

**THE IMPACT OF TERRESTRIAL NOISE ON THE
DETECTABILITY AND RECONSTRUCTION OF
GRAVITATIONAL WAVE SIGNALS FROM
CORE-COLLAPSE SUPERNOVAE**

A Dissertation Presented

by

JESSICA L MCIVER

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

DOCTOR OF PHILOSOPHY

September 1 2015

Physics

© Copyright by Jessica L McIver 2015

All Rights Reserved

ACKNOWLEDGEMENTS

Thanks to Laura Cadonati, who has supported me through many years of graduate school and been a strong role model I'm proud to aspire to. Thanks to my committee members, Alex Pope, Ben Brau, and David Kastor, who each offered a unique perspective on this dissertation that has helped shape it for the better. Special thanks to Laura Nuttall, Josh Smith, James Clark, Sarah Gossan, Ben Brau, and Laura Cadonati for extensive comments that have greatly improved this dissertation.

This work would not have been possible without:

Josh Smith, who has offered me invaluable advice and guidance for many years. While working on LIGO detector characterization as an undergraduate, Josh's enthusiasm and guidance inspired me to pursue a physics Ph.D., and he has been a steadfast mentor ever since. Thank you so much. Your students and colleagues are very lucky to have your guidance.

James Clark, who offered key ideas about the figure of merit used to evaluate waveform reconstruction performance and gave me insight on the inner mechanics of burst searches. Thank you for scores of useful conversations and always pointing me to exactly what I need in the giant libraries of LIGO code.

Duncan Macleod. Evidence of his coding talent can be found in most of the plots in Chapter 5. Thank you for your immeasurable support in navigating existing code, building new code, teaching me about proper pedantic python code formatting, and

making my visits to LLO possible. And most of all, for making doing detector characterization something to look forward to.

Sarah Gossan, whose crucial guidance on supernova waveforms and physical processes influenced much of the work in Chapter 4. Thanks for answering my many questions and sharing your expertise to motivate the supernova waveform reconstruction studies.

Jeff Kissel, who first introduced me to seismic isolation instrumentation. Thank you for investing your time and energy in teaching me and the rest of the detector characterization group. Thanks also for making a visit to LHO possible for me, and for helping to take the seismic isolation data used in Chapter 5.

Florent Robinet, who fielded several key requests to change Omicron's implementation. Thanks especially for adding a reported amplitude parameter for triggers and enabling Omicron to search below 1Hz; two non-trivial changes vital to the seismic transient study in Chapter 5. Thanks also for answering my many questions about Omicron's implementation and inner mechanics.

Claudia Lazzaro, who helped me set up the coherent wave burst tuning runs used and work out some pretty strange bugs. Thanks for your continual help and answering my many cWB questions.

Jonah Kanner and Tyson Littenberg, who were integral to getting me involved with BayesWave. Thanks for spending so much time teaching me about the algorithm, sharing your code, advising me about BayesWave parameters and tuning for supernova waveforms, and helping to debug when things didn't go as planned.

Chris Pankow, who offered insight on ETG performance and implementation over many useful conversations, and generated the Gaussian noise and injections used for

the ETG performance study in Chapter 3. Thank you for sharing your time, code, and ideas.

Thanks to the ETG developers who produced triggers and provided documentation for the ETG study in Chapter 3: John Zwezig (DMT-Omega), Chris Pankow (Excess Power), Joey Key (BayesWave) and Daniele Trifiro (PCAT).

And thanks to many others who have contributed to this work in ways big and small:

Brian Lantz, who helped the detector characterization team integrate with the seismic team and taught me some about filter inversion (I still have a lot to learn). Fabrice Matichard, who has been a valuable resource for doing offsite seismic characterization, and who helped me learn to communicate data analysis for the instrumentation-minded scientist. Rich Mittleman, who was the first person to introduce me to Bode plots, and who helped provide supplementary information for the study in Chapter 5. Hugh Radkins, who showed me how to reset a HEPI pump, and helped to take some of the data used in Chapter 5. Celine Ramet, Sebastian Biscans, Arnaud Pele, and the rest of the seismic isolation team, who have contributed to the design and implementation of the seismic ODC state vectors referenced for taking the data used in the seismic transient propagation study, as well as the seismic summary pages, referenced in Chapter 5. Matt Heintze, who generously showed me a BSC ISI and payload before it was installed, and took pictures used in the defense presentation of this dissertation. Janeen Romie, who made a visit to LLO possible for me, and made it a welcoming place.

Laura Nuttall, who has made working on detector characterization for the last six years a joy, and who has contributed to many fruitful joint studies that don't happen to be within the scope of this thesis. Andy Lundgren, who first interested me in the importance of evaluating the performance of ETGs, and who can be counted on to

think up an aptly amusing name for every glitch type. TJ Massinger, who taught the detector characterization group about cavity locking and how Advanced LIGO reaches full lock, and has also contributed to many key instrumental characterization investigations complementary to those included in this dissertation. Amber Stuver, who produced the injection set used for the core-collapse supernova waveform reconstruction studies. Ryan Fisher, who first developed the seismic ODC state vectors used for data collection. Marissa Walker, who produced the ROC curves used in Chapter 3, is conducting the suspensions transient study complementary to the study in Chapter 5, and ran cWB with me for the very first time. Thomas Abbott and Nutsinee Kibunchoo, who contributed to seismic and suspensions characterization supplementary to this work. Dan Hoak, who also conducted detector characterization studies complimentary to this work. Gaby Gonzalez and Joe Giaime, who have made suggestions that have improved the detector characterization studies in this dissertation.

Thanks also to the entire LIGO detector characterization, burst, supernovae, and seismic isolation and suspensions working groups for comments and suggestions on the included studies.

Special thanks to Peter Saulson, who first inspired me to get involved with physics research as an undergraduate and introduced me to LIGO. Thanks also for your continued mentorship.

To Sarah Zuraw and Kalina Nedkova, thanks for your hard work and dedication, it was a pleasure to work with you over the years. To Jane Knapp, thank you for your incredible support of me and all physics graduate students. It makes all the difference.

Thank you Alex Lombardi, Preema Pais, Szu-Chia Chen, Dana Algaier, and Brendan Gavin for forming Team Thesis and helping me power through writing.

Thanks to my family and friends for their support.

And lastly, thank you Tom. For all the things.

ABSTRACT

THE IMPACT OF TERRESTRIAL NOISE ON THE DETECTABILITY AND RECONSTRUCTION OF GRAVITATIONAL WAVE SIGNALS FROM CORE-COLLAPSE SUPERNOVAE

SEPTEMBER 1 2015

JESSICA L MCIVER

B.S. Physics, SYRACUSE UNIVERSITY

B.S. Magazine Journalism, SYRACUSE UNIVERSITY

M.S. Physics, UNIVERSITY OF MASSACHUSETTS AMHERST

Ph.D. Physics, UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Laura Cadonati

Gravitational waves, small perturbations in spacetime produced by accelerating massive bodies, were first predicted in 1915 by Einstein's Theory of General Relativity. Gravitational waves interact weakly with matter, making them very difficult to detect. However, this property also enables gravitational waves to propagate through

dense matter, such as the outer layers of a collapsing star, making them an ideal probe of astrophysical systems that are difficult to observe via electromagnetic radiation.

A new era of gravitational wave astronomy is approaching as the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) is about to complete its construction and begin its data acquisition in 2015; the centennial of Einstein's General Relativity. Terrestrial gravitational wave observatories are engineered to use laser interferometry to measure very small reverberations in spacetime as strain, the change in length over length of their two perpendicular arms. The Laser Interferometer Gravitational-wave Observatory (LIGO) consists of two of these instruments in the United States, each with two 4km arms. The Advanced LIGO instruments are expected to improve sensitivity over prior LIGO measurements by a factor of ten at 100-400 Hz, the most sensitive frequency region. This improvement will expand the volume of space that the LIGO detectors are able to observe by a factor of 1000.

Amongst the most exciting potential astrophysical sources for Advanced LIGO are galactic core-collapse supernovae (CCSN). A key open question for these events is to discover the mechanism that drives an explosion by re-energizing the matter shock wave, which is widely believed to stall shortly after bouncing off of the dense collapsed core. Detecting and reconstructing gravitational waves produced during these events would yield insight into their explosion mechanism as well as other related physics, such as the characteristics of the neutron star produced after core collapse. However, confident detection and accurate waveform reconstruction is made much more difficult by ambient terrestrial noise.

This work details an exploration of the ability of LIGO software to reconstruct waveforms produced by core-collapse supernovae. The study injects modeled supernovae signals into simulated Gaussian noise at the design sensitivity of Advanced LIGO and Advanced Virgo. The key metric to evaluate algorithm performance is the normalized

overlap between the injected waveforms and the reconstructed waveforms, which measures the degree to which they match. Results show that the analyzed core-collapse models are only recovered a short distance away; up to 1 kpc for the most realistic analyzed set of waveforms. Of the two burst signal recovery algorithms studied, neither produce accurate reconstructions of injected CCSN waveforms. Reconstructed waveforms generally have an overlap of well below 100% with the corresponding injected waveform at astrophysically realistic signal-to-noise ratios. Additionally, injecting the same waveforms in realistic non-Gaussian noise recolored to the same anticipated instrument sensitivity has a significant adverse impact on both efficiency and waveform reconstruction relative to performance in Gaussian noise.

The second portion of this dissertation focuses on characterizing the impact of noise in the Advanced LIGO interferometers on searches for gravitational wave bursts, including supernovae. Terrestrially induced transient noise from environmental sources such as excess ground motion had a detrimental effect on the performance of transient gravitational wave searches during prior LIGO observation runs. This work examines the effect of active seismic isolation in mitigating transient ground motion and projected the impact of seismic noise on the sensitivity of transient astrophysical gravitational wave searches during the first Advanced LIGO observing run. Results show that the upgraded seismic isolation instrumentation installed for Advanced LIGO is very effective in reducing transient noise amplitude in the searchable gravitational wave frequency band, but not in reducing the rate of transient noise during periods of elevated ground motion.

Key to the detector characterization effort is the identification of transient noise events that produce excess power in a given signal, generally called *glitches*. Event trigger generators are algorithms that identify transient events above some loudness threshold and characterize them with an associated time, frequency, and loudness, or

significance. Understanding the detection efficiency and accuracy with which these algorithms recover event parameters is critical not only for characterizing the instruments, but also to ensure that glitch correlation algorithms used to improve data quality perform well. On the whole, these algorithms maximize the potential for confident signal detections while carefully minimizing the likelihood of unintentionally removing a true gravitational wave signal from analysis. The results presented here show that event trigger generators developed for use during Advanced LIGO perform well, and will motivate the use of the ETG Omicron for measuring transient behavior in the active seismic isolation subsystem.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iii
ABSTRACT	viii
GLOSSARY	xiv
 CHAPTER	
INTRODUCTION	1
1. GRAVITATIONAL WAVES	6
1.1 General Relativity and gravitational waves	6
1.1.1 Interaction with matter	9
1.1.2 Gravitational wave emission	9
1.2 The gravitational wave sky	11
1.2.1 Core-collapse supernovae	13
2. THE LIGO DETECTORS	17
2.1 The LIGO detectors	17

2.2	Gravitational wave astronomy with LIGO	21
2.3	Advanced LIGO	31
2.3.1	The global interferometer network	41
2.3.2	Common sources of noise	47
3.	RECOVERY OF BURST TRANSIENT EVENTS	50
3.1	Event trigger generators	51
3.1.1	Parameterization technique and tiling	51
3.1.2	Implementation for Omicron	62
3.1.3	Estimation of trigger SNR and amplitude	76
3.2	Performance of Event Trigger Generators	81
3.2.1	Considered algorithms	82
3.2.2	Study technique	84
3.2.3	Results	89
3.2.4	Recommendations based on results	110
4.	DETECTING AND RECONSTRUCTING CORE-COLLAPSE SUPERNOVAE WITH ADVANCED LIGO	114
4.1	Models for the CCSN explosion mechanism	114
4.2	Waveform reconstruction performance study	116
4.2.1	Overview of reconstruction methods	117
4.2.1.1	cWB2G	117
4.2.1.2	BayesWave	119

4.2.2	Comparing cWB2G and BayesWave results	123
4.2.3	Included waveforms	124
4.2.4	Technique	134
4.3	Waveform reconstruction algorithm tuning	136
4.3.1	cWB2G configuration	136
4.3.2	BayesWave configuration	149
4.4	Results	150
4.4.1	Efficiency	150
4.4.2	Waveform reconstruction	153
4.4.2.1	Trends in waveform reconstruction performance by CCSN model	153
4.4.2.2	Comparison of cWB2G and BayesWave performance	159
4.4.3	Summary and future work	165
5.	CHARACTERIZING AND MITIGATING TERRESTRIAL NOISE SOURCES	166
5.1	Terrestrial noise: impact on detectability, parameter estimation, and waveform reconstruction	166
5.2	Combating transient noise	170
5.2.1	Data quality vetoes	170
5.3	The impact of seismic noise on prior GW transient searches	177
5.4	Advanced LIGO seismic isolation instrumentation	181
5.5	Advanced LIGO seismic noise transient propagation	193

5.5.1	Technique	194
5.5.2	Results	198
5.6	Summary and future work	204
6.	THE IMPACT OF TERRESTRIAL SEISMIC NOISE ON CCSN DETECTABILITY AND RECONSTRUCTION	207
6.1	Technique	207
6.2	CCSN detectability in recolored non-Gaussian data	208
6.3	CCSN reconstruction in recolored non-Gaussian data	210
6.4	Summary and future work	213
7.	CONCLUSIONS AND FUTURE PROSPECTS	217
 APPENDICES		
A.	FULL WAVEFORM RECONSTRUCTION RESULTS FOR CWB2G AND BAYESWAVE	221
B.	SUPPLEMENTARY ANALYSIS OF BURST PARAMETER RECOVERY	225
C.	MAXIMUM ENTROPY AND CCSN WAVEFORM RECONSTRUCTION	237
D.	OMICRON PARAMETERS IN DEPTH	256
	BIBLIOGRAPHY	263

GLOSSARY

A brief glossary of terms and acronyms for quick reference.

aLIGO - Advanced LIGO - two terrestrial gravitational wave interferometers in the United States, recently equipped with upgraded instrumentation.

ALS - Arm Length Stabilization - a system used to acquire and control cavity resonance in each of the two interferometer arms. This procedure is largely independent of status of the rest of the interferometer.

BayesWave - a follow-up Bayesian parameter estimation and signal classification algorithm. BayesWave estimates a posterior distribution for a recovered signal waveform inferred from the data and the interferometer network response. To fit an excess power event as closely as possible, BayesWave uses a reversible-jump Markov Chain Monte Carlo algorithm to explore the application of a distribution of wavelets and wavelet parameters to the signal fit.

Burst - A gravitational wave burst event, or a burst of signal power associated with a gravitational wave source that is localized in time.

BNS range - Binary Neutron Star (BNS) range - the distance from Earth at which the coalescence of binary neutron stars, both $1.4 M_{\odot}$, produces a signal-to-noise ratio (SNR) of 8 in a single detector using matching filtering, averaged over sky location and source orientation. Also referred to as *BNS inspiral range* or informally as just the *inspiral range*.

BSC - Basic Symmetric Chamber - the ultra-high vacuum chambers designed to house LIGO's core optics: the beam splitter, the inner and end test masses, the reaction suspension chains used for quiet actuation of the test masses (and thermal control of the input test masses), and the transmission monitor optics used for the Arm Length Stabilization system installed in the end-arm chambers.

CCSN - Core-collapse supernova are produced after the gravitational collapse of a star core that has expended its fuel.

cWB2G - Coherent wave-burst 2G (second generation) - the flagship burst search algorithm for LIGO and Virgo gravitational wave data. cWB2G is a coherent all-sky search that identifies candidate gravitational wave burst events by tiling the data in

time and frequency via a wavelet-basis projection and extracting significant events with a coherent likelihood statistic maximized over all potential sky locations. Also informally referred to as ‘cWB’.

ETG - Event Trigger Generator - also known as a single interferometer burst algorithm. ETGs identify excess power in the signal that is localized in time and frequency, and generally characterize these events with a characteristic time, frequency, and measure of loudness in the event signal-to-noise ratio (SNR) or amplitude.

HAM - Horizontal Access Module - the ultra-high vacuum chambers designed to house LIGO’s auxiliary optics: the optics forming the input mode cleaner (IMC), power recycling cavity (PRC), signal recycling cavity (SRC), output mode cleaner (OMC) as well as other smaller mirrors and components.

iLIGO - Initial LIGO, which collected data for 5 science runs (S1-S5) over the span of 2002 - 2007.

IMC - Input Mode Cleaner - a three mirror folded cavity used to improve beam mode quality by filtering out higher order spatial modes. The IMC also serves as a high-quality reference cavity for input light frequency stabilization.

ISI - Internal Seismic Isolation - a passively isolated and actively controlled stage or series of stages housed within the vacuum system that supports the optic suspensions.

HEPI - Hydraulic External Pre-Isolator - an actively isolated stage external to the vacuum system that supports the internal isolation (ISI) stages.

MICH - Michelson (interferometer degree of freedom) - defined by the difference in length between the beam splitter (BS) and each of the input test mass optics (ITMX and ITMY). Used to maintain a small asymmetry between these lengths, known as Schnupp asymmetry

OMC - Output Mode Cleaner - a four-optic bowtie cavity designed to strip out all higher order beam shape modes and any injected sidebands used for cavity control so that only good-quality carrier frequency light imprinted with the gravitational wave signal reaches the final readout photodiodes.

Omicron - the primary Event Trigger Generator, or single interferometer burst identification algorithm, intended for transient noise (glitch) detector characterization studies on Advanced LIGO and Advanced Virgo data. Omicron identifies and parameterizes bursts of excess power by projecting conditioned data onto a sine-Gaussian basis.

Overlap - Waveform overlap, or the inner product between two waveforms normalized to 1 for a perfect match. For this dissertation, overlap is defined as the time-optimized, noise-weighted, normalized overlap between an injected waveform and recovered waveform. Explicitly defined in Equation 4.9.

PRC - Power Recycling Cavity - an optic cavity between the input optics and the main interferometer that increases the power circulating in the interferometer arms by reflecting recycled light such that it is in phase with light coming from the input optics.

PSL - Pre-Stabilized Laser - a Nd:YAG laser capable of emitting up to 180W of power at 1064nm. The input power provided by the PSL is conditioned to be extremely stable in intensity and frequency.

Q - quality factor of an oscillating system defined as the ratio between the central frequency and the bandwidth of the oscillation.

Recoloring - a data manipulation technique useful for generating realistic noise for simulations. The data is first whitened to have a flat PSD, and then colored to have some specified PSD, usually one of the Advanced LIGO configurations shown in Figure 2.2. Noise re-coloring was commonly used in conjunction with simulated waveform injections to more accurately test how transient gravitational wave detection and parameter estimation algorithms will perform with realistic Advanced LIGO data.

SEI - Seismic isolation - a subsystem providing active mechanical isolation from local ground motion to the interferometer components in a series of stages.

SG - sine-Gaussian - a sinusoid made discrete in time with a Gaussian envelope. The shape of sine-Gaussian waveforms is characterized by a central frequency and quality factor Q . These waveforms are used as the basis of a family of event trigger generators, or single interferometer signal burst detection algorithms, and are also injected into real or simulated interferometer noise for tests of burst detection algorithms.

SNR - Signal-to-noise ratio. In the context of this thesis, SNR for a particular frequency is the ratio of the power associated with an identified event to the nominal (mean-median) power of the spectrum of the data over the analyzed time segment. For an identified burst event, the SNR is generally the ratio outlined above for a single specific frequency integrated over all frequencies contributing to the excess power event.

SRC - Signal Recycling Cavity - an optic cavity between the main interferometer and output mode cleaner that recycles light that is resonant in the cavity. The interferometer's response and the shape of the strain sensitivity curve can be changed by varying the reflectivity of the signal recycling mirror (SRM) and the signal recycling cavity length. This is potentially very useful to target particular sources of gravitational waves by gaining performance in targeted frequency bands at the expense of a more broadband interferometer response.

SUS - Suspensions - used to support and passively isolate all major optics from the motion of the optics table. Each major optic has an individual suspension (except for the OMC bowtie cavity, which shares a single suspension) mounted on an actively isolated optics table. Active actuation is also used at the top stage of suspensions for

mechanical resonance damping, and at lower stages to control the optic position and orientation for global interferometer control.

TCS - Thermal Compensation System - compensates for thermal distortion of the optics due to absorbed power by thermally inducing the desired radius of curvature of the optics.

WNB - White Noise Burst - a discrete, or time-localized, burst of white noise, which is random signal with a flat frequency response. For the context of this document all white noise burst injections are assumed to be band-limited white noise bursts, with power in only a certain band-limited range of frequency. For this work, WNBs are used to test the response of event trigger generators, or single interferometer signal burst detection algorithms.

INTRODUCTION

Gravitational waves, tiny ripples in the fabric of spacetime produced by accelerating masses, will offer insight into some of the most energetic astrophysical events, such as black hole coalescences and core-collapse supernovae. The Advanced Laser Interferometer Gravitational wave Observatory (aLIGO) is on the brink of discovery, about to begin observations with newly upgraded instrumentation.

Since the prediction of gravitational waves by the theory of General Relativity a century ago [36], there has been a rich history of experiments aimed at detecting them. To date gravitational waves have not been directly observed [94], [8], [16],[14]. However, an exciting experimental result published in 1982 strongly supports the theoretical prediction of gravitational wave emission. First discovered in 1974 by Russell Hulse and Joseph Taylor, the Hulse-Taylor binary pulsar 1916+13 is a binary neutron star system where one of the partner neutron stars is a pulsar¹. When the pulsar beam emission was used to measure the orbit of the binary system to high precision, the observed orbital decay closely matched the prediction of the loss of energy due to the emission of gravitational waves [92]. Subsequent observations of binary pulsar systems have also confirmed orbital decay as expected from General Relativity [58].

¹A pulsar is a radiating neutron star that emits a directional beam of electromagnetic radiation much like a lighthouse.

Gravitational waves interact very weakly with matter and will stretch and squeeze spacetime with a tiny amplitude as they pass². In the 1960s Joseph Weber attempted to make the first direct observation of gravitational waves with series of resonant mass detectors or *bars* designed to measure small vibrations of massive aluminum cylinders as they resonated with passing gravitational waves [83]. Weber announced in 1969 that his detectors observed coincident gravitational wave signals, but during the 1970s this was discredited by other independent resonant bar experiments. It is now widely thought that the bar detector technology was not sensitive enough to register a gravitational wave signal [83].

Laser interferometry is a promising technology that measures the spacetime strain of a passing gravitational wave very precisely using the relative phase change induced in two perpendicular arms. Interferometers have largely replaced resonant bar instruments because they offer a better strain sensitivity over a wider frequency range [83]. Following more than a decade of technology prototypes, in the late 1990s three large-scale gravitational wave interferometers were built called the Laser Interferometer Gravitational-wave Observatory (LIGO) project. One LIGO detector with 4 km long perpendicular arms was established in Livingston, Louisiana, and the other site in Hanford, Washington housed two interferometers³, one with 4 km arms and the other with 2 km arms.

Initial LIGO, the first version of the LIGO detectors, began observing in 2002. The instrumentation was upgraded to Enhanced LIGO with improvements including increased laser power, and improved in-vacuum signal readout [40]. The Enhanced

²For example, the induced gravitational strain, or the fractional change in length, produced by coalescence of two compact binary objects with a total mass of $10 M_{\odot}$ at 100 Mpc would produce a strain amplitude of 10^{-21} .

³Hanford no longer houses two interferometers. Instrumentation for the second (the third LIGO interferometer) is currently packed and waiting for potential installation at a future site in India, yet to be determined [33].

LIGO interferometers were used for the 2009-2010 Science Run 6 (S6) in conjunction with other ground-based interferometric gravitational wave detectors Virgo [18] and GEO600 [97]. While no direct gravitational wave detection has yet been made, the Enhanced LIGO instruments contributed to several astrophysically interesting results including setting an upper limit on the stochastic gravitational wave background [94], setting upper limits on the gravitational wave emission of known pulsars [8], setting an upper limit on the rate of compact binary coalescences [16], and setting an upper limit on the rate of strong gravitational wave bursts [14].

In this dissertation I will focus on one potential gravitational wave source for Advanced LIGO; core-collapse supernovae, violent star deaths that are some of the most energetic events in the Universe. Blocked by the outer layers of the star, the inner mechanics of supernovae are obscure to electromagnetic observations. The only way to gain insight into the explosion mechanism driving these events is to infer the intricate interior dynamics through direct observation of gravitational waves, or learn about the pre-collapse thermodynamics with neutrino observations.

The new Advanced LIGO hardware promises a ten-fold improvement in LIGO's astrophysical observation range [2], but also presents the challenge of understanding unexpected sources of noise and how they will affect searches for gravitational wave transients. A crucial component to gravitational wave detection and astrophysical inference is the Gaussianity of the noise, and efforts to characterize and mitigate transient noise behavior in the interferometers are critical.

Chapter 1 provides an overview of the promise and potential of gravitational wave astronomy, as well as a brief derivation of gravitational waves as predicted by General Relativity, with some insight into the properties most relevant to this work.

Chapter 2 contains the basic principles of interferometers, outlines how they measure a response to passing gravitational waves, and summarizes the instrumentation of the Advanced LIGO interferometers.

Chapter 3 details the parameterization of burst-like excess power events in time and frequency and presents a study evaluating the performance of single-interferometer burst recovery algorithms, or event-trigger-generators, in efficiency and event parameter estimation.

Chapter 4 presents an overview of recent core-collapse supernova models and details a performance study that gauged the efficiency and waveform reconstruction accuracy of current LIGO burst analysis software using supernova waveforms injected into aLIGO-colored Gaussian noise.

Chapter 5 gives an overview of the impact of terrestrial noise on the recovery of astrophysical gravitational wave transients with Advanced LIGO. Much focus is devoted to elevated levels of transient ground motion - a particularly egregious problem in past science runs. This chapter presents a brief overview of the Advanced LIGO seismic isolation instrumentation and summarizes the results of a series of studies characterizing the mitigation of transient ground motion by the actively controlled isolation stages.

Chapter 6 details the results of a performance study similar to that outlined in Chapter 4. This study uses realistic non-Gaussian noise from previous data collection runs⁴ instead of Gaussian noise to gauge the impact of noise transients on the detectability and reconstruction of CCSN.

⁴The data from previous science runs is recolored to expected Advanced LIGO sensitivity to simulate how realistic Advanced LIGO data may behave.

Finally, **Chapter 7** summarizes the results of this work, outlines the broader potential impact of improving the ability of current algorithms to accurately reconstruct rich, complex gravitational wave signals, and suggests future supplementary work.

CHAPTER 1

GRAVITATIONAL WAVES

In this section we introduce gravitational waves and lay the basic foundation of gravitational wave interaction with matter in the context of terrestrial interferometers.

1.1 General Relativity and gravitational waves

Gravitational waves are reverberations in spacetime generated by accelerated mass. They were first predicted as a consequence of Einstein's Theory of General Relativity, published in 1915 [36].

General relativity states that matter curves spacetime and the motion of matter is dictated by spacetime curvature. The relationship governing this interaction between matter and spacetime is described by Einstein's field equations (Equation 1.1),

$$G_{\mu\nu} = 8\pi T_{\mu\nu} \tag{1.1}$$

where

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} \tag{1.2}$$

and $g_{\mu\nu}$ is the general description of the spacetime metric. ($G = c = 1$).

To extract a gravitational wave solution from here, we follow a similar treatment to Saulson [83] and Baumgarte and Shapiro [23].

We first derive the effect of introducing a small perturbation to the simple, flat Minkowski spacetime metric, $\eta_{\mu\nu}$:

$$\begin{bmatrix} -t & 0 & 0 & 0 \\ 0 & x & 0 & 0 \\ 0 & 0 & y & 0 \\ 0 & 0 & 0 & z \end{bmatrix}$$

of the form:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{1.3}$$

where $h_{\mu\nu}$ is a small metric strain perturbation deviation from flat Minkowski spacetime.

It greatly simplifies the calculation of the Riemann tensor in Equation 1.2 to introduce the ‘trace reversed’ perturbation \bar{h} :

$$\bar{h}_{ab} = h_{ab} - \frac{1}{2} \eta_{ab} h \tag{1.4}$$

with a trace in four-dimensional spacetime equal to $-h$.

After imposing the Lorentz gauge condition of

$$\nabla_a \bar{h}^{ab} = 0 \tag{1.5}$$

Einstein's equation simplifies to a wave equation:

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0 . \tag{1.6}$$

This wave equation has the usual solutions of the form:

$$h = Ae^{i(2\pi ft - \mathbf{k} \cdot \mathbf{r})} \tag{1.7}$$

where A is the wave amplitude, f is the wave frequency, k is the wave number (proportional to $\frac{f}{c}$), and the wave is propagating in $\hat{\mathbf{k}}$.

Further imposing the *transverse traceless* or ‘TT’ gauge on the small spacetime perturbation strain h yields two independent polarization states: cross and plus [23].

With these gauge condition imposed, the gravitational wave strain perturbation takes a very simple form for a wave propagating in the \hat{z} direction in Minkowski space (i.e. the weak field limit):

$$\mathbf{h} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

where h_+ is the amplitude of the plus polarized part of the gravitational wave and h_\times is the amplitude of the cross polarized part. Note that for a plus polarized wave propagating in \hat{z} , this indicates an equal and opposite response along x and y, whereas cross polarized waves simultaneously affect both the x and y degrees of freedom.

1.1.1 Interaction with matter

The effect of a passing gravitational wave on matter is most straightforward to understand using a linearly plus-polarized sinusoidal gravitational wave propagating in the \hat{z} direction as an example.

Then the strain becomes:

$$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h & 0 & 0 \\ 0 & 0 & -h & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

which induces a strain effect on matter as seen in the left panel of Figure 1.1. As the wave passes, it's effect on a ring of matter first stretches the ring along the x direction with a strain amplitude of h , then relaxes it back to its original shape, then stretches the ring along y direction with a strain amplitude of h , and so on.

This behavior of matter affected by a passing gravitational wave motivates the design of the LIGO detectors: an interferometer composed of two perpendicular arms, described in Chapter 2.

1.1.2 Gravitational wave emission

The Einstein equations are effectively impossible to solve analytically in strong-field gravity, or close to a massive object. In order to estimate useful properties of gravitational wave emission, we will make use of the post-Newtonian approximation, where

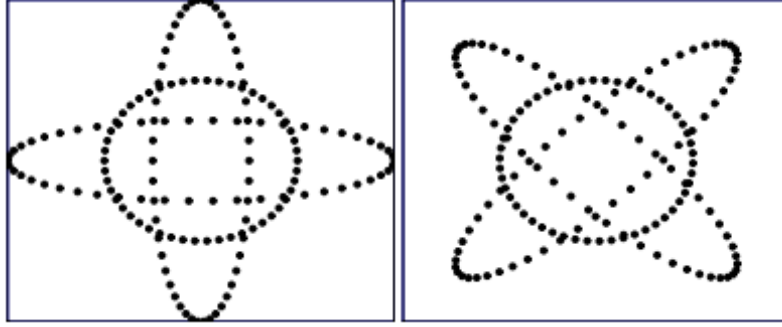


Figure 1.1: The effect of plus (left) and cross (right) gravitational wave polarizations on a ring of matter. As a gravitational wave propagates perpendicular to a circular ring of matter, the ring will be first stretched along one degree of freedom and squeezed along its counterpart, then relaxed back to neutral, then squeezed and stretched along the counterpart degree of freedom. The cross and plus polarizations are separated by a 45 degree rotation, but they are orthogonal. A gravitational wave can have an arbitrary polarization composed of a linear combination of cross and plus polarization. Reproduced from a living review by Sathyaprakash and Schutz [81].

the strain induced by a gravitational wave propagating in a locally flat spacetime metric a distance r from the source is given by:

$$h_{ij} = \frac{2}{r} \frac{d^2 I_{ij}}{dt^2} \quad (1.8)$$

where I_{ij} is a tensor describing the quadrupole moment of the system [81].

Using this approximation we can estimate the total gravitational wave energy emitted by a system, the gravitational wave luminosity, and the special case of the characteristic frequency of an orbiting binary object.

The total energy emitted in the form of gravitational waves by a system in terms of its quadrupole moment is:

$$E_{GW} = \frac{G}{5c^5} \int_{-\infty}^{+\infty} \left(\ddot{I}_{ij} - \frac{1}{3} \dot{\ddot{I}} \right)^2 \quad (1.9)$$

where I is the trace of the quadrupole moment tensor I_{ij} .

The gravitational wave luminosity is expressed as [81]:

$$L = \frac{c^5}{5G} \sum_{i,j} \ddot{I}_{ij} \ddot{I}_{ij} - \frac{1}{3} \ddot{I}^2 . \quad (1.10)$$

Note that the factor $\frac{c^5}{5G}$ is quite large - equal to 7.2×10^{51} Watts!

The characteristic gravitational wave frequency of the radiation of a compact object of mass M and radius R can be roughly approximated as [81]:

$$f = \frac{1}{4\pi} \left(\frac{3M}{R^3} \right)^{1/2} . \quad (1.11)$$

This sets an upper limit on the total mass of a system that could produce quadrupole radiation within the expected Advanced LIGO frequency band, detailed in Chapter 2, of about $10^4 M_\odot$ [81].

1.2 The gravitational wave sky

Our current understanding of the Universe is based on electromagnetic observations. However, there are many open questions that are difficult or impossible to address using light as a probe, such as investigating properties of neutron stars, supermassive black hole formation, the physical engine driving core-collapse supernova explosions, or properties of the very early Universe. Weakly interacting gravitational waves carry

valuable information about such astrophysical phenomena, and if observed, will enable inferences about their origin and behavior.

The gravitational wave sky offers a rich, diverse array of sources that Advanced Era ground-based interferometers may observe, with neutron star binaries the most likely source. The gravitational radiation produced by the coalesce of these systems is well modeled: the frequency and amplitude both increase until the final merger, producing a signal pattern called a *chirp*. Other potential transient gravitational wave sources include core-collapse supernovae, described in great detail in Section 1.2.1, and inspiraling binary systems with total mass ranging from neutron star pairs to black hole pairs, potentially orbiting with some eccentricity. Advanced detectors may also be sensitive to gravitational wave emission from neutron star pulsars spinning with some asymmetry, which produce a predictable periodic signal. Another potential source is a stochastic gravitational wave background, perhaps generated from the superposition of signals from distant sources or the remnants of quantum gravitational effects in the very early Universe manifested in non-uniformities in spacetime [83].

The work described in this document focuses mainly on generic transient, or *burst* sources; core-collapse supernovae in particular. See [1, 52] for recent results of gravitational wave searches for long duration continuous wave sources and [9] for long duration stochastic background sources. See also recent results for searches for the coalescence of compact binary objects like binary black holes [6], intermediate mass black holes [10, 11], and binary neutron stars [4].

In Chapter 3, recent burst search results using interferometric data such as the most recent all-sky search [14], the search for GWs associated with a pulsar timing glitch [13], and the search for GWs associated with observed gamma-ray bursts [15] are summarized and the resulting implications for this work are discussed.

1.2.1 Core-collapse supernovae

Core-collapse supernovae are produced by mid-range mass stars (about 8-50 M_{\odot}). When electron degeneracy pressure can no longer stably support the iron core against gravity, the core collapses within a few seconds into a proto-neutron star, releasing an enormous amount of energy. Neutrons and neutrinos are formed via reversed beta-decay, releasing $\sim 10^{46}$ joules in short burst lasting only a few seconds [19]. Once the density of the proto-neutron star exceeds neutron degeneracy, the outer layers of the collapsed core bounce and generate a shock wave [24].

Understanding the explosion mechanism driving core-collapse supernovae after the core bounce to an explosion (see Figure 1.2) is a key open question in astrophysics. Early models have showed that the initial core-collapse launches a shock wave that bounces off of the very dense proto-neutron star (PNS), but this shock wave stalls in the infalling matter of the outer layers of the star [53]. The mechanism that reignites the shock wave and drives the star to explode, as observed via electromagnetic radiation, may be inferred from gravitational wave and neutrino emission of the event.

In 1987, supernova 1987A was observed in the Large Magellenic Cloud, roughly 50 kpc (approximately 163,000 light years) from earth. In addition to being one of the brightest observed stellar events, 1987A was also one of the first instances of multi-messenger astronomy. Three independent neutrino observatories detected short neutrino bursts less than three hours before visible light from the event was detected (Kamiokande II, in Japan, the Irvine-Michigan-Brookhaven (IMB) detector in the United States, and the Baksan Neutrino Observatory in Russia) [50], [25], [20]. This groundbreaking observation confirmed that core-collapse supernovae form proto-neutron stars and radiate more than 99% of the gravitational energy in the form of neutrinos [39].



Figure 1.2: The electromagnetic signature of the remnant of supernova G292.0+1.8, observed by the Chandra X-ray Observatory. In the center is a pulsar, a rapidly rotating neutron star and remnant of the original star's core.

Although the observation of 1987A does not offer insight into the explosion mechanism problem, core-collapse supernovae are predicted to produce gravitational waves generated by violent asymmetric matter flows [62]. Advanced LIGO is predicted to be sensitive to gravitational waves from galactic core-collapse supernovae and may be key to deducing the interior physical processes of collapsing stars. The predicted event rate for galactic supernovae is low, a few per century at most, but the science return potential is great. If a galactic supernova occurs during the Advanced LIGO observation, then for the first time joint electromagnetic radiation, neutrino, and gravitational wave observation will be possible for an astrophysical event. Observing gravitational waves will be crucial to determine the internal dynamics of these events.

The currently favored models for the mechanism driving core-collapse supernovae explosions include neutrino heating and rotating core collapse, both projected to produce a burst of detectable gravitational waves lasting less than 2 seconds in close proximity to a neutrino signal, as shown in Figure 1.3. Each mechanism is predicted




Figure 1.3: A timeline showing the relative expected timing of the different astrophysical messengers produced by a core-collapse supernova. A short gravitational wave burst expected to last no more than a few seconds is predicted to be emitted at roughly the time of the core bounce. A burst of neutrinos lasting roughly ten seconds is emitted from the collapsed star starting at core bounce. Lastly, an electromagnetic signature is expected to be observed when the shock first breaks out of the outer layers of the star, which could take on the order of minutes up to days after the initial core collapse [19].

to produce gravitational wave signals distinct in morphology that can be uniquely mapped to their corresponding physical processes, detailed in Chapter 4.

Accurately reconstructing the waveform produced by a gravitational wave burst event is part of the parameter estimation effort following the initial event identification. Two algorithms are planned to be used for this purpose: cWB2G and BayesWave. Their ability to reconstruct waveforms produced by core-collapse supernovae will be discussed in Chapter 3.

For core-collapse supernovae in particular, there are several potential benefits to optimizing the accuracy of burst waveform reconstruction algorithms:

- Enable accurate interpretations of core-collapse supernovae physics from established mappings between expected gravitational wave signal and associated physical mechanism, such as the delayed burst of excess power due to neutrino-driven convection [46].

- Facilitate precise timing measurements between physical features in the gravitational wave signal and other forms of information from a core-collapse event, especially neutrino emission, which would be very powerful in constraining the likely explosion mechanism models [76].
- Lay the foundation to accurately infer the most likely physical model for explosion using current core-collapse supernovae model selection algorithms [62].

CHAPTER 2

THE LIGO DETECTORS

2.1 The LIGO detectors

The Laser Interferometer Gravitational-wave Observatory (LIGO) detectors sense gravitational waves by measuring *strain*, or the change in length over the length of two 4 km perpendicular arms. To accomplish this, a 1064 nm laser beam is split down each of the arms, reflected back via highly reflective suspended mirrors, and recombined. This is shown in Figure 2.1, along with the Fabry-Perot cavities¹ that form each arm of the interferometer. These resonant cavities build up power in each arm to increase the effective response of the interferometers to a shift in phase produced by a passing gravitational wave. Any change in relative length between the arms will cause a relative phase shift in the circulating light and alter the interference pattern of the recombined light. We can infer the dynamics of an astrophysical gravitational wave source from the space-time perturbation manifested in this phase shift, measured as strain.

As a result of this fundamental interferometric design, a range of fundamental limitations to strain sensitivity are expected; including seismic noise [65], shot noise [27],

¹Fabry-Perot optical cavities store light reflected between two parallel optics, generally with highly reflecting coatings [83].




Figure 2.1: A simplified interferometer layout. An infrared input beam of laser light is split so that equal parts resonate in two perpendicular arm cavities composed of an input and end test mass optic. A relative change in length between these two perpendicular cavities is sensed as a change in relative phase of the light, manifested in an interference pattern at the output port. Fabry-Perot cavities that increase circulated power in the arms are indicated by thicker red laser power lines. Not shown in this simple layout is the power recycling mirror installed between the input laser and the beam splitter used to recycle the power exiting the beam towards the input laser.




Figure 2.2: Fundamental limiting sources of noise for interferometric terrestrial gravitational wave detectors, depicted here for the nominal Advanced LIGO mode of operation (zero de-tuned signal recycling cavity, high power). The Advanced LIGO instruments are designed to be nearly quantum noise limited everywhere above 10 Hz. The most potentially limiting noise sources (quantum noise, seismic noise, gravity gradients, suspensions thermal noise, and coating Brownian noise) are described in the text. Reproduced from [2].

radiation pressure [32], gravity gradients [34], suspension thermal noise [30], and optic coating Brownian noise [38]. Figure 2.2 shows the expected contributions of these fundamental sources to the Advanced LIGO design sensitivity curve².

Here we summarize the most potentially limiting fundamental noise sources for Advanced era interferometers. For ground-based gravitational interferometers in general, at low frequencies³ the detectors are most limited by seismic motion, or local perturbations in the ground that propagate through the stages supporting the optic suspensions and the suspensions themselves, causing the optics to move. Shot noise is a quantum noise effect due to the uncertainty in detecting photons due to the quantized nature of light. Shot-noise-limited strain sensitivity improves with increased

²Shot noise and radiation pressure are combined and shown as ‘Quantum noise’ in Figure 2.2.

³For Advanced LIGO, seismic noise is expected to dominate below roughly 15 Hz.

laser power. Shot noise is dominant at higher frequencies, above the interferometers' most sensitive frequency region, 100-500 Hz. Radiation pressure occurs when photons impart momentum on a body, usually a very small effect. However, the higher laser power planned for Advanced LIGO is expected to exert a significant force on the suspended optics via radiation pressure in the frequency range between where shot noise is dominant and seismic noise starts to become dominant. Gravity gradients are fluctuations in the local gravitational field due to variations in the distribution of matter near the interferometer, which is potentially dominant below about 1Hz but not expected to be limiting for Advanced LIGO [34]. Optical coating Brownian noise is displacement noise due to mechanical dissipation in the coatings on the surface of the optics, which in Initial LIGO was the dominant source of noise in the interferometers' most sensitive region, at about 100Hz [48]. Suspension thermal noise is a related effect due to dissipation of energy through the fibers that connect the suspension stages [45], [82]. The amplitude of both optical coating Brownian noise and suspension thermal noise depends largely on the properties of the materials used.

There are several other expected noise sources for ground-based interferometric gravitational wave detectors, such as fluctuations in laser power or frequency, angular drift of the optics, or the noise floor of various sensors and actuators. These are known as *technical noise* and are largely accounted for in the interferometer design. There are also many other sources of terrestrial noise that may couple to the gravitational wave strain readout of the interferometers and reduce the sensitivity of the gravitational wave astrophysical searches. Some of these are expected and unavoidable, such as earthquakes, local anthropogenic seismic activity, or environmental effects like high winds [35]. Some are not anticipated. A set of known examples of noise artifacts are summarized in Section 2.3.2 and discussed at length for the most recent LIGO science run in [3].

2.2 Gravitational wave astronomy with LIGO

The interaction of a LIGO-like ground-based gravitational wave interferometer with a gravitational wave is presented in this section, following the treatment of Saulson [83] and a recent review by Sathyaprakash and Schutz [81].

The LIGO instruments are advanced versions of a simple Michelson interferometer, which is essentially shown in Figure 2.1 minus the input test masses that enhance the light circulation in the arm cavities. Michelson interferometers measure the relative time the light beams take to complete a round trip in each of the arms [83]. Light enters the interferometer from the coherent laser source, is split, and enters the arms exactly in phase. If each of the two arms has the same unaltered length⁴, L , the light from the arms will return to the beam splitter exactly in phase⁵.

When the amount of time it takes light to make a round trip is changed due to the stretching and squeezing of space caused by a passing gravitational wave, the effective travel time of the light beams in each arm will change correspondingly and the light beams will return to the beam splitter out of phase. It is this relative phase difference, or difference in the relative round-trip travel time of light in the two arms, that the LIGO interferometers measure.

To understand this, we can trace the beam path shown in Figure 2.1. The input laser light entering the interferometer has some electric field of the form

$$E = E_0 e^{i(2\pi ft - kr)} \tag{2.1}$$

⁴plus or minus some integer value of the light wavelength.

⁵Assuming an infinitely thin beam splitter, for simplicity.

where E_0 is the wave amplitude, k is the familiar wave number governing the travel time in a medium, f is the frequency of the light, and t and r are varying time and distance [83].

The light then encounters the beam splitter, which has 50 percent transmission. The light transmitted to the x-arm then has an electric field of the form

$$E_x = i \frac{E_0}{\sqrt{2}} e^{i(2\pi ft - k_x x)} \quad (2.2)$$

and the light reflected to the y-arm has an electric field of the form

$$E_y = \frac{E_0}{\sqrt{2}} e^{i(2\pi ft - k_y y)} \quad (2.3)$$

where the k_x and k_y wave numbers depend on the effective travel time of the light traveling in that direction [83].

Ignoring the input test masses for this simple case, the light next is reflected by the end test mass optics, which have very near 100 percent reflectivity for the 1064 nm laser light. After reflection at the end of the cavity the light next encounters the 50 percent transmissive beam splitter again, which has the same effect on the light incoming from the arms.

Just taking the part of the light that reaches the output port, the light transmitted through the beam splitter from the y-arm and the light reflected by the beam splitter from the x-arm, the total electric field at the output port after re-combining the beams is:

$$E_{tot} = \frac{iE_0}{2} e^{i(2\pi ft - k_x x)} + \frac{iE_0}{2} e^{i(2\pi ft - k_y y)} \quad (2.4)$$

where the total effective round trip distance made in the x-arm is $2L_x$ and the total round trip in the y-arm is $2L_y$. Substituting this for x and y and simplifying yields:

$$E_{tot} = iE_0 e^{(2\pi ft - k_x L_x - k_y L_y)} \cos(k_x L_x - k_y L_y) \quad (2.5)$$

where the quantity $k_x L_x - k_y L_y$ is $\Delta\phi$, the difference in phase acquired by the light traveling the round trip to the end test mass and back in each of the two arms.

The electric field can be expressed in terms of $\Delta\phi$ as:

$$E_{tot} = iE_0 e^{(2\pi ft - \Delta\phi)} \cos(\Delta\phi). \quad (2.6)$$

From this equation we can see that for the case of the simple Michelson we have described, the most power reaches the output port when the relative phase difference between the two arms $\Delta\phi$ is zero, and no power reaches the output port when this phases difference is maximized to an integer multiple of $\pi/2$.

With a mathematical description of how interferometers measure relative phase using interfered laser light, we can relate this measured phase difference to gravitational wave strain.

Assuming we make use of the same coordinate system described in Chapter 1, using the Transverse Traceless (TT) gauge so that our spatial coordinate system is fixed to the geodesics of freely falling masses, we consider the round trip travel time of light traveling between two freely falling masses. This is analogous to considering the case of one arm of a simple gravitational wave interferometer.

Recall our spacetime metric is a small perturbation of the simple Minkowski metric (Equation 1.3). From the same Special Relativity framework, we know that the spacetime interval, ds^2 , between two points connected by a ray of light is zero.

Setting the spacetime interval to zero:

$$0 = ds^2 = (\eta_{\mu\nu} + h_{\mu\nu})dx^\mu dx^\nu \quad (2.7)$$

and focusing on the case of the x degree of freedom in particular we see:

$$0 = -c^2 dt^2 + (1 + h_{x_{dir}})dx^2 \quad (2.8)$$

where $h_{x_{dir}}$ is the induced spacetime strain along the x-direction, analogously the x-arm interferometer beam path, from an arbitrary direction \hat{r} , and is a function of emitted gravitational wave frequency, time, and the dot product between the direction of gravitation wave propagation and the beam path: $\mathbf{r} \cdot \mathbf{x}$.

We can solve Equation 2.8 for the total travel time of the light traveling between two freely falling masses in the x-direction in the presence of gravitational wave perturbation $h_{x_{dir}}$:

$$c^2 dt^2 = (1 + h_{x_{dir}})dx^2. \quad (2.9)$$

Integrating to obtain the total travel time over the length, L , of our beam path in x, we see:

$$t = \frac{2}{c} \int_0^L (1 + h_{x_{dir}})^{\frac{1}{2}} dx. \quad (2.10)$$

And performing the same analysis of Equations 2.8 through 2.10 yields an identical result for the y-arm path, with the appropriate $x \rightarrow y$ substitutions.

Making the assumption that the gravitational wave is plus polarized with a known frequency f , known amplitude $|h|$, and propagating in the \hat{z} direction simplifies the calculation and offers insight into the response of interferometric detectors to gravitational waves.

Making the additional assumption that the induced spacetime strain is constant for the duration of the light's round trip yields the observation that the total round trip travel time for a photon is:

$$t = \frac{2L}{c} + \frac{2L|h|}{c} \quad (2.11)$$

where the travel time difference due to the passing gravitational wave t_{diff} is:

$$t_{diff} = |h| \frac{2L}{c} \quad (2.12)$$

for both the x and y arms.

The LIGO interferometers measure the relative phase shift between the two measured paths along the x and y arms. The phase shift induced from the relative time travel difference in each of the arms is given by:

$$\Delta\phi = 2\pi f \Delta t \quad (2.13)$$

where

$$\Delta t = t_x - t_y. \quad (2.14)$$

For this case of a gravitational wave propagating perpendicular to the arms of the interferometer, the strain amplitude will be equal and opposite for each arm, mirroring the left panel of Figure 1.1 so that $h_{xdir} = |h|$ and $h_{ydir} = -|h|$.

Then Δt becomes:

$$\Delta t = \frac{2L}{c} + \frac{2L|h|}{c} - \frac{2L}{c} + \frac{2L|h|}{c} \quad (2.15)$$

$$\Delta t = |h| \frac{4L}{c}. \quad (2.16)$$

Therefore the phase shift of light observed for photons with this altered travel time can be written as:

$$\Delta\phi = |h| \frac{8\pi Lf}{c}. \quad (2.17)$$

The observed phase shift in each arm directly depends on the amplitude and frequency of the gravitational wave, as well as the total path length traveled, L , which is significantly increased from the simple round trip case through the use of Fabry-Perot resonance cavities.

Considering the more complex and realistic case that the induced local gravitational wave strain is not constant will reveal more about the interferometer's sensitivity to gravitational wave sources of arbitrary polarization and sky location relative to the interferometer.

Keeping our assumptions that a gravitational wave of frequency f is plus polarized and propagating in the \hat{z} direction, we now assume a gravitational wave strain function of

$$h(t) = |h|e^{i2\pi ft} \quad (2.18)$$

instead of a constant strain amplitude $|h|$. Then performing the analogous integration seen in Equation 2.10 yields for the difference in travel time between the two arms, Δt [83]:

$$\Delta t = |h| \frac{4L}{c} e^{i2L/c} \text{sinc}\left(\frac{\pi f 2L}{c}\right). \quad (2.19)$$

This yields a corresponding phase change between the light traveling along each arm:

$$\Delta\phi = |h| \frac{8\pi Lf}{c} e^{i2L/c} \text{sinc}\left(\frac{f2L}{c}\right). \quad (2.20)$$

To calculate this interferometer's response for an arbitrary gravitational wave propagation direction, we can repeat the same calculation integrated over arbitrary values of $|h|$ for the x and y arms. These are the directional interferometer response functions for gravitational wave beam detectors with the same basic design as LIGO, often denoted F_+ and F_x instead of $\Delta\phi$. This notation prevents confusion of the antenna

pattern itself with the angles used to relate the detector coordinate frame to the most natural sky coordinate frame, described in Figure 2.3.

The antenna pattern for plus polarization is written as:

$$F_+ = \frac{1}{2}(1 + \cos^2\theta) \cos 2\phi \cos 2\psi - \cos\theta \sin 2\phi \sin 2\psi \quad (2.21)$$

and for cross polarization as:

$$F_x = \frac{1}{2}(1 + \cos^2\theta) \cos 2\phi \sin 2\psi - \cos\theta \sin 2\phi \cos 2\psi \quad (2.22)$$

for gravitational waves from some arbitrary direction described with the angles θ and ϕ [81]. These antenna patterns relate the detector frame and the direction of the gravitational wave propagation (also called the source *sky position*) and the angle ψ , which defines the *polarization angle*, or the rotation between the source coordinate frame and the projection of the detector coordinates onto the sky frame.

The most commonly used function to describe interferometer antenna pattern is the rms response of a gravitational wave detector to pure plus polarization:

$$\bar{F} = \left(\int F_+^2 d\psi \right)^{\frac{1}{2}}. \quad (2.23)$$

For a LIGO-like gravitational wave interferometer this is:

$$\bar{F} = \frac{1}{2}(1 + \cos^2\theta) \cos 2\phi + \cos\theta \sin 2\phi. \quad (2.24)$$

The function in Equation 2.24 for the LIGO interferometric detectors, often referred as simply the *antenna pattern*, is shown in Figure 2.4, with the arms of the interferometer




Figure 2.3: The definition of angles θ , ϕ , and ψ , which translate the position and orientation of the gravitation wave source in the sky to the frame of the detector. On the left, angles θ and ϕ define the vector \hat{N} between the detector frame, defined as usual with the arms oriented along the x and y axes, and the location of the gravitational wave source in the sky. On the right, the polarization angle ψ is defined as the rotation of the sky frame relative to the detector frame. This is sometimes called an *internal* angle, because it is generally defined by the orientation of the gravitational wave source., i.e. the orientation of the axis of rotation of a spinning source, relative to the basis acquired by projecting the detector frame onto the sky along the vector \hat{N} (resulting in the pictured basis vectors \hat{e}_x^R and \hat{e}_y^R in the figure). Reproduced from [81].




Figure 2.4: Antenna pattern, \overline{F}^2 , of a LIGO-like gravitational wave detector with arms oriented in a 45 degree rotation relative to the x and y axes of the plot. Sensitivity is strongest when gravitational waves are traveling perpendicular to the two interferometer arms, and shows a significant nonzero response for most directions. Reproduced from [81].

laid along the two lower axes. Sensitivity is strongest when gravitational waves are traveling perpendicular to the two interferometer arms. There are only four directions of gravitational wave propagation relative to the interferometer frame (along x and y) close to which the interferometers have no sensitivity.

The sensitivity of ground-based interferometric detectors to gravitational waves from a broad swath of the sky is a great advantage in that passing gravitational waves are unlikely to be missed due to their direction of travel relative to the interferometers. However, this same trait also makes identifying the sky location of a candidate event very challenging, as discussed in Section 2.3.1.

2.3 Advanced LIGO

As its name may suggest, Advanced LIGO is a significant technological improvement upon a previous generation of terrestrial gravitational wave interferometers. The initial LIGO detectors were constructed beginning in 1992, and collected data over a series of five separate science runs spanning from 2002-2007. No gravitational wave signals were detected, but some interesting astrophysical upper limits were set [94], [12]. The initial LIGO instrumental design included a Fabry-Perot resonance cavity in each interferometer arm as seen in the simplified interferometer layout in Figure 2.1 plus a power recycling mirror used to recycle light exiting the beam splitting toward the input laser. The input laser light for Initial LIGO was stabilized with a series of input optics that led to the Advanced LIGO input optic design [17].

From 2007 to 2009, Enhanced LIGO was improved over Initial LIGO with increased laser power, improved signal readout, and the addition of output optics that removed undesirable beam shape features and conditioned the signal [2]. Enhanced LIGO collected data at LIGO's initial design sensitivity during the most recent LIGO science run, S6, in 2009-2010. Again, no gravitational wave signals were observed, but further interesting upper limits were set [8], [16], [14].

Beginning in 2010, the initial and enhanced LIGO instrumentation was uninstalled from the vacuum chambers. The next five years marked an intense effort to construct, install, and commission the Advanced LIGO instrumentation: all new and improved technology for the optics, laser, seismic isolation, electronics, and other components. The Advanced LIGO design and components are the result of decades of research and prototyping conducted by scientists and engineers around the world [2].

The Advanced LIGO instrumentation improves detector sensitivity over Initial LIGO in all frequency regimes. Upgrades include the installation of active seismic isolation




Figure 2.5: An illustration of the gain in low frequency sensitivity achieved with the new Advanced LIGO seismic isolation instrumentation. The strain sensitivity of the LIGO Hanford interferometer during initial LIGO, science run 5, is shown in the red curve. The Advanced LIGO design sensitivity curve is in black. The expected strain sensitivity improvement is over six orders of magnitude at 10 Hz. Reproduced from [54].

and improved suspensions systems for each mirror as well as improvements to the optic coatings, suspensions, and input laser [48]. The new Advanced LIGO instrumentation is designed to reduce seismic noise by six orders of magnitude at 10Hz and thermal noise by a factor of 10 around 100Hz [2]. Once at design sensitivity, Advanced LIGO data is expected to be used for gravitational wave searches down to about 10 Hz; a 30 Hz improvement over initial LIGO, as shown in Figure 2.5 [2].

The Advanced LIGO layout differs from the simple interferometer scheme presented in Section 2.1 with the addition of a Power Recycling Cavity (PRC), Signal Recycling Cavity (SRC), as well as many other auxiliary optics and support subsystems. The power and signal recycling cavities are both three-mirror folded cavities. The power recycling cavity increases the power circulating in the interferometer arms. The signal recycling cavity increases the effective signal bandwidth, or enables a tunable signal




Figure 2.6: The anticipated Advanced LIGO noise curves in three different configurations. The red curve shows the expected noise in low power mode, with an input power of 25W. In black is the nominal broadband operational mode, with a ‘high power’ input power of 125W. In blue is a Binary Neutron Star (BNS) optimized mode which makes use of the tuneable signal recycling cavity to target the expected signal frequencies of these sources at the cost of greatly reducing broadband sensitivity. This is achieved by lowering the Signal Recycling Mirror (SRM) transmission and detuning the cavity, or introducing some some offset so that only some target frequency is resonant. Reproduced from [2].

bandwidth which can achieve enhanced strain sensitivity in a targeted frequency range at the expense of a more broadband frequency response as shown in Figure 2.6 [2].

In general, a signal recycling cavity recycles light that is resonant in the cavity. A *detuned* signal recycling cavity has a cavity resonance that is shifted away from the carrier light frequency. The bandwidth of the signal recycling cavity resonance depends on the reflectivity of the signal recycling mirror (SRM), and the *tuning frequency*, or the frequency of maximum response of the cavity, is determined by the

cavity length⁶. In this way, by changing the length of the signal recycling cavity or the transmissivity of the signal recycling mirror, we are changing the interferometer's response and the shape of the strain sensitivity curve. This is potentially very useful to target particular sources of gravitational waves with known expected frequencies, as shown for the example of binary neutron stars in Figure 2.6 [49].

Figure 2.7 gives a summary overview of the major interferometer elements. All interferometer components shown are housed within an ultra-high vacuum containment system and seismically isolated from local ground motion, with the exception of the input laser. Each chamber and major optic is color coded to indicate a different configuration of active seismic isolation and suspension, respectively. For perspective on the complexity of the Advanced LIGO project, several major subsystems detailed below are omitted from this overview schematic entirely⁷.

The brief descriptions below should provide the reader enough familiarity with Advanced LIGO instrumentation to understand the work detailed later, especially in Chapter 5. However, there are other major components of Advanced LIGO design that are not summarized here, including DC readout technique, calibration of the $h(t)$ strain signal, angular and length sensing and control, scattered light mitigation, automation of interferometer control, and data acquisition. These important design components as well as those summarized below are described in [2].

⁶A *tuned* signal recycling cavity has a resonance centered at the carrier light frequency. In some Advanced LIGO configuration documentation, such as [5], this is referred to as the 'zero-detuned' Advanced LIGO configuration. The difference between the resonance frequency of the signal recycling cavity and carrier frequency is called the *detuning frequency*.

⁷Including thermal compensation, light scattering control, and physical environment monitoring.

Advanced LIGO
 Optical Layout, L1 or H1
 with Seismic Isolation and Suspensions
 G1200071-v3
 J. Kissel Nov 4 2013




Figure 2.7: A schematic of the Advanced LIGO configuration, its many optic cavity components, and different seismic isolation apparatuses applied to each optic and/or optic table [55]. Of the 11 different optic chambers shown, ‘HAM’ chambers house auxiliary optics used for power recycling, mode cleaning, and other purposes. ‘BSC’ chambers house the core interferometer’s optics: the beam splitter, two sets of inner and end test masses that comprise each arm, the reaction suspension chains used for quiet actuation of the test masses, and transmission monitor optics used for arm length stabilization, described later. See 5.4 for a description of seismic isolation configurations.

Advanced LIGO auxiliary subsystems:

- The **pre-stabilized laser (PSL)** is a Nd:YAG⁸ 1064nm laser that is capable of emitting up to 180W of power. The input power provided by the PSL is conditioned to be extremely stable in intensity and frequency via a series of nested sensors, reference cavities, and control loops. A bow-tie pre-mode cleaner cavity is also used to select for a fundamental Gaussian beam mode shape (TEM_{00}).
- The **Input Mode Cleaner (IMC)** is a three mirror folded cavity used to enforce beam mode quality by filtering out higher order spatial modes. The IMC also serves as a high-quality reference cavity for input light frequency stabilization, being much longer than the PSL's internal reference cavity.
- The **Output Mode Cleaner (OMC)** is a four-optic bowtie cavity designed to remove all higher order beam shape modes and any injected sidebands used for cavity control so that only good-quality carrier frequency light imprinted with the gravitational wave signal reaches the final readout photodiodes.
- The **seismic isolation (SEI)** subsystem provides active mechanical isolation between local ground motion and the interferometer components. The SEI configuration is discussed in detail in Chapter 5.
- The **suspensions (SUS)** are used to support and passively isolate all primary optics from the motion of the optics table. The different suspension types can be seen in Figure 2.7. The number of suspension stages ranges up to four for the test mass optics and their partner suspension chains. The passive isolation

⁸Nd:YAG is a triply-doped neodymium-doped yttrium aluminum garnet solid state laser crystal developed in the 1960s known for good spectral purity.

and active motion control provided by the suspensions is described in Chapter 5.

- The **Thermal Compensation System (TCS)** is intended for use with high power interferometer configurations, as will be needed to reach Advanced LIGO's design sensitivity⁹. At high power, the arm cavities are expected to have a stored optical power of 800 kW. This intense Gaussian-shaped beam is absorbed non-uniformly by the test mass optics, leading to thermal lensing optical distortion effects. The Thermal Compensation System is designed to compensate for these effects by thermally inducing some correction to the optic's radius of curvature. The TCS has three major components: a circular ring heater for heating the test masses outside of the main Gaussian beam profile, a CO₂ laser capable of projecting various annular shapes onto the compensation plates (CP) indicated in Figure 2.7, and a Hartmann wavefront sensor used to measure the current optic shape and diagnose any thermal distortion. Each test mass optic (end and input test masses of both arms) are equipped with a ring heater, which is made from nichrome wire wrapped around two semi-circular glass rods, and a Hartmann wavefront sensor. The two input test mass optics are also equipped with CO₂ laser systems designed to project heating patterns configured for two different tasks. One is an annular pattern around the main Gaussian beam which can add power and finesse to thermal corrections made by the ring heaters. The other is a central pattern used to keep the optics warm during a loss of interferometer resonance so that the instruments can be quickly restored to the high power operating point [2].

⁹The nominal mode interferometer configuration seen in Figure 2.6 requires high power (125W) input to the interferometer at the Power Recycling Mirror (PRM).

- The **Arm Length Stabilization (ALS)** system is used to acquire and control resonance in each of the two arm cavities independent from the rest of the interferometer. The ALS system is an important new feature of the process to reach the instrument’s operating point where the full array of interferometer cavities are in resonance. This is known as acquiring *lock*¹⁰. This system uses Nd:YAG lasers to inject 532nm (green) light at both end test masses. The surfaces of all four test mass optics are treated with specialized coatings that allow for a high arm cavity finesse for 1064nm light, and relatively low finesse for 532nm light. This allows for more robust lock acquisition of the 4km long arm cavities using the green laser light without loss of circulated power for the main beam. The lock acquisition process involves many steps and is fully described in [90].
- The **Physical Environment Monitoring (PEM)** system is a suite of environmental noise sensors positioned at critical points around the interferometer. Seismometers measure local ground motion, accelerometers installed inside and outside the vacuum chamber measure the motion of key structures, magnetometers monitor disturbances in local electromagnetic fields, microphones record acoustic noise, and radio receivers are used to monitor frequencies near to injected modulation frequencies used for optic cavity locking and control [35].

As seen in Figure 2.7, there are many distinct optic cavities that must be locked in concert in order to operate the full interferometer, including the input mode cleaner (IMC), power recycling cavity (PRC), signal recycling cavity (SRC), the output mode

¹⁰Analogously, global interferometer control loops are used to maintain *lock*. When these loops develop some instability and the full interferometer resonance cannot be maintained, this is referred to as *lock loss*.

cleaner (OMC), and each of the arm cavities formed by the inner and end test masses, used to measure the differential change in length.

The **interferometer cavity degrees-of-freedom** used for full interferometer control are shown in Figure 2.8 and defined in Table 2.1.

- The **Michelson interferometer (MICH)** is defined by the difference in length between the beam splitter (BS) and each of the input test mass optics (ITMX and ITMY) referenced in Figure 2.8 as l_x and l_y . This degree of freedom does not contain any resonating cavities, but is used to maintain a small asymmetry between l_x and l_y , known as Schnupp asymmetry, so that sideband power used for signal recycling cavity control is transmitted even when the carrier power is not [2].
- The **differential arm (DARM)** degree of freedom is defined by the difference between the two interferometer arm lengths, referenced in Figure 2.8 as L_x and L_y . This mode is the basis for the measured gravitational wave signal.
- The **common arm degree of freedom (CARM)** is the average length of the two arms, L_x and L_y . The CARM mode control loop actuates on the input laser frequency so that this frequency follows the common arm length (average arm length), which is very stable.
- The **Power Recycling Cavity Length (PRCL)** is controlled so that the light transmitted from the pre-stabilized laser through the Power Recycling Mirror (PRM) is in phase with the light reflected from the PRM after circulating in the interferometer and passing through the beam splitter ¹¹.

¹¹A Faraday Isolator is implemented between the PRM and the input optics (IMC) so that the light transmitted from the beam splitter cannot pass and interfere with the input light conditioning.

- The **Signal Recycling Cavity Length (SRCL)** signal is controlled to optimize the length of the signal recycling cavity such that the desired frequency of light is resonant. Here light travels between the input test masses and the signal recycling mirror (the final mirror in the signal recycling cavity). Note from Figure 2.8 that the PRCL, SRCL, and MICH degrees of freedom all share the beam splitter as part of their light path, so any position or orientation change made to the beam splitter will affect all three, for an additional degree of complexity.




Figure 2.8: The five degrees of freedom necessary for full interferometer control: MICH, the interior Michelson formed by the input test masses and the beam splitter, DARM, the differential interferometer arm degree of freedom used for gravitational wave sensing, CARM, the common arm degree of freedom calculated by the average arm length, SRCL, the signal recycling cavity length degree of freedom, and PRCL, used to control the power recycling cavity length. Definitions are listed in Table 2.1.

Table: Advanced LIGO interferometer degrees of freedom

DARM	$\frac{L_x - L_y}{2}$
CARM	$\frac{L_x + L_y}{2}$
MICH	$l_x - l_y$
SRCL	$l_s + \frac{l_x + l_y}{2}$
PRCL	$l_p + \frac{l_x + l_y}{2}$

Table 2.1: The definitions of the interferometer degrees of freedom shown in Figure 2.8: Differential Arm (DARM), Common Arm (CARM), Michelson (MICH), Signal Recycling Cavity Length (SRCL), and Power Recycling Cavity Length (PRCL).

To install, test, and commission such a complex instrument with many internal dependencies, interferometer components were installed and tested in configurations of increasing complexity, gradually building up to the full interferometer. Important configurations that are referenced in later results are laid out in Figure 2.9. Of particular note, the Power Recycled Michelson Interferometer (PRMI) configuration consists of the power recycling cavity and vertex Michelson interferometer, and the Dual Recycled Michelson Interferometer (DRMI) consists of both the power and signal recycling cavities and the Michelson.

Advanced LIGO performance is climbing toward design sensitivity even faster than anticipated. At the time of writing, the Livingston interferometer has already performed long, stable periods controlled at its operating point with an estimated Binary Neutron Star range¹² in excess of 60 Mpc.

¹²Binary Neutron Star (BNS) range is the distance from Earth at which the coalescence of a system of two $1.4 M_{\odot}$ neutron stars produces a signal-to-noise ratio (SNR) of 8 in a single detector (using matching filtering) averaged over sky location and source orientation.



Figure 2.10: The global ground-based gravitational wave interferometer network of the Advanced detector era is expected to include the two US LIGO sites (in Hanford, Washington and Livingston, Louisiana), the LIGO India interferometer, the German-British GEO600 600m interferometer in Germany, the French-Italian Virgo 3km interferometer in Italy, and the cryogenic interferometer KAGRA currently being built in Japan. Image credit: C. Mayhew and R. Simmon (NASA/GSFC), NOAA/NGDC, DMSP Digital Archive.

2.3.1 The global interferometer network

The anticipated global interferometer network of Advanced, or second generation, detectors shown in Figure 2.10, consists of five detectors with comparable arms lengths and sensitivities (3-4 km) and one with smaller arms and a lower expected sensitivity (GEO600). Together, this network of advanced interferometers will increase sensitivity to gravitational waves tenfold and enable significantly improved sky localization [44].

There are currently two LIGO interferometers, each with 4 km long arms, in Livingston, Louisiana and Hanford, Washington. Installation of Advanced LIGO hardware is now complete at both sites and commissioning to mitigate noise sources and improve lock stability is ongoing. The anticipated sensitivity ranges for likely Advanced LIGO observing scenarios are shown in the left side of Figure 2.11, reaching up to 200 Mpc for binary neutron star sources as early as 2019 [5]. The expected commissioning progression from the first Advanced LIGO observing run to design




Figure 2.11: The left side shows the predicted sensitivity ranges for the most likely observing scenarios for Advanced LIGO in blue, green, and red, with the anticipated dates and BNS ranges inset in the legend. For reference, the design sensitivity of the ‘nominal’ mode is shown with a black trace and the design sensitivity of the BNS-optimized mode, which would require re-configuration of the signal recycling cavity, is shown in pink. On the right, a very similar set of curves are shown for Virgo. Reproduced from [5].

sensitivity requires moving to high-power configurations, which will improve sensitivity at higher frequencies, and improving the seismic isolation performance, which will push the *seismic wall*, or the lowest viable frequency astrophysical gravitational wave searches can be conducted, to lower frequencies. The interferometer commissioning teams of LIGO and collaboration partner Virgo have the significant advantage of having commissioned four interferometers¹³ during the Initial and Enhanced versions of both projects, and expect a similar evolution in sensitivity.

The French-Italian 3 km Virgo detector near Pisa, Italy is currently undergoing hardware upgrades to Advanced Virgo similar to those designed for Advanced LIGO. Like the Advanced LIGO project, the upgraded hardware for Advanced Virgo is being in-

¹³Three initial LIGO interferometers and Virgo.

stalled within the same infrastructure¹⁴ as the first-generation Virgo experiment. The design for Advanced Virgo includes largely the same elements as Advanced LIGO, including two perpendicular arms with Fabry-Perot cavities to boost light circulation, an Nd:YAG solid state input laser, power and signal recycling cavities, input and output conditioning optics, thermal compensation systems, and advanced suspension and seismic isolation for optics. Installation is currently projected to be complete in late 2015, with the first one-hour stable lock occurring sometime in early 2016 [33]. The anticipated sensitivity ranges for likely Advanced Virgo observing scenarios are shown in the right side of Figure 2.11.

The underground cryogenic detector KAGRA, designed with 3km long arms, is currently being constructed in the tunnels of the Kamioka mine in Japan and is expected to come online sometime later in the decade [33]. Including KAGRA in the global network of gravitational wave interferometers in addition to the two LIGO detectors in the United States and Advanced Virgo is expected to give a sky localization accuracy of 10 degrees² for half of binary neutron star coalescence events. Including a LIGO India interferometer would further improve sky location accuracy to 5 degrees² for 30 percent of BNS events [33].

The British-German GEO600 detector in Germany, with 600m long arms, is a successful testing ground for technology developed for second and third generation¹⁵ ground-based GW interferometer upgrades like light squeezing [73]. GEO600 has been in *Astrowatch mode*, where the interferometer is left in a stable data collection configuration and only taken offline for maintenance sparingly, since the LIGO and Virgo detectors went offline for Advanced era upgrades. It is expected to continue in

¹⁴Here infrastructure means the site, building, and vacuum chambers.

¹⁵Second generation interferometers are ‘Advanced Era’ interferometers: Advanced LIGO, Advanced Virgo. Third generation interferometers refer to projects not yet funded for construction, such as the Einstein Telescope [78], currently in the conceptual design phase.

Astrowatch mode at least until the first Advanced LIGO observing run. GEO600 will likely have a much lower sensitivity than Advanced LIGO or Advanced Virgo below roughly 1000 Hz, and at least ten times lower at 100 Hz. However, it is still expected to be sensitive to nearby astrophysical events such as the core-collapse of Betelgeuse, a red supergiant star in the Milky Way¹⁶ [33].

To illustrate the benefits of multiple gravitational wave interferometers around the globe, Figure 2.12 shows a sky map of the network antenna pattern for the two LIGO interferometers and Virgo based on their position and orientation, assuming a flat frequency response for each interferometer. This three interferometer network will enable constraint of the likely sky location of an arbitrary gravitational wave source. The addition of detectors in strategic locations optimally far from current detectors, as is proposed for LIGO-India, would significantly improve the ability of transient gravitational wave parameter estimation algorithms to confidently resolve sky location. At design sensitivity, the joint LIGO and Virgo three detector network is expected to be able to resolve the sky position of a binary neutron star event with a network SNR of 12 to within 5 degrees² for 8 percent of events, and to within 20 degrees² for 28 percent of events [5]. The addition of a LIGO detector operating at design sensitivity in India is expected to dramatically improve sky localization performance to within 5 degrees² for 17 percent of events, and to within 20 degrees² for 48 percent of events [5].

A second generation interferometer science run with the two Advanced LIGO interferometers is expected as soon as fall 2015 with Advanced Virgo following shortly after in 2016. Negotiations are in progress for the hardware belonging to a third LIGO interferometer to be transferred to India with an anticipated data collection run at Advanced LIGO design sensitivity early next decade [5].

¹⁶Betelgeuse is roughly 130 pc, or about 425 light years, away from Earth




Figure 2.12: A sky map of the network antenna pattern for the two LIGO interferometers and Virgo based on their position and orientation. The antenna pattern was calculated for the root-mean-square average of F_+ and F_\times , as defined for a single interferometer in Equations 2.21 and 2.22. Generated with cWB2G macro software.

2.3.2 Common sources of noise

Even when considering a broad overview of potential noise couplings, it is important to emphasize that the Advanced LIGO interferometers, as laid out in Section 2.3, are extremely complex instruments. The detectors are expected to have complicated, often non-linear coupling mechanisms between the gravitational wave readout channel and environmental noise sources such as magnetic fields, ground motion, acoustic noise, or internal interferometric noise such as mechanical resonances, digital saturations in data processing or control loops, or light scattering due to vacuum contamination.

Generally, noise is classified into three major types: lines, stochastic, and non-Gaussian transient noise known as *glitches*.




Figure 2.13: A spectrum of calibrated strain, $h(t)$, for the LIGO Livingston detector. The green curve shows the average interferometer performance during a lock stretch on December 16, 2014, during Engineering Run 6 (ER6). For comparison, a spectrum characteristic of performance during the previous LIGO science run (S6) is shown in grey. Distinct noise lines can be seen for both data sets. Reproduced from the daily summary pages generated by D. Macleod [63].

Noise lines are peaks in the frequency spectrum, as seen in Figure 2.13. This figure compares strain spectra of the LIGO Livingston interferometer between a recent Advanced LIGO Engineering Run¹⁷ (in green) and the previous science run (in grey). The 60Hz line from the U.S. electric power lines is common to both, and each curve shows different mechanical resonances from suspensions, called violin modes, and their harmonics, as well as different dither frequencies used for cavity locking. Constant noise lines are problematic for searches for continuous gravitational waves, and are dealt with by notching these frequency bands out of the data. Noise lines that vary in amplitude or frequency over time can also adversely affect a broader subset of gravitational wave searches, including transient gravitational wave burst searches if the variance is on the scale of a few seconds or less.

¹⁷A series of Engineering Runs was conducted ahead of the first Advanced LIGO observation run to prepare data analysis tools and pipelines as well as provide periods of stable data taking for data quality characterization studies.

Correlated noise that is coherent among a network of interferometers clouds the sensitivity of the search for stochastic gravitational waves. Schumann magnetic field resonances [96] are one common cause. The stochastic gravitational wave searches can measure, subtract, and mitigate coherent noise effects to a significant degree [95].

Although past science runs have given us valuable insight and experience with complex noise couplings and their impact on the gravitational wave searches, the Advanced LIGO detectors are new instruments. At this point it is largely unknown how these interferometers will respond to familiar, expected problems like high wind at Hanford, microseismic ground motion at Livingston, or anthropogenic noise such as trains or traffic.

Seismic noise associated with ground motion, including trains and truck traffic, is addressed in detail along with its potential impact on transient gravitational wave searches using Advanced LIGO data in Chapter 5.

For a detailed description of data quality mitigation techniques of common noise sources see Chapter 5. For a summary of noise sources that limited the gravitational wave searches during the most recent LIGO science run, see [3].

CHAPTER 3

RECOVERY OF BURST TRANSIENT EVENTS

Arbitrary transient gravitational waves events that are distinct in time are known as *bursts*. The methods used to recover generic burst signals have a broad range of applications across gravitational wave data analysis.

In addition to the identification and parameter estimation of burst gravitational wave events, burst analysis is also an essential tool for evaluating transient noise in the detectors. This is generally accomplished through the use of event-trigger-generators (ETGs) which are single-channel burst recovery pipelines used to identify localized excess power in a signal. ETGs are used to characterize non-Gaussian behavior in a multitude of channels, from the gravitational wave strain channel $h(t)$, to auxiliary channels such as seismometers, microphones, or laser power incident on a photodiode used for beam steering.

Trigger events identified by ETGs are used to correlate bursts of transient noise in the gravitational wave strain channel with transients in auxiliary channels. ETG triggers are also used for extremely important detector characterization purposes, such as evaluating the effectiveness and safety¹ of data quality vetoes designed to remove known noise sources from the data.

¹Veto safety is a measure of the likelihood of the veto rejecting a true gravitational wave signal.

3.1 Event trigger generators

In addition to identifying excess power events localized in time and frequency, ETGs also estimate the characteristic time, frequency, and measure of loudness called signal-to-noise ratio (SNR) for each recovered event.

3.1.1 Parameterization technique and tiling

To identify arbitrary burst events, ETGs typically construct a series of two dimensional time-frequency maps by projecting the data onto a set of functions, or *basis*, that covers a targeted region of parameter space. The targeted parameter space is generally defined by the length of time and range of frequencies² to be analyzed, as well as the range of any other characteristic parameter of our chosen basis. In what follows, we will see this is the parameter Q for the ETG Omicron.

Based on the Omega pipeline [28], Omicron uses an identical sine-Gaussian basis to resolve arbitrary burst events in time and frequency.

In the following section, the parameterization technique of Omicron is described and generalized to all ETGs where possible, making use of the same notation used by Shourov Chatterji in the description of the Q-pipeline [28].

²The range of frequency analyzed is generally the searchable range defined by our detector response described in Chapter 2

Parameterizing arbitrary bursts of signal

The *strain amplitude* of a gravitational burst of arbitrary form $h(t)$ can be expressed as:

$$|h| = \sqrt{|h|^2} = \left(\int_{-\infty}^{+\infty} |h(t)|^2 dt \right)^{(1/2)} \quad (3.1)$$

where $h(t)^2$ is the signal energy, with units of power spectral density.

Following the notation of [28], the original waveform, $h(t)$, is equal to the a normalized time function, $\psi(t)$, multiplied by a strain amplitude factor $|h|$:

$$h(t) = |h|\psi(t) \quad (3.2)$$

and in the frequency domain:

$$\tilde{h}(f) = |h|\tilde{\psi}(f) \quad (3.3)$$

where $\psi(t)$ and $\tilde{\psi}(f)$ are normalized to 1:

$$\int_{-\infty}^{+\infty} |\psi(t)|^2 dt = \int_{-\infty}^{+\infty} |\tilde{\psi}(f)|^2 df = 1. \quad (3.4)$$

The characteristic central time, t_0 , is the expectation value:

$$t_0 = \int_{-\infty}^{+\infty} t |\psi(t)|^2 dt \quad (3.5)$$

and the characteristic central frequency, f_0 , is:

$$f_0 = \int_{-\infty}^{+\infty} f |\tilde{\psi}(f)|^2 df. \quad (3.6)$$

The characteristic duration and bandwidth of the event are calculated as the variance in time and frequency respectively of the normalized waveforms.

$$\sigma_t^2 = \int_{-\infty}^{+\infty} (t - t_0)^2 |\psi(t)|^2 dt \quad (3.7)$$

$$\sigma_f^2 = \int_{-\infty}^{+\infty} (f - f_0)^2 |\tilde{\psi}(f)|^2 df \quad (3.8)$$

Lastly, the parameter Q is defined as the ratio of the characteristic central frequency of the event to its bandwidth, or variance in frequency:

$$Q = \frac{f_0}{\sigma_f}. \quad (3.9)$$

This Q is same *quality factor* of mechanical, electrical, or optical systems that characterizes how damped an oscillator is via the ratio of bandwidth to frequency. Lower Q systems will have few oscillations in response to some impulse, whereas high Q systems will have many oscillations for the same impulse.

Sine-Gaussians exhibit the expected behavioral dependence on Q , as illustrated in Figure 3.1, which contains examples of sine-Gaussian waveforms³ of a set frequency and varying Q . For sine-Gaussians of the same central frequency, higher Q values

³The waveforms shown are the same as used as Omicron's basis, defined by Equations 3.11 and 3.12.

yield longer duration waveforms and more narrowband signals, whereas lower Q values yield the inverse.




Figure 3.1: The effect of varying Q on sine Gaussians with a frequency of 20 Hz. The top row shows the real part of each sine Gaussian waveform as defined in Equation 3.11; note the different time axis scales. The middle row shows the Gaussian envelope in time for the same waveforms on a uniform time axis scale, and the bottom row shows the corresponding Gaussian envelope in frequency on a uniform frequency axis scale. From low to high Q , the left column contains examples of $Q = 3$, the center of $Q = 18$, and the right of $Q = 50$. Sine-Gaussians with a higher Q tend to have a longer duration and be strongly peaked in frequency.

It is useful to use Q as a characteristic parameter for arbitrary burst events since their bandwidth and duration cannot be known individually - only to the uncertainty [41]:

$$\sigma_t \sigma_f \geq \frac{1}{4\pi}. \quad (3.10)$$

Then any burst event⁴ can be uniquely described with the parameters time, frequency, Q , and an additional parameter estimating the event loudness relative to the non-stationary background noise: signal-to-noise ratio, or SNR [28].

Why a sine-Gaussian basis?

When choosing a set of basis functions, the goal is to achieve the minimum possible time-frequency area $\sigma_t \sigma_f$, or the best possible resolution in time-frequency space. Omicron uses sinusoids modulated with Gaussian envelopes, or *sine-Gaussians*, which have been shown to have minimal resolution in time and frequency [41].

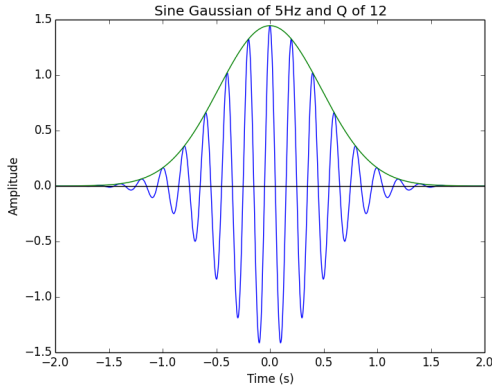
For basis functions, Omicron uses sine-Gaussians of the form [28, 80]:

$$Y(t, t_0, f_0, Q) = \left(\frac{8\pi f_0^2}{Q^2} \right)^{\frac{1}{4}} \exp \left[\frac{-4\pi^2 f_0^2}{Q^2} (t - t_0)^2 \right] \exp [-i2\pi f_0 (t - t_0)] \quad (3.11)$$


or equivalently in the frequency domain:

$$\tilde{Y}(f, t_0, f_0, Q) = \left(\frac{Q^2}{2\pi f_0^2} \right)^{\frac{1}{4}} \exp \left[\frac{-Q^2}{4f_0^2} (f - f_0)^2 \right] \exp [-i2\pi t_0 (f - f_0)]. \quad (3.12)$$

⁴Any burst event shorter in duration than our basis function cover (generally 1 second or less, with notable exceptions such as resolving low frequency seismic motion transients).



(a) A sine Gaussian of 5Hz and a Q of 12



(b) A sinusoid with a frequency of 5Hz.

Figure 3.2: A comparison of a single basis function from a sine-Gaussian basis and a single basis function from the conventional Fourier basis with the same frequency of 5Hz. The sine-Gaussian on the left is identical to the sinusoid on the right, but wrapped in a Gaussian envelope defined by the first exponential term in Equation 3.11 and also a different amplitude governed by the first fractional term of Equation 3.11.

where t and f represent varying time and frequency, t_0 is the central time of the sine-Gaussian, f_0 is the central frequency, and Q , described above, roughly translates to the number of oscillations in the sine-Gaussian. The resulting functions are Gaussian-windowed sinusoids, as shown in Figure 3.2a.

To emphasize the importance of the discreteness of sine-Gaussians in time, let's consider the classic Fourier transform as an alternative basis. This basis would decompose any time series signal into a set of composite sinusoids by projecting the data onto a basis of sinusoids, as shown in Figure 3.2b.

However, a signal that is well localized in time can be approximated using far fewer sine-Gaussian waveforms relative to infinite sinusoids, making sine-Gaussians a much more computationally efficient choice of basis function.

Other time-discrete functions are also known to be effective bases for recovering arbitrary short⁵ bursts of signal, such as Meyer or Morlet-Gabor wavelets [29], the Wilson-Daubechies-Meyer wavelet transform [72], or Hann windows that do not oscillate like wavelets, but effectively apply a bandpass filter⁶ in frequency [22]. However, sine-Gaussians are thought to provide the smallest uncertainty in time and frequency [28].

Placing time-frequency tiles using a sine-Gaussian basis

To optimize both detection efficiency and computational efficiency, a balance must be struck between using more or fewer basis functions. For any arbitrary burst event to be detected, which requires a close match with a basis function, a very large number of basis function should be used. However, this is computationally expensive. An optimal compromise is to define a threshold for the maximum acceptable signal energy lost due to a mismatch between an assumed sine-Gaussian signal and the basis functions used to tile the time-frequency parameter space [28].

For Omicron, this mismatched energy is calculated assuming the signal to be recovered is a sine-Gaussian that is offset from the best-match basis function in time, frequency and Q by δt , δf , and δQ .

The recovered amplitude of the offset sine-Gaussian by the closest basis function is calculated as the inner product between the two waveforms (both of the form given by Equation 3.11) [28, 80]:

⁵Short bursts targeted by event trigger generators and burst identification algorithms are often assumed to be less than 1 seconds long in time.

⁶A bandpass filter is a filter that rejects signal outside of a specified band, or range, of frequencies [47].

$$\alpha = \int_{-\infty}^{+\infty} Y(t, t_0, f_0, Q) Y'(t, t_0 + \delta_t, f_0 + \delta_f, Q + \delta_Q) dt. \quad (3.13)$$

The threshold of acceptable fractional energy loss that will define the minimal number and optimal distribution of sine-Gaussian basis functions is:

$$\mu = 1 - \alpha^2 \quad (3.14)$$

where α^2 is the recovered signal energy of the offset sine-Gaussian by its closest match basis function.

This threshold is used to generate a metric that governs a distribution of basis functions in Q , t_0 , and f_0 that optimally satisfies both the condition for acceptable signal energy loss due to mismatch and the minimization of the number of basis functions used. This optimal tiling of time-frequency parameter space is shown to lead to a basis distributed logarithmically in Q , logarithmically in central frequency f_0 , and linearly distributed in central time t_0 [28], as shown for multiple Q values in Figure 3.3.

This distribution of Q , f_0 and t_0 is a *multi-resolution* basis, which yields a series of time-frequency tiles of different aspect ratios, governed by Q , as demonstrated in Figure 3.4.

From Equations 3.9 and 3.10 we can write the corresponding uncertainty function for the bandwidth of a tile, Δf , and the duration of a tile, Δt , and use this to relate Q to Δt :

$$\begin{aligned} \Delta t \Delta f &= 1/4\pi \\ Q &= f_0/\Delta f = 4\pi f_0 \Delta t. \end{aligned} \quad (3.15)$$




Figure 3.3: An example of Omicron’s time-frequency tiling for multiple Q values. Each distinct position in time and frequency is encompassed by a tile. If the signal bounded by a tile contains excess power relative to the surrounding signal above some threshold, the bounds and center of the tile in time and frequency dictates the central time, central frequency, bandwidth, and duration of the detected event. The placement of tiles is governed by the distribution of sine-Gaussian basis functions in time and frequency. As expected for tiles of the same Q value, as frequency increases, the bandwidth, or frequency range, of the tile also increases. As Q increases, the same range of frequency covered by a tile set with lower Q will have lower bandwidth and higher duration tiles. Note the logarithmic distribution in Q is not captured with only two Q tile sets.




Figure 3.4: Relative aspect ratios of time-frequency tiles by Q value, as defined by Equation 3.15. Tiles with low associated Q tend to have a broad bandwidth and low duration, whereas tiles with high Q tend to have a short bandwidth and long duration, mirroring the behavior shown in Figure 3.1

This demonstrates an inverse relationship between Q and the bandwidth of a tile and a linear relationship between Q and tile duration, as shown in Figure 3.4. This is exactly as we would expect based on the behavior of our proposed Gaussian window, shown for a particular frequency in Figure 3.1. For tiles of the same Q and increasing central frequency, f_0 , we also expect duration to decrease and bandwidth to increase, as illustrated in Figure 3.3.

The benefits of using an overcomplete basis

The sine-Gaussian basis used by Omicron is an *overcomplete* basis, meaning the basis functions are not orthogonal and the number of basis functions may exceed the dimension of the signal space. The major drawback to an overcomplete basis is that signal waveforms may not be reconstructed accurately, as the signal will be oversampled in regions of the parameter space. Thus it is not possible to perform the inverse projection and recover the signal waveform. However, the goal of the Omicron algorithm is not to reconstruct signals, but to identify them and accurately characterize their bulk parameters (time, frequency, duration, bandwidth, and SNR). Including more basis functions to cover the same time-frequency space greatly improves the performance of the algorithm in efficiency and bulk parameter recovery [28].

3.1.2 Implementation for Omicron

Now that we have defined a structure for tiling data in time and frequency with the best possible resolution while also maximizing computational efficiency, we outline the implementation of this method via a series of digital filters and windows. In this section we first outline the general data transform technique, and then follow an input time series through all the major processing steps to the production of event triggers (or simply *triggers*).

The key step is the projection of the data onto the sine-Gaussian basis defined in Section 3.1.1. This projection, identical to that implemented for the Omega pipeline [28], tiles the data in time and frequency by:

1. Applying a series of bandpass filters centered on discretized f_0 values dictated by the optimal tiling of the analyzed parameter space in time and frequency and chosen range of Q values⁷.
2. Discretizing the data linearly in time.

As stated in [28], this is analogous to matching the data to a train of sine-Gaussians lined up end-to-end in time, as shown in Figure 3.5 for multiple frequencies governed by a single Q -plane.

To achieve this transform, Omicron follows Omega in the use of bisquare windows instead of true Gaussian windows, which as described in Section 3.1 are not realizable for a data projection that is also periodically discretized in time [28]. Bisquare windows are a close approximation of true Gaussian windows, achieving a minimum uncertainty in time and frequency that is only 4.5% greater. They also have little spectral leakage in the form of sidebands observed in the time domain after the window is applied to an impulse function, behaving very similarly to Hann windows [28].

This difference between bisquare and Gaussian windows introduces a small correction to our original description of Q in terms of central frequency and bandwidth in Equation 3.9:

$$Q = \frac{1}{\sqrt{11}} \frac{f_0}{\Delta f}. \quad (3.16)$$

⁷Recall the optimal trade-off between minimal mismatch and computational efficiency is to tile the data such that Q values are distributed logarithmically, and the f_0 values for each Q plane are also distributed logarithmically.




Figure 3.5: A series of sine-Gaussians of $Q=12$ using three example frequencies (2Hz, 5Hz, and 10Hz) normalized to have the same amplitude integrated over time. Omicron’s tiling of the data in time and frequency is equivalent to projecting the data onto a train of sine-Gaussians like those shown here, which would be akin to a single Q plane as shown in Figure 3.3. Note that the distance in time between sine-Gaussians isn’t necessarily as neat as shown depending on the maximum allowed mismatch between an assumed sine-Gaussian signal and a sine-Gaussian tiles. Also note the sine-Gaussians in the figure are nicely truncated so that all of the waveform falls between the time segment of interest, which also does not necessarily mirror the behavior of Omicron tiling. Any edge effects are mitigated by requiring significant overlap (on the order of a few seconds) between analyzed time segments.

The bisquare windows used by Omicron (and Omega) are of the following form:

$$w = \begin{cases} \left(\frac{315 Q}{128\sqrt{11} f_0} \right)^{1/2} \left(1 - \left(\frac{fQ}{\sqrt{11} f_0} \right)^2 \right)^2 & |f| < \frac{f_0\sqrt{11}}{Q} \\ 0 & \text{for all other } f \text{ values} \end{cases} \quad (3.17)$$

where the bisquare window has been normalized so that the integral of the window in frequency from 0 to ∞ is always 1. This normalization ensures the tile energy contained in windowed data is straightforward to extract as a coefficient of a sum over the tile bounds⁸ of the corresponding basis function multiplied by the data [28]. This will be very useful in estimating tile SNR and amplitude.

Ensuring that 100% of the bandwidth of any tile is always above 0 Hz and below the Nyquist frequency⁹ for N of samples in a discrete time series yields a minimum possible Q and maximum possible frequency¹⁰ [28]:

$$Q_{min} = \sqrt{11} \approx 3.3166 \quad (3.18)$$

$$f_{max} = \frac{N}{2} \left(1 + \frac{\sqrt{11}}{Q} \right)^{-1}. \quad (3.19)$$

⁸Here the bounds of a tile are minimum and maximum time and frequency.

⁹The Nyquist frequency is half of the sampling rate of a signal that is discretely sampled in time. Digital signals should not be manipulated above this frequency to avoid aliasing effects, which occur when a different function can produce the same set of samples.

¹⁰Note that frequency values are discretized by the applied discrete data projection, they are not continuous.

Figure 3.6 outlines the end-to-end process of the Omicron ETG.

1. Initialization

Just after starting an Omicron run, the time-frequency tiling is computed using a series of bisquare windowed sinusoids optimized over the user requested time segment(s), frequency range, Q range, sampling rate, and data segmentation length¹¹, for a specified amount of acceptable signal mismatch, α [80]. Elongating the data segmentation length is critical for resolving arbitrary burst events with long durations, as we will see for seismic transient studies details in Chapter 5.

2. Data conditioning

Next the data is whitened, which means the data is divided by the average power spectral density (PSD) such that the PSD of the resulting data is flat across all analyzed frequencies. To mitigate any contamination in the calculated average PSD due to loud transients, the *mean-median* method is used. This method breaks the data up into smaller segments that span the analyzed time, calculates the median PSD of each of these smaller segments, and then takes the mean of the median PSDs. This technique is effectively robust to large amplitude excess power events that are well localized in time [21]. For signals with an expected long duration, such as low

¹¹Data is segmented into both ‘chunks’, which govern PSD estimation, and ‘segments’, which govern data tiling and are not related to data quality or other requested input time segments. See Appendix D for a full explanation of the Omicron input parameters for the described version of Omicron.

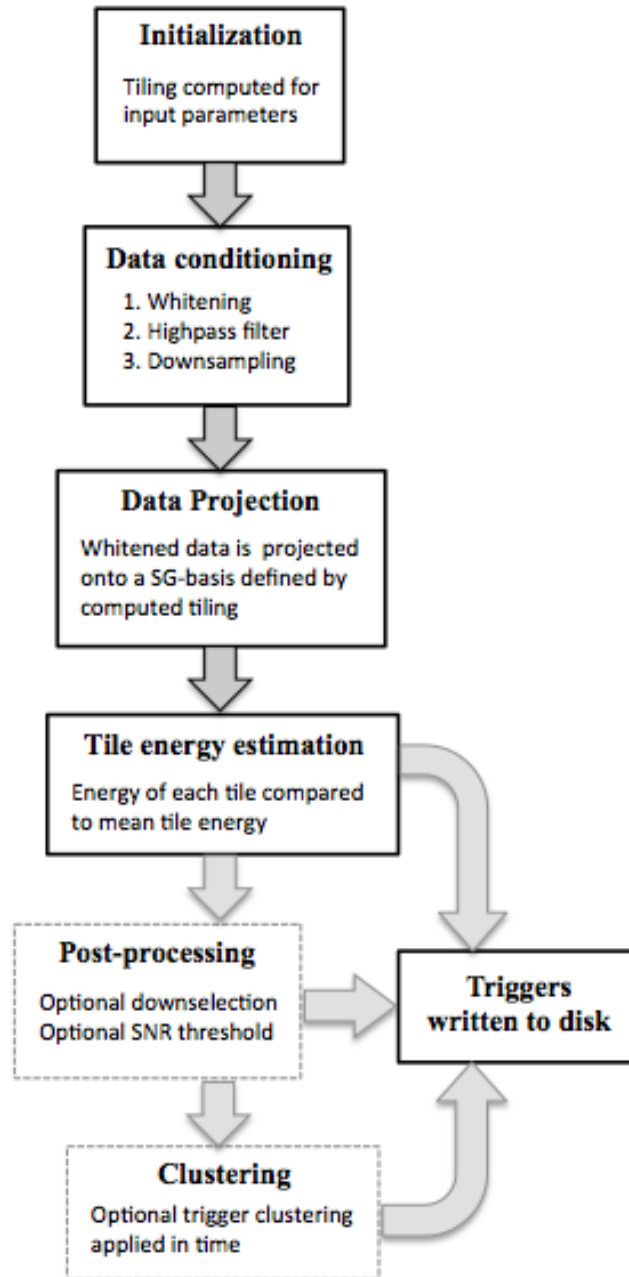


Figure 3.6: A flowchart of the implemented Omicron process, from the the submission of a run through the final step of writing produced triggers to disk. Each process is explained in detail in the text. Core processes that occur for every run are enclosed with a solid black line. Optional processes requested by the user are enclosed with a dashed line. The parameters needed to achieve different post-processing and clustering output are discussed in detail in Appendix D. Note that for the version of Omicron described in this work, all implemented steps are different from the Q-pipeline, described in [28].

frequency seismic events, the length of data used to estimate the average PSD must be longer than the expected transient event(s). This is key to the seismic study described in Chapter 5.

After whitening, a high pass filter is applied to the data to prevent phase distortion using a cutoff frequency equal to the lower limit of the requested frequency range [79].

Downsampling is also used to speed up processing and take advantage of performing shorter fast inverse Fourier transforms. The downsampled frequency is specified by the user and must be compatible with the requested frequency range [80, 28].

3. Data projection onto sine-Gaussian basis

The conditioned data is then projected onto a close approximation to the sine-Gaussian basis detailed in Section 3.1.1. This is achieved by tiling the data in time and frequency with multiple resolutions using the bisquare windowed sinusoids described earlier in Section 3.1.2. As described in Section 3.1.1, multiple resolutions, or different time-frequency planes each with some particular Q value, can be thought of as tiles covering the same time frequency space with different aspect ratios.

4. Tile energy computation

To identify events, Omicron takes all tiles with positive excess energy Z for each Q plane. For each tile, $\text{tile}[i]$, Z is computed as:

$$Z = \frac{|X_i|^2}{\langle |X|^2 \rangle} - 1 \quad (3.20)$$

where $\langle X \rangle$ is the estimated mean of normalized tile energy for all considered tiles in the Q plane containing tile[i] and X_i is the normalized energy of tile[i]. Subtracting 1 makes it so that tiles with energy exactly equal to the mean tile energy have a excess energy Z of zero, and all positive Z values indicate a tile energy greater than the mean. For an arbitrary burst event in the gravitational wave strain channel, the quantities $\langle X \rangle$ and X_i both have units of $|h|^2/\sqrt{S(f)}$ where $S(f)$ is the power spectral density and $|h|$ is the strain amplitude.

Tile energy X^{12} is computed using the coefficients of the sine-Gaussian basis data transform that were normalized for this purpose in Section 3.1.1.

Following the same notation as [28], X is the magnitude of the real and imaginary transform coefficients resulting from the projection of a sampled time series $x[n]$ onto a series of windowed sinusoids:

$$X = \sqrt{X_{real}^2 + X_{im}^2} \quad (3.21)$$

$$X_{real} = \sum x[n] w[n - m, l, Q] \cos\left[\frac{2\pi nl}{N}\right] \quad (3.22)$$

where X is summed over the total number of data samples N . n is the discrete sampled data index (in time), ranging from 0 (the beginning of the time series) to N . m and l index the central time and frequency, respectively, of each included sine-Gaussian in the projection. Quality factor Q governs the included values of time and frequency indices m and l for that Q plane. The w term is the bisquare window term

¹²Tile energy is analogous to signal energy noted in Equation 3.1 but summed over the tile instead of integrated over the entire analyzed parameter space.

defined in Equation 3.17 and the cosine term is the real part of the complex sinusoid of frequency index l .

And similarly for the imaginary term of X ,

$$X_{imag} = \sum x[n] w[n - m, l, Q] \sin \left[\frac{2\pi nl}{N} \right] \quad (3.23)$$

where the window term, w , is identical, and the imaginary term of the complex sinusoid is sine of the frequency index l . X_{imag} is also summed over the total number of data samples, N .

The mean tile energy is estimated by computing the energy, X , of all analyzed tiles in the relevant Q plane. For this version of Omicron, outliers in this distribution are excluded above a normalized tile energy of 2σ and assuming tile energy below this threshold follows a Gaussian distribution, as depicted in Figure 3.7. The mean of the resulting distribution with high tile energy outliers removed is taken as the mean tile energy for the analyzed parameter space. This **outlier rejection** is key to a close estimation of mean tile energy due to nominal background noise, and eliminates any skew that loud excess power events might cause. In contrast to Omega, the mean tile energy is computed on-the-fly to avoid the computationally expensive sorting of tiles by energy [80].

In excluding all tiles above 2σ some noise events not due to very loud outliers are also excluded. This introduces a small bias in overestimating the normalized tile energy. We will see in section 3.1.3, which states that SNR corresponds to the square root of Z , this causes Omicron to overestimate SNR by no more than 10% [80].

5. Postprocessing




Figure 3.7: An example of the rejection of high-energy tiles - a normal distribution is first fit to the energy of all tiles in a particular Q plane. Tiles with energy of the calculated mean $+ 2\sigma$ or above are excluded, shown in red. The tiles below this cut are assumed to be Gaussian, and a Gaussian distribution is re-fit to the remaining tiles. The mean of this new distribution is taken as the mean tile energy of an analyzed data segment, and used to calculate the normalized tile energy defined in Equation 3.20. Note that this method was used for the version of Omicron used to produce the data in Chapter 5, and has subsequently been improved.

After tiles with positive normalized energy are identified, these triggers¹³ can be optionally further processed with downselection and/or an applied SNR threshold.

Optional downselection enforces that whenever multiple tiles overlap significantly in time and frequency, only the loudest tile, or the tile with the highest normalized energy, is kept as a trigger and written to disk.

An optional SNR threshold requires that only tiles above a specified SNR are kept as triggers and written to disk, where SNR translates to the square root of normalized energy, as explained in Section 3.1.3. This is equivalent to setting a threshold on

¹³Note after this point the terms ‘trigger’ and ‘tile’ are used interchangeably, and are essentially identical objects during postprocessing and beyond.

normalized energy and could be readily extended to setting a threshold on trigger significance, or likelihood of observing a tile with a certain energy given the distribution of tile energies¹⁴.

6. Clustering

The end product of unclustered triggers produced by Omicron is nominally many different triggers produced for the same burst power event. These triggers are produced with different time frequency resolutions and have some offset in central time, central frequency, and/or Q value with respect to one another. Given that there is generally significant overlap in trigger duration and bandwidth for triggers associated with the same event, it is often most useful to *cluster* these triggers together to form a clustered event with time and frequency edges defined by the most extreme contributing tiles, as shown in Figure 3.8.

Due to the discretized nature of the sine-Gaussian basis, Omicron tends to resolve events into discrete ‘islands’ of power, especially if they do not have strong oscillating features as in the case of white noise bursts¹⁵, as shown in Figure 3.9. For this reason, Omicron clusters triggers using a default time coincidence window of 0.1 seconds, which is wide enough to collect power associated with the same event and narrow enough to exclude power from non-associated events.

The reported elements of clustered Omicron triggers, illustrated in Figure 3.8, are as follows:

¹⁴Trigger significance is the quantity that is thresholded for the ETG ExcessPower.

¹⁵A white noise burst is a discrete burst of white noise, which is random signal with a flat frequency response. These are discussed in greater detail in Section 3.2.2.




Figure 3.8: An example demonstrating Omicron’s clustering technique. Each black rectangle represents an unclustered tile (trigger) produced by Omicron, with a clustered trigger outlined with a red dashed line. Tiles forming a cluster may be separated in frequency by any frequency difference, and separated in time by up to the specified cluster window, t_c . This separation in time is defined from the end time of a tile to the start time of another, which is shown explicitly for the tiles labeled 1 and 2. A clustered trigger spans the duration and frequency range (bandwidth) of the most extreme edges of the tiles within the cluster. The central time, frequency, and SNR of the most significant (highest SNR) tile are also reported. Note this is a contrived example to illustrate the clustering process - generally there is much more overlap in time and frequency between tiles contributing to a clustered trigger.




Figure 3.9: Omegascans, or spectrograms generated with the same sine-Gaussian basis used by Omicron, of two white noise burst events. On the left is a white noise burst with an injected frequency of 2564.6 Hz, a bandwidth of 225 Hz, a duration of 0.10 s, and an injected SNR of 32. On the right is a white noise burst with an injected frequency of 2279.5 Hz, a bandwidth of 712 Hz, a duration of 0.12 s, and an injected SNR of 36. Note that for both events Omega resolves pockets of signal energy discretized in time, as may be expected from the discretized nature of a sine-Gaussian basis. This effect is more pronounced in the case of white-noise-burst-like events with broader signal bandwidth, as seen in the plot on the right. Appropriate clustering is important for accurately resolving noise-burst-like events with sine-Gaussian basis or wavelet basis event identification algorithms, as addressed in Section 3.2.4.

- **Start time and end time of the clustered trigger** defined by the most extreme edges of tiles composing the clustered trigger.
- **Lower and upper frequency bounds of the clustered trigger**¹⁶ defined by the most extreme edges of the tiles contributing to the clustered trigger.
- **Peak time** - the central time of the most significant tile in the cluster, where most significant translates to highest SNR.
- **Peak frequency** - the central frequency of the most significant tile in the cluster.
- **SNR** - the estimated SNR of the most significant tile in the cluster.
- **Also recorded** are the central frequency of the clustered trigger, as well as the time and frequency bounds and amplitude of the most significant, or highest SNR, trigger in the cluster.

Since Omicron uses an overcomplete basis, summing the SNR contributions from each tile would overestimate the clustered event SNR. Currently, the SNR of the most significant tile in the cluster is reported as the clustered trigger SNR. This is a good estimate for well localized events where a single sine-Gaussian waveform is a good match for the event and can collect most of the signal energy. However, for white-noise-burst-like events where a single sine-Gaussian is not a good match for the whole event, as seen in Figure 3.9, this is not a good estimate. Recommendations to address this in the future are made in Section 3.2.4.

¹⁶The upper and lower frequencies of a clustered trigger are reported by Omicron as *flow* and *fhigh*.

3.1.3 Estimation of trigger SNR and amplitude

Two quantities of great interest in interpreting the trigger output of ETGs are the reported trigger SNR and amplitude, each a different measure of a trigger’s loudness.

SNR is preferred as a measure of event loudness in astrophysical gravitational wave searches as it includes by definition some description of how much louder that tile, or cluster of tiles, is compared to surrounding tiles. This gives some indication of the significance of the event in terms of likelihood of observing it in the surrounding noise.

Amplitude is sometimes preferred by those using ETGs to study the non-stationarity of instrumental signals, such as seismometers. Trigger amplitude gives a better intuition for the absolute energy of the event, without consideration of the surrounding noise at the time. This makes it easier to compare quantities like motion amplitude between periods where the noise may be intrinsically different.

The ‘true’ signal amplitude of a burst event in the frequency domain, making use Parseval’s theorem and Equation 3.1, is given by:

$$|h|^2 = \int_{-\infty}^{+\infty} |\tilde{h}(f)|^2 df. \quad (3.24)$$

The ‘true’ SNR, ρ , of a burst event is proportional to the signal amplitude, $|h|$ ¹⁷.

$$\rho^2 = \frac{|h|^2}{\int_0^{+\infty} S(f) df} \quad (3.25)$$

¹⁷Note the quantity $|h|$ describes the amplitude of any arbitrary burst event, not necessarily a gravitational wave signal.

where $S(f)$ is the ‘true’ power spectral density of the noise as a function of frequency.

We can further simplify this by making the strain amplitude integral also one-sided, which introduces a factor of two in the numerator:

$$\rho = 2 \left[\int_0^{+\infty} \frac{\tilde{h}^2(f)}{S(f)} df \right]^{1/2} \quad (3.26)$$

The differences between SNR and amplitude for an example burst event are illustrated in Figure 3.10, which highlights the different integrated frequency-amplitude areas between the two quantities.

In the limiting case of a very narrowband burst event centered at some frequency f_0 , the signal-to-noise ratio is the ratio of the amplitude of the signal $|h(f_0)|$, to the amplitude of the noise, $[S(f_0)]^{1/2}$. The trigger SNR quantity sums this ratio over the frequency range of the tile, and the total event SNR, or true SNR, integrates this ratio over all frequencies.

To estimate both SNR and amplitude of a particular tile, Omicron and other ETGs compare the energy of each tile to the tile energy mean.

In the case of Omicron, Equation 3.20 states that the excess energy, Z , of a tile contained in a particular Q plane is the tile energy¹⁸ divided by the calculated mean tile energy of that Q plane, minus 1 such that Z is 0 when the the tile energy is equal to the mean tile energy, and positive whenever the tile energy is greater than the mean.

The SNR, ρ , of a particular tile is estimated as:

¹⁸ Tile energy is defined in Equations 3.21 through 3.23.

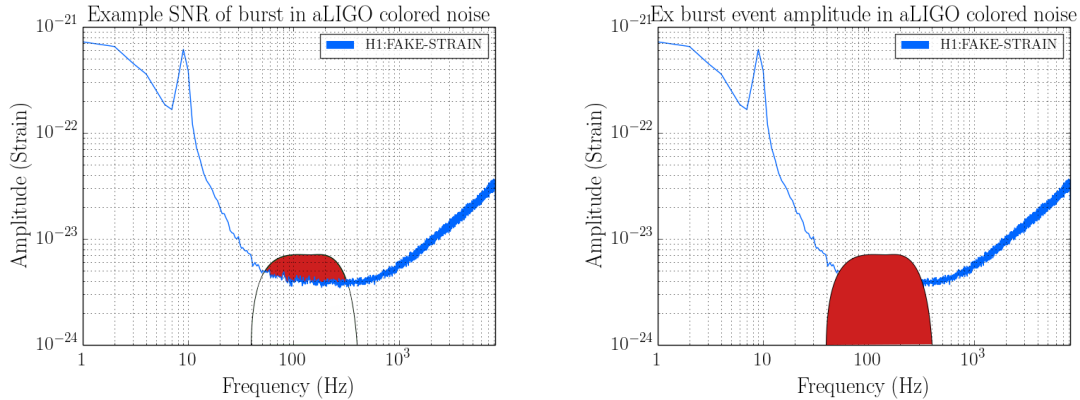


Figure 3.10: On the left is an illustration of the parameter space area used to compute the SNR of a burst event. The amplitude spectral density (ASD) of an arbitrary burst event is shown in a black trace over the amplitude spectral density of aLIGO-design colored Gaussian noise in a blue trace. The event SNR is defined as the amplitude of the event (in black) at some particular frequency f_0 , divided by the amplitude of the noise (in blue) at some particular frequency f_0 , integrated (or summed) over all frequencies, as defined in Equation 3.26. Note that one particular Omicron tile may not capture the entire power of such an event - the trigger SNR in this case is only summed over the frequency bounds of that tile and will be some fraction of the total event SNR.

On the right is an illustration of the parameter spaced used to compute the amplitude of a burst event. The ASD of the same burst event is shown again with a black trace. The ASD of the same background noise is shown for reference, but is irrelevant to the event amplitude. The total event amplitude is defined as the amplitude of the event at some particular frequency f_0 , integrated (or summed) over all frequencies, as defined in Equation 3.24. The total event amplitude is related to the total event SNR by Equation 3.26, and Omicron estimates trigger amplitude using the estimated tile SNR and the calculated average PSD.

Note that SNR (left) is unitless (as the ratio of signal to noise) whereas amplitude has units of the signal.

$$\rho = \sqrt{Z}. \quad (3.27)$$

Intuitively, this follows our definition of ‘true’ SNR as the ratio of signal amplitude to noise amplitude integrated over all frequencies. The calculated tile energy is of whitened data that is normalized by the average PSD curve, and this is summed over all frequencies bounded by the tile.

To estimate tile amplitude, Omicron follows Equation 3.25, making use of the estimated tile SNR, ρ_{tile} , and the estimated average PSD curve for the data:

$$|h|_{tile} = \rho_{tile} \left[\sum_{f_{min}}^{f_{max}} P(f) \right]^{1/2} \quad (3.28)$$

where P is a weighted average of the noise over the frequency bounds of the tile, calculated with the estimated average PSD, $S(f)$, as:

$$P = \frac{\sum w(f) * w(f) S(f)}{\sum w(f) * w(f)} \quad (3.29)$$

where $w(f)$ is the familiar bisquare window defined in Equation 3.17 [80]. Both $|h|_{tile}$ and P are summed over the frequency bounds of the tile.

It is important to note that when interpreting trigger amplitude in physical units, for example a seismometer signal trigger reported in units of nm/s, that the **trigger amplitude is not a good approximation of the maximum amplitude of the motion or other physical signal unless the event is narrowband**. This is illustrated in Figure 3.11. For example, it is not generally correct to interpret a trigger with a reported amplitude of 2000 nm/s produced by a seismometer signal sensing in a particular direction as ‘the platform/surface the sensor is mounted on moved with a velocity of 2000 nm/s in this direction at that time’. Remember that

trigger amplitude is summed over all contributing frequencies of a tile, as is trigger SNR.

3.2 Performance of Event Trigger Generators

Event trigger generators are critical tools for identifying burst events both as part of a gravitational wave detection scheme and also for characterizing transient noise in the detector. ETG triggers are generally cross-correlated in time, frequency, and SNR either between interferometers or between auxiliary channels that measure the behavior of the detector and local environment. ETG triggers are also used for statistical studies of the transient quality of detector channels, or for testing the effectiveness, efficiency, and safety of data quality vetoes¹⁹, as discussed in Chapter 5.

In order for event trigger generators to fulfill these roles, it is critical that they be both highly efficient in identifying excess power events and very accurate in characterizing the bulk properties of these events; especially time, frequency, and SNR.

In the following study, the performance of multiple ETGs was assessed by measuring their efficiency and accuracy in parameter recovery. Participating ETGs produced triggers for a series of burst events injected into simulated Gaussian noise. The output triggers for each ETG were matched to the true injections to test efficiency, and for detected events, the difference between recovered parameters and injected parameters.

¹⁹Data quality vetoes indicate periods of time that the gravitational wave strain channel is known or indicated to be contaminated with noise. Generally veto criteria are determined through the analysis of one or more auxiliary channels. Assuring the *safety* of these channels means verifying they would not respond to a passing GW and accidentally veto a true signal.




Figure 3.11: A figure showing the amplitude spectral densities of several example burst events laid over the amplitude spectral density of aLIGO colored Gaussian noise in blue. The example events are drawn cartoonishly boxy here for easy reference relative to the grid lines.

Note that relative to the aLIGO noise curve, the event with lowest central frequency (in red) has the same signal-to-noise ratio as the event with the middle central frequency (in orange). The event with lowest central frequency (in red) also has the same amplitude as the event with highest central frequency (in purple). This illustrates that the quantities SNR and amplitude as defined in terms of burst analysis depend very strongly on the bandwidth, or equivalently, duration, of the event. Event or trigger amplitude is generally not an accurate gauge of the maximum time series amplitude of a signal with the exception of very narrowband events.

The novelty of this study is not only the direct comparison of multiple burst detection and parameter estimation algorithms, but also in the inclusion of non-oscillating white noise burst²⁰ injections in the mock data challenge set. Previously, internal ETG performance checks mainly focused on recovering well-localized sine-Gaussian waveforms, which do not necessarily reflect the true behavior of gravitational wave or noise transients.

3.2.1 Considered algorithms

ETGs intended for arbitrary burst recovery²¹ all use the same basic burst recovery principles laid out in Figure 3.6. Multiple ETGs have been developed with different goals in mind. Some ETGs employ different basis functions to achieve less discretized recovery of events or improved computational efficiency. Some employ different code bases for improved processing time, and some differ in data processing techniques intended to improve aspects such as PSD estimation or the consolidation of trigger output for faster post processing.

Each algorithm considered in this study is summarized below:

Omicron - a sine-Gaussian basis ETG intended for signal burst detection and bulk parameter recovery, described in great detail in Section 3.1. Triggers for this study were produced with the configuration in Appendix D.

ExcessPower - an orthogonal Hann-window basis ETG that effectively uses a series of bandpass filters to isolate different bands of frequency. Similarly to Omicron and Omega, ExcessPower also tiles data into time-frequency tiles of varying aspect ratios.

²⁰A white noise burst is a discrete burst of random signal with a flat (white) frequency response.

²¹For the purposes of this study, ETGs intended for arbitrary burst recovery are Omicron, ExcessPower, and DMT-Omega. BayesWave and PCAT have different applications, and used this study to test the recovery of parameters for identified events rather than efficiency.

Significant events are identified by thresholding the probability of obtaining a given tile power from Gaussian noise [22].

DMT Omega - a sine-Gaussian basis ETG very nearly identical to Omicron (and Omega) in design, and with some differences in implementation such as data conditioning and post-processing reporting of triggers. DMT-Omega was developed in parallel with Omicron to be integrated into the low-latency Data Monitoring Tool code infrastructure and used for low latency glitch identification and characterization. Note the input parameters, including range of Q , were not identical between Omicron and DMT-Omega for this study; each ETG used the preferred ‘standard’ parameters of the developers. Therefore, some minor differences in performance are expected.

BayesWave* - a Morlet-Gabor wavelet basis burst parameter estimation algorithm described in detail in Chapter 4 [29]. BayesWave was not designed to serve as a conventional ETG for independent identification of burst events²². However, BayesWave participated in this study to help develop the algorithm’s accurate and precise recovery of bulk parameters.

PCAT* - Principal Component Analysis for Transients - a transient noise, or *glitch*, classification algorithm for loud noise transients. To identify glitches, PCAT first whitens the input data time series, then applies an amplitude threshold on the signal based on the standard deviation of the whitened time series. PCAT was not developed to serve as a conventional ETG [77]. A reduced efficiency compared to ETG algorithms such as Omicron, ExcessPower and DMT-Omega, is to be expected, especially at lower SNRs.

²²BayesWave is intended to run as a follow-up to candidate events identified by the all-sky coherent burst search algorithm cWB2G.




Figure 3.12: The distribution of injection waveforms in SNR, with white noise burst (WNB) event on the left and sine-Gaussian (SG) events on the right.

From this point on the term ‘ETG’ is at times used to collectively refer to all of these algorithms out of convenience. However, note that only a subset of these, Omicron, Excess Power, and DMT-Omega, are true ETGs that are designed to have high a detection efficiency. The others, BayesWave and PCAT, are indicated in the list above with an asterisk.

3.2.2 Study technique

The injection set

The burst injection set included both sine-Gaussian and band-limited²³ white noise burst waveforms to test ETG performance in resolving both types.

Both sine-Gaussian and white noise burst waveforms were generated to have a flat distribution in SNR from 3 - 50 as shown in Figure 3.12. This SNR range covers very quiet events, below the SNR threshold of the study, through extremely loud events beyond the astrophysical feasibility of a burst signal in a single detector.

²³White noise bursts with some definite frequency range. For this study, the frequency range of BLWNBs was enforced to be less than half of the frequency range analyzed by the ETGs.

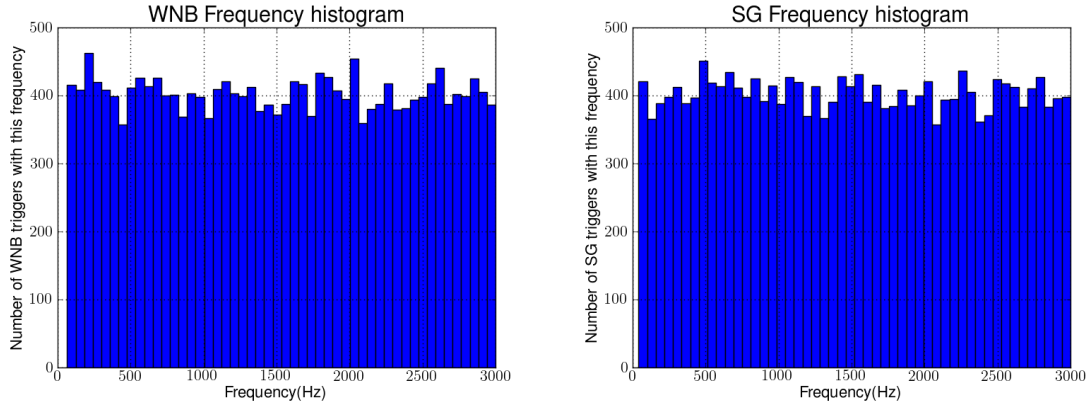


Figure 3.13: The distribution of injection waveforms in frequency, with white noise burst (WNB) event on the left and sine-Gaussian (SG) events on the right.

The frequencies of both sine-Gaussian and white noise burst injections were also distributed evenly in a frequency range of 30-3000 Hz, illustrated in Figure 3.13. This covers the majority of the expected observable frequency range for Advanced LIGO burst sources.

The sine-Gaussian waveforms were also distributed uniformly in Q , with a requested Q range of 3-9, where Q governs the bandwidth and duration of these signals as explained in Section 3.1.1. The white noise burst waveforms were distributed linearly in bandwidth between 50 and 1000 Hz, and bi-modally in duration between 0.01 and 0.32 seconds, as shown in Figure 3.14. An emphasis was placed on shorter duration signals because transient noise is expected to be more commonly shorter in duration.

These waveforms were injected into generated Gaussian noise colored to have the same spectral density as the nominal Advanced LIGO design curve shown in Figure 2.2. Injections were spaced 10-20 seconds apart so as to not have any overlap between events.

Each ETG produced unclustered triggers with ‘vanilla’ parameters as it would on an Advanced LIGO gravitational wave strain channel and:

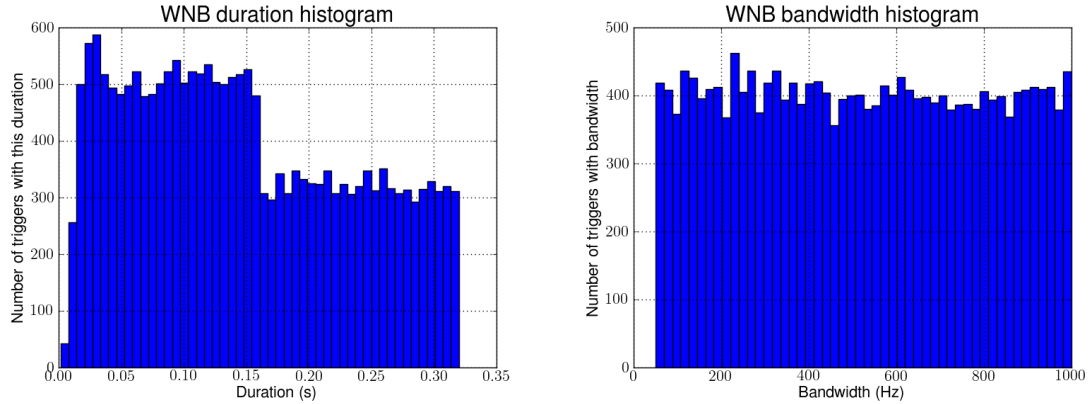


Figure 3.14: On the left, the mock data challenge distribution of white noise burst (WNB) waveforms in duration, and on the right, the distribution of WNB waveforms in bandwidth.

- over the same segment of time (the first 100,000 seconds of the data set)
- over the same frequency range (30 - 3000 Hz)²⁴
- with the same SNR threshold (5.5)

Similarly to BayesWave’s exception to the frequency range constraint, ExcessPower’s SNR estimation was too wildly far off to use a hard SNR threshold constraint and get an accurate measure of efficiency. Therefore, the total ExcessPower efficiency may be overestimated and a few triggers matched to injections that are not correlated with injected events. However, the ROC curves, explained under Figures of Merit, should be accurate. PCAT much more precisely underestimated SNR to a large degree, but as a correction factor has not yet been estimated PCAT was also exempt from the SNR constraint with similar consequences.

²⁴BayesWave was still in a tuning stage at this point and only ran up to 2000 Hz.

Clustering triggers

After each ETG had produced unclustered triggers using the simulated data and standardized study parameters, each set of triggers was clustered in time and frequency²⁵. The criteria for clustering triggers together required overlap between the triggers in time and frequency. Note this is *significantly* different than the version of clustering employed by Omicron, explained in Section 3.1.2, which does not require overlap in either time or frequency. This clustering algorithm was used because it employed a popular clustering technique and was straightforward to apply to all triggers sets in a standardized way.

For each ETG, the resulting clustered trigger set was the final form to be tested.

Matching triggers to injections

Each injection was matched in time to the corresponding clustered trigger in the analyzed trigger set, if any. The matching code first requires overlap in time between the duration of an injection and the duration of a clustered trigger. If more than one clustered trigger matches to an injection, the clustered trigger with the highest SNR is used for parameter recovery analysis.

Frequency was not considered when matching triggers to injections for this study due to the ineffective reporting of frequency for clustered events, as described in Section 3.2.4. For future work examining ETG performance in non-Gaussian noise, also matching the frequency of clustered triggers to the frequency of injections would be an effective way to reduce the number of false matches from transient noise not

²⁵Using the algorithm `ligolw bucluster`.

associated with injected events. However, the reporting of frequency for clustered events would first need to be improved, as discussed in Section 3.2.4.

Performance figures of merit

Two different metrics were used to gauge the performance of each ETG in efficiency. The first: Receiver-Operating-Characteristic (ROC) curves showing the efficiency versus false alarm rate (FAR) of each trigger set. The efficiency was calculated as the number of recovered events divided by the total number of injections. The FAR was calculated as the expected rate of false event detection by chance due to the rate of all triggers produced in that time. The second metric: the overall efficiency for injections above an SNR of 5.5.

The standard deviation in the difference between injected and detected event parameters was used to gauge the performance of bulk parameter recovery. Percent difference in frequency vs. frequency plots were also helpful in identifying tiling artifacts in frequency recovery.

The recovered event time was defined to be the central time (start time of tile + duration/2) of the highest SNR trigger in the clustered trigger. The recovered frequency was defined as the peak frequency of the clustered trigger, which is the central frequency of the highest SNR trigger in the clustered trigger. The recovered SNR was defined as the estimated SNR of the highest SNR trigger in the clustered trigger.

3.2.3 Results

Results are presented in two parts: efficiency and parameter recovery, using the figures of merit defined above.

Note that it is reasonable to expect sine-Gaussian or wavelet basis ETGs to have a higher efficiency and better parameter estimation for sine-Gaussian injections relative to other ETGs.

Efficiency

In the ideal performance case we expect an ETG to have a high efficiency and low false alarm rate. In terms of an ROC curve, we expect an ETG to lie far above the efficiency = false alarm rate line, which is shown in yellow in Figure 3.15.

At high false alarm rate nearly all ETGs approach 100% efficiency, as seen in Figure 3.15. PCAT is the expected exception, with each identified event depending on the time series amplitude statistics of only a short period of surrounding data (8 seconds). For this figure, false alarm rate was estimated by time-shifting triggers relative to the injections and checking for accidental coincidence between offset triggers and injections. Increasingly wide time windows were applied around the injection times to generate each point of the curves in Figure 3.15.

A summary of average efficiency of all ETGs for sine-Gaussian injections, white noise burst injections, and all injections is shown in Table 3.1. All ETGs have lower efficiency for white noise burst injections than for sine-Gaussians.

Note that ExcessPower and PCAT gained an advantage over other ETGs in the calculation of overall efficiency in Table 3.1 because imposing an SNR threshold of

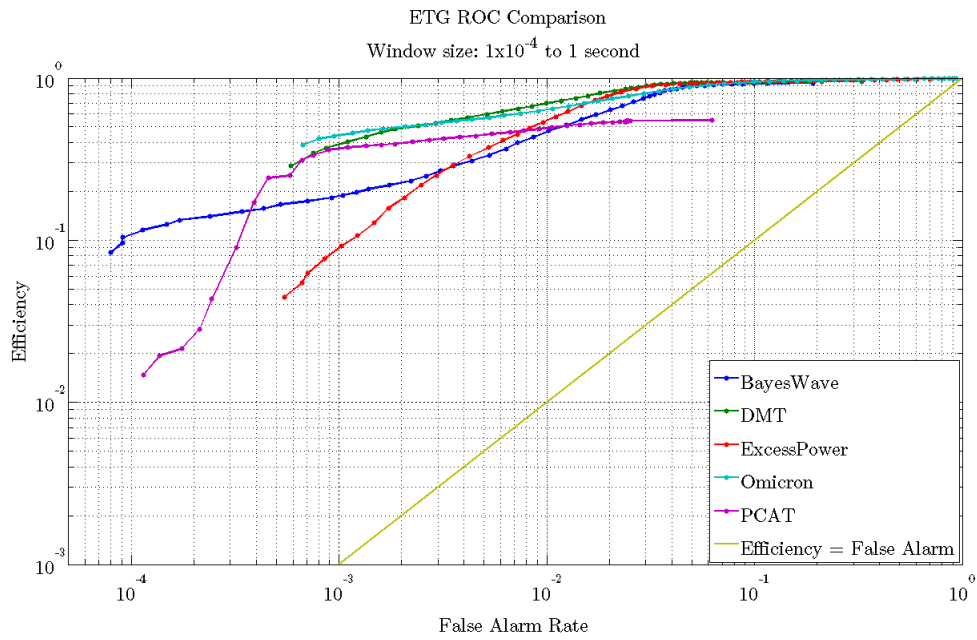


Figure 3.15: A Receiver-Operating-Characteristic (ROC) curve characterizing the efficiency vs. false alarm rate of ETGs included in the performance study. Each point on an ETG curve represents a time window of a different duration around injection times and the false alarm rate is calculated using the number of offset triggers captured by chance within this time. Plot produced by Marissa Walker.

ETG efficiencies for SNR >5.5			
ETG	All MDCs	SGs	WNBs
Omicron	97.5%	99.8%	95.2%
ExcessPower	97.0%	98.5%	95.4%
DMT Omega	94.0%	99.4%	88.6%
BayesWave	90.2%	94.8%	85.7%
PCAT	49.1%	60.7%	37.5%

Table 3.1: A table showing the average efficiency for each ETG for all analyzed injections with an SNR of 5.5 or greater. On the left the average efficiency for both sine-Gaussian and band-limited white noise bursts is listed. In the center are results for only sine-Gaussians, and on the right are results for only white noise bursts. All ETGs show a better efficiency for sine-Gaussian waveforms, including ExcessPower and PCAT. Omicron has the best overall efficiency. Here MDC is the acronym Mock Data Challenge, used as shorthand to mean an injection.

5.5 was not feasible²⁶. However, both are effectively penalized for this in the ROC curve as the relatively higher trigger rate results in a higher false alarm rate.

The efficiency of each ETG was also evaluated using trends in injected frequency or SNR via a series of scatter plots of missed and found injections, as seen in Figure 3.16. In these plots each red x represents a missed injection and each blue dot represents a found injection. There is no obvious correlation between efficiency and frequency for any ETG analyzing white noise burst injections.

The efficiency of the sine-Gaussian basis ETGs is shown in Figure 3.16. Unsurprisingly, Omicron and DMT-Omega had the best efficiencies for sine-Gaussian injections, as listed in Table 3.1. Omicron approaches very near 100%, missing only 5 low SNR injections of 3130 total. DMT-Omega’s efficiency is also above 99%, with three times more low SNR injections missed than Omicron, as seen in the Figure 3.16 scatter plot.

²⁶Because estimation of trigger SNR was too inaccurate.

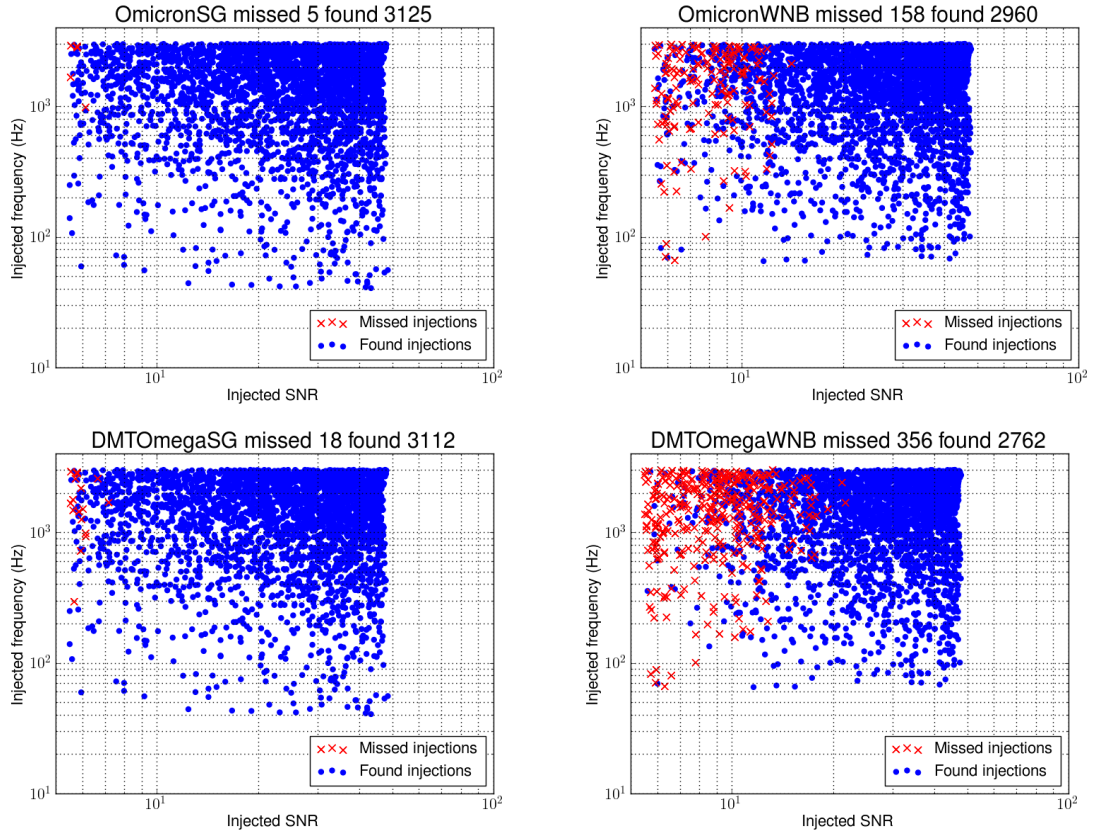


Figure 3.16: A series of scatter plots showing efficiencies for the sine-Gaussian basis ETGs Omicron and DMT-Omega for both types of injections. Each injected event is plotted by injected SNR (loudness) vs. injected frequency. A blue dot indicates a found event and a red x indicates a missed event. See Table 3.1 for the overall efficiency of each ETG.

Omicron also has a nearly equivalent efficiency to the best performing ETG on white noise burst injections, ExcessPower. Both Omicron and DMT-Omega miss higher SNR white noise burst injections than observed for sine-Gaussians, but much more so for the latter ETG. DMT-Omega misses events at significantly higher SNRs, with a few missed events beyond an SNR of 20. Why DMT-Omega, which matches Omicron so closely in sine-Gaussian efficiency, falls so far behind in white noise burst efficiency is not yet understood. Future work should ensure that the input parameters used between the two sine-Gaussian ETGs are as similar as possible.




Figure 3.17: A series of scatter plots showing ExcessPower efficiencies for both types of injections. Each injected event is plotted by injected SNR (loudness) vs. injected frequency. A blue dot indicates a found event and a red x indicates a missed event. See Table 3.1 for the overall efficiency of each ETG.

The efficiency of the Hann-window basis ETG ExcessPower is shown in Figure 3.17. ExcessPower had the next best efficiency for sine-Gaussians, of over 98%. ExcessPower only missed sine-Gaussian injections below an SNR of 10. ExcessPower also has high efficiency for white noise burst events, above 95%. However, ExcessPower also missed white noise burst injections with higher SNRs, with a few missed outliers above an SNR of 20.

The BayesWave and PCAT algorithms did not participate in the study with testing efficiency as a primary goal as they will not be used for event identification in Advanced LIGO's first observing run. However, it is still worth commenting on their current efficiency performance, shown in Figure 3.18, since event detection is a long term future goal of the BayesWave algorithm and it is useful to understand the types of events PCAT does not identify.

Similar to the true ETGs, BayesWave has a high efficiency for sine-Gaussian injections - nearly 95%. However, there are some very high SNR outliers seen in the scatter plot, particularly a few at low frequency (below 60 Hz).




Figure 3.18: A series of scatter plots showing efficiencies for BayesWave and PCAT for both types of injections. Each injected event is plotted by injected SNR (loudness) vs. injected frequency. A blue dot indicates a found event and a red x indicates a missed event. See Table 3.1 for the overall efficiency of each ETG.

PCAT shows an unexpected sharp cutoff in efficiency for sine-Gaussians, being very unlikely to detect sine-Gaussian events below an SNR of 20, and very likely to detect them above an SNR of 20. PCAT also struggles to identify low frequency sine-Gaussian waveforms, even at SNRs above 20.

The efficiency of BayesWave drops by roughly 10% for white noise burst injections. PCAT efficiency drops to under 50% for white noise bursts. This is expected for events that don't show a clear amplitude peak, which is how PCAT identifies events. Even so, this finding will likely prove useful to keep in mind as not all transient noise

is very well localized in time, and these types of events will not be identified for glitch classification by PCAT.

Parameter estimation

Timing

To evaluate timing performance, the distribution of the difference between the injected and recovered time for detected events was evaluated for each ETG. Table 3.2 shows the standard deviation in this timing distribution for each ETG by injection type. All ETGs show a marked decrease in timing precision for white noise bursts relative to sine-Gaussians.

Standard deviation in time difference		
ETG	SGs (s)	WNBs (s)
Omicron	3.4e-4	0.014
ExcessPower	0.014	0.069
DMT Omega	5.4e-4	0.018
BayesWave	0.053	0.044
PCAT	7.9e-4	0.012

Table 3.2: A table showing the standard deviation for the recovered time difference distributions by injection waveform type for each ETG.

To fully illustrate the behavior of each timing distribution and to identify any systematic timing offsets, the distributions are also plotted by waveform type for each ETG.

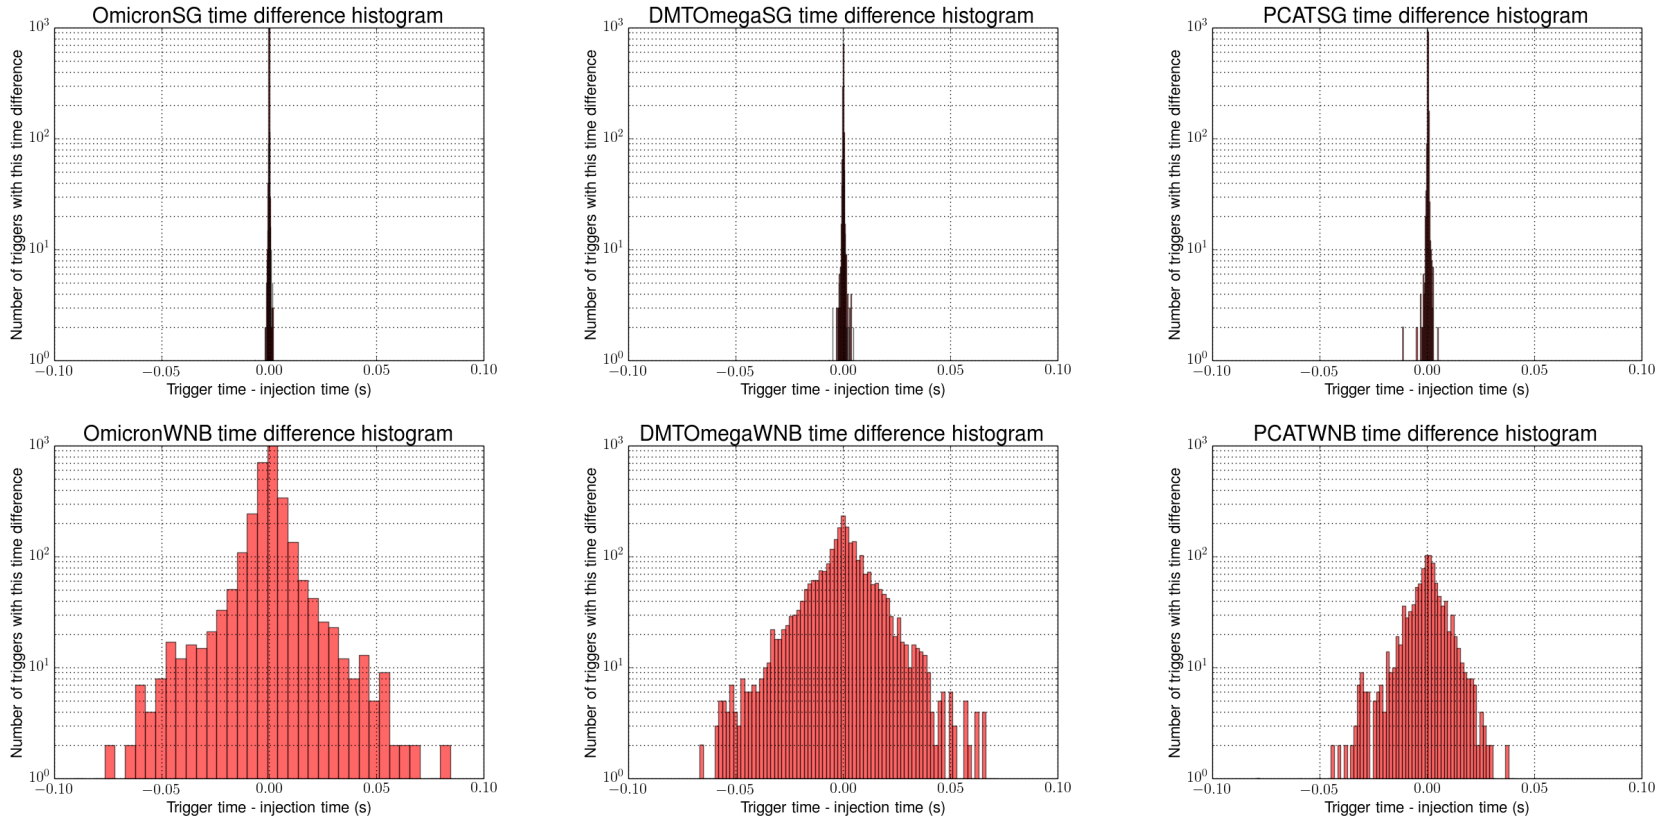


Figure 3.19: Histograms showing the distribution of timing difference between recovered and injected events for Omicron, DMT-Omega, and PCAT, each with a very narrow distribution. See Table 3.2 for the timing distribution standard deviation of each ETG.

Omicron, DMT-Omega, and PCAT all had very narrow distributions centered around zero for sine-Gaussian injections, seen as very sharp peaks in Figure 3.19. The standard deviation of timing difference for all three algorithms is less than 0.001 seconds, shown in Table 3.2. Omicron has the lowest standard deviation of the three, over DMT-Omega by a factor of 1.6 and over PCAT by a factor of 2.9.

Omicron, DMT-Omega, and PCAT again have very similar timing distributions for white noise bursts. The standard deviation of time difference increased over sine-Gaussian performance for Omicron by a factor of 41, for DMT-Omega by a factor of 33, and for PCAT by a factor of 15. Less precise time recovery for white noise bursts is expected for both sine-Gaussian basis ETGs, where the basis does not match the waveforms well, and PCAT, which identifies the central time of an event by the peak amplitude in a time series.




Figure 3.20: Histograms showing the distribution of timing difference between recovered and injected events for ExcessPower. See Table 3.2 for the timing distribution standard deviation of each ETG.

ExcessPower has much less precise timing for sine-Gaussian waveforms, seen in Figure 3.20. Although the timing difference is desirably centered at zero, the standard deviation is 41 times greater than Omicron’s timing distribution for recovered sine-Gaussians.

ExcessPower’s timing resolution increases only by a factor of 4 for white noise bursts, but shows the least precise time resolution of white noise bursts over all considered ETG trigger sets.




Figure 3.21: Histograms showing the distribution of timing difference between recovered and injected events for BayesWave. See Table 3.2 for the timing distribution standard deviation of each ETG.

BayesWave had a consistent timing offset toward reporting sine-Gaussian event times later than they are injected²⁷, which is not observed for white noise bursts, as shown in Figure 3.21. The standard deviation of the time difference for white noise bursts is notably worse for BayesWave relative to Omicron, DMT-Omega and PCAT, but significantly better than ExcessPower.

Frequency

To evaluate frequency recovery performance, the distribution of the difference between the injected and recovered frequency for detected events was evaluated for each ETG. Table 3.3 shows the standard deviation in this frequency difference distribution for each ETG by injection type. All ETGs had significantly less precise frequency recovery for white noise bursts, with the surprising exception of PCAT.

²⁷BayesWave’s timing offset is thought to be understood, and a new trigger set is in the process of development at the time of writing.

Standard deviation in frequency difference		
ETG	SGs (Hz)	WNBs (Hz)
Omicron	81.6	198.6
ExcessPower	196.4	616.4
DMT Omega	92.9	121.2
BayesWave	83.9	138.8
PCAT	142.0	120.6

Table 3.3: A table showing the standard deviation in the distribution of frequency difference (recovered frequency - injected frequency) for each ETG, by injected waveform type. The sine-Gaussian and wavelet-basis ETGs (Omicron, DMT Omega, BayesWave) have a significantly smaller standard deviation away from accurate recovery for the sine-Gaussian injections. ExcessPower has overall the largest standard deviation.

To investigate any tiling artifacts that might cause bias in frequency reconstruction, the percent difference in frequency vs. frequency was also plotted for each ETG by injection type. In these plots, each detected event is represented as a green x. For reference, a percent frequency difference of 10% is shown as two dashed lines.

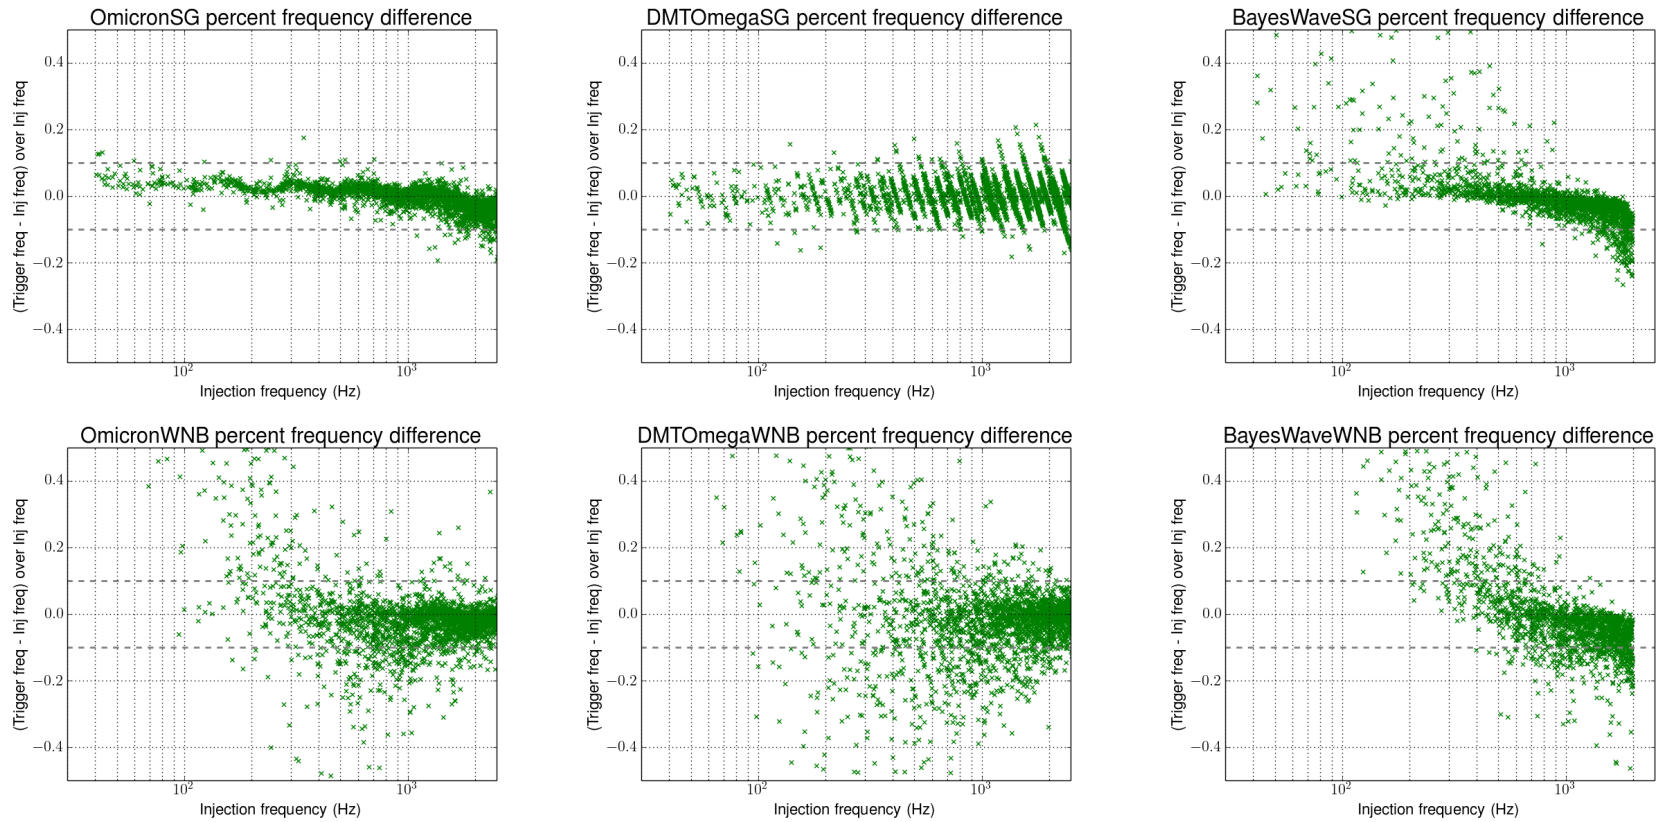


Figure 3.22: The percent difference in frequency vs. frequency for Omicron, DMT-Omega, and BayesWave. See Table 3.3 for the standard deviation in frequency difference distribution for each ETG.

Omicron, BayesWave and DMT-Omega all have accurate frequency recovery, as shown in 3.22, with standard deviations in frequency difference less than 100 Hz, as listed in Table 3.3.

Omicron shows the most precise frequency recovery across all frequency bands, with a subtle oscillating tiling artifact feature and a small bias toward underestimating event frequency above 2 kHz. BayesWave has the next most precise frequency recovery for sine-Gaussians across the analyzed range, however, BayesWave also has some greatly overestimated outliers with injected frequencies below 1 kHz, and a strong trend toward underestimation of frequency above 1kHz. DMT-Omega also has a fairly precise frequency difference distribution for SGs, but with a very apparent tiling feature where the frequency difference can be distinctly seen walking down from one time-frequency tile set to the next as injected frequency increases.

Apart from ExcessPower, Omicron sees the biggest increase in frequency difference standard deviation using white noise bursts, with an increase of over 100 Hz. This is likely due to a skew toward overestimating event frequency at lower injected frequencies not observed for sine-Gaussian injections. BayesWave behaves in a similar way to Omicron for white noise bursts, with a much more exaggerated overestimation of frequency at lower frequencies. The tiling artifact seen in DMT-Omega's frequency recovery of sine-Gaussians is not observed for white noise bursts. DMT-Omega shows a much wider distribution in frequency difference at lower injected frequencies than Omicron or BayesWave, which tend to skew toward overestimation.

PCAT has the next most precise standard deviation of 142 Hz. There are no obvious tiling features seen in Figure 3.23, which is expected since PCAT does not tile events. Instead, PCAT takes the Fourier transform of each identified event and reports the frequency with the maximum amplitude. This process becomes less reliable and




Figure 3.23: The percent difference in frequency vs. frequency for PCAT. See Table 3.3 for the standard deviation in frequency difference distribution for each ETG.

precise at higher injected frequencies for sine-Gaussians and lower injected frequencies for white noise bursts.

The frequency difference standard deviation for PCAT decreases by 21 Hz using white noise bursts instead of sine-Gaussian injections. This is surprising for an algorithm that recovers frequency using the frequency of the maximum amplitude of a Fourier-transformed time series because sine-Gaussian waveforms are generally more peaked and less broadband in frequency than white noise burst events. Generally, PCAT produced the most precise frequency recovery of the analyzed ETGs for white noise bursts.

ExcessPower has the least precise resolution of the frequency of sine-Gaussian injections, with a very strong bias toward overestimation at lower frequencies and then a strong oscillatory tiling artifact at high frequencies²⁸, as shown in Figure 3.24.

The same dominating bias and oscillating artifact is again obvious in ExcessPower's recovery of frequency for white noise burst events. Although the shape is very similar,

²⁸ExcessPower's behavior in frequency reconstruction is thought to be understood, and is currently in development at the time of writing.




Figure 3.24: The percent difference in frequency vs. frequency for ExcessPower. See Table 3.3 for the standard deviation in frequency difference distribution for each ETG.

the standard deviation of the frequency difference increases dramatically to over 600 Hz, which is 20% of the frequency range of the study.

SNR

To evaluate SNR recovery performance, the distribution of the difference between the injected and recovered SNR for detected events was evaluated for each ETG. Table 3.4 shows the standard deviation in this SNR difference distribution for each ETG by injection type. All clustered ETG triggers tend to significantly underestimate the SNR of white noise bursts with the exception of BayesWave.

ETG	Standard deviation in SNR difference			
	SG mean	SG σ	WNB mean	WNB σ
Omicron	2.4	1.4	-13.0	9.1
ExcessPower	-11.7	23.3	-24.0	11.0
DMT Omega	-1.1	1.3	-15.6	8.9
BayesWave	1.4	1.9	-1.0	6.0
PCAT	-27.3	6.5	-30.0	6.4

Table 3.4: A table showing the mean and standard deviation in the distribution of SNR difference (recovered SNR - injected SNR) for each ETG by injected waveform type.

To fully illustrate the behavior of each SNR difference distribution and to identify any systematic SNR estimation offsets, the distributions are also for plotted by waveform type for each ETG.

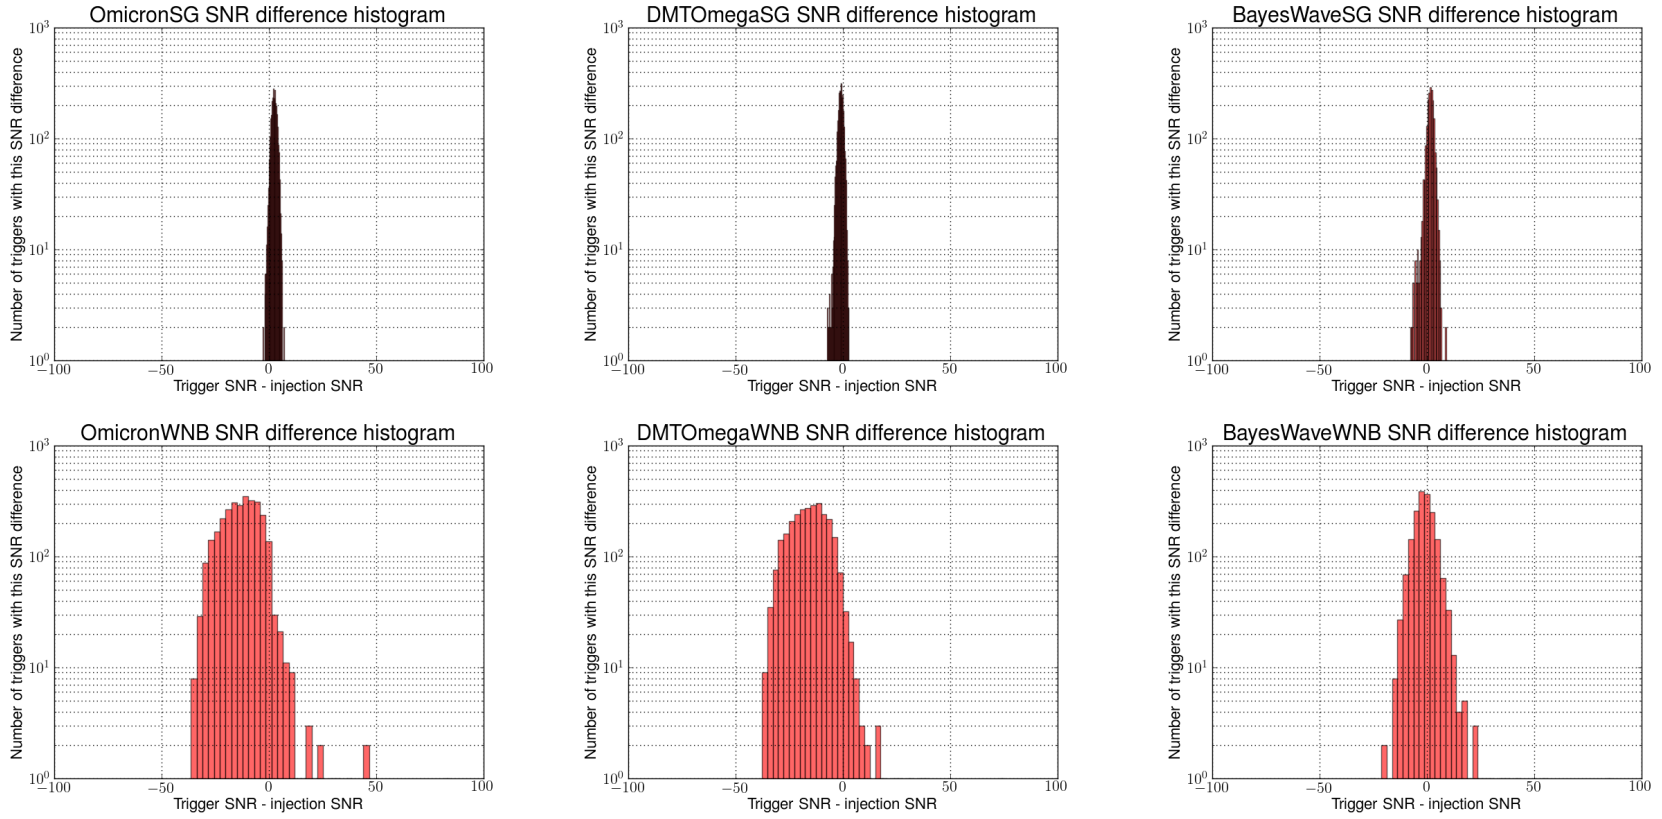


Figure 3.25: Histograms showing the distribution of SNR difference between recovered and injected events for Omicron, DMT-Omega, and BayesWave. See Table 3.4 for the standard deviation of the SNR distribution in SNR difference of each ETG.

Omicron, DMT-Omega, and BayesWave all show tight SNR difference distributions centered near zero for sine-Gaussians, as shown in Figure B.7. Omicron’s slight overestimation of SNR for these waveforms is expected because of the way outliers are excluded in the estimation of excess tile energy, described in Section 3.1.3. The SNR difference distribution for DMT-Omega is closest to zero in mean and also has the smallest standard deviation, by a small margin. BayesWave performs very similarly to the sine-Gaussian basis ETGs in SNR estimation of sine-Gaussians.

The sine-Gaussian and wavelet basis algorithms Omicron, DMT-Omega, and BayesWave all have broader distributions for the recovery of SNR using white noise burst waveforms, indicating less precision. Omicron and DMT-Omega both tend to underestimate the total event SNR by roughly 15 on average. This is expected for clustered Omicron and DMT-Omega triggers, which report only the SNR of the most significant tile, which is not likely to capture the entire power of white noise burst events. The SNR difference distribution for BayesWave changes only slightly using white noise bursts - BayesWave shows the most accurate mean and most precise standard deviation. This indicates that the total event power is better captured by a single unclustered BayesWave trigger relative to Omicron and DMT-Omega.

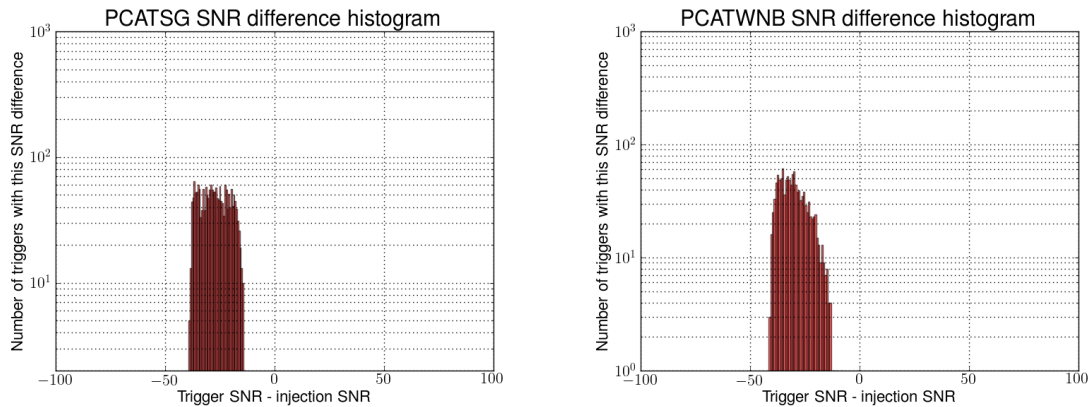


Figure 3.26: The distribution of SNR difference between recovered and injected events for PCAT. See Table 3.4 for the standard deviation of the standard distribution in SNR difference of each ETG.

PCAT has a large, consistent underestimation of SNR, an average offset of 27, for sine-Gaussian waveforms, as shown in Figure 3.26. The standard deviation of the SNR difference distribution is reasonably small at 6.5, indicating a fairly consistent SNR underestimation. This is confirmed with supplementary plots showing injected SNR vs. recovered SNR in Appendix B.

PCAT's underestimation of SNR using white noise bursts waveforms is quite similar to performance using sine-Gaussians. The mean of the SNR difference distribution grows slightly less accurate, to an average offset of 30, and the standard deviation remains nearly the same. Comparing the two plots in Figures 3.26, the shape of the distribution shifts slightly toward further underestimating SNR, but otherwise looks very similar.




Figure 3.27: The distribution of SNR difference between recovered and injected events for ExcessPower. See Table 3.4 for the standard deviation of the standard distribution in SNR difference of each ETG.

ExcessPower tends to underestimate the SNR of sine-Gaussian waveforms by 12 on average. However, ExcessPower also produces enormously loud overestimations of SNR far above any possible injected SNR, with one event in excess of 10,000, as illustrated in Figure 3.27. This effect is not observed for ExcessPower using white noise bursts.

ExcessPower underestimates the SNR of white noise burst events by about a factor of two more than for sine-Gaussians on average. Contrary to performance using sine-Gaussians, the SNR of all white noise bursts is underestimated. The standard deviation of the SNR distribution becomes much more compact, decreasing by about a factor of two compared sine-Gaussian performance.

Discussion

Differences in observed ETG performance aren't necessarily due to any fundamental limit of a chosen basis. A relative lapse in performance may be explained by differences in implementation of data conditioning, event identification criteria, or trigger post-processing such as downselection.

Note that these results represent performance at the current time. There has been significant progress in ETG performance made since previous tests. A striking example is the efficiency of DMT-Omega using white noise burst waveforms less than six months prior to the time of writing, shown in Figure 3.28. This version of the DMT-Omega configuration consistently missed significant portions of power using white noise burst events. A previous iteration of this study diagnosed this to be aggressive downselection of triggers in post-processing. This process selected only the loudest trigger in any instance of trigger overlap in time, which caused much of the power of white noise burst events to not be represented in the resulting clustered trigger set.

This study has proven to be a useful and effective tool in ensuring the efficiency and bulk parameter recovery of all participating ETGs and burst algorithms is well understood. Many improvements to ETGs have already been implemented as a result, and more are expected before the beginning of Advanced LIGO's first observing run.




Figure 3.28: The efficiency of a prior iteration of DMT-Omega triggers on white noise burst waveforms using a broader subset of the generated injection data set. This poor efficiency was the result of over-aggressive downselection of triggers in post-processing, before clustering, that selected only the loudest trigger in any instance of trigger overlap in time.

The full code used for this study can be found in the Detector Characterization software git repository, and instructions for running the code can be found here: <https://wiki.ligo.org/DetChar/TrigMatchingInstructions>.

Additional supplementary plots exploring other perspectives of bulk parameter recovery, in particular the accuracy of parameter recovery versus injected event SNR, can be found in Appendix B.

3.2.4 Recommendations based on results

The inclusion of white noise burst injections in this study highlighted the need to update two current ETG practices.

1. Cluster ETGs using windows

The commonly used clustering algorithm used for this study, `ligolw_bucluster`, clusters events that overlap in time and frequency. For events that are very well localized, like sine-Gaussian injections, this often work well. However, for events like white noise bursts that are not well localized, this will cause excess power related to the same event to be clustered disjointly. All considered ETGs use sinusoids or frequency-band-limited data discretized in time to tile data in time and frequency (Omicron, DMT-Omega, ExcessPower), which naturally tends to discretize power for events where power is dispersed.

The need for a extended window in clustering has also previously been demonstrated in Figure 3.9 and discussed in Section 3.1.2, which details Omicron’s window-based cluster algorithm. Omicron’s clustering approach is expected to deal with disparate events well while allowing minimal contamination from unrelated noise. It is strongly recommended this approach, or perhaps a similar approach also including a frequency window, be adopted for other ETGs in the future.

2. Improve parameter reporting for clustered triggers

As stated in the SNR recovery results in Section 3.2.3, the reported SNR of the clustered triggers produced by ETGs Omicron, DMT-Omega, and ExcessPower all consistently underestimate the SNR of white noise burst events. This is expected, as: 1) no single time-frequency tile would be a good match for a white noise burst event with power spread out in time and frequency, and 2) clustered trigger SNR is currently

reported as the SNR of the most significant trigger in the cluster²⁹. This effect is clearly demonstrated by comparing the SNR recovery of sine-Gaussians to white noise bursts for each ETG. The reported trigger SNR is generally a trusted quantity that other detector characterization algorithms or burst analysis pipelines will use as a measure of event significance, so it is important it be reported as accurately as possible.

It would be extremely difficult to recover the SNR of a non-localized burst event very accurately with the ETGs Omicron or DMT-Omega, which use an overcomplete basis. However, it would be possible to collect more of the power from different regions of the time-frequency parameter space that do not directly overlap with the most significant tile in the cluster. It is left to future work to design and implement such a scheme.

For ExcessPower, which uses an orthogonal basis, it should be possible to recover the total event SNR from the tiles forming the clustered trigger. However, ExcessPower's basic estimation of trigger SNR requires significant improvement before embarking on improving the SNR recovery of clustered triggers.

The recovery of event central frequency could also be significantly improved with a weighted average scheme for clustered triggers. As shown in the frequency recovery results in Section 3.2.3, the standard deviation of the frequency difference distribution grew significantly for the ETGs Omicron, DMT-Omega, and ExcessPower when using white noise burst injections. Similarly to SNR reporting, this is because the full power of a white noise burst event is not well recovered by a single tile, but the peak frequency of the clustered trigger is reported as the frequency of the cluster's single most significant tile.

²⁹Reporting cluster SNR as the SNR of the loudest trigger in the cluster does work very well for localized burst events that are closely related to the basis function, such as sine-Gaussians.




Figure 3.29: A histogram of the difference between injected frequency and the reported frequency of a clustered Omicron trigger by different measures: peak frequency (in blue), or the central frequency of the loudest trigger in a clustered trigger, and central frequency (in red), or the average of the minimum and maximum frequency bounds of a clustered trigger. Clearly, peak frequency is the better measure, however, as shown in Section 3.2.3, this measure is not consistently accurate in recovering the central frequency of white noise burst events.

The accurate recovery of frequency is not improved by using the central frequency of a clustered trigger instead. As demonstrated in Figure 3.29, the peak frequency, although not a great measure, is better than the average of the minimum and maximum frequency of the cluster associated with the event. This is likely due in part to the clustering method used in the study, which used no window to group related islands of signal energy together in time and frequency. However, an algorithm assigning an average cluster frequency based on the frequency of each tile composing the cluster, perhaps weighted by loudness, may improve the accuracy of the reported cluster frequency.

The accurate recovery of event frequency is becoming increasingly more important to detector characterization algorithms that rely on time, frequency, and event SNR to classify or correlate detector noise, as mentioned in Chapter 5. It is a worthwhile

endeavor left to future work to improve it for burst events that are not well localized, as some noise transients (or gravitational wave events) are sure to be.

CHAPTER 4

DETECTING AND RECONSTRUCTING CORE-COLLAPSE SUPERNOVAE WITH ADVANCED LIGO

4.1 Models for the CCSN explosion mechanism

A core-collapse supernova (CCSN) is the final stage of the life cycle of a $8\text{-}130 M_{\odot}$ main sequence star as the fusion in the iron core burns out and the degeneracy pressure can no longer support the star against gravity. The subsequent gravitational core collapse releases $\sim 10^{53}$ ergs of energy, 99% in the form of neutrinos of all flavors [62], and forms a proto-neutron star. Neutrinos observed during the 1987 supernova have confirmed this general model [24].

In 1980, Goldrieich and Weber calculated a family of stable core collapse configurations and showed that as the star core collapses, it separates into a slowly collapsing inner core and a rapidly collapsing outer core. When the inner core reaches nuclear density, it bounces into the rapidly collapsing outer core, generating a shock wave that quickly stalls from losing energy in the interaction with the outer core [43].

Explaining the explosion mechanism that re-energizes this shock wave and drives an explosion is a current definitive problem in core-collapse theory research.

Two of the core-collapse models presently accepted in the literature as likely candidates are the *neutrino* and *magnetorotational* mechanisms [75]:

- The neutrino mechanism assumes that some small fraction of the energy emitted in the form of neutrinos is absorbed by the shock, re-launching the stalled shock wave and producing an explosion. Recent 2D and 3D axisymmetric simulations of the neutrino mechanism have produced reliable explosions in the observed range of emission energies [75], [62].
- An explosion produced by a rotating core-collapse follows from conservation of angular momentum and magnetorotational instability: the collapse of the core induces a 1000x increase in rotational velocity [75] in the outer layers of the star. The differential rotation between the turbulently and rapidly rotating outer layers and slower uniform rotation of inner core could induce magnetic-field amplification that yields jet-like explosions along the axis of rotation [62]¹.

Gravitational radiation will be emitted by core-collapse supernova events depending on the amount of asymmetry involved. According to [75], the gravitational wave features predicted by the rotating core-collapse and neutrino-driven explosion models are mutually exclusive. Thus, detecting a gravitational wave signal and classifying it by its supernovae emission process could powerfully constrain the physics of the core collapse supernovae explosion mechanism [62].

Examples of key waveform differences between these models are shown in Figure 4.1. Essentially, it is possible to confidently map the distinctive gravitational waveform features to an explosion mechanism with an accurately reconstructed signal [75].

It's also worth noting that core-collapse SNe are rare events. The typical CCSN rate for observed Milky Way equivalent galaxies is ~ 3 per century [19]. Advanced LIGO may be able to see rotating core-collapse signals up to ~ 100 kpc [75], [46].

¹Additionally, post-bounce rotational instabilities could yield asymmetric deformations of the rotating compact star, producing gravitational waves for up to hundreds of milliseconds after the initial collapse [62].




Figure 4.1: Examples of supernovae waveforms corresponding to two leading models for core-collapse supernovae explosion. Modified from [62].

4.2 Waveform reconstruction performance study

The following study was designed to evaluate the ability of current gravitational wave burst analysis software to accurately reconstruct modeled core-collapse supernovae waveforms in simulated Advanced LIGO and Advanced Virgo data.

The two major goals were to characterize:

- The current ability of the pipelines to reconstruct modeled waveforms generated with different explosion mechanisms, and how different key features between waveform families affects the accuracy of their recovery².
- The relative strengths and weaknesses of each considered algorithm in reconstructing modeled CCSN waveforms.

²Any group of waveforms that are generated using the same underlying set of assumptions and mechanics, perhaps with differences in progenitor stars, is referred to as a ‘waveform family’.

4.2.1 Overview of reconstruction methods

The study targeted two burst analysis algorithms intended to follow-up externally triggered supernova candidate events during the Advanced detector era: cWB2G and BayesWave [93].

4.2.1.1 cWB2G

Coherent WaveBurst (2nd Generation) is the flagship all-sky burst search algorithm. cWB2G identifies burst candidate events by tiling the data in time and frequency via a wavelet transform [72] in a similar manner to the method described in Chapter 3 [56] and extracting significant events using a coherent likelihood statistic maximized over all potential sky positions.

First, data from an arbitrary number of interferometers is sampled and whitened with linear prediction filters. Then, the likelihood of observing the data assuming a gravitational wave signal is maximized over a grid of sky locations using the following techniques [57]:

- Constraints called *regulators* are used to constrain the likelihood to physically likely configurations of the interferometer network antenna pattern (as shown in Figure 2.3) and can also be used to assume or weigh the likelihood of a particular polarization.
- The discrete sampled data is mapped to time-frequency tiles using the fast Wilson-Daubechies Meyer wavelet transform [72].

- Time-delay filters corresponding to possible positions across the sky³ are applied to the data.
- Excess power triggers are identified for each time delay filter and time-frequency resolution used.
- For every tile in an excess power trigger event, the coherent energy (E_{coh}) and incoherent energy (E_{incoh}) between detectors is used to calculate a network correlation coefficient [57]:

$$C_{net} = \frac{E_{coh}}{E_{incoh} + |E_{coh}|}$$

which is an effective metric to select true signals from background noise events as transient noise is unlikely to produce energy that happens to be coherent between detectors in the network.

For each event, the cWB significance ranking statistic, κ ,⁴ is also calculated:

$$\kappa = \left(\frac{E_{coh}}{N - 1} \right)^{1/2} \quad (4.2)$$

After selection cuts based on significance, κ , and network correlation coefficient, C_{net} , the computed trigger properties for candidate events are written to disk. This includes the estimated event SNR in each detector, which is defined in Equation 3.26. The network SNR, an estimation of the event SNR across the analyzed detector network, is also reported. ρ_{net} is computed as:

³Or some subset of the sky, if a sky mask is applied for a signal with a known associated electromagnetic signal.

⁴In internal cWB notation the significance is called ρ , but as we are using ρ to indicate SNR we use κ here instead.

$$\rho_{net} = \kappa\sqrt{N} \tag{4.3}$$

where N is the number of interferometers in the analyzed network.

In addition to identifying burst candidate events, cWB2G also produces useful information on detected events, including a reconstructed waveform as recovered in each interferometer.

4.2.1.2 BayesWave

BayesWave is parameter estimation and signal vs. noise classification algorithm that is intended to be used to follow up events identified by the all sky burst search (cWB2G) or other interesting times.

BayesWave estimates a posterior distribution for a recovered gravitational waveform inferred from the data and network antenna pattern. To match the signal as closely as possible, BayesWave uses a reversible jump⁵ Markov-chain Monte Carlo algorithm that explores the application distribution of Morlet-Gabor wavelets and wavelet parameters to the signal fit [29].

BayesWave also includes a powerful Bayesian classification technique that distinguishes between gravitational wave events, Gaussian noise, and noise transients, or glitches⁶ [29]. However, this study focuses on the signal reconstruction ability of BayesWave. All BayesWave runs produced for this study did not enable signal classification.

⁵Reversible jump MCMC allows the sampling of a posterior distribution where the number of model dimensions is also allowed to vary. For BayesWave, these additional dimensions are wavelets added to or subtracted from the signal fit.

⁶Morlet-Gabor wavelets are used to fit both true signals and transient noise.

Following the same notation as the BayesWave publication [29], the gravitational wave signal properties are inferred using Bayes' theorem:

$$p(h|d, M) = \frac{p(h|M) p(d|h, M)}{p(d|M)} \quad (4.4)$$

which states that the posterior distribution $p(h|d, M)$ of a gravitational wave h , given data d and a model M , is equal to the prior distribution of gravitational wave h in model M , $p(h|M)$, times the likelihood of observing the data d for the wave h given model M , $p(d|h, M)$, over the evidence for the model given the data, $p(d|M)$.

The data is assumed to contain some combination of gravitational wave signal h and detector noise such that:

$$d = Rh + g + n_G \quad (4.5)$$

where R is the network response, or antenna pattern, of the interferometer network, h is the gravitational wave signal, g represents transient noise glitches, and n_G is Gaussian noise. Note that BayesWave does not assume that instrument noise is Gaussian, but explicitly accounts for noise transients in this data model.

The likelihood function is computed based on the requirement that the residual, r , left after subtracting the gravitational wave signal model from the data is consistent with the model for the noise, i.e.

$$r = d - Rh \quad (4.6)$$

is consistent with noise model $n = g + n_G$.

The PSD curve of the Gaussian noise component of the noise model is computed with the BayesLine estimation algorithm [61].

The gravitational wave signal is modeled with a bank of Morlet-Gabor wavelets, where number of included wavelets in the distribution used for signal fitting is a tunable parameter to optimize the best possible signal fit and computational efficiency. The signal model fits data averaged over all interferometers in the analyzed network⁷. Changes to each intrinsic parameter of a wavelet; time, frequency, amplitude, Q, and phase, as well as the ‘birth’ and ‘death’ of new wavelets are proposed and explored by the RJMCMC algorithm. The wavelets modeling the gravitational wave signal are also constrained by extrinsic parameters: ϵ , the polarization eccentricity described below, ψ , the polarization angle, and θ and ϕ , which correspond to sky location as described in Figure 2.3. Any proposed changes to extrinsic parameters must be internally consistent with the interferometer network to be accepted by the RJMCMC algorithm.

Instead of using regulators like cWB2G, BayesWave currently assumes an elliptical polarization, such that the signal is modeled with a single set of wavelets describing h_+ . h_\times is assumed to be related to h_+ in the frequency domain by ellipticity ϵ :

$$h_\times = \epsilon h_+ e^{i\pi/2}. \tag{4.7}$$

This is not expected to have an impact on the first Advanced LIGO observing run, where having only two detectors with comparable sensitivity operation will make reconstructing h_+ and h_\times unfeasible. However, the following analysis used a network of three interferometers. As a result of assumed polarization, BayesWave did strug-

⁷The glitch model, also modeled with the same bank of Morlet-Gabor wavelets, independently fits only data in each detector.

gle to reconstruct the waveform of the interferometer where the fit due to assumed polarization was the worst. This happened to be Virgo in most cases, which is why BayesWave results are shown primarily for the LIGO interferometers L1 and H1.

The BayesWave priors used for generic signals are flat and evenly distributed across all possible values of: ϵ , which governs the assumed polarization, polarization angle ψ , and intrinsic wavelet parameters central time, central frequency, Q , and phase. The prior for the number of wavelets used to fit the signal and glitch models is also flat across the entire requested range. The prior for the amplitude of wavelets differs between signal and glitch models, with very high SNR wavelets more likely for the glitch model, as shown in Figure 1 of the BayesWave publication [29]. The clustering prior prefers clusters of wavelets over isolated wavelets, which we see later may be improved for following up a candidate CCSN signal.

BayesWave explores and calculates the evidence for all considered models: a combination of all possible Gaussian noise curve estimates, and all possible combinations of wavelets included in the glitch and signal models. RJMCMC is used to marginalize over all the parameters in each of these models using parallel tempering, where multiple chains of the RJMCMC are run simultaneously at different ‘temperatures’, T . ‘Hot’ chains explore the entire prior volume of each model and ‘cool’ chains explore the peaks of the posterior distribution. Exchanges of suggested model changes between chains helps to ensure the entire posterior distribution is explored.

Making use of the ‘temperature ladders’ intended to promote the full exploration of the posterior distribution, thermodynamic integration is used to calculate the evidence for each model as a function of chain temperature T :

$$\ln p(d|M) = \int_0^1 d\beta E_\beta [\ln p(d|\lambda, M)] \quad (4.8)$$

where β is $\frac{1}{T}$, λ is the model index, and E_β is the expectation value of the log likelihood of a chain with inverse temperature β , based on the number of transitions to each model.

The recovered waveform is extracted directly from the wavelets used to model the signal at the peak of the calculated posterior distribution.

4.2.2 Comparing cWB2G and BayesWave results

BayesWave produces a posterior distribution for each analyzed event that can be used to infer the relative confidence in signal recovery for events over a wide range of injected SNRs. Because the BayesWave noise model assumes Gaussian noise, this is analogous to the distribution obtained using cWB2G output to calculate the overlap between injected and recovered waveforms for a large number of injections, provided they are embedded in Gaussian noise. Additionally, BayesWave is very computationally expensive to run, making analyzing all events included in the injection set impossible, and analyzing as few as possible strongly desirable. For these reasons, a small subset of injections, about 20 per CCSN model, spanning a broad range of SNR was chosen for BayesWave analysis.

Note that for non-Gaussian noise the assumption that the BayesWave distribution produced for a small subset of events is analogous to the distribution obtained with many events using cWB2G does *not* hold. Choosing a small subset of events for BayesWave to analyze then introduces the possibility the result will be strongly influenced by a chance series of transient noise events in the detector network. For a future study including non-Gaussian noise we will need to change the focus of our

comparison to characterizing how well each algorithm deals with times with known transient noise contamination.

4.2.3 Included waveforms

For easy comparison to other ongoing CCSN recovery and classification work and the benefit of using already existing data, a subset of the modeled waveforms generated for the 2007⁸ supernova search were used for this study.


These waveforms were injected into simulated Gaussian noise colored to have the nominal Advanced LIGO design sensitivity spectrum, shown in Figure 2.2 and expected Advanced Virgo design spectrum shown in Figure 2.11. Waveforms were injected over all polarization angles, ψ in Figure 2.3, and at a fixed sky position of the Milky Way galactic center; the most likely direction of a potentially observable CCSN gravitational wave signal. In terms of the sky position angles defined in Figure 2.3, this corresponds to a uniform distribution over all possible ϕ , but a fixed theta equal to $\theta = 52.69$ degrees.

The injected waveforms were also scaled to different distances from Earth: 0.2, 0.4, 0.67, 1, 2, 4, and 10 kpc. Note that the Earth is roughly 8 kpc away from the Milky Way galactic center [42]. Depending on the orientation of the source relative to the Earth, the event SNRs produced by the same waveform and injected at the same distance from Earth varied greatly, as shown in Figure 4.2.


The injections were spaced 100 +/- 10 seconds apart to avoid cross-contamination.

Dimmelmeier


⁸2007 refers to the year of the LIGO and Virgo data used, corresponding to LIGO science run 5 and Virgo science run 1. The waveforms used were published significantly after 2007.



(a) Dimmelmeier



(b) Yakunin



(c) Mueller

Figure 4.2: A series of histograms showing the vast range of network SNR measured by cWB2G for the different included CCSN waveform families injected 0.4 kpc away from Earth. The wide range of SNRs observed within each waveform family is largely due to different internal source orientations relative to the global interferometer network configuration. The differences in SNR range between histograms is largely due to differing maximum gravitational wave strain amplitude produced by the three waveform sets relative to the Advanced detectors' predicted sensitivity.

Dimmelmeier waveform	Characteristic differential rotation length scale (A)	Precollapse core angular velocity (Ω)
s15a2o05	10^8 cm	2.4 rad/s
s15a2o09	10^8 cm	4.56 rad/s
s15a3o15	0.5×10^8 cm	13.31 rad/s


Table 4.1: The differential rotational length scale and pre-collapse central core angular velocity for each Dimmelmeier 2008 simulated model included in this study. The characteristic differential rotation length scale, denoted as A in [31], parameterizes differential rotation between the core center and a point some distance away from the core center. This parameterization is described by angular momentum $J = A^2(\Omega_c - \Omega)$ where Ω_c is the angular velocity at the core, reported in this table, and Ω is the angular velocity at some distance away, governed by the length scale A [31].

The Dimmelmeier 2008 catalog simulated 2D numerical models of axisymmetric rotating core collapse supernovae [31]. Three of the modeled waveforms, shown in Figure 4.3, were included in the bank of CCSN injections for the 2007 SN search and used for the following study.


The progenitor star models, taken from the Woosley 2002 catalog [99], had different rates of rotation and rotation profiles imposed. The evolution of the general relativistic hydrodynamics was simulated using a spherical grid, including shock modeling the effects of neutrino emission.

The waveforms included in the 2007 supernova search injection set were progenitor stars with total mass $15 M_\odot$, an iron core of mass $1.55 M_\odot$, and core radius of 1.96×10^6 m. Each included model rotates with a different rate and rotation profile as shown in Figure 4.3 and Table 4.1. All three included Dimmelmeier models were simulated using the same nuclear equation of state⁹ [31].


⁹Waveforms corresponding to the softer Lattimer and Swesty equation of state (1991) were selected [59] as opposed to more stiff Shen et al. (1998) [85].



(a) s15a2o05



(b) s15a2o09



(c) s15a3o15

Figure 4.3: The three modeled rotating core collapse waveforms included from the Dimmelmeier 2008 catalog [31]. The time of core bounce corresponds to the time of maximum gravitational wave amplitude. All waveforms are whitened by the Advanced LIGO design noise curve to show their form as they would be observed in ideal noise conditions. All included Dimmelmeier models use a progenitor star with a mass of $15M_{\odot}$. Waveform s15a2o05 is a star with moderate differential rotation and a slowly rotating core. Waveform s15a2o09 has moderate differential rotation with a moderately rotating core. Waveform s15a3o15 has strong differential rotation and a rapidly rotating core. The differential rotation length scales and angular velocity of the star core pre-collapse can be found in Table 4.1.

Dimmelmeier waveform	Maximum GW strain amplitude at 10kpc
s15a2o05	3.94×10^{-21}
s15a2o09	7.83×10^{-21}
s15a3o15	4.53×10^{-21}

Table 4.2: A table showing the estimated maximum GW strain amplitude for each waveform for a source 10kpc away from Earth. Referenced from Table III of [31].

The included Dimmelmeier waveforms are predicted to be detectable by Advanced LIGO and Advanced Virgo, all producing a maximum gravitational wave strain of greater than 10^{-21} at a distance of 10 kpc from Earth, which includes the Milky Way galactic center. The estimated maximum gravitational wave strain amplitudes are listed in Table 4.2.

It is also important to note that because these are 2-dimensional axisymmetric simulations, all gravitational waveforms taken from the Dimmelmeier 2008 set are linearly polarized (h_+ only).

Yakunin

The Yakunin 2010 catalog simulated 2D axisymmetric non-rotating core-collapsed stars with neutrino-driven supernova explosions. Like the Dimmelmeier simulations, the Yakunin simulations made use of the Lattimer and Swesty equation of state [59]. Progenitor stars were taken from from the Woosley et al. models described in [100]. One of the produced waveforms was included in this study, corresponding to a progenitor star of mass $15M_{\odot}$, shown in Figure 4.4 [101].

The Yakunin 2010 paper identified key features of these simulated waveforms, each corresponding to a different physical mechanism, that can be seen in Figure 4.4 [101]:

- **Prompt signal** - generated by convection inside the newly formed proto-neutron star and asymmetries in-falling matter deflected by the expanding shock




Figure 4.4: The waveform with a $15M_{\odot}$ progenitor star included from the Yakunin 2010 simulation set [101]. Time = 0 seconds corresponds to the time of the core bounce after the initial collapse. This waveform is whitened by the Advanced LIGO design noise curve to illustrate its form as would be observed under ideal noise conditions.

wave. The prompt signal can be seen in Figure 4.4 from bounce ($t=0$ seconds) up to roughly 70 ms post-bounce.

- **Quiescent stage** - associated with declining amplitude of initial signal as convection in the proto-neutron star stops. The quiescent stage can be seen in Figure 4.4 from the end of the prompt signal (roughly 70ms) up to about 170 ms post-bounce.
- **Strong signal** - due to time-varying matter asymmetries produced by modulations in the shock radius as standing accretion shock instability (SASI) develops and SASI-induced accretion matter flows collide with the proto-neutron star. The strong signal portion of the waveform can be seen in Figure 4.4 from the end of the quiescent stage (about 170 ms) up to roughly 400 ms post-bounce. This waveform feature has the highest amplitude and is the most likely to be detected by Advanced ground-based gravitational wave interferometers.

The $15M_{\odot}$ progenitor was chosen from the Yakunin 2010 waveform set because it has the most power above the Advanced LIGO design sensitivity spectrum, as shown in Figure 2 of [101]. This model would be detectable by Advanced LIGO at a distance of 10 kpc with a peak strain amplitude of 10^{-21} at that distance, and a total expected emitted gravitational wave energy of $7 \times 10^{-7}M_{\odot}c^2$.

Like the Dimmelmeier 2008 waveforms, the included Yakunin waveform is the result of a 2-dimensional axisymmetric simulation that produced linearly polarized gravitational waveforms (h_+ only).

Mueller

The Mueller 2012 catalog simulated 3D non-rotating collapsing stars with neutrino-driven supernova explosions [71]. A subset of these was included in this study such that each of the included progenitor stars was represented, as shown in Figure 4.5.

The Mueller simulations used three progenitor stars from three different progenitor star model evolutions. The W15 model is a $15M_{\odot}$ progenitor taken from the 1998 Woosley and Weaver catalog [98]. The L15 model is a $15M_{\odot}$ progenitor from Limongi et al [60]. The N20 model is a $20M_{\odot}$ progenitor designed to have the same metallicity and evolution as SN 1987A, produced by Shigeyama and Nomoto [86].

The initial symmetry of the models was perturbed with random 0.1% amplitude seeds, and then the models were evolved from core bounce through up to 1.4 seconds using the 1996 Janka and Mueller equation of state [71].

Note that Mueller waveforms have no assumed symmetry. The full waveforms were generated with a grid of internal orientation angles α and β seen in the right panel of Figure 2.3 [71]. The relative amplitude over time generated by each progenitor star may vary based on the orientation of the source. Keep in mind that the examples shown in Figure 4.5 are for a particular sky location and source orientation, although




Figure 4.5: The three Mueller 2012 [71] non-rotating CCSN models producing neutrino-driven explosions included in the following waveform reconstruction study. The time of the core bounce and approximate time of the onset of the explosion are indicated with vertical dashed lines. Each different model, L15, W15, and N20 corresponds to a different progenitor star described in the text. These waveforms are whitened by the Advanced LIGO design noise curve to show their form as would be observed under ideal noise conditions.

the general features of each waveform should be consistent for the same progenitor star.

The paper describing the Mueller 2012 waveforms identifies distinct phases common to the progenitor star models [71]:

- **Shock expansion** - during which the shock wave rapidly propagates from the collapsed core to a radius of ~ 2 km, where it stalls. This expansion is roughly spherical, thus no significant gravitational wave strain is produced. As seen in Figure 4.5, this phase lasts from core bounce, marked with a vertical dashed line, until 80-150ms afterward, depending on the progenitor star.
- **Pre-explosion** - a period of post-shock convection and standing accretion shock instability (SASI) growth. The matter flow in this phase begins as quasi-spherical, but grows increasingly more non-radial as SASI develops. This phase can be seen in Figure 4.5 from roughly 80-150ms after core bounce through the onset of the explosion, also marked in a dashed line.
- **Post-explosion accretion** - the stalled shock wave is revived by neutrino heating, launching an explosion and accelerating the shock wave to the outer layers of the star. A significant gravitational wave signal is generated by the non-radial flow of matter in the rapid evolution of the shock radius, violent post-shock convection, and accretion of gas by the proto-neutron star. The asymmetric flux of neutrinos also contributes to the total gravitational wave signal during this phase, but to a lesser degree. For reference in Figure 4.5, this phase lasts from the onset of the explosion until roughly 500 ms after core bounce for models W15 and N20 and about 700 ms after core bounce for model L15.

- **Post-accretion** - the gas accretion on the proto-neutron star (PNS) has ended, and the PNS continues to lose mass in an approximately spherical neutrino wind. Gravitational wave signal amplitude decreases significantly. This phase lasts from 700ms after core bounce until the end of the produced waveform for the L15 progenitor. It lasts from 500 ms after core bounce until 1.25 seconds after bounce for the N20 model, and until 0.8 seconds after bounce for the W15 model.
- **Delayed non-radial flows** - this fifth phase, observed by only the W15 and N20 progenitor stars, is thought to be due to convection in the proto-neutron star¹⁰. This late phase can be seen in Figure 4.5 beginning roughly 1.25 seconds after core bounce in the N20 model, and roughly 0.8 seconds after bounce in the W15 model.

Both non-radial matter flow and asymmetric neutrino emission contribute to the total energy emitted in the form of gravitational waves, but asymmetric matter flow dominates by roughly two orders of magnitude. The total E_{GW} emitted for model L15 is roughly $1.5 \times 10^{-11} M_{\odot} c^2$, $0.8 \times 10^{-11} M_{\odot} c^2$ for model W15, and $0.5 \times 10^{-11} M_{\odot} c^2$ for model N20.

Unlike the 2D axisymmetric Dimmelmeier and Yakunin waveforms, Mueller 2012 waveforms are unpolarized. As a result, we will see that BayesWave, which currently assumes elliptical polarization¹¹, struggles to reconstruct Mueller waveforms relative to the other waveform families.

¹⁰PNS convection is thought to develop differently between the progenitor models based on the proximity of the unstable regions of convection to the PNS radial boundary defined in the simulation code.

¹¹This will not be the case in the future. By the time Advanced Virgo is also online and assuming polarization will detract from performance, BayesWave will no longer assume polarization.

It is also worth emphasizing that Yakunin et al 2010 [101] and Mueller et al 2012 [71] both explicitly note that direct measurement of a gravitational wave produced by a core-collapse supernova may yield insight into the conditions inside proto-neutron stars.

4.2.4 Technique

Emulating how the pipelines are intended to be used in the future, cWB2G was first run over the data set of waveforms injected into simulated Gaussian noise. The injections identified by cWB2G were used to generate statistics on the detectability of these waveforms via efficiency versus their signal-to-noise ratio.

All detected waveforms were reconstructed cWB2G, and a subset of detected events that spanned a broad range of SNR was chosen for comparative analysis by BayesWave.

For detected injections, the key measure of waveform recovery performance is the overlap, α , between the injected and recovered waveforms. For all detected events, the time-optimized, noise-weighted, normalized overlap between the injected waveform and recovered waveform was calculated as:

$$\alpha = \frac{\langle \psi_{rec} | \psi_{inj} \rangle}{\sqrt{\langle \psi_{rec} | \psi_{rec} \rangle \langle \psi_{inj} | \psi_{inj} \rangle}} \quad (4.9)$$

where the inner product between waveforms is weighted by the power spectral density of the data, $S(f)$ as:

$$\langle \psi_a | \psi_b \rangle = \int \frac{\psi_a^*(f) \psi_b(f)}{S(f)} df \quad (4.10)$$

and α is normalized to be 1 for a perfect match between injected and recovered waveforms, and 0 for no match. α is also optimized over time to ensure that the recovered and injected waveforms are lined up optimally.

For cWB2G, the overlap between recovered and injected waveforms was calculated for each detected event using the pyCBC software package to optimize α over different timing offsets applied to the recovered and injected waveforms [66]. The overlap was calculated separately for each interferometer for every detected CCSN injection.

For BayesWave, the overlap of the recovered waveform and 90% confidence interval based on the posterior distribution was reported for the each analyzed event; a small subset of those analyzed by cWB2G.

4.3 Waveform reconstruction algorithm tuning

Before performing the full analysis, both cWB2G and BayesWave were tuned to best resolve CCSN. This preparatory study examined performance using all the considered waveform families and varied the algorithm parameters most likely to impact accurate waveform recovery of CCSN signals.

4.3.1 cWB2G configuration

The major parameters tuned for cWB2G recovery of CCSN were:

- Multi-resolution configuration vs. single best resolution
- Time and frequency clustering

Multi-resolution configuration vs. single best resolution

One feature new to cWB2G from cWB1G is the multi-resolution configuration, which takes the most significant tile from each analyzed time-frequency plane and projects it onto a single time-frequency plane for reconstruction. The single best resolution configuration reconstructs the waveform based on the single time-frequency plane that produces an event with the highest coherent significance, as defined in Equation 4.2.

Burst events with signal power that is best reconstructed by wavelets of different shapes are expected to be better resolved with the multi-resolution configuration. However, if most of the power is best represented by multiple tiles on the same time-frequency plane, say for a waveform with strong wavelet-like features, then using the single best resolution configuration should better resolve the event.

Both the multi-resolution (MR) and single-best resolution (SR) configurations were tested for each waveform family by considering the relative efficiency, waveform reconstruction measured by α , and also event recovery in the frequency domain.

Dimmelmeier

For Dimmelmeier waveforms, the shortest and most single-wavelet-like of the considered CCSN models, it is reasonable to expect that the single-best resolution configuration of cWB2G may better resolve these events.

As a first pass test, cWB2G was run over a subset of Dimmelmeier events injected into simulated Gaussian noise colored to expected Advanced LIGO and Advanced Virgo sensitivities. Each recovered event was analyzed in both time and frequency.

Figure 4.6 shows a representative example. The top row shows the injected (in black) and recovered (in red) waveforms in the time domain, and the bottom row shows the same in the frequency domain. As shown for this particular event, the performance of multi-resolution (MR) and single best resolution (SR) performance was roughly equivalent in time or frequency for most of the subset of analyzed Dimmelmeier waveforms.

For the next test, the entire set of Dimmelmeier waveforms included in the 2007 supernova search was analyzed in both configurations. The MR and SR configuration performance was compared in both efficiency, shown in Figure 4.7 and also in waveform recovery measured by overlap, shown in Figure 4.8.

Figure 4.7 shows the cWB2G efficiency vs. network SNR for each Dimmelmeier rotating core-collapse model, which differ primarily by rotation rate. In the network SNR range of 5-12, all waveforms show an efficiency improvement using the multi-resolution configuration. The most rapidly rotating model, s15a3o15, had the largest

Tuning cWB2G for Dimmelmeier waveforms - s15a2o05

s15a2o05 waveforms

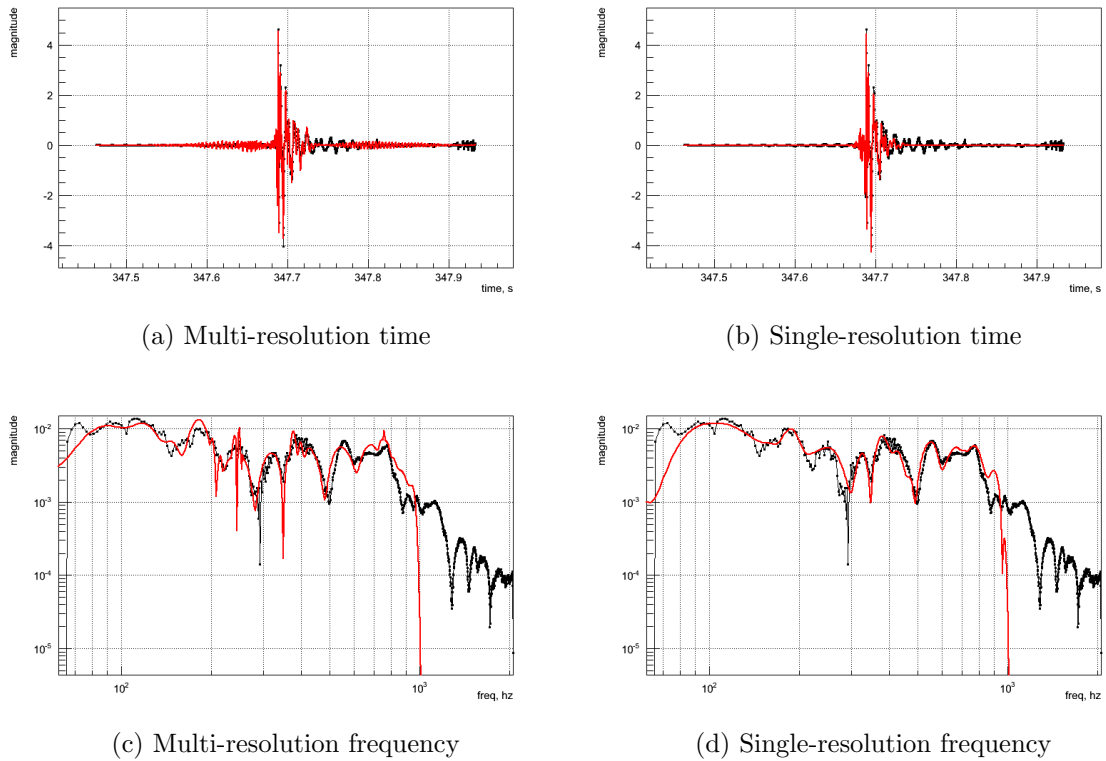




Figure 4.6: A series of plots illustrating the recovery of a Dimmelmeier waveform by cWB2G in time and frequency in both multiple-resolution (MR) and single best resolution (SR) configurations, using s15a2o05 as a representative example. The recovered waveform is plotted in red, and the injection shown as a black curve. The performance is very similar between MR and SR configurations for both time and frequency recovery plots, with each configuration capturing the event peak in time, and most of the frequency content (up to 1 kHz).

Tuning cWB2f for Dimmelmeier waveforms by efficiency



(a) Multi-resolution time

(b) Single-resolution time



(c) Multi-resolution frequency

Figure 4.7: A series of plots showing the efficiency of cWB2G (with tuned clustering) vs. network SNR for each included Dimmelmeier model. In blue is cWB2G using the multi-resolution (MR) configuration, and in red is the single best resolution (SR) configuration. All Dimmelmeier waveforms show an efficiency improvement using the multi-resolution configuration in network SNR range of 5-12.

efficiency gain using MR configuration, up to a 20% improvement at a network SNR of 7.

Figure 4.8 shows a series of 2D histograms of the time-optimized overlap between injected and recovered waveforms, as defined in Equation 4.9, vs single interferometer SNR. The more slowly rotating models, s15a2o05 and s15a2o09, show clear improvements in waveform reconstruction using the single-best resolution configuration above a single interferometer SNR of 10. The rapidly rotating model s15a3o15 also shows a clear improvement above a single interferometer SNR of roughly 70.

Given that the waveform overlap is the figure of merit for the waveform reconstruction performance study and the multiple resolution cWB2G configuration appears to offer no clear benefit in reconstructing the waveform of identified events, the single best resolution configuration was chosen as the tuned configuration for Dimmelmeier events. However, it bears keeping in mind that the multiple resolution configuration could offer a significantly better efficiency for a supernova search in the key SNR range (5-12).

It would also be interesting to study the effect of transient noise on efficiency for wavelet-like Dimmelmeier waveforms using the different cWB2G resolution configurations. Given that a single resolution should capture most of the Dimmelmeier event power, the single best resolution configuration may act to filter unrelated interferometer noise best captured at a different resolution. This is left to future work.

Tuning cWB2f for Dimmelmeier waveforms by efficiency

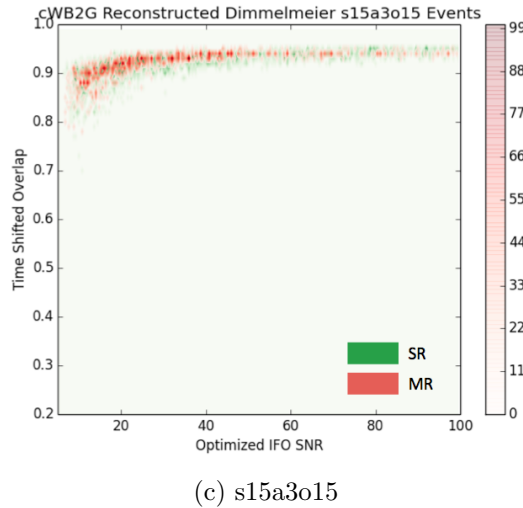
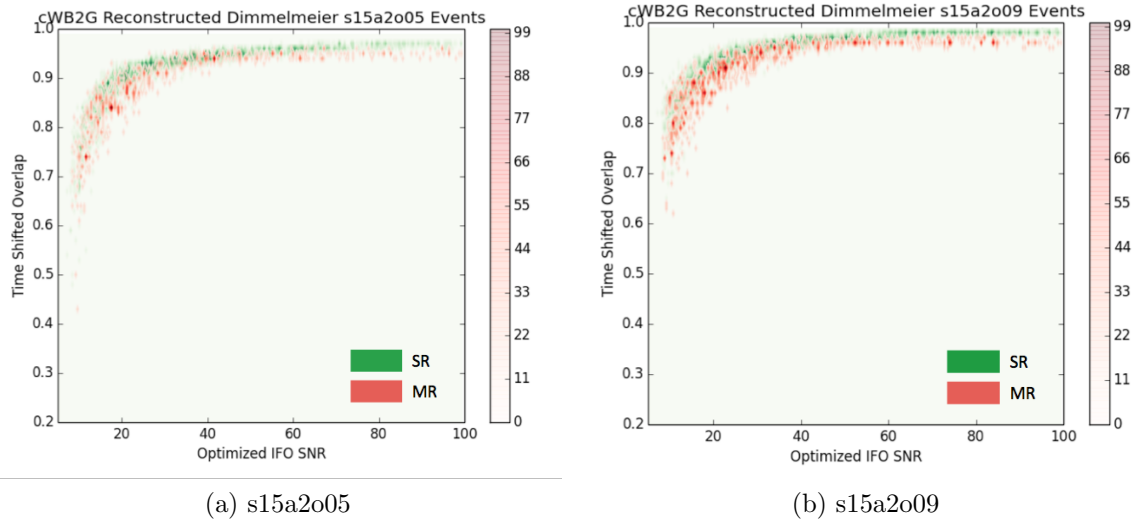


Figure 4.8: A series of 2D histograms showing the time-optimized overlap between injected and recovered waveforms vs single interferometer SNR. To optimize performance over the LIGO and Virgo interferometer network, only the overlap corresponding to the interferometer with the highest single interferometer SNR for that event is plotted. Results for the single best resolution cWB2G configuration are plotted in green, and in red for the multi-resolution configuration.

Tuning cWB2G for Yakunin waveforms

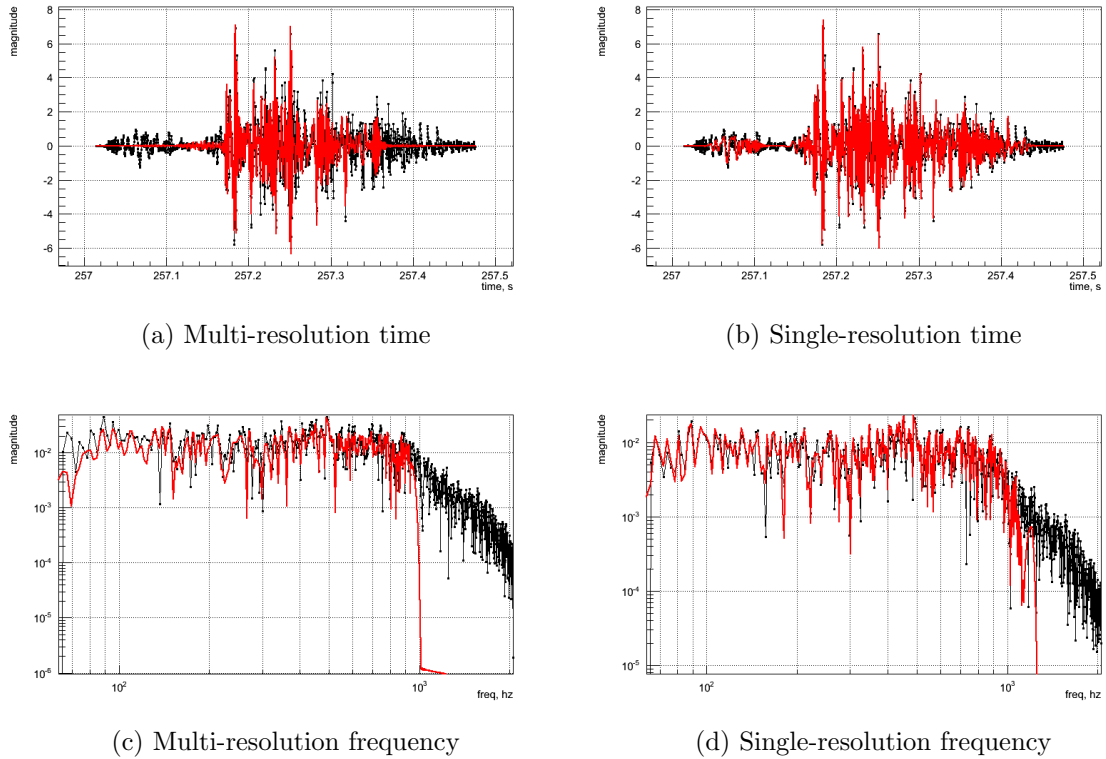


Figure 4.9: The recovery of a representative Yakunin waveform by cWB2G in time and frequency in both multiple-resolution (MR) and single best resolution (SR) configurations. The recovered waveform is plotted in red, and the injection shown as a black curve. The single best resolution configuration appears to better resolve the subtle waveform features in time, both before 257.1 seconds and after 257.3 seconds. The SR configuration also appears to capture more high frequency content of the waveform above 1kHz.

Yakunin

The same tuning procedure was used for Yakunin injections; cWB2G was first run over a small subset of Yakunin events and recovered events were analyzed in time and frequency.

Figure 4.9 shows a single Yakunin event analyzed by the multiple resolution configuration (left) and the single best resolution (right) with the waveform recovery shown

both in time (top) and frequency (bottom). For both configurations, the injected event is in black, and the recovered event is in red.

The single-resolution configuration clearly better captures the example event’s low amplitude features before 257.1 seconds and after 257.3 seconds, and also resolves more high frequency content than the multiple resolution configuration. This is representative of most Yakunin injections analyzed in the small scale tuning test.

Again following the same tuning procedure as for Dimmelmeier waveforms, the entire pool of Yakunin injections was then analyzed using both configurations. The performance of each configuration was compared in efficiency, shown in Figure 4.10 and recovered waveform overlap, shown in Figure 4.11.




Figure 4.10: The efficiency of cWB2G (with tuned clustering) vs. network SNR on Yakunin waveforms. In blue is cWB2G using the multi-resolution (MR) configuration, and in red is the single best resolution (SR) configuration. At a network SNR range of 5-12 the multi-resolution configuration is significantly more efficient in identifying injected Yakunin events, with up to a 10% gain at a network SNR of 7.

As seen for Dimmelmeier waveforms, the efficiency of cWB2G on Yakunin injections improves significantly for the critical network SNR range of 5-12 when using the multiple resolution configuration, shown in Figure 4.10. This result is counterintuitive to the small scale tests, which show a clear advantage in using the single best resolution to recover fine low amplitude features and high frequency content of Yakunin waveforms.




Figure 4.11: Heatmap comparison of cWB2G reconstruction of Yakunin waveforms vs. single interferometer SNR. Only the overlap corresponding to the highest single interferometer SNR in the LIGO and Virgo network is shown for each event. The cWB2G single-best resolution configuration is shown in red and multi-resolution in green.

The results for waveform recovery shown in Figure 4.11 are as expected from initial small scale test results. Below a single interferometer SNR of 40, the single best resolution configuration shows a slight improvement over the multiple resolution configuration in overlap performance. Above a single interferometer SNR of 40, this performance gain increases significantly, growing to 10% above an SNR of 80.

Since the multiple resolution configuration was not found to improve overlap performance for identified Yakunin events, the single best resolution configuration was chosen as the optimal tuned configuration. But again, as for Dimmelmeier events, it is worth noting that the multi-resolution configuration does show a notable improvement in efficiency at a critical network SNR range of 5-12.

Mueller

Following the initial tuning procedure used for Dimmelmeier and Yakunin injections, cWB2G was run over a subset of Mueller events and recovered events were analyzed in time and frequency. A series of plots showing the recovery of a representative waveform for each progenitor star model is shown in Figures 4.12 - 4.14.

Unlike for the Dimmelmeier and Yakunin waveforms, the multiple-resolution cWB2G configuration clearly better recovers Mueller waveforms for all three progenitor stars. The specific advantages for each progenitor star model are discussed in the caption of the corresponding figure.

The multiple resolution (MR) configuration was chosen for cWB2G analysis of Mueller injections for the CCSN waveform reconstruction study.

Tuning cWB2G for Mueller waveforms - I

W15 waveforms

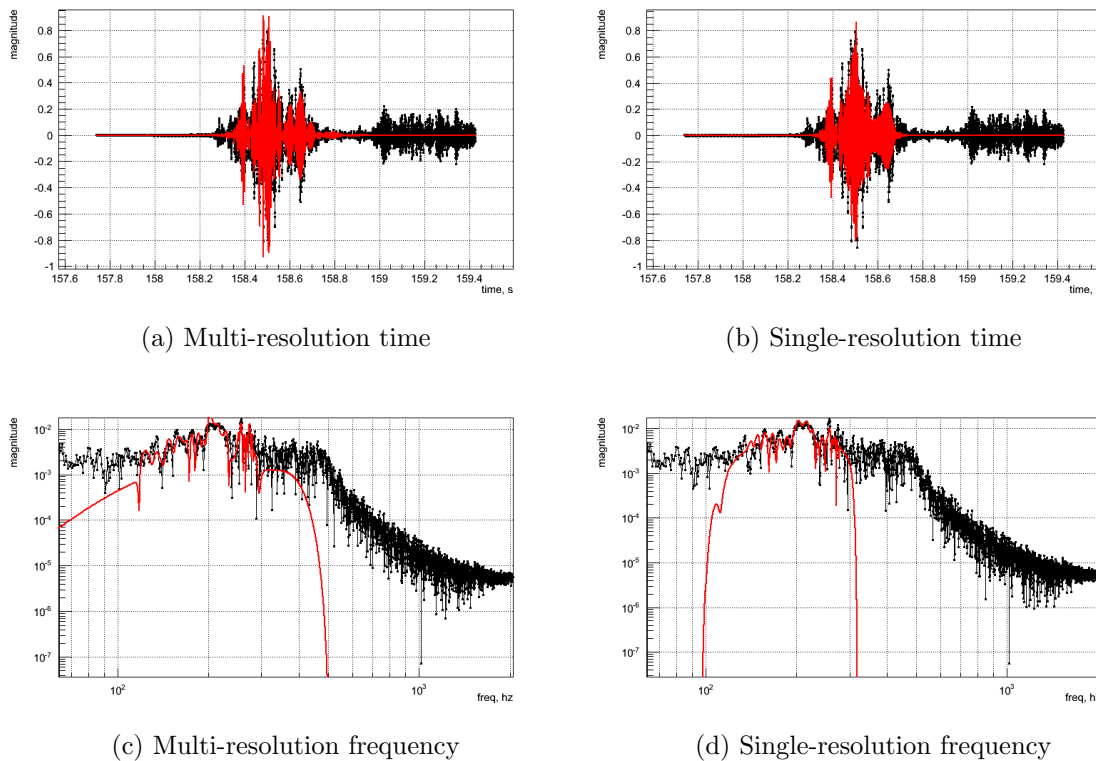


Figure 4.12: The recovery of a representative W15 progenitor waveform by cWB2G in time and frequency in both the multiple-resolution (MR) and single best resolution (SR) configurations. The recovered waveform is plotted in red, and the injection shown as a black curve. The MR configuration clearly more accurately resolves more of the fine waveform features than the SR configuration, especially around 158.6 seconds. The MR configuration also recovers more low and high frequency waveform content, both below 100 Hz and above 300 Hz. Note that neither configuration captures the late burst of signal power after 159 seconds.

Tuning cWB2G for Mueller waveforms - II

L15 waveforms

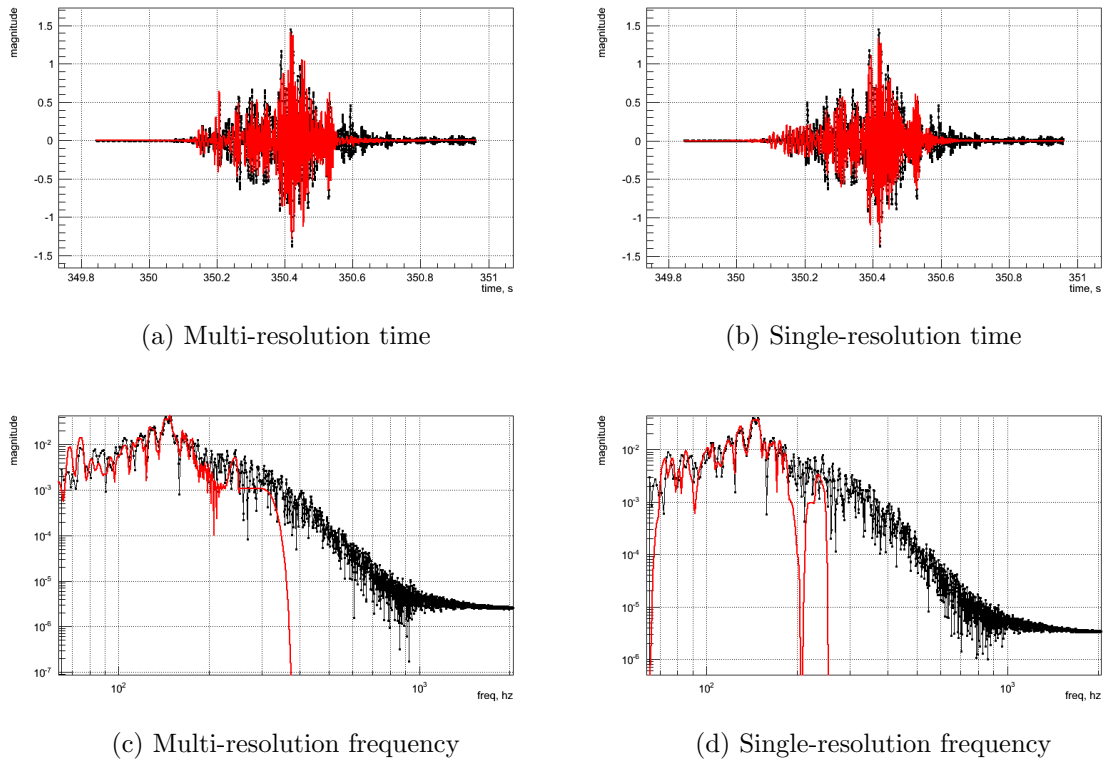


Figure 4.13: The recovery of a representative L15 progenitor waveform by cWB2G in time and frequency in both the multiple-resolution (MR) and single best resolution (SR) configurations. The recovered waveform is plotted in red, and the injection shown as a black curve. Both configurations appear to resolve features in time at about the same performance. However, as seen for the W15 progenitor, the MR configuration does recover more significantly more frequency content both below 70 Hz and above 300 Hz.

Tuning cWB2G for Mueller waveforms - III

N20 waveforms

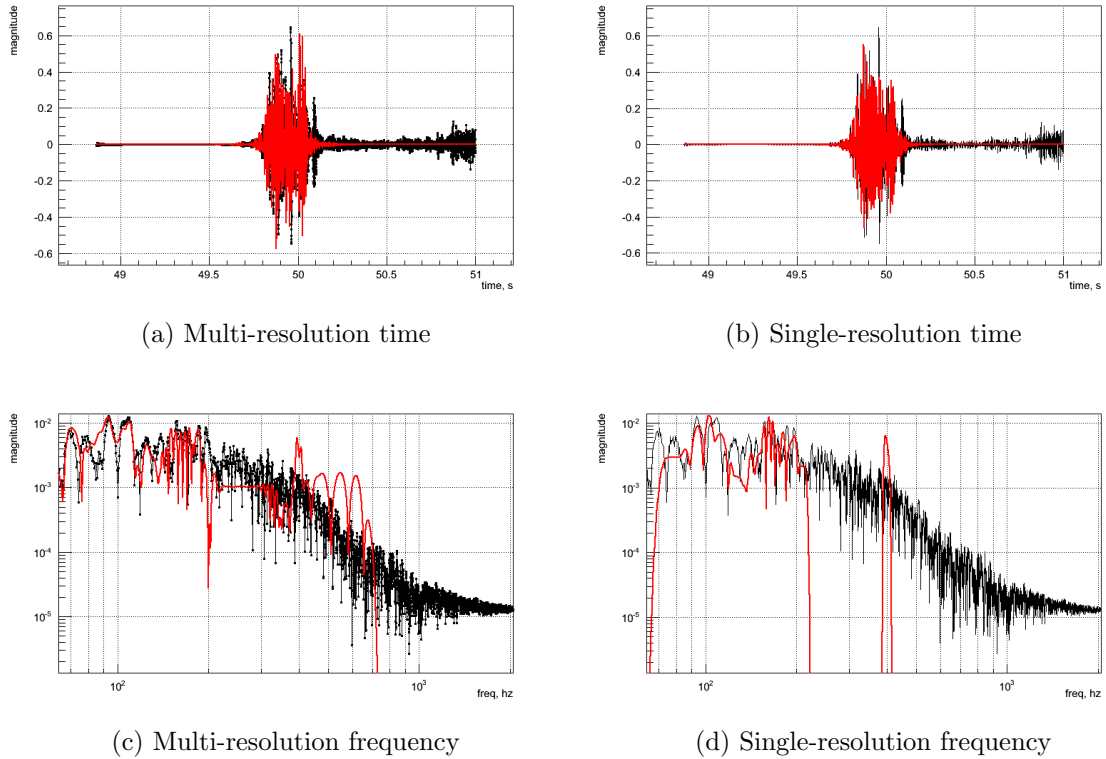


Figure 4.14: The recovery of a representative N20 progenitor waveform by cWB2G in time and frequency using both the multiple-resolution (MR) and single best resolution (SR) configurations. The recovered waveform is plotted in red, and the injection shown as a black curve. The MR and SR configuration appear to resolve the waveforms similarly in time. However, the MR configuration clearly recovers more frequency content both below 70 Hz, between 200 and 400 Hz, and above 400 Hz. Note that neither configuration is able to resolve the burst of signal power after 50.5 seconds.

Clustering in time and frequency

To avoid cWB2G resolving long CCSN waveforms containing multiple bursts of signal as multiple separate events, the clustering settings were tuned. The time window was expanded from the default value of 0 seconds to 5 seconds and the frequency window was expanded from the default 0 Hz to 400 Hz. These values were based on a conservative estimate to ensure the inclusion of the full power of Mueller waveforms.

Although this clustering window expansion improves waveform overlap performance, it also introduces an opportunity for transient noise pollution that would otherwise be rejected by requiring time-frequency overlap for events. The influence of non-Gaussian transient noise on efficiency and waveform reconstruction is discussed in detail in Chapter 6.

It is left to future work to improve this clustering window tuning to balance the influence of noise and enable cWB2G to capture all power in longer CCSN models with multiple discrete bursts.

4.3.2 BayesWave configuration

The default BayesWave configuration was also tuned for optimal recovery of CCSN waveforms. The number of wavelets included in the wavelet distribution was increased to a generous 125. This adversely impacts computational efficiency, but also ensures a large number of wavelets are available to fit the long and complex CCSN injections, if needed. In the future, different configurations using only the maximum necessary number of wavelets for the target CCSN model could be employed.

The analyzed time window was increased from the default to 5.0 seconds in order to ensure the full event duration was captured for all waveforms. This also impacted

computational efficiency, and it is possible this may be softened for shorter CCSN models in the future.

Additionally, BayesLine was not used for PSD estimation for BayesWave runs on simulated Gaussian noise. The input PSD used for coloring the Gaussian noise was used instead for the Gaussian noise estimation model.

In the future, the BayesWave configuration will be further tuned to recover CCSN events by altering the clustering prior. The clustering prior currently favors concentrated clusters of power over isolated wavelets, and this could be tuned to favor discrete islands of power akin to what is expected from CCSN models.

4.4 Results

4.4.1 Efficiency

For each waveform family, the chosen tuned configuration of cWB2G was used to detect modeled signals injected into Gaussian noise, as described in Section 4.2.3. Efficiency in detection of CCSN events varied greatly depending on the model, as shown in Figure 4.15.

The efficiency for all CCSN models approaches 100% by a network SNR of 20. At lower network SNRs, the wavelet-like Dimmelmeier waveforms are the most detectable, by nearly 40% more than Mueller waveforms at a network SNR of 8. Below a network SNR of 8, the long, complex Mueller waveforms are the least detectable at roughly 10% efficiency at a network SNR of 7. Above a network SNR of 8, the Yakunin waveforms are least detectable, at roughly 68% efficiency at a network SNR of 10.






Figure 4.15: A plot showing the cWB2G efficiency vs. network SNR for a representative set of included CCSN models. The L15 progenitor represents efficiency for Mueller waveforms in red, the rapidly spinning model represents Dimmelmeier injections in blue, and Yakunin injections are shown in yellow. Note the efficiency for all CCSN models converges to 100%. The most wavelet-like model plotted, Dimmelmeier s15a3o15, is significantly more detectable at network SNRs lower than 12.

The effective distance reach for each analyzed CCSN model is shown in Figure 4.16. All waveforms were discretely injected at 0.2, 0.4, 0.67, 1, 2, 4, and 10 kpc. No Mueller models were detected at a distance greater than 1 kpc. Most resolved Yakunin waveforms were detected only at a distance of 2 kpc or less. Some Dimmelmeier injections were detected out to 10kpc, with most detected at a distance of 4 kpc or less.


Note that efficiency is expected to decrease for all models with the addition of transient noise, as shown in Chapter 6.



(a) Dimmelmeier



(b) Yakunin



(c) Mueller

Figure 4.16: A series of histograms showing the number of injections detected by cWB2G vs. injected distance. All CCSN models were injected at discrete distances of 0.2, 0.4, 0.67, 1, 2, 4, and 10 kpc. Note there are far fewer Yakunin events because these were contained in the same injection set as all three Dimmelmeier models (and another waveform that was not included in the study) each spaced 100 +/- 10 seconds apart. There are roughly the same number of analyzed Yakunin injections as each of the three Dimmelmeier models.

4.4.2 Waveform reconstruction

The waveform reconstruction results are presented in two parts: trends in waveform family reconstruction, and a comparison between cWB2G and BayesWave performance. The former provides insight on how well each CCSN model is able to be recovered based on general gravitational wave features. The latter identifies strengths and weaknesses between each algorithm in recovering gravitational waves produced by CCSN, and should help guide pipeline development and CCSN follow-up efforts during the Advanced detector era.

4.4.2.1 Trends in waveform reconstruction performance by CCSN model

Two key metrics in quantifying waveform reconstruction performance as evaluated by the waveform overlap defined in Equation 4.9 are the maximum achievable overlap and the average overlap achieved at a realistically detectable SNR of 8.

Maximum achievable overlap is the highest overlap between injected and recovered waveforms observed for a given model at arbitrarily large SNR values. This was calculated by taking the average overlap value of L1 and H1 events greater than an SNR threshold that varied depending on the model¹².

The average overlap at an SNR of 8 was calculated by averaging the overlap of events within an SNR window of 8. For cWB2G events, this window was 2, and for BayesWave, with far fewer events, this window was 4.

¹²For cWB2G events, the maximum overlap was calculated using an SNR threshold of 80 for Dimmelmeier models, 60 for Yakunin and Mueller L15, 50 for Mueller W15, and 40 for Mueller N20. For BayesWave, slightly lower thresholds were used for some models because analyzed events only spanned an SNR range just beyond maximum overlap values. For BayesWave, the maximum overlap thresholds were set at 70 for Dimmelmeier waveforms, 50 for Yakunin and Mueller models L15 and W15, and 30 for Mueller N20.

Summary

Each waveform family, Dimmelmeier, Yakunin, and Mueller, shows a distinctive behavior in the key overlap metrics, as shown in Figure 4.17 and summarized in Table 4.3 for injections analyzed by cWB2G.




Figure 4.17: Summary heatmap of cWB2G overlap performance versus single interferometer SNR for CCSN injections detected and reconstructed by cWB2G. Dimmelmeier models are shown in green, Mueller in red, and Yakunin in blue. The legend also includes a reminder of the tuned cWB2G configurations employed (single best resolution, SR, or multi-resolution, MR). The colorscale indicates the number of reconstructed events per 2 dimensional overlap-SNR bin.

Figure 4.17 shows a summary heatmap for all CCSN models of time-optimized overlap between injected and recovered waveforms versus single interferometer SNR. Dimmelmeier and Mueller performance is shown for all waveforms in these families. These results were optimized over network interferometers L1, H1, and V1 (the two LIGO interferometers and the Virgo interferometer) in that only the overlap and corresponding SNR is plotted for the interferometer with the highest SNR for that injection. The color scale represents the number of recovered injections per 2D overlap-SNR bin. For

Summary of waveform reconstruction performance by model

	Dimmelmeier average	Yakunin	Mueller average
Maximum overlap	0.97	0.91	0.92
Overlap at SNR 8	0.78	0.42	0.57

Table 4.3: A table of the maximum achieved overlap and average overlap at an SNR of 8 for each waveform family, as shown in Figure 4.17. The Dimmelmeier average represents the average performance of all included Dimmelmeier models: s15a2o05, s15a2o09, and s15a3o15. Similarly, the Mueller average values are the average performance of the L15, W15, and N20 models. These reconstructions were produced by cWB2G.

full results for all individual CCSN models in a different colorscale scheme, see Appendix A.

Unsurprisingly, the wavelet-like 2D rotating core-collapse Dimmelmeier waveforms achieve the highest maximum overlap, and at the lowest SNR. As noted in Table 4.3, the recovered waveform match for Dimmelmeier models at an SNR of 8 is greater than the neutrino-driven models by 20% or more. The maximum achieved overlap for Dimmelmeier waveforms approaches close to 1 on average, the highest maximum overlap of the waveform families. Note also in Figure 4.17 three distinct green lines at high SNR - these are different maximum overlap values achieved between Dimmelmeier models discussed as *fine structure* features later in this section.

The longer, complex Mueller waveforms are significantly less well reconstructed relative to the Dimmelmeier waveforms over the entire range of recovered SNR. These models achieve a lower maximum overlap at much higher SNR relative to the Dimmelmeiers, not asymptoting to a maximum overlap of 92% until an SNR of 40.

Yakunin waveforms have the lowest overlap at SNRs less than 60, with only an overlap of 42% at an SNR of 8. This is likely because power is more evenly distributed in time for this model - there is no short distinctive region of peak amplitude. However, at

SNRs greater than 60 the Yakunin injection recovery approaches that of the Mueller models, reaching a maximum overlap of 91%.

Dimmelmeier fine structure

The differences in waveform reconstruction performance between the different Dimmelmeier waveforms were also explored, as summarized in Figure 4.18 and Table 4.4. As explained in Section 4.2.3, these models rotate at different rates and with different scales of differential rotation between the core and the outer layers of the star.




Figure 4.18: Summary heatmap of cWB2G overlap performance versus single interferometer SNR for Dimmelmeier injections detected and reconstructed by cWB2G. s15a2o05 models are shown in green, s15a2o09 in red, and s15a3o15 in blue. The colorscale indicates the number of reconstructed events per 2D overlap-SNR bin.

The Dimmelmeier events that were plotted as a waveform family in Figure 4.17 are shown by model in Figure 4.18. The slowly rotating model with moderate differential rotation, s15a2o05, is plotted in green. The moderately rotating model with moderate

Summary of Dimmelmeier waveform reconstruction performance

	s15a2o05	s15a2o09	s15a3o15
Maximum overlap	0.98	0.99	0.95
Ave overlap at SNR 8	0.71	0.76	0.86

Table 4.4: A table showing the maximum overlap and overlap at SNR 8 achieved for each Dimmelmeier model as recovered by cWB2G.

differential rotation, s15a2o09, is in red. The rapidly rotating model with strong differential rotation, s15a3o15, is in blue.

As shown in Table 4.4, the rapidly rotating model has the lowest maximum overlap, 95%, but achieves this at a significantly lower SNR than the more slowly rotating models. This model reaches a maximum overlap of 86% at an SNR of 8 - the highest average overlap at SNR=8 for any CCSN model analyzed by cWB2G in this study.

The more slowly rotating models perform very similarly to each other, as shown in Figure 4.18. The moderately rotating model achieves a slightly larger maximum overlap, the highest maximum overlap achieved by any CCSN model analyzed by cWB2G, and a 5% higher overlap at an SNR of 8 than the slowly rotating model.

Mueller fine structure

The differences in the reconstruction performance between the different Mueller models were also explored, as summarized in Figure 4.19 and Table 4.5. As explained in Section 4.2.3, each set of Mueller waveforms, L15, W15, and N20, is based on a different progenitor star.

A 2D histogram comparing the recovered waveform overlap vs. SNR by model is shown in Figure 4.18. Reconstructed waveforms corresponding to progenitor star L15 are shown in green, progenitor W15 in red, and progenitor N20 in blue.




Figure 4.19: Summary heatmap of cWB2G overlap performance versus single interferometer SNR for Mueller injections detected and reconstructed by cWB2G. Models for the L15 progenitor star are shown in green, W15 in red, and N20 in blue. The colorscale indicates the number of reconstructed events per 2D overlap-SNR bin.

Summary of Mueller waveform reconstruction performance

	L15	W15	N20
Maximum overlap	0.92	0.89	0.89
Ave overlap at SNR 8	0.61	0.56	0.53

Table 4.5: A table showing the maximum overlap achieved and overlap at SNR 8 for each Mueller progenitor star. Injections were analyzed by cWB2G.

As shown in Table 4.5, of the Mueller injections, the L15 progenitor models achieve the highest maximum overlap, at 92%, and overlap at SNR 8, at 61%. This is expected, as the L15 models are the shortest Mueller waveforms, with nearly all of the produced event power contained within 0.5 seconds.

The W15 and N20 models perform very similarly, each having a delayed burst of power toward the end of the waveform that is difficult for cWB2G to resolve (and as we will see later, that BayesWave also struggles to resolve). Explicit examples of this effect can be seen in the cWB2G tuning investigation detailed in Section 4.3.

However, both the W15 and N20 models achieve an overlap of more than 50% at an SNR of 8, and approach 90% at maximum overlap.

4.4.2.2 Comparison of cWB2G and BayesWave performance

A subset of injections for each CCSN model was selected for analysis by BayesWave to compare its waveform reconstruction performance with cWB2G, as shown in Figures 4.20 - 4.22 and summarized in Table 4.7.

Table 4.7 shows the familiar quantities maximum overlap and overlap at SNR=8, reprinted here for cWB2G results in order to easily compare with the corresponding BayesWave values. These quantities are calculated for BayesWave in the same way as described in Section 4.4.2.1, using the average of the waveforms reconstructed as in the L1 and H1 interferometers.

The key summary of Table 4.7 is that BayesWave achieves a higher or equivalent maximum overlap than cWB2G for all CCSN models, but tends to return a lower waveform overlap at lower SNR values for the complex neutrino-driven CCSN waveforms. Results are discussed in detail by waveform family below.

To simplify the algorithm comparison for each considered CCSN model, the following figures show results for the recovered waveforms as reconstructed in the Livingston LIGO (L1) interferometer by BayesWave and cWB2G¹³. The full BayesWave results for the selected injections can be found in Appendix A.

Events analyzed by BayesWave return a recovered value and a confidence interval for the overlap between the reconstructed and injected waveforms based on the calculated posterior distribution. In the following plots, the recovered overlap value of

¹³Note this may make the values reported in Table 4.7 appear slightly off of what is seen in the plots since the table averages results for the waveform as recovered in both LIGO interferometers.

cWB2G			BayesWave		
Model	Maximum overlap	Ave overlap at SNR 8	Model	Maximum overlap	Ave overlap at SNR 8
Ave Dim	0.97	0.78	Ave Dim	0.99	0.84
s15a2o05	0.98	0.71	s15a2o05	0.99	0.79
s15a2o09	0.99	0.76	s15a2o09	0.99	0.84
s15a3o15	0.95	0.86	s15a3o15	0.99	0.94
Yakunin	0.91	0.42	Yakunin	0.91	0.00
Ave Mueller	0.92	0.57	Ave Mueller	0.96	0.49
L15	0.92	0.61	L15	0.97	0.66
W15	0.89	0.56	W15	0.96	0.66
N20	0.89	0.53	N20	0.92	0.24

Table 4.6: The maximum overlap and overlap at SNR=8 is restated here for injections analyzed by cWB2G. This table contains an average of all injections as recovered in the L1 and H1 interferometers.

Table 4.7: The maximum overlap and overlap at SNR=8 for injections analyzed by BayesWave. This table contains an average of all injections as recovered in the L1 and H1 interferometers.

each injection analyzed by BayesWave is represented with a circle, and bounded by triangles indicating the 90% confidence interval of the result.

For comparison with BayesWave, the cWB2G overlap vs. SNR distributions presented in Section 4.4.2.1 were translated to quantities analogous to the results produced by BayesWave under the assumptions outlined in Section 4.2.2. The overlap vs. SNR distribution for each model analyzed by cWB2G was binned by SNR, and the overlap values contained in each SNR bin were fit to a Gaussian curve. The mean of the Gaussian fit and bounds containing 90% of the Gaussian distribution are reported for each CCSN model for comparison with BayesWave.

Dimmelmeier

Figure 4.20 shows a comparison of cWB2G and BayesWave for each 2D rotating core-collapse Dimmelmeier model. The BayesWave recovered values and 90% confidence

Comparison of waveform reconstruction algorithms - Dimmelmeier





Figure 4.20: A comparison of cWB2G and BayesWave waveform overlap vs. L1 interferometer SNR for each analyzed Dimmelmeier waveform. BayesWave recovered overlap values and 90% confidence intervals are shown in red, and the cWB2G mean overlap and 90% distribution are in blue.

Comparison of waveform reconstruction algorithms - Yakunin




Figure 4.21: A comparison of cWB2G and BayesWave waveform overlap vs. L1 interferometer SNR for Yakunin injections. BayesWave recovered overlap values and 90% confidence intervals are shown in red, and the cWB2G mean overlap and 90% distribution are in blue.

intervals are shown in red, and the cWB2G mean overlap and 90% distribution limits are in blue.

BayesWave has the most clear advantage over cWB2G for the rapidly rotating Dimmelmeier model s15a3o15, more than any other considered CCSN model. For this model BayesWave produces an overlap of at least 5% greater than cWB2G at all SNR ranges, including below an SNR of 20.

BayesWave also consistently achieves equal or more effective waveform reconstruction performance than cWB2G for the other two less rapidly rotating Dimmelmeier models.

Yakunin

Figure 4.20 shows a comparison of cWB2G and BayesWave for 2D neutrino-driven explosion Yakunin injections.

Below an SNR of 12, BayesWave produces a reconstructed waveform with 0% overlap with the Yakunin injection¹⁴. This is also observed for the waveforms as recovered in the H1 interferometer at single interferometer SNRs of less than 12.

This is a surprising result, as an SNR of 10 is fairly loud, and should be resolvable. For reference, cWB2G reconstructs Yakunin waveforms with an average overlap of 50% at SNR=10.

At higher SNRs BayesWave matches, but does not exceed, cWB2G in overlap performance for analyzed Yakunin injections.

Mueller

Figure 4.22 shows a comparison of cWB2G and BayesWave for each progenitor star of 3D neutrino-driven explosion Mueller CCSN models.

The relative reconstruction of the N20 model, the progenitor with the largest time difference between bursts of signal energy, is fairly similar to the comparative performance for Yakunin injections. BayesWave recovers waveforms with 0% overlap at SNRs less than 8, but then matches cWB2G by an SNR of 20 and eventually exceeds cWB2G performance with a slightly higher maximum achievable overlap.

Comparative performance between BayesWave and cWB2G is fairly similar for progenitors W15 and L15. Recall from Figure 4.5 the W15 progenitor produces a gravitational wave signature with a smaller time difference between discrete bursts signal and the L15 progenitor produces only a single discrete burst. For these two Mueller models BayesWave does comparably well with cWB2G at lower SNRs, and then exceeds cWB2G at higher SNRs with a higher maximum overlap, listed in Table 4.7.

¹⁴With 90% confidence intervals of 20% and 40% overlap, for the L1 interferometer.

Comparison of waveform reconstruction algorithms - Mueller

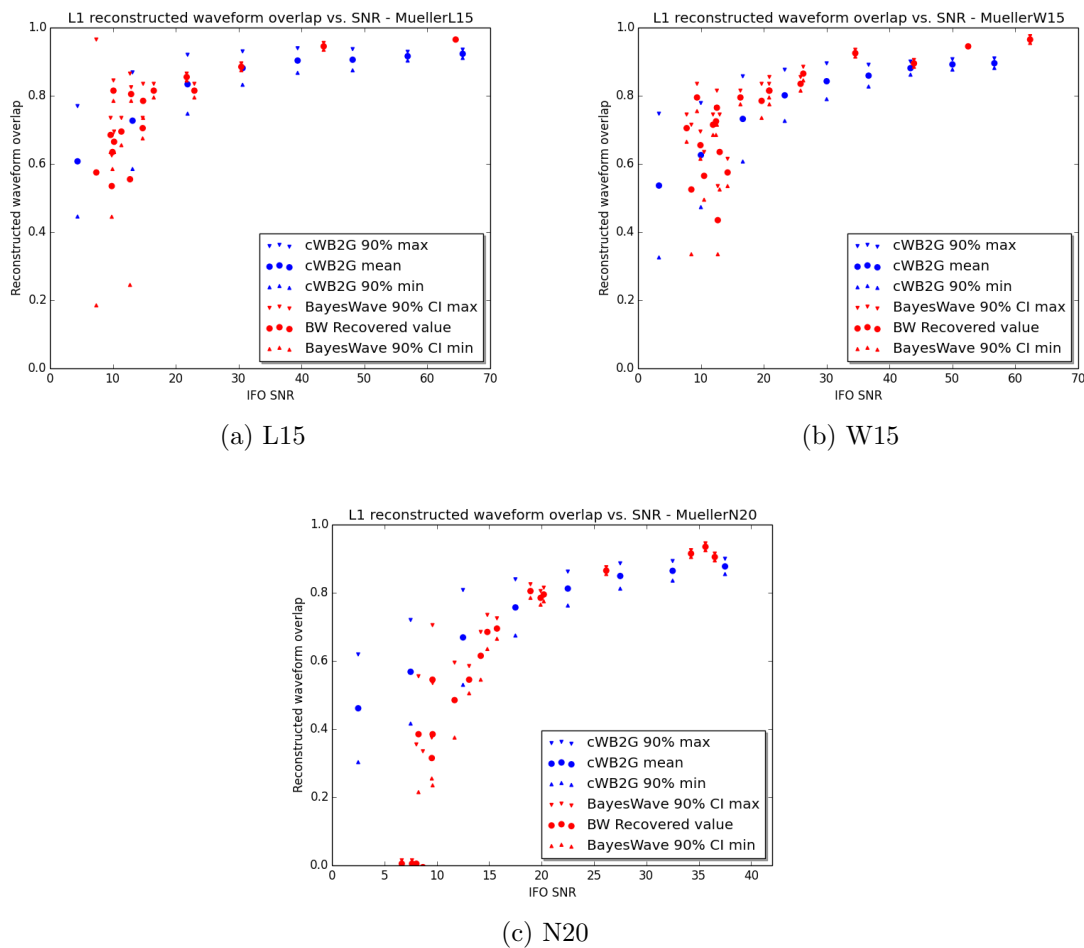


Figure 4.22: A comparison of cWB2G and BayesWave waveform overlap vs. L1 interferometer SNR for each Mueller progenitor star. BayesWave recovered overlap values and 90% confidence intervals are shown in red, and the cWB2G mean overlap and 90% distribution are in blue.

4.4.3 Summary and future work

The efficiency and waveform reconstruction of simple 2D rotating core-collapse supernovae waveforms is good, approaching 1 in efficiency and waveform overlap by both algorithms at SNRs of 20 or less.

However, for the more complex neutrino-driven CCSN models, efficiency is low at realistic SNRs; roughly 50% at a network SNR of 10. Additionally, the current ability of burst waveform reconstruction algorithms to recover these waveforms at realistic single interferometer SNRs is poor, in the range of 50% at a single interferometer SNR of 8 as reconstructed by cWB2G. For 3D neutrino-driven models, astrophysically interesting lower amplitude gravitational wave bursts emitted after the highest amplitude part of the signal are completely unrecovered by both algorithms at realistic SNRs.

It is also interesting that BayesWave, an algorithm with more flexibility in wavelet use and placement when fitting a signal, tends to outperform a time-frequency tiling algorithm (cWB2G) in reconstructing simpler wavelet-like waveforms, but not necessarily for more complex waveforms, such as the neutrino-driven CCSN models.

cWB2G tends to more reliably reconstruct the more complex waveforms at lower, more astrophysical realistic SNRs. This result should help guide CCSN event follow-up efforts and future tuning for the BayesWave algorithm.

It would be very useful as future work to repeat this study using more recent CCSN models that were either omitted out of the convenience of using existing injection data or have been published since the start of this study. Additional suggested models for future studies are Ott 2013 [74], Scheidegger 2010 [84] and Mueller 2013 [70].

Additionally, to repeat or build on this work the reader may find the relevant code detailed in [66].

CHAPTER 5

CHARACTERIZING AND MITIGATING TERRESTRIAL NOISE SOURCES

5.1 Terrestrial noise: impact on detectability, parameter estimation, and waveform reconstruction

Searches for generic GW transients, or *bursts*, in data from previous LIGO science runs made new and astrophysically meaningful upper limit statements ([13], [14], [15]), but their reach was limited by non-Gaussian transient noise. Figure 5.1 shows single interferometer burst pipeline (ETG) triggers generated for the Hanford LIGO interferometer during the first epoch of the previous science run (S6) in colored histograms after different levels of data quality cuts are applied. The gray histogram shows background events from the coherent unmodeled all sky burst search algorithm cWB after all data quality cuts. Comparing the distribution of events to the Gaussian noise expectation, shown by the black trace, there is a long tail of background events that inhibits statistical confidence in a true signal of a similar significance.

Essentially, chance coherent terrestrially-induced excess power noise in multiple detectors produces false candidate events in the detection pipelines. The projected rate at which we would expect to see false alarms versus the significance of the event is calculated using time slides, where real gravitational wave strain channel data is analyzed with interferometers offset artificially in time by some interval longer than it would take a gravitational wave to travel between detectors. Any events identified

by the detection pipeline must then necessarily be due to noise. An analysis of the science run data over many time slides produces *background triggers* used to generate false alarm rate statistics that candidate events (events produced with no time offset) are compared to.

A *tail* of background events, or a non-Gaussian distribution spanning up to large significance, has a detrimental impact on the sensitivity of the burst search in particular. The greater the tail, the closer or more energetic the astrophysical event must be to produce a gravitational wave signal that has a low false alarm rate.

In Figure 5.1, triggers from the first epoch of LIGO's previous science run (S6A) produced by the ETG Omega, which identifies excess power events in a single channel, are shown in color (red, green, blue) after various stages of data quality vetoes are applied, as described in Section 5.2.1. The histogram of events in red shows triggers for the Hanford interferometer left after all data quality cuts were applied. A long tail of triggers at a high signal-to-noise ratio remains. Similarly, the all-sky burst search triggers for this period, produced by cWB, are shown in gray. There are significantly fewer events relative to Omega triggers due to the multiple-interferometer time coincidence and coherence requirements, but there is still a significant tail remaining after all data quality cuts relative to the ideal Gaussian behavior of an ideal detector - shown by the solid black line.

Notably, unlike matched filter searches for binary gravitational wave sources which were most affected by very loud events that ring up the template bank and pollute chunks of data as long as the longest template, Figure 5.1 shows that the coherent burst searches were more affected by a high glitch rate of relatively low, 8-20 SNR events.





Figure 5.1: A histogram of burst event trigger rate during the first epoch of LIGO’s sixth science run (S6A). The colored (red, green, blue) histogram shows triggers generated by the ETG Omega after different levels of data quality vetoes have been applied. In gray is a rate histogram of background triggers of the all-sky coherent waveburst search. The expected trigger rate due to Gaussian noise is shown as a black trace.

Even for transient searches that are well modeled, transient noise is a critical issue. A striking example comes from the NINJA-2 project, which studied the ability to detect and reconstruct the parameters of binary black holes in Advanced LIGO and Advanced Virgo data [7]. Post-Newtonian and numerical hybrid waveforms were injected into real data from previous LIGO and Virgo science runs (S6 and VSR2) recolored to have the same spectral density as projected curves for early Advanced LIGO and Advanced Virgo (circa 2015-2016). This project gave insight into how the parameter estimation pipelines would perform when the waveforms were injected into realistic non-Gaussian data.

A key finding of the study is well illustrated in Figure 5.2 which shows details of a triple-interferometer (LIGO Livingston, LIGO Hanford, and Virgo) coincidence event. This event was identified by the templated search for high mass coalescences with a false alarm rate of one in 6000 years, and also by the unmodelled all-sky burst search cWB with a false alarm rate of one in 40 years. Part (a) of Figure 5.2 shows a



(a) A spectrogram of a transient noise (glitch) event in Hanford recolor noise at the same time that an event was injected into the data (across multiple interferometers).

(b) The 95% credible region for mass estimation based on EOBNRv2 analysis using realistic recolor noise is shown in blue. Analysis of the same event injected into Gaussian noise is shown in dashed pink. The actual injection parameters are marked with the red X.

Figure 5.2: A key figure from the NINJA2 analysis that demonstrates the negative impact of transient noise on extracting physical parameters of modeled signals [7].

spectrogram of a transient noise event in the Hanford interferometer that happened to overlap in time with this injected event. Part (b) of Figure 5.2 shows the output of the EOBNRv2 parameter estimation analysis of the event, with the actual injected mass parameters marked as as ‘X’. The 95% credible region of parameter space when the injection is analyzed in the realistic non-Gaussian data is shown by the solid blue line, and an analysis of the event when injected into Gaussian data is shown in dashed pink. A transient noise event in one interferometer badly skews the parameter estimation, even for a modeled search.

Evidently, parameter estimation for all transient gravitational wave searches depends critically not on the inherent signal strength of an astrophysical event, but the ratio of the signal to the noise in all interferometers, independently, at the time of the event. Good data quality, in particular a low rate of transient noise, is crucial to parameter estimation efforts for all transient gravitational wave searches.

In particular, characterizing and mitigating sources of terrestrial transient noise in the Advanced Era gravitational wave interferometers will likely be vital to confidently identifying and accurately reconstructing any transient signal, including a galactic core-collapse supernova.

5.2 Combating transient noise

Ideally, it is preferable to identify noise sources that pollute astrophysical search performance and improve the hardware to clean the data. However, this is difficult or impossible if the input noise cannot be controlled, as is the case for ground motion, or if the cause of the noise is unidentified. For known, unmitigable noise sources, affected data can be *vetoed* from the astrophysical search analyses. Although efficiently vetoing data associated with known noise couplings does improve the statistical significance of candidate gravitational wave events, this outcome is less desirable than mitigation via hardware improvements since it reduces the overall time available for analysis.

5.2.1 Data quality vetoes

Noise transients targeted for mitigation via hardware improvement or vetoing are those that most impact the astrophysical search background, generally far outside Gaussian noise behavior. However, due to the tremendous complexity of the interferometers, even noise couplings that produce extremely loud transient behavior in the gravitational wave strain channel can be difficult to diagnose. The following abbreviated set of examples gives an idea of the wide, dynamic range of potential noise sources: seasonal ocean waves that induce excess ground motion, electromagnetic signal coupling between nearby wires, loose wire contacts that cause shorts or signal spikes, flaws in optical components that cause light scattering with even very small

optic angular motion, or poor whitening of an in-loop signal used for optic cavity resonance control.

Several tools and techniques have been developed to aid in the challenge of detangling the multitude of interferometer signals and diagnose noise sources. All look for correlation between the noise event in the gravitational wave strain channel and the interferometer auxiliary channels which measure the behavior of the interferometer's subsystems, optic cavities, and local environment.

One such tool is *Omegascan*¹, which produces a set of spectrograms for the gravitational wave strain channel and any auxiliary channels with excess power that exceeds a variable threshold. These scans can be used to correlate transient noise events with auxiliary channels based on timing, frequency, and even morphology, or the shape of the glitch in time and frequency.

Algorithms such as Hierarchical Veto [89], Used Percentage Veto (UPV) [51], and Ordered Veto List (OVL) [37], produce statistical results ranking the significance of auxiliary channels that produce noise transients coincident with transients in the gravitational wave strain channel. These tools suggest which channels or families of channels are the most closely associated with noise transients during an analyzed time. Note that all of these statistical tools depend on the performance of ETGs in identifying transient events and accurately resolving their timing, frequency, and loudness. The current performance of ETGs is summarized in Chapter 3.

Many other detector characterization tools are used for noise diagnostics. Some correlate noise based on coherence between channels, others characterize noise line behavior. A growing subfield is the classification of glitches in the gravitational wave strain

¹Omegascan uses a sine-Gaussian basis to produce the spectrograms, detailed in Chapter 3. Similar scans produced by Omicron, called Omiscans, are also used.

channel based on morphology [77]. Information gleaned from these tools often jointly leads investigations to noise mitigation solutions or data quality flag development.

Key noise investigations in the prior LIGO science run are detailed in [3]. There have also been many successes in identifying and mitigating noise in the Advanced LIGO detectors ahead of the first observing run, to be described in an upcoming paper.

Assuming identifying the noise source is possible, even when fixes to the instrument are made they may happen after weeks or months of polluted data. In this case, or when instrument improvement is impossible or impractical, vetoing is the only option to clean the data.

Data quality vetoes can take the form of segments of time during the occurrence of a noise phenomenon known or suspected to couple to the gravitational wave strain channel. For example, a data quality veto criteria might be to flag times that an auxiliary channel or set of auxiliary channels exceed some amplitude threshold, such as wind speed or multiple local magnetometers. Vetoes can also be defined based on the output of the gravitational wave strain channel itself. For example, an event loudness that exceeds an impossibly high astrophysical limit, or a time segment with a very high glitch rate that would make any surrounding candidate events suspect. The application of such strain-channel-based vetoes is not appropriate for all astrophysical analyses.

In addition to time segments², vetoes can also be applied to specific events based on the output of the statistical noise transient correlation tools hveto, UPV, and OVL.

For a given bank of produced data quality vetoes, each astrophysical search assigns categories to vetoes that govern how the vetoes will be applied to the analysis [88]:

²In S6 and prior science runs, data quality veto flags used integer second time segments.

- **Category 1** vetoes indicate a serious problem with the detector, or that the detector is not operating near its nominal sensitivity. This data is removed prior to running any gravitational wave analysis.
- **Category 2** vetoes flag times associated with a known problem that has been demonstrated to couple to the gravitational wave strain channel. These events are removed before the analysis determines the statistical significance of any candidate events, so added deadtime is also weighted heavily in considering which vetoes to include in category 2.
- **Category 3** vetoes indicate intervals coinciding with a noise source with an incompletely understood coupling mechanism, or a known noise source with high deadtime. Vetoes in category 3 are used to achieve a cleaner search background for improved statistics when setting upper limits, and may also cast doubt on any candidate events.

The vetoed time associated with each category, defined individually for each search, is then hierarchically removed from analysis time when the detectors are resonating at their nominal operating point. Category 1 time is considered unusable. Removing category 2 vetoes allows the searches to cast a wide net in the search for potential candidate events, but still remove data that is clearly suspect. Removing category 3 vetoes enables the searches to significantly clean the signal background for setting upper limits, but sometimes at the cost of a high deadtime. If searches require long or continuous data segments, using many data quality vetoes comes at an even greater cost of total analysis time.

Data quality veto performance

The performance of a data quality veto is measured by its effectiveness in removing noise transients while incurring minimal downtime, and also its safety, or likelihood of vetoing a real gravitational wave event.

To characterize performance, an ETG³ is used to produce gravitational wave strain channel event triggers for the entire duration the interferometer was locked at its nominal operating point. The data quality veto is then applied to the event triggers, and its performance is gauged by the number of glitches removed, the total downtime, and the number of hardware injections vetoed, as described below.

The figure of merit used to evaluate DQ veto effectiveness is the ratio:

$$\kappa = \frac{\epsilon}{\tau} \tag{5.1}$$

where the efficiency ϵ is the percentage of total event triggers removed by the flag during the analyzed time period, and the downtime τ is the percentage of total analyzed time removed by the flag.

A κ of around 1 indicates a data quality veto performance consistent with random chance. During the most recent LIGO science run, S6, data quality vetoes with a κ of roughly 8 or greater were considered to be effective vetoes.

Safety testing and hardware injections

Before a data quality veto is used in a gravitational wave analysis it must first be demonstrated to be *safe*, meaning it does not systematically veto gravitational wave signals at a rate greater than expected by chance. To systematically test safety, hardware injections are used to simulate real signals by inducing motion of the test masses

³For Advanced LIGO, search pipeline output will likely be used to test the effectiveness of data quality vetoes to better tune veto use to different analyses.

that mimic a passing gravitational wave [26]. This technique is used to simulate a wide range of waveforms and SNRs to fully explore the realistic response of the detector to a potential gravitational wave signal, as well as test gravitational wave search algorithms.

Data quality flags that veto hardware injections at a rate greater than that expected by chance are considered *unsafe* and not recommended for use by the astrophysical searches. Note that a large number of hardware injections is needed to determine veto safety with statistical confidence, but this must be balanced with the risk of contaminating the time of a real gravitational wave signal with an artificially injected event.

A proposed veto's κ ratio, safety, and other considerations such as knowledge of the underlying physical coupling and the severity of the targeted glitch population for data analysis, are used to determine the veto's category, if any, for each search.

Burst data quality in LIGO's sixth science run

Data quality vetoes had a significant impact on the all-sky burst search in the previous LIGO science run, as shown in Figure 5.3. Category 2 and 3 vetoes removed most of the high SNR background events from the search, enabling the analyses to set more interesting astrophysical upper limits. However, even after category 3 vetoes were applied, the coherent waveburst background triggers show great non-Gaussianity with a tail of high SNR events. This is extremely limiting to the significance of any potential candidate event.

The cumulative deadtime introduced by category 2 and 3 vetoes on the coherent waveburst search was roughly 8% for H1 and 8.5% for L1 for the S6 search[3].

Two examples of effective data quality vetoes common to both interferometers during LIGO science run 6 are time just prior to loss of interferometer resonance (or *lock*

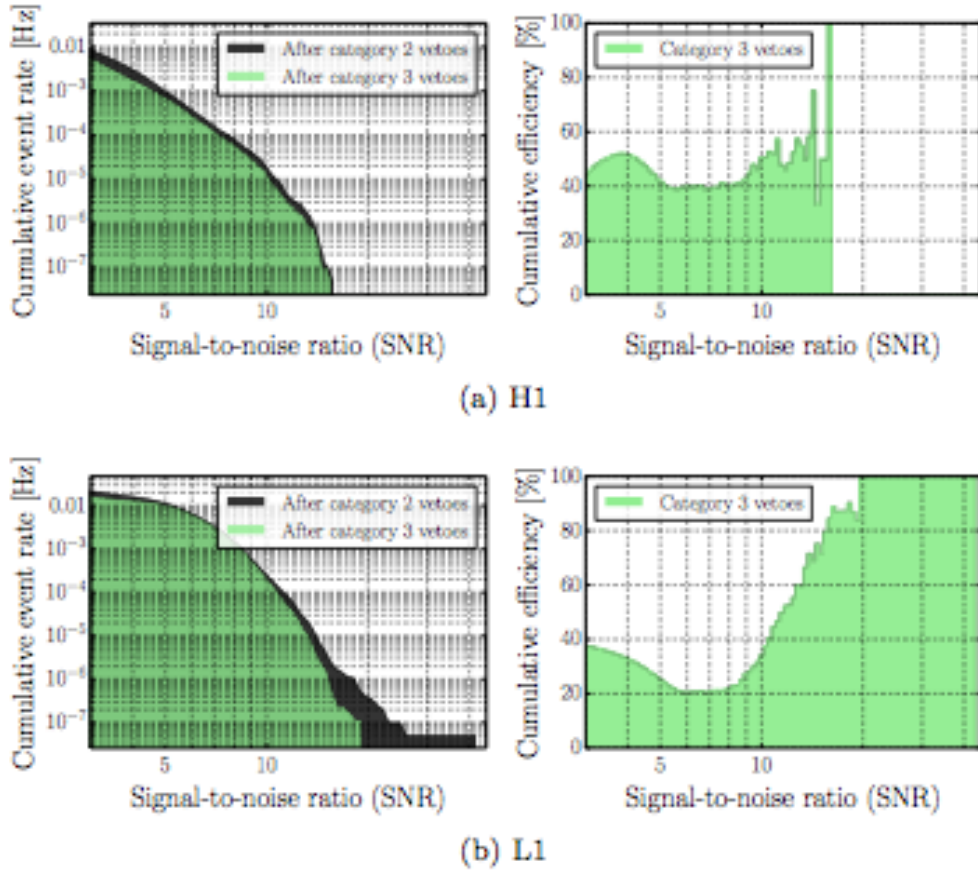


Figure 5.3: The effect of data quality vetoes on the S6 all-sky burst search. On the left is a histogram showing cWB background events after category 2 and 3 vetoes are applied vs SNR as reproduced in each LIGO interferometer. Note the non-Gaussianity of the triggers remaining after category 3. On the right is the total efficiency in removing background events from the analysis vs. single interferometer SNR for both LIGO detectors. The veto efficiency tends to be much higher for louder background events, as higher SNR noise transients tend to be both easier to diagnose and more likely to be targeted for mitigation. Reproduced from [3].

loss) and periods when a control loop managing some interferometer signal such as angular cavity alignment exceeded its actuation limit.

During S6 there were also some data quality vetoes generated in near-real time⁴ based on low latency auxiliary channel data and trends, generated mostly with simple thresholds. However, these were generally not very effective, removing only a small fraction of noise transients.

The most effective S6 data quality vetoes were generated offline to target noise transients with a large impact on astrophysical search performance. The most effective veto was generated by applying a tuned low frequency configuration of the ETG Omega to identify transient ground motion [64]. This is one of the first examples of the targeted application of burst data analysis techniques to instrumental data, and was the a basis for the study of the effect of transient ground motion on the motion of the optics, detailed in Section 5.5.

5.3 The impact of seismic noise on prior GW transient searches

It may seem nonintuitive that excess seismic noise, which according to Figure 2.2 limits the search range only at frequencies below 10 Hz, would interfere with transient astrophysical searches above 40 Hz. However, local ground motion was one of the most limiting noise sources to transient searches during previous LIGO science runs (2005–2010) [3, 69]. During Initial and Enhanced LIGO, seismic noise affected the transient searches via *up-conversion*, which up-converted low frequency seismic noise to significantly higher frequencies in the gravitational wave strain channel, as shown in Figure 5.4 [64, 3].

The top panel of Figure 5.4 shows a normalized spectrogram of local ground motion at the Hanford LIGO interferometer over a 24 hour period. After hour 12 there is

⁴For Advanced LIGO, online state vectors called Online Detector Characterization (ODC) channels will be produced in low latency, providing the low latency searches with detailed information about the status of the interferometer and its subsystems relative to their expected operating points.

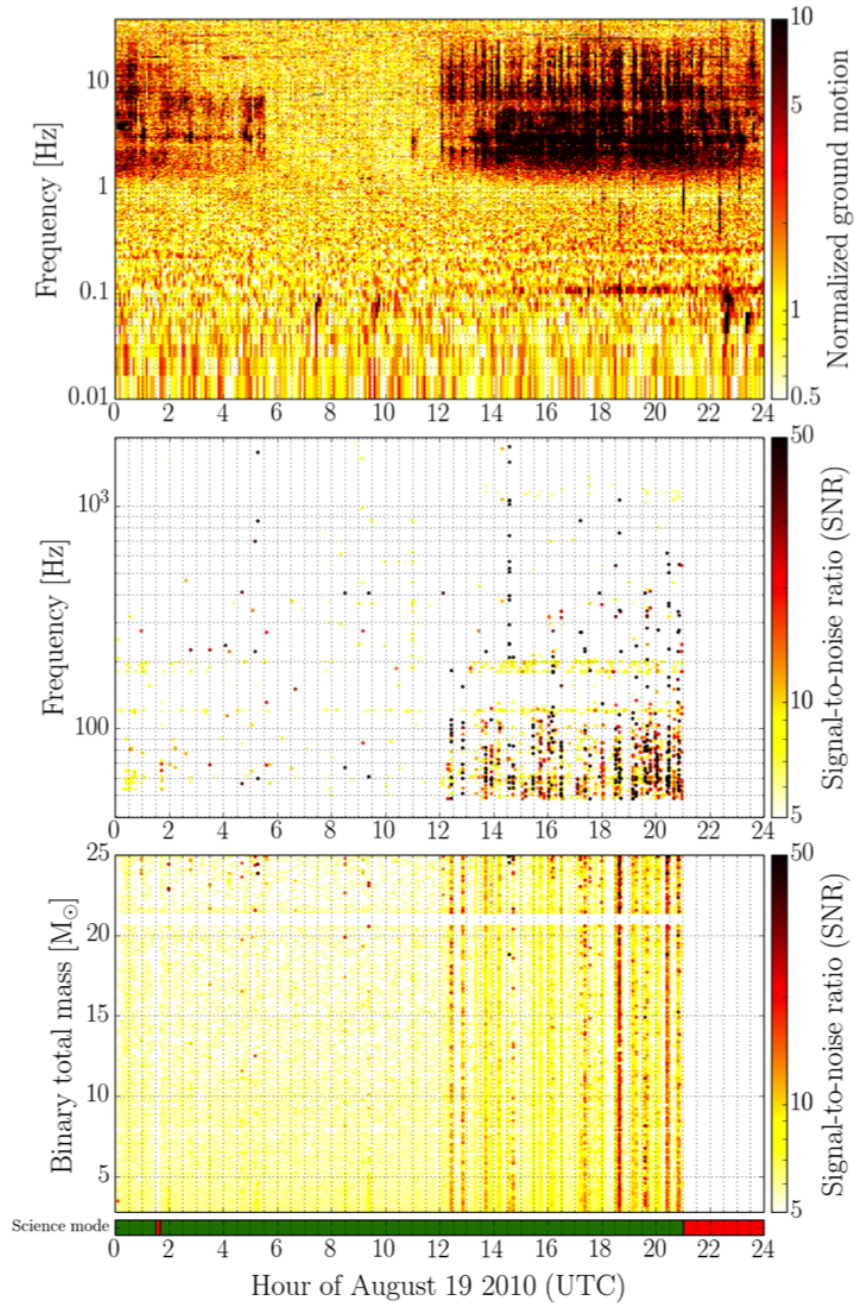


Figure 5.4: The top plot is a normalized spectrogram of low frequency ground motion at the Hanford interferometer during S6, where darker color shows excess motion. The middle plot shows triggers produced by the ETG Omega on Hanford $h(t)$ data during the same time period. A much higher Omega trigger rate coincides with the excess ground motion seen in the spectrogram. The lower plot shows triggers produced by the single interferometer compact binary coalescence pipeline hope on Hanford $h(t)$ data. There is also higher hope trigger rate during elevated ground motion, including times the full range of the mass template bank is rung up during very loud seismic events. Reproduced from [3].

a distinct period of elevated ground motion between 1-12 Hz. The next panel down shows Omega⁵ triggers measuring transient behavior in the Hanford gravitational wave strain channel over the same day. During the period of excess ground motion, there is also an excess rate of Omega triggers ranging from 60Hz to a few hundred Hz; well into the most sensitive regions of the astrophysical search frequency band. The bottom panel shows transient behavior in the gravitational wave strain channel during the same time period as seen by the single interferometer inspiral pipeline ihope. Again, there is a distinct excess of triggers during the period of excess low frequency ground motion, with some events spanning the entire template mass range.

Another metric that illustrates the impact of seismic noise on the transient astrophysical searches during prior science runs is the effect of vetoing times coincident with low frequency ground motion transient events. During the most recent LIGO science run, a novel veto method was introduced that used the ETG Omega, tuned to resolve low frequency, long duration transients, to identify periods of excess ground motion [64].

This seismic transient identification technique, known as SeisVeto, was very efficient in vetoing events identified by Omega in the nominal burst search frequency range of the gravitational wave strain channel. During the last epoch of LIGO science run 6, SeisVeto removed 85% of extremely loud events, 55% of very loud events, and 27% of all loud events from Hanford data, as shown in Table 5.1.

The effect of SeisVeto on Hanford $h(t)$ burst triggers over time is illustrated in Figure 5.5. The veto is clearly effective throughout the duration of the epoch on a broad range of event SNRs. However, the veto did introduce a 56 hour deadtime to the analysis,

⁵Omega is an ETG, or single interferometer burst pipeline, described in Chapter 3.

Performance of SeisVeto on high SNR events during S6D

Event SNR range	Percent of events vetoed
SNR > 20	27%
SNR > 100	55%
SNR > 1000	85%

Table 5.1: The percentage of transient noise events in the LIGO Hanford $h(t)$ channel successfully vetoed by SeisVeto during the last epoch of LIGO science run 6 [69].

omitting 4.7% of all analyzed time during S6D. This high deadtime motivated the burst search to use the flag at category 3 instead of category 2 [69].




Figure 5.5: The effect of SeisVeto on Hanford $h(t)$ Omega triggers over the duration of the last LIGO science run 6 epoch, S6D. Triggers are plotted by SNR vs. time. Triggers removed by SeisVeto are in red, and unvetoed triggers are in blue. Any gaps in triggers are times the interferometer was out of lock or not operating in a nominal configuration.

Notably, this successful application of the principles of burst data analysis to instrument characterization produced one of the most effective data quality vetoes of the entire science run.

Elevated seismic noise also had an adverse impact on all gravitational wave searches by limiting network uptime, or the total time multiple interferometers were jointly resonating at their operating point. For example, short duration seismic events such as earthquakes or local truck traffic commonly caused lock loss, or loss of interferometer resonance, during S6 by perturbing the finely controlled optic motion. Longer duration seismic noise such as periods of high wind at Hanford and microseism and anthropogenic noise at Livingston also interfered with achieving and maintaining interferometer resonance during prior LIGO science runs⁶.

The upgraded seismic isolation instrumentation designed and installed for Advanced LIGO is expected to mitigate the impact of seismic noise on the transient searches by all measures. At the time of writing, these systems have already been shown to provide robust lock stability and improved strain sensitivity below 30 Hz. This instrumentation is also expected to improve the propagation of transient motion from ground to the optics, as explored in Section 5.5.

5.4 Advanced LIGO seismic isolation instrumentation

Overview

Advanced LIGO seismic isolation infrastructure is expected to improve the detectors' sensitivity by many orders of magnitude at 10Hz. This effectively moves the *seismic wall*, or the lowest frequency it is practical to search for gravitational wave signals, from roughly 30-40 Hz in Initial LIGO to roughly 10-15 Hz in Advanced LIGO [2].

The design performance of Advanced LIGO seismic isolation is achieved in two major systems combined in series: a set of actively actuated and passively isolated plat-

⁶This was especially true at Livingston before the installation of the external HEPI isolators.

forms that compromise the seismic isolation subsystem, and the optic suspensions supported by these platforms. The interferometer test mass optics are isolated by seven individual cascaded stages, each passively isolated from the preceding stage, and three actively isolated.

Each optic chamber shown in Figure 2.7 is actively isolated, and each type of chamber has a motion performance requirement. There are two chamber configurations: one for BSC chambers⁷, which house the core optics, shown in Figure 5.6, and one for HAM chambers, which house auxiliary optics, shown in Figure 5.7. The active seismic isolation systems for both chamber configurations are composed of several mechanically isolated platforms, known as *stages*.

The performance criteria of the seismic isolation subsystem from the perspective of the transient gravitational wave searches are:

- The effectiveness of the active isolation control loops in suppressing the average motion of the platforms⁸. The control loops are designed to meet or exceed the Advanced LIGO performance requirement for each type of chamber [65].
- The stability and robustness to strong seismic impulses of the control loops for every platform (2-3 per chamber) and every chamber (11 per interferometer), which individually all directly impact the total effective observation time of the instruments.

⁷BSC stands for Basic Symmetric Chamber, but this acronym is not descriptive. Similarly for HAM - Horizontal Access Module. The important distinction is that BSC chambers contain the core optics and HAM chambers contain auxiliary optics, and thus have different sizes and seismic isolation requirements by design.

⁸Figure 5.10 show spectra of average motion along the beam direction of the optic table for each optic chamber (circa January 2015).

- The mitigation of transient motion (excess motion that is distinct in time) between the ground and the platform supporting the optic suspensions, or optics table.

HEPI

An external isolation stage called Hydraulic External Pre-Isolator (HEPI), is common to both BSC and HAM chambers. An example of HEPI can be seen for BSC chambers in Figure 5.6, and for HAM chambers in Figure 5.7. HEPI differs between BSC and HAM chambers only in support structure geometry and payload, or the platforms it supports.

HEPI platforms are actively isolated from the ground via quiet hydraulic actuators, which act to damp the HEPI structure's internal resonance modes. As shown in Figure 5.6, the isolated HEPI stage, stage 0, is supported with four sets of steel springs. Each of the four points of contact between the HEPI support pier and the isolated stage are also equipped with a set of position sensors and inertial sensors, used for active isolation and positioning of stage 0.

The HEPI platforms are designed to target very low frequency motion, generally below 5 Hz, and can also be used for optic position offset control [65].

BSC ISI

For BSC chambers, there are two internal stages of seismic isolation, illustrated in Figure 5.6. The first stage is mechanically isolated from the stage below (stage 0, HEPI) via a set of blade springs and flexure rods, which respectively provide vertical and horizontal passive isolation. The mechanical resonance of this passive isolation is actively damped.




Figure 5.6: A schematic showing the internal active isolation (ISI) stages, the external active isolation stage (HEPI) and double quadruple suspension for a BSC chamber. Only one of four HEPI support piers is shown, each with a set of hydraulic actuators, positions sensors, and inertial sensors. Stage 0 transitions from out of vacuum to in-vacuum via bellows, not pictured. One of three sets of blade springs and flexure rods that passively isolate stage 1 from stage 0 and stage 2 from stage 1 are shown in blue. Reproduced from [2].




Figure 5.7: The cascade of seismic isolation stages for an example HAM chamber. In dark blue is the HEPI support structure. In light purple is ‘stage 0’, the stage actively and passively isolated from the ground by HEPI. The large hexagonal structure is the HAM ISI, or stage 1, isolated actively and passively from stage 0. The top surface of the HAM ISI, in teal, is the optics table, which supports the various auxiliary optic suspension cages, in red. Reproduced from [54].

Each BSC ISI supports a quadruple monolithic suspension and reaction chain⁹, except the beam splitter, which hangs from a triple pendulum without a reaction chain. The end station BSC ISIs also support transmission mirrors, as shown in 2.7. These suspensions are hung from the underside of the optics table, as shown in Figure 5.6.

The BSC ISIs are expected to provide more than three orders of magnitude of seismic isolation at about 10 Hz [65].

HAM ISI

⁹These are fully described under *Suspensions*.

For HAM chambers, there is one stage of internal seismic isolation, as shown in Figure 5.7. Similarly to the BSC ISIs, the HAM ISI stage is passively isolated from the preceding stage with a set of three blade springs and flexure rods, which are largely hidden underneath the ISI structure in the figure.

Each HAM ISI supports a different set of major auxiliary optics. Figure 5.7 shows a HAM ISI supporting three triple suspensions, shown in red. These suspensions are hung from cages mounted on the top of the optics table.

The HAM ISI is expected to provide more than an order of magnitude of isolation at 1 Hz [65].

Active control

The general isolation control scheme is common to all stages. The motion of each stage is measured with inertial and position sensors and controlled with precise, powerful actuators to remain as stationary as possible in target frequency ranges, as shown in Figure 5.8 [54].

The control loop for an actively isolated platform is shown in more detail in Figure 5.9. The resulting output motion X_1 is the sum of the input motion from the lower stage, X_0 , as filtered by the passive blade and flexure rod isolation, plus the motion induced by the loop actuators. These two paths are indicated in the diagram as P_S and P_F for seismic path and control force path, respectively.

The motion of the platform is sensed with two sets of sensors: inertial sensors which measure the absolute motion of the platform, U_a and position sensors which measure the relative motion of the platform, U_r . The noise of these sensors is indicated in the diagram as N_a and N_r .

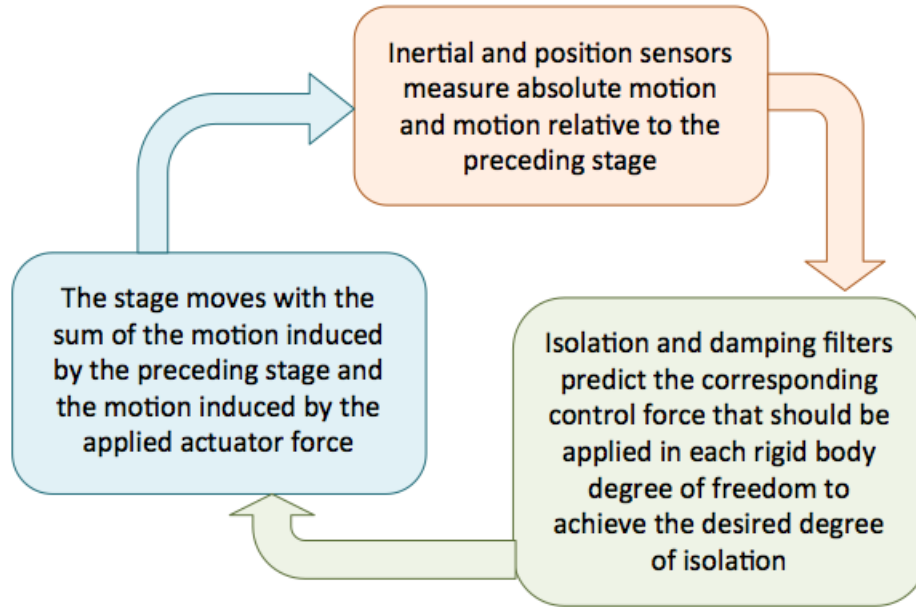


Figure 5.8: This simplified flowchart illustrates the basic feedback loop of the aLIGO active seismic isolation stages.

Since the inertial sensors are more noise limited below roughly 1 Hz and the position sensors are more noise limited above roughly 1 Hz, the inertial and position sensors are blended together into a *super sensor*. This is achieved by applying a high pass filter to the inertial sensors, applying a low pass filter to the position sensors, integrating the inertial sensor signal, and adding the signals together. The frequency at which the position and inertial sensors contribute equally to the blend is known as the *blend frequency*, and is generally marginally below 1 Hz. The high and low pass filters are labeled H and L respectively in the diagram [65].

The blended super sensor is used to determine the actuation applied to the stage, shown as the feedback controller, C_{FB} in the diagram. The applied actuation contributes to the output motion, X_1 , and the control loop is complete, or *closed*.

Also shown in Figure 5.9 are feedforward and sensor correction options in the active seismic isolation control loop. Feedforward control uses the measured input motion

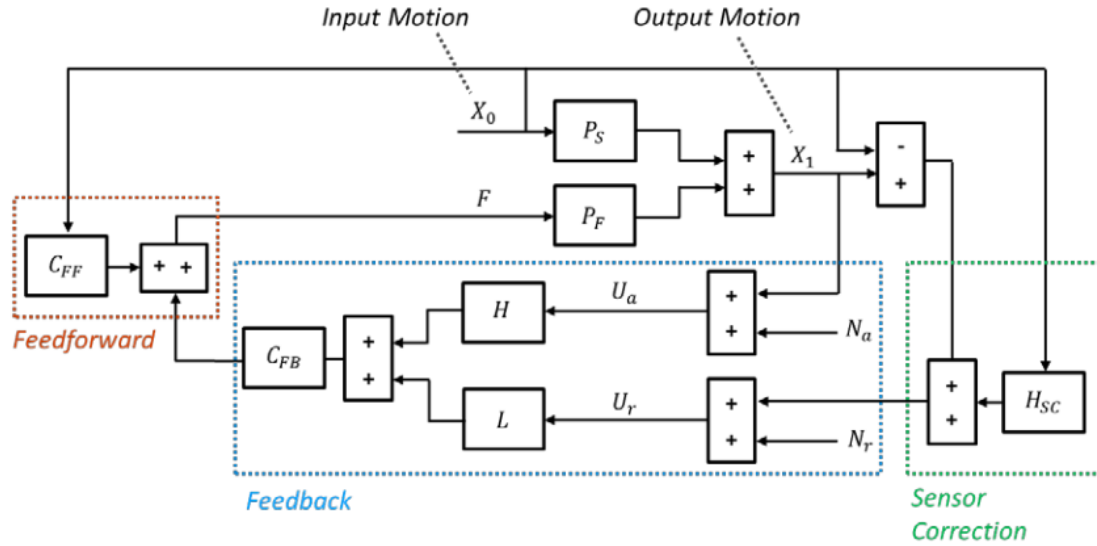


Figure 5.9: The control scheme for the output motion, X_1 , of an actively isolated platform in response to input motion from the prior stage, X_0 . Inertial sensors, U_a , and position sensors, U_r are fused to feed the control signal C_{FB} , which drives the platform through the control force path, P_F . Reproduced from [65].

of the platform below to generate a control signal that is added to the actuation of the platform, feeding the signal forward in the loop. Sensor correction is a technique designed to address the limitation of the position sensor noise. This is accomplished by high passing the input ground motion, measured by an inertial sensor, and adding this signal to the position sensors, which measure the relative motion of that stage with respect to the prior stage [65]. Neither feedforward control or sensor correction were used in the seismic configuration for the study in Section 5.5.

A minimum of six inertial sensors and six position sensors per stage is required to define a complete basis to sense all rigid body motions of the stage. However, on some stages more sensors are used to sense stage deformation, or as an input signal for feedforward control for a subsequent isolation stage.

Current performance

The active seismic isolation systems perform very well during quiet seismic times by either coming very close to meeting or exceeding the Advanced LIGO design requirements for average optics table motion. Figure 5.10 shows the current performance of each chamber measured by optics table motion along a cavity beam¹⁰, the direction most closely coupled with noise in the gravitational wave strain measurement. These curves are compared to both the ground motion, measured in two different degrees of freedom, and the design requirement for that chamber type, shown in the dashed line.


Although the seismic isolation systems perform very well in average motion and have been shown to support long, robust interferometer lock stretches, one of the crucial performance metrics of the seismic isolation subsystem is how effectively the active seismic isolation systems mitigate transient seismic noise.

Suspensions


There are several different types of suspensions used in Advanced LIGO, as shown in Figure 2.7: quadruple suspensions and partner reaction chains used for the test mass optics, a special triple suspension for the beam splitter, two large HAM triple suspensions for the largest optics in the power and signal recycling cavities, seven small HAM triple suspensions for the smaller power and signal recycling cavity optics and the three input mode cleaner optics, plus an assortment of smaller double and single stage suspensions for auxiliary optics with less stringent isolation requirements.

For active seismic isolation system characterization, it is important to consider the active seismic isolation and suspension for each optic as complementary parts of a broader system with the same goal: to isolate the optics from local ground motion.

¹⁰For the HAM ISIs, the auxiliary optic layouts are not as straightforward as X and Y are not quite lined up with a cavity beam direction. The closest parallel Cartesian degree of freedom is shown.)



(a) The average performance of the Livingston BSC chambers, which hold the core interferometer optics (the beam splitter and test masses) on February 7th, 2015. Each chamber was operating in the nominal seismic isolation configuration for the Livingston interferometer circa January 2015: the HEPI was locked into place mechanically coupled to the ground, and the internal isolation stages were actively isolated. The BSC chamber requirements curve is shown as a dashed line.



(b) The average performance of the Livingston HAM chambers on February 7th, 2015. The HEPI supporting each HAM chamber was locked and all chambers were actively isolated. The requirements curve for Power Recycling Cavity HAM chambers, which house the Input Mode Cleaner, Power Recycling Cavity, and Output Mode Cleaner optics, is shown as a dashed line. (Note the Signal Recycling Cavity HAM chambers, which house the Signal Recycling Cavity optics, have a slightly different, more stringent requirements curve [65]).

Figure 5.10: An illustration of the current performance of the Advanced LIGO active seismic isolation, targeting frequencies up to roughly 10Hz. The average measured ground motion is shown for two degrees of freedom in darker shades of blue. The average motion of the most isolated stage, the optics table, is shown in other colors for each chamber along the direction of the beam path. See Figure 2.7 for a reminder of chamber names. Reproduced from the DetChar daily summary pages [63].⁹⁰

There may sometimes be undesirable excess transient motion observed in the optics table sensors, but if this motion does not couple through the suspensions to the motion of the optic, then it will not affect the astrophysical gravitational wave searches.




Figure 5.11: Motion attenuation provided by the Advanced LIGO suspensions, showing strong isolation at higher frequencies (i.e. above roughly 10 Hz). Each stage of passive isolation provides $\frac{1}{f^2}$ isolation, with the quadruple suspensions that hold the core optics (the beam splitter and test masses) providing $\frac{1}{f^8}$ motion attenuation above a few Hz. Figure produced by B. Shapiro.

The passive quadruple and triple suspensions provide extremely effective seismic isolation at frequencies above roughly 10 Hz, as shown in Figure 5.11. Although the actively isolated platforms only target frequencies below roughly 10 Hz, any motion above 10 Hz is dramatically suppressed by the optic stage; already by 4 orders of magnitude at 10 Hz.

Although the suspensions are largely passive stages in contributing to seismic isolation, active control is used to damp mechanical resonances and to control and stabilize the interferometer degrees of freedom, listed in Table 2.1.

Actuation of the uppermost stage of the suspensions is used to actively damp the low frequency rigid body modes of the suspension ($< 9\text{Hz}$). The lower stages are




Figure 5.12: A schematic showing different views of the suspension of a test mass. On the left is the full quadruple suspension shown with a partner reaction chain. On the right the finer details of the bottom two stages of the test mass suspensions are shown. Reproduced from [2].

actuated for global interferometer control - controlling both cavity length and the angular orientation of the optics relative to their counterpart optics to achieve and maintain interferometer resonance at the optimal operating point [2].

The auxiliary suspensions are actuated against the suspension cage, which is mounted directly on the optics table and does not benefit from the passive isolation of the suspension. However, for test mass optics, a partner suspension chain provides a passively isolated set of masses for quieter actuation [2].

Actuation of the actual test mass optics, the lowest stage of the quadruple suspensions, is achieved via an electrostatic drive: an annular pattern of gold electrodes laid on the surface of the lowest stage of the partner suspension chain facing the test mass¹¹.

5.5 Advanced LIGO seismic noise transient propagation

Considering the significant impact of seismic transient events on the gravitational wave transient searches during S6, characterizing the effectiveness of active isolation systems in mitigating transient ground motion is of great interest.

Overview

The following studies compared the rate, amplitude, and frequency distribution of transient events between active isolation stages. The ETG Omicron, described in detail in Section 3.1, was used to measure the transient behavior of the ground motion and the motion of each isolated stage of a set of BSC chambers via inertial sensors.

¹¹The lowest stage of the partner suspension chain, the stage on which the electrostatic drive (ESD) is placed, is called the Compensation Plate (CP) for the input test masses (stemming from use by the Thermal Compensation System) and actuation reaction mass (ERM) for end test masses, as labeled in Figure 2.7.

- **Study I** characterized the effectiveness of transient mitigation using different levels of isolation aggression.
- **Study II** evaluated the performance of the Advanced LIGO active seismic isolation during periods of elevated local ground motion: high wind at Hanford, and nearby anthropogenic noise (logging) at Livingston.

5.5.1 Technique

The following inertial sensors were used: STS2s for local ground motion, L4Cs for HEPI motion, L4Cs for ISI stage 1 motion, and GS13s for ISI stage 2 (optics table) motion. All sensors were calibrated in units of nm/s¹².

Tuning Omicron for seismic signals

Omicron tuning for this study included two major capability upgrades. The first improved Omicron’s data whitening algorithm to enable analysis below 2 Hz [80]. The second introduced the output of amplitude in physical units as a trigger parameter, as defined in Section 3.1.2.

Omicron parameters were also tuned to resolve low frequency transient events. Figure 5.13 demonstrates the need for specialized tuning for very low frequency events, as shown in [64]. The left panel shows excess power in a normalized spectrogram of ground motion, and the right panel shows an absence of Omicron triggers below 1 Hz. This indicates that very low frequency transient events were not recovered by Omicron in a pre-tuned configuration.

¹²The L4C and GS13 signals were not inverted because Omicron does not yet have the capability to apply filters on the fly. This means results reported in physical units are unreliable below 1 Hz, but the reported outcomes are unaffected.

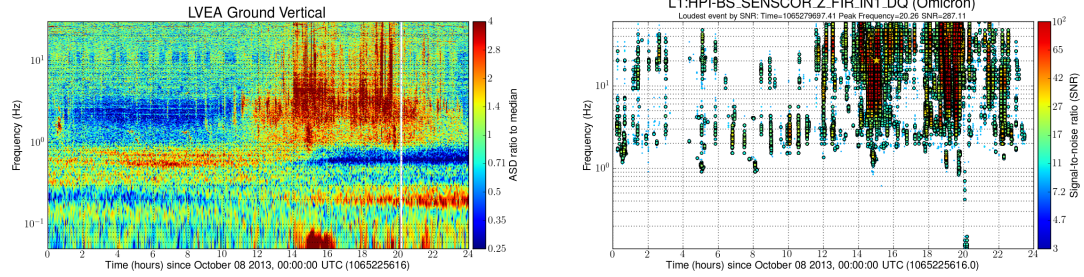


Figure 5.13: A comparison of a recent normalized spectrogram of a Livingston ground motion sensor (left) and Omicron triggers for the corresponding time *before* parameter tuning (right). Note the low frequency structure seen below 1Hz in the normalized spectrogram is not resolved by Omicron.

The Omicron frequency range was set from 0.1 to 64 Hz, in order to capture higher frequency transients that might be well resolved by the L4C and GS13 sensors.

Comparing the low-frequency-tuned Omicron configuration to the nominal configuration used for $h(t)$ data: The ‘chunk’ duration, which is the length of time used for PSD estimation, was lengthened to from 484 to 8192 seconds to enable analysis down to 0.1 Hz. The ‘block’ duration, which is the length of time used to calculate each set of time-frequency tiles, was lengthened from 64 to 8192 seconds. The overlap duration, which governs how much data overlaps between adjacent segments, was expanded to 1280 seconds to avoid poor resolution due to edge effects. The permitted mismatch between a sine-Gaussian event and the sine-Gaussian template was raised from 0.2 to 0.35. The upper limit of the Q range was lowered from 121 to 100 as very high Q events are less likely in the well damped active isolation systems.

A complete example of the configuration used can be found in Appendix D. Note that the very long data segments used for low frequency analysis introduces a very long latency that’s impractical for producing online or automated triggers for data quality purposes.

Study I: active isolation configuration

For this study the input ground motion was kept roughly constant in order to isolate the impact of varying the isolation aggression on the transient motion of the actively isolated stages of the ETMY (end-Y) BSC ISI at Hanford. This was accomplished by using only ‘quiet’ night times, 07:30:00 - 15:00:00 GMT, unpolluted by earthquakes or other seismic noise for each analysis.

The isolation aggression was varied by changing the number of platforms actively isolating and also by adjusting the loop gain used in the isolation filters that contribute to the control signal C_{FB} shown in Figure 5.9.

The two key configurations compared were those with maximum and minimum isolation. For the minimally isolated configuration HEPI was controlled with the least aggressive isolation filters and the ISI stages were controlled to damp the mechanical resonance modes only, with no isolation engaged. This configuration is denoted as ‘Damped’. For the maximally isolated configuration, the HEPI control was unchanged and both ISI stages were controlled with the most aggressive isolation filters.

Other analyzed configurations and additional details can be found in the study documentation [68].

Study II: input ground motion configuration

For this study series, the active isolation configuration was kept constant and the input ground motion was varied to isolate the impact of elevated ground motion on the transient motion behavior of the optics tables.

The high wind analysis at Hanford compared the ETMX chamber response to a period of wind speeds of roughly 30 MPH to a period half a day prior with nominally low wind speeds of roughly 5 MPH. Both chambers were configured so that ISI stage 1 was actively isolated using the most aggressive isolation filter, and ISI stage 2 was only damping resonance modes.

Additional analyzed chambers and details can be found in the study documentation [67].

The high anthropogenic noise, or logging, analysis at Livingston took advantage of a period of distinctly elevated ground motion following relative quiet during a long, stable full interferometer lock stretch as shown in Figure 5.14.

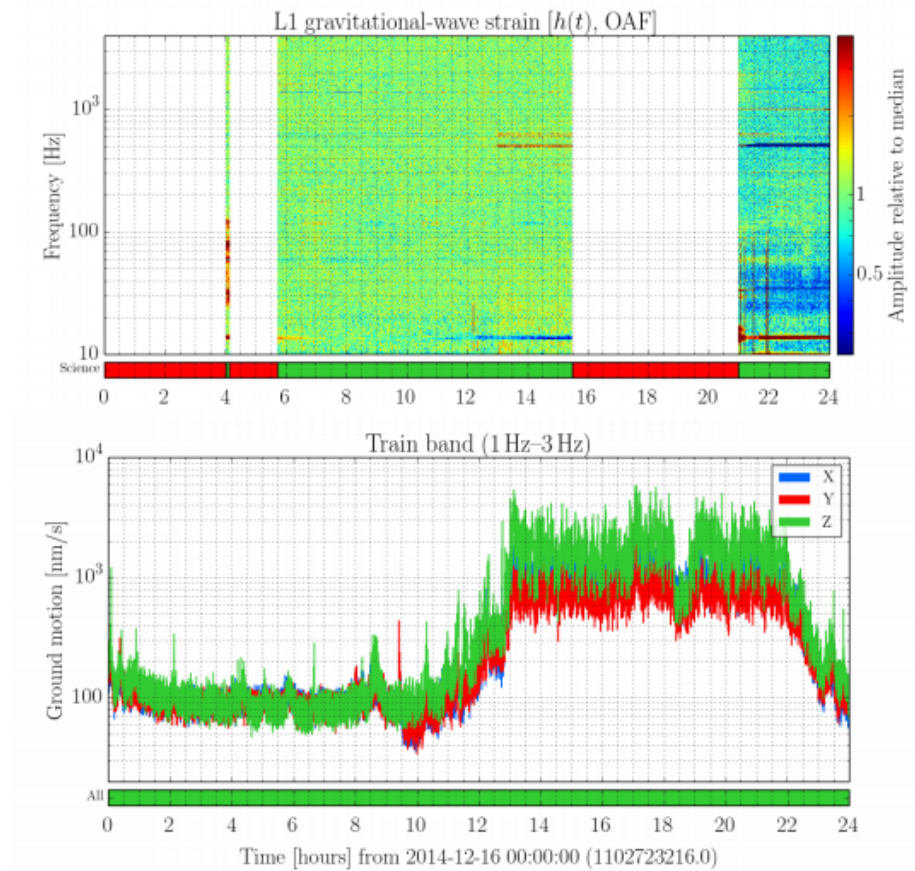


Figure 5.14: The top panel shows a normalized spectrogram of the $h(t)$ channel at LIGO Livingston for a lock stretch spanning from hour 06:00 to 15:30. The bottom panel shows the band-limited root-mean-square (BLRMS) trends for 1-3 Hz local ground motion during the same day. The order of magnitude increase 1-3 Hz ground motion starting at roughly 13:00, identified to be nearby logging activities, clearly rings up the suspension violin mode of 508 Hz and induces greater non-Gaussian behavior in $h(t)$. Reproduced from the DetChar Summary Pages [63].

This study was able to compare the relative glitch rate increase between quiet and ‘logging’ times for the ground, the ITMY BSC optics table, and the differential arm interferometer degree of freedom used to generate calibrated $h(t)$.

5.5.2 Results

For reference in interpreting results, Figure 5.15 shows a simplified diagram of the active and passive seismic isolation stages of a BSC ISI housing a test mass. Generally, for transient motion to be observed in the optics table, ISI stage 2, excess ground motion must propagate through the HEPI stage as well as ISI stage 1.




Figure 5.15: A simplified diagram of the passive and active seismic isolation stages of a BSC ISI chamber housing a test mass with a quadruple suspension. Reproduced from [65].

Study I: varied seismic configuration

Study I compared the mitigation of transient motion between the damped and fully isolated configurations. Figure 5.16 shows the frequency vs. amplitude trends of Omicron triggers for each stage.

Collectively, the aggressive isolation configuration reduces the amplitude of transients as recovered by Omicron by roughly four orders of magnitude below 10 Hz. However, the amplitude of transient motion in the optics table above 10 Hz, approaching the astrophysical search frequency band, is largely unchanged relative to the least aggressive configuration. This is consistent with the active isolation stages targeting frequencies below 10 Hz.

Although it seems unlikely that the type of transient noise seen during quiet times would not be abated via the suspensions, the results of study II show this is indeed the case. Using recent seismic configurations, excess ground motion is either known to affect $h(t)$ or has a great potential to affect $h(t)$ well above 10 Hz.

Study II: varied input ground motion

Study II kept the isolation configuration constant and compared transient mitigation performance between quiet ground motion and elevated ground motion.

High wind at Hanford

This analysis compared a period of high wind to a time just prior with low wind, labeled ‘quiet time’ in Figure 5.17. These plots show rate histograms of Omicron trigger amplitudes by isolation stage for both windy and quiet times.

The median Omicron trigger amplitude for all isolation stages is shifted higher during windy time. Notably, the median amplitudes for ground motion and optic table




Figure 5.16: The frequency vs. amplitude (nm/s) distribution of Omicron triggers for each stage. Each discrete transient motion event is represented as a dot. The top plot shows results for the least aggressive isolation configuration (Damped), and the bottom shows the most aggressive (Fully isolated) during equivalently quiet seismic times. Omicron triggers for the inertial sensor measuring the ground are shown in dark blue, the HEPI stage triggers in green, ISI stage 1 triggers in red, and the ISI stage 2 triggers in teal.

motion are both increased by roughly a factor of 8. Figure 5.18 shows frequency vs. amplitude trends of Omicron triggers for each stage.

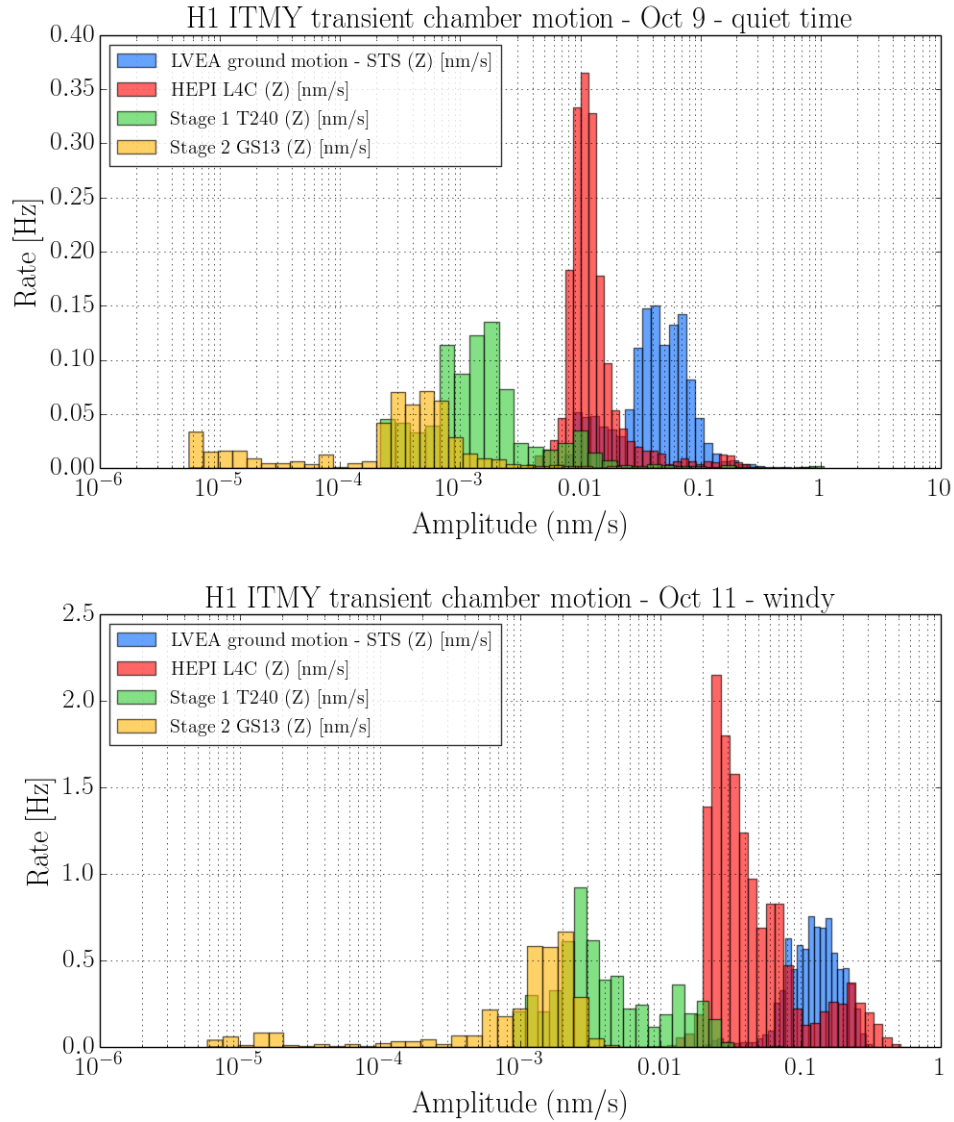


Figure 5.17: Rate histograms of Omicron trigger amplitude for each isolation stage. Omicron triggers for local ground motion are in blue, for the HEPI stage in red, for ISI stage 1 in green, and for ISI stage 2 in yellow. Quiet time, with a wind speed of 5 MPH, is shown in the top plot and windy time, with a wind speed of 30 MPH, is shown in the bottom.

Relative to the quiet seismic time, the high wind increases the average transient event amplitude by up to an order of magnitude from 0.1-10 Hz. However, similarly to what was observed in Study I, the amplitude of Omicron triggers for optics table motion

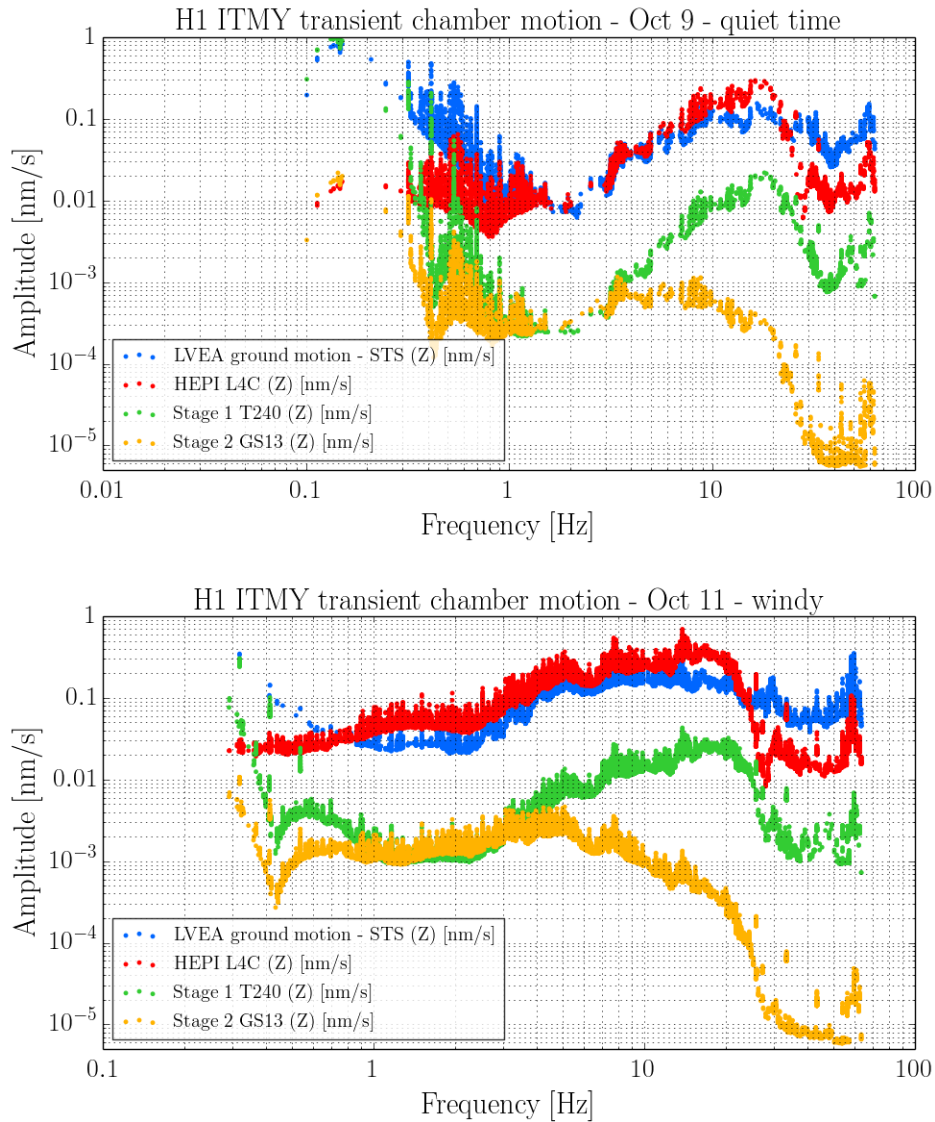


Figure 5.18: The frequency vs. amplitude (nm/s) distribution of Omicron triggers for each stage. Each discrete transient motion event is represented as a dot. The top plot shows results for quiet time, and the bottom shows windy time. Omicron triggers for the inertial sensor measuring the ground are shown in blue, the HEPI stage triggers in red, ISI stage 1 triggers in green, and the ISI stage 2 triggers in yellow.

above 10 Hz is not increased, even with very high wind. This is an interesting and important result.

Figure 5.19 shows Omicron triggers of optics table motion in time vs. frequency, where the SNR, or relative loudness, of each event is indicated by the color scale.




Figure 5.19: The distribution of Omicron triggers for the optics table motion in time vs. frequency. The top plot shows the transient behavior of the ITMY optics table during quiet seismic time, and the bottom plot shows windy time. The color scale indicates the SNR of each event. The event rate is clearly greatly elevated during windy time.

Clearly, the rate of transient events is greatly elevated during periods of high wind relative to nominally quiet times, although the amplitude of transient events in optics table motion above 10 Hz is not. If this elevated noise transient effect were to couple

to $h(t)$ even at a very low rate, it would be critically harmful for the gravitational wave burst search, which is most affected by a high rate of noise events.

As the $h(t)$ channel is still relatively new at Hanford, investigations into the effect of high wind on the transient behavior of $h(t)$ are still underway.

Logging at Livingston

Analysis of quiet time vs. logging time at Livingston showed very similar results between stage performance to the Hanford wind study. Trigger amplitudes were elevated in same range of frequency: 0.1-10 Hz. The transient event rate was also elevated in stages more mechanically isolated from the ground.

Focusing on the relative behavior of the $h(t)$ channel, Figure 5.20 shows the frequency vs. amplitude of Omicron triggers produced on $h(t)$ during quiet time and logging time during the same lock stretch.

After the logging started during the lock, the $h(t)$ glitch rate increased by a factor of 10. Additionally, this glitch increase was very broadband; clearly seen across the entire analyzed frequency range Figure 5.20, well above 1 kHz.

This is exactly the response to seismic noise that is detrimental to burst search performance.

5.6 Summary and future work

Although the Advanced LIGO active and passive seismic isolation stages are very effective in mitigating the average motion of the optics, we do see strong evidence that transient ground motion will impact the transient gravitational wave searches using the current seismic isolation configurations. This is especially true during times of high wind at Hanford or strong anthropogenic noise at Livingston.

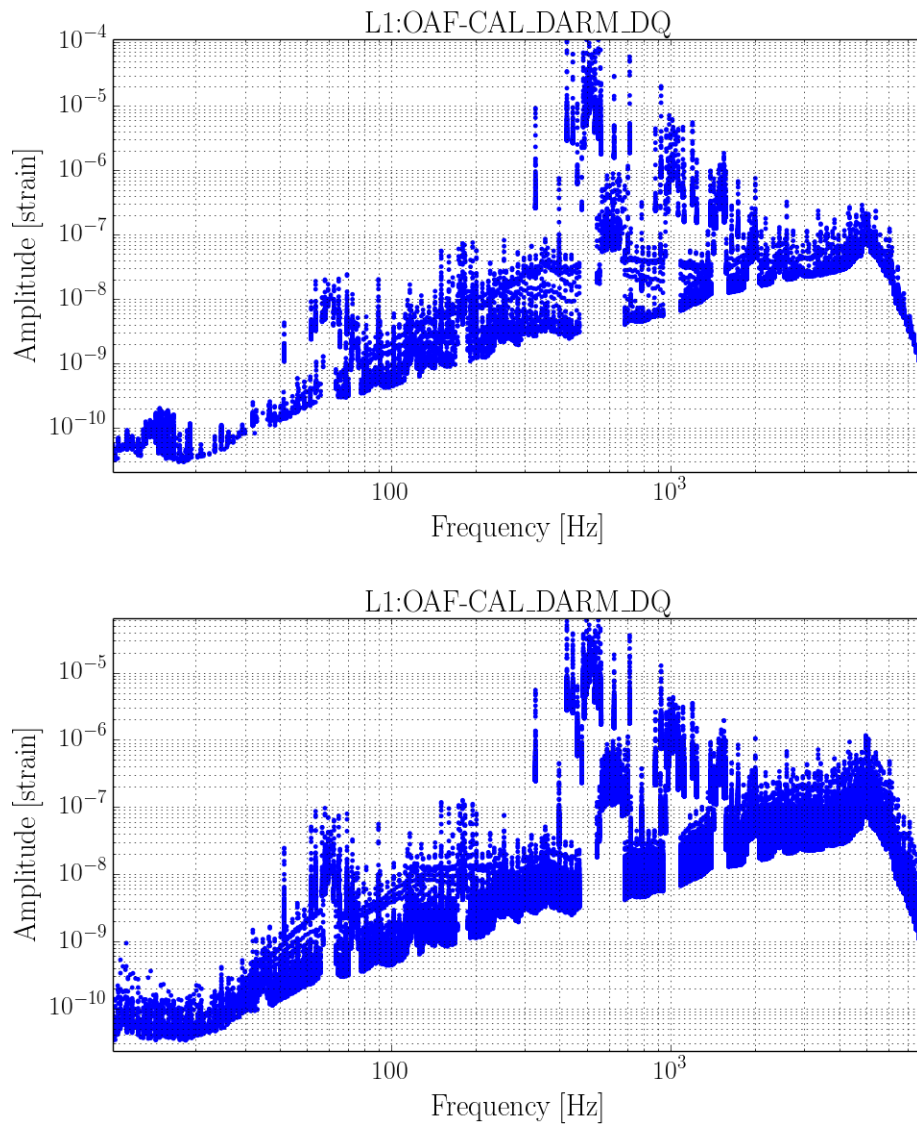


Figure 5.20: The frequency vs. amplitude of Omicron triggers for the Livingston calibrated DARM signal, used to generate $h(t)$. The top plot shows transient events during quiet time, and the bottom plot shows transient events during seismically loud logging that occurred during the same lock stretch as the quiet time. Note that strain is not exactly calibrated, unitless strain, but the relative amplitudes should still hold.

The actively isolated stages greatly mitigate the amplitude of glitches, by up to four orders of magnitude above 10 Hz. However, excess ground motion still produces a greatly elevated rate of transient motion for the optics table, even above 10 Hz. An

elevated glitch rate has been shown also for Livingston $h(t)$ during a period of loud anthropogenic noise.

Clearly there will still be some impact on the astrophysical searches if no further action is taken, albeit likely greatly abated compared to prior LIGO science runs. A useful next step would be to quantify the impact of common sources of elevated ground motion at both LIGO sites on transient search performance, which grows more viable as both detectors increase the number and length of full interferometer locks.

As future work, recommendations for control loop tuning targeting transient motion mitigation should be made using transient ground motion events known to couple to $h(t)$.

It is also worth emphasizing that outside of the seismic isolation subsystem are many other potential glitch sources. Noise couplings depend on roughly a dozen complex subsystems and a myriad of variables. Some of these interact with excess seismic motion, such as mirror imperfections which cause light scattering. However, the Detector Characterization group has made great progress identifying potentially troublesome DQ features in the data as the LIGO interferometers have come online, and the data on the whole looks relatively clean at both sites at the time of writing.

CHAPTER 6

THE IMPACT OF TERRESTRIAL SEISMIC NOISE ON CCSN DETECTABILITY AND RECONSTRUCTION

Studies presented in Chapters 4 and 5 have explored both the detectability and reconstruction of core-collapse supernova events in unrealistically well-behaved Gaussian noise and also the nature of terrestrial noise and its prior impact on the gravitational wave burst searches.

This chapter presents a follow-up study to that outlined in Chapter 4 that characterizes the impact of realistic detector noise on the detectability and reconstruction of core-collapse supernova events, tying all of the work detailed thus far together.

6.1 Technique

This study followed the same procedure outlined in Chapter 4, except instead of injecting modeled CCSN waveforms into Gaussian noise, the waveforms were injected into LIGO science run 5 (S5) and Virgo Science Run 1 (VSR1) data recolored to the expected design sensitivities of Advanced LIGO and Advanced Virgo. The data from the joint science run prior to the most recent was chosen for convenience as the injected times of the modeled waveforms used line up well with the non-Gaussian data.

Category 1, 2, and 3 data quality flags were applied to the data, as described in Chapter 5. Additionally, all hardware injection times were removed from the analysis.

Keep in mind the following results are after meticulous data quality cuts, the result of many noise investigations.

The same set of modeled waveforms described in Chapter 4 were used and cWB2G used the same CCSN-tuned configuration for each waveform set.

6.2 CCSN detectability in recolored non-Gaussian data

Dimmelmeier

The relative recovery rate by cWB2G of Dimmelmeier waveforms injected into colored Gaussian vs. recolored non-Gaussian noise is shown in Figure 6.1.

All Dimmelmeier models show some decrease in efficiency above a network SNR of 7 when the waveforms are injected into realistic recolored non-Gaussian noise. This difference is most pronounced for the rapidly rotating model, s15a3o15, which has a significantly lower efficiency in recolored noise up to a network SNR of over 15.


The efficiency using both recolored non-Gaussian and Gaussian noise converges to 100% for all Dimmelmeier models above a network SNR of 20.

Yakunin


The relative efficiency for Yakunin waveforms injected into colored Gaussian vs. recolored non-Gaussian noise is shown in Figure 6.2.

Similarly to the Dimmelmeier waveforms, the efficiency for Yakunins decreases using recolored non-Gaussian noise in the 8-22 network SNR range, but converges to 1 above a network SNR of 22.


Mueller



(a) s15a2o05



(b) s15a2o09



(c) s15a3o15

Figure 6.1: Efficiency curves for each Dimmelmeier model, comparing efficiency in colored Gaussian noise in blue to re-colored realistic non-Gaussian data from a prior science (S5/VSR1) run in red.

The relative efficiency for Mueller waveforms injected into colored Gaussian vs. re-colored non-Gaussian noise is shown in Figure 6.3.

Like Dimmelmeier and Yakunin waveforms, the efficiency for Mueller waveforms decreases using re-colored non-Gaussian noise in the 8-22 network SNR range. However, at a network SNR of 10 this drop in efficiency is particularly pronounced for the N20




Figure 6.2: Efficiency curves for Yakunin injections, comparing efficiency in colored Gaussian noise in blue to re-colored realistic non-Gaussian data from a prior science run (S5/VSR1) in red.


progenitor, which produces the waveform with the longest time gap between bursts of power.

The efficiency for all Mueller injections in both types of noise also converges to 1 above a network SNR of 22.


6.3 CCSN reconstruction in recolored non-Gaussian data

Waveform reconstruction performance was also evaluated based on the noise-weighted, time optimized overlap between the reconstructed and injected waveforms, as defined in Chapter 4.


Results for the Mueller N20 model are illustrated in Figure 6.4 with a set of 2D histograms of SNR vs. time optimized waveform overlap. The top plot shows Mueller



(a) L15



(b) W15



(c) N20

Figure 6.3: Efficiency curves for each Mueller progenitor star, comparing efficiency in colored Gaussian noise in blue to re-colored realistic non-Gaussian data from a prior science run (S5/VSR1) in red.

N20 waveforms injected into recolored realistic non-Gaussian noise and the bottom shows the same set of N20 progenitor waveforms injected into colored Gaussian noise, as first shown in Chapter 4¹. These plots are representative of the behavior in non-Gaussian noise for all included CCSN waveforms.

¹A different color scheme than the one seen in Chapter 4 is used for the 2D histograms to help make the outliers more visible.

Colored Gaussian noise

Model	Maximum overlap	Ave overlap at SNR 8
Ave Dim	0.97	0.78
s15a2o05	0.98	0.71
s15a2o09	0.99	0.76
s15a3o15	0.95	0.86
Yakunin	0.91	0.42
Ave Mueller	0.92	0.57
L15	0.92	0.61
W15	0.89	0.56
N20	0.89	0.53

Table 6.1: The maximum overlap and overlap at SNR=8 is restated here for injections analyzed by cWB2G in colored Gaussian noise. This table contains an average of all injections as recovered in the L1 and H1 interferometers.

Recolored non-Gaussian noise

Model	Maximum overlap	Ave overlap at SNR 8
Ave Dim	0.98	0.73
s15a2o05	0.98	0.68
s15a2o09	0.99	0.68
s15a3o15	0.96	0.78
Yakunin	0.92	0.38
Ave Mueller	0.88	0.49
L15	0.90	0.54
W15	0.86	0.51
N20	0.83	0.38

Table 6.2: The maximum overlap and overlap at SNR=8 for injections analyzed by cWB2G in recolored non-Gaussian S5 Gaussian noise. This table contains an average of all injections as recovered in the L1 and H1 interferometers.

The bulk of injections still follow the same pattern in recolored noise, achieving nearly the same maximum overlap as in Gaussian noise at 83%. However, the average overlap at an SNR of 8 is significantly lower in recolored non-Gaussian noise; only 38% relative to 53% in colored Gaussian noise. This behavior is characteristic of all analyzed CCSN waveforms, as shown in Table 6.2.

Also seen in Figure 6.4 are extremely low overlap outliers across a wide SNR range - even up to over an SNR of 65. These are injections that coincide with transient noise that confuses the cWB2G algorithm and interferes with the accurate recovery of signal energy. This is also characteristic of the behavior for all analyzed CCSN waveforms.

An example of a waveform reconstruction of a relatively loud (SNR=17) Mueller N20 injection polluted by transient noise is shown in Figure 6.5. The time and

frequency reconstruction plots, above, show that excess power not associated with the waveform is resolved by cWB2G. This additional power is consistent in time and frequency with a coincident glitch in the LIGO Livingston detector, shown below the reconstruction plots. This example demonstrates how a glitch in one interferometer of a three detector network may pollute astrophysical parameter estimation of CCSN.

It should also be noted that for this recolored realistic noise study cWB2G produced some reconstructions with very low overlap with the corresponding injected waveform that are not understood. These reconstructions show unphysical behavior such as more resolved bursts than are apparent in the combination of the injected data and the noise, or reconstructed time series that are flat. Less than 5% of reconstructed events fit into this category. Future work is to understand these anomalies and why realistic non-Gaussian noise might induce this behavior in cWB2G.

6.4 Summary and future work

The use of realistic recolored noise in place of Gaussian noise significantly reduces the rate of recovery for all analyzed CCSN waveforms, as expected.

Waveform reconstruction analysis indicates that if a gravitational wave burst signal aligns with a transient event by chance, cWB2G will not be able to recover the signal power accurately, even for a loud event. Limiting the glitch rate at all interferometers in the global network will be crucial to ensuring accurate burst parameter estimation.

For a future study, it would be a very interesting test of BayesWave’s signal vs. glitch classification ability to investigate whether CCSN events polluted by transient noise are classified as signal, glitch, or perhaps signal plus a glitch.

Additionally, the wide clustering windows in time and frequency used for cWB2G (detailed in Chapter 4) likely admit unrelated transient noise pollution for some

events. These windows may be narrowed to limit this effect and still allow cWB2G to recover the event power, particularly for analysis of the shorter Dimmelmeier and Yakunin waveforms.




Figure 6.4: 2D histograms of SNR vs. time optimized waveform overlap for Mueller N20 injections. The top plot shows waveforms injected into recolored realistic non-Gaussian noise and the bottom shows the same set of waveforms injected into colored Gaussian noise. The colorscale indicates the number of recovered events per 2D overlap-SNR bin. These results were optimized over the interferometer network such that for each recovered event, the waveform overlap and single interferometer SNR of only the interferometer with the loudest signal is represented.

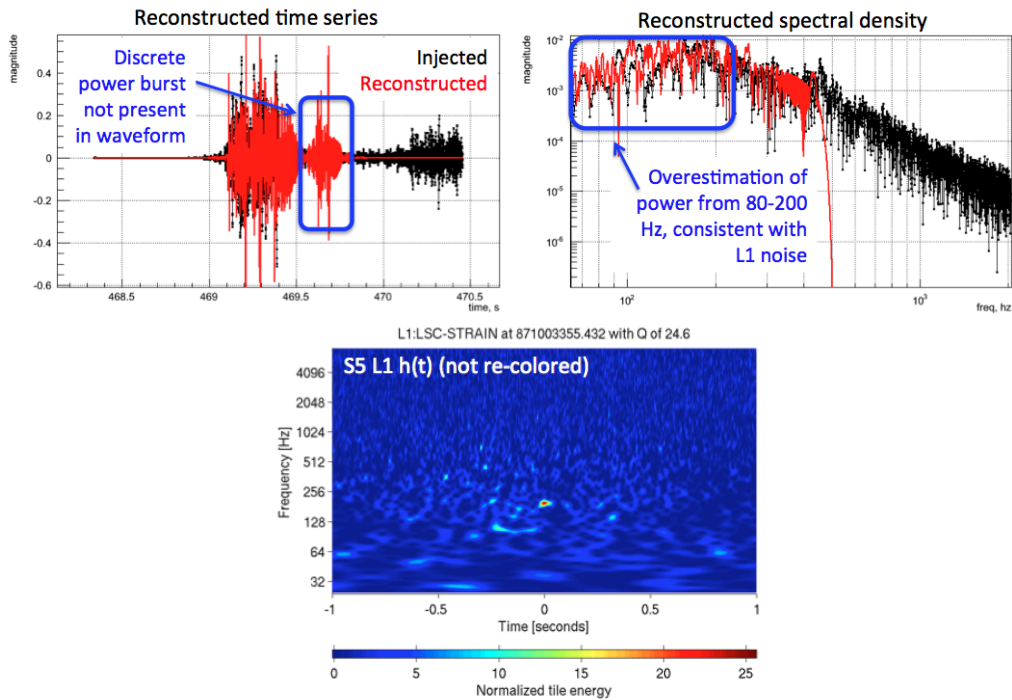


Figure 6.5: An example of a noise transient in a single detector polluting a waveform reconstruction. Above are two plots depicting the time and frequency reconstruction by cWB2G of a loud (17 SNR) Mueller N20 progenitor waveform injected into aLIGO and AdVirgo recolored noise from LIGO science run 5 (S5) and Virgo Science Run 1 (VSR1). Injected power is shown in black and reconstructed event power is shown in red. On the left, the time reconstruction shows a burst of discrete power that is not present in the original waveform. This burst is associated with the overestimation of power from 80-200Hz shown in the frequency reconstruction on the right. Below is an Omegascan of the $h(t)$ channel of LIGO Livingston at the time of the injection, showing a glitch consistent with the excess power in time and frequency.

CHAPTER 7

CONCLUSIONS AND FUTURE PROSPECTS

Depending on the orientation and internal conditions, a supernova signal must be quite nearby to produce a detectable signal, even assuming Advanced LIGO and Advanced Virgo design sensitivity curves and ideal Gaussian noise conditions. Chapter 4 showed the most detectable analyzed waveforms, the wavelet-like 2D rotating core-collapse Dimmelmeier models, are recovered up to 10 kpc, which includes the center of the Milky Way. The most realistic analyzed waveforms, the 3D neutrino-driven Mueller models, are only detectable up to 1 kpc away.

The analysis of waveform reconstruction performance in Chapter 4 projects the ability of the flagship burst search, cWB2G, to recover waveforms using the Advanced interferometers under ideal Gaussian noise conditions. At a realistic SNR of 8, 3D neutrino-driven Mueller waveforms were recovered with an average overlap between injected and reconstructed waveforms of 57%, 2D neutrino-driven Yakunin waveforms with 42%, and 2D rotating core-collapse Dimmelemeier waveforms with 78%. Even at high SNR, reconstructions of the neutrino-driven models tend to omit key waveform features associated with interesting physical processes, such as delayed bursts of gravitational wave energy due to proto-neutron star convection.

The Bayesian follow-up algorithm BayesWave has more flexibility in the placement and parameters of the wavelets used to fit the signal than the time-frequency tiled cWB2G algorithm, as described in Chapter 4. BayesWave does tend to achieve better

waveform reconstruction of CCSN injections at higher, astrophysically unrealistic SNRs. However, at a realistic SNR of 8, BayesWave tends to produce a lower overlap between reconstructed and injected waveforms for the more complex, neutrino-driven waveforms. BayesWave achieved an average overlap of 49% for Mueller waveforms, 0% for Yakunin waveforms, and 84% for wavelet-like Dimmelmeier waveforms at an SNR of 8.

The injection of the same CCSN waveforms into realistic detector noise from a prior LIGO and Virgo joint science run (S5/VSR1) as detailed in Chapter 6 decreased the detection efficiency in the network SNR range 8-20 for all analyzed waveforms. In the most extreme case, the Mueller N20 progenitor model, efficiency was reduced by 20% at a network SNR of 10.

The introduction of recolored non-Gaussian noise also impacted the average recovered overlap at an SNR of 8, as summarized in Table 6.2. For the same Mueller N20 model, the average overlap between injected and recovered waveforms was reduced from 53% to 38% at an SNR of 8. Chance coincidence of an injected supernova event with excess transient noise in any detector had a detrimental impact on the waveform reconstruction, even for high SNR events, as illustrated in Figures 6.4 and 6.5.

Waveform reconstruction is fundamentally limiting to any study of gravitational waves produced by CCSN, including parameter estimation and model selection. This work identifies two major arenas that should be improved in order to confidently extract interesting physics from a CCSN gravitational wave signal: the recovery of CCSN by burst waveform reconstruction algorithms, and the mitigation of transient detector noise.

To that effect, there are promising ongoing efforts to improve performance on both fronts.

The introduction of upgraded seismic isolation instrumentation combats transient ground motion, known to be associated with the majority of high SNR noise transients in the last LIGO science run. Active seismic isolation new to Advanced LIGO was shown in Chapter 5 to reduce the amplitude and rate of transient motion between the ground and the optics tables. Although even with this improvement, ground motion has been shown to impact the $h(t)$ channel during periods of elevated ground motion at the Advanced LIGO sites. Future work will investigate tuning the active control loops to better mitigate transient motion events. The work should inform the development of additional hardware upgrades intended to further improve the sensitivity of the Advanced LIGO instruments, currently in progress [65].

The refinement of ETG performance described in Chapter 3 will enable the design and testing of effective DQ vetoes to improve the performance of astrophysical searches on Advanced LIGO data. The use of ETGs in the characterization of transient noise will also help guide commissioning efforts to target noise sources limiting to gravitational wave detection and astrophysical parameter estimation. As shown in Chapters 5 and 6, high rates of transient noise, even at relatively low SNR, will have a significant impact on CCSN detectability and waveform reconstruction, as well as all other transient burst searches. This must be a top priority for detector characterization efforts.

There is also a great potential for improvement in BayesWave’s ability to reconstruct CCSN waveforms with the design and implementation of a set of priors for CCSN models. BayesWave may also be able to distinguish between injected CCSN events and noise events using signal classification, which should be explored in the future. Additionally, cWB2G performance may improve by applying a sky mask that limits the potential sky position of any recovered event to within a location range identified by EM observations [93].

Future iterations of the waveform reconstruction performance study detailed in this dissertation should target more recent CCSN modeled waveforms, which may provide more nuanced insight for these tuning efforts.

In general, studying and improving the reconstruction of core-collapse supernovae waveforms from realistic interferometer data has a much broader beneficial impact on the extraction of physics from other gravitational wave progenitors that produce longer, richly complex waveforms; perhaps including as of yet unknown burst signal sources.

In conclusion, significant improvements to software are required in order to achieve a good probability of accurately recovering the gravitational wave signal produced by a nearby core-collapse supernova with Advanced LIGO and Advanced Virgo. Additionally, a high rate of noise transients in any single interferometer will be extremely limiting to detection confidence and parameter estimation of these events. However, the future is bright, with many promising improvements in sight.

APPENDIX A

FULL WAVEFORM RECONSTRUCTION RESULTS FOR CWB2G AND BAYESWAVE

The appendix is supplementary to Chapter 4, which detailed the CCSN waveform reconstruction performance of cWB2G and BayesWave.

Here are additional plots showing the full results for each considered waveform analyzed by both algorithms, in separate plots so that it is easier to see the detailed distribution in overlap vs. SNR for each model.

The first set of plots, Figure A.1, shows 2D heatmaps of all injections detected by cWB2G. In the top row are reconstructed overlap performance heatmaps for each Dimmelmeier model. In the center row are plots showing this behavior for all three progenitor stars used for the 3D neutrino-driven CCSN Mueller waveforms. The plot in the bottom row shows results for the 2D neutrino-driven Yakunin model.

This series of heatmap plots shows the time-optimized overlap between injected and recovered waveforms versus single interferometer SNR. These results were optimized over interferometers L1, H1, and V1 (the two LIGO interferometers and the Virgo interferometer) in that only the overlap and corresponding SNR is plotted for the interferometer with the highest SNR for that injection. The color scale represents the number of recovered injections per 2D overlap-SNR bin.

cWB2G reconstruction of CCSN waveforms: optimized overlap vs. SNR





Figure A.1: Heatmaps of the distribution of CCSN events reconstructed by cWB2G showing reconstructed waveform overlap optimized over local interferometer SNR. Plots a-c show Dimmelmeier events, with different rates of core and differential rotation. Plots d-f show Mueller events, with different progenitor stars. Plot f shows Yakunin events.

The second set of plots, Figure A.2, shows overlap vs. SNR results for the subset of event analyzed by BayesWave for each model. The recovered value for the overlap between injected and recovered waveforms for each injection is represented with a circle, and bounded by triangles indicating the 90% confidence interval of the result.


Instead of optimizing over SNR as for the cWB2G results, the overlap and SNR are shown for waveforms as recovered in the Hanford detector, in red, and the Livingston detector, in green ¹.

¹V1, or Virgo, is omitted because for the unpolarized Mueller waveforms, BayesWave's assumption of elliptical polarization adversely affected results for this interferometer.


BayesWave reconstruction of CCSN waveforms: optimized overlap vs. SNR




(a) s15a2o05




(b) s15a2o09




(c) s15a3o15




(d) L15



(e) W15



(f) N20



(g) Yakunin

Figure A.2: Heatmaps of the distribution of a subset of CCSN events reconstructed by BayesWave. As before, plots a-c show Dimmelmeier events, with different rates of core and differential rotation. Plots d-f show Mueller events, with different progenitor stars. Plot f shows Yakunin events.

APPENDIX B

SUPPLEMENTARY ANALYSIS OF BURST PARAMETER RECOVERY

The appendix is supplementary to Section 3.2, which detailed the performance of five Event Trigger Generators (ETGs).

Here are additional plots showing the recovery of bulk parameters by different metrics, including a study of recovered parameter accuracy versus SNR.

The first set of supplementary recovered vs. injected frequency for injections detected by each ETG. Figure B.1 illustrates the behavior for sine-Gaussians and Figure B.2 for white noise bursts.

The artifacts and behavior described in Section 3.2 can be seen here, but the scope of the axes puts the overall performance in a broader perspective. For example, for sine-Gaussian injections most ETGs recover the frequency to within or very near 10% for most events.

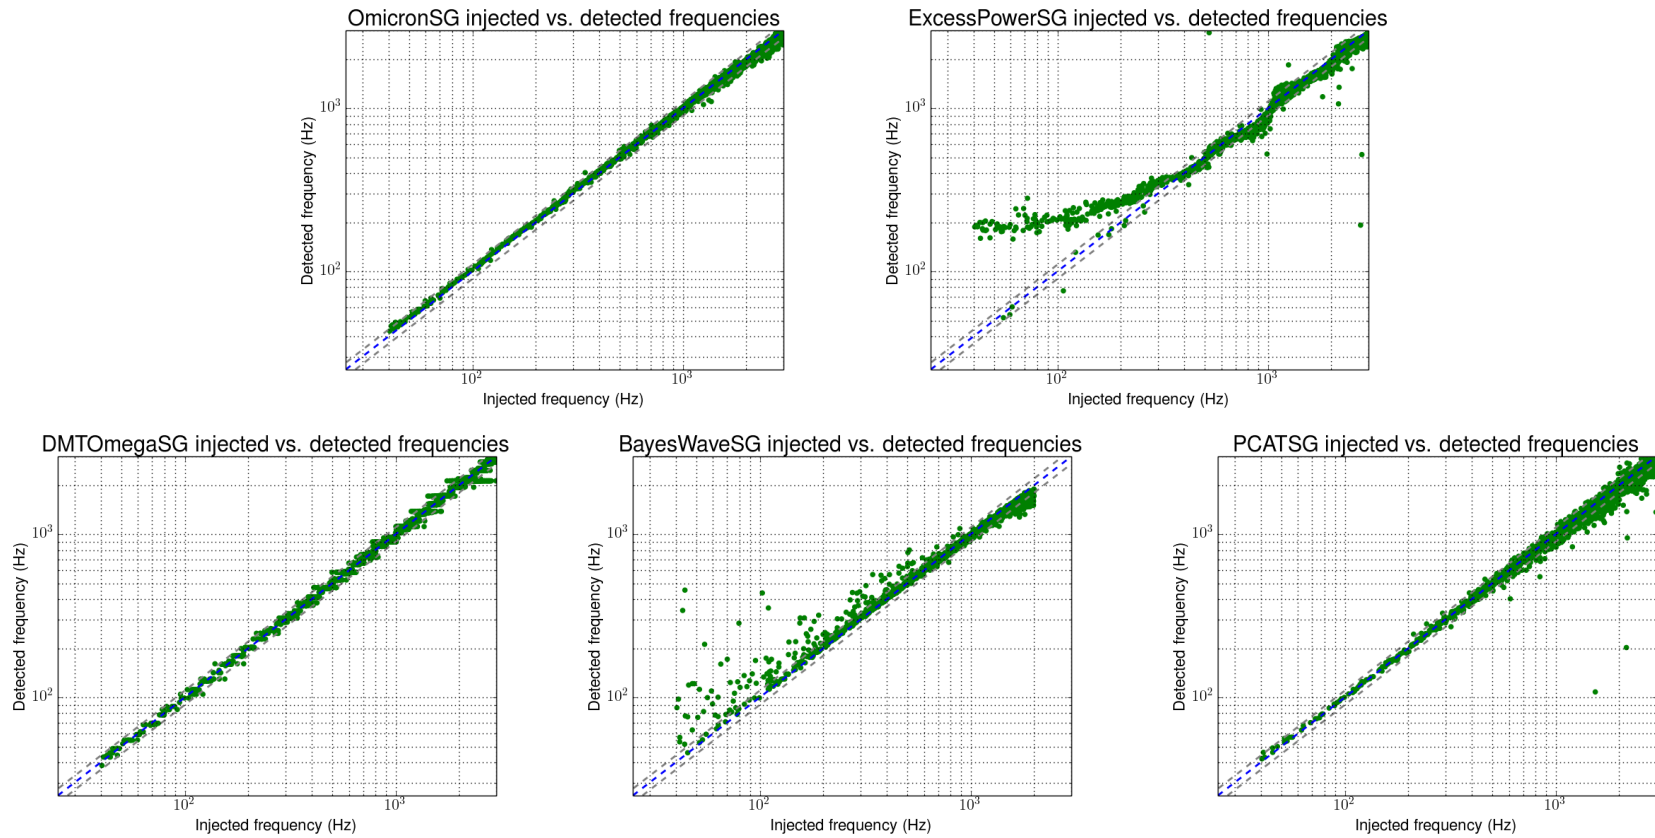


Figure B.1: A series of scatter plots showing the injected frequency versus detected frequency of recovered sine-Gaussian events for each ETG. The blue dashed line represent the ideal injected frequency = recovered frequency. The grey dashed lines represent the area within 10% of ideal performance.




Figure B.2: A series of scatter plots showing the injected frequency versus detected frequency of recovered white noise burst events for each ETG. The blue dashed line represent the ideal injected frequency = recovered frequency. The grey dashed lines represent the area within 10% of ideal performance.

The second set of plots depict the difference between recovered and injected frequency vs. the injected event SNR. Figure B.3 illustrates this for sine-Gaussian events and Figure B.4 illustrates it for white noise bursts.

It might be expected that the recovery of frequency becomes increasingly accurate with event SNR, but this does not appear significant for any ETG for either sine-Gaussian or white noise burst waveforms. Most frequency difference distributions appear more or less uniform across SNR.

The small exceptions are Omicron and BayesWave, which both produce fewer overestimated frequencies for sine-Gaussian waveforms at higher SNRs. The broad frequency difference distributions also appear to narrow at higher SNRs for these ETGs using white noise burst injections.

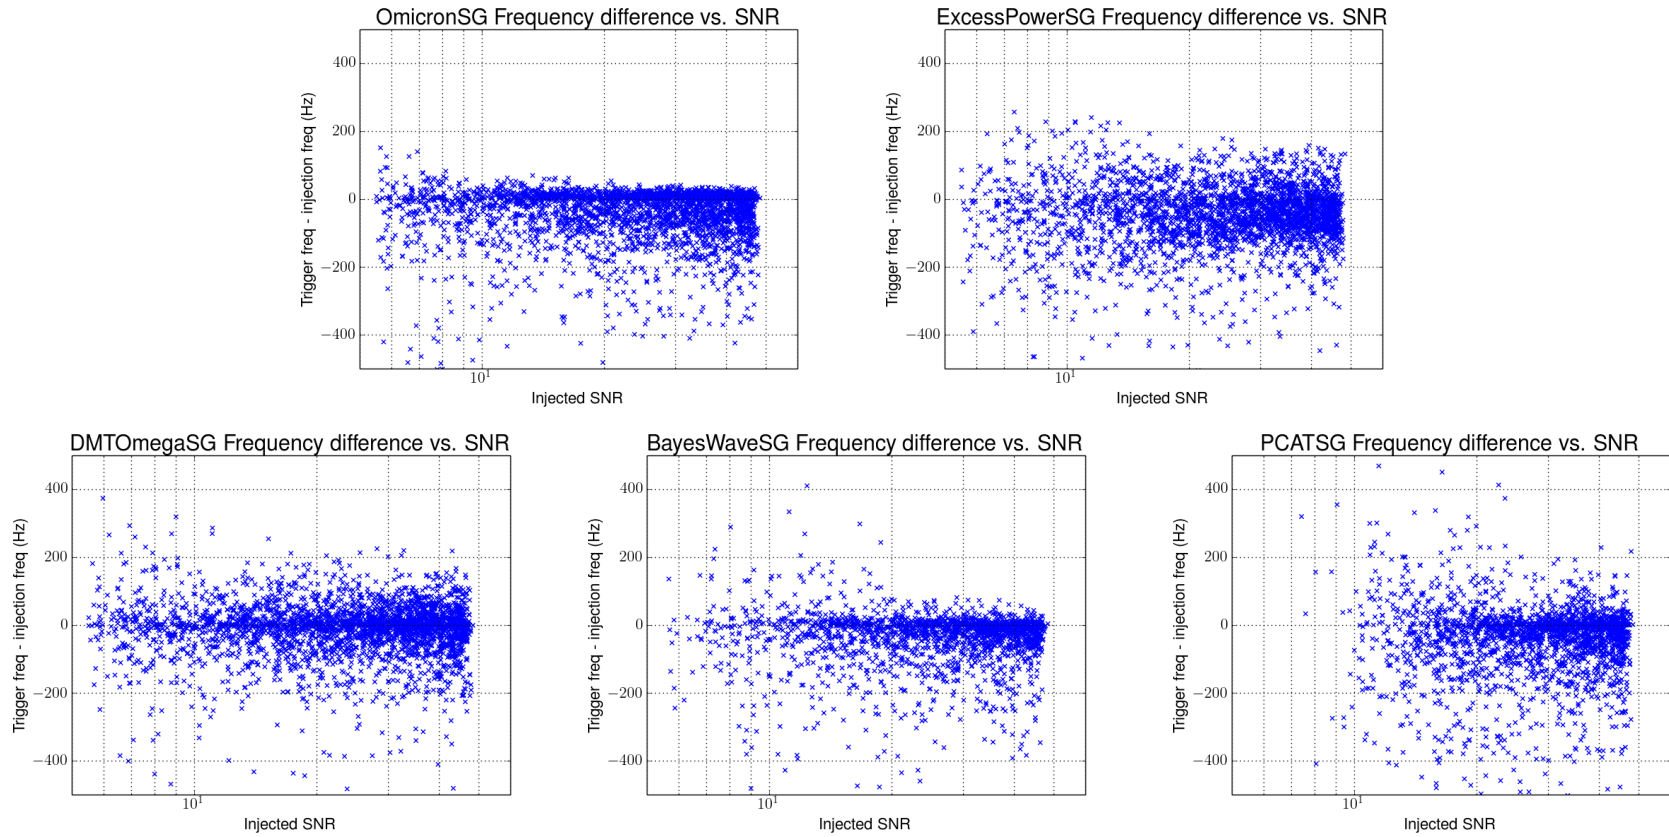


Figure B.3: A series of scatter plots showing the difference in recovered and injected frequency vs. injected SNR for each ETG using only sine-Gaussian injections.

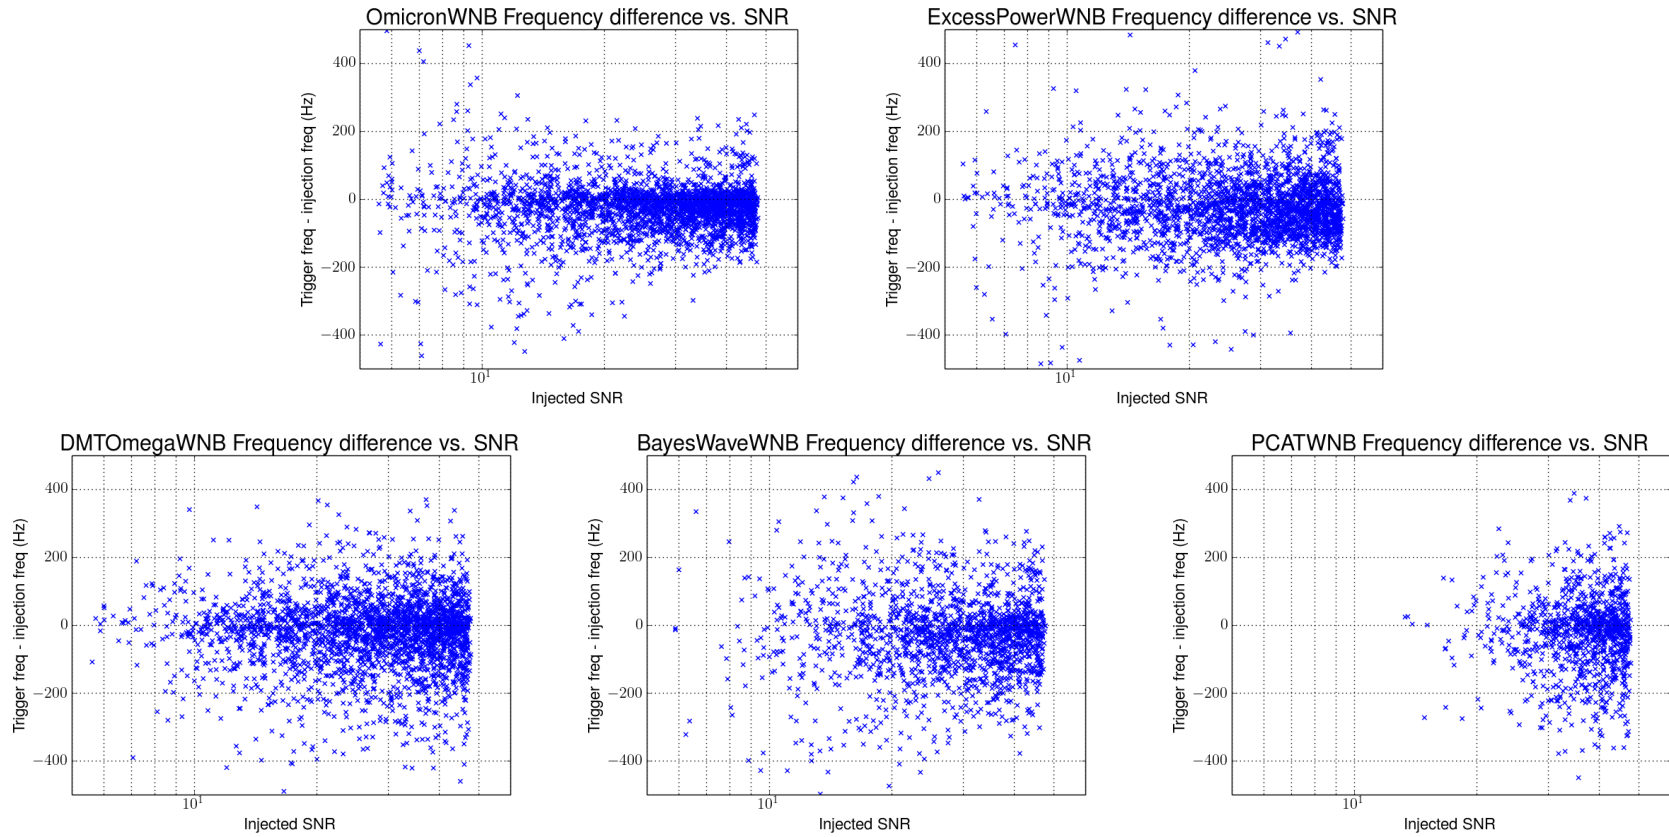


Figure B.4: A series of scatter plots showing the difference in recovered and injected frequency vs. injected SNR for each ETG using only white noise burst injections.

The third series of plots shows the timing difference versus injected SNR for each ETG. Figure B.5 illustrates this for sine-Gaussian injections and Figure B.5 for white noise bursts.

Again, it may be expected that the timing of recovered events would improve with injected SNR.

ExcessPower produces fewer large timing offsets for sine-Gaussian waveforms at higher SNR.

The timing difference distributions of Omicron and BayesWave using white noise bursts also seem to grow more compact at higher injected SNR.

However, for PCAT, timing appears to become less accurate at higher SNR for both sine-Gaussians and white noise bursts. This is not yet understood.

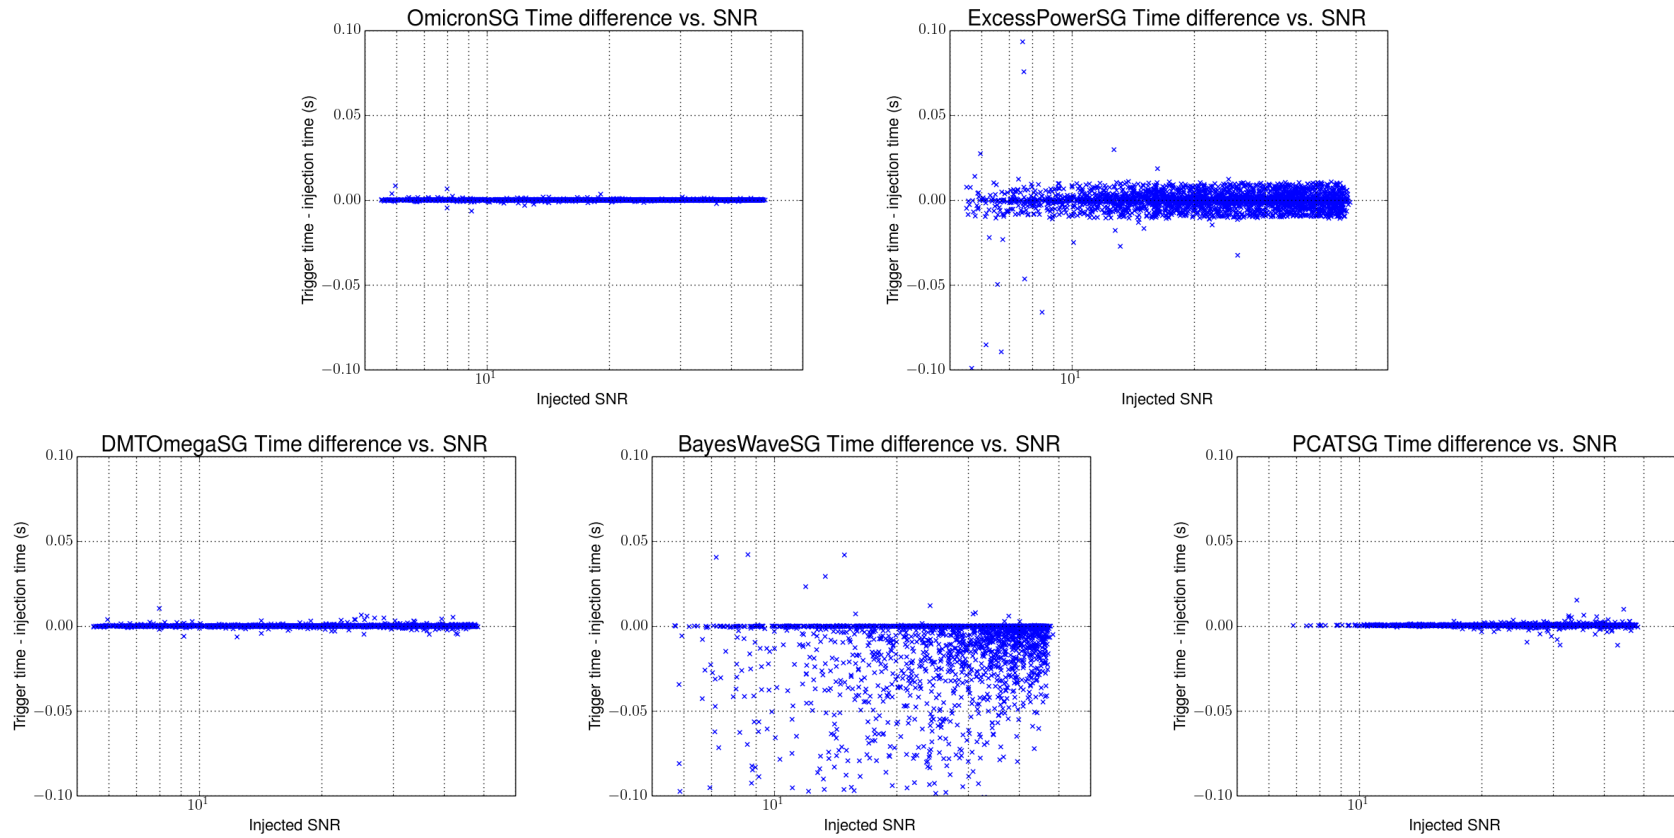


Figure B.5: A series of scatter plots showing the difference between recovered and injected time vs. injected SNR for each ETG using only sine-Gaussian injections.

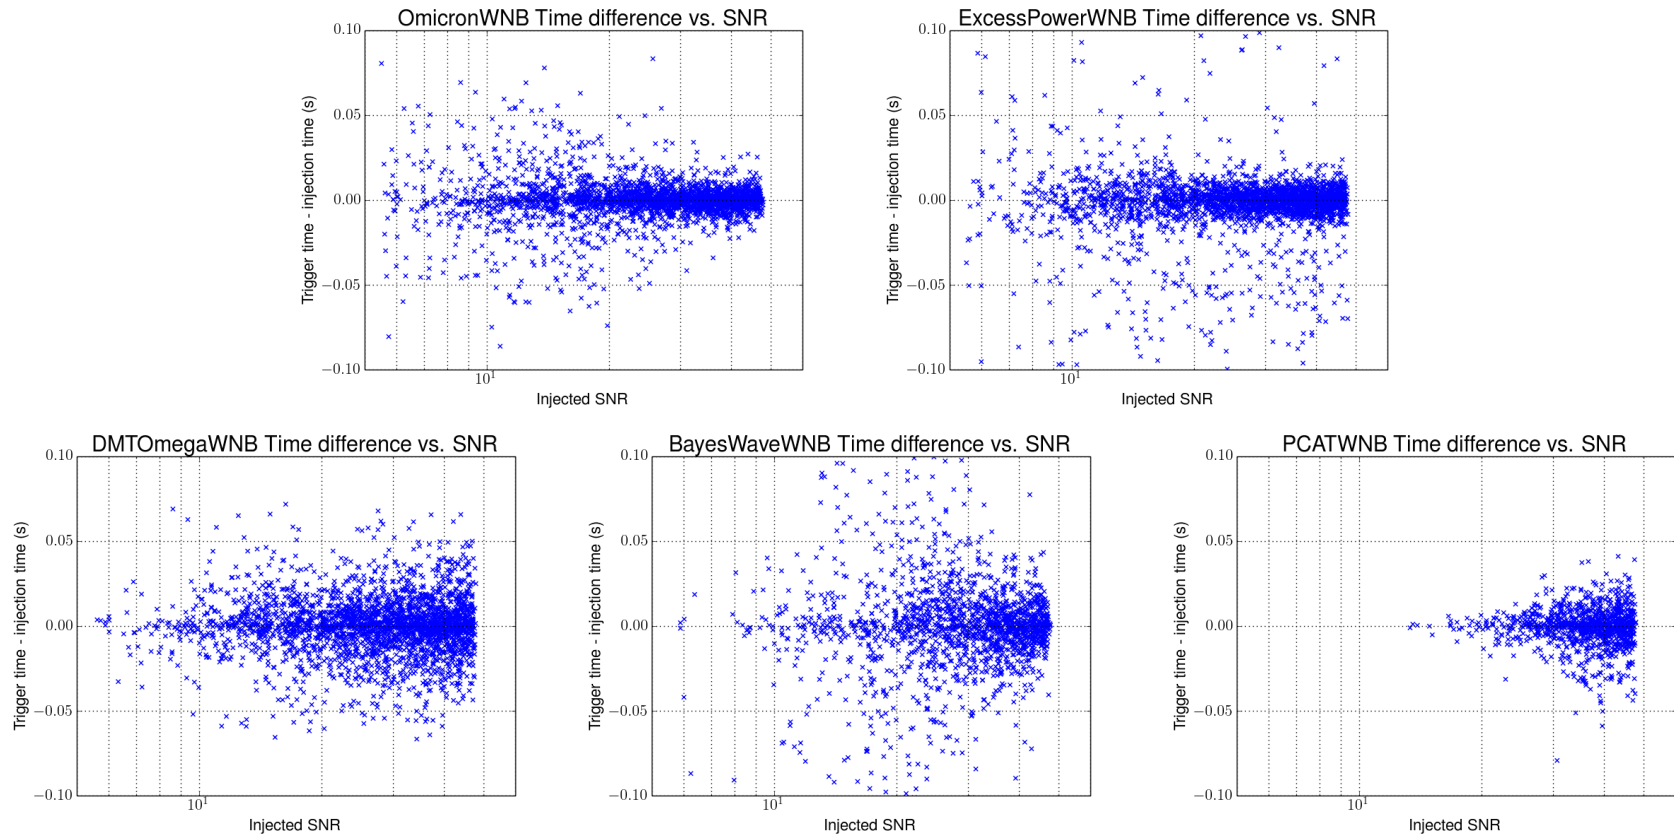


Figure B.6: A series of scatter plots showing the difference between recovered and injected time vs. injected SNR for each ETG using only white noise burst injections.

The final series of plots show recovered vs. injected SNR. The blue dashed lines represent the ideal behavior (recovered SNR = injected SNR), and the grey dashed lined indicated +/- a factor of $\sqrt{2}$.

The sine-Gaussian basis ETGs Omicron and DMT-Omega and wavlet-basis algorithm BayesWave have fairly precise SNR estimations, with Omicron and BayesWave tending to overestimate SNR and DMT-Omega tending to underestimate SNR. ExcessPower's estimation of SNR is far from the ideal behavior, which was known before this study. PCAT tends to greatly, but consistently, underestimate the SNR of sine-Gaussian waveforms.

All ETGs have significantly broader distributions in recovered SNR relative to injected SNR for white noise bursts, with the exception of ExcessPower, which exhibits similar behavior to sine-Gaussian SNR recovery far from the ideal. The sine-Gaussian basis ETGs Omicron and DMT-Omega tend to significantly underestimate the SNR of white noise bursts, which is likely due to the way SNR is reported for clustered triggers. BayesWave shows the most precise and accurate recovery of white noise burst injection SNR.

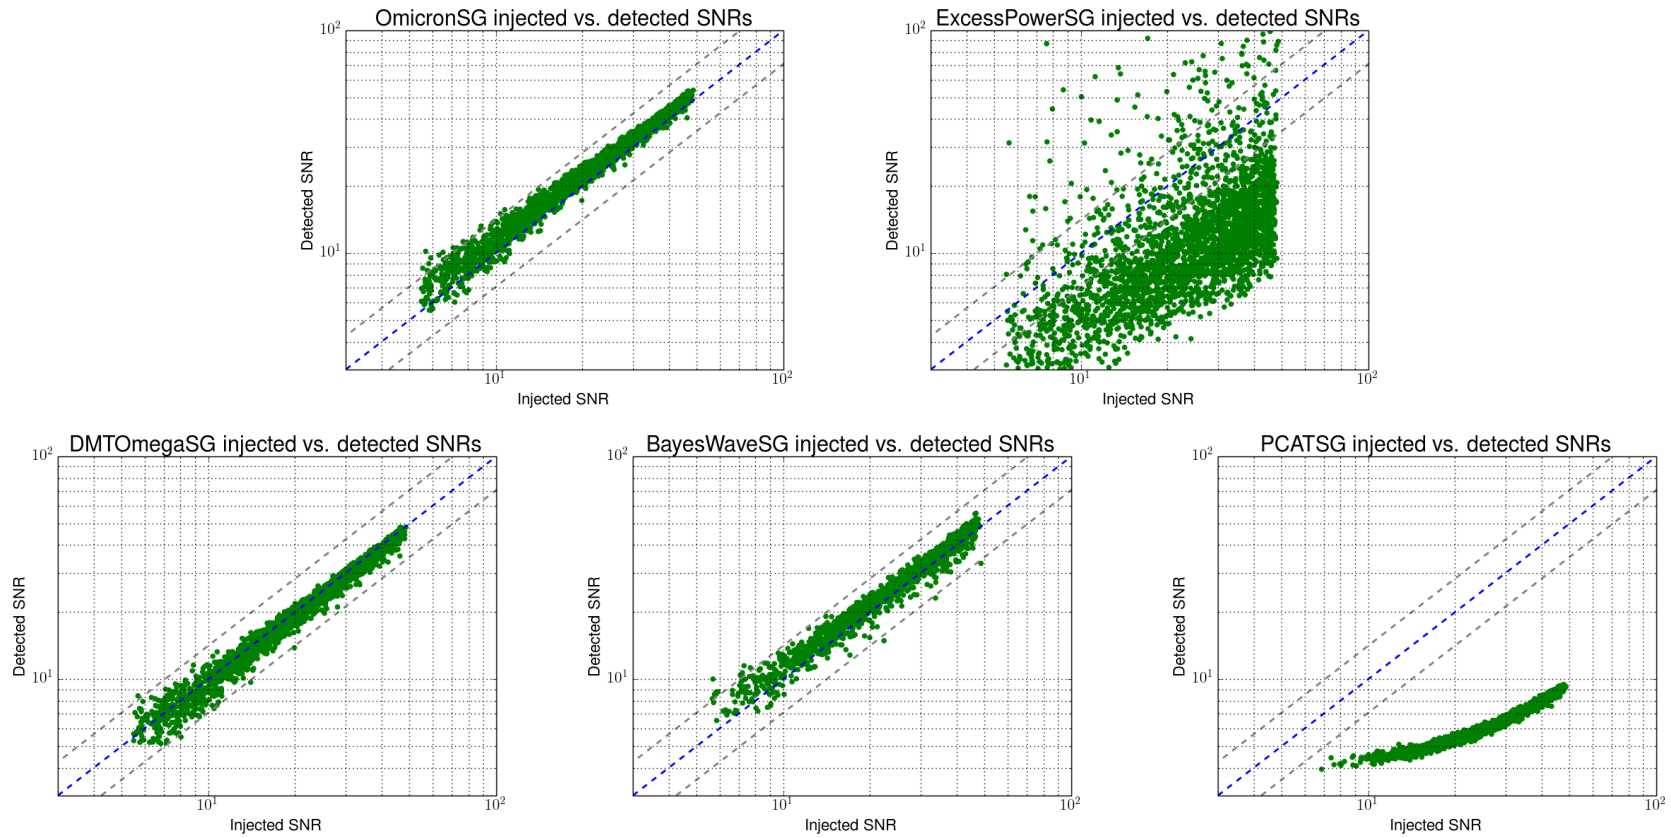


Figure B.7: A set of scatter plots showing recovered SNR vs. injected SNR for every sine-Gaussian event recovered by each ETG. The grey dashed lines indicate \pm a factor of $\sqrt{2}$, which is a common difference in SNR definition between ETGs, and serve as a guide for the eye. The blue dashed line represents the ideal behavior, where the recovered SNR is equal to the injected SNR.




Figure B.8: A set of scatter plots showing recovered SNR vs. injected SNR for every white noise burst injection recovered by each ETG. The grey dashed lines indicate \pm a factor of $\sqrt{2}$, which is a common difference in SNR definition between ETGs, and serve as a guide for the eye. The blue dashed line represents the ideal behavior, where the recovered SNR is equal to the injected SNR.

APPENDIX C

MAXIMUM ENTROPY AND CCSN WAVEFORM RECONSTRUCTION

C.1 Overview

Summerscales et. al. [91] details the application of the principle of maximum entropy to the analysis of gravitational waves (GWs) produced by core-collapse supernovae (CCSN). The authors motivate a Bayesian approach to extracting an unmodeled GW signal from unpredictably noisy data, describe the MaxEnt technique, show examples of its application to simulated CCSN waveforms injected into the noisy signal of a network of GW detectors, and discuss the physics of CCSN that could be inferred from GW signals reconstructed in this way.

C.2 Introduction to maximum entropy : a Bayesian approach

C.2.1 The principle of maximum entropy

The *principle of maximum entropy* states that when many probability densities functions satisfy the constraints on our system, of these we should choose the probability density function with the most entropy S , where entropy is defined as [87]:

$$S = - \sum_{i=1}^M p_i \log(p_i) \quad (\text{C.1})$$

where p_i is the probability distribution of for each of M set of states.

C.2.2 The monkey example

Why is this true? Here is a brief argument following Sivia [87].

Suppose there are M distinct possibilities (X_i) to be considered. How do we ascribe truth-values to them (i.e. some PDF $p_i = \text{prob}(X_i|I)$), given some testable information I ?

- Have a bunch of monkeys throw a large amount of coins into a set of M boxes representing the possibilities X_i .
- Collectively, the fraction of coins found in each box represents a distribution, and every distribution not meeting our constraints (our known/measured testable info) is discarded.
- The distribution that occurs most often is the sensible choice for our PDF, $\text{prob}(X_i|I) = p_i$.

For M total boxes and N total coins such that,

$$N = \sum_{i=1}^M n_i \quad (\text{C.2})$$

where n_i is the number of coins in box i , we have a probability to describe each possibility X_i , $p_i = \frac{n_i}{N}$, that occurs with some expected frequency $F(p_i)$.

$$F(p_i) = \frac{K}{M^n} \quad (\text{C.3})$$

Where K is the number of ways of obtaining n_i for some box i , and M^n is the number of different ways of scattering N coins among M boxes.

$$K = \frac{N!}{n_1!n_2!\dots n_m!} \quad (\text{C.4})$$

Taking the log,

$$\log(F) = -N\log(M) + \log(N!) - \sum_{i=1}^M \log(n_i)! \quad (\text{C.5})$$

Use the Stirling approximation to reduce the last two terms:

$$\log(F) = -N\log(M) + N\log(N) - \sum_{i=1}^M n_i \log(n_i) \quad (\text{C.6})$$

Substituting $p_i = \frac{n_i}{N}$ and using $\sum p_i = 1$,

$$\log(F) = -N\log(M) - N \sum_{i=1}^M p_i \log(p_i) \quad (\text{C.7})$$

Now we see that maximizing how often the monkeys obtain a candidate PDF

$$\log(F) = -N\log(M) - N \sum_{i=1}^M p_i \log(p_i) \quad (\text{C.8})$$

is equivalent to maximizing the entropy, S

$$S = - \sum_{i=1}^M p_i \log(p_i) \quad (\text{C.9})$$

Because this quantity, S , looks just like thermodynamic entropy, it is known as *information entropy*.

C.3 The MaxEnt Technique

C.3.1 Overview

A major problem for unmodeled searches for transient gravitational waves is extracting an unknown signal from noise. The MaxEnt algorithm [91] is designed to infer a gravitational wave signal from the time series data of a network of gravitational wave detectors.

Formulating this problem in a notation that closely follows [91],

$$\mathbf{d} = \mathbf{R}\mathbf{h} + \mathbf{n} \tag{C.10}$$

where \mathbf{d} is the measured time series data of a detector, \mathbf{R} is the detector's response, \mathbf{h} is the strain we seek, and \mathbf{n} is the unpredictable detector noise.

We want to determine \mathbf{h} without assuming anything about detector noise \mathbf{n} or the form of the GW signal \mathbf{h} . To this end, we can find a probability distribution f that \mathbf{h} is a GW signal given \mathbf{d} , \mathbf{R} , noise characterization N , and some limited knowledge of GWs (aside from their form) I :

$$f(\mathbf{h}'|\mathbf{d}, \mathbf{R}, N, I) \tag{C.11}$$

We want to find \mathbf{h}' , the best guess of the incident GW strain, that maximizes this probability density function f .

Using Bayes' law, we can write Equation C.11 as:

$$f(\mathbf{h}'|\mathbf{d}, \mathbf{R}, N, I) = \frac{g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I)u(\mathbf{h}'|I)}{v(\mathbf{d}|\mathbf{R}, N, I)} \quad (\text{C.12})$$

which describes the probability density function f in terms of: 1) the likelihood g for \mathbf{d} given h' , 2) prior knowledge u of gravitational waves as applied to \mathbf{h}' , and 3) the posterior v , which is the probability of observing \mathbf{d} given \mathbf{R} , N , and I .

Ignoring the posterior normalization $v(\mathbf{d}|\mathbf{R}, N, I)$, which is the same is regardless of our guess of \mathbf{h}' , now our problem is to determine the form of the likelihood $g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I)$, and the form of the prior $u(\mathbf{h}'|I)$. Once these are known, the MaxEnt algorithm's task is to find the \mathbf{h}' that maximizes their product. To find the forms of g and u , we use the principle of maximum entropy, where entropy is defined as in Equation C.9 from Sivia [87].

C.3.2 Determining the likelihood

The principle of maximum entropy tells us that the least presumptive distribution is a Gaussian when nothing more than the noise distribution's mean and covariance are known. Here is a short derivation, following section 5.3 of Sivia [87].

Let's say we have some distribution of variable x , and we happen to know (perhaps by measurement) it's mean, μ , and the variation, σ , defined by:

$$\sigma^2 = \langle (x - \mu)^2 \rangle = \int (x - \mu)^2 dx \quad (\text{C.13})$$

and we also know that the sum of the probability of all possible states of x must equal 1:

$$\sum_i p_i = 1 \quad (\text{C.14})$$

To find the probability of x given the constraints of equations C.13 and C.14, we maximize the entropy of the probability of x being in state i , p_i , over all possible states i .

$$S = - \sum_i p_i \log(p_i) + \lambda_0 \left(1 - \sum_i p_i \right) + \lambda_1 \left(\sigma^2 - \sum_i (x_i - \mu)^2 p_i \right) \quad (\text{C.15})$$

where the latter two terms are Lagrangian multipliers governed by equations C.13 and C.14.

To maximize S , we set $\frac{\partial S}{\partial p_i} = 0$ and solve for p_i to get:

$$p_i = e^{-(1+\lambda_0)} e^{-\lambda_1(x_i-\mu)^2} \quad (\text{C.16})$$

We can substitute equation C.16 into equations C.13 and C.14 in the continuum limit (from $-\infty$ to $+\infty$) to solve for the normalization constant $e^{-(1+\lambda_0)}$ and the constant λ_1 . We then see the probability of x when μ and σ are known takes the form of a Gaussian distribution:

$$P(x|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (\text{C.17})$$

Therefore, the principle of maximum entropy tells us that the likelihood we are interested for our signal extraction, knowing information analogous to the mean and variance of the interferometer noise, has the following form:

$$g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I) = \frac{1}{\sqrt{2\pi} \|C\|} e^{\frac{1}{2}(\mathbf{d}-\mathbf{R}\mathbf{h})^T C^{-1}(\mathbf{d}-\mathbf{R}\mathbf{h})} \quad (\text{C.18})$$

where $(\mathbf{d} - \mathbf{R}\mathbf{h})$ is equivalent to the detector noise, by Equation C.10 and C is the covariance matrix, corresponding to the autocorrelation of the interferometer noise.

MaxEnt calculates the covariance matrices for each of the three (or more than two) interferometers by analyzing the amplitude of each samples (where the number of samples, N , is determined by the length of the data and the sampling rate). By assigning the value of each sample to an entry in a $1 \times N$ matrix, the covariance matrix is filled in with the entry (i,j) being the computed covariance between the i^{th} and j^{th} entries.

The MaxEnt code assumes the mean of the noise distribution is zero, which is generally a valid assumption, and calculates the covariance via these lines (in MATLAB):

```
% make the signal a column matrix
signal = signal(:);

% calculate the number of averages
nAverage = length(signal)-n+1;

% initialize covariance matrix
covMatrix = zeros(n);

% compute the covariance matrix
for iAverage = 1:nAverage
    signalSegment = signal(iAverage:iAverage+n-1);
    covMatrix = covMatrix + signalSegment*signalSegment';
end
```

```
% normalize covariance matrix
covMatrix = covMatrix/nAverage;
```

C.3.3 Determining the prior

Again following Summerscales et al [91], the first step in determining the prior is enumerating the assumptions that we make about our GW signal:

Assumptions about GWs:

1. We cannot assume any knowledge of sky location (i.e. $h(t)$ is invariant under a rotation of x or $+$ polarization)
2. There is no favored arrival time of GWs (i.e. $h(t) = h(t + \tau)$)
3. There is no favored degree of smoothness of waveforms

From General Relativity and assumptions 1-3, we know we may write strain as a Fourier expansion:

$$h_k = \sum_j A_j \cos(\omega_j t_k) + B_j \sin(\omega_j t_k) \quad (\text{C.19})$$

which could also be written as

$$h_k = \sum_j a_j (\omega_j t_k + \theta_j) \quad (\text{C.20})$$

$$a_j^2 = A_j^2 + B_j^2 \quad (\text{C.21})$$

$$\tan \theta_j = \frac{B_j}{A_j} \quad (\text{C.22})$$

We also note equation 5.38 of Sivia, [87], which shows that the overall prior is the product of priors of identical form such that

$$P'(\mathbf{a}, \theta | I) = \prod_k P'(a_k, \theta_k | I) \quad (\text{C.23})$$

From our second assumption about GWs, any θ is equally likely, yielding

$$P'(a, \theta | I) = \frac{p'(a | I)}{2\pi} \quad (\text{C.24})$$

and whether we write the Fourier expansion as in equation C.19 or equation C.20 should make no difference to the prior, so we can set these prior forms $Q(A_k, B_k)$ and $P(a, \theta)$ equal to each other:

$$Q(A|I)Q(B|I)dAdB = P(a, \theta|I)dad\theta \quad (\text{C.25})$$

Substituting equation C.24 for $P(a, \theta|I)$:

$$Q(A|I)Q(B|I)dAdB = \frac{p(a|I)}{2\pi}dad\theta \quad (\text{C.26})$$

We can find the form of the prior on A by taking the special case that B=0:

$$Q(A|I)Q(0|I)dAdB = \frac{p(a|I)}{2\pi}dad\theta \quad (\text{C.27})$$

and substituting $p(a|I)dad\theta = 2\pi Q(A|I)Q(0|I)dAdB$ into equation C.26 and a for equation C.21 we see

$$Q(A|I)Q(B|I) = Q(\sqrt{A^2 + B^2})Q(0|I) \quad (\text{C.28})$$

The solution is priors of the form:

$$Q(A|\sigma^2, I) = \frac{e^{-\frac{A^2}{2\sigma^2}}}{\sqrt{2\pi\sigma^2}} \quad (\text{C.29})$$

For some unknown parameter, σ .

We can repeat the same process to see that priors on B, $Q(A|\sigma^2, I)$ have an equivalent form. We also know from Sivia that equation C.23 for P is analogous to Q [87]. Thus,

$$Q(\mathbf{A}, \mathbf{B}|I) = \prod_k Q'(A_k|I)Q'(B_k|I) \quad (\text{C.30})$$

Therefore our prior, u , can be expressed as a family of priors, $u_\sigma(\mathbf{h}|I)$,

$$u_\sigma(\mathbf{h}|I) = \prod_k u'_\sigma(A_k|I)u'_\sigma(B_k|I) = \frac{e^{-\left(\frac{1}{2\sigma^2}\right)\sum_{k=0}^{N-1} h_k^2}}{(2\pi\sigma^2)^{N/2}} \quad (\text{C.31})$$

summed over N data samples.

Now we can see that we need to find an \mathbf{h}' that maximizes a posterior probability density proportional to a likelihood of the known form in equation C.18 and a prior parametrized in terms of an unknown parameter σ as in equation C.31.

C.3.4 Determining the form of unknown parameter σ

We have re-parametrized our signal extraction problem in terms of σ , giving us a posterior pdf, f , of the form

$$f(\mathbf{h}'|\mathbf{R}, \mathbf{N}, \sigma, I) \propto g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I)u(\mathbf{h}'|\sigma, I) \quad (\text{C.32})$$

Filling in what we know of the forms of g and u from sections C.3.2 and C.3.3:

$$f(\mathbf{h}'|\mathbf{R}, \mathbf{N}, \sigma, I) \propto e^{[-\frac{1}{2}(\mathbf{d}-\mathbf{R}\bullet\mathbf{h}')^T C^{-1}(\mathbf{d}-\mathbf{R}\bullet\mathbf{h}') - \frac{\sum_k h'_k{}^2}{2\sigma^2}]} \quad (\text{C.33})$$

We see that we must find the best guess for \mathbf{h}' that minimizes the function in the exponential of equation C.33.

$$f'(h') = \frac{1}{2}(d - R \bullet h')^T C^{-1}(d - R \bullet h') - \frac{\sum_k h'_k{}^2}{2\sigma^2} \quad (\text{C.34})$$

From here, we must find the form of σ in order to maximize entropy for the best guess of \mathbf{h}' .

We can write a probability distribution, $w(\sigma)$, for σ being the appropriate choice to use in equation C.34. First, we can integrate our prior, u , over σ to find a distribution independent of σ .

$$u_\sigma(\mathbf{h}|I) = \int d\sigma^2 w(\sigma^2) u(\mathbf{h}|\sigma^2, I) \quad (\text{C.35})$$

Then note that our observations \mathbf{d} also provide us with a probability distribution for σ , the regularization constant, and that making use of Bayes' theorem results in an expression for $w(\sigma)$ similar in form to equation C.32.

$$w(\sigma|\mathbf{d}, N, \mathbf{R}, I) \propto \int D\mathbf{h} g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I) u_\sigma(\mathbf{h}'|I) \quad (\text{C.36})$$

Since we already know the forms of $g(\mathbf{d}|\mathbf{h}', \mathbf{R}, N, I)$ and $u_\sigma(\mathbf{h}'|I)$ from the discussion in sections C.3.3 and C.3.2 we can straightforwardly evaluate these Gaussian integrals to find the form of $w(\sigma)$.

Here the resulting form of $w(\sigma)$ has been simplified by taking the log:

$$\begin{aligned} \log w &= \frac{1}{2} \left[-(d - R \bullet h'_0)^T C^{-1} (d - R \bullet h'_0) \right]^2 \\ &\quad - \frac{h'_0{}^2}{\sigma^2} - C_d \log 2\pi + 2C_h \log \sigma \\ &\quad - \log \det \| \sigma^{-2} \mathbf{I} + \mathbf{R}^T \mathbf{N}^{-1} \mathbf{R} \| \end{aligned} \tag{C.38}$$

We can now maximize the resulting form of the log of $w(\sigma)$. This is equivalent to maximizing the probability that σ is an appropriate choice for equation C.34.

$$\sigma'^2 = \frac{\sum_k h_k^2}{\| C \| + - \operatorname{tr}[I + \sigma'^2 R^T \| C \|^{-1} R]^{-1}} \tag{C.39}$$

C.3.5 Summary of method

Our problem breaks down maximizing two equations simultaneously.

We solve equation C.34 for h' given a guess of σ , and then using the generated h' to make a new guess of σ' in equation C.41, iteratively until the solution converges.

Summary of simultaneously maximized equations:

$$f'(h') = \frac{1}{2}(d - R \bullet h')^T C^{-1}(d - R \bullet h') - \frac{\sum_k h'_k{}^2}{2\sigma^2} \quad (\text{C.40})$$

$$\sigma'^2 = \frac{\sum_k h_k^2}{\|C\| + -tr[I + \sigma'^2 R^T \|C\|^{-1} R]^{-1}} \quad (\text{C.41})$$

APPENDIX D

OMICRON PARAMETERS IN DEPTH

Below is an example Omicron parameters file with example entries and detailed explanations of each parameter, as referenced in Section 3.1.2.

Notes are commented out following '//**', so this example could be used and altered copied from this form.

An example using a simulated gravitational wave strain channel

```
//*****
//***** Omicron configuration file *****
//*****

//*****
//***** INPUT DATA *****
//*****

//** full path to CACHE file pointing to frame files - this must be a cache
//** file
//** Supported file formats: LCF (lalcache)
```

```
DATA    LCF                ./ETG_test_MDCs.cache
```

```
/** list of channels you want to process
```

```
/** This list can be arbitrarily long, as long as the channels are all
```

```
/** contained in the same gwf file
```

```
/** Can be listed on one or several lines
```

```
DATA    CHANNELS          H1:FAKE-STRAIN
```

```
/** working sampling (one value for all channels)
```

```
/** This is the down-sampled frequency rate, which should be less than
```

```
/** or equal to the native sampling rate.
```

```
/** Should also be an integer and a power of two.
```

```
DATA    SAMPLEFREQUENCY 8192
```

```
/*******
```

```
/*******          SEARCH PARAMETERS          *****
```

```
/*******
```

```
/** chunk duration in seconds (must be an integer)
```

```
/** This is the length of time that is used to estimate the PSD
```

```
PARAMETER    CHUNKDURATION    484
```

```
/** segment duration in seconds (must be an integer)
```

```
/** This is the length of time that is used to calculate time-frequency
```

```
/** tiling
```

```
PARAMETER    BLOCKDURATION    64
```



```
/** overlap duration between segments in seconds (must be an integer)
/** The last N seconds of a segment and first N segments of the next
/** that will be effectively analyzed twice to
/** avoid edge effects.
```

```
PARAMETER      OVERLAPDURATION  4
```

```
/** search frequency range
/** The user requested frequency range - lower limit must be greater
/** than zero and the upper limit
/** must be less than or equal to Nyquist.
```

```
PARAMETER      FREQUENCYRANGE  32      4096
```

```
/** search Q range
/** The range of considered Qs that will be used to generate a
/** logarithmically spaced series of Q-planes
/** tiled in time and frequency.
```

```
PARAMETER      QRANGE          3.3166  141
```

```
/** maximal mismatch between 2 consecutive tiles (0<MM<1)
/** for dark-fringe channel, 0.2 is a good value
/** for auxiliary channels, 0.4 is a good value (faster)
/** Note this computed assuming the mismatched signal is a
/** sine-Gaussian.
```

```
PARAMETER      MISMATCHMAX     0.2
```

```
/*******
/******* TRIGGERS *****
```

```

//*****

/** tile SNR threshold
TRIGGER          SNRTHRESHOLD    5.5

/** maximum number of triggers per file
TRIGGER          NMAX             500000

/** clustering on
/** This should be commented out for unclustered triggers.
//TRIGGER        CLUSTERING      TIME

//*****
//*****                               OUTPUT                               *****
//*****

/** full path to output directory
OUTPUT DIRECTORY          ./triggers

/** full path to output directory
/** Also supported: root
OUTPUT FORMAT             xml

/** verbosity level (0-1-2-3)
OUTPUT VERBOSITY         0

/** flag to save PSDs

```

```
OUTPUT WRITEPSD          0
```

```
/** flag to save data time series
```

```
OUTPUT WRITETIMESERIES   0
```

```
/** flag to save whitened data time series
```

```
OUTPUT WRITEWHITENEDDATA 0
```

textbfAn example using seismometer channels

Here the exact Omicron parameter file used to generate triggers for the seismic platform transient study detailed in Chapter 5 is given in full, without comments¹.

```
/*******  
/******* Omicron configuration file *****  
/*******
```

```
DATA LCF /home/jlmciver/public_html/SEI_transient_wind/Oct9-  
Oct12.cache
```

```
PARAMETER QRANGE 3.3166 100
```

```
PARAMETER MISMATCHMAX 0.35
```

```
TRIGGER SNRTHRESHOLD 5.5
```

¹Note the 'Native frequency' no longer needs to be given in the current version of Omicron.

OUTPUT FORMAT xml
OUTPUT VERBOSITY 0
OUTPUT WRITEPSD 0
OUTPUT WRITETIMESERIES 0
OUTPUT WRITEWHITENEDDATA 0

DATA CHANNELS H1:ISI-GND_STS_ETMX_X_DQ
H1:ISI-GND_STS_ETMX_Y_DQ H1:ISI-GND_STS_ETMX_Z_DQ
H1:ISI-ETMX_ST1_BLND_X_T240_CUR_IN1_DQ
H1:ISI-ETMX_ST1_BLND_Z_T240_CUR_IN1_DQ
H1:ISI-ETMX_ST2_BLND_X_GS13_CUR_IN1_DQ
H1:ISI-ETMX_ST2_BLND_Z_GS13_CUR_IN1_DQ
H1:HPI-ETMX_BLND_L4C_X_IN1_DQ
H1:HPI-ETMX_BLND_L4C_Z_IN1_DQ
H1:ISI-GND_STS_ITMY_X_DQ
H1:ISI-GND_STS_ITMY_Y_DQ
H1:ISI-GND_STS_ITMY_Z_DQ
H1:ISI-ITMY_ST1_BLND_X_T240_CUR_IN1_DQ
H1:ISI-ITMY_ST1_BLND_Z_T240_CUR_IN1_DQ
H1:ISI-ITMY_ST2_BLND_X_GS13_CUR_IN1_DQ
H1:ISI-ITMY_ST2_BLND_Z_GS13_CUR_IN1_DQ
H1:HPI-ITMY_BLND_L4C_X_IN1_DQ
H1:HPI-ITMY_BLND_L4C_Z_IN1_DQ
DATA NATIVEFREQUENCY 512 512 512 512 512 2048 2048 1024
1024 512 512 512 512 512 2048 2048 1024 1024
DATA SAMPLEFREQUENCY 512

PARAMETER CHUNKDURATION 8192
PARAMETER BLOCKDURATION 8192
PARAMETER OVERLAPDURATION 1280
PARAMETER FREQUENCYRANGE 0.1 64
TRIGGER NMAX 500000

OUTPUT DIRECTORY ./triggers

BIBLIOGRAPHY

- [1] Aasi, J. et al. (The LIGO Scientific Collaboration). Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys. Rev. D.* *87*, 042001 (2013).
- [2] Aasi, J. et al. (the LIGO Scientific Collaboration). Advanced LIGO. *Classical and Quantum Gravity* *32*, 7 (2015).
- [3] Aasi, J. et al. (The LIGO Scientific Collaboration). Characterization of the LIGO detectors during their sixth science run. *Class. Quant. Grav.* *32*, 115012 (2015).
- [4] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO's Sixth and Virgo's Second and Third Science Runs. *arXiv 1203.2674* (2012).
- [5] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *arXiv:1304.0670* (2013).
- [6] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Search for gravitational waves from binary black hole inspiral, merger, and ringdown in LIGO-Virgo data from 2009-2010. *Phys. Rev. D.* *87*, 022002 (2013).
- [7] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). The NINJA-2 project: detecting and characterizing gravitational waveforms modelled using numerical binary black hole simulations. *Class. Quantum Grav.* *31*, 115004 (2013).
- [8] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Gravitational waves from known pulsars: results from the initial detector era. *Astrophysics J* *785*, 2 (2014).
- [9] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009-2010 LIGO and Virgo Data. *Phys. Rev. Lett.* *113*, 231101 (2014).

- [10] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Search for gravitational radiation from intermediate mass black hole binaries in data from the second LIGO-Virgo joint science run. *Phys. Rev. D.* *89*, 122003 (2014).
- [11] Aasi, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Search for gravitational wave ringdowns from perturbed intermediate mass black holes in LIGO-Virgo data from 2005-2010. *Phys. Rev. D.* *89*, 102006 (2014).
- [12] Abadie, J. et al. Directional Limits on Persistent Gravitational Waves Using LIGO S5 Science Data. *Phys. Rev. Lett.* *107*, 271102 (2011).
- [13] Abadie, J. et al. (The LIGO Scientific Collaboration). Search for gravitational waves associated with the August 2006 timing glitch of the Vela pulsar. *Phys. Rev. D* *83*, 042001 (2011).
- [14] Abadie, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). All-sky search for gravitational-wave bursts in the second joint LIGO-Virgo run. *Phys. Rev. D.* *85*, 122007 (2012).
- [15] Abadie, J. et al. (The LIGO Scientific Collaboration and the Virgo collaboration). Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3. *Astrophys J* *760*, 1 (2012).
- [16] Abadie, J. et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys Rev D* *85*, 082002 (2012).
- [17] Abbott, B.P. et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics* *72*, 7 (2009).
- [18] Accadia, T. et al. Virgo: a laser interferometer to detect gravitational waves. *Journal of Instrumentation* *7*, 1088 (2012).
- [19] Adam, S.M. et al. Observing the next galactic supernova. *AstroJ* *778* (2013).
- [20] Alekseev, E.N. et al. Properties of the Supernova 1987A Neutrino Signal Recorded by the Baksan Underground Scintillation Telescope. *Soviet Astronomy Letters* *14* (1988).
- [21] Allen, Bruce et al. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys Rev D* *85*, 122006 (2012).
- [22] Anderson, Warren G. et al. An excess power statistic for detection of burst sources of gravitational radiation. *Phys. Rev. D* *63*, 042003 (2001).

- [23] Baumgarte, T.W., and Shapiro, S.L. *Numerical Relativity*. World Scientific Publishing, 1994.
- [24] Bethe, H. A. Supernovae mechanisms. *Reviews of Modern Physics* 62 (1990).
- [25] Bionta, R.M. et al. Observation of a neutrino burst from SN 1987A in IMB. *Neutrino masses and neutrino astrophysics* (1987).
- [26] Brown, D.A. et al. Testing the LIGO inspiral analysis with hardware injections. *Class. Quant. Grav.* 21, 5 (2004).
- [27] Buonanno, A., and Chen, Y. Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors. *Phys. Rev. D* 64, 042006 (2001).
- [28] Chatterji, Shourov K. *The search for gravitational wave bursts in data from the second LIGO science run*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [29] Cornish, Neil J., and Littenberg, Tyson B. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Submitted* (2015).
- [30] Cumming, A. et al. Design and development of the advanced LIGO monolithic fused silica suspension. *Class. Quant. Grav.* 29, 035003 (2012).
- [31] Dimmelmeier, D. et al. Gravitational wave burst signal from core collapse of rotating stars. *Phys Rev D* 78, 064056 (2008).
- [32] Dooley, K. et al. Angular control of optical cavities in a radiation pressure dominated regime: the Enhanced LIGO case. *J. Opt. Soc. Am.* 30 (2013).
- [33] Dooley, K. et al. Status of advanced ground-based laser interferometers for gravitational-wave detection. *Publication in progress, arXiv 1411.6068* (2014).
- [34] Driggers, J. et al. Subtraction of Newtonian noise using optimized sensor arrays. *Physical Review D* 86, 102001 (2012).
- [35] Effler, A. et al. Environmental influences on the LIGO gravitational wave detectors during the 6th science run. *Classical and Quantum Gravity* 32, 3 (2015).
- [36] Einstein, Albert. Zur allgemeinen Relativittstheorie. *Preussische Akademie der Wissenschaften, Sitzungsberichte* (1915).
- [37] Essick, R. et al. Optimizing vetoes for gravitational-wave transient searches. *Class. Quant. Grav.* 30, 15 (2013).
- [38] Evans, M. et al. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Phys Rev D* 78, 102003 (2008).

- [39] Fransson, C. et al. Twenty Years of Supernova 1987A. *The Messenger* 127 (2007).
- [40] Fricke, Tobin et al. DC readout experiment in Enhanced LIGO. *Class. Quantum Grav.* 29, 065005 (2012).
- [41] Gabor, D. Theory of Communication. *Journal of the Institution of Electrical Engineers* 93, 26 (1945).
- [42] Gerhard, Ortwin. Mass distribution in our Galaxy. *Space Science Reviews* 100 (2002).
- [43] Goldriech, and Weber. Homologously collapsing stellar cores. *Astrophys. J* 238 (1980).
- [44] Gonzalez, Gabriela et al. Gravitational wave astronomy. *Frontiers of Physics* (2013).
- [45] Gonzalez, G. Suspensions thermal noise in the ligo gravitational wave detector. *Class. Quantum Grav.* 17, 21 (2000).
- [46] Gossan, Sarah. Personal communication.
- [47] Harris, Fredric J. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE* 66, 1 (1978).
- [48] Harry, Gregory M. for the LIGO Scientific Collaboration. Advanced LIGO: the next generation of gravitational wave detectors. *Class. Quantum Grav.* 27, 084006 (2010).
- [49] Hild, S. et al. Demonstration and comparison of tuned and detuned signal recycling in a large-scale gravitational wave detector. *Class Quant Grav* 24, 1513 (2007).
- [50] Hirata, K.S. et al. Observation in the Kamiokande-II detector of the neutrino burst from supernova 1987A. *Phys. Rev. D.* 38, 2 (1988).
- [51] Isogai, T. et al. Used percentage veto for LIGO and Virgo binary inspiral searches. *J. Phys.: Conf. Ser* 243, 012005 (2010).
- [52] J. Aasi et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). First all-sky search for continuous gravitational waves from unknown sources in binary systems. *Phys. Rev. D.* 90, 062010 (2014).
- [53] Janka, et al. Theory of core-collapse supernovae. *Physics Reports* 442 (2007).
- [54] Kissel, Jeffery. *Calibrating and Improving the Sensitivity of the LIGO Detectors.* PhD thesis, Louisiana State University, 2010.

- [55] Kissel, Jeffery. aLIGO Seismic Isolation and Suspensions Schematic. *LIGO DCC* (2013).
- [56] Klimentenko, S. et al.). A wavelet method for detection of gravitational wave bursts. *Class Quant Grav* 21, 20 (2004).
- [57] Klimentenko, Sergey et al.). A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity* 25, 114029 (2008).
- [58] Kramer, M. et al. Tests of general relativity from timing the double pulsar. *Science* 314, 5796 (2006).
- [59] Lattimer, J.M., and Swesty, F.D. A generalized equation of state for hot, dense matter. *Nucl Phys A* 535 (1991).
- [60] Limongi, M. et al. Massive Stars in the Range 13-25 M_{\odot} : Evolution and Nucleosynthesis. II. The Solar Metallicity Models. *Astro J Supplement Series* 129 (2000).
- [61] Littenberg, Tyson B., and Cornish, Neil J. Bayesian inference for spectral estimation of gravitational wave detector noise. *Phys Rev D* 91, 084034 (2015).
- [62] Louge, J. et al. Inferring core-collapse supernova physics with gravitational waves. *Phys. Rev. D* 86, 044023 (2012).
- [63] Macleod, Duncan. The Livingston Summary Pages. <https://ldas-jobs.ligo-la.caltech.edu/~detchar/summary/> (2015).
- [64] Macleod, Duncan et al. Reducing the effect of seismic noise in LIGO searches by targeted veto generation. *Class. Quant. Grav.* 29, 055006 (2012).
- [65] Matichard, F. et al. Seismic isolation of Advanced LIGO gravitational waves detectors: Review of strategy, instrumentation, and performance. *Submitted. arXiv 1502.06300* (2015).
- [66] McIver, J. CCSN waveform reconstruction code. <https://wiki.ligo.org/Bursts/ReconstructionComparison#Code>, 2015.
- [67] McIver, J. Documentation of SEI configuration for high wind study. https://wiki.ligo.org/DetChar/SEITransient#Hanford_45_October_2014_wind_studies.
- [68] McIver, J. Documentation of SEI configuration for varied loop aggression. https://wiki.ligo.org/DetChar/SEITransient#Varying_the_aggressiveness_of_actuation_45_Jul
- [69] McIver, Jessica. Data quality studies of enhanced interferometric gravitational wave detectors. *Class. Quantum. Grav.* 29, 124010 (2012).
- [70] Mueller, B. et al. A New Multi-dimensional General Relativistic Neutrino Hydrodynamics Code of Core-collapse Supernovae. III. Gravitational Wave Signals from Supernova Explosion Models. *Astro. J.* 766, 1 (2013).

- [71] Mueller, E. et al. Parameterized 3D models of neutrino-driven supernova explosions. *Astronomy and Astrophysics* 537, 63 (2012).
- [72] Necula, V. et al.). Transient analysis with fast Wilson-Daubechies time-frequency transform. *J Phys: Conf. Ser.* 363, 012032 (2012).
- [73] Oelker, E. et al. Squeezed light for advanced gravitational wave detectors and beyond. *Opt Express* 22, 21106 (2014).
- [74] Ott, C. et al. General-relativistic Simulations of Three-dimensional Core-collapse Supernovae. *Astro. J.* 768, 2 (2013).
- [75] Ott, Christian. The gravitational wave signature of core-collapse supernovae. *Class. Quantum Grav.* 26, 063001 (2009).
- [76] Pagliaroli, G. et al. Neutrinos from supernovae as a trigger for gravitational wave search. *Phys. Rev. Lett.* 103, 031102 (2009).
- [77] Powell, J. et al. Classification methods for noise transients in advanced gravitational-wave detectors. *In progress - LIGO DCC P1500055* (2015).
- [78] Punturo, M. et al. The Einstein Telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity* 27, 19 (2010).
- [79] Robinet, Florent. Official omicron documentation. <http://virgo.in2p3.fr/GWOLLUM/v2r1/index.html?Friends/omicron.html> (2015).
- [80] Robinet, Florent. Personal communication.
- [81] Sathyaprakash, B.S., and Schutz, Bernard F. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Reviews in Relativity* 2, 12 (2009).
- [82] Saulson, P. Thermal noise in mechanical experiments. *Phys Rev D* 42, 2437 (1990).
- [83] Saulson, Peter R. *Fundamentals of Interferometric Gravitational Wave Detectors*. Cambridge University Press, 2010.
- [84] Scheidegger, S. et al. Gravitational waves from supernova matter. *Class. Quan. Grav.* 27 (2010).
- [85] Shen, H. et al. Relativistic equation of state of nuclear matter for supernova explosion. *Prog. Theo. Phys.* 100 (1998).
- [86] Shigeyama, T., and Nomoto, K. Theoretical light curve of SN 1987A and mixing of hydrogen and nickel in the ejecta. *Astrophysical Journal* 360 (1990).
- [87] Sivia, D.S. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, 2006.

- [88] Slutsky, J. et al. Methods for Reducing False Alarms in Searches for Compact Binary Coalescences in LIGO Data. *Class Quant Grav* 27, 165023 (2010).
- [89] Smith, Joshua R. et al. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Class. Quantum. Grav.* 28, 235005 (2011).
- [90] Staley, A. et al. Achieving resonance in the Advanced LIGO gravitational-wave interferometer. *Class. Quantum Grav.* 31, 245010 (2014).
- [91] Summerscales, T.Z. et al. Maximum entropy for gravitational wave data analysis: Inferring the physical parameters of core-collapse supernovae. *Astrophys.J* 678, 1142 (2008).
- [92] Taylor, J.H. et al. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *Astrophys.J* 253 (1982).
- [93] The LIGO Scientific Collaboration. The LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics (2015-2016 edition). *LIGO DCC T1500055* (2015).
- [94] The LIGO Scientific Collaboration and the Virgo Collaboration. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature* 460 (2009).
- [95] Thrane, Eric et al. Correlated magnetic noise in global networks of gravitational-wave interferometers: observations and implications. *Phys Rev D* 87, 123009 (2013).
- [96] Thrane, Eric et al. Correlated noise in networks of gravitational-wave detectors: subtraction and mitigation. *Phys Rev D* 90, 023013 (2014).
- [97] Willke, Benno. The GEO600 gravitational wave detector. *Proceedings: Conference on Lasers and Electro-Optics*, 1557527709 (2005).
- [98] Woosley, S.E. et al. The evolution and explosion of massive stars II. explosive hydrodynamics and nucleosynthesis. *Astro J Supplement Series* 101 (1995).
- [99] Woosley, S.E. et al. The evolution and explosion of massive stars. *Rev Mod Phys* 74 (2002).
- [100] Woosley, S.E. et al. Nucleosynthesis and remnants in massive stars of solar metallicity. *Physics Reports* 442 (2007).
- [101] Yakunin, K.N. et al. Gravitational Waves from Core Collapse Supernovae. *Class Quant Grav* 27, 194005 (2010).