with the earthquakes and anthropogenic activities. In the initial era of LIGO, seismic motion had a large effect on the quality of the data; nearby trains and trucks on the highway frequently created noticeable effects on the detectors' sensitivity. Advanced LIGO detectors incorporate stronger isolation of the optics from ground motion, using layers of active seismic isolation platforms, from which are hung multi-stage pendulums. The input and end mirrors of the two arms are the bottom stage of quadruple level suspensions, in order to approximate the behavior of free test masses. Chapter 5 discusses these suspensions in detail, and presents an investigation into their behavior during the first Advanced LIGO observing run.

In the mid frequency range, key noise limitations are the thermal effects from the wires of the suspensions and the Brownian noise in the test mass bulk material and mirror coatings. The design and materials of the mirrors and suspensions have been chosen with the goal of minimizing this noise source.[29, 30] The test masses and the fibers that are used to suspend them are made of fused silica, which has a high quality factor to lower the thermal noise. Coatings of the optics were specifically designed not only to create the high reflectivity needed but also to reduce the Brownian noise.

At higher frequencies, quantum effects are the limiting noise source, specifically from shot noise. Shot noise is the uncertainty that arises when trying to use a photodetector to sense fluctuations in power. Because of the quantum nature of photons, there is some natural variation in the number of photons that will arrive at the output of the detector, with a Poissonian distribution, regardless of the presence of a gravitational wave. In Advanced LIGO, the power of the laser has been increased




Figure 1.4: Estimated noise levels predicted to limit the final Advanced LIGO design sensitivity. Note that this sensitivity has not yet been achieved.[27]




Figure 1.5: Sensitivity of the Advanced LIGO detectors during the first observing run. [28]

(and will continue to be raised in the later observing runs) in order to lower the effects of shot noise, since at higher laser power the fractional fluctuation of the number of photons decreases. In the first observing run, the initial laser power was generally at 20 W, but because of the power amplification in the Fabry-Pérot cavities, the laser power circulating in the arms was 100 kW.

In addition to the constant noise sources that determine the average displacement sensitivity of the detectors, there are transient noise sources of various origins, varying with seismic motion, weather conditions, local anthropogenic activity, and the stability of the many electronics and servo loops used to control the interferometer. Short-duration noises, or *glitches*, may not have a sizeable impact on the average sensitivity spectrum, but affect the background of transient gravitational wave searches. If all of the noise in the instrument were stationary, then a short burst of energy caused by an incoming gravitational wave would be more easily distinguishable from the steady rate of background noise. With the level of transient noise in the detectors, any gravitational wave candidate in the data must be carefully examined against the background transients to determine how confident we can be that the signal is actually astrophysical in origin, rather than an instrumental glitch. The main topic of this dissertation is the process of studying transient noise and removing glitches from the gravitational wave searches, in order to reduce the transient background and improve confidence in gravitational wave detections.

CHAPTER 2

SEARCHING FOR UNMODELED GRAVITATIONAL WAVES

2.1 Unmodeled sources of gravitational waves

There are various astrophysical phenomena that could produce short duration bursts of gravitational waves, on the order of seconds or fractions of seconds within LIGO’s sensitive frequency band. For some of these sources, the exact form of the waves are not well known, due to uncertainties in physical models and limits in numerical simulations. Examples include supernovae and neutron star instabilities. Even for well-modeled sources, such as binary black hole mergers, the influence of high eccentricity and precession could drastically alter the waveforms.

Core collapse supernovae are one potential source of unmodelled transient gravitational waves.[31, 32] The highly energetic nature of supernova explosions makes them a possible source for gravitational radiation during several parts of the process, provided that there is some asymmetry involved. Gravitational waves could be emitted from the bounce as neutron degeneracy pressure increases in the core, when the matter in the star’s core experiences its highest accelerations. If the star is rotating, the accelerations at this stage will produce the rapidly varying quadrupole moment necessary for gravitational radiation. After the core bounce, gravitational waves could also be produced by the highly convective currents that provide energy to the supernova explosion, or by instabilities in oscillation modes of the resulting neutron star as it cools. Neutron stars also undergo energetic events such as pulsar glitches or magnetar flares, which could possibly produce gravitational radiation.

Although there are many scenarios in which supernovae and neutron stars are expected to produce gravitational radiation, the complicated dynamics of potential gravitational wave emission mechanisms are difficult to simulate due to numerical constraints. For example, accurate models need to include correct equations of state, models of the angular momentum of progenitor systems, and account for the effects of magnetic fields. Uncertainties in the theoretical factors of neutron

star evolution and seismology similarly lead to difficulties in understanding the possible effects of oscillation mode instabilities. The complexity of modeling all of these factors leads to large uncertainties in the gravitational wave signatures expected, both in waveform and amplitude.

2.2 Generic searches for unmodeled sources

With the possibility of detectable gravitational radiation from such exciting astrophysical events, it is important to perform gravitational wave searches that can accommodate all the uncertainty involved. These generic transient searches place minimal constraints on the expected waveforms. Rather than restricting the algorithms to identify particular waveforms, they search for any burst of excess power in the strain signal and use the coherence between the different detectors to constrain the results to potential astrophysical signals.

2.2.1 Coherent WaveBurst search method

One of the main methods of searching for unmodeled transient sources is Coherent WaveBurst (cWB).[33, 34] Developed during initial LIGO and upgraded to a newer version for Advanced LIGO, cWB is an algorithm designed to detect signals coincident between two detectors with minimal prior assumptions about the waveforms. Although its primary purpose is to search for transients with unknown waveforms, it also is able to find black hole binary mergers.[35] In fact, it was the first search to identify the black hole merger signal GW150914.[5]

The first step of the cWB search is conditioning the data from each of the sites by whitening the gravitational wave strain data to normalize the power across the frequency range of the search.[33] Filters are also applied in this stage to reduce the influence of predictable, stationary lines in the data, such as known mechanical resonances. The whitened time series is transformed into a time-frequency representation using the Wilson-Daubechies-Meyer wavelet basis, using several different time and frequency resolutions to more effectively cover the parameter space and to allow for the optimal characterization of a variety of signals.

Individual time-frequency tiles are clustered together, and the most energetic samples of the resulting time-frequency maps are chosen based on a power threshold. For example, Figure 2.1 shows the time-frequency pixels selected for the GW150914 event at two of the different resolution layers.

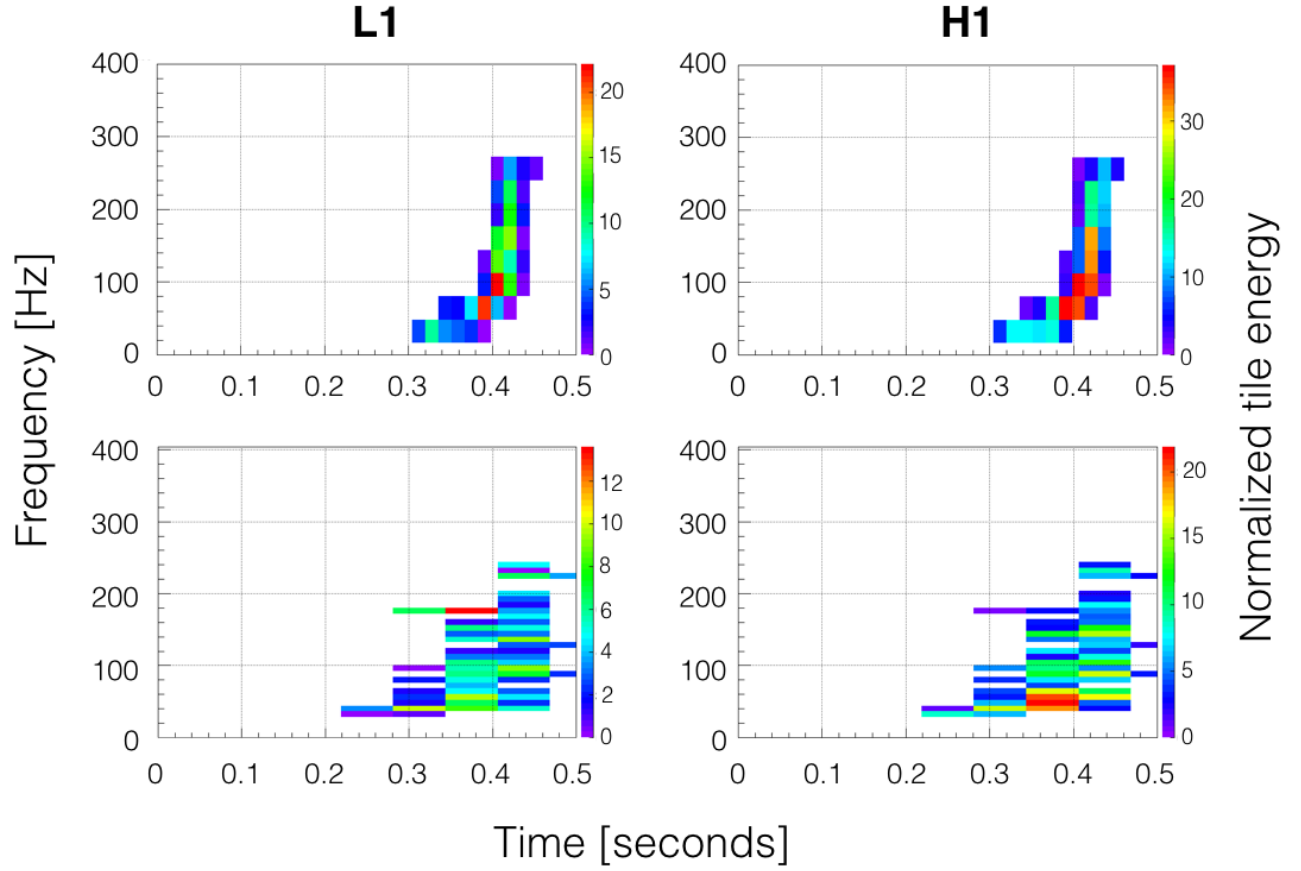


Figure 2.1: GW150914 as seen by each detector in different time-frequency resolutions. Each plot shows the pixels of maximum normalized energy selected for the different resolutions, with the color representing the tile energy. The top plots show the frequency resolution of 32 Hz and time resolution of 15.6 milliseconds. The frequency resolution of the bottom plots is 8 Hz, and their time resolution is 62.5 milliseconds. Plots produced by Gabriele Vedovato.

The waveforms of coincident events are then reconstructed using the maximum likelihood statistic, which is based on the correlation between the different detectors and energy of the waveforms. For each time-frequency cluster, the likelihood is calculated for all sky positions to determine the most likely source direction.

The noise-weighted coherent energy is the cross-correlated energy between different detectors, which is calculated from sum of the components of the likelihood statistic that involve two different detectors. The correlation coefficient, c_c , is defined as the ratio of the coherent energy to the total energy – residual noise added to the coherent energy:

$$c_c = \frac{E_c}{E_c + E_n} \quad (2.1)$$

For a real gravitational wave signal, the normalized coherent energy, E_c , should be high, while the residual noise energy E_n , which is the energy left after subtraction of the waveform from the data, should be quite low, leading to a correlation coefficient close to one. By setting a threshold on the correlation coefficient, the events recovered by the search are the most likely astrophysical events.

The gravitational wave candidates are further ranked by a significance statistic, η_c , which is dependent on the correlation coefficient and the amount of coherent energy. Defined by the following equation, this quantity is analagous to a signal-to-noise ratio, and can help select events with signals that are coherent between the sites with enough energy above the background noise of the detectors:

$$\eta_c = (2c_c E_c)^{1/2} \quad (2.2)$$

Any coherent burst event that passes the required thresholds on the correlation coefficient and coherent significance statistic is recorded by the algorithm as a single *event trigger*, which contains a summary of the key information about the event, such as its peak time, frequency, amplitude, and duration.

During the first observing run, the statistic ρ was also frequently used to represent the significance, which is just $\eta_c/\sqrt{2}$. Both η_c and ρ will be used in this dissertation to show the amplitude of cWB events.

2.2.2 Estimating false alarm rates

When a gravitational-wave search is performed, it is crucial to know whether we can be confident that a signal detected by the algorithm is actually astrophysical, rather than random noise, or a coincidence caused by the nonstationary instrumental or environmental noise of the detectors.

Therefore, in addition to searching for the coincident signals between detectors, the burst searches must also estimate how likely it is that noise in the detectors would create a false positive, called the background rate. To estimate this rate, the search must be performed over data that certainly does not contain gravitational waves, ideally for a long period of time, to find how often the algorithm detects a signal that is not astrophysical.

In the cWB search as well as other gravitational wave searches, the background rate is calculated by artificially introducing a time shift of one detector's data with respect to each other, by an amount greater than the light travel time between the interferometers. This ensures that if a gravitational wave is present in the data, it will no longer appear at the same time in each detector. The search algorithm is then run over this new data set, and the resulting events represent the chance that uncorrelated noise could produce a coherent signal between the detectors. One important underlying assumption in this analysis is that the transient noise that appears in the resulting background is a good representative of all of the glitches that could cause false alarms in the search results. Therefore, many time shifted analyses are performed to increase the available statistics to give false alarm rates over an analysis period equivalent to thousands of years. Additionally, it is vital to carefully characterize the background noise, to understand whether the kinds of glitches present occur throughout the analysis period or whether new types of noises appear.

After a thorough background analysis, any gravitational wave candidate can be compared to the background events to interpret its significance and state a false alarm rate.

2.2.3 Effects of non-stationary noise on the background

If the noise in the detectors followed Gaussian distributions, the rate of events would decrease for higher amplitudes. In such an idealized case, a real gravitational wave signal is easily distinguishable from the background, because its correlated energy will be highly significant.

However, the real background is quite different, due to large non-Gaussian glitches present in both detectors. Therefore the resulting backgrounds from initial and Advanced LIGO have tails in the distribution, loud outliers in significance, as shown in Figure 2.2. These glitches particularly affect the lower frequency range of the search, below a few hundred hertz. While modeled searches have the ability to reject glitches based on whether their waveforms conform to theoretical expectations, the generic transient searches are by definition meant to detect any potential gravitational waves. Louder events can only be excluded based on a confidence that they are glitches caused by an instrumental origin rather than an astrophysical event. To analyze the origins of instrumental noise, thorough characterization of the detectors themselves is necessary.




Figure 2.2: The coherent WaveBurst background distribution from O1, compared with a background of Gaussian noise. At each amplitude, the rate shown is the total rate of background events with a ρ greater than or equal to that amplitude. The amount of background time produced for the Gaussian set is less than for the O1 set (which is why the grey curve stops at a higher event rate) but even with the time shown the real distribution has diverged from the ideal stationary background. Rather than continuing to decrease with higher ρ , the actual rate begins to plateau due to high amplitude outliers. This tail in the distribution limits the false alarm rate that can be claimed for a gravitational wave detection.

2.3 Detector Characterization

To reduce the false alarm rate of gravitational wave candidates, it is vital to remove the events that form the tails of the background distribution, either by instrumental improvements to eliminate noise sources, or by the application of data quality vetoes to remove times when the detector data is known to contain glitches. In any case, it is necessary to understand the behavior of the detectors to diagnose transient problems. The LIGO detector characterization working group is a team of scientists dedicated to the analysis of the instrumental and environmental noise sources that affect the sensitivity of the detectors to gravitational waves.

When a noise source is better understood, the instrument can be improved or *data quality flags* can be produced to cut out compromised data from the search results. For example, in initial LIGO as well as in the first observing run of Advanced LIGO, vetoes were produced to eliminate times when there were errors in the calibration of the strain channel, excess seismic noise, saturations in control signals, and other environmental and instrumental disturbances.[36] Some of the methods for characterizing noise are introduced here, while specific O1 noise sources, vetoes, and their effects on the search will be discussed in greater detail in Chapter 3.

2.3.1 Finding instrumental glitches

In addition to the gravitational wave data, thousands of auxiliary channels are recorded to monitor the status of LIGO’s subsystems and the physical environment. Signals are recorded from sensors all around the interferometer: outside the building to record the local environmental conditions, in the multi-layered seismic isolation platforms, suspended test masses and other optics, laser stabilization systems, and the feedback control loops keeping various cavities of the detector locked and aligned. The detector characterization group uses these auxiliary channels to trace data quality issues to particular subsystems of the interferometer.

Glitches are detected in the auxiliary channel data using event trigger generators, algorithms designed to report short duration bursts of excess energy. Several different algorithms have been developed and used to search for transient noise.[37, 38, 39] The Omicron algorithm is the main one currently in use for detector characterization, after an extensive comparison of its performance with

others determined that it is one of the most efficient at finding glitches.[40] Omicron is not only a tool for detector characterization, but is also part of one of the main gravitational wave searches, *Omicron-LIB*, so it is inherently designed to locate events of the type that would affect the burst searches.[41] When used as a gravitational wave search, Omicron is run over both sites' gravitational wave channels to find events with excess power, and then it is combined with a parameter estimation algorithm, LIB, to find coherent events with reconstructed waveforms. Similar to the first steps of the cWB algorithm, Omicron takes whitened strain data and transforms it into the time-frequency domain with various time and frequency resolutions, using a sine-Gaussian basis rather than the wavelet transform.[37, 42] The resolutions used are defined by their quality factor, or Q , which is the central frequency divided by the bandwidth. In each different resolution, or Q -plane, the energy of all of the points, or Q -tiles, is computed and compared to the mean energy, and events with excess energy above a specified signal-to-noise ratio are selected. Figure 2.3 displays the Omicron triggers in both detectors in the time surrounding GW150914. Although the Omicron triggers associated with GW150914 are among the highest SNR in the short time period shown, many instrumental glitches cause Omicron triggers with much higher SNR, on the order of hundreds or even thousands. For this reason, it is important not only to find excess power in one detector, but to search in both detectors and find coherence between them.

For each resulting trigger, Omicron estimates various parameters of each event, including the central frequency and bandwidth, time and duration, integrated amplitude over the signal, and signal-to-noise ratio (SNR). This information given is useful for characterizing populations of glitches in the gravitational wave channel as well as the auxiliary channels. With the main properties of the events identified by Omicron, we can identify particular periods of time or frequency ranges that tend to have a high rate of glitches that require further investigation.

2.3.2 Analyzing glitch morphology

Although Omicron is a useful way of quickly finding the main properties of glitches, it is also often helpful to look at the specific morphology of different glitches. While Omicron can provide the peak frequency, bandwidth, central time, and estimated duration of a transient signal, it can




Figure 2.3: Single detector Omicron triggers in the 10 minutes surrounding GW150914. The asterisks represent the Omicron triggers in each detector at the time of the gravitational wave signal.

be useful to visualize how the frequency and amplitude of the signal evolve over time. Omicron and the related Omega algorithm both provide short-duration detailed spectrograms to show the features of the data in time and frequency, using the information from multiple Q -planes to display the evolution of the signal over time. As Figure 2.4 shows, the Omega scans of the gravitational wave signal GW150914 exhibit the distinct morphology of increasing frequency over time, which is characteristic of gravitational waves from binary mergers. Unlike the cWB figures shown in Figure 2.1, the scan shown in Figure 2.4 combines various resolutions to highlight the most prominent tiles from each one. While the beginning of the signal is better represented by showing the narrowband tiles with less time resolution, the more rapidly-evolving chirp towards the end of the merger is better captured by the short-duration tiles with a broader frequency range.

Using the Omega algorithm to visualize transient noise has also been a useful tool for unearthing unique characteristics of glitch populations and providing additional evidence to suggest potential origins of the noise. Several types of glitches appeared in both detectors during early Advanced LIGO engineering runs, with distinctive frequency shapes that characterize particular kinds of transients. For example, during several days in June 2015, strain Omicron trigger distributions in




Figure 2.4: Omega scan images of GW150914. *Top row*: Normalized energy of tiles of different durations. *Bottom row*: Whitened spectrograms from the Q -plane containing the highest energy tile (Q of 5.7)




Figure 2.5: Effects of glitches near 500 Hz and harmonics on the Livingston strain Omicron triggers. Each point shown here gives the time and central frequency individual Omicron trigger, with the color indicating the SNR. The hours marked in red at the bottom of the figure are times when the interferometer was not operating, so no triggers are shown during those periods. The analysis segments towards the end of the day show a large number of points near the harmonics of 500 Hz, distinct from the structure of the glitches earlier in the day. The large number of glitches at high frequencies during the segment between 5:30 and 10:30 were due to an unrelated instrumental problem.

the Livingston detector began to exhibit excessive glitches near 500 Hz and harmonics, as shown in Figure 2.5.

Examination of particular glitches using the Omega algorithm, such as the time shown in Figure 2.6, showed that these glitches were short duration, and they occurred at the fundamental and harmonics all at the same times. The main feature of the detectors' spectra in those frequency ranges are the *violin modes*, the resonances of the fibers of the suspension chains used for the optics. On the days when the glitches were occurring most frequently, the amplitude of the suspension violin modes were elevated compared to previous days. The combination of information from the Omicron distributions, Omega scan visualization, and knowledge of the detectors led to the hypothesis that the violin mode excitation was causing these glitches, and a greater effort was therefore implemented to develop better damping of the violin modes. After these improvements, the level

of excitation was reduced, and the glitches were not seen again, though the exact mechanism of the glitch was not fully understood.

Even though this occurred during an engineering run in the summer rather than the official observing period that began in September 2015, the search algorithms were running over the data in preparation for the observing run, allowing us to witness the effects of these glitches on the transient search backgrounds, in addition to observing the glitches in the individual detectors. Figure 2.7 shows a comparison of the cWB background distribution during the times when the violin modes were excited and later damped. While the violin mode amplitudes were elevated, the cWB background distribution contained populations of triggers near 500 Hz and 1500 Hz, which were reduced after the violin mode damping was used.

As another example of the usefulness of visualizing glitch morphology, Figure 2.8 shows a glitch in the gravitational wave strain channel. This glitch was caused by a radio frequency oscillator used for cavity control wandering in frequency near to another radio frequency line, and the two frequencies beating together. As the wandering frequency approaches the fixed line, the beating frequency becomes lower. As the wandering frequency moves farther away again, the beating frequency increases again. This decreasing and increasing of the beating frequency as the oscillator frequency changes creates the “V” shapes shown in Figure 2.8. The distinct pattern of this glitch in time and frequency allowed the detector characterization group to hypothesize the cause of the glitch, and identify some of the fixed frequency lines that were problematic for the oscillators.

During initial engineering runs of Advanced LIGO, these radio frequency beating glitches were quite frequent in the both detectors in the main sensitive frequency range used for gravitational wave searches. Since there are many oscillators used to control the various cavities as well as many lines of constant frequency, it is difficult to pinpoint all of the problematic frequencies and place all of the radio frequency lines far enough apart to prevent beating. However, the worst offenders were discovered, and RF oscillators were moved to carefully chosen frequencies so that during O1, almost all of the beating occurred above 2 kHz, at a higher frequency than the primary searches.




Figure 2.6: Omega scan of one of the short duration glitches near 500 Hz and harmonics which affected the instrumental background in early engineering runs. The frequencies of the glitches and appearance of the harmonics in the scan suggested a relationship with violin modes.




Figure 2.7: A comparison of the frequency distribution of the background events before and after violin mode damping was modified on June 6. The fraction of glitches near the violin modes and harmonics, most notably near 500 Hz and 1500 Hz, were drastically reduced when the amplitudes of the violin modes decreased, supporting the hypothesis that they were related to violin mode excitation. Note that the feature near 420 Hz is unrelated to this issue and persists after the violin mode damping takes effect.




Figure 2.8: Example glitch caused by the beating of different radio frequency controls.

2.3.3 Analyzing correlations between glitches

We can also use Omicron triggers or the results from other event trigger generators to statistically analyze coincidences between glitches in different channels, which is especially useful when trying to pinpoint the instrumental origin of a glitch population. Various methods can be used to find significant correlations between Omicron triggers in the gravitational wave strain (GW strain) channel and in key auxiliary channels. Triggers found to correlate with high statistical significance show that the strain triggers are likely glitches, as auxiliary channels should not detect real astrophysical signals. These times can then be used as vetoes to eliminate the glitches from the gravitational wave data. Various statistical algorithms have been developed to identify correlations between a set of the most important auxiliary channels and the gravitational wave strain data.[43, 38, 39]

Hierarchical veto, or *hveto*, was the primary method used for development of data quality vetoes in O1. *Hveto* is performed in several rounds to remove different populations of glitches coincident with various channels.[43] *Hveto* uses the rates of triggers present in each channel being

compared to determine the number of expected random coincidences within each time window. Then using a Poisson probability distribution, it determines the significance of the number of observed coincidences.

The number of expected coincidences, μ , between the triggers in the strain channel and auxiliary channel is estimated using the numbers of strain and auxiliary triggers, N_h and N_{aux} , the time window defining coincidence, T_{win} , and the total amount of time used for the analysis, T_{tot} :

$$\mu = \frac{N_h N_{aux} T_{win}}{T_{tot}} \quad (2.3)$$

If an either channel has a high trigger rate, the expected number of random coincidences naturally increases, as it does if the allowed coincidence window is increased.

The probability of finding a given number k of coincidences purely by chance is estimated by the Poisson distribution function,

$$P(\mu, k) = \frac{\mu^k e^{-\mu}}{k!} \quad (2.4)$$

And finally, a number called the *significance* is defined,

$$S = -\log_{10} \left(\sum_{k=n}^{\infty} P(\mu, k) \right) \quad (2.5)$$

For a given number of coincidences identified between the two trigger sets, the significance is the negative logarithm of the probability that one could find that number of coincidences (or more - hence the summation) purely by chance based on the trigger rates.

For every auxiliary channel used, the number of coincidences and significance are calculated for several time windows, and the most significant time window is selected. The channels are then ranked by their significance, and hveto then removes the GW strain triggers coincident with triggers in the most significant channel. The process is then repeated for more rounds until the method no longer finds any channels above a given significance threshold. The most significant

channel and potential veto segments from each round are recorded. The effects of hveto results applied to cWB in O1 are discussed in Chapter 3.

Hveto and other algorithms showing the correlations between transients different channels are important not only for applying as a veto to the searches, but also for providing useful insights about problems in the detector. The results sometimes show channels which might not be the most efficient veto, but suggest potential instrumental causes of glitches or avenues to pursue to develop efficient vetoes. For example, the one time that the radio frequency beating caused problems during O1, hveto was helpful in identifying what was causing the glitches.[44]

CHAPTER 3

DATA QUALITY IN O1: TRANSIENTS

After four years spent installing and commissioning the upgraded detectors, Advanced LIGO began its first observation run (O1) in September 2015 and collected data until January 19, 2016. Although the interferometers have still not yet reached their design sensitivity, they were about three times more sensitive during O1 than in the initial LIGO era.[45] Even with such a significant decrease in the average level of noise, the non-stationarity of the noise creates many short-duration glitches. These transient disturbances can be especially harmful to the background of searches for generic transient gravitational waves, such as Coherent WaveBurst (cWB). Without any restrictions on the waveform of potential signals, any glitch in one detector has the potential to create a cWB event, if there is coherence with a similar glitch in the other detector. In order to reduce the effects of instrumental noise on the gravitational wave search background, we must create vetoes, which are lists of segments of time with poor data quality which should be removed from the gravitational wave search. Ideally, the implementation of these vetoes to the searches would result in a background close to the Gaussian noise background shown in Chapter 2.

To create vetoes that eliminate the glitches most likely to affect the transient searches, the transients comprising the search background must be investigated. However, the search background is not ideal for performing all instrumental investigations because it does not fully represent all of the transient noise. The background events are only a subset of the glitches in each detector, which depends on the number of time shifts performed and whether or not there happen to be coherent glitches in both interferometers. It is important to find a way to understand and develop vetoes for the transient instrumental noise in each detector individually. These investigations therefore rely on single detector event trigger generators, such as the Omicron algorithm introduced in Chapter 2, to represent the transient instrumental noise likely to affect the gravitational wave search background. Since Omicron triggers are often used to represent the transient data quality, Section 3.1 presents a

comparison of the Omicron triggers of each detector and the cWB background events. The degree of correlation between the triggers demonstrates that Omicron is a valuable tool characterizing the glitches that produce the outliers of the search background.

Detector characterization techniques discussed in Chapter 2 were used to characterize the transient noise in both detectors and to develop data quality flags. Vetoes are applied to the transient searches in multiple categories, depending on the severity and duration of the instrumental problem. Section 3.2 presents the data quality flags chosen for the unmodeled searches in O1 and their effects on the instrumental noise in each detector, as represented by the Omicron trigger distribution.

3.1 Comparison between cWB and Omicron

As described in Chapter 2, Omicron triggers from auxiliary channels as well as the GW strain channel are used for data quality investigations, and the Omicron algorithm is part of a burst search (oLIB).[41] The key difference between single detector strain Omicron triggers and the coherent Waveburst events is that cWB requires coincident signals with coherence between the two detectors, whereas single detector Omicron triggers only depend on what is happening in one individual detector. Omicron is therefore well suited to investigate populations of glitches specific to one detector. It will not only show events that may be coherent with the other detector, but will also show features that do not contribute as clearly to the cWB background. Since Omicron is used for detector characterization and veto development, it is important to know how well correlated the Omicron triggers are with the cWB background. Here I present investigations that I have performed to investigate the relationship between single detector Omicron triggers and the cWB background to make sure that Omicron is able to identify the problematic populations of glitches most likely to produce significant background events. I have used Omicron and cWB background triggers from four weeks of the analysis of GW150914, from September 21, 2015 to October 19, 2015. These studies give us insight into whether there are subsets of Omicron triggers that are most likely to indicate potential problems for the burst searches.

The correlation between the two sets of triggers has been studied using a figure of merit called a Receiver Operating Characteristic (ROC) curve, which is a measure of the optimal detection of a signal in noise.[46] In this particular case, we are testing whether the cWB background signals are present in the Omicron single detector trigger sets. First we test the efficiency of a match between two sets of data: the percentage of cWB background triggers that are coincident with an Omicron trigger in a given time window. Then we estimate the false alarm rate, which is the number of coincidences that would occur by chance even in uncorrelated data. For this study, the false alarm rate is approximated using ten one-second time shifts in the data, finding coincidences each time, and averaging the false alarm rate from each time shift. This simulates a set of uncorrelated triggers, and is analagous to the background estimation process in transient searches. The efficiency and false alarm rate are calculated for varying time coincidence windows and plotted together such that each point along the curve represents a different time window. As the time window is increased, the probability of overlap between the trigger sets increases, in both the coincident analysis and the time shifted data. If the two data sets are strongly correlated, one would expect to see an efficiency much greater than the false alarm rate for some choice of coincidence window. If they are uncorrelated, the ROC curve would show approximately equal efficiency and false alarm rates.

The most important background events to eliminate are the highly significant triggers which make up the tail of outliers, as described in Chapter 2, so a minimum Rho of 8 was chosen for the cWB trigger set. Figure 3.1 shows the resulting correlation found between cWB background triggers and each of the two single detector Omicron trigger sets. Using the central time of both sets of triggers as comparison, 61.3% of the cWB triggers with a Rho over 8 are coincident within 0.1 seconds with a Livingston (*L1*) single detector Omicron trigger, and 65.2 % are coincident with a Hanford (*H1*) Omicron trigger, with false alarm rates of only 2.25% and 2.78% respectively. The efficiency is significantly greater than the false alarm rate for each set: 27.2 times greater for the L1 trigger set and 23.5 times greater for H1, showing that Omicron triggers are a good diagnostic set for investigating the transient GW search background.




Figure 3.1: Receiver Operator Characteristic (ROC) curve showing correlation between Livingston (L1) and Hanford (H1) single detector Omicron triggers and cWB background, before vetoes. The efficiency is shown as the percentage of cWB background triggers which are coincident with a single detector Omicron trigger. The asterisks mark the efficiency and false alarm for time windows of 0.1 and 0.5 seconds on either side of the cWB trigger time, as reference points.

It is also important to examine the properties of the cWB triggers which are not coincident with single detector triggers. Showing that the two data sets are strongly correlated would not prove that Omicron is a useful tool for characterizing the burst background if the fraction of triggers that were not coincident were highly significant for the cWB search. Fortunately, most of the remaining cWB events that are not coincident with single detector Omicron triggers from either detector have a relatively low significance, as shown in Figure 3.2.

Not only are most cWB background triggers coincident with a single detector Omicron trigger, but the parameters of the events recovered by each of the methods show similar features, such as frequency and amplitude, providing further indication that the two trigger populations are strongly related. Figure 3.3 shows the correlation between the Signal-to-Noise Ratio (SNR) of single detector Omicron triggers and cWB background events, illustrating the fact that the larger amplitude background events are coincident with larger SNR single detector triggers. Given the strong correlation between cWB background trigger populations and single detector strain Omicron




Figure 3.2: Properties of cWB background triggers not coincident with an Omicron trigger from either detector. The top panel shows the distribution in cWB coherent amplitude (ρ) of all the background events studied as well as the events not coincident with Omicron triggers with different time windows. The significance and frequency of the coincident and non-coincident triggers are shown below, using a time window of 0.1 seconds.

triggers, we can investigate Omicron triggers to gain a deeper understanding of the potential effects of instrumental artifacts on the burst searches.

Of course, the fact that most cWB events also have a corresponding Omicron trigger does not imply that any glitch that causes an Omicron trigger will also produce a cWB event. Rather, the cWB background represents a subset of the single detector events, only allowing those that happen to coincide coherently with a glitch in the other detector in one of the time slides. To illustrate this, Figure 3.4 shows ROC curves using the single detector Omicron triggers as the primary set, to demonstrate the fraction of Omicron triggers that are represented by the cWB background trigger set from the analysis period from September 12 to October 20, 2015. For this background set and taking Omicron triggers with a minimum SNR of 8, 19% of the H1 Omicron triggers and 25% of the L1 Omicron triggers are coincident with a cWB background trigger. The coincidence rate is much less for lower amplitude triggers, and the total number of cWB background events is also highly dependent on the number of time slides performed when creating the cWB background.

Therefore, a veto that does a fantastic job with a particular population of Omicron triggers in one detector may not have as significant an impact on the cWB background, and it is important to test vetoes on the cWB background as well as on Omicron to take both trigger sets into account before deciding whether or not to implement a flag. Some instrumental issues, however, are serious enough that we have decided to excise the untrustworthy data from the burst searches, regardless of whether the issue is shown to cause an excess of cWB triggers. The following sections will show various vetoes used for the burst searches during O1 and the impact of the data quality flags on the single detector Omicron trigger populations as well as the cWB background.

3.2 O1 Data Quality

The quality of the data is verified using various means to detect instrumental and environmental noise sources potentially polluting the interferometric data. The methods used are online checks of the state of the instrument, statistical algorithms for finding instrumental or environmental correlations with the strain signal, and in-depth investigations to uncover noise sources. Times of dubious data quality are then vetoed – removed from the analysis to create a cleaner set of data.




Figure 3.3: Comparison of cWB properties with single detector Omicron triggers coincident within 0.1 seconds. Above, the cWB detection statistic ρ for each background event that coincides with a single interferometer Omicron trigger from either detector is plotted against the corresponding Omicron SNR. The more significant cWB background events tend to coincide with higher SNR single detector Omicron triggers. Below, the frequencies reported by cWB and Omicron are compared, with dashed lines included as a reference indicating a relative frequency difference of 50%. Although there is a wide distribution of frequency differences and several outliers, the general trend shows that coincident Omicron triggers tend to be in the same frequency range as their cWB counterparts.




Figure 3.4: ROCs of cWB and Omicron triggers, showing the fraction of Omicron triggers coincident with cWB background events. The asterisks mark the efficiency and false alarm rate of coincidences using a 0.1 second window on either side of the central time of the Omicron trigger.

During operation of the interferometer, online systems ensure that all subsystems of the detector are in their nominal configuration, including the correct laser power level, good alignment of all cavities, and full seismic isolation of all optics. Each component of the interferometer must be functioning before the data is even considered ready for analysis. Automated monitors also check for well understood issues that are known to be harmful to the data, such as the saturation of the interferometers control signals.

In addition to the strain signal, the detectors record thousands of other channels, measuring environmental and instrumental signals.[28] Several algorithms are used to determine statistical coupling between events in these auxiliary channels and the gravitational wave channel, to show whether the strain signal is significantly correlated with channels unlikely to be affected by a real gravitational wave. The vetoes produced by these methods are short duration, and can be used to veto fractions of a second per trigger, rather than cutting out longer segments of time.

Data quality information is also recorded by experts working at the site of the detectors, who document local disturbances that may corrupt the GW strain signal. Furthermore, times of poor data quality can be flagged after an in-depth investigation reveals a new noise source, possibly days or weeks after collection of the data.

Vetoed are classified under three categories that are applied to the search differently, depending on the severity of the data quality issue. Category 1 vetoes are implemented before any analysis of the data is performed because there is certainty of a problem with the detector or calibration of the data. Category 2 includes vetoes that have a well known instrumental coupling into the data, such as a photodiode saturation or the time preceding a lock loss. These times are removed after data whitening, but before any triggers are produced. Finally, Category 3 vetoes are used to remove candidates after trigger production. These vetoes are generally short duration (on the order of a second), based on statistical coupling with auxiliary channel triggers.

All vetoes are evaluated on the background of the burst search to determine which to apply before final analysis. The determination of which data quality flags to apply must not be biased by whether or not they would veto specific gravitational-wave candidates, which is why all the testing is performed on background triggers or on single detector burst triggers. For every veto applied, there is a trade-off that must be considered. It is vital to create vetoes that will remove any data that would not be trusted if the algorithm were to find a gravitational-wave signal at that time. On the other hand, it is possible to be over-zealous in removing any potential glitches, and thereby throw away analysis time that could actually contain some viable signals. Therefore the main metrics used for this decision are efficiency and dead-time: efficiency is the percentage of background triggers removed by the veto, and dead-time is the percentage of analysis time lost. The ideal vetoes should have a high efficiency with a low dead-time, though it is not always clear how to choose acceptable ratios between these factors.

The following sections describe the vetoes in each of these categories that were used for the O1 analysis, some rationale as to why they were chosen, and their impact on the strain Omicron triggers and cWB background. The cWB background analysis was performed separately for two parts of the observing period: September 12 to October 20, 2015, which was used for the analysis of GW150914, and the remainder of the observing run, October 20, 2015 to January 19, 2016. On December 26, 2015, the gravitational wave signal GW151226 was detected with high significance by the searches for compact binary coalescences, but not by the burst searches, due to the lower masses of the black

holes.[47] Since the latter part of the observing period did not contain gravitational wave signals significantly detected by the cWB algorithm, fewer time shifts were performed over this period, and the background analysis was mainly used to verify consistency with previous results.[48] Therefore, the information about the data quality flags presented in the following sections will describe the impact of the vetoes over the two different time periods.

3.2.1 Category 1

The times of worst data quality issues were vetoed as Category 1, and these times were removed from the data before any processing in cWB. The only Category 1 flag used for both detectors was to remove periods of excess non-stationarity as optical cavities begin to lose resonance. The moment of loss of resonance in the cavities always automatically ends an analysis segment, but a few times during the observing run, some additional time near the end of the analysis segment was corrupted and needed to be removed. All other Category 1 flags were specifically created to respond to a severe instrumental issue particular to a certain time for each interferometer.

- Livingston flags

The data from both detectors is calibrated using photon calibrators, *Pcals*, which use auxiliary lasers directed at the test masses to induce motion at specific frequencies because of the radiation pressure of the photons from the laser. By comparing the known amount of force acting on the mirrors and the resulting fluctuations at the readout at those frequencies, the Pcal system can perform the strain calibration. [49] In Livingston, there was a period of three days during O1 when the photon calibrator laser occasionally fluctuated in power beyond its nominal range, causing the laser to push on the test masses with excessive force, creating glitches in the data.

Another Category 1 veto for Livingston was defined to cut out a period of time on December 25 when the voltage control oscillator (VCO) for the laser was in not in a nominal frequency range, causing it to beat against other radio frequency lines, creating glitches like the RF beating described in Chapter 2. During most of O1, the VCO offset was set so that any RF beating glitches occurred at frequencies above the main search range of 30 to 1000 Hz. On December 25, however, the bad VCO offset caused a high rate of glitches above 200 Hz for an analysis segment of 11 hours.

[44] This anomalous glitch rate would cause a transient background substantially different from the rest of the analysis period, so all 11 hours were removed from the data.

During the last month of O1, the Livingston detector began to show excess non-stationarity at lower frequencies (mostly below 30 Hz, but affecting the data up to 100 Hz at times), which had a negative effect on the sensitivity of the instrument and at times caused problems severe enough to merit a Category 1 flag. The low frequency noise was investigated by Livingston instrumental scientists after the observing run was completed, and traced to a failure in a temperature sensor on one of the seismic platforms. [50, 51]

One veto was added to remove four hours when the excess noise in the low frequency range was too severe and nonstationary to keep the data. Another flag was created for two different periods when the strain noise spectrum below 100 Hz suddenly changed. Since all transient analyses require an estimation of the background power spectral density for whitening the data and searching for periods of excess power, it is important that the spectrum remain stationary during an analysis segment. Removing these sudden transitions allows each segment to have a spectrum that is more stable.

During the analysis period used for estimating the background for GW150914 (September 12, 2015 to October 20, 2015), the Category 1 flags only removed 1.1×10^{-3} percent of the operational time of the interferometer. During the rest of O1 (October 20, 2015 to January 19, 2016), however, the low frequency non-stationarity in the detector had a higher impact on the available analysis time, and the Category 1 flags removed 1.4% of the data. Figure 3.5 shows the impact of all of the Category 1 flags applied to the single detector Omicron triggers from Livingston from O1. Although most of the transient glitches falling within the Category 1 segments are low SNR, the flags do remove a few notable populations of loud triggers.

- Hanford flags

At the Hanford site, several Category 1 flags were specifically created for various hardware problems. One Hanford flag removed several hours on September 15, when there was a hardware failure in a controls system, causing numerous additional spectral lines in the data and severe




Figure 3.5: The impact of Category 1 data quality vetoes on single interferometer Omicron triggers from Livingston. The frequency values reported by Omicron are discretized values based on the resolution of the most significant Q -tile of a particular signal. For each Q value, the frequencies used are logarithmically spaced. The distinctive lines of high SNR triggers shown here are primarily due to populations of low Q glitches. The spacing between these lines is the frequency resolution of the Q -plane that best matches these glitches. Although many of the highest SNR transients are not within the Category 1 segments, it is important to apply these flags because the interferometer is in a demonstrably poor state, and not producing trustworthy data.




Figure 3.6: Spectrum comparison between a typical time during O1 and the hours during the Beckhoff system error.

glitching. The clear excess in noise and a prominent 2 Hz comb in the data can be seen by comparing the spectrum of that time to a typical time from O1, shown in Figure 3.6.

Another flag was produced to remove a small amount of time from one day during the run when high excitation of the second harmonic of the violin modes caused severe glitching, similar to the effects described in Chapter 2.

Finally, a recurring instrumental problem throughout the analyzed time was excessive noise coupling into the strain channel from an electro-optic modulator, which is used to modulate the input beam with radio frequency sidebands at 45 MHz, used for locking and controlling the interferometer. Periods of severe glitching were observed to correlate with glitches in the modulator's drive signal. The exact source of this noise remained unsolved throughout the observing run, but times showing this behavior were removed from the data. The flag was created by identifying a threshold on the amplitude of a signal from a monitor of the modulator's drive signal, above which the interferometer began to show the severe glitching behavior.[28] After the threshold had been established, times when the monitor's signal exceeded that level were then removed.




Figure 3.7: The lefthand plot shows the impact of Category 1 data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford for the first analysis period of O1 (September 12 to October 20). This is the period of analysis for GW150914. The righthand plot shows the distribution from September to January. Although the vetoes do not remove high SNR transients, it is important to apply these flags because the interferometer is in a demonstrably poor state, and not producing trustworthy data.

Since Category 1 vetoes are applied before the full transient analyses, the cWB background results do not include these times at all. However, the initial background analysis was performed online without certain flags applied, which allows us to analyze the flags' performance not only on single detector triggers, but also on the analysis. During the development of the veto for Hanford 45 MHz radio frequency modulation glitches, the flag was tested over a small set of cWB background triggers as well as the single detector Omicron triggers, in order to find the best threshold to remove the corrupted data. Over a four-day span when these glitches were particularly prominent, the flag removed 2.94% of the analyzable cWB time, but 15.95% of the cWB background triggers, including six out of the eight triggers with coherent amplitude greater than 8. With a ratio of efficiency to deadtime of over 25, the flag was shown to be highly significant for the cWB background. Although no testing of the flag was needed to prove that there were significant problems with the instrument, the veto testing results showed that the tuning of the flag was done accurately enough to remove the worst instrumental glitches while preserving as much analyzable data as possible.

Because of the major problems with the Beckhoff system and the 45 MHz modulation during the first part of the observing period, the Category 1 flags remove 4.4% of the analysis time from

September 12 to October 20, substantially more deadtime than the Livingston vetoes. However, only 1.2% of the analysis time needed to be removed from the remainder of the run. The impact of all of the Category 1 flags on the Hanford strain Omicron triggers can be seen in Figure 3.7. Just as with the Livingston flags, the Category 1 vetoes do not remove most of the loudest times, but all of the times chosen were periods in which the instruments were known to be in a bad state, which would cast doubt on results from any analysis of the data. Therefore all transient analyses, whether the unmodeled or modeled searches, applied the same Category 1 vetoes.

3.2.2 Category 2

Shorter duration problems in the instrument are typically vetoed in Category 2. These glitches tend to be on the order of a few seconds and should have a negligible effect on the initial processing and whitening of the data, so they are removed from the analysis after trigger production. Many of the data quality vetoes in this category are applied to the burst analyses as well as the CBC searches, although some different choices were made based on the testing of the flags on the relevant background triggers.

The most important Category 2 flag for both detectors correspond with times when the control signals sent to actuate on the test masses saturate the digital-to-analog converter (DAC). These are often huge glitches in the strain data that occurred several times a day during O1, corrupting the surrounding data for several seconds each time, and they are by far the loudest transients in both detectors. The glitches in the strain channel are not necessarily only the result of the saturations, but could also be the cause: as the interferometer feedback loop responds to a large glitch in the gravitational wave strain channel, it requires a large actuation on the test mass, which exceeds the amount of force that can actually be applied.

The exact origin of this type of glitch remains unknown, but it is definitely not an astrophysical signal, due to the lack of coincidences between both detectors. Regardless of whether this DAC overflow is the original cause of the transient, any such saturation in the control sent to actuate on the test masses would cause a glitch in the data, and these times must be removed from the analysis. The largest of these glitches were found to cause problems for the burst searches even




Figure 3.8: Omega plots showing H1 DAC saturations. The time shown on the left is a DAC overflow with an associated Omicron trigger of SNR 85. To the right, a much louder saturation is shown, corresponding to an Omicron trigger with SNR of 9990. Note the difference in time scales. With the larger saturation, more time around the glitch is affected and needs to be removed.

outside of the segment where the actual saturation time was recorded, so an additional padding was added to saturations corresponding with an Omicron trigger of SNR greater than 200. Figure 3.8 shows two glitches of this type with different SNRs. While the lower SNR glitch is shorter than one second in duration, the louder glitch affects the surrounding data for several seconds.

The additional padding is especially important considering that the Category 2 vetoes are applied to the cWB analysis after triggers have already been produced. The trigger list that the vetoes then act on typically uses only the central time of the transient calculated by cWB. If that time falls within a veto segment, the trigger is vetoed. But it is possible that the nonstationary data near a saturation could be part of a cWB background event whose central time technically falls outside the saturation time. Such triggers have the potential of being influenced by untrustworthy data, so they should be removed.

Figure 3.9 shows the effects of these two flags on the cWB background triggers from the analysis of GW150914, from September 12 to October 20. Although the additional padding around larger DAC overflows does not remove many background events that were not already captured by the original veto, the flag is important for taking out the non-stationary data surrounding these large glitches. Although only the Hanford vetoes are shown here, the flags from each detector veto most




Figure 3.9: The impact of the ETMY saturation flags and ETMY saturation with Omicron SNR greater than 200 on cWB background triggers from September 12 to October 20.

of the same background events, because those events tend to be caused from the coincidence of two saturation glitches in the time-shifted background.

A related issue is the saturations of the photodiodes of the output mode cleaner, a serious but rare problem only occurring a few times throughout the run. The flag was created because any time the output mode cleaner photodiodes are saturated, the strain signal is certainly corrupted, and these few times must be removed from the data.

Finally, towards the end of the observing period, a new kind of glitch was observed at the Hanford observatory that affected all transient searches. The noise showed the same characteristics as the 45 MHz modulation glitches, but was not witnessed by the same channels. Rather, it was found that these glitches correlated with increased fluctuations in the 36 MHz modulation.

- Summary of Category 2 vetoes

The Category 2 vetoes had the largest effect on both the Omicron triggers and the cWB background, mostly as a result of the DAC overflows. Figure 3.10 shows the SNR distribution of the Omicron triggers after Category 1, and then after the application of each flag in Category 2.




Figure 3.10: The impact of Category 2 data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford.

- Additional flags not used for the burst searches

Two flags for the Hanford interferometer were used as Category 2 vetoes in the CBC search but were not applied to the burst searches, after testing them on the cWB background showed that they did not significantly reduce the background.

One of the unused CBC Category 2 vetoes flagged times when the seismic noise between 10 and 30 Hz was elevated, which correlated with glitches found in the GW strain data. Figure 3.11 shows one example of a time period in which the seismic noise increased above the threshold for this flag, and the corresponding strain Omicron triggers. This flag does show some interesting nonlinear coupling from the low frequency range of the seismic motion to strain triggers above 100 Hz, which should be further investigated. However, this flag was not highly efficient on single detector triggers, vetoing 6.5% of all the Omicron triggers with a deadtime of 1.2%. Furthermore, the Omicron triggers vetoed are quite low in SNR, so the cWB events potentially vetoed would also be low significance.

The other flag was a specially tuned flag for less severe 45 MHz sideband fluctuations than the veto used in Category 1. While the original Category 1 flag used band-limited root-mean-square values of the 45 MHz sideband monitor using 60 second averages to detect large fluctuations, the second flag used the variation of the same signal over every second to capture shorter-duration fluctuations. The additional 45 MHz sideband flag demonstrated a significant coupling to the




Figure 3.11: Elevated seismic motion corresponding to increased glitch rate in GW strain. These ten minutes of data show two periods where the band-limited RMS of one of the ground seismometers increases above the 600 nm/s threshold (top plot). During these times, the trigger rate in GW strain increases.




Figure 3.12: These flags were used for the CBC search but not implemented for the generic searches. single detector strain triggers. During three weeks of O1 used for testing, this flag vetoed 8.5% of the Omicron triggers (without applying any other vetoes), with a deadtime of only 0.03% of the time the interferometer was taking data. Although this flag certainly demonstrates a strong instrumental coupling into the gravitational wave strain channel in H1, its impact on the cWB background would not be as significant. During the same time period, the flag would only have vetoed 0.5% of the cWB background triggers with ρ greater than 8, none of which were the loudest outliers in the distribution.

The impact of both flags on three weeks of H1 Omicron triggers is shown in Figure 3.12. The set of triggers shown here includes Omicron triggers with an SNR over 8 and a frequency under 1000 Hz, to be most relevant to the low frequency burst search. All of the triggers during the seismic flag segments have an SNR below 20, and most are below 10. Most of the triggers vetoed by the additional RF45 flag have an SNR below 30, but there are several with higher SNRs, even above 100. A case could be made for including the flag based on its significant impact on Omicron and the demonstrated correlations between cWB and Omicron described in Section 3.1, but it was eventually decided to leave the flag out because of its negligible impact on the cWB background.

3.2.3 Category 3

For the final category of vetoes applied to the data, we used the Hierarchical Veto (hveto) method for finding statistically significant coupling between glitches in auxiliary channels with the strain data. Channels found to affect the gravitational wave data were used to create vetoes.

These veto segments were created by running hveto over the single detector strain Omicron triggers and a set of the most important auxiliary channels as selected by the detector characterization group. Before choosing the channels to use, hardware injections are performed to induce a differential arm motion in the interferometer, simulating the effect of a gravitational wave. If the auxiliary channels witness the injected waveform, they are considered *unsafe* to use, as they could potentially eliminate real gravitational wave signals. After running hveto, vetoes were chosen using the most significant channels from each round that had a significance above 20 and an efficiency to deadtime ratio of over 10, when applied to strain triggers with an SNR above 6.

As the instrumental coupling can change over the course of several months, the hveto analysis was performed over ten different segments of about two weeks instead of over the whole observing period, in order to capture the instrumental noise most relevant to each time period. Each of these analyses identified multiple auxiliary channels statistically correlated above the chosen thresholds.

The dozens of channels used will not be explored in detail, but as an example I show here the most significant channel from the first segment, September 12 to 26, 2015. With an hveto significance of over 500 (as defined in [43] and Chapter 2), a seismometer from one of the end stations was able to veto 5.87% of the GW strain Omicron trigger set. Veto segments of 0.2 seconds were created from each of the auxiliary channel triggers that had an SNR above 20 (0.1 seconds on either side of the triggers' central times).

Using the method introduced in Section 3.1, an ROC curve was created to demonstrate the correlation between this seismometer and the strain channel, shown in Figure 3.13. Although the hveto significance method and the ROC estimation of false alarm rate are not identical, both methods show that the two channels are well correlated. Comparing the three different SNR thresholds shows how the SNR threshold of 20 is the best. If the threshold were decreased to




Figure 3.13: ROC Curve showing the efficiency of the most significant auxiliary channel in the hveto analysis from September 12 to 26, at three different SNR thresholds placed on the auxiliary triggers. The asterisks show the results for the selected coincidence window of 0.1 seconds on either side of the auxiliary triggers' central time.

10, the efficiency would increase slightly while significantly increasing the false alarm rate and deadtime. Increasing the threshold to 30 would decrease the false alarm rate slightly, but at the cost of reducing the efficiency of the veto. The strain glitches vetoed by this channel were mostly low in amplitude, but other rounds were able to remove some of the highest SNR single detector triggers remaining after Category 2 vetoes.

3.2.4 Summary of veto impact on Omicron triggers

The impact of each category of vetoes from all of O1 on the SNR of single detector triggers is shown in Figure 3.14, and the specific signal-to-noise ratio and frequency of the vetoed triggers can be seen in Figure 3.15. While Category 1 vetoes are necessary to eliminate untrustworthy data from the interferometer, their impact is primarily on transients with SNR less than 100. The application of Category 2 to the data significantly reduces the rate of higher amplitude Omicron triggers in both detectors, mostly due to the ETMY saturation flags. Implementing veto segments from hveto removes a smaller number of additional transients including a few of the outliers in both detectors, with most of its impact on low frequency triggers.

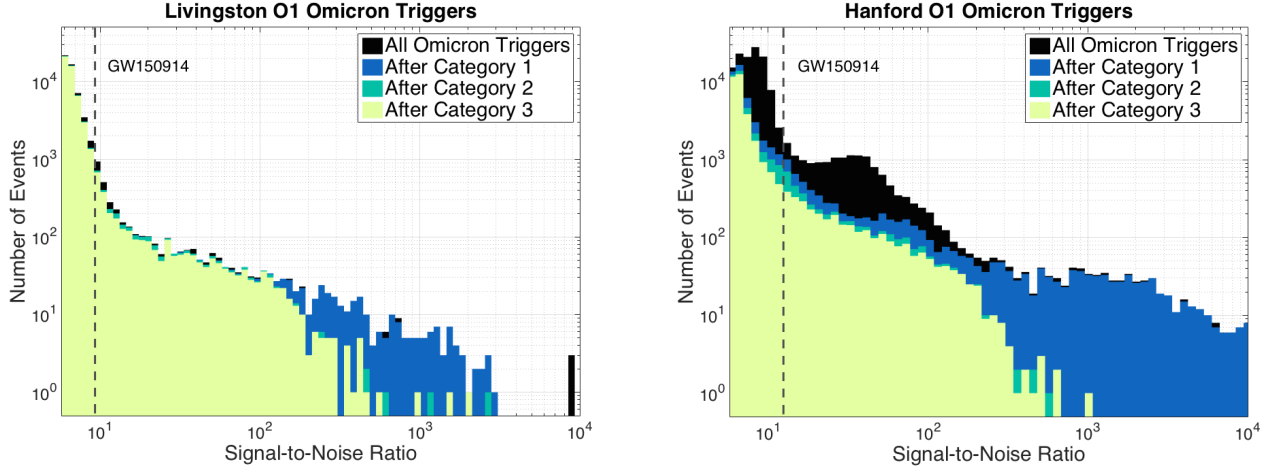


Figure 3.14: The impact of data quality vetoes on single interferometer Omicron triggers from Livingston and Hanford during O1, shown by the distribution of transients in Signal-to-Noise Ratio (SNR). The dashed lines mark the SNR of GW150914 as seen by the Omicron algorithm in each detector. The existence of many transients with higher SNRs in both detectors emphasizes the need for a coherent search to find correlated signals in both detectors at the same time.




Figure 3.15: The impact of data quality vetoes from each category on single interferometer Omicron triggers from Livingston and Hanford, showing the SNR and frequency of the triggers vetoed.

The data quality has a much more visible impact on the transients in the Hanford detector than in Livingston, since there were several severe instrumental problems in that interferometer with clear coupling to the strain data, and a much greater percentage of the time was removed in Category 1 and 2.

Table 3.1 gives a summary of the fractional deadtime and efficiency of each category, when applied to Omicron single detector triggers with an SNR greater than 10. Although the hveto segments were only used if the efficiency to deadtime ratio was over 10, that ratio was calculated using Omicron triggers with an SNR threshold of 6. When the ratio was calculated using a threshold of 10, the Livingston hveto segments are much less efficient. However, for the other categories in Livingston and for all categories in Hanford, the ratio of efficiency to deadtime is large enough to show that the vetoes are effective at removing known instrumental problems without unnecessarily removing extra time from the analysis.

Table 3.1: Summary of vetoes applied to the single detector Omicron triggers with SNR greater than 10. Deadtime and efficiency are given as a percentage of the remaining time and triggers left after the previous category.

Category	Livingston			Hanford		
	Efficiency (%)	Deadtime (%)	$\frac{\text{Efficiency}}{\text{Deadtime}}$	Efficiency (%)	Deadtime (%)	$\frac{\text{Efficiency}}{\text{Deadtime}}$
1	8.38	1.01	8.30	64.4	2.30	28.0
2	9.07	4.12×10^{-2}	220.1	37.0	6.63×10^{-1}	55.8
3	8.51	2.29	3.72	25.5	1.78	14.3

CHAPTER 4

SEARCH RESULTS IN O1: GRAVITATIONAL WAVES

As shown in Chapter 3, even though the data quality flags remove a large proportion of high SNR events in each detector, the remaining distribution of single detector triggers is still far from Gaussian. Unfortunately, the same is true for the cWB background after all categories of vetoes. The large amplitude outliers limit the potential of the cWB search to confidently detect an arbitrary signal. Section 4.1 presents an investigation into the remaining high amplitude glitches, in an attempt to identify instrumental correlations.

The cWB algorithm running online during O1 was the first search to identify the gravitational wave signal GW150914, within three minutes of the time the signal arrived at the detectors.[52] Due to the persistence of a non-Gaussian tail of outliers in the background, the cWB search background was divided into different bins based on event morphology, in order to separate out most of the problematic glitches from the rest of the background, which was significantly cleaner. The different backgrounds, explained in Section 4.1.2, allowed the cWB search to detect GW150914 with high confidence when compared against the appropriate background set. To rule out instrumental or environmental origins of the signal, extensive data quality checks were performed, described in Section 4.2. The final results of the cWB analysis of GW150914 are shown in Section 4.2.

4.1 Remaining glitches

The remaining high SNR outliers are due to one particular class of glitches: short duration (10-20 ms) low frequency transients called blip glitches. In the time-frequency representations, they have a symmetrical tear-drop shape, as shown in Figure 4.1. These transients appear in both detectors (though never simultaneously yet, ruling out an astrophysical origin) at a rate of about once per hour.

The similarity of the blip glitches in time and frequency between detectors and between different times creates a problem for the burst search backgrounds. Figure 4.2 illustrates how alike the blip




Figure 4.1: A typical blip glitch time-frequency morphology. The left-hand plot [28] shows the typical teardrop shape in time and frequency, and the on the right is the whitened time series of the same glitch.

glitches to one another are by showing several of the glitches from each detector that occurred at different times during O1. Although the blip glitches were not coincident in real time, blips from the two detectors do coincide in the time-shifted backgrounds, causing a tail of highly significant events.

4.1.1 Blip glitch investigations

Multiple investigations have been pursued to discover the instrumental source of these glitches; one particular study was to use the hierarchical veto (hveto) algorithm [43] to compare glitch times with auxiliary channel Omicron triggers to find any significant correlations. For this study, a list of blip glitches was created from the cWB background events identified to have similar properties to the blips, during the analysis period from September 12 to October 20. The list was further narrowed down using frequency and signal-to-noise ratio as estimated by Omicron. The resulting 476 times were used as the input to hveto to find coincidences with Omicron triggers from auxiliary channels. There were several channels with high statistical significance, including several accelerometers and channels relating to alignment in the corner station. Table 4.1 shows a few of these channels that were found to be correlated with this list of times after an SNR threshold of 8.5 was applied.




Figure 4.2: Several blips from each detector, at different *non-coincident* times during O1. The similarity of the blip glitches creates a problematic tail in the distribution of cWB background triggers.

Table 4.1: A few of the most significant channels correlating with blip times from cWB. The columns show the hveto significance (explained in [43], the number of coincidences with blip times, the number of expected random coincidences based on the number of auxiliary triggers, and the ratio of expected coincidences to actual coincidences. Although the significance of these correlations is high, only about 20 of the glitch times could be vetoed by these channels.

Subsystem	Significance	Ncoinc	Nexpected	Ncoinc/Nexpected
Arm length stabilization	63.8	19	3.48×10^{-3}	5459
Alignment sensing and control	59.5	21	1.27×10^{-2}	1655
Pre-Mode Cleaner	48.8	21	4.11×10^{-2}	511
Accelerometer	40.6	20	7.75×10^{-2}	258
Pre-Mode Cleaner	38.5	23	0.20	115
Accelerometer	37.5	16	3.1×10^{-2}	517

Although there were a number of channels found to be possibly correlated with the list of cWB background triggers, further investigation showed that these would not be effective vetoes for the blip glitches. The triggers in the significant channels all coincided with the same subset of blip times, with the possibility of vetoing only about 20 of the 476 times, even if segments from multiple channels were combined. Additionally, the glitches coinciding with these channels actually have a time-frequency evolution distinct from the typical blip shape. This points to these being a separate glitch class, not the main one plaguing the search. Furthermore, the events potentially vetoed by these channels were relatively low significance events, and the loudest of them would already be vetoed after the application of the regular hveto segments that had been produced for all the background. Figure 4.3 shows the Omicron triggers that would be vetoed by using the most significant channel.

Similar studies have been performed by others to search for correlations between the blip glitches and transients in auxiliary channels, but no instrumental cause has been found yet. The auxiliary triggers used in this study are from the set of channels deemed most important by experts in detector characterization, but future investigations may need to expand the set of auxiliary channels searched. There could also be clues to the instrumental origin of the glitches from other

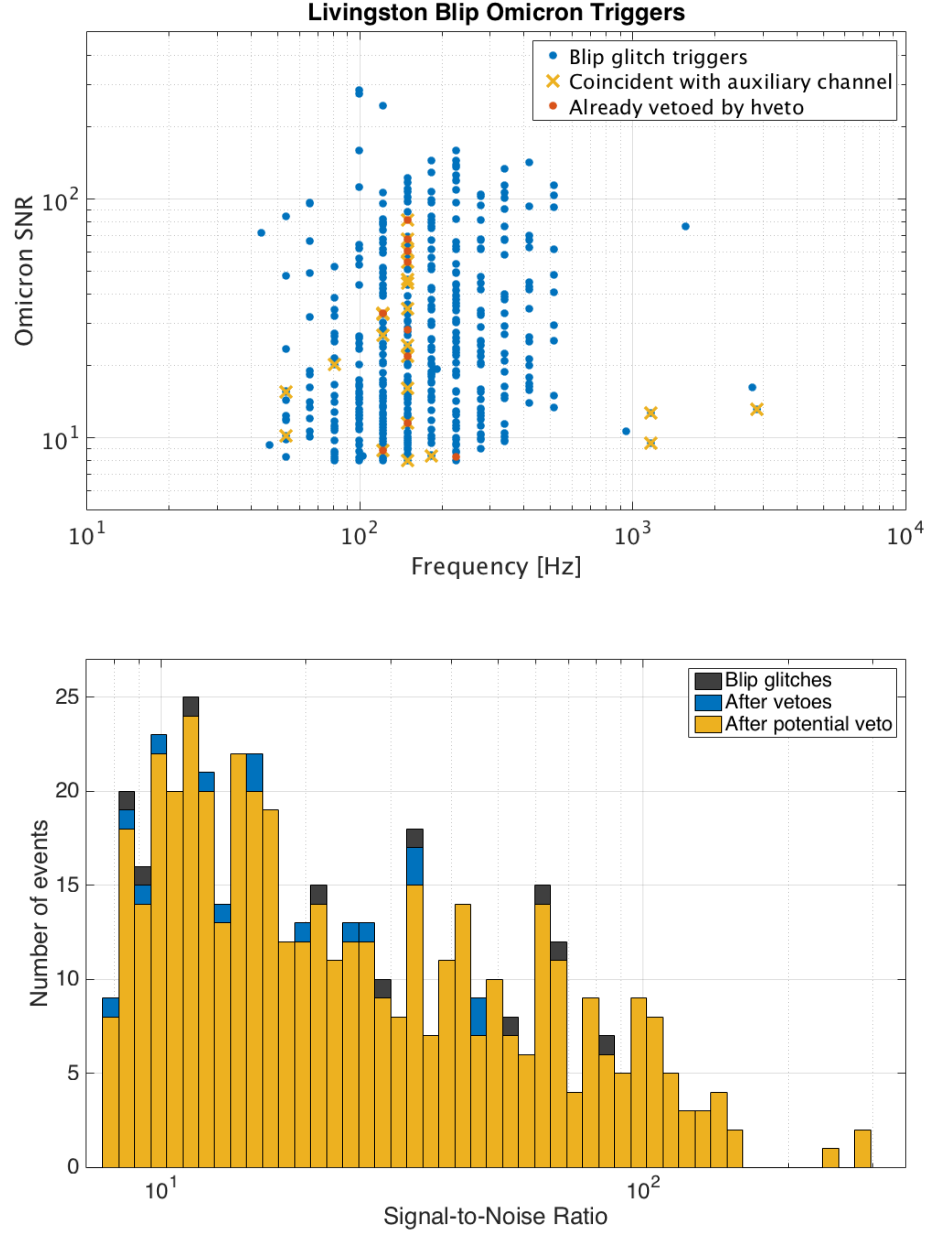


Figure 4.3: The potential effects of applying a veto based on the auxiliary channel with the highest significance in correlations to blip glitches. The top plot shows the Omicron SNR and frequency of all blip glitches in the study, the glitches that would be vetoed using the most significant channel found to correlate with the set of blip glitches, and those that would already be vetoed by the implementation of hveto segments. Below, the SNR distribution of these triggers is shown. Although the significance of this channel was higher than others, the number of additional glitches that would be vetoed is very low, and most of them were low in SNR.

avenues besides auxiliary glitches. Analysis of the rates of the glitches over time could uncover interesting patterns and correlations with instrumental behavior that might not show up as short-duration glitches in the auxiliary channels. For example, the glitch rate significantly increased for a period of a couple of weeks during O1, and a possible correlation with the relative humidity in the buildings was found.[53]

One possible limitation of the studies conducted so far is the problem of glitch classification, and identifying a complete set of the blip glitches. Several techniques are being developed to provide a better method of sorting glitches into different classes, automating the process through machine learning and employing citizen science to utilize the efforts of interested members of the public.[54, 55] In the future, the lists of blip glitches collected by these methods may lead to a more complete set of the glitches and better statistics to analyze.

4.1.2 cWB search bins

Because of the persistence of blip glitches even after all vetoes have been applied, the coherent WaveBurst search implemented an algorithm to separate these events into a separate search bin by using the time and frequency properties of the waveform. This creates one search background dominated by events with waveforms similar to the blips, and another “clean” set. Therefore any gravitational wave candidate with a signal unlike the blip glitches could be compared to a blip-free background.

The division into search bins is based on the quality factor, Q , of the reconstructed waveforms, which is the ratio of the frequency to the bandwidth. The blip glitches tend to have a low Q -factor, so placing a Q threshold on the signals efficiently removed most of them from the clean data set. Figure 4.5 shows the results of applying this algorithm to create a cleaner search background. Although the majority of the outliers are removed from the constrained set, there remain a few blip glitches that managed to evade the Q -based cut.

Additional signal-based cuts were performed in this manner to remove narrow-band events due to narrow non-stationary lines in the data creating populations of cWB background events around wandering lines.[52] For example, Figure 4.4 shows the frequencies of a population of cWB




Figure 4.4: A nonstationary line in the strain data creates excess cWB background triggers around its frequencies as the feature moves between 520-530 Hz. Each point plotted over the spectrogram is a cWB background trigger. There are segments without any background events because the triggers were only produced for times when both detectors were operational. Populations of triggers like these prompted the need to create a signal-based cut on the data to separate out such narrow-band events.

background triggers superimposed on a strain spectrogram from an engineering run in June 2015. At the time, a nonstationary line between 520 and 530 Hz was persistent in the data, and the cWB background events could be seen to correlate with the movement of the frequency of the line. It is difficult to create appropriate veto segments for events such as these, since they are not caused by discrete events in time but rather a constant nonstationarity that makes the cWB algorithm more likely to find the variations in power around these frequencies. So instead of applying vetoes, the algorithm included a “line veto” to separate events that small bandwidth, calculated by whether 80% of the power in the signal is within a 5 Hz band.

Figure 4.5 shows the result of applying these signal-based cuts to the unmodeled background, before any veto application. While the unmodeled background set contains a large population of high amplitude triggers, the constrained background distribution is much closer to the ideal Gaussian background, even without the data quality flags.




Figure 4.5: Coherent WaveBurst background before and after signal-based cuts to separate blips into a separate bin, and the constrained background after the application of data quality vetoes. The dashed line shows the coherent amplitude of GW150914.

Figure 4.6 shows the impact of vetoes on these two search bins, which removes some of the outliers in both distributions. While the background after signal constraints have been applied is significantly cleaner, this method of dealing with data quality issues is not ideal. The usage of multiple search bins requires the inclusion of a trials factor, which reduces final reported significance of an event. Additionally, the separation based on signal morphology means that the constrained search is not truly able to detect any arbitrary gravitational waveform. However, until the problem of the blip glitches is solved, the searches for generic transients are forced to divide up their parameter space to enable detections.

4.2 O1 search results

On September 14, 2015 the cWB search running in low latency detected a gravitational wave event, GW150914, finding the signal in the data only three minutes after it reached the detectors.[52] With a cWB ρ of 14.1 and time-frequency spectrograms showing a clear chirp in both detectors, the event was immediately seen to be quite promising. However, as the first direct




Figure 4.6: The effects of vetoes on the ρ distribution before and after signal constraints, during the analysis period of GW150914, September 12 to October 20, 2015. The lefthand plot shows the distribution of all background events, and the righthand plot shows only the events left after applying the Q threshold and line veto.

detection of a gravitational wave event, great care was taken to fully understand the data quality around the period and to ensure that the false alarm rate of the event could be properly estimated.

4.2.1 Data quality checks

The calibration of the data from the signal readout to gravitational-wave strain was undertaken carefully, and found to have an uncertainty of less than 10% in magnitude in the relevant frequency range, 20 Hz to 1 kHz.[56] Although the uncertainty in calibration leads to some uncertainty in recovery of the parameters of the system, it does not diminish the capability of the detectors to make a confident detection. Since the determination of the sky localization depends on a precise measurement of the arrival time difference between the sites, the timing systems used in each detector were also thoroughly investigated, and found to be working properly.

No data quality flags from any category were active at the time of the event, but numerous additional checks of data quality were conducted by the detector characterization group to investigate the effects of any conceivable noise source on the detectors' strain signal at the time of GW150914.[28] Omega scans of thousands of auxiliary channels were visually inspected by multiple people to search for any evidence of a glitch that could have caused the signal, but no anomalous behavior was found in any auxiliary channel at the time in either detector. Other strain glitches

were investigated to determine whether any known detector glitches possessed similar chirp-like morphology.

While the background produced by the transient searches using time shifts yields an estimate of the false alarm rate caused by uncorrelated noise in the detectors, it is also important to check that there could not be any environmental noise that could affect both sites simultaneously.[28] Potential noise sources considered include the electromagnetic effects of lightning strikes, radio frequency communication, and cosmic ray showers. The coupling between magnetometer signals and GW strain was thoroughly analyzed to determine that none of these noise sources were likely to influence the gravitational wave signal at the time of GW150914.

In addition to checking the data quality immediately surrounding GW150914, the transient background during the week of the event was compared to the results from the following weeks to ensure that the background estimation was consistent throughout the period of the analysis. If the background triggers near the event had exhibited distinctly different character from the rest of the run, it would have raised questions about whether the total background from the six-week period was actually representative of the noise sources at the time of the event.

Fortunately, the types of instrumental noise remained fairly constant throughout the time represented. By comparing the frequency and amplitude of background events from each week, it can be seen that the distribution of events is similar for the duration of the analysis period, as shown in Figure 4.7. The stability of the background over time provides support for using a larger amount of detector time to estimate the false alarm rate of the gravitational wave candidate.

4.2.2 GW150914 and search background

After thoroughly checking the quality of the data and investigating numerous possible noise sources, no plausible instrumental or environmental cause could be identified.[28] The gravitational wave signal was therefore concluded to be astrophysical in origin. Searches specifically designed to detect compact binary coalescences detected the signal with a false alarm rate of less than 1 in 203,000 years, corresponding to a statistical significance of greater than 5.1σ . [57] The properties of the signal are consistent with relativistic predictions of a merger of two black holes, with masses




Figure 4.7: The cWB background distribution over time during the analysis of GW150914.[28] The lefthand histogram shows the frequency distribution of cWB triggers for each calendar week of the analysis, as a fraction of the total triggers present. The righthand plot shows the individual triggers plotted over time, with color representing the cWB coherent amplitude. The key features of the distribution remain constant throughout the analyzed period, providing support for the usage of the entire time to represent the background of GW150914.

$36^{+5}_{-4}M_{\odot}$ and $29^{+4}_{-4}M_{\odot}$.[58] Figure 4.8 shows the gravitational wave signal seen by both detectors and the comparison with numerical relativity waveforms.

Because of the high level of coherence between detectors and amplitude of the event, the searches for generic transients also detected GW150914 with high confidence, after the division of search bins was applied to separate most of the blip glitches from the rest of the background. The level of confidence in this search would have been convincing, even if the waveform had not aligned so well with relativistic predictions. An additional cWB search bin was added to compare GW150914 against a background comprised only of signals with increasing frequency, similar to the chirp of a binary inspiral.

For this analysis, the detectors collected data for 38 days, with a total analyzable time of about 16 days. The cWB background estimation was performed by using over 1.6×10^6 time shifts, resulting in a total background time of 67400 years. Figure 4.9 shows the significance of GW150914 when compared against two different background distributions: one with signal-based cuts applied to remove noise transients (labeled as C2 + C3), and the other with a further restriction placed on the frequency evolution (labeled C3). The blip-dominated bin (C1) is not shown in this figure, since it did not contain the gravitational wave signal.

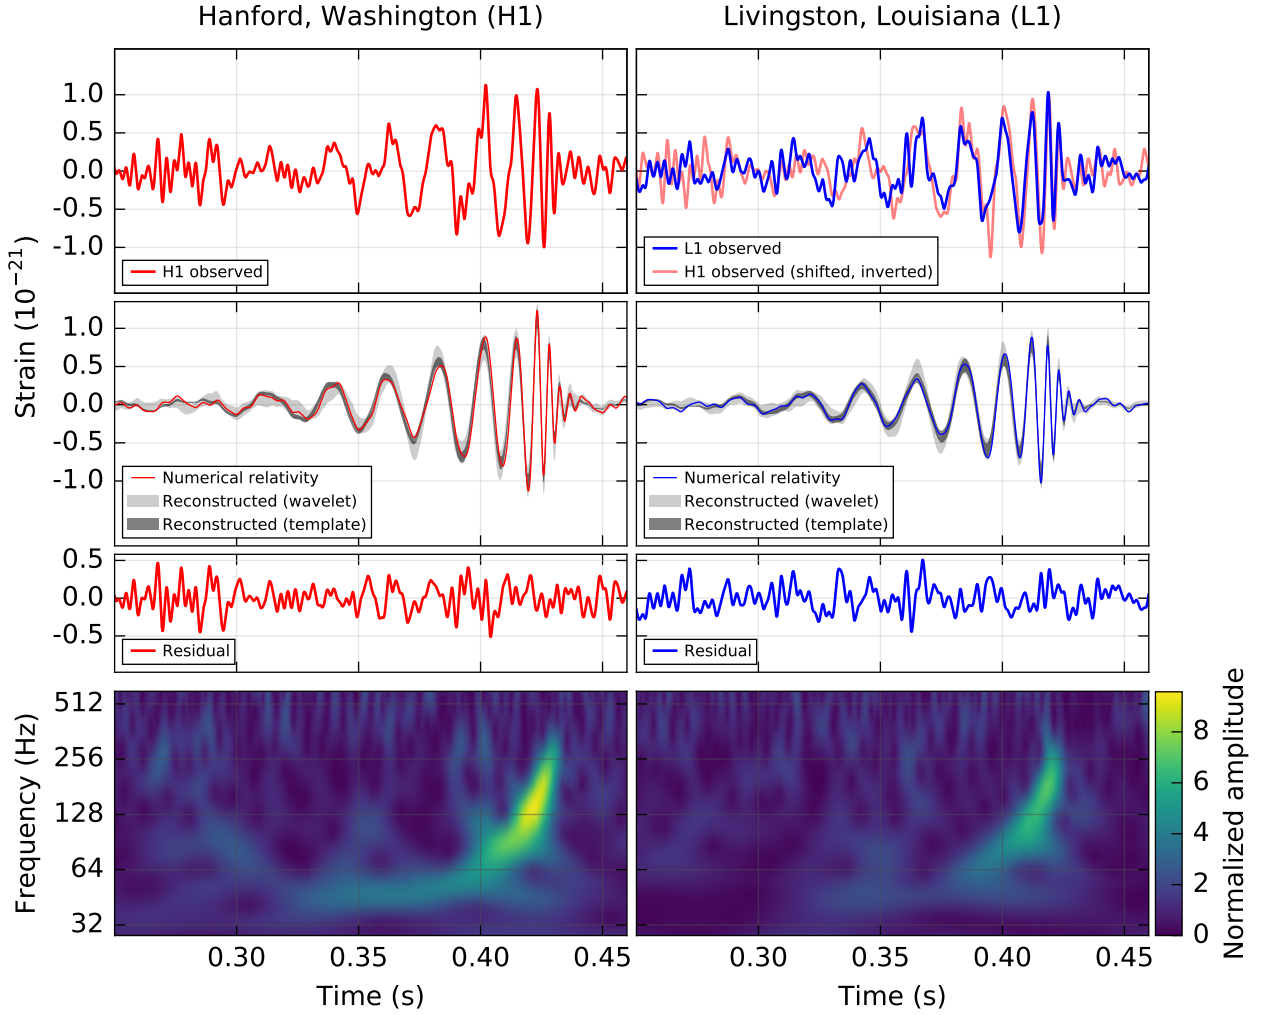


Figure 4.8: GW150914 as seen by the Livingston and Hanford detectors.[5] The top left plot shows the time series of the gravitational wave strain data seen in Livingston, after a bandpass filter has been applied. The top right plot shows the comparison between the signal seen at both sites, after the appropriate time delay has been applied and the Hanford signal has been inverted to account for the difference in physical orientations of the two detectors. The second row shows the waveforms predicted by numerical relativity, superimposed with the reconstructions of the waveform produced by the search algorithms. The third row shows the residuals of the differences between the numerical relativity calculations and the search reconstructions. Finally, the bottom row shows the spectrograms of the signals seen in both sites, normalized by the average spectrum.

The false alarm rate of the event can be calculated from the rate of background triggers with an equivalent or greater amplitude. In the constrained background (C2 + C3), there were four events louder than GW150914, so the rate at which one could expect to get an event with that amplitude by chance is 4 events per 67400 years, or once per 16850 years. Applying the trials factor of two to account for the other search bin (C1), the false alarm estimate of GW150914 would be 1 per 8400 years, equivalent to 4.4σ . By only considering the chirp-like background C3, there can only be an upper limit placed on the false alarm rate, since there are no background events in that bin with a larger amplitude than GW150914. So after the appropriate trials factor of three has been applied, the upper limit of the false alarm rate is 1 per 22000 years, greater than a 4.6σ confidence.

4.3 A new window to the universe

The detection of GW150914 was an important discovery not only because it proved the existence of gravitational waves and the LIGO's ability to detect them, but also because it marked the beginning of a new era of gravitational-wave astronomy. Gravitational waves provide insights outside the realm of electromagnetic observations, and the first detection already deepened our understanding of the universe.

Analysis of the binary black hole merger signal allowed the first tests of general relativity in the strong-field regime, and the signal was shown to be consistent with the theory of general relativity.[59] Electromagnetic followup of the event was conducted by 25 collaborating astronomy teams.[60] A black hole merger is unexpected to produce electromagnetic radiation, so the non-detection of an electromagnetic counterpart was unsurprising, but the exercise was a useful process which will play an important role in future detections.

GW150914 even has some interesting implications for the searches for the stochastic background of gravitational waves.[61] The masses of the black holes indicates that there could a population of binary black hole mergers with higher masses than was previously known. The stochastic background from the superposition of signals from all of these merging binaries could therefore be greater than prior estimates, possibly even detectable by the Advanced LIGO and Virgo instruments operating at their full design sensitivity.




Figure 4.9: The cWB search background and results from September 12 to October 20, 2015.[5] The C2+C3 histogram shows the background distribution after applying the Q -threshold and line veto, while the C3 histogram shows the events after requiring an increasing frequency through the signal. The number of events is normalized to the time of the coincident analysis, and the search results plotted with red and blue squares are the distribution of coincident events. The detection statistic η_c is the $\sqrt{2}$ times the cWB ρ , which is analogous to the SNR calculated by the CBC searches. GW150914 is a clear outlier, but all other coincident events are consistent with the background rate. Above the histograms are shown the equivalent statistical significance of detection for each bin. After the chirp cuts, only a lower bound of 4.6σ on the significance can be stated, since there are no events louder than GW150914 in the background.

The future of gravitational-wave astronomy is quite promising, as detectors continue to improve and more advanced interferometers will soon join the network. The LIGO detectors have now made two confident detections of signals from black hole mergers, GW150914 and GW151226, with more detections expected in the future. Additionally, there will be new types of signals, from neutron stars, continuous waves from pulsars, supernovae, or unpredicted transient signals.

CHAPTER 5

TRANSIENTS IN SUSPENSION SYSTEMS

The main low frequency noise source for LIGO is seismic activity. The LIGO detectors are affected by earthquakes from around the world, windy weather that shakes the buildings housing the interferometer instrumentation, microseismic vibrations from ocean waves crashing on the shores of the Pacific Ocean, Atlantic Ocean, and the Gulf of Mexico, and local anthropogenic activity. [28] The study of the effects of seismic activity was especially important for improving the data quality of transient gravitational wave searches in the initial LIGO era.[62] One of the key improvements from initial to Advanced LIGO is the implementation of a much more sophisticated seismic isolation system, which includes stages of active and passive isolation for all of the cavity optics.

With this entirely new seismic isolation system, it is important to check that the new system provides better isolation as expected, without introducing any new types of transient noise that could add to the noise background for searches of short duration gravitational waves. Here we present a special investigation into the transient motion of the Livingston suspension systems as measured by local sensors on each suspension, specifically looking at the displacement of the quadruple stage pendulums in the longitudinal degree of freedom, which is the direction of the optical path used to sense spacetime strain induced by passing gravitational waves.¹

5.1 Advanced LIGO Suspensions

In Advanced LIGO, all optical cavities use optics suspended from multi-stage pendulums. The input and end mirrors (the most important optics to be isolated from seismic disturbances) are all hung from quadruple-stage suspensions, with each stage providing additional isolation at frequencies above the suspension resonances, which range from 0.4 Hz to 14 Hz. The quadruple pendulum is suspended at the top from maraging steel blade springs, with two further sets of

¹The material presented in this chapter has also been prepared as a journal article, posted on the arXiv database,[63] and has been submitted for publication.




Figure 5.1: Quadruple suspensions design. The left image shows the suspension systems with the blades, fibers, and reaction chain. On the right the whole structure is shown, with the four masses labeled. [29]

springs incorporated into the top two masses, thus providing three stages of enhanced vertical isolation. The two lower masses of the quad are cylindrical silica masses connected by fused silica fibers to reduce thermal noise. Another similar quadruple suspension is hung next to the test mass suspension, so the actuation on lower stages can be done from a similarly isolated reaction chain. [29] Figure 5.1 gives an overview of the design of the Advanced LIGO quadruple suspensions.

The motion of any stage of the suspensions due to a force at another stage can be predicted using the modeled suspensions transfer functions. These models show the frequency response of each stage and degree of freedom to a force at any stage or degree of freedom. For example, Figure 5.2 shows the transfer functions that compare displacement at the top stage of the quadruple suspension with the next two stages, both in the amplitude of the displacement caused and in the phase delay. The shown model assumes a sinusoidal longitudinal force directly on the top stage of the suspension, propagating downwards. Additional motion from different degrees of freedom or forces on other stages would complicate the relationship between the stages, but this model provides an accurate depiction of the behavior of transient motion originating at the top stage.




Figure 5.2: Bode diagram for the modeled longitudinal displacement from the top stage of the quadruple suspension to the two subsequent stages. The upper portion of the figure shows the relationship between the amplitude at the different stages, while the lower plot shows the phase delay. Below the resonant frequencies, the amplitude is expected to be equivalent, at zero decibels. Various resonant peaks lie between 0.4 and 3 Hz, above which the $1/f^2$ amplitude suppression of each stage becomes evident.




Figure 5.3: OSEM schematic

The local displacement of each stage of the suspensions is measured using Optical Sensor and ElectroMagnetic actuators, or OSEMs, which are electromagnetic sensors and actuators used for damping the suspensions' resonances and controlling the mirrors to keep cavities aligned and locked. [29, 64] OSEMs sense displacement using a shadow sensor: a flag attached to a magnet that moves between an LED and a photodiode. The motion sensed is the relative motion of wherever the magnet is attached (one stage of the suspensions) and where the LED and photodiode are mounted. The actuation is done by sending a current through a coil actuator that pushes on the magnet. Figure 5.3 shows the OSEM schematic. On double and triple suspensions, the coil actuators are mounted to the seismic isolation frame so the displacement sensed is relative to the frame. In the input and end test mass suspensions, only the top stage displacement is measured relative to the seismic isolation frame, while the lower stage OSEMs are mounted on the reaction chain, so that actuation can be done from the stages with lower seismic noise.

The OSEMs can sense suspension motion at low frequencies where the displacements are relatively large. At frequencies above 5 Hz the suspension motion has typically fallen below the sensitivity level of the OSEMs such that the resulting spectra are dominated by electronics noise. Multiple OSEMs on each stage allow the calculation of the mass's motion in each degree of freedom using a linear combination of the sensors' signals. Figure 5.4 shows the coherence between different sensors on the same mass from a Livingston suspension. At low frequencies, the signal comes from




Figure 5.4: OSEM coherence

real motion in the suspension and therefore the sensors are coherent. The coherence decreases above 2 Hz, and by 10 Hz the signals are incoherent.

The sensed displacement of the top stage is used in a feedback loop to actuate on that stage of the suspension in order to damp the mechanical resonances of the suspension. The sensors at lower stages are only used as witnesses of the optics' displacement, for the purposes of diagnosing problems in the suspensions. The actuators on lower stages use interferometer and cavity signals to keep various degrees of the interferometer precisely on resonance.

The lowest stage on each quadruple suspension does not have the OSEM sensors, but instead has the possibility of being controlled by an electrostatic drive. The angular motion of the test masses is sensed using optical levers, but since these signals have a different set of transient noises and do not measure the longitudinal motion, they are not used in this study.

5.2 Motion transients in suspensions

It is important to understand the origins of non-astrophysical noise transients in the gravitational wave data in order to eliminate false positives from the gravitational wave searches. Each

subsystem of the detector itself is therefore investigated in great detail to fully study all potential noise sources. [28]

In this article, we characterize transients in the displacement of the suspensions' stages as measured by the OSEMs, the propagation of transients between different stages, and their effects on the gravitational wave strain data. Specifically, transients in the longitudinal degree of freedom were studied in the top three stages of the quadruple suspensions.

Motion transients seen by the local displacement sensors have a few potential sources. For example, they could be caused by motion that is intrinsic to the suspension systems themselves, from the crackling in the suspension wires or the steel blades. [65, 66] Transients could also come from excess seismic motion by propagating through each stage of active seismic isolation and then down through the suspension stages. Above the suspension's resonance f_0 , the seismic transients should decrease in amplitude by a factor of $(f/f_0)^2$ at each stage and be less likely to appear above the sensor noise at lower stages. Therefore any seismic transients that affect multiple stages should appear mostly at low frequencies. Another source of transients seen in the local sensors is the actuation on the suspensions from the feedback loops used to control the interferometer. These glitches involve local displacement of the mirror, but the transient motion is not intrinsic to the suspension itself but is rather caused by a glitch in a different part of the interferometer, and the global control signals respond by pushing on the suspension. The main suspension used for the differential arm length control is the Y-end mirror, so that suspension is the most likely to be affected by glitches elsewhere in the interferometer. However, all suspensions are globally controlled in some way, either for the length or angular controls for different interferometer degrees of freedom.

5.2.1 Typical suspension behavior

The typical spectrum of the suspension motion monitors is characterized by several peaks near the low frequency pendulum resonances between 0.4 to 5 Hz, and the smoother noise above 5 Hz due to the sensors' electronics noise. The main resonances in the longitudinal degree of freedom are modeled for the quadruple suspensions to be at 0.435 Hz, 0.997 Hz, 2.006 Hz, and




Figure 5.5: Typical amplitude spectral density (ASD) for the ITMY longitudinal motion monitors from O1. Many of the low frequency features of the ASD correspond to the pendulum resonances of the suspension. The smooth portion of the ASD above 5 Hz shows where the electronics noise of the OSEM dominates the spectrum. The noise in the penultimate stage at higher frequencies is slightly higher because it has a different kind of OSEM [64].

3.416, but coupling from other degrees of freedom and the active seismic isolation system creates additional peaks in the spectrum. Figure 5.5 shows a typical spectrum of the Y-arm input quadruple suspension (Input Test Mass Y, or *ITMY*) motion in O1 along with estimated sensor noise levels.

Ideally, the measured displacement would be stationary and the average spectrum would statistically characterize the noise level, but in actuality there are non-stationary disturbances at different frequencies. To demonstrate this, Figure 5.6 shows the time series of several minutes of data from the top stage of one suspension with two different bandpass filters applied to select for 4-5 Hz and 10-11 Hz. Above 10 Hz, the sensor noise dominates the signal and the resulting time series is Gaussian distributed, but the lower frequency shows large non-Gaussian transients.

5.2.2 Suspension transients in O1

The Omicron algorithm is used to find transients in the data, producing *triggers* that indicate the time, frequency, amplitude, and signal-to-noise ratio of the transient noise. [42, 37] Figure 5.7 shows distribution in frequency and signal-to-noise ratio of Omicron triggers for the longitudinal




Figure 5.6: An example of non-stationary noise in the suspension monitors. The panel on the left shows a five minute time series from the top stage of ITMY, with bandpass filters applied between 4 to 5 Hz and 10 to 11 Hz. The distribution of time series amplitude over the same time is shown on the right, with dashed lines to indicate a Gaussian distribution. While the Gaussian-distributed stationary sensor noise dominates the higher frequency band shown, the time series from 4 to 5 Hz exhibits large excursions from the average noise level.

degree of freedom, using data from the Y-arm input suspension over one week of the observing run. While stationary noise would produce a background of low SNR triggers across all frequencies, the actual data from the suspension monitors shows a varying structure in different frequency bands. This suggests the presence of non-stationary noise sources.

- Nonphysical data glitches

For the most part, the glitches above 10 Hz are low amplitude and uncorrelated between suspension stages. The transients from both Livingston end test masses (ETMX and ETMY), however, show a population of high frequency triggers that are above the typical noise level. Figure 5.8 shows the amplitude of the ETMY high frequency triggers plotted over time, showing that these higher amplitude glitches are actually clustered in a very few short periods of time.

Further inspection of the times showed that there was not a significant increase in gravitational wave strain noise at the times. Then an investigation of individual glitches uncovered an unusual specific pattern in the data, pointing towards a data recording issue rather than an actual glitch in suspension motion. Figure 5.9 shows this behavior. Every time this glitch has been seen, it is the




Figure 5.7: Omicron trigger distributions from the ITMY suspension monitors from one week of data during O1. Above, the Signal-to-Noise Ratio (SNR) and central frequencies of Omicron triggers of ITMY suspension during one week of O1. Below are histograms showing the SNR distributions of the penultimate stage triggers at three selected frequency ranges (note the different SNR scales). While the distribution of higher frequency triggers fall off much like Gaussian noise, the lower frequency ranges contain more outliers.




Figure 5.8: ETMY high frequency triggers in O1

same pattern: the data is repeated for one second, and then again for half a second before going back to normal. In the periods of high frequency high amplitude Omicron triggers in the end test masses, this glitch actually only happens once, but since it is a large glitch in the data it changes the estimation of the background spectrum that Omicron uses, creating an artificially high number of glitches around that time.

Not only do these glitches appear in the end test masses, but several glitches from the input test masses also showed this type of feature. In ITMY, five out of the loudest eight triggers with a frequency near 1 Hz or higher demonstrated this pattern of data repetition. The few times that the data was corrupted by this feature can be found by checking the data acquisition system. The triggers used for the rest of the analysis do not include those times.

- Interferometric controls

At both sites, the Y-end test mass suspension is actuated upon in response to interferometric variations to keep the arm cavities aligned and in resonance, through a feedback loop used to control the differential arm length (DARM). Therefore, triggers from the ETMY suspension show




Figure 5.9: Time series around loud high frequency glitch in ETMX

distinct features at different frequencies from the other suspensions at both sites, particularly in the frequencies between 0.9 and 4 Hz.

To investigate the relationship between global interferometer control and the glitches seen in suspensions channels, Omicron triggers were produced for the signals used for longitudinal actuation on the Livingston suspensions from the interferometer feedback. Of all of the longitudinal displacement transients in the upper intermediate stage of the main signal used for actuation from the DARM degree of freedom (ETMY), 57% of the signals coincided with a trigger from the actuation signal. For ITMY, which is used only in the alignment feedback loops, only 2.2% of the displacement triggers occurred during an actuation trigger.

Since the ITMY suspension is not part of the DARM length feedback loop, its transients are more dependent on the characteristics of the suspension. For this reason, the ITMY suspension is used for the later analysis. Even so, some of the louder ITMY longitudinal displacement triggers are found to correlate with the interferometer feedback sent to actuate on the lower stages of the suspension. For example, Figure 5.10 shows a bandpassed time series of the longitudinal




Figure 5.10: A filtered time series in ITMY displacement and locking channels, showing a correlation between actuation from the control signals and the resulting displacement of the suspension. Longitudinal displacement are channels shown above, while the lower stage locking channels are plotted below. The time series has been bandpassed around the frequency of the Omicron trigger at the top stage, 2.5 Hz.

displacement channels as well as the locking signals used, around the time of a glitch found by Omicron at the top stage with a frequency of 2.5 Hz and an SNR of 22.7.

5.2.3 Correlations between suspensions stages

An ROC curve was used to compare the transients found by Omicron in longitudinal degrees of the different stages of the ITMY suspension. Figure 5.11 shows these correlations, plotted by varying time coincidence window from 0 to 10 seconds. The percentages shown in this plot are the fraction of the lower stage suspensions that are correlated with a glitch from the higher stages. The correlation between the upper intermediate stage with the penultimate stage is significantly higher than the correlation between either stage with the top stage. The main reason for this is that the displacement sensed at the top stage is relative to the seismic isolation frame, whereas the lower stages' motion are measured with respect to the reaction chain suspension. As one can




Figure 5.11: ROC curves showing correlations between Omicron triggers in different stages of the ITMY suspension. Since the efficiency is significantly higher than the false alarm rate in each comparison, these curves demonstrate that the coinciding transients correspond to physical motion that is propagated through the suspension chain.

see from the ROC curve, a significant correlation is found even with smallest time window. With a 0.1 second window, 66 % of the Omicron triggers from the third stage (L2) are correlated with a trigger at the stage immediately above it (L1), with a rate of only 1.8% accidental coincidence. These coincident triggers also make up a significant fraction (61%) of the total events at the L1 stage. This shows that these transients show real motion in the suspensions rather than just sensor noise, which would be uncorrelated between different stages.

5.3 Propagation of motion transients

To characterize the effects of short duration disturbances in the upper stages of the suspensions on the motion lower in the suspension chain, we need more than just the frequency domain models usually used to characterize the suspensions. Simulink was used to model the response of a simple pendulum to a sine-Gaussian input signal. It is important to note that the input signal is not a pure sine wave at a single frequency, but rather a sine-Gaussian characterized by a peak frequency while also containing broader frequency content. Therefore, the pendulum response more strongly

attenuates the higher frequencies of the input signal, and the peak frequency of the resulting motion is a mixture of the driving and the resonance frequencies, as demonstrated in Figure 5.12.

To examine the propagation of transients in Advanced LIGO suspensions and compare with expected behavior, sine-Gaussian waveforms were physically injected in the Y-end quadruple suspension (End Test Mass Y, or *ETMY*) in the longitudinal direction using the top mass actuators, with central frequencies ranging from 2 to 10 Hz.

The Omicron algorithm was used to characterize the resulting transients caused in the top stage as well as in lower stages. Using a time window of 0.1 seconds, the number of injections coincident with Omicron triggers in the lower stages was calculated. Figure 5.13 shows the amplitudes of the injections that caused high enough amplitude motion in the lower stages to create an Omicron trigger. Of the injections with an amplitude above $3.3 \times 10^{-3} \mu m$ in the top stage, 91% were coincident with a trigger at the next stage, while only 66% were coincident with a glitch in the third stage. However, of the injections at the top stage above an amplitude of $8.4 \times 10^{-3} \mu m$, 97% of the injections coincided with a trigger in the upper intermediate stage, and 90% correlated with a trigger in the penultimate stage. As expected, only the higher amplitude injections propagated down the suspensions chain at an amplitude large enough to be detected above the sensor noise by the Omicron algorithm.

Figure 5.14 shows the ratio of the Omicron trigger amplitudes of the lower suspension stages to the top stage for different frequencies, using the Omicron frequency estimate of the top stage trigger. The solid lines show the frequency response of the suspensions as predicted by the quadruple suspension models. One reason for apparent discrepancies from the model is variation of the transient motion frequency between stages, as well as the fact that the motion at each stage is not characterized by only a single frequency. To analyze this effect, time series of the injections were examined individually to characterize the frequencies and amplitudes of the signals at each stage, similar to the process used in analysis of the Simulink model shown in Figure 5.12.

Using a bandpass filter with a 1 Hz window around the injection frequency and finding the local maxima and minima of the resulting time series, the peak frequency of the induced transient motion

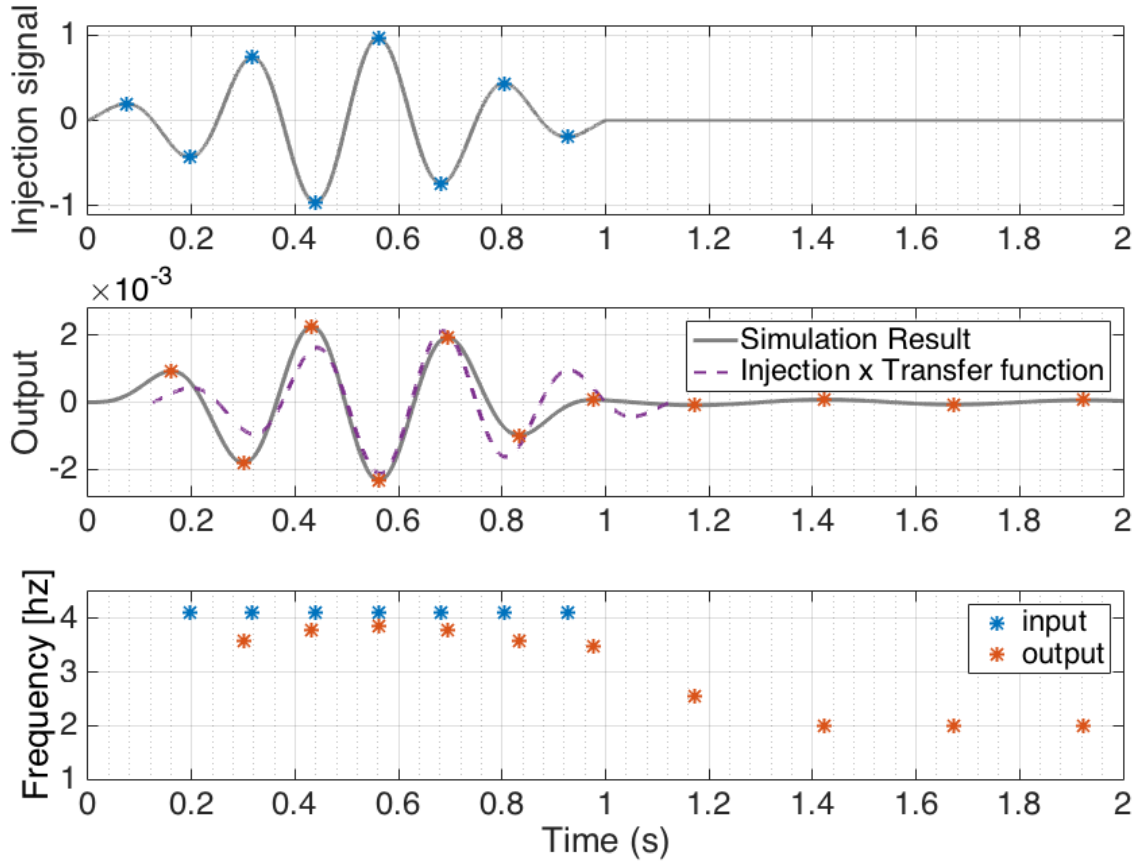


Figure 5.12: Modelled transient response of a simple pendulum to a sine-Gaussian injection. Simulink was used to model a simple pendulum with a resonance at 2 Hz. A one-second 4 Hz sine-Gaussian signal (top panel) was used as the input to show the response of the system (middle panel) compared with the input signal multiplied by the transfer function of the system at 4 Hz. Local maxima and minima of the time series can be used to calculate the frequency for each half-cycle (bottom panel). This simulation shows that even for this simple model, the transient response of the system deviates from the steady state frequency response at 4 Hz.

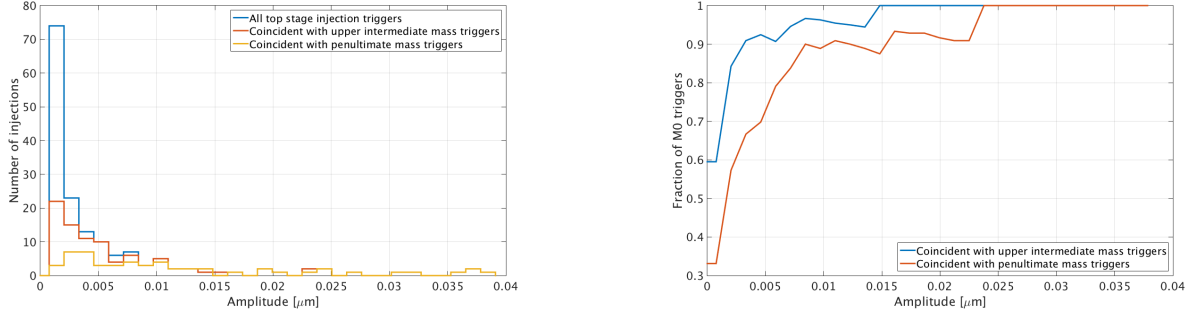


Figure 5.13: Amplitude distribution of injections at the top stage of the suspension, for all injections and for those coincident with lower stages. Histogram of the amplitude of top stage injections found to be coincident with Omicron triggers at each stage (left) and the fraction of coincident triggers above each amplitude threshold (right).




Figure 5.14: Ratio of amplitudes of injection Omicron triggers at lower stages to the top stage, plotted against the peak frequency estimated by Omicron for the top stage motion, compared with the modeled transfer function. At lower frequencies the propagation of the transients is close to the model, but above a few Hz, the motion at the lower stages is smaller than the sensor noise, and the amplitude ratio is not as close to the model. There are fewer Omicron triggers at the penultimate stage, since the motion at that stage is at a lower amplitude and is not great enough at higher frequencies to be seen above the sensor noise.




Figure 5.15: Time series from one injection at 4.1 Hz, after application of a bandpass filter with a window of 1 Hz around the injection frequency. Similar to the simple pendulum model analysis, the frequency shifts throughout the time series. The period of the cycles in the top stage lengthens slightly, from 0.25 seconds (4.0 Hz) at the peak of the transient to 0.28 seconds (3.6 Hz) a few cycles later.

was estimated with each cycle. Similar to the simulation performed in Simulink, the suspension's response is not exactly at the peak sine-Gaussian frequency, and when the injection is finished the suspension's motion begins to ring down with a frequency approaching the nearest resonance. As the motion propagates downwards, the pendulum filter response attenuates the signal more in the frequency range farther from the resonance, resulting in a slight frequency shift towards the resonance at the lower stage. Figure 5.15 shows the period increasing in the bandpassed time series from one of the injections.

The bandpass filter reduces the noise so that the time series cycles can be clearly determined. The frequency of the resulting motion at each stage was then estimated by taking the mean frequency, weighting each cycle by the amplitude of its maximum or minimum. The amplitude of motion was calculated using Omega, a multi-resolution technique for studying transients related to Omicron. [42, 37] The Omega amplitude calculation was used because it gives the peak amplitude, whereas Omicron gives an integrated amplitude over the most significant tile. Using the weighted

average frequency and the amplitudes calculated by Omega, ratio of motion transient amplitudes between suspension stages for each frequency can be better compared to the suspension model.

Figure 5.16 displays the results of the comparison for the propagation of the motion from the top stage to the second and third stages. Since the transient amplitude is much smaller at higher frequencies for each successive stage, the higher frequency injection measurements are farther from the model, due to the sensor noise at the lower stages. In both lower stages we see a shift in the frequency away from the frequency at the top stage, generally closer to the nearest suspension resonance, a pitch mode at 2.7 Hz.

Having understood the propagation of short transients in the suspension stages, we turn now to studying the effect of the actual suspension transients on the LIGO gravitational wave strain data during the first Advanced LIGO observing run.

5.3.1 Correlations with gravitational wave strain data

Taking the data from the first week of November (the same week shown in Figure 5.7), Omicron was used to identify transients in the gravitational wave (GW) strain data in the same frequency range as used to produce the suspension motion triggers (0.1 to 60 Hz), as well as at higher frequencies to check for any nonlinear coupling. Figure 5.17 shows the correlations between transients in the ITMY longitudinal displacement data and the gravitational wave strain data in both frequency ranges. The figures shown are Receiver Operator Characteristics (ROC) curves, which show the time coincidence rate between the two sets of triggers, with various coincidence windows from 0.1 to 10 seconds. This rate is compared with the number of time coincidences that would occur by chance (false alarm rate), using a number of time shifts between the two data sets. In both cases, the small number of coincidences between the sets of data are consistent with the number that would be expected by random chance. The observed transients in the ITMY suspension motion monitors did not show any significant correlation with GW strain noise transients, at any frequency.

We can now place upper limits on the level of noise that would be caused in GW strain from the observed transients in suspension monitors. The amplitude of the Omicron triggers from each of the upper stages of ITMY is multiplied by the suspension transfer function to estimate the

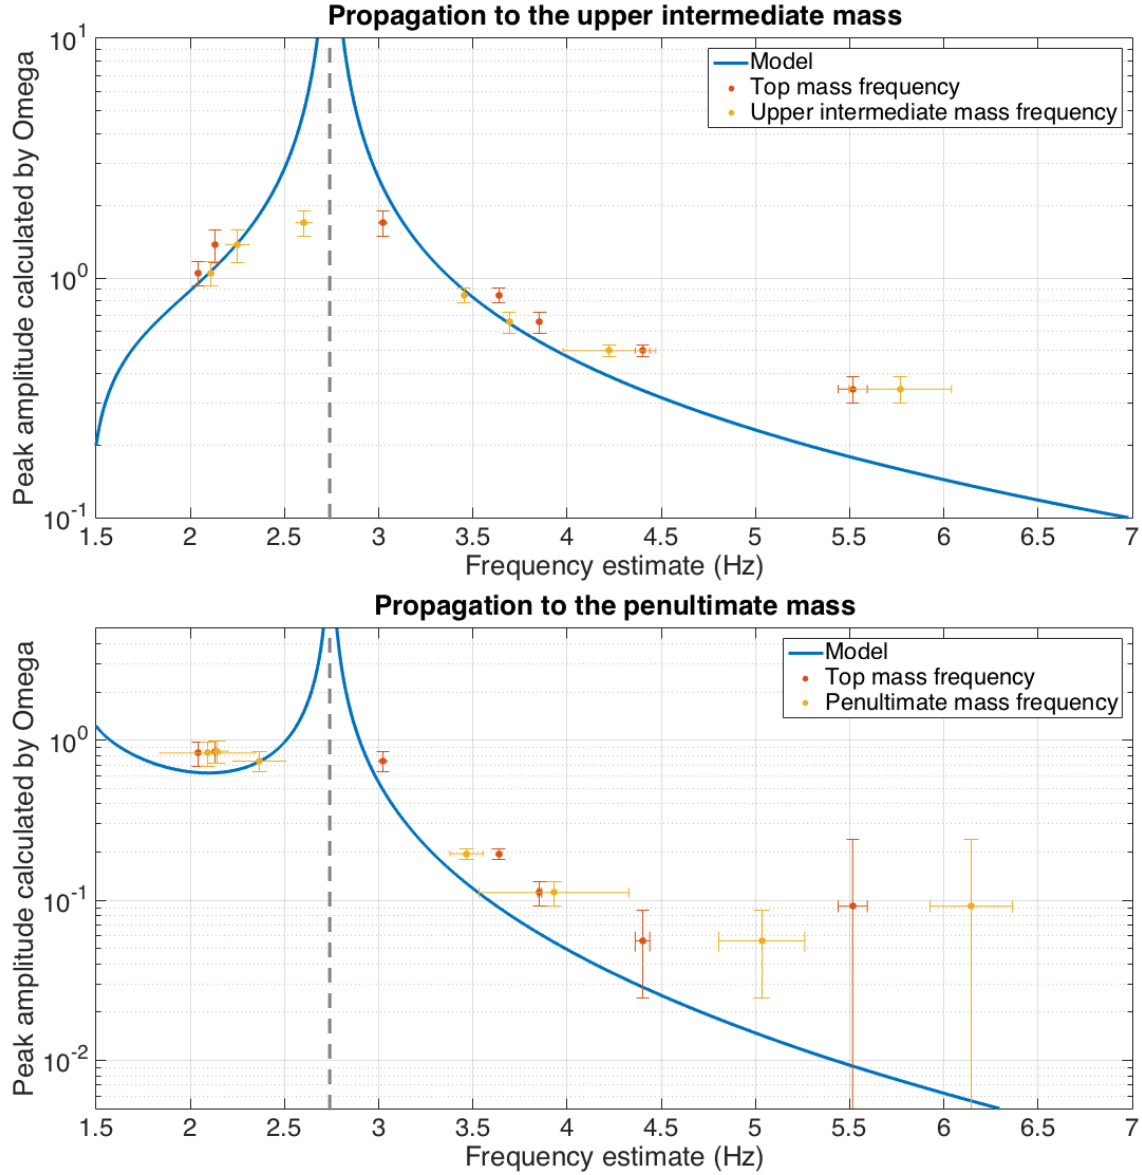


Figure 5.16: Amplitude ratios as calculated by the Omega algorithm, and frequency estimated using the maxima and minima of bandpassed time series. Errors are greater as sensor noise becomes dominant at higher frequencies. Red and yellow points represent the same amplitude ratios between stages, but red points show the frequency estimate at the top stage while yellow points show the frequency estimate at the lower stages. Error bars shown are the standard deviation of the measurement among the various injections of the same frequency, weighted by the amplitude of the injection at the top stage.

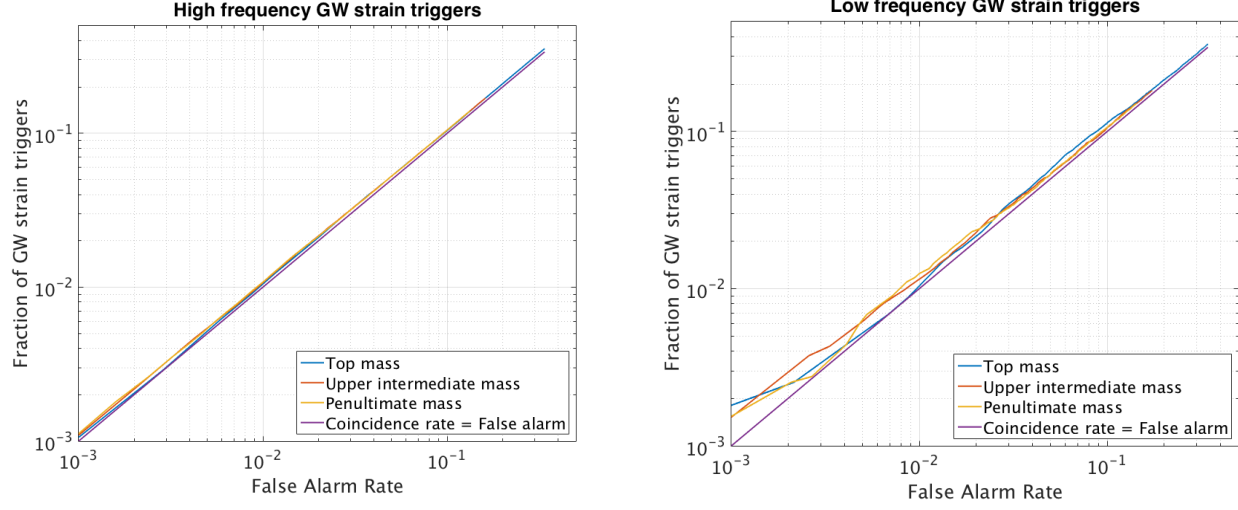


Figure 5.17: Receiver Operator Characteristic (ROC) curves showing the correlation between noise transients in the GW strain and ITMY suspension data from November 1 to 8, 2015. The lefthand plot shows the correlation with higher frequency GW strain triggers (above 60 Hz), while the plot on the right shows the correlation with GW strain triggers below 60 Hz. The y axis shows the fraction of triggers coincident between the two sets of data for varying time windows. The x axis represents the number of coincidences that would appear by chance, estimated by repeating the analysis at each time window with different time shifts between the two data sets. For both sets of GW strain triggers, the coincidence rate is approximately equal to the false alarm rate, whereas a significant correlation would have a much greater efficiency than false alarm rate.




Figure 5.18: Livingston ITMY triggers from a week in the first observing run, multiplied by the transfer function to the lowest stage and divided by the arm length to convert the displacement into equivalent strain amplitude. The strain calibration of the gravitational wave data is only accurate above 10 Hz, at frequencies where the OSEM signals are dominated by sensor noise. Therefore, this calculation can only give us the upper limit of transient motion from each stage that could appear in the GW strain data without also appearing in the local displacement sensor. Where there are GW strain noise transients above one of these levels, we can rule out an origin in a particular stage of the suspension chain. A large number of the GW strain triggers are above the noise level from the top stage, eliminating the origin of the noise at the top of the suspension chain. However, only the very loudest GW strain triggers are above the level of the second stage, and no GW strain triggers are higher than the level of the third stage.

amplitude of noise transients that would be caused in the test mass by a physical displacement of that amplitude. Figure 5.18 shows the resulting projections in equivalent GW strain amplitude, alongside the GW strain triggers from the same time. The sensor noise at the lower stages is much higher than the expected amplitude of motion at those stages, so the upper limit of motion at the lowest stage is above most of the GW strain triggers. The noise level predicted by the top stage triggers, however, is below most of the GW strain triggers up to 37 Hz, so if noise originating in that stage caused high amplitude transients in the GW data, it would be expected to also be seen by the top stage sensors.

Since the top stage triggers are not statistically correlated with any of the GW strain triggers, we can conclude that transient noise originating at the top stage of the suspension is not a signif-

ificant contribution to the transient noise in the interferometer. We cannot, however, rule out the possibility of GW strain noise transients caused by motion originating in the lower stages of the suspension, since there are a significant portion of GW strain triggers that fall below the level of transient noise caused by the local sensor noise.

5.4 Conclusions

Using short duration hardware injections in the top stage of the suspension, we have studied the propagation of transient motion down the suspension chain. The difference of transient amplitudes at different stages is consistent with the models, although slight variations in frequency must be taken into account. The frequency of the transients shifts because the injected waveform is not a pure sine wave but a sine-Gaussian, and after the short duration injection, the suspension motion oscillates with a decreasing amplitude and frequency that shifts toward the closest mechanical resonance frequency. Transients at different stages of the suspension therefore show slightly different frequencies from the same initial sine-Gaussian injection.

Statistical comparisons of the times of transients in the OSEMs and in the GW strain data during O1 show that transients seen by the local displacement sensors of the suspensions are not a significant source of background transient noise in the interferometer. However, this does not rule out transient suspension motion that is below the local sensor noise as a possible source of background noise. Using the suspension models to propagate the sensor noise into the motion at the test mass, upper limits can be placed on the level of noise that could be caused in the GW strain data from transients in suspension motion at each stage. Most GW strain triggers are above the sensor noise level of the top stage of the suspension, but below the noise level of the third stage. Transient noise that originates in the lower stage of the suspension could therefore be a cause of noise in the GW data while not being loud enough to appear above the local sensor noise. [25]

CHAPTER 6

CONCLUSIONS

Through the efforts of many people working together from detector characterization, commissioning, and data analysis, numerous instrumental noise sources were eliminated from the instrument or removed from the data during the first Advanced LIGO observing run. Some transient instrumental issues, such as the glitching caused by radio frequency beating and excited violin modes, were somewhat mitigated before O1 by making improvements to the detector. Other serious problems were effectively removed from the data using vetoes, including the 45 MHz modulation glitches and the ETMY saturations.

Much work remains to be done to reduce instrumental noise sources to allow a confident detection of any arbitrary waveform by the generic transient gravitational wave searches. The transient distribution contains a large non-Gaussian tail of high amplitude glitches in both the Livingston and Hanford detectors. The outstanding problem in both interferometers is the population of blip glitches, which has been investigated extensively but remains a mystery. I have used the hveto algorithm to compare a set of blip glitches with transients in auxiliary channels, but this study did not identify the source of the noise. Eliminating the blip glitch problem would make the searches for generic transients even more powerful. Even with the necessary division of the parameter space into separate search bins, the success of the burst searches in confidently detecting GW150914 proves that they are capable of finding unmodeled gravitational waves.

As part of the effort to thoroughly characterize every subsystem of the complex instruments, I performed a study to analyze the behavior of transient disturbances in the quadruple pendulums used to suspend the test mass optics. The propagation of transients through the suspension chains was found to be consistent with the predictions based on computer modeling, after taking into account frequency variation between stages. Statistical comparisons of the transients in the suspensions with the gravitational wave strain channel show that they are not a dominant source of background noise for the transient gravitational wave searches. However, transient noise that

originates in the lower stage of the suspension cannot be ruled out as a potential cause of noise in the GW data because of the limits of the local displacement sensors.

The sensitivity of the detectors in the first Advanced LIGO observing run allowed the detection of GW150914 with high significance, ushering in a new era of gravitational wave astronomy, but the progress made in this run is only the beginning. We will continue to investigate new noise sources and develop a deeper understanding of the detectors as Advanced LIGO pursues further upgrades and higher laser power to increase its sensitivity in future observing runs. The LIGO detectors will continue to improve and will soon be joined in observation by other advanced detectors, extending the possible range of detections.

More gravitational waves will be seen, from black holes and from new types of sources such as neutron star mergers or the continuous waves emitted from pulsars. Especially exciting would be a detection by the generic transient searches of an unmodeled source, such as a supernova or a neutron star instability. The information gained from these detections would provide great insights into the physics of stars and the mechanisms involved in the emission of gravitational radiation. The searches for unmodeled sources may even detect totally unexpected signals, which could bring up new questions and require deeper theoretical analysis of potential sources. With continued careful detector analysis and application of the data quality principles used in the first observing run, the searches for generic transient gravitational waves may shed new light on many mysteries of the universe.

REFERENCES

- [1] A. Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, Seite 844-847., 1915.
- [2] Albert Einstein. Approximative Integration of the Field Equations of Gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1916:688–696, 1916.
- [3] Peter G. Bergmann. Summary of the chapel hill conference. *Rev. Mod. Phys.*, 29:352–354, Jul 1957. doi: 10.1103/RevModPhys.29.352. URL <http://link.aps.org/doi/10.1103/RevModPhys.29.352>.
- [4] J. H. Taylor and J. M. Weisberg. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *The Astrophysical Journal*, 253:908–920, February 1982. doi: 10.1086/159690.
- [5] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016. doi: 10.1103/PhysRevLett.116.061102. URL <http://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [6] Bernard F. Schutz. Gravitational waves on the back of an envelope. *American Journal of Physics*, 52(5):412–419, 1984. doi: 10.1119/1.13627. URL <http://dx.doi.org/10.1119/1.13627>.
- [7] Wm. Robert Johnston. Gravitational wave-related images'. <http://www.johnstonsarchive.net/relativity/pictures.html>, 2005. Accessed 2017-02-10.
- [8] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing, 1994.
- [9] K. Riles. Gravitational Waves: Sources, Detectors and Searches. *Prog.Part.Nucl.Phys.*, 68: 1–54, 2013. doi: 10.1016/j.ppnp.2012.08.001.
- [10] J. Aasi et al. Gravitational waves from known pulsars: Results from the initial detector era. *The Astrophysical Journal*, 785(2):119, 2014. URL <http://stacks.iop.org/0004-637X/785/i=2/a=119>.
- [11] P. C. Peters and J. Mathews. Gravitational radiation from point masses in a keplerian orbit. *Phys. Rev.*, 131:435–440, Jul 1963. doi: 10.1103/PhysRev.131.435. URL <http://link.aps.org/doi/10.1103/PhysRev.131.435>.
- [12] Jade Powell, Sarah E. Gossan, Joshua Logue, and Ik Siong Heng. Inferring the core-collapse supernova explosion mechanism with gravitational waves. *Phys. Rev. D*, 94:123012, Dec 2016. doi: 10.1103/PhysRevD.94.123012. URL <http://link.aps.org/doi/10.1103/PhysRevD.94.123012>.

- [13] B. P. Abbott et al. Binary black hole mergers in the first Advanced LIGO observing run. *Phys. Rev. X*, 6:041015, Oct 2016. doi: 10.1103/PhysRevX.6.041015. URL <http://link.aps.org/doi/10.1103/PhysRevX.6.041015>.
- [14] B. P. Abbott et al. Astrophysical implications of the binary black-hole merger gw150914. *The Astrophysical Journal Letters*, 818(2):L22, 2016. URL <http://stacks.iop.org/2041-8205/818/i=2/a=L22>.
- [15] B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration). The rate of binary black hole mergers inferred from Advanced LIGO Observations surrounding GW150914. *The Astrophysical Journal Letters*, 833(1):L1, 2016. URL <http://stacks.iop.org/2041-8205/833/i=1/a=L1>.
- [16] B. P. Abbott et al. Binary black hole mergers in the first Advanced LIGO observing run. *Phys. Rev. X*, 6:041015, Oct 2016. doi: 10.1103/PhysRevX.6.041015. URL <http://link.aps.org/doi/10.1103/PhysRevX.6.041015>.
- [17] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysics Journal*, 195:L51–L53, 1975. doi: 10.1086/181708.
- [18] J. Weber. Detection and generation of gravitational waves. *Phys. Rev.*, 117:306–313, Jan 1960. doi: 10.1103/PhysRev.117.306. URL <http://link.aps.org/doi/10.1103/PhysRev.117.306>.
- [19] Barry C. Barish and Rainer Weiss. LIGO and the detection of gravitational waves. *Physics Today*, 52(10):44–50, October 1999.
- [20] B. P. Abbott et al. LIGO: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, 72(7):076901, 2009.
- [21] Alex Abramovici, William E. Althouse, Ronald W. P. Drever, Yekta Gürsel, Seiji Kawamura, Frederick J. Raab, David Shoemaker, Lisa Sievers, Robert E. Spero, Kip S. Thorne, Rochus E. Vogt, Rainer Weiss, Stanley E. Whitcomb, and Michael E. Zucker. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256(5055):325–333, 1992. ISSN 0036-8075. doi: 10.1126/science.256.5055.325. URL <http://science.sciencemag.org/content/256/5055/325>.
- [22] B. P. Abbott et al. LIGO: The Laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.*, 72:076901, 2009. doi: 10.1088/0034-4885/72/7/076901.
- [23] T. Accadia et al. Virgo: a laser interferometer to detect gravitational waves. *JINST*, 7:P03012, 2012. doi: 10.1088/1748-0221/7/03/P03012.
- [24] James Hough, B.J. Meers, Gavin Newton, Norna Robertson, Henry Ward, G. Leuchs, T.M. Niebauer, A. Rudiger, R. Schilling, L. Schnupp, H. Walther, W. Winkler, B.F. Schutz, J. Ehlers, P. Kafka, G. Schafer, M.W. Hamilton, I. Schutz, H. Welling, J.R.J. Bennett, I.F. Cor-

bett, B.W.H. Edwards, R.J.S. Greenhalgh, and V. Kose. Proposal for a joint german-british interferometric gravitational wave detector. Other, Max-Planck-Institut, Germany, September 1989. URL <http://eprints.gla.ac.uk/114852/>. Document number GWD/137/JH(89).

- [25] The LIGO Scientific Collaboration. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015. URL <http://stacks.iop.org/0264-9381/32/i=7/a=074001>.
- [26] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015. doi: 10.1088/0264-9381/32/2/024001.
- [27] S. Hild. Beyond the Second Generation of Laser-Interferometric Gravitational Wave Observatories. *Class. Quant. Grav.*, 29:124006, 2012. doi: 10.1088/0264-9381/29/12/124006.
- [28] B P Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Classical and Quantum Gravity*, 33(13):134001, 2016. URL <http://stacks.iop.org/0264-9381/33/i=13/a=134001>.
- [29] S M Aston, M A Barton, A S Bell, N Beveridge, B Bland, A J Brummitt, G Cagnoli, C A Cantley, L Carbone, A V Cumming, L Cunningham, R M Cutler, R J S Greenhalgh, G D Hammond, K Haughian, T M Hayler, A Heptonstall, J Heefner, D Hoyland, J Hough, R Jones, J S Kissel, R Kumar, N A Lockerbie, D Lodhia, I W Martin, P G Murray, J O'Dell, M V Plissi, S Reid, J Romie, N A Robertson, S Rowan, B Shapiro, C C Speake, K A Strain, K V Tokmakov, C Torrie, A A van Veggel, A Vecchio, and I Wilmut. Update on quadruple suspension design for Advanced LIGO. *Classical and Quantum Gravity*, 29(23):235004, 2012. URL <http://stacks.iop.org/0264-9381/29/i=23/a=235004>.
- [30] Gregory M. Harry, Helena Armandula, Eric Black, D. R. M. Crooks, Gianpietro Cagnoli, Jim Hough, Peter Murray, Stuart Reid, Sheila Rowan, Peter Sneddon, Martin M. Fejer, Roger Route, and Steven D. Penn. Thermal noise from optical coatings in gravitational wave detectors. *Appl. Opt.*, 45(7):1569–1574, Mar 2006. doi: 10.1364/AO.45.001569. URL <http://ao.osa.org/abstract.cfm?URI=ao-45-7-1569>.
- [31] Chris L. Fryer and Kimberly C.B. New. Gravitational waves from gravitational collapse. *Living Reviews in Relativity*, 14(1), 2011. doi: 10.1007/lrr-2011-1. URL <http://www.livingreviews.org/lrr-2011-1>.
- [32] Nils Andersson et al. The transient gravitational-wave sky. *Classical and Quantum Gravity*, 30(19):193002, 2013.
- [33] S Klimenko, I Yakushin, A Mercer, and G Mitselmakher. A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 25(11):114029, 2008. URL <http://stacks.iop.org/0264-9381/25/i=11/a=114029>.
- [34] B. P. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev. D*, 80:102001, Nov 2009. doi: 10.1103/PhysRevD.80.102001. URL <http://link.aps.org/doi/10.1103/PhysRevD.80.102001>.

- [35] G. Mazzolo, F. Salemi, M. Drago, V. Nacula, C. Pankow, G. A. Prodi, V. Re, V. Tiwari, G. Vedovato, I. Yakushin, and S. Klimenko. Prospects for intermediate mass black hole binary searches with advanced gravitational-wave detectors. *Phys. Rev. D*, 90:063002, Sep 2014. doi: 10.1103/PhysRevD.90.063002. URL <http://link.aps.org/doi/10.1103/PhysRevD.90.063002>.
- [36] J Slutsky, L Blackburn, D A Brown, L Cadonati, J Cain, M Cavaglià, S Chatterji, N Christensen, M Coughlin, S Desai, G González, T Isogai, E Katsavounidis, B Rankins, T Reed, K Riles, P Shawhan, J R Smith, N Zotov, and J Zweizig. Methods for reducing false alarms in searches for compact binary coalescences in LIGO data. *Classical and Quantum Gravity*, 27(16):165023, 2010. URL <http://stacks.iop.org/0264-9381/27/i=16/a=165023>.
- [37] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Class. Quant. Grav.*, 21:S1809–S1818, 2004. doi: 10.1088/0264-9381/21/20/024.
- [38] Tomoki Isogai, the Ligo Scientific Collaboration, and the Virgo Collaboration. Used percentage veto for ligo and virgo binary inspiral searches. *Journal of Physics: Conference Series*, 243(1):012005, 2010. URL <http://stacks.iop.org/1742-6596/243/i=1/a=012005>.
- [39] R. Essick, L. Blackburn, and E. Katsavounidis. Optimizing Vetoes for Gravitational-Wave Transient Searches. *Class. Quant. Grav.*, 30:155010, 2013. doi: 10.1088/0264-9381/30/15/155010.
- [40] Jessica L McIver. *The impact of terrestrial noise on the detectability and reconstruction of gravitational wave signals from core-collapse supernovae*. PhD thesis, University of Massachusetts Amherst, 2015.
- [41] Ryan Lynch, Salvatore Vitale, Reed Essick, Erik Katsavounidis, and Florent Robinet. An information-theoretic approach to the gravitational-wave burst detection problem. 2015.
- [42] Florent Robinet. Omicron: An algorithm to detect and characterize transient noise in gravitational-wave detectors. <https://tds.virgo-gw.eu/?content=3&r=11318>, 2015.
- [43] Joshua R Smith, Thomas Abbott, Eiichi Hirose, Nicolas Leroy, Duncan MacLeod, Jessica McIver, Peter Saulson, and Peter Shawhan. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Classical and Quantum Gravity*, 28(23):235005, 2011. URL <http://stacks.iop.org/0264-9381/28/i=23/a=235005>.
- [44] Joshua Smith, Andrew Lundgren, and Patrick Myers. December 25th whistles (rf beats). <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=23869>, December 2015.
- [45] D. V. Martynov et al. Sensitivity of the advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93:112004, Jun 2016. doi: 10.1103/PhysRevD.93.112004. URL <http://link.aps.org/doi/10.1103/PhysRevD.93.112004>.

- [46] William Wesley Peterson and Theodore G Birdsall. The theory of signal detectability. Technical Report 13, Engineering Research Institute, University of Michigan Ann Arbor, June 1953.
- [47] B. P. Abbott et al. Gw151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.*, 116:241103, Jun 2016. doi: 10.1103/PhysRevLett.116.241103. URL <http://link.aps.org/doi/10.1103/PhysRevLett.116.241103>.
- [48] B. P. Abbott et al. All-sky search for short gravitational-wave bursts in the first advanced ligo run. *Phys. Rev. D*, 95:042003, Feb 2017. doi: 10.1103/PhysRevD.95.042003. URL <http://link.aps.org/doi/10.1103/PhysRevD.95.042003>.
- [49] S. Karki, D. Tuyenbayev, S. Kandhasamy, B. P. Abbott, T. D. Abbott, E. H. Anders, J. Berliner, J. Betzwieser, C. Cahillane, L. Canete, C. Conley, H. P. Daveloza, N. De Lillo, J. R. Gleason, E. Goetz, K. Izumi, J. S. Kissel, G. Mendell, V. Quetschke, M. Rodruck, S. Sachdev, T. Sadecki, P. B. Schwinberg, A. Sottile, M. Wade, A. J. Weinstein, M. West, and R. L. Savage. The Advanced LIGO photon calibrators. *Review of Scientific Instruments*, 87(11):114503, 2016. doi: 10.1063/1.4967303. URL <http://aip.scitation.org/doi/abs/10.1063/1.4967303>.
- [50] Ryan Derosa. Back to the usual excess noise. <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=24541>, January 2016.
- [51] Andrew Lundgren. Could PEM have tracked down the SEI RTD noise? <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=24625>, January 2016.
- [52] B. P. Abbott et al. Observing gravitational-wave transient gw150914 with minimal assumptions. *Phys. Rev. D*, 93:122004, Jun 2016. doi: 10.1103/PhysRevD.93.122004. URL <http://link.aps.org/doi/10.1103/PhysRevD.93.122004>.
- [53] Jordan Palamos Paul Schale, Robert Schofield. A class of o1 blip glitches is correlated with low relative humidity in the buildings. <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=28534>, July 2016.
- [54] Jade Powell, Daniele Trifir, Elena Cuoco, Ik Siong Heng, and Marco Cavagli. Classification methods for noise transients in advanced gravitational-wave detectors. *Classical and Quantum Gravity*, 32(21):215012, 2015. URL <http://stacks.iop.org/0264-9381/32/i=21/a=215012>.
- [55] Michael Zevin et al. Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science. 2016.
- [56] B. P. Abbott et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. 2016.

- [57] B. P. Abbott et al. Gw150914: First results from the search for binary black hole coalescence with advanced ligo. *Phys. Rev. D*, 93:122003, Jun 2016. doi: 10.1103/PhysRevD.93.122003. URL <http://link.aps.org/doi/10.1103/PhysRevD.93.122003>.
- [58] B. P. Abbott et al. Properties of the binary black hole merger GW150914. *Phys. Rev. Lett.*, 116:241102, Jun 2016. doi: 10.1103/PhysRevLett.116.241102. URL <http://link.aps.org/doi/10.1103/PhysRevLett.116.241102>.
- [59] B. P. Abbott et al. Tests of general relativity with GW150914. *Phys. Rev. Lett.*, 116:221101, May 2016. doi: 10.1103/PhysRevLett.116.221101. URL <http://link.aps.org/doi/10.1103/PhysRevLett.116.221101>.
- [60] B. P. Abbott et al. Localization and broadband follow-up of the gravitational-wave transient gw150914. *The Astrophysical Journal Letters*, 826(1), 2016.
- [61] B. P. Abbott et al. GW150914: Implications for the stochastic gravitational-wave background from binary black holes. *Phys. Rev. Lett.*, 116:131102, Mar 2016. doi: 10.1103/PhysRevLett.116.131102. URL <http://link.aps.org/doi/10.1103/PhysRevLett.116.131102>.
- [62] D M Macleod, S Fairhurst, B Hughey, A P Lundgren, L Pekowsky, J Rollins, and J R Smith. Reducing the effect of seismic noise in LIGO searches by targeted veto generation. *Classical and Quantum Gravity*, 29(5):055006, 2012. URL <http://stacks.iop.org/0264-9381/29/i=5/a=055006>.
- [63] M. Walker et al. Effects of transients in LIGO suspensions on searches for gravitational waves. 2017. URL <https://arxiv.org/abs/1702.04701>.
- [64] L Carbone, S M Aston, R M Cutler, A Freise, J Greenhalgh, J Heefner, D Hoyland, N A Lockerbie, D Lodhia, N A Robertson, C C Speake, K A Strain, and A Vecchio. Sensors and actuators for the Advanced LIGO mirror suspensions. *Classical and Quantum Gravity*, 29(11):115005, 2012. URL <http://stacks.iop.org/0264-9381/29/i=11/a=115005>.
- [65] I A Bilenko, V B Braginsky, and N Yu Markova. Thermal and excess noise in suspension fibres. *Classical and Quantum Gravity*, 19(7):2035, 2002. URL <http://stacks.iop.org/0264-9381/19/i=7/a=410>.
- [66] G. Vajente, E. A. Quintero, X. Ni, K. Arai, E. K. Gustafson, N. A. Robertson, E. J. Sanchez, J. R. Greer, and R. X. Adhikari. An instrument to measure mechanical up-conversion phenomena in metals in the elastic regime. *Review of Scientific Instruments*, 87(6):065107, 2016. doi: 10.1063/1.4953114. URL <http://dx.doi.org/10.1063/1.4953114>.

VITA

Marissa Walker spent the first years of her life in Shakawe, Botswana. After moving to Abilene, Texas, she studied at Abilene High School, where she graduated as valedictorian in 2006. She studied physics at Abilene Christian University, and earned a B.S. in physics and a teaching certification in 2010. In August 2011 she moved to Baton Rouge to pursue her Ph.D. in physics at Louisiana State University. During her time at LSU, she has done research with the LIGO scientific collaboration and was involved in the analysis of the historic first direct detection of a gravitational wave. She plans to graduate from LSU in May 2017 and continue to work on LIGO research.