

Syracuse University

SURFACE

Dissertations - ALL

SURFACE

August 2016

Detector Characterization for Advanced LIGO

Thomas James Massinger
Syracuse University

Follow this and additional works at: <https://surface.syr.edu/etd>



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

Recommended Citation

Massinger, Thomas James, "Detector Characterization for Advanced LIGO" (2016). *Dissertations - ALL*. 633.

<https://surface.syr.edu/etd/633>

This Dissertation is brought to you for free and open access by the SURFACE at SURFACE. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

ABSTRACT

The first observing run of Advanced LIGO spanned 4 months, from September 12, 2015 to January 19, 2016, during which gravitational waves were directly detected from two binary black hole systems, namely GW150914 and GW151226. Confident detection of gravitational waves requires an understanding of instrumental noise transients and artifacts that can reduce the sensitivity of a search for gravitational waves. Studies of the quality of the detector data yield insights into the cause of instrumental artifacts and data quality vetoes specific to a search are produced to mitigate the effects of problematic data.

This dissertation provides an overview of the methods used to characterize noise in the LIGO interferometers and provides examples of successful removal of transient noise. The data set used in the first observing run is validated. Further, the systematic removal of noisy data from analysis time is shown to improve the sensitivity of searches for compact binary coalescences. The output of the PyCBC pipeline is used as a metric for improvement.

The first direct detection of gravitational waves, GW150914, was a loud enough signal that removing data with excess noise did not improve its significance. However, the removal of data with excess noise decreased the false alarm rate of GW151226 by a factor of 567, from 1 in 320 years (3.9σ) to 1 in 183000 years ($> 5.3 \sigma$).

DETECTOR CHARACTERIZATION FOR ADVANCED LIGO

By

Thomas J. Massinger

B.S. Physics, Utica College, Utica, NY 13502

DISSERTATION

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHYSICS

Syracuse University

August 2016

Copyright © 2016 Thomas J. Massinger
All rights reserved.

ACKNOWLEDGEMENTS

I would first like to thank my advisor, Peter Saulson, for his continued mentorship. Your experience, knowledge, and passion for teaching have inspired me. You taught me to be excited about my work and to have fun doing it. You always make sure that everyone in the room understands the big picture and I hope to carry that philosophy with me.

The Syracuse University Gravitational Wave Group is a truly unique group of people. Diverse scientific interests, expertise, kindness, willingness to teach, and group meetings full of bagels create an environment that fosters the development of graduate students and I wouldn't trade my experiences here for any other group in the LSC.

To Duncan Brown and Stefan Ballmer, thank you for helping me broaden and develop my knowledge of gravitational wave astronomy. You have both taught me a tremendous amount and have encouraged all of the members of the Syracuse group to diversify their interests and understand the experiment from top to bottom.

I would like to thank my officemates: Prayush Kumar, Jim Lough, Chris Biwer, Alex Nitz, Soumi De, and Steve Reyes. Thank you for all of the companionship and for putting up with my horrible jokes. I expect you to continue them once I leave.

To Ryan Fisher, thank you for making sure we enjoyed every minute of our conference travel and for all of the advice and mentorship you've provided.

To Laura Nuttall, thank you for working alongside me, battling through O1 with me, and proofreading so many things. Syracuse DetChar was a lot of fun with you here.

To the Detector Characterization group, Josh Smith, Andy Lundgren, Gaby Gonzalez, Peter Saulson, David Shoemaker, Jess McIver, Laura Nuttall, Ryan Fisher, Duncan Macleod, Marissa Walker, OG Thomas Abbott, Nutsinee Kijbunchoo, Joe Areeda, and everyone else. It's been a pleasure to work with all of you for the last

4 years and you've all taught me something valuable along the way. I always look forward to the next face-to-face meeting.

I would like to thank my committee members, Howard Blair, Matt LaHaye, Alan Middleton, Duncan Brown, Ryan Fisher, and Peter Saulson, for taking time out of their schedules to serve on my thesis defense committee.

To Nathan, Forrest, and Michael, thanks for helping me survive the first few years in Syracuse. It would've been infinitely more difficult without you guys.

To Matt, Danica, Will, Alli, and Joel, thanks for making Syracuse a fun place to live. You are all awesome and I hope we don't end up too far away from each other.

To Mom, Dad, Tim, and Lauren, thank you for always supporting me in everything I've done.

To all of the Millers and Meszlers, thank you for being a second family to me.

To Mary, thank you for everything.

Contents

List of Tables	ix
List of Figures	xi
Preface	xiii
Acronyms	xiv
1 Introduction	1
1.1 The Advanced LIGO Interferometers	5
1.1.1 DC Readout	7
1.1.2 The aLIGO Noise Curve	8
1.2 Sources of Gravitational Waves	11
1.2.1 Compact binary coalescences	11
1.2.2 Burst signals	12
1.2.3 Continuous waves	13
1.2.4 Stochastic background	13
1.3 The Advanced Detector Network	14
2 Searching for Compact Binary Coalescences	15
2.0.1 The matched filter	15
2.0.2 Waveform templates	16
2.0.3 χ^2 signal consistency test	18
2.0.4 Searching for signals	20
2.0.5 Gating	22

3	The First Observing Run	24
3.1	The First Direct Detection of Gravitational Waves	24
3.2	Foreground Events	24
3.2.1	GW150914	24
3.2.2	GW151226	26
3.2.3	LVT151012	26
3.3	CBC Results	26
3.4	Summary	29
4	Detector Characterization	32
4.1	Methods of Detector Characterization	33
4.2	Tools and algorithms	34
4.2.1	Omicron	34
4.2.2	Hierarchichal Veto	36
4.3	Instrumental Detector Characterization Studies	39
4.3.1	Analog-to-Digital Conversion Overflows	39
4.3.2	Suspension DAC calibration glitches	43
4.3.3	RF beatnote whistles	46
4.4	Validation of O1 Data	51
4.5	Validation of Gravitational Wave Signals	53
5	IMC Upconversion	57
5.1	PDH locking	57
5.2	Upconversion noise in aLIGO	63
5.3	Conclusions	65
6	Detector Characterization Subsystem Lead	67
6.1	Length Sensing and Control	69
6.1.1	Online Detector Characterization	70
6.1.2	MEDM screens	73
6.1.3	Summary pages	74
6.2	Alignment Sensing and Control	76
6.2.1	Online Detector Characterization	77
6.2.2	MEDM screens	77

6.2.3	Summary pages	80
6.3	ODC Results	80
6.3.1	MICH ODC as a witness of RF45 glitches	80
7	Data Quality Vetoes	83
7.1	Veto categories	84
7.2	Quantifying the effects of data quality vetoes	84
8	Effects of Data Quality Vetoes on the Analysis Containing GW150914	87
8.1	BNS bin	88
8.2	Bulk bin	89
8.2.1	LVT151012	92
8.3	Edge bin	93
8.3.1	GW150914	94
9	Effects of Data Quality Vetoes on the Analysis Containing GW151226	96
9.1	BNS bin	97
9.2	Bulk bin	98
9.2.1	GW151226	99
9.3	Edge bin	100
10	Limiting Noise Sources in the PyCBC Search	102
10.1	Loud transients	102
10.2	Blip transients	103
10.3	60-200 Hz noise	106
11	Conclusion	110
A	Data Quality Vetoes in O1	112
A.1	Data Quality Vetoes	112
	Bibliography	123

List of Tables

1	Table of GW sources	11
2	Table of foreground events	27
3	HVeto results for ETMY DAC glitches	46
4	Table of aLIGO subsystems	68
5	Table of LIGO degrees of freedom	70
6	BNS bin FAR - GW150914 analysis	90
7	Bulk bin FAR - GW150914 analysis	92
8	LVT150914 FAR	92
9	Edge bin FAR - GW150914 analysis	94
10	GW150914 FAR	95
11	BNS bin FAR - GW151226 analysis	98
12	Bulk bin FAR - GW151226 analysis	98
13	GW151226 FAR	100
14	Edge bin FAR - GW151226 analysis	101

List of Figures

1	Plus and cross polarizations	4
2	Layout of Advanced LIGO	5
3	Sidebands and OMC cavity pole	9
4	Advanced LIGO noise budget	10
5	BNS signal in detector noise	17
6	Transient with BBH waveform	19
7	PyCBC SNR and re-weighted SNR histograms	23
8	GW150914 timeseries	25
9	PyCBC result histograms for GW150914	29
10	PyCBC result histograms for GW151226	30
11	Omicron time-frequency-SNR plot	36
12	Example of coincident noise	37
13	Example of a feedback loop	40
14	ETMY saturation	42
15	DAC glitches in PRCL	44
16	Vetoed DARM triggers from DAC calibration	45
17	Spectrograms of RF whistles	48
18	DARM glitch histograms with and without RF whistles	49
19	Vetoed whistles from Hveto	50
20	Median noise ASD in O1	52
21	Inspirational horizon distance in O1	53
22	PyCBC trigger rates in O1	54
23	H1 Y-end air compressor	56
24	Reflection at the IMC	60
25	Example of a PDH error signal	62

26	PDH response to asymmetric cavity motion	64
27	PDH response to symmetric cavity motion	65
28	Spectral comb in IMC control signal	66
29	LIGO length degrees of freedom	71
30	LSC ODC SIMULINK Model Example	73
31	LSC ODC bits example	74
32	LSC ODC Overview Screen	75
33	MICH control signal on summary pages	76
34	ASC ODC Overview Screen	78
35	ASC ODC Photodiode Monitor in MEDM	79
36	ODC threshold on MICH pitch	82
37	Rate histogram of PyCBC triggers vs. Gaussian noise	86
38	Histogram of BNS bin triggers	89
39	BNS bin histograms - GW150914 analysis	90
40	Bulk bin histograms - GW150914 analysis	91
41	Edge bin histograms - GW150914 analysis	93
42	BNS bin histograms - GW151226 analysis	97
43	Bulk bin histograms - GW151226 analysis	99
44	Edge bin histograms - GW151226 analysis	101
45	Re-weighted SNR histogram with SNR > 20 cut	103
46	Time-frequency plot blip transient	104
47	Time-series plot of blip transient	105
48	Total mass vs. effective spin in GW150914 analysis	107
49	Template duration vs. peak frequency in GW150914 analysis	108
50	Time-frequency plot of 60-200 Hz noise	109
51	Zoomed in time-frequency plot of 60-200 Hz noise	109

Well that was fun.

Preface

The work presented in this thesis stems from my participation in the LIGO Scientific Collaboration (LSC). This work does not reflect the scientific opinion of the LSC and it was not reviewed by the collaboration.

Acronyms

- ADC: Analog-to-Digital Converter.
- ASC: Alignment Sensing and Control.
- ASD: Amplitude Spectral Density.
- BBH: Binary Black Hole.
- BNS: Binary Neutron Star.
- CARM: Common Arm length degree of freedom.
- CBC: Compact Binary Coalescence.
- CW: Continuous Waves.
- DAC: Digital-to-Analog Converter.
- DARM: Differential Arm length degree of freedom.
- ETMX/Y: X- or Y-Arm End Test Mass.
- FAR: False Alarm Rate.
- FAP: False Alarm Probability.
- GW: Gravitational Wave.
- IMC: Input Mode Cleaner.
- ITMX/Y: X- or Y-Arm Input Test Mass.
- LSC: Length Sensing and Control.

- LVT: LIGO-VIRGO Trigger.
- MEDM: Motif Editor and Display Manager.
- MICH: Michelson degree of freedom.
- NSBH: Neutron Star-Black Hole binary.
- O1: The first Advanced LIGO Observing run.
- ODC: Online Detector Characterization.
- OMC: Output Mode Cleaner.
- PDH: Pound-Drever-Hall cavity control technique.
- PRCL: Power Recycling Cavity Length degree of freedom.
- PSD: Power Spectral Density.
- RF: Radio Frequency.
- SNR: Signal-to-Noise Ratio.
- SRCL: Signal Recycling Cavity Length degree of freedom.

Chapter 1

Introduction

In 1915, Albert Einstein published his General Theory of Relativity, a geometric theory of gravitation that sought to expand upon Newtonian mechanics and provide a complete description of gravity and its relationship with space and time. Einstein theorized that space and time were deeply related and existed together as a manifold called spacetime. Matter with energy and momentum existing in this manifold creates curvature in spacetime. Gravitational forces are the result of matter following geodesic curves in spacetime. This concept can be summarized in the Einstein field equation,

$$G_{\mu\nu} = 8\pi T_{\mu\nu} \quad (1.1)$$

where $G_{\mu\nu}$ is the Einstein tensor, which describes the curvature of spacetime, $T_{\mu\nu}$ is the stress-energy tensor, which describes the energy and momentum in spacetime, and $G = c = 1$. The Einstein tensor is defined as,

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} \quad (1.2)$$

where $R_{\mu\nu}$ is the Ricci curvature tensor and $g_{\mu\nu}$ is the metric tensor for the manifold.

An interesting result that arises in this theory is the existence of gravitational waves [1, 2], which are perturbations in spacetime caused by certain types of time-varying mass distributions. To describe gravitational waves, we consider a Minkowski metric with a small perturbation. The Minkowski metric is a flat spacetime metric

defined as

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1.3)$$

where $\mu = 0$ corresponds to the time coordinate and $\mu = 1, 2, 3$ correspond to the spatial coordinates. In examples, we will use the coordinate convention $(x^0, x^1, x^2, x^3) = (ct, x, y, z)$. The full spacetime metric, $g_{\mu\nu}$, is then constructed as a linear perturbation on the Minkowski metric,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.4)$$

where $h_{\mu\nu}$ is the metric perturbation and $|h_{\mu\nu}| \ll 1$.

To explore the effects of this perturbation, it is very useful to move into the transverse traceless gauge where coordinates on the manifold are defined by the geodesic motion of freely-falling test masses [3]. In this gauge, the weak field vacuum solution of the Einstein field equation becomes a wave equation:

$$\square h_{\mu\nu} = 0 \quad (1.5)$$

where \square is the d'Alembert operator,

$$\square = -\frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}. \quad (1.6)$$

The solutions to this differential equation will be plane waves of the form

$$h_{\mu\nu} = C_{\mu\nu} e^{i(2\pi ft - \vec{k} \cdot \vec{x})} \quad (1.7)$$

where $C_{\mu\nu}$ is the wave amplitude, f is the frequency, and \vec{k} is the wave vector which indicates the direction of propagation [4].

For example, consider the case of a gravitational wave propagating along the \hat{z} -axis. When the conditions of the transverse traceless gauge are applied, the resulting form of $h_{\mu\nu}$ is

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.8)$$

where the diagonal and off-diagonal terms represent two polarizations of the resulting gravitational wave, called “h-plus” and “h-cross” respectively. We can see the effects of this perturbation by observing the spacetime interval on the manifold. The spacetime interval is defined as

$$ds^2 = dx^\mu g_{\mu\nu} dx^\nu. \quad (1.9)$$

Using the coordinate convention of (ct, x, y, z) , the unperturbed spacetime interval is given as

$$ds^2 = -c^2 dt^2 + dx^2 + dy^2 + dz^2, \quad (1.10)$$

which is the standard line element in a Minkowski metric. Substituting in our perturbed metric for $g_{\mu\nu}$, we find that the spacetime interval can be broken up into a standard Minkowski line element and a perturbation due to $h_{\mu\nu}$.

$$ds^2 = dx^\mu (\eta_{\mu\nu} + h_{\mu\nu}) dx^\nu \quad (1.11)$$

$$ds^2 = dx^\mu \eta_{\mu\nu} dx^\nu + dx^\mu h_{\mu\nu} dx^\nu \quad (1.12)$$

As an example, we present the case of a plus-polarized gravitational wave propagating in the \hat{z} direction and observe the effect of the perturbation on the spacetime interval. The perturbation will have the form

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & 0 & 0 \\ 0 & 0 & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.13)$$

Since the perturbation is spatially transverse to the direction of propagation, the ct - and \hat{z} -coordinates will not be modulated by the gravitational wave. The \hat{x} - and \hat{y} -coordinates will be modulated according to equation 1.12. The resulting spacetime interval is

$$ds^2 = -c^2 dt^2 + (1 + h_+) dx^2 + (1 - h_+) dy^2 + dz^2. \quad (1.14)$$

This shows that a gravitational wave propagating along the \hat{z} -axis will differentially stretch and squeeze spacetime in the transverse axes. The exact form of h_+ will depend on the source of the gravitational waves. A visualization of this stretching and squeezing is shown in Figure 1 [5]. The cross polarization stretches and squeezes

at a 45 degree angle relative to the plus polarization. The strain, h , imparted by a gravitational wave is typically extremely small by the time it reaches Earth, producing relative length changes on the order of 10^{-22} .

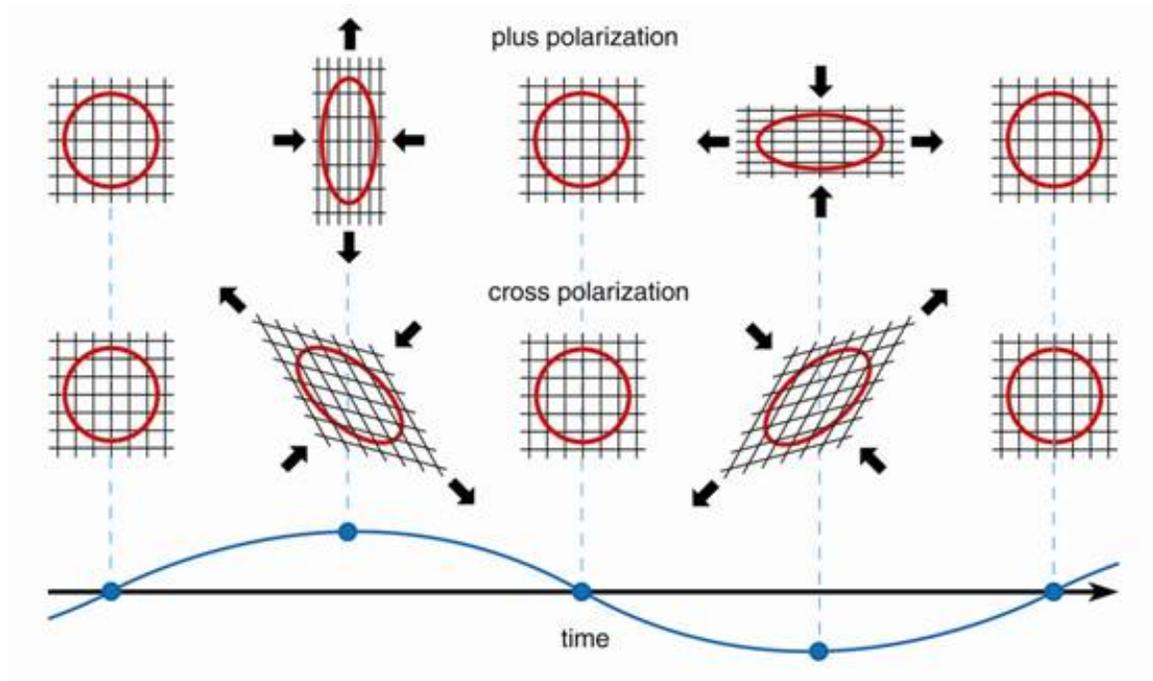


Figure 1: Plus and cross polarizations of a gravitational wave. As a gravitational wave propagates through spacetime, there is a stretching and squeezing effect that changes the relative length between points in spacetime. The strain produced by the cross polarization is at a 45 degree angle relative to the strain produced by the plus polarization.

The Advanced LIGO interferometers [6] are designed to be sensitive to this differential stretching and squeezing by constructing orthogonal optical cavities, referred to as the X- and Y-arms. When a gravitational wave passes through a LIGO interferometer, the length of the arms is modulated, causing the light to have a longer or shorter travel time as it traverses the optical cavities. Since gravitational waves expand space in one direction while the orthogonal direction contracts, the X- and Y-arms will experience differential changes in length. When light from the arms is recombined, there will be a difference in phase between the two beams as they have traveled different paths. The result of this phase mismatch is a change in optical power at the output

port of the interferometer. This output optical signal can be searched for evidence of interactions with gravitational waves. The layout and gravitational wave readout scheme of the interferometers is discussed below.

1.1 The Advanced LIGO Interferometers

The Advanced LIGO (aLIGO) interferometers are a pair of dual-recycled Michelson interferometers that employ 4km long Fabry-Perot cavities in their arms to increase the interaction time with a gravitational wave signal [6]. Figure 2 shows a simplified layout of an aLIGO interferometer.

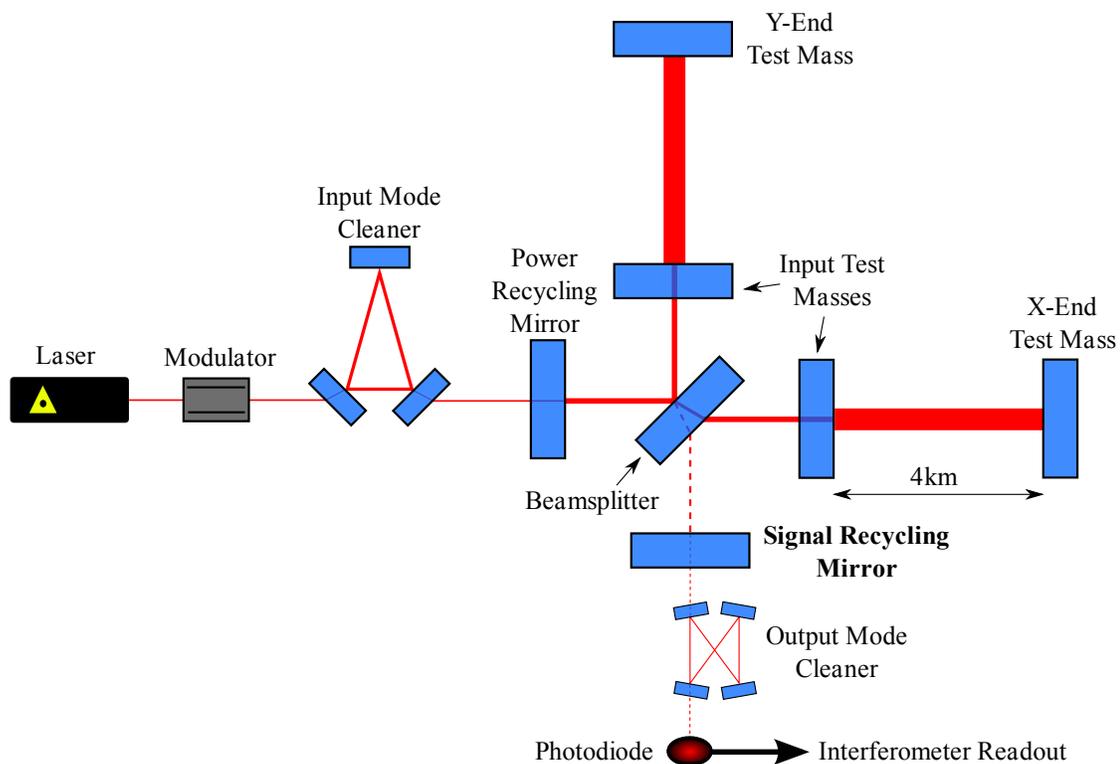


Figure 2: Layout of Advanced LIGO

At the input to an aLIGO interferometer is a solid-state Nd:YAG laser that provides laser light at a wavelength of 1064 nm. Not included in Figure 2 are frequency and intensity stabilization control loops designed to provide as stable a laser source as possible for the experiment. This stabilized laser is called the pre-stabilized laser

(PSL). The laser light is passed through a series of electro-optic modulators (EOM) where radio-frequency (RF) sidebands are generated and imparted onto the light. These RF sidebands are used to sense auxiliary optical degrees of freedom in the interferometer. The beam is then passed through the input mode cleaner (IMC), which rejects higher order spatial modes of the beam and transmits a circular TEM00 mode to be used in the instrument.

Once the beam has been stabilized in frequency and intensity and the higher order optical modes have been stripped away, it is transmitted through the power recycling mirror and enters the vertex of the interferometer. In the vertex, the beam is split 50/50 by the beamsplitter. Half of the light is directed toward the input test mass (ITM) of the X-arm and half of the light is directed toward the ITM of the Y-arm. As mentioned previously, the aLIGO arms are not single bounce cavities; they are comprised of Fabry-Perot cavities that allow the light to circulate in the arm cavities multiple times. The light is stored in the arm cavities for ~ 1 ms, trapped between the highly reflective surfaces of the ITM and the end test mass (ETM), before it is transmitted back through the ITM and into the vertex.

As mentioned above, it is the interaction of the arm cavities with gravitational waves that allows the optical field to be imparted with a gravitational wave signal. The increased light storage time provided by Fabry-Perot cavities increases the interaction between the optical field and a gravitational wave by increasing the effective optical length of the arm cavities. An incident gravitational wave differentially modulates the arm cavities, resulting in a difference in path length for the beams traveling in each arm.

When light from the arms is recombined at the beamsplitter, there will be a difference in phase between the two beams if they have traveled different paths. The resulting light from this recombination of phase shifted beams will be divided based on how much of a phase offset was accumulated as the beams traversed the arms. In the absence of a gravitational wave, most of the light will be directed back toward the power recycling mirror. This is called the symmetric port of the interferometer. A small amount of light is directed toward the signal recycling cavity. This is called the antisymmetric port of the interferometer. The optical power at these ports will fluctuate in the presence of a gravitational wave. It is the antisymmetric port optical field that is used to search for gravitational wave signals.

At the symmetric port, the beam will be sent back toward the power recycling mirror. The power recycling mirror forms a resonant cavity with the ITMs, allowing for light at the symmetric port of the beamsplitter to be added coherently to incoming light from the PSL and increasing the effective power in the vertex. This increase in effective power is known as the power recycling gain [6].

At the antisymmetric port, the beam is sent toward the signal recycling mirror. The signal recycling cavity is used to increase the sensitivity of the interferometer in a band of frequencies by adjusting the effective finesse of the coupled cavity formed by the signal recycling cavity and the arm cavities [6]. If the light returning from the arms has accumulated some differential amount of phase as it traveled along the arms, perhaps from a gravitational wave modulating the length of each arm differentially, it will be transmitted through the signal recycling cavity and into the output mode cleaner (OMC). The OMC behaves similarly to the IMC, stripping away higher order optical modes and isolating the TEM00 mode of the beam. The transmitted, mode cleaned signal is then read out using a homodyne detection scheme on a DC photodiode.

1.1.1 DC Readout

When a gravitational wave modulates the length of an arm cavity, the light traveling in that arm experiences a phase modulation. This phase modulation can be visualized by picturing the beam in frequency space. In Figure 3, the carrier beam frequency is designated as f_0 . The phase modulation due to a gravitational wave signal introduces a frequency sideband at the gravitational wave frequency, which is in the kHz range for signals that LIGO is sensitive to. The RF sidebands used for auxiliary optical cavity control are offset from the carrier frequency by 9, 24, and 45 MHz. In a heterodyne detection scheme, the interferometer would operate at the 'dark fringe', meaning that the output port would not transmit light until there was differential arm motion. In this scheme, the RF sidebands would be detected on the same photodetector as the gravitational wave sidebands. The RF sidebands would be used to demodulate the photodetector signal, leaving behind a gravitational wave signal. This is the method that was used in initial LIGO.

In Advanced LIGO, a homodyne detection, or 'DC Readout', scheme is employed [7]. In a homodyne detection scheme, the RF sidebands are not used to extract the

gravitational wave signal. Instead, the carrier beam itself is used. Instead of aligning the instrument on the dark fringe, a differential offset is introduced to the arm cavities to allow a small amount of light into the output port. The RF sidebands, which if not used for demodulation would only contribute noise to the output signal, are rejected by the output mode cleaner (OMC). The gravitational wave sidebands, however, are at a low enough frequency offset that they are within the pass band of the OMC and are allowed to transmit through the cavity. The pass band of the OMC is defined by the pole frequency of the optical cavity.

Since the OMC DC photodiode measures power, it measures the square of the incident optical field and witnesses beat frequencies between different components of the light. If the RF sidebands have been filtered out by the OMC, the only remaining beat note will be that of the carrier beam (f_0) beating against the gravitational wave sideband ($f_0 + f_{GW}$). This beat note will appear as the difference in frequency between the two optical fields, leaving behind a signal in the 30-2000 Hz range (f_{GW}) and providing a natural demodulation inherent to the measurement process. The process of recovering the gravitational wave sideband using the carrier field as a reference is known as homodyne detection. The advantage in this method lies in the fact that the carrier beam has been passed through the arm cavities. The cavities act as a low pass filter and remove high frequency noise relative to the carrier beam frequency. The RF sidebands are not filtered by the arm cavities and are quite noisy in comparison. This noise is propagated forward when they are used for demodulation.

1.1.2 The aLIGO Noise Curve

The LIGO interferometers are among the most precise measuring devices that have ever been constructed, sensitive to a differential displacement on the order of 10^{-19} m/ $\sqrt{\text{Hz}}$ at their most sensitive frequencies. When operating at such small length scales, the interferometers are susceptible to a number of very subtle noise sources. Figure 4 shows the limiting noise sources for the LIGO interferometers [8]. Each curve represents a known source of noise in the interferometer output.

At frequencies below 10 Hz, the limiting noise at design sensitivity comes a combination of seismic noise and suspension thermal noise. Seismic noise is caused by vibrations coupling from the Earth onto the mechanical structures supporting the aLIGO optics. Seismic noise is attenuated using a multi-tiered active feedback seismic

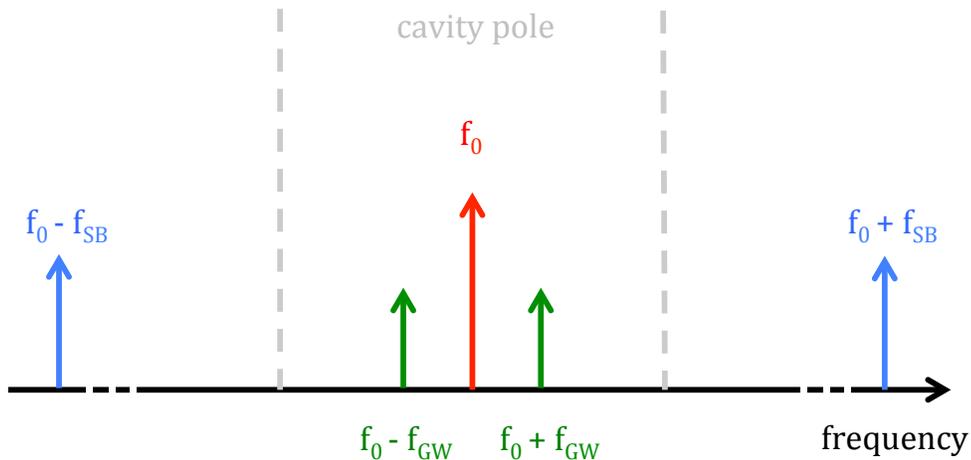


Figure 3: Frequency domain visualization of beam at OMC. Grey dotted lines indicate the cavity pole. The gravitational wave sidebands are within the cavity pole and are transmitted through the OMC. The RF sidebands are in the MHz range and are rejected by the OMC.

isolation system [9, 10]. At higher frequencies, any residual seismic noise is passively attenuated by the suspension systems, which use multiple stages of pendula to reduce displacement noise from the suspension point to the optics [11]. Suspension thermal noise is dominated by the mechanical losses of the fused silica fibers used to suspend the test masses. Since the interferometers have not yet been fully commissioned, the current limiting noise below 10 Hz is driven by noise coupling from auxiliary degrees of freedom [8].

From 30 Hz onwards, there will eventually be two dominant noise sources: quantum noise (the blue curve in Figure 4) and coating Brownian noise (the green curve in Figure 4). Quantum noise is the combination of two sources. The first is shot noise, which is a photon counting noise when light is measured on a photodiode. Shot noise is the dominant noise source above ~ 300 Hz and can be further improved by increasing laser power. The second is radiation pressure noise, which is a fluctuating force on the test masses based on fluctuations in photon number in the cavity. Radiation pressure noise will increase with increasing laser power as a higher photon number implies a higher uncertainty in the momentum imparted onto the test mass optics.

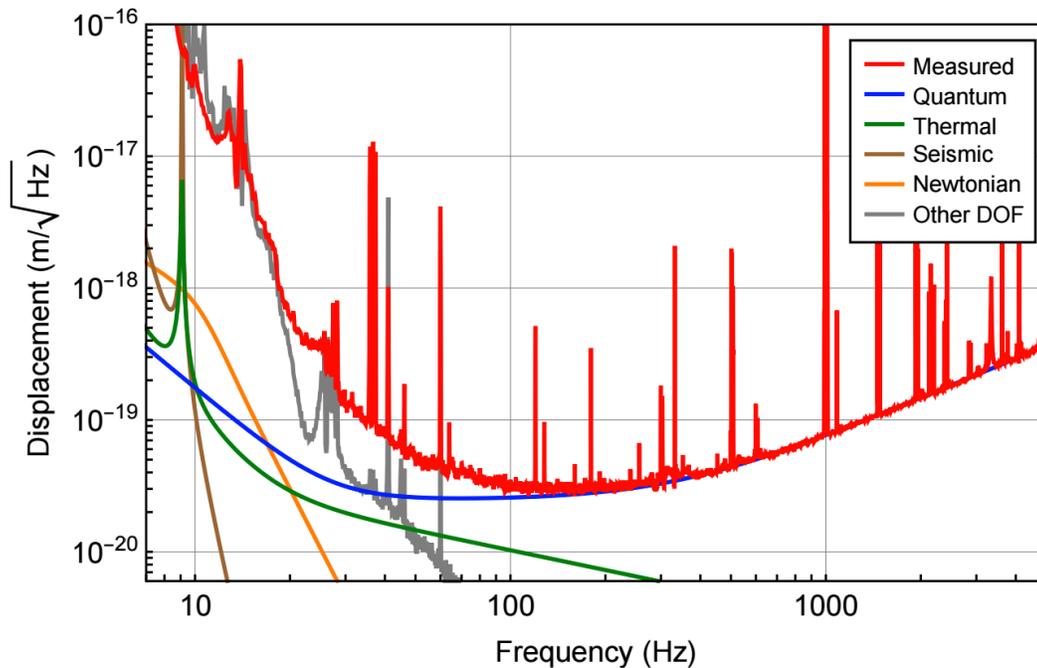


Figure 4: Advanced LIGO noise curve during the first observing run with understood noise sources, reproduced from [8]. The red curve is the measured instrumental noise at Hanford during the first observing run. The blue curve is quantum noise, which is a combination of photon shot noise at the output photodetector and radiation pressure noise on the test mass optics. The green curve is the modeled thermal noise, which is a combination of thermal noise from the suspensions, optics, and optical coatings. The brown curve is seismic noise coupling into the optics, which is highly attenuated at high frequencies. The orange curve is Newtonian noise, which is driven by perturbations in the density of the ground. The grey curve labeled 'Other DOF' is the sum of the noise coupling from auxiliary optics into the output of the interferometer, mainly driven by optical misalignments.

Coating Brownian noise is due to thermally driven mechanical losses in optical coatings. Figure 4 shows that the interferometer noise is limited by the combination of thermal noise and quantum noise from 100 Hz onward.

1.2 Sources of Gravitational Waves

Gravitational wave signals have highly varying characteristics depending on the source of the gravitational waves. The gravitational wave strain produced depends on the distribution of mass-energy at the source and is given as

$$h_{ij} = \frac{2}{r} \ddot{I}_{ij} \quad (1.15)$$

where \ddot{I}_{ij} is the quadropole tensor describing the mass-energy distribution and $c = G = 1$. This equation tells us that a source distribution requires an accelerating quadropole moment to generate gravitational waves. Depending on the dynamics of the source system, the resulting gravitational wave signals will vary greatly in both duration and morphology.

There are a number of potential sources of gravitational waves that are searched for in aLIGO data. Astrophysical searches include both modeled and unmodeled searches for both transient and continuous signals. Table 1 gives an example of each of these sources. Each of these categories are discussed in the following sections.

	Transient	Continuous
Modeled	Compact binary coalescences	Rotating neutron stars
Unmodeled	Core-collapse supernovae	Stochastic GW background

Table 1: Table describing sources of gravitational waves.

1.2.1 Compact binary coalescences

Compact binary coalescences (CBC) are a primary search target of the Advanced LIGO interferometers. These signals are the result of two compact objects, such as neutron stars or black holes, orbiting each other in a binary system. The two objects will lose orbital energy as they orbit around each other and deform spacetime, generating gravitational waves. As the binary system loses energy, the orbit decays until they merge and coalesce into one final compact object. The orbital frequency of such systems increases monotonically as the orbit decays, resulting in a gravitational wave signal that sweeps upwards in frequency known as a 'chirp' [12]. Since the gravitational wave signal from such a system is known, this model is incorporated

into the search algorithm and the search is referred to as a modeled search. CBC waveforms will have a duration from ~ 0.1 -60s in the frequency range that aLIGO is sensitive to and as such they are considered transient signals. Searches for CBC signals are discussed in Section 2.

Of particular interest to aLIGO are binary neutron star systems, which are known to exist from astronomical observation. The 1993 Nobel Prize in Physics was awarded to Hulse and Taylor for the indirect detection of gravitational waves from a binary neutron star system [13, 14]. The orbital period of the Hulse-Taylor binary neutron star system has been measured since 1974. The decay of its orbital period matches the expected orbital decay based on energy loss due to the emission of gravitational waves [15].

Beyond binary neutron stars, the search for compact binary coalescences is expanded to search for binary black hole (BBH) systems and neutron star-black hole (NSBH) systems. The discovery of gravitational waves from binary black holes was accomplished in the first observing run with the discovery of two binary black hole systems, GW150914 and GW151226 [16]. Further discussion of the results from the first observing run is presented in Chapter 3.

1.2.2 Burst signals

A predicted population of signals is from unmodeled gravitational wave transients, or 'bursts'. These signals can come from core-collapse supernovae, cosmic string cusps, and binary black hole mergers [17, 18, 19]. The Advanced LIGO burst search is carried out by a number of search pipelines [20, 21, 22]. In a burst search, there are two primary methods for signal identification. In a standard burst search, potential signals are identified on a single detector basis using an excess power ranking statistic and then checked for coincidence between the two interferometers. In a coherent burst search, the data streams from both interferometers are combined into one coherent ranking statistic based on a maximum likelihood analysis. Each analysis uses their own methods to distinguish potential signals from noise.

1.2.3 Continuous waves

Continuous sources of gravitational waves are those that are constantly emitting gravitational waves. The primary expected source of continuous gravitational waves are rotating neutron stars that have some asymmetry with respect to their axis of rotation. If a neutron star has a mountain on it, its rotation will produce a time-changing quadrupole moment and constantly generate periodic gravitational waves. These waves are not expected to be as loud as the gravitational waves generated from more violent, transient events. As such, the strategy to discover them is different than that of a transient search.

To search for gravitational waves from continuous sources, the data are transformed into the frequency domain and integrated for long periods of time [23, 24]. If there is a constant, periodic gravitational wave signal in the data, it will manifest as a peak in the frequency spectrum of the data. This peak will accumulate signal and grow over the integration period relative to the noise floor.

1.2.4 Stochastic background

In addition to sources that can be directly detected in the data, there is also a search designed to discover a stochastic background of gravitational waves [25, 26]. This stochastic background is the superposition of gravitational wave signals that are too weak to be detected directly, such as distant compact binaries and supernovae. The stochastic background is a statistical background based on the rate and distribution of gravitational wave sources in the universe. To resolve the stochastic background, the data from the two interferometers are cross-correlated over long periods of time. Since the cross-correlation is an integral in the time domain, the effects of quiet, correlated gravitational wave signals are accumulated over time until a statistically significant signal-to-noise ratio can be quoted [27]. This is an interesting search from a detector characterization point of view, as it requires an understanding of correlated noise sources between the two interferometers.

1.3 The Advanced Detector Network

The Advanced Laser Interferometer Gravitational-Wave Observatory is part of a worldwide effort to detect gravitational waves from astrophysical sources. The two LIGO interferometers, one in Hanford, WA and one in Livingston, LA, are part of a growing network of ground-based interferometric gravitational wave detectors. The LIGO detectors are currently the largest and most sensitive interferometric gravitational wave detectors in the network.

There are a number of other interferometric gravitational wave detectors being built and commissioned for future use in collaboration with LIGO. The Advanced VIRGO detector is being built and commissioned in Cascina, Italy and will be joining LIGO in observing runs soon [28]. The VIRGO interferometer will provide enough sensitivity to aid in detection and triangulation of astrophysical sources.

The GEO600 detector [29], located in Hanover, Germany is an interferometer built in collaboration between Germany and the United Kingdom. GEO600 is an extremely valuable test bed for interferometric technologies, including quantum optics [30] and homodyne detection [7]. However, with 600m arms, GEO600 is unlikely to be sensitive enough to witness expected astrophysical sources.

The KAGRA detector [31], located underground in the Kamioka mine in Japan, is in its commissioning phase. KAGRA has 3 km long arms and, unlike other gravitational wave interferometers, employs cryogenics to reduce thermal noise in its optics. When complete, KAGRA should be sensitive enough to contribute to the worldwide detector network.

A third LIGO interferometer, IndIGO [32], is in the planning stages and will be constructed in India. The position and sensitivity of IndIGO will allow for confident triangulation of astrophysical sources of gravitational waves.

Chapter 2

Searching for Compact Binary Coalescences

Since a significant portion of this thesis uses the performance of CBC searches as a metric for the quality of the data, a more thorough discussion of how a CBC search works is necessary. This thesis will focus on the output of the PyCBC pipeline, which is a Python-based software packaged used to search for gravitational waves from compact binary coalescences [33, 34].

CBC search pipelines are designed to search for gravitational wave transients from compact binary coalescences [33]. The signals expected to be measured in the LIGO interferometers are extremely quiet, with gravitational wave strains on the order of 10^{-22} . On these scales, most signals will not be able to be extracted from the background noise with simple filtering. Figure 5 shows the gravitational wave strain from a $1.4\text{-}1.4M_{\odot}$ binary neutron star system at a distance of 20 Mpc overlaid on real detector noise from the Livingston interferometer. The signal has a peak strain roughly two orders of magnitude lower than the peak strain of the detector noise. For this reason, the CBC searches employ a matched filter algorithm, which correlates expected CBC waveforms with detector data and assigns a ranking statistic, the signal-to-noise ratio (SNR), to every event that it finds.

2.0.1 The matched filter

The matched filter calculates the correlation of the detector data with expected CBC waveforms in the frequency domain. The detector data and expected waveform are

multiplied together and their product is divided by the background noise in the detector. The fundamental operation of the matched filter is defined as an inner product of the detector data and the CBC waveform [35],

$$(s|h)(t) = 4\text{Re} \int_{f_{\text{low}}}^{f_{\text{high}}} \frac{\tilde{s}(f)\tilde{h}^*(f)}{S_n(f)} e^{2\pi i f t} df, \quad (2.1)$$

where \tilde{s} is the Fourier transformed detector data, \tilde{h} is the Fourier transformed gravitational waveform, and $S_n(f)$ is the power spectral density of the detector data averaged over 2048 seconds, which represents the average noise in the detector. The maximum bounds of the integral are set to span the frequency space for which the interferometers are sensitive enough to detect gravitational waves, typically 30 - 2000 Hz. However, if a waveform merges at a lower frequency, the upper bound on the integral can be set to capture the frequencies where the signal has power and nothing higher.

As described in Equation 1.8, a gravitational wave is comprised of two polarizations, “h-plus” and “h-cross”. When the data are searched for a CBC signal, a waveform is generated for each polarization and the matched filter is computed separately for each polarization. The SNR of a CBC waveform at any given time is defined as the weighted quadrature sum of the SNR measured for each polarization [35],

$$\rho^2(t) = \frac{(s|h_p)^2 + (s|h_c)^2}{(h_p|h_p)}, \quad (2.2)$$

where h_p and h_c are the plus and cross polarizations of the modeled gravitational waveform respectively and s is the detector data. When the SNR time-series defined in equation 2.2 crosses a certain threshold, the waveform is considered to have significant overlap with the detector data and an event is generated at the time of the SNR peak. These events are called “triggers” and are used to generate populations of potential gravitational wave events for analysis.

2.0.2 Waveform templates

To perform a search, the matched filter algorithm needs to know what to search for. A collection of expected CBC waveforms is generated using the formalism of general relativity before the analysis [36, 37]. Each of the expected waveforms is called a template and the full collection of waveforms is referred to as the template

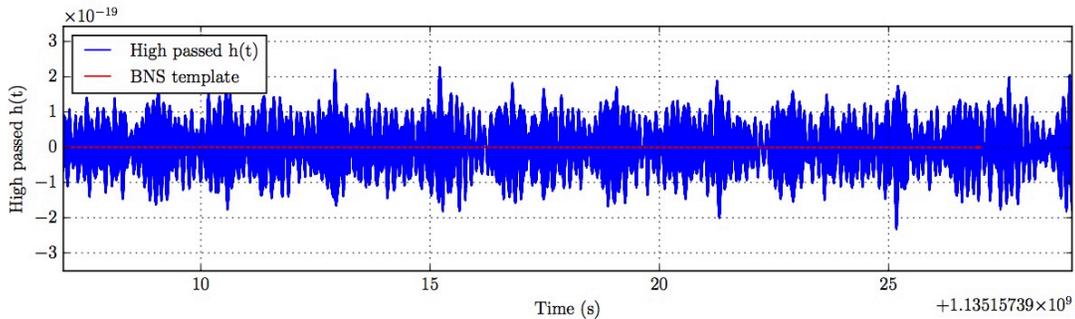


Figure 5: A simulated gravitational wave signal from a binary neutron star system overlaid on real detector noise from L1. The blue curve, labeled $h(t)$, represents the detector data. The red curve represents the gravitational wave strain expected from a $1.4\text{-}1.4M_{\odot}$ binary neutron star system at 20 Mpc. The peak strain of the binary neutron star waveform is 8×10^{-22} . The detector data have been high pass filtered with a corner frequency at 20 Hz and show a peak strain of 2.2×10^{-19} . The signal is buried in the detector noise and requires a matched filter algorithm to be recovered. At this time, the inspiral range for a $1.4\text{-}1.4M_{\odot}$ BNS system was 60 Mpc, indicating that the same system originating at 60 Mpc would be recovered with $\text{SNR} = 8$.

bank. This template bank is constructed to span the astrophysical parameter space included in the search [12]. This parameter space is constrained by the noise spectrum of the interferometers. As shown in Figure 4, the LIGO interferometers are sensitive enough to detect gravitational waves in the region from roughly 30 - 2000 Hz. This rules out detection of sources that are expected to coalesce at very low frequencies, such as supermassive black hole binaries [38] and binary white dwarf systems [39]. The template bank used in Advanced LIGO's first observing run consisted waveforms representing binary neutron stars, binary black holes, and neutron star-black hole binary systems [12]. The total masses of these systems ranged from $2\text{-}100M_{\odot}$. This reflects the set of systems that will have merger frequencies above 30 Hz and will have detectable power in LIGO's sensitive bandwidth.

Each waveform is defined by the mass and spin of each compact object in the binary system. It is convenient to combine the component masses into a new variable, chirp mass, which is used to parameterize gravitational wave signals in general

relativity. Chirp mass is defined as [40]

$$M_{chirp} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (2.3)$$

where the m_i are the component masses of the compact objects in the binary system. Each compact object in the binary system has its rotation represented by a dimensionless spin parameter,

$$\chi_i = \frac{c S_i}{G m_i^2}, \quad (2.4)$$

where S_i is the spin angular momentum of the compact object, m_i is the mass of the compact object, c is the speed of light, and G is the gravitational constant. It is convenient to combine the effects of each object's spin into one parameter called effective spin, χ_{eff} , which is the mass-weighted spin of the system [41]. χ_{eff} is defined as

$$\chi_{eff} = \frac{\chi_1 m_1 + \chi_2 m_2}{m_1 + m_2} \quad (2.5)$$

where the χ_i are the dimensionless spin parameters [42] and the m_i are the masses for each compact object in the binary system.

2.0.3 χ^2 signal consistency test

If the data produced by the interferometers were Gaussian, the matched filter would be sufficient for running a search pipeline and recovering gravitational wave signals. Unfortunately, the data are non-Gaussian, containing noise transients of varying durations [43, 44]. These noise transients, or “glitches”, can have significant amplitude and, when multiplied with a waveform template in the matched filter, can cause loud triggers to be generated. However, a significant advantage of performing a modeled search for gravitational waves is that we know we’re looking for. With this information, the SNR can be refined into a more robust ranking statistic for significant events in the data. This is done using the χ^2 signal consistency test [45].

The SNR produced by the matched filter is an integral in the frequency domain which reports the total accumulated SNR over a given bandwidth. If a noise transient has significant amplitude, it can generate a high SNR trigger by overlapping with the waveform template in the matched filter. However, these noise transients typically have a duration on the order of 0.1s. This type of transient is easily distinguished from a chirp signal that increases monotonically in frequency over the span of many

seconds. Figure 6 shows an example of such a noise transient with the waveform of a binary black hole system overlaid on top of it. Although the peaks of each time-series are aligned, the noise transient clearly has a more localized power distribution that does not match that of the CBC waveform.

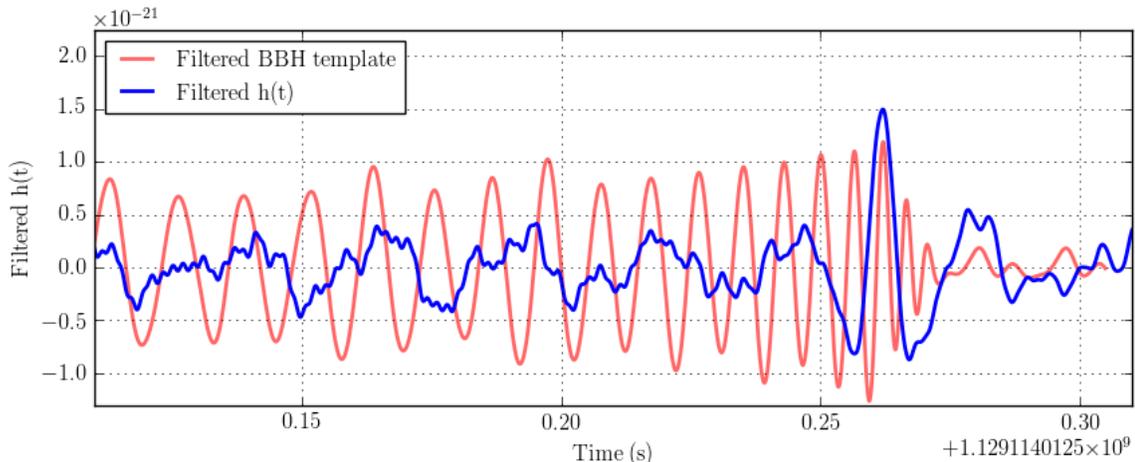


Figure 6: A time domain representation of a noise transient in the L1 detector with a binary black hole waveform overlaid on top of it. Both time-series have been filtered to isolate the frequencies where the glitch and the waveform have significant power. The peak amplitude of the binary black hole waveform lines up with the peak amplitude of the noise transient in the detector data, $h(t)$. However, the binary black hole waveform contains many more cycles than the noise transient. It is clear that this transient does not have the same distribution of power in time and frequency as the CBC signal.

The χ^2 test divides each CBC waveform into frequency bins of equal power, checking that the SNR is distributed as a function of frequency as expected from an actual CBC signal. For a signal divided into p frequency bins, each bin should contain $\frac{1}{p}$ of the power in the signal. In the χ^2 calculation, the SNR is calculated for each frequency bin and compared to the expected amount. The χ^2 statistic is calculated as [33]

$$\chi^2 = p \sum_{l=1}^p \left[\left(\frac{\rho_p^2}{p} - \rho_{p,l}^2 \right)^2 + \left(\frac{\rho_c^2}{p} - \rho_{c,l}^2 \right)^2 \right], \quad (2.6)$$

where ρ_p^2 is the SNR of the plus polarization of the waveform, ρ_c^2 is the SNR of the cross polarization of the waveform, p is the number of frequency bins, and $\rho_{i,l}^2$ is the

calculated SNR for the l^{th} frequency bin. The χ^2 statistic is then normalized such that a real signal will be reported with a value of 1. This normalized χ^2 is called the reduced χ^2 and is denoted by χ_r^2 .

In the PyCBC search, each trigger that comes out of the matched filter search is weighted based on the results of the χ^2 test. This is folded into a new ranking statistic for CBC triggers, which is called re-weighted SNR and is denoted by $\hat{\rho}$. The re-weighted SNR is calculated as [33]

$$\hat{\rho} = \begin{cases} \rho / [(1 + (\chi_r^2)^3)/2]^{\frac{1}{6}}, & \text{if } \chi_r^2 > 1, \\ \rho, & \text{if } \chi_r^2 \leq 1, \end{cases} \quad (2.7)$$

where ρ is the measured SNR and χ_r^2 is the reduced χ^2 . It is important to note that if a real signal has a power distribution that matches the template waveform, it will not be down-weighted by the χ^2 test.

This test is extremely powerful, as shown in Figure 7, which shows the distribution of single detector PyCBC triggers generated from September 12 to October 20, 2015. Figure 7a shows the distribution of triggers in SNR. The extensive tail of triggers with high SNR, which is generated when high amplitude noise transients are processed by the matched filter, extends beyond SNR 100. These high SNR triggers are down-weighted in the re-weighted SNR distribution, leaving behind a tail that extends to $\hat{\rho} \approx 11.5$ as seen in Figure 7b. Keeping in mind that a real signal will be reported at the same value in each plot, the re-weighting of triggers has lowered the noise floor, allowing for signals with SNR > 11.5 to stand out as the loudest events in their respective interferometers rather than being buried beneath a population of high SNR triggers.

The remaining tail of re-weighted SNR triggers represents the loudest background triggers in the CBC search. Investigating this set of loudest background triggers guides data quality efforts in defining the current limiting noise sources to the CBC search. This process is detailed in Chapter 4.

2.0.4 Searching for signals

The matched filter algorithm is run separately on each interferometer's data using the same bank of template waveforms. The output of the matched filter, the SNR time-series, is scanned for peaks and a set of single interferometer triggers is generated. The

two sets of single interferometer triggers are then compared to search for any events that were recorded within 15ms of each other with the same mass and spin parameters. Gravitational waves are predicted by general relativity to travel at the speed of light. The light travel time between the two interferometers is 10 ms and 5ms is added to the coincidence window due to uncertainty in signal arrival time [12]. As such, any triggers that are found within 15 ms of each other and are recovered with the same source parameters are considered to be coincident between the two interferometers. These coincident triggers represent potential gravitational wave signals and are referred to as candidate events.

The ranking statistic for coincident events in the PyCBC search is the network re-weighted SNR, $\hat{\rho}_c$, which is the quadrature sum of the re-weighted SNR from each interferometer.

$$\hat{\rho}_c = \sqrt{\hat{\rho}_L^2 + \hat{\rho}_H^2} \quad (2.8)$$

where $\hat{\rho}_L$ is the re-weighted SNR measured in the Livingston detector and $\hat{\rho}_H$ is the re-weighted SNR measured in the Hanford detector.

Most of these foreground events will be chance coincidences between noise in each interferometer, which is expected given the number of events in each data set. A large number of foreground events will be generated due to the detector noise, but ideally the distribution of foreground events will fall off sharply in $\hat{\rho}_c$ as shown in Figure 9, allowing for genuine signals to be recognized. To understand how stastically significant a foreground event is, a background distribution must be generated to calculate how often such a signal will be produced based on instrumental noise.

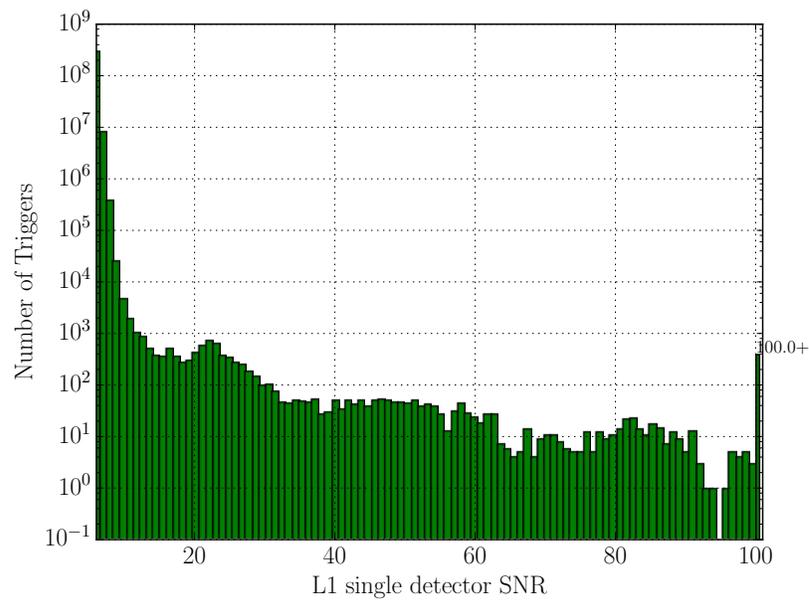
To generate the background distribution, we return to the set of single interferometer triggers that were generated when the matched filter was initially run. Since we want to understand the distribution of triggers based on detector noise, all of the triggers that were found to be coincident between the two interferometers are removed from the data sets, effectively removing all potential gravitational wave signals. The remaining triggers are then due to fluctuations in the background noise in each interferometer. These two sets of triggers, one from each interferometer, are then time shifted by a duration longer than the light travel time between the interferometers. Since gravitational waves are predicted to travel at the speed of light, this time shift ensures that the two sets of triggers are astrophysically uncorrelated

and do not contain any gravitational wave signals. The coincidence test is then performed again with the time shifted triggers, resulting in a coincident trigger set which represents background noise. A full distribution of background triggers is generated by performing this timeslide technique every 0.1 seconds and iterating over all of the data used in the search [33]. This distribution tells us how often we should expect to see coincident triggers with the same waveform parameters at a given value of SNR.

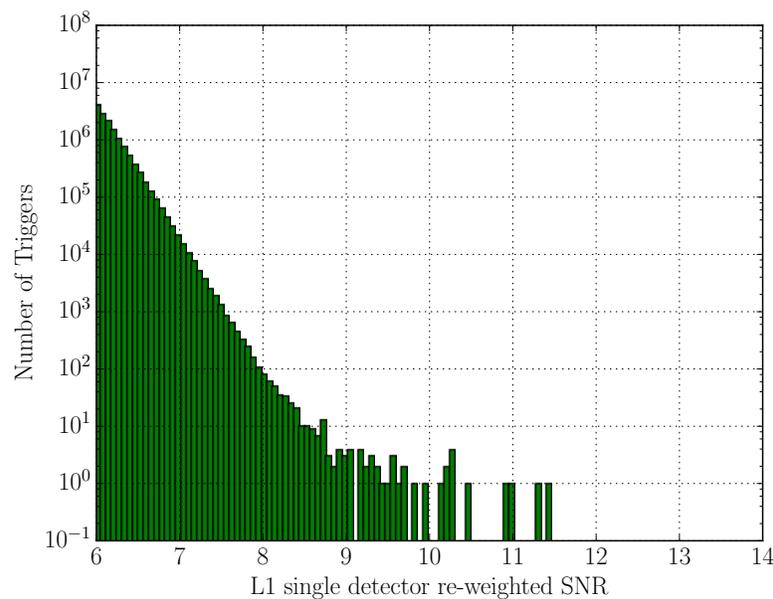
The statistical significance of any candidate gravitational wave is evaluated by calculating the rate of background events from detector noise that are at least as loud as the candidate event [12]. The results of the first observing run and details on the significance of foreground events is presented in Chapter 3. Any loud triggers that appear as the result of instrumental transients will contribute to the tail of the background distribution and influence the statistical significance of a recovered foreground event. The process of performing data quality investigations and data validation are detailed in Chapter 4

2.0.5 Gating

The PyCBC search includes a data conditioning stage that applies preventative cuts to remove large transients from the input data stream. This is done in a process called gating [33], which uses a window function to remove times containing large transients from the input data stream. This windowing function smoothly rolls the problematic section of data to zero, excising the large transient. The gating process is tuned by modifying the selection criteria for transients to be removed and by adjusting the time window to remove around each transient. Chapters 8 and 9 contain details about the gating thresholds used for the analyses containing GW150914 and GW151226 respectively.



(a)



(b)

Figure 7: Histograms of single interferometer PyCBC triggers from the Livingston (L1) interferometer. These triggers were generated from September 12 to October 20, 2015. These histograms contain triggers from the entire template bank, but exclude any triggers found in coincidence between the two interferometers. (7a) A histogram of single interferometer triggers in SNR. The tail of this distribution extends beyond $\text{SNR} = 100$. (7b) A histogram of single interferometer triggers in re-weighted SNR. The chi-squared test down-weights the long tail of SNR triggers in the re-weighted SNR distribution. Note that the x-axis has different limits in each plot.

Chapter 3

The First Observing Run

3.1 The First Direct Detection of Gravitational Waves

Advanced LIGO's first observing run (O1) lasted from September 12, 2015 - January 19, 2016. In this observing run, the first direct detection of gravitational waves was achieved with the discovery of two binary black hole mergers, GW150914 and GW151226 [46, 47]. In total, 51.5 days of coincident analysis data were recorded in O1. After data with excess noise were removed from the analysis, the total amount of coincident data was 49.8 days.

Along with the publication detailing the first direct detection of gravitational waves, several companion papers were released that provide a complete description of the O1 analyses and the state of the interferometers during the run [18, 12, 48, 49, 50, 51, 26, 52, 44, 8, 53, 54, 55].

3.2 Foreground Events

3.2.1 GW150914

The first signal discovered in O1, GW150914, marked the first direct detection of gravitational waves [46]. Figure 8 shows a filtered time domain representation of the first detection, GW150914, with the best estimated waveform overlaid on top. Both the signal and the waveform have been bandpass filtered to isolate the frequency range where the signal has power. Notch filters were used to remove noise sources with a static frequency, such as the 60 Hz power line frequency and interferometer

calibration signals. The signal demonstrates the characteristic “chirp”, increasing in frequency and amplitude as a function of time as expected from a compact binary coalescence.

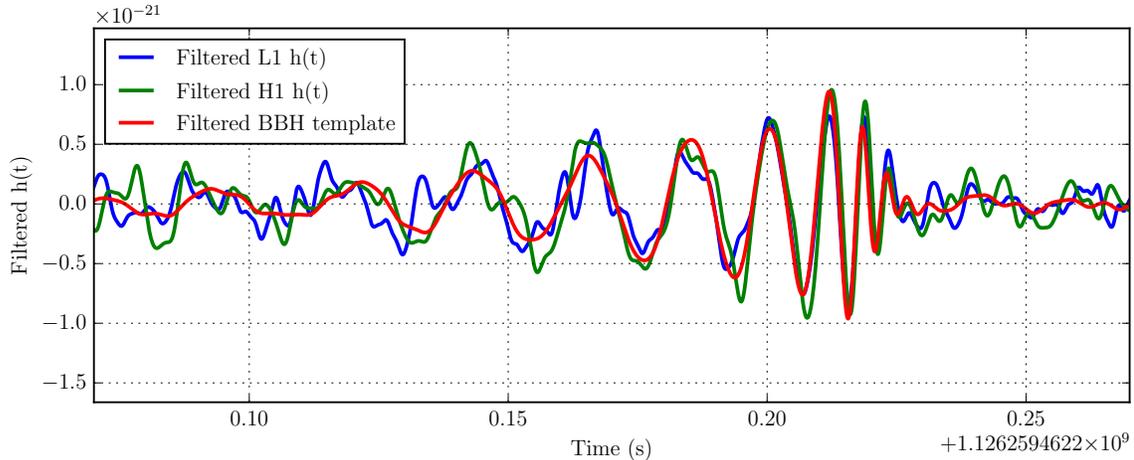


Figure 8: Time domain representation of H1 and L1 gravitational wave strain at the time of GW150914. The blue and green curves are detector strain, $h(t)$, zero-phase bandpass filtered to isolate the frequencies that contain signal. The red curve is a CBC waveform generated using the best estimated parameters. The CBC waveform has been filtered in the same way as the strain curve. The overlap between the three curves is significant, demonstrating many cycles of clear coherence and demonstrating the expected ‘chirp’ signal.

It is exceptional that GW150914 is visible in the detector data with such simple filtering. Due to the high total mass of the system, which is detailed in Table 2, the black holes of GW150914 coalesced quickly and at a low frequency, spending about 0.2 seconds in the frequency range that aLIGO is sensitive to. The signal was also tremendously loud due to its high total mass and relatively close distance. As a result, the power in the signal is highly localized in time, producing a short, loud waveform that is readily visualized. (A discussion of the data validation process relevant to GW150914 is found in Section 4.5.)

3.2.2 GW151226

The other binary black hole signal discovered in O1, GW151226, has a rather different morphology. The system that produced GW151226 was roughly three times less massive than that of GW150914 and merged at a similar distance (see Table 2). Due to the lower total mass, GW151226 has an overall lower amplitude than GW150914 and has its power distributed more broadly in time. GW151226 spent about 2 seconds in the frequency band that LIGO is sensitive to, which is a factor of 10 longer than the duration of GW150914. For these reasons, it is not feasible to generate a time domain visualization of the signal. Thus, this is a dramatic demonstration of the value of a matched-filter search for CBC signals, which is designed to identify modeled signals buried in noise.

3.2.3 LVT151012

The third loudest foreground event in the analysis, LVT151012, stands out from the background distribution but is not statistically significant enough to be labeled as a gravitational wave detection [16]. Its statistical significance is calculated to be just under 2σ . While it is not being claimed as a gravitational wave detection, there is no obvious reason to believe that it is a noise artifact based on detector performance. It is possible that LVT151012 is part of a larger population of gravitational waves that is expected to contain quiet, threshold signals as well as clear detections.

3.3 CBC Results

The search for compact binary coalescences was performed by two search pipelines: PyCBC [34, 33] and GstLAL [56]. In addition to these searches for modeled sources, an unmodeled burst search, Coherent Wave Burst (CWB), was run to search for coherent transient signals in the two Advanced LIGO interferometers [18]. All three of these analyses were able to recover GW150914. GW151226, having its power spread out over a longer period of time, requires a matched filter search and was not recovered by CWB. The two matched filter search produced consistent results. For brevity, we will focus on the results of the PyCBC search pipeline.

Event	Time(UTC)	FAR (yr^{-1})	$m_1 (M_\odot)$	$m_2 (M_\odot)$	χ_{eff}	D_L (Mpc)
	14 September					
GW150914	2015 09:50:45	$< 5.8 \times 10^{-7}$	36_{-4}^{+5}	29_{-4}^{+4}	$-0.06_{-0.18}^{+0.17}$	410_{-180}^{+160}
	26 December					
GW151226	2015 03:38:53	$< 5.8 \times 10^{-7}$	14_{-3}^{+9}	8_{-3}^{+2}	$0.20_{-0.10}^{+0.21}$	490_{-210}^{+180}
	12 October					
LVT151012	2015 09:54:43	0.44	23_{-5}^{+18}	13_{-5}^{+4}	$0.0_{-0.2}^{+0.3}$	1100_{-500}^{+500}

Table 2: Table of foreground events found in the first observing run. The quoted false alarm rates are calculated by the PyCBC search pipeline. The GstLAL search pipeline reported similar results. The astrophysical parameters are further explained in the parameter estimation companion paper [48]. Two binary black hole systems, GW150914 and GW151226, were discovered with a false alarm rate $< 5.8 \times 10^{-7}$, which is the upper limit on false alarm rate set by the amount of time used in the analysis. This corresponds to a statistical significance $> 5.3\sigma$. A third event, LVT151012, was an interesting foreground event that was not statistically significant to be claimed as a detection, but could be part of a larger gravitational wave population that includes weaker signals.

Figure 9 shows the results of the PyCBC search over the whole of the first observing run [47]. The black curve shows the number of expected foreground events at a given $\hat{\rho}_c$ based on background noise for the analysis. For this curve, GW150914 is allowed to remain in the data when generating a background from timeslides. This answers an interesting question: if GW150914 is considered to be a chance coincidence due to noise, could a combination of GW150914 in one detector and background noise in the other detector generate a signal as loud as GW150914?

The blue curve shows the search background when GW150914 is removed from the analysis and not used when generating a background from timeslides. Since we believe that GW150914 is a real gravitational wave signal, using it in background calculations no longer provides a search background that is a realization of detector noise alone when evaluating the significance of quieter signals. If GW150914 is allowed

to produce background events, the significance of GW151226, which is represented by the orange square at $\hat{\rho}_c = 12.6$, is highly diminished. This can be seen by comparing the blue and black curves. The differences between them, including the extension of the background to $\hat{\rho}_c = 21$, are the result of GW150914 combining with background noise.

The orange squares indicate the number of foreground events that were actually recovered by the search pipeline. The statistical significance of a given foreground event is determined by the rate at which detector noise produces background events with a detection statistic higher than that of the signal [12]. The false alarm rate for a given foreground event is defined as the rate of background events with a $\hat{\rho}_c$ greater than or equal to that of the foreground event. This rate can be converted into a false alarm probability by assuming that the background events follow a Poisson distribution.

GW150914 was an exceptionally loud signal and is the loudest event in the analysis. Since there are no background events as loud as GW150914, its statistical significance has a lower limit of 5.3σ but is not exactly calculated. The associated statistical significance is listed on the horizontal bars on the top of the plot. The color of each bar corresponds to the background from which the statistical significance was measured.

With GW150914 removed from the search background, we can correctly evaluate the statistical significance of GW151226. Figure 10 shows a zoomed in version of the search background with GW150914 removed. GW151226 is the loudest event in the analysis once GW150914 and its associated background triggers have been removed. Since there are no background events as loud as GW151226, its false alarm rate can be bounded to 1 per the entire analysis time. The associated statistical significance has a lower limit of 5.3σ but can not be directly calculated. The blue curve in this plot shows the search background with GW151226 removed from the analysis. Any quieter foreground triggers, such as LVT151012, will have their false alarm rate and statistical significance determined by this background distribution. LVT151012, which is the second loudest foreground event in Figure 10, was recovered at $\hat{\rho}_c = 9.6$ and assigned a statistical significance of just under 2σ as it was not louder than all background events. Since there are many background events that are louder than LVT151012, there is no significant change in the background distribution

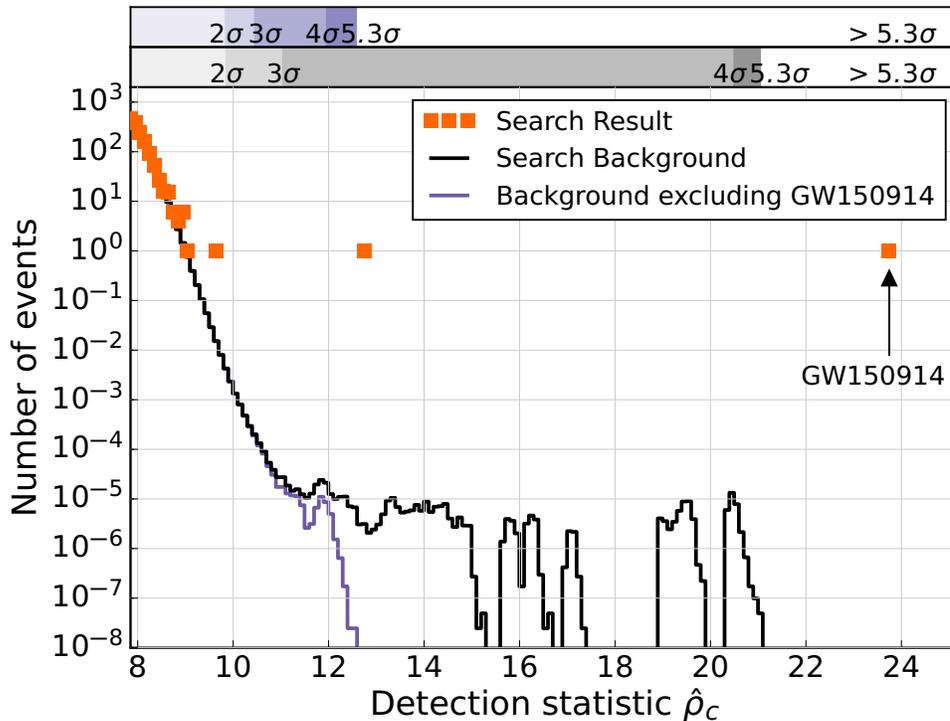


Figure 9: PyCBC search results for the first observing run. The black curve is the search background relevant to GW150914. The blue curve is the search background relevant to GW151226 where GW150914 has not been included in the search background calculation. GW150914 was the loudest event in the first observing run and was reported with a significance $> 5.3\sigma$. Figure 10 provides a better visualization of the significance of GW151226.

if LVT151012 is removed before performing time slides to generate the background.

3.4 Summary

The search results of the first observing run are summarized in Table 2. False alarm rates are quoted as estimated by the PyCBC search pipeline. The two discovered binary black hole signals, GW150914 and GW151226, differed by about a factor of 3 in total mass and originated at similar distances from Earth, which is responsible for the higher SNR of GW150914. Both events are estimated to have occurred at

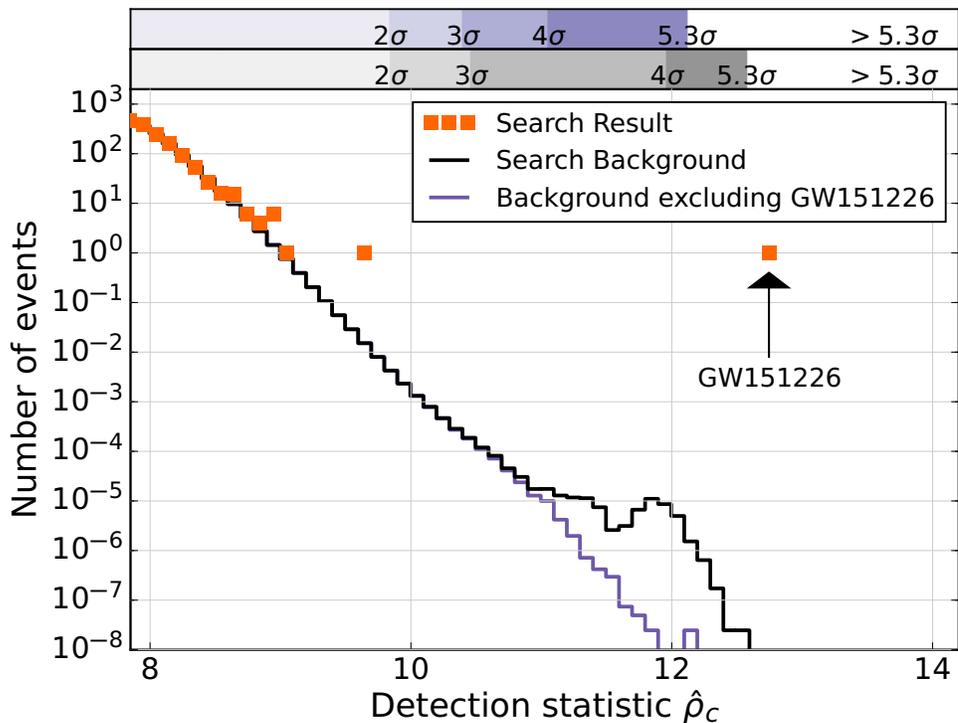


Figure 10: PyCBC search results for the first observing run with GW150914 removed. The black curve is the complete search background. The blue curve is the search background when GW151226 is removed and not allowed to combine with noise to generate background events. In both cases, GW151226 is the loudest event in the analysis. The statistical significance of GW151226 is bounded to be $> 5.3\sigma$. The second loudest event in this plot is LVT151012, which is assigned a statistical significance of just under 2σ .

similar distances, with their error regions having significant overlap. The third interesting foreground event, LVT151012, is estimated to have a total mass greater than GW151226, but its distance is estimated to be much further away.

The estimated distance for GW150914, 410 Mpc, corresponds to 1.3 billion light years. Since general relativity predicts that gravitational waves travel at the speed of light, this means that the signal we measured as GW150914 was emitted 1.3 billion years ago, before complex life existed on Earth. For reference, the Andromeda Galaxy is less than 1 Mpc away from Earth.

Both GW150914 and GW151226 were the loudest events when compared to their

respective background distributions, resulting in a false alarm rate $< 5.8 \times 10^{-7} \text{yr}^{-1}$, which corresponds to a statistical significance of $> 5.3\sigma$. LVT151012 has a false alarm rate of 0.44yr^{-1} , which corresponds to a statistical significance of just under 2σ .

Chapter 4

Detector Characterization

The Detector Characterization (DetChar) group works at the interface between the instrument science and data analysis groups. The purpose of the group is to understand the effects of instrumental noise sources on the output of astrophysical searches and mitigate them if possible.

As mentioned in Section 2.0.3, the detector data contain non-Gaussian noise transients that adversely affect the output of astrophysical searches. The signal consistency tests built into the analysis pipelines are powerful, but there are still certain types of transients in the data that are not fully suppressed by them. This is demonstrated in Figure 7b, which shows a tail of loud events in the re-weighted SNR distribution from the PyCBC search pipeline. In Gaussian noise, this distribution does not extend beyond a re-weighted SNR of 8 in the time scales used for the O1 analysis. (For more discussion on the output of the PyCBC search pipeline using Gaussian noise, see Section 7.2).

Detector Characterization studies are focused on ensuring that the astrophysical searches are using nominal, stationary detector data in their background noise estimations, suppressing the rate of loud events that will pollute both the background and the foreground distributions of astrophysical searches, and validating instrumental performance at the time of a gravitational wave candidate.

4.1 Methods of Detector Characterization

The first step in a detector characterization study is identifying noisy or problematic data. These studies can be initiated in a number of ways. The three most common are the appearance of loud background events in an astrophysical search pipeline as discussed above, a message from the commissioning team regarding instrument performance, or excess noise flagged by data quality monitoring software. Section 4.2 discusses the data quality monitoring software further.

There are a large number of recorded signals used to monitor and control the interferometers that are not used in astrophysical searches. These auxiliary channels are considered safe to use for noise characterization because they are not sensitive to gravitational wave signals. Analyses of auxiliary channels allow for the identification of systematic noise sources [57, 58], such as environmental disturbances [59] or excess motion of auxiliary optics in the interferometer [8, 60].

Once data with excess noise have been identified, they must be characterized in order to track down the source of the noise. A number of questions can be asked to characterize the noise. Is the noise transient or a slow drift? What is the typical frequency and bandwidth of the noise? Does the noise follow a power law in frequency? Does the noise have a characteristic shape in the time-frequency plane? Are the noisy frequencies of the signal coherent with other signals in the instrument such as environmental monitors and optical control signals? Is the noise source localized to a specific chamber or does it exist at multiple physical locations in the interferometer? Does the characteristic frequency match any of the known mechanical resonances in the interferometer? If the noise is a slow drift, does it correlate with the slow drift of other signals in the interferometer? Does the noise seem highly digital or discretized? After gathering all available information about the character of the noise and its coupling mechanisms, efforts shift toward attempting to mitigate the effects of the noise on search pipelines.

There are two primary ways to mitigate the effects of instrumental noise on the output of a search pipeline. The first option, which is highly preferred, is to track down the source of the noise in the interferometer and fix the problem at its origin. If investigations provided enough information that a problem can be traced back to a specific piece of electronics or a specific control loop, the problem can be fixed at the

source. However, this is not always possible since instrument noise can be difficult to pin down and hardware repairs are often too invasive to perform during an observing run.

If the problem cannot be fixed at the source, the second option is to remove the problematic data from the astrophysical analyses. When a significant noise source has been identified using auxiliary channels and cannot be repaired immediately, a data quality veto can be generated to indicate times when the output data from the interferometer are not nominal [43, 61, 44, 62]. Data quality vetoes are discussed further in Chapter 7. If possible, it is always preferable to fix a problem at the source.

4.2 Tools and algorithms

Identifying and characterizing instrument noise is facilitated by a suite of software tools and algorithms designed to flag data with excess noise and to help correlate this noise with other signals in the interferometer. The major tools required for understanding the data quality investigations in this thesis are discussed below.

4.2.1 Omicron

One way to quantify the amount of excess noise in $h(t)$ is to look for times where the signal contains excess power using Omicron, a burst algorithm [63]. The first stage of the Omicron pipeline applies a set of signal conditioning processes to $h(t)$, including a whitening filter, a high pass filter, and a downsampling process. Once the data have been whitened, they are projected into a sine-Gaussian basis. Each sine-Gaussian basis function is defined by a central time, t_0 , a central frequency, f_0 , and a Q-factor, which is defined as [64]

$$Q = \frac{f_0}{\Delta f} = 4\pi f_0 \Delta t, \quad (4.1)$$

where Δt is the time duration of the sine-Gaussian and Δf is the bandwidth of the sine-Gaussian. Using these parameters, each sine-Gaussian basis function can be represented as a tile in the time-frequency plane centered around t_0 and f_0 , where the width of the tile is determined by the time duration and height of the tile is determined by the frequency bandwidth.

For each of these tiles, the energy is measured and compared to the median tile

energy. If there is an excess of energy in a given tile relative to the median tile energy, a signal-to-noise ratio (SNR) is calculated and a trigger is generated to annotate the event. For each trigger, the SNR, central time, central frequency, duration, bandwidth, and Q of the tile are recorded.

Once the data have been decomposed into the full set of basis functions, the resulting set of triggers is sent through a clustering algorithm. This is necessary because the set of sine-Gaussians is an overcomplete, non-orthogonal basis and a single event in the data can generate multiple triggers corresponding to different values of t_0 , f_0 , and Q . The resulting clustered triggers define the peak time, peak frequency, and SNR of a cluster as the central time, central frequency, and SNR of the most significant tile in the cluster [64].

The most useful way to visualize the output of Omicron is in the time-frequency-SNR plane, sometimes referred to as a ‘glitchgram’, where each trigger is represented as a point in a scatterplot. Figure 11 shows an example set of Omicron triggers in the time-frequency-SNR plane. Each dot represents a trigger at a certain peak time and peak frequency. The color of each dot represents the SNR of that trigger.

In Gaussian noise, the SNR of a given trigger is not expected to exceed 8. In this example, there are a number of triggers with $\text{SNR} > 8$, some with noticeable structure and some that seem more randomly scattered. These represent noise in the output of the interferometer. For example, there are numerous triggers between 10-20 Hz that represent excess noise at these frequencies, likely due to scattered light in the interferometer. There is a line of triggers at just above 2kHz that indicates a noise source with a constant peak frequency whose amplitude is being modulated and a high SNR is being reported. There is also a scattering of points with high SNR that are not as structured as the previous two examples, each one likely due to an individual loud glitch rather than a constant, systematic noise source.

The results of Omicron are a commonly used and extremely valuable tool for characterizing the noise in the instrument. A cursory glance at Figure 11 identifies 3 populations of noise in the instrument, each of which can be followed up on individually to discover both the source of the noise and its effect on astrophysical searches. Omicron triggers can also be used in statistical analyses to find correlated noise between auxiliary channels and the output of the interferometers. An often used example of this, Hierarchical Veto, is discussed below.

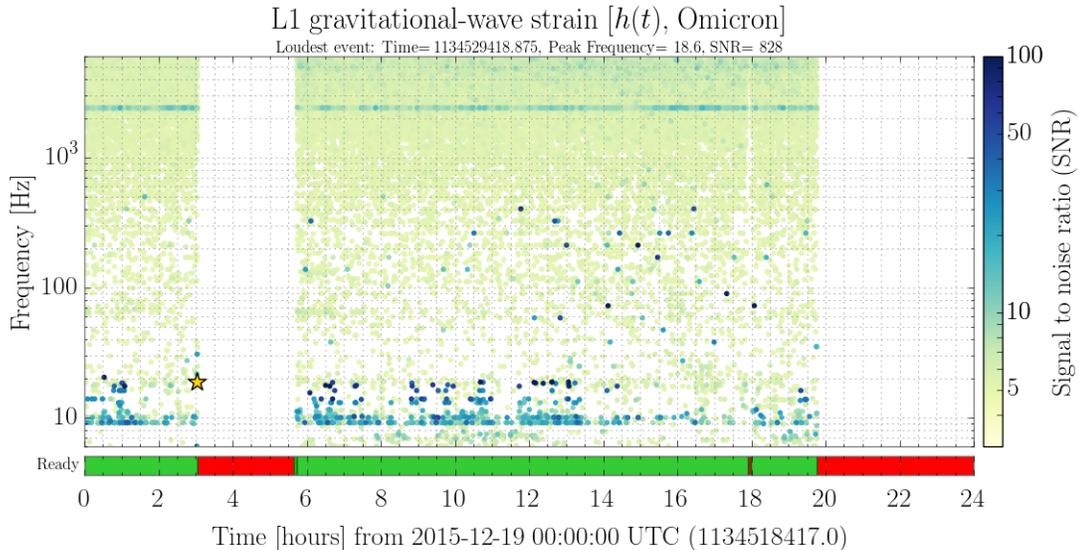


Figure 11: Time-frequency-SNR plot of Omicron triggers, often referred to as a ‘glitch-gram’. Each dot on this plot represents an event in the interferometer output, denoted $h(t)$, that was recorded with a peak time, peak frequency, and SNR. In Gaussian noise, all triggers on this plot would have an $\text{SNR} < 8$. Since this plot is generated using real detector data from O1, there are structures of loud triggers that indicate populations of noise transients. For example, the clusters of triggers between 10-20 Hz that likely represent scattered light in the interferometer.

4.2.2 Hierarchical Veto

One tool that is often used to look for time coincidence between noise transients in auxiliary channels and the output of the interferometer is Hierarchical Veto (Hveto) [57]. Typically, Hveto is used to compare a channel that potentially contains gravitational wave signals, denoted $h(t)$, and an auxiliary channel that does not have direct astrophysical implications. Hveto counts the number of coincident triggers between two time series using a user-defined time window centered around each trigger in the auxiliary channel. Figure 12 shows an illustration of auxiliary channels with noise transients that are coincident with noise in $h(t)$. Hveto iterates over all auxiliary channels to search for noise that is coincident with noise in $h(t)$ in a statistically significant way.

The figure of merit returned by Hveto for each auxiliary channel after comparison

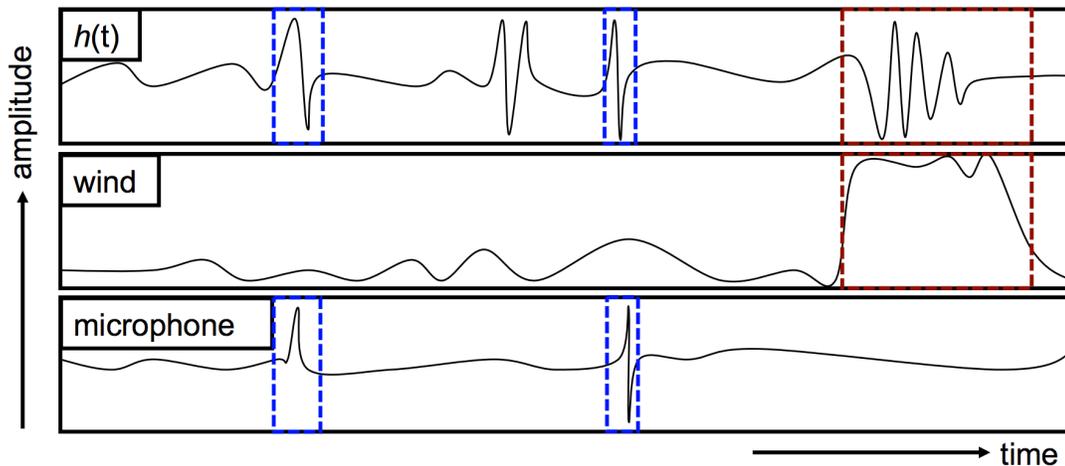


Figure 12: A time-series illustrating coincident noise between auxiliary channels and $h(t)$. The top panel is $h(t)$, which contains multiple noise artifacts of varying duration. The middle panel is a readout of wind speed on site, which shows an elevated period coincident with a longer duration burst of noise in $h(t)$. The third panel is a readout of a microphone on site, which shows two glitches that are coincident with bursts in $h(t)$. If noise in these auxiliary channels are coincident with noise in $h(t)$ in a statistically significant way, the noisy data in $h(t)$ can be removed. Figure reproduced from [57].

to $h(t)$ is called significance. Significance answers the following question: how unlikely is it that the coincident triggers in these two channels were the result of two arbitrary Poisson processes occurring in each channel? More specifically, given two arbitrary Poisson processes, how unlikely is it that we measure n or more coincident triggers given that expected number of coincidences from random chance is μ ?

Significance is calculated as [57],

$$S = -\log_{10}\left(\sum_{k=n}^{\infty} P(\mu, k)\right), \quad (4.2)$$

where n is the number of coincidences found between the two channels during the total analysis time and $P(\mu, k)$ is the Poisson probability distribution function. The Poisson probability of witnessing k coincidences when the expected average number of coincidences is μ is

$$P(\mu, k) = \frac{\mu^k e^{-\mu}}{k!}. \quad (4.3)$$

The expected number of coincidences between triggers in $h(t)$ and the auxiliary channel based solely on chance, μ , is estimated as [57],

$$\mu = \frac{N_h N_{aux} T_{win}}{T_{tot}}, \quad (4.4)$$

where N_h and N_{aux} are the number of triggers in $h(t)$ and a given auxiliary channel respectively, T_{tot} is the total analysis time, and T_{win} is the length of the coincidence window used.

A high value of significance indicates that the triggers in the channels were very often coincident in time and that there is a very small probability that their intersection is a product of random chance. This is a very useful measure when we are searching for auxiliary channels that might have some noise coupling into $h(t)$. A significance value of up to 5 is often observed in channels with no causal relationship to $h(t)$ [57], which is a useful threshold for identifying effective vetoes.

Another interesting figure of merit used for a given comparison Hveto is the ratio of $\frac{efficiency}{deadtime}$. Efficiency is defined as the percent of triggers vetoed from $h(t)$ during a round of vetoes. Deadtime is defined as the percent of total analysis time removed from $h(t)$ during a round of vetoes. A ratio of 1 is what we would expect from vetoing time at random, indicating no strong time correlation between triggers in the two channels. A high value of this ratio, which is ideal, indicates that we are vetoing a large number of triggers while maintaining a high percentage of our analysis time. This means that the triggers are often close enough in time that we can catch a large number of triggers using a small time window.

The deeper utility of Hveto is made evident when a channel is found to have a strong correlation with $h(t)$. When Hveto discovers an auxiliary channel that has a strong correlation with $h(t)$, which is called the round winner, it removes all of the time windows surrounding auxiliary channel glitches and recalculates the significance of the list of auxiliary channels. If a channel's significance has dropped after this removal of time, it must have had a large amount of glitches coincident with the round winner. The change in significance of each channel is displayed on a figure called a 'drop-plot'. This is one of the most powerful features of Hveto - the ability to find families of channels that often glitch at the same time.

Ideally, the list of significant channels displayed on the drop-plot will be able to localize the issue to a specific subsystem or area of the IFO. For example, if a channel

representing the alignment of the input mode cleaner has glitches that are strongly correlated to $h(t)$, it would be interesting to look at the drop-plot and find out what other channels are glitching at the same time (suspensions, laser power, etc.). From there, the issue can be investigated and brought to the attention of commissioners for repair or physical inspection. This is not always possible as sometimes the cause of the glitches is unclear, but identifying times of poor data quality is still useful.

Using Hveto, we can monitor auxiliary channels to find and remove glitches in $h(t)$ that would otherwise pollute a gravitational-wave analysis. Removing these glitches serves multiple purposes for the search pipelines. Removing high SNR glitches cleans up search backgrounds and allows the search pipelines to claim a lower SNR threshold for potential detections. A lower SNR threshold implies a larger search volume for astrophysical analysis. Removing glitches reduces the potential for false alarms in the search pipelines, which in turn increases the confidence of eventual detections.

4.3 Instrumental Detector Characterization Studies

The tools described above have been used in numerous studies to characterize and understand the transient noise in the LIGO interferometers. This section provides an overview of some of the more critical data quality investigations that occurred leading up to and during the first observing run. My personal contributions to the following sections were to develop the monitoring software used in Section 4.3.1 and generate data quality vetoes to remove these times from analysis, to develop the monitoring software used in Section 4.3.2 and follow up using Hveto to indicate which suspensions were showing DAC calibration glitches, and to help find the critical frequencies that generated whistle glitches in Section 4.3.3.

4.3.1 Analog-to-Digital Conversion Overflows

Advanced LIGO interferometers are controlled in real-time using a digital control system installed on a series of computers referred to as front end computers. This system overall is referred to as the Front End Control (FEC) subsystem. This subsystem oversees the operation of a series of feedback loops used to control various systems in the interferometers. Figure 13 shows a basic example of a feedback loop used to control the alignment of an optic. When the optic is misaligned, the position

of the laser on the sensor shifts, generating an error signal. In an aLIGO feedback loop, the front end computers must be capable of reading in an analog signal from the interferometer (position measurements, photodiode currents, etc), which is digitally sampled using an analog-to-digital converter (ADC), using these now digital values in a series of control filters, and outputting an analog control signal to send back into the interferometer via a digital-to-analog converter (DAC). This control signal is sent back into the interferometer to actuate on the alignment of the optic and complete the feedback loop.

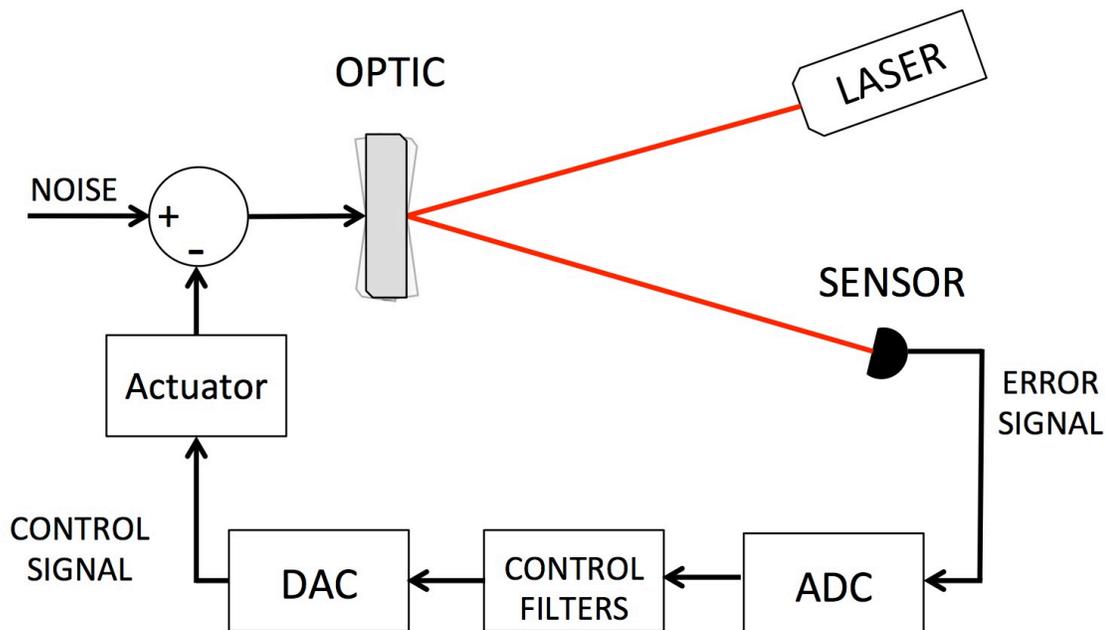


Figure 13: A simplified visualization of a feedback loop designed to control the alignment of an optic. Misalignment of the optic causes the beam to move on the face of the sensor. The output of the sensor is used as an error signal, which is digitally sampled, filtered, and converted back into an analog signal to correct the alignment of the optic. The analog signal is sent through an analog actuator that physically interacts with the optic to correct its motion.

The process of digital sampling is handled by an analog-to-digital converter and the process of analog output is handled by a digital-to-analog converter. Since these converters are linearly mapping a continuous signal onto a discrete range, they are

limited by their digital bit depth. For example, a 16 bit ADC is only capable of representing 2^{16} discrete values, or a range from zero to 65536. This range is often centered around zero, giving the ADC the capability to handle a range of ± 32768 . An incoming analog signal is mapped onto this range and converted into a digital signal. For an analog signal with a range of $\pm 10V$, 10V would be mapped to 32768 digital counts and -10V would be mapped to -32768 digital counts with all of the intermediate voltage values being linearly mapped to the range. This means our digital system would recognize a discrete step size of $10V/32768$ counts $\approx 305\mu V/\text{count}$.

Looking at the digital sampling described above, we must be aware of how our system is going to react when our analog input signal exceeds the intended maximum value of 10V (e.g., an 11V input). The ADC has already assigned its maximum digital value to 10V. In this case, the ADC will continuously output its maximum value as it has no way to map 11V into a discrete value. This is called a digital overflow. The same process can occur in a DAC when a digital signal is sent out at the maximum allowed digital value. The resulting analog signal will be railed at the maximum output value of the DAC, creating a sharp corner in the output signal as it flattens out.

If the digital system is not able to correctly sample and understand an analog error signal, it is easy to imagine a scenario where the response of the digital system and the output control signal are not able to complete the control loop as designed. This may cause glitches or misalignments in the interferometer. We must also consider the fact that many ADCs are calibrated to reflect the intended dynamic range of an optic. If a saturation is occurring, there is a good chance that an optic has moved beyond this intended dynamic range, which also may cause glitches or misalignments.

The ADCs and DACs are monitored by a series of auxiliary channels, which are automatically generated in the front-end system. These auxiliary channels monitor each ADC and DAC channel and note when any of the channels has reached its digital limit. These channels can be used to generate flags that mark ADC and DAC overflows, which can be compared with glitches in $h(t)$ to search for glitch mechanisms driven by overflows. These channels can also be used to flag any large glitches that cause digital overflows so that they can be removed from astrophysical searches.

Figure 14 shows an example of a large glitch that caused a digital overflow and was removed from gravitational wave analyses. Figures 14a and 14b show a large

glitch in $h(t)$ and the response of a control signal that corrects the motion of ETMY respectively. The signal in 14b, which is supposed to be controlling the motion of ETMY, exceeds the digital limit during this glitch. Figure 14c shows the auxiliary channel that monitors this digital overflow incrementing as it witnesses the digital overflow.

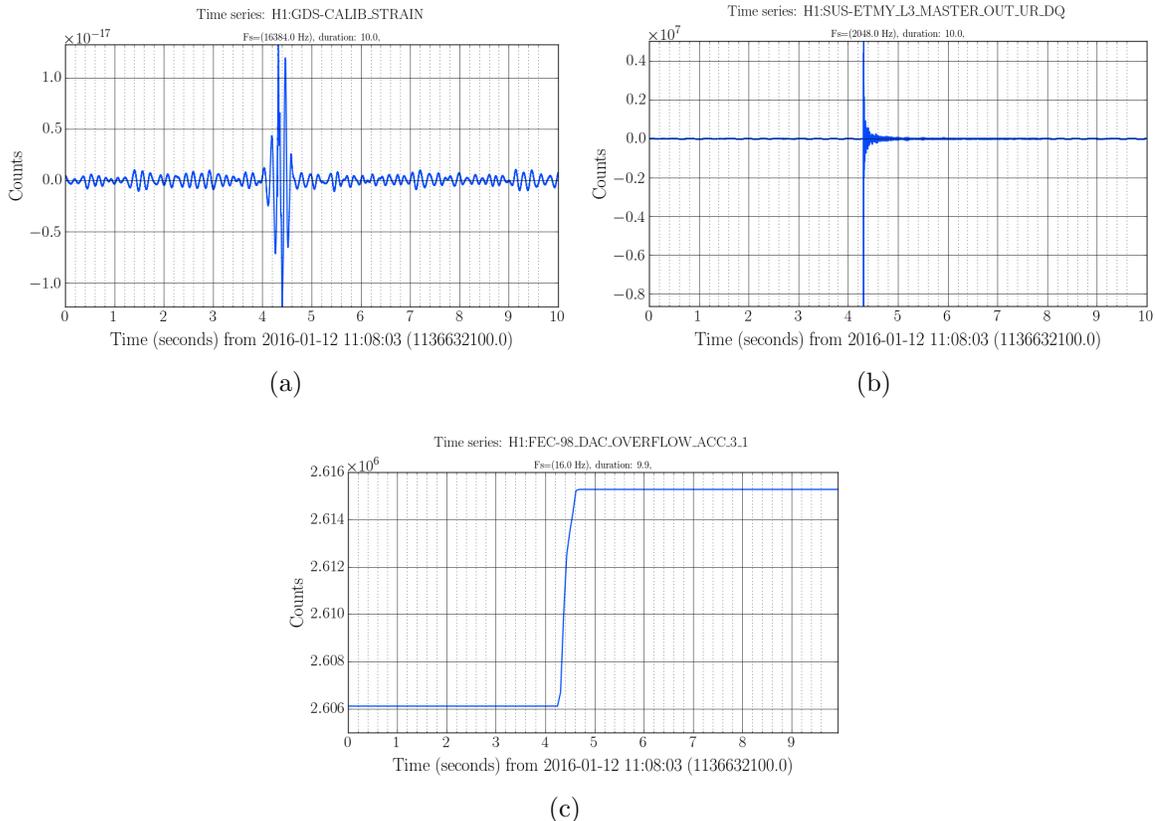


Figure 14: Time-series of an ETMY control signal saturation in the H1 detector. Figure 14a shows a glitch in the calibrated $h(t)$ channel. Figure 14b shows the response to this glitch in the control signal used to control the bottom stage of ETMY and actuate on the DARM degree of freedom. This signal exceeds its digital overflow point at its peak, reaching a value on the order of 10^6 when the DAC can only handle digital values up to 65536. Figure 14c shows the front end channel responsible for monitoring digital overflows of this particular ETMY control signal. Since the witness channel is cumulative, overflows can be identified by flagging any time in which this witness channel is increasing.

This method was used throughout O1 to generate data quality vetoes that were distributed to the Burst and CBC searches. The first veto that was generated this way was used to flag DAC overflows of the ETMY control signal, as demonstrated in Figure 14. The other veto generated in this framework was used to flag ADC overflows in the OMC DC photodiode used as the error point of the DARM control loop.

4.3.2 Suspension DAC calibration glitches

A common glitch mechanism leading up to the first observing run was due to calibration errors in digital-to-analog converters (DACs) responsible for providing analog signals to the aLIGO suspensions. The aLIGO suspension subsystem uses 18-bit DACs to interact with the optics in the interferometer. These 18-bit DACs are created by combining a 16-bit DAC with a 2-bit DAC inside of the same electronics box. The 2-bit DAC is responsible for the two highest order bits of the output, while the 16-bit DAC is responsible for the 16 lowest order bits of the output. If the 16-bit DAC and 2-bit DAC have not had their output voltages carefully calibrated, there will be a voltage discontinuity at the output of the DAC when engaging the 2 highest order bits.

Since these DACs use the two's complement representation for signed binary numbers, there are two critical points where the two highest order bits of the DAC become necessary. The highest order bit is used to indicate negative numbers, so an output discontinuity is expected when transitioning from a positive number to a negative number, that is, crossing through a value of zero. The other bit from the 2-bit DAC is used to represent large output values and engages when the DAC needs to express a value which is unable to be represented by a 16-bit DAC alone. As such, we also expect to see discontinuities when the DAC output crosses $\pm 2^{16}$.

The fact that this discontinuity existed in suspension subsystem was particularly problematic, as the suspension DACs are used to directly actuate on mirror positions and optical cavity lengths. Any time a suspension DAC crossed one of these problematic output values, it would actuate on the optics with a step function and cause a glitch in the optical cavity length. Figure 15 shows an example of this issue where the DAC providing actuation signals to the power recycling mirror (PRM) is crossing through zero and there are associated glitches visible in the length readout of the

power recycling cavity [43].

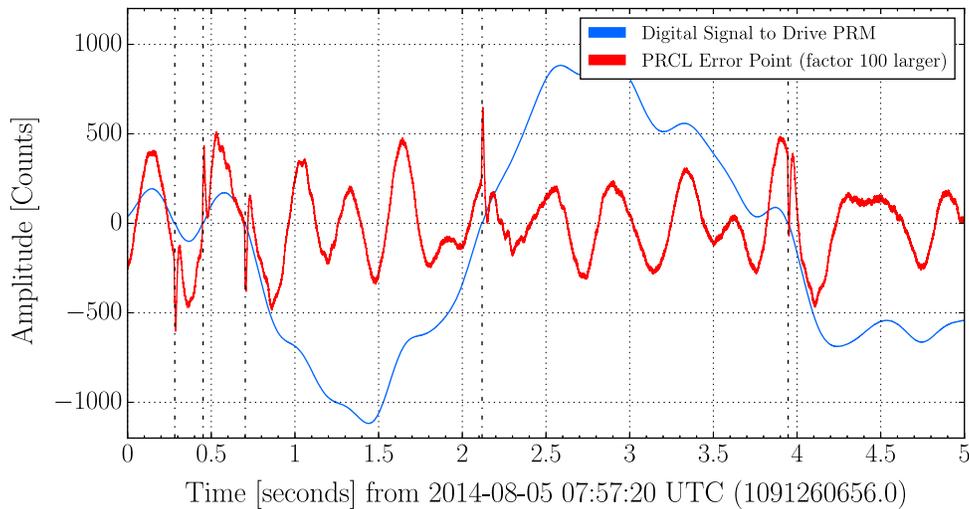


Figure 15: A timeseries plot showing the effects of DAC calibration glitches. The blue trace shows the digital control signal being sent to the digital-to-analog converter. The red shows the resulting power recycling cavity motion rescaled by a factor of 100. When the control signal crosses through a value of zero, the output of the DAC experiences a discontinuity, leading to a glitch in the power recycling cavity length.

The effects of this issue were visible in the $h(t)$ channel in engineering runs leading up to the first observing run. The most problematic culprit was the DAC that applied actuation directly to the optics of the ETMs, effectively pushing directly on the DARM degree of freedom and causing glitches in $h(t)$. These calibration errors manifested themselves as a population of glitches in $h(t)$ recovered by Omicron in the 20-100 Hz range. This is a very damaging frequency range for CBC searches, which hope to accumulate significant SNR in the region from 30-500 Hz. This population of low frequency glitches was obvious in an Omicron time-frequency scatter plot and was considered a significant noise source throughout the sixth engineering run.

Figure 16 shows the result of an Hveto run that looked for time correlations between Omicron triggers in $h(t)$ and times when the ETMY control signal crossed through a value of 2^{16} . The blue dots represent all Omicron triggers in $h(t)$. The red crosses indicate those that were coincident with the ETMY control signal crossing 2^{16} . The population of low frequency glitches with $\text{SNR} > 8$ was shown to be coincident

with the control signal transitions. This veto was very statistically significant, as shown in Table 3. The significance of 192.5 indicates that the probability of these coincidences being due to chance is negligible. The efficiency:deadtime ratio of 27 indicates that these glitches were removed with very small time windows (0.2s) and very little instrument uptime was removed in the process.

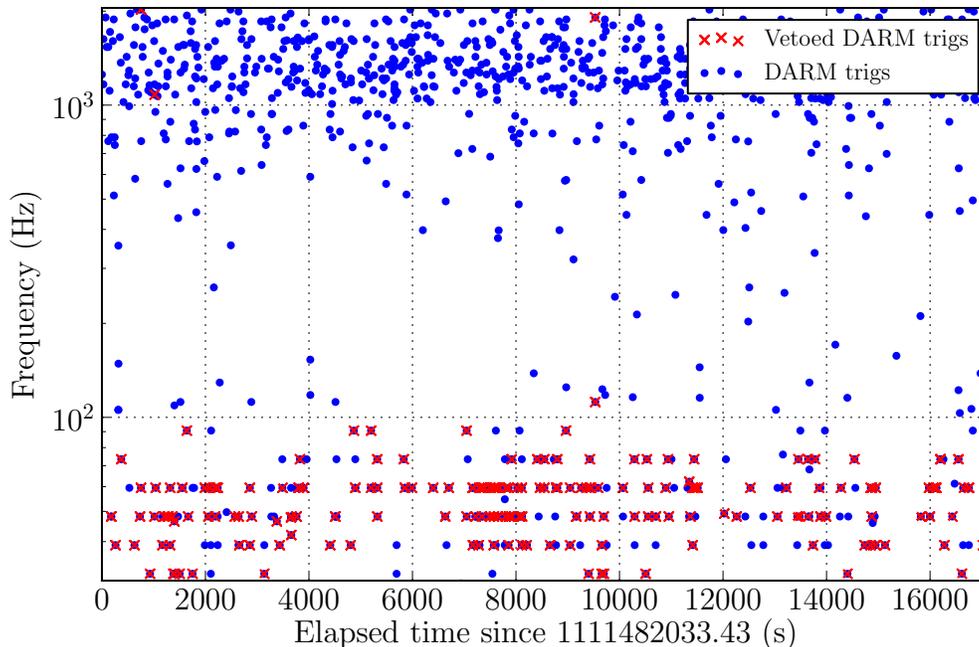


Figure 16: A time-frequency visualization of Omicron triggers in the H1 $h(t)$ channel. The blue dots indicate glitches in the DARM degree of freedom, each with a central time and central frequency. The red crosses indicate that a given trigger was vetoed by an auxiliary channel trigger which was found to be statistically significant using Hveto. The auxiliary channel triggers in this case indicate that the control signal on the bottom stage of ETMY has crossed a value of 2^{16} . The population of glitches between 20 - 100 Hz is highly coincident with these crossings of 2^{16} , indicating that they are caused by DAC calibration errors on this optic.

To fully understand the scope of this problem, software was developed that searched through the output of all suspension DAC digital output signals and marked times when they crossed 0 or $\pm 2^{16}$. These marked times were converted into trigger files and sent through Hveto to look for correlations between crossings of critical values and

Channel	Time window (s)	SNR threshold	Significance	Efficiency %	Deadtime %
ETMY control signal crosses 2^{16}	0.2	8	192.5	18.3	0.674

Table 3: Hveto results for ETMY DAC glitches

glitches in DARM as identified by Omicron. Through this method, we were able to identify which optics were experiencing DAC calibration glitches that had a coupling mechanism into DARM.

There were two approaches taken in an effort to mitigate these DAC glitches. The first was to introduce offsets into the suspension control signals so that they did not cross through a value of zero. This did solve the problem temporarily, but at the cost of a significant portion of the dynamic range of the output actuation. The more permanent fix was to run a calibration routine that resolved the issue between the 16-bit and 2-bit DACs. This was successful, though it had to be run on a weekly basis during site maintenance because the calibration tended to drift away from its nominal point after 2-3 weeks of operation.

During the first observing run, the systematic check of all suspension DAC digital output signals was performed again and the resulting triggers were sent through Hveto. This study revealed that the calibration process was successful; there was no evidence of residual DAC calibration glitches that had any noticeable coupling into $h(t)$. The only signal that had any significant correlation with glitches in $h(t)$ was not causally sensible.

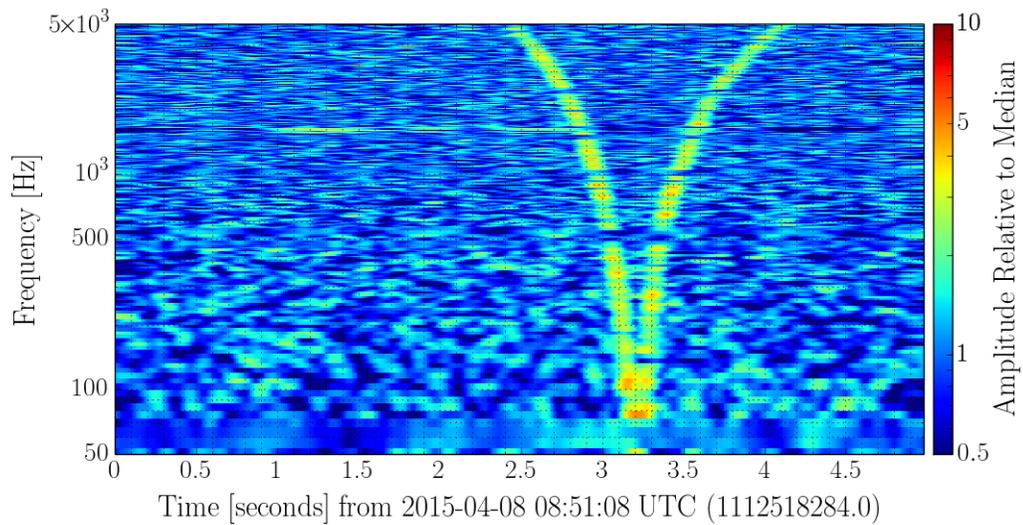
4.3.3 RF beatnote whistles

During Advanced LIGO’s commissioning, a population of glitches appeared in both the L1 and H1 interferometers which came to be known as ‘whistles’ or ‘RF whistles’. These glitches were the intermodulation products of two radio frequency oscillators, which produce sine waves at radio frequencies. A nonlinear mixing between two RF oscillators produced a beatnote signal whose frequency was equal to the difference in frequency between the two oscillator signals. Whistle glitches occur when two oscillator signals drift and cross each other in frequency. If one oscillator is drifting

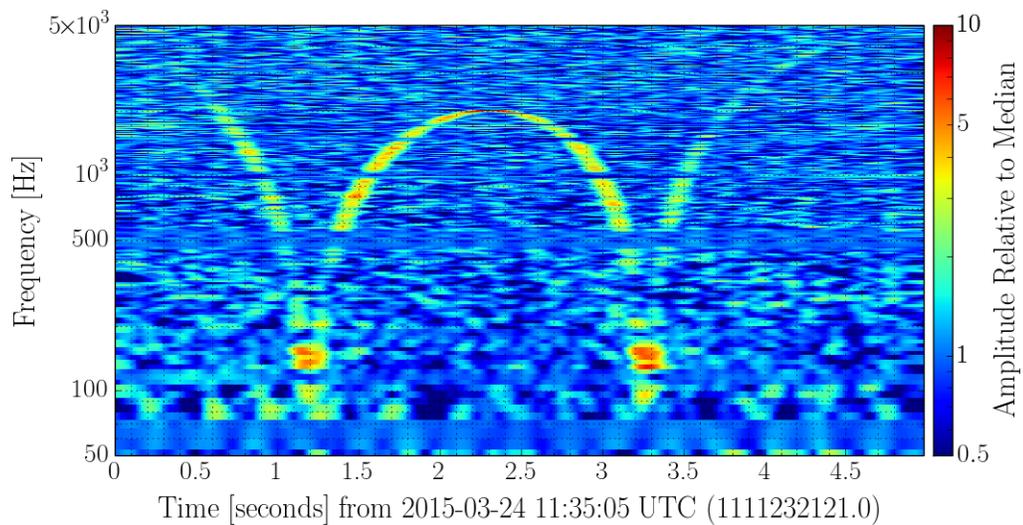
in frequency and another oscillator is at a fixed frequency, the beatnote generated between them will decrease in frequency as the oscillator signals become closer in frequency and then increase in frequency as they cross each other and drift away. As such, these glitches had a characteristic shape, beginning at high frequency and sweeping down in frequency through the detection band before turning around and sweeping back up to high frequencies. Figure 17 shows a time-frequency representation of whistle glitches in both the L1 and H1 interferometers [43]. These show the characteristic ‘V’ or ‘W’ shape produced when two oscillators drift past one another and have nonlinear mixing.

Voltage controlled oscillators (VCOs) are oscillators whose frequency can be tuned using an input voltage. These oscillators are used in control loops throughout the LIGO interferometers. One particular example of this is the control loop which locks the frequency of the input laser to the length of the input mode cleaner to guarantee a resonant optical cavity and effective mode cleaning. A signal representing the changing length of the input mode cleaner is read out using the Pound-Drever-Hall technique, which is discussed in depth in Chapter 5. This signal is used as the input to a VCO, which produces a signal whose frequency is a proxy for the length of the input mode cleaner. This signal is used as the set point in the frequency stabilization loop that controls the frequency of the input laser light. Through this path, the length of the input mode cleaner is used to set the frequency of the input laser light.

The signal that represents the length of the input mode cleaner was found to be a good witness for RF whistle glitches. Figure 18 shows the rate of Omicron triggers, which represent generic transient noise in $h(t)$, as a function of the length of the input mode cleaner [43]. The length of the input mode cleaner is in kHz as it is an error signal used to set the frequency detuning of the VCO. The red bars represent the distribution in the absence of whistle glitches. There is no value for which noise transients in $h(t)$ seem more likely to occur; the rate of glitches seems Gaussian distributed as the length of the input mode cleaner fluctuates about the set point of the control loop. The blue bars represent the rate of transients when RF whistle glitches are occurring. In this case, there are three preferred frequencies where it seems that noise transients in $h(t)$ have a tendency to occur. This indicates that there is a relationship between specific values of the length of the input mode cleaner and the presence of whistle glitches in the interferometer.



(a)



(b)

Figure 17: Time-frequency spectrograms of RF whistles at both the L1 and H1 interferometers. Figure 17a shows a whistle at L1 sweeping down from the kHz range and into the detection band where it interferes with searches for gravitational waves. Figure 17b shows a double whistle whistle at H1 where the two oscillators drifted back and forth across one another and caused two glitches in the detection band.

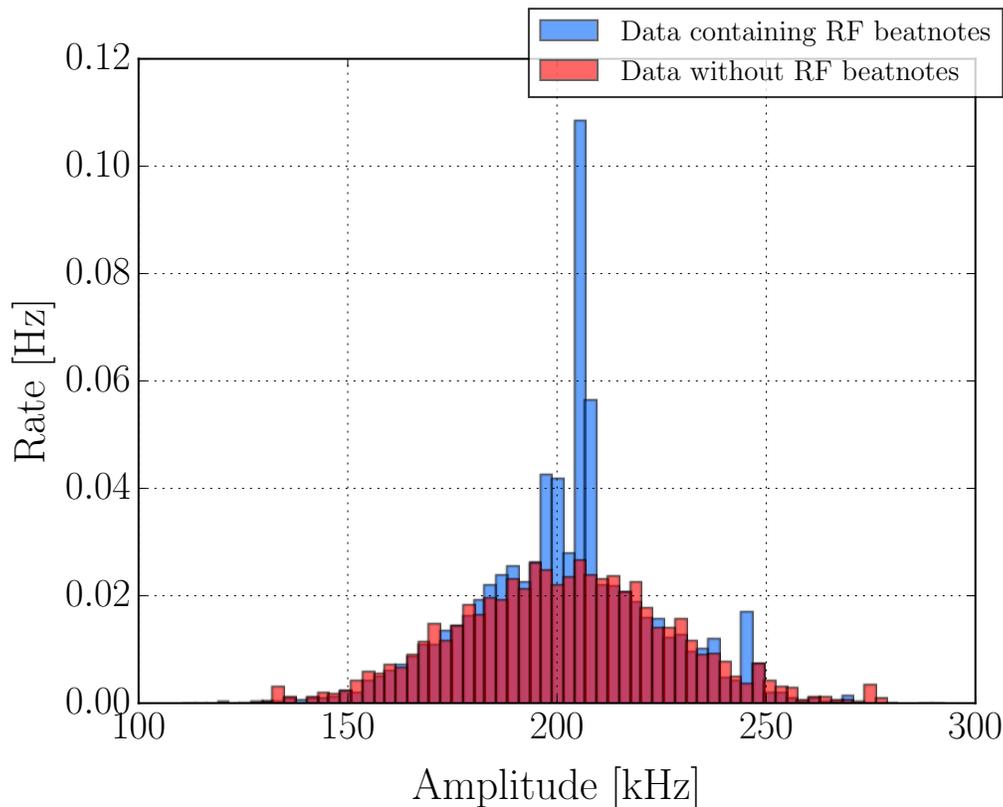


Figure 18: The red distribution shows the rate of Omicron in triggers in $h(t)$ when RF whistles are not present. The blue distribution shows the rate of Omicron triggers in $h(t)$ when RF whistles are occurring. The x-axis is the value of a channel that represents the length of the input mode cleaner. When there are no whistle glitches, there is no channel value for which Omicron triggers are more likely to occur. When there are whistle glitches in $h(t)$, specific values of the input mode cleaner length seem more likely to be coincident with Omicron triggers in $h(t)$.

The VCO that acts as a proxy to the length of the input mode cleaner is nominally set to 80 MHz. As the length of the input mode cleaner drifts, the oscillator frequency can be tuned by ± 1 MHz to track the length, resulting in a signal with a frequency of 79 - 81 MHz. It was found that these whistle glitches occurred when the VCO frequency swept through 79.2 MHz, which is the same frequency as an oscillator used to drive an acousto-optic modulator. The variable oscillator that was tracking the length of the input mode cleaner was drifting past the static oscillator at 79.2 MHz

and creating whistle glitches that were visible in $h(t)$. To reduce the number of whistle glitches in $h(t)$, the oscillator frequencies were moved away from one another so that the static oscillator was outside of the range of the tunable oscillator.

Figure 19 demonstrates how prevalent whistle glitches were before the oscillator frequencies were shifted to avoid them [43]. Figure 19 is a time-frequency scatter plot of Omicron triggers in $h(t)$. The blue dots represent all Omicron triggers generated for $h(t)$ over this stretch of time. The red crosses indicate that a given Omicron trigger was found to be coincident with an Omicron trigger generated for an auxiliary channel that was a capable witness for whistle glitches. Approximately 90% of the glitches in $h(t)$ are vetoed by the witness channel, indicating that whistles were the dominant source of transient noise in $h(t)$ in this time period.

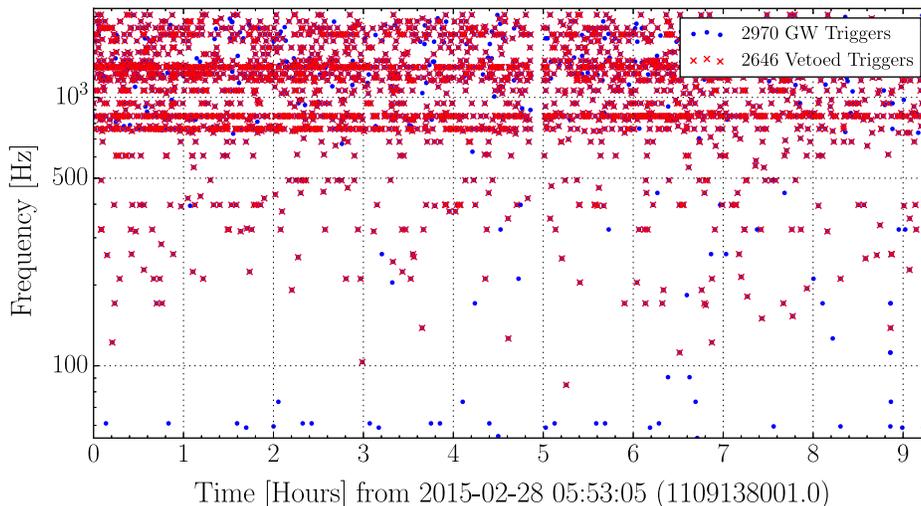


Figure 19: A time-frequency scatter plot of Omicron triggers. The blue dots represent all triggers found for the $h(t)$ channel. Red crosses indicate that a trigger was determined to be coincident with an RF whistle and vetoed. This veto is responsible for removing 90% of the glitches in this time period. The majority of the high frequency glitches were due to RF beatnote whistles

The shape of the whistle glitches was also very problematic for CBC searches since the second half of a whistle is a sinusoid with a monotonically increasing frequency, not unlike the characteristic ‘chirp’ signal produced by a CBC event. Certain CBC waveforms matched the shape of the whistle glitches well enough to fool the χ^2 signal

consistency test and produced loud background triggers in early CBC searches. The whistle glitches were fixed before the first observing run, so they were not a limiting noise source to CBC searches during observation.

4.4 Validation of O1 Data

Data were marked as suitable to be used in a gravitational wave search based on a set of conditions applied to each interferometer. The first condition indicates that the interferometer is in its nominal configuration or observation state according to software monitors used to control the instrument. The second condition indicates that no excitations are being applied to the instrument. This condition is set by the on-duty instrument operator on site who is continuously monitoring the detector performance.

To estimate the stability of the background noise used in the O1 analysis, the time period of September 12 - October 20, 2015 is studied. This is the stretch of time used for an extended background analysis of GW150914. My personal contributions to this study were to generate and analyze the CBC trigger rate plots seen in Figure 22.

The first visualization of background noise stationarity is the amplitude spectral density (ASD) of $h(t)$. Figure 20 shows the median ASD of the detector data from each interferometer over this time period as a function of frequency [44]. The shaded regions indicate the 5th and 95th percentile in that particular frequency bin. The most noticeable feature is the variance in the low frequency noise at L1, which was still not fully understood at the end of the run. Most of this variance is below 30 Hz, which is the low frequency cutoff used in the matched filter search, and would not have an effect on the output of the CBC search. The L1 data also had a higher variance in the 60-200 Hz region, which was a limiting noise in the PyCBC analysis (see Chapter 10). Other than these two regions, the spectra appear to be consistent and have reasonable variance at all frequencies [44].

A useful figure of merit for detector noise stationarity is the inspiral horizon distance. This is the distance at which a given source of gravitational waves could be recovered at SNR 8 in current detector noise assuming optimal orientation and sky location. Since GW150914 is the centerpiece of the analysis, the inspiral range was

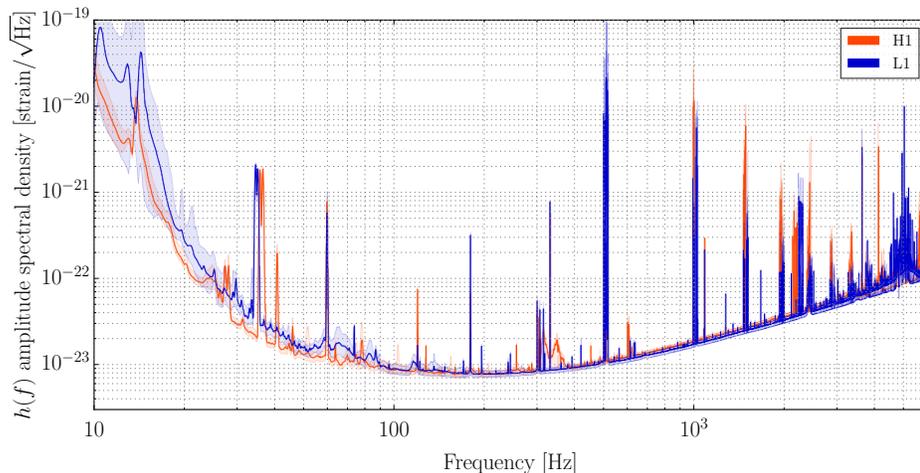


Figure 20: Median noise amplitude spectral density (ASD) in the first analysis period of O1. The shaded regions indicate the 5th and 95th percentile in a given frequency bin. The background for H1 was very stable throughout the run, which very few high variance bins. The background at L1 showed deviations from the median behavior in two regions: below 30 Hz and between 50-200 Hz. The strong lines in the noise spectrum are due to injected calibration lines, environmental sources such as the 60 Hz power line harmonics, and mechanical resonances. This plot is constructed by taking the Fourier transform of $h(t)$ and averaging several measurements of the power spectral density (PSD). The average PSD is then square rooted to calculate the ASD.

calculated using its parameters as the gravitational wave source. Figure 21 shows the inspiral range over the time used for an extended background analysis of GW150914 [44]. The inspiral horizon distance was between 1500-2000 Mpc for both instruments throughout the run. The L1 data, which were noisier overall, were slightly less sensitive to gravitational wave signals throughout the run. The two vertical lines represent GW150914 (dashed) and LVT151012 (dash-dotted) [44].

The final figure of merit for establishing that the detector noise was stationary is to look at the rate of events, or ‘triggers’, produced by the PyCBC pipeline between September 12 - October 20, 2015. This test tells us whether or not it is reasonable to combine the background of the PyCBC search over many weeks. Figure 22 shows the rate of single interferometer triggers produced by PyCBC for both H1 and L1 in this time span [44]. The circles indicate the rate of triggers with a re-weighted SNR

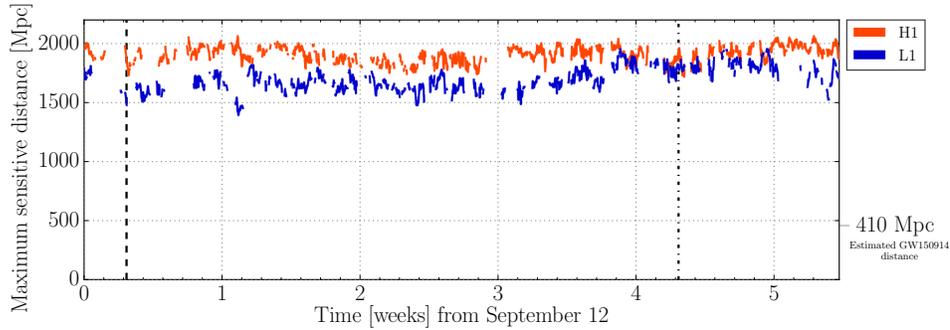


Figure 21: Inspiral horizon distance for sources with the same parameters as GW150914. This is the distance at which a signal could be recovered at SNR 8 if the binary system was at the optimal sky location, which is directly above the detector, and optimal orientation, which means that the orbital plane of the binary system was parallel with the plane of the detector. The inspiral horizon distance for GW150914-like sources was stable through the first analysis period of O1. The inspiral range for H1 was typically higher due to excess noise at L1. Overall, the inspiral range was between 1500 - 2000 Mpc. Each data point was calculated using a 2048 second stride. For reference, GW150914 is estimated to have been generated at a distance of 410 Mpc.

≥ 6.5 . The crosses indicate the rate of triggers with a re-weighted SNR ≥ 8 , which represents the rate of loud, non-Gaussian triggers in the analysis. The overall rate of triggers was consistent throughout the observing run, typically reported between 0.1 - 1.0 Hz. The rate of triggers with a re-weighted SNR ≥ 8 was typically < 0.01 Hz. The dashed and dash-dotted lines indicate the times of GW150914 and LVT151012 respectively [44].

4.5 Validation of Gravitational Wave Signals

The Detector Characterization group was responsible for characterizing the noise in the interferometers in order to validate the gravitational wave signals GW150914 and GW151226 [44]. The first part of this analysis, which studied the stationarity of the background noise, is discussed in Section 4.4. As a further check, the transient noise in the interferometer was also studied so that a confident detection claim could be made regarding GW150914. My personal contributions to this study included a

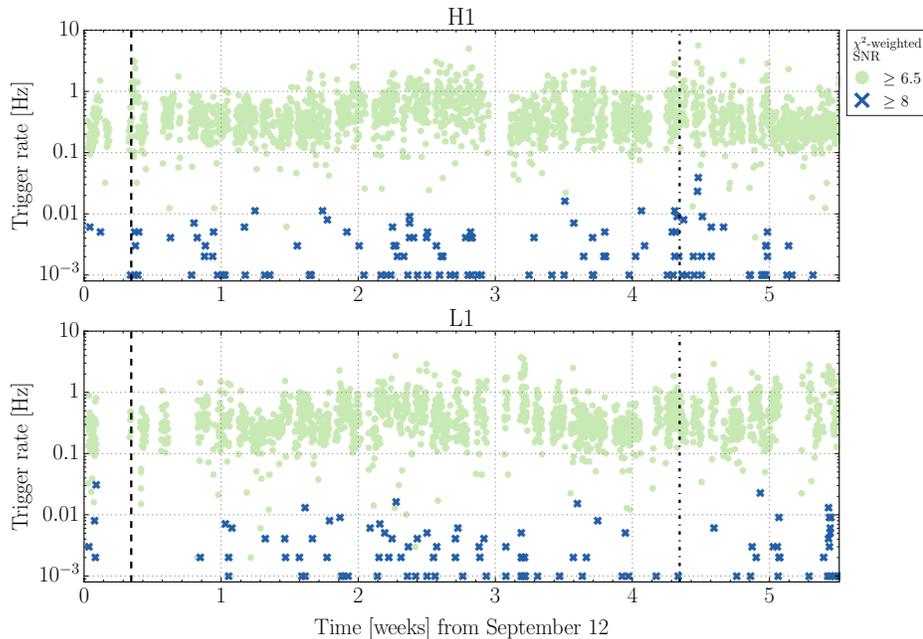


Figure 22: PyCBC single interferometer background trigger rates in O1. For these rates, any triggers found in coincidence were removed from the data set. Rates were calculated over a 2048s stride. The rate of triggers with re-weighted SNR ≥ 6.5 was stable throughout the analysis, typically reported at 0.1 - 1.0 Hz. The rate of louder triggers, at re-weighted SNR ≥ 8 , was typically at < 0.01 Hz. The dotted and dash-dotted lines represent the times of GW150914 and LVT151012 respectively.

thorough check of all instrumental monitors related to the length sensing and control and alignment sensing and control subsystems, a check of the monitors for DAC calibration glitches and ADC overflows that were described above, and to search for noise transients in the data that had a similar time domain morphology to GW150914.

In the engineering runs leading up to the first observing run, a great deal of work was done to understand as much as possible about the noise coupling mechanisms from auxiliary channels into $h(t)$. Among the most important of these was a set of signal injections to test the sensitivity of the physical and environmental monitoring (PEM) subsystem [44]. The PEM subsystem is comprised of a series of sensors that measure the ambient environmental noise at the interferometers [59]. This subsystem is comprised of seismometers, accelerometers, magnetometers, radio antennae, microphones, temperature sensors, and voltage monitors for the power lines supplying the

building. While the aLIGO detectors are extremely sensitive to external perturbations, they were built to be shielded against as many environmental disturbances as possible. In contrast, the PEM subsystem is comprised of extremely sensitive sensors that are more sensitive to environmental disturbances than the interferometer. By injecting signals into the interferometer enclosure, such as magnetic fields or acoustic vibrations, the relative sensitivity of the interferometer and the PEM sensors to environmental disturbances was established.

For both GW150914 and GW151226, a review the PEM subsystem reported that any environmental disturbances were at least 1 order of magnitude too weak to produce such an event. This includes electromagnetic transients, such as lightning strikes, that have the potential to generate coincident electromagnetic transients at L1 and H1.

In addition to the checks performed in the PEM subsystem, a series of standard checks were done to ensure nominal performance in the interferometers. These checks are organized in a detection checklist, which gathers all of the relevant questions about interferometer performance that may influence gravitational wave detection. This list includes checks for the DAC calibration glitches mentioned in Section 4.3.2, the ADC and DAC digital saturation glitches mentioned in Section 4.3.1, coincidence with generic transients as reported by Omicron, time-frequency scans of all auxiliary channels to be investigated by DetChar subsystem leads, injections and test signals, and GPS or digital system timing errors. Each category was investigated and followed up on by the DetChar group and none of them gave significant reason to doubt the validity of GW150914 or GW151226.

For both GW150914 and GW151226 there were small sets of auxiliary channels that showed excess power coincident with the gravitational wave signals, which is expected given the breadth of the auxiliary channel network, but after further investigation none of them had the necessary amplitude and frequency to generate an event similar to a CBC signal. An example of one such auxiliary channel is shown in Figure 23, which is a time-frequency spectrogram of the H1 Y-end seismometer signal that showed excess noise at the time of GW150914 [44]. The excess noise is due to an air compressor turning on roughly 75 seconds before GW150914. The noise is a 14 Hz line with 28, 42, and 56 Hz harmonics visible. This level of ground motion with a 175 second duration and a static frequency distribution was not capable of producing

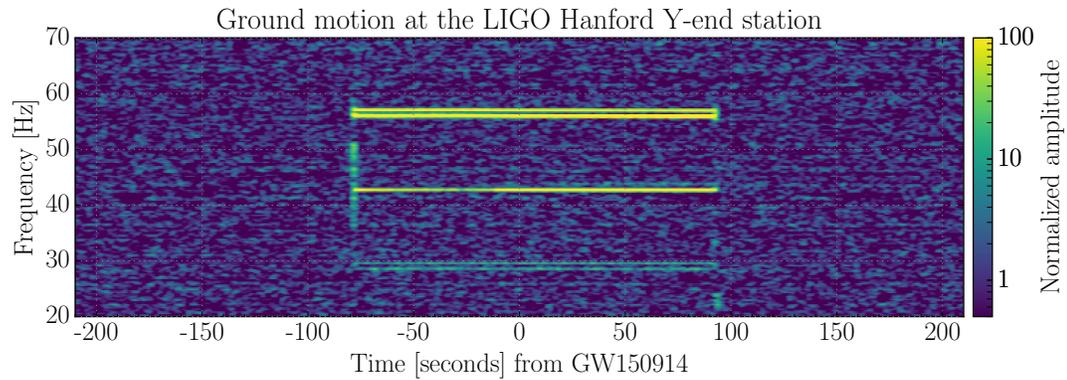


Figure 23: A time-frequency spectrogram of the H1 Y-end seismometer signal near the time of GW150914. An air compressor turns on at -75 seconds and off at +100 seconds, creating ground motion. This level of ground motion with a 175 second duration and a static frequency distribution was not capable of producing a 0.2 s chirp signal with the amplitude of GW150914 in the interferometer output.

a 0.2 s chirp signal with the amplitude of GW150914 in the interferometer output.

Chapter 5

IMC Upconversion

LIGO interferometers use several high finesse optical cavities for gravitational wave detection. The lengths of these cavities are controlled using radio frequency (RF) modulation-demodulation techniques in a Pound-Drever-Hall (PDH) locking scheme [65]. This scheme provides an error signal that is linear to cavity length over a specific range. This study examines the specific case of the triangular ring cavity uses in LIGO interferometers for input mode cleaning. When the length of the cavity approaches the boundaries of the PDH error signal linear range, our model of the input mode cleaner PDH response shows that the resulting error signal contains non-linear spectral artifacts. This study is done in an attempt to understand noise found in LIGO input mode cleaner during engineering runs in the commissioning phase of Advanced LIGO.

5.1 PDH locking

Resonance in an optical cavity is achieved when the round-trip length of the cavity is equal to an integer number of wavelengths of the input beam, that is,

$$L = N\lambda = \frac{Nc}{\nu} \quad (5.1)$$

where L is the round-trip length of the optical cavity, λ is the wavelength of the light, ν is the frequency of the light, and c is the speed of light. Under these conditions, the light circulating in the cavity will be in phase and add constructively, resulting in an optical gain that increases the intracavity power. This is the state in which the

LIGO optical cavities are intended to operate. If we invert this equation, the allowed frequencies of light for which resonance will occur is then

$$\nu = N \frac{c}{L}. \quad (5.2)$$

The spacing between these allowed frequencies is called the free spectral range,

$$\nu_{\text{FSR}} = \frac{c}{L}. \quad (5.3)$$

When the frequency of the light is equal to an integer multiple of the free spectral range, the system will be on resonance. This is, however, a delicate condition to maintain. If the frequency of the light changes while the cavity length is stable, the optical field will no longer overlap perfectly within the cavity and the incident light will be reflected. If the length of the optical cavity changes but the frequency is stable, the geometry of the cavity and the optical field will once again be mismatched and resonance will be lost. To solve this problem, LIGO employs feedback loops that use a PDH error signal to maintain the resonance condition. We will use the LIGO input mode cleaner as an example of PDH locking.

The Advanced LIGO input mode cleaner is a resonant triangular ring cavity used to isolate the TEM00 mode of the input beam. The geometry of the cavity is designed such that higher order modes of the optical field will be reflected and not transmitted to the rest of the interferometer. The carrier beam, however, will be resonant in the input mode cleaner and will be transmitted. To control the input mode cleaner, the reflected light incident on the input mode cleaner is read out on a photodiode. The reflected part of the carrier beam, which on resonance should be highly transmitted, is compared to the reflected part of an RF sideband which should be highly reflected.

The first necessary piece of information to generate the PDH error signal is the reflectivity of the optical cavity as a function of frequency. This function will have minima at integer multiples of the free spectral range, where the cavity is on resonance and light is circulating in the cavity. As the frequency of the light drifts, the reflectivity of the cavity will increase, rejecting more of the incident light. For the IMC, the reflection function is

$$F(\omega) = \frac{r(1 + e^{-i\phi})}{1 + r^2 e^{-i\phi}} = \frac{r(1 + e^{-i(\frac{\omega}{\nu_{\text{FSR}})})}}{1 + r^2 e^{-i(\frac{\omega}{\nu_{\text{FSR}})}}} \quad (5.4)$$

where ω is the frequency of the light, r is the reflection coefficient of the input mirror, ϕ is the round-trip phase accumulated as the light propagates through the cavity, and ν_{FSR} is the free spectral range of the cavity [66]. This function returns a complex value, the amplitude and phase represent the amplitude and phase of the reflected optical field relative to the optical field incident to the cavity.

Figure 24 shows the reflected amplitude and phase of the carrier beam and the 24 MHz RF sideband relative to the incident optical fields. For this demonstration, we will assume that the frequency of the light is stabilized and the x-axis represents a displacement of the optical cavity length from the resonance length. When the cavity length matches the input beam, the reflectivity is minimized and the carrier is transmitted through the IMC. The sideband, which is at a higher frequency, is by design not resonant in the cavity and is fully reflected. As the cavity length deviates from the resonance length, the reflected amplitude and phase of the carrier beam will change along with it. However, the reflected amplitude and phase of the RF sideband, which is not resonant in the IMC, are not sensitive to changes in cavity length around the resonance point.

Using the function for the complex reflection coefficient, the reflected light can be read out on a photodiode and used to generate an error signal that is linear to the length of the cavity within a certain range. Since a photodiode measures optical power, it will measure the square of the optical field, which is comprised of light at both the carrier and sideband frequencies. The signal on the photodiode measuring the reflected light is given as [65],

$$\begin{aligned}
P_{ref} = & P_c |F(\omega)|^2 + P_s \{ |F(\omega + \Omega)|^2 + |F(\omega - \Omega)|^2 \} \\
& + 2\sqrt{P_c P_s} \{ \text{Re}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)] \cos(\Omega t) \\
& + \text{Im}[F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)] \sin(\Omega t) \} \\
& + (\sin(2\Omega t) \text{ terms})
\end{aligned} \tag{5.5}$$

where F is the reflection function of the optical cavity, ω is the carrier beam frequency, $\omega + \Omega$ is the sideband frequency, P_c is the carrier beam power and P_s is the sideband power.

The first line of Equation 5.5 is the DC power in the beam, which is the carrier beam power and the sideband power multiplied by their respective reflection coefficients. The DC power acts as an overall offset in the photodiode signal and can be

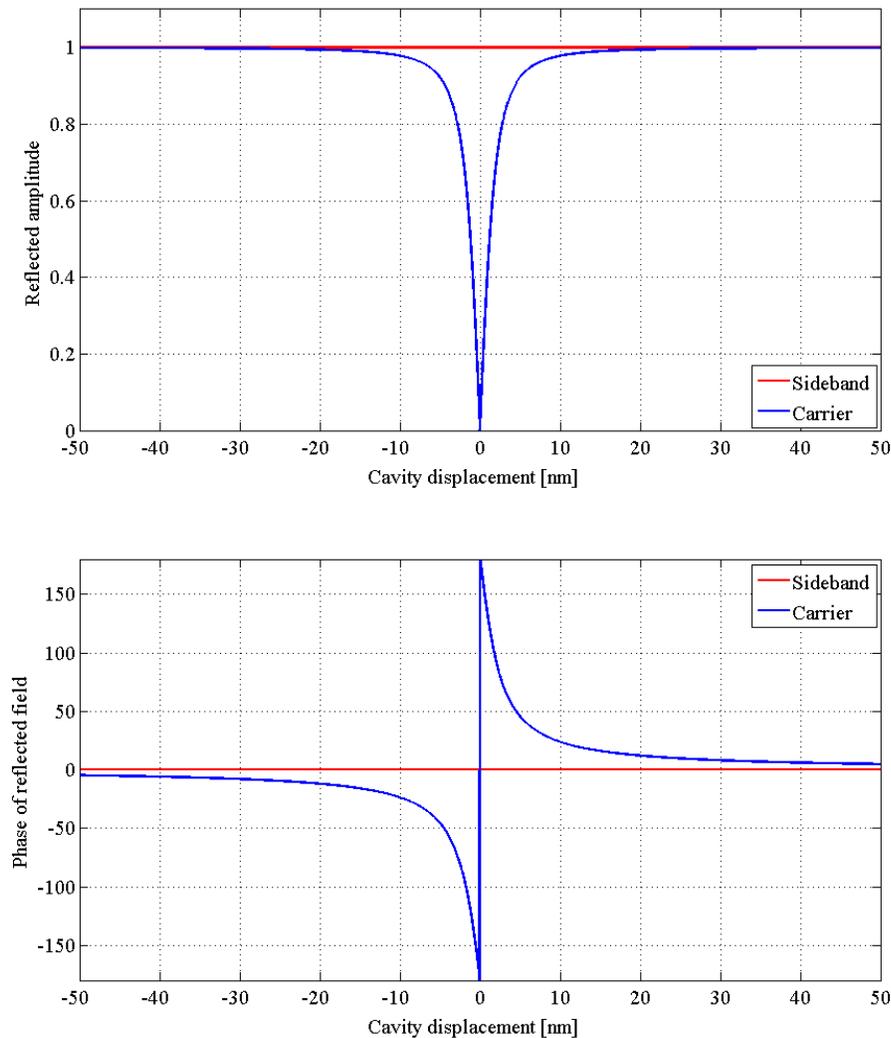


Figure 24: Amplitude and phase of light reflected from the IMC relative to the incident optical field. The zero point of the x-axis represents the resonance point of the cavity. The amplitude reflectivity is at a minimum when the cavity is on resonance, allowing the carrier beam to be transmitted into the interferometer. The amplitude and phase of the carrier beam will change as the cavity length changes. The amplitude and phase of the RF sideband are not sensitive to changes in the cavity length. This information can be used to generate an error signal that represents the length of the input mode cleaner.

subtracted off.

The second and third lines of Equation 5.5 are the most important. These two

lines represent the carrier field and the sideband field beating against each other. The beatnote formed by these two fields will have a frequency equal to the difference in frequency between the two fields, which is Ω . The information about the length of the optical cavity, which is derived from the relationship between the carrier field and sideband field, is encoded in this beatnote. Either the real or imaginary parts of the field can be extracted and used as an error signal; in this demonstration we will use the imaginary part of the field.

A demodulation process is used to extract this information from the photodiode signal. The photodiode signal is multiplied by a factor of $\sin(\Omega t)$ and then low-pass filtered. Since

$$\cos(\Omega t) \times \sin(\Omega t) = \frac{1}{2}\sin(2\Omega t), \quad (5.6)$$

the real part of the beatnote, which is proportional to $\cos(\Omega t)$, will only contain high frequency signal once it has been multiplied by $\sin(\Omega t)$ and will be removed by the low-pass filter. Since

$$\sin(\Omega t) \times \sin(\Omega t) = \frac{1}{2}(1 - \cos(2\Omega t)), \quad (5.7)$$

the imaginary part of the signal, which is proportional to $\sin(\Omega t)$, will have two terms, one which is at high frequency and will be removed by the low-pass filter and one which is multiplied by a scalar and is not oscillating. It is this non-oscillating signal, which has been generated by demodulating the beatnote between the carrier field and the sideband field, that is linear to the length of the optical cavity.

The higher order order terms in Equation 5.5, which vary as $\sin(2\Omega t)$, are generated by beating together the two RF sidebands and are removed by the low-pass filter.

In a situation where the carrier beam is resonant in the cavity and the RF sidebands are high enough in frequency that they are fully reflected, the PDH error signal is then [65],

$$\epsilon(\omega) = -2\sqrt{P_c P_s} \text{Im}\{F(\omega)F^*(\omega + \Omega) - F^*(\omega)F(\omega - \Omega)\}. \quad (5.8)$$

The PDH response of the cavity was modeled using measured values of optical reflectivity and free spectral range of the Livingston input mode cleaner. The input beam was the nominal LIGO carrier beam with a frequency of $\omega = 281.8$ THz ($\lambda = 1064$ nm) and modulation sidebands of $\Omega = \pm 24$ MHz. Figure 25 shows the resulting

PDH error signal as a function of the free spectral range with an overlaid straight line as a reference for linearity. Looking at a zoomed in view of the error signal around the linear part, we can see that the PDH signal matches the linear reference very well up to ± 0.5 nm, or $\sim \lambda/1000$, of cavity displacement.

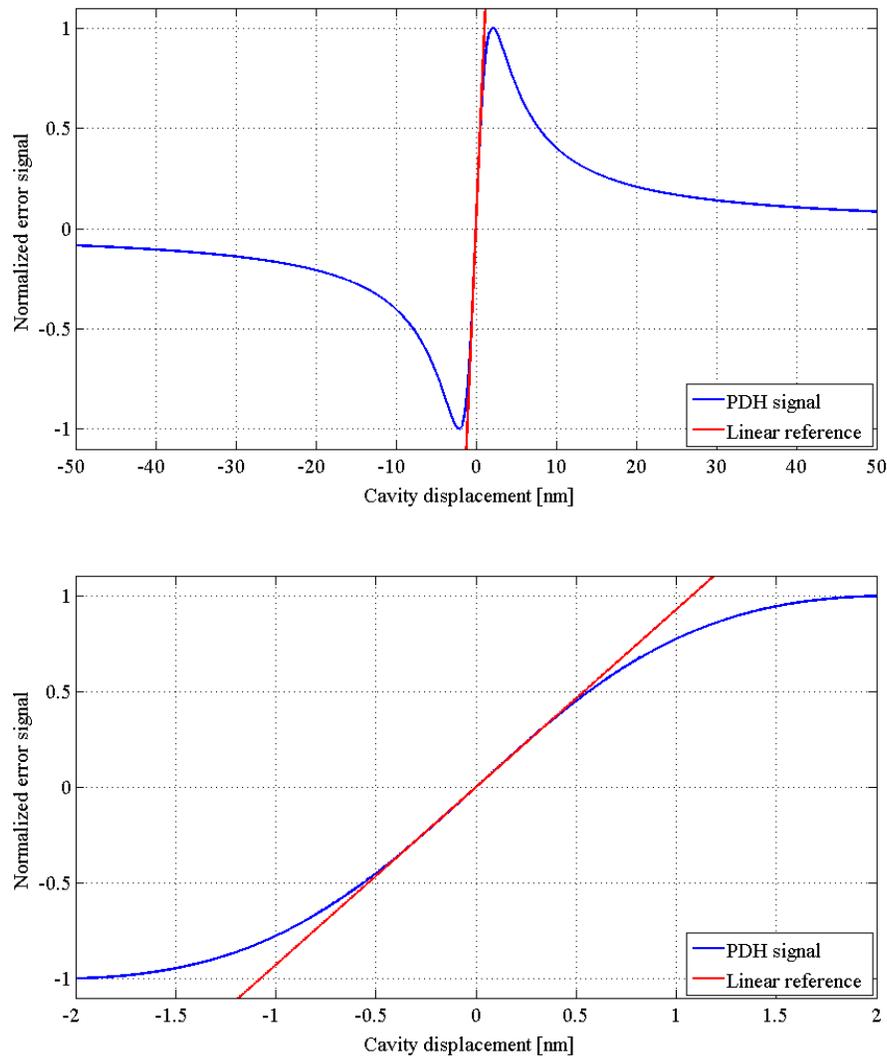


Figure 25: Example of a PDH error signal. The x-axis in this plot is linearly related to the length of the input mode cleaner. The red line is a straight line reference to estimate the linearity of the error signal. The error signal is linear to the length of the input mode cleaner up to ± 0.5 nm of cavity displacement, or $\sim \lambda/1000$. Motion beyond this point will begin to contain non-linear artifacts and eventually reach a turning point where control of the optics is lost.

If the cavity motion exceeds this linear range, the error signal will contain non-linear artifacts which will bleed into the control signal used to actuate on the cavity optics. To explore this non-linearity, we injected a sinusoidal cavity motion into our model and observed the resulting error signal. The frequency of the sine wave was selected in an attempt to model noise seen in the output of the interferometers.

I explored two specific cases in this study. Figure 26 shows the power spectral density of the injected sinusoidal cavity motion (green) and the resulting non-linear error signal (blue). This motion was injected asymmetrically about the nominal cavity locking point ($\epsilon = 0$). The effect of this non-linearity is to take the injected sine wave and produce an error signal that looks like a sine wave with a flattened top, resembling a mixture of a pure sine wave with a square wave. Thus, we see both even and odd harmonics of the injection frequency when the signal is observed in the frequency domain.

Figure 27 shows the power spectral density of the injected sinusoidal cavity motion (green) and the resulting non-linear error signal (blue). However, this time the motion was injected symmetrically about the nominal cavity locking point. The resulting error signal was similar to a square wave and as a result we only see odd harmonics of the fundamental frequency.

5.2 Upconversion noise in aLIGO

Each of the three mirrors in the input mode cleaner cavity is staged as the bottom mass of a triple suspension in order to passively isolate the mirrors from noise. In addition, the chambers holding the IMC mirrors are isolated from ground motion by two stages of active seismic isolation. This isolation, however, is not completely impervious to external excitations. During periods of time with excess ground motion, we can see seismic noise coupling into the cavity length and its control signal.

Specifically, when we see excess seismic noise in the 1-5 Hz anthropogenic band, believed to be caused by a commercial railroad a few kilometers from the LIGO Livingston, we see highly structured noise in the IMC control signal in the 10-100 Hz band. This physical mechanism is consistent with the model of a non-linear PDH error signal. If excess seismic motion reaches the suspension and the optics begin swinging around, it's feasible that they could start to saturate the linear range of

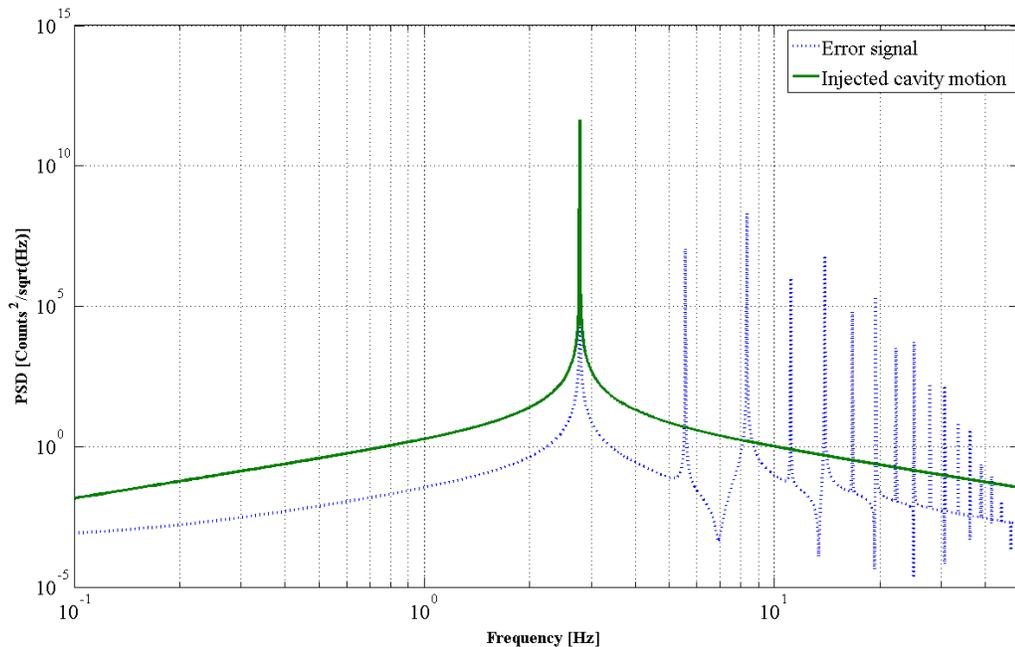


Figure 26: Sinusoidal cavity motion with frequency 2.78 Hz injected asymmetrically about the locking point of the cavity results in a PDH error signal containing non-linear spectral artifacts at harmonics of the injected cavity motion.

the PDH loop. Figure 28 shows the IMC length control signal during a time with excess seismic noise. The noise in the IMC control signal takes a form very similar in structure to the non-linear PDH signal, displaying strong odd harmonics and weaker even harmonics. The IMC control signal has an associated noise floor that obscures parts of these peaks. The theoretical model uses sinusoids with a highly specified frequency and thus displays very sharp peaks in its spectrum. It should be noted that the peaks in the IMC control signal are the manifestation of a physical process, not digitally generated, and have some natural width to them.

While we have demonstrated that this mechanism is consistent with IMC upconversion noise, it has not yet been fully proven. We are currently looking for a better way to look at the IMC error point, which is generated using an analog servo board, during times of excess seismic motion instead of the control signal. We think the source of the excitation may be a vertical resonance of the triple pendulum suspension that houses the IMC optics being rung up by the excess motion.

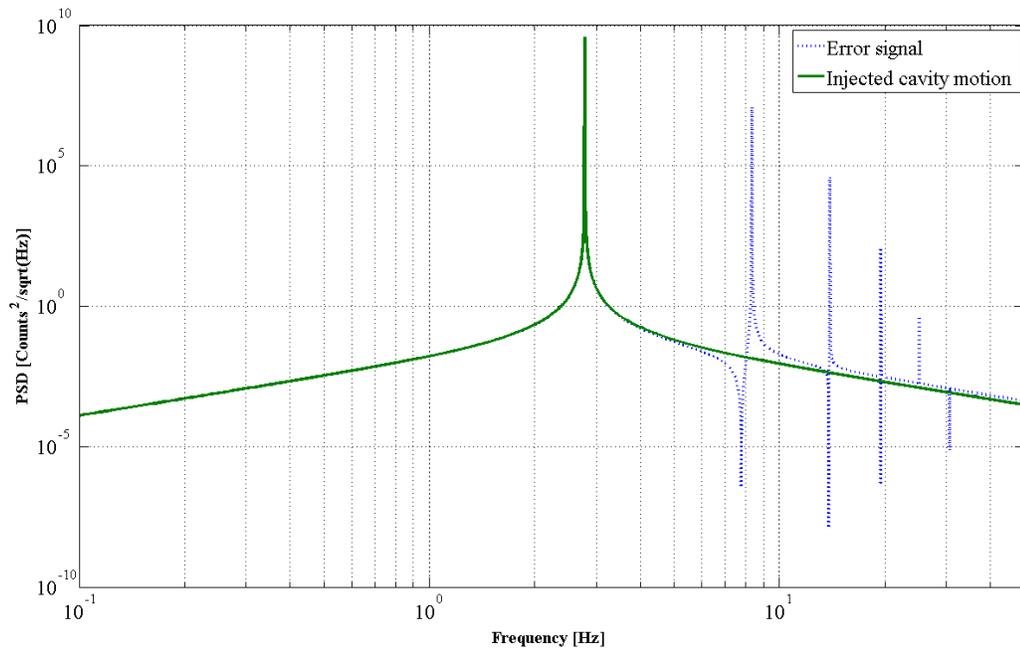


Figure 27: If the motion is symmetric about the cavity locking point, we see only odd harmonics of the injection frequency.

5.3 Conclusions

We found that injecting sinusoidal cavity motion into our input mode cleaner PDH model generates an error signal with non-linear spectral artifacts, specifically harmonics of the injection frequency, if the cavity motion exceeds the linear PDH range. For cavity motion that is symmetric about the locking point of the error signal, we find that the error signal contains only odd harmonics. For asymmetric cavity motion we find both even and odd harmonics, where the odd harmonics are typically higher in amplitude. In such a case, the amplitude of the even harmonics increases as the offset from the nominal locking point increases, that is, as the cavity motion is more asymmetric.

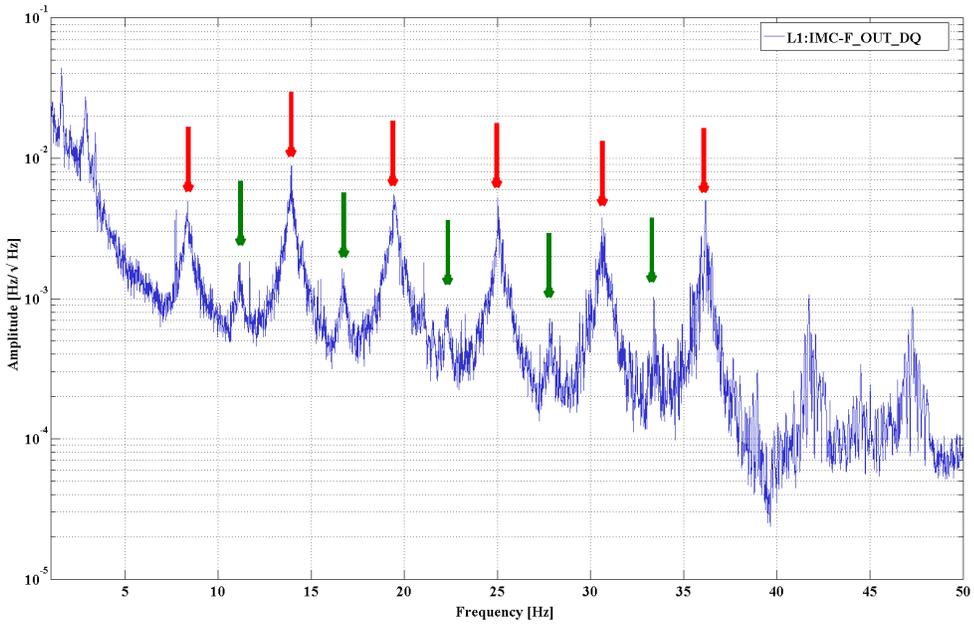


Figure 28: Spectral comb with a fundamental frequency of 2.78 Hz in the IMC control signal. Red arrows indicate odd harmonics, green arrows indicate even harmonics.

Chapter 6

Detector Characterization Subsystem Lead

The aLIGO interferometers are highly complex, high precision devices. Their operation depends on the careful interaction of a series of subsystems, each with its own purpose. In an effort to better understand the operation and output of the interferometers, the Detector Characterization group has been designed to mirror this subsystem approach. Table 4 lists the aLIGO subsystems. Each of these subsystems is assigned a data quality liaison from the DetChar group.

There are 5 main responsibilities assigned to a subsystem liaison. The first is to fully understand the operation and installation of the subsystem so that they can facilitate data quality investigations and act as a point of contact for commissioners assigned to this subsystem.

The second responsibility is to take this knowledge and use it to populate the channel information system (CIS), which is a database that stores information about how to parse and understand the various auxiliary channels that are monitored in each subsystem. This database also contains information about calibration and valid frequency ranges for these channels. This allows newcomers to the collaboration to more easily familiarize themselves with the LIGO naming conventions and facilitates their involvement in data quality investigations.

The third responsibility is to check for signal fidelity, which means to make sure that all of the channels are working as intended and don't contain artifacts from signal conditioning processes.

Subsystem	Description
LSC	Length Sensing and Control
ASC	Alignment Sensing and Control
SUS	Suspensions
IMC	Input Mode Cleaner
OMC	Output Mode Cleaner
PCAL	Photometric Calibration
PEM	Physical Environmental Monitoring
SEI	Seismic Isolation
PSL	Pre-Stabilized Laser
TCS	Thermal Compensation System

Table 4: Table describing aLIGO subsystems

The fourth responsibility is to develop summary pages that monitor important channels and figures of merit for each subsystem. The summary pages are generated every day from a configuration file designed by the subsystem liaisons. The purpose of the summary pages is to gather all of the potentially useful information about a subsystem in an organized way so that the subsystem leads can efficiently evaluate the performance of the subsystem. They are also a useful launch point for data quality investigations since they provide various overviews of instrumental performance ($h(t)$ spectrograms, Omicron triggers, CBC search sensitivity, etc.) that make it easy to identify persistent or egregious data quality issues.

The fifth and final responsibility is to develop and build real-time data quality monitors in the Online Detector Characterization (ODC) framework. The Online Detector Characterization (ODC) system is an infrastructure designed to extract and record metadata describing the state of the aLIGO interferometers. This state information has two main purposes: to inform data quality investigations and to serve as a real-time monitor of the interferometer state that can be accessed in the control room. Each subsystem is monitored using an ODC monitor.

The ODC system is unique in that it runs in real-time in the system that is used to control the LIGO interferometers. Each set of ODC monitors is built in

Simulink to directly interface with the models that control the interferometers. This has several distinct advantages. Since the monitors are run in real-time, they operate in parallel with the control loops that are sensing the various degrees of freedom of the interferometer and are able to achieve highly precise timing. The ODC monitors can also create their own test points, which means an ODC monitor can perform a check on any signal that exists in the front end at its full rate instead of relying on the information that is downsampled and stored on disk. These full rate test points operate at the full sample rate of the model (16384 Hz) and any information recorded in the ODC channel is written at the same rate. In contrast, many channels are only recorded at 16 Hz if they aren't accessed as a test point in the digital system.

The information generated by each ODC monitor can be extracted and sent to a segment generation process, where the most useful information is catalogued and represented by segments of time that indicate when a given condition was considered to be true or false.

Throughout the first observing run, I served as the detector characterization liaison for the Length Sensing and Control (LSC) and Alignment Sensing and Control (ASC) subsystems. I designed and built the ODC models that monitor these subsystems, designed and built the MEDM screens that are used to interact with the ODC models, and designed the summary pages that archive information about these subsystems. I also contributed to the study at the end of this chapter that used the MICH ODC channel as a witness for severe electronics noise.

6.1 Length Sensing and Control

The Length Sensing and Control (LSC) subsystem is used to monitor and control the lengths of the various optical cavities in the aLIGO interferometers. Figure 29 shows the layout of the aLIGO interferometer with the lengths of individual optical cavities labeled. The LSC subsystem is responsible for controlling 5 global degrees of freedom, which are linear combinations of these individual cavity lengths. Table 5 describes the degrees of freedom controlled by the LSC subsystem. The most important degree of freedom in the LIGO interferometers is the Differential Arm length (DARM), which is the degree of freedom that is sensitive to the differential stretching and squeezing of spacetime caused by a gravitational wave. The Common

Arm length (CARM) is the average arm length of the two interferometers. This degree of freedom is important for ensuring that both arms are held at their resonance length and is used to counteract effects such as slow drift from tidal forces. The Michelson (MICH) degree of freedom measures the differential arm length of the small Michelson interferometer formed by the two ITMs and the beamsplitter. This is an important degree of freedom because noise in this degree of freedom will result in excess light at the readout port of the interferometer. The final two degrees of freedom are the Power Recycling Cavity Length (PRCL) and the Signal Recycling Cavity length (SRCL). These are resonant cavities formed between the recycling mirrors and the ITMs. PRCL must be controlled to maintain a steady level of optical power in the interferometer, as modulating this length will modulate the power recycling gain in the interferometer. The signal recycling cavity forms a coupled cavity with the arms and must be controlled to maintain the frequency response of the interferometer. (These optical cavities are controlled using the Pound-Drever-Hall (PDH) technique described in Chapter 5.)

Degree of Freedom	Description
Differential Arm (DARM)	$L_x - L_y$
Common Arm (CARM)	$(L_x + L_y)/2$
Michelson (MICH)	$l_x - l_y$
Power Recycling Cavity Length (PRCL)	$l_p + (l_x + l_y)/2$
Signal Recycling Cavity Length (SRCL)	$l_s + (l_x + l_y)/2$

Table 5: Table of LIGO length degrees of freedom

6.1.1 Online Detector Characterization

The ODC model for the LSC subsystem is designed to monitor the feedback loops used to control optical cavities. A series of test points are placed in the control loop and compared to user-set threshold values to determine whether or not they are in their nominal range. Figure 6.1.1 shows the implementation of an LSC ODC model in SIMULINK. The numbered ovals on the left of the image are test point signals that are being read in from the higher level model that is used to control the lengths of

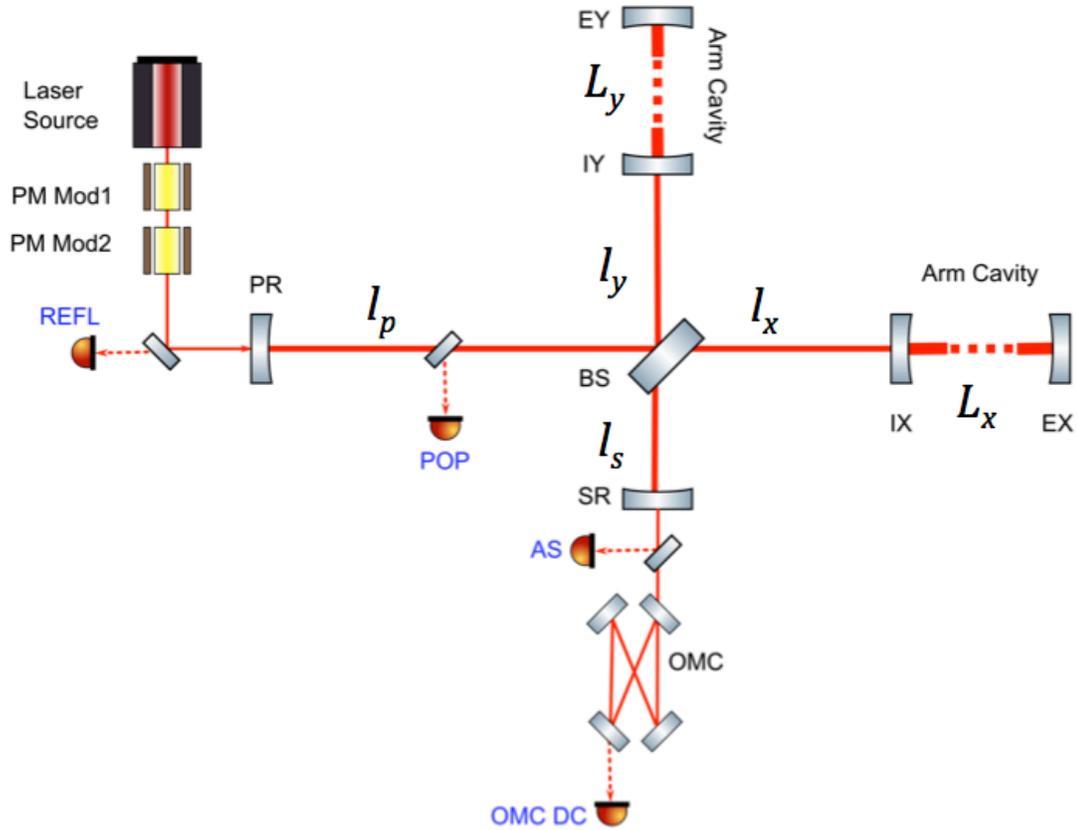


Figure 29: A simplified layout of a LIGO interferometer with the lengths of optical cavities labeled. These optical cavities combine to form 5 global degrees of freedom that are controlled by the Length Sensing and Control subsystem. Reproduced from [67].

the optical cavities. The signals are carried along the wires connecting each box. The green boxes are where the user set thresholds are saved. The white boxes perform operations on input signals.

As an example, we'll follow the path of inputs 46 and 47, which are signals from the photodiode that measures the reflected light at the input to the power recycling cavity. The signals are read in at the ovals labelled 46 and 47, their absolute values are calculated at the boxes labeled 'Abs19' and 'Abs20'. The absolute values of these signals are then fed into boolean comparison boxes, 'Operator39' and 'Operator40'. These boolean operator boxes are also connected to the green box which defines a threshold for the signals to be compared to. If the input signals are less than the designated threshold, the boolean operator passes a value of True. If they have exceeded the intended threshold, the boolean operator passes a value of False. The outputs of the two boolean operators are fed into one last check, which performs an AND operation. If both signals have passed their tests and reported True, the AND block reports a True and this photodiode signal is considered to be in a good state. If one or both of the signals has failed their test, the AND block reports a False and the photodiode signal is considered to be in a bad state.

Some of the checks implemented in the ODC models are more complicated. For example, when the states of multiple control loops are stored in a vector they can be compared to a series of state masks that select which degrees of freedom to check. In this way, the same vector of information can be used to perform hierarchical checks on the state of the instrument. One test can check that the core optics are performing nominally, a more broad test can include checks on both the core optics and recycling cavities, and then an overall test can be done to check that feedback loops are performing as intended. Each of these tests will report its own true or false answer.

Once the ODC model has performed all of its checks and reported a True or False answer, the information is stored in an overall state vector that can be parsed to learn the state of the LSC subsystem at any time. Figure 31 shows a visual representation of the state of the length degrees of freedom in the H1 interferometer over the course of a day. Each horizontal bar represents the state of a length degree of freedom as reported by ODC. In this particular day, the control signals for the Michelson (MICH) and signal recycling cavity (SRCL) degrees of freedom exceeded their nominal range

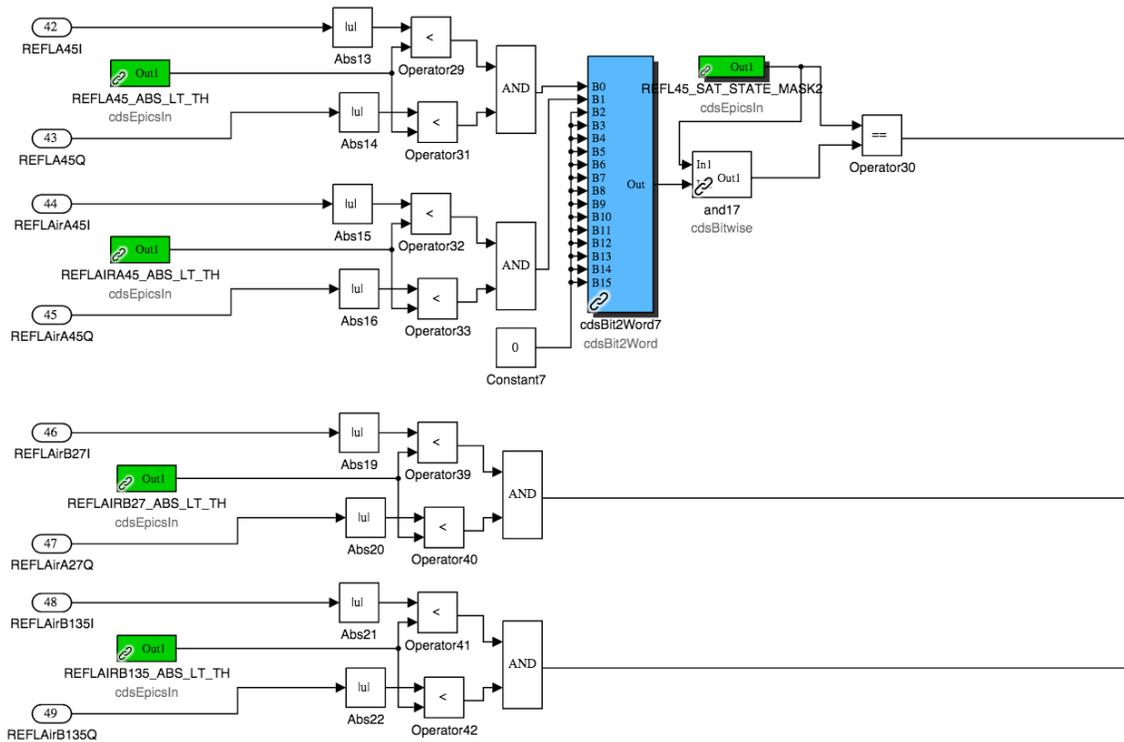


Figure 30: LSC ODC checks implemented in SIMULINK. The numbered ovals on the left indicate signals used in real-time control of the interferometer. The green boxes represent user-defined threshold values to be used in boolean comparisons. The white boxes represent operations such as computing the absolute value of a signal or performing boolean comparisons (less than, greater than, AND, OR, etc.)

while the interferometer was in its nominal operating state. This is indicated by the color of the bars switching between 6:00 - 8:00 UTC and between 14:00 - 16:00 UTC. All of the times when the state changes are recorded as time segments which can be used to correlate excursions in the cavity control signals with transients in the output of the interferometer.

6.1.2 MEDM screens

For real-time use of these monitors, a software package called MEDM is used to display and interact with the ODC models. MEDM can be used to update the thresholds and state masks used to determine the status of a given photodiode or degree of freedom.

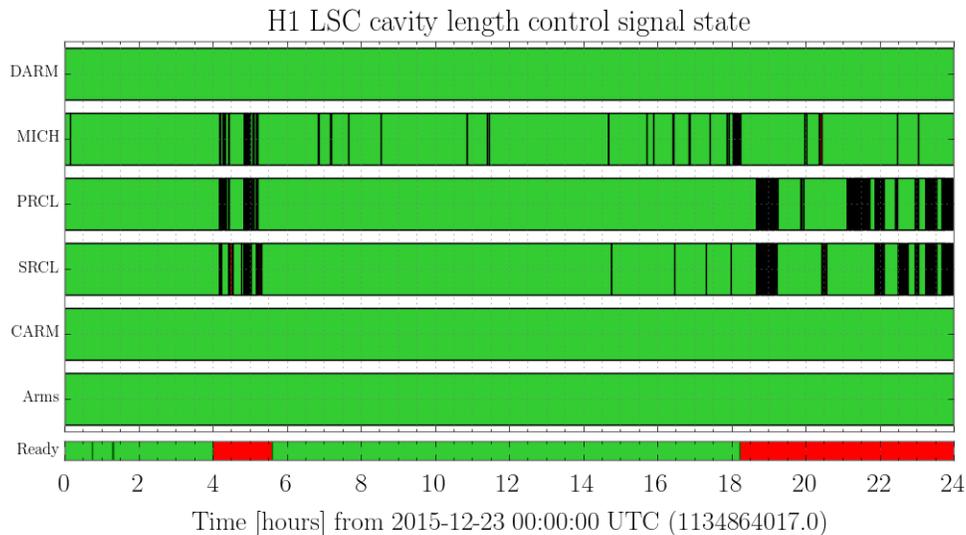


Figure 31: ODC bits representing states of length degrees of freedom. Each horizontal bar represents a length degree of freedom that is controlled in the LSC subsystem. When the bar is green, the control signal for that degree of freedom is in its nominal range. When the bar is not green, the control signal for that particular length degree of freedom has exceeded the threshold set in the ODC model and is reported as out of range.

Figure 32 shows the LSC ODC overview screen in MEDM. The top panel summarizes the overall state of the subsystem, showing the state of each ODC bit and a bitmask that indicates whether or not a given bit is used in determining the overall state of the subsystem. The leftmost panel is used to monitor the state of each length degree of freedom in the interferometer. The rest of the panels are used to monitor the states of the various photodiodes used for sensing length degrees of freedom. These include DC power monitors and the values of RF demodulated photodiode signals. The grey boxes containing numerical values indicate user-set thresholds that can be updated from this screen.

6.1.3 Summary pages

While the MEDM screens are useful for real-time readout of the ODC models, they do not have an easily accessible history. For this reason, summary pages were built

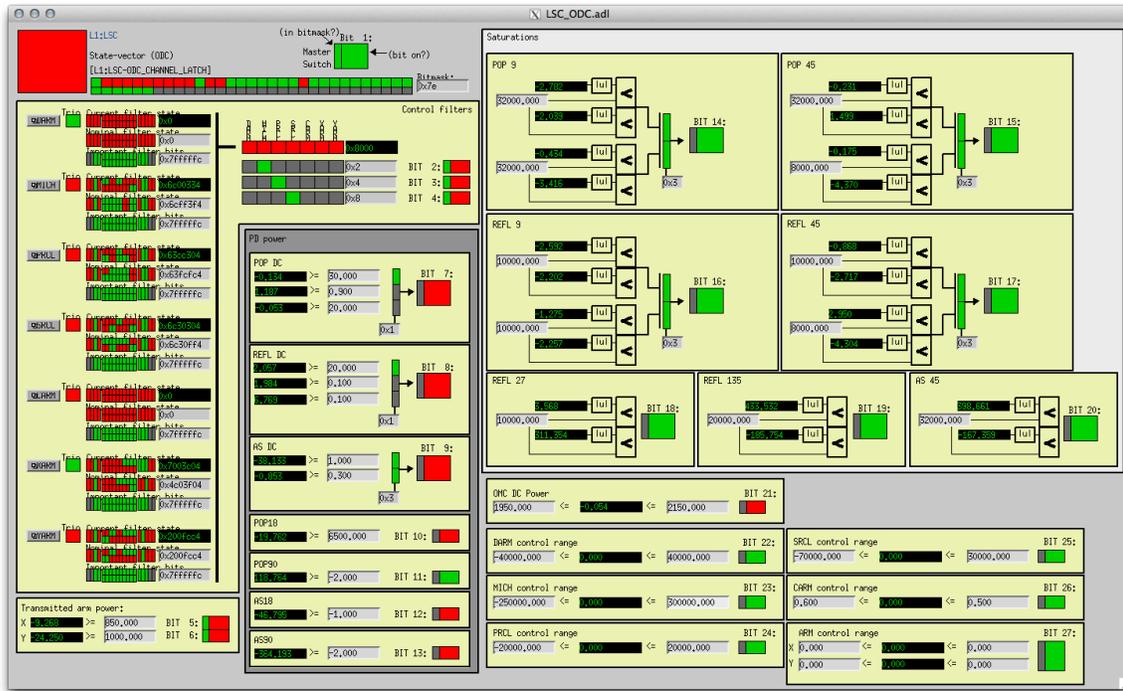


Figure 32: MEDM screen used to interact with the LSC ODC model. This screen contains information regarding the overall state of the LSC subsystem, the state of control loops pertaining to specific length degrees of freedom in the interferometer, and the state of photodiodes used to sense length degrees of freedom in the interferometer.

that contain the most important information from each ODC model. The summary pages are generated multiple times per day and are accessible through a web browser, which allows easy, organized access to past interferometer data when performing a data quality investigation.

Figure 31, which shows the status of the length degrees of freedom of the interferometer over the course of a day, was taken from the summary pages. In this figure, the MICH degree of freedom was seen to move into a bad state during a locked state. As an example, we can look at the summary page visualization of the MICH degree of freedom during this time. Figure 33 shows the control signal for the MICH degree of freedom with the ODC threshold overlaid as a dashed red line. The solid blue line indicates the median value of this signal over the course of 1 minute and the shaded regions indicate the maximum and minimum values over this same stretch of

time. When the control signal exceeds the ODC threshold, such as at about 14:35 UTC, the corresponding bit in Figure 31 flashes red to indicate the excess noise in this channel.

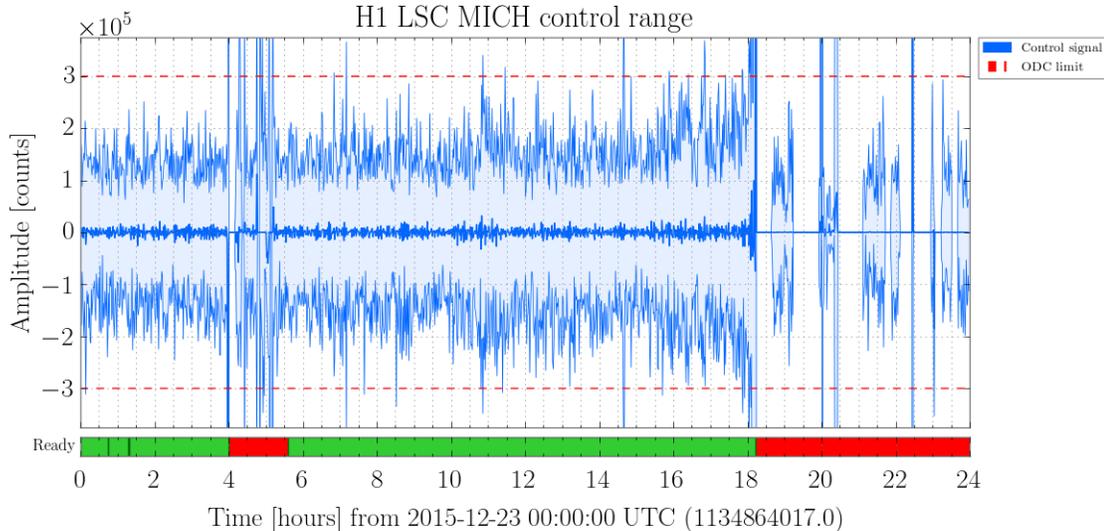


Figure 33: Readout of the MICH degree of freedom as displayed on the summary pages. The dark blue curve indicates the median value of the MICH control signal over the course of 1 minute. The shaded regions indicate the maximum and minimum values of the control signal over the same stretch of time. The dashed red line indicates the ODC threshold set to monitor this control signal.

6.2 Alignment Sensing and Control

The alignment sensing and control subsystem is used to control the alignment of optical cavities as well as the input pointing of light into those cavities. The control loops work in a similar way to the length sensing and control subsystem, using the reflected light from optical cavities to control the optics in a PDH scheme. The same set of RF sidebands that are used for length control are also used for alignment control. The detail that allows length fluctuations to be decoupled from alignment fluctuations is that alignment fluctuations generate higher order modes in the optical field, which are used to generate an error signal.

A length control loop compares the TEM00 mode of the carrier beam to the

TEM00 mode of the RF sidebands. An alignment control loop will compare the TEM00 mode of the carrier beam with the TEM10 and TEM01 modes of its sidebands, which are generated by angular misalignments of optical cavities. Since each optic has two alignment degrees of freedom that are directly controlled, pitch and yaw, and each cavity is comprised of multiple optics, the reflected light from each cavity is read out on a pair of quadrant photodiodes. This is visualized in Figure 35. The four quadrants allow the pitch and yaw error signals to be decoupled. Since higher order optical modes have a different Gouy phase as they propagate through space, a pair of photodiodes separated by a Gouy phase telescope are used to determine the origin of the misalignment [68].

6.2.1 Online Detector Characterization

Given that the alignment sensing control subsystem is designed similarly to the length sensing and control subsystem, the general layout of the ODC model is very similar. The photodiodes signals used to generate alignment error signals are checked against a saturation threshold. The control signals that are sent to the optics are checked to ensure that they aren't exceeding their nominal range. Each degree of freedom is represented as one bit in a state vector, which can be compared to a series of state masks to check for a series of valid states.

6.2.2 MEDM screens

The MEDM screens for the ASC subsystem are similar to those built to monitor the LSC subsystem. Figure 34 shows the ASC ODC overview screen in MEDM. The top panel once again describes the overall state of the subsystem and shows which ODC bits are used to determine that state. The left panel shows the status of the control signals used for each of the alignment degrees of freedom in both pitch and yaw. The bottom right panel, labeled 'QPD Saturations', checks for saturations in each of the quadrant photodiodes used in the ASC subsystem. Since there are many more checks that need to be made for the quadrant photodiodes, each one has a dedicated subscreen.

Since the ASC subsystem uses quadrant photodiodes, each quadrant must be checked for saturation. Figure 35 shows a photodiode monitor screen in the ASC

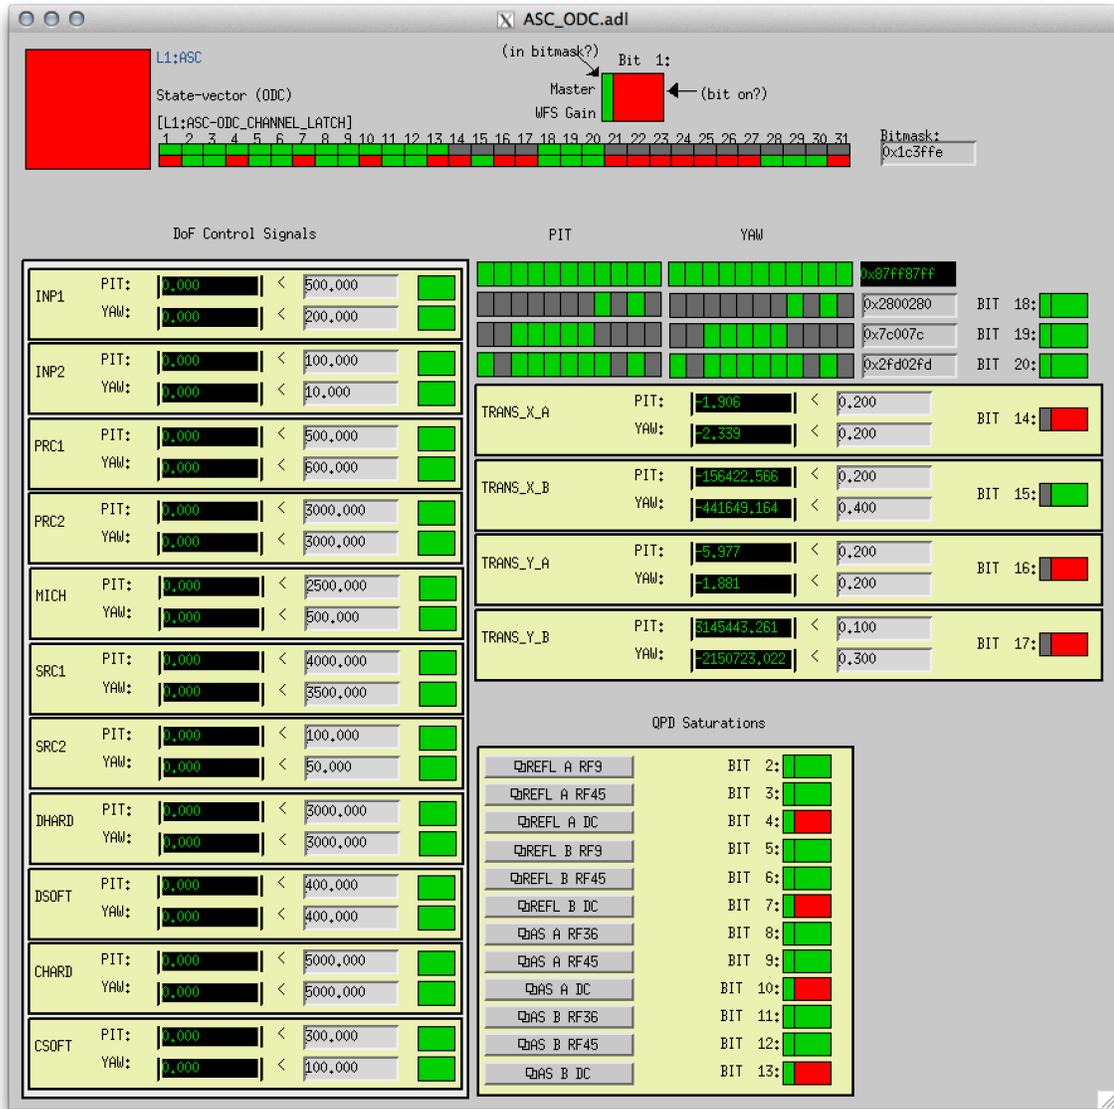


Figure 34: MEDM screen used to interact with the ASC ODC model. This screen shows the overall state of the ASC subsystem, the state of each alignment degree of freedom in the interferometer, and the state of each quadrant photodiode used to sense misalignments in the interferometer.

ODC. The first and second panels show the readouts of each quadrant of a quadrant photodiode in I- and Q-phase respectively. The absolute value of each signal is calculated and compared to a threshold value to see if any quadrants are approaching a saturation limit. The third and fourth panels perform checks on the associated

pitch and yaw readouts from these photodiodes to check for excursions beyond the nominal threshold. The nominal values for the pitch and yaw degrees of freedom are determined by trending these values over long durations of good interferometer performance.

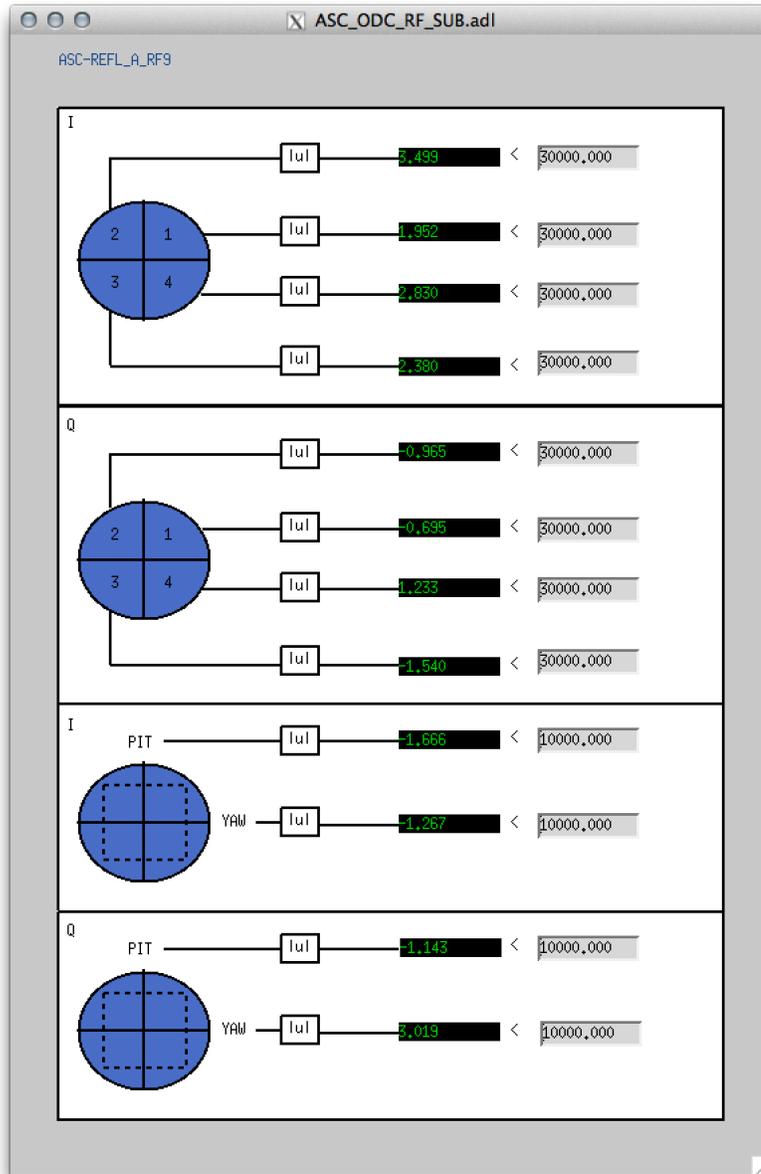


Figure 35: ODC monitor for ASC photodiode in MEDM

6.2.3 Summary pages

The ASC subsystem also has an accompanying set of summary pages that keep a running record of its state. These summary pages are designed similarly to the LSC summary pages, including time-series of the control signals for each alignment degree of freedom, saturation monitors for the ASC photodiodes, and ODC plots indicating the state of each degree of freedom.

6.3 ODC Results

The ODC system was designed and implemented in the engineering runs that preceded the first observing run. During the first observing run, the first efforts were made to use the information reported by the ODC models to flag and understand noisy data. The most successful test used the excess noise in the Michelson pitch degree of freedom to flag upstream electronics issues that caused loud glitches in $h(t)$. There is also evidence that the overall alignment status reported by the ASC ODC can be used as an early warning that the interferometer is going to drop out of its nominal operational state.

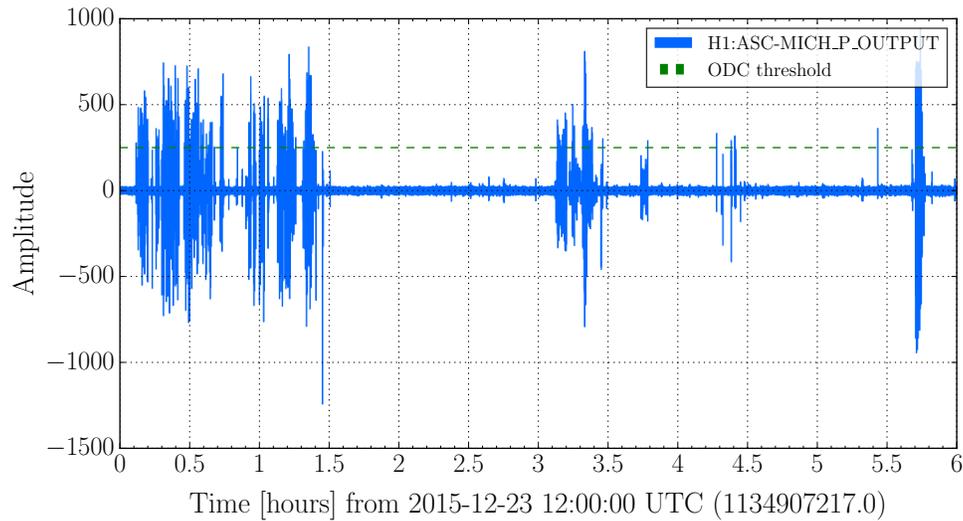
6.3.1 MICH ODC as a witness of RF45 glitches

The ODC channel built to monitor the Michelson (MICH) pitch degree of freedom was used to generate vetoes used in O1 analyses. Throughout O1, the H1 interferometer was prone to a glitch mechanism driven by malfunctions in RF electronics used to generate frequency sidebands on the carrier beam. These RF sidebands are used to control auxiliary degrees of freedom in the interferometer, including the length of the small Michelson interferometer formed by the beamsplitter and the two ITMs. When the RF electronics glitched, the error signals of these cavities would also glitch, causing excess motion in the auxiliary degrees of freedom that was witnessed by ODC monitors set up to monitor the control signal of the MICH alignment control loops.

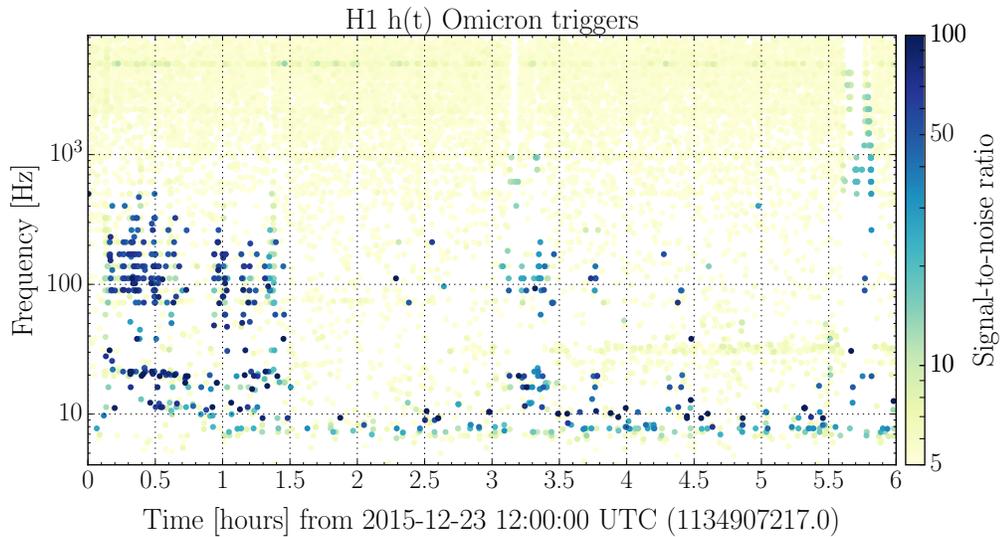
Figure 6.3.1 shows the correlation between the witness channel for this ODC channel and glitches in $h(t)$ as identified by Omicron. Figure 36a shows a time-series of the control signal of the MICH pitch control loop. The ODC threshold, set at a value of 250 for this particular channel, is indicated by the green dotted line. Any

time the control signal crosses this threshold, a time segment is created to indicate that the control loop is not in a nominal operating state. Figure 36b shows the $h(t)$ Omicron triggers over the same duration. When the MICH pitch control point has a high variance, for example in the first 1.5 hours of the plot, there is an overall increase in the rate of high SNR Omicron triggers, indicating that this ODC channel is witnessing alignment fluctuations that couple into the output of the interferometer.

This coupling can be quantified by removing these times from the output of the interferometer and calculating how efficiently this removal of time captures transient noise in $h(t)$. The time segments generated by this ODC channel are very efficient at vetoing high SNR Omicron triggers. Removing these segments of time from $h(t)$ removes Omicron triggers with $\text{SNR} > 8$ with an $\frac{\text{efficiency}}{\text{deadtime}}$ ratio of 47.16, indicating that a large number of high SNR Omicron triggers are removed from $h(t)$ while removing very little time from the analysis. These time segments, when used in the search, are called data quality vetoes. These vetoes were distributed to the CBC and Burst searches in O1 to indicate time that should not be analyzed. (Further discussion of data quality vetoes can be found in Chapter 7.



(a)



(b)

Figure 36: An example of an ODC channel witnessing RF electronics issues, which manifested as angular fluctuations in the vertex degrees of freedom at H1. Figure 36a shows the ODC threshold marking fluctuations in the MICH pitch degree of freedom. Figure 36b shows the associated Omicron triggers from $h(t)$ at the same time. The storms of loud triggers between 10 - 400 Hz are coincident with times flagged by this ODC monitor.

Chapter 7

Data Quality Vetoes

As discussed in Chapters 2 and 4, the data at the output of the LIGO interferometers contain non-Gaussian, transient noise artifacts. Chapter 4, 5, and 6 detail the efforts that have been made to understand and remove transient noise in the LIGO interferometers. When it was possible, transient noise sources were repaired at the source and the noise was not able to impact the output of astrophysical searches. Unfortunately, this is not always possible. If a source of transient noise can't be repaired, the noisy data will be processed by astrophysical search pipelines.

Transient noise artifacts are known to cause loud events at the output of astrophysical searches for gravitational waves, which manifest as a tail in the re-weighted SNR distribution as seen in Figure 7b. The data that comprise this tail of loudest events are the primary target of data quality investigations, such as those discussed in Chapter 4. If the noise sources can be linked to a systematic instrumental cause or a period of highly irregular instrumental performance, they can be flagged and removed from the analysis in the form of a data quality veto.

It is important to note that data quality vetoes are produced for all analysis time based on systematic instrumental conditions without any regard for the presence of gravitational wave signals. All data are treated equally; the removal of data with excess noise has the ability to remove real gravitational wave signals as well as background events. There are two types of vetoes implemented in the PyCBC search: category 1 and category 2.

7.1 Veto categories

Category 1 (CAT1) vetoes are intended to mark times when significant instrumental issues are present and the data should not be used in any analysis. CAT1 flags often indicate time when the character of the data has drastically changed and should not be combined with noise estimations from times of nominal performance. An example of this from O1 is an electronics failure that dramatically changes the character of the background noise and creates noise transients at a very high rate. As such, CAT1 vetoes remove time at the input to the PyCBC search pipeline. This ensures that severely problematic data are not used for background noise estimations and that no triggers will be generated at these times.

Category 2 (CAT2) vetoes are intended to mark short, noisy times that should not be treated as clean data. CAT2 flags are often used to flag transients that could potentially generate loud triggers, but do not corrupt the surrounding data badly enough that they need to be excluded at the input to the pipeline. An example of this is from O1 is a transient electronics saturation that only impacts the output data for 1 second. Times designated as CAT2 will still be used to compute background noise estimations for the matched filter search, but any triggers generated during those times will be excluded before background trigger distributions are calculated.

Further details on the application of CAT1 and CAT2 vetoes in the first observing run are available in a paper detailing the transient noise in the interferometers at the time of GW150914 [44].

7.2 Quantifying the effects of data quality vetoes

Does removing noisy data with data quality vetoes improve the output of the PyCBC search pipeline?

To test the effects of data quality vetoes, the PyCBC search pipeline was run multiple times with varying levels of noisy data removed. The first analysis, labeled “All vetoes applied”, used the full range of relevant data quality vetoes. The second analysis, labeled “No CAT2 applied”, omitted category 2 data quality vetoes. The third and final analysis, labeled “No CAT1 or CAT2 applied”, omitted all data quality vetoes. Gating is internal to the search pipeline and was applied in all of the analyses.

The bank of CBC waveform templates used in the PyCBC search is divided into three bins [12]. The significance of any candidate gravitational wave found in coincidence between the two interferometers is calculated relative to the background in its bin. Waveforms with different parameters will respond to instrumental transients in different ways. This binning is performed so that any foreground triggers are compared to a background generated from similar waveforms. As such, the effects of removing data from the PyCBC search are variable depending on which bin is considered. It should be noted that the actual gravitational wave signals discovered in the PyCBC search, GW150914 and GW151226, were part of a full search that was broken into 3 bins but reported as a single table of results. Because of this, their reported false alarm rates include a trials factor of 3. The comparisons made in Chapters 8 and 9 are done on a bin-by-bin basis, so the quoted rates have not been divided by 3. This is an overall factor of 3 change that does not affect the relative change in significance between the various analysis configurations.

The first bin is called the binary neutron star (BNS) bin and contains all waveforms with an $M_{chirp} < 1.74$. The second bin is the edge bin, which is defined based on the peak frequency (f_{peak}) of each CBC waveform. These are waveforms that are typically rather short in duration and are comprised of both binary black hole (BBH) and neutron star-black hole (NSBH) binary waveforms with high masses and anti-aligned effective spins. In the analysis containing GW150914, the edge bin was defined by $f_{peak} < 220$ Hz. In the analysis containing GW151226, the edge bin was defined by $f_{peak} < 100$ Hz. The third bin is the bulk bin, which contains all remaining waveforms needed to span the parameter space of the search. This contains BBH and NSBH waveforms with a variety of mass ratios and spins.

The following study, which is detailed in Chapters 8, 9, and 10 is my original research.

To study how the removal of noisy data affects the background distributions, we consider a hypothetical detection candidate at $\hat{\rho}_c = 11.3$, which corresponds to $\hat{\rho} = 8$ in each detector. Figure 37 compares the distribution of re-weighted SNR in Gaussian noise to that of real detector data. In Gaussian noise, the re-weighted SNR distribution is limited to $\hat{\rho} < 8$ for single detector triggers. A single detector trigger with $\hat{\rho} = 8$ would stand out against this background distribution as the loudest event in each detector. However, this is well within the region of the re-weighted SNR

distribution that is commonly obscured by instrumental transients and non-Gaussian features in the data. Figure 37 shows significant tails of loud triggers in each detector beyond $\hat{\rho} = 8$. It is this region, where the re-weighted SNR histograms contain non-Gaussian features, that data quality studies aim to improve. The hypothetical detection candidate at $\hat{\rho} = 8$ is a very useful threshold case: loud enough to be interesting, but quiet enough that its significance should be impacted by instrumental transients and data quality vetoes.

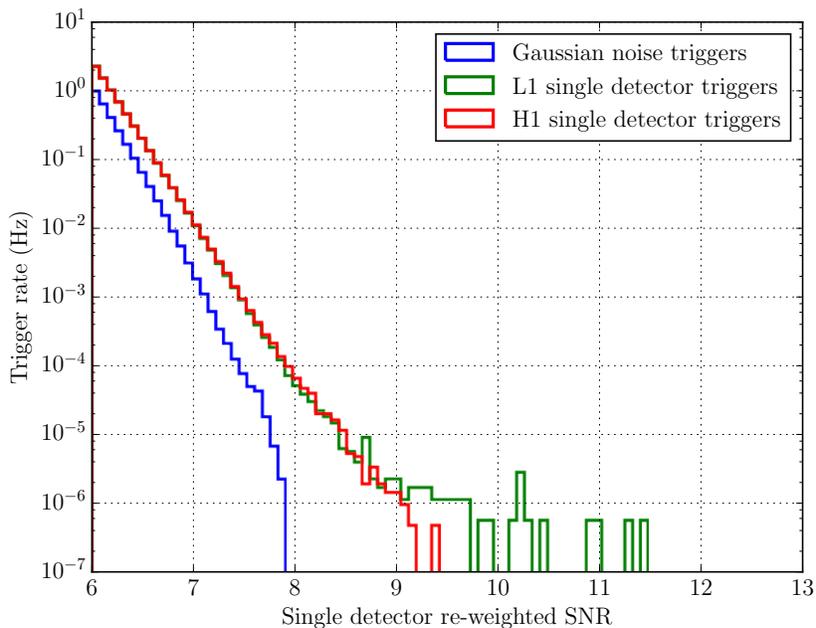


Figure 37: A rate histogram comparing the re-weighted SNR distribution in Gaussian noise to that of 16 days of real detector data. This histogram contains triggers generated over the entire template bank. In Gaussian noise, the re-weighted SNR distribution falls off before reaching $\hat{\rho} = 8$. The distributions from real detector noise show extensive tails beyond $\hat{\rho} = 8$ in addition to having an overall higher trigger rate. The tail of loud triggers in the L1 data is worse due to a higher rate of noise transients in the data, which are further discussed in Chapter 10.

Chapter 8

Effects of Data Quality Vetoes on the Analysis Containing GW150914

This analysis lasted from September 12 - October 20, 2015 and contained a total of 18.4 days of coincident detector data. After category 1 vetoes were applied, 17.9 days of coincident data remained. After category 2 vetoes were applied, 17.8 days of coincident data remained to be used in the final analysis. There were two interesting events that occurred in this analysis period. The first is GW150914, a gravitational wave signal from a binary black hole merger that marked the first direct detection of gravitational waves[46]. The second is a marginal detection candidate, LVT151012, which stands out from the background distribution but does not have enough statistical significance to be quoted as a confident detection.

The only major setting that was different between the analysis containing GW150914 and the analysis containing GW151226 was the gating threshold. As discussed in Section 2.0.5, gating is used at the input to the search pipeline to remove large transients from the data before the analysis is run. For the analysis containing GW150914, Omicron was used to build the list of gating windows. A conservative threshold was set at Omicron SNR > 300 to indicate a transient that should be gated. At each of these times a 1 second time window was constructed, centered around the time of the transient, that indicated the stretch of data to be removed from the analysis. This is the gating strategy that was employed in the analysis that included GW150914.

GW150914, which is considered to be a strong gravitational wave signal, was recovered by Omicron with an SNR of 13 and 9 at Hanford and Livingston respectively, well beneath the gating threshold.

8.1 BNS bin

The BNS bin already has such clean statistics that it does not benefit from the removal of data with excess noise. Binary neutron star systems have the longest waveforms in the template bank, often spanning up to 60 seconds in duration. With such long waveforms, the χ^2 test is effective at reducing the impact of transients on the BNS search. Typical instrumental transients have a small number of cycles and a duration of less than 1 second. As such, the overlap between a transient and a BNS signal is a small fraction of the total duration of the BNS waveform and is easily distinguished as noise in the re-weighted SNR calculation. This is demonstrated in Figure 38, which shows the distribution of single detector triggers in SNR and re-weighted SNR. The tail of high SNR triggers is fully down-weighted, resulting in a re-weighted SNR distribution that extends to $\hat{\rho} \approx 8.3$.

Since the χ^2 test is so effective in this bin, it is rare to see strong outliers in the re-weighted SNR distribution. Figure 39 shows the background distribution of the BNS bin in the PyCBC search for the analysis containing GW150914. The cumulative rate of background events in a given bin indicates the rate of false alarms expected in that bin for a given re-weighted SNR. In this bin, there is no substantial improvement for any value of $\hat{\rho}_c$. In both cases, the loudest background event is reported at $\hat{\rho}_c < 11$. The hypothetical detection candidate at $\hat{\rho}_c = 11.3$ would be the loudest event in this bin whether or not noisy data are removed. See Sections 8.2 and 8.3 for a contrary case.

Table 6 shows the change in false alarm rate and probability at $\hat{\rho}_c$ for multiple tiers of data removal in the BNS bin. The false alarm rate and probability both decrease if no data are removed from the analysis, indicating that the hypothetical detection candidate appears more significant in the presence of noisy data. Since there are few loud triggers to be vetoed, the overall effect of applying data quality vetoes is to reduce the overall analysis time, which restricts the significance that can be assigned to a loud detection candidate.

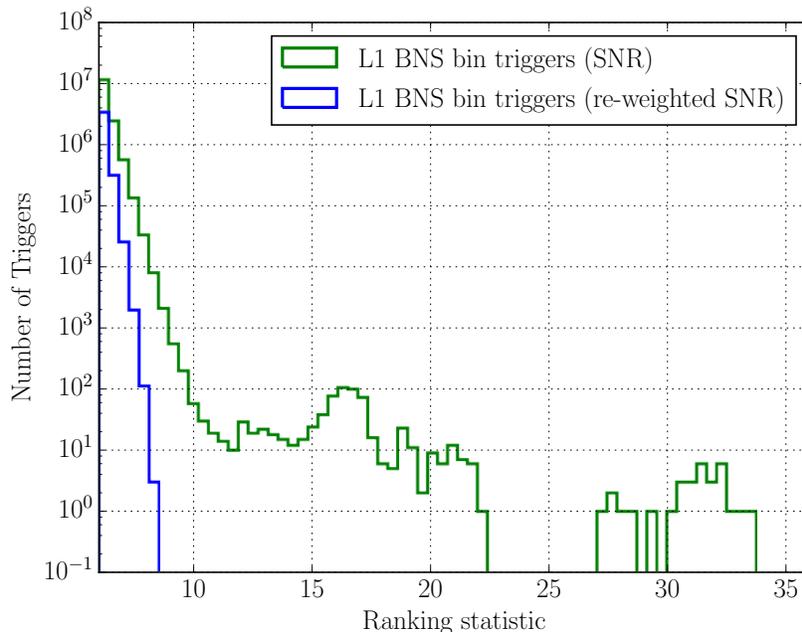


Figure 38: Histograms of Livingston (L1) single interferometer triggers found in the BNS bin. The green curve shows the distribution of BNS bin triggers in SNR and the blue curve shows the distribution of BNS bin triggers in re-weighted SNR. The tail of high SNR triggers have all been down-weighted by the χ^2 test, leaving behind a re-weighted SNR distribution that has a shoulder at just over $\hat{\rho} = 8$. The total number of triggers in each histogram is different, which is an artifact of the χ^2 test down-weighting some triggers so severely that they end up at $\hat{\rho} < 6$. When the analysis is run, triggers found with $\rho < 6$ are discarded because of their low significance.

8.2 Bulk bin

The bulk bin benefits much more from the application of data quality vetoes compared to the BNS bin. Figure 40 shows the background distribution in the bulk bin for the analysis containing GW150914. The first noticeable change is that the loudest background event is at $\hat{\rho}_c = 12.6$ in the presence of noisy data compared to 11.5 when all data quality vetoes are applied. This new loudest event does not show up as a small outlier; there is a significant shoulder in the distribution that persists up to $\hat{\rho}_c = 12$ before falling off. Considering the two distributions as a whole, there is a separation between the two curves beginning at $\hat{\rho}_c = 9$, which reaches an order of

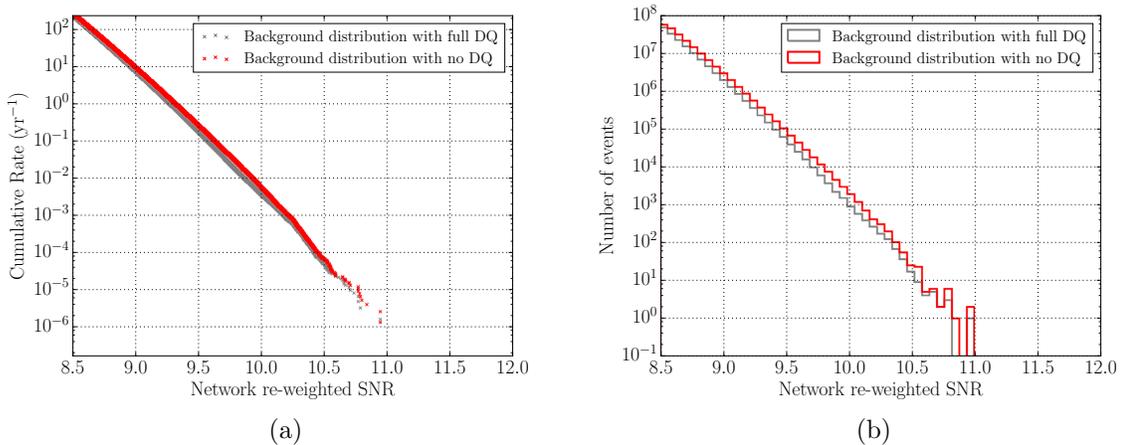


Figure 39: The background distribution in the BNS bin before and after applying data quality (DQ) vetoes. (39a) The cumulative rate of background triggers in the BNS bin as a function of re-weighted SNR. (39b) A histogram of background triggers in the BNS bin. The red traces indicate the distribution of background triggers without noisy data removed, the gray traces indicate the distribution of background triggers with all data quality vetoes applied. The BNS bin shows minimal improvement in cumulative rate. A foreground event at $\hat{\rho}_c = 11.3$, representing the hypothetical detection candidate, would be the loudest event in this bin.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	1.65×10^{-6}	7.22×10^{-8}
No CAT2 applied	1.63×10^{-6}	7.20×10^{-8}
No CAT1 or CAT2 applied	1.33×10^{-6}	6.50×10^{-8}

Table 6: Table of BNS bin false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. Removing noisy data has a negligible effect on false alarm rates in the BNS bin. The small differences in each configuration are due to data quality vetoes changing the total amount of time used in background estimation.

magnitude discrepancy at roughly $\hat{\rho}_c = 10$ and continues to diverge at higher values of network re-weighted SNR.

To quantify the difference between these two distributions, the hypothetical detection candidate at $\hat{\rho}_c = 11.3$ is once again considered. In the analysis with noisy data removed, the hypothetical detection would be amongst the loudest events in the distribution and with a false alarm rate of 3.2×10^{-6} as seen in Table 7. If data quality vetoes are not applied, the false alarm rate and false alarm probability increase by a factor of 1188, severely diminishing the statistical significance of such an event. Table 7 also shows that the majority of the improvement in this bin is the result of applying CAT1 vetoes. The difference between an analysis with all vetoes applied and with CAT2 vetoes removed is negligible in the bulk bin.

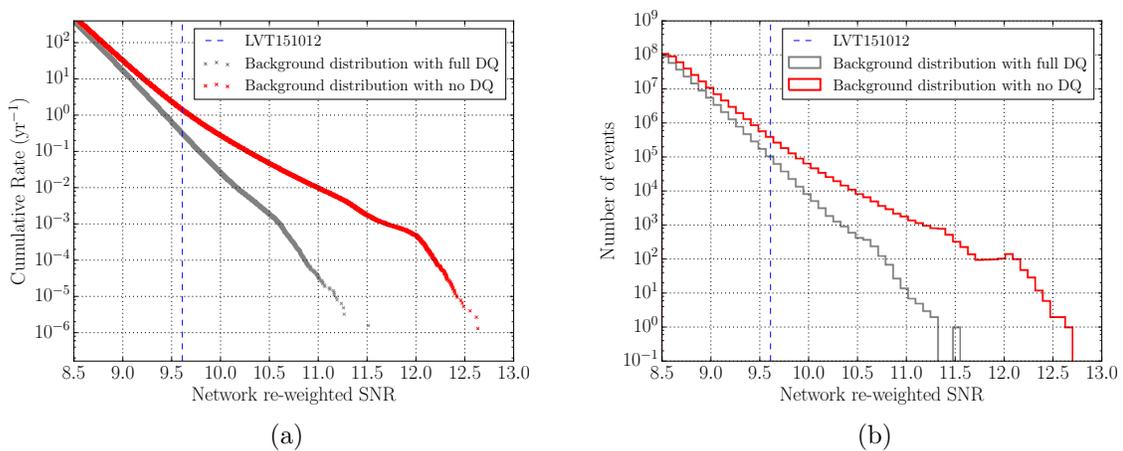


Figure 40: The background distribution in the bulk bin before and after applying data quality (DQ) vetoes. (40a) The cumulative rate of background triggers in the bulk bin as a function of re-weighted SNR. (40b) A histogram of background triggers in the bulk bin. The red traces indicate the distribution of background triggers without noisy data removed and the gray traces indicate the distribution of background triggers with all data quality vetoes applied. The bulk bin sees significant improvement when noisy data are removed, with over 3 orders of magnitude of improvement at $\hat{\rho}_c = 11.3$. The dotted line indicates the network re-weighted SNR of LVT151012, an interesting candidate that was not statistically significant enough to be claimed as a detection. The significance of LVT151012 improved by a factor of 7.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	3.29×10^{-6}	1.44×10^{-7}
No CAT2 applied	3.27×10^{-6}	1.44×10^{-7}
No CAT1 or CAT2 applied	3.91×10^{-3}	1.91×10^{-4}

Table 7: Table of bulk bin false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. The effects of removing data with excess noise are considerable in the bulk bin. The hypothetical detection candidate shows over 3 orders of magnitude of improvement in false alarm rate when all vetoes have been applied.

8.2.1 LVT151012

The second most significant trigger in the analysis containing GW150914 was LVT151012, recorded on October 12, 2015. The trigger was recovered in the bulk bin with $\hat{\rho}_c = 9.6$, which lies in the region where data quality vetoes are expected to begin having a significant impact on false alarm rates, as discussed in Chapter 7.

LVT151012 provides an interesting test case for this study. It had a false alarm probability of 2%, making it an interesting candidate but not statistically significant enough to be claimed as a detection. Though slightly quieter, it is analagous to the hypothetical detection candidate considered in Chapter 7.

The application of data quality vetoes increased the significance of LVT151012. The false alarm rate and probability both decrease by a factor of 7 when data quality vetoes are applied, as shown in Table 8.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	0.44	2.03×10^{-2}
No CAT2 applied	0.51	2.35×10^{-2}
No CAT1 or CAT2 applied	3.09	0.14

Table 8: Table of bulk bin false alarm rates and probabilities for LVT151012 using several different analysis configurations. The false alarm rate and probability of LVT151012 were reduced by a factor of 7 when noisy data were removed from the analysis.

8.3 Edge bin

The edge bin is significantly impacted by the application of data quality vetoes, although in a different way than the bulk bin. Figure 41 shows the background distribution in the edge bin before and after data with excess noise have been removed from the analysis. If noisy data are not removed from the analysis, there is a noticeable extension of the tail of loudest events. The loudest background event with no data removed from the analysis is at $\hat{\rho}_c = 19.3$ compared to $\hat{\rho}_c = 13.3$ when all vetoes are applied. When data with excess noise is removed, any trigger with $\hat{\rho}_c > 13.3$ would be the loudest event in the analysis with a false alarm rate of $O(10^{-6}\text{yr}^{-1})$. When no data are removed, the region between $\hat{\rho}_c = 13.3$ and $\hat{\rho}_c = 18$ is constrained to a false alarm rate of $O(10^{-3}\text{yr}^{-1})$ due to the long tail of background triggers.

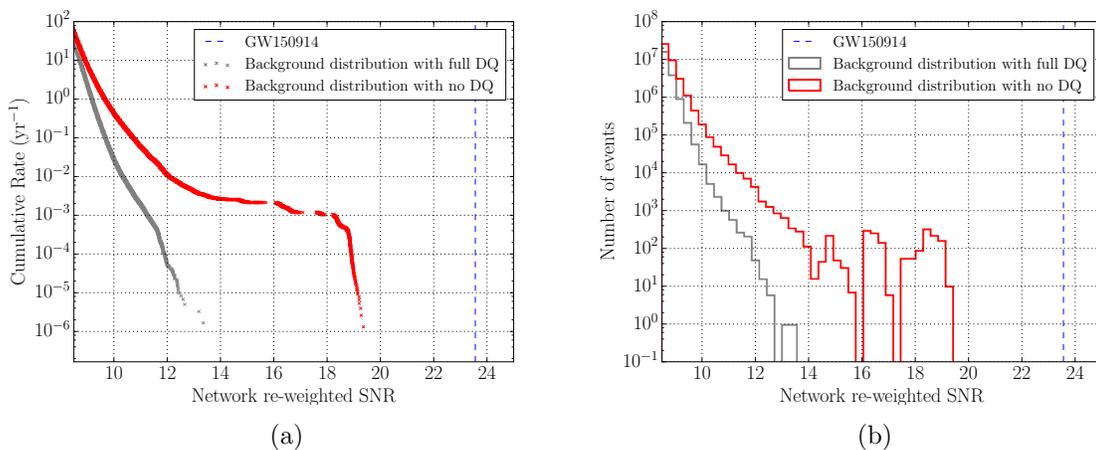


Figure 41: The background distribution in the edge bin before and after applying data quality (DQ) vetoes. (41a) The cumulative rate of background triggers in the edge bin as a function of re-weighted SNR. (41b) A histogram of background triggers in the edge bin. The red traces indicate the distribution of background triggers without noisy data removed from the analysis and the gray traces indicate the distribution of background triggers with all data quality vetoes applied. The dotted line indicates GW150914, which was recovered with $\hat{\rho}_c = 23.56$. When noisy data are not removed, the tail of the distribution extends to $\hat{\rho}_c = 19.5$, which severely impacts the ability to make a significant detection in the $\hat{\rho}_c > 13.3$ region. In both cases, GW150914 is the loudest event in this bin.

Consideration of the hypothetical detection candidate at $\hat{\rho}_c = 11.3$ reveals an order of magnitude improvement in false alarm rate and probability when data quality vetoes are used to remove data. This is not as dramatic as the improvement shown in the bulk bin, but it is still significant. This improvement is quantified in Table 9. It is interesting to note that both the bulk and the edge begin to show improvement around $\hat{\rho}_c = 9$ and show very similar levels of improvement at $\hat{\rho}_c = 10$, but their behavior begins to diverge after this point.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	9.13×10^{-4}	4.00×10^{-5}
No CAT2 applied	2.24×10^{-3}	9.84×10^{-5}
No CAT1 or CAT2 applied	3.74×10^{-2}	1.82×10^{-3}

Table 9: Table of edge bin false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. The false alarm rate is decreased when noisy data are removed, resulting in a factor of 40 improvement when data quality are applied. CAT2 vetoes have a greater impact in this bin, though they are still not as impactful as CAT1.

8.3.1 GW150914

The gravitational wave signal GW150914, produced from the inspiral and merger of a binary black hole system, was detected on September 14, 2015 and was recovered by the PyCBC search with $\hat{\rho}_c = 23.6$ [46]. GW150914 was louder than any background event in the analysis regardless of what data were considered. This being the case, data quality vetoes do not improve the false alarm rate for GW150914. It should be noted that GW150914 was an exceptionally loud event that sits well above the search background and is not the type of event that is expected to benefit from data quality vetoes. This is quantified in Table 10.

Analysis configuration	False alarm rate (yr^{-1})	False alarm probability
All vetoes applied	4.94×10^{-6}	2.17×10^{-7}
No CAT2 applied	4.90×10^{-6}	2.16×10^{-7}
No CAT1 or CAT2 applied	3.99×10^{-6}	1.95×10^{-7}

Table 10: Table of edge bin false alarm rates and probabilities for GW150914 for several analysis configurations. GW150914 is loud enough that its false alarm probability is not strongly affected by removing noisy data from the analysis. The slight change seen in each column is due to each analysis configuration allowing different amounts of analysis time.

Chapter 9

Effects of Data Quality Vetoes on the Analysis Containing GW151226

The extended analysis containing GW151226 lasted from December 3, 2015 - January 19, 2016 and contained a total of 16.7 days of coincident detector data. After category 1 vetoes were removed, 16 days of coincident data remained. After category 2 vetoes were removed, 15.8 days of coincident remained and were used in the final analysis. This analysis time provided an extended background estimation for the binary black hole merger GW151226 [47], which was detected by the aLIGO interferometers on December 26, 2015.

The gating process used in the analysis containing GW151226 was different than that used in previous analyses. Instead of relying on an external process, such as Omicron, to identify large transients, an internal test is used to find large excursions in the input data [33]. The time domain input data are Fourier transformed into the frequency domain and whitened using the measured power spectral density. The data are then inverse Fourier transformed back into the time domain and compared to a threshold value set to perform comparably to the Omicron gating process. If the whitened data have excursions that exceed this threshold, a gating window is constructed to remove these data from the input to the search. Tests were done to ensure that loud signals would not be removed by the auto-gating process.

9.1 BNS bin

The BNS bin shows a slight improvement when data quality vetoes are included. Figure 42 shows the background distributions in the BNS bin before and after removing data with vetoes. The most significant effect is that the loudest event in the background is at a higher re-weighted SNR; the loudest background event is at $\hat{\rho}_c = 10.96$ rather than $\hat{\rho}_c = 10.6$.

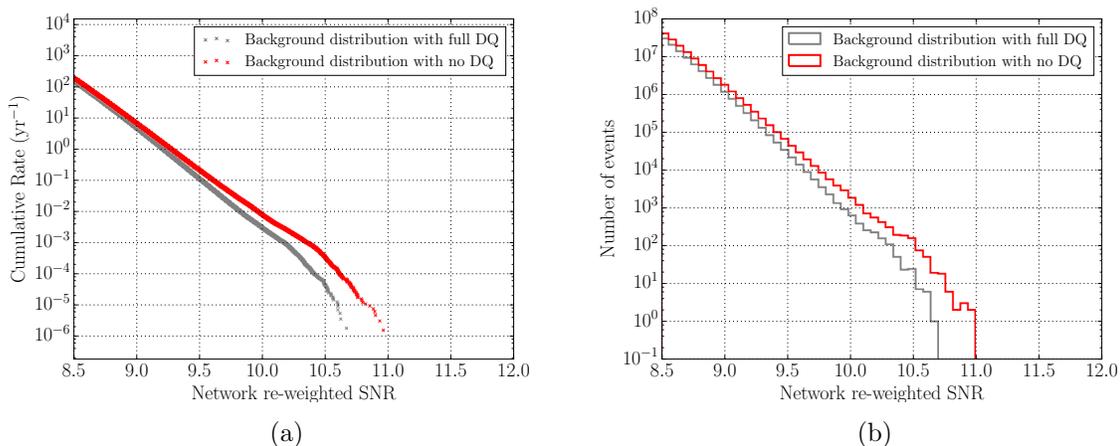


Figure 42: The background distribution in the BNS bin before and after applying data quality (DQ) vetoes. (42a) The cumulative rate of background triggers in the BNS bin as a function of re-weighted SNR. (42b) A histogram of background triggers in the BNS bin. The red traces indicate the distribution of background triggers without noisy data removed and the gray traces indicate the distribution of background triggers with all data quality vetoes applied. The BNS bin shows minimal improvement. With noisy data removed, the loudest background event is at $\hat{\rho}_c = 10.6$. Without removing any data, the loudest background event is at $\hat{\rho}_c = 10.96$.

Although the removal of problematic data has a visible impact on the background of the BNS bin, the resulting distribution is still not limiting to the search. In both cases, the hypothetical detection candidate at $\hat{\rho}_c = 11.3$ would be the loudest event in the analysis. The difference in false alarm rate is quantified in Table 11.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	1.82×10^{-6}	7.60×10^{-8}
No CAT2 applied	1.77×10^{-6}	7.48×10^{-8}
No CAT1 or CAT2 applied	1.63×10^{-6}	7.18×10^{-8}

Table 11: Table of false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. The difference in false alarm rate between the different configurations is negligible in the BNS bin.

9.2 Bulk bin

The bulk bin benefited from the application of data quality vetoes. Figure 43 shows the bulk bin background distribution before and after data quality vetoes applied. The first notable effect is that the loudest background event is at $\hat{\rho}_c = 14.8$ rather than $\hat{\rho}_c = 11.8$. This effect limits the values of $\hat{\rho}_c$ for which a significant detection could be claimed. The second effect is the visible separation between the two curves, indicating an increase in false alarm rate for any trigger with $\hat{\rho}_c > 9$. An example of this is quantified by once again considering a hypothetical detection candidate at $\hat{\rho}_c = 11.3$. At this value of $\hat{\rho}_c$, there is a factor of 800 reduction in false alarm rate when data quality vetoes are applied, as detailed in Table 12. The application of CAT2 vetoes has a positive impact on this bin, providing a factor of 54 reduction in false alarm rate compared to an analysis where only CAT1 vetoes are applied.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	3.10×10^{-5}	1.29×10^{-6}
No CAT2 applied	1.70×10^{-3}	7.23×10^{-5}
No CAT1 or CAT2 applied	2.48×10^{-2}	1.10×10^{-3}

Table 12: Table of bulk bin false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. At $\hat{\rho}_c = 11.3$, the false alarm rate decreases by a factor of 800 when data with excess noise is removed from the analysis. Category 2 vetoes are more effective in this bin, providing a factor of 54 reduction in false alarm rate.

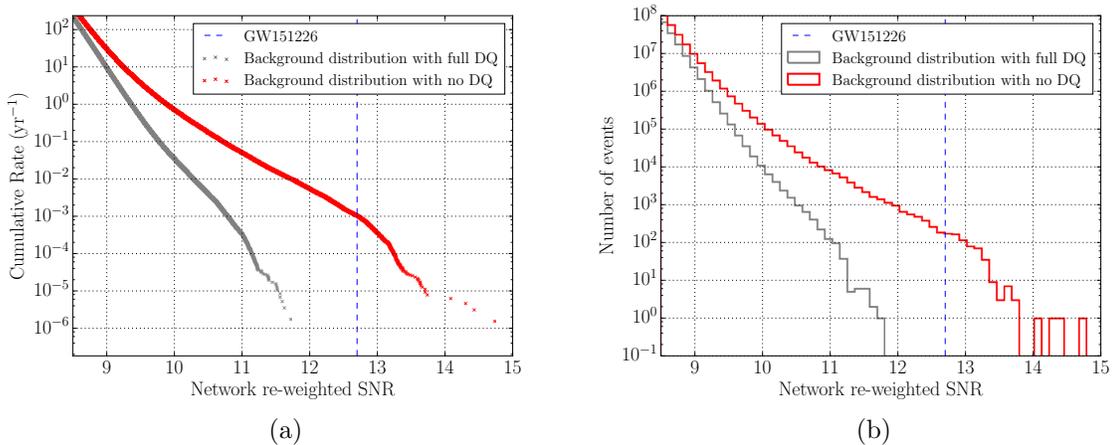


Figure 43: The background distribution in the bulk bin before and after applying data quality (DQ) vetoes. (43a) The cumulative rate of background triggers in the bulk bin as a function of re-weighted SNR. (43b) A histogram of background triggers in the bulk bin. The red traces indicate the distribution of background triggers without noisy data removed, the gray traces indicate the distribution of background triggers with all data quality vetoes applied. The dotted line indicates GW151226, which was recovered with $\hat{\rho}_c = 12.7$. If no data quality vetoes are applied, the tail of loud background triggers extends to $\hat{\rho}_c = 14.8$ instead of $\hat{\rho}_c = 11.8$. The impact of this change is apparent when considering GW151226, which is no longer the loudest event in this bin (see Section 9.2.1).

9.2.1 GW151226

The binary black hole system GW151226 was recovered by the PyCBC pipeline in the bulk bin with $\hat{\rho}_c = 12.7$ [47]. The significance of GW151226 was improved by the application of data quality vetoes. These changes in significance are quantified in Table 13. When all data quality vetoes were applied to the analysis, GW151226 was the loudest event in the bulk bin and has a false alarm rate of 1 per 183000 years. If noisy data are not removed from the analysis, GW151226 is no longer the loudest event in the bulk bin and its false alarm rate increases by a factor of 567 to 1 in 320 years. This increase takes a clear detection ($>5\sigma$) and reduces its significance to that of a more marginal detection candidate (3.9σ).

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	5.46×10^{-6}	2.28×10^{-7}
No CAT2 applied	1.06×10^{-5}	4.49×10^{-7}
No CAT1 or CAT2 applied	3.1×10^{-3}	1.38×10^{-4}

Table 13: Table of bulk bin false alarm rates and probabilities of GW151226 for several analysis configurations. The false alarm rate of GW151226 increases by a factor of 567, from 1 in 183000 years to 1 in 320 years, if data with excess noise is not removed from the analysis.

9.3 Edge bin

The background distribution in the edge bin looks dramatically different if data quality vetoes are not applied. This is not surprising, given that this tuning of the edge bin is restricted to contain the shortest waveforms with $f_{peak} < 100$ Hz, which will have a very short template duration and be susceptible to instrumental transients. Figure 44 shows the background distribution in the edge bin before and after data quality vetoes have been applied. The loudest event in this bin with all vetoes applied was at $\hat{\rho}_c = 15$, which was already inconveniently loud for a search hoping to recover a signal in this bin. When noisy data are not removed from the analysis, the loudest background event is at $\hat{\rho}_c = 18.3$, which further restricts the region where a confident detection could be made.

Further, there is a notable separation between the two background distributions at all values of $\hat{\rho}_c$. This separation can once again be quantified using the hypothetical detection candidate at $\hat{\rho}_c = 11.3$. Table 14 shows the false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for various analysis configurations. When data with excess noise is removed from the analysis, the false alarm rate at $\hat{\rho}_c = 11.3$ is reduced by a factor of 64. The application of CAT2 vetoes has an impact in this bin, providing a factor of 12 reduction in false alarm rate compared to the analysis with only CAT1 vetoes applied.

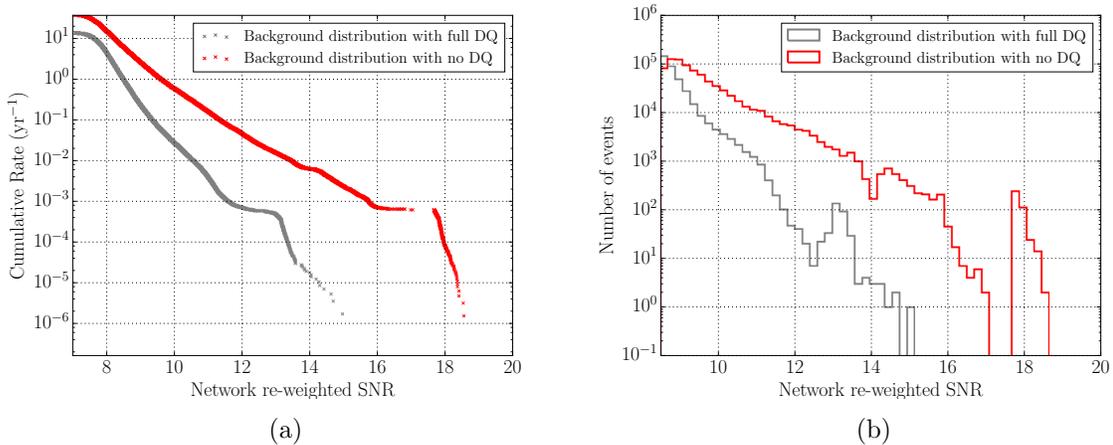


Figure 44: The background distribution in the edge bin before and after applying data quality (DQ) vetoes. (44a) The cumulative rate of background triggers in the edge bin as a function of re-weighted SNR. (44b) A histogram of background triggers in the bulk bin. The red traces indicate the distribution of background triggers without removing noisy data and the gray traces indicate the distribution of background triggers with all data quality vetoes applied. If noisy data are not removed from the analysis, the tail of loud background extends to $\hat{\rho}_c = 18.3$.

Analysis configuration	False alarm rate (yr ⁻¹)	False alarm probability
All vetoes applied	1.70×10^{-3}	7.08×10^{-5}
No CAT2 applied	2.16×10^{-2}	9.16×10^{-4}
No CAT1 or CAT2 applied	0.108	4.77×10^{-3}

Table 14: Table of edge bin false alarm rates and probabilities at $\hat{\rho}_c = 11.3$ for several analysis configurations. The application of CAT2 vetoes has an effect in this bin, providing a factor of 12 reduction in false alarm rate. Applying all data quality vetoes provides a factor of 64 reduction in false alarm rate.

Chapter 10

Limiting Noise Sources in the PyCBC Search

After applying data quality vetoes, there are still noticeable tails in the bulk and edge bin background distributions that limit the sensitivity of the search. This section aims to identify the types of instrumental features that are causing triggers with a high re-weighted SNR and acting as limiting noise sources. This section studies the analysis containing GW150914, which was detailed in Section 8.

10.1 Loud transients

A reasonable hypothesis is that the search is limited by loud transients with an SNR below the gating threshold. There are certainly a number of loud transients in the data that cause triggers with very high values of SNR, as seen in Figure 7a. It is sensible to check if the χ^2 test down-weights some of these high SNR triggers into the tail of the background distribution.

To test this, a cut was applied to the CBC triggers to exclude all triggers with an SNR > 20 . The histograms of Livingston re-weighted SNR triggers in Figure 45 show the results of this test. The green histogram in the foreground of the plot has had all single detector triggers with an SNR > 20 removed. The yellow histogram plotted in the background contains all single detector triggers from the analysis. Any points where the yellow histogram is visible indicate that triggers have been removed by the SNR cut. The cut does remove a small number of triggers with $\hat{\rho} > 8$, but the

overall structure of the tail is not significantly affected. None of the triggers with $\hat{\rho} > 10$ are removed by this cut. Most of the high SNR triggers are down-weighted below $\hat{\rho} = 6$ and are not visible on this histogram. A powerful veto that eliminates all CBC triggers with an SNR > 20 does not significantly improve the tail of the re-weighted SNR distribution.

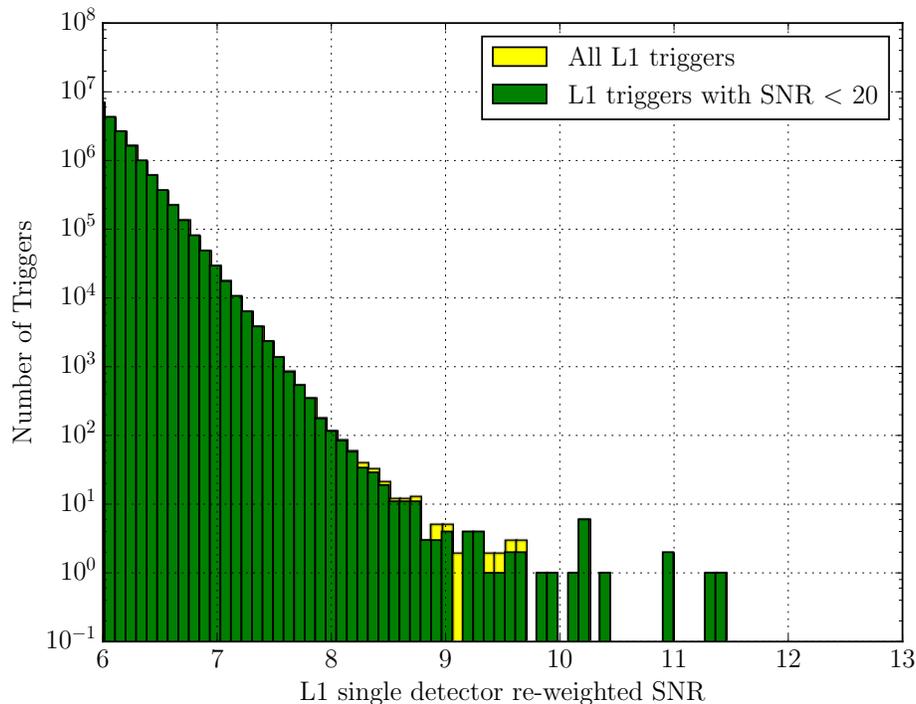


Figure 45: A histogram of single detector re-weighted SNR triggers for the Livingston (L1) detector. The green bins indicate triggers with an SNR < 20 . The yellow bins indicate all triggers in the data set. The triggers removed by the SNR cut do not significantly impact the loudest events which form a tail in the re-weighted SNR distribution. A small number of triggers at $\hat{\rho} > 9$ are removed by the SNR cut, but the population is not fully removed. The majority of the distribution is unchanged.

10.2 Blip transients

The transients that are able to pass the χ^2 test and populate the tail in the re-weighted SNR distribution are in fact those with a specific morphology which resembles that

of certain CBC waveforms. The most common and problematic source of transient noise that causes high re-weighted SNR triggers are called “blip transients.” These transients are often the source of the highest re-weighted SNR triggers at both the Livingston and Hanford interferometers. Although blip transients are seen in both interferometers, they are not found as coincident triggers and do not represent gravitational wave signals.

Blip transients show up as short duration, band-limited impulses that have power in the ~ 30 -300 Hz frequency range (see Figure 46). They don’t couple into any auxiliary channels that are used to monitor interferometer performance. Blip transients aren’t particularly loud, often recovered by Omicron with an SNR of 10-100, well below the gating threshold applied in the PyCBC search.

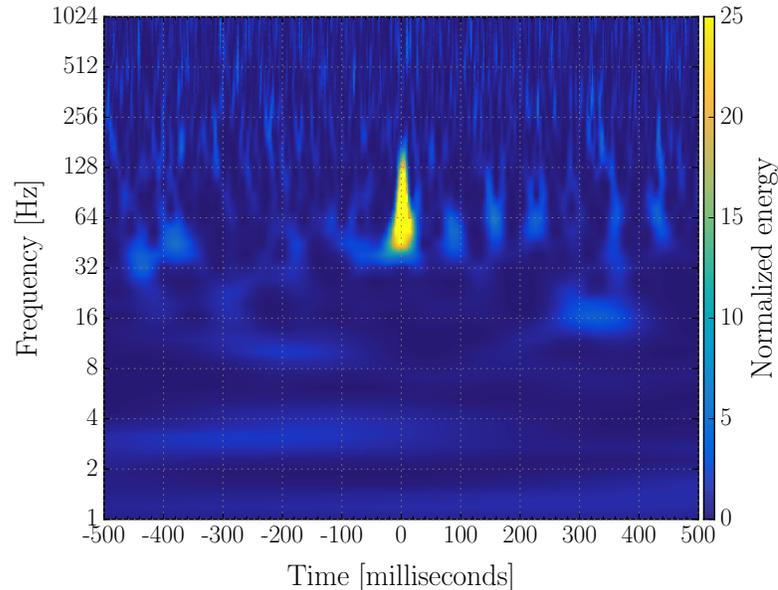


Figure 46: A time-frequency representation [69] of the Livingston strain channel at the time of a blip transient. This visualization of a blip transient demonstrates their typical features: band-limited, short duration, very little visible frequency structure.

A time-domain analysis reveals why these are so damaging to the CBC searches. Figure 47 shows a filtered time-domain representation of a blip transient in the Livingston strain channel. The data have been filtered with a bandpass filter with notch

filters to attenuate strong lines in the strain spectrum, double-passed to be zero-phase. Overlaid on top of the strain data is a CBC waveform that reported a high re-weighted SNR value at the time of the blip transient under study. The two curves show significant overlap in the few cycles where the template has appreciable amplitude.

The CBC template that reported a high re-weighted SNR when filtered against the blip transient in Figure 47 represents a neutron star-black hole binary system with a total mass (M_{total}) of $98.34M_{\odot}$ and a highly anti-aligned effective spin of -0.97 , resulting in a very short template duration. This system will coalesce very quickly and at a relatively low frequency compared to a lower mass binary system. The waveform spends less than 0.1 seconds at the frequencies that aLIGO is sensitive to, which, as shown in Figure 47, is the approximate time scale of some instrumental transients. This time scale is in stark contrast to that of a binary neutron star waveform, which can have a duration on the order of 1 minute and contain ample signal for use in the χ^2 test.

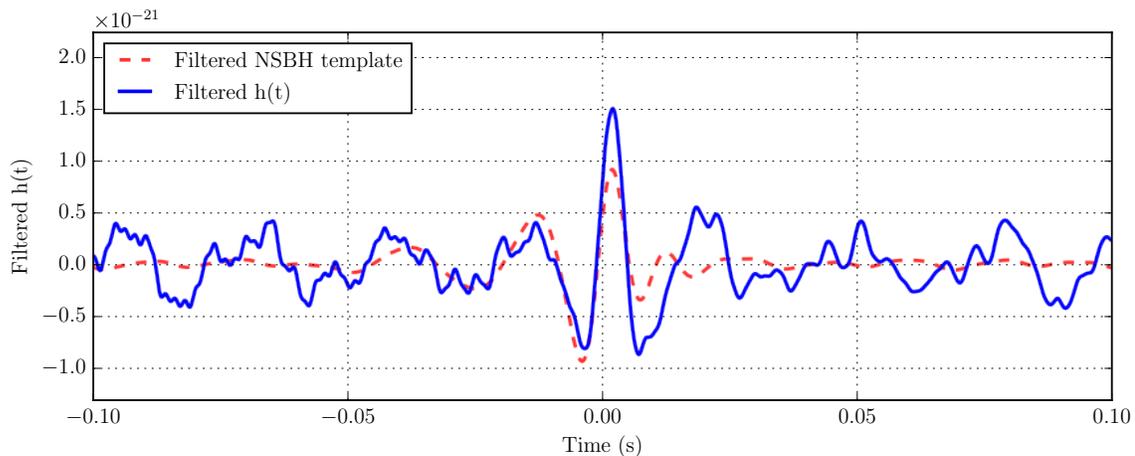


Figure 47: A filtered time-domain representation of the Livingston strain channel, $h(t)$, at the time of a blip transient. Overlaid on the strain plot is a filtered CBC waveform that reported a high re-weighted SNR value at the time of the blip transient. Both sets of data have been zero-phase bandpass filtered to isolate the frequency range that aLIGO is sensitive to. The two curves show significant overlap in the few cycles where the template has appreciable amplitude. The similarity between these two curves causes the χ^2 test to be ineffective at down-weighting these transients.

Although blip transients are capable of creating high re-weighted SNR triggers, their effects are constrained to a fairly small region of the CBC parameter space. Figure 48 shows single interferometer triggers from Livingston binned by total mass and effective spin. The bottom right corner of the plot, bounded by $M_{total} > 80M_{\odot}$ and $\chi_{eff} < -0.5$, contains all of the shortest duration templates and the highest re-weighted SNR triggers. This represents a small fraction of the CBC parameter space, containing only 65 waveform templates out of 249077 total. The very loudest triggers in the plot are even further constrained, corresponding to waveform parameters similar to those in Figure 47.

A further investigation reinforces the notion that the loudest triggers correspond to the templates with the shortest duration. Figure 49 shows single interferometer triggers from Livingston as a function of template duration and peak frequency of the CBC template. There is a systematic clustering of loud triggers below a template duration of 0.1 seconds, which is the timescale of typical instrumental transients. Constraining the loudest triggers using the peak frequency of the waveform template is not as successful. While the region corresponding to $f_{peak} < 100$ Hz does include the templates that are most susceptible to instrumental transients, it also includes numerous templates with a duration between 0.1 - 1.0 seconds that do not report any triggers with a high re-weighted SNR.

10.3 60-200 Hz noise

Another limiting noise source for the CBC search is present only at Livingston and has commonly been referred to as the “60-200 Hz” noise. This noise occurs in storms that can last multiple minutes and are typically comprised of a series of individual flares of noise that seem to last about 10-100 seconds each. These storms of noise correlate visibly with dips in the inspiral range, a figure of merit for the CBC searches that estimates the effective range at which detection of a binary neutron star inspiral is possible based on the shape of the noise curve. This noise contributes to the tail of loudest background triggers in the PyCBC search, including the cluster of loud triggers with a template duration of 4.4 s in Figure 49. Figure 50 shows the time-frequency representation of this noise on a 20 minute timescale.

A more focused look at these noisy periods reveals a structure that is reminiscent

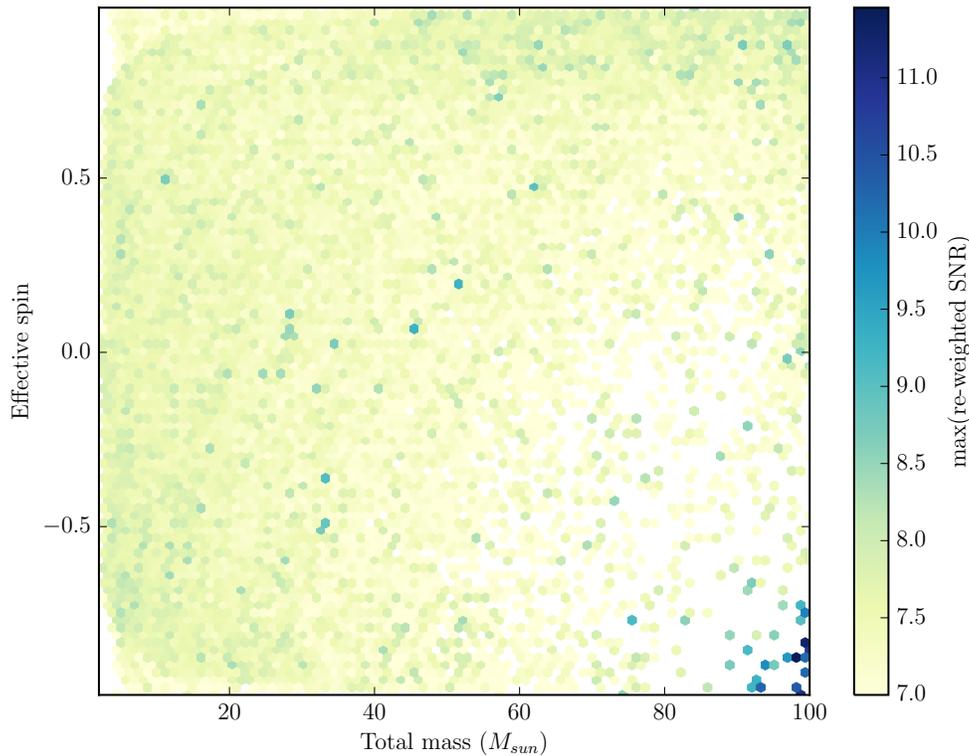


Figure 48: A plot of single interferometer triggers from the Livingston detector binned by total mass and effective spin. The color of each bin indicates the highest re-weighted SNR trigger found in that bin. The highest re-weighted SNR triggers are constrained to the bottom corner of the plot, bounded by $M_{total} > 80$ and $\chi_{eff} < -0.5$. This corner contains the shortest duration templates and is susceptible to instrumental transients such as blip transients.

of scattered light, appearing as arc-like traces in the time-frequency plane as seen in Figure 51. However, the frequency of this noise is higher than is typically expected from scattered light and investigations have not been able to find an associated source of scattered light during these noisy periods. This noise was a common source of high re-weighted SNR triggers in the Livingston data throughout the first observing run, second only to blip transients.

The two transient noise sources discussed in Sections 10.2 and 10.3 populate the tail of loudest events in PyCBC and are the current limiting noise sources. Although

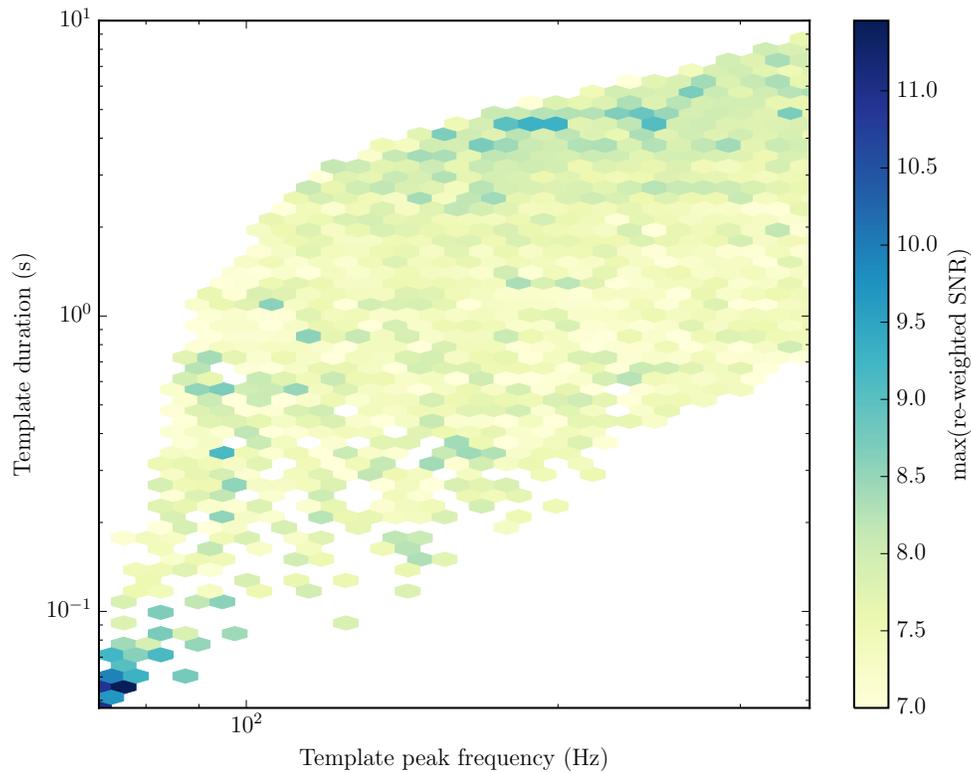


Figure 49: A plot of single interferometer triggers from the Livingston detector binned by template duration and waveform template peak frequency. The loudest triggers in re-weighted SNR are constrained to the area of the parameter space with template durations < 0.1 s, which is the timescale of typical instrumental transients, most notably blip transients. The small cluster of loud triggers with a template duration of roughly 4.4 s corresponds to the 60-200 Hz noise discussed in Section 10.3.

early investigations have not found causes, these noise sources will be the focus of the data quality group heading into upcoming observing runs.

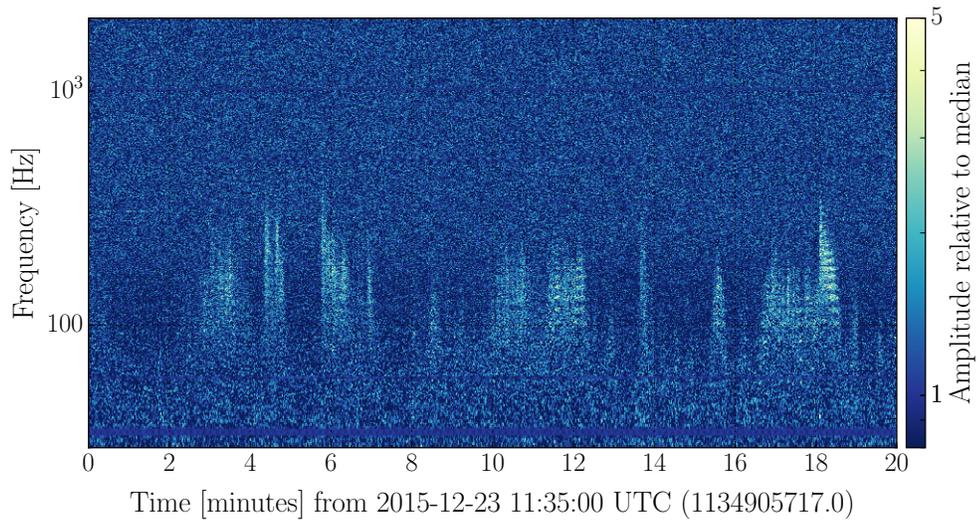


Figure 50: A time-frequency spectrogram of the 60-200 Hz noise. This noise appears in storms that often last for many minutes. This time scale and frequency range is damaging to CBC searches and has often been found responsible for loud background events.

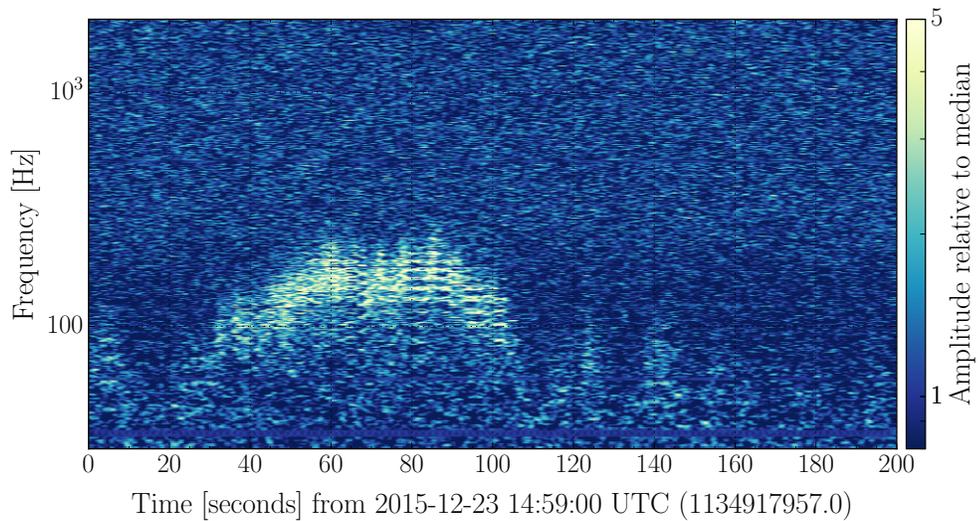


Figure 51: A zoomed in time-frequency spectrogram of the 60-200 Hz noise. This period of noise caused a loud trigger in the PyCBC background. The arc-like shape of the noise is reminiscent of noise due to scattered light, but the frequency of the noise is higher than expected.

Chapter 11

Conclusion

Advanced LIGO's first observing run was highly successful, resulting in the first direct detection of gravitational waves and further tests of Einstein's General Theory of Relativity. Gravitational waves from two binary black hole mergers, GW150914 and GW151226, were measured at the LIGO interferometers and recovered from the data using matched filter search algorithms.

Searching for gravitational waves requires an understanding of instrumental features and artifacts that can adversely affect the output of a gravitational wave search pipeline. Throughout the observing run, data quality vetoes were produced to ensure that the analysis pipelines analyzed clean data [44].

Data quality vetoes improved the sensitivity of the PyCBC search in Advanced LIGO's first observing run. Although the BNS bins were not dramatically affected, the distribution of background events was notably improved in the bulk and edge bins. In both bins, a significant tail of loud background triggers appeared if noisy data were not removed from the search.

In all 3 bins, it is evident that CAT1 vetoes had a more significant impact on false alarm rates than CAT2 vetoes, often providing 2-3 orders of magnitude of improvement in false alarm rate in the bulk and edge bins. This is expected, given that CAT1 vetoes are used to remove the most egregious data from the analysis. CAT2 vetoes had the greatest impact in the bulk and edge bins from the analysis containing GW151226, providing at least one order of magnitude reduction in false alarm rate compared to analyses using CAT1 vetoes only. This is due to a particularly effective CAT2 flag that was implemented in the analysis containing GW151226, but was not

relevant during the analysis containing GW150914.

The black hole binary system GW150914 was a strong enough signal that it was louder than all background events regardless of what data were removed from the search. As such, data quality vetoes did not improve its significance. The significance of LVT151012 was improved when data with excess noise were removed. Its false alarm rate was improved from 3.09 yr^{-1} to 0.44 yr^{-1} when data quality vetoes were applied to the PyCBC search.

The significance of the second binary black hole system discovered in O1, GW151226, was significantly increased by the application of data quality vetoes. The false alarm rate of GW151226 decreases by a factor of 567 when data quality vetoes are applied, which results in a clear detection ($>5 \sigma$) from a marginal detection candidate (3.9σ).

Appendix A

Data Quality Vetoes in O1

This document describes all the data quality (DQ) vetoes which were applied to the analysis of GW150914. For each DQ flag the definition of the veto is given, the interferometer this veto is applicable to, the category the veto was applied to the Burst and Compact Binary Coalescence (CBC) searches and the total amount of downtime associated to each DQ veto. This document has been created as a supplement to LIGO-P1500238.

A.1 Data Quality Vetoes

Missing Data Veto

Purpose: This veto captures any data dropouts at either interferometer.

Definition: Customized software indicate when the recalibrated data frames were unable to be produced either due to missing raw interferometer data or data in the raw data frames that are marked as invalid.

Veto Category: Burst - 1, CBC - 1

Deadtime: LIGO-Hanford - 0%, LIGO-Livingston - 0%

Burst Hardware Injection Veto

Purpose: This veto indicates whenever a burst hardware injection has been performed.

Defintion: The times of transient hardware injections labelled as burst type are recorded by the online detector characterization (ODC) system by monitoring the

state of the calibration injection model. Deadtime quoted includes the padding used in the analyses (± 4 seconds).

Veto Category: Burst - 4¹, CBC - 2

Deadtime: LIGO-Hanford - 0.003%, LIGO-Livingston - 0%

CBC Hardware Injection Veto

Purpose: This veto indicates whenever a CBC hardware injection has been performed.

Defintion: The times of transient hardware injections labelled as CBC type are recorded by the ODC system by monitoring the state of the calibration injection model. Deadtime quoted includes the padding used in the analyses (± 8 seconds).

Veto Category: Burst - 4, CBC - 3²

Deadtime: LIGO-Hanford - 0.052%, LIGO-Livingston - 0.072%

DetChar Hardware Injection Veto

Purpose: This veto indicates whenever a DetChar hardware injection has been performed.

Defintion: The times of transient hardware injections labelled as DetChar type are recorded by the ODC system by monitoring the state of the calibration injection model. Deadtime quoted includes the padding used in the analyses (± 16 seconds).

Veto Category: Burst - 4, CBC - 2

Deadtime: LIGO-Hanford - 0%, LIGO-Livingston - 0%

Stochastic Hardware Injection Veto

Purpose: This veto indicates whenever a stochastic hardware injection has been performed.

Defintion: The times of hardware injections labelled as stochastic type are recorded by the ODC system by monitoring the state of the calibration injection model.

Veto Category: Burst - 1, CBC - 1

Deadtime: LIGO-Hanford - 0%, LIGO-Livingston - 0%

¹Burst veto category 4 is reserved for transient hardware injections only.

²CBC veto category 3 is reserved for CBC hardware injections only.

Beckhoff Hardware Problems

Purpose: To capture times when the Beckhoff system (a slow control system which is used to control a subset of hardware in the interferometer) suffered a hardware failure at the LIGO-Hanford Y-end.

Defintion: The veto was created by hand, where the start time was recorded as 4 seconds before excess non-stationary data started due to the hardware failure and finished 3 seconds after the interferometer dropped out of observing mode.

Veto Category: Burst - 1, CBC - 1

Deadtime: LIGO-Hanford - 1.50%

45 MHz Sideband Fluctuations

Purpose: This veto identifies times when the amplitude of the 45 MHz optical sideband, which is used to generate error signals for optical cavities, has excess noise. If the amplitude of the 45 MHz optical sideband fluctuates, excess noise will be injected in to the associated optical cavities which has been seen to couple to the gravitational wave channel.

Defintion: An auxiliary channel which monitors amplitude fluctuations in the signal used to generate the 45 MHz optical sideband was found to be the optimum witness of non-stationary behaviour seen in the gravitational wave channel data. This veto was designed to capture long duration (on the order of one minute) non-stationary behaviour. Various thresholds on the band limited root-mean-square of this witness channel were investigated to see which threshold proved most effective (in terms of efficiency and deadtime) at removing non-stationary data. Custom software was implemented to automatically capture this behaviour over the analysis period (and throughout the first observing run).

Veto Category: Burst - 1, CBC - 1

Deadtime: LIGO-Hanford - 2.95%

Less severe 45 MHz Sideband Fluctuations

Purpose: See above veto, 45 MHz sideband fluctuations, for description.

Defintion: This veto was designed to capture less severe, short time scale (on the order of 1 second), non-stationary data. This veto was created in a similar manner as the previous veto - a study of different thresholds on the band limited root-mean-square

of the witness channel were investigated to give the optimal efficiency and deadtime. Custom software was implemented to automatically capture this behaviour over the analysis period (and throughout the first observing run).

Veto Category: Burst - not applied, CBC - 2

Deadtime: LIGO-Hanford - 0.014%

Saturations in the SUSETMY model channels

Purpose: This veto captures time when the Y-end test mass actuator saturates. This is due to a relatively fast transient that is on the main carrier beam, and therefore directly on/at the readout, which gets amplified by the differential-arm digital filters sufficiently to cross the digital-to-analog converter limits.

Defintion: This veto was created automatically by monitoring the interface between the computers and the analog electronics that they control on the Y-end test mass.

Veto Category: Burst - 2, CBC - 2

Deadtime: LIGO-Hanford - 0.067%, LIGO-Livingston - 0.021%

Saturations in the SUSETMY model channels with an SNR > 200

Purpose: See veto above - saturations in the SUSETMY model channels. This veto however is aimed specifically to identify very loud saturations.

Defintion: This veto was created automatically by monitoring the interface between the computers and the analog electronics that they control on the Y-end test mass. A subset of these saturations is kept based on their severity as determined by an algorithm that is designed to witness transient power in a given signal. This veto is specific to the Burst search where ± 3 seconds of padding is applied.

Veto Category: Burst - 2, CBC - not applied

Deadtime: LIGO-Hanford - 0.146%, LIGO-Livingston - 0.047%

Output Mode Cleaner (OMC) Photodiodes Analog to Digital Overflows

Purpose: This veto captures times when the signal on the OMC photodiodes exceeds the limit of the analog-to-digital converter at the interface to the computers that control the instrument.

Defintion: This veto was created automatically by monitoring the interface between the OMC photodiodes analog signal and the computers.

Veto Category: Burst - 2, CBC - 2

Deadtime: LIGO-Hanford - 0.002%, LIGO-Livingston - 0.003%

Non-Stationary Data prior to Loss of Resonant Power in the Optical Cavities

Purpose: To veto times when the data became non-stationary before the state of the interferometer reported the end of an observation segment.

Defintion: These times were found by hand by monitoring an algorithm, run over the gravitational wave channel, that is designed to witness transient power.

Veto Category: Burst - 1, CBC - 1

Deadtime: LIGO-Hanford - 0.0004%, LIGO-Livingston 0.001%

Glitches due to DC Power Fluctuations of the Photon Calibrator Laser

Purpose: This veto captures times when the photon calibrator has power fluctuations which exceed 20% of the nominal level.

Defintion: A threshold placed on a witness channel which monitors the power levels of the photon calibrator laser was used to flag times when the power fluctuated beyond the 20% level. These times were then padded by -10 and +20 seconds to capture the full behaviour.

Veto Category: Burst - 1, CBC -1

Deadtime: LIGO-Livingston - 0.058%

Seismic Glitches

Purpose: This veto was created to identify times of strong excess seismic noise that coupled in to the output of the interferometer.

Defintion: The 10-30 Hz band limited root-mean-square of the ground seismometer, located at the input test mass on the Y-arm, in the vertical degree of freedom was found to correlate with excess noise in the output of the interferometer. Different thresholds on this witness channel were tested to find the optimal efficiency and deadtime that captured these effects.

Veto Category: Burst - not applied, CBC - 2

Deadtime: LIGO-Hanford - 0.431%

Bibliography

- [1] A. Einstein. Approximative Integration of the Field Equations of Gravitation. *Preuss. Akad. Weiss. Berlin*, page 688, 1916.
- [2] A. Einstein. Über Gravitationswellen. *Sitzungsber. K. Preuss. Akad. Wiss.*, 1:154–167, 1918.
- [3] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, Singapore, 1994.
- [4] Sean Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Addison Wesley, 2004.
- [5] <http://www.itp.uzh.ch/~chuwyler/gravwaves/polarisations2.png>.
- [6] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.
- [7] T. Fricke et al. DC readout experiment in Enhanced LIGO. *Class. Quant. Grav.*, 29, 2012.
- [8] B. P. Abbott et al. GW150914: The Advanced LIGO Detectors in the Era of First Discoveries. 2016. <https://dcc.ligo.org/LIGO-P1500237/public/main>.
- [9] F. Matichard et al. Seismic isolation of Advanced LIGO: review of strategy, instrumentation and performance. *Class. Quant. Grav.*, 32, 2015.
- [10] S. Wen et al. Hydraulic external pre-isolator system for LIGO. *Class. Quant. Grav.*, 31, 2014.
- [11] L. Carbone et al. Sensors and Actuators for the Advanced LIGO Mirror Suspensions. *Class. Quant. Grav.*, 29, 2012.

- [12] B. P. Abbott et al. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. 2016. <https://dcc.ligo.org/LIGO-P1500269/public/main>.
- [13] J. H. Taylor and J. M. Weisberg. A new test of general relativity : gravitational radiation and the binary pulsar psr 1913 +16. 253:908, 1982.
- [14] R. A. Hulse. Nobel lecture: The discovery of the binary pulsar. 66:699, 1994.
- [15] J. H. Taylor and J. M. Weisberg. Further experimental tests of relativistic gravity using the binary pulsar 1913+16. 345:434, 1989.
- [16] B. P. Abbott et al. Observation of Black Hole Binary Mergers in Advanced LIGO's First Run. 2016. <https://dcc.ligo.org/LIGO-P1600088/public/main>.
- [17] Thibault Damour and Alexander Vilenkin. Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows. *Phys. Rev. D*, 71:063510, 2005.
- [18] B. P. Abbott et al. Observing gravitational-wave transient GW150914 with minimal assumptions. 2016. <https://dcc.ligo.org/LIGO-P1500229/public/main>.
- [19] Chris L. Fryer and Kimberly C.B. New. Gravitational waves from gravitational collapse. *Living Reviews in Relativity*, 14(1), 2011.
- [20] R. Essick E. Katsavounidis R. Lynch, S. Vitale and F. Robinet. An information-theoretic approach to the gravitational-wave burst detection problem. 2015. arXiv:1511.05955.
- [21] N. J. Cornish and T. B. Littenberg. Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches. *Class. Quant. Grav.*, (32(13):135012), 2015.
- [22] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher. A coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 25(11):114029, 2008.

- [23] J. Aasi et al. First all-sky search for continuous gravitational waves from unknown sources in binary systems. *Phys. Rev. D*, 90(062010), 2014.
- [24] J. Aasi et al. Einstein@home all-sky search for periodic gravitational waves in ligo s5 data. *Phys. Rev. D*, 87(042001), 2013.
- [25] B. Abbott et al. Searching for a Stochastic Background of Gravitational Waves with LIGO. *Astrophysical J.*, 659:918–930, 2007.
- [26] B. P. Abbott et al. GW150914: Implications for the stochastic gravitational wave background from binary black holes. 2016. <https://dcc.ligo.org/LIGO-P1500222/public/main>.
- [27] B. Allen and J. Romano. Detecting a stochastic background of gravitational radiation: signal processing strategies and sensitivities. *Phys. Rev. D*, 59(102001), 1999.
- [28] F. Acernese et al. Advanced Virgo: a 2nd generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32, 2015.
- [29] K. L. Dooley et al. GEO 600 and the GEO-HF upgrade program: successes and challenges. 2016. arXiv:1510.00317.
- [30] H. Vahlbruch et al. The GEO600 squeezed light source. *Class. Quant. Grav.*, 27, 2010.
- [31] Y. Aso et al. Interferometer design of the KAGRA gravitational wave detector. *Phys. Rev. D*, 88, 2013.
- [32] C. S. Unnikrishnan. IndIGO and LIGO-India: Scope and Plans for Gravitational Wave Research and Precision Metrology in India. *International Journal of Modern Physics D*, 22, 2013.
- [33] S. Usman et al. An improved pipeline to search for gravitational waves from compact binary coalescence. 2015. arXiv:1508.02357.
- [34] Alexander H. Nitz, Ian W. Harry, Joshua L. Willis, Christopher M. Biwer, Duncan A. Brown, Larne P. Pekowsky, T. Dal Canton, Andrew R. Williamson,

- Thomas Dent, Collin D. Capano, Thomas T. Massinger, Amber K. Lenon, Alex Nielsen, and Miriam Cabero. PyCBC Software. <https://github.com/ligo-cbc/pycbc>, 2016.
- [35] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys.Rev.*, D85:122006, 2012.
- [36] Yi Pan, Alessandra Buonanno, Luisa T. Buchman, Tony Chu, Lawrence E. Kidder, et al. Effective-one-body waveforms calibrated to numerical relativity simulations: coalescence of non-precessing, spinning, equal-mass black holes. *Phys. Rev. D*, 81:084041, 2010.
- [37] Michael Pürrer. Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins. *Phys. Rev. D*, 93:064041, Mar 2016.
- [38] Andrea Merloni. Cosmological evolution of supermassive black holes and AGN: a synthesis model for accretion and feedback. *Mem. Soc. Astron. Italiana*, 79:1310, 2008.
- [39] K. A. Postnov and L. R. Yungelson. The Evolution of Compact Binary Star Systems. *Living Reviews in Relativity*, 9:6–+, December 2006.
- [40] P. C. Peters and J. Mathews. Gravitational radiation from point masses in a Keplerian orbit. *Phys. Rev.*, 131:435–439, 1963.
- [41] S. Privitera et al. Improving the sensitivity of a search for coalescing binary black holes with nonprecessing spins in gravitational wave data. *Phys. Rev.*, D89(2):024003, 2014.
- [42] Lawrence E. Kidder. Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects. *Phys. Rev.*, D52:821–847, 1995.
- [43] L. K. Nuttall et al. Improving the Data Quality of Advanced LIGO Based on Early Engineering Run Results. *Class. Quant. Grav.*, 32(24):245005, 2015.

- [44] B. P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500238/public/main>.
- [45] B. Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Phys.Rev.*, D71:062001, 2005.
- [46] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, 2016. <https://dcc.ligo.org/LIGO-P150914/public/main>.
- [47] The LIGO Scientific Collaboration and the Virgo Collaboration. GW151226: Observation of a Binary Black Hole Merger. *Phys. Rev. Lett. (in LIGO review)*.
- [48] B. P. Abbott et al. Properties of the binary black hole merger GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500218/public/main>.
- [49] B. P. Abbott et al. The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500217/public/main>.
- [50] B. P. Abbott et al. Astrophysical Implications of the Binary Black-Hole Merger GW150914. *Astrophys. J.*, 818(2):L22, 2016. <https://dcc.ligo.org/LIGO-P1500262/public/main>.
- [51] B. P. Abbott et al. Tests of general relativity with GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500213/public/main>.
- [52] B. P. Abbott et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500248/public/main>.
- [53] B. Abbott et al. Localization and broadband follow-up of the gravitational-wave transient GW150914. 2016. <https://dcc.ligo.org/LIGO-P1500227/public/main>.
- [54] S. Adrian-Martinez et al. High-energy Neutrino follow-up search of Gravitational Wave Event GW150914 with ANTARES and IceCube. 2016. <https://dcc.ligo.org/LIGO-P1500271/public/main>.

- [55] B. Abbott et al. Wf accuracy. 2016. <https://dcc.ligo.org/P1500259/public/main>.
- [56] C. Messick et al. Low-latency analysis framework for the prompt discovery of compact binary mergers in gravitational wave data. 2016. <https://dcc.ligo.org/LIGO-P1600009>.
- [57] J. R. Smith et al. A Hierarchical Method for Vetoing Noise Transients in Gravitational-Wave Detectors. *Class. Quant. Grav.*, 28, 2011.
- [58] T. Isogai et al. Used Percentage Veto for LIGO and Virgo Binary Inspiral Searches. *Journal of Physics: Conference Series*, 243, 2010.
- [59] A. Effler, R. M. S. Schofield, V. V. Frolov, G. Gonzalez, K. Kawabe, J. R. Smith, J. Birch, and R. McCarthy. Environmental Influences on the LIGO Gravitational Wave Detectors during the 6th Science Run. *Class. Quant. Grav.*, 32(3):035017, 2015.
- [60] D. Martynov et al. The sensitivity of the advanced ligo detectors at the beginning of gravitational wave astronomy. Technical Report LIGO-P1500260, LIGO Project, 2016.
- [61] S6 Detector Characterization Paper; in preparation for publication.
- [62] J. McIver. Data Quality Studies of Enhanced Interferometric Gravitational Wave Detectors. *Class. Quantum. Grav.*, 29(124010), 2012.
- [63] F. Robinet. Omicron: An Algorithm to Detect and Characterize Transient Noise in Gravitational-Wave Detectors. <https://tds.ego-gw.it/ql/?c=10651>, 2015.
- [64] J. McIver. The impact of terrestrial noise on the detectability and reconstruction of gravitational wave signals from core-collapse supernovae. 2015. LIGO-P1500072.
- [65] E. D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *Am. J. Phys.*, 69:79–87, 2001.
- [66] C. Mueller. The Advanced LIGO Input Mode Cleaner, 2014. LIGO-G1400096.

- [67] K. Kokeyama. DRMI LSC tutorial for detchar working group, 2013. LIGO-G1301236.
- [68] N. Mavalvala. Alignment issues in laser interferometric gravitational-wave detectors, 1997. <https://dspace.mit.edu/bitstream/handle/1721.1/10769/36962416-MIT.pdf?sequence=2>.
- [69] S. Chatterji et al. *Class. Quantum Grav.*, 21:S1809, 2004.

Curriculum Vitae

Title of Dissertation

Detector Characterization for Advanced LIGO

Full Name

Thomas J. Massinger

Place and Date of Birth

Rochester, NY

July 22, 1989

Colleges and Universities,

Utica College

Syracuse University

Years attended and degrees

2007–2011, B.S.

2011–2016, Ph.D.

Major Department

Physics