

SEARCHES FOR GRAVITATIONAL WAVES FROM PERTURBED BLACK HOLES IN
DATA FROM LIGO DETECTORS

A Dissertation

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

in

The Department of Physics and Astronomy

by

Sarah Elizabeth Caudill

B.S., Stetson University, 2006

M.S., Louisiana State University, 2010

August 2012

To Josh for inspiration and relief, to Mom and Dad for support and encouragement, and to Mick for being awesome

Acknowledgments

I am very thankful to all those who have helped me on this path, guiding when I lost sight, helping with any number of mini crises, or just listening. I must start at the very beginning and thank my parents who spent endless hours talking with me as a kid about Stephen Hawking's *Universe* or Carl Sagan's *Cosmos* and planting that seed of wonder about our world. And in high school, thank you to Mr. Young for always challenging us and stepping outside the rigid curriculum. Finally, thank you to Dr. Riggs for being the most dedicated physics teacher I've ever known.

I must also thank the LIGO project and the LSC for the amazing opportunities and the wonderful people that I have met along the way. I am sincerely thankful for the opportunity to have worked as a SURF student at Caltech with Duncan Brown and Peter Shawhan, two great advisors who introduced me to the world of gravitational wave science and continue to provide advice and encouragement. Furthermore, without this opportunity, I would never have met my advisor Gaby Gonzalez who has been a wonderful, patient guide through all sorts of problems and who understands that perfect is the enemy of good. Without her, I am not sure I would have made it through. I am also thankful for the support and advice of many other experienced scientists in the LSC along the way including Warren Anderson, Patrick Brady, Jolien Creighton, Ilya Mandel, Xavi Siemens, and Alan Weinstein who have also become friends.

Thank you also to Lisa Goggin for taking a chance on a young graduate student. I am so very lucky to be a part of the ringdown group and work with such wonderful people. Thank you to Paul Baker, Matt Benaquista, Sukanta Bose, Neil Cornish, and Dipongkar Talukder for making Thursday my favorite day of the week. And a huge thank you to Collin Capano and Kari Hodge for immeasurable help at all hours of the day. The ringdown search would not have been possible without all of you. Also, thanks to Tom Dent for having an ever-watchful eye and catching many steps that were missed along the way.

I would also like to thank Rahul Biswas, Chad Hanna, Nick Fotopoulos, Satya Mohaptra, and Matt West for immense help along the way. I am grateful for all that you have taught me and the hours you've spent with me through computer crises. This work would not have been possible without you.

I am also very lucky to have had some awesome office mates over the past five years including Jake Slutsky, Tobin Fricke, Kate Dooley, Jeff Kissel, Garrett May, Joe Bowers, Anamaria Effler, and Marissa Walker. You guys made the basement a little brighter with your comic relief. Also, I am lucky for the help, guidance, and friendship of the post-docs at LSU including Romain Gouaty, Cristina Torres, Cesar Costa, Keiko Kokeyama, and Adam Mullavey. Also, without my trusty companions, Ashley Pagnotta and Zach Byerly, I never would have made it past the qualifier exam. And thesis writing probably would have been delayed as well without our thesis writing parties...thanks to Azadeh Keivani too.

I have also had the opportunity to work with some amazing people outside the LSC. Thanks to Manuel Tiglio, Scott Field, Frank Herrmann, and Chad Galley for allowing me to work with you on an exciting project. Your work ethic continues to amaze me. I am very lucky to have worked with you all and to now call you my friends. Also, thank you to Vitor Cardoso, Emanuele Berti, and Marco Cavaglia for extremely valuable discussions on black hole ringdowns. I am looking forward to further collaborations with all of you.

Finally, to Josh, who first questioned why I wanted to go to medical school when I spent all my time taking physics classes; thank you for always helping me to see the light.

Preface

The work presented in this thesis was carried out within the LIGO Scientific Collaboration (LSC). The methods and results presented here are under review and are potentially subject to change. The opinions expressed here are those of the author and not necessarily those of the LSC. The author gratefully acknowledges the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, which provided support for this work.

Table of Contents

Dedication	ii
Acknowledgments	iii
Preface	iv
List of Tables	vii
List of Figures	viii
Abstract	xviii
1. Introduction	1
2. Gravitational Waves	5
2.1 Gravitational Waves in General Relativity	5
2.1.1 Linearized Gravity	5
2.1.2 Newtonian Limit of General Relativity	6
2.1.3 Properties of Gravitational Waves In Vacuum	7
2.1.4 Properties of Gravitational Waves With Source	8
2.1.5 Detection of Gravitational Waves	9
2.2 Sources of Gravitational Waves	12
2.2.1 Compact Binary Coalescence	13
2.2.2 Bursts	14
2.2.3 Continuous Sources	14
2.2.4 Stochastic Background	16
2.3 Gravitational Wave Detectors	17
2.3.1 Toward a Real Detector	19
2.3.2 Progress	28
3. Black Holes	30
3.1 Black Holes in General Relativity	30
3.2 Observations of Black Holes	31
3.3 Intermediate Mass Black Holes	32
3.4 Perturbed Black Holes	34
3.5 Ringdown Waveforms	37
3.5.1 Ringdown Waveforms from Compact Binary Coalescence	39
4. Searches for Gravitational Wave Ringdowns	48
4.1 Gravitational Wave Data Analysis	48
4.1.1 The Optimal Detection Statistic	48
4.1.2 Template Banks	50
4.2 The Search Pipeline	52
4.2.1 Summary of Main Changes Since S4	53

4.2.2	Automation with <code>ihope</code>	55
4.2.3	Background Estimation	61
4.2.4	Detection Efficiency	62
4.2.5	Coincident Trigger Analysis	63
4.2.6	Post-processing with Pipedown	67
4.3	Tuning the Pipeline	79
4.3.1	Injection Analysis	79
4.3.2	SNR Threshold	81
4.3.3	Coincidence Windows	84
4.3.4	Detection Statistic	84
4.4	S5 Search Results	98
4.4.1	S5 Data	99
4.4.2	S5 Playground Results	100
4.4.3	S5 Search Results	104
4.5	Parameter Recovery	123
4.5.1	IMR Parameter Recovery	123
5.	Multimode Ringdown Searches	140
5.1	Motivation	140
5.2	The Multimode Ringdown Signal	141
5.3	The Reduced Basis Technique	141
5.4	One Mode Ringdown Catalogs	145
5.4.1	Metric-based Catalogs	145
5.4.2	One-mode Reduced Basis	145
5.4.3	Reduced Basis Representation of One-mode Waveforms	147
5.5	Multi-Mode Ringdown Catalogs	150
5.5.1	Constrained Two-mode Reduced Basis	153
5.5.2	Unconstrained Multi-mode Reduced Basis	156
5.6	Final Remarks	159
6.	Conclusion	160
	Bibliography	163
	Appendix A: Playground Configuration Files	172
	Appendix B: Assignment of Copyright	200
	Appendix C: S5 Search Configuration Files	201
	Vita	225

List of Tables

1.1	LIGO and Virgo Science Runs from the initial, enhanced, and projected advanced eras.	4
2.1	Physical parameters of the Crab pulsar. Data taken from [76].	16
3.1	A list of the first few (complex) quasi-normal mode frequencies $GM\omega_{nl}/c^3$ of oscillation for a non-spinning Schwarzschild black hole [127].	34
3.2	A list of the first few (complex) quasi-normal mode frequencies $GM\omega_{lm}/c^3$, eigenvalues A_{lm} , and dimensionless spin parameter $c^2a/(GM)$ for the fundamental $(l, m, n) = (1, m, 2)$ mode of a rotating Kerr black hole.	35
4.1	Tests of different SNR thresholds for the purpose of finding the best setting to increase the triple coincidence background. An asterik indicates too many triggers to process.	84
4.2	Parameters of loudest triple coincident playground candidate in months 9-10.	101
4.3	Identically ranked H1L1 coincidences associated with one ringdown injection.	104
5.1	Number of RB waveforms (N_{RB}) needed to represent 1-mode training spaces for the $(\ell, m, n) = (2, 2, 0)$ and $(3, 3, 0)$ QNMs with different minimal matches MM. The training space representation error is taken to be $\epsilon^2 = 10^{-12}$. The number of metric-based templates scales with MM as $N_{\text{metric}} \propto (1 - \text{MM})^{-1}$ for the 2-dimensional, 1-mode QNM catalog [170].	149
5.2	Number of reduced basis waveforms (N_{RB}) needed to represent 2-mode training spaces with $(\ell, m, n) = (2, 2, 0)$ and $(3, 3, 0)$ for different minimal matches MM. The training space representation error is $\epsilon^2 = 10^{-12}$. The “2-mode, GR” and “2-mode” cases are discussed in Sections 5.5.1 and 5.5.2, respectively. The number of metric-based templates in those two cases scales with MM as $N_{\text{metric}} \propto (1 - \text{MM})^{-d/2}$, with $d = 3$ and $d = 5$, respectively [170]. For the “2-mode, GR” entries, MM refers to the minimal match associated with the training space for the $(2, 2, 0)$ mode, which is used to generate the points in the (f_{330}, Q_{330}) plane as described in the text. For the “2-mode” case the training space representation error of 10^{-12} corresponds to each mode separately and the corresponding total error is bounded from above by 4×10^{-12} , as discussed in Section 5.5.2. We are grateful to V. Cardoso for his help in generating the values of N_{metric} in the second column for the constrained, 2-mode case.	156

List of Figures

1.1	Aerial views of the three detector sites in the LIGO-Virgo network taken from References [14] and [15]. (a) Hanford, WA site. Home of the 4 km H1 and former home of the 2 km H2 interferometer. (b) Livingston, LA site. Home of the 4 km L1 interferometer. (c) Cascina, Italy site. Home of the 3 km Virgo interferometer.	2
2.1	The effects of plus and cross polarization on a ring of test masses. The plus polarization alternately compresses and stretches the x- and y-separations. The cross polarization has the same effect only rotated by 45°.	12
2.2	This figure shows the cumulative periastron period shift in seconds of the famous Hulse-Taylor binary pulsar PSR B1913+16 from 1975 to 2005. The red data points match exceptionally well with the theoretical blue curve prediction for the period shift due to gravitational wave emission as predicted by General Relativity [83]. The discovery and analysis of this binary star system earned its discoverers, Hulse and Taylor, the 1993 Nobel Prize in Physics.	15
2.3	Optical layout of a Michelson interferometer. Diagram taken from [88].	18
2.4	Optical layout for the LIGO interferometers. The power recycling cavity is formed by the power recycling mirror (PRM), the beam splitter (BS), and input test masses (ITMX,ITMY). The Fabry-Perot cavities are formed by the ITMs and the end test masses (ETMs). This diagram is taken from [88].	20
2.5	The differential arm feedback control loop of an interferometric detector.	23
2.6	Cartoon demonstrating relative positions and orientations of a source frame (x, y, z) , polarization frame of the gravitational wave in the direction of the detector (x'', y'', z'') , and detector frame (x', y', z')	24
2.7	Noise sources of initial LIGO. Figure taken from [90].	27
2.8	The best amplitude spectral density of the displacement noise of the LIGO detectors during the S5 science run. Taken from [88].	29
3.1	The frequency of the $(2, 2, 0)$ mode as a function of final black hole mass for different dimensionless spin values.	36
3.2	The quality factor of the $(2, 2, 0)$ mode as a function of dimensionless spin parameter.	37
3.3	Ringdown waveforms with $(f_{220}, Q_{220})=(2000 \text{ Hz}, 20)$ on top and $(f_{220}, Q_{220}) = (50 \text{ Hz}, 2.1187)$ on bottom.	38

3.4	The amount of energy emitted as gravitational waves during the ringdown associated with compact binary coalescence as a function of the symmetric mass ratio η . The data and fit are taken from [131].	40
3.5	An example chirp waveform.	40
3.6	Ringdown horizon distance as a function of final black hole mass assuming $\epsilon = 0.01$, a dimensionless spin of $a = 0.9$ and an SNR of 8 in the H1, H2, and L1 interferometers during the S5 science run and in the Virgo detector V1 during its first science run.	41
3.7	Ringdown horizon distance as a function of final black hole mass assuming $\epsilon = 0.01$ and an SNR of 8 in the H1 interferometer during the S5 science run. The dimensionless spin parameter is allowed to vary.	42
3.8	IMR waveform for a non-spinning compact binary with component masses of $10 M_{\odot}$	43
3.9	IMR waveform for a non-spinning compact binary with component masses of $150 M_{\odot}$	44
3.10	Final black hole mass as a function of initial mass for different values of the symmetric mass ratio ν	45
3.11	Final black hole spin as a function of a_2 for different a_1 for a $(100, 100) M_{\odot}$ binary black hole system.	46
3.12	Final black hole spin as a function of a_2 for different a_1 for a $(100, 25) M_{\odot}$ binary black hole system.	47
4.1	The S5 ringdown template bank. The settings used to construct this bank were the following: maximum allowed mismatch $ds^2 = MM_{\max} = 0.03$, minimum frequency $f_{\min} = 50$ Hz, maximum frequency $f_{\max} = 2000$ Hz, minimum quality factor $Q_{\min} = 2.1187$, and maximum quality factor $Q_{\max} = 20$	53
4.2	Contours of constant mismatch $ds^2 = MM_{\max} = 0.03$ around each template. The plot on the left demonstrates how these contours overlap for templates at a single frequency. The plot on the right demonstrates the overlap for a snippet of the template bank between $f = 50$ Hz and $f = 57.5$ Hz.	54
4.3	The ringdown search pipeline.	56
4.4	A snippet of calibrated LIGO $h(t)$ data sampled at 16384 Hz that acts as the raw input to the search pipeline. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment.	58

4.5	A snippet of the data shown in Figure 4.4 after resampling from 16384 Hz to 4096 Hz. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment. Any frequency content above about 2 kHz is aliased so a low-pass filter is applied at this frequency.	58
4.6	A snippet of the data shown in Figure 4.5 after applying a high-pass filter to suppress data below the seismic wall at approximately 40 Hz. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment. Any frequency content above about 2 kHz is aliased so a low-pass filter is applied at this frequency.	59
4.7	A snippet of the data shown in Figure 4.6 after discarding 8 s of padding at the start and end of the time series before analysis. This is needed since the previously applied filters corrupt the data at the start and end. Thus time $t = 0.0$ s on these plots correspond to time $t = 8.0$ s in Figures 4.4, 4.5, and 4.6. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment.	59
4.8	The calibrated power spectral density and inverse calibrated power spectral density.	60
4.9	Time of arrival difference versus frequency in S4 injections in H1H2.	64
4.10	Time of arrival difference versus frequency in S4 background in H1H2.	65
4.11	Efficiency versus effective distance and background events per slide in H1H2 for the S4 coincidence test statistic.	66
4.12	Efficiency versus effective distance and background events per slide in H1H2 for the new S5 coincidence test statistic. Note that the efficiency is the same as in Figure 4.11 for the old statistic but the number of H1H2 background coincidences per time slide has consistently decreased.	66
4.13	Application of H1H2 consistency cut to H1L1 background coincidences in triple time. Only H1L1 coincidences in triple time with an H1 SNR less than 16 were retained.	68
4.14	The ringdown Pipedown post-processing pipeline.	78
4.15	Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the component mass plane m_1 - m_2 and in the component spin plane a_1 - a_2 . The color bar represents the value of the $\chi = (m_1 a_1 + m_2 a_2) / (m_1 + m_2)$ parameter in the left plot and the mass ratio in the right plot.	80

4.16	Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the final black hole frequency and quality factor plane. The color bar represents the value of the final black hole mass in the left plot and the final black hole spin in the right plot.	80
4.17	Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the final black hole mass and spin plane. The color bar represents the value of the final black hole frequency in the left plot and the final black hole quality factor in the right plot.	81
4.18	Distribution of ringdown-only injections used for the tuning of the search pipeline parameters in the final black hole mass and spin plane. The color bar represents the value of the frequency in the left plot and the quality factor in the right plot.	81
4.19	Distribution of ringdown-only injections used for the tuning of the search pipeline parameters in the final black hole frequency and quality factor plane. The color bar represents the value of the mass in the left plot and the spin in the right plot.	82
4.20	ROC plots demonstrating the performance of different ds^2 values for EOBNR and PhenomB injections found in triple coincidence.	85
4.21	The H1 and H2 single detector SNRs of triple coincident events. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines. The shaded region marks the space where the H1/H2 SNR ratio would be less than one. Any triggers in this region are excluded in S5.	87
4.22	The H1 and L1 single detector SNRs of triple coincident events. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines.	88
4.23	The H2 and L1 single detector SNRs of triple coincident events. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines.	89
4.24	Contours of constant background trigger density in the H1H2 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.	90
4.25	Contours of constant background trigger density in the H1L1 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.	91
4.26	Contours of constant background trigger density in the H2L1 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.	92

4.27	A scatter plot of the new S5 triples statistic versus the S4 triples detection statistic. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers.	93
4.28	The fraction of injections found louder than the loudest background versus injected effective distance in the H1 detector. A zoomed plot is shown on the right. The black curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S4 statistic. The red curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S5 statistic and when the H1H2 SNR ratio cut was implemented. The brown curve shows the fraction of total found injections, the theoretical best that we could do.	94
4.29	An ROC plot showing the search sensitivity as a function of false alarm rate for the S4 triples' statistic in black and the S5 triples' statistic in red.	94
4.30	An ROC plot demonstrating that MVSC performs better than the new S5 chopped-L-like statistic on two months of S5 data using the 20 input parameters described in Section 4.3.4.2.	98
4.31	An example of the MVSC "bump" in the IFAR plots seen in the S6 high mass search for LIV1 coincident time.	99
4.32	Months 9-10 playground uncombined and combined cumulative IFAR plots. The cyan triangle in the left plot that falls to the far right with an arrow pointing to the right indicates a triple coincident playground candidate that was louder than any background coincidences in all of months 9-10.	102
4.33	Time-frequency whitened spectrograms of the loudest triple coincident playground candidate in months 9-10 in both the H1 and L1 gravitational wave channels.	102
4.34	H1 time-frequency whitened spectrogram of the blind injection performed in LIGO and Virgo data on September 16th, 2010. The signal was from a spinning binary system with component masses of $1 M_{\odot}$ and $24 M_{\odot}$. This injection had an SNR of 14 in H1. For more information on this blind injection challenge, please refer to [161].	103
4.35	Missed/found plots for EOBNRv2, PhenomB, and Ringdown injections in Months 9-10.	106
4.36	IFAR plot of H1H2L1 time in Months 9-10.	107
4.37	List of loudest coincidences in H1H2L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	108
4.38	IFAR plot of H1L1 time in Months 9-10.	109

4.39	List of loudest coincidences in H1L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	110
4.40	IFAR plot of H2L1 time in Months 9-10.	111
4.41	List of loudest coincidences in H2L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	112
4.42	Missed/found plots for EOBNRv2, PhenomB, and Ringdown injections in Months 15-16.	113
4.43	IFAR plot of H1H2L1 time in Months 15-16.	114
4.44	List of loudest coincidences in H1H2L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	115
4.45	IFAR plot of H1L1 time in Months 15-16.	116
4.46	List of loudest coincidences in H1L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	117
4.47	IFAR plot of H2L1 time in Months 15-16.	118
4.48	List of loudest coincidences in H2L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.	119
4.49	The efficiency of detecting ringdown injections in triple coincidence in the 70-140 Hz band in S5 using the new triples chopped-L statistic.	120
4.50	The efficiency as a function of injected distance of ringdown injections in triple coincidence in the 70-140 Hz band in S4. Taken from Reference [26].	121

4.51	The efficiency as a function of injected effective distance of PhenomB injections in triple coincidence with mass ratios less than 4:1 and total masses between 100-200 M_{\odot} in S5.	122
4.52	PhenomB injection distribution in the m_1 - m_2 plane on the left and in the spin1z-spin2z plane on the right. All injections made are shown in red and those that were recovered in triple coincidence are overlain in green.	124
4.53	Histogram of PhenomB injection distributions by mass ratio in the plot on the left and by final black hole spin in the plot on the right. The distributions of all injections are shown in red and the distributions of those recovered in triple coincidence are overlain in green.	125
4.54	Injected versus detected frequency for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.	127
4.55	Injected versus detected quality factor for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.	128
4.56	Differences in the recovered frequency for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal.	129
4.57	Differences in the recovered quality factor for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal. The discreteness of the recovered values is due to the discreteness of quality factor values available in the ringdown template bank.	129
4.58	Histogram of detected minus injected GPS start times for triple found PhenomB injections in H1, H2, and L1.	130
4.59	Difference in GPS start time for triple found PhenomB injections as a function of injected frequency for H1, H2, and L1. The color bar denotes injected dimensionless spin.	131
4.60	Histogram of difference in GPS start time for triple found PhenomB injections in H1 and H2.	132
4.61	Histogram of difference in GPS start time for triple found PhenomB injections in H1 and L1.	133
4.62	Difference in GPS start time for triple found PhenomB injections in H1 and H2 on the left and in H1 and L1 on the right as a function of H1 injected frequency. The color bar denotes injected quality factor in the range from 2.1187-6.0.	133

4.63	Injected versus detected final black hole mass for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.	134
4.64	Injected versus detected final black hole mass for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the final black hole dimensionless spin from 0.0-1.0. The black line marks the diagonal.	135
4.65	Injected versus detected final black hole dimensionless spin for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.	136
4.66	Differences in the recovered mass of the final black hole for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal.	137
4.67	Differences in the recovered spin of the final black hole for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal. The discreteness of the recovered values is due to the discreteness of quality factor values available in the ringdown template bank. Recall that the spin of the final black hole is a function only of quality factor.	137
4.68	Injected versus detected effective distance for triple found PhenomB injections for H1, H2, and L1. The color bar demonstrates the recovered SNR of these injections in each interferometer.	139
5.1	Representation error as a function of the number of reduced basis waveforms for a single mode catalog and all possible choices of seeds (see Section 5.4.2). The dark line shows the average and the shaded area the maximum dispersion around it. Clearly, the method is robust and does not require any fine tuning. Notice that the exponential rate of convergence of the error is present from the outset.	146
5.2	One-mode metric placement catalog (gray dots), which is taken to be the training space, and the subset of points selected by the greedy algorithm (red bars) for the (2, 2, 0) mode. The settings are discussed in Sections 5.4.1 and 5.4.2.	147
5.3	Greedy points selected from a 45×45 uniform training space. The uniform training space is laid out in the plot on the left and the points selected by the greedy algorithm are shown in the plot on the right.	148

5.4	Greedy points selected from a 132×132 uniform training space on the left and from a 300×300 uniform training space on the right. Note how using a more densely packed training space (as opposed to the 45×45 uniform training space in Figure 5.3) reduces the spread in f_{22} at lower values of Q_{22} . The favored greedy points are along the extrema of frequency values except at the highest Q_{22} . As we can see in Figure 5.5, the points selected via denser and denser uniform training spaces asymptote to the points selected when using a typical metric training space.	148
5.5	Greedy points selected from a metric training space with $1 - \text{MM} = 10^{-3}$. The metric training space is laid out in the plot on the left and the points selected by the greedy algorithm are shown in the plot on the right. This is the selection of greedy points to which all other training spaces explored approach.	149
5.6	Waveform representation error of randomly selected QNMs that are not necessarily present in the training space. Distributions of waveform representation errors for 1-mode (gray) and 2-mode (white) using reduced bases built from a training space with $\text{MM} = 0.99$. In all cases, the error is found to be $< 9 \times 10^{-10}$ for single modes and $< 9.8 \times 10^{-10}$ for two modes. See Section 5.4.3 and 5.5.1 for details.	151
5.7	Waveform representation error of randomly selected QNMs that are not necessarily present in the training space. The waveform representation error [defined in (5.9)] as a function of the randomly chosen values of (f_{220}, Q_{220}) for the 1-mode case with the same setup as in Figure 5.7. The color map indicates the value of the error on a \log_{10} scale. The worst errors occur for the largest Q_{220} values and are mostly independent of f_{220} . The crosses indicate the parameters selected by the greedy algorithm for this reduced basis. . . .	152
5.8	Values for the relative amplitude parameter \mathcal{A} in the two-mode constrained waveforms [Eq. (5.14)] selected by the greedy algorithm for each reduced basis waveform (indexed on the horizontal axis). Here there are $n_{\mathcal{A}} = 1,000$ samples for $\mathcal{A} \in [0, 1]$ but essentially either the $(2, 2, 0)$ mode ($\mathcal{A} = 0$) or the $(3, 3, 0)$ mode ($\mathcal{A} = 1$) are selected, with a higher density of points associated with the latter due to its larger maximum quality factor for any given black hole spin parameter value. About 3% of the selected amplitudes are <i>equal</i> to zero (not just nearly so). All of the remaining points selected are either <i>equal</i> to 1 (about 79%) or between $\mathcal{A} = 0.93$ and 1 (about 18%).	154
5.9	The metric-based training space (points) for the same constrained, two-mode catalog of Fig. 5.8 and the subset of parameter values selected by the greedy algorithm (bars) for a training space representation error of $\epsilon^2 = 10^{-12}$. Red: $(2, 2, 0)$ mode. Blue: $(3, 3, 0)$ mode. The distribution of points closely resembles a subset of the two one-mode greedy points due to the linearity of the problem and to the “decoupling” of selected amplitude points (see Fig. 5.8).	155

5.10	Waveform representation errors, as defined by Eq.(5.9), of constrained, two-mode waveforms for a random sample of $(f_{220}, Q_{220}, \mathcal{A})$ triples. The training space used to build the reduced basis is relatively coarse yet the waveform representation error is found to be, in all cases, smaller than 9.8×10^{-10} . See Section 5.5.1 for more details.	157
6.1	Assignment of copyright form from IOP Publishing Limited. Section 3.2.2 grants the author rights to include the article (all or part) in a research thesis or dissertation. This article can be found online: http://iopscience.iop.org/0264-9381/29/9/095016/	200

Abstract

Black hole perturbation theory predicts that a perturbed black hole will emit gravitational waves in a superposition of quasi-normal modes. Various astrophysical processes can produce such a black hole including the merger of two compact binary neutron stars or black holes. The final form of the waveform from such a system is known as a ringdown. We discuss the search through data from LIGO's fifth science run for ringdown gravitational waves from intermediate mass black holes using a matched filtering pipeline. We outline the improvements to the pipeline since LIGO's fourth science run including the creation of a fully automated post-processing pipeline for coincident triggers, updated waveform simulation code, a new 3D coincidence test to check simultaneously for coincidence in frequency, quality factor, and time, and the use of a new detection statistic through a multi-variate statistical classifier. Results from four months of S5 data have been obtained so far and no gravitational wave candidates were found. The results of the search are ongoing. We demonstrate the improvement in the upper limit on the rate of black hole ringdowns in the local universe over the upper limit of the fourth science run. We investigate parameter recovery of full inspiral-merger-ringdown waveforms using a ringdown matched filter. Finally, we explore the Reduced Basis approach which provides very compact and high-accuracy representations of multi-mode ringdown gravitational waves.

1. Introduction

Albert Einstein's theory of General Relativity describes how the curvature of spacetime reacts to the presence of matter and energy. In the weak field limit, the Einstein equation simplifies to a wave equation, providing a theoretical basis for the existence of gravitational waves. These waves are essentially ripples in the fabric of spacetime, propagating outward from the source at the speed of light. Currently, we have only indirect evidence of gravitational radiation. The canonical example is provided by measurements of the orbital period of the Hulse-Taylor binary pulsar. However, through the use of large scale laser interferometric detectors, we could detect these waves directly by measuring the influence of a passing gravitational wave on the fringes of a large Michelson interferometer. A detection would usher in the era of gravitational wave astronomy. The Laser Interferometer Gravitational-wave Observatory (LIGO) [1] would provide a new window on the universe through which we could learn about the sources of these waves including binary black hole and neutron star collisions, collapsing supernova, rotating stars and pulsars, the Big Bang itself, and more. With LIGO, we can begin to probe astrophysical sources that cannot be observed with traditional electromagnetic methods and test General Relativity in extreme relativistic conditions [2].

The work presented in this thesis represents the culmination of five years of projects and investigations as a member of the LIGO Scientific Collaboration (LSC). The main focus of this thesis will be the work I completed as the leader of the search for gravitational waves from perturbed black holes in LIGO data. As a member of the LSC, I have also had the opportunity to participate in other projects including searches for gravitational waves from binary neutron stars and black holes, detector characterization, gravitational wave candidate follow-up, and studies of false alarms in the data. This work will not be presented in this thesis.

The LSC is composed of more than 800 members from over 70 institutions across the globe who aid in the detection of gravitational waves from LIGO. Since 2007, LIGO has operated jointly with the Virgo interferometric gravitational-wave detector through what is known as the LIGO-Virgo Collaboration (LVC). Several other interferometers which have contributed to the network or will hopefully contribute in the future include GEO600¹, TAMA², KAGRA³, and INDIGO⁴. The LIGO project is jointly run by Caltech and MIT and is funded by the National Science Foundation. The observatory consists of large-scale interferometric detectors at two sites: Hanford, WA and Livingston, LA. The Virgo detector, with 3 km arms, is located in Cascina, Italy. Figure 1.1 provides an aerial view of the three sites.

The LIGO and Virgo detectors have seen several phases. Initial LIGO (iLIGO) began with construction of the interferometric detectors in 1996. During this initial phase, the Hanford site operated two collocated interferometers: one detector with 4 km long arms (H1) and another with 2 km long arms (H2). The Livingston site operated a second 4 km

¹GEO600 is a 600 m arm length interferometer located in Germany. It has had several joint science runs with LIGO and continues to be operational [3–8].

²TAMA was a 300 m interferometer located in Tokyo and operated by the University of Tokyo. It had several joint searches with LIGO but has been decommissioned as work focuses on KAGRA [9–11].

³KAGRA, previously known as the Large Scale Cryogenic Gravitational Wave Telescope (LCGT) is a proposed detector for the advanced detector era located in the Kamioka mine. It will host two 3 km arm length interferometers [12].

⁴INDIGO, or INDIan Interferometric Gravitational-wave Observatory, is a proposed advanced generation detector to be located in India [13]




Figure 1.1: Aerial views of the three detector sites in the LIGO-Virgo network taken from References [14] and [15]. (a) Hanford, WA site. Home of the 4 km H1 and former home of the 2 km H2 interferometer. (b) Livingston, LA site. Home of the 4 km L1 interferometer. (c) Cascina, Italy site. Home of the 3 km Virgo interferometer.

detector (L1). The iLIGO phase saw five data runs beginning with the first LIGO science run (S1) on August 23, 2002 that lasted a little over two weeks and ending with the fifth science run (S5) that collected data for two years starting on November 4, 2005. The S5 run marked the time when iLIGO reached the sensitivity for which it was designed [1]; the goal was to collect data at this sensitivity until one year’s worth of triple coincident data was collected. Virgo’s first science run (VSR1) overlapped with the end of S5, enabling a four-detector network, H1H2L1V1 [16]. The second phase, *enhanced* LIGO (eLIGO), began with a series of upgrades to the initial detectors to improve sensitivity and culminated with the sixth LIGO science run (S6) that collected data for over a year starting on July 7, 2009 [17]. However, only H1 and L1 were operational; H2 was not. Virgo’s second and third science runs (VSR2 and VSR3) overlapped with S6. Thus, the network for S6 included H1L1V1. At the time of writing, both LIGO and Virgo detectors are undergoing upgrades in preparations for the third phase of development, the advanced detector era. Advanced LIGO (aLIGO) may begin taking data as early as 2014 [18]. Advanced Virgo is following the same schedule [16]. The dates of all completed science runs for the LIGO-Virgo network are shown in Table 1.1.

The iLIGO and eLIGO phases were prototypes for the development of the advanced detectors. While these initial detectors were sensitive enough to detect gravitational waves from strong sources out to several megaparsecs (in some cases, hundreds of megaparsecs), no detections were made although searches are ongoing. In fact, the detection rates for coalescing binary neutron stars for the initial LIGO interferometers were likely 0.02 per year [19]. However, this likely rate jumps to 40 per year for the advanced detector network [19]. Scientists searching for gravitational waves within LIGO data broadly work within four categories based on the signal character of the most likely astrophysical sources [1]:

1. **The Compact Binary Coalescence (CBC) Group:** This group searches for gravitational waves from the coalescence of binary neutron stars and black holes. The signal is composed of an inspiral, merger, and ringdown as the binary objects orbit each other, collide, and form a single perturbed black hole. This signal is well-modeled and thus, the technique of matched filtering is utilized. Previous CBC search results are outlined in [11, 20–32].

2. **The Burst Group:** This group searches for gravitational waves from transient gravitational wave sources that often are not well-modelled such as core-collapse supernova, as well as other unknown sources. Previous burst search results are outlined in [4, 5, 7, 8, 10, 33–52].
3. **The Continuous Wave Group:** This group searches long stretches of data for signals from sources that emit gravitational waves continuously such as pulsars and rotating stars in general. Previous continuous wave search results are outlined in [6, 53–66].
4. **The Stochastic Group:** This group searches for a stochastic background of broadband gravitational waves by correlating noise at different detector sites. A background is expected from gravitational waves from the early universe and possibly, an excess of unresolved sources. Previous stochastic search results are outlined in [67–74].

Each of these search groups often conduct multiple searches tuned for particular mass ranges, targeted time windows, specific parts of the expected signal, etc. Within the CBC group, searches are generally split into low mass and high mass categories. The low mass search looks for gravitational waves from CBCs with total mass $M_{\text{total}} < 35M_{\odot}$ and component mass $\geq 1 M_{\odot}$. LIGO will only be sensitive to the inspiral part of the signal for low mass CBCs; thus, low mass searches utilize matched filtering with an inspiral template. The high mass search looks for gravitational waves from CBCs with total mass between $25 - 100 M_{\odot}$ and component mass between $1 - 99 M_{\odot}$. LIGO will be sensitive to the last part of the inspiral, merger, and ringdown signal for these high mass sources; thus, high mass searches utilize matched filtering with an inspiral, merger, ringdown template. For the highest mass CBCs visible by LIGO, the only part of the signal that can be detected will be the ringdown. Thus, for binary black holes with total masses in the intermediate range, $100 - 450 M_{\odot}$, a matched filter search using only a ringdown template is useful. The ringdown search also can be utilized to look for perturbed black holes in the intermediate mass black hole range that may not have been the product of a compact binary merger but some other process such as a head-on collision or direct collapse of a massive star.

This thesis presents the results of the search in LIGO data for ringdown gravitational waves from the coalescence of intermediate mass black hole binaries (IMBHB) in the mass range $100\text{-}450 M_{\odot}$ and from perturbed intermediate mass black holes in the mass range $50\text{-}800 M_{\odot}$. The details of the search’s sensitivity to these particular sources will be outlined in the following chapters. Evidence for intermediate mass black holes has historically been difficult to come by although recent electromagnetic observations have breathed life into the possibility, and possible abundance, of such objects [75]. However, unambiguous evidence for their existence will only come from the detection of gravitational waves from these objects. In Chapter 2, I review the theory behind the generation of gravitational waves in General Relativity, the most likely sources visible to current large-scale interferometric detectors, and the current state of gravitational wave astronomy. In Chapter 3, I review the theory and properties of black holes with a focus on intermediate mass black holes, the likely source for the ringdown search. In Section 3.5, I also explore the ringdown waveform and how it can provide details of the mass and spin of its source black hole. In Chapter 4, I outline the basic principles of gravitational wave data analysis using the matched filter technique, specifically for ringdown searches and discuss the development of the S5 ringdown search pipeline. I also detail my work to implement a post-processing stage to the pipeline to analyze potential ringdown gravitational wave candidates. The steps taken to properly tune the parameters of the S5 search pipeline are also outlined. Finally, in Section 4.4.3, the results of the S5 search and an analysis of the parameter recovery ability of the ringdown single-mode template are reviewed. Finally, in Chapter 5, I motivate the need for a multi-dimensional, multimode

Table 1.1: LIGO and Virgo Science Runs from the initial, enhanced, and projected advanced eras.

Phase	Run	Start Date	End Date	Duration
iLIGO	S1	23 Aug 2002	8 Sept 2002	17 days
	S2	14 Feb 2003	14 Apr 2003	60 days
	S3	31 Oct 2003	9 Jan 2004	71 days
	S4	22 Feb 2005	24 Mar 2005	31 days
	S5	4 Nov 2005	Oct 2007	2 years
Virgo	VSR1	May 2007	Dec 2007	6 months
eLIGO	S6	7 Jul 2009	20 Oct 2010	1 year
Virgo	VSR2	7 Jul 2009	11 Jan 2010	6 months
Virgo	VSR3	11 Aug 2010	20 Oct 2010	2 months

ringdown search and the possibility of such a search using the reduced basis method to generate compact waveform catalogs.

2. Gravitational Waves

2.1 Gravitational Waves in General Relativity

The central idea behind General Relativity can be summarized in the following schematic form:

$$\left(\begin{array}{c} \text{a measure of local} \\ \text{spacetime curvature} \end{array} \right) = \left(\begin{array}{c} \text{a measure of} \\ \text{matter energy density} \end{array} \right). \quad (2.1)$$

Here we have the profound relationship between the geometry of spacetime and the matter and radiation that fill it, one of the most important consequences of the general theory of relativity published by Albert Einstein in 1916. The relation essentially describes the physical implications of the Einstein field equations which can be stated more precisely as

$$G_{\alpha\beta}(g_{\alpha\beta}) = \frac{8\pi G}{c^4} T_{\alpha\beta} \quad (2.2)$$

where $G_{\alpha\beta} = R_{\alpha\beta} - 1/2 g_{\alpha\beta} R$ is the Einstein tensor constructed from the Ricci tensor $R_{\alpha\beta}$ and the Ricci scalar R which encode information about the curvature of spacetime and $T_{\alpha\beta}$ is the stress-energy tensor which describes the mass density of matter and the current and internal stresses of that matter. The factor $8\pi G/c^4$ is required to match the Newtonian limit for slow motion and weak gravitational fields. The constant G is Newton's constant and c is the speed of light. Note that $G_{\alpha\beta}$ is a function of the metric tensor of spacetime $g_{\alpha\beta}(x^\alpha)$, a symmetric tensor with ten independent components that are functions of the spacetime coordinates x^α . We choose the index α to run over the four values in four-dimensional spacetime such that x^0 is the single time coordinate and x^1 , x^2 , and x^3 are the three spatial coordinates. Thus the Einstein field equations are actually a set of ten nonlinear partial differential equations for the ten components of the metric. From these field equations, we find that the physical effects of "gravity" are actually a result of the curvature of spacetime!

The Einstein field equations do not have a general solution. However, it is possible to give a complete analysis of the solutions of the Einstein equation for space-times that do not differ significantly from flat space-time. We can find the form of the Einstein field equations in the limit of weak gravitational fields, known as the linearized gravity approximation, and in the limit of slow motion, known as the Newtonian limit of General Relativity. The derivations outlined in the following sections are based on derivations found in References [76–78]. For some seminal references on gravity and General Relativity, see [79–81].

2.1.1 Linearized Gravity

From the relationship between the geometry of spacetime and the physical effects of gravity, we can infer that gravity will be weak when the spacetime metric is nearly flat. Hence in the limit of weak gravity, we can assume that the metric will be that of the flat Minkowski spacetime, $\eta_{\alpha\beta}$, plus a small perturbation, $h_{\alpha\beta}$:

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta} \quad (2.3)$$

In rectilinear coordinates (t, x, y, z) , the Minkowski metric is

$$\eta_{\alpha\beta} = \begin{matrix} & t & x & y & z \\ \begin{matrix} t \\ x \\ y \\ z \end{matrix} & \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}. \quad (2.4)$$

From examining the components of the Minkowski metric, we find that the perturbation $h_{\alpha\beta}$ will be small if $|h_{ij}| \ll 1$, $|h_{0i}| \ll c$, and $|h_{00}| \ll c^2$.

We can compute the linearized Riemann tensor, the linearized Ricci tensor, and the Ricci scalar to obtain the linearized Einstein tensor $G_{\alpha\beta}$. We can reduce the number of terms in $G_{\alpha\beta}$ by expressing it in terms of the trace-reversed metric perturbation $\bar{h}_{\alpha\beta}$ instead of the actual metric perturbation $h_{\alpha\beta}$:

$$\bar{h}_{\alpha\beta} := h_{\alpha\beta} - \frac{1}{2}\eta_{\alpha\beta}h \quad (2.5)$$

where $h = -\eta^{\alpha\beta}\bar{h}_{\alpha\beta}$. Further simplification can be achieved by selecting an appropriate coordinate system in which the divergence of the trace-reversed metric perturbation is zero, a condition very similar to the Lorentz gauge in electromagnetism. We then find that the linearized Einstein equations take the very simple form

$$-\square\bar{h}_{\alpha\beta} = \frac{16\pi G}{c^4}T_{\alpha\beta} \quad (2.6)$$

where \square is the d'Alembertian operator, the flat spacetime wave operator:

$$\square = \nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}. \quad (2.7)$$

Thus we find that the linearized Einstein equations become a wave equation with the matter stress-energy tensor acting as a source term. The metric perturbations are actually waves, or ripples in spacetime, and if matter is present, it can generate the waves. These are what we call gravitational waves.

2.1.2 Newtonian Limit of General Relativity

The Newtonian limit refers to conditions that are close to what we normally observe, i.e. slowly moving mass distributions with small internal stresses. Thus we can find a coordinate system in which the components of the stress-energy tensor take familiar forms. The mass energy density component must satisfy

$$T_{00}/c^4 = \rho \quad (2.8)$$

where ρ is the mass density of the material. The components describing the flux of mass must satisfy the condition for non-relativistic slow motion $v \ll c$ such that

$$|T_{0i}/c^3| \sim \rho(v/c) \ll T_{00}/c^4 \quad (2.9)$$

The shear stress components must satisfy the condition for small internal stresses

$$|T_{ij}/c^2| \sim p/c^2 \ \& \ \rho(v/c)^2 \ll T_{00}/c^4 \quad (2.10)$$

where p is the isotropic pressure for a material acting orthogonal to its element of area. Since the matter is slowly moving, the metric perturbation will also be slowly moving and we can replace the d'Alembertian operator with a spatial Laplace operator in the field equations:

$$\nabla^2 \bar{h}_{00} = -16\pi G\rho \quad (2.11)$$

$$\nabla^2 \bar{h}_{0i} = 0 \quad (2.12)$$

$$\nabla^2 \bar{h}_{ij} = 0 \quad (2.13)$$

where the last two sets of equations have trivial solutions $\bar{h}_{0i} = 0$ and $\bar{h}_{ij} = 0$.

2.1.3 Properties of Gravitational Waves In Vacuum

In order to study the properties of gravitational waves, we may begin by exploring equation 2.6 outside the source where $T_{\mu\nu} = 0$ such that

$$\square \bar{h}_{\alpha\beta} = 0. \quad (2.14)$$

In general, we may solve the wave equation above by assuming a solution of the form

$$\bar{h}_{\alpha\beta} = a_{\alpha\beta} e^{ik_\mu x^\mu} \quad (2.15)$$

where $a_{\alpha\beta}$ is a constant, symmetric tensor, $k_\mu = [-\omega, \vec{k}]$ is a constant wave four-vector with $\omega = v\|\vec{k}\| = c\|\vec{k}\|$, and $x^\mu = [t, \vec{x}]$ is the flat-space position four-vector. Clearly, the perturbation travels at the speed of light c . Physical solutions are obtained by taking the real part of equation 2.15 so that we have a monochromatic plane wave

$$\bar{h}_{\alpha\beta} = |a_{\alpha\beta}| \cos(k_\mu x^\mu). \quad (2.16)$$

The Lorentz gauge condition requires that $\partial \bar{h}_{\mu\alpha} / \partial x^\mu = 0$. Thus, from equation 2.16, we have

$$0 = \frac{\partial \bar{h}_{\mu\alpha}}{\partial x^\mu} = -k_\mu |a^{\mu\alpha}| \sin(k_\nu x^\nu) = -k_\mu |a^{\mu\alpha}| \sin(\vec{k} \cdot \vec{x} - \omega t) \quad (2.17)$$

which is satisfied in every case only if $k_\mu |a^{\mu\alpha}| = 0$. The wave four-vector describes the direction of propagation of the wave. For example, the unit vector in the direction of propagation is $\hat{n} = c\vec{k}/\omega$. And the tensor $a_{\alpha\beta}$ describes amplitudes of the various components of the wave. Thus, the Lorentz condition implies that the direction of propagation of the wave is orthogonal to its amplitude components. Hence, we have a transverse wave.

The Lorentz condition also acts to reduce the 10 components of the metric perturbation to six. However, we can reduce this even further by exploiting more gauge freedom and choosing a gauge in which the perturbation is traceless and purely spatial. Then we find that there are only two independent $a_{\alpha\beta}$. If we orient the spatial coordinates so that the direction of propagation is along the z -axis, the only non-zero components will be the 2×2 symmetric matrix in the $x-y$ subspace whose trace vanishes. Thus, we can write the solution of the linearized Einstein equation as

$$h_{\alpha\beta}^{TT} = \begin{matrix} & t & x & y & z \\ \begin{matrix} t \\ x \\ y \\ z \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} & \cos[\omega(t - z/c)] \end{matrix} \quad (2.18)$$

where $h_{\alpha\beta}^{TT}$ is known as the transverse traceless, or TT, gauge metric perturbation. The two independent degrees of freedom h_+ and h_\times are the amplitudes of the “plus” and “cross” polarizations of the wave. The general solution of the linearized Einstein equation is a superposition of waves of different amplitudes of the two polarizations. The consequence of these polarizations for a wave traveling in the z -direction is that only transverse (x-y) separations will change with time.

2.1.4 Properties of Gravitational Waves With Source

We can learn more about the properties of gravitational waves by exploring their production at the source where the near-field metric perturbation may not be small. Equation 2.6 is the linearized version of the Einstein field equations and thus neglects terms that are second order in the metric perturbation. However, we could write

$$\square \bar{h}^{\alpha\beta} = -\frac{16\pi G}{c^4} T^{\alpha\beta} + O(h^2) \quad (2.19)$$

where $O(h^2)$ includes those second order terms and acts as a contribution of the gravitational radiation itself to the overall stress-energy tensor. Then, we can define the effective stress-energy tensor $\tau^{\alpha\beta}$ so that the exact field equations are

$$\square \bar{h}^{\alpha\beta} = -\frac{16\pi G}{c^4} \tau^{\alpha\beta}. \quad (2.20)$$

This is essentially an inhomogeneous wave equation with known solution

$$\bar{h}^{\alpha\beta}(t, \vec{x}) = \frac{4G}{c^4} \int \frac{\tau^{\alpha\beta}(t - \|\vec{x} - \vec{x}'\|/c, \vec{x}')}{\|\vec{x} - \vec{x}'\|} d^3\vec{x}' \quad (2.21)$$

where $t_{\text{ret}} \equiv t - \|\vec{x} - \vec{x}'\|/c$ is the retarded time at which the integrand should be evaluated. This equation can be evaluated in both the far- and near-field zones. In the far-field zone, one makes the approximation that the distance from the source to the field point r is much greater than the gravitational-wave wavelength λ which in turn is much greater than the size of the source R . When one uses these approximations with the TT-gauge, the solution for radially traveling waves is

$$h_{ij}^{TT}(t) \simeq \frac{2G}{c^4 r} \ddot{I}_{ij}^{TT}(t - r/c) \quad (2.22)$$

where I_{ij}^{TT} is the quadrupole tensor in the the TT-gauge. In general the quadrupole tensor may be defined as

$$I^{ij}(t) = \int x^i x^j \tau^{00}(t - r/c, x) d^3\vec{x}. \quad (2.23)$$

If we consider only the matter component of the effective stress energy-tensor, then $\tau^{00} = \rho$ and equation 2.23 is the second mass moment.

In the near-field zone, one makes the approximation that the distance from the source to the field point r is much smaller than the gravitational-wave wavelength λ but is still much greater than the size of the source R . In the Newtonian limit with the TT-gauge, the solution is found to be

$$h_{ij}^{TT}(t) \simeq \frac{2G}{c^4 r} \ddot{\iota}_{ij}^{TT}(t - r/c) \quad (2.24)$$

where ι_{ij} appears as the quadrupolar term in the expansion of the Newtonian potential:

$$\iota^{ij} := \int \left(x^i x^j - \frac{1}{3} r^2 \delta^{ij} \right) \tau^{00}(\vec{x}) d^3 \vec{x}. \quad (2.25)$$

In the TT-gauge, we find that equation 2.23 and equation 2.25 are identical since the two expressions only differ by a trace. Thus, the near-field and far-field solutions have been found to be consistent. In section 2.2, I will demonstrate how we can use these expressions to calculate the gravitational radiation from various astrophysical sources.

2.1.4.1 Energy Carried by Gravitational Waves

A passing gravitational wave produces a real effect on nearby particles indicating that the wave must be carrying energy. The idea of a unique local energy for a gravitational wave cannot be defined in general relativity. However, one can compute an integral average of the effective stress-energy tensor over a region of spacetime that is small compared to background curvature but still large enough to contain several gravitational-wave oscillations. The energy for a plane wave is found to be

$$T_{00}^{GW} = \frac{c^4}{16\pi G} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle. \quad (2.26)$$

The gravitational wave flux and luminosity, respectively, can also be calculated from this energy expression in terms of the quadrupole tensor in equation 2.25 as

$$\frac{dE}{dt dA} = -c^{-1} T_{00}^{GW} = \frac{c^3}{16\pi G} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle = -\frac{G}{8\pi c^5 r^2} \langle \ddot{\iota}_{TT}^{ij} \ddot{\iota}_{ij}^{TT} \rangle \quad (2.27)$$

$$L_{GW} = -\frac{dE}{dt} = \frac{1}{5} \frac{G}{c^5} \langle \ddot{\iota}^{ij} \ddot{\iota}_{ij} \rangle. \quad (2.28)$$

2.1.5 Detection of Gravitational Waves

In order to understand the interaction of gravitational waves with a detector and what sort of detector we might need, we need to select a reference frame to describe the waves and the detector. As we have already seen, gravitational waves take a particularly simple form in the TT-gauge and choosing a gauge is equivalent to selecting a reference frame. Thus we should understand which reference frame corresponds to the TT-gauge. Also, it will be instructive to investigate the interaction in the detector proper frame.

One important tool for this investigation is the geodesic equation, the classical equation of motion of a test particle in the curved spacetime described by the metric $g_{\alpha\beta}$ in the absence of any other external forces:

$$\frac{d^2x^\alpha}{d\tau^2} + \Gamma_{\beta\rho}^\alpha(x) \frac{dx^\beta}{d\tau} \frac{dx^\rho}{d\tau} = 0 \quad (2.29)$$

where τ is the proper time measured by a clock carried along the trajectory of the geodesic and the coefficients $\Gamma_{\beta\rho}^\alpha$ are the Christoffel symbols constructed from the metric and its first derivatives. Another important tool is the equation of geodesic deviation which describes the difference between two nearby geodesics, one parameterized by the curve $x^\alpha(\tau)$ and the other by $x^\alpha(\tau) + \xi^\alpha(\tau)$ where $|\xi^\alpha|$ is much smaller than typical gravitational field variations. This equation takes the form

$$\frac{D^2\xi^\alpha}{D\tau^2} = -R_{\beta\rho\sigma}^\alpha \xi^\rho \frac{dx^\beta}{d\tau} \frac{dx^\sigma}{d\tau} \quad (2.30)$$

where the term on the left is what is known as the covariant derivative of the vector field $\xi^\alpha(x)$ along the curve $x^\alpha(\tau)$ and $R_{\beta\rho\sigma}^\alpha$ is the Riemann curvature tensor. The equation of geodesic deviation basically shows that two nearby geodesics will experience a tidal gravitational force determined by the Riemann tensor.

In the TT frame of reference, the geodesic equation 2.29 takes a very simple form. If we assume that the test particle is at rest at $\tau = 0$, then the derivatives $\frac{dx^i}{d\tau} = 0$ since the particle is at rest. Furthermore, the Christoffel symbol vanishes because of the TT gauge conditions. Thus, in the TT gauge,

$$\frac{d^2x^\alpha}{d\tau^2} = 0 \quad (2.31)$$

so that a particle initially at rest will remain at rest during the passing of a gravitational wave. Furthermore, for similar reasons, the geodesic deviation equation 2.30 will also vanish in the the TT frame. Thus the coordinate separation between two test particles will remain constant as a gravitational wave passes through. However, while the coordinates do not change in the TT gauge, we do find that the proper distances and proper times do change. Consider the TT metric perturbation in equation 2.18 in terms of the plus and cross polarizations. This can be used to write the invariant interval $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$ as

$$ds^2 = -c^2 dt^2 + dz^2 + \{1 + h_+ \cos[\omega(t - z/c)]\} dx^2 + \{1 - h_+ \cos[\omega(t - z/c)]\} dy^2 + 2h_\times \cos[\omega(t - z/c)] dx dy. \quad (2.32)$$

For two events at $(t, x_1, 0, 0)$ and $(t, x_2, 0, 0)$ with coordinate separation $x_2 - x_1 = L$, equation 2.32 gives the proper distance s between the events as

$$s = (x_2 - x_1) \sqrt{[1 + h_+ \cos(\omega t)]} \\ \simeq L \left[1 + \frac{1}{2} h_+ \cos(\omega t) \right]. \quad (2.33)$$

Thus, the proper distance is changing periodically in time.

However, the TT frame is not very applicable to an experimentalist since positions in a lab are not marked by test particles. Instead, the preferred reference frame is the proper detector frame in which the test particle is free to move because of a passing gravitational

wave. Positions can then be measured with respect to a rigid ruler and origin. Earth-based detectors will essentially be in a frame that accelerates (due to Earth's gravity) and rotates (like a Foucault pendulum). The path of a test particle can then be described by Newtonian equations of motion in terms of forces. There are terms proportional to the Riemann curvature tensor from the gravitational field of the Earth and gravitational waves but also terms from static gravitational forces, Coriolis forces, etc. In order to combat these other forces and leave only the part proportional to the Riemann tensor, experimentalists use several tricks. Gravitational waves with low frequencies (less than a few Hz) will have to compete with slowly varying Newtonian noises. Thus Earth-based detectors look for higher frequency gravitational waves. Furthermore, the acceleration in the z -direction due to Earth's gravity can be combated with suspension systems allowing the test mass to move freely (i.e. in a freely falling frame) in the $x - y$ plane. The equation of geodesic deviation in the proper detector frame is

$$\ddot{\xi}^i = \frac{1}{2} \ddot{h}_{ij}^{TT} \xi^j \quad (2.34)$$

which allows us to immediately write the effect of gravitational waves on a test particle of mass m in terms of a Newtonian force without need for general relativity:

$$F_i = m \ddot{\xi}^i = \frac{m}{2} \ddot{h}_{ij}^{TT} \xi^j. \quad (2.35)$$

The effects of gravitational waves cannot be seen in isolated bodies. This is a result of the fact that a single test mass, in a frame freely falling with it, will remain at rest. At least two test masses are required to measure the effects of gravitational waves. This is also the case when one wants to measure any curvature of spacetime. Consider the ring of test masses in figure 2.1. Let the ring lie in the $x - y$ plane and the gravitational wave propagate in the z -direction. For the plus polarization at $z = 0$ and initial conditions $h_{ij}^{TT} = 0$ at $t = 0$, equation 2.18 reduces to

$$h_{\alpha\beta}^{TT} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} h_+ \sin \omega t \quad (2.36)$$

If the displacement between geodesics is $\xi_\alpha(t) = (x_0 + \delta x(t), y_0 + \delta y(t))$, then equation 2.36 substituted into equation 2.34 gives us

$$\delta \ddot{x} = -\frac{h_+}{2} (x_0 + \delta x) \omega^2 \sin \omega t, \quad (2.37)$$

$$\delta \ddot{y} = \frac{h_+}{2} (y_0 + \delta y) \omega^2 \sin \omega t \quad (2.38)$$

where we can neglect δx and δy . Then integration gives us

$$\delta x(t) = \frac{h_+}{2} x_0 \sin \omega t, \quad (2.39)$$

$$\delta y(t) = -\frac{h_+}{2} y_0 \sin \omega t. \quad (2.40)$$

Similarly, for the cross polarization at $z = 0$ and initial conditions $h_{ij}^{TT} = 0$ at $t = 0$, we find

$$\delta x(t) = \frac{h_x}{2} y_0 \sin \omega t, \quad (2.41)$$

$$\delta y(t) = \frac{h_x}{2} x_0 \sin \omega t. \quad (2.42)$$

Equations 2.39-2.42 describe the changes in the x and y components for a passing gravitational wave. The plus polarization alternately stretches and compresses the ring of test masses in the x and y directions while the cross polarization exhibits the same behavior rotated by 45° in the $x - y$ plane. This feature of gravitational wave interaction is exploited in the detection methods which will be explored in more detail in Section 2.3.




Figure 2.1: The effects of plus and cross polarization on a ring of test masses. The plus polarization alternately compresses and stretches the x - and y -separations. The cross polarization has the same effect only rotated by 45° .

2.2 Sources of Gravitational Waves

Many astrophysical processes will result in the production of gravitational waves because the dynamic motion of the system will not result in perfect spherically symmetric changes in the mass distribution. Since the simplest form of non-spherically symmetric deformation is the quadrupole moment, this means the system will have a time-varying quadrupole moment and we can evaluate the expression for the metric perturbation in Equation 2.22. We can compute the amplitude and frequency of the gravitational waves for various sources. Then,

with these values in mind, we can explore the sensitivity of gravitational wave detectors and their prospects for detecting waves from these astrophysical sources.

We can very roughly estimate the amplitude of the perturbation from Equation 2.22 in terms of the non-symmetric part of the kinetic energy $E_{\text{kin}}^{\text{non-symm}}$ of a dynamic astrophysical source as

$$h \simeq \frac{2G}{c^4 r} \ddot{I} \sim \frac{G (E_{\text{kin}}^{\text{non-symm}}/c^2)}{c^2 r}. \quad (2.43)$$

Rather optimistically, we can set $E_{\text{kin}}^{\text{non-symm}} \sim 1 M_{\odot}$. Most sources will radiate significantly less energy as gravitational waves. Then, for a source in the Virgo cluster at an intergalactic distance of 16.5 Mpc away, the amplitude of the perturbation on Earth would be $h \sim 10^{-21}$. And for a source at a Hubble distance of 4200 Mpc away, the amplitude of the perturbation would be $h \sim 10^{-23}$ [82].

Also, interferometric gravitational wave detectors are broadband detectors and are sensitive to a range of frequencies as we will explore in more detail in Section 2.3. Thus it is important to understand the frequency of emitted gravitational waves to ascertain whether they will fall within a detector's frequency sensitivity band. In order to set an upper limit on the frequency, we can assume the source can be no smaller than its Schwarzschild radius, $R_S = 2GM/c^2$; otherwise, it will collapse into a black hole. We will explore this concept more in Chapter 3. Furthermore, we can estimate that the shortest period of gravitational wave emission cannot be less than the light travel time around its circumference, $\tau = 4\pi GM/c^3$. Then the maximum frequency of the gravitational waves from a source is approximately

$$f \leq \frac{c^3}{4\pi GM} \sim 10^4 \text{ Hz} \frac{M_{\odot}}{M}. \quad (2.44)$$

Thus, we see that the frequency of the emitted gravitational waves is inversely proportional to the mass of the source [82].

2.2.1 Compact Binary Coalescence

A binary system can be analyzed fairly simply as two point particles with masses m_1 and m_2 in a circular orbit in the $x - y$ plane so that the orbital angular momentum vector points along the z -axis. The distances of the two particles from the z -axis are r_1 and r_2 , respectively, so that their orbital separation is $a = r_1 + r_2$ and the orbital phase is $\phi = \omega t$. Other useful quantities to define are the total mass $M := m_1 + m_2$ and reduced mass $\mu := m_1 m_2 / M$.

For an observer at a distance r and angle of inclination ι from the binary, the plus and cross polarizations of the waveform are:

$$h_+ = -\frac{2G\mu}{c^2 r} (1 + \cos^2 \iota) \left(\frac{v}{c}\right)^2 \cos 2\phi(v) \quad (2.45)$$

$$h_{\times} = -\frac{4G\mu}{c^2 r} (\cos \iota) \left(\frac{v}{c}\right)^2 \sin 2\phi(v) \quad (2.46)$$

where $v = (\pi GM f)^{1/3}$ and f is the gravitational wave frequency [76]. An important thing to note about a binary system is that the loss of energy from the system will, over time, cause the orbit to decay, the frequency to increase, and the amplitude of the radiation to increase. As the orbit decays, the phase ϕ does not increase uniformly with time which is why the above equations show the phase as a function of v . The characteristic gravitational

waveform is known as a chirp or inspiral. Eventually, the two objects will be close enough to merge together and the calculation of their merger waveforms requires computationally expensive simulations on a supercomputer. Finally after the objects have merged, the new single object will pulsate in an excited state as it tries to return to equilibrium. This phase, known as the ringdown phase, is well-modeled as a series of quasi-normal modes. Chapter 3 is dedicated to a thorough discussion of this inspiral-merger-ringdown waveform. As we will see, LIGO will be sensitive to a binary system composed of either two neutron stars, two black holes, or one of each. These objects are known as compact objects, hence the term *compact* binary coalescence.

We can use Equation 2.43 to estimate the strain expected from a 1.4-1.4 M_{\odot} binary neutron star system in the Virgo cluster that has a high orbital frequency of 500 Hz. In general, for a binary of total mass M , separation a , and orbital period T , the non-symmetric part of the kinetic energy is $E_{\text{kin}}^{\text{non-symm}} = M (\pi a/T)^2$ and using Kepler's law, the perturbation amplitude becomes

$$h \sim 10^{-20} \left(\frac{6.3 \text{kpc}}{r} \right) \left(\frac{M}{2.8 M_{\odot}} \right)^{5/3} \left(\frac{T}{1 \text{s}} \right)^{-2/3} \sim 10^{-22} \quad (2.47)$$

for our 1.4-1.4 M_{\odot} neutron star binary example.

The Hulse-Taylor binary pulsar was the first binary neutron star system discovered. General Relativity predicted an orbital decay of the system of $\dot{P} = 2.402 \times 10^{-12}$, which matches very well with the observed decay rate as shown in Figure 2.2 providing good indirect evidence for the existence of gravitational waves!

2.2.2 Bursts

Unmodeled bursts describe sources of gravitational waves which produce a short duration of excess strain. Many of these sources do not have a well-modeled form including asymmetric supernova collapse. However, the merger signal of a compact binary coalescence, which would be considered a burst, has recently been modeled for a variety of systems via numerical relativity simulations.

As an example, consider the axisymmetric collapse of an initially rotating stellar core to form a neutron star which will have radius $R \sim 15$ km and mass $M \sim 1 M_{\odot}$. During the collapse, the neutron star will form into an axisymmetric ellipsoid with eccentricity $e = 0.1$ and principal moments of inertia $I_{11} = I_{22} = 2/5 MR^2 (1 - 1/2e^2)$ and $I_{33} = 2/5 MR^2$. An observer at a distance r on a line perpendicular to the line of collapse would measure a perturbation amplitude of

$$h \sim 10^{-20} \left(\frac{10 \text{kpc}}{r} \right) \quad (2.48)$$

where 10 kpc is a typical distance to supernovae in the Milky Way [76]. An observer along the line of collapse would not see any gravitational waves.

2.2.3 Continuous Sources

A source that emits gravitational waves for a long period of time (typically more than a few minutes, oftentimes much more than a year) is known as a continuous gravitational wave source. Such sources would typically be high frequency rotating neutron stars with a non-axial deformation or low frequency binary systems composed of white dwarfs or black




Figure 2.2: This figure shows the cumulative periastron period shift in seconds of the famous Hulse-Taylor binary pulsar PSR B1913+16 from 1975 to 2005. The red data points match exceptionally well with the theoretical blue curve prediction for the period shift due to gravitational wave emission as predicted by General Relativity [83]. The discovery and analysis of this binary star system earned its discoverers, Hulse and Taylor, the 1993 Nobel Prize in Physics.

holes far from merger. Typically, these sources can be modeled by nearly fixed frequency sinusoids.

In the high frequency regime, rapidly rotating, distorted neutron stars could be a source of high frequency gravitational waves. Very simply, we could model such a system as a rotating triaxial ellipsoidal bar spinning about its z -axis, principal axis of inertia I_3 . We can assume that the quadrupole tensor is diagonal at $t = 0$ and that the bar is spinning with angular velocity ω ; but for an observer at an inclination angle ι , we need to rotate the second time derivative of the quadrupole tensor, take the TT projection, and use equation 2.22 to compute the gravitational wave perturbation. This perturbation takes the form:

$$h_{ij}^{TT} = \frac{2G}{c^4 r} \ddot{I}_{ij}^{TT} = \frac{4G\epsilon I_3 \omega^2}{c^4 r} \begin{pmatrix} -\frac{1}{2}(1 + \cos^2 \iota) \cos 2\omega t & \cos \iota \sin 2\omega t & 0 \\ \cos \iota \sin 2\omega t & \frac{1}{2}(1 + \cos^2 \iota) \cos 2\omega t & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2.49)$$

where $\epsilon = (I_1 - I_2)/I_3$ is the ellipticity. It is a function of the three principal axes of inertia around the x -, y -, and z -axis, respectively. Since this perturbation is transverse and traceless, we can immediately read off the two polarizations [76]:

$$h_+ = -\frac{4G\epsilon I_3 \omega^2}{c^4 r} \frac{1 + \cos^2 \iota}{2} \cos 2\omega t, \quad (2.50)$$

Table 2.1: Physical parameters of the Crab pulsar. Data taken from [76].

Rotation period	$P = 0.0333$ s
Spin-down rate	$\dot{P} = 4.21 \times 10^{-13}$
Mass	$M = 1.4 M_{\odot}$
Radius	$R = 10$ km
Mass quadrupole	$I_3 \approx \frac{2}{5} MR^2 = 1.1 \times 10^{38} \text{m}^2 \text{kg}$
Inclination of spin axis	$\iota = 62^\circ$
Distance	$r = 2.5$ kpc

$$h_{\times} = \frac{4G\epsilon I_3 \omega^2}{c^4 r} \cos \iota \sin 2\omega t. \quad (2.51)$$

Consider, for example, the Crab pulsar, a young, rapidly rotating neutron star. It has the approximate physical properties listed in Table 2.1. If the listed spin-down rate (or slow-down rate) assumes loss of angular momentum and energy only due to gravitational radiation (which is not a good approximation), these parameters give an over-estimated ellipticity of $\epsilon \simeq 7.2 \times 10^{-4}$. The largest allowed neutron star deformity from models is $\epsilon \simeq 10^{-6}$. In the case that a perturbation takes the form of a periodic plane wave, as it does in this case of a distorted neutron star, we can read the plus and cross amplitudes, A_+ and A_{\times} , directly from the polarization equations 2.50 and 2.51. Thus, for the over-estimated ellipticity, we would receive gravitational waves from the Crab pulsar with over-estimated amplitudes:

$$A_+ = -\frac{4G\epsilon I_3 \omega^2}{c^4 r} \frac{1 + \cos^2 \iota}{2} = 7.3 \times 10^{-25}, \quad (2.52)$$

$$A_{\times} = \frac{4G\epsilon I_3 \omega^2}{c^4 r} \cos \iota = 5.6 \times 10^{-25}. \quad (2.53)$$

2.2.4 Stochastic Background

When thinking about a stochastic background of gravitational waves, we can no longer focus on individual sources but rather a number of sources that overlap in time and frequency and lead to source confusion. The sum of the signals at any given time and frequency will be a random variable drawn from a Gaussian distribution, hence the term stochastic background.

One possible source of a stochastic background might be the cosmic gravitational wave background (CGWB) analogous to the cosmic microwave background (CMB). This would be caused by density fluctuations in the early universe and would connect us to about 10^{-35} s after the Big Bang. However, to measure waves from this epoch, we would need much more sensitive detectors than the ground-based interferometers available [76].

Another possible background could come from astrophysical sources. These possible sources include a population of rotating neutron stars and a population of white dwarf binaries that would be important mostly for a space-based interferometers such as LISA [76].

2.3 Gravitational Wave Detectors

The first gravitational wave detectors were resonant-mass cylindrical bar detectors. The first was developed and built by Joseph Weber in the 1960s [84]. Over the course of the next several decades, more bar detectors were built that were at least four orders of magnitude more sensitive than Weber's original design [85]. These ~ 2 ton detectors would be set into oscillation at their resonant frequency by passing gravitational waves near that resonant frequency. These detectors were sensitive to gravitational waves with relatively high frequency (~ 1 kHz) and in a narrow frequency band. Some bar detectors are still operating and some have been decommissioned, including ALLEGRO at LSU [86]. To have a chance at detecting signals in a broader range of frequencies and out to farther distances, we have moved to large-scale interferometric detectors. The idea originated with the Russian theorists, M. Gertsenshtein and V. I. Pustovoit in 1962 [87]. But the strong push for using interferometers came in the late 1960s with R. Forward, R. Weiss, R. Drever, and others. From the early 2000s, several kilometer-scale ground-based interferometers operated in the frequency band from ~ 10 Hz to ~ 1 kHz at a sensitivity that allowed the potential for detection from a variety of sources at large extragalactic distances as outlined in Section 2.2. More description of these detectors will be given in section 2.3.2. These broadband interferometric detectors are currently the state-of-the-art. Thus, they will be the focus of the rest of this chapter.

An interferometric detector is based on the simple principles of a Michelson interferometer. In general, an interferometer uses laser light as a displacement measuring device. Two highly reflecting mirrors and a central beam splitter (a 50% reflecting mirror) act as test masses. The three masses are suspended as pendulums and are allowed to swing freely in the horizontal direction. Thus, they can be considered freely falling masses in that direction. The reflecting mirrors are placed equidistant from the central beam splitter but in orthogonal directions. In Figure 2.3, these are labeled Test Mass for the x - and y -directions and Beam Splitter, respectively. A beam of laser light is incident on the beam splitter and is split into two beams along the x and y arms. The light is then reflected by the end mirrors, recombined at the beam splitter, and detected at the photodiode. When the beams recombine, they will interfere constructively if the lengths of the two arms differ by an integral number of wavelengths and interfere destructively if the lengths differ by an odd number of half wavelengths:

$$\Delta L \equiv L_x - L_y = n\lambda, \quad \text{constructive interference} \quad (2.54)$$

$$\Delta L \equiv L_x - L_y = \left(n + \frac{1}{2}\right)\lambda, \quad \text{destructive interference} \quad (2.55)$$

where $n = 0, 1, 2, \dots$

Following the discussion in section 2.1.5, we saw that the effect of the plus and cross polarizations, h_+ and h_\times , on a ring of test masses is to alternately stretch and compress the ring in the $x - y$ plane as the wave propagates in the z direction. This introduces a relative temporary change in the x length with respect to the y length and vice-versa. It is this relative change in distance between test masses on a ring that can be measured with gravitational wave detectors.




Figure 2.3: Optical layout of a Michelson interferometer. Diagram taken from [88].

Suppose we reduce our ring of test masses to an interferometer with two test masses equidistant from a beamsplitter at the origin- one at a distance x from the origin and the other at a distance y from the origin. Consider only a plus polarized gravitational wave traveling in the z direction. The proper distance between the test masses in the curved spacetime of a passing gravitational wave is given by equation 2.32:

$$\begin{aligned}
 ds^2 &= -c^2 dt^2 + dz^2 + \{1 + h_+ \cos[\omega(t - z/c)]\} dx^2 + \\
 &\quad \{1 - h_+ \cos[\omega(t - z/c)]\} dy^2 + 2h_\times \cos[\omega(t - z/c)] dx dy \\
 &= -c^2 dt^2 + \{1 + h\} dx^2 + \{1 - h\} dy^2.
 \end{aligned} \tag{2.56}$$

For interferometric detectors, the light path difference is affected. The geodesic for a photon in this light path is a null geodesic, meaning $ds^2 = 0$. Thus, equation 2.56 simplifies to give two sets of equations for the changing path length of an interferometer's x and y arms, respectively:

$$dx = \frac{cdt}{\sqrt{1 + h}} \approx c\left(1 - \frac{h}{2}\right)dt = L_x dt, \tag{2.57}$$

$$dy = \frac{cdt}{\sqrt{1 - h}} \approx c\left(1 + \frac{h}{2}\right)dt = L_y dt. \tag{2.58}$$

where we have used an approximation for the square roots in the above equations since we can assume that the perturbation h is very small. We can then compute the difference in path length between the x and y arms:

$$L_x - L_y = chdt = Lh \quad (2.59)$$

where $L = cdt$ is the unperturbed path length. Thus we see that a gravitational wave traveling in the z direction produces a *strain* (a term borrowed from engineering), or fractional change in distance, in the length of the x and y arms of an interferometric detector: $h = \Delta L/L$. There will be no change in separation between two masses lying along the z axis. However, a realistic detector will not necessarily be perpendicular to an incoming gravitational wave. The actual change in length for an arbitrary orientation can be calculated as well and will be smaller than for the orthogonal case.

Thus, equation 2.59 indicates that a gravitational wave will cause a strain in the detector and equations 2.54 and 2.55 indicated that such a strain will cause a phase shift that can be detected with a photodetector.

2.3.1 Toward a Real Detector

The basic design of an interferometric detector described above is a good starting point for understanding the layout of a real gravitational wave detector. In general, if the arms of the detector are the same length (up to an integer wavelength), the light, when recombined at the beamsplitter, will interfere constructively and return entirely towards the laser (the symmetric output direction). However, if the arm length differs by an amount other than an integer wavelength, then there will be some constructive interference in the direction opposite the laser (the antisymmetric output direction). A photodiode placed at the antisymmetric output will detect some of this light. A gravitational wave that changes the arm lengths will cause some light to reach the antisymmetric output photodiode. This is how gravitational waves are detected. We will assume initially that a measurable change in the interference pattern would be on the order of the wavelength of the laser light, $\Delta l \sim \lambda_{\text{laser}}$, a very conservative change. Thus, for an infrared laser with wavelength $1 \mu\text{m}$, with this simple design, we could measure perturbations in a kilometer-scale interferometer on the order of

$$h := \frac{\Delta l}{l} \sim \frac{\lambda_{\text{laser}}}{l} \sim \frac{10^{-6}}{10^3} = 10^{-9}. \quad (2.60)$$

Recall that a gravitational wave detector should be able to detect strains on the order of $h \sim 10^{-21}$ in order to have a chance at finding the sources described in Section 2.2. This is much too large! But we can do better with the following improvements.

Increase the optical path length

Examining the above equation, it is clear that if we increase the optical path l , we can increase our strain sensitivity. Due to physical constraints, an earth-based interferometric detector cannot be much larger than $l \sim 1 \text{ km}$. However, by creating an optical cavity in each arm, we can increase the effective length of the arm by forcing the light to travel many times up and down an arm before it recombines with light from the other arm. This type of cavity, known as a Fabry-Perot cavity, is created by two sets of mirrors along an optical axis- the partly transmissive input test mass (ITM) and the highly reflective end test mass (ETM). With careful tuning, there is a buildup of light in the cavity. See Figure 2.4. This is helpful to our sensitivity until the effective arm length l_{eff} approaches the gravitational wave wavelength $\lambda \sim 1000 \text{ km}$. Thus we could measure perturbations on the order of




Figure 2.4: Optical layout for the LIGO interferometers. The power recycling cavity is formed by the power recycling mirror (PRM), the beam splitter (BS), and input test masses (ITMX, ITMY). The Fabry-Perot cavities are formed by the ITMs and the end test masses (ETMs). This diagram is taken from [88].

$$h \sim \frac{\Delta l}{l_{\text{eff}}} \sim \frac{\lambda_{\text{laser}}}{\lambda_{\text{GW}}} \sim \frac{10^{-6}}{10^6} = 10^{-12}. \quad (2.61)$$

Still too large!

Improve optical path length measurement

Again, examining the strain equation, we see that if we can improve our ability to measure changes in the optical path lengths beyond the wavelength of the laser light, we can improve our strain sensitivity. Better sensitivity comes with a better photodiode. The primary limitation in measuring a small change in light is shot noise, or the natural fluctuations in the rate of photons arriving at the photodiode, that follow a Poisson process. Let a photodiode collect N_{photons} over time τ with a natural fluctuation of $\sim \sqrt{N_{\text{photons}}}$. In order to detect a change in light above the natural fluctuations, we require

$$\Delta l \sim \frac{\sqrt{N_{\text{photons}}}}{N_{\text{photons}}} \lambda_{\text{laser}} \quad (2.62)$$

Photons will be collected for about the period of a gravitational wave, $\tau \sim 1/f_{\text{GW}}$. Thus, the total number of photons that will be collected depends on the gravitational wave frequency f_{GW} and also, of course, on the power of the laser, P_{laser} :

$$N_{\text{photons}} = \frac{P_{\text{laser}}}{E_{\text{photon}}} \tau \sim \frac{P_{\text{laser}}}{hc/\lambda_{\text{laser}}} \frac{1}{f_{\text{GW}}} \quad (2.63)$$

where h is Planck's constant. For a laser with $P_{\text{laser}} = 1\text{W}$ and $\lambda_{1\mu\text{m}}$ and a gravitational wave with $f_{\text{GW}} \sim 300\text{Hz}$, the number of photons is found to be

$$N_{\text{photons}} = \frac{1\text{W}}{(6.626 \times 10^{-34}\text{J} \cdot \text{s}) (3 \times 10^8\text{m/s}) / (1\mu\text{M})} \frac{1}{300\text{Hz}} \sim 10^{16} \text{ photons}. \quad (2.64)$$

Then, the perturbation in equation 2.61 can be calculated as

$$h \sim \frac{\Delta l}{l_{\text{eff}}} \sim \frac{\sqrt{N_{\text{photons}}} \lambda_{\text{laser}}}{N_{\text{photons}} \lambda_{\text{GW}}} \sim \frac{\sqrt{10^{16}} 10^{-6}\text{m}}{10^{16} 10^6\text{m}} = 10^{-20}. \quad (2.65)$$

This is very close to the required sensitivity!

Increase laser power

From equation 2.63, it is clear that we can increase the N_{photons} by increasing the laser power P_{laser} . Of course, there are practical limits to how powerful (and stable) a laser can be but it is possible to increase the effective power of a laser by several tricks. In Section 2.3.1, it was explained that if the arms of the interferometer are the same length, all the recombined light heads toward the symmetric output direction. However, if there is an arm length difference, some of that light will head toward the antisymmetric output direction. But a significant amount still enters the symmetric output direction and this light can be recovered and reflected back into the interferometer. This is known as power recycling. See Figure 2.4 where we have labeled this mirror the PRM. Too much power built up in the arms can lead to too much radiation pressure on the mirrors. This can be combated with larger mirrors. Currently, the improvement in strain sensitivity due to power recycling is about one order of magnitude. This puts the sensitivity right where we need it: $h \sim 10^{-21}$.

Advanced interferometry

Advanced detectors including Advanced LIGO (aLIGO) are currently being built and tested. Scientists have even begun looking beyond the advanced detector era to third generation detectors. While aLIGO and beyond are outside the scope of this thesis, there are several interesting features to help with noise reduction that should be mentioned. A further discussion of noise sources can be found in Section 2.3.1.

Quantum squeezing of light is expected to reduce the effect of shot noise. More massive mirrors with lower absorptivity will combat radiation pressure allowing more power in the arms to further reduce shot noise. Improved seismic isolation will decrease low frequency noise. Higher quality optics will help with thermal noise. Additional signal recycling will be implemented using reflected light leaving the interferometer from the antisymmetric output direction changing the nature of shot noise and allowing for easier tuning of the interferometer. Third generation detectors aim to combat seismic noise by operating underground and by monitoring gravity gradient noise so that it may be actively cancelled. Also, cryogenic cooling of mirrors and suspension components will help with thermal noise.

Calibration of the Detectors

Interferometric detectors must maintain resonance in the various cavities so that the antisymmetric port remains dark. Environmental noise constantly moves the mirrors and acts

to destroy the resonance (i.e. the detector “loses lock”). Thus a feedback loop was designed to keep the differential arm length constant. However, the antisymmetric port is where the effects of a passing gravitational wave will be seen so it is important to determine the effect of this feedback loop on a gravitational wave signal.

The control loop is diagrammed in Figure 2.5. In the absence of a feedback loop, the sensing function $C(f)$ describes how the interferometer filters an external strain $s_{\text{ext}}(f)$ to produce a readout signal known as the error signal $e(f)$. This error signal is read out by the channel `DARM_ERR` in LIGO which is sampled in counts at a rate of 16384 Hz. In order to control the effects of low frequency seismic noise on the mirrors, the error signal $e(f)$ is low-pass filtered by a digital filter $D(f)$ to produce a control signal $d(f)$. It is sent to the mirrors as an actuation. This control signal is known as `DARM_CTRL` in LIGO. Finally, an actuation function $A(f)$ relates the actuation force to the motion of the mirrors to produce a control strain $s_c(f)$ that is fed back to the interferometer. When working in the frequency domain, it is straightforward to express this feedback loop mathematically.

The residual strain $s_{\text{res}}(f)$ experienced by the interferometer is the difference between the external strain $s_{\text{ext}}(f)$ and the control strain $s_c(f)$:

$$s_{\text{res}}(f) = s_{\text{ext}}(f) - s_c(f). \quad (2.66)$$

The error signal $e(f)$ results when the residual strain $s_{\text{res}}(f)$ is filtered by the sensing function (or Fabry-Perot cavity transfer function) $C(f)$:

$$e(f) = C(f)s_{\text{res}}(f) \quad (2.67)$$

which is then filtered by the digital filter $D(f)$ to produce the differential arm control signal $d(f)$:

$$d(f) = D(f)e(f). \quad (2.68)$$

Finally, the actuation function $A(f)$ acts on this to produce the control strain:

$$s_c(f) = A(f)d(f). \quad (2.69)$$

Ultimately, we want to express the external strain in terms of the transfer functions of the feedback loop and the readout error signal that we measure. We can use Equations 2.66-?? to express the error signal:

$$\begin{aligned} e(f) &= C(f)s_{\text{res}}(f) = C(f)[s_{\text{ext}}(f) - s_c(f)] \\ &= C(f)s_{\text{ext}}(f) - C(f)A(f)d(f) \\ &= C(f)s_{\text{ext}}(f) - C(f)A(f)D(f)e(f). \end{aligned} \quad (2.70)$$

Then, solving this for the external strain yields

$$s_{\text{ext}}(f) = \frac{1 + C(f)A(f)D(f)}{C(f)}e(f) = R(f)e(f) \quad (2.71)$$

where $R(f)$ is what is known as the response function

$$R(f) = \frac{1 + G(f)}{C(f)} \quad (2.72)$$

and $G(f)$ is known as the open loop gain:

$$G(f) = C(f)A(f)D(f). \quad (2.73)$$




Figure 2.5: The differential arm feedback control loop of an interferometric detector.

In Chapter 4, I will discuss how one can extract a gravitational wave signal $h(f)$ from the external strain $s_{\text{ext}}(f)$ that will also contain noise $n(f)$.

Antenna Response

In general, we cannot expect that a gravitational wave will propagate from directly above or directly below the interferometric detector. This particular configuration would be used to describe an “optimally positioned and oriented” source, one that provides the maximum strain possible in the detector for a source of that type. But, in fact, we expect most sources to *not* be optimally positioned and oriented but rather, to be propagating from some random direction in the sky and from a random orientation. In such cases, the detector will only capture part of the signal. Consider, for example, the generic case in Figure 2.6. A gravitational wave source in a frame measured by (x, y, z) emits waves, a portion of which travel along the red line connecting source to detector. The source itself is rotating with angular momentum pointing in the \hat{z} direction. This makes an angle ι with the \hat{z}' direction in the detector frame. This angle is a measure of the inclination of the source and is known as the inclination angle. The wave propagates along the direction \hat{z}'' from the source to the detector in a frame measured by (x'', y'', z'') . The perpendicular plane (x'', y'') makes an angle Ψ with the (x', y') plane of the detector frame. This is called the polarization angle. The detector arms lie along (x', y') and the propagation direction makes an angle ϕ with the x' axis and an angle θ with the z' axis.

We can write the actual strain produced at the detector using this arbitrary orientation. The plus and cross detector beam functions are

$$F_+ = -\frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi \quad (2.74)$$




Figure 2.6: Cartoon demonstrating relative positions and orientations of a source frame (x, y, z) , polarization frame of the gravitational wave in the direction of the detector (x'', y'', z'') , and detector frame (x', y', z') .

$$F_{\times} = \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \sin 2\psi - \cos \theta \sin 2\phi \cos 2\psi. \quad (2.75)$$

Then, the strain produced at the detector for an arbitrary configuration is

$$h(t) = h_{+}(t)F_{+}(\theta, \phi, \Psi) + h_{\times}(t)F_{\times}(\theta, \phi, \Psi). \quad (2.76)$$

Note that for $\theta = 0^\circ$ and $\Psi = 45^\circ$, the effect in both arms will be the same and no phase shift is produced.

From these expressions, we can define an effective distance D_{eff} . For any an arbitrarily located source with a position, polarization, and inclination at a distance D , we can define the optimally located and positioned source at distance

$$D_{\text{eff}} = \frac{D}{\sqrt{F_+^2 (1 + \cos^2 \iota)^2 + F_\times^2 \cos^2 \iota}}. \quad (2.77)$$

This definition will prove useful in Chapter 4.

The Noise Spectral Density

The various noises of the detector can be conveniently characterized by a spectral strain sensitivity with dimensions of $1/\sqrt{\text{Hz}}$. Recall from Section 2.3.1 that the detector output $s(t)$ is composed of instrumental noise $n(t)$ arising from naturally occurring random processes and a potential strain signal $h(t)$

$$s(t) = n(t) + h(t). \quad (2.78)$$

The detection problem then becomes how to distinguish $h(t)$ from $n(t)$ when $h(t) \ll n(t)$. In a way, $n(t)$ provides a measure of how small an $h(t)$ we can detect. Thus we take $n(t)$ as the detector's noise and have a convenient way to compare performances of different detectors. Instrumental noise is a random process. In the case that we have a stationary random process (which can be the case for detector noise for short periods of time), the expectation value of n at some time t becomes a long time average (as opposed to an ensemble average):

$$\langle n \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} n(t) dt. \quad (2.79)$$

For simplicity, let us assume that $\langle n \rangle = 0$. We can still define the power in the signal by integrating $n^2(t)$ over some duration T and then dividing by T . Again, the expectation value of $n^2(t)$ will just be the time average:

$$\langle n^2 \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} n^2(t) dt \quad (2.80)$$

If $n(t) = 0$ outside of $-T/2 < t < T/2$, then we can extend the integral limits to infinity to take advantage of Parseval's theorem:

$$\begin{aligned} \langle n^2 \rangle &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\infty}^{\infty} n^2(t) dt \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\infty}^{\infty} |\tilde{n}(f)|^2 df \\ &= \lim_{T \rightarrow \infty} \frac{2}{T} \int_0^{\infty} |\tilde{n}(f)|^2 df \\ &= \int_0^{\infty} S_n(f) df \end{aligned} \quad (2.81)$$

where S_n is what is known as the power spectral density of the noise process $n(t)$ and $\tilde{n}(f)$ is the Fourier transform of $n(t)$. Thus, in general, the power spectral density of a stationary random process $n(t)$ is defined as

$$S_n(f) := \lim_{T \rightarrow \infty} \frac{2}{T} \left| \int_{-T/2}^{T/2} n(t) e^{-2\pi i f t} dt \right|^2 \quad (2.82)$$

It can also be shown that, for stationary processes, the power spectral density is twice the Fourier transform of the autocorrelation function $R_n(\tau) = \langle x(t)x(t+\tau) \rangle$:

$$S_n(f) = 2 \int_{-\infty}^{\infty} R_n(\tau) e^{-2\pi i f \tau} d\tau. \quad (2.83)$$

Another useful expression for the power spectral density can be found by finding the expectation value of the frequency components of the noise:

$$\langle \tilde{n}^*(f') \tilde{n}(f) \rangle = \frac{1}{2} S_n(f) \delta(f - f'). \quad (2.84)$$

Since the noise is real, $\tilde{n}(-f) = \tilde{n}(f)$ and thus, $S_n(-f) = S_n(f)$. If $n(t)$ has no dimensions, then $S_n(f)$ has units of Hz^{-1} . If $\langle \tilde{n}^*(f') \tilde{n}(f) \rangle$ is obtained by integrating only over the physical range of frequencies, $f > 0$, then $S_n(f)$ is known as the *one-sided spectral density*.

Finally, on further definition that will come in handy is the inner product of a and b :

$$(a|b) = 2 \int_{-\infty}^{\infty} \frac{\tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} df = 4\Re \left[\int_0^{\infty} \frac{\tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} df \right] \quad (2.85)$$

where \Re takes the real part.

We will make use of these definitions in the following sections as we describe the dominant noise sources of the LIGO interferometers and as we try to characterize the sensitivity. Furthermore, we will encounter these definitions again in Section 4 as we explore gravitational wave data analysis. For a thorough description of these sources, refer to [89].

Dominant Noise Sources

Here we list the common noise sources in ground-based interferometers. Figure 2.7 demonstrates how these noise sources limit the strain sensitivity of LIGO. The three dominant noise sources for the initial detectors were seismic, thermal, and shot noise.

Photon Shot Noise

As described in Section 2.3.1, the primary limitation in measuring a small change in light is shot noise, or the natural fluctuations in the rate of photons arriving at the photodiode, that follow a Poisson process. The noise will decrease with increasing laser power, recycling cavity gain, arm cavity gain, and arm length.

Radiation Pressure Sensor Noise

Radiation pressure noise is associated with the photons from the laser striking the mirror and causing a force on the mirror. Of course, increasing the laser power to combat shot noise will actually result in an increase of radiation pressure. This noise source dominates the sensor noise at low frequencies while shot noise dominates at high frequencies. Thus one must optimize the laser power to give the smallest sensor noise overall.




Figure 2.7: Noise sources of initial LIGO. Figure taken from [90].

Seismic Noise

At low frequencies, below about 40 Hz for initial LIGO, ground motion is the dominant source of noise. It shakes the optics and produces strain signals that mask gravitational wave signals. It would actually dominate at higher frequencies as well if it were not for a series of active- and passive-isolation systems. The first line-of-defense is suspending the optics as pendulums. For ground-motion frequencies much larger than the pendulum frequency ($f_{pend} = 0.76$ Hz for initial LIGO), motion of the optics will be suppressed. Further suppression occurs from four alternating mass-spring layers which provide an attenuation factor $\propto f^{-2}$.

Thermal Noise

At frequencies of about 40-100 Hz, the random Brownian motion of the molecules on the surface of the mirrors and wires dominates. Initial LIGO's two most important sources of

thermal noise are the pendulum suspension system and the internal vibration modes of the mirrors [91].

Gravity Gradient Noise

Gravity gradient noise is actually due to the gravitational attraction between the test masses and density fluctuations in the atmosphere and earth in the local vicinity of the detector. It was not a dominant source of noise for initial detectors but will be for advanced detectors at low frequencies [92].

2.3.2 Progress

Several first-generation detectors have been in operation since the early 2000s. The Laser Interferometer Gravitational-wave Observatory (LIGO) currently has two sites in the United States separated by 3030km. Two interferometers share a vacuum tube at the Hanford, Washington site although the second smaller detector may be relocated to India for the advanced detector era. The large interferometer at Hanford (H1) has 4km arms, and the smaller (H2) had 2km arms. At the Livingston, Louisiana site, there is one interferometer (L1) with 4km arms. All LIGO detectors are power-recycled Michelson interferometers with Fabry-Perot cavities as arms. The worldwide network of observatories also includes the German-British GEO600 with 600 m long arms, the Italian-French Virgo (V1) with 3 km long arms, and the Japanese TAMA with 300 m arms (although TAMA was damaged in the recent Japanese earthquake and will likely be replaced by KAGRA with 3 km long arms). This thesis will focus mainly on data collected using only the LIGO detectors during its fifth science run (S5).

LIGO is looking forward to the first detection of gravitational waves. Analysis of data from (S5) is continuing with some search results already published. No gravitational waves have been found. Data was taken for two years during this science run. Figure 2.8 compares the amplitude spectral density of the displacement noise of the LIGO detectors during the S5 science run. The LIGO detectors are most sensitive to gravitational waves between ~ 100 Hz and ~ 300 Hz. The initial LIGO goal was a strain sensitivity of 10^{-21} , or equivalently, a change in length on the order of 10^{-18} m. From Figure 2.8, we see that this goal was met and surpassed for the most sensitive frequency range.




Figure 2.8: The best amplitude spectral density of the displacement noise of the LIGO detectors during the S5 science run. Taken from [88].

3. Black Holes

3.1 Black Holes in General Relativity

As early as the eighteenth century, French mathematician Laplace and Cambridge professor John Mitchell speculated on objects with such strong gravitational pull that not even light could escape [93, 94]. But it was not until after Einstein formulated his theory of General Relativity in 1915 that German astronomer Karl Schwarzschild found an exact solution for the geometry of such an object [95]. The Schwarzschild metric describes a non-rotating, uncharged black hole with the line element

$$ds^s = - \left(c^2 - \frac{2GM}{r} \right) dt^2 + \left(1 - \frac{2GM}{c^2 r} \right)^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \quad (3.1)$$

where M is the mass of the black hole. These objects are known as Schwarzschild black holes. They are without angular momentum, or spin, and are spherically symmetric. An important characteristic of these objects is their Schwarzschild radius dependent only on the mass of the black hole:

$$R_S = \frac{2GM}{c^2}. \quad (3.2)$$

This radius describes a one-way boundary called the event horizon that separates the inside of the black hole from the rest of the universe. All events that occur inside this radius are forever hidden from us. Radiation from within the event horizon is infinitely red-shifted and the object appears dark, hence the name “black hole” coined by John Wheeler in 1968 [96]. At the center, within the horizon, is a singularity, or point of infinite density, where all the mass of the black hole is concentrated. The curvature of spacetime at this point is infinite.

Surprisingly, despite the extreme nature of these objects, black holes can be completely characterized by just three parameters: mass, angular momentum, and electric charge. This concept is known as the “no-hair theorem” [97]. The Schwarzschild black hole just described has no angular momentum or charge so it is completely characterized by its mass alone! A discussion of charged black holes is outside the scope of this thesis since it is not expected that black holes with significant charge will be formed in nature; the electromagnetic force in compressing a charged mass is much greater than the gravitational attraction, requiring tremendous amounts of energy [98]. However, a black hole with angular momentum is highly relevant as it is the key source for searches of gravitational wave ringdowns as will be explored in Section 3.4.

In 1963, Roy Kerr described the spacetime geometry around an uncharged black hole with angular momentum J , or spin $a = J/cM$ [99]. The Kerr spacetime describes a rotating, uncharged black hole with the line element

$$ds^s = - \left(1 - \frac{2GMr}{c^2 \rho^2} \right) c^2 dt^2 - \frac{4GMa r \sin^2 \theta}{c \rho^2} d\phi dt + \frac{\rho^2}{\Delta} dr^2 + \rho^2 d\phi^2 + \left(r^2 + a^2 + \frac{2GMra^2 \sin^2 \theta}{c^2 \rho^2} \right) \sin^2 \theta d\phi^2 \quad (3.3)$$

where again, a is the spin of the black hole, $\rho^2 \equiv r^2 + a^2 \cos^2 \theta$, and $\delta \equiv r^2 - 2GMr/c^2 + a^2$. Further discussions about the spin of a black hole will refer to the dimensionless spin parameter $\hat{a} = Jc/GM^2$. This quantity is useful since it ranges from $\hat{a} = 0$ for a non-spinning black hole to $\hat{a} = 1$ for an extreme Kerr black hole that has reached its maximum allowed angular momentum. This limit can be understood by examining how the innermost stable circular orbit (ISCO) for a black hole changes with increasing angular momentum. This circular orbit is defined at a radius inside of which particles will fall into the black hole within about one orbital period. For a Schwarzschild black hole, this radius is

$$R_{\text{ISCO}} = \frac{6GM}{c^2}. \quad (3.4)$$

But increasing angular momentum acts to shrink the event horizon and ISCO radius. For a maximally spinning black hole, these orbits are closer to the hole:

$$R_{\text{eh}} = R_{\text{ISCO}} = \frac{GM}{c^2}. \quad (3.5)$$

For $\hat{a} \geq 1$, the black hole would be rotating too fast to have an event horizon. At this point, the singularity becomes naked!

While naked singularities are highly speculative, the existence of black holes is not. As described in the next section, extremely good evidence exists for two classes of black holes: stellar mass, with several solar masses, and supermassive, with millions of solar masses. A third class, known as intermediate mass black holes between hundreds and thousands of solar masses, has historically had weak evidence although new observations look promising.

3.2 Observations of Black Holes

The identification of a black hole, by its very nature, is a difficult job. A black hole cannot be observed directly because light cannot escape it. Thus, its effects on matter in the vicinity must be used to infer its presence. Luckily, black holes are extremely efficient at converting the energy of infalling material into emitted X-ray radiation due to the extremely high temperatures the material reaches as it approaches the ISCO of the black hole. Thus, if a black hole has a donor companion star, then it will be a source of X-rays in a high-mass X-ray binary (HMXB) or low-mass X-ray binary (LMXB). In fact, even infalling gas clouds will produce observable flares. Furthermore, since black holes are small and the processes occurring in the inner parts of the accretion disk can be quite complicated, their brightness can vary on short-timescales. Also, magnetic fields in the vicinity may cause material to be ejected from the black hole in jets, for example in the nearby galaxy M87. Thus, the process of black hole identification involves observations at several different wavelengths as well as careful analyses such as monitoring the donor star's motion in a binary.

The first strong evidence for a black hole came from X-ray and optical observations of the HMXB Cygnus X-1 in 1964 [100]. Initial observations indicated that the binary contained a blue supergiant star and a hidden X-ray source [101–105]. The supergiant was found to move with a velocity of $\sim 75 \text{ km s}^{-1}$ in a 5.6-day orbit about its unseen companion. Using one of the most important expressions in X-ray astronomy, the mass function, it was shown that the mass of the object in Cyg X-1 should be $\geq 4 M_{\odot}$. The mass function is an expression that is derivable simply from the binary's period P_{orb} and the orbital speed of the companion v_C :

$$f(M_C, M_X, i) \equiv \frac{(M_X \sin i)^3}{(M_C + M_X)^2} = \frac{v_C^3 P_{\text{orb}}}{2\pi G} \quad (3.6)$$

where M_C is the mass of the companion, M_X is the mass of the X-ray source, and i is the inclination angle between the normal to the orbital plane and the line of sight. The reason Cyg X-1 became the first solid candidate was that, in addition to knowing P_{orb} and v_C , astronomers also knew that the eccentricity was very small (≤ 0.02) and that the orbit was nearly circular. Furthermore, the companion was a relatively normal blue supergiant for which masses can be as high as $20 M_\odot$ but due to the likely mass-transfer and measured luminosity at Earth, a lower limit was found to be $\sim 10 M_\odot$ providing a lower limit to the mass of the X-ray source of $M_X \geq 4 M_\odot$. The importance of this result is that neutron stars, the second most compact astronomical object in the universe, have a mass upper limit of about $\sim 3 M_\odot$ [106]. With a mass upper limit of $\sim 4 M_\odot$, the unseen object in Cyg X-1 became a very compelling candidate for the most compact astronomical object, a black hole.

To date, there are many black hole candidates. Due to the assumptions that must be made to extract a mass from the mass function and our inability to directly observe a black hole, unambiguous evidence for their existence will only come from the detection of gravitational waves. Approximately two dozen stellar mass black hole candidates exist. These objects, in the mass range $3 \leq M/M_\odot \leq 20$, are the end points of stellar evolution when the collapsing part of the star is more massive than the mass upper limit of a neutron star ($\sim 3 M_\odot$). In addition, several dozen supermassive black hole candidates exist. These objects have masses in excess of $10^6 M_\odot$ and are found in the centers of galaxies. Little is known about the formation of these impressive black holes [107].

A keen observer will note that there is a large gap in the mass ranges of stellar and supermassive black holes. This mass range, from $10^2 \leq M/M_\odot \leq 10^4$, is predicted to be populated by a class of elusive black holes known as intermediate mass black holes [108]. Evidence for these objects has historically been difficult to come by although recent observations have breathed life into the possibility, and possible abundance, of such objects [75]. Their existence, and furthermore, their existence in binary systems, is very important for ringdown searches as intermediate mass black hole binary systems produce black hole ringdown gravitational waves that fall within LIGO's sensitivity band and thus form an important source for the ringdown search.

3.3 Intermediate Mass Black Holes

Two types of data provide compelling evidence for the existence intermediate mass black holes (IMBHs). The first are ultraluminous X-ray sources (ULXs) that have angle-averaged fluxes many times that which would be expected from a stellar mass black hole ($M < 20 M_\odot$) accreting at the Eddington limit. The second are cores of globular clusters showing excesses of dark mass. For an excellent review of intermediate mass black holes, see Reference [109].

The first possible source of IMBHs can be cited as an explanation for the apparent super-Eddington luminosities observed in some ULXs. As we have already seen, the X-ray luminosity of a black hole is generated by accretion. For a source of mass M that accretes and radiates isotropically, the luminosity cannot exceed the Eddington luminosity:

$$L_E = \frac{4\pi G M m_p c}{\sigma_T} \sim 3 \times 10^{13} \left(\frac{M}{10^{-9} M_\odot} \right) L_\odot \quad (3.7)$$

where $\sigma_T = 6.65 \times 10^{-25} \text{ cm}^2$ is the Thomson scattering cross section and m_p is the mass of a

proton. At greater luminosities than L_E , the pressure of radiation exceeds the gravitational force and the gas is blown away. Furthermore, the observed flux, if isotropy holds, can be used to place a lower limit on the mass of the black hole. If one considers a stellar mass black hole to have a maximum mass of $\sim 20M_\odot$, then from Equation 3.7 we see that this corresponds to a maximum Eddington luminosity of $\sim 3 \times 10^{39} \text{ erg s}^{-1}$. However, several dozen ULXs have $L_X > 3 \times 10^{39} \text{ erg s}^{-1}$, some even more than $L_X > 10^{40} \text{ erg s}^{-1}$. If these objects are radiating isotropically, the only explanation is that these are accreting black holes with masses larger than any known stellar mass black hole. However, there is still doubt as to whether these could be IMBHs because there is no way to discern whether the radiation is isotropic or anisotropic. If it is anisotropic, for example the radiation is beamed in a particular direction, then there is no fundamental reason why L_X cannot be greater than L_E . Thus ULXs could be explained as a different phenomena than IMBHs. For a review of ULX sources, see Reference [110]. Some recent observations of ULX sources, some of which could host IMBHs, include the Compact Blue Dwarf galaxy MRK996 [111], the Local Group starburst galaxy IC 10 with the most massive known stellar mass black hole [112, 113], the ultraluminous X-ray sources M82 X-1 [114, 115] and NGC 1313 X-2 [116], and an eclipsing X-ray binary in M101 [117, 118].

The second possible sources of IMBHs may be globular clusters in the halo of our galaxy [119–121]. These clusters are spherical collections of stars that are much older than the galactic clusters found in the disk of the galaxy. They have been found to have an excess of dark mass in their centers that could be explained by the presence of a $\sim 10^3 M_\odot$ IMBH. However, it has been shown that the excess of dark mass could just as easily be explained by a cluster of lower mass objects such as white dwarfs, neutron stars, or stellar mass black holes.

Further speculation about the existence of IMBHs comes from questions about the formation mechanisms of such objects. Black holes in the mass range from $10^2 - 10^4 M_\odot$ are more massive than the most massive stars forming in our current universe. Even the most massive stars, certainly no larger than $\sim 200 M_\odot$, will lose significant mass due to winds and pulsations leaving only a fraction for remnant black holes. Thus is unlikely that an IMBH would form directly from stellar collapse in the current universe.

However, one interesting idea postulates that Population III stars could have been progenitors. These are the first generation of stars that would have formed in the very low metallicity environments (low metallicity in astronomy implies a composition of mostly hydrogen and helium) of the early universe [122, 123]. Such an environment would allow for much higher initial masses. Above about $250 M_\odot$, these very massive stars would have collapsed directly into black holes with very little mass loss [123]. Perhaps these isolated IMBHs would have been gravitationally captured in a star cluster, sunk to the center of the cluster, and acquired a companion.

Also, it is reasonable to expect that the mergers of a collection of compact objects like black holes, neutron stars, and binaries of these object may combine to give an IMBH. However, the probability of multiple of these events is generally very small. Nevertheless, in the center of old, dense stellar clusters where most stars would have evolved off the main sequence, we would expect to find their massive remnants (black holes, neutron stars, etc.) sunk to the center. Here it may be possible for multiple interactions via binary mergers, direct capture, or another type of n-body encounter [119, 120].

Of course, in order for LIGO to have a chance at detecting IMBHs, there must be gravitational wave production. This can be accomplished via the merger of compact binary systems containing an IMBH or the direct, asymmetrical collapse of a very massive star directly into an IMBH. In the case of compact mergers, Reference [19] sets a plausible optimistic estimate of $0.007 \text{ GC}^{-1}\text{Gyr}^{-1}$ and a rate upper limit of $0.07 \text{ GC}^{-1}\text{Gyr}^{-1}$ (where

Table 3.1: A list of the first few (complex) quasi-normal mode frequencies $GM\omega_{nl}/c^3$ of oscillation for a non-spinning Schwarzschild black hole [127].

l	n	$GM\omega_{nl}/c^3$
2	1	$\pm 0.3737 + 0.0890i$
	2	$\pm 0.3467 + 0.2739i$
	3	$\pm 0.3011 + 0.4783i$
3	1	$\pm 0.5994 + 0.0927i$
	2	$\pm 0.5826 + 0.2813i$
	3	$\pm 0.5517 + 0.4791i$

GC stands for globular cluster) assuming that $\sim 10\%$ of young dense clusters have sufficient mass and a binary fraction that exceeds $\sim 10\%$. This translation to an iLIGO detection rate for IMBH-IMBH mergers of $10^{-4} - 10^{-3}$ per year and an aLIGO detection rate of 0.1-1 per year. These values are calculated from rates estimated in Reference [124].

3.4 Perturbed Black Holes

We can study the effects of a small mass on a black hole by examining the reaction of black hole spacetimes to perturbations as was first done by Regge and Wheeler in 1957 [125]. In the case of a Schwarzschild spacetime, we find that a complete description of the outgoing gravitational wave (ie that experienced by an observer at a large distance), will be contained in a single component of the Riemann tensor,

$$\Psi_4 = c^{-2} \left(\ddot{h}_+ - i\ddot{h}_\times \right) \quad (3.8)$$

where Ψ_4 is a complex Weyl scalar which provides a description of the curvature of spacetime [125]. Solving the Einstein equation provides a wave equation for Ψ_4 known as the Teukolsky equation for a Schwarzschild black hole [126]. This equation is separable into a radial function $R_l(r, \omega)$ and spin-weighted spherical harmonics ${}_sY_{lm}(\theta, \phi)$ of spin-weight $s = -2$. The radial equation can only be solved for certain complex values of the eigenfrequency ω . These are known as the quasi-normal modes of the Schwarzschild black hole, a set of complex natural frequencies of oscillation for the black hole [125]. In 1985 Leaver calculated the fundamental $l=2$ and $l=3$ modes as well as the first 62 overtones, indexed by n [127]. The first of these are shown in Table 3.1. These quasi-normal modes (QNMs) are complex frequencies which implies that they decay exponentially with time. Over time, the black hole will settle to a stationary state as the different modes of oscillation caused by the perturbation is radiated away as gravitational waves. In Table 3.1, we see that each QNM has a characteristic complex angular frequency $GM\omega_{ln}/c^3$ where (l, n) label the different modes. The real part characterizes the spectrum of gravitational wave frequencies and equals $f_{ln} = 2\pi\Re\{\omega_{ln}\}$. The imaginary part is the inverse of the damping time $\tau_{ln}^{-1} = \Im\{\omega_{ln}\} = \pi f_{ln}/Q_{ln}$ where Q_{ln} is referred to as the *quality factor*. From Table 3.1, note that $(2, 1)$ has the smallest imaginary piece, or longest damping time. This mode is also typically the most strongly excited mode of oscillation of the black hole.

Similarly, for a generic stationary rotating black hole, we can obtain a decoupled curvature

Table 3.2: A list of the first few (complex) quasi-normal mode frequencies $GM\omega_{lm}/c^3$, eigenvalues A_{lm} , and dimensionless spin parameter $c^2a/(GM)$ for the fundamental $(l, m, n) = (1, m, 2)$ mode of a rotating Kerr black hole.

l, m	$c^2a/(GM)$	A_{lm}	$GM\omega_{lm}/c^3$
2, +2	- 0.98	4.7108 + 0.1950i	+ 0.2927 + 0.0881i
	- 0.60	4.4827 + 0.1277i	+ 0.3168 + 0.0889i
	0	4	\pm 0.3737 + 0.0890i
	+ 0.60	3.1454 + 0.1567i	- 0.4940 + 0.0838i
	+ 0.98	1.3336 + 0.1500i	- 0.8254 + 0.0386i
2, +1	- 0.98	4.3892 + 0.0787i	+ 0.3439 + 0.0837i
	- 0.60	4.2558 + 0.0577i	+ 0.3490 + 0.0876i
	0	4	\pm 0.3737 + 0.0890i
	+ 0.60	3.6125 + 0.0835i	- 0.4360 + 0.0846i
	+ 0.98	3.0797 + 0.1022i	- 0.5642 + 0.0516i
2, 0	0	4	\pm 0.3737 + 0.0890i
	\pm 0.60	3.9730 \pm 0.0126i	\mp 0.3881 + 0.0860i
	\pm 0.98	3.9127 \pm 0.0315i	\mp 0.4222 + 0.0735i

perturbation Ψ_4 . We again obtain a separable Teukolsky equation for the Kerr spacetime in Equation 3.3 with angular $S_{lm}(\cos\theta, \omega)$ and radial $S_{lm}(r, \omega)$ function pieces and a separation constant $A_{lm}(a\omega/c)$. The solutions to the angular functions $S_{lm}(\cos\theta, \omega)$ are the spin-weighted spheroidal harmonics of spin-weight $s = -2$. These are again QNMs of oscillation of the Kerr black hole where the frequencies now depend on the spin parameter a and the angular mode number m [126]. Table 3.2 lists the eigenvalues A_{lm} and QNM frequencies $GM\omega_{lm}/c^3$ for the fundamental $(l, m, n) = (1, m, 2)$ mode. Again, these were calculated by Leaver in 1985 [127]. The lists are separated for various values of the dimensionless spin parameter $\hat{a} = c^2a/(GM)$. We again have the gravitational wave frequency and quality factor relationships $f_{lmn} = 2\pi\Re\{\omega_{lmn}\}$ and $\tau_{lmn}^{-1} = \Im\{\omega_{lmn}\} = \pi f_{lmn}/Q_{lmn}$. For an excellent review of black hole perturbation theory, see Reference [128].

The frequency and quality factor for each multipole can be related to the black hole mass M and dimensionless spin parameter \hat{a} . Echeverria was the first to find a fit to Leaver's calculations [129]. The fitting formulae have since been updated with better accuracy by Berti, et al. in [130]:

$$f_{lmn} = \frac{1}{2\pi} \frac{c^3}{GM} \frac{f_1 + f_2(1 - \hat{a})^{f_3}}{2\pi M}, \quad Q_{lmn} = q_1 + q_2(1 - j)^{q_3} \quad (3.9)$$

where values for the (l, m, n) dependent constants f_i and q_i can be found in Reference [130]. For example, for the $(2, 2, 0)$ and $(3, 3, 0)$ modes considered in the following sections and in

Chapter 5 we have:

$$\begin{aligned} f_{220} &= \frac{1}{2\pi} \frac{c^3}{GM} [1.5251 - 1.1568(1 - \hat{a})]^{0.1292}, \\ Q_{220} &= 0.7000 + 1.4187(1 - \hat{a})^{-0.4990} \end{aligned} \quad (3.10)$$

and

$$\begin{aligned} f_{330} &= \frac{1}{2\pi} \frac{c^3}{GM} [1.8956 - 1.3043(1 - \hat{a})]^{0.1818}, \\ Q_{330} &= 0.900 + 2.3430(1 - \hat{a})^{-0.4810}. \end{aligned} \quad (3.11)$$

In Figures 3.1 and 3.2, we can see graphical representations of these relationships for the $(2, 2, 0)$ mode. Figure 3.1 demonstrates the inverse relationship of f_{220} and final black hole mass for different dimensionless spin parameter values. High spins shift the curve to higher frequencies. Figure 3.2 demonstrates how Q_{220} varies with dimensionless spin. The minimum value that Q_{220} can take is 2.1187; otherwise, the dimensionless spin becomes negative.




Figure 3.1: The frequency of the $(2, 2, 0)$ mode as a function of final black hole mass for different dimensionless spin values.

Conceptually, these (m, \hat{a}) and (f, Q) relationships make sense. We can roughly think of the “wavelength” of a black hole as its circumference at its Schwarzschild radius from Equation 3.2, $\lambda = 2\pi (2GM/c^2)$. Then, the “frequency” of the black hole is found to be

$$f = \frac{c}{\lambda} = \frac{c^3}{4\pi GM}. \quad (3.12)$$




Figure 3.2: The quality factor of the $(2, 2, 0)$ mode as a function of dimensionless spin parameter.

Thus the inverse of the mass of a black hole is inherently tied to its frequency. Similarly, the quality factor is inversely related to the rate of energy dissipation. From Equations 3.10 and 3.11, we see that the quality factor increases with increasing spin. Thus, a highly spinning black hole will have a smaller energy dissipation rate and will take a longer time to damp the perturbation. Thus, the spin actually works to amplify the perturbation.

3.5 Ringdown Waveforms

In order to perform a search for ringdown waveforms, we need to restrict ourselves to the simplest possible waveform. Generally, it is assumed that the gravitational wave far from the source will be dominated by the least-damped $l = m = 2$ mode and that contributions from higher order modes can be ignored. As we have already discussed, at a large distance r , the gravitational waveform will be given by Ψ_4 :

$$\Psi_4 = \frac{A}{r} e^{-\pi f_{220}(t-t_0)/Q_{220}} \times \Re \left\{ e^{2i\phi} e^{2\pi i f_{220}(t-t_0)} \times [S_{22}(\cos \theta, 2\pi f_{220}) + S_{22}^*(-\cos \theta, 2\pi f_{220})] \right\} \quad (t > t_0) \quad (3.13)$$

where A is the initial amplitude, f_{220} is the frequency of the $(2, 2, 0)$ mode, Q_{220} is the quality factor of the $(2, 2, 0)$ mode.

In its simplest form which is a form that has been found to be sufficient for a search, a single $(2, 2, 0)$ QNM waveform has the form

$$h_{lmn}(t) = h_0(t) = A_{220}(\Omega) \frac{M}{r} e^{-(\pi f_{220}/Q_{220})t} \cos(2\pi f_{220}t) \quad (3.14)$$

where $A_{220}(\Omega)$ is the orientation-dependent dimensionless amplitude of the $(2, 2, 0)$ mode, r is the distance to the source, M is the black hole mass, and for simplicity we have set the phase and arrival time to zero and $G = c = 1$ unless otherwise stated. In Figure 3.3, we see two scaled ringdown waveforms with different (f_{220}, Q_{220}) values. The top plot shows a ringdown with $(2000 \text{ Hz}, 20)$ and the bottom plots shows a ringdown with $(50 \text{ Hz}, 2.1187)$.




Figure 3.3: Ringdown waveforms with $(f_{220}, Q_{220})=(2000 \text{ Hz}, 20)$ on top and $(f_{220}, Q_{220}) = (50 \text{ Hz}, 2.1187)$ on bottom.

Of course, the waveform for an arbitrary position and orientation takes the form of Equation 2.76 where the plus and cross polarizations are given by

$$\begin{aligned}
h_+(t) &= (1 + \cos^2 \iota) h_0(t) \\
h_\times(t) &= 2 \cos \iota h_0(t)
\end{aligned}
\tag{3.15}$$

where again, ι is the inclination of the source with respect to the detector frame.

The amplitude of the wave can be evaluated from the stress-energy tensor in Equation 2.26:

$$A = \sqrt{\frac{5}{2}} \epsilon \left(\frac{GMM_\odot}{3.085 \times 10^{22} c^2 r} \right) Q^{-\frac{1}{2}} \left(1 + \frac{7}{24Q^2} \right)^{-\frac{1}{2}} (1.5251 - 1.1568(1 - a)^{0.1292})^{-\frac{1}{2}}. \tag{3.16}$$

where ϵ is the amount of energy emitted as gravitational waves during a ringdown, r is the distance to the source, f and Q are the frequency and quality factor of the (2,2,0) mode, M is the mass of the final black hole, and a is the dimensionless spin parameter.

The amount of energy emitted as gravitational waves during a ringdown will depend on the strength of the perturbation. For a ringdown associated with compact binary coalescence which will be explored in more detail in Section 3.5.1, Reference [131] found that ϵ is proportional to the symmetric mass ratio η by the relation

$$\epsilon_{\text{CBC}} = 0.44\eta^2 \tag{3.17}$$

where $\eta = (m_1 m_2) / (m_1 + m_2)^2$ and m_1 and m_2 are the component masses of the two compact objects in the binary. The quantity η is proportional to the mass ratio $q = m_2 / m_1 < 1$ such that if $q = 1$, then $\eta = 0.25$. In Figure 3.4, we see how ϵ varies over an η range of 0.16-0.25 or equivalently, a q range of 0.25-1.0. For purposes of computing the ringdown amplitude in the ringdown search, we have taken $\epsilon = 0.01$. This is a reasonable value over the range of mass ratios shown in Figure 3.4 but for extreme mass ratios, it is incorrect to assume $\epsilon \sim 0.01$.

3.5.1 Ringdown Waveforms from Compact Binary Coalescence

As discussed in section 2.2.1, one type of astrophysical source that would produce gravitational waves within iLIGO's and eLIGO's frequency sensitivity band is the coalescence of two compact binary objects. These compact objects could be neutron stars, stellar to intermediate mass black holes, or a combination of the two. The process of binary coalescence is generally divided into three parts: inspiral, merger, and ringdown.

The *inspiral* phase marks the time when the objects are in orbit around each other, radiating away energy and angular momentum in the form of gravitational waves. The waveform is generally referred to as a ‘‘chirp’’ waveform as shown in Figure 3.5. The waves from this phase carry information about the masses and spins of the system. This stage is long, marking the time from initial separation to the ISCO.

The *merger* phase marks the coalescence of the two objects and generally emits the peak of gravitational wave emission. For two black holes, this would be marked by the coalescence of the horizons. The waves from this phase carry information regarding the fully general relativistic dynamics of the system.

The *ringdown* phase marks the period after coalescence when the two objects have formed a single perturbed black hole that rings down in a superposition of quasi-normal modes (QNMs). The waves provide information on the mass and spin of the final perturbed black hole.




Figure 3.4: The amount of energy emitted as gravitational waves during the ringdown associated with compact binary coalescence as a function of the symmetric mass ratio η . The data and fit are taken from [131].




Figure 3.5: An example chirp waveform.

As we will discuss in more detail in Chapter 4, the ringdown search can measure its sensitivity to various sources by determining how well it can detect various simulated waveforms that are added to detector data. We expect that the ringdown template will match well with the ringdown part of the IMR waveform. In fact, from the relationship between the amount of energy radiated as ringdown gravitational waves, ϵ , in Equation 3.17 for IMR signals, we can predict a theoretical ringdown horizon distance for CBC signals with different parameters and different detector sensitivities. Figure 3.6 shows the ringdown horizon distance as a function of final black hole mass assuming $\epsilon = 0.01$, a dimensionless spin of $a = 0.9$ and an SNR of 8 in the H1, H2, and L1 interferometers during the S5 science run and in the Virgo detector V1 during its first science run. The features of the horizon distance curves can be mapped to features in the LIGO and Virgo strain sensitivity curves. In Figure 3.7, we can

see how the horizon distance varies with dimensionless spin. For a near-maximally spinning final black hole, H1 could detect a 600-700 M_{\odot} IMBH binary out to 300 Mpc with an SNR of 8. The horizon distance at high masses decreases as dimensionless spin decreases.




Figure 3.6: Ringdown horizon distance as a function of final black hole mass assuming $\epsilon = 0.01$, a dimensionless spin of $a = 0.9$ and an SNR of 8 in the H1, H2, and L1 interferometers during the S5 science run and in the Virgo detector V1 during its first science run.

Of course, the real test of the search’s sensitivity will come from the recovery of signals injected into detector data. Previous searches for binary black holes used templates and simulated waveforms generated with post-Newtonian analytical techniques using the Stationary Phase Approximation (SPA) [25, 132–136]. This is essentially a phase evolution approximation from a given low frequency up to the frequency of the inner most stable circular (ISCO) orbit. This has been shown to be a good approximation for initial LIGO and binary systems with low total masses, for example NS-NS or NS-BH systems because only the inspiral part of the signal will be visible. The merger and ringdown, which are not included in the SPA model, would be buried in the detector shot noise at high frequencies. However, the inspiral, merger, and ringdown shift to lower frequencies for higher mass binaries thus making the inclusion of merger and ringdown important. Eventually, the merger and ringdown will be the only part of the signal with frequencies above the seismic wall and the only contributors to the total SNR. Thus, for a matched filter pipeline, it is important to include these parts of the signal for detection of high mass binaries. Two such inspiral-merger-ringdown (IMR) waveforms are shown in Figures 3.8 and 3.9. Figure 3.8 shows an IMR waveform for a non-spinning compact binary with component masses of 10 M_{\odot} and Figure 3.9 shows an IMR waveform for a non-spinning compact binary with component masses of 150 M_{\odot} . Note that for the lower mass binary, there is significant contribution to the signal from the inspiral of




Figure 3.7: Ringdown horizon distance as a function of final black hole mass assuming $\epsilon = 0.01$ and an SNR of 8 in the H1 interferometer during the S5 science run. The dimensionless spin parameter is allowed to vary.

the two objects. However, for the higher mass system, there are many fewer cycles in the inspiral and the ringdown phase contributes a larger fraction to the overall signal.

Recently, due to the successes of numerical relativity in evolving the binary black hole system through its merger and ringdown [137], full parameterized IMR waveforms have become available for gravitational wave data analysis [135, 138–142]. While it is still computationally expensive to generate these NR waveforms and thus cover the whole parameter space, methods have been developed to help fill in the parameter gaps not covered yet by NR. Typically, the inspiral portion generated by post-Newtonian techniques is connected smoothly to merger and ringdown in order to cheaply incorporate and generate waveforms from NR over a larger parameter space.

One such family of NR-inspired waveforms are those formulated via the Effective One Body (EOB) approach [143–145]. This method evolves the EOB Hamiltonian for a binary black hole system up to the “plunge” of the two black holes right before merger. The plunge effects are calibrated based on NR results [135, 140, 146]. The gravitational waveform is extracted in the far-field limit. Then the “plunge-merger” portion of the waveform is connected to the merger-ringdown portion. These waveforms, as discussed in more detail in Chapter 4, are used in ringdown pipeline sensitivity measurements. The initial waveform family was known as EOBNR [144, 147–149]; these were used in the tuning of the pipeline as outlined in Section 4.3. However, recently higher order terms were introduced and calibrated against numerical relativity in the EOBNRv2 family. We have used these new waveforms in the final S5 analysis runs as described in Section 4.4.3.

Another family of NR-inspired waveforms are the parameterized, phenomenological bank of binary black hole waveforms known as IMRPhenom [136, 150]. These are constructed in the Fourier domain by the smooth connection of inspiral, merger, and ringdown evolution stages. The initial waveform family was known as IMRPhenomA. However, these waveforms have also been adapted to include spin effects for initially spinning black holes aligned with the angular momentum of the binary [151]. These waveforms are known as PhenomB. We




Figure 3.8: IMR waveform for a non-spinning compact binary with component masses of $10 M_{\odot}$.

have used these both in tuning the search pipeline and in the final S5 analysis runs as described in Section 4.4.3.

The method we use to calculate the mass of the final black hole after the two objects have merged assumes that we have an initially non-spinning binary as calculated in (Buonanno, et al). The final mass from Table I in this paper is found to be:

$$M_{\text{final}} = 1 + \left(\sqrt{8/9} - 1 \right) \eta - 0.498\eta^2. \quad (3.18)$$

Figure 3.10 shows a graphical representation of this relationship for different values of the symmetric mass ratio η .

Numerical relativity waveforms have shown us that equal mass non-spinning binary black hole systems will often evolve to a final black hole with a final spin close to 0.7 [152, 153]. Furthermore, studies of unequal mass non-spinning black holes have shown that the amplitude of the $l = m = 2$ mode of the ringdown waveform decreases as the mass ratio of the binary components becomes more extreme while the amplitude of the $l = m = 3$ mode increases [154]. Also, we know that if the binary black holes are initially spinning, then the spin of the final black hole and the amount of radiation depend on the magnitude and inclination of the spins with respect to the orbital angular momentum [132–134, 155]. We will see how these characteristics affect the sensitivity of the ringdown pipeline to particular types of sources in Section 4.5.

We can calculate the norm and direction of the final spin of a black hole if we know




Figure 3.9: IMR waveform for a non-spinning compact binary with component masses of $150 M_{\odot}$.

the initial spins and masses of the two black holes as well as components of the orbital angular momentum \mathbf{L} . However, for the purposes of simulated waveform recovery in the ringdown pipeline, we use PhenomB waveforms for which we know that the orbital angular momentum \mathbf{L} is along the positive \hat{z} direction and the dimensionless spins of the black holes before merger a_1 and a_2 will be aligned or anti-aligned with \mathbf{L} and thus with each other too. Using results from numerical relativity, Reference [156] shows that the final spin norm is

$$|a_{\text{final}}| = \frac{1}{(1+q)^2} [|a_1|^2 + |a_2|^2 q^4 + 2 |a_2| |a_1| q^2 \cos \alpha + 2 (|a_1| \cos \beta(r) + |a_2| q^2 \cos \gamma(r)) |\ell| q + |\ell|^2 q^2]^{1/2} \quad (3.19)$$

with mass ratio $q \equiv m_2/m_1$ and $\ell \equiv \tilde{\ell}/(m_1 m_2)$ where $\tilde{\ell}$ is the residual orbital angular momentum contributing to the final spin. The angles α , β and γ are defined as follows:

$$\cos \alpha \equiv \hat{a}_1 \cdot \hat{a}_2, \quad \cos \beta \equiv \hat{a}_1 \cdot \hat{\ell}, \quad \cos \gamma \equiv \hat{a}_2 \cdot \hat{\ell}. \quad (3.20)$$

Thus, the values of these expressions for PhenomB waveforms will be either 1 or -1 depending on whether we have aligned or anti-aligned systems. Using fits to the results of NR calculations, specifically aligned binary data, Reference [156] provides an expression for the norm of ℓ used in equation 3.19:




Figure 3.10: Final black hole mass as a function of initial mass for different values of the symmetric mass ratio ν .

$$|\ell| = 2\sqrt{3} + t_2\eta + t_3\eta^2 + \frac{s_4}{(1+q^2)^2} (|a_1|^2 + |a_2|^2q^4 + 2|a_1||a_2|q^2 \cos \alpha) + \left(\frac{s_5\eta + t_0 + 2}{1+q^2} \right) (|a_1| \cos \tilde{\beta}(r_in) + |a_2|q^2 \cos \tilde{\gamma}(r_in)) \quad (3.21)$$

where $\eta \equiv m_1 m_2 / (m_1 + m_2)^2$ and the constants have been calculated to be the following: $s_4 = -0.1229 \pm 0.0075$, $s_5 = 0.4537 \pm 0.1463$, $t_0 = -2.8904 \pm 0.0359$, $t_2 = -3.5171 \pm 0.1208$, $t_3 = 2.5763 \pm 0.4833$. In Figure 3.11, we can see the final black hole spin as a function of a_2 for different a_1 for a (100, 100) M_\odot binary black hole system and in Figure 3.12 for a (100, 25) M_\odot binary black hole system.




Figure 3.11: Final black hole spin as a function of a_2 for different a_1 for a $(100, 100) M_\odot$ binary black hole system.




Figure 3.12: Final black hole spin as a function of a_2 for different a_1 for a $(100, 25) M_\odot$ binary black hole system.

4. Searches for Gravitational Wave Ringdowns

4.1 Gravitational Wave Data Analysis

In Chapter 2, we saw how a gravitational wave can produce a fractional strain in the lengths of the x and y arms of an interferometric detector. Furthermore, in Section 2.3.1, we saw that the response of the detector to a gravitational wave will be influenced by the control loop that corrects for external noise. In this chapter, we will explore how one can extract a gravitational wave signal $h(f)$ from the external strain $s_{\text{ext}}(f)$ that will also contain noise $n(f)$.

Gravitational wave data analysis is essentially an attempt to extract a very weak signal buried in detector noise. When the signal is known, the optimal method for extracting the signal from Gaussian noise is matched filtering. While detector noise is not Gaussian, the matched filtering technique can still be a powerful technique as outlined below. Further methods for distinguishing true gravitational wave signals from non-Gaussian detector noise are also employed and will be outlined below as well.

The Fourier transform convention for a signal $x(t)$ that we will use is

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t)e^{-2\pi ift} dt \quad (4.1)$$

and for its inverse

$$x(t) = \int_{-\infty}^{\infty} \tilde{x}(f)e^{2\pi ift} df. \quad (4.2)$$

4.1.1 The Optimal Detection Statistic

From Section 2.3.1, we know that the detector output may or may not contain a weak signal. Thus, the strain data $s(t)$ will be composed of random noise $n(t)$ and possibly a gravitational wave of known form $h(t)$:

$$s(t) = \begin{cases} n(t), & \text{null hypothesis : no signal} \\ n(t) + h(t), & \text{alternative hypothesis : signal present, } t > 0. \end{cases}$$

Let us assume that the detector noise is a stationary, random process with zero mean, a Gaussian probability distribution, and a one-sided power spectrum $S_n(f)$ given by Equation 2.84 in Section 2.3.1. Furthermore, we know the form of the signal $h(t)$ we are looking for (i.e. ringdown). In general, if we know the waveform, then the best way to find that signal in the data $s(t)$ is to pass the data through a filter $K(t)$ that will produce a large output when a signal is present and a small output when there is no signal. For our purposes, the output Z of the filter will be

$$\begin{aligned}
Z &\equiv \int_{-\infty}^{\infty} K(t)s(t)dt = \int_{-\infty}^{\infty} K(t)n(t)dt + \int_{-\infty}^{\infty} K(t)h(t)dt \\
&= N + H
\end{aligned} \tag{4.3}$$

where N is the filtered noise and H is the filtered signal. Since $n(t)$ is a random process, we can use the results from Section 2.3.1, particularly Equations 2.81, to see that

$$\langle N^2 \rangle = \int_0^{\infty} |\tilde{K}(f)|^2 S_n(f) df. \tag{4.4}$$

Similarly, we can use the the results from Section 2.3.1 to write the filtered signal H as

$$H = 2 \int_0^{\infty} \tilde{K}^*(f)\tilde{h}(f)df. \tag{4.5}$$

Equation 4.4 and 4.5 can be used to construct a statistic that we can maximize by the appropriate choice of a filter K so that a large value indicates signal. Thus, we seek to find the filter K that maximizes this statistic ξ which we will call the signal-to-noise ratio (SNR):

$$\begin{aligned}
\xi &\equiv \frac{H^2}{\langle N^2 \rangle} \\
&= \frac{\left| 2 \int_0^{\infty} \tilde{K}^*(f)\tilde{h}(f)df \right|^2}{\int_0^{\infty} |\tilde{K}(f)|^2 S_n(f) df} \\
&= \frac{4 \left| \int_0^{\infty} \left[\tilde{K}^*(f)\sqrt{S_n(f)} \right] \left[\tilde{h}(f)/\sqrt{S_n(f)} \right] df \right|^2}{\int_0^{\infty} |\tilde{K}(f)|^2 S_n(f) df}.
\end{aligned} \tag{4.6}$$

At this point, we should recall our specialized definition of the inner product between quantities a and b in Equation 2.85. The last step in the above equation is important because it lets us take advantage of the Cauchy-Schwarz inequality for the expression in the numerator. Recall that this inequality says that for two arbitrary functions $C(f)$ and $D(f)$,

$$\left| \int C(f)D(f)df \right|^2 \leq \int |C(f)|^2 df \int |D(f)|^2 df. \tag{4.7}$$

Thus, in order to maximize Equation 4.6, we need to force the denominator to be as large as it can be (i.e. the equality in Equation 4.7 has to hold). The only solution to such an integral equation is when C and D differ only by a multiplicative factor A . Thus we can find a relationship between the filtered noise and the filtered signal:

$$\begin{aligned}
\tilde{K}(f)\sqrt{S_n(f)} &= A \frac{\tilde{h}(f)}{\sqrt{S_n(f)}} \\
\tilde{K}(f) &= A \frac{\tilde{h}(f)}{S_n(f)}
\end{aligned} \tag{4.8}$$

to obtain the optimal filter \tilde{K} for distinguishing a known signal $\tilde{h}(f)$ from Gaussian noise with a power spectrum of $S_n(f)$.

Now we can apply this result to an actual gravitational wave search, one where we know the form of the signal such as the ringdown search (damped sinusoid waveform) or the inspiral search (chirp waveform). In the language of matched filtering searches, we call the known waveform a *template*. Looking again at the form of the filter output in Equation 4.3 and the optimal filter expression in Equation 4.8, we find a time domain gravitational wave filter of the form

$$Z(t) = A \int_{-\infty}^{\infty} \frac{\tilde{s}(f)\tilde{h}^*(f)e^{2\pi ift}}{S_n(f)} df. \quad (4.9)$$

This is the quantity on which we could set a threshold value such that a $|Z|$ above the threshold would indicate a signal and below would indicate no signal. However, we choose first to normalize the filter output by the variance of the optimal filter σ^2 for which $A = 2$:

$$\sigma^2 = 2 \int_0^{\infty} \frac{\tilde{h}(f)\tilde{h}^*(f)}{S_n(f)} df. \quad (4.10)$$

Then, our new thresholding statistic, which we will call ρ is the SNR of the normalized output of the optimal filter:

$$\rho(t) = \frac{|Z(t)|}{\sigma}. \quad (4.11)$$

We can set any threshold on this value that we want to distinguish signal from noise. However, in reality, the threshold we set has to be computationally feasible. Too low of a threshold may admit more signals but will also admit more false alarm signals from the background noise. We have to tune this threshold so that we can admit as many signals as possible while still keeping the false alarm rate low. However, even then, we must employ further techniques to distinguish whether we have a signal or not; until we know for sure, we say that we have a candidate signal. This will be discussed further in Section 4.2.6. Once we have chosen a value ρ^* on which to threshold, we say if

$$\rho \begin{cases} < \rho^* & \text{candidate signal is absent} \\ \geq \rho^* & \text{candidate signal is present.} \end{cases}$$

with some associated confidence in this cut. If no signal is present but $\rho \geq \rho^*$, then we have a false alarm. Also, we can have a case where a signal is present but $\rho < \rho^*$. This is a false dismissal.

Finally, following the definition in Equation 2.85, we can express the variance of the optimal filter as $\sigma^2 = \langle h|h \rangle$ and hence the SNR as

$$\rho(t) = \frac{1}{\sigma} \langle s|h(t) \rangle = \frac{\langle s|h(t) \rangle}{\sqrt{\langle h|h \rangle}}. \quad (4.12)$$

4.1.2 Template Banks

As described in Section 4.1.1, the optimal statistic when the signal is known is the Fourier transform of the signal $\tilde{h}(f)$ weighted by the inverse of the power spectrum and normalized by the variance. However, in gravitational wave searches, we do not often know the *exact*

form of the signal. Often the signal will be characterized by several intrinsic parameters λ_i that can affect the signal in a non-trivial way. In the case of the ringdown search, these intrinsic parameters are the central frequency and quality of the $l = m = 2, n = 0$ mode of oscillation of the perturbed black. The ringdown template itself has the form

$$h(t) = \cos(2\pi ft) e^{-\frac{\pi f}{Q}t}. \quad (4.13)$$

Thus, we actually create an array of filters, or a template bank, to search over the range of intrinsic parameters that we are interested in for our search. We then filter the data with each of the templates in the bank. We then obtain an estimate of the intrinsic parameters for the source of the signal from the template that provided the largest SNR. However, one can imagine that it would be impossible to have a template for every possible set of intrinsic parameters. The templates actually have to be placed discretely in the parameter space in such a way that any signal that has parameters in the space will still produce an SNR above threshold by matching with the ‘‘closest’’ template.

In order to understand how to place templates in a parameter space, we can define two templates with different parameters as $\tilde{u}(f; \mu, \lambda)$ and $\tilde{u}(f; \mu + \Delta\mu, \lambda + \Delta\lambda)$ where μ are extrinsic parameters such as time of arrival and phase. We can define the match M between these two templates as their inner product maximized over the extrinsic parameters:

$$M(\lambda, \Delta\lambda) \equiv \max_{\mu, \Delta\mu} (u(\mu, \lambda) | u(\mu + \Delta\mu, \lambda + \Delta\lambda)). \quad (4.14)$$

This is the fractional loss in SNR that we can expect when filtering a signal with a template of the same form but slightly different parameters. We can expand this expression in a power series about $\Delta\lambda = 0$ to get

$$M(\lambda, \Delta\lambda) \approx 1 + \frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right)_{\Delta\lambda^k=0} \Delta\lambda^i \Delta\lambda^j. \quad (4.15)$$

The metric in the template bank is defined by the mismatch $MM = (1 - M)$ between two ringdown templates. Thus, we find a general expression for the mismatch which is just the square of the proper distance between two templates:

$$ds_{ij}^2 = g_{ij} \Delta\lambda^i \Delta\lambda^j \quad (4.16)$$

and for the metric in the template bank:

$$g_{ij}(\lambda) = -\frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right)_{\Delta\lambda^k=0}. \quad (4.17)$$

For two ringdown templates differing in frequency by df and in quality factor by dQ , the mismatch is

$$ds^2 = \frac{1}{8} \left[\frac{3 + 16Q^4}{Q^2 (1 + 4Q^2)^2} dQ^2 - 2 \frac{3 + 4Q^2}{fQ (1 + 4Q^2)} dQ df + \frac{3 + 8Q^2}{f^2} df^2 \right]. \quad (4.18)$$

The mismatch is traditionally defined in terms of $\phi = \log(f)$ so that the metric coefficients are independent of f :

$$ds^2 = \frac{1}{8} \left[\frac{3 + 16Q^4}{Q^2 (1 + 4Q^2)^2} dQ^2 - 2 \frac{3 + 4Q^2}{Q (1 + 4Q^2)} dQ d\phi + (3 + 8Q^2) d\phi^2 \right] \quad (4.19)$$

$$= g_{QQ} dQ^2 + g_{Q\phi} dQ d\phi + g_{\phi\phi} d\phi^2. \quad (4.20)$$

For any given metric, we can generally define the proper lengths of segments. For example, for a diagonal metric of the form

$$ds^2 = g_{00} (dx^0)^2 + g_{11} (dx^1)^2 + g_{22} (dx^2)^2 + g_{33} (dx^3)^2 \quad (4.21)$$

the proper lengths of two segments will be $dl^1 = \sqrt{g_{11}}dx^1$ and $dl^2 = \sqrt{g_{22}}dx^2$. The ringdown template bank metric is not diagonal. However, currently, the population algorithm uses only the diagonal terms containing df^2 and dQ^2 . Starting with the smallest central frequency and quality factor, the algorithm places templates along Q in steps of

$$dQ = \frac{ds_{\text{eff}}Q(1 + 4Q^2)}{\sqrt{3 + 16Q^4}} \quad (4.22)$$

and along $d\phi = d\log(f)$ in steps of

$$d\phi = \frac{ds_{\text{eff}}}{\sqrt{3 + 8Q^2}}, \quad (4.23)$$

where $ds_{\text{eff}} = 4\sqrt{(1 - \text{MM})}$ and MM is the specified mismatch between template and signal.

The templates are placed in $\log(f)$ and Q space such that the mismatch between any signal and its closest template is 3% or less. Figure 4.1 shows an example of such a bank where frequency ranges from 50 Hz to 2000 Hz and quality factor ranges from 2.1187 to 20. The frequency spacing of the templates is finer for lower frequencies and for higher quality factors. One can begin to understand this placement by examining the contours of constant mismatch around each template. These contours form overlapping ellipses that cover the entire template parameter space so that any signal in this parameter space will match a template by no less than 3%. The ellipses are shown in Figure 4.2 for a single frequency to demonstrate the behavior of the ellipses and for a snippet of the template bank to demonstrate the overlap of the ellipses. High quality factor templates have narrow minor axes; hence these templates are placed finely in frequency.

4.2 The Search Pipeline

As discussed in Section 4.1.1, the optimal statistic when searching for a known signal in Gaussian noise is the method of matched filtering. However, detector noise is both non-stationary and non-Gaussian. Thus, matched filtering is useful but not sufficient for extracting a true gravitational wave signal from LIGO data. A number of other tests and checks must be implemented in order to distinguish the noise in the detector that can sometimes mimic the signal from a true signal. We say that these extra tests help to improve our detection confidence. The automated set of steps form our search pipeline.

The ringdown search pipeline was developed during the S4 LIGO data run and used to search S4 data [26, 157]. We have improved and expanded the S4 search pipeline infrastructure for use in the S5 search. The main changes are outlined in Section 4.2.1. The steps of the pipeline will be explained in detail in the following sections. However, here we briefly outline the steps of the pipeline:

- Analysis times and data quality veto files are obtained to segment the data.
- Detector data is read in from gravitational-wave frame files and conditioned.




Figure 4.1: The S5 ringdown template bank. The settings used to construct this bank were the following: maximum allowed mismatch $ds^2 = MM_{\max} = 0.03$, minimum frequency $f_{\min} = 50$ Hz, maximum frequency $f_{\max} = 2000$ Hz, minimum quality factor $Q_{\min} = 2.1187$, and maximum quality factor $Q_{\max} = 20$.

- Template banks are generated and used to filter the data. GPS times of local maxima that exceed the SNR threshold are recorded as triggers.
- Triggers are clustered, data quality vetoes are applied, and then coincidence of triggers from each detector is checked.
- Triggers that fail coincidence are discarded. Triggers that pass coincidence are subjected to further tests in post-processing.

4.2.1 Summary of Main Changes Since S4

Here we outline some of the main changes to the ringdown search pipeline and post-processing since the S4 ringdown search. The changes focus on two aspects of the search: the further automation and enhancement of the search pipeline and the development of a ringdown post-processing code. Details will be explored in the following sections.

Automation

- Automated the ringdown search pipeline through the coincidence stage using the In-spiral Hierarchical PipelinE (HIPE)




Figure 4.2: Contours of constant mismatch $ds^2 = MM_{\max} = 0.03$ around each template. The plot on the left demonstrates how these contours overlap for templates at a single frequency. The plot on the right demonstrates the overlap for a snippet of the template bank between $f = 50$ Hz and $f = 57.5$ Hz.

- Automated the processing of coincident triggers; automation is important to allow for efficient tuning of the pipelines

Data Handling and Conditioning

- Changed format of strain data on which search pipeline was run
- Added ability to read Virgo data
- Added ability to add both ringdown and full inspiral-merger-ringdown (IMR) simulations into Virgo data

Simulation Code

- Updated mass, spin, frequency, and quality factor relationships based on expressions in [130]
- Added option to record IMR simulation parameters to help with parameter recovery studies

Coincidence Code

- Implemented a new 3D coincidence test to check simultaneously for coincidence in frequency, quality, and time
- Added Virgo handling to the coincidence test in `lalapps_rinca`

Post-processing of Coincident Triggers

- Developed new post-processing code

- New post-processing code now calculates the detection statistic, clusters coincidences and applies consistency cuts, ranks coincidences, does injection finding, and produces plots
- Implemented a Multi-Variate Statistical Classifier (MVSC) as a detection statistic
- Implemented a coherent stage to the search pipeline (although this stage was not incorporated in the work presented here)

4.2.2 Automation with `ihope`

The `ihope` pipeline is the automated search pipeline originally designed to detect gravitational waves from compact binary coalescences. It is a more sophisticated version of the Inspirational Hierarchical PipelinE (HIPE) which was the initial CBC pipeline set up to run a series of dependent executables on large computer clusters using the Condor high-throughput computing system. The infrastructure of the `ihope` pipeline was easily adapted for ringdown matched filter searches as the hierarchical nature of the search is essentially the same as the CBC search. A diagram of the ringdown `ihope` search pipeline and the executables used in each step can be seen in Figure 4.3.

The `ihope` pipeline is actually run via a directed acyclic graph (DAG) that is generated by the code `lalapps_ihope` to be submitted to a Condor pool. The DAG is essentially a set of instructions to Condor so that the steps of the pipeline are run in the correct order. For instance, one could not run a coincidence step before all matched filtering steps were done. The DAG breaks the steps into individual jobs that are scheduled to run by the Directed Acyclic Graph Manager (DAG Man). The jobs are placed in the Condor pool queue and executed on various compute nodes in the pool as those nodes become available. Condor provides various methods of monitoring jobs and tracking progress. The input to `ihope` is an extensive configuration file that outlines the options used in each step of the pipeline. An example configuration file for the S5 ringdown search can be found in Appendix ???. Also required are a GPS start time and GPS end time over which to search.

The code to run `ihope` and any of the steps of the ringdown search can be found in the LIGO Algorithm Library (LALSuite), a set of gravitational wave analysis routines documented here. The code is maintained via the `git` version control system and all development code for this search can be found on the `ringdown` branch of LALSuite. The commands used to run the `ihope` pipeline are also outlined in Appendix ???. The steps of the pipeline are described in more detail below.

4.2.2.1 Preliminaries

As described above, one sets up a DAG to run the search pipeline using the code `lalapps_ihope`, a configuration file such as the one in Appendix ???, and a set of GPS times to search. The code also sets up the following products to be used by the filtering pipeline. This stage was previously done by hand in S4 but has now been automated for the ringdown search within `ihope`.

Segment Lists

First, `ihope` identifies the GPS times which are available to be analyzed. These times are generally when the interferometer was “locked” as described in Section 2.3.1 and declared in *Science Mode* by operators in the control room. The interferometers do often “lose lock” and come out of *Science Mode* so the list of GPS times to analyze will be composed of many truncated GPS segments.

Category 1 vetoes are applied to these segment lists as well. These vetoes are lists of GPS times when the data was declared un-analyzable due to data quality or other reasons. This is the first time vetoes are applied in the search pipeline. Further category vetoes are applied at the coincidence stage as described in Section 4.2.5. However, the veto lists themselves are generated next.



Figure 4.3: The ringdown search pipeline.

Veto Lists

Veto files are lists of GPS times structured very similarly to the segment lists. However, veto lists mark times that are removed from the analysis due to poor data quality or other reasons that make gravitational wave detection even more difficult than it already is. Category 1 vetoes are removed from the segment list right away; these are times of very poor data quality and should not be analyzed. However, Category 2, 3, and 4 vetoes are applied directly before coincidence as described in Section 4.2.5. Category 2 vetoes mark times when there is a known correlation between a noise source and the gravitational wave channel. Category 3 vetoes mark times when there is a statistical correlation between a noise source and the gravitational wave channel. Category 4 vetoes mark times when there is a

weak correlation with false alarms [158]. Coincidences will be written to file directly after Category 2 vetoes have been applied, and then written to separate files after Category 3 vetoes have been applied, and so on. Thus we can look at the results of the search after each application of vetoes to ensure that we did not veto a time with a potential gravitational wave signal.

4.2.2.2 The Filtering Stage

After running `ihope`, one can then submit the DAG to the condor pool via the `condor _submit_dag` command to run the steps of the pipeline as outlined in Figure 4.3. The first stage of the pipeline utilizes the filtering code `lalapps_ring`. This is actually a sophisticated piece of code that has many steps including data conditioning, template bank generation, matched filtering operations, and more. These steps are outlined in more detail below although, for full details, one should refer to [157] and [26]. We also highlight where changes have been made since S4.

Read and Condition Data

The first step is the further segmentation of the segment lists into basic analysis segments. These are 2176 s long segments that overlap each other by 128 s on each end with corrections for non-integer numbers of 2176 s in the science segments. Then the appropriate calibrated LIGO $s(t)$ data is read in from frame files using a call to `ligo_data_find`. In S4, uncalibrated `DARM_ERR` data was used so there was a need to calculate the response function $R(f)$ as described in Section 2.3.1. However, this is no longer the case for S5. The $s(t)$ data is sampled at 16384 Hz but is resampled to 8192 Hz to reduce computational costs. This sets the Nyquist frequency at $f_{\text{Nyquist}} = 1/2 \times (\text{sampling rate}) = 4096$ Hz. Then, to prevent aliasing, a low-pass, time-domain digital filter is applied to remove any power above f_{Nyquist} . Then a high-pass filter is used to remove any power below 35 Hz where the detectors have little sensitivity due to seismic noise. In S4, this was set at 40 Hz. Padding of 8 s at the start and end of the 2176 s long basic analysis segments is used to avoid corruption of the data from the filters. The transformation of the data through these various conditioning steps is outlined in Figures 4.4, 4.5, 4.6, and 4.7.

Calculate the Power Spectrum

In order to calculate the gravitational wave filter from Equation 4.9, we need to calculate the one-sided power spectral density $S_n(f)$. This calculation is performed for each 2176 s segment using the median-mean method described in detail in [157] in the Fourier domain. The calibrated power spectral density has units of $s \cdot \text{count}^2$. However, looking at Equation 4.9, we actually want the inverse calibrated power spectrum S_n^{-1} with units of $s^{-1} \cdot \text{count}^{-2}$. Figure 4.8 shows an example calibrated power spectral density and inverse calibrated power spectral density for a particular analysis segment during S5.

Fourier Transform the Data

In S4, the uncalibrated data was read in as 2176 s segments, divided into 256 s overlapping blocks, Fourier transformed, and calibrated. In S5, we use calibrated $s(t)$ so we only perform the Fourier transform to obtain \tilde{s}_k . In preparation for the calculation in Equation 4.9, we then multiply by the inverse calibrated spectrum. As outlined in [157], this is a discrete quantity calculation that results in a frequency series of the form:

$$S_n^{-1} \Delta t \tilde{s}_k \tag{4.24}$$

with units of strain^{-1} . The quantity Δt is the time step at which the discrete time series is




Figure 4.4: A snippet of calibrated LIGO $h(t)$ data sampled at 16384 Hz that acts as the raw input to the search pipeline. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment.




Figure 4.5: A snippet of the data shown in Figure 4.4 after resampling from 16384 Hz to 4096 Hz. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment. Any frequency content above about 2 kHz is aliased so a low-pass filter is applied at this frequency.

sampled. It is the inverse of the resampling rate: $\Delta t = 1/8192$ s.

Generate the Template Bank

As detailed in Section 4.1.2, the template bank metric dictates that templates be laid in Q and $\phi = \log(f)$ space using steps given by Equations 4.22 and 4.23. One must specify the allowed mismatch MM between template and signal and the frequency and quality factor ranges. The template bank algorithm begins at the user-specified f_{\min} and Q_{\min} , then steps along ϕ in steps of $d\phi$ until f_{\max} is reached. Then Q is incremented by dQ . The algorithm then places templates along this new Q in steps of $d\phi$ until f_{\max} is again reached. This process repeats until the user-specified (f_{\max}, Q_{\max}) is reached.

In S4, these values were the following: $MM = 3\%$, $f_{\min} = 50$, $Q_{\min} = 2$, $f_{\max} = 2000$, $Q_{\min} =$




Figure 4.6: A snippet of the data shown in Figure 4.5 after applying a high-pass filter to suppress data below the seismic wall at approximately 40 Hz. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment. Any frequency content above about 2 kHz is aliased so a low-pass filter is applied at this frequency.




Figure 4.7: A snippet of the data shown in Figure 4.6 after discarding 8 s of padding at the start and end of the time series before analysis. This is needed since the previously applied filters corrupt the data at the start and end. Thus time $t = 0.0$ s on these plots correspond to time $t = 8.0$ s in Figures 4.4, 4.5, and 4.6. Figure (a) shows a 2176 s long analysis segment. Figure (b) shows a zoom of the first 0.3 s of the segment, and Figure (c) shows a zoom of 7.0-7.4 s of the segment.

20. However, in S5, due to the updated mass, spin, frequency, and quality factor relationships from [130], we could no longer set $Q_{\min} = 2$ as this value results in a negative spin value. The smallest value that Q can now take is 2.1187. Thus, for S5, the values were the following: $MM = 3\%$, $f_{\min} = 50$, $Q_{\min} = 2.1187$, $f_{\max} = 2000$, $Q_{\min} = 20$.

Create the Template

As outlined in [157], the discrete time domain ringdown template is given by

$$h_j = e^{-\frac{\pi f_0 j \Delta t}{Q}} \cos(2\pi f_0 j \Delta t) \quad (4.25)$$




Figure 4.8: The calibrated power spectral density and inverse calibrated power spectral density.

where again, $\Delta t = 1/8192$ s. The length of each template was set to be 10 e-folding times where an e-folding time is the interval τ over which a quantity increase by a factor of e. From the form of the template, we see that this is $\tau = Q/\pi f_0$. The discrete Fourier transform of the template time series is the frequency series:

$$\tilde{h}(f_k) = \Delta t \tilde{h}_k \quad (4.26)$$

in units of strain/Hz.

Filter the Data

The discrete form of the matched filter operations in Equations 4.9-4.11 are evaluated for each of the sixteen 256 s overlapping blocks of data in the 2176 s long analysis segment. As outlined in [157], the filter output is

$$z(t_j) = \frac{2}{\Delta t N} \Re \left\{ \sum_{k=0}^{\frac{N}{2}-1} S_n^{-1}(\Delta t \tilde{s}_k^*) (\Delta t \tilde{h}_k) e^{i2\pi jk/N} \right\} \quad (4.27)$$

where N is the number of sampled data points. The variance of the template is

$$\sigma_T^2 = \frac{4}{N \Delta t} \sum_{k=0}^{\frac{N}{2}-1} S_n^{-1} \left[\Re \left(\Delta t \tilde{h}_k \right)^2 + \Im \left(\Delta t \tilde{h}_k \right)^2 \right]. \quad (4.28)$$

Thus, the SNR threshold ρ^* that we define can be used to calculate the equivalent threshold on the filter output z^* :

$$z^* = \rho^* \sigma. \quad (4.29)$$

Then we compare the time series $z(t)$ to z^* for each template on each 256 s block of data. In S4, the SNR thresholds were set to the following for each interferometer's data:

$\rho_{H1}^* = 5.5, \rho_{H2}^* = 5.5, \rho_{L1}^* = 5.5$. However, in S5, we chose to lower the SNR threshold of our least sensitive detector H2 so that the thresholds for each interferometer were the following: $\rho_{H1}^* = 5.5, \rho_{H2}^* = 4.0, \rho_{L1}^* = 5.5$. The details of this choice will be outline in Section 4.3.2.

Trigger Management

Due to the sheer numbers of triggers that we find above threshold in the filtering stage, it is advantageous to cluster the triggers from each interferometer in time before coincidence is applied. In S5, we choose to use the same 1 s sliding window clustering method that was used in S4. The method is outlined in detail in [157]. However, the main result of this clustering is that the trigger with the loudest SNR in a cluster of triggers will be marked as *the* trigger to be recorded.

These clustered triggers are written to file with their associated SNR, recovered GPS time, template parameters f and Q , amplitude A and effective distance D_{eff} . The amplitude A is calculated assuming that the amount of energy emitted as gravitational waves during the ringdown phase is $\epsilon = 0.01$ and the distance to the source is $r = 0.5$ Mpc. In S4, this distance was set to $r = 1$ Mpc. The expression used to calculate the amplitude in the S5 search is given by Equation 3.16. Furthermore, the effective distance D_{eff} is a derived quantity based on the amplitude and the sensitivity of the detector to a signal at 1 Mpc. This sensitivity is calculated as

$$\sigma^2 = (1\text{Mpc})\sigma_T^2 A^2 \quad (4.30)$$

where σ_T^2 is the variance of the template from Equation 4.28. Thus the effective distance of a trigger is calculated as

$$D_{\text{eff}} = (1\text{Mpc})\frac{\sigma}{\rho}. \quad (4.31)$$

The physical interpretation of this effective distance is the distance to an optimally positioned and oriented source that produces the same strain in the detector as a source at 1 Mpc.

4.2.3 Background Estimation

As with any experiment, there will always be a background of false alarms that will mimic the behavior of the signal through the various stages of the search pipeline. In fact, in LIGO data, false alarms are quite common compared to the occurrence of an actual signal in the data. Fortunately, we have many methods of distinguishing false alarms due to detector noise from an actual signal. These methods include the coincidence stage which will be discussed in Section 4.2.5 as well as other ranking statistics that will be discussed in Section 4.2.6 that allow us to assign a high rank to a likely signal and a low rank to a likely false alarm. However, with each high-ranked candidate signal, we need to be able to assign a significance to that rank. Thus, we need to be able to attach a false alarm rate (FAR) to a candidate's rank in order to assess the actual significance of a candidate. In other words, how many false alarms will occur in a given period of time with a rank as high as the candidate signal's rank? The hope is that the false alarm rate will be very low implying that the significance of that candidate is very high.

Typically, a detector's background false alarm rate can be determined simply by turning off the source or changing the orientation of the detector. However, these are not options for LIGO. Our sources are many parsecs away and the detector is firmly attached to the earth. Thus, in order to estimate the LIGO detectors' background false alarm rate, we take advantage of the fact that a real signal will be coincident in time between triggers in multiple detectors. By shifting the data stream of one of the detectors in time by an interval longer

than the expected length of the signal and light travel time between sites, we can then do a full analysis, and any coincidences found are guaranteed *not* to be gravitational waves. We then count the number of false alarm coincidences in the time-slide search, record their ranking values, and determine the analysis time of the time-slide search. This allows us to assign false alarm rates for candidate signals at various ranks. In reality, there are limitations to this method based on how many time-slides are done and how loud (or how high of a rank) a candidate signal has. However, for our purposes here, we can assume that the time-slide method is sufficient for estimating the background of false alarms in the S5 ringdown search. For more details on the method, see [157].

In S4, L1 data was slid on a ring fifty times in steps of 5 s and fifty times in steps of -5 s, H2 data in steps of ± 10 s, and H1 data was left in place. In S5, we decided not to slide H2 with respect to H1 because the time-slide method does very poorly at background estimation for colocated detectors. It tends to eliminate known correlations in noise and results in an under-estimated rate of false alarms. Thus, in S5, we chose to slide L1 data on a ring fifty times in steps of 5 s and fifty times in steps of -5 s while both H1 and H2 data were left in place. Time-slides are implemented by the coincidence code `lalapps_rinca`.

4.2.4 Detection Efficiency

In order to assess the sensitivity of a search pipeline to a particular source, we can conduct searches on data that has had simulated gravitational waveforms added. These simulated waveforms are called injections. When combined with time-slide analyses, we can ascertain how well we detect certain types of sources by determining false alarm rates for the injections that we recover and noting which injections are not recovered at all from our pipeline. Studies of injection recovery also aid with tuning different parameters of the search pipeline and setting upper limits on the rates of particular sources.

The recovery of lone ringdown waveform injections and also ringdown waveforms attached to inspiral-merger CBC signals is studied for the ringdown search. These types of waveforms were discussed in detail in Section 3.5. The sensitivity of the ringdown search pipeline differs depending on whether it is recovering IMR signals or lone ringdown signals. The amount of energy emitted as gravitational waves for these two distinct types of signals can differ by orders of magnitude depending on the source, thus affecting the distance to which we can observe as outlined in Section 3.5. Furthermore, the ringdown pipeline recovers the parameters of the two waveforms with different accuracies. The details of parameter recovery of these two types of waveforms is discussed in Section 4.5.

One method of determining sensitivity to particular injections is the efficiency ϵ , or fraction of found injections, as a function of injected physical distance. An example of an efficiency plot is shown in Figure 4.11. Traditionally, found injections are binned in logarithmic distance and the efficiency is calculated in each bin. The uncertainty in each bin is binomial,

$$\sigma_\epsilon^2 = \frac{\epsilon(1 - \epsilon)}{N} \quad (4.32)$$

where N is the total number of injections made. In S5, we choose to count injections missed due to vetoes as a loss in live-time rather than a loss in efficiency. This differs from the S4 method, where the efficiency at small injected distances would never be unity because missed injections due to vetoes were counted as a loss in efficiency. Thus, at small injected distances in S5, we do see efficiencies go to unity. At larger injected distances, we see that the efficiencies fall off gradually due to injections being quieter in the detector and other characteristics that are not obvious in efficiency plots. For instance, efficiency is a strong

function of frequency. It is informative to examine efficiency plots in different frequency bands. Also, another method of determining a search’s sensitivity to a particular source is to examine the injected distance as a function of frequency (or total mass) for missed and found injections. An example plot is shown in Figure 4.35. Here we see injections plotted as red crosses if they were not found as triple coincidences, as blue stars if they were found as triple coincidences with zero combined FAR, and as dots colored by combined FAR if they were found as triple coincidences with a non-zero combined FAR. It is interesting to note the transition region on these plots between injections found with zero combined FAR and missed injections. These are the quietest found injections (those with non-zero FAR) and are the most influential during tuning; we watch how the FARs of these quiet injections change with different tunings of the pipeline.

4.2.5 Coincident Trigger Analysis

The coincidence code is `lalapps_rinca`. It gathers triggers from multiple interferometers and compares them to see if any of the triggers fall within the allowed range of one another. This code also has the capability to apply Category 2-4 vetoes, cluster single detector triggers across the template bank, and apply an H1H1 consistency test although, for S5, we chose to only utilize this code for the first to capabilities. The application and types of vetoes will be discussed in more detail in Section 4.4.1.

The coincidence test utilizes information from all interferometers to determine whether coincident triggers in each are consistent. Of course, the signals must be close in time, within the light travel time (10 ms) between Hanford and Livingston. We also require that these signals be close in frequency and quality parameters. If triggers fail the coincidence test, they are discarded. Otherwise, they are recorded as a coincidence and are passed on to the next stage of the pipeline. For the S5 analysis, there are four possible types of coincidence: a triple coincidence consisting of triggers from H1, H2, and L1 and three types of double coincidence, H1H2, H1L1, and H2L1. For reasons that will be explained below, we chose to only consider triple coincident triggers from H1H2L1 operation time, double coincident H1L1 triggers from H1L1 and H1H2L1 operation times, and double coincident H2L1 triggers from H2L1 operation time.

In S4, the ringdown coincidence test was composed of two separate stages, a test to check for time coincidence and then another which made use of the metric to measure the distance between the triggers’ templates. However, in S5 we have utilized a new test that combines information into a single coincidence check. We will first motivate the move to this new test and then describe its properties.

S4 Coincidence Test

Time coincidence is the required coincidence time window of δt for co-located detectors and $\delta t + 10\text{ms}$ for Hanford-Livingston pairs. Any pair of triggers outside this window were not considered close enough in time to be from a real signal. The value of δt was chosen as 2 ms by tuning the width of the window to detect as many injections as possible while keeping the false alarm rate low. Thus, the test required $\delta t \leq 2\text{ms}$ for H1H2 pairs and $\delta t \leq 12\text{ms}$ for H1L1 pairs.

Parameter coincidence, for two triggers in a coincidence, is the metric distance ds^2 between their best match templates was calculated. It was assumed that the template closest in f and Q to an actual signal would be the one with the loudest SNR. Of course, the effects of detector noise will change this somewhat. Recall that the templates’ metric distance in the template bank is calculated as

$$ds^2 = \frac{1}{8} \left[\frac{3 + 16Q^4}{Q^2(1 + 4Q^2)^2} dQ^2 - 2 \frac{3 + 4Q^2}{fQ(1 + 4Q^2)} dQdf + \frac{3 + 8Q^2}{f^2} df^2 \right]. \quad (4.33)$$

Thus, this can be used to define a window (or contour of constant ds^2 as described in Section 4.1.2) that varies the size of df and dQ depending on the location in the bank. For two triggers a and b , the frequency difference is computed as $df = f_a - f_b$ and the quality factor difference is computed as $dQ = Q_a - Q_b$. A value of $ds^2 \leq 0.05$ was chosen based on tuning.

However, it was shown in S4 that the time of arrival difference of recovered coincidences is actually frequency dependent. See Figures 4.9 and 4.10. It was determined that a frequency dependent time cut (instead of the flat dt time cut) could eliminate a large amount of background in the mid-to-high frequency range and $dt > 2\text{ms}$.




Figure 4.9: Time of arrival difference versus frequency in S4 injections in H1H2.

S5 Coincidence Test

In order to include information about dt , df , and dQ in a single coincidence test, a 3D metric was constructed to calculate distances in (f, Q, t) -space. The goal of this new test was to reproduce the injection recovery efficiency of the S4 coincidence test while greatly reducing the background. The following line element is used:




Figure 4.10: Time of arrival difference versus frequency in S4 background in H1H2.

$$\begin{aligned}
 ds^2 = & \frac{1 + 6Q^2 + 16Q^4}{4f^2(1 + 2Q^2)} df^2 + \frac{1 + 28Q^4 + 128Q^6 + 64Q^8}{4Q^2(1 + 6Q^2 + 8Q^4)^2} dQ^2 + \\
 & \frac{f^2\pi^2(1 + 4Q^2)}{Q^2} dt^2 - \frac{1 + 2Q^2 + 8Q^4}{2fQ(1 + 6Q^2 + 8Q^4)} df dQ - \\
 & \frac{2\pi Q(1 + 4Q^2)}{1 + 2Q^2} dt df + \frac{2\pi f(1 - 2Q^2)}{(1 + 2Q^2)^2} dt dQ.
 \end{aligned} \tag{4.34}$$

to define the metric distance. However, for non-collocated detectors, a temporal offset is caused by light travel time introducing an unknown time of arrival discrepancy dependent on the direction of the origin of the signal. Thus, the apparent dt may not reflect the true dt . This is not a problem for collocated detectors like H1 and H2 so this test recovers injections just as well the S4 coincidence test but with a much reduced H1H2 background. See Figures 4.11 and 4.12 for a demonstration of this improvement. But for non-collocated detectors, this 3D test would naively miss many injections. Thus in order to combat the signal travel problem, a small brute force search is done through all allowed corrected time differences for the dt that gives the minimum ds^2 . The search selects the smallest ds_{\min}^2 for a range of corrected time differences. It is this ds_{\min}^2 that is compared to the user-specified ds^2 threshold. While this minimization method would help fix time discrepancies, it does add a bias but we still see a reduced H1L1 background compared to S4. See Figures 4.11 and 4.12 for a demonstration of this improvement in the S5 H1L1 background.




Figure 4.11: Efficiency versus effective distance and background events per slide in H1H2 for the S4 coincidence test statistic.




Figure 4.12: Efficiency versus effective distance and background events per slide in H1H2 for the new S5 coincidence test statistic. Note that the efficiency is the same as in Figure 4.11 for the old statistic but the number of H1H2 background coincidences per time slide has consistently decreased.

For any two triggers, ds^2 will be a scalar. We can define a ds^2 threshold to declare whether triggers will be coincident or not. For S5, we set the coincidence threshold at $ds^2 = 0.4$. See Section 4.3.3 for details regarding this choice. This value was tuned concurrently with the maximization interval, or the time period over which single detector triggers were clustered

across the template bank. In S4, this window was set to 1 ms. However, in S5, we chose to shorten this window to 0.5 ms.

In S4, the coincidence test was followed by an H1H2 consistency cut, coincidence clustering, effective distance cuts, detection statistic calculation, and more using both `lalapps_rinca` and `lalapps_coincringread`. In S5, however, we have moved all of these operations, except the H1H2 consistency cut, to the PIPEDown post-processing code. See Section 4.2.6 for more details on PIPEDown.

The H1H2 consistency cut relies on the fact that if H1 sees a trigger that H2 does not see, then if H1 is above a certain SNR threshold and H2 is not vetoed, then H2 should have seen that trigger as well. Since H1 and H2 are collocated and H2 is half as sensitive as H1, for any H1 trigger, we would expect an H2 trigger at about half the H1 SNR. Of course, if the SNR in H1 is low, then it may be the case that the SNR in H2 does not fall above our pre-set SNR thresholds. However, for a loud H1 trigger, we would expect to see a trigger in H2; otherwise, the lack of a trigger in H2 indicates that we do not have a real signal and the coincidence can be vetoed and removed from the pipeline. Similarly, if H2 sees a trigger that H1 does not see, then we do not have a real signal. This test obviously applies only to H1L1 and H2L1 double coincidences in H1H2L1 triple time.

A small study was done to explore the H1 SNR threshold at which we could apply a cut for H1L1 coincidences in triple time. For an H2 SNR threshold of 5.5, the S4 pipeline's threshold for storing triggers, we would expect an H1 trigger with $\text{SNR} = 2 \times 5.5 \times \text{padding} \approx 16$. This conservative cut was found to reduce the H1L1 background by a small amount. For example, see the plot of H1L1 background triggers in triple time in Figure 4.13. The triggers in red are those that would be removed by an H1 SNR cut of 16. Thus, there is a reduction in the long tail of the H1 background triggers.

For our initial S5 playground runs, we did not use the H1H2 consistency cut. See Section 4.4.2. However, after examining playground results, it was determined that the test was needed and it was turned on in the final S5 analysis. See Section 4.4.3.

4.2.6 Post-processing with Pipedown

In S4, the ringdown search analyzed coincident triggers with a post-processing pipeline that was run step-by-step by hand. Many of the plotting scripts were lost and the entire post-processing analysis would be difficult to reproduce. Thus, in S5, we opted to standardize our post-processing pipeline and fully automate the steps in order to reduce the chance of human error and speed up the analysis. We chose to adapt the pre-existing CBC post-processing pipeline for the ringdown search.

The post-processing pipeline serves several purposes. The pipeline is run after the coincidence stage of the `ihope` pipeline has completed. It combines data (zero-lag, time-slides, injection runs) from all the coincident `RINCA` files into a single SQLite database. These databases are versatile and computationally less expensive to manipulate than `LIGO_LW XML` files, the products of the `ihope` pipeline. These databases simplify the calculation of new coincident statistics, false alarm rates and the plotting of multiple coincident characteristics. The details of the advantages of SQLite databases and their use within Pipedown are outlined in [159]. Below, we will outline the main parts of the Pipedown post-processing pipedown and its specific uses for the S5 ringdown search. Figure 4.14 outlines a simple flowchart of the the post-processing pipeline.

Store Coincidences

After the coincidence test has been performed, any single detector triggers found to be in coincidence with a trigger from at least one other detector are stored in an XML file in the `sngl_ringdown` table. An XML file is created for each `lalapps_rinca` job. Then




Figure 4.13: Application of H1H2 consistency cut to H1L1 background coincidences in triple time. Only H1L1 coincidences in triple time with an H1 SNR less than 16 were retained.

time slides are performed and higher category vetoes are applied (ie higher than Category 1). All coincident information from zero-lag, time slides, and injections were stored in the `sngringdown` table with an associated `event_id`:

- `ifo`- the instrument associated with a single trigger. In the S5 ringdown search, this could be either H1, H2, or L1.
- `frequency`- the frequency of the *best-match* template for that trigger.
- `quality`- the quality factor of the *best-match* template for that trigger.
- `mass`- the mass of the final black hole derived from the frequency and quality factor of the *best-match* template for that trigger.
- `spin`- the spin of the final black hole derived from the frequency and quality factor of the *best-match* template for that trigger.

- `start_time`- the GPS start time of the ringdown trigger.
- `start_time_ns`- the GPS start time nanosecond piece of the ringdown trigger. This value is stored to allow nanosecond timing of triggers.
- `snr`- the raw single detector SNR as defined in Equation 4.11.
- `eff_dist`- the distance to an optimally positioned and oriented source as defined in Equation 4.31
- `ds2_H1H2`- the 3D distance between H1 and H2 as described in Section 4.2.5.
- `ds2_H1L1`- the 3D distance between H1 and L1 as described in Section 4.2.5.
- `ds2_H1V1`- the 3D distance between H1 and V1 as described in Section 4.2.5. This column is not used in the S5 search.
- `ds2_H2L1`- the 3D distance between H2 and L1 as described in Section 4.2.5.
- `ds2_H2V1`- the 3D distance between H2 and V1 as described in Section 4.2.5. This column is not used in the S5 search.
- `ds2_L1V1`- the 3D distance between L1 and V1 as described in Section 4.2.5. This column is not used in the S5 search.
- `event_id`- a unique identifier for each single trigger in the `sngl_ringdown` table. This id is then used to map entries in `sngl_ringdown` to entries in other tables with the same id such as an entry in the `coinc_event_map` table to map to a coincidence in the `coinc_ringdown` table.

The `event_id` was meant to act as a unique identifier to pair up triggers in a coincidence. The id was 18 digits long- the first 9 digits were the GPS time of the trigger, the next 4 digits gave the slide number (0000 denotes zero-lag), and the last 5 were a counter for each coincidence. Thus, any rows in the `sngl_ringdown` table with the same last 9 digits would be considered coincident triggers. However, one can imagine scenarios where this counting system might run into trouble. For example, no more than 99999 coincidences or 9999 timeslides could be stored in a single file. Furthermore, as is often the case, multiple coincidences are formed with the same triggers and a whole new entry would need to be created despite redundant information. One can imagine the heavy-lifting required to extract coincidences from the `sngl_ringdown` tables.

To remedy these issues in S5, we implemented the coincident event tables developed in [159] and used by the CBC searches in S5 and S6. These tables include the `coinc_ringdown`, `coinc_event`, `coinc_event_map`, and `coinc_definer` tables. The uses of these tables are identical to the uses by the CBC searches as outlined in [159]. The only differences are that we have renamed the `coinc_inspiral` table as the `coinc_ringdown` table and the columns in the `coinc_ringdown` table are specific to the ringdown search. The creation of these tables is done through the code `ligolw_rinca_to_coinc`. Below, we give a brief description of each table and its contents.

The `coinc_ringdown` table stores coincident information rather than single trigger information. Each row in the `coinc_ringdown` outlines a single coincidence. The values stored in this table are from combined statistics rather than from a single trigger as is done in the `sngl_ringdown` table. The columns stored for the S5 search are outlined below. The values that we choose to store are used for multiple tasks including sorting, clustering, plotting, and ranking with the multi-variate statistical classifier.

- **ifos**- a list of the interferometers forming the coincidence. The values that this column could take for S5 would be { H1,H2,L1}, { H1,L1}, or {H2,L1}.
- **frequency**- stores the SNR-weighted mean of the `sngl_ringdown` table frequencies f_i from N triggers in the coincidence:

$$\bar{f}_{\text{coinc}} = \frac{\sum_{i=1}^N f_i \cdot \text{SNR}_i}{\sum_{i=1}^N \text{SNR}_i} \quad (4.35)$$

- **quality**- stores the SNR-weighted mean of the `sngl_ringdown` table quality factors Q_i from N triggers in the coincidence:

$$\bar{Q}_{\text{coinc}} = \frac{\sum_{i=1}^N Q_i \cdot \text{SNR}_i}{\sum_{i=1}^N \text{SNR}_i} \quad (4.36)$$

- **mass**- stores the SNR-weighted mean of the `sngl_ringdown` table final black hole masses M_i from N triggers in the coincidence:

$$\bar{M}_{\text{coinc}} = \frac{\sum_{i=1}^N M_i \cdot \text{SNR}_i}{\sum_{i=1}^N \text{SNR}_i} \quad (4.37)$$

- **spin**- stores the SNR-weighted mean of the `sngl_ringdown` table final black hole spins \hat{a}_i from N triggers in the coincidence:

$$\bar{a}_{\text{coinc}} = \frac{\sum_{i=1}^N \hat{a}_i \cdot \text{SNR}_i}{\sum_{i=1}^N \text{SNR}_i} \quad (4.38)$$

- **start_time**- the time taken from the start time of the first detector's trigger when sorted alphabetically. For example, for an H1L1 coincidence, the coincident start time will be the value of the H1 start time.
- **start_time_ns**- the GPS start time nanosecond piece; see above.

- **kappa**- stores the effective distance ratio for H1/H2. This value is not currently used but would be useful if a cut on effective distance ratio was applied. Instead, we apply a cut on the value stored in the **snr_ratio** column in the coincidence clustering stage. If a coincidence does not have both an H1 trigger and an H2 trigger, a default value of 1.0 is stored.
- **snr**- stores the coincidence statistic from S4. This is two separate equations depending on whether one has a triple coincidence $\rho_{\text{DS,trip}}$ or a double coincidence $\rho_{\text{DS,doub}}$ as shown below. For more information regarding this statistic, see Section 4.3.4. Note that $\rho_{\text{DS,trip}}$ is actually a squared coincident statistic while $\rho_{\text{DS,doub}}$ is not. This was done to ensure that triples would be ranked above doubles in S4 so the two equations are not directly comparable:

$$\rho_{\text{S4_DS,trip}} = \rho_{\text{H1}}^2 + \rho_{\text{H2}}^2 + \rho_{\text{L1}}^2, \quad (4.39)$$

$$\rho_{\text{S4_DS,doub}} = \min \{ \rho_{\text{IFO1}} + \rho_{\text{IFO2}}, 2.0\rho_{\text{IFO1}} + 2.2, 2.0\rho_{\text{IFO2}} + 2.2 \}. \quad (4.40)$$

- **snr_sq**- stores the sum of the SNRs squared. This value is used in the coincidence clustering stage in order to break the degeneracy of identically ranked coincidences found at the same GPS time. See 4.4.2.
- **snr_ratio**- stores the SNR ratio for H1/H2. A cut on this value is applied in the coincidence clustering stage and removes any coincidences with an H1H2 snr ratio of less than 1.0. If a coincidence does not have both an H1 trigger and an H2 trigger, a default value of 2.0 is stored thus ensuring that the coincidences will not be removed by the cut.
- **chopped1_snr**- stores the coincidence statistic from S5. The name denotes the fact that both the triple statistic and double statistic use a chopped-L-like statistic. For more information regarding this statistic, see Section 4.3.4. Again, though, note that $\rho_{\text{DS,trip}}$ is actually a squared coincident statistic.

$$\rho_{\text{S5_DS,trip}} = \min \left\{ \sqrt{\rho_{\text{H1}}^2 + \rho_{\text{H2}}^2 + \rho_{\text{L1}}^2}, \rho_{\text{H1}} + \rho_{\text{H2}} + 0.75, \rho_{\text{H1}} + \rho_{\text{L1}} + 0.75, \right. \\ \left. \rho_{\text{H2}} + \rho_{\text{L1}} + 0.75 \right\}^2, \quad (4.41)$$

$$\rho_{\text{S5_DS,doub}} = \min \{ \rho_{\text{IFO1}} + \rho_{\text{IFO2}}, 2.0\rho_{\text{IFO1}} + 2.2, 2.0\rho_{\text{IFO2}} + 2.2 \}. \quad (4.42)$$

- **eff_coh_snr**- stores the values calculated from the coherent ringdown pipeline. The complex SNRs of coincident triggers from each interferometer are combined with appropriate weights (ie template norms, time delays, antenna factors) to compute the coherent SNR and also the null stream in the next column. These values were not used for the S5 search due to the need for further investigation. These columns will likely be used in the S6 ringdown search. Further references can be found in [160].
- **null_stat**- see above.
- **false_alarm_rate**- calculated from time-slides and associated detection statistic as described in Section 4.2.3. More details in the following description of the post-processing code

- `combined_far`- the combined false alarm rate. More details in the following description of the post-processing code
- `coinc_event_id`- a unique identifier for each coincidence providing a quick way to extract coincident information. Every entry in the `coinc_ringdown` table will have a unique `coinc_event_id`. This id is then used to map entries in `coinc_ringdown` to entries in other tables with the same id such as a single entry in the `coinc_event` table or multiple entries in the `coinc_event_map` table. This is the coincidence analogy to the `event_id` in the `snl_ringdown` table.

The `coinc_event_map` table maps coincident triggers in the `coinc_ringdown` table to the triggers forming the coincidence in the `snl_ringdown` table. This table has only three columns which are outlined below:

- `coinc_event_id`- the unique identifier that maps back to the `coinc_ringdown` table.
- `event_id`- the id of the single trigger in the `snl_ringdown` table belonging to a coincidence. There may be multiple different `event_id`'s mapped to a single `coinc_event_id` if there are multiple participants in the coincidence. Furthermore, a single `event_id` may be mapped to multiple different `coinc_event_id`'s if it takes part in multiple coincidences.
- `table_name`- the name of the table from whence a single event came. It could be `snl_ringdown` but it could also draw mappings between coincident events and injections.

The `coinc_event` stores more information about coincident triggers. Every coincident event has one entry in this table via the `coinc_event_id`.

- `instruments`- gives the coincident time that the coincident event occurred in for computing false alarm rates.
- `likelihood`- column populated by MVSC likelihood statistic. See Section 4.3.4 for further details.
- `coinc_event_id`- the unique identifier that maps back to the `coinc_ringdown` table.
- `time_slide_id`- an id that allows a mapping to the `time_slide` table. This table provides information on the offset and detector that was analyzed.
- `process_id`- an id that maps back to the `process` table to provide information about the program that created the entry in the table.
- `coinc_def_id`- the unique identifier that maps back to the `coinc_ringdown` table.

The `coinc_definer` defines what types of coincidences we can have. For instance, we can have the normal coincidence between single detector triggers. However, we could also have a coincidence between a `coinc_ringdown` entry and an injection with parameters stored in the `sim_ringdown` table. Thus, coincidence types can take an abstract form. The mappings used for the S5 ringdown search are shown below:

- `snl_ringdown<-->snl_ringdown` coincidences- all coincident events mapped to `snl_ringdown` triggers

- `sim_ringdown<-->snl_ringdown` coincidences- all `snl_ringdown` events mapped to injections.
- `sim_ringdown<-->coinc_event` coincidences (nearby)- all coincident events mapped to injections from our original injection-finding code, `lalapps_cbc_injfind`
- `ring_exact`- a subset of coincident events mapped to injections that are subject to strict coincidence conditions from our specialized injection-finding code, `ligolw_dbinjfind`. This set is used for the training set of MVSC. See Section 4.3.4 for more details.
- `ring_nearby`- all coincident events mapped to injections from our specialized injection-finding code, `ligolw_dbinjfind`. This mapping acts as a replacement for `lalapps_cbc_injfind` mappings. See Section 4.2.4 and 4.3.4 for more details.

Convert to SQLite Databases

Once we have created specialized tables to store coincidence information in the `RINCA_TO_COINC` XML file using `ligolw_rinca_to_coinc`, we convert those files to SQLite databases using the code `ligolw_sqlite`. This combines multiple XML files into one database including both zero-lag and time-slide data.

Clean Databases

Using the code `ligolw_cbc_dbsimplify`, we remove redundant entries in tables introduced by the addition of multiple `RINCA_TO_COINC` files. This is mostly redundant information that describes which interferometers were on, time-slide information, and `coinc_definer` mappings. This information remains the same across multiple databases and thus, only needs to be logged once per SQLite database.

Calculate Coincident Statistics

After the SQLite databases have been created and tidied up, it is straightforward to add new columns to the `coinc_ringdown` table using the code `ligolw_cbc_repop_coinc`. These new columns generally store coincident statistics that are calculated from the single detector parameters in the `snl_ringdown` table. The `coinc_ringdown` columns were listed above and from that list, those that were populated using `ligolw_cbc_repop_coinc` are the following.

The `coinc_ringdown` frequency, quality factor, mass, and spin columns are calculated and populated using a call to the function `snr_mean` in `ligolw_cbc_repop_coinc`. This function calculates the SNR-weighted mean of any single detector parameter p_i over N detectors and stores the value as the coincident parameter \bar{p}_{coinc} :

$$\bar{p}_{\text{coinc}} = \frac{\sum_{i=1}^N p_i \cdot \text{SNR}_i}{\sum_{i=1}^N \text{SNR}_i}. \quad (4.43)$$

These values are stored for plotting purposes and parameter recovery investigations.

The `coinc_ringdown` kappa and snr ratio columns are calculated using a call to the function `ratio` in `ligolw_cbc_repop_coinc`. This function takes as input interferometer 1, interferometer 2, a default value, and either effective distances or SNRs. It then calculates the ratio of either the effective distances or SNRs of interferometer 1 to interferometer 2. If a particular coincidence does not have either interferometer 1 or interferometer 2, then a

default value is stored. These values are stored in order to apply cuts during the coincidence clustering stage as described next. In S5, as in S4, we exploit the fact that H1 is collocated and twice as long as H2 so we expect $\rho_{\text{H1}} \approx 2\rho_{\text{H2}}$ and $\sigma_{\text{H1}}^2 \approx 2\sigma_{\text{H2}}^2$. From this information and Equation 4.31, we then expect that H2 SNR should be about half of H1 SNR, and H1 and H2 effective distances should be approximately the same. Thus the kappa column was calculated as

$$\kappa = \begin{cases} \frac{d_{\text{eff,H1}}}{d_{\text{eff,H2}}} & : \text{H1 and H2 exist} \\ 1.0 & : \text{H1 or H2 does not exist} \end{cases}$$

and the snr ratio column was calculated as

$$\kappa = \begin{cases} \frac{\rho_{\text{H1}}}{\rho_{\text{H2}}} & : \text{H1 and H2 exist} \\ 2.0 & : \text{H1 or H2 does not exist.} \end{cases}$$

The `coinc_ringdown` snr squared column is calculated using a call to the function `snr_sq` in `ligolw_cbc_repop_coinc`. This function simply calculates the quadrature sum of the single SNRs in the coincidence. The value is used to break the degeneracy of identically ranked coincidences at the same GPS time as described in Section 4.4.2. If there are N interferometers in the coincidence, then this value is calculated as

$$\rho_{\text{combined}}^2 = \sum_{i=1}^N \rho_i^2. \quad (4.44)$$

The `coinc_ringdown` snr column is actually calculated using the S4 ranking statistic using a call to the function `chopped1` in `ligolw_cbc_repop_coinc`. This column is used in the multivariate statistical classifier. The S4 statistic, which is described in more detail in Section 4.3.4, is calculated as

$$\rho_{\text{S4_DS,trip}} = \rho_{\text{H1}}^2 + \rho_{\text{H2}}^2 + \rho_{\text{L1}}^2, \quad (4.45)$$

$$\rho_{\text{S4_DS,doub}} = \min \{ \rho_{\text{IFO1}} + \rho_{\text{IFO2}}, 2.0\rho_{\text{IFO1}} + 2.2, 2.0\rho_{\text{IFO2}} + 2.2 \} \quad (4.46)$$

where IFO1 and IFO2 can be H1L1 or H2L1 for S5.

Finally, the `coinc_ringdown` chopped-L snr column is calculated using the new S5 triples' statistic using a call to the function `chopped1forall` in `ligolw_cbc_repop_coinc`. This statistic, which was developed to deal with the characteristics of triple coincident background in S5 data, uses a chopped-L-like statistic for both doubles and triples:

$$\rho_{\text{S5_DS,trip}} = \min \left\{ \sqrt{\rho_{\text{H1}}^2 + \rho_{\text{H2}}^2 + \rho_{\text{L1}}^2}, \rho_{\text{H1}} + \rho_{\text{H2}} + 0.75, \rho_{\text{H1}} + \rho_{\text{L1}} + 0.75, \right. \\ \left. \rho_{\text{H2}} + \rho_{\text{L1}} + 0.75 \right\}^2, \quad (4.47)$$

$$\rho_{\text{S5_DS,doub}} = \min \{ \rho_{\text{IFO1}} + \rho_{\text{IFO2}}, 2.0\rho_{\text{IFO1}} + 2.2, 2.0\rho_{\text{IFO2}} + 2.2 \}. \quad (4.48)$$

For more details on this S5 statistic, refer to Section 4.3.4.

Cluster Coincidences

Even after single event clustering in the `ihope` pipeline, we still need to do a round of coincidence clustering in the `pipedown` post-processing stage. As a new feature of the S5 ringdown post-processing pipeline, we have implemented the ability to do multiple clustering stages. This is all done with the code `ligolw_cbc_cluster_coincs` which will cluster coincidences with a sliding time window of 10 s based on combined statistics.

The first clustering stage which we have denoted `choppedl_snr_cluster` performs multiple tasks. First, it removes any coincidences with an `snr_ratio` value less than 1.0. These are coincidences where H2 SNR was greater than H1 SNR which should not happen considering the two detectors are collocated and H2 is half as sensitive as H1. We can remove these types of coincidences with high confidence that they are due to detector noise. Furthermore, we choose to remove all H1H2 coincidences due to the difficulty in estimating the background for these types of coincidences. We also choose to remove H2L1 coincidences in H1H2L1 triple time. Finally, the code clusters coincidences with a 10 s sliding window by the maximum S5 chopped-L statistic in that window. This clustering is split by interferometers that took part in the coincidence so that triples are clustered together and doubles are clustered together separately. In other words, H1L1, H2L1, and H1H2L1 coincidences are all clustered by maximum chopped-L statistic separately. Furthermore, separate clustering jobs are performed for time-slides and different data-types (playground, zero-lag, simulation) which all considered distinct experiments.

The second clustering stage which we have denoted `snr_sq_cluster` will remove identically ranked coincidences at the same GPS time. This was an issue uncovered by our initial examination of S5 playground results in Section 4.4.2. If identically ranked coincidences are found at the same GPS time, they will not be clustered together in the first clustering stage. However, we wish to select only one of these identically ranked coincidences since one coincidence is all we need to have a detection. The problem is due to the discreteness of the chopped-L statistic; thus, by introducing another statistic that takes into account the unique single SNRs of each identically ranked coincidence, we can break the degeneracy and cluster again in a 10 s window by this new statistic. This stage is run after the `choppedl_snr_cluster` stage so that the only thing that remains unclustered are identically ranked coincidences. Thus, we cluster again by the maximum of the quadrature sum of the SNRs (the `snr_sq` column) in a 10 s window.

Include Injection Information

At this point, `pipedown` adds information about injections so that mappings can be made between injections and coincidences. First `ligolw_cbc_dbaddinj` adds the injection files that were created with `lalapps_inspinj` to the `sim_inspirals` and `sim_ringdown` tables and with `lalapps_rinj` to the `sim_ringdown` table. Then, we initially ran `lalapps_cbc_injfind` to perform the injection finding, or the mapping of injections to triggers that were caused by injections. This was done with a 1 s time window so that any coincident trigger that occurs with ± 1 s of an injection is considered to have been caused by that injection. This code also drew mappings between the `sim_inspirals` and `sim_ringdown` tables so that we could extract an IMR injection's ringdown parameters.

However, with the adoption of the multivariate statistical classifier MVSC as our detection statistic, we needed a more sophisticated injection finding code. The reason we required this was because MVSC's power as a signal-glitch classifier is dependent on how well we train it with a well-constructed training set of signal coincidences and background coincidences. The *signals* that we give to MVSC are recovered software injections and the *background* is composed of time-slide coincidences. However, a 1 s time window for injection finding is actually a loose requirement and leads to multiple events being mapped to a single injection. And oftentimes, random noise events are mapped to injections. While we do not care about

such things when calculating upper limits (it does not matter how we detected the software injection as long as we detected it), it is very important that MVSC not be taught that these random noise events are injections. Otherwise, when MVSC is ranking events in zero-lag and encounters random noise events, it will give them high rankings because we taught it that random noise is signal-like.

In order to combat this issue with MVSC, we adopted a new injection-finding code called `ligolw_dbinjfind`. This code allows us to make multiple mappings for different types of injection-finding. For example, for the MVSC training set, we need to define a strict, clean set of found injections. In order to do this, we require that triggers be associated with injections within a much smaller time window than 1 s. Furthermore, we throw-out any injections that could possibly be mapped to random noise. The code does this by requiring a time window to check zero-lag data for triggers. If zero-lag data is found to have a trigger in any of the interferometers within the time window around the injection time, then that injection is not used for the MVSC training set. These time windows had to be tuned to optimize performance as outlined in Section 4.3.4.1. We designated the MVSC training injection set with the map label *ring_exact*. The injection finding window was tuned so that a trigger’s start time must be within ± 0.01 s of the injection’s start time. Furthermore, we required that no trigger be found in zero-lag data within ± 0.5 s of the injection’s start time. If so, that injection was thrown out of the training set.

In order to compute upper limits, we also needed to define mappings for another injection set that gave a more realistic view of our injection finding ability. We designated this injection set with the map label *ring_nearby*. This set was meant to reproduce the set from our original injection finding code, `lalapps_cbc_injfind`. The only requirement is that a trigger’s start time must be within ± 1.0 s of the injection’s start time.

Compute False Alarm Rates

In Section 4.2.3, we discussed how one generates a background of gravitational waves by adding time offsets to each detector and looking for coincidences across the time-slid detectors. The offsets should be large enough that each slide can be thought of as an independent experiment. As we discussed, in S5, we chose to slide L1 data on a ring fifty times in steps of 5 s and fifty times in steps of -5 s while both H1 and H2 data were left in place. Thus, we performed a total of 100 time slide experiments and one zero-lag experiment. The coincident time-slide triggers are then binned by coincidence type and time.

First we compute the uncombined false alarm rate (FAR) for every zero-lag trigger in each bin. This is done by counting the number of time slide coincidences with a combined detection statistic greater-than-or-equal-to the combined detection statistic of the given zero-lag coincidence. In S5, the detection statistic that we chose to rank candidate with was the MVSC likelihood Λ . The FAR will be the count of these coincidences with a MVSC likelihood $\Lambda \geq \Lambda^*$, where Λ^* is the detection statistic value we are interested in, divided by the total slide duration T_b for all slides performed:

$$\text{FAR} = \frac{\sum_{k=1}^{N_t} N_k (\Lambda \geq \Lambda^*)}{T_b} \quad (4.49)$$

where N_k is the measured number of coincidences with $\Lambda \geq \Lambda^*$ in the k^{th} slide and N_t is the total number of slides performed. The total slide duration T_b is the sum of the length of each slide which is the intersection of every detector’s analysis time after being slid. The same method is then used to compute uncombined FARs for time-slide triggers except that the time-slide of interest is treated as zero-lag and every other time-slide is examined.

Then the final combined false alarm rate ($\text{FAR}_{\text{combined}}$) is calculated across all bins. This is done for the zero-lag triggers by counting the number of time-slide triggers across all bins than have an uncombined FAR less-than the given trigger's FAR*:

$$\text{FAR}_{\text{combined}} = \frac{N_{\text{tot}}(\text{FAR} \leq \text{FAR}^*)}{T_b}. \quad (4.50)$$

Directly after injection information is included in the database, veto segment information is included. At this point the databases have all the information needed to compute the FARs. The code `ligolw_cbc_compute_durations` is run to compute the analysis times in both zero-lag and time-slides. Then, both the uncombined and combined FARs are calculated using `ligolw_cbc_cfar`. This code is run twice. The first run calculates the uncombined false alarm rate based on the MVSC `likelihood` column in the `coinc_event` table and the calculation is done by turning on the option `--group-by-multiplicity` (all doubles are grouped together and all triples are grouped together). This then populates the `false_alarm_rate` column in the `coinc_ringdown` table. Then the code is run again to calculate the combined false alarm rate based on the uncombined false alarm rates. There is no grouping of coincidences in this step. The `combined_far` column is then populated in the `coinc_ringdown` table.

The calculated $\text{FAR}_{\text{combined}}$ is actually the final detection statistic of the ringdown search. Thus, it is these values that form our figure-of-merit when determining whether we have a significant outlier or not. These figure-of-merit plots will be presented in Section 4.4.3 where we examine the results of the S5 ringdown search. Here, we will briefly describe the features of these plots which are produced with the code `ligolw_cbc_plotifar`. These plots are cumulative histograms of the coincidences as a function of *inverse* false alarm rate, IFAR. We could have chosen to plot cumulative histograms as function of false alarm rate but due to historical reasons, it was desirable to show our most significant candidates to the far right of the plot. Since significance is inversely proportional to false alarm rate, these plots are made as a function of inverse false alarm rate. The plots are generally made for both inverse uncombined FAR and inverse combined FAR as shown in Section 4.4.2.1. Also shown is a black dashed line which is meant to illustrate the mean of the expected background distribution which is simply $1/\text{IFAR}$. Yellow regions mark the 1 and 2σ significance bands in a Gaussian distribution, $\pm\sqrt{N}$ and $\pm\sqrt{N}$, respectively, where N is the number of counts. A grey line is plotted for every time-slide experiment thus demonstrating the actual background distribution. Colored triangles mark the zero-lag distributions. In the uncombined plot, these colored triangles represent different detector combinations. The dashed colored lines mark points of minimum background IFAR for a particular detector combination based on the color. If a coincidence is louder than all the background the FAR is set to zero and, on the IFAR plots, we plot it at $x = T_b$ with an arrow pointing to the right. These are the coincidences that we are most interested in.




Figure 4.14: The ringdown Pipedown post-processing pipeline.

4.3 Tuning the Pipeline

The settings for the ringdown pipeline were chosen via “blind analyses” with a variety of different settings for multiple parameters that we will explore below. The purpose of blind analyses is to ensure that the constraints are decided upon before we conduct the search and determine whether we have detected a gravitational wave. We allow ourselves three tools for blind analyses: time-slides, simulated gravitational waveforms, and a small segment of real data known as the playground. We discussed the method of performing time-slides and simulated gravitational waveforms, or injections, into the data in Sections 4.2.3 and 4.2.4 in order to determine detection efficiency and false alarm rates. We will discuss the results of the playground analyses in Section 4.4.2. The parameters that needed to be tuned before running the S5 search were SNR thresholds, the coincidence test threshold, a new detection statistic and its various settings. Our goal in choosing the correct setting for the parameters is to maximize the number of recovered injections while minimizing the false alarm rate at which they are recovered. Once we have settled on final tunings, we can run the search and “open the box,” or examine the full data search for detections.

4.3.1 Injection Analysis

The injection sets used for the tuning of the pipeline including both ringdown-only injections and full inspiral-merger-ringdown injections. In each case, the simulated waveform is generated in the software and added to the data stream. This can be done many times. For example, for two months of data, we can make hundreds of injections. This is done by specifying a time-step at which to place injections. We specify 724.077 s. However, to randomize the placement, we then provide a time interval of 300 s around each time-step in which the injection is actually placed. Then a search can be performed over the data with the added injections just as one would normally do with the real data. Furthermore, one can choose to set up multiple injection searches with different injection sets. Thus, we can study the recovery of thousands of injections to decide on the best tuning for a search.

For IMR injection runs, we performed three sets of PhenomB injections. All of the sets had total masses between 50-450 M_{\odot} and mass ratios between 0.1-1.0. The total masses and mass ratios were distributed uniformly. Two sets were given component dimensionless spin parameters between 0.1-0.85; the third set had dimensionless spin set to 0.0. The inclination and polarization angles as described in Section 2.3.1 and the initial phase were uniformly distributed. Finally, the non-spinning set was distributed linearly between 10-10⁶ kpc and the two spinning sets were distributed uniformly in logarithmic distance between 10 – 10⁴ kpc and 10³ – 10⁶ kpc, respectively.

In Figure 4.15, we see how the PhenomB injections were distributed in the component mass plane m_1 - m_2 and in the component spin plane a_1 - a_2 . The color bar represents the value of the $\chi = (m_1 a_1 + m_2 a_2) / (m_1 + m_2)$ parameter in the left plot and the mass ratio in the right plot. In Figures 4.16 and 4.17, we can see how these values translate to injections made in the frequency and quality factor plane and in the final black hole mass and spin plane, respectively.

We also performed three sets of ringdown-only injections. All three sets had the amount of energy radiated as gravitational waves during the ringdown set to 1%. One ringdown injection set was distributed across the template bank uniformly in quality factor between 2.1187-20 and \log_{10} (frequency) between 50-2000 Hz. The other injection sets were distributed uniformly in final black hole mass and spin. The first set had masses between 200-800 M_{\odot} and dimensionless spin parameters between 0.6-0.99. The second set had masses between 50-400 M_{\odot} and dimensionless spin parameters between 0.1-0.99. All sets were distributed uniformly in logarithmic distance between 10 – 10⁶ kpc.




Figure 4.15: Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the component mass plane m_1 - m_2 and in the component spin plane a_1 - a_2 . The color bar represents the value of the $\chi = (m_1 a_1 + m_2 a_2) / (m_1 + m_2)$ parameter in the left plot and the mass ratio in the right plot.




Figure 4.16: Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the final black hole frequency and quality factor plane. The color bar represents the value of the final black hole mass in the left plot and the final black hole spin in the right plot.

In Figure 4.18, we see how the ringdown-only injections were distributed in the final black hole mass and spin plane. The color bar represents the value of the frequency in the left plot and the quality factor in the right plot. In Figures 4.19, we can see how these values




Figure 4.17: Distribution of PhenomB injections used for the tuning of the search pipeline parameters in the final black hole mass and spin plane. The color bar represents the value of the final black hole frequency in the left plot and the final black hole quality factor in the right plot.

translate to injections made in the frequency and quality factor plane.




Figure 4.18: Distribution of ringdown-only injections used for the tuning of the search pipeline parameters in the final black hole mass and spin plane. The color bar represents the value of the frequency in the left plot and the quality factor in the right plot.




Figure 4.19: Distribution of ringdown-only injections used for the tuning of the search pipeline parameters in the final black hole frequency and quality factor plane. The color bar represents the value of the mass in the left plot and the spin in the right plot.

4.3.2 SNR Threshold

In order to compute any type of false alarm probabilities to rank gravitational wave candidates, one needs a good estimation of the background. In the S4 ringdown search, no triple coincident zero-lag events were found and the background of triple coincidences was very low. Preliminary tunings of the ds^2 coincidence window showed that this was also likely to be the case for S5 triple background as well. Thus, in an effort to increase the rate of triple coincidence false alarms and also improve the sensitivity of the search, it was decided to lower the SNR threshold on our least sensitive detector, H2, and also decrease the maximization interval (the window used to cluster triggers over the template bank) from 1 ms to 0.5 ms.

The SNR settings for the S4 ringdown search were based on the settings used by the CBC searches: $\rho_{H1}^* = 5.5$, $\rho_{H2}^* = 5.5$, $\rho_{L1}^* = 5.5$. However, the ρ^* in the CBC search has a slightly different meaning than the ρ^* in the ringdown searches. In the CBC search, the two orthogonal phases of the binary inspiral waveform can be written as a cosine and sine piece

$$h_c(t) = \frac{2}{c^2} \left(\frac{\mu}{M_\odot} \right) [\pi G M f(t)]^{2/3} \cos(2\phi(t) - 2\phi_0), \quad (4.51)$$

$$h_s(t) = \frac{2}{c^2} \left(\frac{\mu}{M_\odot} \right) [\pi G M f(t)]^{2/3} \sin(2\phi(t) - 2\phi_0). \quad (4.52)$$

Then, the normalized filter-output of data $s(t)$ against a circular-polarized template $h(t) = h_c(t) - ih_s(t)$ is expressed by

$$\begin{aligned}
\rho(t) &= \frac{|\langle s, h(t) \rangle|}{\sqrt{\langle h, h \rangle}} \\
&= \frac{|\langle s, h_c(t) \rangle + i\langle s, h_s(t) \rangle|}{\sqrt{\langle h_c - ih_s, h_c - ih_s \rangle}} \\
&= \frac{|\langle s, h_c(t) \rangle^2 + i\langle s, h_s(t) \rangle^2|}{\sqrt{\langle h_c, h_c \rangle \langle h_s, h_s \rangle}} \\
&= \frac{1}{\sqrt{2}} \frac{|\langle s, h_c(t) \rangle^2 + i\langle s, h_s(t) \rangle^2|}{\sqrt{\langle h_c, h_c \rangle}} \\
&= \frac{1}{\sqrt{2}} \rho_{\text{CBC}}^*
\end{aligned} \tag{4.53}$$

Thus, in the CBC searches, it is the quantity $\rho_{\text{CBC}}^* \geq 5.5$, which is what we call the SNR, for which triggers are kept. It is not the direct normalized filter output but rather a factor of $\sqrt{2}$ times the filter output that we threshold on. However, in the ringdown searches, the waveform is

$$h_0(t) = \frac{A}{r} e^{-\frac{\pi ft}{Q}} \cos(2\pi ft). \tag{4.54}$$

The normalized filter output of the data $s(t)$ against $h_0(t)$ is written as

$$\rho_0(t) = \frac{|\langle s, h_0(t) \rangle|}{\sqrt{\langle h_0, h_0 \rangle}}. \tag{4.55}$$

Thus, the SNR threshold for S4 ringdown searches was actually $\rho_0 = \rho_{\text{RD}}^* \geq 5.5$ which is the direct normalized filter output. Thus, there is no reason to expect that the SNR thresholds for the two searches should be the same. In fact, an SNR threshold of $\rho_{\text{CBC}}^* \geq 5.5$ is actually equivalent to a normalized filter output of $\rho_{\text{CBC}} \geq 3.9$. This threshold turned out to be too low to be computationally feasible for ringdown searches, as described below. However, it does indicate that ringdown SNR thresholds could be lower than CBC SNR thresholds.

When conducting a search, we generally do want to see some triple coincident background not just because we need to calculate false alarm rates for gravitational wave candidates but because a search with no triple background indicates that we are not digging deep enough into the noise. In Table 4.1, we can see that several two month runs were set up testing a variety of SNR thresholds with and without the H1H2 consistency test and the maximization interval clustering performed in the coincidence stage as we saw in Section 4.2.5. The most triple coincidences found before coincidence clustering, without yielding too many triggers to process, resulted from a run setting $\rho_{\text{H1}} = \rho_{\text{L1}} = 5.5$ and $\rho_{\text{H2}} = 4.0$ with the maximization interval turned off and the H1H2 consistency test applied. We decided to use these SNR settings although we ultimately chose to turn on the maximization interval and set it to 0.5 ms. This was done to ensure that we would not run into computational problems for particularly glitchy times. Furthermore, we chose not to look at all at H2L1 coincidences in triple time. Not only would the number of these types of triggers shoot through the roof with the lowering of the H2 SNR threshold but any H2 trigger in triple time should be accompanied by an H1 trigger. Thus these types of coincidences are not useful to gravitational wave searches.

Table 4.1: Tests of different SNR thresholds for the purpose of finding the best setting to increase the triple coincidence background. An asterik indicates too many triggers to process.

H1 SNR	L1 SNR	H2 SNR	# H1H2L1	# H1L1/100	# H2L1/100
5.5	5.5	5.5	5	620	480
9	9	4.5	0	0.6	141
6	6	3	26	150	4800
8	6	5	1	18	490
Add the H1H2 consistency test					
9	6	5	0	7	493
5	5	5	21	4032	30
5.5	5.5	4.5	16	631	33
5.5	5.5	5	6	631	9
Turn off the maximization interval					
5.5	5.5	5	59	1241	19
5.5	5.5	4.5	174	1240	67
5	5	4.5	*	*	*
5	5	4	*	*	*
5.5	5.5	4	403	1240	173

4.3.3 Coincidence Windows

The coincidence window sets the threshold value, ds^2 for the 3D coincidence test. Triggers found with a ds^2 value greater than the specified threshold would be removed from the analysis. This threshold proved particularly difficult to tune because the lack of triple coincidences in the runs did not give us enough statistics to discern the correct tuning easily. In general, increasing the ds^2 threshold will improve the efficiency of the search. However, it will also admit more background coincidences. A series of ds^2 values were tested from 0.1-0.6. In the ROC plots in Figure 4.20, we can see that $ds^2 = 0.4$ performs better than $ds^2 = 0.6$ but its performance in comparison to $ds^2 = 0.2$ is not so clear cut since $ds^2 = 0.2$ performs slightly better for PhenomB injections but it admits few triple coincidences; this indicates that it may be too tight of a constraint. Thus, we chose $ds^2 = 0.4$ as our coincidence test threshold.

4.3.4 Detection Statistic

A detection statistic is the value that we use to rank candidate gravitational waves. It must posses several characteristics. It must be a monotonic function. This allows us to simply say that high values are significant while low values are not or vice versa. Furthermore, the values that indicate significance should be correlated with signal-like coincidences while the values that do not indicate significance should be correlated with the coincidences associated




Figure 4.20: ROC plots demonstrating the performance of different ds^2 values for EOBNR and PhenomB injections found in triple coincidence.

with non-Gaussian transient noise in the detectors. It also must assign appropriate values to both double coincident candidates and triple coincident candidates. In other words, in most cases, a triple coincident candidate is more significant than a double coincident candidate so the ranking should reflect that characteristic. Finally, it must simply model the characteristic distributions of signal and noise. We should be able to explain the mathematical expression for the statistic in a straight-forward manner based on our understanding of the physical nature of the noise and signal characteristics.

In Gaussian noise, the optimal statistic described in Equation 4.10 used to calculate an SNR would meet these requirements. The combined SNR ρ_c for triggers in the coincidence could then be calculated as the sum of the SNR's squared:

$$\rho_c = \sqrt{\sum_{i=1}^N \rho_i^2} \quad (4.56)$$

where N is the number of interferometers contributing to the coincidence. However, due to the non-Gaussianity of detector noise, despite vetoes and cuts designed to remove them, a small number of coincident noise triggers will remain at the end of the pipeline. And if ranked by combined SNR alone, these noise coincidences are often the loudest (ie most significant) thing found at the end of the pipeline. Thus, it is imperative that we utilize more information to devise a statistic that will downgrade the significance of these noise coincidences. After the preliminary analyses outlined in Section 4.4.2 of the types of coincidences that were found as loudest at the end of the ringdown pipeline, it was apparent that the detection statistic used in S4 would not be sufficient for S5 data.

In S4, two separate statistics were used to rank double and triple coincidences. Triple coincidences were ranked using the sum of the squares of the SNRs,

$$\rho_{S4_DS,trip} = \rho_{H1}^2 + \rho_{H2}^2 + \rho_{L1}^2. \quad (4.57)$$

Note that this statistic is actually the square of Equation 4.56. This was done in order to ensure that triple coincidences would be louder than double coincidences. Double coincidences were ranked using what is known as a chopped-L statistic:

$$\rho_{S4_DS,doub} = \min \{ \rho_{IFO1} + \rho_{IFO2}, a\rho_{IFO1} + b, a\rho_{IFO2} + b \} \quad (4.58)$$

where a and b are tuned parameters [25]. As outlined in [157], this test prevents a high SNR glitch in one interferometer from unfairly influencing significance. In fact, only a rare source orientation could produce a very loud SNR in one interferometer and very quiet SNR's in the two other interferometers. Thus, those types of coincidences should be downgraded compared to triple coincidence with similar, moderate SNR's in all three interferometers. In S4, the parameters were tuned such that $a = 2.0$ and $b = 2.2$.

The detection statistic is used to rank all types of coincidences. Values are assigned to triple and double coincidences in zero-lag, in timeslides, and in simulation runs. As described in Section 4.2.6, we can compute an uncombined false alarm rate for every zero-lag trigger by counting the number of timeslide triggers with a detection statistic value ρ_{DS} greater than or equal to the detection statistic of the given zero-lag. This allows us to determine the significance of any of the loudest zero-lag coincidences in our search and also, to determine how well we detect different types of simulated signals. In S5, we still use a combined false alarm rate as our *final* detection statistic. However, it is the MVSC likelihood ranking that determines the false alarm ranking. And the MVSC likelihood is constructed from multiple parameters and statistics including the S4 detection statistic, ρ_{S4_DS} and the new ρ_{S5} statistic described below.

4.3.4.1 Triple Coincident Statistic

As discussed in Section 4.4.2, it was determined that we needed a new detection statistic, or at the very least a better statistic to provide to MVSC, that would effectively rank coincidences with a loud SNR in only one interferometer at a low value compared to coincidences with comparable SNR in each interferometer. Essentially, this problem is related to the S4 detection statistic whereby triple coincidences were ranked according to Equation 4.79 and thus a single loud SNR in the triple would give a very large detection statistic ρ_{S4_DS} . Consider, for example, the distribution of triple coincident triggers in the raw SNR-SNR scatter plots in Figures 4.21, 4.22, and 4.23. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. In Figure 4.21, we have plotted the H1 and H2 single detector SNRs of triple coincident events, in Figure 4.22, we have plotted the H1 and L1 single detector SNRs of triple coincident events, and in Figure 4.23, we have plotted the H2 and L1 single detector SNRs of triple coincident events.

Figure 4.22 in the H1L1 SNR plane is the plot we must focus on in order to understand the construction of a new triple ranking statistic. Note that the majority of injections fall along the diagonal while background triggers are clustered at low SNRs in each detector with the exception of a few very loud background triggers in L1 that are quiet in H1. We call this the L1 *tail* of background triggers although these tails for triple background are significantly less pronounced than they are for double background. At the far reaches of the L1 tail, there is an blue star playground trigger with an SNR of 49 in L1 and an SNR of 5.5 in H1. As we discuss in Section 4.4.2, this playground trigger was found as the loudest coincidence in all




Figure 4.21: The H1 and H2 single detector SNRs of triple coincident events. The red pluses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines. The shaded region marks the space where the H1/H2 SNR ratio would be less than one. Any triggers in this region are excluded in S5.

of our preliminary studies with the S4 detection statistic on S5 data. However, its suspicious location along the L1 tail, where background triggers tend to lie and few injections are found, leads us to believe that this loud playground trigger should probably be ranked amongst the other background triggers clustered at low SNRs.

The S4 detection statistic for doubles accomplished just such a feat using a chopped-L statistic that ranks triggers along the tails of the SNR distributions with other triggers clustered at low SNRs while ranking triggers along the diagonal, where injections tend to be found, above these tail triggers. This idea served as motivation for constructing a chopped-L-like statistic for triple coincident triggers that would perform a chopped-L ranking in three 2D SNR planes instead of one. The goals in constructing a new statistic were the following:

- Reweight coincident background triggers that are loud in only one interferometer to a ranking consistent with the rest of the Gaussian distributed background (ie implement a version of chopped-L for triples).
- Keep the values of the new statistic similar to the old sum of the squares values except for pathological triggers.
- Select the statistic that gives us the best improvement in efficiency above the loudest background while still using a model that can be explained via the physical distributions of triggers.




Figure 4.22: The H1 and L1 single detector SNRs of triple coincident events. The red pluses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines.

To address the first item, we designed a statistic that has contours of constant rank shown in Figures 4.21- 4.23. The contours are the same in each SNR plane even though the injection and background distributions are not quite the same (for instance, H2 is less sensitive and has smaller SNRs than the other detectors). However, it is important that the statistic be designed as symmetric; otherwise, we start getting into complicated discussions of the triggers' distributions. The simplest physical model we can find is the best. This triples' statistic that we chose with best improvement in injection-finding efficiency (when combined with an H1H2-consistency cut) has the form:

$$\rho_{S5_DS,trip} = \min \left\{ \sqrt{\rho_{H1}^2 + \rho_{H2}^2 + \rho_{L1}^2}, \rho_{H1} + \rho_{H2} + 0.75, \rho_{H1} + \rho_{L1} + 0.75, \rho_{H2} + \rho_{L1} + 0.75 \right\}^2. \quad (4.59)$$

It can down-weight glitches in the tail of one interferometer significantly. Furthermore, it can account for glitches in the SNR plane of two interferometers, not just one, as long as the difference between the loudest and quietest trigger in the coincidence is significant. For example, an SNR~20 glitch in H1 coincident with an SNR~20 glitch in L1 with a quiet SNR~4.0 in H2 would be down-weighted by a small amount. This becomes more significant as the difference increases. However, by accounting for glitches in the plane, we may also




Figure 4.23: The H2 and L1 single detector SNRs of triple coincident events. The red pluses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Contours of constant S5 triples' statistic are shown as purple lines.

down-weight the rare signal with just two out of three SNRs being outliers. However, this problem cannot be fully avoided without the addition of more information, such as from a signal-based veto. And considering that this statistic will simply be an input to MVSC, it does not have to be a perfect classifier itself. Finally, in other cases where there are not large discrepancies in the SNRs, this new statistic will remain the same as the sum of the squares from S4. In Figures 4.24, 4.25, and 4.26, we show contours of constant background trigger density. These density contours motivated the design of the ranking contours; the density of the background tails of the SNR distributions should be matched with the appropriate background density at low SNRs.

In Figure 4.27, we see a plot of the new S5 triples statistic on the y-axis versus the S4 triples detection statistic on the x-axis. Again, the red pluses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers. Note that the loud blue star playground trigger had an S4 statistic of over 2000 while it is down-weighted significantly with the S5 statistic to a little over 100. The loudest black cross background trigger when ranked with the S4 statistic had a statistic value of ~ 450 while the loudest black cross background trigger when ranked with the S5 statistic had a statistic value of ~ 270 . Thus, the S5 statistic succeeds at reducing the significance of background triggers. Note that, for S4 statistics less than 100, the S5 statistic remains the same. Furthermore, as we move to higher statistic values, the largest value the S5 statistic can take is its prior value of the S4 statistic, hence the injection coincidences along the diagonal of the plot. Note that many injections are also down-weighted with the new S5 statistic. To determine




Figure 4.24: Contours of constant background trigger density in the H1H2 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.

whether this hurts the search efficiency, we have plotted the efficiency of the new statistic in Figure 4.28 and made ROC plots in Figure 4.29 described in more detail below.

Figure 4.28 shows the fraction of injections found louder than the loudest background. In S4, this was the background trigger with S4 statistic of ~ 450 . In S5, this was the background trigger with S5 statistic of ~ 270 . The efficiencies are shown as a function of the injected effective distance in the H1 detector. Large effective distances indicate harder-to-detect injections. There are three efficiency curves in the plot. The black curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S4 statistic. The red curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S5 statistic and when the H1H2 SNR ratio cut was implemented. The brown curve shows the fraction of total found injections; these are found as triple coincidences above the SNR thresholds in each detector regardless of their FAR value. Theoretically, this brown curve is the best injection-finding efficiency that we could have. It would indicate that every found injection had been ranked above the loudest background. Thus, we seek to find the statistic that gives the best injection-finding efficiency at large effective distances. At small effective distances, injections are exceptionally loud and given that a particular ranking statistic is not terrible, we expect the statistics to behave nearly the same in this regime with an efficiency of 1. Indeed, this is what we observe. At large effective distances, we see that the S5 statistic with an H1H2 SNR ratio cut, performs better than the S4 statistic.

Another method of quantifying the improvement of the S5 triples statistic over the S4 statistic is seen in Figure 4.29. This plot is a *receiver operator characteristic* (ROC) plot. These plots are common in signal processing and traditionally show a *true positive* rate on




Figure 4.25: Contours of constant background trigger density in the H1L1 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.

the y-axis and a *false positive* rate on the x-axis. In LIGO searches, the true positive rate would be analogous to a detection probability, or search sensitivity, that could be determined by the detection efficiency of injections. The false positive rate would be analogous to a false alarm probability determined from time slides. Thus, ROC plots show the search sensitivity as a function of false alarm rate which can be plotted by retrieving all injections that are found by the injection-finding code with a FAR less-than or equal-to the given FAR. In Figure 4.29, we have plotted the ROC curve for the S4 statistic in black and for the S5 statistic in red. We see that at low false alarm probabilities, where we are likely to find signals, the S5 statistic has a better detection probability. Thus, we see that the use of the S5 triples' statistic allows for a more sensitive ringdown search.

4.3.4.2 Multi-Variate Statistical Classifier

The fundamental issue we face in gravitational wave data analysis is how to best separate a coincidence caused by a real gravitational wave from background coincidence. We have many tools available for this purpose, although none are perfect. For instance, we discussed how combined SNR can be used in Gaussian noise to separate signal from background. However, since LIGO detector noise is not Gaussian, we discussed the development of a chopped-L statistic for S4 and the development of an even better chopped-L statistic for S5. Furthermore, there are other quantities such as the coincidence test value ds^2 , the timing difference between triggers in the coincidence, the template differences, etc. that may help us to distinguish signal from background. None of these parameters by itself can separate signal from background but perhaps if we examined these parameters in three or more dimensions,




Figure 4.26: Contours of constant background trigger density in the H2L1 SNR plane. These density contours motivated the design of the ranking contours in Figures 4.21- 4.23.

the separation would be more distinct. With a multivariate analysis, we seek to use all (or as many as possible) of these numbers to classify a coincidence as signal or background, essentially a two-class problem. We can train the classifier that software injections are signal-like and time-slides are background-like. Then given a coincidence, which we can think of as a multidimensional vector with many distinguishing parameters, the classifier will tell us whether the vector points to a signal-like or background-like region of multidimensional parameter space. In fact, the classifier can assign a probability to coincidences to span the range from signal to background.

Previous studies of multivariate statistical classifiers for gravitational wave searches have used a random forest of bagged decision trees. In its basic form, a decision tree is composed of a series of binary splits on input data. Given an input vector, for instance a coincidence with multiple associated parameters, the decision tree will direct the vector through different branches by querying whether one of its dimensions is greater than or less than some critical value. Each query marks a node. The final branch, which is called a leaf, determines the classification of the vector.

However, a single decision tree is not suitable for LIGO searches and will lead to the misclassification of many coincidences. Instead, a forest of decision trees is used. This forest is created by the method of bagging, or bootstrap aggregating, which essentially constructs multiple decision trees by repeatedly resampling the training data with replacement so that different training sets can be used to make multiple decision trees. Then, for the classification of a vector, each decision tree in the forest is used and the final leaves are averaged for a final classification.

For example, consider a training set composed of N training cases of signal and back-




Figure 4.27: A scatter plot of the new S5 triples statistic versus the S4 triples detection statistic. The red plusses denote software injections, the black crosses denote time-slide triggers, and the blue stars denote playground triggers.

ground. Let there be M parameters that make up the decisions at each node in the classifier. The decisions themselves are cuts on parameters chosen to best separate signal from background. There are many different algorithms that can be used to decide the best cuts. We have chosen to use the binary information entropy function:

$$H(p) = -p \log_2 p - (p - 1) \log_2 (p - 1) \quad (4.60)$$

where p is the fraction of events in a node that are signals and $q = p - 1$ are the fraction of events in the node that are background. If a particular cut gives $p = q = 1/2$, then $H(1/2) = 1$. However, in the case of a good sort, if a particular cut gives $p = 1$ and $q = 0$ or vice versa, then $H(1) = H(0) = 0$. Thus, we can minimize the binary information entropy function to find the optimal sort. We must decide, by tuning, the number of parameters ($m \ll M$) that will be used to determine the decision at a node. A training set is randomly chosen of n samples from the N available training cases. At each node, we then randomly choose m out of M possible parameters on which to base the decision. When a node meets one of two conditions, it becomes a leaf and is no longer divided. This means that the node is either perfectly sorted (with only signal or only background), or there are fewer events in the node than the user-specified limit, a tunable number called “leaf size.” When there are no more nodes to divide, the decision tree has been created. This is repeated for as many trees as one wants to use in the forest; the number will be limited by computational power and must be tuned.

Finally, when one has a coincidence with M parameters that needs to be ranked as signal-like or background-like, we start at the initial node and follow all the cuts along the tree until




Figure 4.28: The fraction of injections found louder than the loudest background versus injected effective distance in the H1 detector. A zoomed plot is shown on the right. The black curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S4 statistic. The red curve shows the fraction of injections found louder than the loudest background when coincidences were ranked with the S5 statistic and when the H1H2 SNR ratio cut was implemented. The brown curve shows the fraction of total found injections, the theoretical best that we could do.




Figure 4.29: An ROC plot showing the search sensitivity as a function of false alarm rate for the S4 triples' statistic in black and the S5 triples' statistic in red.

the coincidence arrives at a leaf. The probability that the coincidence is a signal is given by

the fraction of signal events in the leaf, p_s . Then, the MVSC rank R_{MVSC} is characterized as the ratio of probabilities of being a signal and background

$$R_{\text{MVSC}} = \frac{p_s}{1 - p_s} \quad (4.61)$$

which is essentially a Bayes-factor which tells us how much one model is favored over another. The coincidence will have a signal probability from its final leaf from each tree in the forest. Its probability of being a signal is then the fraction of signals from all leaves

$$p_{s,\text{forest}} = \frac{\sum_{\text{leaves}} s_i}{\sum_{\text{leaves}} s_i + b_i} = \frac{1}{n} \sum_{\text{leaves}} s_i \quad (4.62)$$

where n is the total number of signals and background in all leaves. Then, the overall MVSC rank is given by

$$R_{\text{forest}} = \frac{p_{s,\text{forest}}}{1 - p_{s,\text{forest}}}. \quad (4.63)$$

For the S5 search, we used the StatPatternRecognition (SPR) package originally developed for high energy physics and used previously for the S6 highmass search. As implemented, MVSC looks only at two pairs of triggers at a time. Thus, it will examine three pairs of doubles for a triple coincidence and just one double for a double coincidence. Then, the MVSC rank of each double in a triple is combined by simple multiplication:

$$R_{\text{triple}} = R_{\text{D1}} \cdot R_{\text{D2}} \cdot R_{\text{D3}}. \quad (4.64)$$

where D1-D3 denote the double pairs. Below, we briefly review the tunings:

- number of trees in forest- increasing the number of trees provides a more stable MVSC rank but increases computation time. We found that using 200 trees gives good performance.
- leaf size- minimum number of events allowed in a node before becoming a leaf. We chose a leaf size of 5.
- number of sampled parameters- number of randomly selected parameters at each node to find the optimal cut. For 20 total parameters, we found that a random selection of 12 performs best.
- optimization criterion- the method used to choose the best cuts at each node. We chose the continuous criteria cross-entropy since it performs fairly well and also does not cause excessive fanning in the IFAR plots.

The 20 parameters that we chose to characterize coincidences with MVSC are listed below. As we develop more statistics, including coherent SNR and a null statistic from the coherent pipeline, we will add these to our list, with the S6 ringdown search in mind. The subscripts a and b denote the two different detectors in one of the pairs of triggers that MVSC ranks:

- dt - the difference between time of arrival (`start_time+10-9start_time_ns`) in detector a and detector b :

$$dt = |t_a - t_b| \quad (4.65)$$

- df - the difference in recovered template frequency in detector a and detector b :

$$df = |f_a - f_b| \quad (4.66)$$

- dQ - the difference in recovered template quality factor in detector a and detector b :

$$dQ = |Q_a - Q_b| \quad (4.67)$$

- ds^2 - the 3D metric distance in (f, Q, t) -space between templates in detector a and detector b . The expression for this metric is provided in Equation 4.34.

- g_{tt} - the t - t template metric coefficient:

$$g_{tt} = \frac{f^2 \pi^2 (1 + 4Q^2)}{Q^2} \quad (4.68)$$

- g_{ff} - the f - f template metric coefficient:

$$g_{ff} = \frac{1 + 6Q^2 + 16Q^4}{4f^2(1 + 2Q^2)} \quad (4.69)$$

- g_{QQ} - the Q - Q template metric coefficient:

$$g_{QQ} = \frac{1 + 28Q^4 + 128Q^6 + 64Q^8}{4Q^2(1 + 6Q^2 + 8Q^4)^2} \quad (4.70)$$

- g_{tf} - the t - f template metric coefficient:

$$g_{tf} = -\frac{2\pi Q(1 + 4Q^2)}{1 + 2Q^2} \quad (4.71)$$

- g_{tQ} - the t - Q template metric coefficient:

$$g_{tQ} = \frac{2\pi f(1 - 2Q^2)}{(1 + 2Q^2)^2} \quad (4.72)$$

- g_{fQ} - the f - Q template metric coefficient:

$$g_{fQ} = -\frac{1 + 2Q^2 + 8Q^4}{2fQ(1 + 6Q^2 + 8Q^4)} \quad (4.73)$$

- ρ_a - the raw SNR as given by Equation 4.11 from detector a .

- ρ_b - the raw SNR as given by Equation 4.11 from detector b .
- $\max\left(\frac{\rho_a}{\rho_b}, \frac{\rho_b}{\rho_a}\right)$ - the ratio of the raw SNRs from detectors a and b where the ratio selected will be greater-than-or-equal-to 1.0.
- ρ^2 - the network SNR described in Section 4.3.4 for N detectors. This is the sum of the squares of the SNRs:

$$\rho^2 = \sum_i^N \rho_i^2 \quad (4.74)$$

- ρ_{S4} - the S4 detection statistic outlined in Equation 4.79 and Equation 4.80:

$$\begin{aligned} \rho_{S4_DS,trip} &= \rho_{H1}^2 + \rho_{H2}^2 + \rho_{L1}^2, \\ \rho_{S4_DS,doub} &= \min\{\rho_{IFO1} + \rho_{IFO2}, 2.0\rho_{IFO1} + 2.2, 2.0\rho_{IFO2} + 2.2\} \end{aligned} \quad (4.75)$$

- ρ_{S5} - the S5 chopped-L statistic outlined in Equation 4.59 with a new expression for triples:

$$\begin{aligned} \rho_{S5_trip} &= \min\left\{\sqrt{\rho_{H1}^2 + \rho_{H2}^2 + \rho_{L1}^2}, \rho_{H1} + \rho_{H2} + 0.75, \rho_{H1} + \rho_{L1} + 0.75, \right. \\ &\quad \left. \rho_{H2} + \rho_{L1} + 0.75\right\}^2, \\ \rho_{S5_doub} &= \min\{\rho_{IFO1} + \rho_{IFO2}, 2.0\rho_{IFO1} + 2.2, 2.0\rho_{IFO2} + 2.2\} \end{aligned} \quad (4.76)$$

- $D_{\text{eff},a}$ - recovered effective distance in detector a .
- $D_{\text{eff},b}$ - recovered effective distance in detector b .
- dD_{eff} - the difference in recovered effective distance in detector a and detector b

$$dD_{\text{eff}} = |D_{\text{eff},a} - D_{\text{eff},b}| \quad (4.77)$$

- $\max\left(\frac{D_{\text{eff},a}}{D_{\text{eff},b}}, \frac{D_{\text{eff},b}}{D_{\text{eff},a}}\right)$ - the ratio of the recovered effective distance from detectors a and b where the ratio selected is the larger of the two.

Figure 4.30 shows an ROC plot demonstrating that MVSC performs better than the new S5 chopped-L-like statistic on two months of S5 data using the 20 input parameters listed above.

One final issue that we needed to resolve in order to use MVSC as a detection statistic was the issue of the ‘‘bump’’ in the IFAR plots seen in the S6 high mass search in Figure 4.31. This bump would seem to indicate a significant deviation from the expected background at low IFAR (or high FAR) values. Essentially, there appears to be an excess of coincidences of




Figure 4.30: An ROC plot demonstrating that MVSC performs better than the new S5 chopped-L-like statistic on two months of S5 data using the 20 input parameters described in Section 4.3.4.2.

low significance in zero-lag. As discussed in Section 4.2.6, we found that this was caused by an injection-finding window that was too large. Thus, oftentimes, random noise coincidences are mapped to injections and MVSC is then trained incorrectly that noise is correlated with injections. Thus, the bump in the IFAR plot is caused by random noise that is given a significance that is too large because MVSC learned that such things were associated with injections. To eliminate such mappings, we implemented the code `ligolw_cbc_dbinjfind` to define a set of “exactly” found injections for the MVSC training set. This exact set was constructed with the following requirements- each found trigger must be within ± 0.01 s of the time the injection was made into that trigger’s detector and if there is a zero-lag trigger within ± 0.5 s of the found injection, the found injection is removed from the training set. Then, we can define another set of injections that are “nearby.” This set only has the requirement that each found trigger must be within ± 1 s of the time the injection was made.

These new injection-finding sets have to be handled carefully within MVSC. Previously, all training events were shuffled, then the first 90% would go in the training set and the remaining 10% would be ranked using that trained forest with rotation of the 10% until all events were ranked on a forest not trained on themselves. Now, both the exact and nearby sets are sorted by GPS time. The exact set is split into 10 chunks and the GPS times of the start and end of these chunks are saved. Then 9 out of the 10 chunks of the exact set are used as a training set to rank injections from the nearby set that do not fall into any of the GPS times from the training set. A rotation of the 9 chunks used ensures that nearby injections will not be ranked with a forest trained on themselves.




Figure 4.31: An example of the MVSC “bump” in the IFAR plots seen in the S6 high mass search for L1V1 coincident time.

4.4 S5 Search Results

In this section, we will first present the results of the playground data set from 14 months of S5. These original playground results were obtained before the new triples statistic or MVSC were implemented. The ranking statistic used was the S4 statistic and as we will see, did not perform very well. In fact, the original playground results highlighted numerous issues with the search and led us to the tuning investigations in Section 4.3. Then, we will have a look at some of the first results obtained from the final pipeline using MVSC as a statistic. The search is ongoing so we do not have results from every month in S5 but we will attempt an order-of-magnitude calculation of the upper limit that we may be able to place on the rate of ringdowns, or more astrophysically motivated, the rate of intermediate mass black hole mergers.

4.4.1 S5 Data

This search uses data from three interferometers, H1, H2, and L1 that were operational during LIGO’s fifth science run. The run took place for approximately two years between November 2005 and October 2007 until ~ 1 year of triple coincidence data was obtained.

4.4.2 S5 Playground Results

When tuning the parameters of an analysis pipeline, we need to be careful not to bias ourselves by examining the output of the pipeline after every new tuning. In order to avoid this, we allow ourselves to look at a small subset of data, known as the playground. Any triggers that have a GPS time t that satisfies $\text{mod}(t - t_{S2} < 600)$. After Category 3 data quality flags are applied, approximately 0.7 years of triple coincidence time remained

$$\text{mod}(t - t_{S2} < 600) \quad (4.78)$$

is considered a playground trigger where t_{S2} is the GPS start time of the second LIGO science run 729273613. The playground times will constitute about 10% of the data. We are blind to the other 90% of the data until tuning has been finalized based on playground, timeslide, and simulation data. At this point, we can run the analysis with the best tunings and “open the box”, or reveal the results of the analysis on the entire data set.

Initial playground runs were completed with the following settings. The template banks were generated for a frequency range of 50-2000 Hz, a quality factor range of 2.1187-20, and a maximum allowed mismatch of 3%. SNR thresholds for H1, H2, and L1 were set to 5.5, 4.0, and 5.5, respectively. We clustered single detector triggers both with a sliding maximize-duration window of 1 s and a maximization-interval over the template bank of 0.5 ms. All coincidence windows were set to $ds^2 = 0.4$. The detection statistic was set to the S4 chopped-L detection statistic:

$$\rho_{DS} = \rho_{H1}^2 + \rho_{H2}^2 + \rho_{L1}^2. \quad (4.79)$$

and

$$\rho_{DS} = \min \{ \rho_{IFO1} + \rho_{IFO2}, a\rho_{IFO1} + b, a\rho_{IFO2} + b \} \quad (4.80)$$

with $a = 2.0$ and $b = 2.2$. A single coincidence clustering window of 10 ms was implemented. All H1H2 coincidences were removed from the analysis and H2L1 coincidences found in triple H1H2L1 time were removed. The H1H2 consistency test to remove H1L1 coincidences in triple time was not applied in this playground run; however, as outlined below, it was determined that this test was needed and thus was used in the full data runs in Section 4.4.3.

The settings of these initial runs were not ideal as will be outlined below. However, the results of these runs indicated many important issues with the settings and motivated new and improved settings and statistics for the full data runs in Section 4.4.3.

4.4.2.1 Issues with Playground Results

Loudest Coincidences

Examining the loudest playground coincidences at the end of the pipeline, it was clear that the S4 chopped-L detection statistic was incorrectly ranking particularly egregious glitch-like coincidences in S5 data as the loudest, most significant candidates. For example, consider the playground results from months 9-10. The uncombined and combined cumulative IFAR plots are shown in Figure 4.32. The cyan triangle (in the left uncombined plot) or the blue triangle (in the right combined plot) with an arrow pointing to the right indicates a triple coincident playground candidate that was louder than any background coincidences in all of months 9-10. This indicates that we have a very significant candidate that has a very low false alarm rate and warrants a full followup.

However, a bit of investigation into the details of this loudest coincidence in months 9-10 revealed that it was not a very good candidate. In Table 4.2, we see that while the time,

Table 4.2: Parameters of loudest triple coincident playground candidate in months 9-10.

CAT 2 cFAR	CAT 3 cFAR	IFOs	GPS time (s)	f (Hz)	Q	SNR
0.0	0.0	H1	839927281.03	152.94	17.95	5.54
		H2	839927281.03	148.83	17.95	4.38
		L1	839927281.04	148.83	17.95	48.88

frequency, and quality factors of the H1, H2, and L1 triggers were similar, the SNR of the L1 trigger was much louder than either H1 or H2 which both had SNRs right above threshold. While it is not impossible for a source to be configured and located in such a way to cause a much louder SNR in L1, it is very unlikely. Furthermore, as discussed in Section 4.3.4.1, this distribution of SNRs (very loud in one interferometer and very quiet in the other two) is actually characteristic of a background coincidence or a false alarm. While it is possible that it could be a very rare configuration, we do not want to rank such coincidences as the loudest in our pipeline. It is likely that when declaring a first detection, we will require the coincidence to have similar SNRs in all three interferometers so that it will be easier to defend. Only later, after we have grown confident in our ability to detect gravitational waves, will we declare detections from rare configurations. Finally, in Figure 4.33, we see a whitened time-frequency spectrogram of the trigger in the H1 gravitational wave channel and in the L1 gravitational wave channel. In the H1 gravitational wave channel, we see an inconspicuous trigger while in L1, we see a very loud suspiciously glitch-like burst of noise. The two interferometers' time-frequency spectrogram are thus inconsistent and the L1 spectrogram is not what we would expect from an inspiral-merger-ringdown gravitational wave. See Figure 4.34 for a time-frequency spectrogram from a inspiral-merger-ringdown hardware injection.

Thus, it was determined that we needed a new detection statistic that would effectively rank such coincidences with a loud SNR in only one interferometer at a low value compared to coincidences with comparable SNR in each interferometer. Essentially, this problem is related to the S4 detection statistic whereby triple coincidences were ranked according to Equation 4.79 and thus a single loud SNR in the triple would give a very large detection statistic $\rho_D S$. The problem and resolution are described in more detail in Section 4.3.4. The S4 search did not encounter this problem since no triple coincident candidates and very little triple coincident background were found and thus, the triples' ranking statistic was never used except for triple found injections. Since we lowered the SNR threshold of H2 in S5, we found a number of triple coincident background and playground triggers. Thus, it was important that we properly rank these coincidences. In Section 4.3.4, we discuss the use of a new triple coincident statistic and the move to a multi-variate statistical classifier as our new detection statistic.

Identically Ranked Coincidences

In other searches, it is very rare to find two (or more) coincidences at the same GPS time with exactly the same ranking statistic. However, because of the ringdown search's ranking statistic, this is actually a fairly common occurrence. The issue is related to the chopped-L test in Equation 4.80. Doubles are ranked with this statistic. However, consider the following case. Say we have two different L1 triggers coincident with the same H1 trigger with an SNR of 30. The first L1 trigger has an SNR of 50. The second L1 trigger has an




Figure 4.32: Months 9-10 playground uncombined and combined cumulative IFAR plots. The cyan triangle in the left plot that falls to the far right with an arrow pointing to the right indicates a triple coincident playground candidate that was louder than any background coincidences in all of months 9-10.




Figure 4.33: Time-frequency whitened spectrograms of the loudest triple coincident playground candidate in months 9-10 in both the H1 and L1 gravitational wave channels.

SNR of 70. The ranking of these two coincidences is the following:

CASE 1

$$\rho_{DS} = \min \{30 + 50, 2 \times 30 + 2.2, 2 \times 50 + 2.2\} = \min \{80, 62.2, 102.2\} = 62.2 \quad (4.81)$$




Figure 4.34: H1 time-frequency whitened spectrogram of the blind injection performed in LIGO and Virgo data on September 16th, 2010. The signal was from a spinning binary system with component masses of $1 M_{\odot}$ and $24 M_{\odot}$. This injection had an SNR of 14 in H1. For more information on this blind injection challenge, please refer to [161].

CASE 2

$$\rho_{DS} = \min \{30 + 70, 2 \times 30 + 2.2, 2 \times 70 + 2.2\} = \min \{100, 62.2, 142.2\} = 62.2 \quad (4.82)$$

Thus, while these are obviously two distinct coincidences, they will both have exactly the same ranking statistic. Normally, clustering by detection statistic will group together any coincidences within a 10 second window of each other. However, in the case that there are multiple “loudest” coincidences with exactly the same “loudest” ranking statistic, clustering is not performed. This issue can be particularly egregious when ranking coincidences associated with injections. Consider, for example, a snippet of H1L1 coincidences associated with ringdown injections in H1L1 time in Table 4.3. Four H1L1 coincidences are associated with the same injection and are given exactly the same ranking!

The main problem regarding this issue is that we will recover more injections than were made. Thus, injection recovery efficiency will be greater than 1 in some cases. Furthermore, due to the new triple ranking statistic outlined in Section 4.3.4 and its similar chopped-L behavior, we expect this redundancy issue to show up for triples as well. Finally, the effects of providing a multi-variate statistical classifier with multiple coincidences associated with the same injection are unknown; it may introduce unknown biases to the training of the classifier.

Thus, in order to fix this issue, we introduced a second clustering stage as outlined to break the degeneracy of identically ranked coincidences.

Table 4.3: Identically ranked H1L1 coincidences associated with one ringdown injection.

cFAR	IFO	Inj GPS time	Rec GPS time (s)	SNR	ρ_{DS}
4.42	H1	827371752.75	827371752.75	15.9	20.93
	L1	827371752.75	827371752.76	9.36	
	H1	827371752.75	827371752.75	20.68	
	L1	827371752.75	827371752.76	9.36	
	H1	827371752.75	827371752.75	20.77	
	L1	827371752.75	827371752.76	9.36	
	H1	827371752.75	827371752.75	19.71	
	L1	827371752.75	827371752.76	9.36	

H1H2 Consistency Test

As discussed in Section 4.2.5, we did not implement an H1H2 consistency cut in the initial playground runs. Recall that this cut relies on the fact that if H1 sees a trigger that H2 does not see, then if H1 is above a certain SNR threshold and H2 is not vetoed, H2 should have seen that trigger as well. Also, we did not use an H1H2 snr ratio cut. Recall that this cut checks that an H2 SNR is not greater than an H1 SNR for a given triple coincidence.

Nevertheless, it was apparent that the application of these cuts would be advantageous for the ringdown search, even improving our detection efficiency. When looking at a table of loudest time slide events in months 3-4 with Category 2 vetoes applied, we found that the third loudest background coincidence in triple time was an H1L1 double coincidence with an H1 SNR of 46.75. If we had required an H1H2 consistency cut, this coincidence would have been removed from the pipeline because H2 should have seen a trigger considering H1 had such a loud SNR. In fact, if we assume that $\rho_{H1} \approx 2\rho_{H2}$, then H2 should have had an SNR of about 23. Instead, H2 saw nothing with an SNR greater than 4.0!

4.4.3 S5 Search Results

The ringdown analysis of S5 split the 24 months into twelve 2-month analysis periods. During the first 18 months of S5, data was collected from three interferometers: H1, H2, and L1. During months 19-24, the Virgo detector also joined the network. However, due to the complications of adding a fourth interferometer to the analysis, the S5 ringdown search was run only on data from H1, H2, and L1. Here we present the results from four analyzed months of S5. The search is ongoing so results from the remaining twenty months of S5 are not yet available. Here we will present results from the box-openings of Months 9-10 (July 12, 2006 to September 13, 2006) and Months 15-16 (January 9, 2007 to March 6, 2007).

The final tuning settings used in these runs were the following. The template banks were generated for a frequency range of 50-2000 Hz, a quality factor range of 2.1187-20, and a maximum allowed mismatch of 3%. SNR thresholds for H1, H2, and L1 were set to 5.5, 4.0, and 5.5, respectively. We clustered single detector triggers both with a sliding maximization window of 1 s and a maximization-interval over the template bank of 0.5 ms. All coincidence windows were set to $ds^2 = 0.4$. An H1H2 consistency cut was applied. We used MVSC as a detection statistic with the input parameters as outlined in Section 4.3.4.2. Two coincidence clustering stages were applied. The first applied a coincidence clustering window of 10 ms and kept only the coincidence with the maximum chopped-L SNR in that

window. It also applied an H1H2 SNR ratio cut and removed all H1H2 coincidences and H2L1 coincidences found in triple H1H2L1 time. The second clustering stage applied a coincidence clustering window of 10 ms and kept only the coincidence with the maximum network SNR in order to remove identically-ranked coincidences. We also performed 12 injection runs for each of these two month segments to be used in the training of MVSC and also to be used for sanity checks. We performed four non-spinning EOBNRv2 runs, four spinning PhenomB runs, and four ringdown-only runs. The configuration files for the search and post-processing are shown in Appendix 6.

4.4.3.1 Months 9-10

In preparation for opening the box, we examine the injection recovery of both IMR and ringdown-only injection sets. In Figure 4.35, we see three sets of plots for EOBNRv2, PhenomB, and Ringdown injections. These are the missed/found injection plots that we discussed in Section 4.2.4. The y-axis marks the second smallest injected distance and the x-axis marks the final black hole mass of the injection. Blue stars are injections found with zero FAR, red crosses are injections that were not found in triple coincidence by the pipeline, and colored circles are injections found with non-zero FAR. In these plots, we hope to see no nearby missed injections except perhaps at the limits of our sensitivity (for example, at high masses for IMR injections or low masses for ringdown injections). Furthermore, we expect a well-defined transition region between missed and found injections. Figure 4.35 demonstrates both of these properties indicating that we recovered the software injections properly and that we can open the box.

In Figure 4.36, we see the uncombined and combined IFAR plots for months 9-10 triple coincident time and a list of the top ten loudest coincidences in Figure 4.37. There was one triple coincident candidate at GPS time 839927281.03 with a combined FAR of 0.45/yr. It had an SNR of 5.54 in H1, 4.38 in H2, and 48.88 in L1. This candidate is in fact, the loudest playground coincidence that we saw in the playground tuning in Section 4.4.2.1 with a glitch-like trigger in L1. However, now we see that MVSC has ranked this candidate so that it has a non-zero FAR; there are time-slide triggers that are louder than this coincidence. Thus, it is not a significant candidate although we will perform follow-up tests on this coincidence to see if we can find an environmental or instrumental cause.

This was the only mildly interesting coincidence in months 9-10. In Figure 4.38, we show IFAR plots for H1L1 coincident time and in Figure 4.39, a list of the loudest coincidences in H1L1 time. Again, there are no statistically significant coincidences in this time. In Figure 4.40, we show IFAR plots for H2L1 coincident time and in Figure 4.41, a list of the loudest coincidences in H2L1 time. There are no statistically significant coincidences in this time, either.

4.4.3.2 Months 15-16

In Figure 4.42, we see three sets of plots for EOBNRv2, PhenomB, and Ringdown injections in Months 15-16. No obvious issues exist; injections are recovered properly.

In Figure 4.43, we see the uncombined and combined IFAR plots for months 15-16 triple coincident time and a list of the top ten loudest coincidences in Figure 4.44. In Figure 4.45, we show IFAR plots for H1L1 coincident time and in Figure 4.46, a list of the loudest coincidences in H1L1 time. In Figure 4.47, we show IFAR plots for H2L1 coincident time and in Figure 4.48, a list of the loudest coincidences in H2L1 time. There are no statistically significant coincidences in any of these times.




Figure 4.35: Missed/found plots for EOBNRv2, PhenomB, and Ringdown injections in Months 9-10.




Figure 4.36: IFAR plot of H1H2L1 time in Months 9-10.

4.4.3.3 Preliminary Rate Upper Limit Calculations

In the event that we do not have a detection in the rest of S5, we can calculate an upper limit on the rate of ringdown injections. Here we attempt a rough upper limit calculation and do a preliminary comparison with the S4 upper limit [26]. Since we do not have all the data yet from the analysis with MVSC, we cannot compare the performance of the final pipeline. However, we can do a preliminary comparison using data from the original playground runs and the new S5 triples' chopped-L statistic. To compare, we will look only at ringdown injections in the 70-140 Hz band (or equivalently, the 85-390 M_{\odot} band) with distances between 0.5-1000 Mpc. In S4, no triple background was found so all injections above threshold were counted. In S5, for the 16 months analyzed with the new chopped-L triples statistic, the loudest event was in the background with a chopped-L statistic of ~ 160 so all injections above this threshold were counted.

Figure 4.49 shows the efficiency curve for triple found ringdown injections in S5. We injected 12,515 signals in total and 9047 were recovered after Category 3 vetoes were applied. The playground runs were only completed for 14 months of S5; this amounted to a total triple analysis livetime of ~ 0.47 years after Category 3 vetoes. Figure 4.50 shows the efficiency curve for triple found ringdown injections in S5. In S4, 5701 signals were injected and 3010 were recovered. The total triple analysis livetime in S4 was only 0.0375 years.

Comparing Figures 4.49 and 4.50, we see that signals in S5 were detected out to nearly 1000 Mpc, almost an order of magnitude further than in S4. Also, note that the efficiency in S4 is not equal to 1 at nearby distances. This is because applied vetoes were treated as a loss of efficiency rather than a loss of livetime. In S5, we treated vetoes as a loss of livetime. Thus, the S5 efficiency does go to 1 at nearby distances. We can use these efficiency ϵ as a function of distance D plots to calculate the effective volume to which we are sensitive to

$$V_{\text{eff}} = 4\pi \int \epsilon(r)r^2 dr. \quad (4.83)$$

In S4, V_{eff} was found to be 2.6×10^6 Mpc³. In S5, we find a preliminary value of

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	0.45	0.0297	0.36	839927281.03	H1	H1,H2,L1	152.94	17.95	194.22	0.99	5.54	24.65
				839927281.03	H2		148.83	17.95	199.59	0.99	4.38	
				839927281.04	L1		148.83	17.95	199.59	0.99	48.88	
2	16.25	0.67	1.0	837326922.99	H1	H1,H2,L1	92.35	17.95	321.66	0.99	8.1	24.65
				837326922.98	L1		93.39	6.24	246.45	0.93	10.87	
3	25.19	0.82	1.0	838171339.59	H1	H1,H2,L1	66.58	17.95	446.18	0.99	8.03	24.65
				838171339.58	L1		65.67	17.95	452.3	0.99	8.2	
4	28.62	0.86	1.0	839804954.67	H1	H1,H2,L1	65.95	10.6	404.45	0.98	6.26	24.65
				839804954.67	L1		68.33	6.24	336.81	0.93	9.8	
5	43.82	0.95	1.0	837207088.94	H1	H1,H2,L1	109.17	6.24	210.82	0.93	8.27	24.65
				837207088.93	L1		96.73	3.66	190.07	0.77	5.78	
6	46.65	0.96	1.0	839251404.59	H1	H1,H2,L1	264.55	2.12	44.98	0.0	6.66	24.65
				839251404.59	L1		295.63	2.12	40.26	0.0	7.49	
				840103823.97	H1		73.88	6.24	311.51	0.93	9.46	
				840103823.97	L1		79.36	3.66	231.68	0.77	6.38	
8	55.59	0.98	1.0	839781979.4	H1	H1,H2,L1	319.37	17.95	93.01	0.99	6.11	24.65
				839781979.4	L1		328.19	17.95	90.51	0.99	8.81	
9	61.56	0.98	1.0	837711642.01	H1	H1,H2,L1	551.14	10.6	48.4	0.98	8.41	24.65
				837711642.01	L1		564.0	10.6	47.3	0.98	5.88	
10	90.32	1.0	1.0	839785111.3	H1	H1,H2,L1	101.6	17.95	292.38	0.99	6.61	24.65
				839785111.3	L1		100.97	6.24	227.94	0.93	53.32	

Figure 4.37: List of loudest coincidences in H1H2L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.

$V_{\text{eff}} = 4 \times 10^9 \text{ Mpc}^3$, three orders of magnitude better than S4. These volumes are then used to compute the 90% confidence upper limit on the rate, an expression derived in Reference ???. The expression is

$$R_{90\%} = \frac{2.303}{TV_{\text{eff}}} \quad (4.84)$$

where T is the analysis livetime. In S4, the upper limit on the rate was found to be 2.4×10^{-5} per year per Mpc^3 . In S5, using a livetime of just 14 months of S5, we find a preliminary upper limit of 1×10^{-9} per year per Mpc^3 . Using the average of the livetimes from each of




Figure 4.38: IFAR plot of H1L1 time in Months 9-10.

the 14 months, we can extrapolate the triple analysis livetime for all 24 months to be about 0.7 years. This would yield a rate upper limit of 9×10^{-10} per year per Mpc^3 .

We can also attempt to make a rough comparison with the rate density upper limit set in Reference [48] for IMBH-IMBH mergers. In this reference, a set of non-spinning EOBNR injections were made with mass ratios less than 4:1 and total masses between $100\text{-}450 M_{\odot}$. This search also included H1H2L1V1 time in the last 6 months of S5. The upper limit for the combined search, averaged over all masses, was estimated to be 0.9 per Myr per Mpc^3 . The most sensitive mass bin ($88 + 88 M_{\odot}$) had a rate limit of 1.3×10^{-7} per yr per Mpc^3 .

We cannot make a direct comparison with the IMBH search because we have used spinning PhenomB injections in our playground runs. One of the PhenomB runs had non-spinning injections but the other two had spinning injections with component spins between 0.1-0.85. Selecting only those with mass ratios less than 4:1 and total masses between $100\text{-}200 M_{\odot}$ (our most efficiently detected bin), we find a total of 11,054 were made and 5784 of those were detected. The V_{eff} calculated from the efficiency curve in Figure 4.51 is similar to that for ringdown injections, $V_{\text{eff}} = 4 \times 10^9 \text{ Mpc}^3$. Thus, the rate upper limit for all of S5 is the same order of magnitude, 9×10^{-10} per year per Mpc^3 .

Before we can make firm statements on the rate upper limits of the ringdown search, we will need to determine the total uncertainty in the effective ranges and the systematic uncertainty due to the accuracy of simulated waveforms and calibration of the data. We will also need to make informed decisions about which types of injection sets will give the most meaning full rate upper limits. With regards to astrophysical rates of IMBH-IMBH coalescences, Reference [19] estimates 0.07 per GC per Gyr where GC is globular cluster. From Reference [162], we have an approximate astrophysical source density of 0.3 GC per Mpc^3 . This translates to about 2.1×10^{-11} per year per Mpc^3 . This is only about an order of magnitude from our rough estimated rates. With the addition of the search through S6 data and using MVSC as a detection statistic, the ringdown search may be able to limit the astrophysical rate of IMBH-IMBH mergers.

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	675.31	0.6	1.0	837405907.17	H1	H1,L1	86.37	6.24	266.47	0.93	6.14	0.5
				837405907.16	L1		77.97	2.12	152.64	0.0	7.2	
2	984.83	0.74	1.0	837385456.11	H1	H1,L1	53.53	17.95	554.95	0.99	5.67	0.5
				837385456.1	L1		52.09	17.95	570.29	0.99	15.11	
3	2215.86	0.95	1.0	837386116.73	H1	H1,L1	271.16	17.95	109.55	0.99	5.6	0.5
				837386116.73	L1		267.49	17.95	111.05	0.99	5.87	
4	2637.93	0.97	1.0	837865658.47	H1	H1,L1	109.58	10.6	243.43	0.98	5.75	0.5
				837865658.46	L1		105.84	17.95	280.66	0.99	6.98	
5	2940.41	0.98	1.0	838109564.46	H1	H1,L1	110.26	17.95	269.41	0.99	5.55	0.5
				838109564.45	L1		107.08	10.6	249.12	0.98	8.35	
6	3840.83	0.99	1.0	839302506.37	H1	H1,L1	779.09	10.6	34.24	0.98	5.8	0.5
				839302506.38	L1		799.13	3.66	23.01	0.77	5.85	
				840195217.5	H1		346.59	17.95	85.71	0.99	5.59	
				840195217.51	L1		341.9	17.95	86.88	0.99	5.95	
				837990145.74	H1		1051.66	10.6	25.36	0.98	5.6	
				837990145.75	L1		1074.73	17.95	27.64	0.99	6.3	

Figure 4.39: List of loudest coincidences in H1L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.




Figure 4.40: IFAR plot of H2L1 time in Months 9-10.

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	108.75	0.56	1.0	839065780.85	H2	H2,L1	55.76	17.95	532.71	0.99	4.67	2.79
				839065780.85	L1		56.21	6.24	409.42	0.93	15.17	
2	166.4	0.72	1.0	837080071.97	H2	H2,L1	50.0	10.6	533.5	0.98	5.8	2.79
				837080071.98	L1		50.0	10.6	533.5	0.98	6.08	
3	180.81	0.75	1.0	836958016.95	H2	H2,L1	151.82	2.12	78.39	0.0	5.07	2.79
				836958016.95	L1		134.54	3.66	136.65	0.77	14.26	
4	276.46	0.88	1.0	836966823.11	H2	H2,L1	108.77	17.95	273.11	0.99	5.43	2.79
				836966823.1	L1		109.17	6.24	210.82	0.93	21.49	
5	651.19	0.99	1.0	837511413.74	H2	H2,L1	153.52	3.66	119.76	0.77	5.04	2.79
				837511413.74	L1		143.71	3.66	127.93	0.77	11.29	
6	862.14	1.0	1.0	840323151.55	H2	H2,L1	149.2	6.24	154.26	0.93	5.4	2.79
				840323151.56	L1		135.86	2.12	87.6	0.0	6.84	
7	1086.19	1.0	1.0	837074176.18	H2	H2,L1	319.37	17.95	93.01	0.99	5.48	2.79
				837074176.19	L1		315.04	17.95	94.29	0.99	8.76	
8	1377.07	1.0	1.0	841029962.23	H2	H2,L1	103.33	3.66	177.93	0.77	4.5	2.79
				841029962.24	L1		121.58	2.12	97.89	0.0	6.65	
9	1420.31	1.0	1.0	840341971.28	H2	H2,L1	234.66	10.6	113.67	0.98	4.47	2.79
				840341971.28	L1		240.14	10.6	111.08	0.98	6.68	
10	1539.54	1.0	1.0	838500653.94	H2	H2,L1	412.52	2.12	28.85	0.0	6.22	2.79
				838500653.94	L1		369.16	2.12	32.24	0.0	6.05	

Figure 4.41: List of loudest coincidences in H2L1 time in Months 9-10. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.




Figure 4.42: Missed/found plots for EOBNRv2, PhenomB, and Ringdown injections in Months 15-16.




Figure 4.43: IFAR plot of H1H2L1 time in Months 15-16.

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	1.8	0.11	0.83	852983295.55	H1	H1,H2,L1	182.06	10.6	146.52	0.98	11.51	24.53
				852983295.56	L1		175.28	17.95	169.47	0.99	8.0	
2	13.47	0.6	1.0	856895791.65	H1	H1,H2,L1	63.2	6.24	364.17	0.93	6.08	24.53
				856895791.65	H2		64.79	17.95	458.51	0.99	5.35	
				856895791.66	L1		65.72	6.24	350.22	0.93	5.92	
3	15.42	0.64	1.0	853490637.42	H1	H1,H2,L1	71.27	17.95	416.77	0.99	5.88	24.53
				853490637.42	H2		72.25	17.95	411.13	0.99	5.22	
				853490637.42	L1		72.25	17.95	411.13	0.99	6.07	
4	16.17	0.66	1.0	856925804.36	H1	H1,H2,L1	55.87	2.12	212.99	0.0	14.58	24.53
				856925804.38	L1		62.44	2.12	190.6	0.0	7.27	
5	20.96	0.76	1.0	856470386.62	H1	H1,H2,L1	110.37	3.66	166.57	0.77	9.68	24.53
				856470386.62	L1		109.17	6.24	210.82	0.93	6.84	
6	24.85	0.81	1.0	856375703.36	H1	H1,H2,L1	52.09	17.95	570.29	0.99	8.07	24.53
				856375703.37	L1		52.8	17.95	562.57	0.99	8.8	
7	36.83	0.92	1.0	853356347.07	H1	H1,H2,L1	66.58	17.95	446.18	0.99	8.23	24.53
				853356347.08	L1		67.49	17.95	440.14	0.99	10.1	
8	39.97	0.93	1.0	853763848.08	H1	H1,H2,L1	134.54	3.66	136.65	0.77	6.54	24.53
				853763848.08	L1		138.02	10.6	193.27	0.98	7.24	
9	41.77	0.94	1.0	856688343.98	H1	H1,H2,L1	77.52	10.6	344.12	0.98	7.94	24.53
				856688343.98	L1		75.75	10.6	352.16	0.98	7.21	
10	43.11	0.94	1.0	852487430.11	H1	H1,H2,L1	338.89	3.66	54.25	0.77	6.17	24.53
				852487430.11	L1		295.63	2.12	40.26	0.0	7.02	

Figure 4.44: List of loudest coincidences in H1H2L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.




Figure 4.45: IFAR plot of H1L1 time in Months 15-16.

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	52.06	0.24	1.0	853028265.43	H1	H1,L1	137.99	6.24	166.79	0.93	13.22	1.9
				853028265.43	L1		135.28	17.95	219.58	0.99	6.07	
2	260.3	0.74	1.0	853036806.73	H1	H1,L1	117.9	3.66	155.93	0.77	26.34	1.9
				853036806.74	L1		122.98	10.6	216.91	0.98	5.7	
3	563.02	0.95	1.0	853031878.26	H1	H1,L1	97.36	2.12	122.23	0.0	7.11	1.9
				853031878.26	L1		100.97	6.24	227.94	0.93	6.02	
4	670.99	0.97	1.0	853349312.11	H1	H1,L1	911.87	3.66	20.16	0.77	6.04	1.9
				853349312.11	L1		803.27	2.12	14.82	0.0	5.87	
5	746.19	0.98	1.0	852505158.33	H1	H1,L1	108.8	2.12	109.38	0.0	5.85	1.9
				852505158.33	L1		96.73	3.66	190.07	0.77	5.51	
6	800.18	0.98	1.0	856795521.32	H1	H1,L1	127.62	6.24	180.33	0.93	5.8	1.9
				856795521.32	L1		132.71	6.24	173.43	0.93	5.78	
7	1097.11	1.0	1.0	853139337.5	H1	H1,L1	324.16	10.6	82.29	0.98	5.6	1.9
				853139337.51	L1		325.76	6.24	70.65	0.93	6.9	
8	1316.92	1.0	1.0	855887442.71	H1	H1,L1	91.11	10.6	292.79	0.98	5.64	1.9
				855887442.7	L1		93.39	6.24	246.45	0.93	13.53	
9	1351.62	1.0	1.0	856502144.8	H1	H1,L1	93.39	6.24	246.45	0.93	6.62	1.9
				856502144.8	L1		89.03	10.6	299.63	0.98	6.58	
10	1554.08	1.0	1.0	856106712.21	H1	H1,L1	62.19	17.95	477.66	0.99	6.2	1.9
				856106712.21	L1		60.78	6.24	378.67	0.93	7.06	

Figure 4.46: List of loudest coincidences in H1L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.




Figure 4.47: IFAR plot of H2L1 time in Months 15-16.

loudest_events												
rank (using combined far)	combined far	fap	fap 1yr	sngl start time	sngl ifo (click for elog)	instruments on	sngl frequency	sngl quality	sngl mass	sngl spin	sngl snr	all data duration (days)
1	398.37	0.63	1.0	855950629.4	H2	H2,L1	174.42	6.24	131.95	0.93	5.37	0.91
				855950629.4	L1		169.65	2.12	70.15	0.0	8.78	
2	864.46	0.89	1.0	852926306.0	H2	H2,L1	189.58	2.12	62.77	0.0	4.76	0.91
				852926305.99	L1		151.82	2.12	78.39	0.0	13.14	
3	1446.08	0.97	1.0	855675517.49	H2	H2,L1	50.0	2.12	238.01	0.0	4.99	0.91
				855675517.49	L1		51.99	6.24	442.67	0.93	8.15	
4	1565.59	0.98	1.0	857214890.78	H2	H2,L1	97.36	2.12	122.23	0.0	6.03	0.91
				857214890.79	L1		117.9	3.66	155.93	0.77	6.65	
5	1669.17	0.98	1.0	853944578.86	H2	H2,L1	79.49	17.95	373.71	0.99	5.7	0.91
				853944578.87	L1		76.82	6.24	299.58	0.93	7.57	
6	1693.07	0.99	1.0	852924191.85	H2	H2,L1	88.65	17.95	335.09	0.99	11.58	0.91
				852924191.86	L1		89.87	17.95	330.55	0.99	6.39	
				856058890.47	H2		86.26	17.95	344.35	0.99	5.0	
				856058890.48	L1		83.94	17.95	353.87	0.99	14.31	
8	2079.49	0.99	1.0	852924431.76	H2	H2,L1	180.13	17.95	164.91	0.99	4.93	0.91
				852924431.76	L1		177.69	17.95	167.17	0.99	9.95	
9	2681.02	1.0	1.0	856058902.7	H2	H2,L1	77.97	2.12	152.64	0.0	4.97	0.91
				856058902.7	L1		76.82	6.24	299.58	0.93	8.33	
10	3051.51	1.0	1.0	853276567.44	H2	H2,L1	296.99	3.66	61.9	0.77	5.09	0.91
				853276567.44	L1		295.63	2.12	40.26	0.0	6.08	

Figure 4.48: List of loudest coincidences in H2L1 time in Months 15-16. The columns in this table are rank with rank of 1 being the most significant, combined false alarm rate, false alarm probability, false alarm probability in one year of data, recovered GPS start time in each interferometer, recovered interferometer, instruments operating at the time, recovered frequency, quality factor, mass, and spin in each interferometer, and duration of this type of coincident time.




Figure 4.49: The efficiency of detecting ringdown injections in triple coincidence in the 70-140 Hz band in S5 using the new triples chopped-L statistic.




Figure 4.50: The efficiency as a function of injected distance of ringdown injections in triple coincidence in the 70-140 Hz band in S4. Taken from Reference [26].




Figure 4.51: The efficiency as a function of injected effective distance of PhenomB injections in triple coincidence with mass ratios less than 4:1 and total masses between 100-200 M_{\odot} in S5.

4.5 Parameter Recovery

In this section, we examine the ability of the S5 ringdown pipeline to recover IMR signals with the correct template. This issue is different from the issue of the efficiency of the search which directly affects our ability to detect. We could detect a signal with the wrong template but it would have still been detected. However, if we are interested in understanding the physics behind the source, then we hope that we can accurately recover the signal with the correct template. In other words, we hope that the template that is closest to the signal's actual frequency and quality factor will give the largest SNR when matched filtering the data. As we will see, this is not always the case for the present ringdown search. In addition to some systematic errors in parameter recovery which we will explore in more detail below, we also encounter the same issue explored in Section 4.3.4.2 that plagued the creation of the MVSC training set. With an injection-finding window of 1 s, it is very possible that a loud noise trigger could occur in that window and get picked up as the recovered injection. And it is unlikely that the noise trigger would have the same parameters as the signal. Furthermore, we are interested to see whether recovered templates are consistent across the various interferometers. For example, for a particular signal, do we expect to find the same template in both H1 and L1?

We will explore the recovery of the ringdown part of an IMR signal from set of PhenomenologicalB waveforms first discussed in Section 3.5.1 and how well we can estimate the physical mass and spin parameters. We will try to understand how the addition of inspiral and merger affect the physical parameter recovery. We will focus only on injections found as triple coincidences. Thus, injections made into double time or found only as double coincidences will not be explored here.

4.5.1 IMR Parameter Recovery

Previous studies of the parameter recovery abilities of the ringdown search pipeline have focused on the recovery of ringdown-only waveforms [26]. However, we now have the ability to inject full inspiral-merger-ringdown simulated signals into the data and test the ringdown pipeline's parameter recovery ability. We seek to understand whether the addition of an inspiral and merger would help or hurt the search efficiency.

In this study, we injected 61,411 PhenomB injections spread over 12 months of S5 data. The distributions included the same mass, spin, and position settings as outlined in Section 4.3.1 for the tuning of the pipeline. This included a mass distribution distributed uniformly in mass ratio q from 0.1-1.0 and total mass from 50-450 M_{\odot} . One third of the injections were nonspinning while the rest were distributed uniformly with component dimensionless spin parameter from 0.1-0.85. The spins were only aligned or anti-aligned with the orbital angular momentum. After collecting only those recovered in triple time after Category 3 vetoes, we found 29,382 of the original 61,411 injected PhenomB injections.

In Figure 4.52 and Figure 4.53, we can see some general patterns of the recovery of these PhenomB injections. In these plots, all injections that were made are shown in red and those that were recovered in triple coincidence are overlain in green. In Figure 4.52 we see the distribution in the component mass m_1 - m_2 plane on the left and in the component spin a_1 - a_2 plane on the right. In the m_1 - m_2 plane, we note that the density of recovered injections for systems with $q < 1$ decreases as q becomes more extreme. This is an effect that we expect based on studies of the amplitude of the $l = m = 2$ mode of unequal mass non-spinning black holes with decreasing mass ratio. As the mass ratio becomes more extreme, studies show that the amplitude of the $l = m = 2$ mode decreases, making the ringdown waveform for these types of systems more difficult to detect. In the a_1 - a_2 plane, we notice no obvious gaps in detectability. In general, the amount of radiation emitted during the ringdown will

depend on the magnitude and inclination of spins of the initial objects. However, since the initial spins of PhenomB injections must always be aligned or anti-aligned with the orbital angular momentum, we cannot explore the full effects of initial spins on ringdown detectability. Furthermore, subtle differences in detectability are not obvious in the a_1 - a_2 plane. In Figure 4.53, we see histograms of all made PhenomB injections in red and those recovered in triple time in green. Again, the histograms of mass ratio show that for an injected uniform q distribution, the detected q distribution is not quite uniform with fewer injections found at low mass ratios. The histogram of final black hole spin a_{final} demonstrates a few things. Despite the uniform distribution of aligned or anti-aligned initial spins in the a_1 - a_2 plane, a_{final} tends toward 0.7. Furthermore, the detected distribution of a_{final} shows that in the range of 0.5-0.7, we do not detect as large a fraction of injections as in the range from 0.7-0.9. This is also an expected effect from what is known as the orbital hangup.




Figure 4.52: PhenomB injection distribution in the m_1 - m_2 plane on the left and in the spin1z-spin2z plane on the right. All injections made are shown in red and those that were recovered in triple coincidence are overlain in green.

4.5.1.1 Frequency and Quality Factor

Examining injected versus detected frequency and quality factor distributions for the different interferometers will allow us to determine whether IMR injections are recovered with the correct template. In Figure 4.54, we see injected versus detected frequency for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0 and the black line marks the diagonal. Here we see that below about 150 Hz, f_{INJ} and f_{DET} are fairly consistent. But this is not the case above 150 Hz where f_{DET} is consistently found to be less than f_{INJ} . Here we find that IMR injections are being recovered with templates with frequencies in the 100-250 Hz band. This is a phenomenon that has been explained in Goggin, et.al. For an IMR signal, as we increase the ringdown frequency, more of the inspiral and merger part of the signal come into LIGO's sensitivity band. As was demonstrated in Goggin, et. al, most of the ringdown templates in the bank will contribute to a trigger in 120 ms around the IMR signal if the inspiral and merger are




Figure 4.53: Histogram of PhenomB injection distributions by mass ratio in the plot on the left and by final black hole spin in the plot on the right. The distributions of all injections are shown in red and the distributions of those recovered in triple coincidence are overlain in green.

visible. The templates' frequency tends to match the frequency of the chirp or merger part of the waveform that is contributing to the signal. And it is often the case that the template that contributes to the loudest SNR trigger is the template with frequency near 100-200 Hz. This is where the LIGO strain sensitivity is best.

Another characteristic related to mass ratio q can be seen in Figure 4.54. The colorbar on the plots denotes mass ratios in the range from 0.1-1.0. Above 150 Hz, where f_{DET} is consistently found to be less than f_{INJ} , we note the general trend that $q \sim 1.0$ injections are found a bit closer to the diagonal than injections with $q \sim 0.1$. This is expected because the amplitude of the $l = m = 2$ mode of the ringdown waveform decreases as the mass ratio of the binary components becomes more extreme. Since this is the mode that the ringdown templates are designed to find, unequal mass systems will suffer a loss of SNR. For a given f_{INJ} , we expect that the ringdown of an equal mass system will contribute more to the SNR than that of an unequal mass system. Thus the f_{DET} of an equal mass system above 150 Hz will shift to slightly higher frequencies in LIGO's most sensitive frequency band.

In Figure 4.55, we see injected versus detected quality factor values for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0 and the black line marks the diagonal. One interesting feature to note is that the PhenomB injections are injected with continuous quality factor values from about 2.1187-6. However, due to the discrete nature of the ringdown template bank, we only detect injections with five possible quality factor values. And the detected values certainly do not lie along the diagonal. In fact, some of the detected injections were found to have quality factor values of ~ 18 , much larger than the maximum injected quality factor of about 6. This is possibly due to the fact that f_{DET} is consistently found to be less than f_{INJ} . If the ringdown template finds a best match at too low of frequency, it will essentially have found a ringdown signal that starts too soon. Thus, the detected ringdown signal will be longer than the injected ringdown signal. Since the quality factor is a measure of how quickly energy is

lost for a system, if the signal is longer than it should be, then the detected quality factor will be larger than it should be, indicating that the signal has a long damping time.

Despite the fact that the ringdown pipeline finds IMR signals above ~ 150 Hz with f_{DET} consistently less than f_{INJ} , we can see in Figure 4.56 that f_{DET} is remarkably similar across interferometers. These plots show f_{DET} for both H1 and H2 on the left and for H1 and L1 on the right. Thus, even though the search picks up the wrong frequency, it does so consistently. We can then expect that bad frequency recovery does not affect the coincidence test. In Figure 4.57, we again see the discrete nature of Q_{DET} .

4.5.1.2 Time Difference

Another indicator of parameter recovery performance is the difference in detected GPS time and injected GPS time of the start of the ringdown of the IMR injection. We can examine the time difference for each interferometer and also between two different interferometers. In Figure 4.58, we have plotted three simple histograms showing the difference between detected and injected start times in each of the interferometers H1, H2, and L1. All three interferometers demonstrate remarkably similar performance, with distributions skewed toward negative Δt . This indicates that the detected GPS time is often found to be smaller than the injected GPS time. In other words, the ringdown template finds the IMR signal's ringdown too soon. This is consistent with the observation that f_{DET} is consistently less than f_{INJ} and Q_{DET} consistently larger than Q_{INJ} . Essentially, the ringdown template will pick up the part of the IMR signal that falls in LIGO's most sensitive band, 100-200 Hz. For IMR signals with ringdown frequency above this band, this means that the template will find the signal too soon, before the actual ringdown begins, at too low of frequency and too large of quality factor. In Figure 4.59, we see detected minus injected time difference as a function of frequency for each interferometer. Largest time differences occur at low frequency and decrease with increasing frequency. This same effect is observed for ringdown injections and was shown to be caused by random phase ringdowns being filtered with templates with a phase of zero. We expect a similar effect for IMR injections that will also have a random phase ringdown. However, above ~ 150 Hz, the templates will match whichever part of the signal falls in the $\sim 100 - 200$ Hz frequency band so time difference will not be random but rather have a peak near zero with a skew toward negative differences. Finally, we see that higher final black hole spin systems are shifted to positive time differences at low frequency. For the very highest spins near $a_{\text{INJ}} = 1$, we even see a separate population emerging. This implies that the detected ringdown start time was significantly later than injected ringdown start time for these high spin injections. We know that the very high final black hole spin injections will have a large quality factor and will have significantly more cycles than non-spinning final black hole injections. This population is due to the template picking up the ringdown signal several cycles late. For example, a ringdown at ~ 150 Hz will take about 13 ms to go through two cycles. In Figure 4.59, we can see that +13 ms is approximately where the second population lies.

In Figure 4.60, 4.61, and 4.62, we explore a different type of time difference between recovered times in different interferometers for triple found PhenomB injections. Figure 4.60 shows the difference in recovered time between H1 and H2 interferometers. This is a sharply peaked distribution about zero as we expect given that the interferometers are collocated. However, in Figure 4.61, we see the difference in recovered time between H1 and L1 interferometers is a much broader distribution where the light travel time of 10 ms between detector sites can be seen. Figure 4.62 shows these time differences as a function of H1 injected frequency for H1 and H2 in the plot on the left and for H1 and L1 in the plot on the right. In Δt_{H1H2} , there is an obvious discrete structure; the lines are separated by the inverse of the sampling rate (1/8192) s. We also see the spread at low frequencies that we saw in the




Figure 4.54: Injected versus detected frequency for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.




Figure 4.55: Injected versus detected quality factor for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.




Figure 4.56: Differences in the recovered frequency for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal.




Figure 4.57: Differences in the recovered quality factor for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal. The discreteness of the recovered values is due to the discreteness of quality factor values available in the ringdown template bank.




Figure 4.58: Histogram of detected minus injected GPS start times for triple found PhenomB injections in H1, H2, and L1.




Figure 4.59: Difference in GPS start time for triple found PhenomB injections as a function of injected frequency for H1, H2, and L1. The color bar denotes injected dimensionless spin.




Figure 4.60: Histogram of difference in GPS start time for triple found PhenomB injections in H1 and H2.

detected versus injected time difference plots. In Δt_{H1L1} , we see a much larger spread in time difference over the whole injected frequency range.

4.5.1.3 Mass and Spin

Based on observations of frequency, quality factor, and time recovery of the ringdown part of IMR injections, we can begin to understand the biases in physical parameter recovery. For instance, consider Figure 4.63, where we have shown M_{DET} versus M_{INJ} for H1, H2, and L1. Consistently, M_{DET} is larger than M_{INJ} . We know that above 150 Hz, f_{DET} is consistently found to be less than f_{INJ} . Since M_{DET} is inversely proportional to f_{DET} , this will act to overestimate M_{DET} . Furthermore, since Q_{DET} is consistently found to be more than Q_{INJ} and since M_{DET} is proportional to Q_{DET} , this will further act to overestimate M_{DET} . Thus the final black hole mass recovered by the ringdown pipeline will be too large. We can also observe that the extent of this bias depends on mass ratio. Again, we see that for equal mass systems, more SNR will be in the ringdown portion of the signal so this will allow for slightly better final mass parameter recovery. In Figure 4.65, we see that again, the detected spin a_{DET} can only take a few values. This arises from the discreteness of the quality factor parameter in the template bank.

Finally, we can see in Figure 4.66 that M_{DET} is similar across interferometers. So again, even though the search picks up the wrong parameters, it does so consistently. In Figure 4.67, we again see the discrete nature of a_{DET} .




Figure 4.61: Histogram of difference in GPS start time for triple found PhenomB injections in H1 and L1.




Figure 4.62: Difference in GPS start time for triple found PhenomB injections in H1 and H2 on the left and in H1 and L1 on the right as a function of H1 injected frequency. The color bar denotes injected quality factor in the range from 2.1187-6.0.




Figure 4.63: Injected versus detected final black hole mass for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.




Figure 4.64: Injected versus detected final black hole mass for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the final black hole dimensionless spin from 0.0-1.0. The black line marks the diagonal.




Figure 4.65: Injected versus detected final black hole dimensionless spin for PhenomB injections for each of H1, H2, and L1 interferometers. The color scale spans the mass ratio range from 0.1-1.0. The black line marks the diagonal.




Figure 4.66: Differences in the recovered mass of the final black hole for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal.




Figure 4.67: Differences in the recovered spin of the final black hole for PhenomB waveforms found in triple coincidence between H1 and H2 for the plot on the left and between H1 and L1 for the plot on the right. The black line denotes the diagonal. The discreteness of the recovered values is due to the discreteness of quality factor values available in the ringdown template bank. Recall that the spin of the final black hole is a function only of quality factor.

4.5.1.4 Effective Distance

Recall from Section 4.2.2.2 that the effective distance D_{eff} is a derived quantity based on the amplitude and the sensitivity of the detector to a signal at 1 Mpc. The physical interpretation of this effective distance is the distance to an optimally positioned and oriented source that produces the same strain in the detector as a source at 1 Mpc. Ideally, we should recover the same effective distance that the signal was injected at. In Figure 4.68, however, we see that this is not quite the case, as most of the injections have a detected effective distance that is larger than the injected effective distance. These plots are also colored by recovered SNR so that we can see that injections with very loud SNR are recovered at nearby effective distances and those with quiet SNRs are recovered at further effective distances. Signals injected into noisy LIGO data will often lose SNR and thus be recovered at further effective distances than they were injected.




Figure 4.68: Injected versus detected effective distance for triple found PhenomB injections for H1, H2, and L1. The color bar demonstrates the recovered SNR of these injections in each interferometer.

5. Multimode Ringdown Searches ¹

5.1 Motivation

We are quickly approaching the era of gravitational wave astronomy with a world-wide network of advanced interferometric detectors. By 2014, all advanced LIGO interferometers are expected to be operational and will be joined by at least three other detectors, the GEO600, advanced Virgo, and KAGRA interferometers. Detection rates for binary coalescences have been estimated to likely be around $\sim 40 \text{ yr}^{-1}$ for binary neutron stars and $\sim 20 \text{ yr}^{-1}$ for binary black holes [19, 164]. With the potential for discovery and rich science, it will be important to accurately detect these signals. The method of matched filtering, which is currently used for searches of known gravitational waveforms in LIGO's S1-S6 runs, is the optimal method to detect signals buried in Gaussian noise [165]. While LIGO's noise is not Gaussian, the method, supplemented with additional signal-based vetoes and multi-detector coincidence requirements, is quite suitable for detection purposes.

Black hole perturbation theory and numerical simulations have shown that, according to General Relativity (GR) and many alternative theories of gravity, perturbed black holes experience an exponentially decaying and oscillatory (or ringdown) phase in which the gravitational wave signal is dominated by a series of quasinormal modes (QNMs) (see [166] for a review). LIGO ringdown searches currently assume that the waveform is dominated by the fundamental $l = m = 2$ mode [26]. However, a single mode ringdown search for black holes of final mass $M \gtrsim 10^2 M_\odot$ can miss more than 10% of events in both LIGO and advanced ground-based detectors [131]. Furthermore, parameter estimation errors can be large for such single mode searches when the actual waveform contains a second mode [131]. Two-mode searches have also been proposed for consistency tests of GR and the no-hair theorem [130, 131, 167] and for inferring information about the progenitors of the final black hole formation [168]. Detection of the ringdown signal from an intermediate mass black hole may also provide information into the formation history of these potential gravitational wave sources via the black hole's spin [124, 169]. Thus, if feasible, it is reasonable to expect that future ringdown searches will look for multi-mode signals.

However, multi-dimensional template banks (or catalogs) can result in such a large number of templates that multi-mode ringdown searches might not be possible in practice. The traditional method for building catalogs is to compute a metric for the parameter space that one uses in an algorithm to place points at a given proper distance away from each other and thereby guarantees a minimal match between a signal and template [170–174]. Using this method for a minimal mismatch of 97%, a two-mode catalog suitable for testing the no-hair theorem with advanced ground-based detectors would require roughly 10^6 templates versus 10^3 for a similar one-mode search. This implies a huge increase in the computational cost of matched filtering searches.

A number of reduced order modeling techniques, such as Proper Orthogonal Decompositions, Singular Value Decompositions, or Principal Component Analysis, can be applied to

¹Reprinted with permission from IOP Publishing Limited. The article reprinted in this section has been published in *Classical and Quantum Gravity* [163]. Please see Figure 6.1 in Appendix B for IOP's assignment of copyright form granting permission to include all or part of the article in this dissertation. The article can be found online: <http://iopscience.iop.org/0264-9381/29/9/095016/>

represent the original template bank by potentially much fewer elements (see, for example, 175–178).

Following Field et al. [179], we show that the Reduced Basis (RB) approach provides a dramatically compact representation of multi-mode QNM catalogs. The main two features responsible for such savings are (1) the exponential convergence of the training space representation error (the square of Eq. (5.6) below) with respect to the number of waveforms selected by the algorithm and (2) the linearity of multi-mode ringdown waveforms.

Key results are summarized in the following sections. Additional details and an exploration of RB representations for ringdown searches and parameter estimation are ongoing studies and will be presented as future work.

5.2 The Multimode Ringdown Signal

As discussed in Section 3.5.1, a single QNM waveform has the form

$$h_{\ell mn}(t) = \mathcal{A}_{\ell mn}(\Omega) \frac{M}{r} e^{-(\pi f_{\ell mn}/Q_{\ell mn})t} \cos(2\pi f_{\ell mn}t) \quad (5.1)$$

where $\mathcal{A}_{\ell mn}(\Omega)$ is the orientation-dependent dimensionless amplitude, r is the distance to the source, M is the black hole mass, and for simplicity in the following discussion, we have set the phase and arrival time to zero. We work in units where $G = c = 1$ unless otherwise stated. We will ignore the angular dependence in the amplitude and regard $\mathcal{A}_{\ell mn}$ as constants since, as we discuss in Section 5.5, the linearity of the problem implies that the reduced bases are valid for all values of the amplitudes and, hence, for all angles.

Defining $h_{\ell mn}(t) = 0$ for $t < 0$, the Fourier transform of (5.1) is

$$\hat{h}_{\ell mn}(f) = \frac{\mathcal{A}_{\ell mn}(f_{\ell mn} - 2if_{\ell mn}Q_{\ell mn})Q_{\ell mn}}{\pi[4(f_{\ell mn}^2 - f^2)Q_{\ell mn}^2 + f_{\ell mn}^2 - 4if_{\ell mn}Q_{\ell mn}]}. \quad (5.2)$$

Also, for our purposes, a multi-mode signal impinging on the detector at time $t = 0$ and its Fourier transform are then

$$h(t) = \sum_{\ell, m, n} h_{\ell mn}(t), \quad \hat{h}(f) = \sum_{\ell, m, n} \hat{h}_{\ell mn}(f),$$

with $h_{\ell mn}(t)$ and $\hat{h}_{\ell mn}(f)$ given by Eqs. (5.1) and (5.2), respectively.

Again, as discussed in Section 3.5.1, we can relate the central frequency and quality factor for each multipole to the black hole mass and dimensionless spin parameter using the fitting formulae that can be found in Refs. [130, 131].

5.3 The Reduced Basis Technique

The RB approach is a framework for efficiently solving parametrized problems, representing the solutions in a compact way, and predicting new ones based on an offline-online decomposition (see [179] and references therein). In the current context, the method identifies a set of parameter points such that the associated waveforms constitute a nearly optimal basis for accurate “spectral” expansions of any other waveform [180]. If the waveforms depend smoothly on the parameters of the problem the expansion is expected to converge very

quickly to the original waveform as the number of basis elements is increased. In [179], it was shown that in the case of gravitational waves from non-spinning compact binary inspirals such convergence is exponential. Below we show that this is also the case for ringdown waveforms (see, for example, Fig. 5.1).

In what follows we denote a generic waveform, which can be in the time or frequency domain, by $h_{\vec{\mu}}$ and the relevant parameters by $\vec{\mu}$, which is in general multi-dimensional. In the current application, $\vec{\mu} = \{(f_{\ell mn}, Q_{\ell mn}, \mathcal{A}_{\ell mn})\}$ ². The scheme starts with a *training space* $\mathcal{T} \equiv \{\vec{\mu}_i\}_{i=1}^P$ of P numerical values of the parameter $\vec{\mu}$ and associated normalized waveforms $\{h_{\vec{\mu}_i}\}_{i=1}^P$, which we call the *training space catalog*. (The underlines denote waveforms or parameters present in the training space.) The natural scalar product used is directly related to the one of Wiener filtering (see, for example, 78). For two functions F, G in frequency space

$$\langle F, G \rangle = \int_{f_L}^{f_U} \frac{F^*(f)G(f)}{S_n(f)} df, \quad (5.3)$$

with $S_n(f)$ the one-sided power spectral density (PSD) of the detector and f_L, f_U frequency limits implied by the structure of the PSD. The scalar product of (5.3) is related to the standard overlap integral [181] by $(F, G) := 4 \operatorname{Re}\langle F, G \rangle$. For definiteness, in this paper we use the Adv LIGO PSD fit of Ajith & Bose [182]. While any PSD may be used, as discussed in Field et al. [179], the high accuracy of the reduced basis implies that it can be used with other PSD's with a marginal increase in the representation error.

The training space can be constructed by any means, including simple random or uniform sampling, more sophisticated stochastic methods [183], the metric approach [170], or those of Manca & Vallisneri [184], for example. Regardless of the method used to populate it, the RB formalism produces a compact and highly accurate representation of the training space catalog, among other things.

Part of the output of the algorithm is a sequential selection of N parameter points $\{\vec{\mu}_1, \vec{\mu}_2, \dots, \vec{\mu}_N\} \subset \mathcal{T}$ and their associated waveforms $\{h_{\vec{\mu}_1}, h_{\vec{\mu}_2}, \dots, h_{\vec{\mu}_N}\} \subset \{h_{\vec{\mu}_i}\}_{i=1}^P$. The set of waveforms $\{h_{\vec{\mu}_i}\}_{i=1}^N$ (or a linear combination of them) constitutes the *reduced basis*. It is sometimes convenient [see Eq. (5.5)], though not necessary, to work with an orthonormal set $\{e_i\}_{i=1}^N$ instead of the $\{h_{\vec{\mu}_i}\}_{i=1}^N$.

The points $\{\vec{\mu}_i\}_{i=2}^N$ are selected by a *greedy algorithm* [185] and are sometimes referred to as *greedy points*. A choice for the first one, $\vec{\mu}_1$ (called the *seed*), is needed to initialize the algorithm. The seed can be chosen arbitrarily, and since the greedy scheme is a global optimization method it is insensitive to its choice; we explicitly illustrate this in Section 5.4.2 (see Fig. 5.1). The algorithm ends when a prescribed tolerance error, defined below in Eq. (5.6), in the representation of the training space catalog is reached. For the salient features of the greedy algorithm see Alg. 1.

²There is actually one less degree of freedom with respect to the number of amplitudes since the waveforms are normalized.

Algorithm 1 Brief description of the Greedy Algorithm

- 1: **Input:** $\{\underline{\mu}_i, \underline{h}_{\underline{\mu}_i}\}_{i=1}^P, \epsilon$
 - 2: **Seed choice** (arbitrary): $\underline{\mu}_1$
 - 3: RB = $\{h_{\underline{\mu}_1}\}$
 - 4: $i = 1$ and $\epsilon_1 = 1$
 - 5: **while** $\epsilon_i \geq \epsilon$ **do**
 - 6: $i = i + 1$
 - 7: $\epsilon_i = \max_{\underline{\mu} \in \mathcal{T}} \|P_{(i-1)} \underline{h}_{\underline{\mu}} - \underline{h}_{\underline{\mu}}\|$
 - 8: $\underline{\mu}_i = \operatorname{argmax}_{\underline{\mu} \in \mathcal{T}} \|P_{(i-1)} \underline{h}_{\underline{\mu}} - \underline{h}_{\underline{\mu}}\|$
 - 9: RB = RB $\cup \underline{h}_{\underline{\mu}_i}$
 - 10: **end while**
 - 11: **Output:** RB and greedy points
-

The reduced basis is used to approximate other waveforms, whether they were in the training space catalog or not, through linear combinations that represent an orthogonal projection (with respect to the scalar product (5.3)) onto its span,

$$h_{\underline{\mu}} = P_N h_{\underline{\mu}} + \delta h_{\underline{\mu}}. \quad (5.4)$$

The RB approximation is $P_N h_{\underline{\mu}}$ and satisfies $\langle P_N h_{\underline{\mu}}, \delta h_{\underline{\mu}} \rangle = 0$, by construction. If using an orthonormal basis $\{e_i\}_{i=1}^N$ then

$$P_N h_{\underline{\mu}_j} = \sum_{i=1}^N \langle e_i, h_{\underline{\mu}_j} \rangle e_i =: \sum_{i=1}^N \alpha_{ij} e_i, \quad (5.5)$$

where we have implicitly introduced the projection coefficients α_{ij} . Another output of the algorithm is precisely the set of α_{ij} projection coefficients for waveforms in the training space catalog, whereas coefficients for any other waveforms may be easily computed.

The greedy algorithm sequentially selects $N = N(\epsilon)$ greedy points and associated waveforms until the *maximum* error to represent *every* element in the training space catalog reaches the specified tolerance ϵ ,

$$\epsilon := \max_{\underline{\mu} \in \mathcal{T}} \|\delta h_{\underline{\mu}}\|, \quad (5.6)$$

where the norm $\|\cdot\|$ is the one induced by the scalar product (5.3). As discussed in Field et al. [179], in the limit of sufficiently dense training spaces the square of ϵ is comparable to the minimal match (MM) through ³

$$\epsilon^2 \sim 1 - \text{MM} \quad \text{as } P \rightarrow \infty. \quad (5.7)$$

³Strictly speaking, standard template placement algorithms use $1 - \text{MM}$ to characterize how well a discrete template bank captures any signal h . On the other hand, ϵ^2 characterizes how well $P_N h$ represents any signal h . Albeit different, these error measurements are appropriate tools for comparison.

Therefore, we refer to ϵ^2 as the *training space representation error*⁴. The exponential convergence of the RB method implies that ϵ can be made arbitrarily small with a *reduced* number N of *basis* elements (thus the Reduced Basis denomination) with $N < P$ and, in many cases, $N \ll P$. The quantity

$$C_r := P/N \tag{5.8}$$

is called the compression ratio [186]. Unless otherwise noted, all the results that we quote are for $\epsilon^2 = 10^{-12}$ in order to avoid roundoff artifacts that gradually appear (otherwise we could choose ϵ^2 around double precision roundoff, $\epsilon^2 \sim 10^{-14}$). Therefore, any waveform $\underline{h}_{\vec{\mu}}$ in the training space catalog equals its RB approximation $P_N \underline{h}_{\vec{\mu}}$ to this level of accuracy.

Another error of interest is associated with the RB representation of *any* waveform $h_{\vec{\mu}}$, not necessarily present in the training space catalog. We call this error the *waveform representation error*, $\delta(\vec{\mu})$, and define it as,

$$\delta(\vec{\mu}) := \|\delta h_{\vec{\mu}}\|^2 = 1 - \text{Re}\langle h_{\vec{\mu}}, P_N h_{\vec{\mu}} \rangle \tag{5.9}$$

We quantify this error through Monte Carlo simulations, where we randomly sample $\vec{\mu}$ and compute $\delta(\vec{\mu})$. Remarkably, as discussed in Sections 5.4.3 and 5.5, a training space with a modest, finite minimal match will produce a reduced basis that represents the *whole space* of QNM waveforms with extremely high accuracy.

The greedy points are hierarchical (i.e. nested) implying that for extra/less accuracy in the RB representation points are simply added/removed. Each greedy sweep is embarrassingly parallel and the computational complexity in going from i to $(i + 1)$ basis waveforms is independent of i . In other words, the total cost of building a RB with N waveforms is linear in N . These features allow the framework to handle large training spaces (cf. Table 5.1). Another salient aspect of the RB approach is that it allows one to efficiently identify on the fly those parameters most relevant for numerical simulations, without any a-priori knowledge of the solutions (see, for example, [187]).

In the simplest, idealized conceptual matched filtering search one would compute the overlap between the signal s and every member of the training space through its RB representation

$$(s, P_N \underline{h}_{\vec{\mu}_j}) = 4 \text{Re} \sum_{i=1}^N \langle s, e_i \rangle \alpha_{ij} + \mathcal{O}(\epsilon^2) . \tag{5.10}$$

Since the α_{ij} have been precomputed offline, the filtering now involves computing significantly fewer integrals if $N \ll P$, as is usually the case.

This idealized picture gets considerably more involved when including external parameters. The standard approach for maximizing the filter output over these parameters involves an extremely efficient Fast Fourier Transform (FFT). However, similar to a Singular Value Decomposition reconstruction, the RB reconstruction computational cost of the same direct FFT strategy would be very high and offset the savings in the reduced order modeling. This is discussed in detail in Cannon et al. [188] so we do not repeat the analysis here since the issue is exactly the same. Strategies such as those proposed in Cannon et al. [188, 189, 190] should be investigated in detail to evaluate the efficiency of RB in actual searches.

If one attempted a matched filter search with a RB catalog by filtering each basis function against the data and maximizing over arbitrary linear combinations of these filter outputs,

⁴Note that in Field et al. [179] it is called ϵ the *greedy error* and here ϵ^2 is called the *training space representation error*

one might encounter high false alarm rates. Instead, it is important to allow only linear combinations that correspond to physical waveforms; see Field et al. [179] for more details.

5.4 One Mode Ringdown Catalogs

5.4.1 Metric-based Catalogs

Searches for gravitational waves from perturbed black holes have so far used template banks with only the fundamental $(\ell, m, n) = (2, 2, 0)$ quasinormal mode (see [26] and references therein). As described in Section 4.1.2, the ringdown analysis pipeline constructs a 2-dimensional lattice template bank in the (f, Q) space where the mismatch between two templates differing in ringdown frequency by df and in quality factor by dQ is given by the metric [26, 171, 191]

$$ds^2 = \frac{1}{8} \left[\frac{3 + 16Q^4}{Q^2(1 + 4Q^2)^2} dQ^2 - 2 \frac{3 + 4Q^2}{fQ(1 + 4Q^2)} dQ df + \frac{3 + 8Q^2}{f^2} df^2 \right], \quad (5.11)$$

which assumes white noise.

An example of a ringdown template bank for

$$\text{MM} = 0.99, \quad (5.12)$$

$$f_{220} \in [10, 4000] \text{ Hz}, \quad Q_{220} \in [2.1187, 20] \quad (5.13)$$

is shown in Fig. 5.2 as the gray points. This catalog has a total of 2,213 templates. From Eqs. (3.10) these particular parameter choices correspond to the spin of the black hole j in the range $[0, 0.9947]$ and mass M in the range $[2.9744, 3025.7]M_\odot$. While the (f_{220}, Q_{220}) parameter space is a rectangle by construction, the corresponding shape in (M, j) space is a “warped rectangle” and so our stated ranges are not inclusive. Except when otherwise noted, the values and ranges in (5.12) and (5.13) are our default throughout the remainder.

5.4.2 One-mode Reduced Basis

We begin by using the single-mode QNM metric template bank of Eqs. (5.12,5.13) and Fig. 5.2 as our default training space for building a reduced basis. In order to explicitly show the robustness of the method we use all 2,213 possible elements of the training space as a seed and run the algorithm 2,213 times, once for each possible seed. The results are summarized in Fig. 5.1, which clearly shows that the method is robust and its accuracy and exponential convergence rate do not depend on any fine tuning of the seed. This is actually expected as a consequence of the greedy algorithm being a *global* optimization method. As a result, in practice one chooses a single, arbitrary value of the seed. Notice also from Fig. 5.1 that the exponential convergence of the error is not asymptotic (i.e. for large number of RB waveforms) but is present from the outset.

Fig. 5.2 shows the points in the training space and the subset selected by the greedy algorithm for our default training space representation error of $\epsilon^2 = 10^{-12}$. The selected




Figure 5.1: Representation error as a function of the number of reduced basis waveforms for a single mode catalog and all possible choices of seeds (see Section 5.4.2). The dark line shows the average and the shaded area the maximum dispersion around it. Clearly, the method is robust and does not require any fine tuning. Notice that the exponential rate of convergence of the error is present from the outset.

points are essentially those with the largest quality factor Q (corresponding to the slowest decaying modes) and a very few extra ones with lower Q and with large and small central frequencies. The figure illustrates the global nature of the algorithm and the redundancy present in sampling via a local criteria.

As the maximum quality factor in the training space is increased, we find that the number of RB waveforms grows linearly with the corresponding number of templates in the training space, leading to a roughly constant compression ratio [as defined by Eq. (5.8)] for any chosen ϵ . The qualitative behavior of selected points remains the same, with the vertical right red line in Fig. 5.2 shifting to higher Q values and the reduced basis asymptotically resembling a Fourier representation with a few fast decaying waveforms to account for the damping of QNMs.

Another interesting feature of the greedy point selection is that it can be training space-dependent. A rough uniform sampling of f_{22} and Q_{22} can lead to a distinctly different selection of points than a metric training space as can be seen in Figures 5.3-5.5. However, finer uniform sampling (or any sampling method), results in a selection of points that approach the selection from a metric training space. This is likely due to the fact that the metric placement algorithm already does an efficient and thorough job of placing templates in f_{22} and Q_{22} space. This observation motivated our choice of a metric training space for these studies. However, our results can be generalized for any choice of a reasonable training space.

Columns 2 and 3 in Table 5.1 show the number of RB waveforms needed to represent training space catalogs with different minimal matches. The first element populated by the metric placement algorithm is used as the seed for the greedy algorithm but the results are insensitive to any other choice (cf. Fig. 5.1 and related discussion).




Figure 5.2: One-mode metric placement catalog (gray dots), which is taken to be the training space, and the subset of points selected by the greedy algorithm (red bars) for the $(2, 2, 0)$ mode. The settings are discussed in Sections 5.4.1 and 5.4.2.

It might appear from Table 5.1 that for $MM = 0.97 - 0.99$ and for a high-accuracy training space representation error of $\epsilon^2 = 10^{-12}$ the compression ratio should be modest ($C_r \sim 2 - 4$) for a single mode catalog and would only be significant for larger MMs and/or smaller RB representation errors. However, as discussed below in Section 5.4.3, it turns out that the reduced basis represents *any* waveform in the given ranges of central frequency and quality factor with extremely high accuracy, leading to huge compression factors even when the reduced basis is built from catalogs with comparatively coarse minimal matches.

As in the case of inspiral waveforms described in [179], it is found that for any finite range of parameters (e.g., of the central frequency and quality factor) we can represent the *whole continuum* of waveforms within any given training space representation error ϵ^2 by a *finite* number of RB waveforms. One way of showing this (as was done in [179]) is by explicitly computing the number of RB waveforms needed to represent training spaces built from different values of MM. An extrapolation to the limit that the training space becomes the continuum space of waveforms (i.e., $MM \rightarrow 1$) turns out to asymptote to a finite number of RB waveforms (see also [189] where a similar result is found for $MM \leq 0.99$ in the context of singular value decompositions for inspiral waveforms). Another way, which actually has broader implications and has not been demonstrated before, is discussed next.

5.4.3 Reduced Basis Representation of One-mode Waveforms

Every waveform from the training space catalog is represented with an arbitrarily small error ($\epsilon^2 \lesssim 10^{-12}$ in all of our cases) by a relatively small number of RB elements. We now discuss the accuracy of the basis in representing waveforms that are *not* necessarily part of the training space.




Figure 5.3: Greedy points selected from a 45×45 uniform training space. The uniform training space is laid out in the plot on the left and the points selected by the greedy algorithm are shown in the plot on the right.




Figure 5.4: Greedy points selected from a 132×132 uniform training space on the left and from a 300×300 uniform training space on the right. Note how using a more densely packed training space (as opposed to the 45×45 uniform training space in Figure 5.3) reduces the spread in f_{22} at lower values of Q_{22} . The favored greedy points are along the extrema of frequency values except at the highest Q_{22} . As we can see in Figure 5.5, the points selected via denser and denser uniform training spaces asymptote to the points selected when using a typical metric training space.




Figure 5.5: Greedy points selected from a metric training space with $1 - \text{MM} = 10^{-3}$. The metric training space is laid out in the plot on the left and the points selected by the greedy algorithm are shown in the plot on the right. This is the selection of greedy points to which all other training spaces explored approach.

Table 5.1: Number of RB waveforms (N_{RB}) needed to represent 1-mode training spaces for the $(\ell, m, n) = (2, 2, 0)$ and $(3, 3, 0)$ QNMs with different minimal matches MM. The training space representation error is taken to be $\epsilon^2 = 10^{-12}$. The number of metric-based templates scales with MM as $N_{\text{metric}} \propto (1 - \text{MM})^{-1}$ for the 2-dimensional, 1-mode QNM catalog [170].

1 - MM	1-mode		
	N_{metric}	$N_{\text{RB}} (2, 2, 0)$	$N_{\text{RB}} (3, 3, 0)$
0.03	999	487	711
10^{-2}	2,213	505	732
10^{-3}	19,900	565	930
10^{-4}	192,747	595	972
10^{-5}	1,903,689	603	987

Starting from a catalog with a given MM to build the basis, the worst possible scenario would be a representation error for some particular waveform $h_{\vec{\mu}}$ of $\|\delta h_{\vec{\mu}}\|^2 \sim (1 - \text{MM})$. However, one expects the accuracy to be much better since the RB framework exploits the global structure of the template space. As we show below, the accuracy is indeed orders of magnitude better than the pessimistic $1 - \text{MM}$ upper bound.

We first built a RB representation of our default training space catalog as specified by Eqs. (5.12) and (5.13). We then randomly sampled more than 6×10^8 single-mode QNM waveforms drawn uniformly from the same ranges of central frequency and quality factor (5.13). We finally evaluated the waveform representation error (5.9) for each sample, which gives the fractional signal-to-noise loss from approximating a waveform by its RB representation. The results of these Monte Carlo computations are shown in Fig. 5.7. The error in representing *any* of the randomly chosen waveforms with its RB representation is found to be smaller than 9×10^{-10} . In order to pinpoint this upper bound we refined the Monte Carlo simulations to draw samples from the regions of parameter space with the largest waveform representation errors (i.e., the areas shaded red in Fig. 5.7b).

The average waveform representation error for a Monte Carlo simulation with 10^7 points, corresponding to Fig. 5.7, is approximately 4.51×10^{-13} , which is less than the training space representation error of $\epsilon^2 = 10^{-12}$, and the most frequent value is $\approx 2.6 \times 10^{-14}$.

In other words, in all of our Monte Carlo simulations, *any* waveform, not just those present in the training space catalog, is found to be represented by the RB with extremely high accuracy.

A simple extrapolation from Table 5.1 shows that $\sim 10^{10}$ metric templates would be needed to achieve a MM comparable to the above waveform representation error of $\sim 10^{-9}$ for our reduced basis with 505 elements, implying an effective compression ratio of $C_r \sim 10^7$. Furthermore, the fact that our simulations strongly indicate that the maximum waveform representation error is strictly bounded (by 9×10^{-10} in the settings discussed) would imply a formally infinite compression ratio. This again reflects another facet of being able to represent the whole spectrum of waveforms with a finite number of basis elements. The key distinction here, compared to the discussion at the end of the previous subsection (5.4.2), is that one does not need to explicitly compute a set of bases with increasingly larger number of training space points to reach a target accuracy limit. This would allow for the *a priori* construction of high accuracy reduced bases for sources with larger number of parameters, such as precessing binary inspirals, at a fraction of the projected computational cost.

We have done exhaustive additional tests to support the conclusion that the representation error is extremely small for *all* waveforms with parameter values in the set ranges [e.g., as in Eq. (5.13)]. For example, since a random, uniform sampling of (f_{220}, Q_{220}) leads to a non-uniform sampling of (M, j) [see Eq. (3.10)] we have performed Monte Carlo simulations with random, uniform sampling of mass and spin such that the corresponding central frequencies and quality factors are within the ranges of Eq. (5.13). In addition, instead of selecting waveforms at random, we have used all those from a metric-based catalog with a minimal match of about 0.9999, corresponding to more than 10^5 templates. We represented the waveforms using the RB generated by our default training space [from Eqs. (5.12) and (5.13)] and computed the corresponding waveform representation errors. In all cases the errors were found to be bounded by the previously quoted value of 9×10^{-10} .

As discussed in the next section, the high accuracy associated with representing arbitrary waveforms by a RB carries over to the case of multiple QNMs.

5.5 Multi-Mode Ringdown Catalogs

In this section we discuss how the RB approach can be used to efficiently build, compress,




Figure 5.6: Waveform representation error of randomly selected QNMs that are not necessarily present in the training space. Distributions of waveform representation errors for 1-mode (gray) and 2-mode (white) using reduced bases built from a training space with $\text{MM} = 0.99$. In all cases, the error is found to be $< 9 \times 10^{-10}$ for single modes and $< 9.8 \times 10^{-10}$ for two modes. See Section 5.4.3 and 5.5.1 for details.




Figure 5.7: Waveform representation error of randomly selected QNMs that are not necessarily present in the training space. The waveform representation error [defined in (5.9)] as a function of the randomly chosen values of (f_{220}, Q_{220}) for the 1-mode case with the same setup as in Figure 5.7. The color map indicates the value of the error on a \log_{10} scale. The worst errors occur for the largest Q_{220} values and are mostly independent of f_{220} . The crosses indicate the parameters selected by the greedy algorithm for this reduced basis.

and represent the space of multiple QNM waveforms. We present two kinds of multi-QNM reduced basis. The first is for a two-mode ringdown waveform wherein the two central frequencies and quality factors are related to each other through a constraining relation provided by GR for distorted black holes. The second is for a ringdown waveform consisting of p modes where the central frequencies and quality factors are unconstrained by any model or theory. In particular, this latter case is amenable for providing a test of the no-hair theorem [131], through a test of the precise relationship between the multiple ringdown modes as predicted by GR.

5.5.1 Constrained Two-mode Reduced Basis

Here we consider two-mode ringdown waveforms with $(\ell, m, n) = (2, 2, 0)$ and $(3, 3, 0)$, which are of the form

$$h = \mathcal{C} [(1 - \mathcal{A})h_{220} + \mathcal{A}h_{330}] \quad (5.14)$$

where $\mathcal{A} \in [0, 1]$ and \mathcal{C} is fixed by the normalization condition $\langle h, h \rangle = 1$. The parameters for this two-mode QNM are (f_{220}, Q_{220}) , (f_{330}, Q_{330}) , and the relative amplitude parameter \mathcal{A} . However, GR provides a relation between $(f_{\ell mn}, Q_{\ell mn})$ and the mass and spin of a perturbed black hole. From Eqs. (3.9) or, specifically for the two-mode case here, Eqs. (3.10) and (3.11) it follows that (f_{330}, Q_{330}) are related to (f_{220}, Q_{220}) so that there are only three independent parameters. In this sense, GR constrains the otherwise 5-dimensional parameter space to a 3-dimensional one, namely, $\{f_{220}, Q_{220}, \mathcal{A}\}$ (or, equivalently, $\{j, M, \mathcal{A}\}$).

For a given MM and for $\mathcal{A} \in [0, 1]$ we build a training space for the $(2, 2, 0)$ mode using the metric approach as described in Section 5.4.1⁵. Next, using the fitting formulae given by Eq. (3.10), the values of mass and spin for each template are determined. Then, using Eq. (3.11), the corresponding values of (f_{330}, Q_{330}) are computed and paired to the starting (f_{220}, Q_{220}) . This procedure is repeated for each template in the $(2, 2, 0)$ catalog and subsequently populates the (f_{330}, Q_{330}) plane with templates that are “inherited” from the $(2, 2, 0)$ training space. The relative amplitude parameter \mathcal{A} is sampled with $n_{\mathcal{A}}$ equally spaced points. In this section and in Table 5.1 for the $(3, 3, 0)$ mode column, MM refers to the minimal match of the $(2, 2, 0)$ starting catalog.

5.5.1.1 Results

In all of our numerical simulations we have found that, for any fixed training space representation error ϵ^2 , the number of reduced basis waveforms is *independent of $n_{\mathcal{A}}$ for $n_{\mathcal{A}} \geq 2$* . This is not too surprising, given the linearity of the two-mode waveform (5.14); see also the discussion in Section 5.5.2. A related observation is that, as shown in Fig. 5.8, the greedy algorithm essentially selects *either* the $(2, 2, 0)$ mode ($\mathcal{A} = 0$) *or* the $(3, 3, 0)$ one ($\mathcal{A} = 1$) as opposed to combinations of them. As a result, the distribution of parameter points selected by the algorithm closely resembles the union of a subset of points for two individual, one-mode reduced bases. This is illustrated in Figure 5.9, which should be compared to Figure 5.2 for a one-mode $(2, 2, 0)$ reduced basis.

The number of constrained, two-mode RB waveforms N_{RB} for different values of MM is shown in Table 5.2 in the column labeled “2-mode, GR.” Notice that these values are only marginally larger than those corresponding to a one-mode $(3, 3, 0)$ basis in Table 5.1.

⁵One could use any method to populate the training space, as mentioned in Section 5.3, but we use the metric approach since it is available for this case and the training space then has a well-defined interpretation in terms of minimal matches.




Figure 5.8: Values for the relative amplitude parameter \mathcal{A} in the two-mode constrained waveforms [Eq. (5.14)] selected by the greedy algorithm for each reduced basis waveform (indexed on the horizontal axis). Here there are $n_{\mathcal{A}} = 1,000$ samples for $\mathcal{A} \in [0, 1]$ but essentially either the $(2, 2, 0)$ mode ($\mathcal{A} = 0$) or the $(3, 3, 0)$ mode ($\mathcal{A} = 1$) are selected, with a higher density of points associated with the latter due to its larger maximum quality factor for any given black hole spin parameter value. About 3% of the selected amplitudes are *equal* to zero (not just nearly so). All of the remaining points selected are either *equal* to 1 (about 79%) or between $\mathcal{A} = 0.93$ and 1 (about 18%).

To understand why, notice from Fig. 5.8 that the $(3, 3, 0)$ mode, corresponding to $\mathcal{A} = 1$ in (5.14), is chosen much more frequently than the $(2, 2, 0)$ mode. In turn, the greedy algorithm’s preference for $\mathcal{A} = 1$ can be understood by the fact that: i) for any maximum value of Q in the starting $(2, 2, 0)$ catalog the corresponding value for the $(3, 3, 0)$ mode for the same black hole spin parameter j is larger (see Eqs. (3.10) and (3.11)), and ii) as we have already discussed in Section 5.5.1, the greedy algorithm for a single-mode RB mostly selects points with the largest value of Q in the catalog (see Figure 5.2).

Table 5.2 also shows the approximate number of templates N_{metric} for constrained two-modes using the metric placement method in the column labeled “2-mode, GR.” The compression ratio with respect to the metric number of templates for $\text{MM} = 0.99$ is $C_r \approx 24$ and dramatically increases as MM increases. The number of RB waveforms asymptotically approaches about 1,000, which suggests that a two-mode search with the RB approach may be easily feasible.

Next, we show that the accuracy of 2-mode reduced bases and the associated compression ratios are much better, like in the 1-mode case of Section 5.4.3, than initially indicated in Table 5.2, which are already very good.




Figure 5.9: The metric-based training space (points) for the same constrained, two-mode catalog of Fig. 5.8 and the subset of parameter values selected by the greedy algorithm (bars) for a training space representation error of $\epsilon^2 = 10^{-12}$. Red: $(2, 2, 0)$ mode. Blue: $(3, 3, 0)$ mode. The distribution of points closely resembles a subset of the two one-mode greedy points due to the linearity of the problem and to the “decoupling” of selected amplitude points (see Fig. 5.8).

5.5.1.2 Reduced Basis Representation of Two-mode Waveforms

In analogy with the 1-mode case, we have found that two-mode, constrained reduced bases built from relatively coarse training spaces turn out to represent *any* waveform with extremely high accuracy. Since the details follow those of Section 5.4.3 our description here is more succinct.

In order to emphasize our finding that for constrained, two-mode ringdown waveforms the number of basis elements saturates at $n_{\mathcal{A}} = 2$ amplitude samples, we describe waveform representation error results for a reduced basis built with exactly $n_{\mathcal{A}} = 2$. That is, from here on the training space and, as a consequence, the resulting reduced basis only includes $\mathcal{A} = 0$ and $\mathcal{A} = 1$ relative amplitudes in the waveforms.

We have uniformly sampled the constrained, 2-mode parameter space $(f_{220}, Q_{220}, \mathcal{A})$ with f_{220} and Q_{220} in the same range used to build the reduced basis, Eq. (5.13), and with $\mathcal{A} \in [0, 1]$. The corresponding values of f_{330} and Q_{330} were computed as described earlier using the constraint from General Relativity that they correspond to the same black hole mass and spin as those of (f_{220}, Q_{220}) . The resulting waveforms were then projected onto the reduced basis built with a training space of $\text{MM} = 0.99$ as described above and the waveform representation error (5.9) for each sample was computed. We have executed many Monte Carlo simulations, with the largest one having 10^9 random triples $(f_{220}, Q_{220}, \mathcal{A})$. In all cases the maximum waveform representation error was found to be below 9.8×10^{-10} , which is remarkably close to our strict numerical upper bound found for the one mode case

Table 5.2: Number of reduced basis waveforms (N_{RB}) needed to represent 2-mode training spaces with $(\ell, m, n) = (2, 2, 0)$ and $(3, 3, 0)$ for different minimal matches MM. The training space representation error is $\epsilon^2 = 10^{-12}$. The “2-mode, GR” and “2-mode” cases are discussed in Sections 5.5.1 and 5.5.2, respectively. The number of metric-based templates in those two cases scales with MM as $N_{\text{metric}} \propto (1 - \text{MM})^{-d/2}$, with $d = 3$ and $d = 5$, respectively [170]. For the “2-mode, GR” entries, MM refers to the minimal match associated with the training space for the $(2, 2, 0)$ mode, which is used to generate the points in the (f_{330}, Q_{330}) plane as described in the text. For the “2-mode” case the training space representation error of 10^{-12} corresponds to each mode separately and the corresponding total error is bounded from above by 4×10^{-12} , as discussed in Section 5.5.2. We are grateful to V. Cardoso for his help in generating the values of N_{metric} in the second column for the constrained, 2-mode case.

1 - MM	2-mode, GR		2-mode	
	N_{metric}	N_{RB}	N_{metric}	N_{RB}
0.03	3.5×10^3	737	3.4×10^6	1,198
10^{-2}	1.8×10^4	751	5.3×10^7	1,237
10^{-3}	5.8×10^5	958	1.9×10^{10}	1,495
10^{-4}	1.8×10^7	1,007	5.3×10^{12}	1,567
10^{-5}	5.8×10^8	1,018	1.9×10^{15}	1,590

described in Section 5.4.3, namely, 9×10^{-10} . As an example, Fig. 5.7a shows a histogram of the waveform representation error for a Monte Carlo simulation with 10^7 points. In this case, the average error is about 7.35×10^{-13} , which is less than the training space representation error of $\epsilon^2 = 10^{-12}$, and the most frequent value is $\approx 2.4 \times 10^{-13}$.

Fig. 5.10 shows the distribution of the waveform representation error as a function of the sample values for $(f_{220}, Q_{220}, \mathcal{A})$. As with the 1-mode case, the largest errors occur for larger Q_{220} values and are mostly independent of f_{220} and the relative amplitude.

5.5.2 Unconstrained Multi-mode Reduced Basis

As mentioned in Section 5.1, searches for multi-mode ringdown gravitational waves have been proposed as a consistency test of GR and, more specifically, the no-hair theorem. The basic idea is not to enforce the constraint that the ringdown frequencies of the different modes should correspond to the same black hole mass and spin, as in the previous subsection (5.5.1), but to check it *a posteriori*, by assuming the standard relationship in Eq. (3.9) between QNM frequencies and black hole mass and spin for *each* mode.

The dimension of the parameter space for p unconstrained fundamental modes⁶ is $(3p-1)$: $2p$ for (f_{lmn}, Q_{lmn}) and $(p-1)$ for the relative amplitudes of the normalized waveform. Already for $p = 2$ this results in a very large bank of roughly $10^6 - 10^7$ templates when using the metric approach for $\text{MM} = 0.97 - 0.99$ (see Table 5.2).

One could build a “full” reduced basis from a $(3p-1)$ -dimensional training space. However, the linearity of the problem and the results of Section 5.5.1 suggest a much simpler approach. Since ringdown waveforms are a *linear* superposition of single modes [in the

⁶As it should become clear, this is just a simplifying assumption and the general argument does not rely on it.




Figure 5.10: Waveform representation errors, as defined by Eq.(5.9), of constrained, two-mode waveforms for a random sample of $(f_{220}, Q_{220}, \mathcal{A})$ triples. The training space used to build the reduced basis is relatively coarse yet the waveform representation error is found to be, in all cases, smaller than 9.8×10^{-10} . See Section 5.5.1 for more details.

following the index I generically labels the (ℓ, m, n) triple],

$$h = \sum_{I=1}^p \mathcal{A}_I h_I, \quad (5.15)$$

then a “simple” RB representation for (5.15) could be given by building a reduced basis for each I -th mode separately from the I -th training space $\mathcal{T}_I = \{\underline{\mu}_{Ii}\}_{i=1}^{P_I}$ and to define the

representation of the multi-mode wave to be the sum of the individual projections

$$P_N h_{\underline{\mu}} := \sum_{I=1}^p \mathcal{A}_I P_{N_I} h_{I, \underline{\mu}}. \quad (5.16)$$

Here P_{N_I} denotes the standard orthogonal projection onto the reduced basis for the single I -th mode built from a training space catalog associated with \mathcal{T}_I . Denoting by ϵ_I^2 the training space representation error for the I -th mode,

$$\epsilon_I^2 := \max_{\underline{\mu}_I \in \mathcal{T}_I} \|\underline{h}_{I, \underline{\mu}_I} - P_{N_I} \underline{h}_{I, \underline{\mu}_I}\|^2, \quad (5.17)$$

then the total training space representation error for (5.15) is bounded by

$$\epsilon_{\text{simple}}^2 := \max_{\underline{\mu} \in \mathcal{T}} \|\underline{h}_{\underline{\mu}} - P_N \underline{h}_{\underline{\mu}}\|^2 \leq \sum_{I=1}^p \epsilon_I^2 + \sum_{\substack{I, J=1 \\ I \neq J}}^p \epsilon_I \epsilon_J,$$

where the full training space \mathcal{T} is the product of the individual training spaces $\bigotimes_{I=1}^p \mathcal{T}_I$. If all the errors are chosen to be comparable (i.e. $\epsilon_I \sim \epsilon$) then the bound is simply

$$\epsilon_{\text{simple}}^2 \lesssim p^2 \epsilon^2. \quad (5.18)$$

In most cases of practical interest for gravitational wave searches, including advanced and third generation earth-based detectors such as the Einstein telescope, $p \leq 4$ [168]. If $\epsilon^2 = 10^{-12}$, as in all of our cases, then the total training space representation error would still be very small, of order $\epsilon_{\text{simple}}^2 \lesssim \mathcal{O}(10^{-11})$. In fact, one could include up to $p = 10^5$ modes and still have a maximum total error $\epsilon_{\text{simple}}^2 \lesssim 0.01$, which is comparable in terms of waveform representation accuracy to a minimal match of about 0.99. This feature is another advantage of having RB representations with very high accuracy, even if the latter exceeds the accuracy of the detector itself or of the physical modeling. Such representations leave room for simplifications while preserving high accuracy in the overall approach.

The total number $N = \sum_{I=1}^p N_I$ of RB waveforms scales *linearly* with p . If $N_I \sim \mathcal{N}$ basis elements are needed to represent each mode then $N \sim p\mathcal{N}$. In contrast, the number of metric-based templates is proportional to $(1 - \text{MM})^{-(3p-1)/2}$ [170] and increases dramatically with p for any fixed MM (see, for example, Tables 5.1 and 5.2).

Furthermore, the representation (5.16) is valid for all values and ranges of amplitudes \mathcal{A}_I . In particular, N is independent of the range or values of the $\{\mathcal{A}_I\}$. This last observation, though trivial, helps to understand the results of Section 5.5.1 for constrained 2-mode waveforms: essentially one mode *or* the other is selected at each step by the greedy algorithm (see Fig. 5.8), the number of RB waveforms saturates at $n_{\mathcal{A}} = 2$ samples ($\mathcal{A} = 0$ and $\mathcal{A} = 1$), and the selected (f, Q) points (Fig. 5.9) resemble a subset of those picked by the two 1-mode reduced bases.

Already for the $p = 2$ case and template banks with minimal matches as low as $\text{MM} = 0.97$ and 0.99 , this simple RB approach for unconstrained multi-mode QNMs implies a compression ratio of around *three and four orders of magnitude relative to the respective training spaces*. In addition, as discussed in Sections 5.4.3 and 5.5.1, the resulting bases actually represent the *whole space of waveforms* with an accuracy comparable to the training space

representation error, resulting in effectively much larger, if not formally infinite, compression ratios. The number of unconstrained 2-mode [(2, 2, 0) and (3, 3, 0)] metric-based templates for different values of MM and the number of RB waveforms needed to represent them with this simple approach are given in Table 5.2.

5.6 Final Remarks

We have shown that the RB approach provides very compact and high-accuracy representations of multi-mode ringdown gravitational waves. For example, the number of RB waveforms needed to represent *any two-mode* General Relativity quasinormal mode with a representation error of 10^{-12} is *smaller* than the number of metric-based templates for a *one-mode* catalog with a minimal match of $\text{MM} \sim 0.99$. The comparison with the number of two-mode, metric-based templates for a catalog with $\text{MM} \sim 0.97 - 0.99$ designed to test the consistency of GR and the no-hair theorem is even more striking: the number of RB waveforms needed to represent the continuum is 3–4 orders of magnitude smaller than the number of metric-placed templates for minimal matches of 0.97–0.99.

We have demonstrated the robustness of the RB method by observing essentially identical training space representation errors regardless of the seed value. In the case of one-mode ringdown waveforms, we showed how the greedy algorithm preferentially selects waveforms with the largest quality factors to construct the basis. The number of RB waveforms for the $(\ell, m, n) = (2, 2, 0)$ one-mode case needed to represent the continuum asymptotes to about 600 for advanced ground-based gravitational wave detectors. We observed that the compression ratio remains constant as a function of the maximum quality factor. Through detailed numerical studies we demonstrated that any waveform that falls within the one- and two-mode training space’s ranges in central frequency and quality factor (and, in the case of 2-modes, for an arbitrary relative amplitude) can be recovered with an error no larger than $\approx 9 \times 10^{-10}$. Thus, we can construct high accuracy reduced basis representations of these ringdown waveforms very easily off-line. We expect this result to extend to RB representations of spinning binary inspirals and other multi-dimensional systems that depend smoothly on their parameters.

In the case of constrained two-mode ringdown waveforms, we find that the RB method, in addition to its exponential convergence, benefits from the linearity of the QNMs. The number of RB waveforms is independent of the number of relative amplitude parameter samples n_A for $n_A \geq 2$ for any fixed training space representation error. The greedy algorithm essentially only selects, for example, the (2,2,0) mode or the (3,3,0) mode. The number of RB waveforms for the constrained two-mode case to represent the continuum asymptotes to about 1,000 for advanced ground-based detectors. Finally, in the case of unconstrained two-mode ringdown waveforms, we again exploit the linearity of the problem to propose the construction of a “simple” RB for each mode separately. Then the representation of the multi-mode wave is simply the sum of the individual projections. In addition, the number of RB waveforms scales linearly with the number of modes indicating a very large reduction in the computational cost of multi-mode ringdown searches.

The results open the possibility of searches of multi-mode ringdown gravitational waves thereby allowing one to test the no-hair theorem, to significantly reduce the event loss rate of these kinds of signals, to improve parameter estimation, and to possibly infer progenitors from the relative amplitudes of the different QNMs. The data is made freely available at the following website: [192].

6. Conclusion

This thesis demonstrates how a matched filtering search for ringdown waveforms can be used to search LIGO data for perturbed intermediate mass black holes, a class of elusive objects for which electromagnetic observational evidence has been historically difficult to come by. Direct detection of gravitational waves from these sources would allow us to constrain the mass firmly within the intermediate regime for which LIGO is sensitive to ringdown waveforms ($\sim 100\text{-}600 M_{\odot}$). Thus, the ringdown search could provide irrefutable evidence for the existence of intermediate mass black holes, a feat that has not yet been possible via electromagnetic observations and the mass function.

In Chapter 3, we explored the theoretical ringdown horizon distance for an optimally oriented source. If we assume that 1% of the total energy is emitted as gravitational waves during the ringdown phase of a perturbed black hole of $200 M_{\odot}$ and a spin that is 90% of maximum possible value, then we could detect the signal at nearly 500 Mpc in H1 with an SNR of 8.0. However, if the amount of ringdown energy emitted as gravitational waves decreases, or the mass or spin decreases, then the horizon distance will be smaller. In Chapter 4, we saw that the S5 ringdown search pipeline did indeed succeed at detecting simulations of ringdown and IMR signals at $200 M_{\odot}$ out to ~ 300 Mpc. We performed a rough estimate of the rate upper limit that we will be able to set for ringdown injections in S5 and found that it was significantly better than that set by the S4 ringdown search.

Furthermore, in Chapter 4, we outlined the ringdown search pipeline and the changes needed in order to run on S5 data at a lower SNR threshold. We examined the implementation of a new coincidence test, the development of a post-processing pipeline, the implementation of a new detection statistic, and the parameter recovery performance of the pipeline. We also examined results from four months of the S5 ringdown search; no significant gravitational wave candidates were found.

Finally, in Chapter 5, we discussed the benefits of upgrading the ringdown search to a multi-mode ringdown search. A single mode ringdown search, where the waveform is dominated by the fundamental $l = m = 2$ mode, for black holes of final mass $M \gtrsim 10^2 M_{\odot}$ can miss more than 10% of events in both LIGO and advanced ground-based detectors [131]. And parameter estimation errors, as we saw in Chapter 4, can be large for such single mode searches when the actual waveform contains a second mode [131]. Furthermore, two-mode searches have been proposed for consistency tests of GR and the no-hair theorem [130, 131, 167] and for inferring information about the progenitors of the final black hole formation [168]. However, a difficulty arises in the sheer number of templates needed to do a multi-parameter ringdown search; the computational cost becomes large. In Chapter 5, we demonstrated the success of the reduced basis method at compressing the space significantly and reducing the number of matched filter operations, opening the possibility of searches of multi-mode ringdown gravitational waves.

To summarize, the main conclusions of this thesis were the following:

- In order to efficiently analyze two years of S5 data, the S4 ringdown search pipeline needed a makeover. Implemented changes included further automation of the search pipeline and the creation of a fully automated post-processing pipeline for coincident triggers from the search pipeline. Relationships for mass, spin, frequency, and quality factor were updated in the injection code. A new 3D coincidence test was implemented to check simultaneously for coincidence in frequency, quality factor, and time. H1H2 consistency and SNR ratio cuts were applied, new coincidence clustering stages were added, multiple new statistics aimed at separating signal from background were

computed and fed to a Multi-Variate Statistical Classifier (MVSC), our new detection statistic. In using MVSC, we also tackled previous issues with improper injection identification that led to the misclassification of detector noise. We hope that our development and successful use of MVSC will prove useful to other search pipelines in the advanced detector era.

- Results from 4 months of S5 data were obtained with the new pipeline and MVSC as a detection statistic. No plausible ringdown gravitational wave candidates have yet been found.
- Rough calculations of the rate density upper limit for both ringdown injections and IMR injections show that the search will likely set a better upper limit for rate of ringdowns and/or IMBH-IMBH mergers than any previous search.
- For simulated signal parameter recovery, we find that the recovered mass is too large. However, since the recovered parameters are similar in each interferometer, these signals will not fail the coincidence test and thus, detection is not affected.
- In order to perform a ringdown search with better parameter estimation, we advocated the move to a multi-mode ringdown search and in order to combat the increased computational load of such a search, we explored the Reduced Basis approach which provides very compact and high-accuracy representations of multi-mode ringdown gravitational waves. We found that the number of RB waveforms needed to represent any two-mode General Relativity quasinormal mode with a representation error of 10^{-12} is smaller than the number of metric-based templates for a one-mode catalog with a minimal match of $MM \sim 0.99$. The comparison with the number of two-mode, metric-based templates for a catalog with $MM \sim 0.97 - 0.99$ designed to test the consistency of GR and the no-hair theorem is even more striking: the number of RB waveforms needed to represent the continuum is 3-4 orders of magnitude smaller than the number of metric-based templates for minimal matches of 0.97-0.99.

Looking toward the future, there are several aspects of the ringdown search that need to be addressed. First and foremost, we will examine the results of the ringdown search on the entire two years of LIGO’s fifth science run. The search has been split into 24 two-month analysis periods. We must run the search pipeline and the post-processing pipeline on each of these 24 periods. After examining results in a small section of data to ensure that the output looks sane, we will “open the box” on each of these periods and determine whether we have a candidate ringdown gravitational wave. In the event of no significant candidates, we will place an upper limit on the rate of black hole ringdowns. We may also place a limit on the rate of intermediate mass black hole mergers depending on whether we have superseded the upper limit placed in the S5 intermediate mass black hole burst search[48]. In the event of a significant candidate, we will need to perform thorough follow-up tests to confirm that we do indeed have a signal. Furthermore, in anticipation of such a candidate in S5 or S6, we should explore methods to improve parameter estimation of one-mode ringdown searches. We then will repeat the procedure with data from LIGO’s sixth science run. In S6, only H1 and L1 were operational; H2 was not. However, Virgo’s second and third science runs (VSR2 and VSR3) overlapped with S6. Thus, the network for S6 included H1L1V1. We may need to re-examine some tunings of the search pipeline to adjust for the replacement of H2 with V1 but we plan to use the same search and post-processing pipeline otherwise.

Finally, as we contemplate the future of the ringdown search in the advanced detector era, it is worthwhile to mention current studies to understand the best search for the highest

mass regime in which LIGO is sensitive. A number of searches for binary black hole coalescences in LIGO-Virgo data have been implemented or are on-going including a matched filtering search using non-spinning, full-waveform EOBNR analytical templates in S5/VSR1 [31]. This search targeted a mass covering 25-100 M_{\odot} . The burst search targeting intermediate mass binary black hole (IMBH) mass has been carried out with the pipeline coherent Waveburst for S5/VSR1 [48] and the burst pipeline Omega was used for an S5 all-sky burst search and in studies for binary black hole detectability studies [193, 194]. Each of these pipelines target a region of parameter space that overlaps with the ringdown search's parameter space. Is one better than the other? Should the pipelines be combined into a joint search? These questions will need to be answered as we converge on a intermediate mass black hole search pipeline in the advanced detector era. Our recent results demonstrating the feasibility of multimode ringdown searches via dramatic reduction in the number of matched filter operations will need to be developed further to understand the performance of such a pipeline in actual LIGO data. It will be important to understand the performance of such a search compared to the other searches just mentioned.

Bibliography

- [1] Abbott, B., et al. 2009, Reports on Progress in Physics, 72, 076901
- [2] Sathyaprakash, B., & Schutz, B. 2009, Living Rev. Relativity, 12
- [3] <http://www.geo600.org/>
- [4] Abadie, J., et al. 2010, Phys. Rev. D, 81, 102001
- [5] Abadie, J., et al. 2011, The Astrophysical Journal Letters, 734, L35
- [6] Abbott, B., et al. 2004, Phys. Rev. D, 69, 082004
- [7] Abbott, B., et al. 2007, Classical and Quantum Gravity, 24, 5343
- [8] Abbott, B., et al. 2008, Classical and Quantum Gravity, 25, 245008
- [9] <http://tamago.mtk.nao.ac.jp/spacetime/tama300e.html>
- [10] Abbott, B., et al. 2005, Phys. Rev. D, 72, 122004
- [11] Abbott, B., et al. 2006, Phys. Rev. D, 73, 102002
- [12] <http://gwcenter.icrr.u-tokyo.ac.jp/en/>
- [13] <http://www.gw-indigo.org/tiki-index.php>
- [14] <http://www.ligo.caltech.edu/>
- [15] <https://wwwcascina.virgo.infn.it/>
- [16] Acernese, F., Amico, P., Alshourbagy, M., et al. 2006, Classical and Quantum Gravity, 23, S635
- [17] Adhikari, R., Fritschel, P., & Waldman, S. 2006, LIGO document, LIGO-T060156-01, <http://www.ligo.caltech.edu/docs>
- [18] Harry, G. 2010, Classical and Quantum Gravity, 27, 084006
- [19] Abadie, J., et al. 2010, Classical and Quantum Gravity, 27, 173001
- [20] Abbott, B., et al. 2004, Phys. Rev. D, 69, 122001
- [21] Abbott, B., et al. 2005, Phys. Rev. D, 72, 082001
- [22] Abbott, B., et al. 2005, Phys. Rev. D, 72, 082002
- [23] Abbott, B., et al. 2006, Phys. Rev. D, 73, 062001

- [24] Abbott, B., et al. 2008, Phys. Rev. D, 77, 062002
- [25] Abbott, B., et al. 2008, Phys. Rev. D, 78, 042002
- [26] Abbott, B., et al. 2009, Phys. Rev. D, 80, 062001
- [27] Abbott, B., et al. 2009, Phys. Rev. D, 79, 122001
- [28] Abbott, B., et al. 2009, Phys. Rev. D, 80, 047101
- [29] Abadie, J., et al. 2010, The Astrophysical Journal, 715, 1453
- [30] Abadie, J., et al. 2010, Phys. Rev. D, 82, 102001
- [31] Abadie, J., et al. 2011, Phys. Rev. D, 83, 122005
- [32] Abadie, J., et al. 2012, Phys. Rev. D, 85, 082002
- [33] Abbott, B., et al. 2004, Phys. Rev. D, 69, 102001
- [34] Abbott, B., et al. 2005, Phys. Rev. D, 72, 042002
- [35] Abbott, B., et al. 2005, Phys. Rev. D, 72, 062001
- [36] Abbott, B., et al. 2006, Classical and Quantum Gravity, 23, S29
- [37] Abbott, B., et al. 2008, The Astrophysical Journal, 681, 1419
- [38] Abbott, B., et al. 2008, Phys. Rev. Lett., 101, 211102
- [39] Abbott, B., et al. 2008, Phys. Rev. D, 77, 062004
- [40] Baggio, L., Bignotto, M., Bonaldi, M., et al. 2008, Classical and Quantum Gravity, 25, 095004
- [41] Abbott, B., et al. 2009, Phys. Rev. D, 80, 102001
- [42] Abbott, B., et al. 2009, Phys. Rev. D, 80, 102002
- [43] Abbott, B., et al. 2009, The Astrophysical Journal Letters, 701, L68
- [44] Abbott, B., et al. 2009, Phys. Rev. D, 80, 062002
- [45] Abbott, B., et al. 2010, The Astrophysical Journal, 715, 1438
- [46] Abadie, J., et al. 2011, Phys. Rev. D, 83, 042001
- [47] The LIGO Scientific Collaboration, Abadie, J., Abbott, B. P., et al. 2012, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2012arXiv1201.4413T>
- [48] Abadie, J., et al. 2012, Phys. Rev. D, 85, 102004
- [49] the LIGO Scientific Collaboration, the Virgo Collaboration: J. Abadie, Abbott, B. P., et al. 2012, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2012arXiv1202.2788T>

- [50] The ANTARES Collaboration, the LIGO Scientific Collaboration, the Virgo Collaboration. THE ANTARES COLLABORATION: S. Adrián-Martínez, et al. 2012, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2012arXiv1205.3018T>
- [51] Briggs, M. S., Connaughton, V., Hurley, K. C., et al. 2012, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2012arXiv1205.2216B>
- [52] Evans, P. A. and Fridriksson, J. K. and Gehrels, N. and Homan, J. and Osborne, J. P. and Siegel, M. and Beardmore, A. and Handbauer, P. and Gelbord, J. and Kennea, J. A. and et al. 2012, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2012arXiv1205.1124E>
- [53] Abbott, B., et al. 2005, Phys. Rev. D, 72, 102004
- [54] Abbott, B., et al. 2005, Phys. Rev. Lett., 94, 181103
- [55] Abbott, B., et al. 2007, Phys. Rev. D, 76, 062003
- [56] Abbott, B., et al. 2007, Phys. Rev. D, 76, 082001
- [57] Abbott, B., et al. 2007, Phys. Rev. D, 76, 042001
- [58] Abbott, B., et al. 2008, Phys. Rev. D, 77, 022001
- [59] Abbott, B., et al. 2008, The Astrophysical Journal Letters, 683, L45
- [60] Abbott, B., et al. 2009, Phys. Rev. Lett., 102, 111102
- [61] Abbott, B., et al. 2009, Phys. Rev. D, 80, 042003
- [62] Abbott, B., et al. 2009, Phys. Rev. D, 79, 022001
- [63] Abbott, B., et al. 2010, The Astrophysical Journal, 713, 671
- [64] Abadie, J., et al. 2010, The Astrophysical Journal, 722, 1504
- [65] Abadie, J., et al. 2011, The Astrophysical Journal, 737, 93
- [66] Abadie, J., et al. 2012, Phys. Rev. D, 85, 022001
- [67] Abbott, B., et al. 2004, Phys. Rev. D, 69, 122004
- [68] Abbott, B., et al. 2005, Phys. Rev. Lett., 95, 221101
- [69] Abbott, B., et al. 2007, Phys. Rev. D, 76, 022001
- [70] Abbott, B., et al. 2007, The Astrophysical Journal, 659, 918
- [71] Abbott, B., et al. 2007, Phys. Rev. D, 76, 082003
- [72] Abbott, B., et al. 2009, Nature, 460, 990
- [73] Abadie, J., et al. 2011, ArXiv e-prints. <http://adsabs.harvard.edu/abs/2011arXiv1112.5004A>

- [74] Abadie, J., et al. 2011, *Phys. Rev. Lett.*, 107, 271102
- [75] Farrell, S., Webb, N., Barret, D., Godet, O., & Rodrigues, J. 2009, *Nature*, 460, 73
- [76] Creighton, J., & Anderson, W. 2012, *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis* (Wiley-VCH)
- [77] Hartle, J. 1998, *Gravity. An introduction to Einstein's General Relativity*. 2003
- [78] Maggiore, M. 2008, *Gravitational Waves, Volume 1: Theory and Experiment* (Oxford)
- [79] Misner, C., Thorne, K., & Wheeler, J. 1973, *Gravitation* (WH Freeman & co)
- [80] Wald, R. 1984, *General relativity* (University of Chicago press)
- [81] Schutz, B. 1985, *A first course in general relativity* (Cambridge Univ Pr)
- [82] Thorne, K. 1987, *Gravitational Radiation*, ed. SW Hawking & W. Israel (Cambridge: Cambridge Univ. Press), 330
- [83] Weisberg, J. M., & Taylor, J. H. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 328, *Binary Radio Pulsars*, ed. F. A. Rasio & I. H. Stairs, 25
- [84] Weber, J. 1973, *DETECTOR OF DYNAMIC GRAVITATIONAL FORCE GRADIENT FIELDS*, uS Patent 3,722,284
- [85] Allen, Z., Astone, P., Baggio, L., et al. 2000, *Physical review letters*, 85, 5046
- [86] Astone, P., Bassan, M., Bonifazi, P., et al. 1999, *Physical review D*, 59, 122001
- [87] Pustovoit, V., & Gertsenshtein, M. 1962, *Zhur. Eksptl'. i Teoret. Fiz.*, 42
- [88] Kissel, J. 2010, PhD thesis, Louisiana State University
- [89] Saulson, P. 1994, *Fundamentals of interferometric gravitational wave detectors* (World Scientific)
- [90] Sigg, D. 1998, in *Proceedings of TASI*, Vol. 98, 6–11
- [91] González, G. 2000, *Classical and Quantum Gravity*, 17, 4409
- [92] Hughes, S., & Thorne, K. 1998, *Physical Review D*, 58, 122002
- [93] Laplace, P. 1835, *Exposition du système du monde*, Vol. 4 (Bachelier)
- [94] Michell, J. 1784, *Philosophical transactions of the Royal Society of London*, 74, 35
- [95] Schwarzschild, K. 1916, *Sitzungsber. Preuss. Akad. Wiss., Phys. Math. Kl*, 189
- [96] Wheeler, J. A. 1968, *American Scientist*, 56, 1
- [97] Ruffini, R., & Wheeler, J. 1971, *Physics Today*

- [98] Reissner, H. 1916, *Annalen der Physik*, 355, 106
- [99] Kerr, R. P. 1963, *Phys. Rev. Lett.*, 11, 237
- [100] Bowyer, S., Byram, E. T., Chubb, T. A., & Friedman, H. 1965, *Science*, 147, 394
- [101] Kristian, J., Brucato, R., Visvanathan, N., Lanning, H., & Sandage, A. 1971, *The Astrophysical Journal Letters*, 168, L91
- [102] Braes, L. L. E., & Miley, G. K. 1971, *Nature*, 232, 246
- [103] International Astronomical Union. 1972, *New directions and new frontiers in variable star research*. International Astronomical Union, Colloquium No. 15, Vol. 9 (Bamberg: Veroeffentlichungen der Remeis-Sternwarte zu Bamberg)
- [104] Webster, B. L., & Murdin, P. 1972, *Nature*, 235, 37
- [105] Bolton, C. T. 1972, *Nature*, 235, 271
- [106] Kalogera, V., & Baym, G. 1996, *The Astrophysical Journal Letters*, 470, L61
- [107] Ferrarese, L., & Ford, H. 2005, *Space Science Reviews*, 116, 3, 523
- [108] Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010, *The Astrophysical Journal*, 714, 1217
- [109] Coleman Miller, M., & Colbert, E. J. M. 2004, *International Journal of Modern Physics D*, 13, 1
- [110] King, A., & Dehnen, W. 2005, *Monthly Notices of the Royal Astronomical Society*, 357, 275
- [111] Georgakakis, A., Tsamis, Y., James, B., & Aloisi, A. 2011, *Monthly Notices of the Royal Astronomical Society*
- [112] Silverman, J. M., & Filippenko, A. V. 2008, *The Astrophysical Journal Letters*, 678, L17
- [113] Bulik, T., Belczynski, K., & Prestwich, A. 2011, *The Astrophysical Journal*, 730, 140
- [114] Okajima, T., Ebisawa, K., & Kawaguchi, T. 2006, *The Astrophysical Journal Letters*, 652, L105
- [115] Casella, P., Ponti, G., Patruno, A., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 387, 1707
- [116] Patruno, A., & Zampieri, L. 2010, *Monthly Notices of the Royal Astronomical Society: Letters*, 403, L69
- [117] Liu, J., Stefano, R. D., McClintock, J., et al. 2006, *The Astrophysical Journal*, 653, 602
- [118] Kuntz, K. D., Gruendl, R. A., Chu, Y.-H., et al. 2005, *The Astrophysical Journal Letters*, 620, L31

- [119] Miller, M. C., & Hamilton, D. P. 2002, *Monthly Notices of the Royal Astronomical Society*, 330, 232
- [120] Vesperini, E., McMillan, S. L. W., D’Ercole, A., & D’Antona, F. 2010, *The Astrophysical Journal Letters*, 713, L41
- [121] Mapelli, M., Huwyler, C., Mayer, L., Jetzer, P., & Vecchio, A. 2010, *The Astrophysical Journal*, 719, 987
- [122] Umeda, H., & Nomoto, K. 2003, *Nature*, 422, 871
- [123] Norman, M. L. 2010, in *American Institute of Physics Conference Series*, Vol. 1294, American Institute of Physics Conference Series, ed. D. J. Whalen, V. Bromm, & N. Yoshida, 17–27
- [124] Fregeau, J. M., Larson, S. L., Miller, M. C., O’Shaughnessy, R., & Rasio, F. A. 2006, *The Astrophysical Journal*, 646, L135
- [125] Regge, T., & Wheeler, J. 1957, *Physical Review*, 108, 1063
- [126] Teukolsky, S. 1973, *The Astrophysical Journal*, 185, 635
- [127] Leaver, E. 1985, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 402, 285
- [128] Kokkotas, K., & Schmidt, B. 1999, *Living Rev. Rel*, 2
- [129] Echeverria, F. 1989, *Phys. Rev. D*, 40, 3194
- [130] Berti, E., Cardoso, V., & Will, C. M. 2006, *Phys. Rev. D*, 73, 064030
- [131] Berti, E., Cardoso, J., Cardoso, V., & Cavaglià, M. 2007, *Phys. Rev. D*, 76, 104044
- [132] Campanelli, M., Lousto, C. O., Zlochower, Y., Krishnan, B., & Merritt, D. 2007, *Phys. Rev. D*, 75, 064030
- [133] Campanelli, M., Lousto, C., Zlochower, Y., & Merritt, D. 2007, *The Astrophysical Journal Letters*, 659, L5
- [134] Baker, J. G., Boggs, W. D., Centrella, J., et al. 2007, *The Astrophysical Journal*, 668, 1140
- [135] Buonanno, A., Pan, Y., Baker, J. G., et al. 2007, *Phys. Rev. D*, 76, 104049
- [136] Ajith, P., Babak, S., Chen, Y., et al. 2007, *Classical and Quantum Gravity*, 24, S689
- [137] Pretorius, F. 2005, *Phys. Rev. Lett.*, 95, 121101
- [138] Baumgarte, T., Brady, P. R., Creighton, J. D. E., et al. 2008, *Phys. Rev. D*, 77, 084009
- [139] Damour, T., Nagar, A., Dorband, E., Pollney, D., & Rezzolla, L. 2008, *Physical Review D*, 77, 084017
- [140] Buonanno, A., Cook, G., & Pretorius, F. 2007, *Physical Review D*, 75, 124018

- [141] Hannam, M., Husa, S., González, J., Sperhake, U., & Brüggmann, B. 2008, *Physical Review D*, 77, 044020
- [142] Hannam, M. 2009, *Classical and Quantum Gravity*, 26, 114001
- [143] Buonanno, A., & Damour, T. 1999, *Physical Review D*, 59, 084006
- [144] Buonanno, A., Pan, Y., Pfeiffer, H., et al. 2009, *Physical Review D*, 79, 124028
- [145] Buonanno, A., & Damour, T. 2000, *Physical Review D*, 62, 064015
- [146] Buonanno, A., Iyer, B., Ochsner, E., Pan, Y., & Sathyaprakash, B. 2009, *Physical Review D*, 80, 084043
- [147] Buonanno, A., Chen, Y., & Vallisneri, M. 2003, *Physical Review D*, 67, 104025
- [148] Damour, T., Iyer, B., Jaranowski, P., & Sathyaprakash, B. 2003, *Physical Review D*, 67, 064028
- [149] Damour, T., & Nagar, A. 2009, *Physical Review D*, 79, 081503
- [150] Ajith, P., et al. 2008, *Physical Review D*, 77, 104017
- [151] Ajith, P., Hannam, M., Husa, S., et al. 2011, *Physical Review Letters*, 106, 241101
- [152] Campanelli, M., Lousto, C. O., Marronetti, P., & Zlochower, Y. 2006, *Phys. Rev. Lett.*, 96, 111101
- [153] Baker, J. G., Centrella, J., Choi, D.-I., Koppitz, M., & van Meter, J. 2006, *Phys. Rev. Lett.*, 96, 111102
- [154] Herrmann, F., Hinder, I., Shoemaker, D., & Laguna, P. 2007, *Classical and Quantum Gravity*, 24, S33
- [155] Campanelli, M., Lousto, C. O., & Zlochower, Y. 2006, *Phys. Rev. D*, 74, 041501
- [156] Barausse, E., & Rezzolla, L. 2009, *The Astrophysical Journal*, 704, L40
- [157] Goggin, L. 2008, PhD thesis, California Institute of Technology
- [158] Slutsky, J. 2010, PhD thesis, Louisiana State University
- [159] Capano, C. 2011, PhD thesis, Syracuse University
- [160] Bose, S., Dayanga, T., Ghosh, S., & Talukder, D. 2011, *Classical and Quantum Gravity*, 28, 134009
- [161] <http://www.ligo.org/news/blind-injection.php>
- [162] Mandel, I., Brown, D., Gair, J., & Miller, M. 2008, *The Astrophysical Journal*, 681, 1431
- [163] Caudill, S., Field, S. E., Galley, C. R., Herrmann, F., & Tiglio, M. 2012, *Classical and Quantum Gravity*, 29, 095016

- [164] Belczynski, K., Dominik, M., Bulik, T., et al. 2010, *The Astrophysical Journal*, 715, L138
- [165] Kumar, B., Mahalanobis, A., & Juday, R. 2005, *Correlation pattern recognition* (Cambridge Univ Pr)
- [166] Berti, E., Cardoso, V., & Starinets, A. O. 2009, *Classical and Quantum Gravity*, 26, 163001
- [167] Dreyer, O., Kelly, B., Krishnan, B., et al. 2004, *Classical and Quantum Gravity*, 21, 787
- [168] Kamaretsos, I., Hannam, M., Husa, S., & Sathyaprakash, B. S. 2012, *Phys. Rev. D*, 85, 024018
- [169] Berti, E., & Volonteri, M. 2008, *The Astrophysical Journal*, 684, 822
- [170] Owen, B. J. 1996, *Phys. Rev. D*, 53, 6749
- [171] Creighton, J. D. E. 1999, *Phys. Rev. D*, 60, 022001
- [172] Nakano, H., Takahashi, H., Tagoshi, H., & Sasaki, M. 2003, *Phys. Rev. D*, 68, 102003
- [173] Tsunesada, Y., Kanda, N., Nakano, H., et al. 2005, *Phys. Rev. D*, 71, 103005
- [174] Nakano, H., Takahashi, H., Tagoshi, H., & Misao. 2004, *Progress of Theoretical Physics*, 111, 781
- [175] Heng, I. S. 2009, *Classical and Quantum Gravity*, 26, 105005
- [176] Brady, P. R., & Ray-Majumder, S. 2004, *Classical and Quantum Gravity*, 21, 1839
- [177] Cannon, K., Chapman, A., Hanna, C., et al. 2010, *Phys. Rev. D*, 82, 044025
- [178] Galley, C. R., Herrmann, F., Silberholz, J., Tiglio, M., & Guerberoff, G. 2010, *Classical and Quantum Gravity*, 27, 245007
- [179] Field, S. E., Galley, C. R., Herrmann, F., et al. 2011, *Physical Review Letters*, 106, 221102
- [180] Binev, P., Cohen, A., Dahmen, W., et al. 2011, *SIAM J. Math. Anal.*, 43, 1457
- [181] Finn, L. S. 1992, *Phys. Rev. D*, 46, 5236
- [182] Ajith, P., & Bose, S. 2009, *Phys. Rev. D*, 79, 084032
- [183] Messenger, C., Prix, R., & Papa, M. A. 2009, *Phys. Rev. D*, 79, 104017
- [184] Manca, G. M., & Vallisneri, M. 2010, *Phys. Rev. D*, 81, 024004
- [185] Cormen, T., Leiserson, C., Rivest, R., & Stein, C. 2009, *Introduction to Algorithms*, Vol. 2 (The MIT Press), 588–592

- [186] Salomon, D., Motta, G., & Bryant, D. 2009, Handbook of data compression (Springer-Verlag New York Inc)
- [187] Quarteroni, A., Rozza, G., & Manzoni, A. 2011, Journal of Mathematics in Industry, 1, 3
- [188] Cannon, K., Hanna, C., Keppel, D., & Searle, A. C. 2011, Phys. Rev. D, 83, 084053
- [189] Cannon, K., Hanna, C., & Keppel, D. 2011, Phys. Rev. D, 84, 084003
- [190] Cannon, K., Cariou, R., Chapman, A., et al. 2012, The Astrophysical Journal, 748, 136
- [191] Allen, B., et al. 2005, LAL Software Documentation. Revision 1.44 (2005)
- [192] <http://www.cscamm.umd.edu/people/faculty/tiglio/RB.html>
- [193] Aylott, B., Baker, J. G., Boggs, W. D., et al. 2009, Classical and Quantum Gravity, 26, 165008
- [194] Fischetti, S., Healy, J., Cadonati, L., et al. 2011, Physical Review D, 83, 044019


```

lalapps_inspiral_hipe
plot                = /bin/false
follow              = /bin/false
datafind            = /home/scaudill/opt/ringdown_branch/bin/
ligo_data_find      = /home/scaudill/opt/ringdown_branch/bin/
segfind             = /home/scaudill/opt/ringdown_branch/bin/
ligolw_segment_query
segs_from_cats      = /home/scaudill/opt/ringdown_branch/
bin/ligolw_segments_from_cats
ligolw_segments_compat = /home/scaudill/opt/ringdown_
branch/bin/ligolw_segments_compat
ligolw_print        = /home/scaudill/opt/ringdown_branch/
bin/ligolw_print
ligolw_add          = /home/scaudill/opt/ringdown_branch/
bin/ligolw_add
tmpltbank           = /bin/false
inspiral            = /home/scaudill/opt/ringdown_branch/bin/
lalapps_ring
; both inspinj and rinj are necessary if you want to run inspiral
; and ringdown injections
inspinj             = /home/scaudill/opt/ringdown_branch/
bin/lalapps_inspinj
rinj                = /home/scaudill/opt/ringdown_branch/bin/
lalapps_rinj
inca                = /bin/false
thinca              = /home/scaudill/opt/ringdown_branch/bin/
lalapps_rinca
trigbank            = /bin/false
cohbank             = /bin/false
frjoin              = /bin/false
chia                = /bin/false
sire                = /bin/false
coire               = /home/scaudill/opt/ringdown_branch/bin/
lalapps_coincringread
tmpltbank32         = /bin/false
inspiral32          = /bin/false
inspinj32           = /bin/false
inca32              = /bin/false
thinca32            = /bin/false
trigbank32          = /bin/false
cohbank32           = /bin/false
frjoin32            = /bin/false
chia32              = /bin/false
sire32              = /bin/false
coire32             = /bin/false
plotinspiral        = /bin/false
plotthinca          = /bin/false
plotnumtemplates    = /bin/false
plotinjn            = /bin/false
plotethinca         = /bin/false
plotinspmiss        = /bin/false

```

```
plotinspfound      = /bin/false
plotinspfound      = /bin/false
plotcoincmissed    = /bin/false
plotinspinj        = /bin/false
plotsnrchi         = /bin/false
plotinspiralrange  = /bin/false
ploteffdistcut     = /bin/false
cohbank           = /bin/false
cohinspbank        = /bin/false
frjoin            = /bin/false
chia              = /bin/false
cohire            = /bin/false
```

```
[condor-max-jobs]
```

```
;maximum number of jobs that condor should run of a given category
datafind= 10
thinca = 20
trigbank = 20
coire = 20
sire = 20
```

```
[ldgsubmitdax]
```

```
sites-file = ldg-sites.xml
properties-file = pegasus-properties.bundle
gsiftp = ldas-grid.ligo.caltech.edu
pool = cit
```

```
[sites]
```

```
eobnr_inj      = local
ringdown_inj   = local
phenom_inj     = local
playground     = local
full_data      = local
```

```
[pipeline]
```

```
; tagging information for the configure script
version =
cvs-tag = $Name$
user-tag =
num-slides = 50
playground-data-mask = all_data
retry-subdag = 1
;Set matplotlibdir to a local directory to avoid NFS issues
;with the TeX part
;of matplotlib. Ensure this directory exists on all nodes!
;matplotlibdir =
```

```
; INSPINJ SECTIONS
```

```
;;;;;;;;;
```

```
[injections]
```

```
eobnr_unid = 11118205
eobnr_logdnear = 66668205
eobnr_logdfar = 99998205

phenom_nospin = 21098205
phenom_spinnear = 54328205
phenom_spinfar = 98768205

ringdown_qlogf = 24688205
ringdown_amhigh = 35798205
ringdown_amlow = 13698205
```

```
[eobnr_unid]
time-step = 724.077
time-interval = 300
d-distr = uniform
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
max-mass2 = 400
min-distance = 10
max-distance = 1000000
waveform = EOBNRpseudoFourPN
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =
```

```
[eobnr_logdnear]
time-step = 724.077
time-interval = 300
d-distr = log10
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
```

```
max-mass2 = 400
min-distance = 10
max-distance = 10000
waveform = EOBNRpseudoFourPN
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =
```

```
[eobnr_logdfar]
time-step = 724.077
time-interval = 300
d-distr = log10
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
max-mass2 = 400
min-distance = 1000
max-distance = 1000000
waveform = EOBNRpseudoFourPN
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =
```

```
[phenom_nospin]
time-step = 724.077
time-interval = 300
d-distr = uniform
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
max-mass2 = 400
min-distance = 10
max-distance = 1000000
```

```
waveform = IMRPhenomBpseudoFourPN
disable-spin =
taper-injection = start
injection-type = PHENOM
write-sim-ring =
```

```
[phenom_spinnear]
time-step = 724.077
time-interval = 300
d-distr = log10
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
max-mass2 = 400
min-distance = 10
max-distance = 10000
waveform = IMRPhenomBpseudoFourPN
enable-spin =
aligned =
min-spin1 = 0.1
max-spin1 = 0.85
min-spin2 = 0.1
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =
```

```
[phenom_spinfar]
time-step = 724.077
time-interval = 300
d-distr = log10
l-distr = random
m-distr = totalMassRatio
i-distr = uniform
f-lower = 10.0
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
min-mass1 = 10
min-mass2 = 10
max-mass1 = 400
```

```
max-mass2 = 400
min-distance = 1000
max-distance = 1000000
waveform = IMRPhenomBpseudoFourPN
enable-spin =
aligned =
min-spin1 = 0.1
max-spin1 = 0.85
min-spin2 = 0.1
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =
```

```
[ringdown_qlogf]
waveform = Ringdown
inj-distr = 0
time-step = 724.077
time-interval = 300
minimum-quality = 2.1187
maximum-quality = 20
minimum-frequency = 50
maximum-frequency = 2000
minimum-distance = 10
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN
```

```
[ringdown_amhigh]
waveform = Ringdown
inj-distr = 2
time-step = 724.077
time-interval = 300
minimum-mass = 200
maximum-mass = 800
minimum-spin = 0.6
maximum-spin = 0.99
minimum-distance = 10
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN
```

```
[ringdown_amlow]
waveform = Ringdown
inj-distr = 2
time-step = 724.077
time-interval = 150
minimum-mass = 50
maximum-mass = 400
minimum-spin = 0.1
maximum-spin = 0.99
```

```

minimum-distance = 10
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN

; HIPE SECTIONS
;;;;;;;;;;;;;

[input]
ligo-channel = LSC-STRAIN
ligo-type = RDS_CO3_L2
ligo-qscan-type = RDS_R_L1
max-thinca-segment = 3600
geo-channel =
geo-type =
virgo-type = HrecV2_16384Hz
virgo-channel =h_16384Hz

[calibration]
; location of the calibration cache and the cache files
; if using h(t) data, then no calibration cache required
path = /home/lgoggin/projects/ringdown/s5/calibration
L1 = L1_CAL_REF_DARM_ERR_S5_V1_U.cache
H1 = H1_CAL_REF_DARM_ERR_S5_V1_U.cache
H2 = H2_CAL_REF_DARM_ERR_S5_V1_U.cache

[datafind]
; type of data to use -- for LDAS-GRID, nothing required here.
;match=localhost
;url-type = file

[segfind]
segment-url=https://metaserver.phy.syr.edu
;file:///home/lppekows/dmt/triggers/DQ_Segments/

[data]
; data conditioning parameters common to tmpltbank and inspiral
pad-data = 8
sample-rate = 8192
block-duration = 2176
segment-duration = 256

[ligo-data]
data-type = ht_real8
highpass-frequency = 35
cutoff-frequency = 39
dynamic-range-factor = 1.0e+20
spectrum-type = median_mean

```

```
[virgo-data]
data-type = ht_real4
highpass-frequency = 35
cutoff-frequency = 39
dynamic-range-factor = 1.0e+20
spectrum-type = median_mean
```

```
[geo-data]
; not needed for ringdowns
```

```
[tmplbank]
; not needed for ringdowns
```

```
[h1-tmplbank]
; not needed for ringdowns
```

```
[h2-tmplbank]
; not needed for ringdowns
```

```
[l1-tmplbank]
; not needed for ringdowns
```

```
[inspiral]
bank-max-mismatch = 0.03
bank-min-frequency = 50
bank-max-frequency = 2000
bank-min-quality = 2.1187
bank-max-quality = 20.0
bank-template-phase = 0
maximize-duration = 1
debug-level = 33
```

```
[h1-inspiral]
threshold = 5.5
```

```
[h2-inspiral]
threshold = 4.0
```

```
[l1-inspiral]
threshold = 5.5
```

```
[v1-inspiral]
threshold = 5.0
```

```
[no-veto-inspiral]
; not needed for ringdowns
```

```
[veto-inspiral]
; not needed for ringdowns
```

```
[coherent-inspiral]
```

```

; not needed for ringdowns

[inca]
; not needed for ringdowns

[thinca]
debug-level = 33
multi-ifo-coinc =
maximization-interval = .5
parameter-test = ds_sq_fQt
h1-ds_sq-accuracy = 0.4
h2-ds_sq-accuracy = 0.4
l1-ds_sq-accuracy = 0.4
v1-ds_sq-accuracy = 0.4

[thinca-slide]
; time slide parameters
h1-slide = 0
h2-slide = 0
l1-slide = 5
v1-slide = 10

[trigbank]
; not needed for ringdowns

[sire]
; not needed for ringdowns

[sire-inj]
; not needed for ringdowns

[coire]
cluster-time = 10000
coinc-stat = bitten_1
h1-bittenl-a = 2
h1-bittenl-b = 2.2
h2-bittenl-a = 2
h2-bittenl-b = 2.2
sort-triggers =
h1-h2-distance-cut =
h1-kappa = 2
h2-kappa = 2

[coire-1]
; not needed for ringdowns

[coire-2]
; not needed for ringdowns

[coire-2-cat-2]
; not needed for ringdowns

```

```

[coire-2-cat-3]

[coire-2-cat-4]
; not needed for ringdowns

[coire-inj]
injection-window = 15

[cohbank]
; not needed for ringdowns

[trigbank-coherent]
; not needed for ringdowns

[inspiral-coherent]
; not needed for ringdowns

[cohinspbank]
; not needed for ringdowns

[chia]
; not needed for ringdowns

[chia-inj]
; not needed for ringdowns

[cohire]
; not needed for ringdowns

[cohire-inj]
; not needed for ringdowns

; PLOTTING SECTIONS
;;;;;;;;;;;;;;;;;;;;;;;;

[plotnumtemplates]
; not needed for ringdowns

[plotnumtemplates-meta]
; not needed for ringdowns

[plotinspiralrange]
; not needed for ringdowns

[plotinspiralrange-meta]
; not needed for ringdowns

[plotinspiral]
; not needed for ringdowns

```

[plotinspiral-meta]
; not needed for ringdowns

[plotthinca]
; not needed for ringdowns

[plotthinca-meta]
; not needed for ringdowns

[plotinjnum]
; not needed for ringdowns

[plotinjnum-meta]
; not needed for ringdowns

[plotethinca]
; not needed for ringdowns

[plotethinca-meta]
; not needed for ringdowns

[plotinspmissed]
; not needed for ringdowns

[plotinspmissed-meta]
; not needed for ringdowns

[plotinspfound]
; not needed for ringdowns

[plotcoinmissed]
; not needed for ringdowns

[ploteffdistcut]
; not needed for ringdowns

[ploteffdistcut-meta]
; not needed for ringdowns

[plotinspinj]
; not needed for ringdowns

[plotinspinj-meta]
; not needed for ringdowns

[plotsnrchi]
; not needed for ringdowns

[plotsnrchi-meta]
; not needed for ringdowns

```

; FOLLOWUP SECTIONS
;;;;;;;;;;;;;;;;;;;;;;;;;
[triggers]
; not needed for ringdowns

[q-datafind]
; not needed for ringdowns

[qscan]
; not needed for ringdowns

[q-hoft-datafind]
; not needed for ringdowns

[qscan-hoft]
; not needed for ringdowns

[plots]
; detail of followup plots

[output]
; details of follow-up output

[seg]
; details of follow-up segments

```

The following configuration file is the file used for the playground post-processing.

```

; online inspiral pipeline configuration script.
;
;
; this is the configuration file for the inspiral DAG generation program
; lalapps_cbc_pipedown that creates a condor DAG to run the cbc
; pipedown dag.

[condor]
;  setup of condor universe and location of executables
universe = vanilla
thinca_to_coinc = /home/scaudill/opt/ringdown_
branch/bin/ligolw_rinca_to_coinc
ligolw_sqlite = /home/scaudill/opt/ringdown_
branch/bin/ligolw_sqlite
dbsimplify = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_dbsimplify
repop_coinc = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_repop_coinc
compute_durations = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_compute_durations
dbaddinj = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_dbaddinj

```

```

injfind = /home/scaudill/opt/ringdown_branch/
bin/lalapps_cbc_injfind
dbinjfind = /home/scaudill/opt/ringdown_branch/
bin/ligolw_dbinjfind
cluster_coins = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_cluster_coins
mvsc_get_doubles = /home/scaudill/opt/ringdown_
branch/bin/mvsc_get_doubles_ringdown
mvsc_train_forest = /usr/bin/SprBaggerDecisionTreeApp
mvsc_use_forest = /usr/bin/SprOutputWriterApp
mvsc_update_sql = /home/scaudill/opt/ringdown_
branch/bin/mvsc_update_sql
mvsc_dag = /home/scaudill/opt/ringdown_branch/
bin/mvsc_dag
cfar = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_cfar
printlc = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_printlc
printsims = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_printsims
printmissed = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_printmissed
minifollowups = /bin/true
plotslides = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_plotslides
plotcumhist = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_plotcumhist
plotifar = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_plotifar
plotfm = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_plotfm
search_volume = /bin/true
search_upper_limit = /bin/true

```

[pipeline]

```

; set node-tmp-dir to be a local directory that exists
; on all the slave nodes
node-tmp-dir = /usr1/scaudill
search-file-tag = RING_FIRST
coinc-file-tag = RINCA_FIRST
coinc-slide-file-tag = RINCA_SLIDE_FIRST

```

[input]

```

; options common to all commands
ihope-segments-directory = /home/scaudill/Projects/ringdown/S5/
s5_ihope_runs/months_3-4/820555271-825955331/segments

```

[plot_input]

```

; options common to plotting commands
enable-output =

```

```

[thinca_to_coinc]
; set options for thinca_to_coinc jobs
;search = low_mass
;lars-id =
statistic=snr

[ligolw_sqlite]
; set options for ligolw_sqlite jobs
;preserve_ids =
;replace =

[dbsimplify]
; set options for dbsimplify jobs
;vacuum =

[repop_coinc]
snr_sq =
ranking_stat =
ranking_stat2 =
freq_param =
q_param =
mass_param =
a_param =
kappa_param =
snr_ratio_param =

[snr_sq]
snrl-table = snrl_ringdown
coinc-table = coinc_ringdown
output-column = snr_sq
statistic-type = snr_sq

[ranking_stat]
; set options for the ranking_stat repop_coinc jobs
snrl-table = snrl_ringdown
coinc-table = coinc_ringdown
output-column = snr
statistic-type = chopped1
chopla = 2.0
choplb = 2.2

[ranking_stat2]
; set options for the ranking_stat repop_coinc jobs
snrl-table = snrl_ringdown
coinc-table = coinc_ringdown
output-column = chopped1_snr
statistic-type = chopped1forall
chopla = 2.0
choplb = 2.2
choplc = 0.75

```

```
[freq_param]
; set options for the frequency repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = frequency
statistic-type = snr_mean
param = frequency
```

```
[q_param]
; set options for the quality repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = quality
statistic-type = snr_mean
param = quality
```

```
[mass_param]
; set options for the mass repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = mass
statistic-type = snr_mean
param = mass
```

```
[a_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = spin
statistic-type = snr_mean
param = spin
```

```
[kappa_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = kappa
statistic-type = ratio
ratio-column = eff_dist
ratio-default = 1.0
ifo1 = H1
ifo2 = H2
```

```
[snr_ratio_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = snr_ratio
statistic-type = ratio
ratio-column = snr
ratio-default = 2.0
```

```

ifo1 = H1
ifo2 = H2

[compute_durations]
; set options for compute_durations jobs
live-time-program = rinca

[dbaddinj]
; set options for dbaddinj jobs
sim-table = sim_ringdown

[injfind]
; set options for injfind jobs
match-algorithm = ringdown

[dbinjfind]
exact_ring =
nearby_ring =
exact_insp =
nearby_insp =

[exact_ring]
simulation-table = sim_ringdown
recovery-table = sngl_ringdown
match-criteria = startTime:startTime:0.005
map-label = ring_exact
rough-match = geocent_start_time:start_time:10
check-all-data = startTime:2
rough-all-data-check = start_time:10

[nearby_ring]
simulation-table = sim_ringdown
recovery-table = sngl_ringdown
match-criteria = startTime:startTime:1.0
map-label = ring_nearby
rough-match = geocent_start_time:start_time:10

[exact_insp]
simulation-table = sim_inspiral
recovery-table = sngl_ringdown
match-criteria = endTime:startTime:0.005
map-label = insp_exact
rough-match = geocent_end_time:start_time:10
check-all-data = startTime:2
rough-all-data-check = start_time:10

[nearby_insp]
simulation-table = sim_inspiral
recovery-table = sngl_ringdown
match-criteria = endTime:startTime:1.0
map-label = insp_nearby

```

```

rough-match = geocent_end_time:start_time:10

[cluster_coincs]
choppedl_snr_cluster =
snr_sq_cluster =

[choppedl_snr_cluster]
; set static options for cluster_coinc jobs
; following are required options
param-name = snr_ratio
param-ranges = (1.0,1000000.0)
ranking-table = coinc_ringdown
ranking-stat = choppedl_snr
rank-by = MAX
time-column = start_time
cluster-window = 10000
; following are optional
exclude-coincs = '[all in h1,h2];[h1,h2 in all];[h2,l1 in h1,h2,l1]'
group-by-ifos =
vacuum =
debug =
verbose =

[snr_sq_cluster]
ranking-table = coinc_ringdown
ranking-stat = snr_sq
rank-by = MAX
time-column = start_time
cluster-window = 10000
; following are optional
group-by-ifos =
vacuum =
debug =
verbose =

[cfar-uncombined]
; set static options for cfar jobs that
; compute uncombined false alarm rates
ranking-table = coinc_event
output-column = false_alarm_rate
output-table = coinc_ringdown
ranking-stat = likelihood
ifos-table = coinc_ringdown
rank-by = MAX
group-by-multiplicity =

[cfar-combined]
; set static options for cfar jobs that
; compute combined false alarm rates
ranking-table = coinc_ringdown
output-column = combined_far

```

```

ranking-stat = false_alarm_rate
rank-by = MIN

[cbc_print]
; set static options common for printlc, printsims, and printmisses
daily-ihope-pages-location = https://ldas-jobs.ligo.caltech.edu/
~cbc/ihope_daily

[printlc]
; set static options for printlc jobs
coinc-table = coinc_ringdown
ranking-stat = combined_far
rank-by = MIN
convert-durations = days
limit = 20
; following are needed for getting single ifo information
get-sngl-info =
sngl-table = sngl_ringdown
time-column = start_time
; if set the output format to anything but xml,
;minifollowups will not run
output-format = html
columns = rank,combined_far,fap,fap_1yr,sngl_start_time,
sngl_ifo__Px_click_for_eelog_xP_,instruments_on,sngl_frequency,
sngl_quality,sngl_mass,sngl_spin,sngl_snr,duration

[printsims]
; set static options for printsims jobs
ranking-stat = combined_far
rank-by = MIN
sort-by = injected_decisive_distance
convert-durations = days
get-sngl-info =
sngl-table = sngl_ringdown
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
time-column = start_time
map-label = ring_nearby
param-name = combined_far
param-ranges = '!0'
output-format = html
include-only-coincs = '[ALLinH1,H2,L1]'
columns = rank,recovered_combined_far,injected_gps_time,
injected_gps_time_ns,injected_decisive_distance,sngl_ifo,
instruments_on,recovered_start_time,sim_tag,
injected_mass1,injected_mass2,injected_spin1z,
injected_spin2z,injected_frequency,
sngl_frequency,injected_quality,sngl_quality,sngl_snr

[printmisses]
; set static options for printmisses jobs

```

```

simulation-table = sim_ringdown
recovery-table = coinc_ringdown
map-label = ring_nearby
livetime-program = rinca
output-format = html
limit = 20
columns = rank,decisive_distance,gps_time,instruments_on,
eff_dist_h,eff_dist_l,frequency,quality,mass,spin,sim_tag

```

[minifollowups]

```

; set static options for minifollowup jobs
; not used for ringdown

```

[plotfm]

```

; establish what plotfm jobs to do; the following options
; must have corresponding sections with options given.
; each section name will be added to the plotfm user_tag

```

```

fm_injected_params_imr =
fm_lin_injected_params_ringdown =
fm_log_injected_params_ringdown =
fm_lin_detected_vs_injected_imr =
fm_log_detected_vs_injected_imr =
fm_lin_detected_vs_injected_ringdown =
fm_log_detected_vs_injected_ringdown =
fm_loglog_detected_vs_injected_ringdown =
fm_lin_accuracy_plots_imr =
fm_lin_accuracy_plots_ringdown =
fm_freq_accuracy_plots_ringdown =
fm_injdist_plots_imr =
fm_log_injdist_plots_ringdown =
fm_loglog_injdist_plots_ringdown =

```

[fm_injected_params_imr]

```

; plot injected imr parameters
variables = 'injected_mass1:Injected Mass 1($M_\odot$);
injected_mass2:Injected Mass 2($M_\odot$)'
--variables 'injected_mass1 +
injected_mass2:Injected Mass1 + Mass2($M_\odot$);
(injected_mass1/injected_mass2):Injected Mass
Ratio($m1/m2$)'
--variables 'injected_mchirp:Injected Chirp Mass
($M_\odot^{\frac{2}{5}}$);
(injected_mass1/injected_mass2):Injected Mass Ratio($m1/m2$)'
--variables 'injected_spin1z:Injected Spin1z; injected_
spin2z:Injected Spin2z'
logz =
colorbar =
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
simulation-table = sim_inspiral
recovery-table = coinc_ringdown

```

```

livetime-program = rinca
map-label = insp_nearby

[fm_lin_injected_params_ringdown]
; plot injected ringdown parameters
variables = 'injected_mass:Injected Ringdown Mass ( $M_{\odot}$ );
injected_spin:Injected Ringdown Spin' --variables 'injected_mass:
Injected Ringdown Mass( $M_{\odot}$ ); injected_quality:
Injected Ringdown Quality Factor'
logz =
colorbar =
ranking-stat = 'combined_far:Combined FAR ( $\text{yr}^{-1}$ )'
rank-by = MIN
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_log_injected_params_ringdown]
; plot injected ringdown parameters
variables = 'injected_frequency:Injected Ringdown
Frequency ( $\text{Hz}$ );
injected_quality: Injected Ringdown Quality Factor'
--variables 'injected_frequency:
Injected Ringdown Frequency ( $\text{Hz}$ ); injected_mass:
Injected Ringdown Mass
( $M_{\odot}$ )' --variables 'injected_frequency:Injected
Ringdown Frequency ( $\text{Hz}$ );
injected_spin: Injected Ringdown Spin'
logx =
logz =
colorbar =
ranking-stat = 'combined_far:Combined FAR ( $\text{yr}^{-1}$ )'
rank-by = MIN
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_lin_detected_vs_injected_imr]
; detected vs injected imr parameters
variables = '(injected_gps_time+injected_gps_time_
ns*1e-9):
Injected GPS Time( $s$ ); (recovered_start_time+recovered_
start_time_ns*1e-9):
Recovered Ringdown Start Time( $s$ )'
ranking-stat = 'combined_far:Combined FAR ( $\text{yr}^{-1}$ )'
rank-by = MIN
logz =
colorbar =
plot-x-function = x

```

```

simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_log_detected_vs_injected_imr]
; detected vs injected imr parameters
variables = 'injected_mass1 + injected_mass2: Injected Mass1 +
Mass2($M_\odot$); recovered_frequency:Detected
Ringdown Frequency($Hz$)'
--variables 'injected_mchirp: Injected Chirp Mass
($M_\odot^{\frac{2}{5}}$);
recovered_frequency:Detected Ringdown
Frequency($Hz$)' --variables
'injected_mass1 + injected_mass2: Injected Mass1
+ Mass2($M_\odot$);
recovered_snr:SNR' --variables 'injected_mchirp:
Injected Chirp Mass
($M_\odot^{\frac{2}{5}}$); recovered_snr: SNR'
--variables 'injected_mass1 +
injected_mass2: Injected Mass1 + Mass2($M_\odot$);
recovered_choppedl_snr:
New Triple SNR' --variables 'injected_mchirp:
Injected Chirp Mass
($M_\odot^{\frac{2}{5}}$); recovered_choppedl_
snr: New Triple SNR'
ranking-stat = 'combined_far:Combined FAR
($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_lin_detected_vs_injected_ringdown]
; detected vs injected ringdown parameters
variables = 'injected_quality:Injected Ringdown
Quality Factor;
recovered_quality:Detected Ringdown Quality
Factor' --variables
'injected_mass:Injected Ringdown Mass($M_\odot$);
recovered_mass:
Detected Ringdown Mass($M_\odot$)' --variables
'injected_spin:
Injected Ringdown Spin; recovered_spin:Detected
Ringdown Spin'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN

```

```

logz =
colorbar =
plot-x-function = x
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
lifetime-program = rinca
map-label = ring_nearby

[fm_log_detected_vs_injected_ringdown]
; detected vs injected ringdown parameters
variables = 'injected_mass:Injected Ringdown
Mass($M_\odot$);
recovered_frequency:Detected Ringdown
Frequency($Hz$)'
--variables 'injected_quality:Injected Ringdown Quality
Factor;
recovered_snr: SNR' --variables 'injected_mass:Injected
Ringdown
Mass($M_\odot$); recovered_snr: SNR' --variables
'injected_spin:
Injected Ringdown Spin; recovered_snr: SNR'
--variables 'injected_quality:
Injected Ringdown Quality Factor; recovered_
choppedl_snr:
New Triple SNR' --variables 'injected_mass:Injected
Ringdown
Mass($M_\odot$); recovered_choppedl_snr: New
Triple SNR'
--variables 'injected_spin:Injected Ringdown Spin;
recovered_choppedl_snr:
New Triple SNR'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
lifetime-program = rinca
map-label = ring_nearby

[fm_loglog_detected_vs_injected_ringdown]
; detected vs injected ringdown frequency
variables = 'injected_frequency:Injected Ringdown
Frequency ($Hz$); recovered_frequency:Detected
Ringdown Frequency ($Hz$)' --variables 'injected_frequency:
Injected Ringdown Frequency ($Hz$); recovered_snr:SNR'
--variables 'injected_frequency:Injected Ringdown Frequency
($Hz$); recovered_choppedl_snr:New Triple SNR'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN

```

```

logx =
logy =
logz =
colorbar =
plot-x-function = x
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_lin_accuracy_plots_imr]
; plot imr accuracy
variables = 'injected_mass1 + injected_mass2: Injected
Mass1 + Mass2($M_\odot$); 1000*((recovered_start_time+
recovered_start_time_ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)):
Recovered - Injected Start Time ($ms$)' --variables
'injected_mchirp:Injected
Chirp Mass($M_\odot^{\frac{2}{5}}$); 1000*((recovered_
start_time+
recovered_start_time_ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)):
Recovered - Injected Start Time ($ms$)' --variables
'injected_spin1z:Injected
Spin1z; 1000*((recovered_start_time+recovered_
start_time_ns*1e-9)-
(injected_gps_time+injected_gps_time_ns*1e-9)):
Recovered - Injected Start Time ($ms$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logz =
colorbar =
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_lin_accuracy_plots_ringdown]
; plot ringdown accuracy
variables = 'recovered_snr: SNR; 1000*((recovered_
start_time+
recovered_start_time_ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)):
Det Start Time - Inj End Time($ms$)' --variables
'recovered_chopped1_snr:
SNR; 1000*((recovered_start_time+recovered_
start_time_ns*1e-9)-
(injected_gps_time+injected_gps_time_ns*1e-9)):
Det Start Time - Inj End
Time($ms$)' --variables '(recovered_quality-injected_

```

```

quality): Det-Inj
Ringdown Quality Factor; 1000*((recovered_start_
time+recovered_start_time_ns*1e-9)-
(injected_gps_time+injected_gps_time_ns*1e-9)):
  Det Start Time - Inj End Time($ms$)'
--variables '(recovered_mass-injected_mass):
Det-Inj Ringdown Mass($M_\odot$);
1000*((recovered_start_time+recovered_start_
time_ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)): Det Start Time -
  Inj End Time($ms$)' --variables
'(recovered_mass-injected_mass):Det-Inj
Ringdown Mass($M_\odot$);
(recovered_quality-injected_quality):Det-Inj
Ringdown Quality Factor' --variables
'recovered_frequency-injected_frequency:
Det-Inj Ringdown Frequency ($Hz$);
recovered_mass-injected_mass: Det-Inj
  Ringdown Mass($M_\odot$)' --variables
'recovered_frequency-injected_frequency:Det-Inj
  Ringdown Frequency ($Hz$);
recovered_quality-injected_quality:Det-Inj Ringdown
Quality Factor'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logz =
colorbar =
plot-x-function = 0
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_freq_accuracy_plots_ringdown]
; plot ringdown frequency accuracy
variables = 'injected_frequency: Injected Ringdown
  Frequency ($Hz$);
1000*((recovered_start_time+recovered_start_time_
ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)): Recovered Start Time -
  Injected End Time($ms$)'
--variables 'injected_frequency: Injected Ringdown
Frequency ($Hz$);
(injected_frequency - recovered_frequency)/((injected_
frequency+
recovered_frequency)/2.0): Fractional Ringdown
  Frequency Accuracy
($\frac{\text{freq}_{\text{inj}}-\text{freq}_{\text{det}}}{(\text{freq}_{\text{inj}}+\text{freq}_{\text{det}})/2}$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logx =

```

```

logz =
colorbar =
plot-x-function = 0
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_injdist_plots_imr]
; plot injected decisive distance v. imr parameters
variables = 'injected_gps_time:gps_days; injected_
decisive_distance:
Injected Decisive Distance ($Mpc$)' --variables
'injected_mass1+
injected_mass2:Injected Mass1+Mass2($M_\odot$);
injected_decisive_distance:
Injected Decisive Distance ($Mpc$)' --variables
'injected_mchirp:Injected Chirp
Mass($M_\odot^{\frac{2}{5}}$);injected_decisive_
distance:Injected Decisive
Distance ($Mpc$)' --variables '(injected_mass1/injected_
mass2):Injected Mass
Ratio; injected_decisive_distance:Injected Decisive
Distance ($Mpc$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_log_injdist_plots_ringdown]
; plot injected decisive distance v. injected ringdown
mass
variables = 'injected_mass: Injected Ringdown
Mass($M_\odot$);
injected_decisive_distance:Injected Decisive
Distance ($Mpc$)'
--variables 'injected_quality:Injected Ringdown
Quality Factor;
injected_decisive_distance:Injected Decisive
Distance ($Mpc$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
simulation-table = sim_ringdown

```

```

recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_loglog_injdist_plots_ringdown]
; plot injected decisive distance v. injected ringdown
  frequency
variables = 'injected_frequency: Injected Frequency
($Hz$);
injected_decisive_distance: Injected Decisive Distance
($Mpc$)'
ranking-stat = 'combined_far: Combined FAR
($yr^{-1}$)'
rank-by = MIN
logx =
logy =
logz =
colorbar =
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[plotslides]
; set static options for plotslide jobs
coinc-table = coinc_ringdown

[plotcumhist]
; set static options for plotcumhist jobs
num-bins = 20
ranking-stat = choppedl_snr
rank-by = MAX
coinc-table = coinc_ringdown

[plotifar]
; set static options for plotifar jobs
plot-uncombined =
group-by-multiplicity =
plot-combined =
plot-slides =
show-min-bkg =
show-max-bkg =
show-two-sigma-error =
coinc-table = coinc_ringdown

[mvsc_dag]
number-of-trees = 200
criterion-for-optimization = 6
leaf-size = 5
sampled-parameters = 12
all-instruments = H1,H2,L1

```

```
[mvsc_get_doubles]
number = 10
exact-tag = ring_exact
nearby-tag = ring_nearby
start-time = 0
end-time = 999999999
```

Appendix B: Assignment of Copyright

IOP Journals

Assignment of copyright form - IOP Publishing Limited

1. IOP Publishing Limited ("IOP") agrees to publish:
Manuscript Title: ("the Article") written by Names of all Authors:
..... ("the Named Authors") in the following journal ("the Journal")

2. Transfer of Copyright Agreement

2.1 On acceptance for publication the undersigned author(s) ("Author") of the Article assigns exclusively to IOP worldwide copyright in the Article for the full term and for all media and formats in all material published as part of the Article, which expression includes but is not limited to the text, abstract, tables, figures, graphs, video abstracts and other multimedia content but excludes any other item referred to as supplementary material.

2.2 If any of the Named Authors are Government employees, on acceptance for publication the Author shall grant IOP a royalty free exclusive licence for the full term of copyright for all media and formats to do in relation to the Article all acts restricted by copyright worldwide.

2.3 On acceptance for publication the Author shall grant IOP a royalty free non-exclusive licence for the full term of copyright for all media and formats to do in relation to any supplementary material deemed to be part of the Article all acts restricted by copyright worldwide.

3. Author Rights

3.1 IOP grants the Named Authors the rights specified in 3.2 and 3.3. All such rights must be exercised for non-commercial purposes, if possible should display citation information and IOP's copyright notice, and for electronic use best efforts must be made to include a link to the on-line abstract in the Journal. Exercise of the rights in 3.3 additionally must not use the final published IOP format but the Named Author's own format (which may include amendments made following peer review).

3.2 The rights are:

3.2.1 To make copies of the Article (all or part) for teaching purposes;

3.2.2 To include the Article (all or part) in a research thesis or dissertation;

3.2.3 To make oral presentation of the Article (all or part) and to include a summary and/or highlights of it in papers distributed at such presentations or in conference proceedings; and

3.2.4 All proprietary rights other than copyright.

3.3 The additional rights are to:

3.3.1 Use the Article (all or part) without modification in personal compilations or publications of a Named Author's own works (provided not created by third party publisher);

3.3.2 Include the Article (all or part) on a Named Author's own personal web site;

3.3.3 Include the Article (all or part) on web sites of the Institution (including its repository) where a Named Author worked when research for the Article was carried out; and

3.3.4 No sooner than 12 months after publication to include the Article (all or part) on third party web sites including e-print servers, but not on other publisher's web sites.

Figure 6.1: Assignment of copyright form from IOP Publishing Limited. Section 3.2.2 grants the author rights to include the article (all or part) in a research thesis or dissertation. This article can be found online: <http://iopscience.iop.org/0264-9381/29/9/095016/>

Appendix C: S5 Search Configuration Files

The following configuration file is the file used for the S5 ringdown search.

```
; ihope configuration script for the ringdown pipeline.
;
;
; GENERAL SECTIONS
;;;;;;;;;;;;;;;;;;;;;;;;;

[segments]
l1-analyze = L1:Science,L1:Injection
h1-analyze = H1:Science,H1:Injection
h2-analyze = H2:Science,H2:Injection
veto-def-server-url = https://www.lsc-group.phys.uwm.edu/ligovirgo/
cbc/public/segments/S5/
veto-def-file = H1H2L1-CBC_RINGDOWN_DQ_VETO_CATEGORIES-
815155213-60076801.xml
;v1-analyze = STATEVEC.5.12,STATEVEC.5.8
; server where veto-definer xml file located
veto-categories = 2,3,4

[ifo-details]
l1-data =
h1-data =
h2-data =
;v1-data =
one-ifo =
two-ifo =
three-ifo =
;four-ifo =

[hipe-arguments]
datafind =
inspiral =
output-segs =
write-script =
inspinj =
ringdown =
coincidence =
;coire-coincidence =
;summary-first-coinc-triggers =

[condor]
; setup of condor universe and location of executables
universe          = standard
hipe              = /home/scaudill/opt/ringdown_branch/bin/
lalapps_inspiral_hipe
plot              = /bin/false
```

```

follow                = /bin/false
datafind              = /home/scaudill/opt/ringdown_branch/bin/
ligo_data_find       = /home/scaudill/opt/ringdown_branch/bin/
segfind              = /home/scaudill/opt/ringdown_branch/bin/
ligolw_segment_query
segs_from_cats       = /home/scaudill/opt/ringdown_branch/
bin/ligolw_segments_from_cats
ligolw_segments_compat = /home/scaudill/opt/ringdown_
branch/bin/ligolw_segments_compat
ligolw_print         = /home/scaudill/opt/ringdown_branch/bin/
ligolw_print
ligolw_add           = /home/scaudill/opt/ringdown_branch/bin/
ligolw_add
tmpltbank            = /bin/false
inspiral             = /home/scaudill/opt/ringdown_branch/bin/
lalapps_ring
; both inspiral and rinj are necessary if you want to run
inspiral and ringdown injections
inspiral             = /home/scaudill/opt/ringdown_branch/bin/
lalapps_inspiral
rinj                 = /home/scaudill/opt/ringdown_branch/bin/
lalapps_rinj
inca                 = /bin/false
thinca               = /home/scaudill/opt/ringdown_branch/bin/
lalapps_rinca
trigbank             = /bin/false
cohbank              = /bin/false
frjoin               = /bin/false
chia                 = /bin/false
sire                 = /bin/false
coire                = /home/scaudill/opt/ringdown_branch/bin/
lalapps_coincringread
tmpltbank32         = /bin/false
inspiral32           = /bin/false
inspiral32
inspiral32           = /bin/false
inca32               = /bin/false
thinca32             = /bin/false
trigbank32          = /bin/false
cohbank32            = /bin/false
frjoin32             = /bin/false
chia32               = /bin/false
sire32               = /bin/false
coire32              = /bin/false
plotinspiral         = /bin/false
plotthinca           = /bin/false
plotnumtemplates    = /bin/false
plotinjn            = /bin/false
plotethinca         = /bin/false
plotinspmiss        = /bin/false
plotinspfound       = /bin/false
plotinspfound       = /bin/false

```

```
plotcoinmissed = /bin/false
plotinspinj    = /bin/false
plotsnrchi     = /bin/false
plotinspiralrange = /bin/false
ploteffdistcut = /bin/false
cohbank        = /bin/false
cohinspbank    = /bin/false
frjoin         = /bin/false
chia           = /bin/false
cohire         = /bin/false
```

```
[condor-max-jobs]
```

```
;maximum number of jobs that condor should run of a
;given category
datafind= 10
thinca = 20
trigbank = 20
coire = 20
sire = 20
```

```
[ldgsubmitdax]
```

```
sites-file = ldg-sites.xml
properties-file = pegasus-properties.bundle
gsiftp = ldas-grid.ligo.caltech.edu
pool = cit
```

```
[sites]
```

```
eobnr_inj = local
ringdown_inj = local
phenom_inj = local
playground = local
full_data = local
```

```
[pipeline]
```

```
; tagging information for the configure script
version =
cvs-tag = $Name$
user-tag =
num-slides = 50
playground-data-mask = all_data
retry-subdag = 1
;Set matplotlibdir to a local directory to avoid NFS issues
; with the TeX part
;of matplotlib. Ensure this directory exists on all nodes!
;matplotlibdir =
```

```
; INSPINJ SECTIONS
```

```
;;;;;;;;;
```

```
[injections]
```

```
eobnr2_lin_a_inj = 11117085
```

```
eobnr2_lin_b_inj = 22227085
eobnr2_log_a_inj = 33337085
eobnr2_log_b_inj = 44447085
phenomb_lin_a_spinj = 55557085
phenomb_lin_b_spinj = 66667085
phenomb_log_a_spinj = 77777085
phenomb_log_b_spinj = 88887085
ringdown_ma_log_a_inj = 99997085
ringdown_ma_log_b_inj = 12127085
ringdown_qf_log_a_inj = 13137085
ringdown_qf_log_b_inj = 14147085
```

```
[eobnr2_lin_a_inj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = EOBNRv2pseudoFourPN
min-distance = 100
max-distance = 600000
d-distr = uniform
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =
```

```
[eobnr2_lin_b_inj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = EOBNRv2pseudoFourPN
min-distance = 100
max-distance = 600000
d-distr = uniform
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
```

```

max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =

[eobnrv2_log_a_inj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = EOBNRv2pseudoFourPN
min-distance = 10
max-distance = 1000000
d-distr = log10
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =

[eobnrv2_log_b_inj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = EOBNRv2pseudoFourPN
min-distance = 100
max-distance = 1000000
d-distr = log10
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0

```

```

disable-spin =
taper-injection = start
injection-type = EOBNR
write-sim-ring =

[phenomb_lin_a_spinj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = IMRPhenomBpseudoFourPN
min-distance = 100
max-distance = 600000
d-distr = uniform
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
enable-spin =
aligned =
min-spin1 = 0.0
max-spin1 = 0.85
min-spin2 = 0.0
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =

[phenomb_lin_b_spinj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = IMRPhenomBpseudoFourPN
min-distance = 100
max-distance = 600000
d-distr = uniform
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450

```

```

min-mratio = 0.1
max-mratio = 1.0
enable-spin =
aligned =
min-spin1 = 0.0
max-spin1 = 0.85
min-spin2 = 0.0
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =

[phenomb_log_a_spinj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = IMRPhenomBpseudoFourPN
min-distance = 10
max-distance = 1000000
d-distr = log10
l-distr = random
i-distr = uniform
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
enable-spin =
aligned =
min-spin1 = 0.0
max-spin1 = 0.85
min-spin2 = 0.0
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =

[phenomb_log_b_spinj]
f-lower = 10
time-step = 724.077
time-interval = 300
waveform = IMRPhenomBpseudoFourPN
min-distance = 100
max-distance = 1000000
d-distr = log10
l-distr = random
i-distr = uniform

```

```
min-mass1 = 10
max-mass1 = 400
min-mass2 = 10
max-mass2 = 400
m-distr = totalMassRatio
min-mtotal = 50
max-mtotal = 450
min-mratio = 0.1
max-mratio = 1.0
enable-spin =
aligned =
min-spin1 = 0.0
max-spin1 = 0.85
min-spin2 = 0.0
max-spin2 = 0.85
taper-injection = start
injection-type = PHENOM
write-sim-ring =
```

```
[ringdown_ma_log_a_inj]
waveform = Ringdown
inj-distr = 2
time-step = 724.077
time-interval = 300
minimum-mass = 50
maximum-mass = 900
minimum-spin = 0
maximum-spin = 0.99
minimum-distance = 1000
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN
```

```
[ringdown_ma_log_b_inj]
waveform = Ringdown
inj-distr = 2
time-step = 724.077
time-interval = 300
minimum-mass = 50
maximum-mass = 900
minimum-spin = 0
maximum-spin = 0.99
minimum-distance = 1000
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN
```

```
[ringdown_qf_log_a_inj]
waveform = Ringdown
inj-distr = 0
time-step = 724.077
```

```

time-interval = 300
minimum-quality = 2.1187
maximum-quality = 20
minimum-frequency = 50
maximum-frequency = 2000
minimum-distance = 1000
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN

[ringdown_qf_log_b_inj]
waveform = Ringdown
inj-distr = 0
time-step = 724.077
time-interval = 300
minimum-quality = 2.1187
maximum-quality = 20
minimum-frequency = 50
maximum-frequency = 2000
minimum-distance = 1000
maximum-distance = 1000000
epsilon = 0.01
injection-type = RINGDOWN

; HIPE SECTIONS
;;;;;;;;;;;;;;;;;;

[input]
ligo-channel = LSC-STRAIN
ligo-type = RDS_C03_L2
ligo-qscan-type = RDS_R_L1
max-thinca-segment = 3600
geo-channel =
geo-type =
virgo-type = HrecV2_16384Hz
virgo-channel =h_16384Hz

[calibration]
; location of the calibration cache and the cache files
; if using h(t) data, then no calibration cache required
path = /home/lgoggin/projects/ringdown/s5/calibration
L1 = L1_CAL_REF_DARM_ERR_S5_V1_U.cache
H1 = H1_CAL_REF_DARM_ERR_S5_V1_U.cache
H2 = H2_CAL_REF_DARM_ERR_S5_V1_U.cache

[datafind]
; type of data to use -- for LDAS-GRID, nothing required here.
;match=localhost
;url-type = file

[segfind]

```

```

segment-url=https://metaserver.phy.syr.edu
;file:///home/lppekows/dmt/triggers/DQ_Segments/

[data]
; data conditioning parameters common to tmpltbank and inspiral
pad-data = 8
sample-rate = 8192
block-duration = 2176
segment-duration = 256

[ligo-data]
data-type = ht_real8
highpass-frequency = 35
cutoff-frequency = 39
dynamic-range-factor = 1.0e+20
spectrum-type = median_mean

[virgo-data]
data-type = ht_real4
highpass-frequency = 35
cutoff-frequency = 39
dynamic-range-factor = 1.0e+20
spectrum-type = median_mean

[geo-data]
; not needed for ringdowns

[tmpltbank]
; not needed for ringdowns

[h1-tmpltbank]
; not needed for ringdowns

[h2-tmpltbank]
; not needed for ringdowns

[l1-tmpltbank]
; not needed for ringdowns

[inspiral]
bank-max-mismatch = 0.03
bank-min-frequency = 50
bank-max-frequency = 2000
bank-min-quality = 2.1187
bank-max-quality = 20.0
bank-template-phase = 0
maximize-duration = 1
debug-level = 33
[h1-inspiral]
threshold = 5.5

```

```

[h2-inspiral]
threshold = 4.0

[l1-inspiral]
threshold = 5.5

[v1-inspiral]
threshold = 5.0

[no-veto-inspiral]
; not needed for ringdowns

[veto-inspiral]
; not needed for ringdowns

[coherent-inspiral]
; not needed for ringdowns

[inca]
; not needed for ringdowns

[thinca]
debug-level = 33
multi-ifo-coinc =
maximization-interval = 0.5
parameter-test = ds_sq_fQt
h1-ds_sq-accuracy = 0.4
h2-ds_sq-accuracy = 0.4
l1-ds_sq-accuracy = 0.4
v1-ds_sq-accuracy = 0.4
h1-h2-consistency =
h1-snr-cut = 15

; only turn on if you want no vetoes applied
;[thinca-2]

[thinca-slide]
; time slide parameters
h1-slide = 0
h2-slide = 0
l1-slide = 5
v1-slide = 10

```

The following configuration file is the file used for the S5 ringdown post-processing.

```

; online inspiral pipeline configuration script.
;
; $Id$
;
; this is the configuration file for the inspiral DAG generation

```

```
; program
; lalapps_cbc_pipedown that creates a condor DAG to run the cbc
; pipedown dag.
```

```
[condor]
; setup of condor universe and location of executables
universe = vanilla
thinca_to_coinc = /home/scaudill/opt/ringdown_branch/
bin/ligolw_rinca_to_coinc
ligolw_sqlite = /home/scaudill/opt/ringdown_branch/
bin/ligolw_sqlite
dbsimplify = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_dbsimplify
repop_coinc = /home/scaudill/opt/ringdown_branch/
bin/ligolw_cbc_repop_coinc
compute_durations = /home/scaudill/opt/ringdown_
branch/bin/
ligolw_cbc_compute_durations
dbaddinj = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_dbaddinj
injfind = /home/scaudill/opt/ringdown_branch/bin/
lalapps_cbc_injfind
dbinjfind = /home/scaudill/opt/ringdown_branch/
bin/ligolw_dbinjfind
cluster_coincs = /home/scaudill/opt/ringdown_
branch/bin/ligolw_cbc_cluster_coincs
mvsc_get_doubles = /home/scaudill/opt/ringdown_
branch/bin/
mvsc_get_doubles_ringdown
mvsc_train_forest = /usr/bin/SprBaggerDecisionTreeApp
mvsc_use_forest = /usr/bin/SprOutputWriterApp
mvsc_update_sql = /home/scaudill/opt/ringdown_
branch/bin/mvsc_update_sql
mvsc_dag = /home/scaudill/opt/ringdown_branch/
bin/mvsc_dag
cfar = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_cfar
printlc = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_printlc
printsims = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_printsims
printmisses = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_printmisses
minifollowups = /bin/true
plotslides = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_plotslides
plotcumhist = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_plotcumhist
plotifar = /home/scaudill/opt/ringdown_branch/bin/
ligolw_cbc_plotifar
plotfm = /home/scaudill/opt/ringdown_branch/bin/
```

```

ligolw_cbc_plotfm
search_volume = /bin/true
search_upper_limit = /bin/true

[pipeline]
; set node-tmp-dir to be a local directory that exists
; on all the slave nodes
node-tmp-dir = /usr1/scaudill
search-file-tag = RING_FIRST
coinc-file-tag = RINCA_FIRST
coinc-slide-file-tag = RINCA_SLIDE_FIRST

[input]
; options common to all commands
ihope-segments-directory = /home/scaudill/Projects/ringdown/S5/
s5_ihope_runs_final/months_15-16/852393970-857232370/segments

[plot_input]
; options common to plotting commands
enable-output =

[thinca_to_coinc]
; set options for thinca_to_coinc jobs
;search = low_mass
;lars-id =
statistic=snr

[ligolw_sqlite]
; set options for ligolw_sqlite jobs
;preserve_ids =
;replace =

[dbssimplify]
; set options for dbssimplify jobs
;vacuum =

[repop_coinc]
snr_sq =
ranking_stat =
ranking_stat2 =
freq_param =
q_param =
mass_param =
a_param =
kappa_param =
snr_ratio_param =

[snr_sq]
snr-table = snr_ringdown
coinc-table = coinc_ringdown
output-column = snr_sq

```

```

statistic-type = snr_sq

[ranking_stat]
; set options for the ranking_stat repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = snr
statistic-type = choppedl
chopla = 2.0
choplb = 2.2

[ranking_stat2]
; set options for the ranking_stat repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = choppedl_snr
statistic-type = choppedlforall
chopla = 2.0
choplb = 2.2
choplc = 0.75

[freq_param]
; set options for the frequency repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = frequency
statistic-type = snr_mean
param = frequency

[q_param]
; set options for the quality repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = quality
statistic-type = snr_mean
param = quality

[mass_param]
; set options for the mass repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = mass
statistic-type = snr_mean
param = mass

[a_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = spin
statistic-type = snr_mean

```

```

param = spin

[kappa_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = kappa
statistic-type = ratio
ratio-column = eff_dist
ratio-default = 1.0
ifo1 = H1
ifo2 = H2

[snr_ratio_param]
; set options for the spin repop_coinc jobs
sngl-table = sngl_ringdown
coinc-table = coinc_ringdown
output-column = snr_ratio
statistic-type = ratio
ratio-column = snr
ratio-default = 2.0
ifo1 = H1
ifo2 = H2

[compute_durations]
; set options for compute_durations jobs
live-time-program = rinca

[dbaddinj]
; set options for dbaddinj jobs
sim-table = sim_ringdown

[injfind]
; set options for injfind jobs
match-algorithm = ringdown
;time-window = 1.
;force =
;comment = None

[dbinjfind]
exact_ring =
nearby_ring =
exact_insp =
nearby_insp =

[exact_ring]
simulation-table = sim_ringdown
recovery-table = sngl_ringdown
match-criteria = startTime:startTime:0.01
map-label = ring_exact
rough-match = geocent_start_time:start_time:10

```

```

check-all-data = startTime:0.5
rough-all-data-check = start_time:10

[nearby_ring]
simulation-table = sim_ringdown
recovery-table = sngl_ringdown
match-criteria = startTime:startTime:1.0
map-label = ring_nearby
rough-match = geocent_start_time:start_time:10

[exact_insp]
simulation-table = sim_inspirial
recovery-table = sngl_ringdown
match-criteria = endTime:startTime:0.01
map-label = insp_exact
rough-match = geocent_end_time:start_time:10
check-all-data = startTime:0.5
rough-all-data-check = start_time:10

[nearby_insp]
simulation-table = sim_inspirial
recovery-table = sngl_ringdown
match-criteria = endTime:startTime:1.0
map-label = insp_nearby
rough-match = geocent_end_time:start_time:10

[cluster_coincs]
choppedl_snr_cluster =
snr_sq_cluster =

[choppedl_snr_cluster]
; set static options for cluster_coinc jobs
; following are required options
param-name = snr_ratio
param-ranges = (1.0,1000000.0)
ranking-table = coinc_ringdown
ranking-stat = choppedl_snr
rank-by = MAX
time-column = start_time
cluster-window = 10000
; following are optional
exclude-coincs = '[all in h1,h2];[h1,h2 in all];[h2,l1 in h1,h2,l1]'
group-by-ifos =
vacuum =
debug =
verbose =

[snr_sq_cluster]
ranking-table = coinc_ringdown
ranking-stat = snr_sq
rank-by = MAX

```

```

time-column = start_time
cluster-window = 10000
; following are optional
group-by-ifos =
vacuum =
debug =
verbose =

[cfar-uncombined]
; set static options for cfar jobs that
; compute uncombined false alarm rates
ranking-table = coinc_event
output-column = false_alarm_rate
output-table = coinc_ringdown
ranking-stat = likelihood
ifos-table = coinc_ringdown
rank-by = MAX
; the following options define a 'category':
; the first 2, what parameter to bin by
; the last, whether or not to bin by coincident
; ifos
group-by-multiplicity =
; following are optional
;ifos-table = coinc_ringdown
;output-table = coinc_ringdown
;param-table = coinc_ringdown
;exclude-coincs = [all in h1,h2];[h1,h2 in all];[h2,l1 in h1,h2,l1]

[cfar-combined]
; set static options for cfar jobs that
; compute combined false alarm rates
ranking-table = coinc_ringdown
output-column = combined_far
ranking-stat = false_alarm_rate
rank-by = MIN
; no categories are defined when combining;
; however, if only want to combine over param-bins,
; uncomment group-by-ifos
;group-by-ifos =
; following are optional
;ifos-table = coinc_ringdown
;output-table = coinc_ringdown
;param-table = coinc_ringdown

[cbc_print]
; set static options common for printlc, printsims, and printmisses
daily-ihope-pages-location = https://ldas-jobs.ligo.caltech.edu/
~cbc/ihope_daily

[printlc]
; set static options for printlc jobs

```

```

coinc-table = coinc_ringdown
ranking-stat = combined_far
rank-by = MIN
convert-durations = days
limit = 20
; following are needed for getting single ifo information
get-sngl-info =
sngl-table = sngl_ringdown
time-column = start_time
; if set the output format to anything but xml, minifollowups will not run
output-format = html
columns = rank,combined_far,fap,fap_1yr,sngl_start_time,sngl_ifo
__Px_click_for_eelog_xP_,instruments_on,sngl_frequency,sngl_quality,sngl_mass,
sngl_spin,sngl_snr,duration

```

[printsims]

```

; set static options for printsims jobs
ranking-stat = combined_far
rank-by = MIN
sort-by = injected_decisive_distance
convert-durations = days
get-sngl-info =
sngl-table = sngl_ringdown
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
time-column = start_time
map-label = ring_nearby
param-name = combined_far
param-ranges = '!0'
output-format = html
include-only-coincs = '[ALLinH1,H2,L1]'
columns = rank,recovered_combined_far,injected_gps_time,injected_gps_time_
ns,injected_decisive_distance,sngl_ifo,instruments_on,recovered_start_time
,sim_tag,injected_mass1,injected_mass2,injected_spin1z,injected_spin2z,
injected_frequency,sngl_frequency,injected_quality,sngl_quality,sngl_snr
;rank-range = !1

```

[printmissed]

```

; set static options for printmissed jobs
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
map-label = ring_nearby
lifetime-program = rinca
output-format = html
limit = 20
columns = rank,decisive_distance,gps_time,instruments_on,eff_dist_h,
eff_dist_l,frequency,quality,mass,spin,sim_tag

```

[minifollowups]

```

; set static options for minifollowup jobs

```

```

[plotfm]
; establish what plotfm jobs to do; the following options
; must have corresponding sections with options given
; each section name will be added to the plotfm user_tag
fm_injected_params_imr =
fm_lin_injected_params_ringdown =
fm_log_injected_params_ringdown =
fm_lin_detected_vs_injected_imr =
fm_log_detected_vs_injected_imr =
fm_lin_detected_vs_injected_ringdown =
fm_log_detected_vs_injected_ringdown =
fm_loglog_detected_vs_injected_ringdown =
fm_lin_accuracy_plots_imr =
fm_lin_accuracy_plots_ringdown =
fm_freq_accuracy_plots_ringdown =
fm_injdist_plots_imr =
fm_log_injdist_plots_ringdown =
fm_loglog_injdist_plots_ringdown =

[fm_injected_params_imr]
; plot injected decisive distance v. chirp mass
variables = 'injected_mass1:Injected Mass 1( $M \odot$ ); injected_
mass2:Injected Mass 2( $M \odot$ )' --variables 'injected_mass1 +
injected_mass2:Injected Mass1 + Mass2( $M \odot$ ); (injected_
mass1/injected_mass2):Injected Mass Ratio( $m_1/m_2$ )' --variables
'injected_mchirp:Injected Chirp Mass( $M \odot^{\frac{2}{5}}$ );
(injected_mass1/injected_mass2):Injected Mass Ratio( $m_1/m_2$ )'
--variables 'injected_spin1z:Injected Spin1z; injected_spin2z:Injected Spin2z'
logz =
colorbar =
; add lines at each mass bin boundary
;plot-y-function = '3.48:$3.48$' --plot-y-function '7.4:$7.4$'
ranking-stat = 'combined_far:Combined FAR ( $\text{yr}^{-1}$ )'
rank-by = MIN
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_lin_injected_params_ringdown]
; plot injected mass2 v. injected mass1
variables = 'injected_mass:Injected Ringdown Mass ( $M \odot$ );
injected_spin:Injected Ringdown Spin' --variables 'injected_mass:Injected
Ringdown Mass( $M \odot$ ); injected_quality: Injected Ringdown Quality Factor'
logz =
colorbar =
; add lines at each mass bin boundary
;plot-y-function = '3.48:$3.48$' --plot-y-function '7.4:$7.4$'
ranking-stat = 'combined_far:Combined FAR ( $\text{yr}^{-1}$ )'
rank-by = MIN
simulation-table = sim_ringdown

```

```

recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_log_injected_params_ringdown]
; plot injected mass2 v. injected mass1
variables = 'injected_frequency:Injected Ringdown Frequency ($Hz$);
injected_quality: Injected Ringdown Quality Factor' --variables 'injected_
frequency:Injected Ringdown Frequency ($Hz$); injected_mass: Injected
Ringdown Mass ($M_\odot$)' --variables 'injected_frequency:Injected
Ringdown Frequency ($Hz$); injected_spin: Injected Ringdown Spin'
logx =
logz =
colorbar =
;plot-y-function = '3.48:$3.48$' --plot-y-function '7.4:$7.4$'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_lin_detected_vs_injected_imr]
variables = '(injected_gps_time+injected_gps_time_ns*1e-9): Injected
GPS Time($s$); (recovered_start_time+recovered_start_time_ns*1e-9):
Recovered Ringdown Start Time($s$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logz =
colorbar =
plot-x-function = x
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_log_detected_vs_injected_imr]
variables = 'injected_mass1 + injected_mass2: Injected Mass1 +
Mass2($M_\odot$); recovered_frequency:Detected Ringdown
Frequency($Hz$)' --variables 'injected_mchirp: Injected Chirp
Mass($M_\odot^{\frac{2}{5}}$); recovered_frequency:Detected
Ringdown Frequency($Hz$)' --variables 'injected_mass1 + injected_
mass2: Injected Mass1 + Mass2($M_\odot$); recovered_snr:SNR'
--variables 'injected_mchirp: Injected Chirp Mass($M_\odot^{\frac{2}{5}}$);
recovered_snr: SNR' --variables 'injected_mass1 + injected_
mass2: Injected Mass1 + Mass2($M_\odot$); recovered_choppedl_
snr:New Triple SNR' --variables 'injected_mchirp: Injected Chirp Mass
($M_\odot^{\frac{2}{5}}$); recovered_choppedl_snr: New Triple SNR'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =

```

```

logz =
colorbar =
simulation-table = sim_inspirar
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_lin_detected_vs_injected_ringdown]
variables = 'injected_quality:Injected Ringdown Quality Factor;
recovered_quality:Detected Ringdown Quality Factor' --variables
'injected_mass:Injected Ringdown Mass( $M \cdot$ );recovered_
mass:Detected Ringdown Mass( $M \cdot$ )' --variables 'injected_
spin:Injected Ringdown Spin; recovered_spin:Detected Ringdown Spin'
ranking-stat = 'combined_far:Combined FAR ( $\text{Yr}^{-1}$ )'
rank-by = MIN
logz =
colorbar =
plot-x-function = x
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_loglog_detected_vs_injected_ringdown]
variables = 'injected_frequency:Injected Ringdown Frequency ( $\text{Hz}$ );
recovered_frequency:Detected Ringdown Frequency ( $\text{Hz}$ )' --variables
'injected_frequency:Injected Ringdown Frequency ( $\text{Hz}$ ); recovered_
snr:SNR' --variables 'injected_frequency:Injected Ringdown Frequency
( $\text{Hz}$ ); recovered_choppedl_snr:New Triple SNR'
ranking-stat = 'combined_far:Combined FAR ( $\text{Yr}^{-1}$ )'
rank-by = MIN
logx =
logy =
logz =
colorbar =
plot-x-function = x
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_lin_accuracy_plots_imr]
; plot linear fm plots
variables = 'injected_mass1 + injected_mass2: Injected Mass1 +
Mass2( $M \cdot$ ); 1000*((recovered_start_time+recovered_start_
time_ns*1e-9)-(injected_gps_time+injected_gps_time_ns*1e-9)):
Recovered - Injected Start Time ( $\text{ms}$ )' --variables 'injected_mchirp:
Injected Chirp Mass( $M \cdot^{\frac{2}{5}}$ ); 1000*((recovered_
start_time+recovered_start_time_ns*1e-9)-(injected_gps_time+
injected_gps_time_ns*1e-9)): Recovered - Injected Start Time ( $\text{ms}$ )'
--variables 'injected_spin1z:Injected Spin1z; 1000*((recovered_start_

```

```

time+recovered_start_time_ns*1e-9)-(injected_gps_time+injected_gps_
time_ns*1e-9)): Recovered - Injected Start Time ($ms$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logz =
colorbar =
simulation-table = sim_inspiral
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_freq_accuracy_plots_ringdown]
; plot linear fm plots
variables = 'injected_frequency: Injected Ringdown Frequency ($Hz$);
1000*((recovered_start_time+recovered_start_time_ns*1e-9)-(injected_
gps_time+injected_gps_time_ns*1e-9)): Recovered Start Time -
Injected End Time($ms$)' --variables 'injected_frequency: Injected
Ringdown Frequency ($Hz$); (injected_frequency - recovered_
frequency)/((injected_frequency+recovered_frequency)/2.0):
Fractional Ringdown Frequency Accuracy ($\frac{freq_{inj}-freq_
{det}}{(freq_{inj}+freq_{det})/2}$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logx =
logz =
colorbar =
plot-x-function = 0
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_injdist_plots_imr]
; plot injected decisive distance v. gps_time in days
variables = 'injected_gps_time:gps_days; injected_decisive_distance:
Injected Decisive Distance ($Mpc$)' --variables 'injected_mass1+injected_
mass2:Injected Mass1+Mass2($M_\odot$); injected_decisive_distance:
Injected Decisive Distance ($Mpc$)' --variables 'injected_mchirp:Injected
Chirp Mass($M_\odot^{\frac{2}{5}}$);injected_decisive_distance:Injected
Decisive Distance ($Mpc$)' --variables '(injected_mass1/injected_mass2):
Injected Mass Ratio; injected_decisive_distance:Injected Decisive Distance
($Mpc$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
; split up by mass-bin
;param-name = injected_mchirp
;param-ranges = [0,3.48];[3.48,7.4];[7.4,20]
simulation-table = sim_inspiral

```

```

recovery-table = coinc_ringdown
livetime-program = rinca
map-label = insp_nearby

[fm_log_injdist_plots_ringdown]
; plot injected decisive distance v. difference in recovered/injected gps-times
variables = 'injected_mass: Injected Ringdown Mass($M\odot$); injected_
decisive_distance:Injected Decisive Distance ($Mpc$)' --variables 'injected_
quality:Injected Ringdown Quality Factor; injected_decisive_distance:Injected
Decisive Distance ($Mpc$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logy =
logz =
colorbar =
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[fm_loglog_injdist_plots_ringdown]
; plot injected decisive distance v. difference in recovered/injected gps-times
variables = 'injected_frequency: Injected Frequency ($Hz$); injected_decisive
_distance:Injected Decisive Distance ($Mpc$)'
ranking-stat = 'combined_far:Combined FAR ($yr^{-1}$)'
rank-by = MIN
logx =
logy =
logz =
colorbar =
simulation-table = sim_ringdown
recovery-table = coinc_ringdown
livetime-program = rinca
map-label = ring_nearby

[plotslides]
; set static options for plotslide jobs
coinc-table = coinc_ringdown

[plotcumhist]
; set static options for plotcumhist jobs
num-bins = 20
ranking-stat = choppedl_snr
rank-by = MAX
;square-stats =
coinc-table = coinc_ringdown
;param-name = mchirp
;param-ranges = [0,3.48);[3.48,7.4);[7.4,20]

[plotifar]
; set static options for plotifar jobs

```

```
; ringdown search doesn't separate by mass bin (yet)
;param-name = mchirp
;param-ranges = [0,3.48);[3.48,7.4);[7.4,20]
plot-uncombined =
group-by-multiplicity =
plot-combined =
plot-slides =
show-min-bkg =
show-max-bkg =
show-two-sigma-error =
coinc-table = coinc_ringdown
;x-min =
;x-max =
;y-min =
;y-max =

[mvsc_dag]
number-of-trees = 200
criterion-for-optimization = 6
leaf-size = 5
sampled-parameters = 12
all-instruments = H1,H2,L1

[mvsc_get_doubles]
number = 10
;apply-weights =
;check-weights =
exact-tag = ring_exact
nearby-tag = ring_nearby
start-time = 0
end-time = 999999999
```

Vita

Sarah Caudill was born in Daytona Beach, Florida to Michael and Peggy Caudill. She graduated from New Smyrna Beach High School in 2002, and earned her B.S. in physics from Stetson University in 2006. As an undergraduate physics student, she participated in several research projects, one of which introduced her to the exciting science of the Laser Interferometer Gravitational-wave Observatory. She chose to continue along this path as a graduate student at LSU in 2007.