

Open Issues in the Search for Gravitational Wave Transients

by

Lindy L. Blackburn

A.B. Physics with Honors and Mathematics, University of Chicago
(2003)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2010

© Lindy L. Blackburn, MMX. All rights reserved.

The author hereby grants to MIT permission to reproduce and
distribute publicly paper and electronic copies of this thesis document
in whole or in part.

Author

Department of Physics

August 31, 2010

Certified by

Erotokritos Katsavounidis

Associate Professor of Physics

Thesis Supervisor

Accepted by

Krishna Rajagopal

Professor of Physics

Associate Department Head for Education

Open Issues in the Search for Gravitational Wave Transients

by

Lindy L. Blackburn

Submitted to the Department of Physics
on August 31, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

The LIGO-Virgo network of kilometer-scale laser interferometric gravitational-wave detectors reached a major milestone with the successful operation of LIGO's fifth (S5) and Virgo's first (VSR1) science runs during 2005–2007. This thesis presents several issues related to gravitational-wave transient detection from the perspective of the joint all-sky, un-triggered burst search over S5/VSR1 data.

Existing searches for gravitational-wave bursts must deal with the presence of non-Gaussian noise transients which populate the data over the majority of sensitive signal space. These events may be confused with true signals, and are the current limiting factor in search sensitivity and detection confidence for any real event. The first part of this thesis focuses on the development of tools to identify, monitor and characterize these instrumental disturbances in LIGO and Virgo data. An automated procedure is developed and applied to the S5/VSR1 search in order to safely remove noise transients from the analysis without sacrificing sensitivity by making use of the wealth of auxiliary information recorded by the detectors.

The second part of this thesis focuses on the interpretation of outlier events in the context of a non-Gaussian, non-stationary background. An extensive follow-up procedure for candidate gravitational-wave events is developed and applied to a single burst outlier from the S5/VSR1 search, later revealed to be a blind simulation injected into the instruments. While the follow-up procedure correctly finds no reason to reject the candidate as a possible gravitational wave, it highlights the difficulty in making a confident detection for signals with similar waveform morphology to common instrumental disturbances. The follow-up also deals with the problem of objectively defining the significance of a single outlier event in the context of many semi-disjoint individual searches. To address this, a likelihood-ratio based unified ranking is developed and tested against the original procedures of the S5/VSR1 burst search. The new ranking shows a factor of four improvement in the statistical significance of the outlier event, and a 12% reduction using fixed thresholds and 38% reduction using a loudest event statistic for a rate upper limit on a mock signal population.

Thesis Supervisor: Erotokritos Katsavounidis
Title: Associate Professor of Physics

Acknowledgments

I'd like to begin by thanking my grandparents—I can't imagine growing up in a more loving and welcoming family. My entire approach to life I owe to my father—thank you for teaching me to think for myself. Thank you Mom for your irreplaceable love and for teaching me to appreciate good food. To Lawrence Lee for being such an incredible family friend. The coolest things were always at your place, and thank you for showing us how to live and love life. I thought maybe if I became a physicist I could be more like you. Jim Riley taught a truly amazing 5–6th grade class at Hawthorne Elementary. Thank you for showing us about the wonder in the world, and the trips to Balboa Park, Mission Bay, the symphony, circuits, magic, combinatorics, puzzles, art, ambigrams, square dancing, ultimate aerobic. . . And finally a late thank you to Mr. Feynman and Mr. Sagan for leaving your written legacy. No doubt you have transformed millions of people, including me.

I've met genuinely great people during my graduate student career at LIGO. I'd like to start with my advisor Erik Katsavounidis whose better judgment I benefited from innumerable times. I'd also like to thank my excellent committee members Scott Hughes and Nergis Mavalvala. Scott, thanks for teaching me G.R., but also, thanks for happily accommodating my rushed defense just before your wedding! When I started at MIT-LIGO, I shared an office with Shourov Chatterji. Thanks for being such a great example (aside from graduating in a timely manner!), but really for everything: attention to detail, an unrelenting drive to figure things out and do them right, amazing patience, Matlab wizardry. I'm not at all surprised to still see influence of your work and style all over the collaboration.

To Peter Shawhan—thanks for always being the first to understand and appreciate my garbled presentations to the burst group. And to Laura Cadonati, thank you for years of support and guidance. It was a pleasure every time I had a chance to work with you. Jolien Creighton—I know you mostly through emails not addressed to me, but there is always great stuff there and a wealth of good ideas, a few of which made their way into this work. Peter Saulson and Alan Weinstein—I always looked forward

to listening to all the valuable things either of you had to say. Josh Smith, Nicolas Leroy—it was great working with you both. Brennan Hughey, John Zweizig, Igor Yakushin—much of this thesis would have been impossible without great help from all of you, thank you! To my lab-mates Jared Markowitz and Leo Stein, it was fun brainstorming through problems and random stuff, best of luck to you both.

Thank you David Deveau for bringing me back to the piano in graduate school. To you and Marcus Thompson, I'll never forget all the great times in AMP and CMS. Thank you for making my life at MIT that much richer. To Rosa, and Paymon—your decade of friendship is invaluable to me. And to Amanda, thank you for your love and patience in seeing this through.

The authors gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, the Science and Technology Facilities Council of the United Kingdom, the Max-Planck-Society, and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector, and the Italian Istituto Nazionale di Fisica Nucleare and the French Centre National de la Recherche Scientifique for the construction and operation of the Virgo detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Council of Scientific and Industrial Research of India, the Istituto Nazionale di Fisica Nucleare of Italy, the Spanish Ministerio de Educación y Ciencia, the Conselleria d'Economia Hisenda i Innovació of the Govern de les Illes Balears, the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research, the Polish Ministry of Science and Higher Education, the FOCUS Programme of Foundation for Polish Science, the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, The National Aeronautics and Space Administration, the Carnegie Trust, the Leverhulme Trust, the David and Lucile Packard Foundation, the Research Corporation, and the Alfred P. Sloan Foundation.

Portions of this work do not reflect the scientific opinion of the LSC and were not reviewed by the collaboration. This document has LIGO number LIGO-P1000093.

Contents

1	Introduction	15
2	Gravitational radiation	19
2.1	Relativity	19
2.2	Weak field limit	22
2.2.1	Transverse traceless gauge	24
2.2.2	Effect of gravitational waves on matter	25
2.2.3	Stress-energy of gravitational waves	26
2.2.4	Generation of gravitational waves	28
2.3	Sources of gravitational radiation	30
2.3.1	Compact binary coalescence	30
2.3.2	Gravitational core collapse	35
2.3.3	Rotating neutron stars	37
2.3.4	Stochastic background	40
3	Detecting gravitational waves	45
3.1	Resonant mass detectors	45
3.2	Interferometric detectors	48
3.3	Laser Interferometer Gravitational-wave Observatory (LIGO)	50
3.3.1	Fabry-Perot cavities	51
3.3.2	Length sensing and control	53
3.3.3	Noise sources	55
3.3.4	Advanced LIGO	56

4	Gravitational wave bursts	59
4.1	Burst parametrization	59
4.2	Detecting bursts in stationary noise	61
4.2.1	Power spectral density	61
4.2.2	Matched filter signal-to-noise ratio	62
4.2.3	Time-frequency decomposition	66
4.3	Multiple detector observation	70
4.3.1	Coincidence in observed parameters	70
4.3.2	Coherent observation	71
4.3.3	Time-shift method of background estimation	73
5	Non-Gaussian noise transients	75
5.1	Transient population in LIGO data	76
5.2	Examples of transients	78
5.3	kleineWelle: wavelet-based identification	82
5.3.1	Wavelet decomposition	83
5.3.2	Data conditioning	85
5.3.3	Event generation	87
5.3.4	Parameter estimation	89
5.3.5	Implementation	90
6	Detector characterization, data quality, and vetoes	93
6.1	On-line detector characterization	93
6.2	Data quality flags	95
6.3	Event-by-event vetoes for burst background	96
6.3.1	Veto effectiveness metrics	99
6.3.2	Veto safety	100
6.3.3	Selection of veto conditions	101
7	Candidate event follow-up	107
7.1	Detection checklist	109

7.2	Blind injections during S5	110
7.3	Follow-up of S5/VSR1 burst outlier	111
7.3.1	Introduction	111
7.3.2	The event	112
7.3.3	State of the instruments	115
7.3.4	All-sky search and background estimates	122
7.3.5	Event reconstruction	127
7.3.6	Astrophysical analysis	130
7.3.7	Other events in S5	135
7.3.8	Conclusion of the follow-up procedure	137
8	Unified ranking for gravitational-wave events	143
8.1	Trial factor issues in S5/VSR1	145
8.2	Inverse false-alarm rate ranking	148
8.3	Likelihood-ratio ranking	150
8.3.1	Signal population	151
8.3.2	Background population	153
8.3.3	Analytic fits	153
8.3.4	Results	156
9	Conclusions	165
A	Statistical significance of Poisson and χ^2 processes	169
A.1	Poisson distribution	169
A.2	Gaussian approximation	170
A.3	Calculating Poisson significance	171
A.3.1	Gamma function approximations	171
A.3.2	Extreme tail of the Poisson CDF	173
A.4	χ^2 distributions	174
B	KleineWelle channel configuration	175
B.1	S5 channels and frequency range	175

B.2	S6 channels and frequency range	178
C	Auxiliary channel veto statistics	183
C.0.1	Veto performance on H1 transients	183
C.0.2	Veto performance on H2 transients	185
C.0.3	Veto performance on L1 transients	187
D	Burst detection checklist	189

List of Figures

2-1	The effect of a gravitational wave on a ring of particles	27
2-2	Binary black hole merger waveform	34
2-3	Gravitational waveform from a stellar collapse	37
2-4	Sources and limits of a stochastic gravitational wave background . . .	43
3-1	Current resonant bar detector strain sensitivities	47
3-2	A two mirror resonant cavity	51
3-3	LIGO optical layout	54
3-4	LIGO S5 sensitivity	57
5-1	Transient population rate vs SNR during S5	77
5-2	Transient population frequency and Q during S5	79
5-3	Transient rate vs day for S5	80
5-4	An event caused by magnetic disturbance at LHO	81
5-5	Discrete wavelet decomposition tree using conjugate mirror filters . .	84
5-6	Example of LPEF+DWT applied to a test signal in noise	86
5-7	Definition of distance in the dyadic wavelet tiling	88
6-1	S5 online kleineWelle trends	94
6-2	Application of event-by-event veto window	96
6-3	Hierarchical veto tuning procedure	103
6-4	Auxiliary channel veto performance on burst background	105
7-1	Filtered time series of LIGO data at the time of the blind injection .	113
7-2	Spectrogram of LIGO calibrated strain data	114

7-3	H1H2 consistency of the blind burst injection	115
7-4	Q Pipeline correlated vs coherent energy	116
7-5	Q Pipeline single-interferometer triggers in ± 8 s of the blind injection	118
7-6	Q Pipeline rate vs. time	119
7-7	Q Pipeline SNR vs. time	120
7-8	Q Pipeline frequency vs. time	122
7-9	Results from the kleineWelle+CorrPower online analysis	123
7-10	Coherent Waveburst S5 H1H2L1 background distribution	124
7-11	Coherent Waveburst S5 H1H2L1 background properties	125
7-12	Results from the Q/ Ω Pipeline analysis	127
7-13	Coherent Waveburst sky statistic for H1H2L1 and H1H2L1V1 networks	128
7-14	Coherent Waveburst best-fit strain waveform	129
7-15	Coherent best fit waveform in data	130
7-16	Ω Pipeline Bayesian statistic sky map	131
7-17	Hilbert-Huang transform of candidate event	132
7-18	Equal mass black hole merger waveform	133
7-19	Time series of equal mass black hole merger simulation in LIGO noise	133
7-20	Spectrogram of equal mass black hole merger simulation in LIGO noise	134
7-21	L10 density up to 100 Mpc	136
7-22	Sine-Gaussian $f=100$ Hz $Q=9$ hardware injection	137
7-23	Gaussian 3 ms hardware injection	137
7-24	Comparison of injected and reconstructed waveforms	140
7-25	Comparison of injected and reconstructed sky location	141
8-1	Coherent Waveburst H1H2L1V1 likelihood-ratio maps	160
8-2	Coherent Waveburst H1H2L1 likelihood-ratio maps	161
8-3	Q/ Ω Pipeline H1H2L1 likelihood-ratio maps	162
8-4	Q/ Ω Pipeline H1H2V1 likelihood-ratio maps	163
8-5	Q/ Ω Pipeline L1V1 likelihood-ratio maps	164

List of Tables

2.1	Predicted rates for binary coalescence	31
6.1	Data quality categories for S5 transient searches	97
7.1	Basic information on burst event 090722	112
7.2	State of the LIGO, GEO600, and Virgo instruments	116
7.3	Auxiliary channel triggers detected by kleineWelle	121
7.4	Coherent Waveburst single-interferometer parameters for candidate event	128
7.5	Coherent Waveburst parameters for reconstructed gravitational wave	130
7.6	Electromagnetic events from the GCN catalog around 2007-09-22 . .	131
8.1	Configurations and live-times for S5/VSR1 burst searches	147
8.2	Comparison of hand-chosen and likelihood-ratio thresholds	157

Chapter 1

Introduction

One of the most fascinating predictions of Einstein’s general theory of relativity is the existence of gravitational radiation—ripples in the geometry of space-time which propagate as waves at the speed of light. Due to the weakness of gravitational interaction, it takes an immense amount of matter accelerating at relativistic speeds to make an appreciable amount of gravitational radiation. While any terrestrial attempt to generate gravitational waves would result in waves much too weak to be detectable, sufficiently extreme conditions do occur in the universe to produce gravitational waves at the Earth which approach the limit of current detector sensitivities, such as the coalescence of two neutron stars or black holes.

The experimental quest to detect astrophysical sources of gravitational radiation began with the work of Joseph Weber in the 1960s at the University of Maryland [1]. Weber’s idea was to use large suspended resonant metal bars whose modes would be excited by a passing burst of gravitational radiation. Today, the most promising detectors are broadband laser interferometers which began development in the late 1960s with the independent work of Weiss [2], Moss, Miller and Forward [3]. The largest and most sensitive interferometric detectors currently in operation include the two US Laser Interferometer Gravitational-wave Observatory (LIGO) 4 km-baseline instruments at Hanford, WA and Livingston, LA [4], and the 3 km French-Italian Virgo detector near Pisa, Italy [5].

Although gravitational waves have yet to be directly observed, strong evidence of

their existence comes from the measured orbital decay of the binary pulsar system PSR 1316+16, which was discovered by Russell Hulse and Joseph Taylor in 1974 [6]. Over decades of observation, the measured energy loss of the double neutron star orbit has proved very consistent with the amount of energy and angular momentum radiated away in gravitational waves predicted by general relativity [7]. The LIGO and Virgo instruments are designed to capture, among other things, the signal produced during the final coalescence and merger of such binary systems. Gravitational waves can provide a unique undisturbed signature of the inner dynamics of systems which are either dark in or opaque to photons and neutrinos, providing an entirely new window through which to observe the universe.

The material in this thesis deals primarily with the search for gravitational-wave transients in data from LIGO's fifth science run (S5) which took place between November 2005 and October 2007. During this time, LIGO collected one full year of data with all three LIGO instruments (including the currently decommissioned 2 km instrument at Hanford) operating at design sensitivity. The latter half of S5 overlapped with Virgo's first science run (VSR1), which provided the opportunity for a coordinated joint search.

An introduction to gravitational waves arising from Einstein's theory of general relativity is presented in chapter 2, accompanied by an overview of astrophysics sources relevant to ground-based detection. A brief description of ground-based detectors, including resonant-mass detectors and interferometric detectors, is presented in chapter 3. Additional detail is given to the LIGO instruments during S5, along with current and future upgrades. The basic techniques for the detection of burst sources (short-duration gravitational-wave transients) in detector noise are presented in chapter 4. This includes burst parametrization, the concept of matched-filter signal-to-noise, time-frequency decompositions and coherent network analysis, and the time-shift method of background estimation.

In chapter 5, we describe the character of the noise transients which populate the gravitational-wave data from the LIGO instruments. Accidental coincidences of noise transients across multiple detectors dominate the background in existing

gravitational-wave transient searches. These noise events are the limiting factor in search sensitivity as well as the ability to detect a real event with confidence. An efficient wavelet-based method called *kleineWelle* is developed to identify and parametrize time-series transients in real-time across the hundreds of test points (channels) recorded across the sites. In chapter 6, a hierarchical method is described which reduces the large amount of auxiliary (non-gravitational) information into that which is most effective at identifying and removing periods of instrumental disturbances. These techniques are ultimately used to reduce background in the S5/VSR1 all-sky un-triggered burst search [8].

The S5/VSR1 burst search resulted in a single outlier event which had an associated false-alarm rate from background of once per 43 years. As a potential detection candidate, the outlier prompted an extensive follow-up, described in detail in chapter 7. The outlier event was ultimately revealed as a blind injection of a simulated ad-hoc burst gravitational waveform. As such, it provided a valuable test of the detection readiness of our search procedures, as well as valuable experience in candidate event follow-up and interpretation for the search teams.

The outlier event follow-up also highlighted the need for an objective procedure to evaluate the global significance of an event (unlikelihood to have arisen from background) in the context of the many semi-disjoint sub-searches which together make up a burst search covering multiple networks, frequencies, and analysis methods (together we define these choices as the search configuration). This is achieved most cleanly through a global ranking of all events. When evaluated for a single population of simulated signals, the likelihood ratio provides a natural ranking which is directly comparable across sub-searches. This unified ranking is applied to events from the S5/VSR1 burst search in chapter 8 using a homogeneous population of ad-hoc burst signals and time-shift sample background, and various metrics of performance are used to compare the new ranking with the hand-tuned fixed thresholds used in the original search.

Chapter 2

Gravitational radiation

2.1 Relativity

The principle of relativity is the idea that the laws of physics do not depend on the inertial reference frame of any particular observer, and so no experiment can be constructed to measure one's velocity with respect to a fundamental notion of space at rest. The additional empirical requirement that light is observed to travel at a fixed speed regardless of reference frame gives us special relativity as described by Einstein in 1905. In special relativity, spacetime is a four dimensional manifold with a Minkowski metric (in Cartesian coordinates t, x, y, z with natural units $c = 1$),

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2.1)$$

which defines the squared distance between two neighboring points, or spacetime interval:

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu \quad (2.2)$$

$$= -dt^2 + dx^2 + dy^2 + dz^2. \quad (2.3)$$

In this and following expressions, we adopt tensor notation used by MTW [9] and many others where Greek indices (μ, ν, \dots) represent space-time coordinates while Latin indices (i, j, \dots) represent spatial coordinates only. We will also use shorthand for partial derivative operators $\partial_\mu = \partial/\partial x^\mu$. Repeated indices in an expression indicate a sum over all possible values. Indices which appear on each side of an equation means the equation holds for all possible values. ds^2 is invariant under the special group of Lorentz transformations which relate coordinates among different inertial observers. For a clock moving along a particular spacetime path, $d\vec{x}^2 = 0$ in the instantaneous moving inertial frame of the clock, so that the that the proper time $d\tau = \sqrt{-ds^2}$ elapsed for the world-line of a particle is also an invariant quantity and can be used to parametrize its path.

To include the effects of gravity into relativistic theory, Einstein made the conceptual leap to describe spacetime as fundamentally curved. In general relativity, gravity is not an invisible and instantaneous force between two masses, but a manifestation of the behavior of locally straight lines (geodesics) in a curved space. This is suggested by the observation that all objects undergo the same acceleration in a gravitational field regardless of mass or composition. At the same time, the local curvature itself is governed by mass-energy density.

The geometry of curved space is characterized by a symmetric metric $g_{\mu\nu}$ which replaces the flat Minkowski metric $\eta_{\mu\nu}$ of special relativity. The metric cannot be arbitrary—the geometry must reduce to the flat spacetime of special relativity at a sufficiently small scale. In a small enough laboratory at free-fall, we should not be able to observe the effects of gravity at finite precision.

The essential nature of curvature is contained in the Riemann tensor,

$$R^\rho{}_{\sigma\mu\nu} = \partial_\mu\Gamma^\rho{}_{\sigma\nu} - \partial_\nu\Gamma^\rho{}_{\sigma\mu} + \Gamma^\rho{}_{\lambda\mu}\Gamma^\lambda{}_{\sigma\nu} - \Gamma^\rho{}_{\lambda\nu}\Gamma^\lambda{}_{\sigma\mu}, \quad (2.4)$$

written here in terms of the affine connection coefficients which relate partial deriva-

tives to covariant derivatives suitable for curved space,

$$\Gamma^{\rho}_{\mu\nu} = \frac{1}{2}g^{\rho\lambda} (\partial_{\mu}g_{\lambda\nu} + \partial_{\nu}g_{\lambda\mu} - \partial_{\lambda}g_{\mu\nu}). \quad (2.5)$$

The Riemann curvature tensor records the way vectors change slightly in curved space when parallel transported about infinitesimal loops of arbitrary orientation. Parallel transport means the vector undergoes no first order rotation in its locally flat neighborhood as it moves along a path. The metric with raised indices is the inverse of the original metric, so that $g^{\mu\lambda}g_{\lambda\nu} = \delta^{\mu}_{\nu}$ where repeated indices indicate a sum over coordinates. The metric can also be used to raise or lower indices: $V_{\mu} = g_{\mu\nu}V^{\nu}$, which allows us to take inner products: $A^{\mu}B_{\mu} = g_{\mu\nu}A^{\mu}B^{\nu}$ and contractions: $A^{\mu}_{\mu} = g_{\mu\nu}A^{\mu\nu}$.

In a local inertial frame for some point in spacetime, the metric $g_{\mu\nu}$ reduces to the Minkowski metric at that point and its first derivatives vanish, setting the affine connection to zero when evaluated at the point as well. This shows us that spacetime curvature and the Riemann tensor are fundamentally characterized by second derivatives of the metric. In general relativity there are 20 second derivatives of the metric which cannot be removed by choice of coordinates, and there are 20 independent components of the Riemann tensor as well.

There is one independent contraction of the Riemann tensor which gives the Ricci tensor

$$R_{\mu\nu} = R^{\lambda}_{\mu\lambda\nu} \quad (2.6)$$

and Ricci scalar

$$R = R^{\lambda}_{\lambda}. \quad (2.7)$$

We can now write down the Einstein equations which relate the dynamics of spacetime curvature to the local stress-energy tensor $T^{\mu\nu}$,

$$R^{\mu\nu} - \frac{1}{2}g^{\mu\nu}R = \frac{8\pi G}{c^4}T^{\mu\nu}. \quad (2.8)$$

The trace-reversed form of the Ricci tensor on the left hand side of the equation is known as the Einstein tensor, $G^{\mu\nu} = R^{\mu\nu} - \frac{1}{2}g^{\mu\nu}R$, which is the unique second-rank

tensor built from linear contributions of the Riemann tensor that allows Equation 2.8 to satisfy conservation of energy-momentum ($\nabla_\mu T^{\mu\nu} = 0$). The constant factor on the right hand side gives the correct Newtonian limit (subsequent equations will use natural units where $c = G = 1$ unless otherwise stated). The equations provide ten independent constraints on the initial conditions and dynamical evolution of the metric. If an initial metric is chosen which satisfies Einstein's equations, their unique evolution (up to a choice of coordinates) will continue to satisfy the equations for all time, thus determining all geometric properties of the spacetime.

2.2 Weak field limit

In the limit of weak curvature, the metric can be considered approximately Minkowskian plus a small perturbation,

$$g_{\mu\nu} \simeq \eta_{\mu\nu} + h_{\mu\nu}, \quad \|h_{\mu\nu}\| \ll 1. \quad (2.9)$$

In this limit, we ignore all contributions to curvature that are more than first order in $h_{\mu\nu}$. The connection coefficients become

$$\Gamma^\rho{}_{\mu\nu} = \frac{1}{2}\eta^{\rho\lambda} (\partial_\mu h_{\lambda\nu} + \partial_\nu h_{\lambda\mu} - \partial_\lambda h_{\mu\nu}) \quad (2.10)$$

and the Riemann tensor becomes

$$R^\rho{}_{\sigma\mu\nu} = \partial_\mu \Gamma^\rho{}_{\sigma\nu} - \partial_\nu \Gamma^\rho{}_{\sigma\mu} \quad (2.11)$$

as the Γ^2 terms drop out. The Riemann tensor depends entirely on second derivatives of $h_{\mu\nu}$ giving rise to the interpretation that $h_{\mu\nu}$ is a small curvature field resting on top of flat background spacetime.

In addition to the usual Lorentz transformations, we are free to choose coordinates

which represent an infinitesimal gauge transformation,

$$x^{\mu'} = x^\mu + \xi^\mu \quad (2.12)$$

$$h_{\mu'\nu'} = h_{\mu\nu} - \partial_\mu \xi_\nu - \partial_\nu \xi_\mu, \quad (2.13)$$

for some small vector $\xi^\mu = \eta^{\mu\nu} \xi_\nu$. Under this transformation the metric maintains the form of Equation 2.9 to first order. There is also no change in the Riemann tensor as $R^\rho{}_{\sigma\mu\nu}$ is already small, and the new coordinates represent the same physical curvature.

Particularly useful is the harmonic or Lorentz gauge $\square x^\mu = 0 \rightarrow g^{\mu\nu} \Gamma^\rho{}_{\mu\nu} = 0$. We have introduced the D'Alembertian wave operator $\square = \partial^\mu \partial_\mu$. In the weak field limit, this requirement becomes

$$\partial_\mu \bar{h}^{\mu\nu} = 0 \quad (2.14)$$

for the trace-reversed metric perturbation

$$\bar{h}^{\mu\nu} = h^{\mu\nu} - \frac{1}{2} \eta^{\mu\nu} h^\lambda{}_\lambda, \quad (2.15)$$

which can be satisfied by absorbing the divergence of $\bar{h}^{\mu\nu}$ into the available coordinate perturbation,

$$\square \xi^\nu = \partial_\mu \bar{h}^{\mu\nu}. \quad (2.16)$$

In the Lorentz gauge, the linearized Einstein's equations take the simple form

$$\square \bar{h}^{\mu\nu} = -\frac{16\pi G}{c^4} T^{\mu\nu}. \quad (2.17)$$

We have reintroduced the factor of G/c^4 in the linearized Einstein equation to emphasize the smallness of the perturbations.

2.2.1 Transverse traceless gauge

For vacuum ($T^{\mu\nu} = 0$), the linearized Einstein equations give a three-dimensional wave equation for metric perturbation $\bar{h}^{\mu\nu}$,

$$\left(-\frac{\partial^2}{\partial t^2} + \nabla^2\right)\bar{h}^{\mu\nu} = 0. \quad (2.18)$$

The equation has general solutions that can be represented as linear combinations of plane waves,

$$\bar{h}^{\mu\nu} = C^{\mu\nu} \exp[ik_\sigma x^\sigma]. \quad (2.19)$$

Here $C^{\mu\nu}$ is a constant complex symmetric tensor orthogonal to the wavevector ($C^{\mu\nu}k_\nu = 0$). The wavevector itself must be null ($k^\sigma k_\sigma = 0$) to satisfy 2.18 so that a single mode represents a wave traveling in some direction at the speed of light.

The orthogonality requirement that $C^{\mu\nu}k_\nu = 0$ comes from the harmonic gauge condition (Equation 2.14). However the gauge is not unique, and any additional coordinate perturbation which satisfies the wave equation will allow us to remain in the harmonic gauge as it does not change $\partial_\mu \bar{h}^{\mu\nu}$. By using a transformation of the form

$$\xi^\mu = B^\mu \exp[ik_\sigma x^\sigma] \quad (2.20)$$

we can, for the correct choice of B^μ [10, 6.50], further restrict $C^{\mu\nu}$ to satisfy two new conditions,

$$C^\mu{}_\mu = 0 \quad (2.21)$$

$$U^\mu C_{\mu\nu} = 0, \quad (2.22)$$

for some arbitrary timelike vector U^μ . If we choose $U^\mu = (1, 0, 0, 0)$ and orient k_σ to represent a wave traveling in the z direction, the new conditions mean that $\bar{h}_{\mu\nu}$ takes

the simple form,

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & C_{xx} & C_{xy} & 0 \\ 0 & C_{xy} & -C_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \exp[-i\omega(t-z)]. \quad (2.23)$$

This is known as the *transverse-traceless* (TT) gauge because the metric perturbation is trace-free (due to 2.21) and only has components transverse to the direction of propagation (due to the orthogonality requirement and 2.22). Because $h_{\mu\nu}^{\text{TT}}$ is already trace-free, we can drop the bar from $\bar{h}_{\mu\nu}$. Also note that $C_{xx} = -C_{yy}$ (trace-free) and that $C_{xy} = C_{yx}$ (symmetry of $C_{\mu\nu}$). Since the wavevector is null we have $\omega = k_{tt} = -k_{zz}$ to give the final form of 2.23. Although we have defined here the TT gauge for a single mode, the TT gauge exists for any radiative wave.

With this in mind, we write the general form of a gravitational plane wave traveling in the z direction,

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+(t-z) & h_\times(t-z) & 0 \\ 0 & h_\times(t-z) & -h_+(t-z) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2.24)$$

h_+ and h_\times are independent real-valued functions representing the two polarizations of a gravitational wave.

2.2.2 Effect of gravitational waves on matter

Free particles will follow geodesics in curved space meaning their paths will satisfy the geodesic equation,

$$\frac{d^2 x^\rho}{d\tau^2} + \Gamma^\rho_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = 0. \quad (2.25)$$

A particle initially at rest with $dx^\mu/d\tau = (1, 0, 0, 0)$ will experience no change in coordinates in the presence of a gravitational wave as $\Gamma^\rho_{00} = 0$ in the TT gauge.

Therefore coordinates in the TT gauge follow geodesics of free particles at rest.

While we have chosen coordinates which do not change for a particle at rest in the presence of a gravitational wave, there is a change in the measurable distance between nearby particles. This is seen by considering the light travel time between two coordinates separated by δx : $\delta t^2 = (1 + h_+) \delta x^2$ (this is just the rearranged space-time interval for a null geodesic $\delta s^2 = 0$ in the perturbed metric with $\delta y = \delta z = 0$). The general effect of a gravitational wave on a region of nearby particles can be quantified in terms of the geodesic deviation which shows how a vector S^μ pointing from a reference geodesic to nearby geodesics changes in a curved spacetime. In this case the geodesics are the world lines of our free test particles at rest with four-velocities $U^\mu = (1, 0, 0, 0)$.

The equation of geodesic deviation is

$$\frac{D^2}{d\tau^2} S^\mu = R^\rho{}_{\sigma\mu\nu} U^\sigma U^\mu S^\nu, \quad (2.26)$$

which gives us for the vector S^μ between free particles [10, 6.62],

$$\frac{\partial^2}{\partial t^2} S^\mu = \frac{1}{2} S^\nu \frac{\partial^2}{\partial t^2} h^\mu{}_\nu. \quad (2.27)$$

Deviation occurs only in spatial directions perpendicular to the direction of wave propagation. For a wave traveling in the z direction we have,

$$S^x = \left(1 + \frac{1}{2} h_+\right) S^x(0) + \left(1 + \frac{1}{2} h_\times\right) S^y(0) \quad (2.28)$$

$$S^y = \left(1 - \frac{1}{2} h_+\right) S^y(0) + \left(1 + \frac{1}{2} h_\times\right) S^x(0). \quad (2.29)$$

2.2.3 Stress-energy of gravitational waves

The concept of local gravitational energy is inconsistent with general relativity. General relativity assumes no preferred coordinate system, and the equivalence principle tells us that it would be impossible to measure any kind of gravitational energy




Figure 2-1: The effect of a gravitational wave on a single ring of particles. The gravitational wave propagates into the page, and causes deviations of the apparent position S^μ of particles at rest from the perspective of an observer in the center as space is stretched and contracted (Equation 2.28). The ring is shown at various stages in its phase evolution. The top row shows the effect of a monochromatic wave at fixed intervals in phase with only h_+ polarization ($h_\times = 0$), and the bottom row shows pure h_\times polarization ($h_+ = 0$). Both waves are linearly polarized, and can be swapped under a $\pm\pi/4$ rotation about the direction of propagation. As in electromagnetism, circularly polarized states can be formed by a superposition of equal amplitude h_+ and h_\times components with a $\pi/2$ phase offset. A general gravitational wave with arbitrary h_+ and h_\times components is invariant under a rotation of π along the direction of propagation which is a reflection of the spin-2 nature of the field.

in a small enough laboratory undergoing free fall. However it is possible to define concepts of stress-energy of fluctuations on a comparatively smooth background curvature. Thus we require that the effective wavelength of fluctuations is much shorter than any background curvature: $\lambda \ll R$. For the same reason, the energy cannot be localized to within a wavelength of a particular fluctuation, but we must average over many wavelengths to get a meaningful measure. Even the energy radiated from a strong non-linear isolated source can be calculated this way as long as space becomes asymptotically flat at some sufficiently far away shell of integration surrounding the source.

Actual calculation of the energy in fluctuations on flat background requires a more careful approach than simple linearized theory presented here as the energy density

is quadratic in the gravitational field amplitude. The *shortwave formalism* [11, 12] handles more general gravitational waves by splitting the metric into background curvature and fluctuations, and keeping contributions up to second order in $h_{\mu\nu}$ in the Einstein curvature tensor. The stress-energy contribution from the fluctuations are identified by their second-order contribution to the background curvature in vacuum. In the transverse-traceless gauge it takes the form [9, 35.70],

$$T_{\mu\nu}^{(\text{GW})} = \frac{1}{32\pi} \langle \partial_\mu h_{jk}^{\text{TT}} \partial_\nu h_{jk}^{\text{TT}} \rangle, \quad (2.30)$$

where quantities between angle brackets $\langle \dots \rangle$ are to be averaged over many wavelengths. We see from Equation 2.30 that the energy density varies quadratically with both amplitude and frequency of a monochromatic wave of the form 2.23.

2.2.4 Generation of gravitational waves

Gravitational radiation is sourced by accelerating mass+energy m much in the same way that electromagnetic radiation is a result of accelerating charge q . Similarly, it is useful to decompose gravitational radiation into multiple moments. The monopole moment is simply the total mass of the system, which does not change and therefore cannot produce any radiation. Next would be mass (“electric”) and mass current (“magnetic”) dipole radiation, which would be caused by changes in the mass distributions,

$$\int \rho(\mathbf{x}) \mathbf{x} d^3x \quad (2.31)$$

$$\int \rho(\mathbf{x}) \mathbf{x} \times \mathbf{v}(\mathbf{x}) d^3x. \quad (2.32)$$

Changes in these distributions would violate conservation of linear and angular momentum, so there can be no gravitational dipole radiation as well. For slowly moving sources, we expect the dominant form of radiation to come from changes in the mass quadrupole moment,

$$I_{ij} = \int \rho(\mathbf{x}) x_i x_j d^3x, \quad (2.33)$$

which is not generally conserved.

Due to the complexity of the Einstein equations, it is difficult to solve a general form for the radiation from an arbitrary source. A useful approximation comes from solving the linearized Einstein equations (Equation 2.17) far from an isolated source which is moving at non-relativistic speeds. In this case we can apply the known general solution of the wave equation with a source and outgoing boundary conditions,

$$\bar{h}^{\mu\nu}(t, \mathbf{x}) = 4 \int d^3x' \frac{[T^{\mu\nu}(t', \mathbf{x}')]_{\text{ret}}}{|\mathbf{x} - \mathbf{x}'|}, \quad (2.34)$$

where the source $-16\pi T^{\mu\nu}$ is evaluated at the retarded time. In the long-wavelength (low velocity) far-field approximation, this becomes

$$\bar{h}^{\mu\nu}(t, \mathbf{x}) \xrightarrow{r \rightarrow \infty} \frac{4}{r} \int d^3x' T^{\mu\nu}(t - r, \mathbf{x}'). \quad (2.35)$$

Here we see that the amplitude of a gravitational wave diminishes linearly with the distance to an isolated source.

Using the flat-space conservation law $\partial_\mu T^{\mu\nu} = 0$ and non-relativistic approximation $T^{tt} = \rho$, it is possible to solve for the spatial components [9, 36.20],

$$h_{ij}^{\text{TT}} \xrightarrow{r \rightarrow \infty} \frac{2}{r} \ddot{I}_{ij}^{\text{TT}}(t - r). \quad (2.36)$$

This simple result is known as the quadrupole formula for gravitational wave generation. In order to project onto transverse traceless coordinates, we have used the transverse traceless projection of the reduced quadrupole moment I_{ij} ,

$$I_{ij} = I_{ij} - \frac{1}{3} \delta_{ij} I^k_k = \int \rho(\mathbf{x}) \left(x_i x_j - \frac{1}{3} \delta_{ij} r^2 \right) d^3x \quad (2.37)$$

$$\mathcal{I}_{ij}^{\text{TT}} = P_{ia} I_{ab} P_{bj} - \frac{1}{2} P_{ij} P_{ab} I_{ab} \quad (2.38)$$

$$P_{ab} = \delta_{ab} - \hat{n}^a \hat{n}^b \quad (\text{projection operator onto } \mathbf{n}) \quad (2.39)$$

The gravitational wave amplitude depends on the second time derivative of the reduced quadrupole moment. By plugging into Equation 2.30 and integrating around

all solid angles, we get for this limit an expression for the total luminosity of a source in gravitational waves,

$$L_{GW} = \frac{1}{5} \langle \ddot{\mathbf{F}}_{ij} \ddot{\mathbf{F}}_{ij} \rangle. \quad (2.40)$$

2.3 Sources of gravitational radiation

2.3.1 Compact binary coalescence

As two gravitationally bound compact objects such as neutron stars (NS) or black holes (BH) orbit each other, they gradually lose orbital energy and angular momentum to gravitational radiation and proceed toward an inevitable collision. During the final highly relativistic stages of this process, an immense amount of gravitational radiation corresponding to as much as 5–10% of the total system mass is released in a very short amount of time making these sources the most distant objects potentially visible to gravitational wave detectors.

The existence of NS/NS binary systems has been confirmed through several galactic observations where one of the objects is a radio pulsar. In this case, orbital parameters can be deduced from the periodic Doppler shifting of the radio signals. The first and most famous pulsar in a binary system found to date is PSR 1913+16 which was discovered in 1974 by Russell Hulse and Joseph Taylor [6]. Careful measurement of the orbital decay over many years [13] showed remarkable agreement with the prediction from general relativity of energy loss due to gravitational radiation. Hulse and Taylor received the 1993 Nobel prize in physics for their discovery. While the Hulse-Taylor binary will take ~ 300 million years to coalesce, the number of observed binary pulsar systems can be used to infer rates of compact object coalescence in Milky Way type galaxies. These predicted rates, along with rates predicted from stellar population synthesis are presented in Table 2.1.

Compact object mergers where at least one member is a neutron star may be the mechanism which drives gamma-ray bursts (GRBs) [15]. The short (<2 seconds) burst of hard gamma-rays are thought to arise from shocks created by the relativistic

Compact binary coalescence rates per L_{10} per Myr.					
Source	R_{low}	R_{likely}	R_{high}	$R_{\text{LIGO I}} [\text{yr}^{-1}]$	$R_{\text{AdvLIGO}} [\text{yr}^{-1}]$
NS/NS	0.6	60	600	0.02	40
NS/BH	0.03	2	60	0.004	10
BH/BH	0.006	0.2	20	0.007	20

Table 2.1: The compact binary coalescence rates (lowest possible rate, likely rate, and highest possible rate given observations) per 10^{10} solar blue light luminosities per million years as estimated from the number of observed binary pulsar systems, as well as from stellar population synthesis models. One Milky Way equivalent galaxy corresponds to approximately $1.7 L_{10}$. Contributions from elliptical galaxies and old stellar populations are not explicitly included. Also shown are predicted likely detection rates for initial and advanced LIGO detectors. The low and high estimates can differ from this likely estimate by as much as three order of magnitude. Quoted rates are taken from [14].

jets produced. That this process happens so quickly and that short GRBs have been associated with old elliptical galaxies [16] with no young massive stars support the merger theory. If this is the case, there is the possibility that a nearby short-GRB could be associated with a characteristic gravitational-wave signature in ground-based detectors, confirming the model and providing unique observation in gravitational waves of the tidal breakup of the neutron star [17, 18].

The adiabatic inspiral phase of a binary coalescence lasts while the objects are in stable orbits which decay only as a result of gravitational wave emission. As the objects lose orbital energy, the period of their orbit shrinks and gravitational wave amplitude increases due to the increased velocities. This gives a characteristic chirp waveform. The gravitational wave emission is well approximated by post-Newtonian (PN) methods which analytically solve the field equations for the two point masses with arbitrary spin in quasi-circular orbit to some order in $(v/c)^2$ or equivalently M/R of the binary by the virial relationship of the gravitationally bound system. At lowest order, the PN formalism gives the quadrupole formula for gravitational wave emission. For binary systems, the PN approximation has been calculated to much higher orders [19] to get the accuracy required at small separations and high velocities.

When the two objects pass their innermost stable circular orbit (ISCO) at $r \simeq 6M$,

$f_{\text{ISCO}} \simeq 4.2(M_{\odot}/M)$ kHz, they inevitably fall into each other regardless of gravitational radiation emitted. This marks the end of the adiabatic inspiral and the beginning of the merger. The highly relativistic conditions provide a unique opportunity to observe strong field gravitational effects [20]. For the same reason, however, post-Newtonian approximations break down and no complete analytic solution of a general merger exists. Fortunately numerical relativity, after decades of work, has recently been able to model complete black-hole binary coalescence [21].

The final phase in binary coalescence is characterized by the ringdown of the resulting black hole as it settles into its stable Kerr state. The dynamics are studied using black hole perturbation theory, and can be described as excitations of the quasinormal modes of the black hole which decay exponentially due to the emission of gravitational waves at the corresponding frequencies [22, 23]. Eventually the lowest mode dominates and we expect all coalescence waveforms to end with an exponentially damped sinusoid at the $l = m = 2$ quasinormal mode fundamental frequency and Q ,

$$f_{\text{QNM}} \simeq (2\pi M)^{-1} [1 - 0.63(1 - a)^{0.3}] \quad (2.41)$$

$$Q \simeq 2(1 - a)^{-0.45}. \quad (2.42)$$

A $10 M_{\odot}$ black hole has a $f_{\text{QNM}} \simeq 3.2$ kHz.

Searches for binary compact object coalescence are typically done with matched filter techniques. This procedure uses a known waveform weighted against the stationary spectral noise as a linear filter to check for the presence of a signal. For a single waveform, the matched filter maximizes the signal-to-noise ratio in the case of a match. A neutron star binary system will chirp outside of the LIGO sensitive band of around 100–150 Hz well before reaching its ISCO, so PN inspiral waveforms are sufficient for a template search. Similarly, a massive black hole system with total mass of several hundred M_{\odot} will have most of its inspiral and early merger hidden by the low frequency noise in ground based detectors. These can be found efficiently by a template of ringdown waveforms. Such searches have been performed on recent data from the LIGO and Virgo instruments [24, 25, 26]. The inspiral searches in LIGO

S5 data can detect sources out to 30–100 Mpc depending on mass, and set 90% rate upper limits for binary coalescence with total mass 2–35 M_{\odot} of 1.4×10^{-2} (NS/NS), 7.3×10^{-4} (BH/BH), and 3.6×10^{-3} (BH/NS) $\text{yr}^{-1} L_{10}^{-1}$ where L_{10} represents 10^{10} solar blue light luminosities. Published results for searches for intermediate mass black hole ringdowns in LIGO S4 data set a limit of $1.6 \times 10^{-3} \text{ yr}^{-1} L_{10}^{-1}$ for binary mergers with total mass 85–390 M_{\odot} , which can be seen out to ~ 85 Mpc in S4 data.

For a large range of total mass between 20–500 M_{\odot} , a significant amount of the detectable signal for current ground based interferometers is present in the binary merger phase [27]. The best sensitivity will be achieved by coherently integrating signal energy across the entire binary evolution. However this is only possible with matched filtering if the waveform is known. In 2005, the first successful equal mass, non-spinning complete black hole merger simulation was achieved [28]. Since then, many groups have run successful simulations using a wide variety of codes and formalisms [21]. The results have been consistent with each other, as well as with PN approximations for the inspiral and perturbative ringdown. The numerical simulations have explored a large portion of the binary black hole parameter space, including precessing orbits with misaligned spin [29]. However, running simulations is computationally expensive, and it is currently not possible to build up a template bank that sufficiently samples the parameter space for a real search.

To be able to make a set of templates suitable for matched filter analysis, hybrid semi-analytic waveforms have been developed. In one method, phenomenological waveforms [30] are created by stitching PN inspirals to numerical mergers. A certain simple parametrization in the frequency domain is then assumed to represent the general hybrid waveform. The parametrization has several coefficients which can be smoothly varied to fit any particular binary configuration. The latest parametrization have been able to include the effects of unequal mass and aligned and anti-aligned spins [31]. The effective one-body (EOB) formalism [32] has also been able to successfully model a complete binary coalescence. In the EOB approach, the dynamics of the binary system are mapped to those of a single test particle in a effective external metric which changes as a result of radiation damping. This gives a resummation of

the PN series approximation, with the advantage of providing a smooth transition to merger and ringdown of a single object. The lower order terms are determined by their correspondence with the known PN expansion, while higher order PN and other adjustable terms are fit to numerical relativity simulations. The EOB-NR technique has been used to model unequal mass black hole mergers with aligned and anti-aligned spins as well [33]. The general case of misaligned spins which cause precession and changes in the orbital plane has yet to be modeled.




Figure 2-2: Numerical simulation of a binary black hole merger using SpEC [34] with EOB-NR waveform superimposed. The initial configuration is equal mass black holes with spins ~ 0.44 opposite the orbital angular momentum. The top plots show the waveforms, where the right plot is a zoom-in on the merger and ringdown. The EOB-NR waveform is an effective one-body PN expansion with 6 adjustable terms calibrated by the numerical simulation. The bottom plots shows the phase and amplitude consistency between the waveforms. This figure is reproduced from [33]

2.3.2 Gravitational core collapse

Gravitational core collapse of massive stars is another highly anticipated source of gravitational radiation. Core collapse is the natural conclusion of main sequence stellar evolution for stars $\gtrsim 9 M_{\odot}$. These central areas of these stars are hot and dense enough to form iron cores as the end product of nuclear fusion. As the degenerate iron core accumulates mass and contracts under gravity, the environment becomes so hot and dense that photodissociation of nuclei and electron capture by protons can occur. These two processes remove thermal and electron degeneracy pressure—the only means of support, and thus trigger a gravitational collapse where the core shrinks to a radius of ~ 20 km in ~ 1 second [35].

The equation of state rapidly stiffens as the core becomes dense enough for neutron degeneracy pressure to become significant. For stars with an initial main sequence mass of $\lesssim 25 M_{\odot}$ [36], the resulting bounce of the inner core and outward shock wave when it hits the falling outer layers, along with neutrino pressure from the hot, convective protoneutron star, carry enough of the released gravitational energy to blow off the surrounding envelope of the star in a Type II supernova (SN). A hot, rapidly-rotating neutron star remnant is left. For heavier stars, the shock plus neutrinos do not transfer enough energy to unbind all the outer material. A weak, Type Ib supernova is expected, and a short lived proto neutron star eventually collapses to a black hole due to matter infall. Stars with initial mass $\gtrsim 40 M_{\odot}$ and low metallicity collapse directly to black holes due to their larger cores at time of collapse. Long-duration gamma-ray bursts, which are the most commonly observed, have been connected to the collapse of very massive stars in Type Ic supernovae [37], and are thought to be caused by relativistic jets created by the accretion disk of stellar matter surrounding a newly formed black hole.

The most notable observational evidence to support the massive stellar core collapse and subsequent explosion model of Type II supernovae is the coincident observation of a burst of 20 neutrinos from the Kamiokande II and IMB water Cherenkov detectors approximately three hours before visible observation of SN1987a [38]. Mod-

ern neutrino detectors should see around 2000 neutrinos from a galactic Type II supernova. Planetary nebula have also been linked to the expelled stellar envelopes of historical supernovae. The Crab nebula has been identified with a bright supernova recorded by Chinese and Arab astronomers in 1054. At the center of the nebula is the Crab pulsar - the rapidly rotating neutron star remnant which was the first of its type to be found in 1968.

While perfectly spherical collapse would not emit any gravitational radiation, the collapse of a rotating core will instead flatten along the axis of rotation due to centrifugal forces. The axisymmetric but aspherical collapse, bounce, and subsequent core oscillations are a reliable source of linearly polarized gravitational radiation. Convection and anisotropic neutrino emission are additional significant sources of collapse asymmetry in existing simulations [39]. Triaxial rotational instabilities can develop in the proto neutron star itself such as the standard bar mode instabilities, or radiation reaction driven r -modes. These draw energy from the core's rotation and can be a powerful source of gravitational radiation. In the collapse of a very massive star, accretion onto a newly formed black hole can excite the quasinormal modes of the black hole, which ring down through gravitational radiation.

Numerical simulation of stellar collapse is a very active field as the exact process of collapse has yet to be definitively answered. The field has provided a number of sophisticated 2+1D axisymmetric and fully 3D models, with the 2D models able to allow for a larger parameter space and more computationally demanding microphysics at the expense of being able to investigate non-axisymmetric behavior which tends to develop after the initial collapse [40]. Many models produce radiation strong enough for a galactic supernova, with a rate of one per 50 years, to be detected by first generation broadband interferometric detectors. Advanced detectors will extend the range by a factor of ten into the local group, but it will likely be third generation detectors which will be able to reliably detect supernovae from the Virgo cluster at a rate of 1 per year.




Figure 2-3: Gravitational waveform from a $20 M_{\odot}$ rotating stellar collapse. At a galactic distance of $D = 10$ kpc, 100 cm on the vertical axis corresponds to a strain of 3.24×10^{-21} . The two polarizations are shown in separate plots, while colors distinguish equatorial observers (black) and those looking down the pole (red). In this 3D simulation, the initial collapse remains axisymmetric emitting a typical linearly polarized bounce waveform perpendicular to the rotation axis, while non-axisymmetric instabilities contribute to circularly polarized radiation along the rotation axis beginning ~ 30 ms post bounce. Approximately $7.5 \times 10^{-8} M_{\odot}$ is released in gravitational radiation with about equal contributions from bounce and late-time 3D dynamics. This figure is reproduced from Ott et al [40]

2.3.3 Rotating neutron stars

Any kind of off-axis asymmetry in a rapidly rotating neutron star will radiate gravitationally. The asymmetry could be due to a misaligned stiff crust, electromagnetic or accretion induced anisotropies, precession, or the excitation of some form of rotational instability. Gravitational radiation, regardless of mechanism, will draw rotational energy and angular momentum out of the star, slowing it down. In many cases this process is quite slow and results in weak radiation over a timescale that is much longer

than a typical search. For this reason, such objects are often referred to as continuous wave sources.

Weak deviations from mass axisymmetry along the rotational z axis can be parametrized by the ellipticity for a triaxial ellipsoid, written here in terms of the moment of inertia tensor $\mathcal{I}_{ij} = \delta_{ij}I^k_k - I_{ij}$ (see Equation 2.33), along principal axes,

$$\epsilon = \frac{\mathcal{I}_{xx} - \mathcal{I}_{yy}}{\mathcal{I}_{zz}}. \quad (2.43)$$

Neutron star crusts are thought to be able to support an ellipticity no larger than $\sim 10^{-6}$. The characteristic strain observed at the Earth at rotational frequency ν is

$$h_0 = \frac{16\pi^2 G}{c^4} \frac{\epsilon \mathcal{I}_{zz} \nu^2}{r}. \quad (2.44)$$

The gravitational wave frequency will be at 2ν since, as typical of rotating systems, the quadrupole moment goes through two cycles for every rotation. While the gravitational wave emission is essentially monochromatic, the observed signal at the Earth is Doppler shifted by non-negligible amounts by the Earth's motion in the galactic frame.

Loss of angular momentum through gravitational radiation is a leading candidate to explain the observed sharp cutoff at ~ 720 Hz in the rotation rate distribution of millisecond accreting X-ray pulsars [41]. The neutron stars are spun up by accretion from a companion star, yet the maximum observed frequency is well below their breakup frequency of ~ 1500 Hz for typical equations of state. Bildsten originally postulated [42] that accreting X-ray pulsars would reach an equilibrium between spin up due to accretion and spin down due to gravitational radiation, making it difficult such systems to reach frequencies greater than 600–800 Hz. If angular momentum loss through gravitational waves dominates over electromagnetic effects, their strength is set to balance the torques due to mass transfer. In this optimistic limit, a few of the brightest sources corresponding to the largest accretion rates may be detectable by Advanced LIGO [43], especially if the instrument is tuned for the frequencies of interest.

This type of upper limit for gravitational radiation can be established for any known pulsar where the frequency ν and spin down rate $\dot{\nu}$ are observed through radio or X-ray timing. Assuming the spin down is due completely to gravitational radiation, the strain at the Earth would be

$$h_{\text{sd}} = \left(\frac{5 G \mathcal{I}_{zz} |\dot{\nu}|}{2 c^2 r^2 \nu} \right)^{1/2}. \quad (2.45)$$

The spin down limit on gravitational radiation depends on the moment of inertia \mathcal{I}_{zz} as this is what determines the rate of angular momentum loss given a measured spin down rate. In many cases, one can assume canonical value of $\mathcal{I}_{zz} = 10^{38} \text{ kg m}^2$ which corresponds to a $1.4 M_{\odot}$ neutron star with uniform density and radius 10 km. To date, the only known pulsar for which a direct observational limit on gravitational radiation is smaller than the limit imposed by spin down is the Crab pulsar. The Crab pulsar has an unusually high spin down limit of $h_{\text{sd}} = 1.4 \times 10^{-24}$ due to its close proximity and high spin down rate, and data from LIGO's fifth science run has been used to beat this limit on gravitational radiation by a factor of 7.2 for the most straightforward model of emission [44]. This is not surprising given the large expected electromagnetic torques slowing down the young supernova remnant.

In addition to rotating mass asymmetry, neutron stars have a rich spectrum of non-radial quasi-normal mode oscillations which can radiate gravitationally [45]. The modes can be easily excited by the core bounce and fall-back during a gravitational collapse, but also by other sources like irregular accretion and crust quakes which could accompany a pulsar glitch or soft gamma repeater (SGR) flare. Unlike the triaxial rotational instabilities, these pulsations do not require rotation to occur, but the rotation profile will affect their frequencies and mixing in non-trivial ways [46]. In neutron star astroseismology, one uses information carried by the gravitational waves from normal mode oscillations to model fundamental properties of the neutron star.

2.3.4 Stochastic background

The stochastic gravitational wave background is a random field of gravitational perturbations throughout all space. The cause of the perturbations may be an ensemble of unresolved astrophysical point sources, or they may be a result of fundamental processes in the early universe. Gravitational waves from the early universe are particularly interesting because they decouple from other particles at extremely high energies. The universe remained opaque to electromagnetic radiation until the epoch of recombination around 3.8×10^5 years after the Big Bang, and the cosmic microwave background (CMB) radiation we detect today gives us a fairly undisturbed snapshot of the structure of the universe at that time. We expect a similar neutrino background from ~ 2 seconds after the Big Bang, although their low energy today of ~ 1.95 K makes direct detection very remote. A gravitational wave background from as early as the Planck era ($\sim 10^{-42}$ s) is also expected to be very difficult to detect although many processes in the very early universe are expected to leave signatures which could be observed by future instruments. Gravitational waves produced in the early universe will have maximum wavelength on the order of the horizon size at time of production, so that waves entering as early as $\sim 10^{-25}$ seconds would be redshifted into the LIGO sensitive band today [47].

Because the spacetime perturbations carry energy, they contribute to the total energy density of the universe. This contribution is generally represented in terms of the gravitational wave energy density per unit logarithmic frequency,

$$\Omega_{\text{GW}}(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}. \quad (2.46)$$

The energy density is written as a fraction of the critical density needed to make the universe flat in standard cosmology given the measured Hubble expansion rate $H_0 \simeq 72 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [48, 49],

$$\rho_c = \frac{3c^2 H_0^2}{8\pi G}. \quad (2.47)$$

The flatness of the universe has been verified to $\sim 2\%$ using CMB measurements.

However, it is often assumed to be exactly flat as a consequence of inflation.

The inflationary stochastic gravitational wave background arises from quantum vacuum spacetime perturbations which were stretched and amplified to classical scales during inflation. Exponential inflation contributes a flat spectrum to $\Omega_{\text{GW}}(f)$ today over a very large range of frequencies [50]. The current best limits on gravitational waves from inflation come from the smallness (10^{-6}) of large scale temperature anisotropies measured in the CMB. Long wavelength gravitational waves create observed temperature anisotropies through gravitational redshift, known as the Sachs-Wolfe effect. The low modes ($2 \leq l \leq 30$) in the CMB set a limit of $\Omega_{\text{GW}}(f) < 10^{-13}$ at 10^{-16} Hz [47], which for a flat $\Omega_{\text{GW}}(f)$ spectrum is well below the noise in other measurements of the gravitational wave spectrum. Future measurements of CMB polarization may be able to decrease this limit or detect gravitational waves from inflation by measuring their unique contribution to polarization B modes [51].

The gravitational wave spectrum is much more poorly constrained at higher frequencies. The total energy in gravitational radiation,

$$\Omega_{\text{GW}}^{\text{Total}} = \frac{\rho_{\text{GW}}}{\rho_{\text{c}}} = \int d(\ln f) \Omega_{\text{GW}}(f), \quad (2.48)$$

cannot exceed 1.5×10^{-5} for gravitational waves which were present at the time of Big Bang nucleosynthesis [52]. A larger amount of gravitational radiation would increase the relative radiation density of the universe, which speeds up expansion and implies an earlier freeze-out for neutron-proton thermal equilibrium. The resulting overabundance of neutrons due to a higher freeze-out temperature is inconsistent with the observed fraction of ${}^4\text{He}$ [50]. The limit only constrains gravitational waves present during Big Bang nucleosynthesis minutes after the Big Bang. The wavelength is limited by the horizon scale at that time, corresponding to frequencies today which are greater than 10^{-10} Hz. A similar bound on the total radiation density during the epoch of recombination can be made by calculating its effect on matter and CMB power spectra. The total energy in gravitational radiation when the CMB decoupled is constrained to be less than 1.3×10^{-5} [53] and includes gravitational waves which

are now greater than 10^{-15} Hz. This limit is projected to improve as the CMB power spectrum becomes better constrained.

Precision pulsar timing of millisecond pulsars has established a bound of $\Omega_{\text{GW}}(f) < 3.7 \times 10^{-8}$ at 4×10^{-9} Hz [54]. A stochastic gravitational wave background contributes to timing residuals of the extremely regular pulses. Measurements of correlations between the two LIGO sites over the course of two years have also been used to set a direct upper limit on the stochastic gravitational wave background of $\Omega_{\text{GW}}(f) < 6.9 \times 10^{-6}$ at ~ 100 Hz [55]. Advanced LIGO will be sensitive to an $\Omega_{\text{GW}}(f)$ which is smaller by a factor of 100 assuming a factor of 10 increase in strain sensitivity. This will test cosmic-string [56] and pre-Big Bang [57] models which predict significant high frequency stochastic gravitational wave backgrounds.

The energy density in gravitational waves is determined by the average amplitude of their fluctuations, and thus we can relate the gravitational wave energy spectrum to an observable power spectrum of strain fluctuations (signal power per unit frequency),

$$S_{\text{GW}}(f) = \frac{3H_0^2}{10\pi^2} f^{-3} \Omega_{\text{GW}}(f). \quad (2.49)$$

The stochastic gravitational wave background contributes directly to the strain power spectrum of a gravitational wave detector such as LIGO. For the LISA space based interferometric detector, the galactic population of unresolved white dwarf binary systems is expected to be a limiting noise source at low frequencies. Better sensitivity can be achieved for two uncorrelated detectors by looking for an excess in correlated noise, which is caused by gravitational wavelengths which are longer than the distance between instruments. The exact parametrization of the degree of correlation is called the overlap reduction function, $\gamma(f)$, which is a function of frequency as well as distance and relative orientation between the detectors [58]. An optimal filter to detect a particular spectrum of gravitational radiation will also weigh contributions by frequency according to the individual detector sensitivities and their overlap. The sensitivity of such a measurement scales with the square root of observation time.




Figure 2-4: Cosmological source models as well as limits of and future sensitivities to a stochastic gravitational wave background. The limits based on BBN constraints as well as from CMB and matter spectra are integral limits on the total energy in gravitational radiation over the frequencies shown, in units of the critical density of the universe. The models shown for standard inflation, cosmic strings, and pre-Big Bang scenarios, as well as remaining limits from large angle CMB anisotropy, millisecond pulsar timing, and current and future ground and space based interferometric detectors are limits on gravitational wave energy density per unit logarithmic frequency: $\Omega_{\text{GW}}(f) = d\Omega_{\text{GW}}^{\text{Total}}/d \ln f$. This figure is reproduced from [55].

Chapter 3

Detecting gravitational waves

3.1 Resonant mass detectors

Passing gravitational waves transfer a small amount of energy to a rigid body as atomic forces fight the stretching and squeezing of space. In a resonant mass (bar) detector, gravitational waves at the fundamental frequency of the bar will cause the greatest amplitude of excitations, which are converted into signals using sensitive electromechanical transducers attached to the bar. The first such devices were designed and tested in the 1960s by Joseph Weber [1, 59]. Weber used two 1.5 ton cylindrical aluminum bars with a strain sensitivity of 10^{-16} at 1660 Hz to search for coincidences from astrophysical gravitational waves. Today's bar detectors are ~ 3 meters with masses $\sim 1,000$ kg and have fundamental longitudinal mode resonant frequencies about 1,000 Hz. When cooled to temperatures as low as 0.1 K, they reach sensitivities of $\sim 10^{-21}$, approaching their quantum limit for zero-point energy fluctuations.

The three primary noise sources for resonant mass detectors are sensor noise, thermal noise, and recently quantum noise. Sensor noise refers to noise in reading the position of the bar endpoints. The sensor noise is the limiting noise source away from resonance where the amplitude of oscillations becomes small. Modern detectors make use of resonant sensors matched to the bar frequency which allow gravitational waves near the resonant frequency to first transfer energy to the bar whose oscillations then drive the resonant sensor. The bandwidth of such detectors is limited by the time it

takes to drive the resonant sensor to peak amplitude. Random thermal fluctuations in the bar also cause noise which can be large compared to an expected gravitational-wave induced perturbation. To minimize thermal noise, resonant bars are cooled to below liquid helium temperatures. This also maximizes the quality factor, Q , of the bars mechanical resonance which keeps the thermal fluctuations restricted to a very narrow bandwidth about resonance. If sensor and thermal noise are made small enough, a bar detector reaches a fundamental quantum uncertainty in any measurement of the position of the bar endpoints. This limits strain sensitivity to $\sim 10^{-21}$ for a typical bar configuration, just at the edge of being astrophysically interesting.

Current cryogenic gravitational wave bar detectors include the ALLEGRO detector in Baton Rouge, Louisiana [60], the AURIGA detector in Lengaro, Italy [61], the EXPLORER detector in Geneva, Switzerland [62], and the NAUTILUS detector in Rome, Italy [62]. The four detectors took part in a coincident run of data-taking from 2005-2007 under the second International Gravitational Event Collaboration (IGEC-2) [63, 64]. The joint analysis, sensitive to short <1 s signals with frequency content above the detector noise, found no gravitational wave candidate events. Interferometric detectors such as LIGO and Virgo have since eclipsed the bar detectors in sensitivity and bandwidth. The ALLEGRO detector ceased operation in 2007. The remaining detectors will provide important coverage during the substantial down-time of interferometric detector upgrades over the next several years, and they will also provide enhanced detection confidence for the case of a loud gravitational-wave event at bar frequencies.

A cylindrical resonant bar is sensitive to strain along its main axis which can excite the bar's fundamental longitudinal mode. The geometric projection of a gravitational wave from arbitrary sky location onto components of strain measured by the detector is represented by the detector antenna response. The two gravitational wave polarizations, $h_+(t)$ and $h_\times(t)$ contribute a measured strain of,

$$h(t) = F_+ h_+(t) + F_\times h_\times(t), \quad (3.1)$$




Figure 3-1: Typical strain noise sensitivities of current cryogenic resonant bar detectors for data taken during 2005-2007. This figure is reproduced from [64]

where F_+ and F_\times are the antenna factors for plus and cross polarization. For a bar detector, these antenna factors are,

$$F_+ = \sin^2 \theta \cos 2\psi \quad (3.2)$$

$$F_\times = \sin^2 \theta \sin 2\psi, \quad (3.3)$$

where θ represents the angle of the source from the bar axis, and ψ represents the orientation of the gravitational wave polarization axes with respect to an orientation defined by the bar axis-source plane. Resonant bars are sensitive to a large area of the sky perpendicular to their axis. Spherical resonant detectors have been proposed that would provide a uniform coverage of the sky with no blind spots. Because they have no preferred oscillation axis, they can also resolve both gravitational wave polarizations, and provide some directional information for a single source. Two prototype detectors include the MiniGRAIL sphere [65] in the Netherlands, and the Schenberg detector[66] in Brazil.

3.2 Interferometric detectors

Interferometry provides very precise measurement of relative length changes. A Michelson interferometer with a single beam splitter and two arms at 90 degrees is particularly well suited to measure the differential strain produced by a passing gravitational wave. Over bar detectors, interferometers have the fundamental advantage of being able to be made much larger. As they do not depend on any resonance between the instrument and signal, they are also broadband detectors sensitive to a much larger range of frequencies than resonant bars. Development of laser interferometric gravitational wave detectors began in the late 1960s with the independent work of Weiss [2], Moss, Miller, and Forward [3]. Only in the last several years with the successful construction of first generation km-scale interferometers have they met and surpassed bar detector sensitivity.

In an idealized Michelson interferometer, a single beam of light is sent through a 50:50 beam splitter at a 45° angle. Half the light intensity is transmitted down the same direction as the incident beam, and half is reflected at 90° . Each beam is reflected back at some distance away by end mirrors and they recombine at the same beam splitter to form a reflected beam back toward the original source direction, and an anti-symmetric output opposite the first reflected beam. A monochromatic laser input will have electric field,

$$E_{\text{in}} = E_0 e^{i(\omega t - kx)}. \quad (3.4)$$

The 50:50 beam splitter may be characterized by amplitude transmission coefficient $i/\sqrt{2}$ and reflection coefficient $1/\sqrt{2}$, while the end mirrors have a reflection coefficient of -1 . The field exiting the antisymmetric port is a superposition of the portion of the beam first transmitted then reflected by the beam splitter, and the portion reflected first then transmitted by the beam splitter on the way back. The amplitude of the field exiting the antisymmetric port is,

$$|E_{\text{antisym}}| = E_0 \cos k(L_x - L_y). \quad (3.5)$$

A photodetector placed at the antisymmetric port will measure the intensity of exiting light which is related to the average squared electric field, and is therefore related to the input power by,

$$P_{\text{antisym}} = P_{\text{in}} \cos^2 k(L_x - L_y). \quad (3.6)$$

Measurement at the antisymmetric port reflects the accumulated phase difference between the light traveling in each arm of the interferometer.

A linearly polarized gravitational wave passing through the plane of the interferometer and aligned to the interferometer arms is optimally oriented to produce a change in the differential arm length. Like bar detectors, the single strain signal measured by an interferometer from a source with arbitrary direction and polarization is subject to a geometric projection factor,

$$h(t) = F_+ h_+(t) + F_\times h_\times(t). \quad (3.7)$$

A detector with arms along \hat{x} and \hat{y} has antenna factors,

$$F_+ = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi \quad (3.8)$$

$$F_\times = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \sin 2\psi + \cos \theta \sin 2\phi \cos 2\psi, \quad (3.9)$$

where (θ, ϕ) represents the standard spherical coordinate system inclination and azimuthal angles from zenith and the $y = 0$ plane respectively, and ψ is a polarization angle which represents the angle from the coordinate axes defined by $(-\hat{\phi}, \hat{\theta})$ at a particular point in the sky to the chosen gravitational wave (x, y) coordinate axes. This is a choice of convention as the gravitational wave can be written with any choice of polarization axis.

The geometric antenna pattern is only valid in the long wavelength limit where the gravitational wavelength is much longer than the total light travel distance in the interferometer so that the instrument is at any particular time measuring an essentially static gravitational field. For shorter wavelengths, or equivalently high gravitational wave frequencies, the light circulating the interferometer arms is both

redshifted and blueshifted before being recombined, thus canceling the signal. This is further complicated by the resonant cavities used in modern interferometers where light is trapped for a variable number of trips in each arm before returning to the beam splitter. Longer arm lengths do result in proportionally greater path length changes at fixed strain amplitude, so the ideal size of an interferometer will balance the two effects in addition to practical considerations.

First generation large scale interferometric gravitational wave detectors have come online within the last decade. These include the two 4 km and one 2 km (currently decommissioned) LIGO detectors in Hanford, Washington and Livingston, Louisiana [4], the 3 km Virgo detector in Cascina, Italy [5], the 600 meter GEO600 detector near Hannover, Germany [67], and the 300 meter TAMA300 detector outside Tokyo, Japan [68]. The LIGO and GEO600 detectors operate within the LIGO Scientific Collaboration, which has entered into close data sharing and joint analysis agreements with the Virgo Collaboration in order to make best use of the scientific data. While no gravitational wave events have yet to be observed (consistent with predicted detection rates at current sensitivities), the successful technological development and operation of first generation detectors is an important step toward the realization of advanced detectors sensitive enough to make regular detections.

3.3 Laser Interferometer Gravitational-wave Observatory (LIGO)

The Laser Interferometer Gravitational-wave Observatory (LIGO) consists of three kilometer-scale gravitational wave interferometers located at two sites. An observatory at Hanford, Washington houses a 4 km baseline (H1) and 2 km (H2) detector which share a common vacuum enclosure and seismic isolation. Another nearly identical observatory three thousand miles away at Livingston, Louisiana has a single 4 km detector. Between November 2005 and September 2007 the detectors engaged in their fifth science run (S5) where they operated at their design sensitivity and collected in




Figure 3-2: A simple two-mirror Fabry-Perot resonant cavity of length L . The reflection and transmission coefficients for each mirror must satisfy energy conservation: $r^2 + t^2 = 1 - \text{losses}$. Light reflecting from the inner (coated) mirror surface suffers a sign change. There are both right (shown) and left circulating fields. Typically the mirrors have high reflectivity. On resonance, light escaping the cavity interferes destructively with the promptly reflected input light from the first mirror causing laser light to build up between the two mirrors.

total a year of science data with all three detectors operating in coincidence. The LIGO instruments H1 and L1 have since undergone minor upgrades for a currently running sixth science run. The upgrades are part of a transition to Advanced LIGO which will require several years of commissioning beginning in late 2010.

The LIGO instruments are modeled after Michelson interferometers with suspended end mirrors that act as freely falling gravitational test masses over the frequency range of interest (40–7000 Hz). Through S5, the instruments operated on a dark fringe with no light exiting the antisymmetric port without the presence of a gravitational wave (or some other disturbance). To increase the effective arm length as well as the total amount of circulating light sampling the gravitational wave, the instruments make wide use of resonant cavities which we will review here briefly due to their importance.

3.3.1 Fabry-Perot cavities

Each arm of a LIGO interferometer is a Fabry-Perot resonant cavity, which traps light between two parallel mirrors. A Fabry-Perot cavity is a kind of interferometer in itself as light in the cavity interferes with other beams of light which have made various numbers of paths between the optics.

The steady state fields for a Fabry-Perot cavity for a constant input laser source are given by the superposition of light from all possible paths between both mirrors. A field E_0 incident on a mirror is split into a reflected field of amplitude $\pm rE_0$ depending on which side of the mirror reflection takes place and transmitted field of amplitude tE_0 . The right circulating field in a two mirror cavity of length L (Figure 3-2) with reflection and transmission coefficients $r_{1,2}$ and $t_{1,2}$ is found by adding up all round trips of the incident light,

$$\frac{E_{\text{circ}}}{E_{\text{in}}} = t_1 \sum_{n=0}^{\infty} (r_1 r_2 e^{-2ikL})^n = \frac{t_1}{1 - r_1 r_2 e^{-2ikL}}. \quad (3.10)$$

The cavity is on resonance when $2kL = 2\pi n$ for some n . For highly reflective mirrors, the denominator becomes very small, and the amount of light trapped can be very large. The transmitted field is the portion of this field that escapes through the second mirror,

$$\frac{E_{\text{trans}}}{E_{\text{in}}} = t_2 e^{-ikL} E_{\text{circ}} = \frac{t_1 t_2 e^{-ikL}}{1 - r_1 r_2 e^{-2ikL}}. \quad (3.11)$$

A cavity with highly reflective mirrors acts as a filter for the light, only passing through modes which are resonant in the cavity. The reflected field is the sum of the promptly reflected light with the light escaping the cavity back toward the source,

$$\frac{E_{\text{refl}}}{E_{\text{in}}} = r_1 - r_2 t_1 e^{-2ikL} E_{\text{circ}} = r_1 - \frac{r_2 t_1^2 e^{-2ikL}}{1 - r_1 r_2 e^{-2ikL}}. \quad (3.12)$$

Off resonance, most of the light is promptly reflected if the first mirror has high reflectivity. The reflected light can also be written in the form,

$$\frac{E_{\text{refl}}}{E_{\text{in}}} = \frac{r_1 - r_2(1 - l_1)e^{-2ikL}}{1 - r_1 r_2 e^{-2ikL}} \quad (3.13)$$

where $l_1 = 1 - r_1^2 - t_1^2$ represents small absorption and scattering losses due to the first mirror. When both mirrors are identical and losses are negligible, a special case arises for resonance where all the light is transmitted by the cavity and none is reflected back.

For the LIGO arm cavities, the end mirrors are made to be as reflective as possible. The input mirrors have small nonzero transmission. In this case, the coefficient for reflected light from the cavity becomes very close to -1 . Near resonance, the reflected light can be expanded to first order about $L + \delta L$,

$$\frac{E_{\text{refl}}}{E_{\text{in}}} = \left(\frac{E_{\text{refl}}}{E_{\text{in}}} \right)_{\text{res}} \left(1 - \frac{r_1 r_2}{1 - r_1 r_2} 2ik\delta L \right). \quad (3.14)$$

The phase shift of the reflected light is larger by a factor of $r_1 r_2 / (1 - r_1 r_2)$ from the phase shift due to a single round trip. Therefore, when applied to a Michelson interferometer, the Fabry-Perot cavities greatly increase the effective arm length. The full gain in phase sensitivity happens as long as δL varies on a time scale longer than the effective light storage time. Above the cavity pole frequency of $(c/4\pi L)(1 - r_1 r_2)$, phase sensitivity falls as f^{-1} . For LIGO 4 km instruments, the reduction in sensitivity begins around 90 Hz.

The LIGO detectors make use of another resonant cavity to increase the effective laser power going to the beam splitter. High laser power is desirable in order to reduce photon counting statistical noise which is the limiting noise factor at high frequency. In this case a partially transmitting power recycling mirror is placed before the Michelson interferometer. Since the interferometer operates on a dark fringe, it acts as a second mirror for the resonant cavity reflecting incident light back toward the input port. Power builds up for this cavity in the same way as for a standard two mirror Fabry-Perot cavity. By making use of power recycling, the incident laser power on the beam splitter is increased by a factor of ~ 100 .

3.3.2 Length sensing and control

The main drawback of resonant cavities is that they must be kept on resonance. The tolerance for cavity length changes can be very small, around $\sim 10^{-10}$ m for LIGO. This requires active control of mirror position which is done by way of small magnets bound to the optics. Variations on Pound-Drever-Hall (PDH) reflection locking are used to construct error signals used to drive these magnets and keep the instrument at

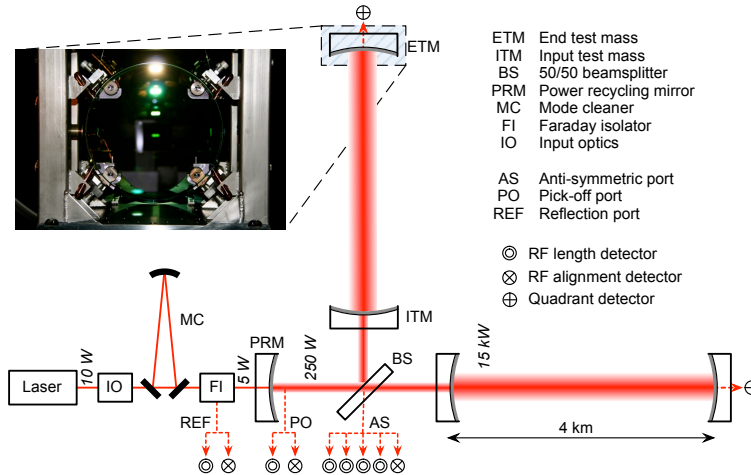


Figure 3-3: Optical layout of a LIGO interferometer. This figure is reproduced from [69].

the ideal operating point. First the 1064 nm ($f \simeq 280$ THz) laser light is modulated at radio frequencies $f_m \sim 25$ MHz by an electro-optical modulator. This has the effect to first order of introducing two new small sideband fields at frequencies $f \pm f_m$. If these sideband fields are not resonant in an optical cavity, they will be promptly reflected. The reflected sidebands will mix with the reflected carrier light forming optical beats at the modulation frequency. The amplitude of this signal is easily read by demodulating at the original modulation radio frequency and is proportional to the phase difference between the reflected carrier and sideband signals. Since this phase difference is itself proportional to small deviations from resonance, the demodulated signal can be used in a feedback loop to keep the cavity locked.

In LIGO two sets of sidebands are introduced. The first set from modulation at 62.5 MHz are not resonant in the power recycling cavity (PRC) and are directly reflected at the recycling mirror (RM). The second set from modulation at 25 MHz are resonant in the PRC and enter the Michelson where they can be used to control the individual arm cavity lengths. No sidebands are resonant in the arm cavities themselves. The final degree of freedom is the differential length of the Michelson interferometer. Since the interferometer operates on a dark fringe (minimum), the amplitude of light at the antisymmetric (AS) port varies quadratically with differential

arm length and thus does not make a good control signal. Instead, a 30 cm Schnupp asymmetry [70] is introduced into the Michelson arm lengths. The asymmetry is an integral number of wavelengths of the carrier light to maintain the dark fringe condition, but it is chosen so that some of the sideband power is channeled to the AS port. As in the PDH scheme, this gives a linear signal appropriate to measure and control the differential degree of freedom.

3.3.3 Noise sources

Three primary noise sources limit the design sensitivity of the initial LIGO detectors. At low frequency, seismic noise couples to random motion of the mirrors. At high frequencies, photon counting statistics limit the measurement accuracy as shot noise. At middle frequencies where the detector is most sensitive, thermal noise which causes random fluctuations in the mirror surfaces dominate. Other important noise sources which must be minimized are laser frequency and amplitude noise as the laser is not perfectly stable, various environmental noise such as acoustic or magnetic disturbances which can couple to mirror motion, and electronics noise. Mirror shape and alignment issues can also reduce sensitivity and must be carefully controlled.

The LIGO mirrors are isolated from ground vibration by four-layer passive seismic isolation stacks which provide f^{-8} suppression of seismic noise above ~ 10 Hz. In addition, an active seismic pre-isolator was deployed at the Livingston site due to the high ground motion there. The mirrors themselves are hung by steel wire which provides another f^{-2} level of isolation above the pendulum frequency of ~ 0.75 Hz. At low frequencies, the amount of ground motion that couples to mirror motion increases dramatically, creating an effective seismic noise wall at ~ 45 Hz.

Thermal noise enters the instrument through excitations of the suspension modes and fluctuations on the mirror surfaces. It can be diminished by choosing materials with very low mechanical loss and high Q , thus concentrating the noise at well defined frequencies which ideally can be placed outside of the detector band. The effects of thermal fluctuations at the mirror surface can be averaged out by using a large beam size. The Japanese LCGT project [68] aims to beat thermal noise by cooling the

optics to cryogenic temperatures.

At high frequency, shot noise limits the detector sensitivities. Even if the mirrors were perfectly undisturbed with zero displacement noise, the random nature of individual photons arriving at the photodetector means that rapid changes in intensity become unresolvable. As for any Poisson process, shot noise has a flat power spectrum which means there are no correlations between one measurement and the next. However the response of an interferometer with Fabry-Perot cavity arms decreases as f^{-1} above frequencies near the inverse light storage time. This must be multiplied by the shot noise spectrum to get the frequency dependent sensitivity to differential motion. The obvious way to reduce shot noise is to increase laser power. High power can lead to increased radiation pressure noise, another quantum noise source which will be important for Advanced LIGO. It also means more heat deposited on the optics which requires better thermal compensation to avoid distortion.

3.3.4 Advanced LIGO

The Advanced LIGO sensitivity goal calls for a factor of ~ 10 increase in sensitivity over initial LIGO across all frequencies. While the advanced interferometers will use the same site and vacuum enclosure as the current instruments, a large number of upgrades will be necessary to reach the new target. Some changes have already been implemented for Enhanced LIGO [71] and the current S6 run. These include an increase in laser power from 10 to 35 W and implementation of a DC readout scheme for the differential error signal.

Advanced LIGO will upgrade the laser to 200 W and add a signal recycling mirror to the antisymmetric port, creating a resonant cavity for gravitational wave induced sidebands near a chosen frequency. This allows tuning of the instrument to reduce shot noise at frequencies of interest. The test masses will be significantly larger and heavier to allow for bigger beam size and a reduction of radiation pressure noise. The use of fused silica rather than steel suspensions will greatly reduce suspension thermal noise, and seismic isolation will be improved by more complicated passive and active isolation systems. The factor of ten increase in strain sensitivity translates




Figure 3-4: Strain noise sensitivities for the three LIGO interferometers during the S5 data run 2005-2007. H1 (red) and L1 (green) have a 4 km baseline and are at or near their initial LIGO design sensitivity (dashed). H2 (blue) has a 2 km baseline reducing sensitivity by a factor of two at low frequency. The instruments reach a peak strain sensitivity around ~ 150 Hz, but are sensitive to a very wide bandwidth less some narrow spectral lines due largely to 60 Hz harmonics and suspension resonances. At the most sensitive frequencies, the 4 km instruments have reach a root-mean-square (rms) strain noise of 3×10^{-22} when restricted to a 100 Hz band. This figure is reproduced from [69].

into a factor of 1,000 in effective search volume, and regular detection of gravitational waves with the advanced detectors is expected.

Chapter 4

Gravitational wave bursts

Short bursts of gravitational radiation are expected from violent astrophysical phenomena such as the core collapse of massive stars and the compact binary merger scenario discussed in section 2.3. A search for the signatures of gravitational-wave bursts in detector data must maintain sensitivity to a large variety signals representing a range in progenitor initial conditions, as well as signals from sources with dynamics that may be poorly modeled or completely un-modeled. Thus, the generic signals targeted by burst analyses are only loosely constrained by their properties in time, frequency, shape, and by detector noise and network antenna response.

4.1 Burst parametrization

Burst signals typically have some natural frequency due to orbital, rotational, or acoustic timescales, and so a parametrization in time-frequency is a useful characterization. The burst signal, as detected by a single instrument, has representations in time and frequency which are related by the Fourier transform,

$$h(t) = \mathcal{F}^{-1}[\tilde{h}(f)] = \int_{-\infty}^{+\infty} \tilde{h}(f) e^{+2\pi i f t} df \quad (4.1)$$

$$\tilde{h}(f) = \mathcal{F}[h(t)] = \int_{-\infty}^{+\infty} h(t) e^{-2\pi i f t} dt. \quad (4.2)$$

We assume the burst signals have finite extent in time and frequency, with detector noise placing a practical limit on frequency extent.

The total signal energy is conventionally defined as the sum squared signal, and is related in time and frequency representations by Parseval's theorem,

$$\|h\|^2 = \int_{-\infty}^{+\infty} |h(t)|^2 dt = \int_{-\infty}^{+\infty} |\tilde{h}(f)|^2 df. \quad (4.3)$$

$|h(t)|^2$ is the density of signal energy in time, while $|\tilde{h}(f)|^2 = \tilde{h}^*(f)\tilde{h}(f)$ is the signal energy per unit frequency, or energy spectral density. In this formula, \tilde{h}^* is the complex conjugate of \tilde{h} . In order to have a representation of gravitational-wave amplitude independent of detector orientation and response, we also define a quantity h_{rss} as the summed contributions from h_+ and h_\times components in the gravitational-wave plane,

$$h_{\text{rss}}^2 = \|h_+\|^2 + \|h_\times\|^2. \quad (4.4)$$

The two amplitude values $\|h\|$ and h_{rss} are equal only for the case of a linearly polarized wave which is optimally oriented ($h_\times = 0$ and $F_+ = 1$).

Central time and frequency of the signal are given by the first moments of the squared signal,

$$t_c = \langle t|h(t)|^2 \rangle = \frac{1}{\|h\|^2} \int_{-\infty}^{+\infty} t|h(t)|^2 dt \quad (4.5)$$

$$f_c = \langle f|\tilde{h}(f)|^2 \rangle = \frac{2}{\|h\|^2} \int_0^{+\infty} f|\tilde{h}(f)|^2 df. \quad (4.6)$$

The factor of 2 for the frequency integral is to compensate for the integration over positive frequencies only as $\tilde{h}(-f) = \tilde{h}^*(f)$ for real signals $h(t)$. Signal duration and bandwidth are given by the second central moments,

$$\sigma_t = \frac{1}{\|h\|^2} \int_{-\infty}^{+\infty} (t - t_c)^2 |h(t)|^2 dt \quad (4.7)$$

$$\sigma_f = \frac{2}{\|h\|^2} \int_0^{+\infty} (f - f_c)^2 |\tilde{h}(f)|^2 df. \quad (4.8)$$

Finally, the quality factor, Q , of the signal is,

$$Q = \frac{f_c}{\sigma_f}. \quad (4.9)$$

The quality factor is a measure of the sharpness of the energy distribution about the central frequency. The nomenclature comes from the quality of low dispersion resonators. Q is useful as a parameter because it is a property of the shape of the waveform, and is invariant under arbitrary scaling in time.

The central time and frequency measurements and their spread are just the mean and variance of the respective values if the normalized squared signal is interpreted as a probability distribution. Like the Heisenberg uncertainty limit on position and momentum wave-functions in quantum mechanics, σ_t and σ_f satisfy a time-frequency volume bound,

$$\sigma_t \sigma_f \geq 1/4\pi. \quad (4.10)$$

Minimal uncertainty waveforms have $\sigma_t \sigma_f = 1/4\pi$ and have the form of a Gaussian windowed exponential wave packet. In this case, high Q wave packets will have longer duration σ_t . From 4.9 and 4.10 one can count $Q/2\pi$ oscillations within $\pm\sigma_t$ about the central time for a Gaussian wave packet regardless of frequency.

4.2 Detecting bursts in stationary noise

4.2.1 Power spectral density

Ideal detectors will have stationary noise $n(t)$ that can be characterized as a stochastic process with a given auto-correlation between points in time,

$$r_n(\tau) = \langle n(t)n(t - \tau) \rangle. \quad (4.11)$$

The Fourier transform of the auto-correlation gives the two-sided power spectral density, a frequency domain representation of the signal power per unit frequency,

$$G_n(f) = \int_{-\infty}^{+\infty} r_n(\tau) e^{-2\pi i f \tau} d\tau \quad (4.12)$$

The power spectral density is the expected energy spectral density per unit time. Signal energy is introduced in Equation 4.3, but because the noise is infinitely long, it has infinite signal energy unless restricted to a finite amount of time. An alternate definition of the power spectral density takes the expected energy spectral density per unit time for a finite interval of duration T ,

$$G_n(f) = \lim_{T \rightarrow \infty} \frac{1}{T} \langle |\tilde{n}(f)|^2 \rangle. \quad (4.13)$$

For real-valued signals, one often encounters the one-sided power spectral density,

$$S_n(f) = \begin{cases} 2G_n(f) & f \geq 0 \\ 0 & f < 0 \end{cases}, \quad (4.14)$$

which is nonzero for non-negative frequencies $f \geq 0$ and contains contributions at each non-negative frequency from both the equal magnitude left (positive frequency) and right (negative frequency) complex oscillators which make up the real signal.

4.2.2 Matched filter signal-to-noise ratio

The noise power spectrum tells us about the contribution from noise to mean-squared amplitude fluctuations over some frequency range. A gravitational wave signal must be above these random fluctuations in order to be detectable. In a general detection scheme, the signal-to-noise ratio refers to the amplitude of the measurement in the presence of a signal divided by the standard deviation on the measurement due to noise fluctuations, assuming the noise leads to Gaussian fluctuations. To characterize the strength of a signal relative to a particular noise spectrum, we use the signal-to-

noise ratio for a linear matched filter,

$$\rho^2 = \int_0^\infty \frac{4|\tilde{h}(f)|^2}{S_n(f)} df. \quad (4.15)$$

The signal-to-noise ratio ρ scales with any measure of signal amplitude including the root-sum-squared amplitude at the detector $\|h\|$ defined in Equation 4.3. If the noise spectrum is nearly constant over the frequency range of the signal, ρ takes the simple form,

$$\rho^2 \simeq \frac{2\|h\|^2}{S_n(f_c)}. \quad (4.16)$$

The matched filter maximizes the signal-to-noise ratio for the case of a linear real filter $b(t)$ in data which is the sum of some isolated signal $h(t)$ plus stationary noise: $x(t) = h(t) + n(t)$. The output $y(t)$ is the filter convolved with the detector data,

$$y(t) = b(t) * [h(t) + n(t)] \quad (4.17)$$

$$\tilde{y}(f) = \tilde{b}(f)[\tilde{h}(f) + \tilde{n}(f)]. \quad (4.18)$$

The asterisk $*$ denotes convolution in the time domain,

$$[f * g](t) = \int_{-\infty}^{+\infty} f(\tau)g(t - \tau) d\tau, \quad (4.19)$$

which is equivalent to multiplication in the frequency domain,

$$\mathcal{F}[f(t) * g(t)] = \mathcal{F}[f(t)]\mathcal{F}[g(t)]. \quad (4.20)$$

If we multiply the frequency domain filter $\tilde{b}(f)$ by $G_n^{1/2}G_n^{-1/2} = 1$, we can see that maximizing the signal-to-noise ratio for $b(t) * x(t)$ is the same as maximizing the signal-to-noise for the filter $\tilde{b}'(f) = \tilde{b}(f)G_n^{1/2}(f)$ applied to the whitened data $\tilde{x}' = G_n^{-1/2}(f)\tilde{x}(f)$. The signal-to-noise ratio is the magnitude of the filter output for filter b' applied to a pure whitened signal $G_n^{-1/2}(f)\tilde{h}(f)$ at the time of the signal

divided by the standard deviation of the filter output for whitened noise $G_n^{-1/2}(f)\tilde{n}(f)$,

$$\rho = \frac{\int_{-\infty}^{+\infty} b'(-t)h'(t) dt}{\sqrt{\int_{-\infty}^{+\infty} |\tilde{b}'(f)|^2 df}}. \quad (4.21)$$

The denominator is constant if we fix a normalization for b' . In that case, the numerator is maximized by choosing $b'(-t) \propto h'(t)$. In the vector analogy, this maximizes the dot product of vectors $b'(-t)$ and $h'(t)$ by choosing them to be in the same direction. In the frequency domain this gives $\tilde{b}'^*(f) \propto \tilde{h}'(f)$, which defines the matched filter,

$$\tilde{b}(f) \propto \frac{\tilde{h}^*(f)}{G_n(f)}. \quad (4.22)$$

We can directly compute the observed signal-to-noise ratio for the matched filter by expanding 4.21 in the case of a match,

$$\rho = \frac{\int_{-\infty}^{+\infty} \tilde{h}^*(f)\tilde{h}(f)G_n^{-1}(f) df}{\sqrt{\int_{-\infty}^{+\infty} \tilde{h}^*(f)\tilde{h}(f)G_n^{-1}(f) df}} = \left[\int_{-\infty}^{+\infty} \frac{\tilde{h}(f)^*\tilde{h}(f)}{G_n(f)} \right]^{1/2}, \quad (4.23)$$

which is the same as Equation 4.15.

By thresholding on the amplitude of the filtered output, one can decide if there is a signal present in the noise. For a standard normalized (such that the denominator in 4.21 is 1) matched filter with a real valued template, contributions to amplitude measurements from stationary noise alone are Gaussian distributed about zero with unity variance. The best estimate of the true signal-to-noise ratio of a signal present in noise is then just the amplitude measurement $y(t_{\text{signal}})$ itself. The statistical significance of a measurement ρ_0 is defined as the negative log probability of observing a higher apparent signal-to-noise from noise alone,

$$Z = -\ln \left[\frac{1}{\sqrt{2\pi}} \int_{\rho_0}^{\infty} e^{-\rho^2/2} d\rho \right]. \quad (4.24)$$

In the case of a gravitational wave from a compact binary inspiral, the orbital dynamics and gravitational radiation may be known, but the phase offset of the

observed signal at the detector depends on the orientation of the source. Since the inspiral evolution assumes adiabatic orbital decay from a dominant quadrupole mode, we can make use of the stationary phase approximation to search over arbitrary phase offset [72, 73]. This requires a complex filter of the form,

$$b(-t) = A(t)e^{i\Phi(t)} = A(t) \cos \Phi(t) + iA(t) \sin \Phi(t), \quad (4.25)$$

where the known amplitude envelope $A(t)$ is assumed to vary slowly relative to the known frequency $d\Phi(t)/dt$. The real and imaginary projections represent two standard matched filters for the waveform with $\pi/2$ phase offset. The magnitude of the complex projection is the same as the magnitude of a standard matched filter using the correct phase offset, and the phase offset can be read off from the phase of the complex projection. The expected signal energy from the projected noise, however, is twice the signal energy of the real-valued case as the equal and uncorrelated contributions from real and imaginary components must be added. Therefore the ratio of filtered signal amplitude to root-mean-square noise fluctuation is reduced by a factor of $1/\sqrt{2}$ due to the use of a complex filter. However, it may be misleading to characterize the signal with a signal-to-noise ratio reduced by the same amount because the noise fluctuations are no longer Gaussian distributed. Rather the squared noise fluctuations arise from the sum squared of two Gaussian random variables and are thus exponentially distributed. Instead we can again use statistical significance for a meaningful comparison with other measurements,

$$Z = -\ln \int_E^\infty \chi_2^2(E) dE = -\ln \int_E^\infty \frac{1}{2} e^{-E/2}, \quad (4.26)$$

where E represents the observed signal energy $|y(t)|^2$. The expected observed energy for a signal present in noise is related to the intrinsic matched filter signal-to-noise,

$$\langle E \rangle = \langle |y(t)|^2 \rangle = \rho^2 + 2, \quad (4.27)$$

where ρ^2 is the contribution from the signal projection and 2 comes from the energy

contribution from both phases of noise. One can then threshold on detected energy E to decide if a signal is present.

The statistics regarding matched filters presented so far only deal with a single template at a single point in time. For gravitational-wave searches, we are interested in searching over a range of times, and most likely a range of many other parameters as well. This requires the use of a bank of matched filter templates which covers the space of signals we want to detect. For continuously varying parameters (time, frequency, mass, mass ratio, etc.), templates are generally chosen to sufficiently sample the space in order to maintain a maximum fractional energy loss due to mismatch for any particular signal. This ensures near optimal sensitivity to signals at a fixed energy threshold. One of the advantages of the complex filter used in the stationary phase approach is that it makes template spacing in time much easier. Computational cost places a practical upper limit on the number of templates used. In addition, the naive false alarm probability from noise fluctuations scales directly with the number of templates. The increase in false alarms from a dense template bank ignores the large overlap between nearby templates, and in practice some sort of ad-hoc down-selection is used to isolate loud matches.

4.2.3 Time-frequency decomposition

The space of signals targeted by a burst search is generally too large and poorly constrained to practically cover with a bank of matched filter templates. Instead of searching coherently for a single waveform, it is possible to integrate signal energy over a number of (generally orthogonal) templates in order to search for excess total signal power over what is expected from stationary noise. Any signal which contains appreciable signal energy within the chosen time-frequency bounds can be detected this way.

To isolate a signal in time-frequency, one option is to use the windowed Fourier

transform defined for a time series $x(t)$ by the integral,

$$X(\tau, \phi) = \int_{-\infty}^{+\infty} x(t)w(t - \tau)e^{-i2\pi\phi t} dt \quad (4.28)$$

$$= \int_{-\infty}^{+\infty} \tilde{x}(f + \phi)\tilde{w}^*(f)e^{+i2\pi f\tau} df. \quad (4.29)$$

The transform projects a time series onto a windowed complex exponential centered at time τ and at frequency ϕ . Represented as an integral in time, we recognize the operation as a Fourier transform of the windowed time series $x(t)w(t - \tau)$. Alternatively the transform can be represented as an integral over all frequencies using Parseval's relation. This representation is useful if the shape of the window function does not vary with time τ but does vary with ϕ , as is the case for a scale-invariant wavelet transform. It is important to note the difference in time-frequency variables for the data and filter (t, f) versus those for the transform $X(\tau, \phi)$.

The degree of time-frequency isolation provided by the projection depends on the narrowness of the windows $w(t)$ and $\tilde{w}^*(f)$. The operation captures the signal energy which overlaps a tile in time-frequency whose geometry is determined by the window extent. In 1944, Gabor investigated the use of Gaussian window functions for $w(t)$ for the windowed Fourier transform. Gabor's filters have minimal time-frequency volume $\sigma_t\sigma_f$ (as in 4.10) and thus best isolate signal power in time-frequency. The Q transform [74] uses approximately Gaussian windowed exponential filters of varying widths to sample the signal power in a space parametrized by τ , ϕ , and $Q = \phi/\sigma_f$. At fixed Q , the transform is scale invariant can be described as projections onto a single function stretched and compressed in time by various amounts.

The windowed Fourier transform and derivatives are highly redundant as they are by definition evaluated continuously over a bank of templates parametrized by at least τ and ϕ . When a signal is thought to contain power over a number of time-frequency elements, it is desirable to have orthogonal basis elements so that the signal content can be captured in a small set of coefficients. The Q transform attempts to do this by a hierarchical method of down selection in which significantly overlapping tiles are discarded. However because there is no clear delineation between Gaussian windowed

complex exponential templates, some information can be lost in the down selection process. The Excess Power pipeline [75] and TFClusters pipeline [76] also make use of the windowed Fourier transform to capture signal energy in time-frequency. Both pipelines directly achieve an orthogonal basis suitable for clustering by using simple disjoint rectangular windows in either time or frequency. The disadvantage is that for a rectangular window in time, frequency isolation (which takes the form of a sinc function) is poor and vice versa.

The wavelet transform [77, 78] provides an alternate time-frequency decomposition of a signal onto an orthogonal wavelet basis. The continuous wavelet transform for a time series $x(t)$ is defined by the integral,

$$W(\tau, s) = \int_{-\infty}^{+\infty} x(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - \tau}{s} \right) dt \quad (\tau, s) \in \mathbb{R} \otimes \mathbb{R}^+ \quad (4.30)$$

$$= \int_{-\infty}^{+\infty} \tilde{x}(f) \sqrt{s} \tilde{\psi}^*(sf) e^{+i2\pi f\tau} df. \quad (4.31)$$

The time series is projected onto scaled versions of a single mother wavelet $\psi(t)$ which is assumed to be time localized and zero-average. The scale factor s determines the amount of stretching or shrinking of the wavelet which otherwise does not change in shape. As is the case for the Q transform at constant Q , the windowed Fourier transform with window size that scales inversely with frequency ϕ can be considered a special case of the wavelet transform, with the main difference being that the phase of the oscillatory component is absolute in time for the Fourier transform. A number of real-valued wavelets have been discovered which have the unique property that when evaluated on a dyadic grid,

$$\psi_{(j,n)}(t) = \frac{1}{\sqrt{2^j}} \psi \left(\frac{t - 2^j n}{2^j} \right) \quad (j, n) \in \mathbb{Z}^2, \quad (4.32)$$

they form an orthonormal basis for a real valued signal. The dyadic (discrete) wavelet transform provides a multiresolution decomposition of the signal with small spacing $\delta\tau = 2^j n$ between elements in time and large spacing $\delta f = 2f$ between elements in frequency at small scale j (large f), and correspondingly poor time resolution

and better frequency resolution at large scale. Alternatively one may continue to decompose the coefficients at a given scale in a wavelet packet decomposition which results in a rectangular array with constant time and frequency resolution at all frequencies similar to the fixed-windowed Fourier transform. This decomposition is used by the Coherent Waveburst pipeline [79]. While the wavelets have larger time-frequency volume than Gaussian wave-packets, many still provide very good time-frequency localization while retaining all the benefits of an orthogonal basis.

Each basis element can be characterized by the parameters defined in section 4.1. The linear transformations for orthogonal templates are energy-preserving so that the total signal energy can be reconstructed from adding up the signal energy (squared coefficients) in the transformed space. Quantities such as central time, frequency and their second moments can also be reconstructed easily. Alternatively one may simply keep the values of the few top coefficients which give a good approximation to the signal in order to retain information about more complicated time-frequency structure.

Time-frequency decomposition is an effective way to isolate a localized signal from noise. Typically data is whitened and normalized prior to any transformation so that the noise power spectrum is flat and individual noise samples are uncorrelated. A projection of noise onto a normalized real template will then be a Gaussian random variable. Such is the case for the dyadic wavelet decomposition. The squared magnitude of a windowed Fourier coefficient is exponentially distributed as the sum of two squared Gaussian random variables corresponding to the real and imaginary phase components. Total noise energy from a cluster of independent coefficients then follows a χ^2 distribution where the number of degrees of freedom is equal to the number of coefficients ($\times 2$ for complex filters). In general the noise energy from a time-frequency area A will be χ^2 distributed with $2A$ degrees of freedom [75]. This result follows most quickly from considering the number of independent samples necessary to represent a signal of duration Δt and bandwidth Δf according to the Nyquist criterion. One can then keep only coefficients or clusters of coefficients whose signal energy is statistically significant.

4.3 Multiple detector observation

One of the generally accepted requirements for a gravitational-wave detection is that the event be observed simultaneously in multiple detectors. For this reason, the LIGO detectors consist of two sites separated by a large distance (3,000 km) in order to minimize environmental coupling. Coincident observation greatly enhances our confidence in an event by dramatically reducing the false alarm probability. It also serves as the principle means for evaluating the background of a search through artificial relative time-shifts applied to data from two detectors.

4.3.1 Coincidence in observed parameters

The most basic criterion for coincidence is that the events from multiple detectors happen at the same time. For two gravitational-wave detectors, the coincidence time window is determined by the light travel time between detectors (~ 10 ms for the two LIGO detectors) plus any uncertainty from noise in arrival time estimation for the search algorithm. Misaligned detectors sensitive to different gravitational-wave polarizations may also contribute to observed timing differences. The coincidence window Δt is generally chosen to be large enough to cover all these effects. If the single detector accidental event arrival times can be modeled as a Poisson process with rate λ and we use a coincidence window $\pm \Delta t$ that is small compared to the time between events, the coincident rate between two uncorrelated detectors is,

$$\lambda_{12} = 2\lambda_1\lambda_2\Delta t. \tag{4.33}$$

The rate decreases rapidly as we add additional detectors so long as $\lambda\Delta t \ll 1$.

In addition to time, other parameters such as frequency and amplitude estimated by a search algorithm can be used as coincidence criteria to further reject coincident noise events. When the inconsistency parameters are approximately independent, the naive likelihood ratio is particularly effective at distinguishing signal from background [80]. In this Bayesian approach, distributions for n inconsistency parameters

$\mathbf{x} = (\delta\mathbf{t}_c/\sigma_t, \delta\mathbf{f}_c/\sigma_f, \text{etc.})$ are formed for a large number of simulated signals as well as sample background. The likelihood ratio for a coincident event is the product of these distributions evaluated at \mathbf{x} for simulations (hypothesis H_1) divided by the distributions for pure noise (hypothesis H_0),

$$\Lambda(\mathbf{x}) = \frac{P(\mathbf{x}|H_1)}{P(\mathbf{x}|H_0)} \simeq \prod_{i=1}^n \frac{P(x_i|H_1)}{P(x_i|H_0)} \quad (4.34)$$

4.3.2 Coherent observation

Coherent techniques make use of the known response of a network of detectors to fit a gravitational waveform to data which can then be checked for consistency with the hypothesis that the data consists of the best fit waveform plus Gaussian noise. A single measurement of gravitational-wave strain (h_+, h_\times) across D detectors can be represented as a matrix equation,

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_D \end{bmatrix} = \begin{bmatrix} F_1^+ & F_1^\times \\ F_2^+ & F_2^\times \\ \vdots & \vdots \\ F_D^+ & F_D^\times \end{bmatrix} \begin{bmatrix} h_+ \\ h_\times \end{bmatrix} + \begin{bmatrix} n_1 \\ n_2 \\ \vdots \\ n_D \end{bmatrix} \quad (4.35)$$

or in compact notation,

$$\mathbf{x} = \mathbf{F}\mathbf{h} + \mathbf{n}. \quad (4.36)$$

Here \mathbf{x} are the measurements from detector data appropriately time-shifted for the assumed source direction of incident gravitational wave \mathbf{h} . $\mathbf{F}(\hat{\Omega})$ represents the directional antenna factors of each detector in the gravitational-wave frame, and \mathbf{n} are the contributions to \mathbf{x} from noise. While the measurements \mathbf{x} could in principle be the measured strain for each detector, we are generally interested in situations where the noise contributions \mathbf{n} are Gaussian random variables with standard deviations σ_i . Real detector data is colored with potentially narrow features. To maintain a Gaussian distribution from noise, the data point must represent some spectral feature, such as a single narrow-band wavelet coefficient, of the data and gravitational

wave. We cannot simply whiten the gravitational wave \mathbf{h} along with the data for a network analysis as 4.35 requires it to be the same for all detectors

The likelihood ratio is defined in Equation 4.34 as the probability density of observing the data \mathbf{x} across all detectors assuming the signal \mathbf{h} is present (H_1) divided by the probability density of observing the data assuming \mathbf{h} is not present (H_0). Assuming the detector noise is uncorrelated, this gives

$$P(\mathbf{x}|H_1) = \prod_{i=1}^D \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(x_i - \mathbf{F}_i\mathbf{h})^2}{2\sigma_i^2}\right), \quad (4.37)$$

$$P(\mathbf{x}|H_0) = \prod_{i=1}^D \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{x_i^2}{2\sigma_i^2}\right). \quad (4.38)$$

For N independent measurements (such as a cluster of wavelet coefficients), we can write the log likelihood ratio,

$$\mathcal{L} = \ln \prod_{k=1}^N \frac{P(\mathbf{x}[k]|H_1)}{P(\mathbf{x}[k]|H_0)} = \sum_{k=1}^N \sum_{i=1}^D \left(\frac{x_i[k]\mathbf{F}_i\mathbf{h}}{\sigma_i[k]^2} - \frac{(\mathbf{F}_i\mathbf{h})^2}{2\sigma_i[k]^2} \right) \quad (4.39)$$

The best-fit waveform (h_+, h_\times) and sky location $\hat{\Omega}$ is that which maximizes \mathcal{L} [81]. This can be solved analytically for a fixed sky location, but generally requires sampling over a grid of locations in order to find the global maximum. At typical frequencies, thousands of sky locations must be sampled for each event which causes coherent methods have high computational cost. For this reason, a coherent stage is usually triggered by the computationally cheaper search for coincident excess power.

One problem with fitting to the maximum value of \mathcal{L} is that it ignores any prior information we may have about the gravitational wave \mathbf{h} . This becomes most apparent when dealing with coincident noise transients from a small number of detectors. In that case, it is not difficult to find arbitrary (h_+, h_\times) which manage to cancel out a large portion of the noise. Various types of regularization [82, 83] and other constraints [84, 85] have been used to steer the search toward physically plausible waveforms. A fully Bayesian approach [86] solves the problem by marginalizing over generic but physically meaningful signal and noise transient priors.

Coherent analyses provide not only excellent network sensitivity to gravitational-wave signals, but also a reconstructed waveform and sky location. The reconstructed waveform can be subtracted from the data in each detector leaving residuals in the case of accidental noise coincidences. This can provide a very powerful test for removing transient background [87]. The probability sky maps based on the likelihood ratio are also useful for coincident searches with electromagnetic (EM) transients as well as for the prospects of a gravitationally-triggered EM follow up.

4.3.3 Time-shift method of background estimation

The time-shift method of background estimation is critical for the ability of a network of detectors to be able to detect gravitational waves with confidence. The procedure relies on the assumption that if an entire analysis is run with data from the various detectors artificially time-shifted with respect to one another, the resulting events will be drawn from the same underlying background distribution as for un-shifted data. Real gravitational waves, however, should be either invisible or negligibly weak compared to background in the time-shifted analysis. These assumptions require that true signal rate is low which is valid for current gravitational-wave searches and is easily verified at the end of the analysis. The time shifts cannot be longer than the timescale of non-stationarity at each instrument, and to maintain independence they must be spaced at least as far apart as a typical coincidence window. For burst searches it is possible to achieve thousands of independent realizations of the background distribution this way, and thus estimate the underlying distribution to a high degree of accuracy.

Chapter 5

Non-Gaussian noise transients

The matched filter, excess power, and coherent maximum likelihood signal analysis techniques presented in the previous chapter are all built on the assumption that the detector noise can be reduced to stationary white Gaussian noise. This means that the noise can be completely characterized by its power spectral density, and that any changes in the underlying power spectrum occur slowly enough so that the spectrum can still be estimated from the data itself to a high degree of accuracy. The backgrounds present in analyses of real data for gravitational-wave transients, however, are typically dominated by noise transients non-Gaussian in character. These short instrumental disturbances are caused by a variety of sources including coupling of environmental disturbances to mirror motion saturation of various control signals or other instrumental artifacts.

The presence of noise transients does not necessarily make the Gaussian noise assumption a poor characterization. A real gravitational-wave signal will most likely rest on top of nearly stationary noise, so the maximum likelihood approach can, for example, still provide valid sky location and waveform reconstruction. However the amplitudes of transient disturbances in the gravitational-wave signal from noise fluctuations can be quite large which corresponds to a vanishingly small probability of arising from Gaussian noise. Any set of noise events ranked by statistical significance under the Gaussian assumption will be dominated by such non-Gaussian transients at sufficiently large amplitudes.

Coincidence requirements across multiple detectors bring the false alarm probability of a transient search down to manageable levels of $O(1)$ per analysis livetime, and coherent waveform consistency checks are particularly good at rejecting events which stand out well above the Gaussian noise. Still, accidental coincidences of noise transients which manage to pass coherence tests are the limiting factor for the sensitivity of current transient searches, and their presence makes it very difficult to assign a real event a background rate sufficiently low enough to claim a detection. This and the following chapters deal with characterizing the noise transient populations in the instruments, and then using information about transients from the hundreds of auxiliary environmental, diagnostic, and control channels being monitored simultaneously with the gravitational-wave channel in order to identify disturbances in the detectors and thus remove potential transient background without using data from the gravitational-wave channel itself.

5.1 Transient population in LIGO data

As we are interested in identifying the noise transients which might be confused with gravitational-wave bursts, the methods to find both types of events are very similar. The time-frequency excess power techniques discussed in subsection 4.2.3 are well suited to the problem as they maintain sensitivity over a very large class of waveforms with time-frequency structure. The parametrization provided by time-frequency methods are also very useful for characterizing the transient population.

The most telling (and problematic) property of noise transients is that their distribution in amplitude falls off much more slowly than what is expected from Gaussian fluctuations. This is shown in Figure 5-1 which shows the rate of single interferometer (not coincident) transients in LIGO S5 gravitational-wave data above a certain threshold in signal-to-noise ratio. The transients occur during science data after removing times representing known problems with the instruments (control channel saturations, calibration line dropouts, large power main glitches, drops in circulating light in the arm cavities). They are identified by the Q Pipeline [74] which matches each




Figure 5-1: Rate of single-interferometer noise transients during the entire S5 run as a function of signal-to-noise ratio threshold as detected by the Q Pipeline [74]. The gravitational-wave data from the LIGO instruments during S5 contains populations of noise transients well in excess of those expected from stationary white noise (dashed curve) which for the Q Pipeline are exponentially distributed in SNR^2 . The search covers the LIGO sensitive band of 64-2048 Hz with $4 < Q < 64$. Since these signals do not survive coincidence requirements across multiple sites beyond what is expected from random coincidences, we know they are overwhelmingly not of astrophysical origin.

transient to a best fit sine-Gaussian, parametrizing it in amplitude, time, frequency, and Q , and they cover the LIGO sensitive band of 64-2048 Hz with $4 < Q < 64$.

Figure 5-2 shows the population of noise transients in terms of their distribution in frequency and Q . The number of transients is represented by a fractional excess above what is expected in ideal stationary noise. The number is slightly different for different templates as longer low frequency or high Q templates are spaced further apart and have correspondingly lower false alarm rates in Gaussian noise. The expected number of events with signal-to-noise ratio greater than 8 is vanishingly small for Gaussian noise, so the relative excess for a real population of non-Gaussian noise transients can be quite large. All instruments show a considerable excess of low-frequency low- Q transients. These are moderately loud signals (Figure 5-1) around 60-150 Hz which last for just a few cycles. Not surprisingly they also dominate the population of accidental coincidences across multiple instruments.

The average rates for the same population of transients (restricted to the LIGO sensitive band 64-2048 Hz) is shown in Figure 5-3 as a function of day in S5. The dramatic variation in rates arises from environmental changes such as increased seismic activity due to storms or earthquakes as well as instrumental changes. The variation in daily rates highlights the fact that not only is the baseline stochastic noise non-Gaussian and non-stationary, but the transient population is also highly non-stationary. For this reason, care must be taken when interpreting aggregate transient statistics over a long period of time.

5.2 Examples of transients

The most straightforward type of transient in the gravitational-wave data is caused by a direct coupling of environmental noise into mirror motion. In LIGO's second science run (S2), a single outstanding event remained at the end of the burst analysis [88]. During follow-up investigations, the source was identified as acoustic coupling from an overhead airplane at the Hanford site. These can couple directly to the mirrors (which do not sit in vacuum) or through the induced ground motion. This relatively

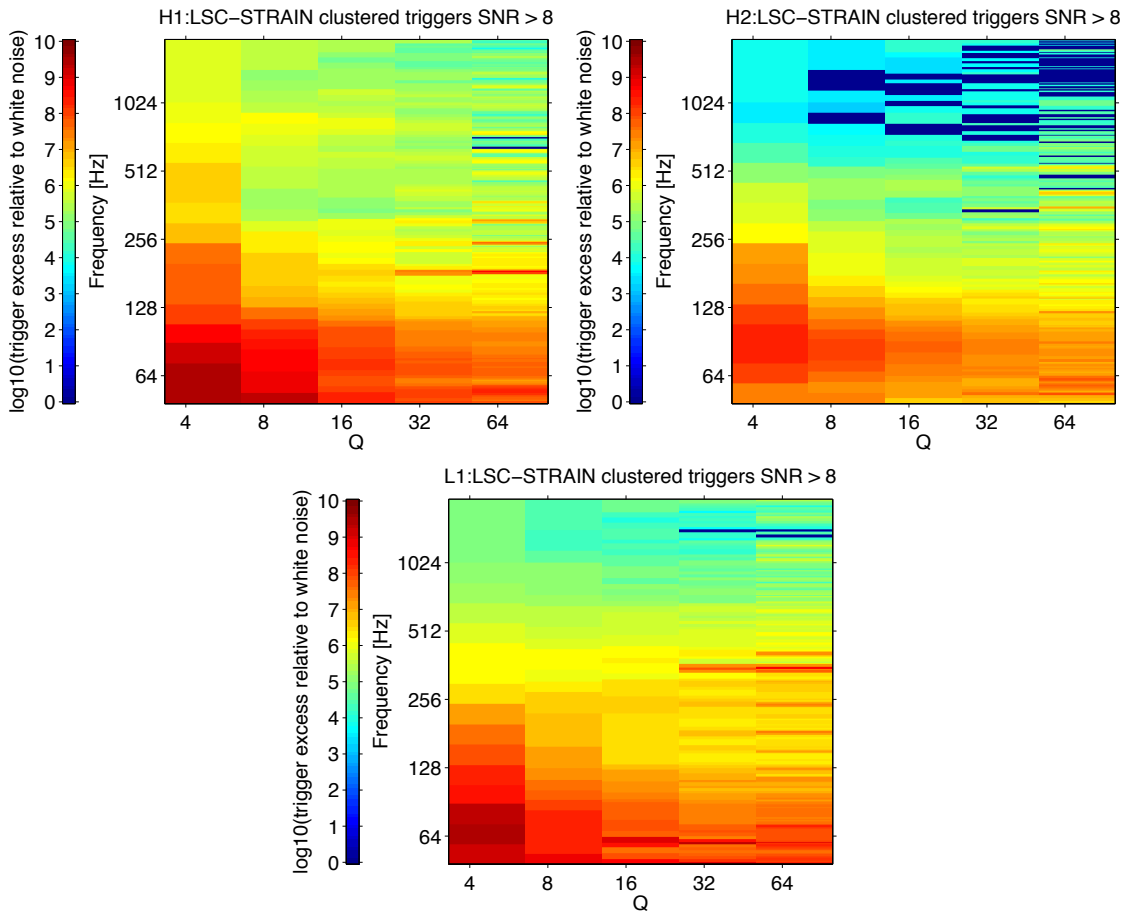


Figure 5-2: Population of noise transients in the LIGO instruments for S5 as identified by the Q Pipeline [74]. The counts are represented in terms of a fractional excess of counts over what is empirically expected at the same SNR threshold from Gaussian noise. All instruments exhibit an excess of low frequency, low Q transients (below ~ 150 Hz), and there is a population of higher Q power transients from nonstationarities in the 60 Hz line as well as mirror suspension resonances. The short-cycle low frequency transients overlap with the most sensitive frequency band of the LIGO instruments, and their accidental coincidences dominate background rates for transient searches.




Figure 5-3: Rate of transients detected in the gravitational-wave data for S5 with $\text{SNR} > 8$. The transients are detected by the Q Pipeline [74] and averaged over each day. Therefore any intra-day variation in rates (which can be very large) are not represented here. Rates can vary by orders of magnitude and are affected by changes in the environment (e.g. storms) as well as changes in the instrument (for example due to commissioning). The Q Pipeline also makes use of a trigger down-selection process which artificially limits the rate to 1.0 Hz. The transients shown come from the same population as in 5-2 except here they are restricted to 64-2048 Hz.

loud and semi-coherent H1H2 signal was in random coincidence with a weak transient at Livingston. Since S2, additional acoustic isolation of the optics has helped mitigate this kind of event, though airplanes are still carefully monitored for their effect on data quality.

The environment can also couple to the mirrors through ground motion transferred through the suspension system, as well as ambient magnetic fields which create forces on the mirror through small magnets attached to the mirrors for active control. Through S5, environmental coupling was particularly problematic at the Hanford site as it caused coincident signals at H1 and H2 with similar spectral properties. Coherent S5 burst analyses in particular relied strongly on waveform consistency between the two co-aligned Hanford instruments in order to reject background. In S5, the most prominent source of these were magnetically induced events caused by glitches in the power mains. This caused semi-coherent low-frequency signals in H1 and H2 as well as large disturbances in ambient magnetic fields measured all around the site. Figure 5-4 shows an example of such an event.

Other sources of instrumental artifacts have been traced back to (among other things) scattered light between the two Hanford interferometers which share a com-

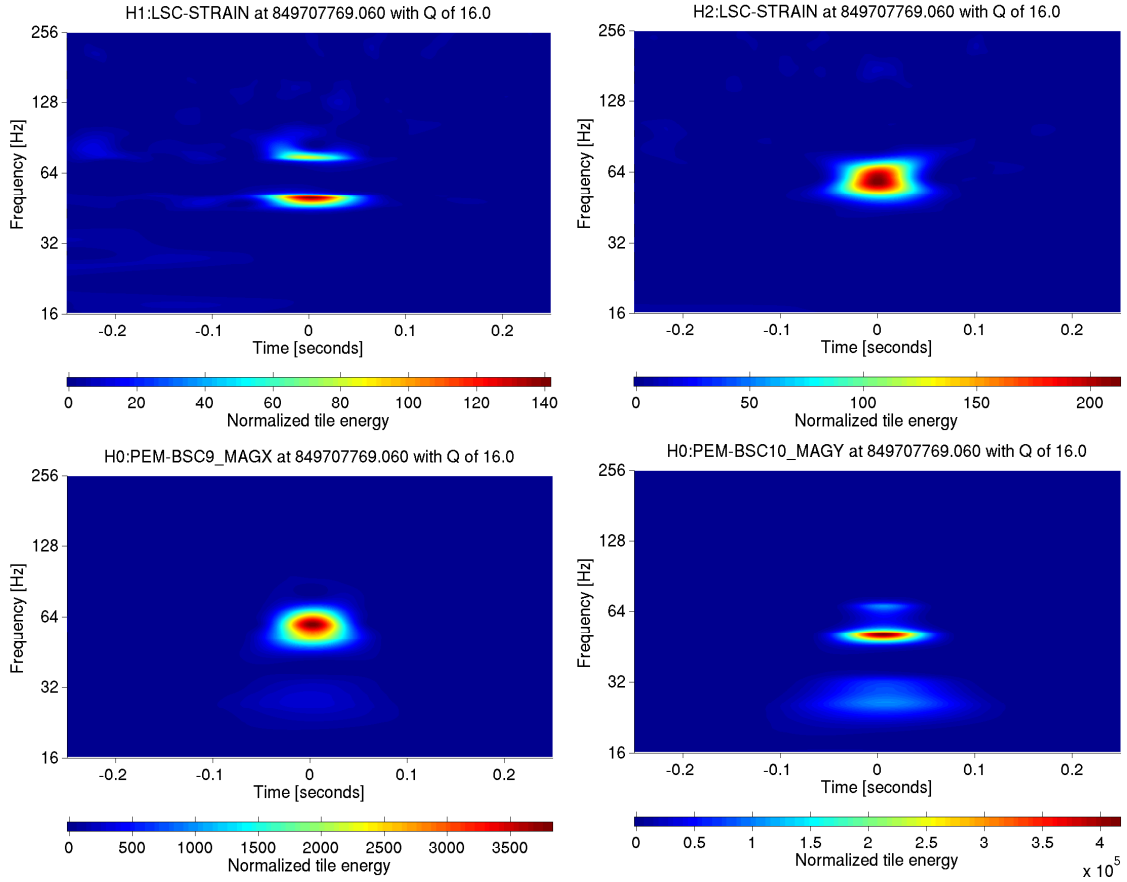


Figure 5-4: Semi-coherent signals in the gravitational-wave data at H1 and H2 during the time of a site-wide magnetic disturbance from a glitch in the power mains. The disturbance is easily picked up in environmental monitors measuring ambient magnetic fields making this particular source easy to remove from the analysis. The signals are centered about 60 Hz (with the narrow line generally suppressed by whitening filters). The magnetometers shown are from the end stations, and similar signals are seen at other locations.

mon vacuum tube, saturation of digital control signals, transient misalignment of the beams which cause a loss of circulating light in the resonant cavities, and glitches in the digital mirror excitation or data acquisition system which happen as a result of computer load or malfunction. Transients which can be clearly associated with a well understood environmental or instrumental disturbance are typically removed from analyses through automated procedures. A great many noise transients, however, are not traced back to a specific source. Statistically they may be associated with extreme conditions in the instrument such as those created by elevated ground noise. They may also be associated with specific instrumental disturbances if coincident with transients observed in auxiliary channels not sensitive to gravitational waves. In such cases their removal must generally be balanced with a non-trivial cost in analysis livetime.

5.3 **kleineWelle: wavelet-based identification**

KleineWelle [89] is a signal analysis pipeline which was developed to find and characterize non-Gaussian transients in an input timeseries. The method consists of the following steps,

- Decimate, high-pass, and whiten the data with a linear predictive error filter (LPEF)
- Decompose the data onto a dyadic wavelet basis and identify amplitude outliers
- Cluster wavelet coefficient outliers based on time-frequency proximity
- Rank clusters by the statistical significance of their excess total signal energy and characterize cluster properties

The procedure is conceptually similar to other methods which search for excess power in time-frequency introduced in subsection 4.2.3. The method has the advantages of speed and ease of tuning, which are largely provided by the properties of the dyadic wavelet transform and scale-invariant clustering. For this reason it is suitable for rapid processing of large amounts of data.

5.3.1 Wavelet decomposition

The wavelet transform 4.30 for timeseries $x(t)$ is defined by the integral,

$$W(\tau, s) = \int_{-\infty}^{+\infty} x(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - \tau}{s} \right) dt, \quad (5.1)$$

where the mother wavelet, ψ , is a time-localized function of zero average. The transform is simply a projection of the series onto stretched and compressed representations of a single function. The coefficients $W(\tau, s)$ are evaluated continuously over times, τ , and scales, s . Our ability to resolve in time and frequency is then determined by the properties ψ assumes at each scale. At large scale, ψ is highly dilated yielding improved frequency resolution at the expense of time resolution. At small scale, we achieve good time resolution with large uncertainty in frequency.

For the case of discrete data, a computationally efficient algorithm exists for calculating wavelet coefficients over scales and times that vary as powers of two: $s \in \{2^j \mid j \in \mathbb{Z}^+\}$, $\Delta t \propto s$ (Equation 4.32). This is the dyadic wavelet transform, which can be implemented for a limited family of wavelets using conjugate mirror filters. The filters consist of a high pass filter, \hat{H} , and low pass filter, \hat{L} , which can be applied in a cascade to obtain the wavelet coefficients. Beginning with the original time series, A_0 , of length N , two sequences of length $N/2$ are obtained by application of the high pass and low pass filters followed by down-sampling. The sequence of detail coefficients, D_j , and approximation coefficients, A_j , are defined at each level, j , of the decomposition by

$$\begin{aligned} D_j &= \hat{H}(A_{j-1}) \quad \text{and} \\ A_j &= \hat{L}(A_{j-1}). \end{aligned} \quad (5.2)$$

The detail coefficients for scale s , where $s = 2^j$, calculated in this manner are the same as the wavelet coefficients obtained from Equation 4.30. If N is a power of two, so that $N = 2^m$, the final approximation sequence will be A_m and will contain one point. The entire decomposition requires $O(N)$ computations.




Figure 5-5: Discrete wavelet decomposition tree using conjugate mirror filters. The original sequence A_0 is decomposed using high pass (left) and low pass (right) conjugate mirror filters and then down-sampled resulting in detail and approximation coefficients at each successive scale. The decomposition ends for a finite series when the approximation series is reduced to a single point. The detail coefficients along with the final approximation series up to any given scale can be used to reconstruct the original sequence A_0 .

The simplest dyadic wavelet is the Haar function [90]:

$$\psi^{\text{Haar}}(t) = \begin{cases} 1 & 0 \leq t < 1/2 \\ -1 & 1/2 \leq t < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (5.3)$$

The corresponding high pass and low pass filters are

$$\begin{aligned} \hat{H}^{\text{Haar}} &= [+1, -1]/\sqrt{2} \quad \text{and} \\ \hat{L}^{\text{Haar}} &= [+1, +1]/\sqrt{2}, \end{aligned} \quad (5.4)$$

from which we see that the detail coefficients are related to the differences of each pair of points in the parent series, while the approximation coefficients are related to the averages of each pair. It is easy to check that for the Haar wavelet, the continuous wavelet integral 5.1 and the detail coefficients of the discrete dyadic decomposition 5.2 give the same values. It is also easy to check that the dyadic sampling of the Haar wavelet provides an orthonormal basis. The Haar wavelet is often not the wavelet of choice for signal processing applications because of its poor frequency localization.

However it has the advantages of being computationally very fast and compact making it appropriate for current applications of the method.

5.3.2 Data conditioning

Instrumental noise is generally colored and may be populated with many narrow lines (Figure 3-4). *KleineWelle* makes use of linear predictive filtering [89] to whiten the data prior to the discrete wavelet decomposition. In linear prediction, the n^{th} sample of a sequence is modeled by a linear combination of the previous L samples,

$$\hat{x}[n] = \sum_{m=1}^L b[m]x[n-m]. \quad (5.5)$$

A common choice of coefficients $b[m]$ involves minimizing the expected squared prediction error,

$$\sigma_e^2 = \text{E}[|x[n] - \hat{x}[n]|^2]. \quad (5.6)$$

Solving the least squares minimization for a limited amount of training data involves estimating the auto-correlation coefficients and solving the Yule Walker matrix equation. Application of this method in the context of gravitational-wave data conditioning is described in detail as a component of the Q Pipeline [74]. By subtracting the prediction sequence $\hat{x}[n]$ from the original data $x[n]$, correlations are removed from the baseline stationary stochastic noise, reducing it to white noise. This process is called linear predictive error filtering (LPEF). Non-stationary transients which cannot be predicted remain, but are shaped by the frequency response of the filter.

The LPEF filter length and training length are constrained by the properties of the data and parameters of the analysis. Narrow spectral lines in the noise represent correlations over long time periods, and the filter length must be comparatively long in order to suppress them. The training length must be long relative to the filter length in order to allow reliable estimates of the auto-correlation of the data for the longest time delay. It also must be long enough so that transient non-stationarities of moderate strength contribute little to the auto-correlation estimates. A high pass prefilter must




Figure 5-6: Decomposition of a $\sigma = 1$ ms Gaussian injection with match filter signal-to-noise ratio of 11 onto simulated LIGO S2 noise. Shown at top is the signal+noise after high pass filtering at 50 Hz and LPEF whitening. Shown at bottom are the normalized Haar discrete wavelet decomposition coefficient amplitudes for the whitened data along with a timeseries of the original unfiltered signal.

be used to remove low frequency content which otherwise cause correlations on the timescale of the filter length. Finally, an upper limit on filter length and training length is set by computational considerations and the maximum timescale for which the assumption of stationarity in the baseline noise is expected to hold.

In the *kleineWelle* pipeline, data is read sequentially in constant strides of length N samples, or N/f_s seconds where f_s is the sampling frequency. The data is first decimated using a low-pass filter along with resampling such that the Nyquist frequency is above the highest requested frequency of the analysis. The decimated series is then high pass filtered using a 6th order Butterworth IIR high pass filter with cutoff fre-

quency at the lowest frequency of the analysis f_{low} . For many channels representing physical degrees of freedom in the instrument, this also removes the dominant seismic noise contribution. IIR filter state is preserved across contiguous blocks of data in order to avoid the initial transient response. In the case of a data discontinuity, a tune-able length of initial data is additionally read, filtered, and discarded to remove the filter transient. The first $L < N$ biased auto-correlation coefficients are calculated from the high passed data stream and used for training the L linear predictor filter coefficients. The predicted sequence of data is then removed from the initial high passed data leaving a whitened data stream. Both the auto-correlation coefficient estimates and FIR filter application can be done efficiently in the frequency domain. The linear predictor filter also keeps a history of the last N data points to avoid filter transients for the case of continuous data. In the case of a discontinuity, the filter transient will be at most L samples long and must be discarded as well.

5.3.3 Event generation

The data after high pass and linear predictor error filtering is a series of white noise of length N . The stride is chosen such that N is a power of two, so that the entire dyadic wavelet decomposition results in $N - 1$ detail coefficients $W_{(n,j)}$ with $N/2^j$ coefficients at each scale where,

$$\begin{aligned} j &\in \{1, 2, \dots, \log_2 N\} \\ n &\in \{0, 1, \dots, N/2^j - 1\}. \end{aligned} \tag{5.7}$$

The final approximation coefficient is zero as the data is zero-mean so that the original sequence can be reconstructed from the $N - 1$ remaining detail coefficients. At each scale, the wavelet coefficients can be assumed to be zero-mean Gaussian distributed with some standard deviation σ_j . By setting a threshold on the absolute normalized amplitude $|W_{(j,n)}|/\sigma_j$ across all wavelet scales, we build a list of outliers which are called black pixels.

A general transient signal isolated in time-frequency will result in a cluster of




Figure 5-7: Definition of distance from a wavelet coefficient (in white) to other coefficients in the time-scale decomposition. Tiles with a distance of 1 (light pink) and 2 (dark pink) from the original tile in white are shown. A typical distance threshold used for clustering outliers is 2.

large amplitude coefficients in the index-scaling wavelet tiling. To capture the total signal energy of such a transient, black pixels nearby in time-frequency are clustered together to form a single event, which we call a trigger. Clustering is done based on a concept of index-scale distance,

$$d = \Delta j + \Delta n \quad (5.8)$$

where Δn is based on the time between adjacent coefficients at the higher scale (larger separation). This definition of distance is symmetric between two wavelet coefficients. Figure 5-7 illustrates the calculated distance between wavelet coefficients in the dyadic tiling. A cluster is made up of the largest set of wavelet coefficients such that the minimum distance from any single wavelet coefficient to all others does not exceed some distance threshold d_{\max} .

The discrete dyadic wavelet decomposition is both orthogonal and energy preserving, so the cluster of N independent wavelet coefficients represents a signal with total

signal energy,

$$E_c = \sum_{i=1}^N |A_i|^2 \quad (5.9)$$

where A_i are the amplitudes of all the normalized wavelet coefficient amplitudes that make up the cluster. Since each normalized wavelet coefficient is drawn from a zero-mean unity-variance Gaussian random variable given pure stationary noise, the total signal energy of a random cluster of N coefficients will be χ^2 distributed with N degrees of freedom for stationary noise without the presence of any transient. The statistical significance of a measure of total signal energy for a trigger is then defined as,

$$Z_{(E,N)} = -\ln \int_{E_C}^{\infty} \chi_N^2(E) dE. \quad (5.10)$$

which is a function of total signal energy E_C and the number of pixels in the cluster N . The significance represents the unlikelihood of measuring a cluster of coefficients with greater total signal energy in Gaussian noise. Appendix A provides calculational details.

5.3.4 Parameter estimation

In addition to the total normalized energy and statistical significance of a trigger, the time-scale wavelet decomposition provides a general time-frequency parametrization of the signal. Each wavelet coefficient represents a projection onto a basis element with normalized amplitude A_i , scale j , index n , and central time $t_{(j,n)} = t_0 + (n+1/2) 2^j / f_s$. For a cluster of N outliers, we use the energy weighted central values to parametrize

the trigger,

$$\text{start time } t_s = \min_i [n_i 2^{j_i} / f_s] \quad (5.11)$$

$$\text{end time } t_e = \max_i [(n_i + 1) 2^{j_i} / f_s] \quad (5.12)$$

$$\text{central time } t_c = \frac{1}{E_c} \sum_i \left[|A_i|^2 \left(n + \frac{1}{2} \right) 2^{j_i} / f_s \right] \quad (5.13)$$

$$\text{central scale } j_c = \frac{1}{E_c} \sum_i |A_i|^2 j \quad (5.14)$$

$$\text{total normalized energy } E_c = \sum_i |A_i|^2 \quad (5.15)$$

$$\text{number of pixels } N = \sum_i 1 \quad (5.16)$$

$$\text{cluster significance } Z = -\ln \int_{E_c}^{\infty} \chi_N^2(E) dE. \quad (5.17)$$

In addition we define the more interpretable approximate central frequency,

$$\text{central frequency } f_c = f_s / 2^{j_c}. \quad (5.18)$$

Because the frequency isolation of the dyadic wavelet transform is poor, particularly when using the Haar transform, we do not attempt to calculate frequency bounds or bandwidth of the signal. The frequency resolution can be improved by making use of a wavelet packet decomposition which no longer uses a multiresolution logarithmic tiling of the time-frequency plane, but a more traditional rectangular tiling or something in between. Such an approach is used by the Waveburst pipeline for the detection of gravitational-wave bursts [79].

5.3.5 Implementation

KleineWelle runs under the Data Monitor Tool (DMT) framework [91] which is a C++ framework for rapid processing of instrumental data and publishing of results. DMT is responsible for the majority of real-time feedback from instrumental data in the control room as well as the production of noise spectra, trends for transient

rates, sensitivity, and other values, and the production of various data-quality flags. In online mode, the DMT environment provides a steady stream of live data from the instrument which is read from a shared memory partition. In offline mode this stream can be simulated from a series of frame data files on disk. The DMT libraries also provide the necessary data structures and many of the signal processing routines needed for standard signal analysis, as well as routines which can publish results to disk or SQL databases. This keeps the *kleineWelle* specific code compact and relatively easy to maintain.

LIGO's fifth science run (S5) occurred between November 2005 and October 2007. During this time, the three LIGO instruments collected one full year data when all three detectors were operating in coincident science data taking mode. The individual duty cycles were 78% for H1 and H2, and 66% at L1. During this time, *kleineWelle* operated in a pseudo-online mode recording transients for a number of environmental and interferometric channels at each site, including the gravitational-wave channel. The *kleineWelle* jobs were launched on fresh data every five minutes under the Onasys infrastructure [92]. In this environment, jobs were set up when the instrument was fully locked. The Onasys daemon handled the submission of jobs to LDAS computing clusters [93] at each LIGO site and monitored their progress. Online triggers were used for prompt detector characterization as well as a low latency (~ 1 day) burst search for outstanding events.

In parallel to the online production, an offline production of triggers was generated using the LDAS computing cluster resources at Caltech. The offline production was necessary to cover holes in the online run due to failed services or other technical issues. Offline triggers were updated every few months with around 99% coverage of science data. In total, 201 interferometric channels and 155 environmental channels with frequencies of interest between 1 Hz and 4096 Hz (Appendix B) were processed across the two sites and three interferometers. These offline triggers were used for the production of event-by-event vetoes for burst and inspiral searches, as well as any other followup studies concerning auxiliary channel transients. A similar production of *kleineWelle* triggers took place at Virgo using the Virgo interferometer V1 and site

environmental data.

The LIGO instruments are currently in their sixth science run (S6) after a number of upgrades lead up to the successful commissioning of Enhanced LIGO [71]. For this run, `kleineWelle` was moved to true online production under the DMT infrastructure, continuously reading and processing live data off shared memory in 32-second intervals. In total, 339 interferometric channels and 158 environmental channels are currently being processed across the two sites and interferometers H1 and L1 (H2 is not being run for S6). At each site, the complete `kleineWelle` production runs on a single 2.6 GHz AMD Opteron core processing multiple channels in parallel. The memory footprint is around 1 GB in addition to a shared 1 GB data input buffer. Under the current configuration, processing runs at a factor of 2.6 faster than real-time resulting in a maximum latency of around 45 seconds between the time data is taken and transients written to disk. Approximately 3 million transients are recorded per site each day. The transient rate at chosen thresholds for the gravitational-wave channel alone is ~ 0.2 Hz. The low latency for online auxiliary channel trigger production is necessary in order to quickly identify disturbances in the instruments. There is a current effort to detect gravitational-wave candidates within minutes of their occurrence in order to trigger a targeted electromagnetic followup [94]. We want to make sure not to trigger on events which can be easily traced back to non gravitational-wave effects.

Chapter 6

Detector characterization, data quality, and vetoes

6.1 On-line detector characterization

An important application of online noise transient detection is continuous monitoring of the quality of gravitational wave data and state of the various instrumental subsystems. While stochastic rms noise leads to poorer resolvability and sensitivity to signals of all types, noise transient rates contribute directly to the background of transient searches and trace the overall stability of the interferometer. The characterization of the transient behavior of the instrument is typically done in terms of the transient rate versus time in the form of a rate trend. For S4, S5, and S6 *kleineWelle* rate trends of gravitational wave and auxiliary channels were made available to operators with minutes latency to aid in tracking changes to the instruments.

Noise transients in the gravitational wave data itself understandably are subject to the closest monitoring and follow-up by the LIGO glitch working group [95]. In addition to transient rate versus time, distributions in strength and frequency are checked periodically. A noisy non-stationary line due to an approximately monochromatic noise source might not be strong enough to cause concern in the noise spectrum, but could lead to a significant excess of detected transients around the resonant frequency. The regularity of trigger times can also point to interesting effects. During S3, ex-




Figure 6-1: Rate of transients vs. time as identified by the kleineWelle pipeline. The rates are calculated every few minutes and available through a web interface. Shown here are correlations at L1 between single-detector noise transient rates in the gravitational wave channel (DARM_ERR: error signal for differential arm length control) and alignment channel transients from the mirror wavefront sensors.

cess triggers were detected exactly at the 1/16 second boundary reflecting glitches in the data acquisition system. During S5, a dramatic excess of triggers occasionally showed up at the top of the hour. This was eventually traced back to scheduled digital snapshots of the detector state whose computation load interfered with the detector operation. Other regularities can show up in the time between successive triggers. The LIGO instruments are extremely sensitive to seismic noise in the 0.1-1 Hz range which is caused mostly by weather and ocean waves. Transients are found at points of extreme mirror motion leading to a preferred time between triggers at multiples of half the inverse of the micro-seismic noise peak frequencies.

6.2 Data quality flags

Data quality (DQ) flags identify epochs in science data which may have a negative impact on the analyses due to errors in data acquisition, poor sensitivity, excessive contribution to the false event rate, or general un-trustworthiness of data. The intervals constructed target known problems with the instrument or environmental conditions. Individual data quality flags are evaluated by their effectiveness at removing single-interferometer and coincident noise transients, especially large amplitude ones, from the gravitational-wave data streams. To make sure data quality flags remain independent of the presence of a true gravitational wave, we check that they are not triggered by hardware injections where simulated signals are physically injected into the instrument by differential actuation of the mirrors.

Data quality flags from individual sources are selected for use by the burst analysis based on their effectiveness at removing non-Gaussian transients from the data while minimally effecting the live-time of the search. Each set of flags is tested over single-interferometer transients found using *kleineWelle* as well as a sample of time-shifted triple-coincident background events. Single-interferometer analysis provides the best statistics and a clear picture of what happens at each instrument, while the time-shifted coincident analysis preferentially targets the sources of background that should appear in the real search.

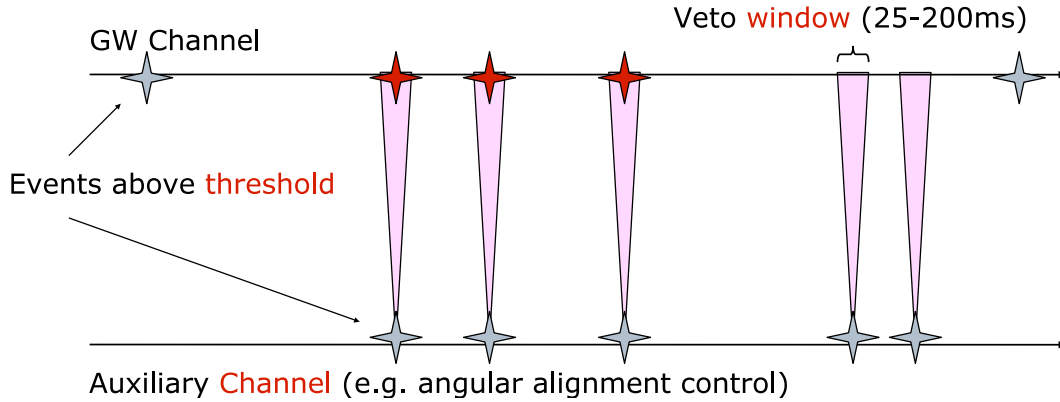


Figure 6-2: Simple veto logic used by auxiliary channel event-by-event vetoes for S5. A channel-specific veto window is constructed around the central times of auxiliary channel triggers above a certain threshold (represented as gray stars on the bottom time-line). Gravitational-wave channel events are removed (red stars) if their central time falls within the veto window. The veto procedure removes a small amount of live-time from the analysis.

To organize the use of data quality flags in transient searches, individual sources are categorized by the LIGO and Virgo detector characterization groups according to their severity and our level of understanding [96, 8, 97, 98]. The categories then determine how the flags are used in an analysis. A description of these categories and examples of the sources of disturbances in each for S5 can be found in Table 6.1.

6.3 Event-by-event vetoes for burst background

Event-by-event vetoes attempt to discard individual gravitational-wave (GW) channel noise events by using information from the many environmental and interferometric auxiliary channels which measure non-GW degrees of freedom. Good vetoes are found by looking for situations in which a short (\sim ms) noise transient in an auxiliary channel often coincides within a short interval (\sim 100 ms) with noise transients in the GW channel. The work, then, is in identifying useful auxiliary channels which are well correlated with noise transients in the GW data, choosing the relevant veto parameters to use, and finally establishing that the veto procedure will not systematically throw out true gravitational waves.

Data quality categories for S5 burst analyses

Category	Description
1	<i>Do not analyze.</i> These flags define the set of data that is processed by search algorithms and exclude unlocked, uncalibrated, or non-science mode times, times affected by data corruption or missing data, photodiode saturations and calibration line dropouts which cause extreme signal transients which can affect detector response and power spectral estimates, scattered light contamination, and the often unstable 30 seconds prior to each lock loss.
2	<i>Unconditional post-processing cut.</i> These flags are applied unconditionally to events surviving the analysis pipeline. Therefore, events during these times are never considered as gravitational-wave candidates. Flagged times cover well understood disturbances in the instrument which have a reliable one-to-one correspondence with loud transients and include saturations in the alignment control system, severe impulsive glitches in the power mains, uncertain calibration, and large glitches in the thermal compensation system.
3	<i>Poor quality data.</i> Data during these times have poor sensitivity or excess rates of loud noise transients and are removed from upper limit searches in order to set a better upper limit with more reliable background. Flagged times cover a full 120 seconds prior to lock loss, noise in power mains, the possibility of cross instrument scattered light contamination from an unlocked Hanford interferometer, very poor rms noise, severe seismic noise or wind speed, earthquakes and hurricanes.
4	<i>Advisory flags.</i> These flags have not shown a particularly strong association with noise transients in the gravitational wave channel but cover times when one has special reason to investigate data quality carefully in the case of a candidate event.

Table 6.1: The organization of data quality flags for LIGO S5 burst searches. The flags are categorized by our level of understanding and the severity of their impact on transient analysis. The categorization then determines how the flags are used. Detail in the data quality procedure for transient searches can be found in [96, 8, 97]. A similar categorization is used for Virgo data quality [98].

In this way, the event-by-event vetoes complement the data quality flags described previously. While data quality flags generally identify a stretch of time when background rates are elevated or the data is in some way less trustworthy or problematic, event-by-event vetoes are meant to specifically target individual transient disturbances in the instrument, so they cover a comparatively short time and are always applied in post-processing. In their application to LIGO data through S5, vetoes have been applied in a very similar fashion to data quality flags where the veto choices are used to create a list of short time intervals covering excluded times. Events which fall within the vetoed time intervals are removed in post-processing. More complicated veto methods (for example those which use additional information from the gravitational-wave channel) may not be easily representable as a direct list of exclusion segments. We choose to use the simple method of veto intervals and veto application based on a single event time alone so that false dismissal probabilities can be estimated in a straightforward manner.

The event-by-event vetoes used for the S5 burst search onward are divided into two categories which follow the same naming convention used for data quality flags (Table 6.1). Category 2 vetoes are a conservative set of vetoes targeting known electromagnetic and seismic disturbances at the LIGO and Virgo sites such as the problematic site-wide magnetic disturbances caused by power main glitches at the Hanford site (Figure 5-4). These are identified by requiring a coincident observation of an environmental disturbance across several channels at a particular site.

Category 3 vetoes make use of all available auxiliary channels shown not to respond to gravitational waves. While the physical source of a particular disturbance may not be well understood or the coupling between the gravitational-wave channel and auxiliary channel poorly modeled, the presence of a transient in an auxiliary channel not sensitive to gravitational waves can still be good reason to reject a transient in the gravitational-wave channel if a strong correlation between noise transients from the two channels can be made.

6.3.1 Veto effectiveness metrics

A collection of non-overlapping veto time intervals (segments) is judged by its effectiveness at flagging gravitational-wave channel noise transients. This will depend on the live-time under consideration and the particular set of noise transients being tested. Given a set of veto segments and noise transients we define the following parameters:

Veto efficiency is the fraction of noise transients from the sample which are removed by our veto method. We use a simple veto logic where an event is vetoed if its central time falls within a particular veto interval. Veto efficiency is dependent on the initial sample of noise transients being tested.

Dead-time fraction is the fraction of live-time flagged by all the veto intervals. Assuming that real gravitational-wave events are randomly distributed in time, dead-time fraction represents the probability of vetoing a true gravitational-wave event by chance. For a realistic population of sources, detected gravitational-wave events are likely to scale with the sensitivity of the instrument which is non-stationary. Still, dead-time fraction is a useful population-independent proxy for false dismissal probability and is easy to calculate. If the veto segments are completely uncorrelated with the noise transients, measured veto efficiency should approach the dead-time fraction. A veto efficiency greater than the dead-time fraction indicates a correlation between the triggers and veto segments.

Veto significance reflects the level of statistical significance of a measurement of excess veto efficiency under the assumption of uncorrelated veto segments and noise transients. Under this assumption, the number of events that fall within the flagged dead-time is Poisson distributed with mean value μ equal to the number of noise events times the fractional dead-time, or equivalently, the noise event rate times the duration of veto segments. We define the statistical significance of actually observing N vetoed events as,

$$Z_{(N,\mu)} = -\ln [P_{\text{Pois}}(x \geq N|\mu)]. \quad (6.1)$$

Appendix A provides calculational details.

Efficiency/dead-time provides a measure of the performance of a set of veto segments. The ratio represents the relative excess rate of noise transients within the veto segments to those outside. All things equal, we want to remove the times corresponding to the highest noise transient rates first. Significance reflects our confidence that the measured performance is not simply a statistical fluctuation. A veto with low efficiency/dead-time of just a small excess over 1.0 may still be considered highly significant if a large enough test sample is used. Similarly a high efficiency/dead-time measurement may be a result of low number statistical fluctuations if the test sample is too small. Requiring a level of statistical significance corresponding to the number of investigations being done helps protect against introducing random vetoes.

6.3.2 Veto safety

Veto safety deals with the possibility that a veto might systematically reject true gravitational waves. Although the many auxiliary channels measuring test points in the instrument and environment are not supposed to be sensitive to a passing gravitational wave, we cannot assume that a strong gravitational wave would not leave a signature in channels other than those in the differential arm motion control loop. The source of the coupling could be from mis-tuning of the interferometer, or something as mundane as electronic pickup. In order to explicitly check for an auxiliary channel response, hardware signal injections simulating the passage of a real gravitational wave by actuating on the end mirrors using magnets are performed in each instrument at scheduled times throughout each run. The impulsive, burst-like injections span the entire sensitive frequency range of the interferometer as well as a very large range of amplitudes including very loud events.

Each auxiliary channel is independently checked for veto safety by determining if it has any measurable response to hardware simulations. First, transients are identified in the auxiliary channel using *kleineWelle* as described in section 5.3). A set of veto segments is formed by taking fixed ± 100 ms time windows about the auxiliary channel triggers at some threshold on trigger significance as in Figure 6-2. The dead-time of the veto segments are evaluated for a local region of time surrounding transient

hardware injections. If there is no correlation between hardware injected signals and the auxiliary channel, we expect the number of hardware injection times which fall within the veto intervals by chance to follow a Poisson distribution with mean value equal to (dead-time fraction \times number of injections).

The veto significance of the actual number of hardware injection times flagged determines whether or not there is a measurable correlation between the injections and auxiliary channel transients at the chosen threshold. In a long run such as S5, there are thousands of injections and it only takes a small correlation to be significant. To ensure veto safety, we require the threshold for auxiliary channel triggers to be large enough such that at the chosen threshold and all higher thresholds there is no significant correlation between the resulting veto segments and hardware injection times to the 1% level.

6.3.3 Selection of veto conditions

Category 2 vetoes

Category 2 vetoes are applied unconditionally to all burst searches so they are chosen from the most conservative and best modeled instrumental disturbances. For S5, this included seismic and magnetic transients with known coupling to the mirror motion. Transients from each environmental channel were correlated with sample background transients in the gravitational-wave data, and veto windows and thresholds were tuned in order to cover efficiently the observed overlap. Three classes of environmental channels were adopted as vetoes. For LHO this included transients recorded in 24 magnetometers and voltmeters with a `kleineWelle` significance threshold of 200 and time window of 100 ms, and 32 accelerometers and seismometers with a threshold on the `kleineWelle` significance of 100 and a time window of 200 ms. For LLO these were 12 magnetometers and voltmeters with a `kleineWelle` threshold of 200 and a time window of 100 ms.

The environmental disturbances which caused the most reliable response from the interferometers were generally large enough to be detected site-wide and therefore

were picked up by multiple channels of a given class. This allows us to reduce the dead-time and increase the reliability of the chosen vetoes by requiring coincidence across several channels in a given class before generating a veto for the gravitational-wave channel. The coincidence step keeps genuine site-wide environmental disturbances while greatly suppressing the contribution from noisy non-stationary channels. A veto segment was only created from the window overlap of three or more channels from a particular environmental class. This procedure removed in total about 0.1% of coincident S5 live-time.

Category 3 vetoes

Category 3 vetoes attempt to efficiently clean the data as much as possible from contributions to transient background and attempt to make use of all the measured transient behavior in the instrument. For S5 this includes `kleineWelle` triggers generated over approximately 200 interferometric channels and 150 environmental channels (subsection 5.3.5) which have been determined to be safe according to the procedure outlined in subsection 6.3.2. In many cases the nature of the transient coupling between channels is poorly known and the causal relationships between coincident transient observations are not established. However statistical correlations between noise transients in auxiliary channels and the gravitational wave channel can still be used to construct vetoes which are effective at systematically flagging background. The vetoes constructed may also flag many times unrelated to actual disturbances in the gravitational-wave data and for this reason they are not used as an unconditional veto in a detection search. Their application does however improve the upper limit by allowing for lower thresholds at the same false-alarm probability and the veto segments themselves are useful for establishing confidence in the case of a detection follow-up.

The task of constructing vetoes from the auxiliary channel transients is complicated by the desire to choose optimal veto frequency ranges, thresholds and windows, and the fact that the veto channels themselves can be highly correlated with each other so that applying one veto channel changes the incremental cost (in additional

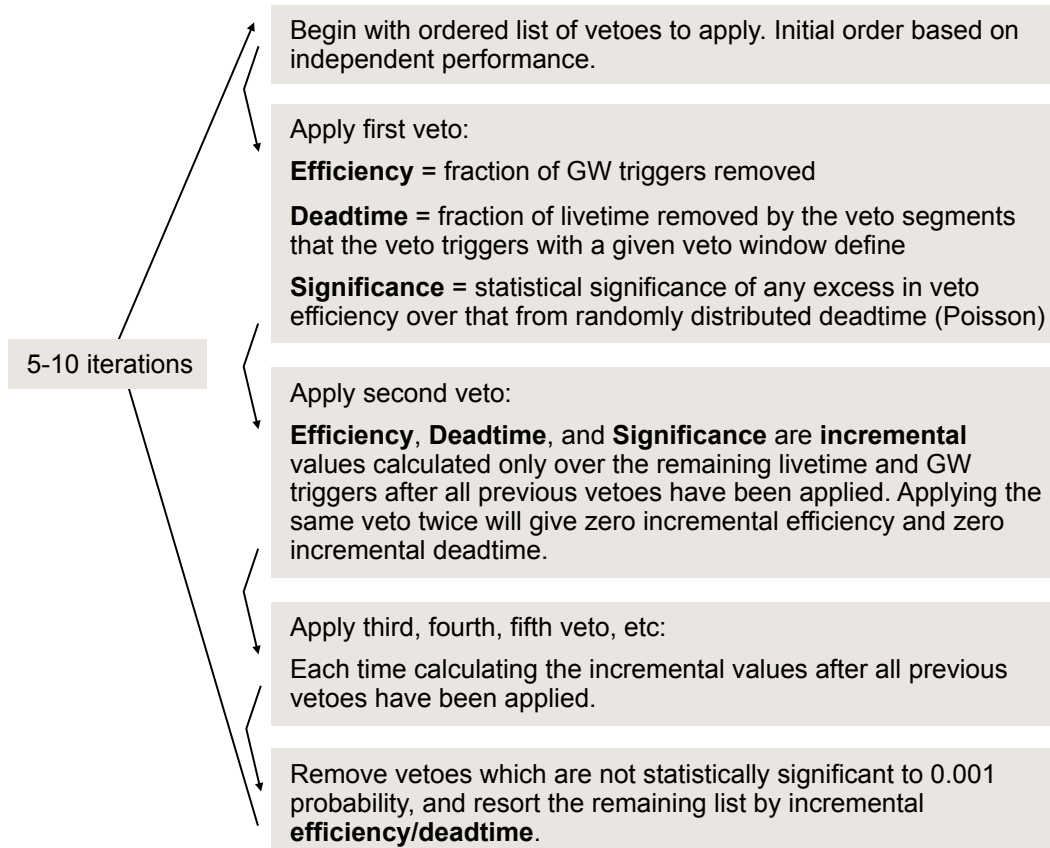


Figure 6-3: Procedure used for ordering a list of veto conditions from best to worst in terms of channel, frequency range, window, and threshold. A set of vetoes can then be chosen for any choice of dead-time threshold by choosing the corresponding top ranked vetoes. The particular method of refining the list order over several iterations is chosen because it minimizes the number of steps which must be done in serial and therefore is easy to implement on a computing cluster.

dead-time) and benefit (in additional veto efficiency) of applying another. Applying all vetoes which perform well by themselves often leads to an inefficient use of dead-time as dead-time continues to accumulate while the same noise events are vetoed over and over.

For a particular set of GW channel noise events, we adopt a hierarchical approach to choose the best subset of all possible veto conditions to use for a target dead-time. This amounts to finding an ordering of veto conditions (veto channel, frequency range, threshold, and window) from best to worst such that the desired set of veto conditions can be made by accumulating from the top veto conditions so long as the dead-time does not exceed our limit, which is typically a few percent.

We begin with an approximately ordered list based on the performance of each veto condition (frequency range, channel, window, and threshold) considered separately. Incremental veto statistics are calculated for the entire list of conditions using the available ordering. This means that for a given veto condition, statistics are no longer calculated over the entire S5 live-time, but only over the fraction of live-time that remains after all veto conditions earlier in the list have been applied. The list is then re-sorted according to the incremental performance metric and the process is repeated until further iterations yield a negligible change in ordering.

The ratio of *incremental veto efficiency* to *incremental dead-time* is used as a performance metric to sort veto conditions. This ratio gives the factor by which the rate of noise events inside the veto segments exceeds the average rate. By adopting veto conditions with the largest incremental efficiency/dead-time ratio, we maximize total efficiency for a target dead-time. We also set a threshold of probability $P < 0.001$ on veto significance (not to be confused with the significance of the triggers themselves). This is particularly important for low-number statistics when large efficiency/dead-time ratios can occasionally result from a perfectly random process.

Vetoes were optimized over several different sets of gravitational-wave channel noise events including low-threshold H1H2L1 coherent WaveBurst time-shifted events, H1H2 coherent WaveBurst playground events (not used in the final analysis), as well as Q Pipeline and KleineWelle single-interferometer triggers. For example, the effect

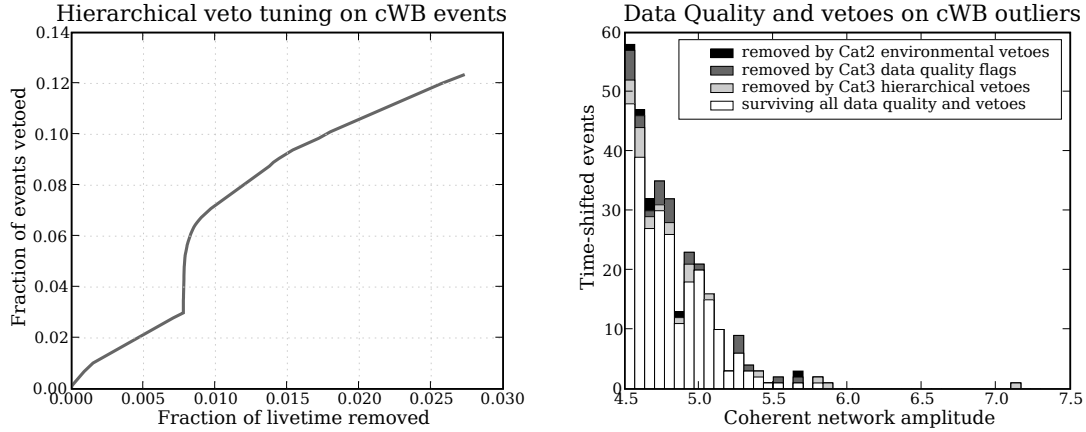


Figure 6-4: Left: Accumulated veto efficiency versus dead-time as vetoes are applied cumulatively down the veto list to coherent WaveBurst time-shift H1H2L1 background during the first calendar year of S5. The best vetoes are applied first, so we see a general decrease in the effectiveness of vetoes at higher dead-time. Vetoes from environmental channels are artificially prioritized over interferometric channels, giving rise to the knee in the plot around 0.8% deadtime where the environmental vetoes are exhausted. Right: Histogram of coherent network amplitude, η , for coherent WaveBurst time-shift background events representing 100 S5 year 1 live-times. The different shades show events removed by data quality cuts and vetoes at various stages in the analysis.

of the S5 data quality flags and event-by-event vetoes on the sample of coherent WaveBurst time-shifted events is shown in Fig. 6-4. A final list of veto segments to apply to the S5 upper limit analysis was generated from the union of these individually-tuned lists.

Chapter 7

Candidate event follow-up

To minimize bias, existing searches for gravitational waves make use of a blind procedure to tune all parameters of the search using Monte Carlo and off-source data before actually running once on on-source data. For burst searches, Monte Carlo methods involve running over simulated ad-hoc signals added to the real data. The signals simulate gravitational waves which sample in time, amplitude, frequency and morphology the region of interest. Off-source data to estimate background is generated for a multi-site search by artificially time-shifting the data stream from one site thereby destroying the coherence of real gravitational waves while ideally preserving the noise properties of the on-source (unshifted) data (subsection 4.3.3).

To implement such a procedure efficiently, all aspects of the analysis must be automated so that the on-source analysis can proceed without any additional human input. The choice of veto conditions (section 6.3), for example, is made using background-dominated single-interferometer transients or time-shift sample background only. The decision about whether or not a particular event in the on-source analysis can be dismissed as an instrumental artifact is then set by predetermined rules. Existing burst searches use information about the non-stationary noise power spectrum, transient background rates and their distributions in strength and frequency, instrumental and environmental disturbances picked up in auxiliary channels, and consistency between the waveforms seen in multiple detectors to separate signals from background. Constructing automated methods of folding this information into the search is a time

consuming process but allows for an objective estimate of the false-alarm probability associated with a given measurement critical for any detection statement. It also allows for the unbiased measurements necessary for proper coverage in an upper limit calculation. While the need for rapid online results has eroded at the concept of total blindness for any subsequent analysis, the methodology of automated blind analysis has not changed.

Events which stand out in the blind analysis are subject to a follow-up procedure [99] which provides a more careful subjective evaluation of the wealth of information relating to the specific event. The follow-up procedure is designed to satisfy the following functions:

- provide a central resource for basic information about the event
- serves as a careful secondary review of the specifics of the end-to-end analysis that relate to the event
- provide an opportunity to check for obvious reasons to dismiss a candidate event (e.g. clear environmental cause) or increase our interest in an event (e.g. optical counterpart)
- outline specific event details which may be tangentially related to interpretation as a gravitational wave candidate (e.g. single-interferometer signal properties)
- provide a subjective but detailed evaluation of information which may enter only crudely (e.g. in the form of hard cuts) in the blind analysis
- provide a test-bed for new ideas which have not been able to make their way into the quantitative detection statistic

Most importantly the follow-up procedure increases our confidence in the validity of the analysis with respect to a specific event outlier by checking the sanity of the many decisions which factor into the automated procedure. This is particularly relevant for gravitational waves where no detection has yet been made.

7.1 Detection checklist

The follow-up procedure for S5 burst searches was implemented in the form of a detection checklist which was to be applied to an event from the blind analysis that had any chance of being reported as a detection. The first time such a checklist was implemented for LIGO burst analysis was during S2 [88] whose search yielded a single H1H2L1 coincident event which narrowly survived all cuts used for the upper limit analysis. These cuts were set to yield a background expectation of $\sim 5\%$, or 1.8 events per year so that an event from background was unlikely but not too surprising. In follow-up, the event was unambiguously linked to the acoustic coupling of an overhead airplane into the mirrors at both Hanford interferometers, thus disqualifying it as a gravitational wave candidate.

The second time the follow-up procedure was exercised was during the end of LIGO's fifth science run. The second calendar year of S5 was analyzed jointly with Virgo's first science run (VSR1) [8]. A single outlier was identified which passed all upper limit cuts, except this time no reason was readily found to clearly reject the event as non-astrophysical in origin. The result was an expanded and more thorough follow-up procedure from that applied to the S2 event. In final form, the detection checklist for S5/VSR1 (Appendix D) consisted of an 80-point list of tests covering the following categories:

- zero-level sanity: reports in detector logs, check hardware injections
- data integrity: frame file check-sum, undocumented injections, check against raw frames
- state of the instrument: obvious disturbances reflected in auxiliary channels, verify coupling for any proposed veto, check by hand against known disturbances: dust, cosmic rays, power fluctuations, acoustic, etc.
- event properties: construct detailed spectrogram, reconstructed waveform and direction, compare background from various methods, check signal consistency across interferometers

- astrophysical interpretation: check for external EM or neutrino events, catalog sources consistent with reconstructed direction, compare waveform against simulations

Ultimately it was decided that while the event passed all follow-up tests, the expected background rate in the blind analysis for events of similar quality (between 1–10% depending on the comprehensiveness of the search) was too high to report a detection. Afterward it was revealed that the source of the event was a blind hardware injection designed to test the end-to-end analysis pipeline.

7.2 Blind injections during S5

Beginning with S5, the burst and inspiral analysis teams were subject to a blind injection challenge. In this test, a small random number (possibly zero) of simulated events are coherently injected into the LIGO interferometers. The data analysis teams know of the possibility of an injection but not whether or not an injection actually occurs nor anything about the morphology of the waveform except that it should be within the target parameters of the search. Only after the completion of the entire analysis, including interpretation of any outliers, are the details of any blind injections revealed. In this way, the blind injection serves as a test for the end-to-end procedures involved in making a detection claim.

The event identified by the burst analysis was the only blind burst hardware injection during S5. There was also a blind inspiral injection which was missed due to data quality issues [100]. The blind injection provided a special opportunity to exercise the follow-up procedure without the bias introduced from looking at a known simulation. The following section summarizes the investigations of the event and was prepared prior to the revelation of the event as a simulation. It has been edited only briefly for clarity in order to maintain the unique perspective.

7.3 Follow-up of S5/VSR1 burst outlier

7.3.1 Introduction

Burst candidate event 070922 (September 9, 2007) was first observed in the burst group KleineWelle-CorrPower S5 online analysis for H1L1, and reported at the next weekly teleconference the burst group holds. The KleineWelle-CorrPower infrastructure consisted of single-detector kleineWelle (section 5.3) triggers generated online at the sites, with a centralized daily coincidence and CorrPower [101] follow-up. The H1L1 event stood out clearly above the expected 1-day background in CorrPower distributions.

The September 22, 2007 event was quickly verified by other burst methods, Astroburst/BlockNormal [102] and Q Pipeline [74] which also generated online triggers. KleineWelle and Astroburst recorded signals in H1 and L1, while Q Pipeline showed an additional weak signal present in H2, visible in Q-scans. Astroburst was looking only for triple-coincident events throughout the run (none were found), while no online coincidence was performed for Q Pipeline triggers. Coherent Event Display [103], which generally presents details of Coherent Waveburst events, showed the signals to be very consistent across the three detectors, and provided the first position and waveform reconstruction of a possible gravitational wave.

The burst group resurrected a detection checklist, unused since the analysis of the S2 data. A large number of items were added to investigate the event in detail, and the checklist would serve as a model for any future candidate events. Items were assigned to relevant group members, and the checklist was gradually completed over the course of a couple months. However it was decided that the event, appearing essentially at the end of S5/VSR1, could only be interpreted properly in the context of a complete S5/VSR1 analysis. Thus the S5 second year all-sky burst search, which was finished over a year later, became a prerequisite for bringing the event to a close. The September 22 event eventually became known in the burst group as the *equinox event* due to its anticipation of the 2007 Autumn equinox.

detector	GPS time	f	Q	SNR
H1	874465554.7158	96.8 Hz	4.7	11.8
H2	874465554.7119	110.9 Hz	22.6	5.4
L1	874465554.7100	118.3 Hz	4.7	11.3

Table 7.1: Basic information about the candidate event 090722. The time, frequency, and Q represent the best-match sine-Gaussian as determined by the Q Pipeline [74, 104].

7.3.2 The event

Basic information

The equinox event occurred on Saturday September 22, 2007 03:05:40.71 UTC (Friday September 21 20:05 PDT, 22:05 CDT, 874465554.71 GPS). Around 100 Hz, the event was a low-frequency, low- Q event similar to common transients seen in the instruments. A basic table of Q Pipeline [74, 104] derived parameters follows. The h_{rss} , about $3 \times 10^{-22} \text{ Hz}^{-1/2}$, is not estimated reliably by the Q Pipeline as it fits the waveform in whitened space.

The measured signal

Figure 7-1 shows whitened plots of the uncalibrated signal from the gravitational-wave channel LSC-DARM_ERR for each instrument over 60–2048 Hz and a restricted 60–140 Hz band. The signal, clearly visible in H1 and L1, is low Q with only a couple cycles.

Figure 7-2 shows time-frequency spectrograms of the calibrated data. The spectrograms are created using the Q Pipeline basis at fixed Q . The space is oversampled in time and frequency corresponding to low energy mismatch between tiles. Spectrograms made in this way are known as Q-scans [104]. Although the Q-scan shown for H2 is generated for $Q=4.2$ for direct comparison to the H1 and L1 Q-scans, the most significant projection onto a minimal-uncertainty waveform in H2 has a Q in the mid-20s.




Figure 7-1: Whitened plots of the uncalibrated signal LSC-DARM_ERR (differential arm length error signal) for each instrument are shown over 60–2048 Hz and a restricted 60–140 Hz band.

H1-H2 consistency

Figure 7-3 shows various combinations of calibrated H1 and H2 strain data which test waveform consistency between the co-located, aligned detectors. The first plot is the coherent sum $(H1/S_{H1})+(H2/S_{H2})$. This combination is designed to maximize signal-to-noise for a true gravitational wave and shows a signal which is more significant than purely H1 alone. From the second plot of the H2–H1 null stream, we see that H2 is not completely consistent with H1, and there is a clear signal preceding the event time at the same frequency present in the H2–H1 residual. However the secondary H2 excess at the time of the event is removed by the subtraction. Strangely, the preceding signal in H2 shows a larger SNR in H1–H2 than in H2 alone. The third plot shows the *incoherent* H1/H2 ‘null’ stream [96] which represents what we would expect to see in the null stream for randomly correlated signals of the same amplitude. The incoherent null stream shows contributions from both the main event and the preceding additional signal from H2, and we see that the actual H1–H2 null stream

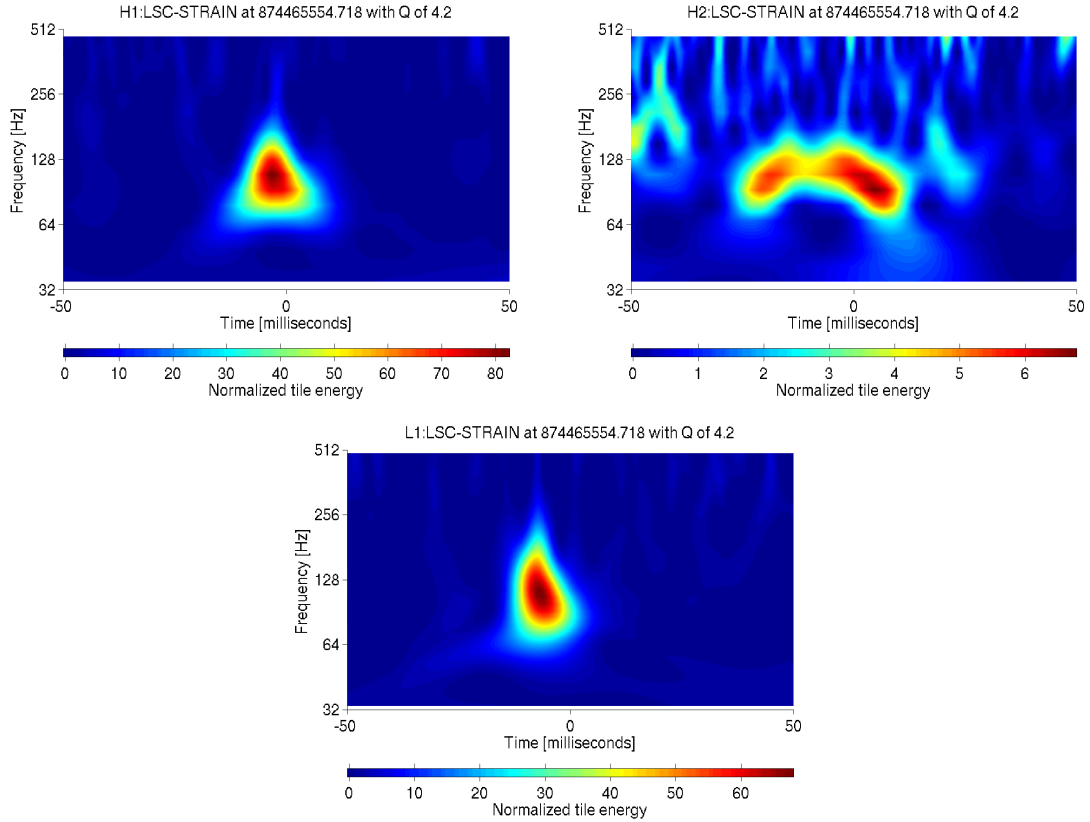


Figure 7-2: Spectrograms of LIGO calibrated strain data for each instrument. Normalized energy Z is related to signal-to-noise $\rho = \sqrt{2Z}$. Note the different color scales in each plot, and the much weaker signal in H2. The spectrograms are generated over the Q Pipeline tiling [74, 104].

is much weaker than expected from random signals.

We can also check where the H1H2 event falls in the space of correlated vs. coherent normalized energy. The coherent energy represents the coherent sum (first plot in 7-3) which maximizes signal-to-noise assuming a common waveform. The correlated energy is the signal energy present in the coherent sum minus that expected from randomly correlated H1 and H2 waveforms of the same strength. The correlated energy goes negative for anti-correlated signals. The left plot shows the equinox event in the context of other time-shift and unshifted H1H2 triggers in the S5 second-year search. It also shows where the simulated burst signals lie. The right plot shows the event in the context of simulations of the same frequency, Q , and amplitude.

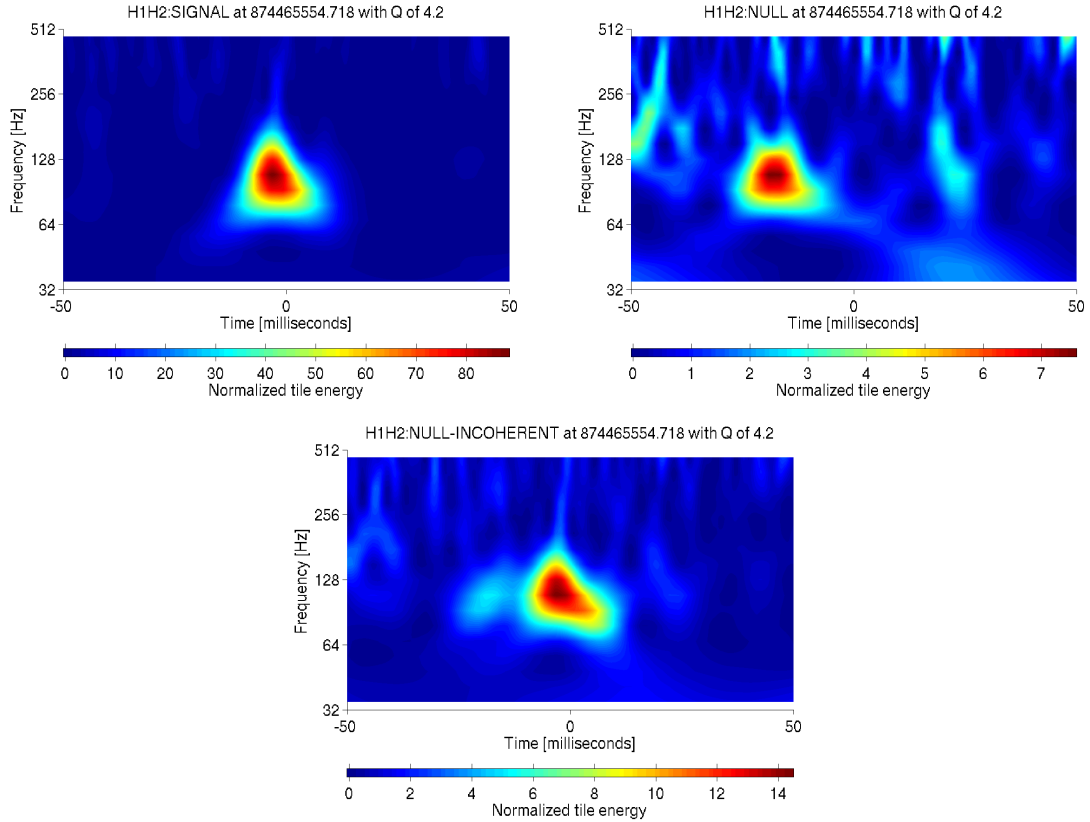


Figure 7-3: Consistency of the measured strain waveforms in the two co-located aligned Hanford detectors. The first is a spectrogram of the frequency-dependent optimal combination $(H1/S_{H1})+(H2/S_{H2})$ designed to maximize signal-to-noise of the combined signal. The second is the H2–H1 null stream which for calibrated data should leave no gravitational-wave signal. The third spectrogram shows the *incoherent* H1/H2 ‘null’ stream which shows what we would expect to see in the null stream for randomly correlated signals of the same amplitude.

7.3.3 State of the instruments

Detector state

All five LSC/Virgo detectors were in the middle of long stretches of science mode (Table 7.2), with no major anomalies or data quality issues.

Hardware injections

Hardware burst injections were scheduled just six minutes after the event, beginning at 874465959 (03:12:25 UTC). These show up as yellow bars in the Q Pipeline plots

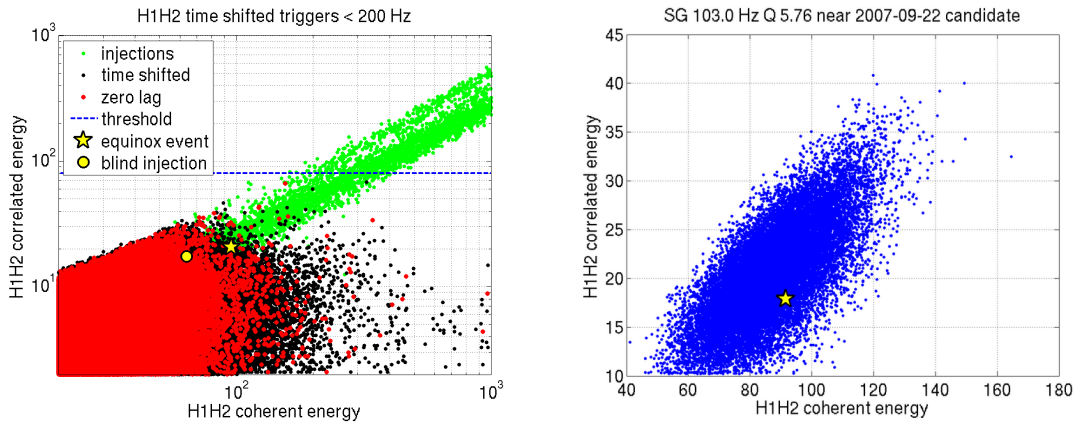


Figure 7-4: Q Pipeline correlated vs coherent energy in the context of expected background (left) and simulated signals (right). The left plot shows the burst event (equinox event) as well as an earlier semi-blind inspiral injection. The equinox event falls within the expected parameters for both all burst simulated signals (green bar) and specifically the closet match sine-Gaussian simulation (blue cloud). Black dots on the left plot show the population of sample time-shift background.

detector	state	start time	relative	stop time	relative
H1	Science Mode	874438904	-26650	874501515	+35961
H2	Science Mode	874441095	-24459	874478798	+13244
L1	Science Mode	874452909	-12645	874488229	+22675
G1	Science Mode	874453140	-12414	874479600	+14046
V1	Science Mode	874449546	-16008	874547216	+81662

Table 7.2: State of the LIGO, GEO600, and Virgo instruments during the time of the detected event. It is remarkable that all five instruments were operating in long stretches of science mode.

shown in section 7.3.3. The burst hardware injections were identical strain waveforms injected in coincidence across the three LIGO instruments. Weak persistent pulsar hardware injections were also running at the time of the event. The persistent injections can occasionally cause transients if they suddenly turn on or turn off causing a discontinuity, but do not appear related to this particular event.

Analysis of excitation channels by `kleineWelle` (excluding the blind injection channel) shows that the signal was not accidentally injected into any of the LIGO instruments else they would have shown up clearly in the excitation channels themselves.

Calibration

The event is seen both in methods which analyze uncalibrated gravitational-wave data (`kleineWelle` and `Astroburst`), as well as methods which use calibrated $h(t)$ strain data (`Q/Ω Pipeline` and `Coherent Waveburst`). There are no artifacts in the calibration channels or problems with calibration at the time of the event.

Single-detector noise transients

A plot (Figure 7-5) of the non-overlapping Q triggers within ± 8 seconds of the event shows weak broad-band transients in the H1 detector. No explanation of these transients has been found.

Q pipeline triggers for each detector are plotted in rate trends covering the hour of the event (Figure 7-6). The event occurs six minutes from the beginning on the plot, minutes before the hardware injections which are marked by the yellow bar. The broad-band transients in H1 seen in the ± 8 seconds surrounding the event show up in the sixth point from the left as a substantial increase in these un-clustered trigger rates. The H1 excess is not seen in clustered rates which only represent the loudest trigger in each second. H2 and L1 rates are flat and quiet.

The strength distribution of the triggers is reflected in scatter plots (Figure 7-7) with weak blue, moderate green, and loud red triggers. The yellow bars containing many loud triggers mark a set of burst hardware injections. The same events are plotted in frequency-vs-time in Figure 7-8.

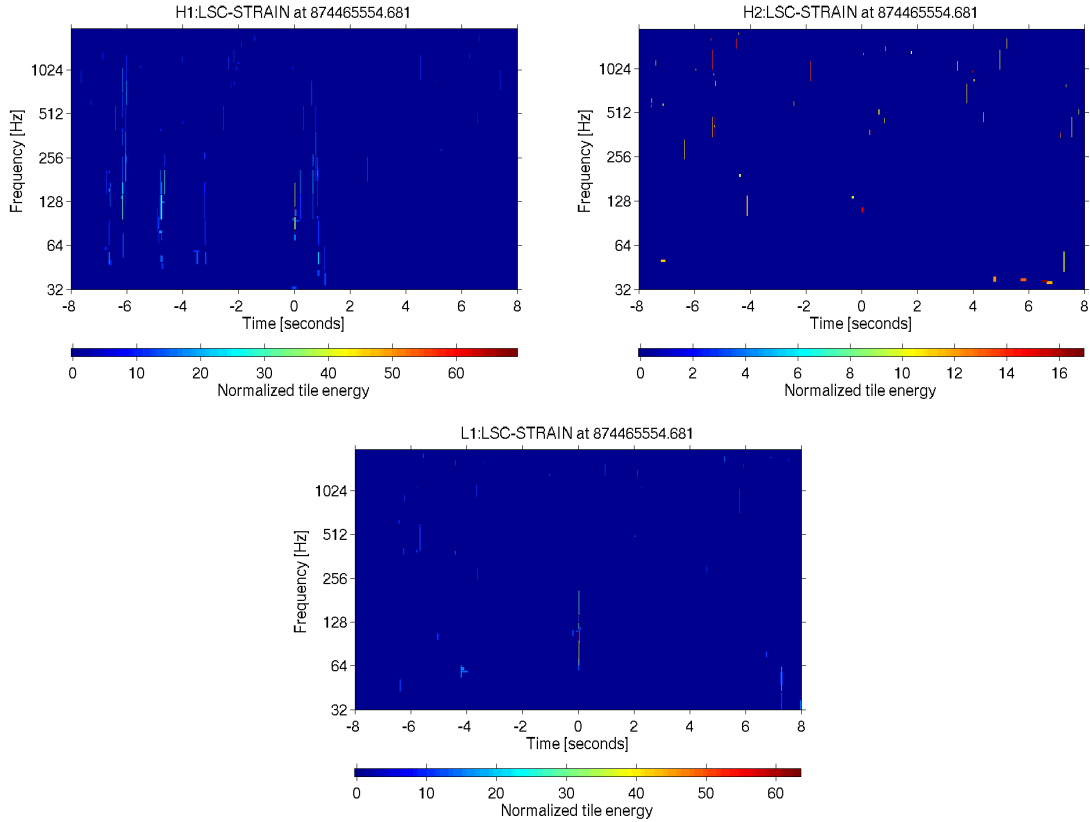


Figure 7-5: Plots showing Q Pipeline single-interferometer non-overlapping triggers in time-frequency space in ± 8 s of the event time. H1 shows a large number of weaker transients in the immediate neighborhood of the event.

In these plots, the equinox event shows up clearly as the first trigger in green for H1 and L1. For H2, the event with an SNR of 5.4 cannot be resolved in this plot. We can also see the series of small transients in H1 that accompany it. Clusters of triggers show up later in the hour for H1, though the frequency of the strongest trigger is no longer at 100 Hz. For L1, however, there are a couple more green triggers which show up at 100 Hz throughout the hour, and many throughout the day. This population is not particularly uncommon for that detector, though usually there are more strong lower-frequency triggers as well. The event in H2 is the third blue dot from the left in the horizontal line of triggers just below 128 Hz.

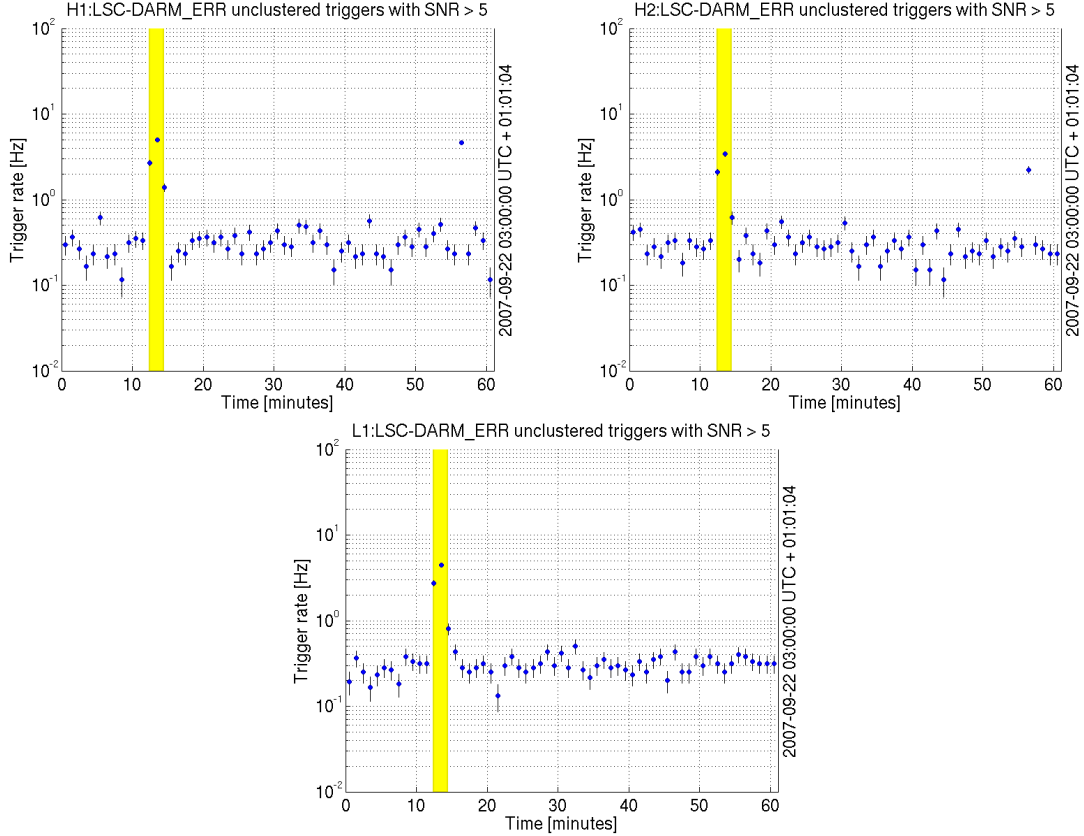


Figure 7-6: Q Pipeline trigger rate vs. time during the hour of the event which occurs at 3:06 UTC. The rate is for non-overlapping (in time-frequency) triggers with $\text{SNR} > 5$ prior to 1 Hz down-selection. The event occurs during an increase in H1 trigger rate. The yellow bar marks a series of hardware injections.

Data Quality and vetoes

There are no active data quality flags in any of the instruments at the time of the event. Vetoes were checked by looking for nearby *kleineWelle* auxiliary channel triggers, as well as looking at full-frame Q-scans at the time of the event. While many auxiliary channels are not particularly quiet, there is no indication of a coincident transient observed in correlation with the gravitational-wave signals. There is no signal observed in ASI for any of the interferometers.

Auxiliary *kleineWelle* triggers found within $[874465553, 874465556)$ and with *kleineWelle* significance above 35 ($\text{SNR} 15$). Peak time is with respect to the equinox event peak GPS time, 874465554.710 . The times of the events are ordered and are in seconds.

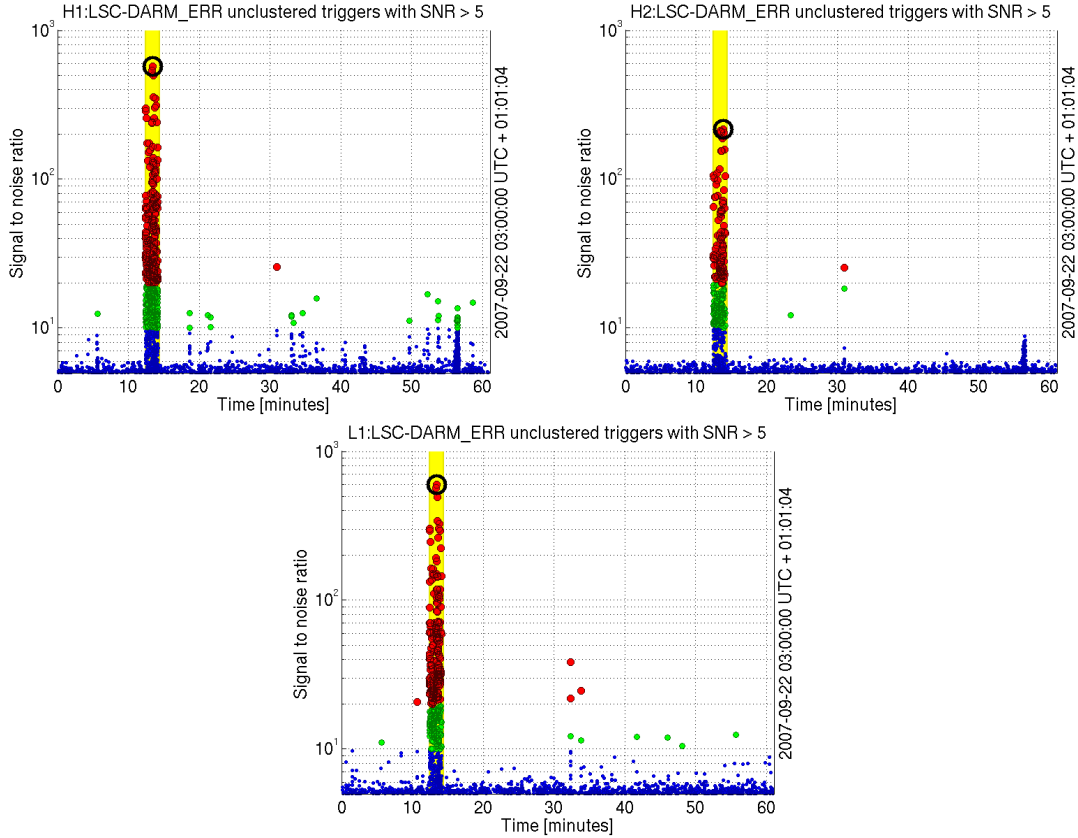


Figure 7-7: Q Pipeline trigger SNR vs. time during the hour of the event which occurs at 3:06 UTC. The event is seen as the first green dot in the H1 and L1 scatter plots. The yellow bar marks a series of hardware injections.

Frequency is in Hz.

At -200 ms there is also a 2.6 Hz excursion in H0:PEM-BSC3_ACCX, an accelerometer sitting on one of the optics tables, which was followed up by the burst and glitch groups. However this channel was noted in the detector e-log as having excess noise inconsistent with other accelerometers for a month surrounding the event, and Q-scans of random times (as opposed to times triggered by a gravitational-wave channel noise transient) often showed coincident disturbances in the channel. None of the environmental channels at Hanford show a signal which would account for that measured in the gravitational-wave channels using the known environmental couplings.

The DMT monitor PdNMon checks to see if the recorded signal was the same in

IFO Channel	Time [s]	Dur. [s]	Freq [Hz]	Significance
h0 coilmagx	-1.651	+- 0.001	488.0	80.68
h1 mcf	-1.577	+- 0.003	402.0	76.40
l0 coilmagz	-0.991	+- 0.003	208.0	52.22
l0 coilmagx	-0.991	+- 0.002	420.0	56.10
h0 coilmagx	-0.944	+- 0.004	136.0	49.35
l1 mcfhifreq	-0.831	+- 0.000	1686.0	37.27
h1 etmxexcdaq	-0.784	+- 0.003	645.0	36.21
h2 etmxexcdaq	-0.784	+- 0.003	565.0	35.95
l1 etmxcal	0.009	+- 0.004	395.0	37.69
h1 mcf	0.071	+- 0.003	274.0	115.87
l1 etmxexcdaq	0.091	+- 0.009	380.0	38.39
h1 asac	0.202	+- 0.001	568.0	54.55
h1 refl1	0.208	+- 0.002	284.0	40.04
h1 asac	0.616	+- 0.001	658.0	90.96
h1 asac	0.665	+- 0.001	411.0	189.00
h1 refldc	0.657	+- 0.002	221.0	41.76
h1 mcl	0.657	+- 0.002	359.0	126.00
h1 pobi	0.673	+- 0.004	263.0	36.63
h1 prcctrl	0.665	+- 0.004	179.0	38.33
h1 refl1	0.665	+- 0.003	219.0	132.23
h0 coilmagx	0.767	+- 0.004	373.0	701.90
h1 asac	0.759	+- 0.001	652.0	195.40
h1 mcl	0.759	+- 0.004	220.0	48.90
h1 refl1	0.755	+- 0.002	443.0	55.73
h1 asac	0.837	+- 0.009	376.0	318.76
h1 mcl	0.837	+- 0.008	131.0	223.63
h1 michctrl	0.868	+- 0.014	125.0	88.52
h1 pobq	0.852	+- 0.011	161.0	95.29
h1 refldc	0.837	+- 0.009	131.0	96.40
h1 refl1	0.868	+- 0.010	186.0	197.75
h1 prcctrl	0.868	+- 0.014	125.0	61.63
h1 pobi	0.868	+- 0.013	135.0	65.71
h1 reflq	0.868	+- 0.013	78.0	64.74

Table 7.3: List of auxiliary channel triggers near the time of the event with kleineWelle significance above 35 ($\text{SNR} \gtrsim 15$). The peak time is displayed relative to the event peak time. None of the detected transients triggered a veto for the event under the predetermined veto choices.

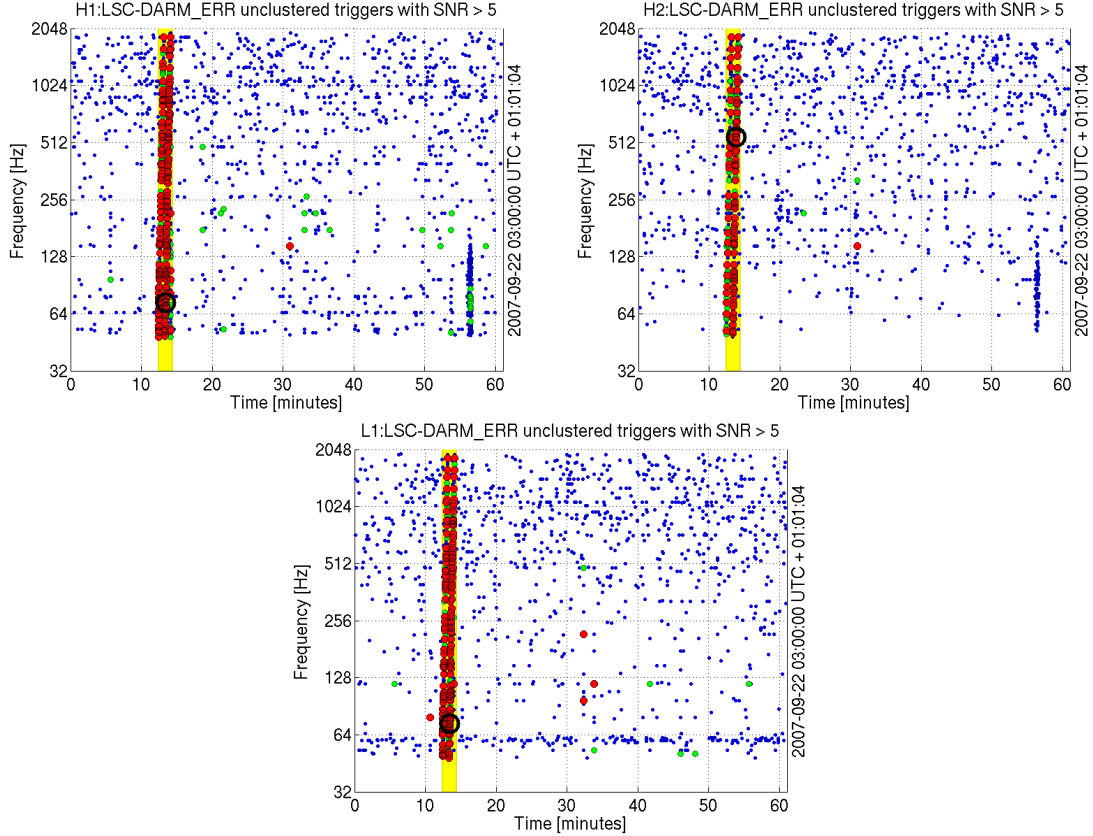


Figure 7-8: Scatter plot of Q Pipeline triggers in frequency vs. time. Green dots have SNR > 10 and red dots have SNR > 20. The yellow bar marks a series of hardware injections.

all four photodiodes at each instrument, useful for distinguishing transients caused by dust particles. There were no PdNMon triggers at the time of the event.

7.3.4 All-sky search and background estimates

KleineWelle+CorrPower online analysis

Figure 7-9 shows unshifted and time-shift background distributions estimated from 100 time-shifts for the H1L1 high threshold kleineWelle+CorrPower online analysis over the entire S5 run. The second plot includes a rudimentary H1H2 null-stream veto based on Q/ Ω Pipeline triggers, and is necessarily only on triple-coincident live-time. The equinox event is the loudest event at zero time-shift with $\Gamma = 9.97$, and the expected background for the analysis above this Γ value is 3.7 events per S5

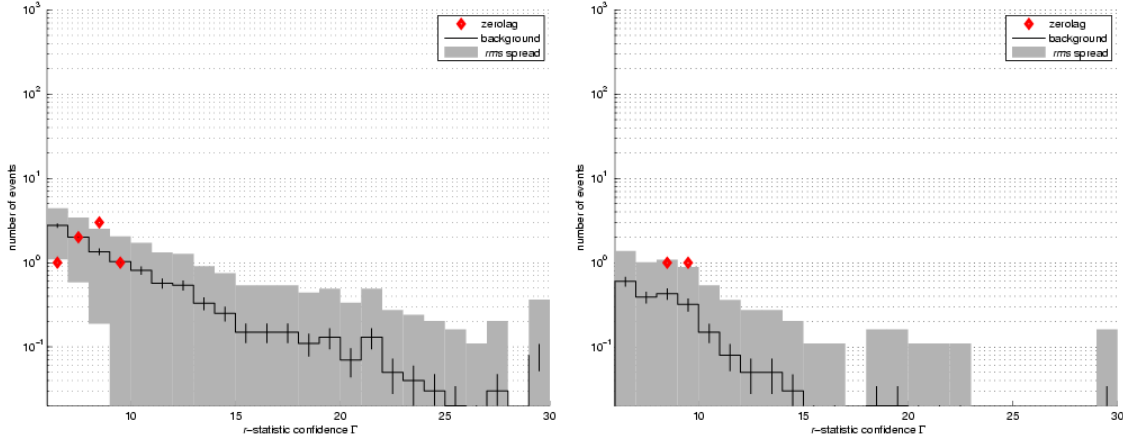


Figure 7-9: Results from an S5 online search using coincident *kleineWelle* 5.3 triggers from the gravitational-wave data with *CorrPower* [101] follow-up. The online search was less sensitive than the offline searches using *Coherent Waveburst* and the Q/Ω Pipeline and was meant to quickly identify extraordinary burst events. On the left is the distribution in *CorrPower* significance Γ of events in red and the expected background in black. On the right is the result after the post-analysis application of an H1–H2 null stream veto which was part of the Q Pipeline analysis for S5. In the online search, the event does not stand out in the context of the entire S5 run.

search without the null stream, and 0.5 events with the null stream veto applied. Originally during the online analysis, the event was compared against its 1-day background which showed a rate of 1 event per 100 calendar days, prompting the extensive follow-up analysis. Otherwise the event does not stand out above background using *kleineWelle*+*CorrPower* over the entire S5 run.

The other event in the post null-stream set is from December 4, 2006 (GPS 849267817.78). It is low frequency (79 Hz in H1, 65 Hz in L1), and barely passes the null-stream veto applied for this analysis. The SNR is 9.6 in H1 and 12.8 in L1. There is no signal in H2. L1 is very noisy, and loses lock 20 minutes later. The event is not recorded by *Coherent Waveburst*.

Coherent Waveburst

Coherent Waveburst is better able than *kleineWelle*+*CorrPower* to resolve GW-like signals from noise transients, and is used along with Q/Ω Pipeline for the offline second-year all-sky burst search below 2048 Hz. *Coherent Waveburst* events are

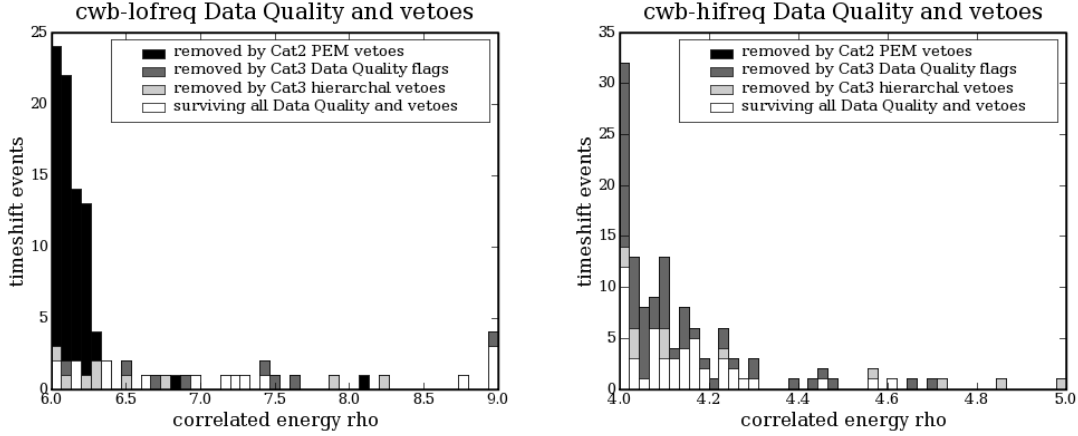


Figure 7-10: Sample background from the Coherent Waveburst pipeline from 1000 time-shifts of H1H2L1 data in the second calendar year of S5. On the left are events from 64–200 Hz, and on the right are events from 200–2048 Hz. Shaded events are removed by various stages of data quality and vetoes. The candidate event at ~ 110 Hz has $\rho = 6.6$.

divided into two frequency bands: 64–200 Hz and 200–2048 Hz. The events in each band are ranked by their effective significance, ρ , and are subject to the detection thresholds of $\rho > 6.0$ for low frequency, and $\rho > 4.2$ for high frequency bands. This gives a combined false-alarm probability of 6% for Coherent Waveburst on second-year data. Category 2 Data Quality and vetoes (chapter 6) are applied before checking for detection candidates, while category 3 data quality and vetoes are used to create the clean data set for an upper limit. The number of low-frequency time-shift events after category 2 cuts stronger than the equinox event with $\rho \geq 6.60$ is 18 in 1000 time-shifts. After category 3 cuts the number is reduced to 10. Some amount of live-time is lost for large time-shifts, and accounting for this loss the probability of observing a low-frequency background event in the detection set of the second year of S5 with $\rho \geq 6.60$ is 2% or once in every 28 years of S5 second-year live-time. After category 3 cuts this is reduced to a probability 1% or once in every 43 years.

The two frequency bands for Coherent Waveburst result from the large population of noise transients at low frequency. Figure 7-11 includes a histogram showing the frequency distribution of time-shift events at various thresholds in ρ . The counts are displayed on a logarithmic scale, and a low threshold they are largely dominated by

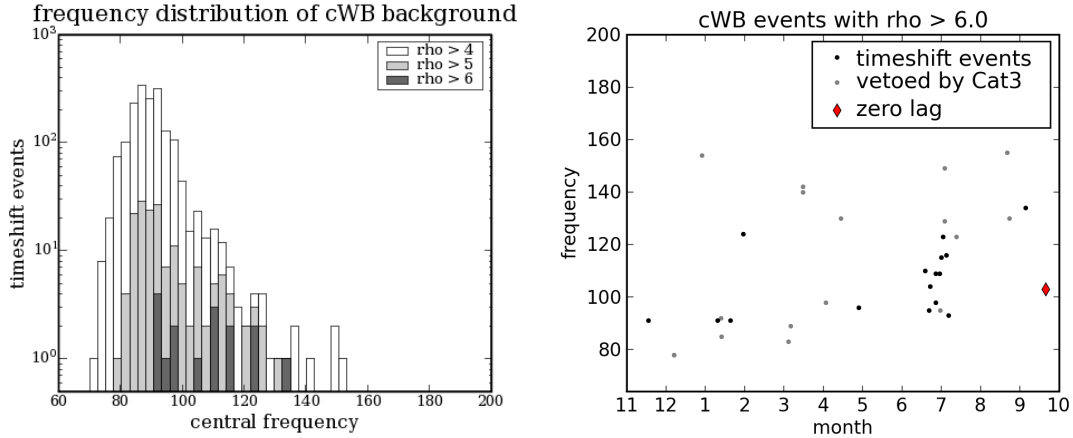


Figure 7-11: Frequency and time-frequency distribution of sample background from the Coherent Waveburst pipeline from 1000 time-shifts of H1H2L1 data in the second calendar year of S5. The frequency histogram on the left shows a broad excess of low frequency background between 80–130 Hz. The right plot shows the individual time-shift events between Nov 2006 and Oct 2007, with the candidate event in red near the end of the run.

events between 85–95 Hz, which come primarily from the first half of the second-year. However the strongest events don't show the same narrow distribution. This can be seen in detail in the second scatter plot of frequency versus time for the time-shift events above our follow-up threshold of $\rho \geq 6.0$. Much of the background comes from a bad period in July 2007. During this time H2 suffered from poor sensitivity making it a much poorer consistency check for moderate transients in H1.

The detection candidates themselves are subject to a myriad of tests by way of the detection checklist. These tests are not applied to the time-shift background events because of the amount of work involved, and not reflected in the calculation of significance because of the ambiguity about what constitutes a definite veto and the possibility of introducing bias in an un-blinded procedure. Without some consideration of the effects of post-analysis, however, rates such as once every 26 years will be misleading about the expected rate at which detection candidates above a certain strength are presented as possible gravitational waves. In theory the best measurement of rates would come from application of the exact same rules to time-shift and unshifted events.

Another important factor when considering significance is the total number of experiments. With the division of Coherent Waveburst into two frequency bands, the above-2048 Hz high-frequency burst search, and separate first and second year analyses, the burst group has at least six independent all-sky searches for S5. In addition, searches such as Q/ Ω Pipeline and the search over LIGO/Virgo VSR1 data cover the same live-time and are not independent from Coherent Waveburst LIGO-only searches in non-trivial ways. The burst group has brainstormed extensively on ways to combine multiple searches, but currently the equinox event can only be directly compared with expected background distributions in its own restricted 64–200 Hz search over S5 year 2. The 2% probability of observing a stronger event in this set should thus be taken in appropriate context.

Q/ Ω Pipeline

The equinox event did not pass the final cut on H1H2 correlated energy for the Q/ Ω Pipeline second year analysis. Figure 7-12 shows event scatter plots of H1H2 correlated vs. coherent energy, and H1H2 correlated vs. L1 energy for time-shift, unshifted events, and simulated signals, as well as the equinox event and semi-blind inspiral hardware injection. The H1H2 plot on the left contains H1H2 events during all time when H1 and H2 are in science mode with category 3 cuts applied, and the cut on H1H2 correlated energy is placed to maintain zero remaining H1H2 events in 10 time-shifts. The H1H2L1 plot on the right contains the small subset of H1H2 events which are also coincident with a trigger in L1. Here a lower threshold can be set for H1H2 correlated energy to maintain a rate of 15 coincident events in 1000 time-shifts.

We obtain a measure of the significance of the equinox event in the H1H2L1 set by relaxing the cut on H1H2 correlated energy to 10, and ranking all events by the joint probability of observing a greater H1H2 correlated energy and a greater L1 energy. Using this ranking, we expect 2.2 events of greater significance in the H1H2L1 set below 200 Hz, or once per six months. This ranking does not take into account the H1H2 correlated vs. coherent energy consistency or the H1H2 correlated energy

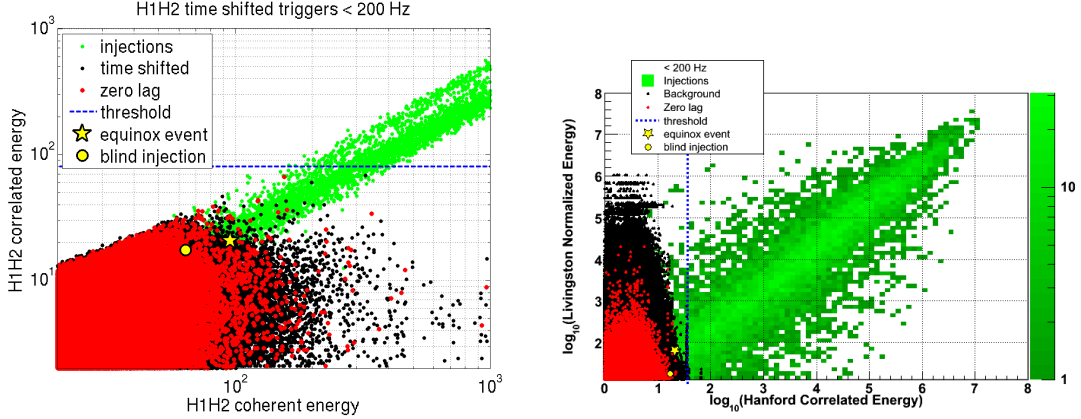


Figure 7-12: Sample background from 10 and 1000 time-shifts (black) and unshifted (red) events from the Q/Ω Pipeline H1H2 and H1H2L1 analyses for the second calendar year of S5. The three major energy measurements are shown: H1H2 correlated energy (may be negative for anti-correlated signals), H1H2 coherent energy (energy of the optimal H1H2 sum), and L1 normalized signal energy. The candidate event (yellow star) is below the final threshold on correlated energy in each search. Also shown is an earlier semi-blind inspiral injection.

vs. L1 energy consistency which are both shown to be quite good for the equinox event itself. The Q/Ω pipeline itself was not tuned to perform well on events below threshold, so it is not surprising that the equinox event, despite its strong waveform consistency, is not well separated from background.

7.3.5 Event reconstruction

Coherent Event Display

The sky-maps in Figure 7-13 show the Coherent Waveburst *sky statistic* for H1H2L1 and H1H2L1V1 detector configurations. The sky statistic is derived from the likelihood statistic and other output from Coherent Waveburst at different sky positions. Although the Virgo interferometer during VSR1 has much worse sensitivity at 100 Hz, it is able to reject regions of the sky where it has a favorable antenna pattern compared to LIGO.

Coherent Event Display [103] is also able to estimate the most likely coherent waveform after choosing the best sky position. Figure 7-14 shows this single best-fit




Figure 7-13: Coherent Waveburst sky statistic for H1H2L1 and H1H2L1V1 networks. The sky statistic is derived from the likelihood statistic and other output from Coherent Waveburst at different sky positions.

detector	GPS time	h_{rss} [$\text{Hz}^{-1/2}$]	SNR
H1	874465554.715	3.16×10^{-22}	13.6
H2	874465554.715	3.16×10^{-22}	5.3
L1	874465554.710	3.04×10^{-22}	11.9

Table 7.4: Parameters of the Coherent Waveburst best-fit gravitational wave when projected onto the interferometer geometry of each detector in local noise.

gravitational wave projected onto the arms of each detector, so that the H1 and H2 waveforms are the same by construction. From left to right are the calibrated strain waveforms for H1, H2, and L1. *Due to a plotting bug, the amplitude of each waveform is twice what it should be.*

Also informative is the estimated waveform shown on top of the calibrated detector output (Figure 7-15). The time series have been whitened with the same filter so that the noise is not dominated by lines. The H1 and L1 waveforms fit very well to the $h(t)$ data. H2 shows some amount of excess noise.

Using the estimated waveform, the per-detector signal properties for Coherent Waveburst are listed in Table 7.4. Other properties of the best-fit gravitational wave are listed in Table 7.5. The addition of Virgo gives a slightly different estimate for galactic coordinates $(\alpha, \delta) = (38.9^\circ, 30.5^\circ)$.




Figure 7-14: Coherent Waveburst best-fit strain waveform projected onto the detector geometry of each interferometer. Because this represents the projection of a coherent gravitational wave, the reconstructed signal for the co-aligned H2 detector is the same as for the H1 detector.

Omega Pipeline

Omega Pipeline’s Bayesian statistic [86] can calculate the probability of observing the data given the presence of an elliptically-polarized sine-Gaussian gravitational wave at a given sky location marginalized over random arrival time, amplitude, frequency, Q , inclination and phase. Figure 7-16 shows arbitrarily normalized H1L1V1 probability distributions on the sky in linear and log scale. Omega Pipeline puts the most probable source location at $(\alpha, \delta) = (87.6^\circ, -54.9^\circ)$, which puts it in the southern hemisphere.

The Bayesian method has yet to produce a detection statistic for the equinox event,

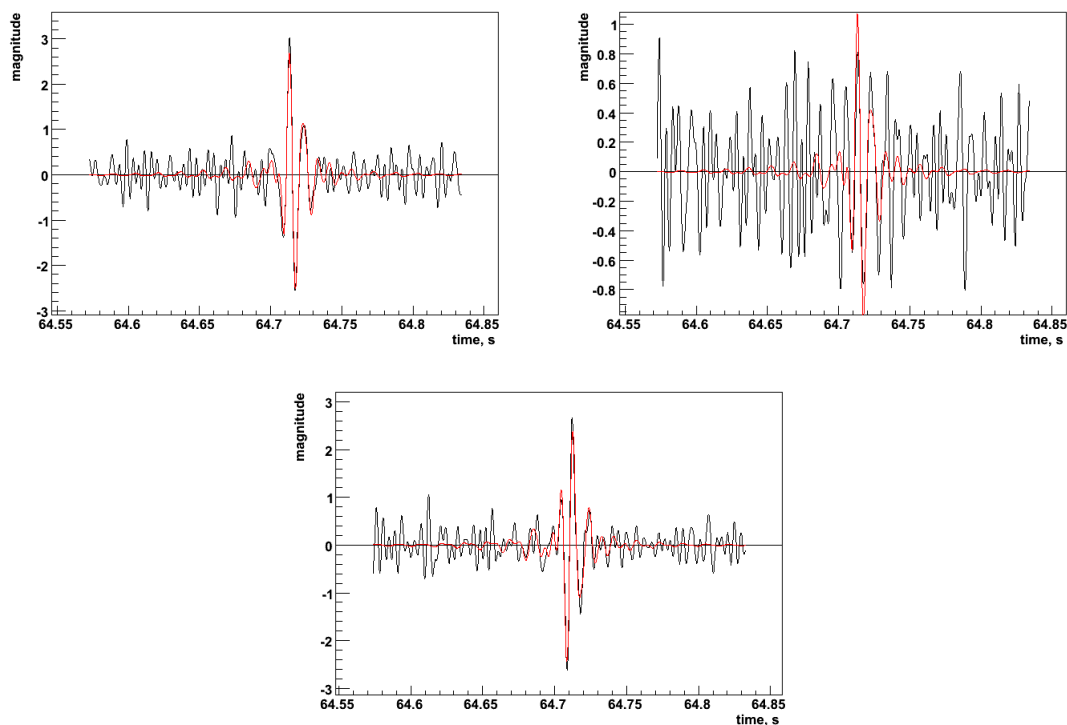


Figure 7-15: Coherent Waveburst best-fit (whitened) waveform plotted on top of the whitened data streams for H1 (top-left), H2 (top-right), and L1 (bottom). The H1 and L1 data show a very good fit.

or a statement about the physical consistency of the signal at the most probable sky location.

7.3.6 Astrophysical analysis

Nearby electromagnetic events

There are a number of public listings of electromagnetic transients available. Table 7.6 lists transients from the Gamma-ray bursts Coordinates Network (GCN) which have

duration	frequency	α	δ	ρ
0.039 s	102.54 Hz	43.9°	24.5°	6.60

Table 7.5: Parameters of the Coherent Waveburst best-fit gravitational wave in Earth-centered coordinates.




Figure 7-16: Ω Pipeline’s Bayesian sky probability distribution for the candidate event. The right plot shows the log probability in order to better see the detailed structure.

GCN	name	time	instrument	type	α°	δ°
6807	GRB 070920	07/09/20 04:00:13	Swift	long	100.968	72.250
6811	GRB 070920B	07/09/21 13:22:10	Swift	long	0.127	-34.844
6821	GRB 070923	07/09/23 19:15:23	Swift	short-hard	184.623	-38.294
6822	SGR 1806-20	07/09/21 11:45:35	Konus-Wind	SGR burst	272.164	-20.411
”	”	07/09/23 12:46:15	”	”	”	”
6823	GRB 070925	07/09/25 17:26:59	Integral	long	253.218	-22.0355

Table 7.6: Electromagnetic events from the GCN catalog around 2007-09-22

been reported within a few days of the equinox event. None of the publicly reported transients are plausible counterparts to the equinox event, mostly being from the wrong sky location. International Astronomical Union (IAU) circulars were also checked for the presence of nearby supernovae or other transients.

BH/BH merger interpretation

Fitting the waveform estimated by Coherent Waveburst to a ring-down gives ring-down parameters $f = 97$ Hz, $Q = 2.7$ for L1. Using the following relations for a perturbed black hole,

$$f = \frac{c^3}{2\pi GM} [1 - 0.63(1 - \alpha)^{0.3}] \quad (7.1)$$

$$Q = 2(1 - \alpha)^{-0.45}, \quad (7.2)$$




Figure 7-17: The Hilbert-Huang transform of the candidate event indicates a decrease in frequency from about 160 to 60 Hz in just a couple cycles.

gives total mass $M = 140 M_{\odot}$ and spin $\alpha = 0.5$.

One problem with the BH/BH merger interpretation is that the frequency seems to decrease with time from about 160 to 60 Hz in just a couple cycles (Figure 7-17). This indication is from the Hilbert-Huang transform, an adaptive data analysis strategy which provides a detailed description of the instantaneous frequency of the waveform [105]. In the following plots, time-frequency maps derived from the HHT are shown for H1 and L1. These HHT-derived maps have a frequency accuracy related to the instantaneous power of the waveform. The uncertainty in frequency over the central 5 ms of the waveform, where the power is high, is ± 5 Hz, while at the edges, where the power is lower, the uncertainty is ± 30 Hz. A trend of decreasing frequency in time is clearly seen. For a compact object coalescence, we would expect a chirp waveform with increasing frequency.

We can compare the LIGO data with recent numerical simulations of non-spinning equal-mass binary black-hole mergers from the GSFC group [106]. The simulated waveforms are taken with total mass $M = 140 M_{\odot}$. At 70 Mpc and optimal orientation, the strain at Earth is shown in Figure 7-18. To match the peak strain observed in the 4km instruments, we reduce the strain by a factor of six and inject into LIGO data one second before the equinox event. We make equivalent plots to Figure 7-1 by plotting the whitened LIGO data plus simulated BH merger signal at 60–2048 Hz




Figure 7-18: 70/70 M_{\odot} equal mass black hole merger waveform at 70 Mpc from the GSFC group [106].




Figure 7-19: Time series of 70/70 M_{\odot} equal mass black hole merger simulation in LIGO noise with arbitrary scaling. The low frequency inspiral and early merger is strongly suppressed as it cannot be seen above the seismic noise wall.

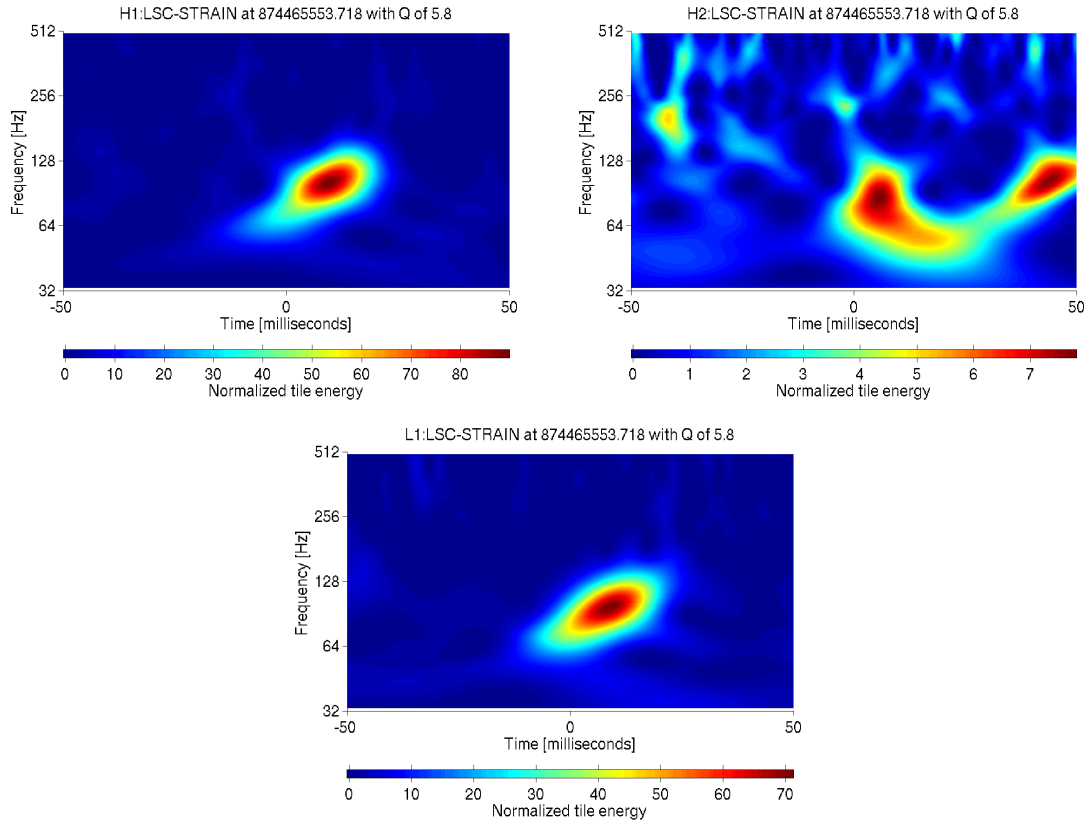


Figure 7-20: Spectrogram of equal mass black hole merger simulation in LIGO noise. and 60–140 Hz.

We see that the initial low frequency signal is quickly swamped by the noise curve, and only the low- Q merger and ring-down show. A Q -scan (Figure 7-20) shows the best match to be at slightly higher Q than for the actual equinox event. We also see evidence of a small chirp. Curiously there a weak 100 Hz transient in H2 which occurs less than 50 ms following the merger, similar to the additional noise preceding the coherent component of the equinox event’s H2 waveform.

Source population study

The sky-maps in Figure 7-21 show blue-light luminosity L10 density per solid angle in the sky over a certain distance range. The sources are taken from the CBC catalog [107], and smoothed over a resolution of 10 degrees. The first plot shows everything within 1 Mpc where the major feature is the Andromeda galaxy M31. We also see two

other major nearby objects, M33 and LMC, intersecting the time-delay sky ring. The fifth plot shows blue-light luminosity between 50 and 80 Mpc where the major feature is the large Perseus-Pisces supercluster (about $10^{16} M_{\odot}$). An optimally oriented 75/75 M_{\odot} non-spinning BH/BH merger at 70 Mpc would give an h_{peak} at the earth of about 2.5×10^{-20} according to numerical simulations [106]. This optimally oriented and zero inclination waveform is six times larger than the best fit waveform projected onto the arms at either site as estimated by Coherent Waveburst.

7.3.7 Other events in S5

Outliers in time-shift analysis

Most of the loud time-shift events in the Coherent Waveburst analysis show similar characteristics of the equinox event, namely low frequency of 70–120 Hz and low- Q . This expectation of similar events in the background does not change the measured significance of observing the event at zero time-shift, but the event would arguably be more significant if it was clearly separate from the dominant background population.

Outliers at zero time-shift

There are no other outliers at zero time-shift in the Coherent Waveburst analysis. The remaining events are all well below threshold.

Similar transients

Low-frequency transients are common in the LIGO instruments. On the day of the event, there were three transients in H1 at the same frequency and Q which were louder than the candidate event, and 30 stronger L1 transients at the same frequency and Q .

More information about the presence of similar single detector transients in the hour of the event can be found in section 7.3.3.

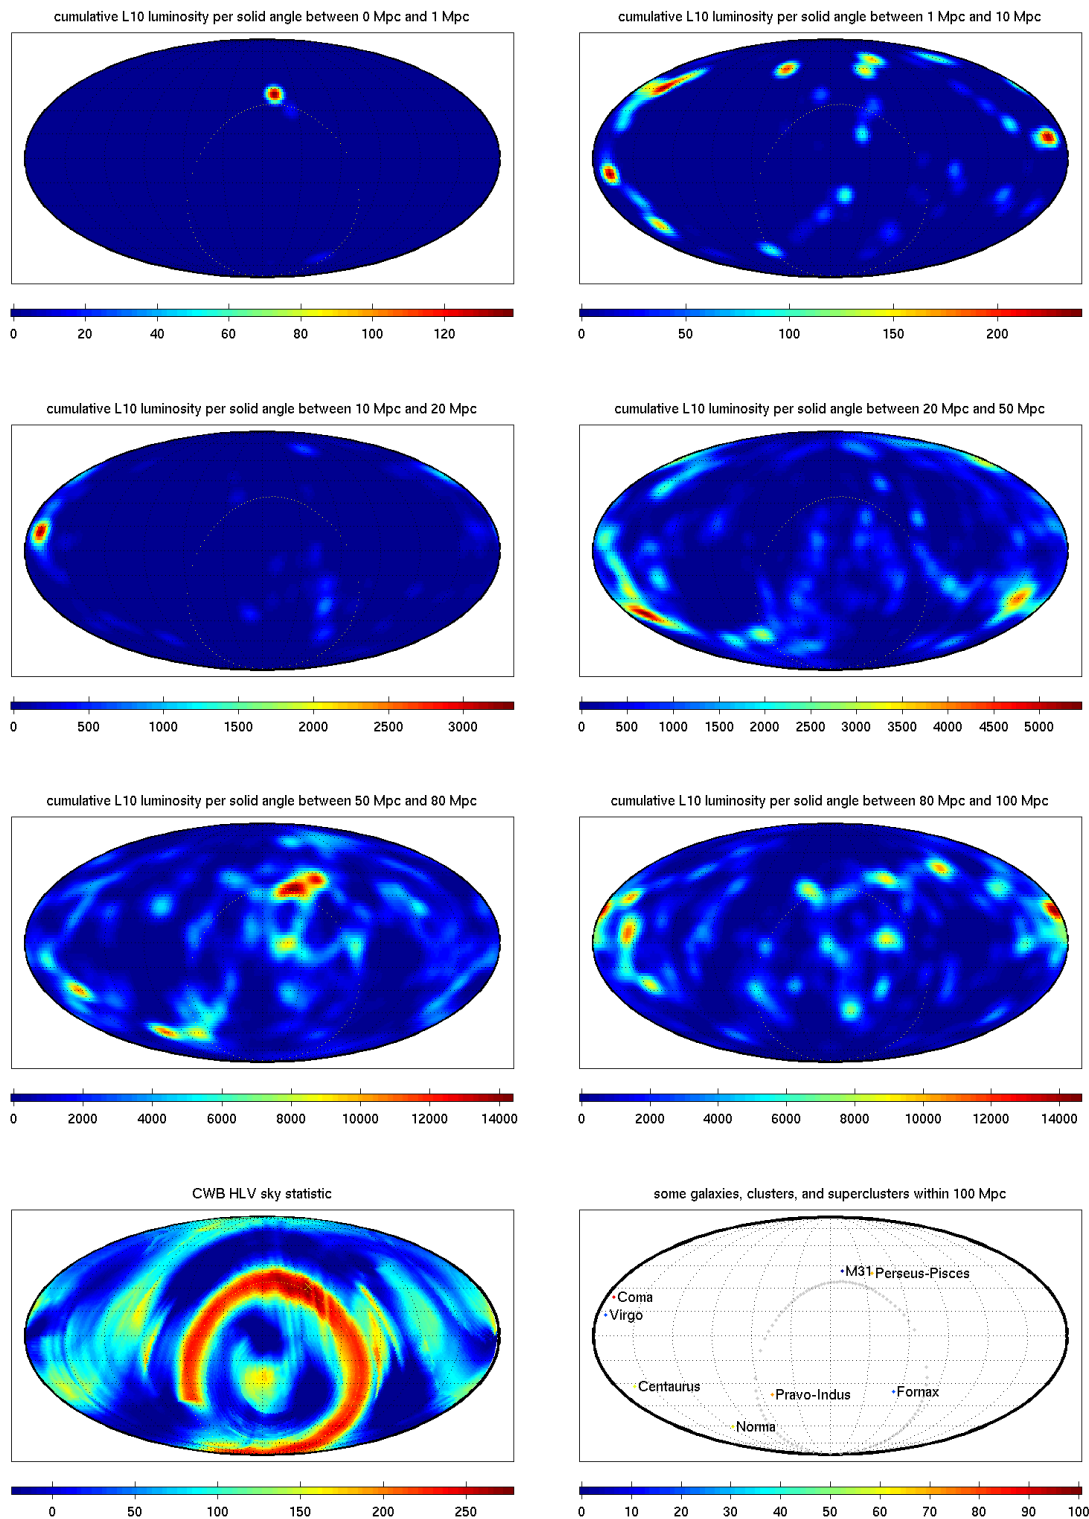


Figure 7-21: L10 blue light luminosity density per solid angle over fixed distance ranges out to 100 Mpc using the catalog of Kopparapu *et al* [107]. Also shown is a sample of well known galaxies and galaxy clusters.




Figure 7-22: Sine-Gaussian $f=100$ Hz $Q=9$ hardware injection.




Figure 7-23: Gaussian 3 ms hardware injection.

Similar hardware injections

The most similar hardware injection in S5 is a sine-Gaussian waveform with $f = 100$ Hz, $Q = 9$, $h_{\text{RSS}} = 6.4 \times 10^{-21}$ Hz $^{-1/2}$ (Figure 7-22). There is such an injection at GPS 874465969.5 in the set six minutes following the equinox event. We also have a 3 ms Gaussian waveform injected at GPS 875253573.5 with $h_{\text{RSS}} = 4.0 \times 10^{-20}$ Hz $^{-1/2}$ (Figure 7-23). Because of the low frequency noise, the whitened signal resembles a sine-Gaussian waveform. Unfortunately we do not have weaker hardware injections for these parameters.

7.3.8 Conclusion of the follow-up procedure

Summary of findings prior to blind injection revelation

As soon as the equinox event was observed, it was clear that the low frequency, low- Q morphology of both H1 and L1 signals made them very similar to the most common transients observed. Also apparent is that the 1-2% chance of a stronger

event originating from background is marginal given the several independent burst searches we run for S5.

The follow-up analysis of the event has brought in a large amount of additional information to consider when evaluating our believability in this event as a true gravitational wave. In favor of the event is,

- We looked very hard and did not find any evidence of an instrumental cause or noise transient
- The event did not happen during noisy times in the run which represent a large contribution to our background
- The impressive consistency between H1L1 and H1H2 (minus the extra content at the beginning of the H2 signal) is very unusual
- The event has passed a very extensive checklist. It is likely that on deep inspection the event is a more convincing gravitational-wave candidate than events in our expected background at the same ρ threshold.

Evidence found in post analysis which decrease our confidence in the event is,

- The dominance of low frequency, low- Q events in our background suggests that an optimal search might rank such events as less significant, allowing events outside of this noisy region a higher ranking because they buy more sensitivity at less false alarm rate. Thus an objectively better search may result in a lower significance for this event. Generally we assume that a better search will provide increased significance for a real event.
- The fact that the frequency and Q of the event is very consistent with our background in morphology leaves open the possibility that observation of the event is consistent with background and we have just made a mistake somewhere and underestimated the rate.
- The marked increase in weak H1 transients immediately around the event is very suspicious. Though there is no clear connection other than time between

these transients and the event, the sudden and coincident excess transient rate makes us wonder if H1 should be trusted during that short interval as it implies an unexplained change in the instrument state as well as an elevated background rate.

These considerations are beyond the scope of our blind analysis, and are highly subjective. However that does not mean that subjective post-analysis will not yield a better separation of background from true gravitational waves. In the future, we of course want to incorporate any useful information into the detection statistic itself. Many could find their way into a likelihood statistic, which by construction provides the best foreground-background separation given the parameters we measure.

Ultimately while very interesting, the event does not qualify as a gold-plated detection candidate. A claim that the event is a true gravitational wave would be subject to too high a chance of misidentification.

Closeout of blind injection challenge for S5

With the completed analysis and follow-up in place, the blind injection challenge was finally closed in March of 2010. It was then that the analysis teams learned that the burst outlier was in fact a hardware injection, leaving no remaining true outliers in the entire S5/VSR1 burst search. The blind injection was a coherent low frequency broadband signal injected into the three LIGO detectors resulting in match-filter signal-to-noise ratios of 11.92 (H1), 6.15 (H2), and 10.31 (L1), close to those measured by the burst pipelines (Table 7.1 and Table 7.4). Figure 7-24 and Figure 7-25 compares the reconstructed waveform and sky location provided by the coherent methods.

The result of the blind injection challenge highlighted the problem of excess low frequency transient background in the LIGO instruments. Even at a signal-to-noise greater than ten in both 4 km instruments, we were not able to resolve the signal sufficiently from background in order to confidently make a detection. Such a signal amplitude in Gaussian noise would have vanishingly small background even without a coincidence requirement (Figure 5-1).

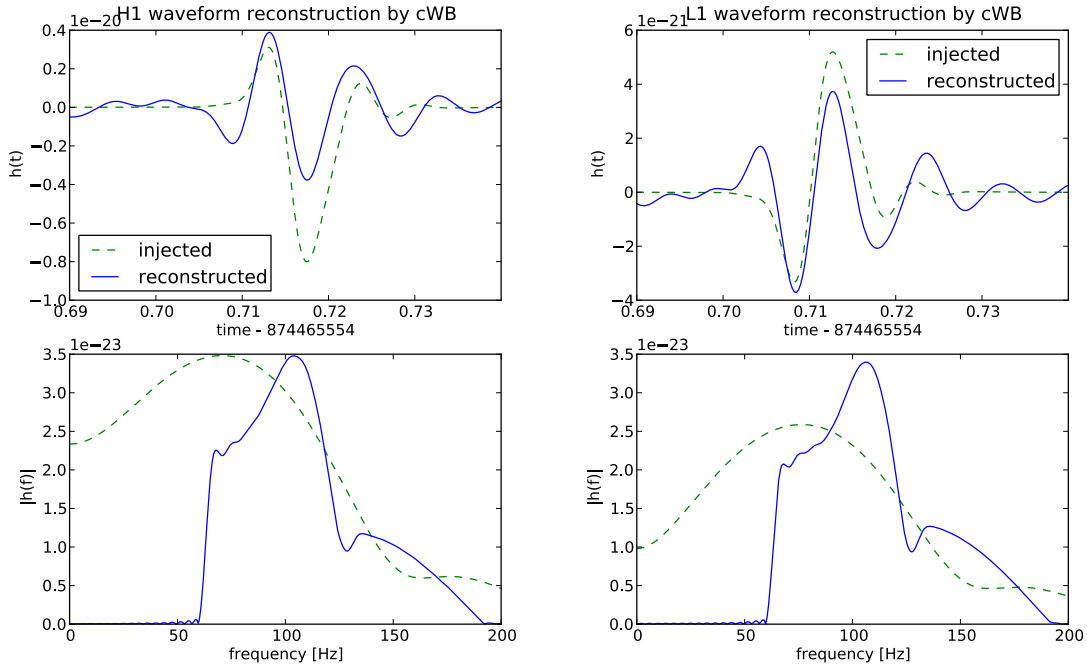


Figure 7-24: Comparison of injected and reconstructed strain waveform projections for the S5 blind burst injection. The simulated burst gravitational wave was injected coherently into the LIGO interferometers H1, H2, and L1 at the end of S5 as part of the blind injection challenge. Coherent Waveburst identified and reconstructed the waveforms in blue using H1, H2, and L1 data. Because H1 and H2 share alignment, they have identical waveforms thus H2 is not shown. Coherent Waveburst is able to reproduce the main part of the wave in the LIGO sensitive band (64-2048 Hz).

The data quality and auxiliary channel-derived vetoes applied to the burst search as well as the signal-based vetoes which are part of the coherent multi-detector analysis go a long way to reducing the background to suitable levels for maintaining sensitivity in an upper limit analysis. However, background levels need to be further reduced by at least a couple orders of magnitude to obtain event significances (for moderate gravitational wave amplitudes) of 10^{-4} to 10^{-5} , approaching a more reasonable standard for detection. Doing so will require either another instrument, such as tight coincidence with an electromagnetic counterpart, or several detectors operating at comparable sensitivity in order to produce a robust null-stream veto. For S5/VSR1, the coherent null-stream vetoes, which rely on testing the residual after




Figure 7-25: Comparison of injection and reconstructed sky location for the S5 blind burst injection. Since the gravitational wave is seen only at two sites, it is largely localized to a ring on the sky. The true location is represented as a white diamond at $(\alpha, \delta) = (80.47, -73.06)$. Coherent Waveburst is run on H1H2L1 data and gives a maximum likelihood sky location at $(\alpha, \delta) = (43.9, 24.5)$. The Omega Bayesian statistic was not part of the S5 blind analysis but ran on the event after it was identified. On H1L1V1 data, the statistic gives a maximum likelihood sky position of $(\alpha, \delta) = (87.6, -54.9)$ though this is not quite applicable because the injection was not present in Virgo data.

removing the best-fit gravitational wave from the data, suffered from the poorer low frequency performance of the Virgo instrument as well as the lower sensitivity of H2 and correlated H1-H2 environmental noise. Therefore a coincident set of moderately loud low-frequency noise transients in H1 and L1 could be fit to the two degrees of freedom in an arbitrary gravitational wave without incurring a null stream penalty from the two less sensitive detectors. Since then, Virgo has made fast progress catching up to the LIGO low-frequency sensitivity, and Advanced LIGO should see H2 upgraded to a 4 km baseline and running at full sensitivity.

The follow-up procedure also highlighted the difficulty for the burst analysis to assign the global significance of an outlier in terms of the probability of detecting a better event from background over the entire search. The S5/VSR1 burst search was particularly aggressive over previous searches in maximizing coverage to reduce the chance of missing a rare event. The result was a large number of separate networks, epochs, frequency bands, and analysis methods each of which had its own background distribution unique to the particular configuration. The significance of an outlier was

only defined in context of the limited configuration or at the global upper limit cuts (which were set at too high a false-alarm-probability to be useful for detection). In the following chapter, we introduce procedures for establishing a unified ranking of events across multiple search configurations to address this problem.

Chapter 8

Unified ranking for gravitational-wave events

A search method has the task of distinguishing signal events from background. In the simplest of cases, this includes assigning each event to one of those categories. Events classified as signal are those which pass the selection cuts, while events classified as background are everything else. If N events from the data are classified as signal and the expected contribution from true background to this measurement (false-alarm expectation μ_{BG}) is known, the significance of this measurement can usually be expressed by assuming Poisson statistics on the event counts (Appendix A),

$$Z_{(N, \mu_{\text{BG}})} = -\ln [\text{P}_{\text{Pois}}(k \geq N | \mu = \mu_{\text{BG}})]. \quad (8.1)$$

The significance of the measurement is the negative logarithm of the false-alarm probability of measuring the observed number of events or more from background alone.

The simple classification of events into signal and background can also be used to set upper limits or confidence intervals on a true rate of signals. For the S5 all-sky burst analysis, classical frequentist upper limits were established for the true rate of events classified as signal (above threshold) under the assumption of Poisson statistics. The frequentist method attempts to set a confidence interval which contains the true

value at least $1 - \alpha$ fraction of the time where $1 - \alpha$ is the confidence level. Assuming Poisson statistics and an observation of N events, this can be achieved by setting an upper limit at the mean value μ_{UL} which solves,

$$\alpha = P_{\text{Poiiss}}(x \leq N | \mu = \mu_{\text{UL}}). \quad (8.2)$$

For an observation $N = 0$, this gives an upper limit of $\mu \leq 2.303$ at 90% confidence, meaning that for any true $\mu > 2.303$, we would have measured zero (or fewer) events less than 10% of the time. Thus the upper limit derived this way contains the true μ in at least 90% of experiments. The frequentist upper limit is a limit on events classified as signal, which includes in general contributions from foreground and background sources. While background may be known and folded into the calculation, the limit can be used as a conservative limit on the number of foreground events only (setting the background contribution to zero).

A unified ranking of events implies a unique and global ordering of events from background-like to signal-like. A ranking can be based on signal-to-noise ratio, for example, in the case of a matched-filter or excess power search (section 4.2). It is also the output of a number of multivariate statistical classifiers. It is easily reduced to the classification case by setting a single cut on the ranking variable r .

The ranking is more flexible than a simple categorization. Rate upper limits can be established from the loudest event statistic of Brady, Creighton, and Wiseman [108]. The loudest event statistic is based on the sensitivity of the search at a threshold set at the rank of the loudest detected event. It therefore removes the need to set a threshold by hand, which can be a daunting process in uncertain background. Such is the case for a search on H1H2 data due to correlated environmental noise affecting the validity of the time-shift background estimate [74].

The loudest event statistic sets the threshold of the analysis at the value of the loudest observed event r_{max} . The upper limit is a statement about the number of expected events above r_{max} and for the frequentist construction of the statistic gives

the same limit as the classical frequentist upper limit for an observation $N = 0$,

$$\alpha = P_{\text{Pois}}(k = 0 | \mu = \mu_{\text{UL}}). \quad (8.3)$$

This is because $1 - \alpha$ is the probability that the loudest event occurs above the threshold value that corresponds to a true $\mu = \mu_{\text{UL}}$. At 90% confidence, the upper limit is 2.303 events with $r \geq r_{\text{max}}$. In 90% of experiments, the loudest event will have higher r than that which would have given a true expectation value of 2.303. In those cases, 2.303 is greater than the actual expected number of events with $r \geq r_{\text{max}}$.

In the event of a single outlier (which is the most likely case given the rarity of strong gravitational waves and shape of the background amplitude distribution), the local significance of the event can be derived from the probability of observing one or more events from background with higher rank using 8.1 with $N = 1$,

$$Z_{(1,\mu)} = -\ln [1 - P_{\text{Pois}}(k = 0 | \mu)] = -\ln [1 - e^{-\mu}]. \quad (8.4)$$

Here μ is the integrated background distribution from the rank of the event to a rank of infinity, and we have used the fact that the Poisson distribution is a probability distribution. The advantage of a single classification is that progressively better events can be assigned correspondingly higher significance.

8.1 Trial factor issues in S5/VSR1

Evaluation of the global significance of an event was an issue in the S5/VSR1 burst search as no universal ranking of events was defined as part of the blind analysis. Instead, the analysis focused on classifying events as signal or background by setting cuts dependent on a particular detector network, frequency range, analysis method, and data quality and veto configuration [8]. Events were only compared to one another in the sense that they either passed or did not pass the predetermined analysis cuts. Each search method generally did have a means to rank events within a single configuration. The ranking was based on a single-to-noise like quantity and applied as

the final fixed cut for that configuration. The significance of an event is thus defined relative to background from the same configuration. To evaluate a global significance, one can evaluate the probability of getting one or more events from any of the M configurations with local background expectation less than μ ,

$$Z_{1,\text{Global}} = -\ln \left[1 - \prod_{i=1}^M e^{-\mu} \right]. \quad (8.5)$$

This is unsurprisingly just the significance of a single event from a background expectation of $\mu_{\text{Global}} = M\mu$. For very small probabilities, the global false-alarm probability is the individual false-alarm probability multiplied by the M independent trials.

Table 8.1 lists the configurations used for the S5/VSR1 burst search [8]. The total live-time of 5.1 years is greater than the observation time because of overlap between configurations. Exclusive networks are times when only the listed detectors are operating. Inclusive networks include time when the unlisted detectors may be operating, but only the detectors listed are part of the analysis. The configurations are distinguished by their search method, frequency range, and detector network. Each configuration has a different set of hand-chosen cuts to classify events for the upper limit analysis. In addition, a distinction is made between events flagged by category 3 data quality and vetoes and those not flagged. Flagged events are cut from the upper limit set but are still considered for detection. In total this represents 40 different “boxes” for the search, or 80 when counting the category 3 flag distinction.

The large number of configurations is a problem for the S5/VSR1 search as a naive application of the trials factor would severely diminish the significance of an event. In this straightforward approach, events which may come from a particularly insensitive network configuration or a configuration with low live-time will rank alongside other events at the same local false-alarm probability. These configurations then contribute to the global background without adding much sensitivity. The cuts chosen by hand at the global false-alarm probability used for the upper limit classification imply a certain fixed rank of for events which live at the position of the cut. At those parameters, they are often more reasonable than the false-alarm probability ranking. Insensitive

pipeline	network	days	notes
cWB 64–200 Hz	H1H2L1V1	68.2	all live-times for cWB 64–2048 Hz are pre Cat2 flags, except for H1H2L1
	H1H2L1	191.6	including V1 time
	H1L1V1	3.9	excluding H2 time
	H1H2V1	14.5	excluding L1 time
	H1H2	56.9	excluding L1 time, including V1 time
	H1L1	10.1	excluding H2 time, including V1 time
	H2L1	3.5	excluding H1 time, including V1 time
cWB 200–2048 Hz	all (7)	348.7	same as below 200 Hz
Q/ Ω 64–200 Hz	H1H2L1	194.5	including V1 time
	H1H2V1	85.9	including L1 time
	L1V1	83.1	including H1H2 time
	H1H2	34.9	excluding L1 and V1 time. Q/ Ω was run on inclusive networks, but it was decided to remove the L1 and V1 time from the H1H2 network for this calculation.
Q/ Ω 200–2048 Hz	all (4)	398.4	same as below 200 Hz
EGC 300–5000 Hz	H1H2L1V1	69.3	all networks are exclusive
	H1H2L1	16.6	
	H1H2V1	16.3	
	H1L1V1	4.6	
	H2L1V1	1.5	
	H1H2	4.8	
	H1L1	1.2	
	H1V1	1.9	
	H2L1	0.2	
	H2V1	9.8	
	L1V1	6.4	
	cWB 1280–6000 Hz	H1H2L1V1	70
H1H2L1		128	
H1H2V1		15	
H1L1V1		4	
H2L1V1		1	
H1H2		35	
H1L1		6	
H2L1		3	
H1V1		1	
H2V1		0	
L1V1		6	
total	all (40)	1881.4	5.1 years

Table 8.1: List of configurations and live-times used for the S5/VSR1 burst search. See text for details.

networks are subject to higher thresholds, and a threshold on signal-to-noise ratio is not sensitive to divisions in live-time.

We can list a few guiding principles to keep in mind when constructing an improved global ranking of events to better handle the trials factor,

- The ranking should be robust against divisions in live-time
- The ranking should be robust against divisions in frequency or other signal parameters
- The ranking should show consistency with cuts chosen by hand, such as the upper limit cuts. That is, the cuts should be able to be described as a single cut on rank.

8.2 Inverse false-alarm rate ranking

Inverse false-alarm rate (IFAR) provides an extremely simple way to rank events that is by construction robust against divisions in live-time. In this ranking, events are ranked by the rate at which louder events arise from background in the search configuration where the event was found. It provides a live-time-scaled version of the false-alarm probability ranking. The IFAR ranking was used to rank events for the S5 inspiral search [24, 25] and was also used to estimate the significance of the burst outlier in the S5/VSR1 burst analysis followup [8].

Rankings such as IFAR that are based on predetermined monotonic functions of local false-alarm probability have a distinct practical advantage. For such a ranking, the global background at a given value of the rank variable is free of any measurement error by construction assuming only that local false-alarm probability could in principle be measured to arbitrary precision. The error instead is contained in the measurement of the rank variable itself for a particular event. The measurement, in turn, is based only on local data. To reduce error in the measurement of global significance for a particular event, one must only better resolve the background distribution for the particular configuration where the event was observed, and not all

the background distributions of all configurations. This is very different from a global ranking which is instead based on some function of signal-to-noise ratio, for example. In gravitational-wave transient searches, the background is estimated by running the full analysis method over time-shifted data multiple times with different time-shifts (subsection 4.3.3). The most sensitive search methods are computationally intensive so there is great benefit from running additional time-shifts over just a subset of data.

The S5/VSR1 burst outlier provides an application of this procedure. The event was detected in the H1H2L1 Coherent Waveburst 64–200 Hz analysis after application of category 2 and 3 data quality and vetoes. In 1000 time-shifts of L1 against H1H2 data, there were 10 time-shift background events with a larger local ranking statistic than the observed event, corresponding to a false-alarm probability of 1% and a false-alarm rate of 1 per 43 years. Rather than trying to gather background events from all 40 configurations using the same number of time-shifts in order to make a global background distribution (most configurations were run with 100), we can instead multiply the false-alarm rate by the total accumulated live-time of 5.1 years (Table 8.1). This gives us a total background expectation of $5.1/43 = 0.12$, or a corresponding false-alarm probability of observing one or more events from background with a false-alarm rate of less than $1/43 \text{ yr}^{-1}$ of 11% (8.4).

The calculation relies on the assumption that all configurations could in principle resolve background to a rate of at least $1/43 \text{ yr}^{-1}$ given a sufficiently loud event observed there. This might not be possible for configurations with a small live-time as there are limits to the total number of independent and valid time-shifts that can be performed. If a configuration cannot produce events resolvable to a sufficiently low false-alarm rate, they must assume the highest resolvable false-alarm rate. Instead of estimating this threshold, we consider the global false-alarm probability to be a conservative estimate based on the assumption that all configurations can contribute. Another assumption which factors into the assumption of Poisson statistics is that the background from various configurations are independent. For S5/VSR1, this was briefly verified by hand for certain overlapping data and background correlations were deemed negligible. Correlations become much more important for the case of multiple

detection.

The quality of a ranking based on false-alarm-probability can be improved by folding in information about the relative expected sensitivities of the search configurations (one example is used in [100]). So long as the ranking is based on fixed monotonic functions of false-alarm-probability (such as one derived from the application of configuration-dependent weight factors), the practical advantages of zero error in the expected background at any given rank remains. IFAR is an example of this where the weights are based on live-time. The general limitation to the approach is that the rank is only guaranteed to be good for the region of interest where the functions (or weights) have been calibrated. For a large search with many diverse configurations, a simple ranking based on weighted false-alarm-probability provides a balance between rank quality (separation of signal and background) and computational as well as human cost. If a population of target signals is defined, a choice of weight factors could, for example, be set in order to provide equal detection efficiency divided by expected background contribution across all configurations when constrained to a single chosen global false-alarm-probability of interest.

8.3 Likelihood-ratio ranking

The inverse false-alarm rate ranking is not explicitly independent of divisions in the search parameters other than time. For example, dividing part of a search into two frequency bands will split the background, reducing the false-alarm rate assigned to events and thus increasing their rank with respect to a different portion of the search which did not undergo any divisions. The frequency bands and other divisions in the S5/VSR1 search which define the search configurations in Table 8.1 are not based on maximizing global efficiency with the use of an IFAR ranking, but rather they are based on practical considerations, such as making easier the implementation of frequency dependent cuts or the tuning and application of specific analysis methods.

The likelihood ratio, introduced in Equation 4.34, provides the optimal ranking of independent events parametrized by a set of measurements \mathbf{x} . It is defined as the

ratio of the probability of measuring \mathbf{x} under one hypothesis H_1 to the probability of measuring \mathbf{x} under an alternate hypothesis H_0 ,

$$\Lambda(\mathbf{x}) = \frac{P(\mathbf{x}|H_1)}{P(\mathbf{x}|H_0)}. \quad (8.6)$$

The Neyman-Pearson lemma states that a threshold on likelihood ratio is the most powerful test to distinguish between the two hypotheses. In our case, H_1 is the hypothesis that the measured parameters arise in the presence of a gravitational-wave signal while for H_0 they are a product of background.

The measurements \mathbf{x} can be any product of the data, but the power of the likelihood-ratio test will depend on the ability of the chosen parameters \mathbf{x} to represent essential features in the two populations as well as our ability to measure $P(\mathbf{x})$ accurately. In the extreme case, \mathbf{x} can be the raw data itself which would provide a very powerful test were the probability densities practical to measure. In our case, we have just two parameters: one discrete parameter i indexing the search configuration, and one continuous parameter x which is the local detection statistic for the search. Optimizing the detection statistic is left to the search method, and our task is in combining events from different search configurations in a meaningful way.

With just one continuous parameter, it is practical to calculate the likelihood ratio directly. We use analytic fits to a sample signal population from Monte Carlo simulations and background from time-shift experiments. The ratio of these functional fits generates a mapping from the local detection statistic to the likelihood ratio $\Lambda_i(x)$ which is then used to globally rank the events.

8.3.1 Signal population

The parameters x measured for hypothesis H_1 should reflect those expected from a realistic signal population. The search for gravitational-wave bursts generally targets unknown or poorly modeled signals covering a large parameter space, and tests detection efficiency over a set of ad-hoc waveforms meant to sample the space sufficiently. Simulated waveforms used for the S5/VSR1 burst analysis between 64–2048

Hz include [8],

- Sine-Gaussian waveforms:

$$h_+(t) = h_0 \sin(2\pi f_0 t) \exp\left[-\frac{(2\pi f_0 t)^2}{2Q^2}\right], \quad (8.7)$$

$$h_\times(t) = 0, \quad (8.8)$$

with discretely sampled $f = 70 \dots 2000$ Hz and $Q = 3 \dots 100$.

- Gaussian waveforms:

$$h_+(t) = h_0 \exp\left[-\frac{t^2}{\tau^2}\right], \quad (8.9)$$

$$h_\times(t) = 0, \quad (8.10)$$

with duration $\tau = 0.1, 1.0, 2.5,$ or 4.0 ms.

- Band-limited white noise signals consisting of Gaussian-windowed white noise limited to frequency bands centered about 100, 250, 1000, and 2000 Hz with bandwidths of 10, 100, and 1000 Hz and durations $\tau = 10$ and 100 ms. The signals have uncorrelated $h_+(t)$ and $h_\times(t)$ with equal rms amplitudes.

There are roughly equal numbers of waveforms injected of each type. Discrete amplitudes are spaced uniformly in $\log(h_{\text{rssi}})$ (4.3) from about 2×10^{-22} to 3×10^{-19} to span the sensitive range of the instrument. The sudden cutoff at high amplitude is somewhat arbitrary. For a more realistic population, we use the amplitude distribution of a homogeneous population of standard candles given $h \propto 1/r$:

$$\frac{dN}{dr} \propto r^2 \quad (8.11)$$

$$dN \propto -r^4 dh \quad (8.12)$$

$$dN \propto -h^{-4} dh. \quad (8.13)$$

The original discrete amplitude distribution is uniform in $\log(h)$:

$$dN \propto -h^{-1} dh, \tag{8.14}$$

so to compensate, we assign a weight of $(2 \times 10^{-21}/h_{\text{rSS}})^3$ to each injection when counting. Other weights could, for example, make the population uniform in gravitational-wave energy, but we choose here to keep the remaining parameter density intact.

Each configuration runs over the same Monte Carlo sample of simulated signals which are defined over the entire run and spaced randomly in time, sky position, and polarization angle. Since we are only concerned about relative and not absolute factors in the likelihood ratio for the purposes of ranking events, we use the total weighted number density of detected signals $dN_i(x)/dx$ for each configuration as the numerator in the likelihood ratio. This amounts to the expected number density of detected events at a certain value of the detection statistic x from a homogeneous population of standard candles with some arbitrary total rate.

8.3.2 Background population

The population of sample background is provided by time-shift experiments performed by each analysis method. Because the number of time-shifts analyzed under each configuration may be different, we normalize the total time-shift background counts so that they represent the expected background distribution from a single live-time. This is approximately equal to dividing by the number of time-shifts, though edge effects create small corrections.

8.3.3 Analytic fits

Analytic fits to the signal and background distributions are used in order to get a smooth mapping from configuration-dependent detection statistic x to the likelihood ratio $\Lambda_i(x)$. This is complicated by the fact that the analytic form of each distribution is not known and that the background distribution transitions at some unspecified point from a bulk region of high statistics to the tail of the distribution with low

statistics. With a functional form which will necessarily be overly simplistic, classical fitting techniques applied to the background distribution will tend to use up all the free parameters fitting the well-sampled bulk of the distribution while it is the tails which are important for most analyses.

We are not so interested in getting the correct analytic form of the distribution, fitting the exact parameters, or getting the best statistical match, however. Instead the goal is to have the smallest error on $\Lambda_i(x)$ over a large range of x with particular focus on the tail of the distribution. Moreover, the statistical errors of the signal distribution are no longer Poisson because of the amplitude-dependent weighting applied to each injection. For this reason, we fit to the discrete binned signal and background distribution based on minimizing the least-squared *fractional* difference between fit and measurements. At the tails of the background distribution, defined by the point at which the bin counts fall below some tune-able threshold value (~ 100), the cost function transitions smoothly to negative log-likelihood assuming Poisson errors.

For the signal population which typically maintains good statistics across the range of x , the cost function for fit $\mu[n]$ to measurements $y[n]$ across the N bins is,

$$g_{\text{signal}} = \sum_{n=1}^N \left[\frac{y[n] - \mu[n]}{\max(y[n], \mu[n])} \right]^2 \quad (8.15)$$

Using the maximum of $y[n]$ and $\mu[n]$ is necessary to make the function symmetric and avoid bias in the fit. For the background population which generally consists of a bulk distribution at low x with high statistics, and a tail at high x that transitions to zero, the cost function is,

$$g_{\text{BG}} = \sum_{n=1}^N \begin{cases} \frac{M}{2} \left[\frac{y[n] - \mu[n]}{\max(y[n], \mu[n])} \right]^2 + \frac{1}{2} \ln 2\pi M & y[n] \geq M \\ -\ln [\text{P}_{\text{Pois}}(y[n], \mu[n])] & y[n] < M \end{cases}. \quad (8.16)$$

At M , the cost function transitions from Poisson/Gaussian statistics to fractional χ^2 ,

$$-\ln [\text{P}_{\text{Poiss}}(M + \Delta M, M)] \simeq -\ln \left[\frac{1}{\sqrt{2\pi M}} \exp \left\{ -\frac{(\Delta M)^2}{2M} \right\} \right] \quad (8.17)$$

$$\simeq \frac{M}{2} \left[\frac{\Delta M}{M} \right]^2 + \frac{1}{2} \ln 2\pi M. \quad (8.18)$$

The parameters for the analytic fits are chosen to minimize the cost functions using the Nelder-Mead Simplex algorithm from SciPy [109].

Next we must choose functional forms for the analytic fits. The detection statistics can usually be cast in a form related to the signal amplitude. For Coherent Waveburst, we use the average network signal-to-noise ρ , and for the Q/ Ω Pipeline we use the square root of the correlated H1H2 normalized energy $\sqrt{Z_{\text{H}^+}^{\text{corr}}}$ for the case of networks including H1H2. For the L1V1 network, we use the the square root of the normalized energy at L1, $\sqrt{Z_{\text{L1}}}$. Since the detection statistic x scales with signal amplitude, we expect the simulated population to trace the h^{-4} power-law distribution of a homogeneous population (8.13). To account for method-dependent effects, we allow the power-law index b to vary and for the detection variable x to be shifted away from the origin. The functional form of the binned signal population is therefore fit to a shifted power-law with three free parameters,

$$\mu_{\text{signal}}[n] = a(x[n] + c)^{-b}. \quad (8.19)$$

The background distribution is fit to a bulk distribution plus a power-law tail. The distribution probably contains the superposition of several different tails from different epochs and sources, but due to limited statistics we try to fit as best as possible to one. The single-interferometer noise transient distribution (Figure 5-1) can give an indication of what to expect from a coincident search. However the coherent methods should be able to systematically eliminate very loud events through null-stream vetoes. Therefore the power-law most likely cuts off at very high amplitude. In practice, more time-shifts should be run if there is the need to rank an event outside of the range of the available time-shift background samples. The bulk of the

background is either best fit to an exponential distribution or a second power-law depending on how far the events go into the noise. We fit to both forms: exponential bulk distribution + power-law tail and power-law bulk distribution + power-law tail. The fit with the lower cost function is used. Thus, the background distribution takes one of the following forms,

$$\mu_{\text{BG}}[n] = e^{a-bx[n]} + cx[n]^{-d} \quad \text{or} \quad (8.20)$$

$$\mu_{\text{BG}}[n] = ax[n]^{-b} + cx[n]^{-d}. \quad (8.21)$$

Guesses are used to guide the signal and background distribution fits. The free variables are also all limited to positive values.

8.3.4 Results

Likelihood ratio maps are generated (Figure 8-1 to 8-5) for the subset of searches in Table 8.1 which cover the 64–2048 Hz frequency range and have reliable background estimates (no H1H2 only search due to correlated environmental effects) and appreciable live-time. As in the IFAR calculation, events are separated by the configurations listed in the table as well as by whether or not they fall within times flagged by category 3 data quality and vetoes which cover about 15% of the total live-time (category 2 and below data quality and vetoes are applied unconditionally). Data passing category 3 data quality and vetoes are referred to as “clean” data, while the exclusive set of flagged data is referred to as “flagged.”

In general the likelihood-ratio mappings match our expectations from relative configuration sensitivities and background rates. The 64–200 Hz band has lower spectral noise density than the 200–2048 Hz frequency range (Figure 3-4), so that at fixed signal-to-noise ratio, the higher frequency simulations have larger strain amplitude and thus smaller signal weight factors (Equation 8.13). However, the relative absence of loud expected background in the upper frequency range more than compensates for this effect from reduced sensitivity so that the high frequency events are assigned higher likelihood ratios. We also see that events from flagged times are generally

pipeline (range [Hz])	network config	pipeline threshold				false-alarm [%]		efficiency [#]	
		UL	UL _f	LR	LR _f	UL	LR	UL	LR
cWB (64–200)	H1H2L1V1	4.50	∞	4.41	7.64	1.07	1.60	28.3	30.8
	H1H2L1	6.00	∞	5.92	8.32	1.68	2.10	56.4	62.3
Q/ Ω (64–200)	H1H2L1	6.08	∞	8.83	9.90	1.50	0.00	8.3	1.3
	H1H2V1	4.69	∞	7.32	28.33	0.90	0.00	1.0	0.3
	L1V1	5.66	∞	45.79	∞	0.40	0.00	0.0	0.0
cWB (200–2048)	H1H2L1V1	4.10	∞	4.12	4.21	0.00	0.53	32.8	37.2
	H1H2L1	4.20	∞	4.15	4.44	0.84	2.00	132.7	155.5
Q/ Ω (200–2048)	H1H2L1	3.61	∞	3.55	4.40	1.10	1.80	119.6	135.7
	H1H2V1	3.74	∞	4.01	5.84	0.00	0.00	8.1	6.1
	L1V1	5.48	∞	∞	∞	0.50	0.00	0.0	0.0
total						7.99%	8.03%	239.4	270.5

Table 8.2: Comparison of the hand-chosen thresholds used for the S5/VSR1 upper limit (UL) and the ones derived from setting a likelihood-ratio threshold (LR) which gives the equivalent total false-alarm probability of $\sim 8\%$. The likelihood ratio allocates background away from insensitive configurations toward more sensitive ones. It is also able to automatically choose higher thresholds (LR_f) to use for data flagged by category 3 data quality and vetoes which were originally excluded from the upper limit analysis (UL_f = ∞). The false-alarm percent is represented as 100 times the background expectation for each configuration rather than a probability of non-zero events (Equation 8.4). The efficiency represents the number of signals above threshold from an ad-hoc homogeneous population of sources with an arbitrarily chosen total rate. The total efficiency is lower than the combined efficiency from all configurations because of overlap which, for counting purposes, is checked for simulations but not for sample background events. At this choice of total false-alarm probability, the likelihood-ratio thresholds give a 13% increase in expected signals from the ad-hoc homogeneous population of sources. It also gives an automatic tuning at any choice of total false-alarm probability by defining a universal ranking of events.

assigned lower rank than those from the corresponding clean data configurations.

There is an occasional exception to this rule at high amplitude where there may not be enough live-time in the flagged data to resolve the tail of the background distribution. In practice, more time-shifts would need to be run should an event occur outside the sampled region. A peculiar case arises when the expected background distribution falls more slowly with signal-to-noise ratio than the homogeneous signal population (e.g. the first three plots in Figure 8-2). This causes the likelihood-ratio mapping to turn over for loud events because we expect true loud signals to be very rare. The best way to correct this behavior is to find ways to further clean the background so that its distribution falls off more quickly with signal strength. Alternatively, adding in different signal populations (such as a disk population) may lead to broader distributions, or some ad-hoc preferential weighting applied to well-resolved signals may be used.

The blind injection from the burst S5/VSR1 analysis has a Coherent Waveburst $\rho = 6.60$ in the H1H2L1 “clean” data set after application of category 3 flags and data quality. The likelihood ratio for $\rho = 6.60$ in this configuration is 26.7. At this threshold, the background expectation is 0.0217 giving a false-alarm probability of 2.2%. The IFAR ranking over the same configurations gives a background expectation of 0.08, which is the same as the false-alarm probability set by the original upper limit cuts (Table 8.2). The likelihood-ratio test shows the event to be more significant because it comes from a network and configuration of high sensitivity and low background.

The expected number of total background events across all configurations at the hand-tuned upper limit cuts defined in the original S5/VSR1 analysis below 2048 Hz is 0.080. A likelihood-ratio threshold of 7.97 (arbitrary units) gives the same background expectation and is set to apply the equivalent thresholds on pipeline statistic presented in Table 8.2. At the original upper limit thresholds, 239.4 weighted injections pass all cuts. 270.5 weighted injections pass the likelihood-ratio cut reflecting a 13% gain in the expected signal rate from the mock population at this particular false-alarm threshold.

After removing the blind injection from the search, the surviving highest-rank un-shifted event is from the Coherent Waveburst H1H2L1 200–2048 Hz search over the clean data set (all data quality and veto cuts applied). The event has $\rho = 3.87$ which corresponds to a likelihood ratio of 0.572 for this configuration (Figure 8-2). The global background expectation at this likelihood-ratio threshold is 1.29 (0.18 contribution from this search configuration alone), so an event of this strength is not unexpected. Using the loudest event prescription of Brady, Creighton, and Wiseman [108] we can set a frequentist upper limit of 2.303 expected events with $\Lambda > 0.572$ at 90% confidence. At this threshold, 385.7 events from the mock population survive, corresponding to a 61.1% increase in population detection efficiency over the fixed thresholds used for the original S5/VSR1 all-sky upper limit (Table 8.2).

As the upper limit for the population itself scales inversely with detection efficiency, a 61.1% increase in efficiency implies a 37.9% reduced upper limit on the intrinsic number of events in the population (such as rate per volume). In the original analysis [8], individual upper limits were established for various populations of signals that shared a single morphology and fixed amplitude at the Earth. Such populations are not explored here. The original upper limit thresholds were also chosen in the context of additional contributions to false-alarm-probability from networks and frequency ranges not used in this study (Table 8.1), and a reduced set of configurations would imply slightly lower thresholds assuming a fixed total false-alarm probability.

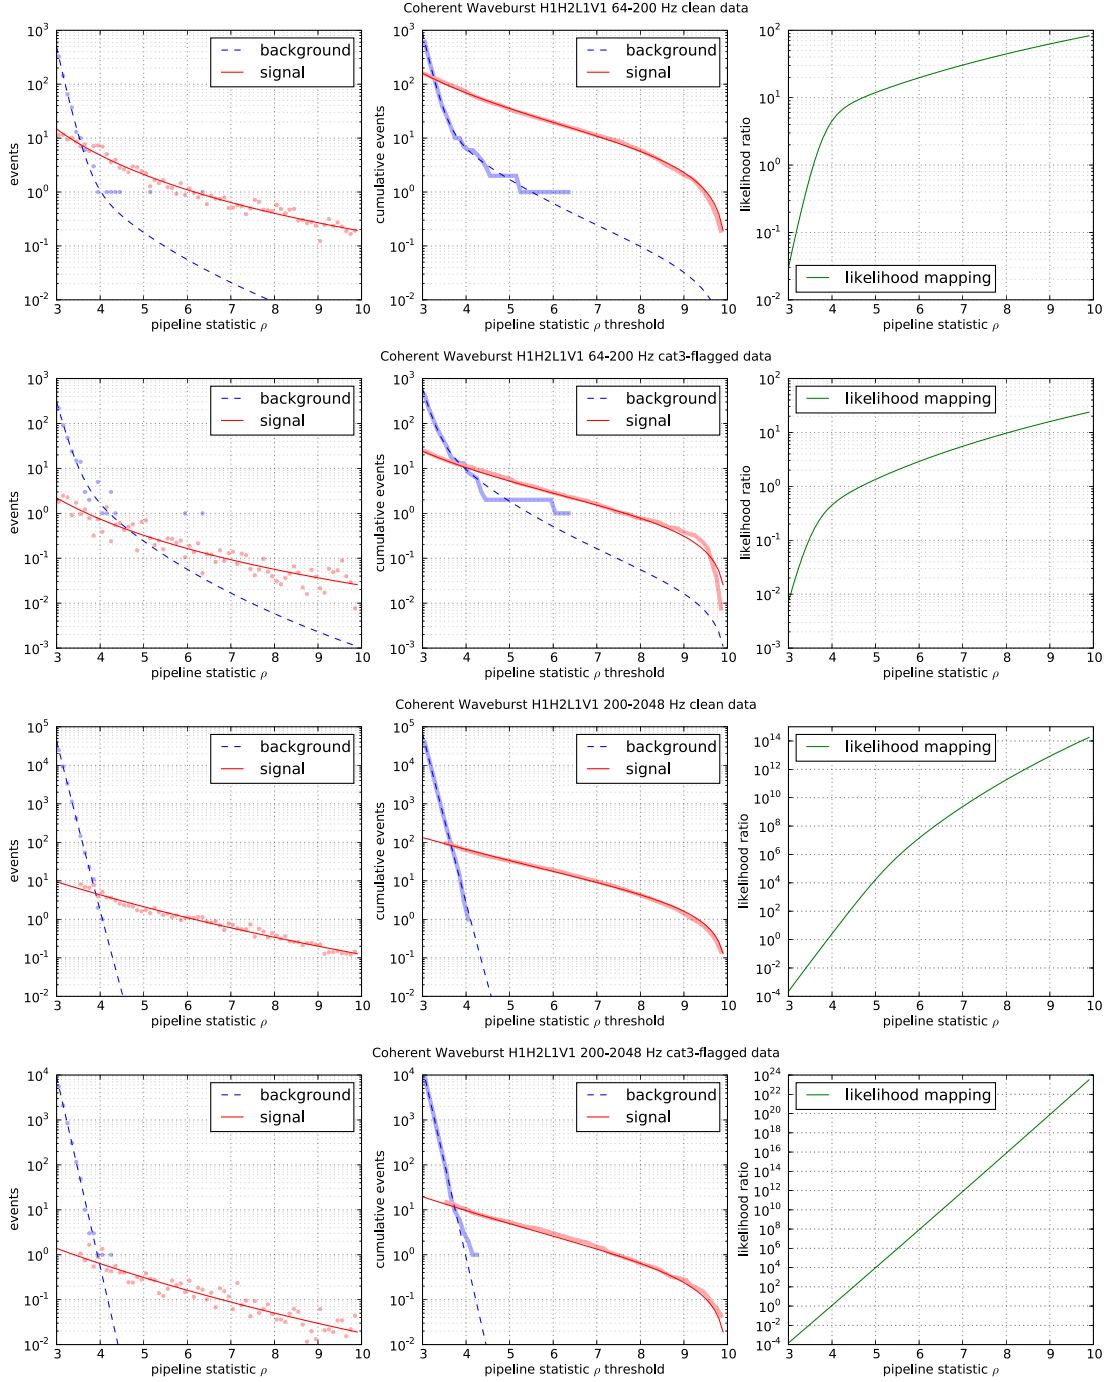


Figure 8-1: Coherent Waveburst H1H2L1V1 likelihood-ratio maps. Dashed blue curves represent the population of sample background events from 200 time-shifts. Red curves represent the number of signals from the single homogeneous mock population, scaled appropriately to the number of time-shifts and relative bin size (for the case of the discrete histogram). Analytic fits to the discrete distribution are used to calculate the likelihood-ratio mapping (signal/background) for the configuration.

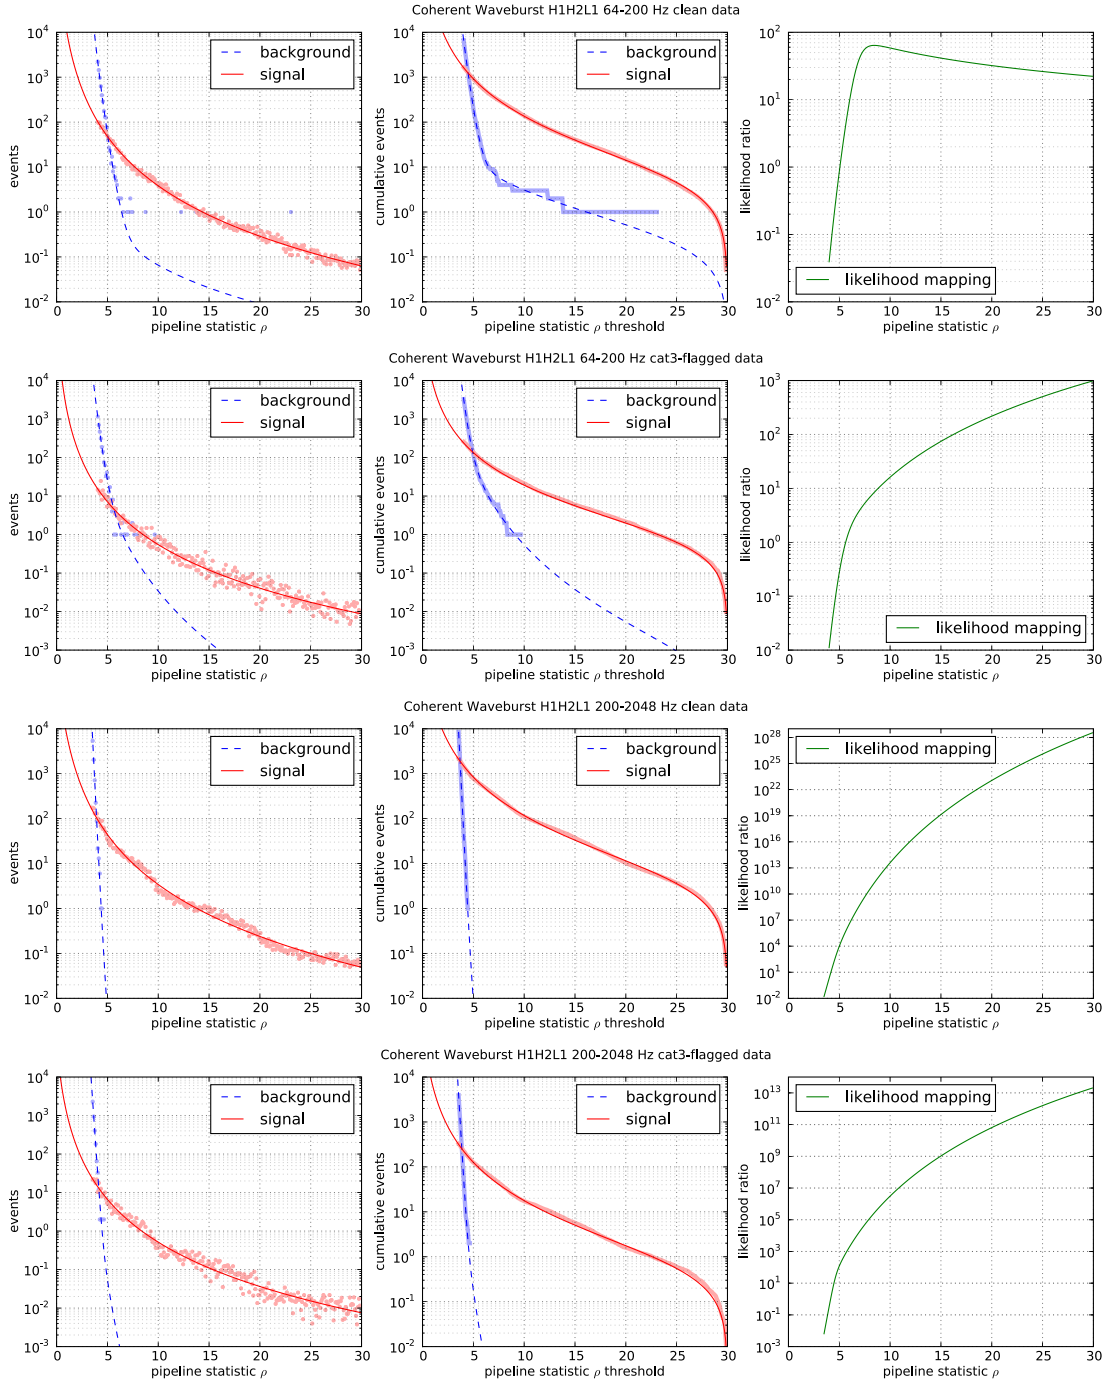


Figure 8-2: Coherent Waveburst H1H2L1 likelihood-ratio maps. Dashed blue curves represent the population of sample background events from 1000 time-shifts. Red curves represent the number of signals from the single homogeneous mock population, scaled appropriately to the number of time-shifts and relative bin size (for the case of the discrete histogram). Analytic fits to the discrete distribution are used to calculate the likelihood-ratio mapping (signal/background) for the configuration.

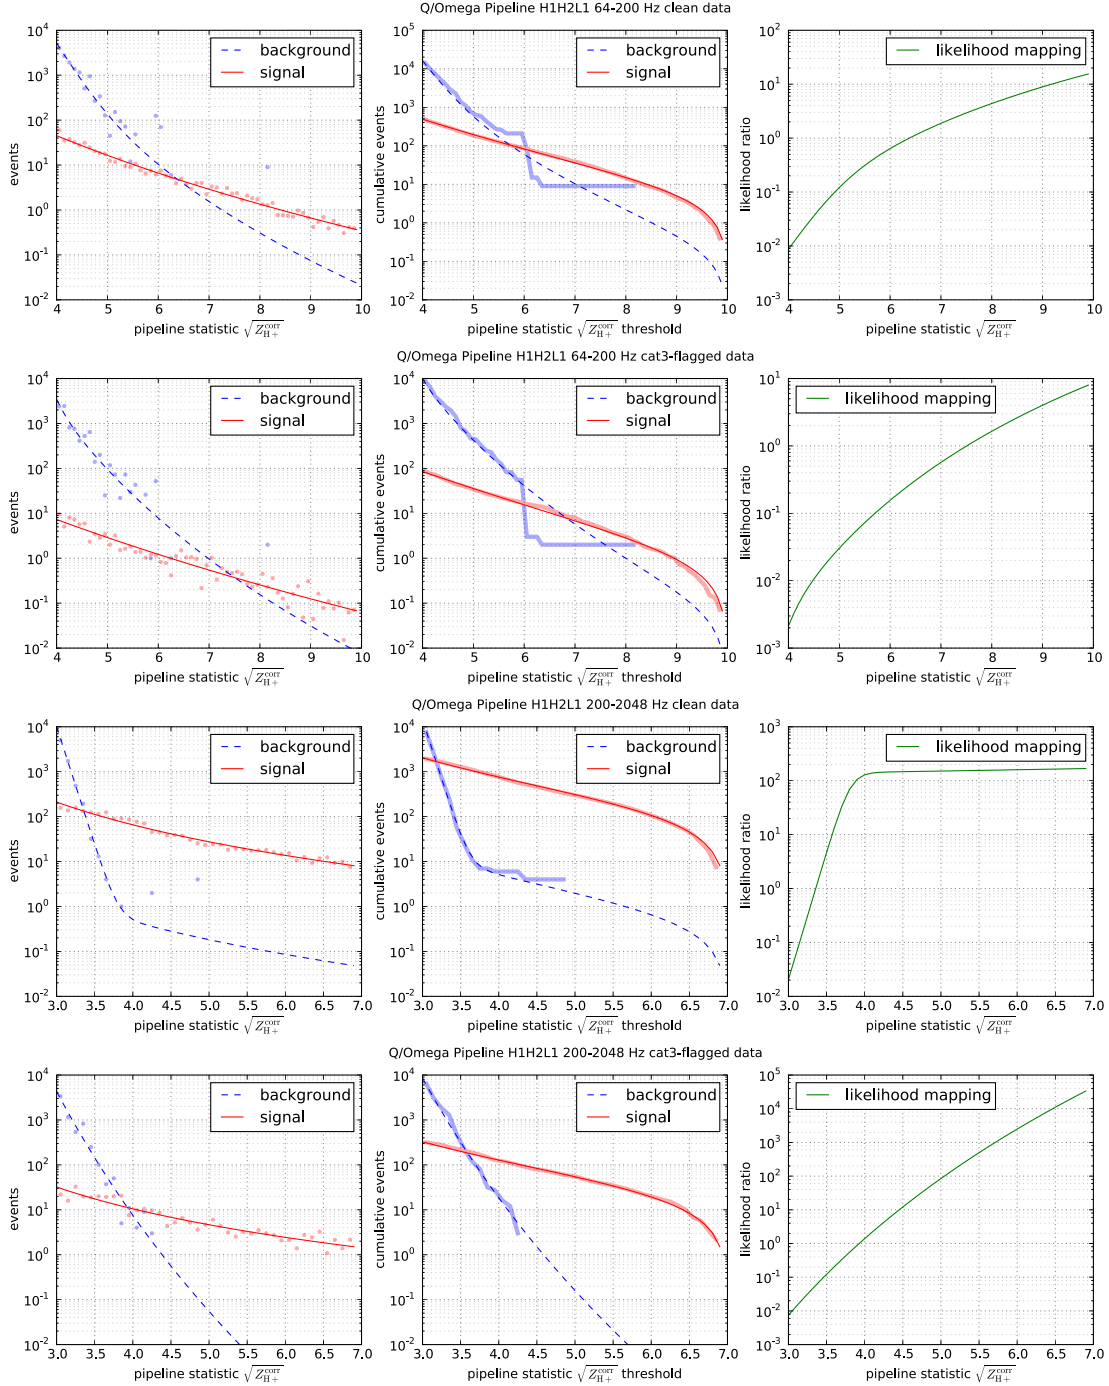


Figure 8-3: Q/ Ω Pipeline H1H2L1 likelihood-ratio maps Dashed blue curves represent the population of sample background events from 1000 time-shifts. Red curves represent the number of signals from the single homogeneous mock population, scaled appropriately to the number of time-shifts and relative bin size (for the case of the discrete histogram). Analytic fits to the discrete distribution are used to calculate the likelihood-ratio mapping (signal/background) for the configuration.

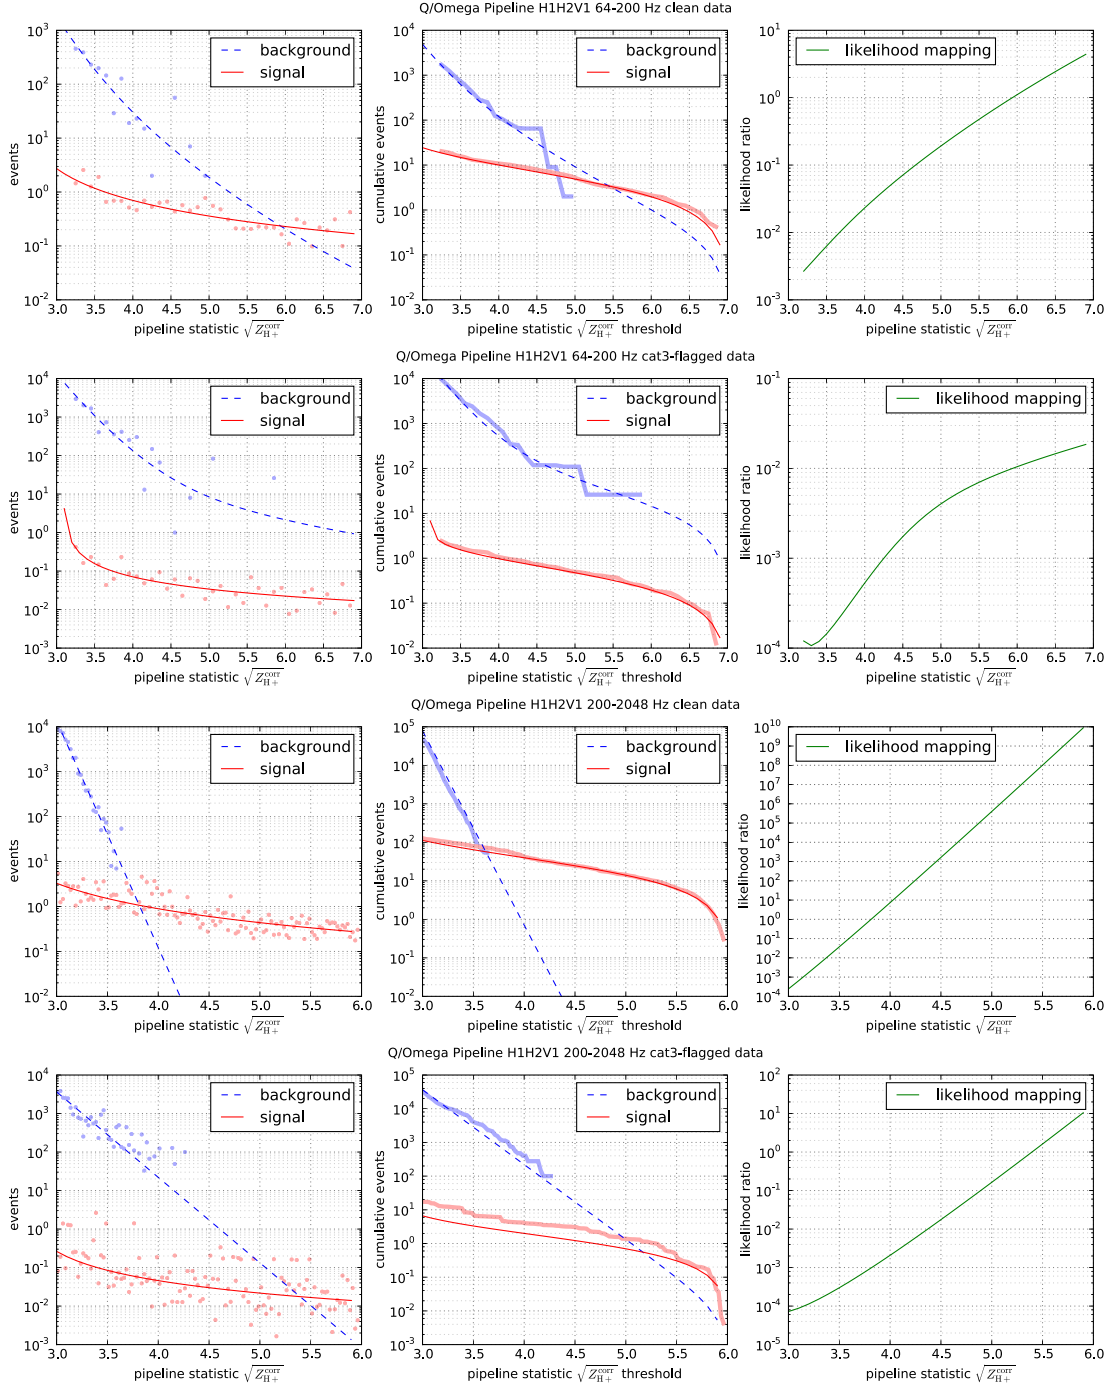


Figure 8-4: Q/ Ω Pipeline H1H2V1 likelihood-ratio maps Dashed blue curves represent the population of sample background events from 1000 time-shifts. Red curves represent the number of signals from the single homogeneous mock population, scaled appropriately to the number of time-shifts and relative bin size (for the case of the discrete histogram). Analytic fits to the discrete distribution are used to calculate the likelihood-ratio mapping (signal/background) for the configuration.

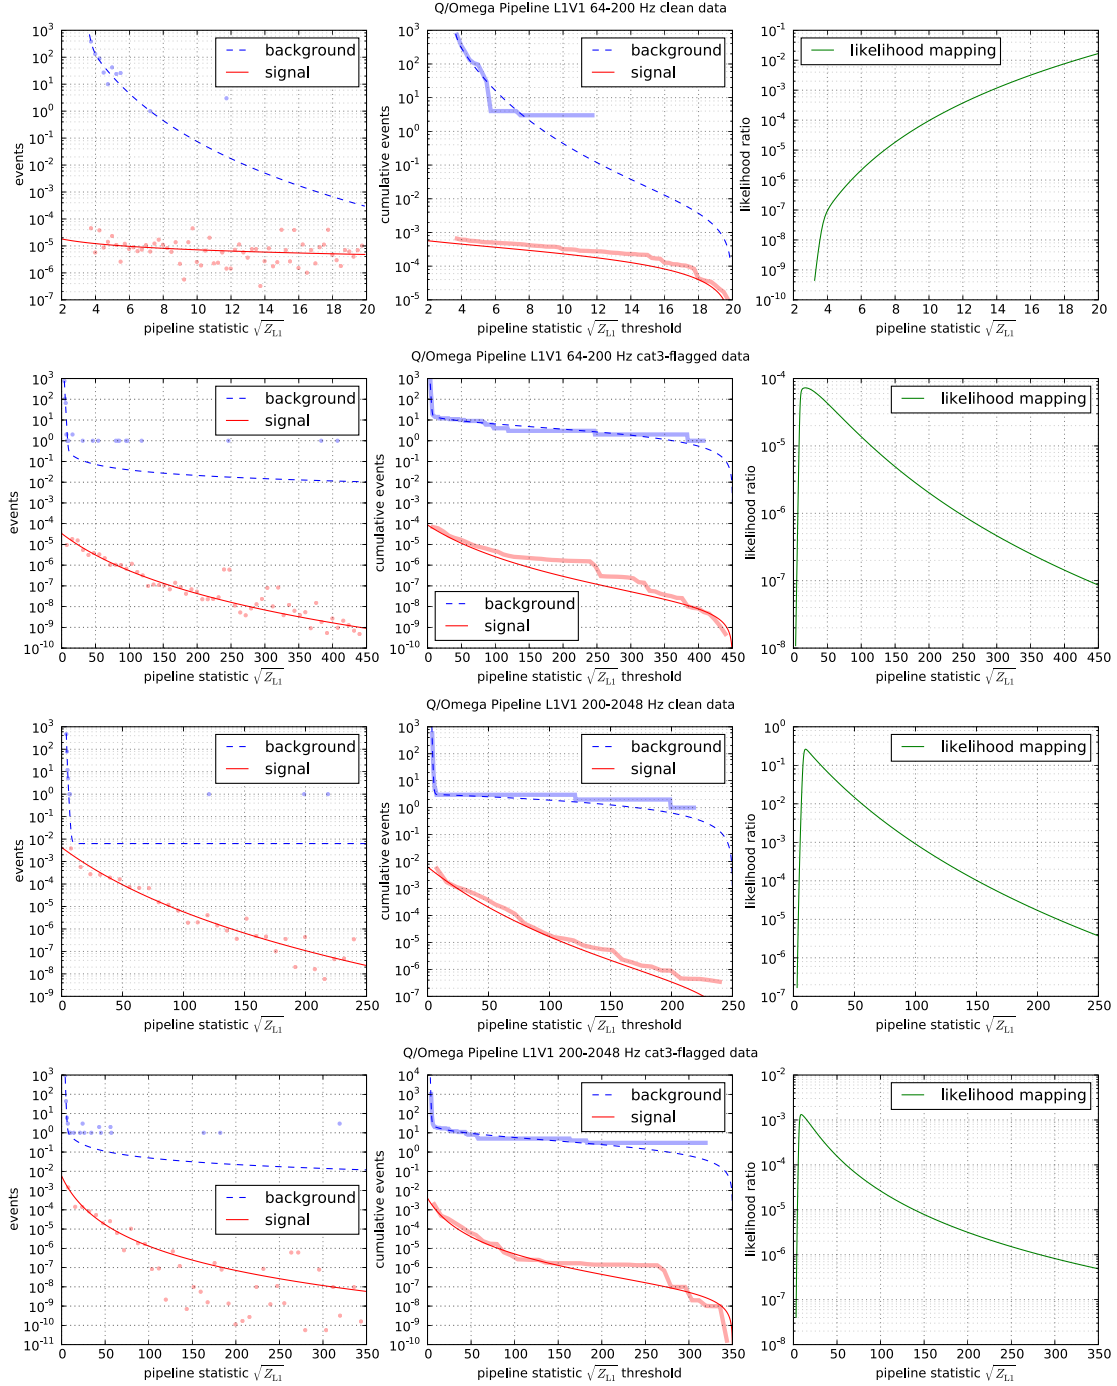


Figure 8-5: Q/ Ω Pipeline L1V1 likelihood-ratio maps Dashed blue curves represent the population of sample background events from 1000 time-shifts. Red curves represent the number of signals from the single homogeneous mock population, scaled appropriately to the number of time-shifts and relative bin size (for the case of the discrete histogram). Analytic fits to the discrete distribution are used to calculate the likelihood-ratio mapping (signal/background) for the configuration.

Chapter 9

Conclusions

The prospects for direct detection of gravitational radiation famously predicted by Einstein's theory of general relativity have improved enormously in the last few years with the successful commissioning of first generation kilometer-scale laser interferometric detectors. Having met its initial target of one full year of three-detector coincident observation at design sensitivity in October 2007, the US led Laser Interferometric Gravitational-wave Observatory (LIGO) is currently running in an upgraded configuration in anticipation of the move to Advanced LIGO. It is joined by the French-Italian Virgo detector which provides comparable sensitivity to LIGO across much of the observation band. Together, the three sites provide not only an ultra wide-field view of the gravitational-wave sky but also the important ability to resolve a source so that we can learn the most about the origins and dynamics of a potential gravitational-wave event.

Being prepared for a detection, then, is of prime importance for current analyses. Searches for gravitational-wave transients, which for ground-based detectors come primarily from the rare, violent motions and interactions of compact objects, are currently limited by the presence of non-Gaussian noise transients which populate the data to a sufficient degree that accidental coincidences dominate the background. Identifying and removing these noise transients are then key to maintaining sensitivity to weak gravitational wave signals. Furthermore, being able to characterize and developing a familiarity for the expected background is an important step to establish

confidence in any outlier event.

In this work, we have developed *kleineWelle*, an efficient method for transient identification based on the dyadic wavelet transform. Due to its simplicity and streamlined data handling, *kleineWelle* is able to scan all the fast and semi-fast channels at each LIGO site in real-time, producing lists of triggers which characterize the transient population at hundreds of test points in the interferometer and local environment. This information is fed back graphically to the instrument operators to aid in monitoring stationarity and correlations between channels, and is also used for real-time identification of instrumental disturbances for rapid rejection of online search background.

An automated procedure for scanning over the large collection of noise transients and selecting and tuning an efficient set of auxiliary channels for use as vetoes for burst search background was developed and applied to the all-sky un-triggered search for gravitational-wave bursts during S5/VSR1. The selection produces a list of veto conditions ordered from most effective to least effective at flagging noise transients observed in the gravitational-wave data using observed transient behavior in auxiliary channels. The automated classification greatly aids in tuning these veto choices, which allowed the S5 search to use a much richer set of auxiliary information to reject instrumental disturbances than what was permitted by time constraints in previous LIGO searches. The “event-by-event” vetoes were applied as a final step for background rejection in the S5/VSR1 burst analysis, and remove around 30–50% of remaining single-detector noise transients with a 1% per-instrument reduction in live-time. The performance on random coincident, coherent background varies considerably from about 10% for weak signals to around 50% for background outliers.

In the second half of this work, we introduced the follow-up procedure used to thoroughly test potential gravitational-wave candidate burst events surviving our otherwise blind analysis. The follow-up built upon an earlier example from the LIGO S2 burst search, but developed into a much more involved investigation with the presence of a single statistically significant outlier near the end of the S5/VSR1 run. The outlier event was subject to an extensive detection checklist where it was deter-

mined that there was no compelling evidence to suggest the signal was caused by an instrumental artifact. Ultimately the level of background in our blind analysis was such that the moderately loud event could not be considered sufficiently significant for a detection claim. The event remained as a marginally significant outlier in the final analysis until it was ultimately revealed to have been a blind hardware injection applied during the run to test the end-to-end search, specifically including follow-up procedures and the readiness to report a detection. The unique opportunity to vet a possibly real signal was a valuable exercise and experience for the analysis teams.

The follow-up procedure for S5 highlighted the importance of having a well defined procedure for evaluating the significance of an outlier in the context of a larger burst search encompassing several detector combinations, methods, and other divisions of the search. These divisions prevent an event from one sub-search from being directly comparable with an event from another. We have introduced two procedures which can be used to handle this trials factor with special consideration to the practicalities of transient gravitational-wave searches. The first is a unified comparison based on local false-alarm-rate, which has already been applied successfully in searches for compact binary coalescence [24, 25]. It is noted that the global background expectation for rankings based on fixed monotonic functions of local false-alarm probability (such as false-alarm-rate) can be readily calculated from local background only, which is of great practical merit. When applied to the burst outlier, there is an 11% chance of an event with lower false-alarm-rate to arise from one of the many burst search configurations, or an 8% chance for the subset of burst searches below 2048 Hz.

The second unified ranking is based on the likelihood ratio for each event when events are assumed to be independent and characterized only by the detection statistic assigned to them by the corresponding search method. The likelihood ratio requires the calculation of a signal and background population. We have implemented a procedure to cast existing burst simulations into a more realistic homogeneous population model, and developed procedures to calculate robust analytic fits for both signal and time-shift background distributions from which the likelihood-ratio mapping is readily derived. When applied to the burst outlier, there is a 2.2% chance of an

event with higher likelihood ratio to arise from background in one of the burst search configurations for events below 2048 Hz.

The likelihood-ratio ranking provides a 13% increase in detection efficiency for the mock signal population at the same total false-alarm probability as the original hand-tuned upper limit thresholds. The unified ranking also allows for the use of a loudest-event based upper limit determined by the highest rank event in the analysis. Applied to the S5/VSR1 burst events, the threshold set at the highest rank event provides a 61% increase in detection efficiency for the mock signal population over the original fixed upper limit thresholds. The frequentist upper limit of 2.303 detectable events at 90% confidence is the same for all three cases (hand-chosen fixed thresholds, fixed likelihood-ratio threshold, threshold at loudest event), so these increases in detection efficiency translate into 12% and 38% reductions in the population rate upper limit respectively.

While the likelihood ratio calculated this way is optimal under the assumption of independence among events, this is a poor assumption for the S5/VSR1 burst configurations. Events from gravitational-wave signals for search configurations which overlap in parameter space are generally dependent, indicating that better rankings could be calculated if the correlations between overlapping searches could be measured. For the time being, this is made difficult due to technical differences in time-shift background generation. In the case of search configurations which are largely independent, the likelihood ratio can give a fast and efficient means to combine events once the appropriate maps are calculated.

As the network of ground-based gravitational-wave detectors grows and improves in sensitivity, the confidence of a coincident observation as well as the ability to resolve an incoming gravitational wave is greatly enhanced. Joint analysis becomes increasingly more complicated, however, due to the additional non-stationarity introduced by each instrument. Automated procedures for characterizing each detector and interpreting and unifying results from a variety of different search configurations will play an important role in the first detection of a gravitational-wave event.

Appendix A

Statistical significance of Poisson and χ^2 processes

A.1 Poisson distribution

The Poisson distribution represents the probability of measuring k discrete occurrences from a Poisson process with expectation value μ . Often times we are dealing with a random process with a mean rate λ . Then, for a given observation time, L , the expectation value is $\mu = \lambda L$. The Poisson distribution has a discrete probability density function,

$$P_{\text{Pois}}(k|\mu) = f(k; \mu) = \frac{\mu^k e^{-\mu}}{k!}. \quad (\text{A.1})$$

The cumulative distribution function is the probability of measuring k or fewer occurrences,

$$g(k; \mu) = \sum_{n=0}^{n=k} f(n; \mu) = \frac{\Gamma(k+1, \mu)}{k!} = \frac{\Gamma(k+1, \mu)}{\Gamma(k+1)} = Q(k+1, \mu) \quad (\text{A.2})$$

where we have defined the gamma function, the upper incomplete gamma function, and the regularized upper incomplete gamma function, Q ,

$$\Gamma(z) = \int_0^{\infty} t^{z-1} e^{-t} dt \quad (\text{A.3})$$

$$\Gamma(z, x) = \int_x^{\infty} t^{z-1} e^{-t} dt \quad (\text{A.4})$$

$$Q(z, x) = \frac{\Gamma(z, x)}{\Gamma(z)}. \quad (\text{A.5})$$

For the purposes of measuring statistical significance of some measurement excess, we are generally interested in the probability of measuring k or more occurrences,

$$\sum_{n=k}^{n=\infty} f(n; \mu) = 1 - g(k-1; \mu) = 1 - \frac{\Gamma(k, \mu)}{\Gamma(k)} = 1 - Q(k, \mu) = P(k, \mu) \quad (\text{A.6})$$

where we have defined the regularized lower incomplete gamma function, P , which can be expressed in terms of the lower incomplete gamma function γ ,

$$\gamma(z, x) = \int_0^x t^{z-1} e^{-t} dt = \Gamma(z) - \Gamma(z, x) \quad (\text{A.7})$$

$$P(z, x) = \frac{\gamma(z, x)}{\Gamma(z)} = 1 - Q(z, x). \quad (\text{A.8})$$

We can then define the *significance* of this excess as the negative natural logarithm of this probability of observing k or more occurrences,

$$Z = -\ln P(k, \mu). \quad (\text{A.9})$$

A.2 Gaussian approximation

For large number statistics ($\mu \gg 1$), we can make use of the Gaussian approximation ($\sigma^2 = \mu$) to the Poisson distribution,

$$f(k; \mu) \approx \frac{1}{\sqrt{2\pi\mu}} \exp \left\{ \frac{-(k - \mu)^2}{2\mu} \right\} \quad (\text{A.10})$$

and use the number of sigmas away from mean as a proxy for $f(k, \mu)$ itself,

$$n_\sigma \equiv \frac{k - \mu}{\sqrt{\mu}}. \quad (\text{A.11})$$

This does not help us integrate the Poisson distribution for calculating a significance. In fact we need to be careful not to compare the discrete Poisson cumulative distribution to a continuous integration of the Gaussian distribution. In addition the approximation breaks down for large sigma regardless of sample size, making it inappropriate for estimating significance in the regime where the sum in equation A.6 is dominated by the first $n = k$ term. In general the Gaussian approximation can give the Poisson probability for large expectation value μ when the observed number of occurrences k is not too far from the mean. Using the error function to estimate the probability of observing k or more occurrences will only be a good approximation in the same regime.

A.3 Calculating Poisson significance

Calculating Poisson significance can be tricky because at small excesses, various approximations break down, and at large excesses the probabilities can become so small that numerical precision or floating-point boundaries get in the way. Although it might not actually matter if we can distinguish between significance values of several hundred, and the ridiculously small probabilities they correspond to, it is still nice to have a method of calculation which behaves well numerically. Because of the extreme probabilities involved, calculating directly the log-probability avoids many problems.

A.3.1 Gamma function approximations

The Lanczos approximation is used to calculate log-gamma,

$$\Gamma(a + 1) = \sqrt{2\pi} \left(a + g + \frac{1}{2} \right)^{a+1/2} e^{a+g+1/2} A_g(a), \quad (\text{A.12})$$

where g is a fixed constant, and A_g is a simple series with precalculated coefficients. Code for calculating $\ln \Gamma(a + 1)$ is available in Numerical Recipes.

The lower incomplete gamma function is approximated using the series

$$\gamma(a, x) = e^{-x} x^a \sum_{n=0}^{\infty} \frac{\Gamma(a)}{\Gamma(a + 1 + n)} x^n \quad (\text{A.13})$$

which can be made efficient with the recursion relation,

$$\Gamma(a + 1) = a\Gamma(a) \quad (\text{A.14})$$

$$\gamma(a, x) = e^{-x} x^a \sum_{n=0}^{\infty} \frac{x^n}{a(a + 1) \dots (a + n)} \quad (\text{A.15})$$

$$\ln \gamma(a, x) = -x + a \ln x + \ln \sum_{n=0}^{\infty} \frac{x^n}{a(a + 1) \dots (a + n)} \quad (\text{A.16})$$

$$\ln P(a, x) = \ln \frac{\gamma(a, x)}{\Gamma(a)} = \ln \gamma(a, x) - \ln \Gamma(a). \quad (\text{A.17})$$

The upper incomplete gamma function is approximated using the continued fraction

$$\Gamma(a, x) = e^{-x} x^a \left(\frac{1}{x+} \frac{1-a}{1+} \frac{1}{x+} \frac{2-a}{1+} \frac{2}{x+} \dots \right) \quad (\text{A.18})$$

$$\ln \Gamma(a, x) = -x + a \ln x + \ln \left(\frac{1}{x+} \frac{1-a}{1+} \frac{1}{x+} \frac{2-a}{1+} \frac{2}{x+} \dots \right) \quad (\text{A.19})$$

$$\ln Q(a, x) = \ln \frac{\Gamma(a, x)}{\Gamma(a)} = \ln \Gamma(a, x) - \ln \Gamma(a). \quad (\text{A.20})$$

The series expansion converges more rapidly on the domain $x < a + 1$ while the continuous fraction representation converges more rapidly for $x > a + 1$. Therefore if we want to calculate the regularized lower incomplete gamma function for $x > a + 1$, we first calculate $Q(a, x)$ and use $P(a, x) = 1 - Q(a, x)$. For calculation of $\ln P$ or $\ln Q$ at the extreme tails of the distribution (where we are often interested), it is necessary to move the logarithm inside the incomplete gamma function when making use of the code presented in Numerical Recipes. This generally requires only making a small change to take the log of the return statement.

A.3.2 Extreme tail of the Poisson CDF

We can revisit the question of what happens at the extreme tail of the upper Poisson CDF; equation A.6 for $k \gg \mu$. Here the upper Poisson CDF is dominated by the first term in the sum. Using the Poisson distribution function we can estimate the significance,

$$-Z \approx \ln \frac{\mu^k e^{-\mu}}{k!} \quad (\text{A.21})$$

$$\approx k \ln \mu - \mu - \ln k! \quad (\text{A.22})$$

$$\approx k \ln \mu - \mu - \ln k - \ln(k-1)! \quad (\text{A.23})$$

$$\approx k \ln \mu - \mu - \ln k - \ln \Gamma(k). \quad (\text{A.24})$$

Stirling's approximation can be used in place of $\ln k!$ in equation A.22 to give,

$$-Z \approx k \ln \mu - \mu - \left(k + \frac{1}{2}\right) \ln k + k - \frac{1}{2} \ln 2\pi - \frac{1}{12k} + \mathcal{O}\left(\frac{1}{k^3}\right). \quad (\text{A.25})$$

We can compare equation A.24 with the series expansion for $\ln P(k, \mu)$,

$$-Z = \ln P(k, \mu) = -\mu + k \ln \mu + \ln \sum_{n=0}^{\infty} \frac{\mu^n}{k(k+1)\dots(k+n)} - \ln \Gamma(k) \quad (\text{A.26})$$

$$= -\mu + k \ln \mu + \ln \left[\frac{1}{k} + \frac{\mu}{k(k+1)} + \frac{\mu^2}{k(k+1)(k+2)} + \dots \right] - \ln \Gamma(k) \quad (\text{A.27})$$

$$= -\mu + k \ln \mu + \ln \frac{1}{k} \left[1 + \frac{\mu}{(k+1)} + \frac{\mu^2}{(k+1)(k+2)} + \dots \right] - \ln \Gamma(k) \quad (\text{A.28})$$

$$= -\mu + k \ln \mu - \ln k - \ln \Gamma(k) + \ln \left[1 + \frac{\mu}{(k+1)} + \frac{\mu^2}{(k+1)(k+2)} + \dots \right] \quad (\text{A.29})$$

$$= -\mu + k \ln \mu - \ln k - \ln \Gamma(k) + \left[\frac{\mu}{(k+1)} + \frac{\mu^2}{(k+1)(k+2)} - \frac{\mu^2}{2(k+1)^2} + \mathcal{O}\left(\frac{\mu^3}{k^3}\right) \right]. \quad (\text{A.30})$$

So our approximation for the significance in equation A.24 is too large by the amount shown by the series. This makes sense because it only includes the dominant term in the sum for the upper Poisson CDF, and thus overestimates the significance.

A.4 χ^2 distributions

The χ^2 distribution often arises in excess power statistics as the values summed are often the squares of Gaussian random variables. A bonus of having robust means for calculating regularized incomplete gamma distributions is that they are readily applicable for calculating significances based on χ^2 statistics. The random value, x , representing the sum of squares of k random normal processes (zero-mean and unity variance) will be χ^2 distributed with k degrees of freedom,

$$f(x; k)dx = \chi_k^2(x)dx \quad (\text{A.31})$$

The probability of observing x greater than some value E is represented by the upper cumulative χ^2 distribution,

$$\text{Prob}(x \geq E; k) = \int_x^E \chi_k^2(x)dx = \frac{\Gamma(k/2, E/2)}{\Gamma(k/2)} = Q(k/2, E/2). \quad (\text{A.32})$$

Thus the significance of a measurement E given k degrees of freedom is

$$Z = -\ln Q(k/2, E/2) \quad (\text{A.33})$$

H0:PEM-ISCT1_ACCY	10	512	H1:ASC-WFS3_IY	4	256
H0:PEM-ISCT1_ACCZ	10	512	H1:ASC-WFS4_IP	4	256
H0:PEM-ISCT1_MIC	10	512	H1:ASC-WFS4_IY	4	256
H0:PEM-ISCT4_ACCX	10	512	H1:IOO-MC_F	64	1024
H0:PEM-ISCT4_ACCY	10	512	H1:IOO-MC_F	1024	2048
H0:PEM-ISCT4_ACCZ	10	512	H1:LSC-AS_AC	64	1024
H0:PEM-ISCT4_MIC	10	512	H1:LSC-AS_I	64	1024
H0:PEM-ISCT7_ACCX	10	512	H1:LSC-AS_I	1024	4096
H0:PEM-ISCT7_ACCY	10	512	H1:LSC-AS_Q	64	1024
H0:PEM-ISCT7_ACCZ	10	512	H1:LSC-DARM_CTRL	64	1024
H0:PEM-ISCT7_MIC	10	512	H1:LSC-DARM_CTRL_EXC_DAO	64	1024
H0:PEM-LVEA2_V1	10	512	H1:LSC-DARM_CTRL_EXC_DAO	1024	2048
H0:PEM-LVEA2_V2	10	512	H1:LSC-DARM_ERR	64	1024
H0:PEM-LVEA2_V3	10	512	H1:LSC-DARM_ERR	1024	4096
H0:PEM-LVEA_MAGX	10	512	H1:LSC-ETMX_EXC_DAO	64	1024
H0:PEM-LVEA_MAGX	0.5	32	H1:LSC-ETMX_EXC_DAO	1024	2048
H0:PEM-LVEA_MAGY	10	512	H1:LSC-MC_L	64	1024
H0:PEM-LVEA_MAGY	0.5	32	H1:LSC-MICH_CTRL	64	1024
H0:PEM-LVEA_MAGZ	10	512	H1:LSC-MICH_CTRL	1024	2048
H0:PEM-LVEA_MAGZ	0.5	32	H1:LSC-POB_I	64	1024
H0:PEM-LVEA_MIC	10	512	H1:LSC-POB_I	1024	4096
H0:PEM-LVEA_SEISX	10	512	H1:LSC-POB_Q	64	1024
H0:PEM-LVEA_SEISY	10	512	H1:LSC-POB_Q	1024	4096
H0:PEM-LVEA_SEISZ	10	512	H1:LSC-POBS_DC	64	1024
H0:PEM-MX_SEISX	10	512	H1:LSC-PRC_CTRL	64	1024
H0:PEM-MX_SEISY	10	512	H1:LSC-PRC_CTRL	1024	2048
H0:PEM-MX_SEISZ	10	512	H1:LSC-REFL_DC	64	1024
H0:PEM-MX_V1	10	512	H1:LSC-REFL_I	64	1024
H0:PEM-MX_V2	10	512	H1:LSC-REFL_I	1024	4096
H0:PEM-MY_SEISX	10	512	H1:LSC-REFL_Q	64	1024
H0:PEM-MY_SEISY	10	512	H1:LSC-REFL_Q	1024	2048
H0:PEM-MY_SEISZ	10	512	H1:SUS-BS_OPLEV_PERROR	4	256
H0:PEM-MY_V1	10	512	H1:SUS-BS_OPLEV_YERROR	4	256
H0:PEM-MY_V2	10	512	H1:SUS-ETMX_OPLEV_PERROR	4	256
H0:PEM-PSL1_ACCX	10	512	H1:SUS-ETMX_OPLEV_YERROR	4	256
H0:PEM-PSL1_ACCZ	10	512	H1:SUS-ETMY_OPLEV_PERROR	4	256
H0:PEM-PSL1_MIC	10	512	H1:SUS-ETMY_OPLEV_YERROR	4	256
H0:PEM-PSL2_ACCX	10	512	H1:SUS-ITMX_OPLEV_PERROR	4	256
H0:PEM-PSL2_ACCZ	10	512	H1:SUS-ITMX_OPLEV_YERROR	4	256
H0:PEM-PSL2_MIC	10	512	H1:SUS-ITMY_OPLEV_PERROR	4	256
H0:PEM-RADIO_CS_1	10	512	H1:SUS-ITMY_OPLEV_YERROR	4	256
H0:PEM-RADIO_CS_2	10	512	H1:SUS-MMT3_OPLEV_PERROR	4	256
H0:PEM-RADIO_LVEA	10	512	H1:SUS-MMT3_OPLEV_YERROR	4	256
H1:ASC-BS_P	4	256	H1:SUS-RM_OPLEV_PERROR	4	256
H1:ASC-BS_Y	4	256	H1:SUS-RM_OPLEV_YERROR	4	256
H1:ASC-ETMX_P	4	256	H1:TCS-ITMX_PD1AC	4	256
H1:ASC-ETMX_Y	4	256	H1:TCS-ITMX_PD2AC	4	256
H1:ASC-ETMY_P	4	256	H1:TCS-ITMY_PD1AC	4	256
H1:ASC-ETMY_Y	4	256	H1:TCS-ITMY_PD2AC	4	256
H1:ASC-ITMX_P	4	256	H2:ASC-BS_P	4	256
H1:ASC-ITMX_Y	4	256	H2:ASC-BS_Y	4	256
H1:ASC-ITMY_P	4	256	H2:ASC-ETMX_P	4	256
H1:ASC-ITMY_Y	4	256	H2:ASC-ETMX_Y	4	256
H1:ASC-QPDX_DC	4	256	H2:ASC-ETMY_P	4	256
H1:ASC-QPDX_P	4	256	H2:ASC-ETMY_Y	4	256
H1:ASC-QPDX_Y	4	256	H2:ASC-ITMX_P	4	256
H1:ASC-QPDY_DC	4	256	H2:ASC-ITMX_Y	4	256
H1:ASC-QPDY_P	4	256	H2:ASC-ITMY_P	4	256
H1:ASC-QPDY_Y	4	256	H2:ASC-ITMY_Y	4	256
H1:ASC-RM_P	4	256	H2:ASC-QPDX_DC	4	256
H1:ASC-RM_Y	4	256	H2:ASC-QPDX_P	4	256
H1:ASC-WFS1_QP	4	256	H2:ASC-QPDX_Y	4	256
H1:ASC-WFS1_QY	4	256	H2:ASC-QPDY_DC	4	256
H1:ASC-WFS2_IP	4	256	H2:ASC-QPDY_P	4	256
H1:ASC-WFS2_IY	4	256	H2:ASC-QPDY_Y	4	256
H1:ASC-WFS2_QP	4	256	H2:ASC-RM_P	4	256
H1:ASC-WFS2_QY	4	256	H2:ASC-RM_Y	4	256
H1:ASC-WFS3_IP	4	256			

H2:ASC-WFS1_QP	4	256	H2:TCS-ITMY_PD2AC	4	256
H2:ASC-WFS1_QY	4	256	L0:PEM-BSC1_ACCX	10	512
H2:ASC-WFS2_IP	4	256	L0:PEM-BSC1_ACCY	10	512
H2:ASC-WFS2_IY	4	256	L0:PEM-BSC1_ACCZ	10	512
H2:ASC-WFS2_QP	4	256	L0:PEM-BSC2_ACCX	10	512
H2:ASC-WFS2_QY	4	256	L0:PEM-BSC2_ACCY	10	512
H2:ASC-WFS3_IP	4	256	L0:PEM-BSC2_ACCZ	10	512
H2:ASC-WFS3_IY	4	256	L0:PEM-BSC3_ACCX	10	512
H2:ASC-WFS4_IP	4	256	L0:PEM-BSC3_ACCY	10	512
H2:ASC-WFS4_IY	4	256	L0:PEM-BSC3_ACCZ	10	512
H2:IOO-MC_F	64	1024	L0:PEM-BSC4_ACCX	10	512
H2:IOO-MC_F	1024	2048	L0:PEM-BSC4_ACCY	10	512
H2:LSC-AS_AC	64	1024	L0:PEM-BSC4_ACCZ	10	512
H2:LSC-AS_I	64	1024	L0:PEM-BSC4_MIC	10	512
H2:LSC-AS_I	1024	4096	L0:PEM-BSC5_ACCX	10	512
H2:LSC-AS_Q	64	1024	L0:PEM-BSC5_ACCY	10	512
H2:LSC-DARM_CTRL	64	1024	L0:PEM-BSC5_ACCZ	10	512
H2:LSC-DARM_CTRL_EXC_DAQ	64	1024	L0:PEM-BSC5_MIC	10	512
H2:LSC-DARM_CTRL_EXC_DAQ	1024	2048	L0:PEM-COIL_MAGX	10	512
H2:LSC-DARM_ERR	64	1024	L0:PEM-COIL_MAGX	0.5	32
H2:LSC-DARM_ERR	1024	4096	L0:PEM-COIL_MAGZ	10	512
H2:LSC-ETMX_CAL	64	1024	L0:PEM-COIL_MAGZ	0.5	32
H2:LSC-ETMX_CAL	1024	2048	L0:PEM-EX_BAYMIC	10	512
H2:LSC-ETMX_CAL_EXC_DAQ	64	1024	L0:PEM-EX_MAGX	10	512
H2:LSC-ETMX_CAL_EXC_DAQ	1024	2048	L0:PEM-EX_MAGX	0.5	32
H2:LSC-ETMX_EXC_DAQ	64	1024	L0:PEM-EX_MAGY	10	512
H2:LSC-ETMX_EXC_DAQ	1024	2048	L0:PEM-EX_MAGY	0.5	32
H2:LSC-ETMY_CAL	64	1024	L0:PEM-EX_MAGZ	10	512
H2:LSC-ETMY_CAL	1024	2048	L0:PEM-EX_MAGZ	0.5	32
H2:LSC-ETMY_CAL_EXC_DAQ	64	1024	L0:PEM-EX_SEISX	10	512
H2:LSC-ETMY_CAL_EXC_DAQ	1024	2048	L0:PEM-EX_SEISY	10	512
H2:LSC-MC_L	64	1024	L0:PEM-EX_SEISZ	10	512
H2:LSC-MC_L	1024	2048	L0:PEM-EX_V1	10	512
H2:LSC-MICH_CTRL	64	1024	L0:PEM-EY_BAYMIC	10	512
H2:LSC-MICH_CTRL	1024	2048	L0:PEM-EY_MAGX	10	512
H2:LSC-POB_I	64	1024	L0:PEM-EY_MAGX	0.5	32
H2:LSC-POB_I	1024	4096	L0:PEM-EY_MAGY	10	512
H2:LSC-POB_Q	64	1024	L0:PEM-EY_MAGY	0.5	32
H2:LSC-POB_Q	1024	4096	L0:PEM-EY_MAGZ	10	512
H2:LSC-POY_DC	64	1024	L0:PEM-EY_MAGZ	0.5	32
H2:LSC-PRC_CTRL	64	1024	L0:PEM-EY_SEISX	10	512
H2:LSC-PRC_CTRL	1024	2048	L0:PEM-EY_SEISY	10	512
H2:LSC-REFL_AC	64	1024	L0:PEM-EY_SEISZ	10	512
H2:LSC-REFL_I	64	1024	L0:PEM-EY_V1	10	512
H2:LSC-REFL_I	1024	4096	L0:PEM-HAM1_ACCX	10	512
H2:LSC-REFL_Q	64	1024	L0:PEM-HAM1_ACCZ	10	512
H2:LSC-REFL_Q	1024	4096	L0:PEM-HAM2_ACCX	10	512
H2:SUS-BS_OPLEV_PERROR	4	256	L0:PEM-HAM2_ACCZ	10	512
H2:SUS-BS_OPLEV_YERROR	4	256	L0:PEM-ISCT1_ACCX	10	512
H2:SUS-ETMX_OPLEV_PERROR	4	256	L0:PEM-ISCT1_ACCY	10	512
H2:SUS-ETMX_OPLEV_YERROR	4	256	L0:PEM-ISCT1_ACCZ	10	512
H2:SUS-ETMY_OPLEV_PERROR	4	256	L0:PEM-ISCT1_MIC	10	512
H2:SUS-ETMY_OPLEV_YERROR	4	256	L0:PEM-ISCT4_ACCX	10	512
H2:SUS-FMX_OPLEV_PERROR	4	256	L0:PEM-ISCT4_ACCY	10	512
H2:SUS-FMX_OPLEV_YERROR	4	256	L0:PEM-ISCT4_ACCZ	10	512
H2:SUS-FMY_OPLEV_PERROR	4	256	L0:PEM-ISCT4_MIC	10	512
H2:SUS-FMY_OPLEV_YERROR	4	256	L0:PEM-LVEA_BAYMIC	10	512
H2:SUS-ITMX_OPLEV_PERROR	4	256	L0:PEM-LVEA_MAGX	10	512
H2:SUS-ITMX_OPLEV_YERROR	4	256	L0:PEM-LVEA_MAGX	0.5	32
H2:SUS-ITMY_OPLEV_PERROR	4	256	L0:PEM-LVEA_MAGY	10	512
H2:SUS-ITMY_OPLEV_YERROR	4	256	L0:PEM-LVEA_MAGY	0.5	32
H2:SUS-MMT3_OPLEV_PERROR	4	256	L0:PEM-LVEA_MAGZ	10	512
H2:SUS-MMT3_OPLEV_YERROR	4	256	L0:PEM-LVEA_MAGZ	0.5	32
H2:SUS-RM_OPLEV_PERROR	4	256	L0:PEM-LVEA_MIC	10	512
H2:SUS-RM_OPLEV_YERROR	4	256	L0:PEM-LVEA_SEISX	10	512
H2:TCS-ITMX_PD1AC	4	256	L0:PEM-LVEA_SEISY	10	512
H2:TCS-ITMX_PD2AC	4	256	L0:PEM-LVEA_SEISZ	10	512
H2:TCS-ITMY_PD1AC	4	256			

L0:PEM-LVEA_V1	10	512	L1:LSC-DARM_ERR	1024	4096
L0:PEM-RADIO_LVEA	10	512	L1:LSC-ETMX_CAL	64	1024
L1:ASC-BS_P	4	256	L1:LSC-ETMX_CAL	1024	4096
L1:ASC-BS_Y	4	256	L1:LSC-ETMX_EXC_DDAQ	64	1024
L1:ASC-ETMX_P	4	256	L1:LSC-ETMX_EXC_DDAQ	1024	2048
L1:ASC-ETMX_Y	4	256	L1:LSC-ETMY_CAL	64	1024
L1:ASC-ETMY_P	4	256	L1:LSC-ETMY_CAL	1024	4096
L1:ASC-ETMY_Y	4	256	L1:LSC-MICH_CTRL	64	1024
L1:ASC-ITMX_P	4	256	L1:LSC-MICH_CTRL	1024	2048
L1:ASC-ITMX_Y	4	256	L1:LSC-POB_I	64	1024
L1:ASC-ITMY_P	4	256	L1:LSC-POB_I	1024	4096
L1:ASC-ITMY_Y	4	256	L1:LSC-POB_Q	64	1024
L1:ASC-QPDX_DC	4	256	L1:LSC-POB_Q	1024	4096
L1:ASC-QPDX_P	4	256	L1:LSC-PRC_CTRL	64	1024
L1:ASC-QPDX_Y	4	256	L1:LSC-PRC_CTRL	1024	2048
L1:ASC-QPDY_DC	4	256	L1:LSC-REFL_AC	64	1024
L1:ASC-QPDY_P	4	256	L1:LSC-REFL_I	64	1024
L1:ASC-QPDY_Y	4	256	L1:LSC-REFL_I	1024	4096
L1:ASC-RM_P	4	256	L1:LSC-REFL_Q	64	1024
L1:ASC-RM_Y	4	256	L1:LSC-REFL_Q	1024	2048
L1:ASC-WFS1_QP	4	256	L1:SEI-ETMX_STS2_X	1	128
L1:ASC-WFS1_QY	4	256	L1:SEI-ETMY_STS2_Y	1	128
L1:ASC-WFS2_IP	4	256	L1:SEI-LVEA_STS2_X	1	128
L1:ASC-WFS2_IY	4	256	L1:SEI-LVEA_STS2_Y	1	128
L1:ASC-WFS2_QP	4	256	L1:SEI-LVEA_STS2_Z	1	128
L1:ASC-WFS2_QY	4	256	L1:SUS-BS_OPLEV_PERROR	4	256
L1:ASC-WFS3_IP	4	256	L1:SUS-BS_OPLEV_YERROR	4	256
L1:ASC-WFS3_IY	4	256	L1:SUS-ETMX_OPLEV_PERROR	4	256
L1:ASC-WFS4_IP	4	256	L1:SUS-ETMX_OPLEV_YERROR	4	256
L1:ASC-WFS4_IY	4	256	L1:SUS-ETMY_OPLEV_PERROR	4	256
L1:IOO-MC_F	64	1024	L1:SUS-ETMY_OPLEV_YERROR	4	256
L1:IOO-MC_F	1024	2048	L1:SUS-ITMX_OPLEV_PERROR	4	256
L1:LSC-AS_AC	64	1024	L1:SUS-ITMX_OPLEV_YERROR	4	256
L1:LSC-AS_I	64	1024	L1:SUS-ITMY_OPLEV_PERROR	4	256
L1:LSC-AS_I	1024	4096	L1:SUS-ITMY_OPLEV_YERROR	4	256
L1:LSC-AS_Q	64	1024	L1:SUS-MMT3_OPLEV_PERROR	4	256
L1:LSC-DARM_CTRL	64	1024	L1:SUS-MMT3_OPLEV_YERROR	4	256
L1:LSC-DARM_CTRL_EXC_DDAQ	64	1024	L1:SUS-RM_OPLEV_PERROR	4	256
L1:LSC-DARM_CTRL_EXC_DDAQ	1024	2048	L1:SUS-RM_OPLEV_YERROR	4	256
L1:LSC-DARM_ERR	64	1024			

B.2 S6 channels and frequency range

channel	f_{low}	f_{high}			
H0:PEM-BSC8_MIC			H0:PEM-BSC8_MIC	8	1024
H0:PEM-BSC9_ACC1X			H0:PEM-BSC9_ACC1X	8	1024
H0:PEM-BSC9_MAGX			H0:PEM-BSC9_MAGX	1	1024
H0:PEM-BSC9_MAGY			H0:PEM-BSC9_MAGY	1	1024
H0:PEM-BSC9_MAGZ			H0:PEM-BSC9_MAGZ	1	1024
H0:PEM-BSC9_MIC			H0:PEM-BSC9_MIC	8	1024
H0:PEM-COIL_MAGX			H0:PEM-COIL_MAGX	1	1024
H0:PEM-COIL_MAGZ			H0:PEM-COIL_MAGZ	1	1024
H0:PEM-EX_PWR1			H0:PEM-EX_PWR1	8	1024
H0:PEM-EX_SEISX			H0:PEM-EX_SEISX	1	8
H0:PEM-EX_SEISY			H0:PEM-EX_SEISY	8	128
H0:PEM-EX_SEISZ			H0:PEM-EX_SEISZ	1	8
H0:PEM-EX_V1			H0:PEM-EX_V1	8	1024
H0:PEM-EY_SEISX			H0:PEM-EY_SEISX	1	8
H0:PEM-EY_SEISY			H0:PEM-EY_SEISY	8	128
H0:PEM-EY_SEISZ			H0:PEM-EY_SEISZ	1	8
H0:PEM-EY_SEISZ			H0:PEM-EY_SEISZ	8	128
H0:PEM-EY_SEISZ			H0:PEM-EY_SEISZ	1	8
H0:PEM-EY_SEISZ			H0:PEM-EY_SEISZ	8	128

H0:PEM-EY_V1	8	1024	H0:PEM-RADIO_CS_1	8	1024
H0:PEM-HAM1_ACCX	8	1024	H0:PEM-RADIO_CS_2	8	1024
H0:PEM-HAM1_ACCZ	8	1024	H0:PEM-RADIO_LVEA	8	1024
H0:PEM-HAM3_ACCX	8	1024	H0:PEM-RADIO_LVEA_H1	8	1024
H0:PEM-HAM6_ACCX	8	1024	H1:ASC-BS_P	8	256
H0:PEM-HAM6_ACCY	8	1024	H1:ASC-BS_Y	8	256
H0:PEM-HAM6_ACCZ	8	1024	H1:ASC-ETMX_P	8	256
H0:PEM-HAM6_MIC	8	1024	H1:ASC-ETMX_Y	8	256
H0:PEM-HAM7_ACCX	8	1024	H1:ASC-ETMY_P	8	256
H0:PEM-HAM7_ACCZ	8	1024	H1:ASC-ETMY_Y	8	256
H0:PEM-HAM9_ACCX	8	1024	H1:ASC-ITMX_P	8	256
H0:PEM-IOT1_MIC	8	1024	H1:ASC-ITMX_Y	8	256
H0:PEM-IOT7_MIC	8	1024	H1:ASC-ITMY_P	8	256
H0:PEM-ISCT10_ACCX	8	1024	H1:ASC-ITMY_Y	8	256
H0:PEM-ISCT10_ACCY	8	1024	H1:ASC-QPDX_DC	8	128
H0:PEM-ISCT10_ACCZ	8	1024	H1:ASC-QPDX_P	8	128
H0:PEM-ISCT10_MIC	8	1024	H1:ASC-QPDX_Y	8	512
H0:PEM-ISCT1_ACCX	8	1024	H1:ASC-QPDY_DC	8	128
H0:PEM-ISCT1_ACCY	8	1024	H1:ASC-QPDY_P	8	128
H0:PEM-ISCT1_ACCZ	8	1024	H1:ASC-QPDY_Y	8	128
H0:PEM-ISCT1_MIC	8	1024	H1:ASC-RM_P	8	256
H0:PEM-ISCT4_ACCPER	8	1024	H1:ASC-RM_Y	8	256
H0:PEM-ISCT4_ACCZ	8	1024	H1:ASC-WFS1_QP	8	256
H0:PEM-ISCT4_MIC	8	1024	H1:ASC-WFS1_QY	8	256
H0:PEM-ISCT7_ACCX	8	1024	H1:ASC-WFS2_IP	8	256
H0:PEM-ISCT7_ACCY	8	1024	H1:ASC-WFS2_IY	8	256
H0:PEM-ISCT7_ACCZ	8	1024	H1:ASC-WFS2_QP	8	256
H0:PEM-ISCT7_MIC	8	1024	H1:ASC-WFS2_QY	8	256
H0:PEM-LSC1_MAGX	1	1024	H1:ASC-WFS3_IP	8	256
H0:PEM-LSC1_MAGY	1	1024	H1:ASC-WFS3_IY	8	256
H0:PEM-LSC1_MAGZ	1	1024	H1:ASC-WFS4_IP	8	256
H0:PEM-LVEA2_V1	8	1024	H1:ASC-WFS4_IY	8	256
H0:PEM-LVEA2_V2	8	1024	H1:IOO-MC_F	32	2048
H0:PEM-LVEA2_V3	8	1024	H1:ISI-OMC_CONT_RX_IN1_DAQ	8	1024
H0:PEM-LVEA_MAGX	1	1024	H1:ISI-OMC_CONT_RY_IN1_DAQ	8	1024
H0:PEM-LVEA_MAGY	1	1024	H1:ISI-OMC_CONT_RZ_IN1_DAQ	8	1024
H0:PEM-LVEA_MAGZ	1	1024	H1:ISI-OMC_CONT_X_IN1_DAQ	8	1024
H0:PEM-LVEA_MIC	32	2048	H1:ISI-OMC_CONT_Y_IN1_DAQ	8	1024
H0:PEM-LVEA_MIC	1024	4096	H1:ISI-OMC_CONT_Z_IN1_DAQ	8	1024
H0:PEM-LVEA_PWR1	8	1024	H1:ISI-OMC_DISPPF_H1_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISX	1	8	H1:ISI-OMC_DISPPF_V1_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISX	8	128	H1:ISI-OMC_GEOPF_H1_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISY	1	8	H1:ISI-OMC_GEOPF_H2_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISY	8	128	H1:ISI-OMC_GEOPF_H3_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISZ	1	8	H1:ISI-OMC_GEOPF_V1_IN1_DAQ	8	1024
H0:PEM-LVEA_SEISZ	8	128	H1:ISI-OMC_GEOPF_V2_IN1_DAQ	8	1024
H0:PEM-MX_SEISX	1	8	H1:ISI-OMC_GEOPF_V3_IN1_DAQ	8	1024
H0:PEM-MX_SEISX	8	128	H1:LSC-AS_AC	8	1024
H0:PEM-MX_SEISY	1	8	H1:LSC-AS_DC	8	128
H0:PEM-MX_SEISY	8	128	H1:LSC-DARM_CTRL	32	2048
H0:PEM-MX_SEISZ	1	8	H1:LSC-DARM_CTRL	1024	4096
H0:PEM-MX_SEISZ	8	128	H1:LSC-DARM_CTRL_EXC_DAQ	32	2048
H0:PEM-MY_SEISX	1	8	H1:LSC-DARM_CTRL_EXC_DAQ	1024	4096
H0:PEM-MY_SEISX	8	128	H1:LSC-DARM_ERR	32	2048
H0:PEM-MY_SEISY	1	8	H1:LSC-DARM_ERR	1024	4096
H0:PEM-MY_SEISY	8	128	H1:LSC-ETMX_CAL	32	2048
H0:PEM-MY_SEISZ	1	8	H1:LSC-ETMX_CAL	1024	4096
H0:PEM-MY_SEISZ	8	128	H1:LSC-ETMX_CAL_EXC_DAQ	32	2048
H0:PEM-OMC1_MAGX	1	1024	H1:LSC-ETMX_CAL_EXC_DAQ	1024	4096
H0:PEM-OUT_PWR1	8	1024	H1:LSC-ETMX_EXC_DAQ	32	2048
H0:PEM-PSL1_ACCX	8	1024	H1:LSC-ETMX_EXC_DAQ	1024	4096
H0:PEM-PSL1_ACCY	8	1024	H1:LSC-ETMY_CAL	32	2048
H0:PEM-PSL1_ACCZ	8	1024	H1:LSC-ETMY_CAL	1024	4096
H0:PEM-PSL1_MIC	8	1024	H1:LSC-ETMY_CAL_EXC_DAQ	32	2048
H0:PEM-PSL2_ACCX	8	1024	H1:LSC-ETMY_CAL_EXC_DAQ	1024	4096
H0:PEM-PSL2_ACCZ	8	1024	H1:LSC-MC_L	32	2048
H0:PEM-PSL2_MIC	8	1024	H1:LSC-MICH_CTRL	32	2048
H0:PEM-RACK_1Y22_MAGZ	1	1024			

H1:LSC-POB_I	32	2048	H1:SUS-ITMX_COIL_LR	8	1024
H1:LSC-POB_I	1024	4096	H1:SUS-ITMX_COIL_UL	8	1024
H1:LSC-POB_Q	32	2048	H1:SUS-ITMX_COIL_UR	8	1024
H1:LSC-POB_Q	1024	4096	H1:SUS-ITMX_OPLEV_PERROR	8	256
H1:LSC-POBS_DC	8	512	H1:SUS-ITMX_OPLEV_YERROR	8	256
H1:LSC-PRC_CTRL	32	2048	H1:SUS-ITMX_SENSOR_SIDE	8	256
H1:LSC-REFL_DC	8	512	H1:SUS-ITMX_SUSPIT_IN	8	32
H1:LSC-REFL_I	32	2048	H1:SUS-ITMX_SUSPOS_IN	8	32
H1:LSC-REFL_I	1024	4096	H1:SUS-ITMX_SUSYAW_IN	8	32
H1:LSC-REFL_Q	32	2048	H1:SUS-ITMY_COIL_LL	8	1024
H1:LSC-SPOB_I	8	128	H1:SUS-ITMY_COIL_LR	8	1024
H1:OMC-DUOTONE_OUT_DAQ	32	2048	H1:SUS-ITMY_COIL_UL	8	1024
H1:OMC-DUOTONE_OUT_DAQ	1024	4096	H1:SUS-ITMY_COIL_UR	8	1024
H1:OMC-HTR_DRV_OUT_DAQ	8	256	H1:SUS-ITMY_OPLEV_PERROR	8	256
H1:OMC-LSC_Q_OUT_DAQ	256	2048	H1:SUS-ITMY_OPLEV_YERROR	8	256
H1:OMC-NULLSTREAM_OUT_DAQ	32	2048	H1:SUS-ITMY_SENSOR_SIDE	8	256
H1:OMC-PD_SUM_OUT_DAQ	32	2048	H1:SUS-ITMY_SUSPIT_IN	8	32
H1:OMC-PD_SUM_OUT_DAQ	1024	4096	H1:SUS-ITMY_SUSPOS_IN	8	32
H1:OMC-PZT_LSC_OUT_DAQ	8	256	H1:SUS-ITMY_SUSYAW_IN	8	32
H1:OMC-PZT_VMON_AC_OUT_DAQ	32	2048	H1:SUS-MMT3_OPLEV_PERROR	8	256
H1:OMC-PZT_VMON_DC_OUT_DAQ	8	256	H1:SUS-MMT3_OPLEV_YERROR	8	256
H1:OMC-QPD1_P_OUT_DAQ	8	1024	H1:SUS-RM_OPLEV_PERROR	8	256
H1:OMC-QPD1_SUM_IN1_DAQ	8	1024	H1:SUS-RM_OPLEV_YERROR	8	256
H1:OMC-QPD1_Y_OUT_DAQ	8	1024	H1:SUS-RM_SENSOR_SIDE	8	256
H1:OMC-QPD2_P_OUT_DAQ	8	1024	H1:SUS-RM_SUSPIT_IN	8	32
H1:OMC-QPD2_SUM_IN1_DAQ	8	1024	H1:SUS-RM_SUSPOS_IN	8	32
H1:OMC-QPD2_Y_OUT_DAQ	8	1024	H1:SUS-RM_SUSYAW_IN	8	32
H1:OMC-QPD3_P_OUT_DAQ	8	1024	H1:TCS-ITMX_PD_ISS_OUT_AC	8	1024
H1:OMC-QPD3_SUM_IN1_DAQ	32	2048	H1:TCS-ITMY_PD_ISS_OUT_AC	8	1024
H1:OMC-QPD3_SUM_IN1_DAQ	1024	4096	L0:PEM-BSC1_ACCX	8	1024
H1:OMC-QPD3_Y_OUT_DAQ	8	1024	L0:PEM-BSC1_ACCY	8	1024
H1:OMC-QPD4_P_OUT_DAQ	8	1024	L0:PEM-BSC1_ACCZ	8	1024
H1:OMC-QPD4_SUM_IN1_DAQ	8	1024	L0:PEM-BSC2_ACCX	8	1024
H1:OMC-QPD4_Y_OUT_DAQ	8	1024	L0:PEM-BSC2_ACCY	8	1024
H1:OMC-TT1_SUSPIT_IN1_DAQ	8	512	L0:PEM-BSC2_ACCZ	8	1024
H1:OMC-TT1_SUSPOS_IN1_DAQ	8	512	L0:PEM-BSC3_ACCX	8	1024
H1:OMC-TT1_SUSYAW_IN1_DAQ	8	512	L0:PEM-BSC3_ACCY	8	1024
H1:OMC-TT2_SUSPIT_IN1_DAQ	8	512	L0:PEM-BSC3_ACCZ	8	1024
H1:OMC-TT2_SUSPOS_IN1_DAQ	8	512	L0:PEM-BSC4_ACCX	8	1024
H1:OMC-TT2_SUSYAW_IN1_DAQ	8	512	L0:PEM-BSC4_ACCY	8	1024
H1:PSL-FSS_MIXERM_F	8	1024	L0:PEM-BSC4_ACCZ	8	1024
H1:SUS-BS_OPLEV_PERROR	8	256	L0:PEM-BSC4_MIC	8	1024
H1:SUS-BS_OPLEV_YERROR	8	256	L0:PEM-BSC5_ACCX	8	1024
H1:SUS-BS_SENSOR_SIDE	8	256	L0:PEM-BSC5_ACCY	8	1024
H1:SUS-BS_SUSPIT_IN	8	32	L0:PEM-BSC5_ACCZ	8	1024
H1:SUS-BS_SUSPOS_IN	8	32	L0:PEM-BSC5_MIC	8	1024
H1:SUS-BS_SUSYAW_IN	8	32	L0:PEM-COIL_MAGX	1	1024
H1:SUS-ETMX_COIL_LL	8	1024	L0:PEM-COIL_MAGZ	1	1024
H1:SUS-ETMX_COIL_LR	8	1024	L0:PEM-EX_BAYMIC	8	1024
H1:SUS-ETMX_COIL_UL	8	1024	L0:PEM-EX_MAGX	1	1024
H1:SUS-ETMX_COIL_UR	8	1024	L0:PEM-EX_MAGY	1	1024
H1:SUS-ETMX_OPLEV_PERROR	8	256	L0:PEM-EX_MAGZ	1	1024
H1:SUS-ETMX_OPLEV_YERROR	8	256	L0:PEM-EX_SEISX	1	8
H1:SUS-ETMX_SENSOR_SIDE	8	256	L0:PEM-EX_SEISX	8	128
H1:SUS-ETMX_SUSPIT_IN	8	32	L0:PEM-EX_SEISY	1	8
H1:SUS-ETMX_SUSPOS_IN	8	32	L0:PEM-EX_SEISY	8	128
H1:SUS-ETMX_SUSYAW_IN	8	32	L0:PEM-EX_SEISZ	1	8
H1:SUS-ETMY_COIL_LL	8	1024	L0:PEM-EX_SEISZ	8	128
H1:SUS-ETMY_COIL_LR	8	1024	L0:PEM-EX_V1	8	1024
H1:SUS-ETMY_COIL_UL	8	1024	L0:PEM-EY_BAYMIC	8	1024
H1:SUS-ETMY_COIL_UR	8	1024	L0:PEM-EY_MAGX	1	1024
H1:SUS-ETMY_OPLEV_PERROR	8	256	L0:PEM-EY_MAGY	1	1024
H1:SUS-ETMY_OPLEV_YERROR	8	256	L0:PEM-EY_MAGZ	1	1024
H1:SUS-ETMY_SENSOR_SIDE	8	256	L0:PEM-EY_SEISX	1	8
H1:SUS-ETMY_SUSPIT_IN	8	32	L0:PEM-EY_SEISX	8	128
H1:SUS-ETMY_SUSPOS_IN	8	32	L0:PEM-EY_SEISY	1	8
H1:SUS-ETMY_SUSYAW_IN	8	32	L0:PEM-EY_SEISY	8	128
H1:SUS-ITMX_COIL_LL	8	1024	L0:PEM-EY_SEISZ	8	128

L0:PEM-EY_SEISZ	1	8	L1:ISI-OMC_DISPPF_H1_IN1.DAQ	8	1024
L0:PEM-EY_SEISZ	8	128	L1:ISI-OMC_DISPPF_V1_IN1.DAQ	8	1024
L0:PEM-EY_V1	8	1024	L1:ISI-OMC_GEOPF_H1_IN1.DAQ	8	1024
L0:PEM-HAM1_ACCX	8	1024	L1:ISI-OMC_GEOPF_H2_IN1.DAQ	8	1024
L0:PEM-HAM1_ACCZ	8	1024	L1:ISI-OMC_GEOPF_H3_IN1.DAQ	8	1024
L0:PEM-HAM2_ACCX	8	1024	L1:ISI-OMC_GEOPF_V1_IN1.DAQ	8	1024
L0:PEM-HAM2_ACCZ	8	1024	L1:ISI-OMC_GEOPF_V2_IN1.DAQ	8	1024
L0:PEM-HAM6_ACCX	8	1024	L1:ISI-OMC_GEOPF_V3_IN1.DAQ	8	1024
L0:PEM-HAM6_ACCY	8	1024	L1:LSC-AS_AC	8	1024
L0:PEM-HAM6_ACCZ	8	1024	L1:LSC-AS_DC	8	512
L0:PEM-HAM6_MIC	8	1024	L1:LSC-DARM_CTRL	32	2048
L0:PEM-ISCT1_ACCX	8	1024	L1:LSC-DARM_CTRL	1024	4096
L0:PEM-ISCT1_ACCY	8	1024	L1:LSC-DARM_CTRL_EXC.DAQ	32	2048
L0:PEM-ISCT1_ACCZ	8	1024	L1:LSC-DARM_CTRL_EXC.DAQ	1024	4096
L0:PEM-ISCT1_MIC	8	1024	L1:LSC-DARM_ERR	32	2048
L0:PEM-LVEA_BAYMIC	8	1024	L1:LSC-DARM_ERR	1024	4096
L0:PEM-LVEA_MAGX	1	1024	L1:LSC-ETMX_CAL	32	2048
L0:PEM-LVEA_MAGY	1	1024	L1:LSC-ETMX_CAL	1024	4096
L0:PEM-LVEA_MAGZ	1	1024	L1:LSC-ETMX_CAL_EXC.DAQ	32	2048
L0:PEM-LVEA_MIC	32	2048	L1:LSC-ETMX_CAL_EXC.DAQ	1024	4096
L0:PEM-LVEA_MIC	1024	4096	L1:LSC-ETMX_EXC.DAQ	32	2048
L0:PEM-LVEA_SEISX	1	8	L1:LSC-ETMX_EXC.DAQ	1024	4096
L0:PEM-LVEA_SEISX	8	128	L1:LSC-ETMY_CAL	32	2048
L0:PEM-LVEA_SEISY	1	8	L1:LSC-ETMY_CAL	1024	4096
L0:PEM-LVEA_SEISY	8	128	L1:LSC-ETMY_CAL_EXC.DAQ	32	2048
L0:PEM-LVEA_SEISZ	1	8	L1:LSC-ETMY_CAL_EXC.DAQ	1024	4096
L0:PEM-LVEA_SEISZ	8	128	L1:LSC-MC_L	32	2048
L0:PEM-PSL1_ACCX	8	1024	L1:LSC-MICH_CTRL	32	2048
L0:PEM-PSL1_ACCY	8	1024	L1:LSC-POB_I	32	2048
L0:PEM-PSL1_ACCZ	8	1024	L1:LSC-POB_I	1024	4096
L0:PEM-PSL1_MIC	8	1024	L1:LSC-POB_Q	32	2048
L0:PEM-RADIO_LVEA	8	1024	L1:LSC-POB_Q	1024	4096
L0:PEM-RADIO_ROOF	8	1024	L1:LSC-PRC_CTRL	32	2048
L1:ASC-BS_P	8	256	L1:LSC-REFL_AC	8	1024
L1:ASC-BS_Y	8	256	L1:LSC-REFL_DC	8	128
L1:ASC-ETMX_P	8	256	L1:LSC-REFL_I	32	2048
L1:ASC-ETMX_Y	8	256	L1:LSC-REFL_I	1024	4096
L1:ASC-ETMY_P	8	256	L1:LSC-REFL_Q	32	2048
L1:ASC-ETMY_Y	8	256	L1:LSC-SPOB_I	8	128
L1:ASC-ITMX_P	8	256	L1:OMC-ASC_ANG_Y_IN1.DAQ	8	1024
L1:ASC-ITMX_Y	8	256	L1:OMC-ASC_POS_X_IN1.DAQ	8	1024
L1:ASC-ITMY_P	8	256	L1:OMC-ASC_POS_Y_IN1.DAQ	8	1024
L1:ASC-ITMY_Y	8	256	L1:OMC-DUOTONE_OUT.DAQ	32	2048
L1:ASC-QPDX_DC	8	128	L1:OMC-DUOTONE_OUT.DAQ	1024	4096
L1:ASC-QPDX_P	8	128	L1:OMC-NULLSTREAM_OUT.DAQ	32	2048
L1:ASC-QPDX_Y	8	128	L1:OMC-NULLSTREAM_OUT.DAQ	1024	4096
L1:ASC-QPDY_DC	8	128	L1:OMC-PD_SUM_OUT.DAQ	32	2048
L1:ASC-QPDY_P	8	128	L1:OMC-PD_SUM_OUT.DAQ	1024	4096
L1:ASC-QPDY_Y	8	128	L1:OMC-PZT_LSC_OUT.DAQ	8	1024
L1:ASC-RM_P	8	256	L1:OMC-PZT_VMON_AC_OUT.DAQ	32	2048
L1:ASC-RM_Y	8	256	L1:OMC-PZT_VMON_DC_OUT.DAQ	8	256
L1:ASC-WFS1_QP	8	256	L1:OMC-QPD1_P_OUT.DAQ	32	2048
L1:ASC-WFS1_QY	8	256	L1:OMC-QPD1_SUM_OUT.DAQ	32	2048
L1:ASC-WFS2_IP	8	256	L1:OMC-QPD1_Y_OUT.DAQ	32	2048
L1:ASC-WFS2_IY	8	256	L1:OMC-QPD2_P_OUT.DAQ	32	2048
L1:ASC-WFS2_QP	8	256	L1:OMC-QPD2_SUM_OUT.DAQ	32	2048
L1:ASC-WFS2_QY	8	256	L1:OMC-QPD2_Y_OUT.DAQ	32	2048
L1:ASC-WFS3_IP	8	256	L1:OMC-QPD3_P_OUT.DAQ	8	1024
L1:ASC-WFS3_IY	8	256	L1:OMC-QPD3_SUM_IN1.DAQ	8	1024
L1:ASC-WFS4_IP	8	256	L1:OMC-QPD3_Y_OUT.DAQ	8	1024
L1:ASC-WFS4_IY	8	256	L1:OMC-QPD4_P_OUT.DAQ	8	1024
L1:IOO-MC_F	32	2048	L1:OMC-QPD4_SUM_IN1.DAQ	8	1024
L1:ISI-OMC_CONT_RX_IN1.DAQ	8	1024	L1:OMC-QPD4_Y_OUT.DAQ	8	1024
L1:ISI-OMC_CONT_RY_IN1.DAQ	8	1024	L1:OMC-TT1_SUSPIT_IN1.DAQ	8	64
L1:ISI-OMC_CONT_RZ_IN1.DAQ	8	1024	L1:OMC-TT1_SUSPOS_IN1.DAQ	8	64
L1:ISI-OMC_CONT_X_IN1.DAQ	8	1024	L1:OMC-TT1_SUSYAW_IN1.DAQ	8	64
L1:ISI-OMC_CONT_Y_IN1.DAQ	8	1024	L1:OMC-TT2_SUSPIT_IN1.DAQ	8	64
L1:ISI-OMC_CONT_Z_IN1.DAQ	8	1024			

L1:OMC-TT2.SUSPOS.IN1.DAQ	8	64	L1:SEI-OUT.RY	8	128
L1:OMC-TT2.SUSYAW.IN1.DAQ	8	64	L1:SEI-OUT.RZ	8	128
L1:PSL-FSS.MIXERM.F	8	1024	L1:SEI-OUT.X	8	128
L1:SEI-BS.RX	8	128	L1:SEI-OUT.Y	8	128
L1:SEI-BS.RY	8	128	L1:SEI-RM.RX	8	128
L1:SEI-BS.RZ	8	128	L1:SEI-RM.RY	8	128
L1:SEI-BS.X	8	128	L1:SEI-RM.RZ	8	128
L1:SEI-BS.Y	8	128	L1:SEI-RM.X	8	128
L1:SEI-BS.Z	8	128	L1:SEI-RM.Y	8	128
L1:SEI-ETMX.RX	8	128	L1:SEI-RM.Z	8	128
L1:SEI-ETMX.RY	8	128	L1:SUS-BS.OPLEV.PERROR	8	256
L1:SEI-ETMX.RZ	8	128	L1:SUS-BS.OPLEV.YERROR	8	256
L1:SEI-ETMX.STS2.X	8	256	L1:SUS-BS.SENSOR.SIDE	8	32
L1:SEI-ETMX.X	8	128	L1:SUS-BS.SUSPIT.IN	8	32
L1:SEI-ETMX.Y	8	128	L1:SUS-BS.SUSPOS.IN	8	32
L1:SEI-ETMX.Z	8	128	L1:SUS-BS.SUSYAW.IN	8	32
L1:SEI-ETMY.RX	8	128	L1:SUS-ETMX.OPLEV.PERROR	8	256
L1:SEI-ETMY.RY	8	128	L1:SUS-ETMX.OPLEV.YERROR	8	256
L1:SEI-ETMY.RZ	8	128	L1:SUS-ETMX.SENSOR.SIDE	8	256
L1:SEI-ETMY.STS2.Y	8	256	L1:SUS-ETMX.SUSPIT.IN	8	32
L1:SEI-ETMY.X	8	128	L1:SUS-ETMX.SUSPOS.IN	8	32
L1:SEI-ETMY.Y	8	128	L1:SUS-ETMX.SUSYAW.IN	8	32
L1:SEI-ETMY.Z	8	128	L1:SUS-ETMY.OPLEV.PERROR	8	256
L1:SEI-ITMX.RX	8	128	L1:SUS-ETMY.OPLEV.YERROR	8	256
L1:SEI-ITMX.RY	8	128	L1:SUS-ETMY.SENSOR.SIDE	8	256
L1:SEI-ITMX.RZ	8	128	L1:SUS-ETMY.SUSPIT.IN	8	32
L1:SEI-ITMX.X	8	128	L1:SUS-ETMY.SUSPOS.IN	8	32
L1:SEI-ITMX.Y	8	128	L1:SUS-ETMY.SUSYAW.IN	8	32
L1:SEI-ITMX.Z	8	128	L1:SUS-ITMX.OPLEV.PERROR	8	256
L1:SEI-ITMY.RX	8	128	L1:SUS-ITMX.OPLEV.YERROR	8	256
L1:SEI-ITMY.RY	8	128	L1:SUS-ITMX.SENSOR.SIDE	8	32
L1:SEI-ITMY.RZ	8	128	L1:SUS-ITMX.SUSPIT.IN	8	32
L1:SEI-ITMY.X	8	128	L1:SUS-ITMX.SUSPOS.IN	8	32
L1:SEI-ITMY.Y	8	128	L1:SUS-ITMX.SUSYAW.IN	8	32
L1:SEI-ITMY.Z	8	128	L1:SUS-ITMY.OPLEV.PERROR	8	256
L1:SEI-LVEA.STS2.X	8	256	L1:SUS-ITMY.OPLEV.YERROR	8	256
L1:SEI-LVEA.STS2.Y	8	256	L1:SUS-ITMY.SENSOR.SIDE	8	32
L1:SEI-LVEA.STS2.Z	8	256	L1:SUS-ITMY.SUSPIT.IN	8	32
L1:SEI-MC1.RX	8	128	L1:SUS-ITMY.SUSPOS.IN	8	32
L1:SEI-MC1.RY	8	128	L1:SUS-ITMY.SUSYAW.IN	8	32
L1:SEI-MC1.RZ	8	128	L1:SUS-MMT3.OPLEV.PERROR	8	256
L1:SEI-MC1.X	8	128	L1:SUS-MMT3.OPLEV.YERROR	8	256
L1:SEI-MC1.Y	8	128	L1:SUS-RM.OPLEV.PERROR	8	256
L1:SEI-MC1.Z	8	128	L1:SUS-RM.OPLEV.YERROR	8	256
L1:SEI-MC2.RX	8	128	L1:SUS-RM.SENSOR.SIDE	8	32
L1:SEI-MC2.RY	8	128	L1:SUS-RM.SUSPIT.IN	8	32
L1:SEI-MC2.RZ	8	128	L1:SUS-RM.SUSPOS.IN	8	32
L1:SEI-MC2.X	8	128	L1:SUS-RM.SUSYAW.IN	8	32
L1:SEI-MC2.Y	8	128	L1:TCS-ITMX.PD_ISS.OUT.AC	8	1024
L1:SEI-MC2.Z	8	128	L1:TCS-ITMY.PD_ISS.OUT.AC	8	1024
L1:SEI-OUT.RX	8	128			

Appendix C

Auxiliary channel veto statistics

KleineWelle auxiliary (AUX) channel individual veto performance on noise transients also identified by kleineWelle running on the gravitational-wave (GW) channel during the second calendar year of S5. For each auxiliary channel, the configuration with the highest statistical significance is reported if it is greater than 10. Columns are,

channel	auxiliary channel transient source
f_{low}	lower frequency for AUX trigger generation
f_{high}	upper frequency for AUX trigger generation
thr	kleineWelle significance threshold for AUX triggers used for veto
$\pm ms$	veto window in milliseconds between AUX and GW-channel transient
dt [%]	fractional dead-time incurred by veto
ϵ_{35} [%]	veto efficiency on GW noise transients with trigger significance > 35
ϵ_{100} [%]	veto efficiency on GW noise transients with trigger significance > 100
signif	statistical significance of veto correlation ($-\ln P$)

C.0.1 Veto performance on H1 transients

channel	f_{low}	f_{high}	thr	$\pm ms$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H1:LSC-AS_AC	64	1024	200	25	0.01	39.89	31.81	141451
H1:LSC-REFL_DC	64	1024	35	25	0.02	39.15	32.02	132768
H1:LSC-PRC_CTRL	64	1024	50	25	0.02	38.98	34.55	130209
H1:LSC-POB_I	64	1024	50	25	0.02	38.91	34.51	130101
H1:LSC-MC_L	64	1024	200	25	0.02	37.23	30.93	121651
H1:LSC-REFL_I	64	1024	200	25	0.02	37.16	30.74	121246
H1:LSC-REFL_Q	64	1024	35	25	0.03	36.91	32.13	116327
H1:LSC-POB_Q	64	1024	50	25	0.03	37.04	31.93	115591
H1:LSC-MICH_CTRL	64	1024	50	25	0.03	36.95	31.81	115330
H1:ASC-ETMX_P	4	256	400	50	0.01	3.76	0.69	9348
H1:ASC-ETMY_P	4	256	400	50	0.01	3.79	0.62	9241

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H1:ASC-ETMX_Y	4	256	50	200	0.02	1.25	0.65	2016
H1:LSC-POBJ	1024	4096	35	50	0.00	0.53	1.54	1900
H1:LSC-PRC_CTRL	1024	2048	35	200	0.00	0.63	1.47	1872
H1:ASC-ITMX_Y	4	256	50	200	0.02	1.21	0.71	1861
H1:ASC-WFS1_QP	4	256	100	200	0.01	0.88	0.86	1855
H1:ASC-ETMY_Y	4	256	50	200	0.02	1.07	0.54	1704
H1:ASC-ITMX_P	4	256	50	200	0.01	0.98	0.66	1651
H1:ASC-ITMY_Y	4	256	50	200	0.02	1.10	0.56	1645
H1:ASC-RM_Y	4	256	50	200	0.02	1.09	0.65	1615
H1:ASC-ITMY_P	4	256	50	200	0.01	0.94	0.65	1547
H1:ASC-RM_P	4	256	100	200	0.00	0.68	0.46	1471
H1:ASC-QPDY_DC	4	256	35	200	0.07	1.33	2.72	1339
H1:ASC-WFS1_QY	4	256	100	200	0.02	0.85	0.66	1299
H1:LSC-REFLJ	1024	4096	35	150	0.00	0.57	0.71	1208
H1:ASC-WFS2_QP	4	256	50	200	0.01	0.69	0.50	1206
H1:ASC-WFS2_IY	4	256	100	200	0.00	0.54	0.20	1191
H1:ASC-WFS3_IY	4	256	35	200	0.11	1.38	1.47	1144
H1:ASC-WFS2_IP	4	256	100	200	0.00	0.44	0.35	1044
H1:LSC-MICH_CTRL	1024	2048	35	200	0.00	0.40	0.93	1004
H1:ASC-WFS2_QY	4	256	50	200	0.07	1.01	0.65	924
H1:ASC-WFS3_IP	4	256	100	200	0.00	0.46	0.52	918
H1:LSC-POB_Q	1024	4096	35	50	0.00	0.27	0.83	828
H1:ASC-QPDY_P	4	256	35	200	0.02	0.62	1.26	760
H1:ASC-QPDY_Y	4	256	35	200	0.01	0.39	0.76	601
H1:LSC-REFL_Q	1024	2048	35	100	0.00	0.15	0.46	527
H1:ASC-WFS4_IY	4	256	50	200	0.04	0.46	0.47	360
H1:ASC-QPDX_DC	4	256	50	200	0.03	0.39	0.51	320
H1:ASC-WFS4_IP	4	256	50	200	0.07	0.52	0.50	299
H1:ASC-BS_Y	4	256	35	200	0.00	0.18	0.12	297
H1:ASC-QPDX_Y	4	256	400	200	0.00	0.09	0.07	278
H1:ASC-QPDX_P	4	256	100	200	0.00	0.13	0.08	192
H1:ASC-BS_P	4	256	35	200	0.01	0.13	0.12	154
H1:SUS-ITMX_OPLEV_YERROR	4	256	35	200	0.00	0.08	0.06	149
H1:SUS-RM_OPLEV_YERROR	4	256	35	200	0.01	0.08	0.02	50
H0:PEM-ISCT4_ACCZ	10	512	1600	25	0.00	0.01	0.01	47
H0:PEM-ISCT4_ACCX	10	512	800	25	0.00	0.01	0.01	46
H0:PEM-ISCT4_ACCY	10	512	400	25	0.00	0.01	0.01	42
H1:SUS-ETMX_OPLEV_YERROR	4	256	35	200	0.00	0.03	0.01	32
H0:PEM-RADIO_LVEA	10	512	50	200	0.06	0.16	0.10	32
H1:SUS-ITMX_OPLEV_PERROR	4	256	35	200	0.00	0.03	0.01	28
H0:PEM-LVEA_MIC	10	512	35	200	0.00	0.03	0.01	24
H1:TCS-ITMY_PD1AC	4	256	35	200	0.53	0.74	0.52	23
H1:SUS-ETMX_OPLEV_PERROR	4	256	35	200	0.01	0.04	0.02	23
H0:PEM-MX_V2	10	512	35	200	4.13	4.68	4.66	21
H0:PEM-BSC9_ACC1X	10	512	400	200	0.00	0.01	0.01	21
H0:PEM-ISCT4_MIC	10	512	200	25	0.00	0.01	0.01	19
H0:PEM-ISCT7_ACCZ	10	512	35	200	0.04	0.11	0.10	18
H0:PEM-BSC7_MIC	10	512	35	200	0.00	0.03	0.01	18
H0:PEM-PSL1_ACCX	10	512	35	200	0.01	0.05	0.02	18
H0:PEM-ISCT7_ACCX	10	512	35	200	0.02	0.07	0.07	17
H1:IOO-MC_F	64	1024	800	25	0.00	0.01	0.01	17
H0:PEM-ISCT1_ACCX	10	512	35	200	0.03	0.08	0.03	16

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H0:PEM-BSC9_MIC	10	512	50	200	0.00	0.02	0.02	15
H0:PEM-LVEA2_V1	10	512	100	200	0.02	0.05	0.10	15
H0:PEM-HAM9_ACCX	10	512	50	150	0.00	0.03	0.01	15
H0:PEM-ISCT7_ACCY	10	512	35	200	0.03	0.07	0.09	14
H0:PEM-EX_SEISZ	10	512	50	200	0.00	0.02	0.04	14
H0:PEM-BSC9_MAGX	10	512	35	50	0.00	0.01	0.02	14
H1:SUS-MMT3_OPLEV_PERROR	4	256	50	200	0.00	0.02	0.00	13
H0:PEM-LVEA_SEISY	10	512	35	200	0.00	0.02	0.01	13
H1:TCS-ITMY_PD2AC	4	256	400	100	0.00	0.00	0.01	13
H1:SUS-BS_OPLEV_PERROR	4	256	35	150	0.00	0.01	0.01	12
H0:PEM-ISCT1_ACCY	10	512	100	200	0.00	0.01	0.00	12
H0:PEM-HAM7_ACCZ	10	512	35	200	0.01	0.03	0.01	12
H0:PEM-BSC5_ACCX	10	512	50	200	0.01	0.04	0.04	12
H1:IOO-MC_F	1024	2048	400	25	0.00	0.01	0.03	11
H0:PEM-COIL_MAGZ	10	512	35	200	4.74	5.14	5.13	11
H0:PEM-PSL2_ACCZ	10	512	35	200	0.03	0.07	0.05	11
H0:PEM-HAM7_ACCX	10	512	35	200	0.01	0.02	0.01	10
H0:PEM-ISCT10_ACCX	10	512	50	200	0.12	0.18	0.23	10
H0:PEM-LVEA2_V3	10	512	200	200	0.02	0.04	0.05	10
H0:PEM-EX_SEISY	10	512	100	200	0.00	0.01	0.00	10

C.0.2 Veto performance on H2 transients

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H2:LSC-PRC_CTRL	64	1024	50	25	0.00	28.90	35.91	62487
H2:LSC-POBJ	64	1024	50	25	0.00	28.66	35.69	61867
H2:LSC-REFL_Q	64	1024	35	150	0.01	16.51	18.81	28677
H2:LSC-MICH_CTRL	64	1024	35	200	0.03	17.17	17.98	26049
H2:LSC-POB_Q	64	1024	35	200	0.03	17.03	17.98	24915
H2:LSC-AS_AC	64	1024	35	200	0.02	4.92	5.21	6291
H2:LSC-POBJ	1024	4096	35	200	0.00	3.10	4.46	5875
H2:ASC-QPDY_Y	4	256	35	150	0.01	3.73	2.03	5302
H2:ASC-ETMX_P	4	256	100	200	0.08	5.78	3.64	5249
H2:ASC-ITMX_P	4	256	200	200	0.02	4.35	2.57	5023
H2:ASC-ETMY_P	4	256	200	200	0.02	3.81	2.23	4386
H2:ASC-ETMY_Y	4	256	200	200	0.03	3.91	2.52	4382
H2:ASC-WFS2_QP	4	256	200	200	0.01	3.57	2.18	4378
H2:ASC-ITMX_Y	4	256	200	200	0.01	3.55	2.02	4355
H2:ASC-ETMX_Y	4	256	100	200	0.05	4.42	2.46	4329
H2:ASC-QPDY_P	4	256	35	200	0.03	4.14	2.26	4320
H2:ASC-ITMY_P	4	256	200	200	0.04	4.18	2.62	4245
H2:ASC-ITMY_Y	4	256	200	200	0.04	4.15	2.77	4109
H2:ASC-QPDY_DC	4	256	50	200	0.03	3.74	2.44	4095
H2:LSC-PRC_CTRL	1024	2048	35	200	0.00	2.12	3.00	4014
H2:ASC-WFS2_QY	4	256	200	200	0.01	3.02	1.79	3599
H2:ASC-WFS1_QY	4	256	100	200	0.04	3.61	2.29	3477
H2:ASC-WFS1_QP	4	256	100	200	0.11	4.57	3.05	3401
H2:ASC-RM_Y	4	256	100	200	0.06	3.86	2.73	3373
H2:ASC-WFS4_IY	4	256	35	200	0.01	2.56	1.84	3289

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H2:SUS-ITMX_OPLEV_YERROR	4	256	35	150	0.01	2.30	1.47	3238
H2:ASC-QPDX_P	4	256	35	200	0.03	3.13	1.82	3073
H2:SUS-ITMX_OPLEV_PERROR	4	256	35	150	0.01	2.29	1.39	3017
H2:ASC-WFS3_IY	4	256	35	200	0.04	3.27	2.44	2981
H2:ASC-WFS4_IP	4	256	35	200	0.01	2.43	1.77	2825
H2:SUS-ITMY_OPLEV_PERROR	4	256	35	200	0.01	2.38	1.48	2769
H2:ASC-BS_P	4	256	100	200	0.06	3.07	1.45	2570
H2:ASC-WFS2_IP	4	256	50	200	0.72	6.85	6.33	2552
H2:SUS-ITMY_OPLEV_YERROR	4	256	35	150	0.01	2.13	1.42	2526
H2:LSC-REFL_AC	64	1024	35	25	0.00	1.48	1.51	2489
H2:ASC-WFS2_IY	4	256	100	200	0.10	3.45	2.49	2477
H2:SUS-ETMY_OPLEV_YERROR	4	256	35	200	0.00	1.64	1.24	2460
H2:ASC-RM_P	4	256	100	200	0.03	2.59	1.77	2432
H2:SUS-ETMX_OPLEV_PERROR	4	256	35	150	0.02	2.25	1.52	2352
H2:ASC-WFS3_IP	4	256	35	200	0.01	1.79	1.38	2027
H2:SUS-ETMX_OPLEV_YERROR	4	256	35	200	0.01	1.71	1.21	2000
H2:LSC-MC_L	64	1024	200	25	0.00	1.08	0.93	1986
H2:LSC-REFL_I	64	1024	200	25	0.00	1.07	0.94	1962
H2:SUS-ETMY_OPLEV_PERROR	4	256	35	200	0.02	1.96	1.31	1835
H2:ASC-BS_Y	4	256	50	200	0.11	2.26	1.22	1300
H2:IOO-MC_F	64	1024	400	25	0.00	0.57	0.40	1083
H2:LSC-REFL_I	1024	4096	35	25	0.00	0.44	0.36	832
H2:TCS-ITMY_PD2AC	4	256	400	100	0.00	0.60	0.13	755
H2:IOO-MC_F	1024	2048	100	25	0.00	0.60	0.43	755
H2:LSC-MC_L	1024	2048	35	25	0.00	0.36	0.32	688
H2:ASC-QPDX_Y	4	256	200	200	0.00	0.46	0.27	683
H2:LSC-REFL_Q	1024	4096	35	25	0.00	0.24	0.35	548
H2:TCS-ITMY_PD1AC	4	256	400	100	0.01	0.64	0.16	538
H2:ASC-QPDX_DC	4	256	400	200	0.00	0.21	0.18	347
H2:LSC-POB_Q	1024	4096	35	25	0.00	0.15	0.24	345
H2:LSC-MICH_CTRL	1024	2048	35	25	0.00	0.07	0.11	164
H2:SUS-RM_OPLEV_YERROR	4	256	35	150	0.00	0.12	0.18	148
H0:PEM-ISCT4_ACCY	10	512	1600	200	0.00	0.08	0.03	61
H0:PEM-ISCT4_ACCX	10	512	1600	200	0.00	0.08	0.03	59
H0:PEM-ISCT4_ACCZ	10	512	1600	200	0.00	0.08	0.03	57
H0:PEM-BSC5_ACCX	10	512	800	200	0.00	0.02	0.01	34
H0:PEM-PSL2_ACCZ	10	512	200	50	0.00	0.02	0.01	34
H2:SUS-RM_OPLEV_PERROR	4	256	35	100	0.00	0.05	0.07	32
H0:PEM-MX_SEISY	10	512	400	200	0.00	0.03	0.01	31
H0:PEM-ISCT10_ACCY	10	512	35	200	0.00	0.04	0.01	30
H0:PEM-MX_SEISZ	10	512	800	150	0.00	0.01	0.00	27
H0:PEM-MX_SEISX	10	512	800	200	0.00	0.01	0.00	24
H0:PEM-ISCT4_MIC	10	512	35	200	0.01	0.06	0.02	22
H2:SUS-FMX_OPLEV_PERROR	4	256	50	200	0.00	0.03	0.03	21
H0:PEM-BSC6_MAGZ	10	512	50	200	1.17	1.58	1.91	20
H0:PEM-BSC5_MIC	10	512	35	200	0.01	0.06	0.04	18
H0:PEM-PSL2_ACCX	10	512	50	200	0.01	0.06	0.05	17
H0:PEM-MY_SEISX	10	512	35	200	0.00	0.04	0.02	16
H0:PEM-BSC10_MAGZ	10	512	50	100	0.10	0.19	0.26	13
H2:SUS-FMY_OPLEV_PERROR	4	256	50	200	0.00	0.01	0.01	12
H0:PEM-MY_SEISY	10	512	200	200	0.00	0.01	0.01	12
H0:PEM-MY_SEISZ	10	512	200	200	0.00	0.01	0.01	10

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
H2:SUS-MMT3_OPLEV_PERROR	4	256	35	200	0.00	0.02	0.01	10

C.0.3 Veto performance on L1 transients

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
L1:LSC-PRC_CTRL	64	1024	35	200	0.08	7.48	8.69	21864
L1:LSC-MICH_CTRL	64	1024	35	100	0.03	4.92	6.14	16964
L1:LSC-POBJ	64	1024	35	25	0.01	3.63	4.77	14773
L1:LSC-POB_Q	64	1024	35	200	0.05	4.79	6.09	14066
L1:ASC-WFS2_QY	4	256	50	200	0.03	3.40	3.07	10514
L1:ASC-WFS2_IY	4	256	100	200	0.01	2.88	3.01	10223
L1:ASC-RM_Y	4	256	50	200	0.04	3.32	3.15	9656
L1:ASC-ITMY_Y	4	256	50	200	0.05	3.41	3.19	9279
L1:ASC-ITMX_Y	4	256	50	200	0.05	3.42	3.23	8982
L1:ASC-ITMY_P	4	256	50	200	0.05	3.11	2.92	8030
L1:ASC-ITMX_P	4	256	50	200	0.05	3.10	2.91	7945
L1:ASC-WFS2_IP	4	256	50	200	0.04	2.85	3.16	7586
L1:ASC-WFS1_QY	4	256	100	200	0.03	2.46	2.58	7250
L1:ASC-RM_P	4	256	35	200	0.13	3.60	3.36	7034
L1:ASC-WFS2_QP	4	256	35	200	0.04	2.55	2.79	6777
L1:LSC-PRC_CTRL	1024	2048	35	200	0.00	1.05	1.64	4906
L1:LSC-REFL_Q	64	1024	35	200	0.07	2.29	3.21	4815
L1:ASC-WFS3_IY	4	256	35	200	0.02	1.39	1.69	3963
L1:LSC-REFL_I	64	1024	800	25	0.00	0.62	1.17	3209
L1:ASC-QPDX_Y	4	256	100	200	0.00	0.69	0.97	2901
L1:ASC-BS_Y	4	256	35	200	0.01	0.67	0.83	2231
L1:ASC-WFS1_QP	4	256	400	200	0.00	0.46	0.75	2208
L1:ASC-ETMY_Y	4	256	200	200	0.00	0.45	0.67	2130
L1:ASC-ETMX_Y	4	256	200	200	0.00	0.36	0.61	1721
L1:LSC-POBJ	1024	4096	35	25	0.00	0.29	0.54	1385
L0:PEM-LVEA_MAGY	10	512	200	25	0.00	0.31	0.58	1372
L1:ASC-BS_P	4	256	35	200	0.01	0.48	0.62	1305
L1:LSC-MICH_CTRL	1024	2048	35	200	0.00	0.35	0.62	1241
L1:LSC-REFL_I	1024	4096	100	25	0.00	0.20	0.50	926
L1:ASC-QPDX_P	4	256	100	200	0.00	0.22	0.36	794
L1:ASC-QPDY_Y	4	256	50	200	0.01	0.35	0.61	714
L1:SUS-ETMY_OPLEV_YERROR	4	256	35	200	0.00	0.24	0.34	693
L1:LSC-AS_AC	64	1024	35	200	0.04	0.47	0.99	624
L1:LSC-POB_Q	1024	4096	35	25	0.00	0.12	0.30	539
L1:SUS-ITMX_OPLEV_YERROR	4	256	35	200	0.00	0.16	0.25	507
L1:ASC-QPDY_DC	4	256	100	200	0.03	0.33	0.53	468
L1:ASC-ETMX_P	4	256	400	200	0.00	0.12	0.16	417
L1:ASC-WFS4_IY	4	256	35	200	0.02	0.27	0.16	417
L1:ASC-WFS3_IP	4	256	35	200	0.14	0.60	0.53	350
L1:SUS-ITMY_OPLEV_YERROR	4	256	35	200	0.01	0.16	0.27	339
L1:ASC-QPDY_P	4	256	50	200	0.01	0.20	0.32	327
L1:ASC-ETMY_P	4	256	400	200	0.00	0.10	0.10	326
L1:LSC-REFL_AC	64	1024	50	200	0.08	0.43	0.31	314
L0:PEM-LVEA_MAGZ	10	512	200	25	0.03	0.28	0.52	303

channel	f_{low}	f_{high}	thr	$\pm\text{ms}$	dt[%]	ϵ_{35}	ϵ_{100}	signif
L1:ASC-WFS4_IP	4	256	400	200	0.03	0.25	0.08	271
L0:PEM-LVEA_MAGY	0.5	32	100	200	0.02	0.18	0.46	217
L0:PEM-LVEA_MAGX	10	512	50	25	0.01	0.14	0.49	210
L1:LSC-REFL_Q	1024	2048	35	25	0.00	0.04	0.17	209
L1:SUS-MMT3_OPLEV_YERROR	4	256	35	200	1.15	1.90	1.99	175
L1:SUS-MMT3_OPLEV_PERROR	4	256	50	200	0.99	1.66	1.66	157
L0:PEM-BSC4_ACCZ	10	512	35	200	0.07	0.29	0.12	144
L1:IOO-MC_F	64	1024	35	50	0.73	1.19	1.22	105
L1:IOO-MC_F	1024	2048	35	50	0.54	0.91	0.86	89
L1:SUS-ETMY_OPLEV_PERROR	4	256	800	200	0.00	0.02	0.01	73
L1:SUS-ETMX_OPLEV_YERROR	4	256	35	200	0.01	0.09	0.12	72
L1:SUS-ITMX_OPLEV_PERROR	4	256	35	150	0.00	0.04	0.03	57
L1:ASC-QPDX_DC	4	256	400	200	0.00	0.01	0.02	50
L1:SUS-RM_OPLEV_YERROR	4	256	35	200	0.01	0.04	0.03	35
L0:PEM-HAM1_ACCX	10	512	35	200	0.01	0.03	0.01	34
L0:PEM-ISCT1_ACCY	10	512	35	200	0.02	0.06	0.03	24
L0:PEM-HAM1_ACCZ	10	512	35	200	0.02	0.06	0.02	22
L0:PEM-BSC4_ACCY	10	512	35	200	0.01	0.04	0.03	20
L0:PEM-BSC1_ACCX	10	512	35	200	0.00	0.01	0.01	19
L1:SUS-BS_OPLEV_YERROR	4	256	35	200	0.01	0.04	0.01	19
L0:PEM-EY_MAGY	10	512	400	25	0.00	0.00	0.00	14
L0:PEM-BSC5_ACCX	10	512	50	100	0.01	0.03	0.01	13
L0:PEM-BSC2_ACCY	10	512	800	100	0.00	0.01	0.00	13
L0:PEM-HAM2_ACCX	10	512	50	200	0.00	0.01	0.00	13
L0:PEM-EY_MAGX	10	512	100	50	0.00	0.01	0.01	13
L0:PEM-LVEA_MAGZ	0.5	32	200	200	0.09	0.14	0.34	12
L0:PEM-EY_MAGZ	10	512	1600	150	0.00	0.00	0.01	12
L0:PEM-RADIO_LVEA	10	512	35	25	0.08	0.12	0.20	11
L0:PEM-ISCT1_ACCX	10	512	50	200	0.02	0.04	0.01	11
L1:SUS-BS_OPLEV_PERROR	4	256	35	200	0.01	0.02	0.01	10
L0:PEM-EX_SEISZ	10	512	35	100	0.00	0.02	0.02	10

Appendix D

Burst detection checklist

Zero-level sanity

- Record GPS/UTC time of event
- Read detector logbook for surrounding days
- Verify instrument state is correctly flagged
- Check baseline sensitivity and range of detectors
- Check times of nearest hardware signal injections
- How close is event to the end of a science segment?

Data acquisition and integrity

- Check timing across instruments
- Record nearest reboots or software/configuration changes
- Check hardware injection excitation channels
- Verify data checksum and consistency between RAW and reduced data sets
- Examine any other test points for stored data

Calibration

- Record calibration constants and errors about event
- Check robustness of event properties to differences in calibration
- Is the event seen in uncalibrated data?
- Check for glitches in calibration excitation channels

Event properties and methods

Construct detailed spectrograms of events in different detectors
Check for time-frequency consistency
Run Coherent Event display
Extract best-fit waveform and sky location
Compare coherent and incoherent extracted parameters
Check all null streams
Record which methods and detector networks saw the event
Check event amplitude and background for different methods
Are observations consistent with method and network sensitivity?
Compare parameter estimation from different methods
Verify method robustness to stride boundary, prefiltering, specific thresholds
Could the event arise from nonstationary lines?
What is the search method sensitivity at the time of the event?
Check signal content at all frequencies at time of event
Is anything seen in the inspiral or ringdown searches?

State of the instrument and instrumental vetoes

Check stationary of the instrument: power spectrum and transient rates
Check levels of upconversion of low frequency noise sources
Run Q-scan (nearby transients and spectrograms) on all channels in full frames
Check for associated instrumental and environmental disturbances
Which data quality flags were in effect at and near the time of the event?
Check records for earthquakes, trucks, trains, wind, storms, airplanes
Look at seismic spectrograms and rms levels at time of event
Check correlations with auxiliary channel transients (kleineWelle)
Establish the significance and safety of any correlation
Are auxiliary channel coincidences consistent with known transfer functions?
Check for fluctuations in thermal compensation laser power
Check for cosmic ray events

Verify signals in four photodiodes are the same (dust veto)
Check for abnormalities in channel trend data
Listen to gravitational-wave stream in audio about event time
Investigate any other possible non-astrophysical causes of the event

Background estimation

Run additional time-shifts on local data to measure local background
Establish how often single-detector waveform arises from background
Check event frequency, bandwidth against known background distribution
Check background isotropy for coherent methods at reconstructed sky location
Check for non-stationarities in background rate and properties near event
Estimate any trials factor, and establish global significance of the event

Other detectors

Any other interferometers or bars online, would they see the event?
Check GCN, IAU, and other public transient catalogs
Any known sub-prime EM or particle events?

Astrophysical analysis

Which sources overlap with the sky location?
How do the extracted waveforms compare with astrophysical waveforms?
Does the amplitude/energy scale make sense?
Characterize waveform morphology, is there a chirp?

Bibliography

- [1] J. Weber, “Detection and generation of gravitational waves,” *Phys. Rev.* **117** (1960) 306–313.
- [2] R. Weiss, “Electromagnetically coupled broadband gravitational antenna,” *Quarterly Progress Report of the Research Laboratory of Electronics* **105** (1972) 54.
<http://www.ligo.caltech.edu/docs/P/P720002-01/P720002-01.pdf>.
- [3] G. E. Moss, L. R. Miller, and R. L. Forward, “Photon-noise-limited laser transducer for gravitational antenna,” *Applied Optics* **10** (1971) 2495–2498.
<http://ao.osa.org/abstract.cfm?URI=ao-10-11-2495>.
- [4] <http://www.ligo.caltech.edu>.
- [5] <http://www.virgo.infn.it>.
- [6] R. A. Hulse and J. H. Taylor, “Discovery of a pulsar in a binary system,” *Astrophys. J.* **195** (1975) L51–L53.
- [7] J. M. Weisberg and J. H. Taylor, “Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis,” *ASP Conf. Ser.* **328** (2005) 25, [arXiv:astro-ph/0407149](https://arxiv.org/abs/astro-ph/0407149).
- [8] **The LIGO Scientific and The Virgo** Collaboration, J. Abadie *et al.*, “All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run,” *Phys. Rev.* **D81** (2010) 102001, [arXiv:1002.1036](https://arxiv.org/abs/1002.1036) [gr-qc].
- [9] C. Misner, K. Thorne, and J. Wheeler, *Gravitation*. W.H. Freeman, San Francisco, 1973.
- [10] S. M. Carroll, “Lecture notes on general relativity,” [arXiv:gr-qc/9712019](https://arxiv.org/abs/gr-qc/9712019).
- [11] R. A. Isaacson, “Gravitational Radiation in the Limit of High Frequency. 1. The Linear Approximation and Geometrical Optics,” *Phys. Rev.* **166** (1967) 1263–1271.
- [12] R. A. Isaacson, “Gravitational Radiation in the Limit of High Frequency. II. Nonlinear Terms and the Effective Stress Tensor,” *Phys. Rev.* **166** (1968) 1272–1279.

- [13] J. H. Taylor and J. M. Weisberg, “A new test of general relativity: Gravitational radiation and the binary pulsar PSR 1913+16,” *Astrophys. J.* **253** (1982) 908–920.
- [14] **LIGO Scientific** Collaboration, J. Abadie *et al.*, “Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors,” *Class. Quant. Grav.* **27** (2010) 173001, arXiv:1003.2480 [astro-ph].
- [15] E. Nakar, “Short-hard gamma-ray bursts,” *Phys. Rept.* **442** (2007) 166–236, arXiv:astro-ph/0701748.
- [16] E. Berger *et al.*, “The afterglow and elliptical host galaxy of the short gamma-ray burst GRB 050724,” *Nature* **438** (2005) 988–990, arXiv:astro-ph/0508115.
- [17] E. E. Flanagan and T. Hinderer, “Constraining neutron star tidal Love numbers with gravitational wave detectors,” *Phys. Rev.* **D77** (2008) 021502, arXiv:0709.1915 [astro-ph].
- [18] J. S. Read *et al.*, “Measuring the neutron star equation of state with gravitational wave observations,” *Phys. Rev.* **D79** (2009) 124033, arXiv:0901.3258 [gr-qc].
- [19] L. Blanchet, “Gravitational radiation from post-Newtonian sources and inspiralling compact binaries,” *Living Rev. Rel.* **5** (2002) 3, arXiv:gr-qc/0202016.
- [20] C. M. Will, “The confrontation between general relativity and experiment,” *Living Rev. Rel.* **9** (2005) 3, arXiv:gr-qc/0510072.
- [21] I. Hinder, “The Current Status of Binary Black Hole Simulations in Numerical Relativity,” *Class. Quant. Grav.* **27** (2010) 114004, arXiv:1001.5161 [gr-qc].
- [22] E. W. Leaver, “An Analytic representation for the quasi normal modes of Kerr black holes,” *Proc. Roy. Soc. Lond.* **A402** (1985) 285–298.
- [23] F. Echeverria, “Gravitational wave measurements of the mass and angular momentum of a black hole,” *Phys. Rev.* **D40** (1989) 3194–3203.
- [24] **LIGO Scientific** Collaboration, B. P. Abbott *et al.*, “Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO’s S5 Data,” *Phys. Rev.* **D79** (2009) 122001, arXiv:0901.0302 [gr-qc].
- [25] **LIGO Scientific** Collaboration, B. P. Abbott *et al.*, “Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO’s fifth Science Run,” *Phys. Rev.* **D80** (2009) 047101, arXiv:0905.3710 [gr-qc].

- [26] **LIGO Scientific** Collaboration, B. P. Abbott *et al.*, “Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data,” *Phys. Rev. D* **80** (2009) 062001, [arXiv:0905.1654 \[gr-qc\]](#).
- [27] E. E. Flanagan and S. A. Hughes, “Measuring gravitational waves from binary black hole coalescences. I: Signal to noise for inspiral, merger, and ringdown,” *Phys. Rev. D* **57** (1998) 4535–4565, [arXiv:gr-qc/9701039](#).
- [28] F. Pretorius, “Evolution of Binary Black Hole Spacetimes,” *Phys. Rev. Lett.* **95** (2005) 121101, [arXiv:gr-qc/0507014](#).
- [29] M. Campanelli, C. O. Lousto, H. Nakano, and Y. Zlochower, “Comparison of Numerical and Post-Newtonian Waveforms for Generic Precessing Black-Hole Binaries,” *Phys. Rev. D* **79** (2009) 084010, [arXiv:0808.0713 \[gr-qc\]](#).
- [30] P. Ajith *et al.*, “Phenomenological template family for black-hole coalescence waveforms,” *Class. Quant. Grav.* **24** (2007) S689–S700, [arXiv:0704.3764 \[gr-qc\]](#).
- [31] P. Ajith *et al.*, “‘Complete’ gravitational waveforms for black-hole binaries with non-precessing spins,” [arXiv:0909.2867 \[gr-qc\]](#).
- [32] A. Buonanno and T. Damour, “Effective one-body approach to general relativistic two-body dynamics,” *Phys. Rev. D* **59** (1999) 084006, [arXiv:gr-qc/9811091](#).
- [33] Y. Pan *et al.*, “Effective-one-body waveforms calibrated to numerical relativity simulations: coalescence of non-precessing, spinning, equal-mass black holes,” *Phys. Rev. D* **81** (2010) 084041, [arXiv:0912.3466 \[gr-qc\]](#).
- [34] L. Kidder, H. Pfeiffer, M. Scheel, *et al.*, “SpEC: Spectral Einstein Code.” <http://www.black-holes.org/SpEC.html>.
- [35] L. Finn, “Supernovae, gravitational radiation, and the quadrupole formula,” in *Frontiers in Numerical Relativity*, C. Evans, L. Finn, and D. Hobill, eds., pp. 126–145. Cambridge University Press, Cambridge, U.K.; New York, U.S.A., 1989.
- [36] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, “How Massive Single Stars End their Life,” *Astrophys. J.* **591** (2003) 288–300, [arXiv:astro-ph/0212469](#).
- [37] J. Hjorth *et al.*, “A very energetic supernova associated with the gamma-ray burst of 29 March 2003,” *Nature*. **423** (2003) 847–850, [arXiv:astro-ph/0306347](#).
- [38] W. D. Arnett, J. N. Bahcall, R. P. Kirshner, and S. E. Woosley, “Supernova 1987a,” *Annu. Rev. Astr. Astrophys.* **27** (1989) 629–700.

- [39] E. Mueller, M. Rampp, R. Buras, H. T. Janka, and D. H. Shoemaker, “Towards Gravitational Wave Signals from Realistic Core Collapse Supernova Models,” *Astrophys. J.* **603** (2004) 221–230, [arXiv:astro-ph/0309833](#).
- [40] C. D. Ott *et al.*, “Computational Models of Stellar Collapse and Core-Collapse Supernovae,” *J. Phys. Conf. Ser.* **180** (2009) 012022, [arXiv:0907.4043](#) [astro-ph.HE].
- [41] D. Chakrabarty *et al.*, “Nuclear-Powered Millisecond Pulsars and the Maximum Spin Frequency of Neutron Stars,” *Nature* **424** (2003) 42–44, [arXiv:astro-ph/0307029](#).
- [42] L. Bildsten, “Gravitational radiation and rotation of accreting neutron stars,” *Astrophys. J.* **501** (1998) L89, [arXiv:astro-ph/9804325](#).
- [43] A. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, “Detecting gravitational wave emission from the known accreting neutron stars,” *Mon. Not. Roy. Astron. Soc.* **389** (2008) 839–868, [arXiv:0803.4097](#) [astro-ph].
- [44] **The LIGO Scientific** Collaboration, B. Abbott *et al.*, “Beating the spin-down limit on gravitational wave emission from the Crab pulsar,” *Astrophys. J.* **683** (2008) L45–L50, [arXiv:0805.4758](#) [astro-ph].
- [45] N. Andersson and K. D. Kokkotas, “Towards gravitational-wave asteroseismology,” *Mon. Not. Roy. Astron. Soc.* **299** (1998) 1059–1068, [arXiv:gr-qc/9711088](#).
- [46] H. Dimmelmeier, N. Stergioulas, and J. A. Font, “Non-linear axisymmetric pulsations of rotating relativistic stars in the conformal flatness approximation,” *Mon. Not. Roy. Astron. Soc.* **368** (2006) 1609–1630, [arXiv:astro-ph/0511394](#).
- [47] B. Allen, “The stochastic gravity-wave background: Sources and detection,” in *Relativistic Gravitation and Gravitational Radiation*, J.-A. Marck and J.-P. Lasota, eds., Cambridge Contemporary Astrophysics, pp. 373–418. Cambridge University Press, 1997. [arXiv:gr-qc/9604033](#).
- [48] O. Lahav and A. R. Liddle, “The Cosmological Parameters 2010,” [arXiv:1002.3488](#) [astro-ph.CO].
- [49] **WMAP** Collaboration, E. Komatsu *et al.*, “Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations:Cosmological Interpretation,” *Astrophys. J. Suppl.* **180** (2009) 330–376, [arXiv:0803.0547](#) [astro-ph].
- [50] M. Maggiore, “Gravitational wave experiments and early universe cosmology,” *Phys. Rept.* **331** (2000) 283–367, [arXiv:gr-qc/9909001](#).

- [51] U. Seljak and M. Zaldarriaga, “Signature of gravity waves in polarization of the microwave background,” *Phys. Rev. Lett.* **78** (1997) 2054–2057, [arXiv:astro-ph/9609169](#).
- [52] R. H. Cyburt, B. D. Fields, K. A. Olive, and E. Skillman, “New BBN limits on Physics Beyond the Standard Model from He4,” *Astropart. Phys.* **23** (2005) 313–323, [arXiv:astro-ph/0408033](#).
- [53] T. L. Smith, E. Pierpaoli, and M. Kamionkowski, “A new cosmic microwave background constraint to primordial gravitational waves,” *Phys. Rev. Lett.* **97** (2006) 021301, [arXiv:astro-ph/0603144](#).
- [54] F. A. Jenet *et al.*, “Upper bounds on the low-frequency stochastic gravitational wave background from pulsar timing observations: Current limits and future prospects,” *Astrophys. J.* **653** (2006) 1571–1576, [arXiv:astro-ph/0609013](#).
- [55] **LIGO Scientific and Virgo** Collaboration, B. Abbott *et al.*, “An Upper Limit on the Stochastic Gravitational-Wave Background of Cosmological Origin,” *Nature* **460** (2009) 990, [arXiv:0910.5772 \[astro-ph.CO\]](#).
- [56] X. Siemens, V. Mandic, and J. Creighton, “Gravitational wave stochastic background from cosmic (super)strings,” *Phys. Rev. Lett.* **98** (2007) 111101, [arXiv:astro-ph/0610920](#).
- [57] V. Mandic and A. Buonanno, “Accessibility of the Pre-Big-Bang Models to LIGO,” *Phys. Rev.* **D73** (2006) 063008, [arXiv:astro-ph/0510341](#).
- [58] E. E. Flanagan, “The Sensitivity of the laser interferometer gravitational wave observatory (LIGO) to a stochastic background, and its dependence on the detector orientations,” *Phys. Rev.* **D48** (1993) 2389–2407, [arXiv:astro-ph/9305029](#).
- [59] J. Weber, “Gravitational radiation,” *Phys. Rev. Lett.* **18** (1967) 498–501.
- [60] M. P. McHugh *et al.*, “Calibration of the ALLEGRO resonant detector,” *Class. Quant. Grav.* **22** (2005) S965–S973.
- [61] **AURIGA** Collaboration, A. Vinante, “Present performance and future upgrades of the AURIGA capacitive readout,” *Class. Quant. Grav.* **23** (2006) S103–S110.
- [62] P. Astone *et al.*, “Status report on the EXPLORER and NAUTILUS detectors and the present science run,” *Class. Quant. Grav.* **23** (2006) S57–S62.
- [63] **IGEC-2** Collaboration, P. Astone *et al.*, “Results of the IGEC-2 search for gravitational wave bursts during 2005,” *Phys. Rev.* **D76** (2007) 102001, [arXiv:0705.0688 \[gr-qc\]](#).

- [64] P. Astone *et al.*, “IGEC2: A 17-month search for gravitational wave bursts in 2005-2007,” *Phys. Rev.* **D82** (2010) 022003, [arXiv:1002.3515 \[gr-qc\]](#).
- [65] L. Gottardi *et al.*, “Sensitivity of the spherical gravitational wave detector MiniGRAIL operating at 5 K,” *Phys. Rev.* **D76** (2007) 102005, [arXiv:0705.0122 \[gr-qc\]](#).
- [66] O. D. Aguiar *et al.*, “The Brazilian gravitational wave detector Mario Schenberg: Progress and plans,” *Class. Quant. Grav.* **22** (2005) S209–S214.
- [67] <http://www.geo600.org>.
- [68] <http://tamago.mtk.nao.ac.jp>.
- [69] **LIGO Scientific** Collaboration, B. Abbott *et al.*, “LIGO: The Laser Interferometer Gravitational-Wave Observatory,” *Rept. Prog. Phys.* **72** (2009) 076901, [arXiv:0711.3041 \[gr-qc\]](#).
- [70] L. Schnupp. Presentation at a European collaboration meeting on interferometric detection of gravitational waves, Sorrento, 1988.
- [71] **LIGO Scientific** Collaboration, J. R. Smith, “The path to the enhanced and advanced LIGO gravitational- wave detectors,” *Class. Quant. Grav.* **26** (2009) 114013, [arXiv:0902.0381 \[gr-qc\]](#).
- [72] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen, “Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform,” *Phys. Rev.* **D59** (1999) 124016, [arXiv:gr-qc/9901076](#).
- [73] **LIGO Scientific** Collaboration, B. Abbott *et al.*, “Search for gravitational waves from galactic and extra- galactic binary neutron stars,” *Phys. Rev.* **D72** (2005) 082001, [arXiv:gr-qc/0505041](#).
- [74] S. K. Chatterji, *The search for gravitational-wave bursts in data from the second LIGO science run*. PhD thesis, MIT, 2005.
- [75] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and E. E. Flanagan, “An excess power statistic for detection of burst sources of gravitational radiation,” *Phys. Rev.* **D63** (2001) 042003, [arXiv:gr-qc/0008066](#).
- [76] J. Sylvestre, “Time-frequency detection algorithm for gravitational wave bursts,” *Phys. Rev.* **D66** (2002) 102004, [arXiv:gr-qc/0210043](#).
- [77] I. Daubechies, *Ten lectures on wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992.
- [78] S. Mallat, *A wavelet tour of signal processing*. Academic Press, 1998.

- [79] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, “Coherent method for detection of gravitational wave bursts,” *Class. Quant. Grav.* **25** (2008) 114029, arXiv:0802.3232 [gr-qc].
- [80] K. C. Cannon, “A Bayesian coincidence test for noise rejection in a gravitational-wave burst search,” *Class. Quant. Grav.* **25** (2008) 105024.
- [81] E. E. Flanagan and S. A. Hughes, “Measuring gravitational waves from binary black hole coalescences. II: The waves’ information and its extraction, with and without templates,” *Phys. Rev.* **D57** (1998) 4566–4587, arXiv:gr-qc/9710129.
- [82] M. Rakhmanov, “Rank deficiency and Tikhonov regularization in the inverse problem for gravitational-wave bursts,” *Class. Quant. Grav.* **23** (2006) S673–S686, arXiv:gr-qc/0604005.
- [83] T. Z. Summerscales, A. Burrows, C. D. Ott, and L. S. Finn, “Maximum Entropy for Gravitational Wave Data Analysis: Inferring the Physical Parameters of Core-Collapse Supernovae,” *The Astrophysical Journal* **678** (2008) no. 2, 1142, arXiv:0704.2157 [astro-ph].
- [84] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, “Constraint likelihood analysis for a network of gravitational wave detectors,” *Phys. Rev.* **D72** (2005) 122002, arXiv:gr-qc/0508068.
- [85] C. Pankow *et al.*, “A burst search for gravitational waves from binary black holes,” *Class. Quant. Grav.* **26** (2009) 204004, arXiv:0905.3120 [gr-qc].
- [86] A. C. Searle, P. J. Sutton, and M. Tinto, “Bayesian detection of unmodeled bursts of gravitational waves,” *Class. Quant. Grav.* **26** (2009) 155017, arXiv:0809.2809 [gr-qc].
- [87] S. Chatterji *et al.*, “Coherent network analysis technique for discriminating gravitational-wave bursts from instrumental noise,” *Phys. Rev.* **D74** (2006) 082005, arXiv:gr-qc/0605002.
- [88] **LIGO Scientific** Collaboration, B. Abbott *et al.*, “Upper limits on gravitational wave bursts in LIGO’s second science run,” *Phys. Rev.* **D72** (2005) 062001, arXiv:gr-qc/0505029.
- [89] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, “Multiresolution techniques for the detection of gravitational-wave bursts,” *Class. Quant. Grav.* **21** (2004) S1809–S1818, arXiv:gr-qc/0412119.
- [90] A. Haar, “Zur theorie der orthogonalen funktionensysteme,” *Math. Annalen* **69** (1910) 331–371.
- [91] J. Zweizig. <http://www.ligo.caltech.edu/~jzweizig/dmt/DMTPProject>.

- [92] K. Cannon. http://www.lsc-group.phys.uwm.edu/daswg/legacy_projects/onasys.html.
- [93] <http://www.lsc-group.phys.uwm.edu/lscdatagrid/overview.html>.
- [94] J. Kanner *et al.*, “LOOC UP: Locating and observing optical counterparts to gravitational wave bursts,” *Class. Quant. Grav.* **25** (2008) 184034, arXiv:0803.0312 [astro-ph].
- [95] L. Blackburn *et al.*, “The LSC Glitch Group : Monitoring Noise Transients during the fifth LIGO Science Run,” *Class. Quant. Grav.* **25** (2008) 184004, arXiv:0804.0800 [gr-qc].
- [96] **LIGO Scientific** Collaboration, B. P. Abbott *et al.*, “Search for gravitational-wave bursts in the first year of the fifth LIGO science run,” *Phys. Rev.* **D80** (2009) 102001, arXiv:0905.0020 [gr-qc].
- [97] J. Slutsky *et al.*, “Methods for Reducing False Alarms in Searches for Compact Binary Coalescences in LIGO Data,” *Class. Quant. Grav.* **27** (2010) 165023, arXiv:1004.0998 [gr-qc].
- [98] **LIGO Scientific and VIRGO** Collaboration, N. Leroy, “Data quality and vetoes for the gravitational-wave burst and inspiral analyses in Virgo’s first science run,” *Class. Quant. Grav.* **26** (2009) 204007.
- [99] **LIGO Scientific** Collaboration, R. Gouaty, “Detection Confidence Tests for Burst and Inspiral Candidate Events,” *Class. Quant. Grav.* **25** (2008) 184006, arXiv:0805.2412 [gr-qc].
- [100] **LIGO Scientific** Collaboration, and others, “Search for Gravitational Waves from Compact Binary Coalescence in LIGO and Virgo Data from S5 and VSR1,” arXiv:1005.4655 [gr-qc].
- [101] L. Cadonati and S. Marka, “CorrPower: A cross-correlation-based algorithm for triggered and untriggered gravitational-wave burst searches,” *Class. Quant. Grav.* **22** (2005) S1159–S1167.
- [102] J. W. C. McNabb *et al.*, “Overview of the BlockNormal event trigger generator,” *Class. Quant. Grav.* **21** (2004) S1705–S1710, arXiv:gr-qc/0404123.
- [103] R. A. Mercer and S. Klimenko, “Visualizing gravitational-wave event candidates using the coherent event display,” *Class. Quant. Grav.* **25** (2008) 184025.
- [104] S. K. Chatterji, “Q-scan homepage.”
<http://www.lsc-group.phys.uwm.edu/lscdatagrid/overview.html>.

- [105] A. Stroeer, J. K. Cannizzo, J. B. Camp, and N. Gagarin, “Methods for detection and characterization of signals in noisy data with the Hilbert-Huang Transform,” *Phys. Rev.* **D79** (2009) 124022, [arXiv:0903.4616](#) [physics.data-an].
- [106] J. G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower, “Comparisons of binary black hole merger waveforms,” *Class. Quant. Grav.* **24** (2007) S25–S31, [arXiv:gr-qc/0701016](#).
- [107] R. K. Kopparapu *et al.*, “Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events,” *The Astrophysical Journal* **675** (2008) no. 2, 1459, [arXiv:0706.1283](#) [astro-ph].
- [108] P. R. Brady, J. D. E. Creighton, and A. G. Wiseman, “Upper limits on gravitational-wave signals based on loudest events,” *Class. Quant. Grav.* **21** (2004) S1775–S1782, [arXiv:gr-qc/0405044](#).
- [109] <http://www.scipy.org>.