The Importance of Spin for Observing Gravitational Waves from Coalescing Compact Binaries with LIGO and Virgo

Thesis by

Stephen Privitera

In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy



California Institute of Technology Pasadena, California

2014

(Defended May 16, 2014)

© 2014 Stephen Privitera All Rights Reserved

DISCLAIMER

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. This work is assigned LIGO document number LIGO-P1400092.

Acknowledgments

I did not exactly take the straight line to the finish. Many people helped me get to the end. This page cannot fit all the names (see any LIGO paper), so I will mention just a few people who have been particularly impactful to my graduate studies. From softball to group meeting, from the face-to-face to pub crawls, there have been so many other wonderful influences in my life, not mentioned below, and I thank you all.

Advisers I wholeheartedly thank my graduate adviser Alan Weinstein for his patience and commitment to my completion of this work. Chad Hanna has also provided an exceptional amount of academic and moral support throughout my studies, for which I am infinitely grateful. Part of the work presented here was conducted in close collaboration with Ajith Parameswaran and Nickolas Fotopoulos, without whom these results would not exist. Thank you all! You have each inspired me and helped me to grow as a scientist and as a professional.

Collaborators I have had the honor to work with many talented scientists within the CBC working group of the LIGO and Virgo collaborations, to whom I am deeply indebted. None of this work could be possible without the hard teamwork of this group in maintaining a common computing infrastructure and exchanging ideas and feedback on work in progress.

Family Thank you, family. To Dad, who would buy me as many books as I could read. To Mom, who cheered me to the top, like I did for her on Echo Mountain. To Joe, who has always been one of my best friends. To my other brothers and sisters: Rachel and Gerald, Greg and Alisha, Andy and Jordan. To my nieces and nephews: Nicholas and Bella, Aiden and Gracie, Sam, Adam and Hannah. You are all my money.

Friends I am particularly indebted to Peter Chung, Jeffrey Kaplan, Richard Norte, and Kari Hodge, all of whom provided an incredible amount of moral support throughout my time here.

Funding I gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, which provided support for this work.

Abstract

General Relativity predicts the existence of *gravitational waves*, which carry information about the physical and dynamical properties of their source. One of the many promising sources of gravitational waves observable by ground-based instruments, such as in LIGO and Virgo, is the coalescence of two compact objects (neutron star or black hole). Black holes and neutron stars sometimes form binaries with short orbital periods, radiating so strongly in gravitational waves that they coalesce on astrophysically short timescales. General Relativity gives precise predictions for the form of the signal emitted by these systems. The most recent searches for theses events used waveform models that neglected the effects of black hole and neutron star spin. However, real astrophysical compact objects, especially black holes, are expected to have large spins. We demonstrate here a data analysis infrastructure which achieves an improved sensitivity to spinning compact binaries by the inclusion of spin effects in the template waveforms. This infrastructure is designed for scalable, low-latency data analysis, ideal for rapid electromagnetic followup of gravitational wave events.

Summary of Results

The work reported here was carried out in collaboration with a large number of contributing scientists. Below, we indicate the specific work for which the author has made significant contributions. For a more thorough summary of the results, see Sec. 8.1.

Refs. [1,2] Secs. 2.5, 4.1

These articles present a characterization of the achieved LIGO and Virgo detector sensitivity to binary neutron star coalescence in LIGO's Fifth and Sixth Science Runs and Virgo's First, Second and Third Science Runs.

Refs. [3,4] Chap. 4

These articles present two searches for compact binary coalescence signals in data collected during LIGO's Sixth and Virgo's Second and Third Science Runs.

Refs. [3–7] Chap. 4

These articles present the implementation and application of the loudest event statistic for interpreting search results in terms of event rates. Refs. [6,7] are not discussed in this thesis.

Refs. [8,9] Secs. 6.2, 7.2

These articles illustrate the application of an extensible, generic platform for generating template banks for compact binary searches.

Ref. [8] Sec. 6.3

This article demonstrates that a three parameter aligned-spin waveform model may be adequate for the detection of a realistic population of compact binary systems with generic component spins.

Ref. [9] Secs. 7.2, 7.3

This article demonstrates an end-to-end pipeline analysis in S5 LIGO detector noise for which spinning templates improve the search sensitivity to aligned spin signals. We also extend these results to simulated aLIGO noise (unpublished).

Contents

A	Acknowledgments				
A	bstra	let	v		
1	Gra	avitational Waves from Compact Binary Coalescence			
	1.1	Gravitational waves in General Relativity	2		
	1.2	Gravitational waves from compact binary coalescence	8		
	1.3	Formation scenarios for compact binaries	17		
		1.3.1 Binary common evolution	17		
		1.3.2 Dynamical capture	21		
	1.4	Coalescence rates of compact binaries	22		
	1.5	Compact binaries with spinning components	26		
		1.5.1 Post-Newtownian reduced spin approximation	28		
		1.5.2 Effective spin inspiral-merger-ringdown approximation	31		
	1.6	Overview of the remainder of this thesis	34		
2	Las	er Interferometric Gravitational Wave Observatories	35		
	2.1	The Michelson interferometer	36		
	2.2	Length sensing in the shot noise limit	38		
	2.3	Fabry-Perot cavities	40		
	2.4	Limiting length stabilization noise	44		

		2.4.1	Seismic noise	47
		2.4.2	Thermal noise	48
		2.4.3	Quantum noise	51
	2.5	Actua	l LIGO and Virgo sensitivities	52
	2.6	Anten	na response of LIGO detectors	53
3	\mathbf{Det}	\mathbf{ecting}	Compact Binary Coalescence Signals	58
	3.1	Detect	ing signals of known form in stationary Gaussian noise	59
	3.2	Coveri	ing the parameter space	61
		3.2.1	Searching over extrinsic parameters	62
		3.2.2	Covering masses and spins	64
		3.2.3	Lattice template placement	66
		3.2.4	Random template placement	68
	3.3	Non-st	tationary and non-Gaussian noise	69
		3.3.1	Data segmentation	69
		3.3.2	Coincidence between detectors	70
		3.3.3	Signal consistency tests	71
		3.3.4	Data quality metrics	74
	3.4	The II	HOPE Pipeline	75
4	Two	o Searc	thes for CBCs in Data from LIGO's Sixth and Virgo's Second and Third	1
	Scie	ence R	uns	79
	4.1	LIGO'	's sixth and Virgo's second and third science runs	80
	4.2	Search	parameters	83
	4.3	Evalua	ating the search sensitivity	85
	4.4	Estima	ating binary coalescence rates	87
	4.5	Margi	nalizing over uncertainties	91

	4.6	Search results	93
		4.6.1 Low mass	94
		4.6.2 High mass	98
	4.7	Comments on spin in these searches	100
5	Tow	vards Advanced LIGO Searches	102
	5.1	Scaling CBC Searches from Initial LIGO to Advanced LIGO	103
	5.2	Multibanding and down-sampling	107
	5.3	Orthogonal template banks	110
	5.4	GSTLAL: A stream based approach to gravitational wave searches	114
	5.5	Low latency trigger generation	115
6	$\mathbf{Th}\epsilon$	e Importance of Spin for the Detection of Gravitational Waves from Compac	t
	Bin	ary Coalescences	119
	6.1	Astrophysical expectations for spin in compact binaries	120
	6.2	SBANK: a generic infrastructure for template bank construction	122
6.3 An aligned spin template bank for the detection of generically spinning binary inspirals in advanced LIGO			123
	6.4	Aligned spin template banks for the detection of merging binary black holes	131
		6.4.0.1 Initial LIGO sensitivity	133
		6.4.0.2 Advanced LIGO sensitivity	134
7	A S	ensitive Search Pipeline for Binary Black Holes with Aligned Spin	135
	7.1	Pipeline implementation challenges	136
	7.2	A sensitive search pipeline for gravitational waves from merging binary black holes	
		with aligned spin	138
	7.3	Prospects for advanced LIGO	142
8	Con	nclusions	145

8.1	Summary of results	145
8.2	Future work	147
8.3	Long term prospects	148

151

Bibliography

List of Figures

1.1	Plus and cross polarizations of a gravitational wave	5
1.2	Coordinate systems for describing gravitational waves	8
1.3	Circularly polarized gravitational waves	13
1.4	Inspiral, merger, and ringdown	16
1.5	Formation of compact binaries by common evolution	18
1.6	The Hulse-Taylor binary pulsar	23
1.7	Effect of spin magnitude and orientation on the inspiral signal $\ldots \ldots \ldots \ldots \ldots$	28
1.8	The IMRPHENOMB waveform approximation for aligned spin binary black holes	29
1.9	The dimensionality problem for numerical relativity simulations of compact binary mergers	32
2.1	Michelson interferometer design	37
2.2	Fabry-Perot optical cavity	41
2.3	Advanced LIGO optical layout	43
2.4	Length sensing of the Michelson and Michelson-Fabry-Perot interferometer designs	45
2.5	Anticipated advanced LIGO noise budget	46
2.6	Recently achieved LIGO and Virgo detector sensitivities	52
2.7	LIGO as a gravitational wave antenna	54
2.8	The importance of multiple detectors for full sky coverage	55
2.9	A developing global network of gravitational wave detectors	56

3.1	Lattice-based placement strategy for non-spinning templates	67
3.2	Breakdown of the chirp time coordinates for lattice template placement	68
3.3	Non-stationarity of real LIGO instrumental noise	70
3.4	Non-Gaussianity of real LIGO instrumental noise	71
3.5	Architecture of the IHOPE analysis pipeline	76
4.1	Detector sensitivities in S6 and $VSR2/3$	81
4.2	Variation of detector sensitivity over the course of the S6/VSR2-3 science runs $\ . \ . \ .$	83
4.3	Defining the inspiral-merger-ringdown boundary	85
4.4	Combining independent experiments with the loudest event statistic $\ldots \ldots \ldots$	91
4.5	Comparison of achieved mean sensitive distance to SenseMon expectations $\ldots \ldots$	94
4.6	Upper limits on the binary coalescence rate as a function of mass	96
4.7	Coalescence rate upper limit rates for BNS, NSBH, and BBH systems	97
4.8	Loudest events in the high mass search	98
5.1	Comparison of initial and advanced LIGO spectra	104
5.2	Dependence of coalescence signal duration on the masses of the system $\ldots \ldots \ldots$	105
5.3	Signal-to-noise available in the low frequency band of aLIGO	105
5.4	Number of templates required to cover a space as a function of low frequency cutoff $% \mathcal{A}$.	106
5.5	Computational benefits of multibanding and the singular value decomposition \ldots .	109
5.6	Illustration of the transformation from correlated physical templates to orthogonal unphysical templates	113
5.7	Architecture of the GSTLAL analysis pipeline	114
5.8	Possible electromagnetic counterparts to compact binary mergers involving matter	116
6.1	Effective fitting factors towards a simulated population of precessing binaries \ldots .	125
6.2	Increase in detection volume at fixed signal-to-noise	125
6.3	Dependence of precession effects on mass ratio	127

6.4	Selection bias against precessing systems	128
6.5	Difference between the T4 and T5 post-Newtonian approximations	129
6.6	Capturing aligned spin effects in high mass template banks	131
6.7	Importance of spin for aLIGO binary black hole searches	134
7 1		190
7.1	Comparison of search sensitivities as a function of false alarm rate threshold	139
7.2	Improvement in parameter recovery with aligned spin templates	140
7.3	Increase in false alarm rate from additional templates	141
7.4	Improved sensitivity to spinning signals in simulated aLIGO noise	143
7.5	The effect of waveform mismatch on the χ^2 statistic	144

List of Tables

1.1	Known relativistic binary pulsars	24
1.2	Known binaries with black hole companions	25
1.3	Event rates predictions for compact binary coalescence	26
3.1	The compact binary parameter space	61
3.2	Data quality metrics for gravitational wave searches	75
4.1	Observational times for the S6/VSR2-3 high mass search $\hfill \ldots \ldots \ldots \ldots \ldots \ldots$	84
4.2	Event rate confidence intervals for a simple counting experiment	89
4.3	Rate upper limits of BNS, BBH, and NSBH coalescence	95
4.4	High mass search sensitivity	99
6.1	Reduced spin template bank parameters	123
6.2	Precessing binary population simulation parameters	124

Chapter 1

Gravitational Waves from Compact Binary Coalescence

Student: Dr. Einstein, aren't these the same questions as last year's [physics] final exam?

Einstein: Yes, but this year the answers are different.

General Relativity (GR) predicts [10] that local changes in a gravitational field are communicated across large distances via gravitational waves (GWs), which carry energy and momentum away from the field at the speed of light. These waves manifest themselves physically as changes in the proper distance between two freely falling test masses, that is, two bodies with negligible self-gravity moving only under the influence of some external gravitational field. The expected size of the effect at Earth from distant astrophysical sources is extraordinarily small even for the most massive and dynamic gravitating systems. A strong source of gravitational waves might optimistically induce at Earth a change in proper distance of peak amplitude $\Delta L \sim 10^{-18}$ m between two masses separated by L = 1 km. Nonetheless, an ambitious worldwide effort [11–14] has led to the development of groundbased detectors with proven sensitivity [1, 2] to such small length changes. Over the next decade, these detectors will continue to improve in sensitivity and bandwidth [15, 16] and hopefully soon make the first direct measurement of GWs.

One of the most promising sources of detectable gravitational waves is the coalescence of *compact binaries*, that is, binaries consisting solely of neutron stars and/or black holes. According to GR, binary systems lose energy through the emission of GWs, shrinking their orbits. If the components of the binary are compact enough, such as with a neutron star or a black hole, the orbit may decay

until the components orbit each other upwards of ten times a second and eventually merge together in what is one of the most energetic astrophysical events thought to occur in Nature, with energy scales exceeding $E_{\rm GW} \gtrsim 10^{53}$ erg. In the following chapters, we report on the results of recent searches for GW signals from compact binary coalescences (CBCs) and demonstrate methods to improve the sensitivity to such signals in future observations.

In this chapter, we motivate our study of GWs from CBCs. We review the basic properties of gravitational waves and show why CBCs are promising targets for GW science. One especially appealing aspect of compact binaries is that their gravitational wave signal can be accurately predicted by General Relativity. Having these waveform models allows for (i) the efficient extraction of the weak GW signal from noisy detector data and (ii) the interpretation of observations in terms of the physical parameters which describe the binary system, such as the component masses and spins. We review two plausible astrophysical scenarios which lead to the formation of coalescing compact binaries. There are large uncertainties in the details of these scenarios and electromagnetic observations of compact binaries are limited. Based on our understanding of the formation channels for these systems, together with what limited observational data we have, we then assess reasonable expectations for the event rates for CBCs.

We close the chapter with a brief review of two waveform models for binary systems with spinning components. While having accurate waveform models for the expected signal in principle improves the ability to extract these signals from the noisy detector data, the successful calculation and search implementation of waveforms which include all possible physical effects is non-trivial. We will later use the waveforms presented at the end of this chapter to study the importance of spin for detection and to demonstrate a search pipeline with improved sensitivity to spinning signals.

1.1 Gravitational waves in General Relativity

In General Relativity, space-time is a four-dimensional manifold M endowed with a metric tensor field \mathbf{g} whose properties determine the motion of free particles. The manifold curves in response to the presence of energy and momentum, and the relationship between space-time curvature and energy-momentum is governed by the Einstein equations

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu},$$
(1.1)

where $R_{\mu\nu}$ is the Ricci curvature tensor, R its trace, $g_{\mu\nu}$ is the field metric, and $T_{\mu\nu}$ is the stressenergy tensor. The local physics defines the stress-energy; the Einstein equations govern how spacetime changes in response to this stress-energy.

The metric tensor \mathbf{g} is the fundamental geometric construct in GR. The metric gives rise to the concept of length and determines the motion of free particles. For two infinitesimally separated points \mathcal{P} and \mathcal{Q} in space-time, the invariant distance ds between them is given in terms of the metric and local coordinates as

$$ds^2 = g_{\mu\nu}dx^{\mu}dx^{\nu}, \qquad (1.2)$$

where $dx^{\mu} = x^{\mu}(\mathcal{P}) - x^{\mu}(\mathcal{Q})$. We adopt the standard convention that repeated indices indicate summation over those indices. For a finite path $\gamma : [a, b] \to M$ described in terms of coordinates as $x^{\mu} = x^{\mu}(\tau)$, its length is defined as

$$L = \int_{\gamma} ds = \int_{a}^{b} \sqrt{g_{\mu\nu} \frac{dx^{\mu}}{d\tau} \frac{dx^{\nu}}{d\tau}} d\tau.$$
(1.3)

These constructs then determine the motion of free particles, i.e., particles moving only under the influence of gravity. If a free particle begins and ends at two finitely separated points $\gamma(a)$ and $\gamma(b)$, then of all paths through M connecting these points, the particle moves along that of minimal length as defined by Eqn. 1.3.

The Einstein equations represent a set of ten coupled, non-linear, partial differential equations for the metric tensor, and exact solutions to these equations are difficult to find. The Ricci tensor on the left hand side of Eqn. 1.1 is defined as sums and products of the Christoffel symbols $\Gamma^{\alpha}_{\beta\gamma}$ and their partial derivatives by the formula

$$R_{\mu\nu} = \partial_{\rho}\Gamma^{\rho}_{\nu\mu} - \partial_{\nu}\Gamma^{\rho}_{\rho\mu} + \Gamma^{\rho}_{\rho\lambda}\Gamma^{\lambda}_{\rho\mu} - \Gamma^{\rho}_{\nu\lambda}\Gamma^{\lambda}_{\rho\mu}.$$
 (1.4)

The Christoffel symbols are themselves derivatives of the metric

$$\Gamma_{\alpha\beta\gamma} = \frac{1}{2} \left(\partial_{\beta} g_{\gamma\alpha} + \partial_{\alpha} g_{\gamma\beta} - \partial_{\gamma} g_{\alpha\beta} \right)$$
(1.5)

and $\Gamma^{\alpha}_{\beta\gamma} \equiv g^{\alpha\delta}\Gamma_{\delta\beta\gamma}$. In the study of gravitational wave sources, we are typically interested in physical systems which involve strong and highly dynamical gravitational fields and for understanding such systems there is no getting around confronting the difficulties imposed by the non-linearity of the

Einstein equations.

However, our primary interest here is in the detection of the relativity weak observable effect of gravitating systems far from the source, for which $T_{\mu\nu} = 0$. In this case, the Einstein equations simplify to

$$R_{\mu\nu} = 0. \tag{1.6}$$

Far away from the gravitational source, we can consider the space-time as a small perturbation to the flat Minkowski space-time, writing

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \qquad (1.7)$$

where

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$
(1.8)

and where $h_{\mu\nu} \ll 1$. Under these assumptions, Eqn. 1.6 admits a transversely-propagating wave solution, which travels at the speed of light and has two independent degrees of freedom. If we choose our coordinates such that the wave travels in the +z direction, we can write the solution in terms of the metric as

$$g_{\mu\nu} = \eta_{\mu\nu} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{+} & h_{\times} & 0 \\ 0 & h_{\times} & -h_{+} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$
(1.9)

where h_+ and h_{\times} are functions of time and space satisfying the wave equation,

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)h = 0, \tag{1.10}$$

where $h = h_+$ or $h = h_{\times}$. The wave equation Eqn. 1.10 is valid for the components of the metric perturbation h in the *transverse traceless* (TT) gauge. In this gauge, the metric perturbation affects only the components of the metric which are orthogonal to the direction of propagation (transverse), and is symmetric about reflection through the direction of propagation (traceless). The amplitudes h_+ and h_{\times} correspond to the two possible polarizations for gravitational waves called *plus* and *cross*,



Figure 1.1: Plus and cross polarizations of a gravitational wave. Gravitational waves have two basic polarizations, termed plus and cross. These polarizations manifest as ellipses with time-varying eccentricities and axes rotated with respect to each other by $\pi/4$ radians. Here we have chosen the axes such that the left figure depicts the plus polarization and the right figure depicts the cross polarization. The third dimension can be used equally well to visualize the effect of the polarizations in time at fixed z or in space at fixed t. Note that for the purpose of demonstration the strains depicted here (h = 0.15) are highly exaggerated compared to those strains ($h \sim 10^{-23}$) detectable by aLIGO instruments.

respectively. For more detail, see Ref. [17], one of the many classic texts on GR.

The general solution to Eqn. 1.10 in free space (without boundary conditions) is given by a continuous spectrum of periodic functions,

$$h = \int h(\omega) \cos(\vec{k} \cdot \vec{x} - \omega t + \delta) d\omega, \qquad (1.11)$$

where $h(\omega)$ is the wave amplitude at the frequency ω , \vec{k} is the wave vector, and $\delta = \delta(\omega)$ is a phase offset which may depend on frequency. The wave vector points in the direction of wave propagation and has magnitude $k = 2\pi/\lambda$ for wavelength λ , which is related to the angular frequency ω through the dispersion relation $\omega = ck$. In general, a gravitational wave will consist of many different frequency components. In applications, we often approximate the gravitational wave by a quasimonochromatic signal that has a single frequency f(t) which evolves over time.

It is also useful to consider only a single frequency component in understanding the physical manifestation of a gravitational wave. Consider an ensemble of test particles situated on a ring of unit radius centered at z = 0. We take first the plus polarization for which $h_{\times} = 0$. The proper distance from the origin to a point with coordinates (x, y) = (dx, dy) is

$$ds^{2} = g_{xx}dx^{2} + 2g_{xy}dxdy + g_{yy}dy^{2}$$

= $(1 + h_{+}\cos(\omega t))dx^{2} + (1 - h_{+}\cos(\omega t))dy^{2}.$ (1.12)

We can write this result in a more suggestive form by dividing through by $(1 - h_{+} \cos(\omega t))(1 + h_{+} \cos(\omega t))$ and using the weak field approximation to discard terms of order $\mathcal{O}(h_{+}^{2})$. In doing so, we find that

$$ds^{2} = \frac{dx^{2}}{1 - h_{+}\cos(\omega t)} + \frac{dy^{2}}{1 - h_{+}\cos(\omega t)}.$$
(1.13)

From this expression, we see that the locus of points at constant proper distance describes an ellipse whose semi-major and semi-minor axes coincide with the x and y coordinate axes and have time-varying lengths $a = \sqrt{1 + h_+ \cos(\omega t)} \approx 1 + \frac{1}{2}h_+ \cos(\omega t)$ and $b = \sqrt{1 - h_+ \cos(\omega t)} \approx 1 - \frac{1}{2}h_+ \cos(\omega t)$.

One can understand the cross polarization in a similar way. One particularly simple way to see the effect of the cross polarization on a ring of test particles is to notice that a rotation of $\pi/4$ about the z-axis transforms the cross polarization into the plus polarization. Suppose that the metric perturbation in a given coordinate system is purely of the cross component, i.e.,

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & h_{\times} & 0 \\ 0 & h_{\times} & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$
 (1.14)

If we then apply a spatial rotation of $\pi/4$ radians about the z-axis

$$\Theta = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$
(1.15)

and use the usual tensor transformation law $h_{\mu'\nu'} = \Theta^{\mu}_{\mu'}\Theta^{\nu}_{\nu'}h_{\mu\nu}$, we find that in these new coordinates

the metric pertubation is given by

$$h_{\mu'\nu'} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{\times} & 0 & 0 \\ 0 & 0 & -h_{\times} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$
 (1.16)

Thus, the physical manifestation of the cross polarization acting on a unit ring of test particles is again an ellipse with time varying eccentricity, but whose axes are rotated by $\pi/4$ radians relative to the plus polarization. These two polarizations are illustrated in Fig. 1.1.

We have so far been working in a specific coordinate system in which the wave travels along the +z direction, which is convenient for understanding the physical nature of these waves. This coordinate system is referred to as the *radiation frame* (see Fig. 1.2). However, we frequently encounter situations in which there is a different natural coordinate system within which to work. It is useful therefore to have a coordinate-independent expression for the metric perturbation. In the radiation frame, we can introduce two unit vectors, $\vec{e_1}$ and $\vec{e_2}$, which lie along the x and y axes, respectively. We then define the plus and cross polarizations as

$$\mathbf{e}_{+} \equiv \vec{e}_{1} \otimes \vec{e}_{2} - \vec{e}_{2} \otimes \vec{e}_{1} \tag{1.17}$$

$$\mathbf{e}_{\times} \equiv \vec{e}_1 \otimes \vec{e}_2 + \vec{e}_2 \otimes \vec{e}_1 \tag{1.18}$$

where \otimes denotes the outer product and bold face font indicates that the object represented is a second rank tensor. With these expressions for the two polarizations, the metric perturbation can then be concisely written as

$$\mathbf{h} = h_+ \mathbf{e}_+ + h_\times \mathbf{e}_\times. \tag{1.19}$$

Although we have used the radiation frame to construct this expression for the metric perturbation, the result is nonetheless coordinate-independent. Having this form for the metric perturbation will be useful later when we need to express \mathbf{h} in other coordinate systems.



Source frame. The z-axis points along the direction of total angular momentum $\vec{J} = \vec{L} + \vec{S_1} + \vec{S_2}$, where \vec{L} is the orbital angular momentum and $\vec{S_i}$ are the spin angular momenta of the component objects. The x and y axes are defined to point along the directions of the semi-major and semi-minor axes of the orbit at some fiducial initial time. The coalescence phase φ_{coal} is defined as the angle between the x-axis and the line joining the two bodies at the last stable circular orbit.

Radiation frame. The z-axis points along direction of propagation of the gravitational wave, the x-axis lies in the plane of the ecliptic, and the y-axis is chosen to make a right-handed coordinate system. Two angles transform from the source frame to the radiation frame: (i) the initial inclination ι of the total angular momentum relative to the line of sight, and (ii) the orbital phase φ_{coal} at coalescence.

Detector frame. The z-axis points towards the local zenith and the x and y axes are chosen to point along the two detector arms such as to form a right-handed coordinate system. Three additional angles transform from the radiation frame to the detector frame: (i) the angular position θ of the binary relative to the local zenith, (ii) the azimuthal anglular position ϕ of the binary relative to the x arm of the detector, and (iii) the polarization angle ψ .

Detector frame

Figure 1.2: *Coordinate systems for describing gravitational waves.* We define three convenient coordinate systems for describing gravitational waves and the related angular parameters which transform from one frame to another. Figure courtesy of P. Ajith.

1.2 Gravitational waves from compact binary coalescence

We have just established some of the general properties of gravitational waves. Now we want to relate the form of a gravitational wave to the properties of the source generating it. We focus specifically on binary sources which consist of only neutron stars and/or black holes. We refer to such binary systems as *compact binaries*; they may be one of binary black holes (BBHs), binary neutron stars (BNSs) and neutron star-black hole binaries (NSBHs). Gravitational waves carry away energy from the source. In a binary system, this loss of energy results in a shrinking orbital separation. If the two objects in a compact binary start out in a sufficiently close orbit, they may radiate so much energy in GWs that the two objects *coalesce* on realistic astrophysical timescales. Compact binary mergers are fortunately relatively simple to model, and the GW signature of these systems can be accurately predicted by solving the field equations of GR. In this section, we illustrate the basic features of the gravitational wave signature of these spectacular events.

We give only a very elementary overview of the basic concepts involved in computing the gravitational wave signals from CBCs. The general treatment of the calculation of gravitational waveforms from compact binary coalescence is far beyond the scope of this thesis. Instead, we will attempt only to provide some intution to the nature of the compact binary coalescence signal by considering these systems in the quasi-Newtonian limit, where analytic expressions are relatively easy to obtain. More accurate analytic and semi-analytic expressions for the gravitational waveform from CBCs to higher post-Newtonian order are well known [18], and we will briefly outline the flavor of these more sophisticated calculations. The post-Newtonian approximation breaks down near the coalescence of the binary when the characteristic velocities approach the speed of light. In the merger regime, one has to confront the full non-linearity of the field equations, and developing waveform models requires insight from numerical simulations [19–21].

The post-Newtonian formalism is a semi-perturbative approach to computing the gravitational waveform in the slow motion limit. We associate with the binary a dimensionless characteristic velocity

$$v \equiv \frac{(GM\omega_{\rm orb})^{1/3}}{c},\tag{1.20}$$

where M is the binary total mass and ω_{orb} is the orbital angular frequency. The *post-Newtonian* order refers to half the largest exponent in v to which the expansion is computed, i.e., the Nth post-Newtonian order approximation to the binary inspiral contains all terms up to v^{2N} . However, we note that there are actually as many as *four* characteristic velocities associated with any given binary – the two individual component orbital speeds and, in the presence of spin, the surface velocities of the spinning bodies. Thus, the accuracy of a post-Newtonian approximation at fixed post-Newtonian order may vary dramatically over the parameter space as the adequacy of the parameter v to characterize the system's degree of relativistic behavior varies.

We make an adiabatic approximation to the binary dynamics, referred to as the *stationary phase approximation*, in which the time scale for the radial inspiral is large compared to the orbital period.

Specifically, we assume

$$\dot{r} \ll r\omega.$$
 (1.21)

This approximation effectively allows us to ignore terms involving \dot{r} in the following calculations. At any moment, we assume that the binary parameters are described by the Keplerian relation

$$\omega^2 = \frac{GM}{r^3},\tag{1.22}$$

where $M = m_1 + m_2$ is the total binary mass. We introduce general relativistic dynamics to the system by considering the energy carried away by the GWs.

The generation of gravitational waves is given to leading order in v by the quadrupole radiation formula [10, 17]

$$h_{ij}^{TT}(t) = \frac{2}{D} \frac{G}{c^4} \ddot{Q}_{ij}^{TT} \left(t - \frac{D}{c} \right),$$
(1.23)

where D is the distance to the source center of mass and Q_{ij}^{TT} is the source mass quadrupole moment computed in the TT gauge. For a system of point masses, which we will take as a first approximation for compact binaries, the mass quadrupole moment is defined as

$$Q_{ij}^{TT} = \sum_{k} m_k (x_i x_j - \frac{1}{3} r^2 \delta_{ij}).$$
(1.24)

where x_i and x_j are the *i*th and *j*th component of the *k*th object's position vector. The sum is taken over the masses in the system. The component-wise definition in Eqn. 1.24 is equivalent to the coordinate-independent tensor definition

$$\mathbf{Q} = \sum_{k} m_{k} (\vec{r}_{k} \otimes \vec{r}_{k} - \frac{1}{3} |\vec{r}_{k}|^{2} \delta), \qquad (1.25)$$

where δ is the Kronecker tensor.

Eqn. 1.23 applies specifically to the coordinates of the radiation frame, as defined in the previous section and in Fig. 1.2. On the other hand, at least for binary systems, it is simpler to compute the quadrupole moment in the *source frame*, which we define as the coordinate system whose origin lies at the center of mass and z-axis points along the direction of the binary total angular momentum $\vec{J} = \vec{L} + \vec{S}_1 + \vec{S}_2$, where \vec{L} is the orbital angular momentum and \vec{S}_i are the spin angular momenta of the component bodies (see Fig. 1.2). The x and y axes are taken to lie in the plane of the binary orbit so as to form a right-handed coordinate system. We note that if the spins are not aligned to the orbital

angular momentum, then there is precession of the orbital plane and the coordinate system is usually specified by the angular momentum at some fiducial reference time. After computing the quadrupole moment in the source frame coordinates with the tensorial definitions given in Eqn. 1.25, we can transform back to the radiation frame and project onto the TT gauge to compute the gravitational radiation coming from the source at an arbitrary direction.

Consider two black holes of masses m_1 and m_2 in a circular Netwonian orbit about their center of mass. In the source frame, the position vectors of the two bodies are given by

$$\vec{r}_1 = \frac{rm_2}{m_1 + m_2} \left(\cos(\omega t), \sin(\omega t), 0 \right)$$
 (1.26)

$$\vec{r}_2 = \frac{rm_1}{m_1 + m_2} \left(-\cos(\omega t), -\sin(\omega t), 0 \right), \qquad (1.27)$$

where r is the orbital separation and ω is the orbital angular frequency. From these expressions, we can then compute the quadrupole moment in the source frame using the definition in Eqn. 1.25, from which we find

$$Q_{ij} = r^2 \frac{m_1 m_2}{m_1 + m_2} \begin{pmatrix} \cos^2(\omega t) & \sin(2\omega t) & 0\\ \sin(2\omega t) & \sin^2(\omega t) & 0\\ 0 & 0 & 0 \end{pmatrix},$$
(1.28)

where we have not yet removed the trace. To the level of approximation we are making here, the trace is essentially constant, and the missing terms disappear when we take time derivatives.

We can now compute the metric perturbation via Eqn. 1.23 together with the binary quadrupole moment computed in Eqn. 1.28. We find

$$h_{ij}^{TT}(t) = \frac{4G^2}{Dc^4} \frac{m_1 m_2}{r} \begin{pmatrix} \cos(2\omega t) & \sin(2\omega t) & 0\\ \sin(2\omega t) & -\cos(2\omega t) & 0\\ 0 & 0 & 0 \end{pmatrix}.$$
 (1.29)

If the source frame happens to coincide with the radiation frame, then we can simply read off the gravitational wave signal from this result using Eqn. 1.23 and Eqn. 1.9, giving

$$h_{+}(t) = \left(\frac{4G^2m_1m_2}{Dc^4}\right)\frac{\cos(2\omega t)}{r}$$
(1.30)

$$h_{\times}(t) = \left(\frac{4G^2m_1m_2}{Dc^4}\right)\frac{\sin(2\omega t)}{r}.$$
 (1.31)

However, in general the radiation frame can be oriented differently with respect to the source frame.

The relative orientation can be fully specified by two angles: the binary inclination ι and the coalescence phase φ_{coal} , as depicted in Fig. 1.2. We can compute the radiation in an arbitrary direction by rotating the source frame by these angles and projecting onto the TT gauge. Rotating first by the inclination angle, we find that

$$h_{+}(t) = \frac{1 + \cos^{2} \iota}{2} \left(\frac{4G^{2}m_{1}m_{2}}{Dc^{4}} \right) \frac{\cos(2\omega t)}{r}$$
(1.32)

$$h_{\times}(t) = \cos \iota \left(\frac{4G^2 m_1 m_2}{Dc^4}\right) \frac{\sin(2\omega t)}{r}.$$
 (1.33)

We have already examined above the effect of an azimuthal rotation of the radiation frame and found that the rotation mixes the plus and plus components, so the full expression for the gravitational wave signal in the radiation frame is given by

$$h_{+}(t) = \frac{1 + \cos^{2} \iota}{2} \left(\frac{4G^{2}m_{1}m_{2}}{Dc^{4}} \right) \frac{\cos(2\omega t - 2\varphi_{\text{coal}})}{r}$$
(1.34)

$$h_{\times}(t) = \cos \iota \left(\frac{4G^2 m_1 m_2}{Dc^4}\right) \frac{\sin(2\omega t - 2\varphi_{\text{coal}})}{r}.$$
(1.35)

We note that the gravitational waveform consists of both + and \times components which are $\pi/4$ radians out of phase. These waveforms are therefore elliptically polarized and the degree of ellipticity depends on the orientation of the binary relative to the observer. In Fig. 1.3, we present a visualization of the effects of a circularly polarized gravitational wave signal.

What we have computed so far is the gravitational radiation emitted by a stationary binary system whose orbital separation never changes. However, gravitational waves carry away energy, causing the radial separation to shrink with time. It can be shown [17] that the energy flux for a gravitational wave is given by

$$\langle \dot{E} \rangle = -\frac{c^3}{32G\pi} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle, \qquad (1.36)$$

where the brackets indicate an average over a single orbit, during which r is nearly constant. Putting our results from Eqns. 1.34 and 1.35 into Eqn. 1.36, we obtain

$$\langle \dot{E} \rangle = -\frac{4G^4 M (m_1 m_2)^2}{r^5 c^5} \frac{\left(\frac{1}{4} (1 + \cos^2 \iota)^2 + \cos^2 \iota\right)}{\pi D^2}.$$
 (1.37)

The gravitational wave flux is in the direction of propagation, and the total energy loss from the system is given by integrating this flux over all angles. Integrating over a sphere of radius D, we



Figure 1.3: Circularly polarized gravitational waves. For a face-on observer, gravitational waves produced by coalescing compact binaries are circularly polarized. In the figure, the rotating arrows indicate the instantaneous polarization of the gravitational wave. When the vector points horizontally, the wave consists purely of the plus component. When the vector points at an angle $\pi/4$ relative to the horizontal, the wave consists purely of the cross component.

find that

$$\frac{16\pi D^2}{5} = 2\pi D^2 \int_0^\pi \left(\frac{1}{4}(1+\cos^2\iota)^2 + \cos^2\iota\right) \sin\iota d\iota, \tag{1.38}$$

and therefore that the total energy loss from the binary system is given by

$$\langle \dot{E} \rangle = -\frac{64}{5} \frac{G^4 M(m_1 m_2)^2}{r^5 c^5}.$$
(1.39)

This loss of energy corresponds to a decrease in the orbital separation r. For a Newtonian orbit, the energy is related to the orbital separation by

$$E = -\frac{1}{2} \frac{Gm_1m_2}{r}.$$
 (1.40)

Simply taking derivatives, we can use this equation to relate the energy loss to the radial evolution, which gives

$$\frac{dr}{dt} = \frac{1}{2} \frac{r^2}{Gm_1m_2} \frac{dE}{dt}.$$
(1.41)

Using Eqn. 1.39, we can express the consequent evolution of the binary orbit as the time derivative of the orbital separation as

$$\frac{dr}{dt} = -\frac{64}{5} \frac{G^3}{c^5} \frac{m_1 m_2 (m_1 + m_2)}{r^3}.$$
(1.42)

One can immediately integrate this equation to obtain the evolution of the orbital separation

$$r(t) = \left(r_0^4 - \frac{256}{5}\frac{G^3}{c^5}m_1m_2(m_1 + m_2)t\right)^{1/4}.$$
(1.43)

Finally, combining this result with Eqn. 1.31, we obtain an expression for the gravitational wave signal emitted by a compact binary system,

$$h_{+}(t) = \frac{1 + \cos^{2} \iota}{2} \frac{G\mathcal{M}}{c^{2}D} \left(\frac{5G\mathcal{M}}{c^{3}}\right)^{1/4} \frac{\cos(2\omega t - 2\varphi_{\text{coal}})}{(t_{\text{coal}} - t)^{1/4}}$$

$$h_{\times}(t) = \cos \iota \frac{G\mathcal{M}}{c^{2}D} \left(\frac{5G\mathcal{M}}{c^{3}}\right)^{1/4} \frac{\sin(2\omega t - 2\varphi_{\text{coal}})}{(t_{\text{coal}} - t)^{1/4}},$$
(1.44)

where \mathcal{M} is the so-called *chirp mass* and is related to the ordinary mass parameters by

$$\mathcal{M} = \mu^{3/5} M^{2/5} = \eta^{3/5} M. \tag{1.45}$$

Here, μ is the reduced mass and $\eta = \mu/M$.

Incidentally, we have also just computed the phase evolution of the binary, since $\omega = \sqrt{GM/r^3}$, which when combined with Eqn. 1.43 gives

$$\phi_{\rm GW}(t) = 2\omega t$$

$$= \frac{2t\sqrt{GM}}{\left(r_0^4 - \frac{256}{5}\frac{G^3}{c^5}m_1m_2(m_1 + m_2)t\right)^{3/8}}.$$
(1.46)

This result, together with Eqn. 1.44, then gives the complete evolution of the plus and cross polarizations emitted by an inspiralling compact binary in the slow motion approximation. The phase $\phi_{\rm GW}$ becomes infinite in a finite amount of time, namely when

$$t = \frac{5}{256} \frac{c^5}{G^3} \frac{r_0^4}{m_1 m_2 (m_1 + m_2)},\tag{1.47}$$

at which point the denominator in Eqn. 1.46 vanishes.

There is a maximal initial separation r_{max} such that the binary will coalesce in less than $T = 10^{10}$ yr (roughly the age of the Universe), and therefore contribute to the population of sources for which we are searching. Using Eqn. 1.47, we find

$$r_{\rm max} = 2 \times 10^6 \times \left(\frac{m_1}{M_{\odot}}\right) \left(\frac{m_2}{M_{\odot}}\right) \left(\frac{m_1 + m_2}{M_{\odot}}\right) \,\,\rm{km.} \tag{1.48}$$

For an equal mass $m_1 = m_2 = 1M_{\odot}$ binary, this maximal separation is equal to about three solar radii. The maximal initial separation is a strong function of mass. For a binary black hole system with $m_1 = m_2 = 10M_{\odot}$, the compact binary may be borne within several thousand solar radii and still merge within relevant astrophysical time scales. In the following section, we will discuss mechanisms by which binaries with such small orbital separations may be created.

Although the calculation we have carried out here is valid only in the Newtonian limit, it in fact illustrates a very general method, reviewed in Ref. [18], for calculation of binary inspiral signals. In the stationary phase approximation, we write the gravitational wave signal as

$$h(t) = A(t)\cos(\phi_{\rm GW}(t)) \tag{1.49}$$

and interpret A(t) to be the instantaneous amplitude and $\omega_{\rm GW} = \dot{\phi}_{\rm GW}$ the instantaneous gravitational wave frequency. Note that the gravitational wave phase is twice the orbital phase $\phi_{\rm GW} = 2\phi_{\rm orb}$, and therefore $f_{\rm GW} = 2f_{\rm orb}$; unless otherwise specified, "frequency" henceforth refers to the gravitational wave frequency and not to the orbital frequency. By definition,

$$\omega_{\rm GW} = \dot{\phi}_{\rm GW}(t) \tag{1.50}$$

and hence we obtain

$$\dot{\phi}_{\rm GW}(t) = \frac{v^3}{GM}.\tag{1.51}$$

Then by simple application of the chain rule for derivatives, one can separate the equation for the time evolution of v into two smaller pieces:

$$\frac{dv}{dt} = \frac{dv}{dE}\frac{dE}{dt},\tag{1.52}$$

where E is the orbital binding energy. We call $\mathcal{F} = -\dot{E}$ the gravitational wave luminosity and write

$$\dot{v} = -\frac{\mathcal{F}}{dE/dv}.\tag{1.53}$$

Eqns. 1.51 and 1.53 form the basis for the flux-energy balance approach to computing the gravitational waveforms for compact binary coalescences. Expanding the Einstein equations in the parameter v, one computes the flux \mathcal{F} and orbital energy E(v) to obtain coupled ordinary differential equations for v and ϕ_{GW} .



Figure 1.4: Inspiral, merger, and ringdown. We show the plus and cross polarizations of the gravitational wave signal from a non-spinning coalescing equal-mass binary with $m_1 = m_2 = 10 M_{\odot}$. The inspiral part of the waveform is long and slowly evolving. In the text, we treat this stage of the coalescence as an adiabatic evolution of a Keplerian orbit with energy loss. In the late inspiral (inset), the orbital behavior becomes highly dynamic. As the orbital separation dips below the last stable circular orbit, the two black holes merge and produce a single, perturbed black hole. Full understanding of the waveforms in these late stages requires careful numerical simulations. Indeed, the calculation of gravitational waveforms for compact binary coalescence is still very much an active and crucial ongoing research activity for LIGO and Virgo searches (see for instance Refs. [18–25]).

All the calculations we have made up to now break down near the merger phase (see Fig. 1.4, inset) of coalescence, where $v \to 1$. We often take the innermost stable circular orbit (ISCO),

$$f_{\rm ISCO} = \frac{c^3}{6\sqrt{6}\pi GM},\tag{1.54}$$

as the cutoff (gravitational wave) frequency for the applicability of the post-Newtonian approximation, although this choice is only a rule of thumb. For a BNS system with $m_1 = m_2 = 1.4 M_{\odot}$, this cutoff frequency is $f_{\rm ISCO} = 1.6$ kHz. For a BBH system with $m_1 = m_2 = 10 M_{\odot}$, this critical frequency is $f_{\rm ISCO} = 220$ Hz. In general, as we can see from Eqn. 1.54, the higher the mass of the system, the smaller the orbital frequency at which the post-Newtonian approximation breaks down. To compute the waveform during the merger phase, we must resort to numerical methods (see for instance Ref. [22]). Numerical simulations of binary coalescence are extremely difficult to carry out, and only a few hundred waveforms have been extracted from such calculations. In Fig. 1.4, we show an example of a gravitational wave which includes each of the inspiral, merger, and ringdown phases of the coalescence. This waveform is derived from a parameterized phenomenological model which essentially interpolates between existing numerical simulations; see Sec. 1.5.2 for more discussion of this particular waveform model, which will be a major player in our later analysis.

1.3 Formation scenarios for compact binaries

We would like to now understand how compact binary systems form and whether they are produced with sufficiently tight orbits¹. There are two commonly discussed formation channels for compact binaries: (1) common binary evolution [26], and (2) dynamical capture resulting from three-body interactions [27]. In both cases, the efficiency of the channel in creating compact binaries is highly uncertain. Our discussion of the formation channels for compact binaries leads naturally to the question of how often these systems form, a question which we will address in the subsequent section. We ultimately want to estimate how often we can expect to observe compact binary coalescence and what we can learn from the observations. As before, the discussion here is necessarily superficial. For a more detailed account of the astrophysics of compact binary formation, see Refs. [26–28] and references therein.

1.3.1 Binary common evolution

The common binary evolution scenario for the formation of a compact binary is depicted in Fig. 1.5. In this scenario, two massive stars are born together in a binary system. The stars must be massive enough $(m_i \gtrsim 10 M_{\odot})$ at birth so that their individual evolution leads to the formation of a neutron star or black hole (due to mass transfer during the binary evolution, this statement is not strictly true). The more massive star evolves off the main sequence first and expands rapidly into a red giant. This stage is referred to as the *Hertzsprung gap*, because the evolution of a star off the main sequence to a red giant occurs extremely fast (thousands of years) compared to the stellar lifetime (millions of years) and stars in the Hertzsprung gap are considerably rarer than main sequence and red giant stars. Consequently, the details of stellar evolution in the Hertzsprung gap are still poorly

 $^{^{1}}$ For short, we will occasionally refer to such systems as *relativistic compact binaries*, meaning compact binaries with sufficiently short orbital periods such that they will coalesce by the emission of gravitational waves on sub-cosmological timescales.



Figure 1.5: Formation of compact binaries by common evolution. We show schematically the critical life stages of a binary which determine whether the binary evolves into a tight compact binary. In order to form a compact binary, the progenitor binary must survive (at the minimum) two supernova explosions, which may impart large kicks to the component stars. Furthermore, the tightening of the orbit requires mass transfer or energy dissipation, which most likely occurs in the poorly understood common envelope phase. The binary evolution is also strongly influenced by the stellar metallicities. Compared to stars with high metallicity, stars with low metallicity have weaker stellar winds and retain more of their mass, favoring the formation of more massive compact remnants. Figure credit [26].

understood.

Binaries are generally born with orbits that are too large to coalesce by the emission of gravitational waves on a sub-cosmological time scale. In order to form a relativistic compact binary from common evolution, some process must work to tighten the orbit. The tightening of the orbit can be achieved through mass transfer and energy dissipation in the Hertzsprung gap phase if the giant star expands beyond its Roche lobe, which is the surface of gravitational equipotential between the two stars. The Roche lobe is tear-drop shaped, with a characteristic size given by

$$R_{\rm Roche} \sim a \left(\frac{m_1}{m_1 + m_2}\right)^{1/3},$$
 (1.55)

where m_1 is the mass of the donor star, m_2 the mass of the companion, and a is the orbital separation. As the more massive star expands into a red giant, it may overflow its Roche lobe and transfer mass and angular momentum to its companion. The orbital angular momentum for a Keplerian binary is given by

$$J = \mu \sqrt{G(m_1 + m_2)a}.$$
 (1.56)

The transfer of mass from the more massive star to the less massive star increases the reduced mass μ , and the conservation of angular momentum then implies a decrease in the orbital separation².

If the primary star expands so much that the two stars form a *common envelope*, the tightening of the orbit can be much more dramatic. As the two stars move through their common envelope, they lose orbital energy to heating of the gas, drawing them closer together. Mass may be lost from the system during this stage, as some fraction of the hot gas becomes gravitationally unbound and is ejected from the system. The outcome of this process is highly uncertain. It is crucial that this mechanism is not too efficient. The survival of the binary through this phase requires the formation of a well-defined helium core structure in the expanding star. If the secondary heats up and ejects the primary's outer layers before a core structure can form, the two stars will simply merge into one. On the other hand, if a core structure does form but the atmosphere is entirely ejected from the system, there may be insufficient mass remaining in the primary to eventually lead to a supernova.

If the primary star is left with sufficient mass following its evolution through the Hertzsprung gap, it will eventually explode in a supernova, forming either a neutron star or black hole. The effect of the supernova event on the binary is also poorly understood. We suspect that the supernova

 $^{^{2}}$ Mass transfer also tends to equalize the masses of the components, which is preferred from a search point of view; equal mass systems are intrinsically louder in GWs compared to high mass ratio binaries, and we have better confidence in our waveforms in this limit.

explosion is aspherical, imparting an impulsive "kick" to the remnant, but the degree of asphericity is unknown. A sufficiently strong kick may potentially disrupt the binary. We know that at least some supernova have large kicks, as we routinely observe pulsars (neutron stars) and γ -ray bursts far outside the plane of the host galaxy. Gamma-ray bursts are highly energetic events, thought to be due to the post-merger accretion of disrupted neutron star matter in a compact binary merger. If this hypothesis is true, then the observation of γ -ray bursts not closely associated with any host galaxy also points to the possibility of large natal kicks from supernovae.

If the binary survives the explosion of the primary, we are halfway towards forming a compact binary. The whole process must then be repeated for the companion star. The secondary star evolves off the main sequence and expands into a red giant. The expanding secondary may overflow its Roche lobe or even form a second common envelope with the primary. If this occurs, then mass, energy, and angular momentum are redistributed, as before, and the transfer may widen or tighten the binary orbit, depending on the details. Eventually the companion star explodes, potentially disrupting the binary. Finally, if the binary survives both supernovae and the components remain in sufficiently close orbit, then we are left with a compact binary which will coalesce by the emission of gravitational waves. Although this mechanism provides a lot of opportunity for the binary to fall off the path to forming a tight orbit compact binary, remarkably there is strong observational evidence that the branching fractions are non-zero. This conclusion is also supported by theoretical simulations of populations of binaries (see Sec. 1.4).

Another important factor in the binary evolution is the *metallicity* (abundance of elements other than hydrogen and helium) of the progenitor stars. The metallicity determines the opacity of the star's outer layers and thereby determines the strength of the stellar winds which push on these layers with photons. Low metallicity stars are less susceptible to mass loss through stellar winds, favoring the formation of higher mass systems. In addition to carrying away precious mass required to create compact remnants, mass loss through stellar winds also *widens* the binary orbit. Thus, low metallicity environments are considered the best birth places for short period compact binaries, and these environments particularly favor the formations of higher mass systems.

There is a strong connection between the component spin magnitudes and orientations and the compact binary formation mechanism. Due to mass transfer during the common evolution stage, the resulting compact binary is expected to have significant spins. The transferred mass rotates in the plane of the orbit, and the resulting spins are therefore expected to be nearly aligned to the orbital angular momentum. The collapse of the star in its supernova stage will further increase the spin of the components. Supernova kicks may also impart a tilt to the spin relative to the orbital angular momentum, and so this formation mechanism does not guarantee any particular alignment of the spin vectors. Observations of these events in gravitational waves will help us to understand the efficiency of the common evolution formation mechanism, and in particular the role of supernova kicks and the common envelope phase in binary stellar evolution.

1.3.2 Dynamical capture

Another possible formation channel is the dynamical capture of two black holes in dense stellar environments as a result of three body interactions. Simulations of dense stellar environments, such as globular clusters, indicate that relativistic compact binaries may form with appreciable rates. Furthermore, relativistic binary neutron stars have been associated with globular clusters [29], where there is little ongoing star formation.

One possible mechanism for forming tight binaries from many body interactions is the *Kozai* mechanism. Consider a tightly bound binary system interacting with a third, loosely bound compact object. Conceptually, we can think of this system as consisting of two binary systems: (1) the tight inner orbit between masses m_1 and m_2 , and (2) the outer orbit of the third mass m_3 (effectively) about the center of mass of the first binary system. One can show then that the conservation of angular momentum in this three body system implies that the quantity

$$L = \sqrt{1 - e^2} \cos \iota \tag{1.57}$$

is a constant of motion. Here e is the eccentricity of the tight inner orbit and ι is the angle between the two orbital momenta. Thus, the inner binary, in its interaction with the third body, can trade orbital plane orientation with eccentricity. Eccentric orbits emit greater power in gravitational waves and evolve quickly compared to binaries in circular orbits (see Eqn. 1.58). The increased emission of gravitational waves from the inner binary due to its high orbital eccentricity may further tighten the inner orbit on a time scale comparable to the period of the outer orbit. Thus, the flyby of the third body repeatedly drives the inner binary into an elliptic orbit, which by the emission of gravitational waves shrinks the orbital separation, leading potentially to the formation of a relativistic compact binary. In contrast to those formed through common evolution, black hole binaries formed directly through dynamical capture in dense stellar environments are expected to have spins arbitrarily oriented relative to each other and to the orbital angular momentum [30]. We therefore may be able to distinguish between the dynamical capture mechanism and the common evolution mechanism by measuring the component spins with gravitational wave observations.

1.4 Coalescence rates of compact binaries

We now give estimates for the intrinsic compact binary coalescence rate and fold in detector sensitivity models to obtain estimates for the aLIGO detection rate [31]. Our knowledge of the rates of CBCs comes from a combination of interpretation of astrophysical observations and predictions based on simulated models of binary stellar evolution and the dynamical interactions in dense stellar environments. Presently, the observational input is growing, but is still insufficient for making tight event rate confidence intervals. Simulations are plagued by large uncertainties in the key processes in the hypothesized formation mechanisms. Consequently, the various methods of estimating the rates often give vastly different predictions, and the event rates for all three compact binary source types are highly uncertain, spanning several orders of magnitude.

The Hulse-Taylor binary pulsar (PSR B1913+16) is the iconic example of a relativistic compact binary system [32]. A pulsar is a rapidly rotating neutron star beaming electromagnetic radiation along an axis tilted relative to its spin axis. The beam sweeps out a cone, and if Earth happens to lie within that cone, we will observe the neutron star as regular pulses of radio waves. The Hulse-Taylor binary was first observed as a millisecond pulsar with period $P = 59.03 \times 10^{-3}$ seconds. It was later observed that the time intervals between pulses from B1913+16 were being modulated at a frequency of 7.75 hr, implying the existence of a second, unseen companion star. The period of the pulse modulation gives a direct measurement of the orbital period for the binary system and indirect measurements for the masses of the components. At present, the mass of the pulsar in B1913+16 is known to be $m_p = 1.4414 \pm 0002 M_{\odot}$ and the unseen companion star has a mass $m_c = 1.3867 \pm 0.0002 M_{\odot}$ [33].

What makes this binary particularly interesting is that it is relativistic, with orbital speeds on the order of 10^{-3} c. Over the course of the last thirty years, observations of the Hulse-Taylor binary system have shown a decrease in the orbital period of the binary. The period decrease due to the


Figure 1.6: The Hulse-Taylor binary pulsar. The Hulse-Taylor binary consists of two neutron stars in an eccentric 7.75 hr orbit. Due to the emission of gravitational waves, the orbital period has decreased by about 40 seconds over 30 years in agreement with the predictions of General Relativity given by Eqn. 1.58 to 0.2% [33]. It stands as one of the cornerstone pieces of observational evidence for the existence of gravitational waves, and is the first known example of a relativistic compact binary object.

emission of gravitational waves is predicted by GR to be given by

$$\frac{dP}{dt} = -\frac{192\pi G^{5/3}}{5c^5} \left(\frac{P}{2\pi}\right)^{-5/3} m_1 m_2 (m_1 + m_2)^{-1/3} f(e), \qquad (1.58)$$

where e is the orbital eccentricity and

$$f(e) = (1 - e^2)^{-7/2} \left(1 + \frac{73}{24}e^3 + \frac{37}{96}e^4 \right).$$
(1.59)

One can check that with e = 0, the expression for the period derivative given above coincides with the results derived in Sec. 1.2. All physical quantities in Eqn. 1.58 can be measured from the pulsar observations and therefore the period decrease prediction is one with *no free parameters*. In a remarkable triumph for the theory of General Relativity, the measurements of the orbital parameters agree with the prediction in Eqn. 1.58 to better than 0.2% (see Fig. 1.6) over thirty years of observation [33].

An indirect piece of evidence for the existence of relativistic binaries with neutron stars (i.e., BNS or NSBH) comes from the observation of short hard γ -ray bursts. These events are highly energetic ($E \sim 10^{51}$ erg) but are very short, typically lasting less than one second. The timescale of the burst itself suggests an origin in compact objects; the corresponding length scale is given by $d \sim c/\Delta t \approx 1R_{\odot}$. This observation, combined with the extreme energy scale of the event, suggests

Table 1.1: Known relativistic binary pulsars. pulsar systems. These systems have orbital periods $f_{\rm orb}$ which are sufficiently short that the time to coalescence τ_{GW} of the system due to gravitational wave emission is small compared to the age of the Universe. Observations of the orbit give an estimate for the total mass $M_{\rm total}$, which determines the post-Newtonian parameter $v \equiv (2\pi G f_{\rm orb}/c^3)^{1/3}$. The short orbital periods imply short orbital separations and indicate that the pulsar companion must be a compact object. Based on component mass estimates, all binary systems listed here are thought to consist of two neutron stars (rather than a neutron star and a black hole). In J0737-3039, both components of the binary are seen as pulsars and in particular both have measured spin magnitudes. The system B2127+11C is unique among these examples in that it is found within the globular cluster M15 [29], suggesting an origin in dynamical capture.

pulsar	$f_{\rm psr} \ ({\rm ms})$	$f_{\rm orb}$ (hours)	$M_{\rm total} \ (M_{\odot})$	$v (10^{-3})$	τ_{GW} (Myr)
J0737-3039A	22.70	2.45	2.58	2.1	87
J0737-3039B	2773	_	_	_	_
B2127 + 11C	30.53	8.04	2.71	1.4	220
J1906 + 073	144.1	2.78	2.61	2.0	300
B1913 + 16	59.03	7.75	2.83	1.5	310
J1756-2251	28.46	7.68	2.57	1.4	1690
B1534 + 12	37.90	10.1	2.75	1.3	2690

that these bursts could be the result of disrupted neutron star matter falling back onto a remnant black hole after a binary merger. If this association is indeed correct, then short hard γ -ray bursts serve as further evidence for the abundance of relativistic binary neutron stars, and suggests the possibility for the existence of yet to be observed neutron star-black hole binaries.

The discovery of the Hulse-Taylor binary and other binary pulsars discovered later (see Tbl. 1.1) provides an empirical estimate for how often relativistic binary systems form and coalesce. In an early paper on the subject, Phinney estimated [34] a merger rate of $R_{\rm MW} \sim 10^{-6} {\rm yr}^{-1}$ in the Milky Way based on the existing pulsar observations at the time. To extrapolate this result beyond the Milky Way, traditionally it has been assumed that the coalescence rate is proportional to the stellar birth rate, which in spiral galaxies is proportional to the galaxy's blue light luminosity L_B . However, there are a number of problems with using blue light luminosity as a proxy for coalescence rate. For one, it ignores contributions to the event rate from old galaxies with little star formation activity but plenty of evolved stars (black holes) to form relativistic compact binaries through dynamical capture. Older galaxies have also been associated with short hard γ -ray bursts [35], which if the compact binary merger progenitor hypothesis is true, indicates the presence of compact binaries involving neutron stars. Furthermore, calculating the blue light luminosity out to a given distance requires a galaxy catalog complete out to that distance with accurate L_B measurements [36], which on the scale of recent and anticipated LIGO sensitivities are not currently available. Fortunately, beyond a distance of roughly 20 Mpc, the local mass inhomogeneities average out to a constant blue light density of $1 \text{Mpc}^3 \approx 0.0198 L_{10}$, where $L_{10} \equiv 10^{10} L_{B,\odot}$ and $L_{B,\odot} = 2.16 \times 10^{33}$ erg/s is the

Table 1.2: Known binaries with black hole companions. We list the binary parameters for several high mass x-ray binaries which are thought to consist of an evolved massive star in orbit with a black hole companion. These systems strongly suggest that the common evolution formation channel (see Fig. 1.5) occurs with non-zero rate for binary black holes and neutron star-black hole binaries.

system	$m_{ m BH}$	$m_{ m c}$	
IC10 X-1	$\geq 25 M_{\odot} \; [37]$	$\sim 35 M_{\odot}$	
NGC300 X-1	$\sim 20 M_{\odot} \ [42]$	$\sim 26 M_{\odot}$	
Cyg X–3	$2 - 5 M_{\odot}$ [39]	$8-14M_{\odot}$	

solar blue light luminosity. This allows us to compute volumetric rate densities, which are void of these assumptions. Using the estimate $L_B = 1.7L_{10}$ for the blue light luminosity of the Milky Way, we can then convert the galactic rate estimates to a volumetric rate density.

The extent of observational knowledge of neutron star-black hole binaries and binary black holes is considerably sparser, and consequently the event rate for these systems is even more uncertain. There are no known relativistic NSBH or BBH systems. However, there are a few examples of massive stars thought to be gravitationally bound to a stellar mass black hole, and which are promising candidate progenitor systems for relativistic black hole binaries [37, 38] or possibly neutron star-black hole binaries [39], as listed in Tbl. 1.2. These systems are observed as strong x-ray sources associated with an evolved massive star [40]. The observed star exhibits periodic Doppler shifts in its spectra, indicating the presence of an unseen companion star. The high energy radiation is attributed to accretion of matter from the primary star onto a compact companion object. Measurements of the binary parameters (in particular, the component masses) imply that the compact object must be a black hole. Regardless of the fate of these particular systems, these systems and others like them [41] are important examples of high mass binary systems which have survived both the common envelope and supernova stages, maintaining a close orbit. After completion of the second (current) mass transfer stage, the remaining star will evolve into a black hole. Provided the second supernova explosion does not disrupt the binary and the two objects remain in a tight orbit, these systems may give rise to relativistic binary black holes. Some theoretical considerations suggest this outcome will indeed be the case, implying quite high rates for binary black hole mergers [38], although these predictions are predicated on assumptions for the uncertain future evolution of the known x-ray binaries.

A complementary approach to estimating the CBC event rate is through simulation of a population of stellar systems. As discussed in the previous section, the details of the processes that drive the creation of relativistic binaries are not well-constrained, and consequently present simulations

Table 1.3: Event rates predictions for compact binary coalescence. We summarize the various estimates for the rate of compact binary coalescences. At the top of the table, we show estimates for the *intrinsic rate* densities. Combining these rate density estimates with estimates for the sensitive distance of various detector configurations, we show below this the estimated detection rate for various detector configurations. The rate estimates given here result from the combination of theoretical simulations and observational input, as discussed in the text. Where such information is available, $R_{\rm re}$ refers to the mean of the posterior probability distribution function for the rates, $R_{\rm low}$ and $R_{\rm high}$ are the 95% confidence intervals, and $R_{\rm min}$ and $R_{\rm max}$ are the ultra-conservative lower and upper limits, respectively. The neutron star mass is assumed $m_{\rm NS} = 1.4 M_{\odot}$ and the black hole mass is assumed $m_{\rm BH} = 10 M_{\odot}$.

Source			$R_{\rm low}$	$R_{\rm re}$	$R_{\rm high}$	$R_{\rm max}$
NS-NS $(Mpc^{-3} Myr^{-1})$			0.01 [43]	1 [43]	10 [43]	50 [44]
NS-BH $(Mpc^{-3} Myr^{-1})$			6×10^{-4} [45]	0.03 [45]	1 [45]	
BH-BH $(Mpc^{-3} Myr^{-1})$			1×10^{-4} [46]	0.005 [46]	0.3 [46]	
Source	IFO	$D_{\rm horizon}$	$\dot{N}_{ m low}$	$\dot{N}_{\rm re}$	$\dot{N}_{\rm high}$	$\dot{N}_{\rm max}$
		Mpc	yr^{-1}	yr^{-1}	yr^{-1}	yr^{-1}
	Initial	30	2×10^{-4}	0.02	0.2	1
NS-NS	Enhanced	50	0.001	0.1	1	5
	Advanced	450	0.4	40	400	2000
	Initial	70	9×10^{-5}	0.006	0.2	
NS-BH	Enhanced	120	7×10^{-4}	0.04	1	
	Advanced	930	0.2	10	300	
	Initial	160	2×10^{-4}	0.009	0.7	
BH-BH	Enhanced	270	0.002	0.07	5	
	Advanced	2200	0.5	20	1000	

are not capable of making precise predictions for the formation rates. Depending on the choice of model parameters, the predicted rates can vary by orders of magnitude. The literature on this subject is vast, and the techniques and results vary dramatically from study to study.

After combining observational and theoretical considerations, there remain significant uncertainties in the astrophysical rates for compact binary coalescences. In a systematic survey of the recent rate prediction literature, we obtain a range of possible rates [31], shown in Tbl. 1.3 (top). With basic assumptions about the sensitive distance for LIGO compact binary searches, we convert these intrinsic event rates into detection rates, also shown in Tbl. 1.3 (bottom). These calculations indicate that a realistic expectation for late stage aLIGO is the order of ten detections in each target category per year of observational time: 40 for binary neutron stars, 20 for binary black holes and 10 for neutron star-black holes.

1.5 Compact binaries with spinning components

The most recent LIGO/Virgo searches for gravitational waves from CBCs [3,4,47,48] have searched with waveforms that neglect the effects of spin. This choice was made on the basis of studies which concluded that including spin effects in the search filters does not improve the search sensitivity relative to a search that neglects spin effects, even to target signals which are spinning [49, 50]. In the case of higher mass binary black hole searches, where all of inspiral, merger, and ringdown are important for detection, models incorporating spin effects were simply not available. We now have available several new technologies for advancing the sensitivity of CBC searches towards spinning systems, which we will describe throughout this work. Here we describe two recently developed waveform models, which include leading order spin effects for aligned spin systems. We will be using these models later to develop a search pipeline which is *more sensitive* to spinning signals when using aligned spin filters, even in an analysis of real LIGO data.

From both a data analysis and astrophysical perspective, it is convenient to separate the possible cases into three categories. In the non-spinning case, both components have zero spin angular momentum, and nothing special happens. In the aligned spin case, the spin vectors are non-zero and point parallel to the orbital angular momentum. The orbital plane remains fixed, and the effect of spin on the dynamics is a systematic lag (or lead) of the amplitude and phase evolution. Finally, if the spin vectors are non-zero but misaligned with the orbital angular momentum, then the spins and orbital angular momenta precess about the total angular momentum and appear in the GW signal as a quasi-periodic modulation of the amplitude and phase. See Fig. 1.7 and Fig. 1.8 for examples of these three cases.

The technical difficulty of developing an effective templated search increases sharply as one transitions from considering non-spinning to aligned spin to misaligned spin templates. In the non-spinning case, only two mass parameters are required to cover the template parameter space. For the aligned spin case, two additional parameters – the spin magnitudes of the two components – are required to specify the waveform. In the misaligned case, we must search over nine parameters: two mass parameters, six spin parameters, and the initial inclination of the binary³. As the number of model parameters increases, so do the number of templates required for filtering, the computational cost of the search, and number of background events.

As discussed in Sec. 1.3, there is a strong connection between the spin orientations and magnitudes and the mechanism by which the binary formed. In particular, compact binaries formed through ordinary binary stellar evolution are expected to have large spins that are nearly aligned with their

³The inclination is constant for non-precessing systems, and otherwise equivalent signals with different inclinations are proportional since the inclination appears only in the overall amplitude. For precessing systems, the inclination evolves in time and the non-trivial amplitude evolution removes the simple relationship between equivalent systems with differing initial inclinations.



Figure 1.7: Effect of spin magnitude and orientation on the inspiral signal. Here we illustrate a post-Newtonian approximation (known in the literature as SPINTAYLORT4) to the inspiral portion of the gravitational wave signal from compact binaries with different spin orientations. In each waveform, the masses of the system are fixed to $m_1 = 15M_{\odot}$ and $m_2 = 3M_{\odot}$ at a distance d = 1 Mpc from Earth. For each waveform, we show the + polarization portion of the signal from $f_{\rm GW} = 40$ Hz to $f_{\rm ISCO}$. In the top example, the spins are aligned with the orbital angular momentum \vec{L} . In the bottom example, one of the spins is orthogonal to \vec{L} and the other is zero. The middle example shows a non-spinning system. We see here the lengthening of the waveform in the (positively) aligned spin case relative to the non-spinning case, and clear amplitude modulation effects in the misaligned spin case.

orbital angular momentum, while those formed through dynamical capture should have arbitrarily oriented spins. Thus, from an astrophysical perspective, when interpreting a CBC event observed in gravitational waves, it is of great interest to know to which of the above categories the event belongs.

1.5.1 Post-Newtownian reduced spin approximation

We consider first waveforms in the post-Newtonian stationary phase approximation, which are wellsuited for the search and analysis of low mass systems (~ $1 - 15M_{\odot}$), which merge at higher frequencies. In this approximation, the coupling between the spin angular momenta \vec{S}_1 and \vec{S}_2 and



Figure 1.8: The IMRPHENOMB waveform approximation for aligned spin binary black holes. We plot the late inspiral, merger, and ringdown portions of the gravitational wave signal computed with the IMRPHE-NOMB approximation for three different aligned spin systems. Across the three examples, we have changed only the aligned spin parameter χ_{eff} . Each waveform consists of the + polarization from an equal mass binary with $m_1 = m_2 = 10 M_{\odot}$ at a distance D = 1 Mpc and is shown starting at $f_{\text{GW}} = 85$ Hz. We note that the positively aligned spin system evolves more slowly compared to the non-spinning system and that the non-spinning system evolves more slowly than the negatively spinning system.

the Newtonian orbital angular momentum \vec{L}_N is governed by the system of differential equations

$$\frac{dS_i}{dt} = \vec{\Omega}_i \times \vec{S}_i \tag{1.60}$$

$$\frac{d\vec{L}_N}{dt} = \frac{-1}{||\vec{L}_N||} \frac{d}{dt} \left(\vec{S}_1 + \vec{S}_2\right)$$

$$(1.61)$$

$$\frac{dv}{dt} = -\frac{\mathcal{F}}{dE/dv}.$$
(1.62)

The vectors $\vec{\Omega}_i$ point along the spin precession axis and have magnitudes equal to the instantaneous precession frequency. They are computed as an expansion in the post-Newtonian parameter v and depend on both the spins and orbital angular momentum vectors. To leading order in v, $\vec{\Omega}_i$ point along the direction of the Newtonian angular momentum \vec{L}_N . The orbital binding energy E(v) and gravitational wave flux \mathcal{F} are also computed as expansions in v and depend on the spins and orbit. These equations contain ten coupled non-linear differential equations, which become increasingly more non-linear as you increase the order of the post-Newtonian expansion. They are non-trivial to solve analytically but can be solved numerically for systems with generic spin magnitudes and orientations [18] up to the available post-Newtonian order for E(v) and \mathcal{F} .

If the spins are parallel to the Newtonian angular momentum, then the vectors $\vec{\Omega}_i$ are both parallel to \vec{L}_N to all post-Newtonian orders. The spin and orbital angular momentum evolution becomes trivial:

$$\frac{d\vec{S}_i}{dt} = \frac{d\hat{L}_N}{dt} = 0, \tag{1.63}$$

leaving only Eqn. 1.62 left to solve. The remaining equation depends on only four parameters: two mass parameter and two spin parameters. Recently, a *reduced spin* template family computed in the post-Newtonian approximation [25] called TAYLORF2REDSPIN was developed to capture leading order spin effects for aligned spin systems with only one additional physical parameter. In this model, the gravitational wave signal is given by

$$\tilde{h}(f) \equiv \mathcal{C} f^{-7/6} \exp\{-i[\Psi(f) - \pi/4]\}, \qquad (1.64)$$

where C is a constant that depends on the relative sky position and orientation of the binary with respect to the detector, and f is the Fourier frequency. We have kept only the leading term in the frequency-domain amplitude. The phase of the GW signal is given by

$$\begin{split} \Psi(f) &= 2\pi f t_0 + \phi_0 + \frac{3}{128\eta v^5} \left\{ 1 + v^2 \left[\frac{55\eta}{9} + \frac{3715}{756} \right] \right. \\ &+ v^3 \left[4\beta - 16\pi \right] \\ &+ v^4 \left[\frac{3085\eta^2}{72} + \frac{27145\eta}{504} + \frac{15293365}{508032} - 10 \sigma_0 \right] \\ &+ v^5 \left[\frac{38645\pi}{756} - \frac{65\pi\eta}{9} - \gamma_0 \right] \left(3\ln(v) + 1 \right) \\ &+ v^6 \left[-\frac{6848\gamma_E}{21} - \frac{127825\eta^3}{1296} + \frac{76055\eta^2}{1728} \right] \\ &+ \left(\frac{2255\pi^2}{12} - \frac{15737765635}{3048192} \right) \eta - \frac{640\pi^2}{3} \\ &+ \frac{11583231236531}{4694215680} - \frac{6848\ln(4v)}{21} \right] \\ &+ v^7 \left[-\frac{74045\pi\eta^2}{756} + \frac{378515\pi\eta}{1512} + \frac{77096675\pi}{254016} \right] \right\}, \end{split}$$
(1.65)

where t_0 is the time of arrival of the signal at the detector, ϕ_0 the corresponding phase, $v \equiv (\pi M f_{\rm GW})^{1/3}$ as defined above, and γ_E is the Euler gamma. Note that spin terms do not appear in

the post-Newtonian expansion until 1.5PN order. The spin effects in the waveform are completely known up to 2.5PN order (v^5) , and are described by the following parameters:

$$\beta = \frac{113 \chi_{\rm red}}{12,}$$

$$\sigma_0 = \left(-\frac{12769 (4\eta - 81)}{16 (76\eta - 113)^2}\right) \chi^2_{\rm red},$$

$$\gamma_0 = \left(\frac{565 (17136\eta^2 + 135856\eta - 146597)}{2268 (76\eta - 113)}\right) \chi_{\rm red},$$
(1.66)

where $\chi_{\rm red}$ is the reduced-spin parameter. Defining the dimensionless spin $\vec{\chi}_i = c\vec{S}_i/Gm_i^2$, the symmetric and antisymmetric combinations of the spins

$$\chi_{\rm s} \equiv \frac{1}{2} (\vec{\chi}_1 + \vec{\chi}_2) \cdot \hat{L}$$
 (1.67)

$$\chi_{\rm a} \equiv \frac{1}{2} (\vec{\chi}_1 - \vec{\chi}_2) \cdot \hat{L},$$
(1.68)

and the asymmetric mass ratio $\delta \equiv (m_1 - m_2)/m$, the reduced spin parameter is given by

$$\chi_{\rm red} \equiv \chi_{\rm s} + \delta \chi_{\rm a} - \frac{76\eta}{113} \chi_{\rm s}.$$
(1.69)

One of the main advantages of this waveform family for data analysis is that it reduces the parameter space of aligned spin systems from four to three dimensions, making it appealing as a first step towards integrating spin effects into our analysis pipelines. We demonstrate in Chap. 6 that these waveforms, using only three parameters, capture greater than 90% of the search volume for generically spinning binaries with $m_1 + m_2 \leq 12M_{\odot}$, and furthermore using these waveforms in a pipeline could give up to a 50% increase in detection rate for binary neutron star systems⁴.

1.5.2 Effective spin inspiral-merger-ringdown approximation

As we have mentioned above, the adiabatic approach to computing the gravitational wave signal from CBCs becomes invalid as $v \to 1$. In this case, we resort to numerical solutions of the Einstein equations. However, these simulations are computationally intensive and the parameter space for compact binaries is very large, including one mass (solutions scale in total mass) and six spin

⁴These numbers are quoted for volume at fixed signal-to-noise ratio. See Chap. 6 for more detail.



Figure 1.9: The dimensionality problem for numerical relativity simulations of compact binary mergers. We show the available numerical relativity simulations from the NINJA2 catalog [51]. Numerical relativity solutions scale in total mass, so there is only one mass parameter to cover, but there are six spin parameters. Here we only show the aligned spin simulations, in which there are only two extra parameters χ_1 and χ_2 . The paucity of coverage in this parameter space makes it difficult to construct even phenomenological models that interpolate between the solutions.

parameters. Covering the whole space with numerical simulations to the level that is useful for gravitational wave searches is at present impractical, as illustrated in Fig. 1.9. However, the numerical solutions can still provide insight into the basic behavior of systems during the merger. For the purpose of using the waveforms in gravitational wave searches, we can use these simulations as inspiration for phenomenological models that cover the parameter space continuously.

First, we have to develop hybrid waveforms, which extend the short numerical relativity solution down to the early inspiral. Hybridization works by choosing a frequency interval $[f_1, f_2]$, which includes the early part of the numerical simulation, and "matching" the NR waveform to a post-Newtonian one in that interval. The parameters for the post-Newtonian system need not be the same as the parameters of the numerically simulated system. Instead, a suitable post-Newtonian signal is chosen by performing a least squares analysis over the chosen frequency band, allowing the post-Newtonian source parameters to vary; this least squares fitting can be performed on phase, amplitude, or some other quantity of interest. Finally, the two waveforms are combined in the transition region $[f_1, f_2]$ by a weighted sum such that at f_1 the waveform exactly matches the post-Newtonian waveform, and at f_2 the waveform exactly matches the numerical relativity waveform. The IMRPHENOMB model [19] takes a collection of such hybrid waveforms to fit a parameterized ansatz for the amplitude and phase of the complete IMR signal, motivated by both the known form of the signal in the post-Newtonian limit, the known form of the signal in the ringdown limit, and the anecdotal understanding of the waveform in the merger regime. Our primary interest in IMRPHENOMB comes from the fact that it models aligned spin effects and does so with only one spin parameter, the *effective spin*, given by

$$\chi_{\text{eff}} = \frac{m_1 \chi_1 + m_2 \chi_2}{m_1 + m_2}.$$
(1.70)

We emphasize that this parameter is not the same as the χ_{red} parameter introduced above (except in the limit that $\eta \to 0$). We then introduce phenomenological parameters and an ansatz for the shape of the signal,

$$A(f) = Cf^{-7/6} \begin{cases} 1 + \alpha_2 v^2 + \alpha_3 v^3 & \text{if } f < f_1 \\ 1 + \epsilon_1 v + \epsilon_2 v^2 & \text{if } f_1 \le f \le f_2 \\ \mathcal{L}(f, f_2, \sigma) & \text{if } f_2 < f < f_3 \end{cases}$$
(1.71)

$$\Psi(f) = 2\pi f t_0 + \phi_0 + \frac{3}{128\eta v^5} \left(1 + \sum_{k=2}^7 \psi_k v^k \right).$$
(1.72)

The three regimes correspond to the inspiral, merger, and ringdown portions of the binary evolution. Note that the form of the amplitude and phase in the $f \leq f_1$ (inspiral) regime are functionally the same as in the post-Newtonian expansion, though the coefficients are now free parameters for the model. The functional form for the merger is an ansatz inspired by available numerical simulations of the binary coalescence. The final stage (ringdown) of the binary evolution is modeled with a one-sided Lorenztian, which asymptotically goes as f^{-2} . The hybrid waveforms used to develop this model were constructed from numerical relativity simulations which spanned the parameter range $1 \leq q \leq 4$ and $-0.85 \leq \chi_{\text{eff}} \leq 0.85$.

As with the TAYLORF2REDSPIN family, the IMRPHENOMB waveform family also has the appeal of the reduced dimensionality of the parameter space, at least from the perspective of designing a search. Below we use this waveform family to implement a search pipeline for coalescing binary black holes with aligned spin, and demonstrate an analysis – in real LIGO detector noise, obtained from LIGO's fifth science run – which improves in sensitivity to signals from spinning systems when spin effects are included in the model waveforms.

34

1.6 Overview of the remainder of this thesis

This chapter has set the stage for the work to follow. We now understand the shape of the gravitational wave signal from compact binary coalescence. We know that there are multiple channels by which these systems may form, but we do not really know which of these channels are actually realized or at what rate. Nonetheless, a combination of observational evidence and theoretical simulations indicates that these systems form frequently enough that coalescence events will very likely be observed in LIGO data.

We have emphasized here the importance of knowing the signal waveform for detection and the dependence of the signal on the binary parameters. In particular, we have seen that the spins of the components in a compact binary can significantly alter the shape of the waveform, and the processes which form compact binaries naturally lead to components with large spins. As we report in Chap. 4, recent searches of LIGO and Virgo data for CBCs have used non-spinning templates. The inclusion of spin effects in the templates used for a search pipeline has never been successfully demonstrated to improve the sensitivity of the pipeline. This issue will form one of the primary foci for this work. We will show that the incorporation of spin effects in CBC search templates can vastly improve the extraction of these signals from the detector data. The returns on these results could turn out to be mission-critical if, as we expect, most compact binaries have highly spinning components.

Chapter 2

Laser Interferometric Gravitational Wave Observatories

Dreams about the future are always filled with gadgets.

Neil deGrasse Tyson

This chapter introduces the basic theory of gravitational wave measurement by laser interferometry in ground based gravitational wave observatories such as LIGO [11], Virgo [12], GEO [13] and KAGRA [14]. We begin with a brief review of the design and properties of a simple Michelson interferometer, which forms the conceptual backbone for these instrument designs. We will find that detecting gravitational waves of typical astrophysical strain magnitudes $h \sim 10^{-21}$ with a Michelson interferometer is exceedingly challenging, even in principle. The interferometry technique, from the perspective of sensing the effects of gravitational waves, is limited by two basic challenges: (i) the ability to sense minute changes in light intensity at a photodiode, and (ii) the ability to construct a device with a sufficiently large effective optical baseline. We discuss each of these problems in turn and show that a 4 km folded-arm Michelson interferometer operating with an input laser power of 10 W can, in principle, achieve an astrophysically interesting strain sensitivity.

The initial discussion lays out a highly idealized instrument design for a LIGO interferometer in order to demonstrate, in simple terms, the *possibility* of using such a device to detect gravitational waves from astrophysical sources. In practice, implementing the design requires many additional techniques for suppressing changes in optical path lengths arising from effects other than gravitational waves. We discuss some of the basic noise sources which limit the sensitivity of these detectors, as well as techniques to minimize the impact of these noise sources. For more details on the vast subject of interferometer design and construction, we point the reader to Ref. [52].

We then report on the *achieved* sensitivity of the LIGO and Virgo detectors in their latest observations. By the end of these observations, the LIGO detectors were operating at a sensitivity which went beyond the initial predicted pen-and-paper design by as much as a factor of two in strain, equivalent to a factor of eight in observable volume. In Chap. 4, we will report on the results of two searches for gravitational waves for CBCs in data from these observations. Finally, we illustrate the antenna nature of a gravitational wave detector, emphasizing the need for multiple detectors for full sky coverage and source localization.

2.1 The Michelson interferometer

At its core, a LIGO detector¹ is a Michelson interferometer (see Fig. 2.1). The Michelson uses interference between two co-originating beams of light traveling along separate paths to measure the difference in length along those two paths. At the front end of the interferometer, a beam of coherent light with wavelength λ is incident upon a beam splitter. At the beam splitter, half of the incident field is reflected (vertically in Fig. 2.1) and the rest is transmitted (horizontally in Fig. 2.1). The split beams travel along separate paths of lengths L_1 and L_2 toward two mirrors with amplitude reflection coefficients r_1 and r_2 , and are partially reflected back to the beam splitter. At the second encounter, the beams are again split and two beams emerge from the Michelson. One beam is sent back toward the light source and another is passed to a photodetector.

The light arriving at the photodetector is a superposition of two wave fields. For an ideal Michelson interferometer with a lossless, infinitesimally thin 50/50 beam splitter, the amplitude is given by

$$\vec{E} = r_1 \frac{\vec{E}_0}{2} e^{2\pi i (2L_1/\lambda)} - r_2 \frac{\vec{E}_0}{2} e^{2\pi i (2L_2/\lambda)}, \qquad (2.1)$$

where \vec{E}_0 is the electric field amplitude of the light injected at the beam splitter. Note that the amplitude reflection coefficient at the beam splitter is opposite in sign when the field is incident from the right side. The intensity of beam is given by

$$P = |\vec{E}|^2 = \frac{|\vec{E}_0|^2}{4} (r_1^2 + r_2^2 - 2r_1 r_2 \cos(4\pi\Delta L/\lambda)), \qquad (2.2)$$

¹In this work, the acronym "LIGO" may refer to either (i) any gravitational wave detector based on the principle of laser interferometry or (ii) the specific observatory sites at Hanford and Livingston. Thus, the instruments at LIGO, Virgo, GEO and KAGRA, are collectively referred to as LIGO detectors.

where $\Delta L = L_1 - L_2$. We see from Eqn. 2.2 that the minimum output power $P_{\text{out}}/P_{\text{in}} = (r_1 - r_2)^2/4$ is obtained if the path lengths $\Delta L = 0 \mod \lambda/2$. On the other extreme, if the $\Delta L = \lambda/4 \mod \lambda/2$, then the maximal intensity $P_{\text{out}}/P_{\text{in}} = (r_1 + r_2)^2/4$ will be observed at the photodetector and none is reflected back to the source. In this way, the relative length of the two optical paths couples directly to the intensity of light observed at the output port and provides a sensitive method for determining length changes.

We make a first estimate of the strain sensitivity of a Michelson interferometer by asserting that signal detection requires that the mirror moves through a quarter wavelength, also referred to in this context as a *Michelson fringe*. In moving through a fringe, the intensity at the photodetector goes from its minimum to its maximum value. The difference between the maximum and minimum output power, given by $\Delta P/P_{in} = r_1r_2$, is maximized when the two mirrors are highly reflective with $r_1 \sim r_2 \sim 1$. Given this observation, we will assume for our estimate the best case scenario



Figure 2.1: *Michelson interferometer design.* At the beam splitter, half of the injected light is reflected (vertical beam) and half is transmitted (horizontal beam). The split beams are then passed to highly reflective end mirrors and sent back to the beam splitter. The light observed at the output port is the sum of the transmitted beam from the vertical arm at the beam splitter and the reflected beam from the horizontal arm at the beam splitter. As depicted by the arrow in the horizontal arm, changes in the position of an end mirror couples to the number of wavelengths traveled by a given beam and therefore couples to the intensity of the recombined beams, which is ultimately used to infer the relative lengths of the two arms.

 $r_1 = r_2 = 1$. Our sensing criterion then requires

$$\frac{4\pi\Delta L}{\lambda} = \frac{\pi}{4},\tag{2.3}$$

and implies a strain sensitivity of

$$h = \frac{\Delta L}{L} = \frac{1}{16} \frac{\lambda}{L}.$$
(2.4)

In order to reach strain sensitivity on the order $h \sim 10^{-21}$ characteristic of astrophysical sources, a Michelson interferometer operating at wavelength $\lambda \sim 1 \mu m$ would require a baseline of length

$$L \approx 10^{11} \text{ km},\tag{2.5}$$

larger than the major axis of Pluto's orbit. Clearly, a baseline of this magnitude is unreasonable for Earth-based observatories (or space-based ones for that matter). So how can the Michelson interferometry technique be useful for the detection of gravitational waves?

2.2 Length sensing in the shot noise limit

In Sec. 2.1, we estimated the strain sensitivity of a Michelson interferometer by assuming that movement through a full Michelson fringe was needed to observe a signal. Such a requirement necessitates absurdly large baselines for typical gravitational wave strains from astrophysical sources. We now refine our sensitivity estimate above by considering what is the smallest measurable change in light intensity, or equivalently, what is the smallest observable fringe movement.

The length sensing of the Michelson design is achieved via the intensity observed at the output port, given by Eqn. 2.2. From a classical point of view, intensity is a property of the electric field and arbitrarily precise measurements of the intensity are possible. In this case, the limit on fringe movement sensing is a *technical* limit: the precision of the measurement depends on the clever design and construction of the photodetector. However, in the quantum picture, light arrives in the form of photons in a Poisson (shot noise) process, and consequently the intensity observed at the photodetector fluctuates, even without any length changes in the optical cavities. In this case, the limit on fringe movement sensing is a *fundamental* limit: a measurement limited by shot noise cannot be improved, even in principle.

We now consider the shot noise limit to fringe movement sensing in a Michelson interferometer.

Let us suppose then that we have at the output port an ideal photodetector which can detect individual photons and therefore is limited only by shot noise. The photon flux \dot{N} at the photodetector is related to the intensity P at the photodetector by

$$\dot{N} = P \frac{\lambda}{2\pi\hbar c},\tag{2.6}$$

where $2\pi\hbar c/\lambda$ is the energy of an individual photon. The photodetector integrates the flux over a short time scale τ , during which the expected number of photons is $N = \dot{N}\tau$ and the shot noise in the photodetector is given by

$$\Delta N = \sqrt{N} = \sqrt{\frac{P\lambda\tau}{2\pi\hbar c}} \tag{2.7}$$

or equivalently, using Eqn. 2.6,

$$\Delta P = \frac{2\pi\hbar c}{\lambda\tau} \Delta N = \sqrt{P\lambda\tau 2\pi\hbar c}.$$
(2.8)

Eqn. 2.8 gives the magnitude of the fluctuations in intensity at the output port simply due to the statistical nature of the photon counting process. Changes in the intensity due to length changes in the arms of the interferometer must exceed the size of these fluctuations in order to be measurable.

Let us now translate this result into fringe movement sensitivity. The optimal operating point for the interferometer is where the intensity at the output port varies most strongly with length changes. From Eqn. 2.2, we find that this criterion gives

$$L_1 - L_2 = (2n+1)\frac{\lambda}{4} \tag{2.9}$$

where n is an integer. Thus, the optimal operating point is with the two arms having slightly different lengths and the power at the output port at the operating point is $P_{\text{out}} = P_{\text{in}}/2$. At this operating point, the intensity-length coupling is given by

$$dP = P_{\rm in} \frac{2\pi}{\lambda} d(\Delta L). \tag{2.10}$$

Equating $dP = \Delta P$, the latter given by Eqn. 2.8, we can determine the smallest resolvable strain movement, which is given by

$$h = \frac{1}{L} \sqrt{\frac{\hbar \lambda c}{4\pi P_{\rm in} \tau}}.$$
(2.11)

For the best sensitivity, we should choose the integration time τ to be as large as possible, but much smaller than the gravitational wave period, so that we are measuring in a sense the instantaneous strain in the detector. We will see later that ground based detectors are insensitive to gravitational waves with $f_{\rm GW} < 10$ Hz, or equivalently, $T_{\rm GW} > 100$ ms. Thus, an integration time $\tau \approx 10$ ms would be appropriate for all gravitational wave frequencies. Applying Eqn. 2.11, we find that with a modest laser power of $P_{\rm in} = 1W$ and $\lambda = 1\mu m$, gravitational wave strains of down to $h \sim 10^{-21}$ could be resolvable with a baseline of a "mere" 500km. Such a baseline is still unreasonably large for Earth-based detectors, and so one additional trick, to be discussed in the next section, is required.

Our discussion here has considered only the sensing capabilities of an ideal Michelson interferometer. We have neglected to consider whether such an interferometer can be held in a sufficiently stable configuration that length changes due to gravitational waves are dominant. We take up this latter topic in Sec. 2.4. Here we simply point out that we have neglected a related noise source arising from the quantum nature of light. Since photons carry with them a momentum $\hbar/2\pi\lambda$, the reflection of photons at the mirrors imparts a *radiation pressure* force, which fluctuates with the statistical fluctuations in the arrival times of the photons. Therefore, in addition to limiting the sensing capabilities of the interferometer, the quantum nature of light also induces length fluctuations not associated with gravitational waves. The radiation pressure effect becomes increasingly important as the circulating power increases; thus, contrary to what one might infer from Eqn. 2.11, one cannot achieve arbitrarily small strain sensitivities simply by increasing the laser power.

2.3 Fabry-Perot cavities

So far we have established that a Michelson interferometer with a baseline of 500 km operating at an input power of 1 W could in principle measure strains of the order $h \sim 10^{-21}$, typical of realistic astrophysical sources of gravitational waves. The problem with this design of course is that it requires one to build an enormous instrument. This problem is solved in LIGO and Virgo instruments by *folding* the interferometer arms. Since it is the optical path length rather than the physical path length that matters for length sensing, one can fold the interferometer arms so that the distance traveled by any single photon is typically many times the physical length of the instrument.

One way of folding the instrument arms is with a Fabry-Perot optical cavity, as depicted in Fig. 2.2. A Fabry-Perot cavity consists of two highly reflective mirrors, aligned such that an input



Figure 2.2: Fabry-Perot optical cavity. A Fabry-Perot optical cavity consists of two highly reflective mirrors facing each other. The use of such optical cavities in LIGO instruments in place of the usual Michelson end mirrors increases the optical path length from the input beam to the output beam at the photodiode, allowing for high sensitivity strain sensing with a baseline over 100 times smaller than would otherwise be required. This cartoon drastically over-simplifies the optical properties of Fabry-Perot cavities used in real LIGO instruments. The laser beam is depicted here for simplicity and clarity by geometric rays incident at an angle to flat cavity mirrors. In real applications, the laser beam has finite width, and great care must be taken to align the mirrors (which are curved) such that the laser beam is resonant in the cavity.

beam incurs multiple reflections before exiting the cavity. At every encounter with a mirror, some part of the light exits the cavity, and each of these transmitted beams interfere with each other. The light that remains in the cavity continues to bounce around, picking up a phase shift

$$\Delta \phi = \frac{2\pi L}{\lambda} \tag{2.12}$$

for each one-way trip down the cavity. The summed beam effectively reflected at the input mirror contains contributions from fields with any integral number of phase shifts $2n\Delta\phi$, and consequently has an intensity which is highly sensitive to changes in the length of the cavity. We will see here that by replacing the end mirrors in a Michelson interferometer with Fabry-Perot cavities, we can achieve the desired strain sensitivity with only kilometer-scale arm lengths.

Consider a Fabry-Perot cavity consisting of two highly reflective mirrors with reflection coefficients r_1 and r_2 and transmission coefficients t_1 and t_2 , corresponding to reflection and transmission from the left at the left and right mirrors, respectively, as shown in Fig. 2.2. A beam of light is incident normally from the left upon the first mirror (the rays are shown at an angle in Fig. 2.2 for clarity). The initial beam is partially reflected and partially transmitted. If the initial electric field is \vec{E}_0 , then the reflected field² is related to the input field by $\vec{E}_1^{\text{refl}} = r_1 \vec{E}_0$ and a field $\vec{E}_1 = t_1 \vec{E}_0$ enters the cavity. The transmitted field crosses the cavity and encounters the second mirror, picking up a

 $^{^{2}}$ For consistency with Sec. 2.1, the reflection coefficient at the initial mirror, which was the only mirror in the Michelson design, must be positive. Within the cavity, the reflection coefficients are then forced to be negative.

phase shift of $\Delta \phi = 2\pi L/\lambda$. Again, part of the beam is reflected and part is transmitted. We obtain a transmitted field at the second mirror $\vec{E}_1^{\text{trans}} = t_2 \vec{E}_1 e^{i\Delta\phi} = t_1 t_2 \vec{E}_0 e^{i\Delta\phi}$. This process continues ad *infinitum.* The reflected beam crosses the cavity, picking up an additional phase $\Delta \phi$. Part of this beam emerges from the left mirror, now with $\vec{E}_2^{\text{refl}} = -t_1^2 r_2 \vec{E}_0 e^{2i\Delta\phi}$, and the rest is reflected back into the cavity. Continuing in this way, we can compute the superimposed reflected and transmitted beams from the cavity via the infinite series:

$$\vec{E}_{1}^{\text{refl}} = r_{1}\vec{E}_{0}$$

$$\vec{E}_{2}^{\text{refl}} = -t_{1}^{2}r_{2}\vec{E}_{0}e^{i2\Delta\phi}$$

$$\vec{E}_{3}^{\text{refl}} = -t_{1}^{2}r_{1}r_{2}^{2}\vec{E}_{0}e^{i4\Delta\phi}$$

$$\vdots$$

$$\vec{E}_{n\geq2}^{\text{refl}} = -t_{1}^{2}r_{2}e^{2i\Delta\phi}\vec{E}_{0}\left(r_{1}r_{2}e^{i2\Delta\phi}\right)^{n-2}$$

$$\vdots$$

$$\vec{E}^{\text{refl}} = r_{1}\vec{E}_{0} - t_{1}^{2}r_{2}e^{2i\Delta\phi}\vec{E}_{0}\sum_{n=0}^{\infty}\left(r_{1}r_{2}e^{2i\Delta\phi}\right)^{n}$$

$$= \vec{E}_{0}\left(r_{1} - \frac{t_{1}^{2}r_{2}e^{2i\Delta\phi}}{1 - r_{1}r_{2}e^{2i\Delta\phi}}\right).$$
(2.13)

Note that reflection at the initial mirror from the right incurs an additional minus sign. In a similar manner, we can compute the electric field transmitted by the Fabry-Perot cavity, which results in

$$\vec{E}^{\text{trans}} = \vec{E}_0 \left(\frac{t_1 t_2 e^{i\Delta\phi}}{1 - r_1 r_2 e^{2i\Delta\phi}} \right).$$
(2.15)

The transmitted power is then

$$\frac{P_{\rm trans}}{P_{\rm in}} = \frac{t_1^2 t_2^2}{1 - 2r_1 r_2 \cos(2\Delta\phi) + r_1^2 r_2^2}$$
(2.16)

and the reflected power is given by $P_{\text{refl}} = P_{\text{in}} - P_{\text{trans}}$.

≓rof

Now imagine that we have replaced the two end mirrors of a Michelson interferometer with two Fabry-Perot cavities (see the boxed out region in Fig. 2.3). What is the intensity of the resulting field at the output port? We neglect the Michelson length degree of freedom, that is, the distance between the beam splitter and the initial mirror of the Fabry-Perot cavity. To simplify the calculations, we will assume that the initial mirrors are identical and the end mirrors are perfectly reflecting. That is,



Figure 2.3: Advanced LIGO optical layout. We show the basic optical subsystems of the aLIGO instrument design. The Michelson-Fabry-Perot model we have considered in this chapter, highlighted by the dashed box, forms only one of the many optical cavities. As mentioned in the caption for Fig. 2.2, our geometric optics approximation to the electric fields in the optical cavities is a gross oversimplification of the physics of these devices. In particular, the mirrors in a Fabry-Perot cavity are typically curved and the lowest resonant mode is not a standing plane wave. Fabry-Perot cavities have not one but infinitely many resonant modes. A typical laser beam has a profile similar to the fundamental mode of the cavity, but contains contributions which may excite higher order modes of the Fabry-Perot. The input and output mode cleaners remove these higher order modes.

we take $r_{1X} = r_{1Y} = r$ and $r_{2X} = r_{2Y} = 1$. Assuming lossless mirrors, the transmission coefficients are determined by conservation of energy $t^2 = 1 - r^2$. We then allow an input beam to pass through the beam splitter. As before, half of the incident light is reflected at the beam splitter and the other half is transmitted. These two beams are then incident upon their respective downstream Fabry-Perot cavities. Using Eqn. 2.14, we treat each Fabry-Perot cavity as a single mirror with effective reflection coefficient

$$r_i = r - \frac{(1 - r^2)e^{4\pi i L_i/\lambda}}{1 - re^{4\pi i L_i/\lambda}},$$
(2.17)

where $i \in \{X, Y\}$ and L_i is the length of the Fabry-Perot cavity in the *i*-th arm. Ignoring the Michelson degrees of freedom (the distance between the beam splitter and the initial mirrors) and applying Eqn. 2.1, we have for the electric field observed at the photodetector

$$\vec{E} = \frac{\vec{E}_0}{2}(r_X - r_Y) \tag{2.18}$$

and the power is given by

$$P = |\vec{E}|^2 = \frac{|\vec{E}_0|^2}{4} (|r_X|^2 + |r_Y|^2 - 2\text{Re}\{r_X r_Y^*\}).$$
(2.19)

Compare this result to Eqn. 2.2, which we derived for a simple Michelson interferometer. Note that since we have assumed lossless, perfectly reflecting end mirrors, the *magnitudes* of the now complex reflection coefficients r_X and r_Y are each unity, and we obtain

$$P = |\vec{E}|^2 = \frac{|\vec{E}_0|^2}{2} (1 - \operatorname{Re}\{r_X r_Y^*\}).$$
(2.20)

The sensitivity to length changes in the cavities comes (as before) in the cross terms in which the relative phase shifts between the two beams of light lead to interference. In Fig. 2.4, we show the dependence of the power observed at the output port on the relative length changes for a simple Michelson, and comparing this setup to a Michelson coupled to Fabry-Perot light storage cavities with perfectly reflecting end mirrors.

2.4 Limiting length stabilization noise

So far we have considered a highly idealized situation in which the only factor limiting our detector sensitivity is the ability to *sense* changes in distance. In this section, we consider some of the noise sources which are unfortunately easy to sense. These noise sources change the state of the detector, and if not properly treated will mask the gravitational wave signal.

We are considering here any mechanism which induces strain in the detector, but is not due to gravitational waves. For example, a seismic disturbance at the detector site could propagate all the way through to the suspensions which hold the test mirrors. The suspensions will then tug on the mirrors and generate a signal at the photodetector. The Poisson noise in photon arrival times contributes to motion of the mirrors through a fluctuating radiation pressure. The atoms of the



Figure 2.4: Length sensing of the Michelson and Michelson-Fabry-Perot interferometer designs. We compare the intensity of light observed at the output port of a simple Michelson interferometer (see Fig. 2.1) compared to a Michelson interferometer with Fabry-Perot light storage cavities (see Fig. 2.3). For a fixed cavity length and input power, the Fabry-Perot cavity clearly buys a significant amount of length sensitivity compared to the simple Michelson. The initial mirror reflectivities here are r = 0.99; for the Fabry-Perot, the end mirrors are assumed to be perfectly reflecting.

mirrors and mirror coatings, being at non-zero temperature, are constantly moving about, leading to fluctuations in the optical properties of the device, and thereby coupling directly to the phase of the circulating light. To further complicate matters, the mirrors and mirror coatings can change in temperature, perhaps by the absorption of energetic photons, leading to other unwanted changes in the optical properties of the mirrors and coatings.

The power spectrum (or power spectral density) of a noise process characterizes the average behavior of the noise. Suppose that h(t) represents the strain in a LIGO detector recorded over some finite stretch of time. We can compute the Fourier transform $\tilde{h}(f)$ and examine its *power spectrum* $|\tilde{h}(f)|^2$. The power spectrum tells you what the noise "sounds" like; literally, how much intensity is at a given frequency³. For a noise process, any particular measurement of h(t) will yield a different power spectrum. To characterize the process, we therefore average over possible realizations of the noise process. Since h(t) is a real-valued time series, its Fourier transform satisfies $\tilde{h}(f) = \tilde{h}^*(-f)$, where * denotes complex conjugation. It is customary to fold the positive and negative frequencies, such that the mean power spectrum is

$$\frac{1}{2}S_h(f) = \langle |\tilde{h}(f)|^2 \rangle.$$
(2.21)

The angle brackets indicate an average over an ensemble of measurements of h(t) running the same

³Note that the sensitive band of LIGO instruments is $\sim 10 - 10,000$ Hz (see Fig. 2.5) and coincides roughly with the sensitive band of the human ear, making the hearing analogy particularly apt.



Figure 2.5: Anticipated advanced LIGO noise budget. Here we show predictions for aLIGO strain sensitivity together with estimates of a few major contributing sources of noise. The power spectral density curves shown here correspond to the high power zero-detuned detector configuration. In this configuration, quantum noise in the readout and radiation pressure is a significant contributor to the noise budget at all frequencies. Mirror thermal noise contributes mostly to the sensitivity limit in the mid-frequency range. Thermal noise in the suspensions determines the low frequency sensitivity limit. The quadruple pendulum with a resonant frequency at ~ 1 Hz and f^{-8} noise suppression above the resonant frequency significantly reduces the impact of seismic noise.

random process.

We now take a look at three limiting sources of noise in LIGO gravitational wave detectors. The impact of these noise sources on our particular source target, the coalescing compact binary, depends on the shape of the signal and the power spectrum of the noise. Gravitational wave signals from CBCs are broadband but spend a large portion of their lives at low frequencies. Recall from Sec. 1.5 that the leading order term in the frequency domain for a GW signal from a CBC goes like $f^{-7/6}$. Thus, the low and mid-range frequencies are of primary interest for CBC searches. The mid-frequency range is particularly important for high mass systems, which merge in this band (and therefore the $f^{-7/6}$ law no longer holds). The mid-frequency band can also be important for parameter estimation or tests of General Relativity. However, since the gravitational wave signal from a CBC spends very little time at high frequencies, the high frequency band is generally irrelevant for the purposes of

our searches.

2.4.1 Seismic noise

Seismic noise was the limiting noise source at low frequency for initial LIGO detectors. In initial LIGO, each of the test mirrors was suspended at the bottom of a single pendulum. The pendulum passively isolates the test mass from seismic disturbances for frequencies above its resonant frequency. In the advanced LIGO detector design, the test mirrors are suspended at the end stage of a system of *four* successive pendulums, and this redesign dramatically reduces the impact of seismic noise on the detectors. Indeed, for the aLIGO instrument design, the seismic noise is so well suppressed that it no longer dominates the low-frequency noise budget (see Fig. 2.5).

The noise suppression mechanism of the pendulum can be understood in terms of a damped driven harmonic oscillator, with equation of motion

$$F_{\text{drive}}(t) = m l \ddot{\theta}(t) - F_{\text{frict}} + mg \sin(\theta(t))$$
(2.22)

$$\approx m l\ddot{\theta}(t) + b\dot{\theta} + m g \theta(t), \qquad (2.23)$$

where g is the local gravity, l is the length of the pendulum, and θ denotes the angular position of the pendulum arm relative to the vertical. We have made the usual small angle approximation in going from Eqn. 2.22 to Eqn. 2.23 and approximated the dissipative forces with a linear velocity damping relation $F_{\text{frict}} = -b\dot{\theta}$, where b is a measured constant. In the detector, the driving force might be motion at the pendulum suspension or a flip in a magnetic domain of one of the actuating mirrors, for example.

The behavior of the pendulum can be easily understood in the frequency domain. By Fourier transforming the equation of motion Eqn. 2.23, recalling that for Fourier transforms $(d/dt)^n = (-2\pi i f)^n = (-i\omega)^n$, we obtain an algebraic relation between the input $\tilde{F}_{drive}(f)$ and the output $\tilde{\theta}(f)$:

$$\tilde{F}_{\text{drive}}(f) = (-ml\omega^2 - ib\omega + mg)\tilde{\theta}(f).$$
(2.24)

For a linear system, such as we have in Eqn. 2.23, the *transfer function* is defined as the ratio of the output to the input in the frequency domain, and it uniquely characterizes that linear system (the transfer function can be reversed to infer the original equations of motion). The transfer function

between the pendulum suspension point and the test mass motion is given by

$$T_1(f) = \frac{1}{-ml\omega^2 - ib\omega + mg}.$$
 (2.25)

Defining $\omega_{\rm res} = \sqrt{g/l}$ and $b_0 = \omega_{\rm res}(b/mg)$, we can also write

$$T_1(f) = \frac{1/mg}{1 + ib_0(\omega/\omega_{\rm res}) - (\omega/\omega_{\rm res})^2},$$
(2.26)

showing that the driving force is suppressed in the pendulum by a factor of $(\omega/\omega_{\rm res})^{-2}$ for frequencies $\omega \gg \omega_{\rm res}$. For a human scale pendulum, this turn-over frequency is approximately

$$f_{\rm turn} \approx \frac{1}{2\pi} \sqrt{\frac{9.8 {\rm m/s}^2}{1 {\rm m}}}$$
(2.27)

$$\approx$$
 1Hz, (2.28)

which is well below the sensitive band of the LIGO detectors. This calculation suggests then that passive isolation from a pendulum suspension would be effective is suppressing seismic noise with LIGO's sensitive band.

In the advanced LIGO detector design, the test mirrors are suspended at the end stage of a system of four successive pendulums. It turns out that a system composed of serially coupled *linear* sub-systems has a transfer function which is equal to the product of the transfer functions of the individual sub-systems. So the transfer function for a quadruple pendulum in the linear approximation is simply $T_{\text{quad}}(f) = T_1(f)^4$, and motion of the mirrors is suppressed by $(f/f_{\text{res}})^{-8}$ at frequencies much greater than the resonant frequency $f_{\text{res}} \sim 1$ Hz. These considerations here assume that the subcomponents of a LIGO instrument couple to each other linearly, which is generally not the case. We therefore need additional techniques to keep the driving force in the linear regime. In addition to the passive damping provided by the pendulum suspension, we also use active sensing of the mirror positions as feedback for pushing back on the mirrors to keep them still.

2.4.2 Thermal noise

At an atomic level, all parts of the interferometer are constantly moving due to the non-zero temperature of its parts. These thermal fluctuations couple to fluctuations in the optical path length traversed by any given photon. We can understand the effects of thermal noise in an approach similar to that taken to understand seismic noise. Here we consider a collection of many damped harmonic oscillators connected to a thermal bath. For example, we may consider the test masses as in thermal contact with themselves. Each of the atoms in the test mass moves about its equilibrium position in the mirror due to random thermal fluctuations, changing the physical and optical properties of the mirror.

The equilibrium behavior of the system is well-described in terms of the Hamiltonian of the harmonic oscillator

$$H(x,\dot{x}) = \frac{1}{2}m\dot{x}^2 + \frac{1}{2}kx^2,$$
(2.29)

which determines its partition function Z, given by

$$Z = \int_{\mathbb{R}^2} e^{-H/k_B T} dx d\dot{x}$$
 (2.30)

$$= \frac{1}{2} \frac{\pi k_B T}{\sqrt{mk}},\tag{2.31}$$

where k_B is the Boltzmann constant and k is an effective spring constant. The equilibrium distribution of the collection of oscillators is determined by the Hamiltonian and partition function as

$$p(x,\dot{x}) = \frac{e^{-H(x,\dot{x})/k_B T}}{Z}.$$
(2.32)

In this harmonic oscillator model, the mean fluctuations about the equilibrium point is $\langle x \rangle = 0$. The mean square fluctuation about the equilibrium point is given in terms of the partition function as

$$\langle x^2 \rangle = -2k_B T \frac{\partial \log Z}{\partial k}$$
(2.33)

$$= \frac{k_B T}{k}.$$
 (2.34)

These results give the total fluctuations in displacement integrated over all fluctuation timescales. Written in terms of energy fluctuations, we can view this result as an instance of the general equipartition theorem for quadratic potentials since

$$\langle E_{\text{thermal}} \rangle = \frac{1}{2} k \langle x^2 \rangle = \frac{1}{2} k_B T.$$
 (2.35)

We now wish to decompose these thermal fluctuations into their frequency components. To do

so, we return to the equation of motion

$$F_{\rm drive} = m\ddot{x}(t) + b\dot{x} + kx, \qquad (2.36)$$

which is just a generalized version of Eqn. 2.23. As before, we determine the fluctuation spectrum

$$\tilde{x}(f) = \frac{\tilde{F}_{\text{drive}}(f)}{-m\omega^2 - ib\omega + k}.$$
(2.37)

What remains is to understand the form of the driving force. If we assume that the thermal force spectrum is independent of frequency $\tilde{F}_{drive}(f) = F_0$, then we immediately obtain the power spectrum

$$S_x(f) = \langle |\tilde{x}(f)|^2 \rangle = \frac{F_0^2}{(k - m\omega^2)^2 + (b\omega)^2}.$$
(2.38)

One can show through integration of Eqn. 2.38 and comparison with Eqn. 2.34 that the driving force F_0 is related to the friction constant b by $F_0^2 = k_B T b$. This relationship is an example of the fluctuation dissipation theorem, and shows that to minimize the displacement fluctuations due to thermal effects one must minimize the dissipation in that system.

We obtain finally the spectrum for thermal noise:

$$S_x(f) = \frac{k_B T b/m^2}{(k/m - \omega^2)^2 + (b\omega/m)^2}.$$
(2.39)

At frequencies $\omega \gg \sqrt{k/m}$ the noise *amplitude* is suppressed by f^{-2} (power by f^{-4}). For an inertiadominated system, specifically a system with $b/m < \omega_{\rm res}$, the power spectrum achieves a maximum near the resonance of the oscillator. The height (in amplitude) of this resonance scales as b^{-1} . So while a low friction system will exhibit less thermal noise at frequencies well above its resonance, it will compensate by concentrating the power near the resonance of the system. If on the other hand the friction dominates the system, $b/m > \omega_{\rm res}$, the power spectrum achieves a maximum at f = 0and the low frequency noise amplitude falls as f^{-1} before turning flat.

Thus, the test mass mirrors, which are engineered to have very low dissipation (small b) and very high fundamental modes ($\sqrt{k/m} \sim 1$ kHz), have thermal spectra that fall with f^{-1} in the LIGO band. In the case of thermal noise in the test mass itself, the oscillating entities are its constituent atoms and the restoring force is provided by electrodynamic interactions between the atoms – small m and large k. On the other hand, thermal noise in the suspension that holds the mirror is inertiadominated. The oscillating entity is a 40kg of fused silica (the test mass) and the restoring force is provided by tension in the suspension – large m and large k. In this case, the resonant frequency is $f \approx 8$ Hz and the noise amplitude falls as f^{-2} but with a large peak near the resonant frequency, the so-called "bounce" mode of the suspension, which ultimately limits our the detector sensitivity at low frequencies. See Fig. 2.5.

2.4.3 Quantum noise

High frequency noise is dominated by shot noise in the readout, which we have already discussed in Sec. 2.2. Shot noise has a flat power spectrum, since the arrival times of individual photons are uncorrelated. Yet in Fig. 2.5, we see that the shot noise *increases* as a function of frequency. This apparent coloring of shot noise actually arises from the frequency response of the detector to gravitational waves.

Consider that the interferometer is being driven by a gravitational wave field

$$h(\vec{x},t) = h e^{i(k\vec{x}-\omega t)}.$$
(2.40)

A detailed calculation [53] shows that in the long (gravitational) wavelength limit the transfer function from displacement (or strain) to optical phase shift of the reflected light is proportional to

$$C(f) = \frac{1}{1 + if/f_0},\tag{2.41}$$

where $f_0 = 91$ Hz is called the *cavity pole*. This result assumes the same configuration described in Sec. 2.3, in which the end mirrors are perfectly reflecting and the initial mirrors have equal reflectivity r. We see from Eqn. 2.41 that the response of the detector falls off as f^{-1} . Therefore, the shot noise goes up as the gravitational wave frequency increases because the high frequency begins to offset against the long light storage time, reducing the accumulated phase shift of the reflected light relative to the no-signal case.

We mentioned briefly above radiation pressure noise, which is also due to the quantum nature of photons. Radiation pressure is a white noise process and acts as a driving force on the test mass, which is suspended from a pendulum. The suspension damps the amplitude of the mirror motion by f^{-2} . Thus, the importance of radiation pressure is greatest for low frequencies. For a high power aLIGO configuration, the radiation pressure can rival that of the seismic and thermal contributions



Figure 2.6: Recently achieved LIGO and Virgo detector sensitivities. We show typical detector strain sensitivities for the LIGO and Virgo detectors during LIGO's fifth (left) and sixth (right) and Virgo's first (left), second, and third (right) science runs. In their fifth science run, the LIGO observatories at Hanford and Livingston achieved design sensitivity above 60Hz [1]. LIGO's sixth science run introduced a higher power laser [55, 56], which led to an improvement of nearly a factor of two in strain sensitivity [2], which exceeded the initial design sensitivity.

to the noise (see Fig. 2.5).

2.5 Actual LIGO and Virgo sensitivities

We have seen that there are a vast number of noise sources that must be kept down in order to detect gravitational waves. Understanding the noise sources is difficult, and what works on pen and paper doesn't always work in the laboratory.

Initial LIGO consisted of five science runs. By the end of the fifth science run, the LIGO instruments had already reached the design sensitivity for gravitational waves with frequencies greater than 60 Hz [1], as shown in Fig. 2.6. In an additional science run, which went beyond the scope of initial LIGO, the successful implementation of a high power laser was demonstrated and led to an improvement in strain sensitivity of up to a factor of two [2]. After the sixth science run, further improvements in sensitivity in the mid-frequency range were demonstrated by the use of quantum squeezing of the vacuum state [54]. The squeezing technique is not a part of the aLIGO design but may be applied later in the instrument construction as an add-on to the core design. In Chap. 3, we will examine the implications of these sensitivities for the astrophysical reach of compact binary coalescence searches.

2.6 Antenna response of LIGO detectors

Now we turn our attention to how the interferometer design we have been studying responds to gravitational waves from a specific source location and orientation with respect to the detector. Recall from Chap. 1 that we have used the *source frame* and the *radiation frame* to describe how gravitational waves go from source to waves. We now wish to understand how these gravitational waves go from waves to strain in the detector. For this purpose, it will be convenient to introduce one more coordinate system, which we call the *detector frame*.

In the detector frame, the x and y coordinate axes point along the detector arms such that the positive z axis points towards the local zenith. This coordinate system motivates the introduction of the *detector tensor*, defined as

$$\mathbf{D} = \hat{x} \otimes \hat{y} - \hat{y} \otimes \hat{x}, \tag{2.42}$$

where \hat{x} and \hat{y} are the unit vectors pointing along the detector frame x and y coordinate axes, respectively. In the detector frame, the coordinate expression for the detector tensor is simply

$$D = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$
 (2.43)

The detector is only sensitive to differential strains along the x and y axes. Thus, the gravitational wave signal in the detector is given by the component of the propagating gravitational wave along the detector tensor, by which we mean the contraction

$$h = \mathbf{D} :: \mathbf{h} \equiv D_{ij} h_{ij}, \tag{2.44}$$

where \mathbf{h} is the spatial part of the propagating gravitational wave (see Eqn. 1.19).

Given this result, we can now easily derive the detector response as a function of the direction of the gravitational wave. We use spherical polar coordinates θ and ϕ to describe the sky position; these angles are given with respect to the detector frame axes. Additionally, we shall need the polarization angle ψ , which describes the orientation of the binary with respect to the line of sight at a fixed inclination angle. In order to compute the contraction, we transform the strain tensor



Figure 2.7: *LIGO as a gravitational wave antenna.* to gravitational waves from the whole sky, but the sensitivity is not isotropic. Above we show the sensitivity of such a detector to the plus (left) and cross (center) polarizations of a gravitational wave as a function of the source position on the sky relative to the detector. The thick black lines indicate the arms of the detector, and the distance of the surface from the origin indicates the relative amplitude response of the detector to a source from that direction in the sky. To the right, we show the two response patterns added in quadrature, which gives a measure of the detector response to randomly polarized gravitational waves. Gravitational waves from compact binary coalescences are circularly polarized, and the average response of the detector to such sources averaged over several cycles is more accurately depicted by the right-most figure. Figure credit [11].

from the radiation frame, in which its coordinate expression is given by

$$h = \begin{pmatrix} h_{+} & h_{\times} & 0\\ h_{\times} & -h_{+} & 0\\ 0 & 0 & 0 \end{pmatrix}, \qquad (2.45)$$

to the detector frame coordinates. To do so, we perform three rotations: (1) a rotation by the angle ψ about the radiation frame z-axis, (2) a rotation by the angle θ about the resulting x-axis, and (3) a rotation by the angle ϕ about the resulting z-axis.

Thus, in the detector frame, we obtain

$$h = R_z(\phi)R_x(\theta)R_z(\psi)h(R_z(\phi)R_x(\theta)R_z(\psi))^{-1}$$
(2.46)

where

$$R_z(\alpha) = \begin{pmatrix} \cos(\alpha) & \sin(\alpha) & 0\\ -\sin(\alpha) & \cos(\alpha) & 0\\ 0 & 0 & 1 \end{pmatrix},$$
(2.47)



Figure 2.8: The importance of multiple detectors for full sky coverage. On the left, we see that a detector network consisting of just the Hanford, Livingston, and Virgo detectors has a significant region of depressed sensitivity. This region is essentially given by the plane of the three detectors, where the individual detector antenna patterns also have nodes. Simply adding a fourth detector, KAGRA in Japan, to the network significantly reduces these deficits and gives more uniform sky coverage, as seen to the right. Figure credit [57].

and

$$R_x(\alpha) = \begin{pmatrix} 1 & 0 & 0\\ 0 & \cos(\alpha) & \sin(\alpha)\\ 0 & -\sin(\alpha) & \cos(\alpha) \end{pmatrix}.$$
 (2.48)

Through straightforward computation, one obtains the following explicit expression for the strain h observed in the detector,

$$h = F_{+}(\theta, \phi, \psi)h_{+} + F_{\times}(\theta, \phi, \psi)h_{\times}, \qquad (2.49)$$

where

$$F_{+} = \frac{1}{2} (1 + \cos(\theta)^{2}) \cos(2\phi) \cos(2\psi) - \cos(\theta) \sin(2\phi) \sin(2\psi)$$
(2.50)

$$F_{\times} = \frac{1}{2} (1 + \cos(\theta)^2) \cos(2\phi) \sin(2\psi) + \cos(\theta) \sin(2\phi) \cos(2\psi)$$
(2.51)

These functions are known as the antenna response of the gravitational wave detector and are illustrated in Fig. 2.7.

A LIGO instrument is sensitive to GWs from all parts of the sky, apart from the four nodes in the plane of the detector. However, the response of a LIGO instrument to GWs is not uniform over the sky. As can be seen in Fig. 2.7, a LIGO instrument is most sensitive to overhead GW sources. For detectors sensitive to GW sources beyond the local group, as are the LIGO and Virgo detectors as of their most recent science runs, signals are expected to be uniformly distributed over



Figure 2.9: A developing global network of gravitational wave detectors. Here we show the five currently planned gravitational wave observatories: LIGO Hanford (US), LIGO Livingston (US), GEO (Germany), Virgo (Italy), KAGRA (Japan), and LIGO India (India). The exact site of the Indian detector has not yet been determined.

the sky. Thus, for optimal detection of gravitational waves, it is crucial to have multiple detectors with differing orientations.

There is currently a developing global network of gravitational wave detectors, as illustrated in Fig. 2.9. Shown in the figure are the two four-kilometer LIGO detectors at Hanford and Livingston in the United States, the three-kilometer Virgo detector in Cascina, Italy, and the 600 meter detector in Hanover, Germany. In addition, there are on-going negotiations to have a third LIGO detector placed in India, at a yet to be determined site. The Japanese detector KAGRA is currently under construction and will be the first gravitational wave detector built underground (to reduce seismic noise) and featuring cryogenic mirrors (to attack thermal noise).

In the case of multiple detectors, the effective antenna response is given by

$$F(\theta, \phi, \psi)^2 = \sum_{\text{detectors } k} F_{+,k}^2 + F_{\times,k}^2.$$
(2.52)

In Fig. 2.8, we show the sky coverage obtained through the use of multiple detectors. Here we are comparing the sky coverage for a network consisting of only LIGO and Virgo to one that includes LIGO, Virgo and the KAGRA detector. Not shown here is the effect of the LIGO India detector, whose site has not yet been chosen. We see that the addition of the KAGRA detector significantly improves deficits in the network sensitivity. Multiple detectors also provide superior coverage in *time* since an individual detector is not in observational mode. One can also use the analysis from independent detectors to reject spurious single-detector events, by requiring that a candidate signal

appear in more than one detector simultaneously.

Another important point regarding multiple detectors is the ability to localize the GW source on the sky. Compact binary coalescence might be followed by a γ -ray burst, and the ability to locate CBC sources on the sky will be crucial for proving this hypothesis and enabling the multi-messenger study of the progenitor. A single detector is sufficiently omnidirectional such that a measurement of a transient signal in one detector gives nearly no information about the location of the source. With two detectors, we can constrain the source location to a ring on the sky based on the time delay between the appearance of the signal in each detector. With three detectors, the source can be localized to two antipodal points on the sky. A fourth detector breaks this final degeneracy, uniquely identifying the location of the source. Several existing low-latency pipelines have rapid sky-localization methods based on the relative arrival times and amplitudes of the signal in different detectors [58, 59]. More generally, the coherent analysis of multiple detector data provides a robust method for determining the source parameters [60].

Chapter 3

Detecting Compact Binary Coalescence Signals

The first principle is that you must not fool yourself, and you are the easiest person to fool.

> Richard P. Feynman in his Caltech commencement address (1974)

This chapter introduces the basic theory and techniques behind the analysis of LIGO instrumental data for gravitational wave signals from compact binary coalescence. We define the matched filter signal-to-noise ratio (SNR), which is the optimal statistic for the detection of a signal of known form added to stationary Gaussian noise. In practice, the exact form of the CBC signal depends on the unknown source parameters, including the masses, spins, orientation, and sky position, and we require methods to maximize the SNR over these parameters. Most of the extrinsic waveform parameters (those that depend on the observer) are easy to handle, entering only into the overall amplitude of the signal. We can maximize the SNR over the coalescence time using a computationally efficient inverse Fourier transform and maximize over the coalescence phase analytically. The remaining intrinsic source parameters, however, require the construction of a *template bank*.

We introduce here two techniques for covering the intrinsic signal parameter space, one based on placing templates on a lattice and another based on stochastic template placement. Additionally, LIGO detector data are in general neither stationary nor Gaussian, violating the assumptions that ensure the optimality of the SNR detection statistic. In actual searches, we employ additional background suppression techniques beyond the simple matched filter. We briefly discuss some of these techniques, which include vetoes based on data quality and signal consistency tests.
After a generic treatment of search techniques for CBCs in gravitational wave data, we describe one specific implementation: the IHOPE pipeline. In the IHOPE pipeline, each of the detector data is searched separately for candidate gravitational wave events with a bank of waveforms from model systems. To suppress the background, a candidate event in one detector must have a corresponding candidate event in another detector with similar measured masses and coalescence time to merit further consideration. We estimate the background rate of coincidence between instrumental noise triggers by sliding the detector data relative to each other in time such that any coincident events would necessarily be of non-astrophysical origin. In the next chapter, we will discuss the results of the analysis of recent LIGO and Virgo data with the IHOPE pipeline using lattice-based template banks with non-spinning waveform templates.

3.1 Detecting signals of known form in stationary Gaussian noise

In this section, we the introduce the vocabulary and tools needed for our later analysis of gravitational wave data for signals from coalescing compact binaries. We define a random process $X(\alpha)$ as a parameterized set of random variables. A random process is fully characterized by specifying the joint probability density $p_{\vec{\alpha}}(\vec{x})$ for every finite sequence of points $\vec{\alpha}$. For our applications, the parameter α will represent either time or frequency. When α denotes time, we assume that the random process is real-valued, as will be the case for any time series that we encounter here. In the frequency domain, the random process will in general be complex-valued, typically arising as the Fourier transform of some real-valued random process in the time domain.

Our baseline framework for gravitational wave data analysis begins with the assumption that the noise is stationary. A random process is said to be *stationary* if the two-point joint probability distribution $p_{\vec{\alpha}}(x_1, x_2)$ for $\vec{\alpha} = (\alpha_1, \alpha_2)$ depends only on the difference $\alpha_2 - \alpha_1$. We consider a noisy process n(t) in the time domain, such as the readout from the detector strain channel. If n(t) is stationary, then the application of the Fourier transform to n(t) removes the correlations in the joint probability density, and we have

$$\mathbb{E}[\tilde{n}(f)\tilde{n}^{*}(f')] = \frac{1}{2}S_{n}(f)\delta(f - f'), \qquad (3.1)$$

where $S_n(f)$ is the power spectral density, also defined in Eqn. 2.21. As before, we include the

factor of 1/2 in the definition of the power spectral density because we have assumed that n(t) is real-valued, from which it follows that $\tilde{n}(-f) = \tilde{n}^*(f)$. This relation implies that there is equal power at negative and positive frequencies, so we fold the negative frequencies onto the positive ones and consider only positive frequencies.

The Fourier transform leaves variances in the transformed process $\tilde{n}(f)$. To remove these variances, we divide by the square root of the power spectral density to obtain

$$\tilde{n}_W(f) = \frac{\tilde{n}(f)}{\sqrt{\frac{1}{2}S_n(f)}}$$
(3.2)

which now has unit power spectral density. When $S_n(f) = 1$, we say that the noise is *white*, and the transformation of a random process n(t) by Eqn. 3.2 is referred to as *whitening*. Although we have defined the transformation here on the random process n(t), we will also apply this transformation to signal and speak of the *whitened signal*.

A remarkable simplification of the signal detection problem arises under the additional assumption that n(t) is Gaussian (see Refs. [61–63] for the details). A random process $X(\alpha)$ is called *Gaussian* if $p_{\vec{\alpha}}(\vec{x})$ is jointly Gaussian distributed for every finite sequence $\vec{\alpha}$. Let s(t) denote the data collected from a channel in the gravitational wave detector, and let h(t) denote a putative signal in the data. We define an inner product between these two times series by

$$\langle s|h\rangle \equiv 2 \int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*(f) + \tilde{s}^*(f)\tilde{h}(f)}{S_n(f)}.$$
(3.3)

It is convenient to normalize the signal waveform by this inner product $\hat{h} = h/\sqrt{\langle h|h\rangle}$. Then, under the assumption that n(t) is Gaussian, the quantity

$$\rho \equiv \langle \vec{s}, \hat{h} \rangle, \tag{3.4}$$

which we call the *signal-to-noise ratio*, is the Neyman-Pearson optimal detection statistic for the signal h(t). That is, among all linear filters on the data s(t), the signal-to-noise ratio ρ gives the maximum detection probability for any false alarm probability. We refer to the template h(t) as a matched filter.

3.2 Covering the parameter space

=

The methods described in the previous section only apply if the form of the signal is known *a priori*. In fact, there are *at least* fifteen parameters required to specify the expected signal from a compact binary coalescence. We enumerate these parameters in Tbl. 3.1. Since the source parameters are not known, we must consider all possible source parameters and maximize the SNR over these parameters.

Many of the extrinsic parameters enter only into the signal amplitude at the detector: the sky position angles, binary plane orientation angles, and the luminosity distance to the source. Since the SNR in Eqn. 3.4 is defined for normalized templates, we do not need to search over these parameters – provided that they are constant through the binary's evolution. However, when the component spins are misaligned with the orbital angular momentum, the inclination and polarization angles are time dependent and lead to a non-trivial modification of the observed signal. In the misaligned case, we must therefore also search over these angles.

Table 3.1: The compact binary parameter space. There are at least fifteen parameters required to specify the orbit of a compact binary (we have ignored parameters associated with eccentricity and the finite size of neutron stars). We refer to the parameters (1–8) as *intrinsic* parameters, while (9–15) are called *extrinsic*. Parameters (9–13) enter only in the overall amplitude of the signal, (14) can be maximized over analytically, and (15) can be efficiently searched over with an inverse Fourier transform.

1 - 2	component masses m_1 and m_2	
3-8	component spin vectors \vec{S}_1 and \vec{S}_2 , each having three components	
9 - 10	sky position: right ascension α and declination δ	
11–12	orientation of the binary relative to the line of sight: inclination ι and polarization angle ψ	
13	luminosity distance D	
14	coalescence phase φ_{coal}	
15	coalescence time $t_{\rm coal}$	

If we restrict our attention to aligned-spin systems, then the inclination and polarization angles are constant and enter only into the overall amplitude of the signal, making them degenerate¹ with the sky position and luminosity distance. After these, we are left with only the masses and spins of the binary, which for non-precessing systems amounts to at most four parameters. Efficient techniques for covering this remaining parameter space have been extensively studied [64–66]. We describe two such techniques below, each of which will be applied in the analysis presented in

 $^{^{1}}$ These parameters are not degenerate for a multi-detector coherent analysis. Here we are considering an incoherent analysis in which each detector data are searched separately. For a single detector, the parameters that affect the overall amplitude can't be disentangled.

subsequent chapters. We will in fact restrict our attention to covering a subset of the aligned-spin parameter space; later we will assess the efficacy of aligned-spin signal models for capturing signals from generically spinning systems.

3.2.1 Searching over extrinsic parameters

As we have just mentioned, parameters of the binary which are extrinsic (depend on the observer) tend to enter into the signal in very simple ways, making the search over these parameters relatively easy. Most of them enter only into the amplitude of the signal and don't have to be searched over at all. The coalescence time and phase, however, are not so simple. Here we review the methods we use to maximize the SNR over these two parameters.

We efficiently maximize the SNR over the coalescence time t_{coal} through the use of an inverse Fourier transform. Recall that the time translation of a signal by offset τ in the time domain becomes a phase drift in the frequency domain,

$$\widetilde{h(t-\tau)}(f) = \tilde{h}(f)e^{2\pi i f\tau}.$$
(3.5)

It follows from Eqn. 3.3 that

$$\rho(\tau) = \langle s(t), h(t-\tau) \rangle \tag{3.6}$$

$$= 4 \operatorname{Re} \int_0^\infty \frac{\tilde{s}^*(f)h(f)}{S_n(f)} e^{2\pi i f \tau} df.$$
(3.7)

Note that the numerator of Eqn. 3.3 is purely real, since the two terms are complex conjugates of each other. Eqn. 3.7 gives a computationally efficient method to maximize the SNR over time. Rather than having to compute the overlap integral in Eqn. 3.3 repeatedly for each time translation of the template, we can compute the SNR for all such templates with a single fast (inverse) Fourier transform, since $\rho(\tau)$ in Eqn. 3.7 is just the inverse Fourier transform of $\tilde{s}^*(f)\tilde{h}(f)/S_n(f)$. We compute this inverse Fourier transform once and search for peaks in $\rho(\tau)$ over τ to obtain candidate events.

The coalescence phase φ_{coal} can be maximized over analytically whenever the stationary phase approximation holds, that is, whenever the GW signal has only a single frequency component at any time. In this approximation, the plus and cross polarizations of the CBC signal are proportional to each other and out of phase by $\pi/2$ radians:

$$\tilde{h}_{+}(f) = \frac{1 + \cos^2 \iota}{2} e^{2i\varphi_{\text{coal}}} \hat{h}_{0}(f)$$
(3.8)

$$\tilde{h}_{\times}(f) = \cos \iota e^{2i(\varphi_{\text{coal}} + \pi/4)} \hat{h}_0(f), \qquad (3.9)$$

where $\hat{h}_0(f)$ is the signal for $\iota = \varphi_{\text{coal}} = 0$ placed at a distance d such that $\langle \hat{h}_0, \hat{h}_0 \rangle = 1$ (see Eqn. 1.44). The signal observed in the detector is given by Eqn. 2.49, and we find with some algebra that

$$\tilde{h}(f) = \left(\frac{D}{D_{\text{eff}}}\right) e^{2i\phi_0} e^{2i\varphi_{\text{coal}}} \hat{h}_0(f).$$
(3.10)

Here, ϕ_0 is the *termination phase*, and is given by

$$\tan 2\phi_0 = \frac{2F_{\times} \cos \iota}{F_+(1 + \cos^2 \iota)},$$
(3.11)

and

$$D_{\rm eff} = D \left(F_+^2 \left(\frac{1 + \cos^2 \iota}{2} \right)^2 + F_\times^2 \cos^2 \iota \right)^{-1/2}$$
(3.12)

is the *effective distance* – the distance at which an optimally located and oriented binary would give the same signal amplitude. Using the waveform $\tilde{h}(f)$ as a template, we obtain the SNR as a function of the coalescence phase

$$\rho(\varphi_{\text{coal}}) = \langle s, h \rangle = \langle s, \hat{h}_0 e^{2i(\varphi_{\text{coal}} + \phi_0)} \rangle.$$
(3.13)

Recall that the SNR defined in Eqn. 3.4 is defined for *normalized* templates, which is why the distances do not appear in the above equation. Defining the *complex* SNR

$$z = 4 \int_0^\infty \frac{\tilde{s}^*(f)\hat{h}_0(f)}{S_n(f)} df,$$
(3.14)

we can also write the SNR as

$$\rho(\varphi_{\text{coal}}) = \cos 2(\varphi_{\text{coal}} + \phi_0) \operatorname{Re} z + \sin 2(\varphi_{\text{coal}} + \phi_0) \operatorname{Im} z.$$
(3.15)

Maximizing $\rho(\varphi_{\text{coal}})$ in Eqn. 3.15 with respect to φ_{coal} , we find

$$\tan 2(\varphi_{\text{coal}} + \phi_0) = \frac{\text{Im } z}{\text{Re } z}.$$
(3.16)

From this condition, one can compute $\cos 2(\varphi_{\text{coal}} + \phi_0) = \text{Re } z/|z|$ and $\sin 2(\varphi_{\text{coal}} + \phi_0) = \text{Im } z/|z|$. Plugging these back into Eqn. 3.15, we find that the maximal SNR is given by

$$\rho(\varphi_{\text{coal}})|_{\text{max}} = \sqrt{(\text{Re } z)^2 + (\text{Im } z)^2}.$$
(3.17)

Thus, to maximize over the coalescence phase, we simply compute the magnitude of the complex SNR z. This maximization comes essentially at no computational cost since $\rho = \text{Re } z$, and we are already computing ρ .

Putting the phase and time maximizations together, we obtain the result

$$z(t) = 4 \int_0^\infty \frac{\tilde{s}^*(f)\hat{h}(f)}{S_n(f)} e^{2\pi i f t} df$$
(3.18)

and the SNR maximized over time and phase is $z = \max_{t} |z(t)|$.

Although the sky position of the binary remains approximately fixed during its evolution in band, the orientation (given by the inclination angle ι and polarization angle ψ) of the binary can evolve in a non-trivial way if the component spins are misaligned with the orbital momentum. In the special case in which the binary spins are aligned to the orbital momentum, the orientation angles of the binary are constant throughout the evolution of the system. We will consider in this work only aligned spin filters, and therefore the discussion above completely solves the problem of searching over the unknown extrinsic parameters. Of course, astrophysical CBCs may very well have significant precession. We emphasize that we are only neglecting precession effects *in the filters*. We will later address the question of whether these aligned spin templates are sufficient for the detection of generically spinning binaries by using the aligned spin filters to search for simulated precessing signals added to detector data.

3.2.2 Covering masses and spins

Finally, we have the problem of searching over the intrinsic binary parameters, which are the masses and spins. These parameters enter in a much more complicated way into the phase and amplitude of the GW signal compared to the extrinsic parameters. We cover this portion of the parameter space using a *template bank*, which is a finite collection B of potential signals. The primary goal for the construction of a template bank is to "cover" a target signal space S. One useful measure for the coverage of a given signal space by a template bank is given in terms of the mean SNR recovery by template bank compared to the optimal case $S \subset B$ (i.e. the case in which the bank contains the actual signal). We formalize this notion below. We also desire the template bank to use as few templates as possible in covering the space so as to minimize the computational cost of the search.

We write the parameters $\vec{\lambda}$ for an arbitrary CBC signal as $\vec{\lambda} \equiv {\{\vec{\lambda}_{intr}, \vec{\lambda}_{extr}\}}$, where $\vec{\lambda}_{intr}$ are the intrinsic parameters and $\vec{\lambda}_{extr}$ are the extrinsic parameters (see Tbl. 3.1). We have just shown how to optimize a search with respect to the extrinsic parameters. Given two signals \hat{s} and \hat{t} with intrinsic parameters $\vec{\lambda}_s$ and $\vec{\lambda}_t$, respectively, we define the *match* \mathcal{M} between them by

$$\mathcal{M} \equiv \max_{\vec{\lambda}_{\text{extr}}} \langle \hat{s} | \hat{h} \rangle.$$
(3.19)

We denote the match by $\mathcal{M}(\vec{\lambda}, \Delta \vec{\lambda})$, where $\Delta \vec{\lambda} = \vec{\lambda}_t - \vec{\lambda}_s$. We obtain a convenient approximate expression for the match between neighboring templates by Taylor-expanding the match about $\Delta \vec{\lambda} = 0$. Since the match function has its maximum value of unity at $\Delta \vec{\lambda} = 0$, there are no linear terms in the expansion, and truncating the expansion at second order, we get

$$\mathcal{M}(\vec{\lambda}, \Delta \vec{\lambda}) \simeq 1 - g_{ij} \,\Delta \lambda^i \Delta \lambda^j \tag{3.20}$$

where

$$g_{ij} \equiv -\frac{1}{2} \left(\frac{\partial^2 \mathcal{M}}{\partial \Delta \lambda^i \, \partial \Delta \lambda^j} \right)_{\Delta \vec{\lambda} = 0} \tag{3.21}$$

can be interpreted as a metric on a manifold in the coordinates of the binary parameter space. The mismatch $1 - \mathcal{M}$ between two neighboring templates has the interpretation of the proper distance in the parameter space [67] and for short distances is given approximately by

$$1 - \mathcal{M} \simeq g_{ij} \,\Delta\lambda^i \Delta\lambda^j. \tag{3.22}$$

The metric coefficients are coordinate dependent. We will see that at the second order post-Newtonian limit, there are certain coordinates in which the metric takes an exceptionally simple form. The use of an approximate semi-analytic expression for the metric, when available, greatly reduces the computational cost of the match calculation between two templates, since it obviates the need for computing the inverse Fourier transform to the overlap over time translations.

The *fitting factor* of a template bank towards a particular signal quantifies the effectiveness of a template bank for SNR recovery [68]. The fitting factor of a template bank B towards a particular

signal $s \in S$ with intrinsic parameters $\vec{\lambda}_s$ is defined by

$$FF(\vec{\lambda}_s; B) = \max_{\hat{h} \in B} \mathcal{M}(\vec{\lambda}_h; \vec{\lambda}_h - \vec{\lambda}_s)$$
(3.23)

where the maximization is over all templates in the bank. When the metric g_{ij} is known, we can approximate the mismatch with Eqn. 3.22 and therefore approximate the fitting factor by

$$FF(\vec{\lambda}; B) \approx 1 - \min_{\hat{h} \in B} |\vec{\lambda}_h - \vec{\lambda}_s|^2, \qquad (3.24)$$

where the norm || is defined with respect to the metric, i.e., $|\vec{\lambda}| = g_{ij}\lambda^i\lambda^j$. The fitting factor quantifies the optimality of the template bank for detecting the signal \vec{s} . Relative to the optimal case in which $\hat{s} \in B$, the use of the sub-optimal template bank for filtering incurs a loss

$$\frac{V_{\text{bank}}}{V_{\text{opt}}} = FF^3 \tag{3.25}$$

in detection rate. Here V denotes the observable volume at a fixed SNR threshold, and neglects the dependence of the background event rate on the particular template bank used for the search.

3.2.3 Lattice template placement

Currently implemented LIGO-Virgo matched filtering searches for CBCs are based on a lattice approach for placing the templates [3, 4], and similar techniques are being developed for advanced generation searches [69–71]. This technique relies on the existence of coordinates $\vec{\lambda}' = f(\vec{\lambda})$ in which the metric coefficients g_{ij} are constant (or nearly so) across the parameter space. One can then place the templates on a regular lattice in these coordinates to guarantee a minimal loss of SNR [66, 72, 73]. The lattice technique is highly computationally efficient, but requires one to determine the appropriate coordinates in which to lay down the templates. Depending on the waveform approximation used, such coordinates may not exist, rendering the lattice approach inapplicable.

In the second order post-Newtonian approximation to the CBC signal and when the spins of the system are negligible, these coordinates are found to be the (dimensionless) *chirp times* [72, 73],



Figure 3.1: Lattice-based placement strategy for non-spinning templates. We show the parameters of a template bank constructed by placing templates in a hexagonal-lattice in the chirptime coordinates τ_0 and τ_3 . Setting up a lattice in this coordinate system covers the physical parameter space without putting too many templates in any one place. For example, the lattice construction tells us that the higher mass systems require fewer templates to detect.

defined by

$$\theta_0 = \frac{5}{128\eta} (\pi M f_{\text{low}})^{-5/3} \tag{3.26}$$

$$\theta_3 = -\frac{\pi}{4\eta} (\pi M f_{\text{low}})^{-2/3}, \qquad (3.27)$$

where M is the binary total mass and $\eta = m_1 m_2/M^2$ is the symmetric mass ratio. The dimensionless chirp time parameters are related to the usual chirp time parameters by $\theta_0 = 2\pi f_{\text{low}}\tau_0$ and $\theta_3 = -2\pi f_{\text{low}}\tau_3$. When merger, ringdown, spin, or higher order post-Newtonian corrections become important for detection, however, no such coordinates are known, and therefore the lattice techniques cannot be reliably applied to place templates.

In Fig. 3.1, we illustrate the coordinate-dependence of the metric (top vs. bottom) and the construction of a template bank in a lattice on the $\tau_0 - \tau_3$ coordinates (bottom). To the extent that



Figure 3.2: Breakdown of the chirp time coordinates for lattice template placement. We compare banks generated by the random placement method (left) and the lattice placement method (right) in the regime $M_{\text{total}} \in [15, 150] M_{\odot}$ where merger and ringdown contribute significantly to the SNR. We see that at the smallest chirp times (highest masses), the random placement technique uses fewer templates than the lattice placement. Both banks are generated with a nominal minimum match of $\mathcal{M} = 0.97$, but the lattice placement over covers at high mass, giving much higher fitting factors.

the post-Newtonian approximation is accurate, this placement technique covers the parameter space using the fewest possible templates for a given tolerance of mismatch, which here is $\mathcal{M}_{\min} = 0.97$. We see in the top figure that the template placement in the ordinary component mass parameters m_1 and m_2 is highly non-uniform, thus emphasizing the importance of knowing the correct parameters for the lattice construction.

3.2.4 Random template placement

As we add more physical effects to the template waveforms, such as merger or spin, the metric coefficients in the chirp time coordinates are no longer constant across the parameter space. In these cases, we may resort to a brute force approach for template placement, recently studied in [64, 65], in which template parameters are chosen at random until the parameter space is fully covered. Templates whose overlap with other templates in the bank are greater than the specified minimal match are discarded. The random placement method is robust if not computationally efficient, and by no means guarantees that the resulting bank is optimal in terms of the number of filters. However, the technique is completely generic, not depending on having any special coordinates with which to work, and straightforward to implement. In certain cases, such as when the filters are

sufficiently short such that the inverse Fourier transform is not prohibitively expensive, or when we have available an analytic approximation to the *non-constant* metric, the random placement approach is an appealing alternative to the lattice-based approach.

The random placement method begins with a seed bank of template parameters $B_0 = {\{\vec{\lambda}_{intr}^1, \vec{\lambda}_{intr}^2, \dots, \vec{\lambda}_{intr}^N\}}$, which may be empty. A set of template waveform parameters $\vec{\lambda}^{prop}$ is proposed randomly, and the bank B_0 is checked to see whether it already contains a template which sufficiently overlaps with the proposed template. We measure the coverage of the template bank via the fitting factor, which gives the fraction of optimal SNR that can be obtained towards the proposed template waveform $\vec{\lambda}^{prop}$ using the existing template bank without including the newly proposed template. If the fitting factor for the proposed template is above a given *minimum match* threshold \mathcal{M}_{min} , then the proposed template is discarded to prevent over-coverage, and we repeat the process with the same bank seed B_0 . Otherwise, the proposed template is added to the bank and we repeat the process using $B_1 \equiv B_0 \cup {\{\vec{\lambda}^{prop}\}}$ as the new bank seed. The process continues until some convergence criterion is satisfied. In our implementation, we terminate the bank construction when the mean number of discarded proposals per accepted proposal (averaged over the last ten accepted proposals) exceeds a specified critical value k_{max} .

We have implemented a generic infrastructure for generating stochastic banks, which we call SBANK [8,9]. See Sec. 6.2 for more details on this implementation. In Fig. 3.2, we compare the random template placement (as implemented in SBANK) and the lattice placement method for systems with $M_{\text{total}} \in [15, 150] M_{\odot}$. We will demonstrate the application of this infrastructure to the construction of template banks from signals that include effects from aligned spin, merger and ringdown. The great utility of SBANK as we have implemented it is that it is easily extensible to other waveform approximations as they become available.

3.3 Non-stationary and non-Gaussian noise

3.3.1 Data segmentation

Our considerations so far have assumed that our background is stationary and Gaussian. In practice, LIGO data are neither of these. Non-stationarities can arise for any number of reasons. For example, as discussed in Chap. 2, seismic noise is a significant source of noise in the low-frequency regime. Any changes in the seismic activity near the detector will register as an evolving background. For



Figure 3.3: Non-stationarity of real LIGO instrumental noise. We show the mean sensitive distance to a canonical $m_1 = m_2 = 1.4 M_{\odot}$ binary neutron star source for the four detectors which operated in S5/VSR1 as a function of time. Each point indicates the average sensitive distance over an interval of seven days and the error bar indicates the variance in the sensitivity during that week. Figure credit [1].

example, anthropogenic sources of seismic activity are clearly seen in the sensitivity of the detector. The detectors tend to be more sensitive during the evening hours, when there is little anthropogenic activity, and less sensitive in the day hours, when more people are on the roads. In the evening hours, the anthropogenic activity is less, improving not just the sensitivity at low frequencies where seismic activity couples most strongly to the detector, but also allowing for more stable detector locks at higher laser power.

On longer time scales, non-stationarities may arise from changes in the detector configuration, such as maintenance work or hardware upgrades. Dealing with slowly evolving non-stationarity is rather straightforward, provided that the signals are short in duration compared to the non-stationarity time scale. In that case, we can filter the data over intervals of time which are small compared to the non-stationary time scale but large compared to the signal time scale (the interval must be larger than the signal time scale to ensure minimal bias in the background estimation due to the presence of a signal). We treat the noise in each interval (for IHOPE this time interval is chosen to be 2048 s) as stationary, apply the usual methods, and take into account the non-stationary by re-estimating the PSD in every segment. In Fig. 3.3, we illustrate the non-stationarity of the LIGO detector noise over the course of an entire science run, in this case S5.

3.3.2 Coincidence between detectors

The events we are looking for will be quite rare, even in the most optimistic scenarios [31]. The background in a single detector can overwhelm the signal. Fortunately, with more than one detector we can severely cut down on the background rate by requiring *coincidence* between detectors for any candidate gravitational wave event. For each detector, we filter the data separately with the template bank and generate a list of *triggers*, that is, events in which some template in the bank generates



Figure 3.4: Non-Gaussianity of real LIGO instrumental noise. We show an example of $\rho - \chi^2$ statistics for triggers in an analysis of background only (black) and simulated signal added to detector noise (red). The χ^2 statistic shown here is the traditional χ^2 described in the text. The blue curves indicate the contours which are used to rank triggers. We see clearly from this figure that a simple threshold on SNR is insufficient for optimal detection in real LIGO data.

a SNR above some predetermined threshold. By saying that we require coincidence, we mean that we require that a trigger in one detector must have a corresponding trigger in another within some time and template parameter windows [74]. This requirement greatly reduces the number of false triggers associated with noise in the detector.

3.3.3 Signal consistency tests

To combat non-Gaussian behavior in the data, we perform additional statistical tests that go beyond the signal-to-noise. The most common of these types of tests are χ^2 tests, which measure the residual power in the data stream after subtracting off the inferred signal. If the residual is consistent with the Gaussian assumption, the χ^2 value should be low; otherwise the χ^2 should be large (see Fig. 3.4). As with the SNR, these χ^2 are sensitive to disagreements between the filter template and the actual signal. Large values of χ^2 can arise from either non-Gaussian fluctuations or mismatch between the true waveform and the closest matching template in the bank. There are several types of χ^2 statistics used in LIGO data analysis (see Chap. 6 in Ref. [75] for a brief review). Here we describe two particular tests, which will appear in pipelines later in this work. The first test (traditional χ^2) is a frequency-domain test which has featured in the most recent LIGO/Virgo searches for GWs from CBCs. The second test (autocorrelation χ^2) is a time-domain test, which is the currently implemented χ^2 statistic used in the new GSTLAL pipeline to be described in Chap. 5.

Traditional χ^2

Recent LIGO/Virgo searches for GWs from CBCs employed a χ^2 statistic [76] based on the consideration that although a detector glitch may generate triggers with the same SNR as a GW signal, the manner in which the SNR is accumulated over time and frequency is likely to be different. For example, a glitch that resembles a delta function corresponds to a burst of signal power concentrated in a small time-domain window, but smeared out across all frequencies. A CBC waveform, on the other hand, will accumulate SNR across the duration of the template, consistently with the chirp-like morphology of the waveform.

To test whether this is the case, the template is broken into p frequency bands with boundaries $f_0 = 0 < f_1 < f_2 < ... < f_{p-1} < f_p = \infty$ in such a way that the SNR accumulated in each frequency interval,

$$z_j = 4 \operatorname{Re} \int_{f_{j-1}}^{f_j} \frac{\tilde{s}(f)\hat{h}(f)}{S_n(f)} df,$$
 (3.28)

has the same expectation value $\langle z_j \rangle = z/p$, where $z = \sum z_j$ is total SNR achieved by the template. The χ^2 statistic is computed from these by

$$\chi^2 = \frac{1}{p} \sum_{j=1}^p (z - pz_j)^2, \qquad (3.29)$$

which is χ^2 -distributed with p degrees of freedom. The quantity pz_j is the expected total SNR based on the SNR in frequency band $[f_{j-1}, f_j]$. This technique is highly effective for suppressing the background when the signal duration is long, but computationally is rather costly, requiring pinverse Fourier transforms per template. Thus, in the IHOPE pipeline, which implements this test, the χ^2 is computed only for triggers that have passed the coincidence stage (see Sec. 3.4).

Autocorrelation χ^2

Another useful signal consistency test is obtained by comparing the SNR time series obtained from filtering the data to the autocorrelation of the template. This statistic has the main advantage that once the SNR time series is computed for a given template, all the required data to compute the statistic are already in memory; the statistic is extremely computationally efficient compared the traditional χ^2 . The autocorrelation χ^2 statistic is closely related to the bank χ^2 , as described in Ref. [75].

Suppose we collect data s = n + Ah we suspect to contain a signal Ah, where A is some amplitude and $\langle h, h \rangle = 1$. Then in our filtering, as we maximize the signal to noise over time translations of the template h, we measure the SNR time series

$$\rho(\tau) = \langle n, he^{2\pi i f\tau} \rangle + A \langle h, he^{2\pi i f\tau} \rangle$$
(3.30)

See Eqn. 3.7. We define

$$\alpha(\tau) \equiv \langle h, he^{2\pi i f \tau} \rangle, \tag{3.31}$$

which is the autocorrelation of the whitened template. The autocorrelation takes its maximal value of 1 at $\tau = 0$. In Eqn. 3.30, τ is to be interpreted as the difference in time between the actual coalescence time and the time-shifted template. In actual analysis, in which we do not know the actual coalescence time, we would take $\tau = 0$ to correspond to the time of the maximum SNR.

These considerations give us an estimate for the amplitude A by maximizing Eqn. 3.30 over time and taking an ensemble average, so that the noise term disappears, giving $A \approx \langle \rho_{\max} \rangle$. Now put all the measurable quantities on the same side of the equality to find that

$$\rho(\tau) - \rho_{\max}\alpha(\tau) = \langle n, he^{2\pi i f \tau} \rangle.$$
(3.32)

On the right, we have only the noise term, which will be Gaussian distributed if the noise n is Gaussian. On the left, we have only quantities that are measurable from the data and templates. We can exploit this to compute a χ^2 statistic from a given trigger, namely

$$\chi^{2} = \int_{0}^{T_{\max}} |\rho(\tau) - \rho_{\max}\alpha(\tau)|^{2} d\tau$$
 (3.33)

$$= \int_0^{T_{\text{max}}} |\langle n, he^{2\pi i f\tau} \rangle|^2 d\tau.$$
(3.34)

Here T_{max} is a tunable parameter that corresponds to the number of degrees of freedom in the χ^2 . If the SNR time series is computed with time resolution Δt , then the number of degrees of freedom is $N = T_{\text{max}}/\Delta t$.

In Fig. 3.4, we show an example of the statistical properties of the background triggers compared to triggers arising from simulated signals added to real detector data. In Gaussian noise, the SNR is the optimal detection statistic, and furthermore the background distribution in SNR falls exponentially. Here we see that the background has long excursions into the high SNR regime where, without the additional χ^2 statistic, it would blend into the signal. However, background triggers usually also have higher χ^2 than one would expect from signals. Therefore, the χ^2 test helps to distinguish between background and signal when the background has non-Gaussian fluctuations. We note that in the high SNR regime, simulated signals also obtain large χ^2 . This fact arises from the discreteness of the template bank. Typically the simulated signal does not exactly match any of the templates in the bank. As the signal gets louder, the small fractional mismatch (e.g. $\mathcal{M}_{\min} = 0.97$) becomes amplified and pulls the signal dangerously close to the background. In particular, since our searches have (so far) used non-spinning templates, a signal from a compact binary with spin could register in the pipeline with a large χ^2 , obscuring it into the background. Thus, while the SNR is optimal in Gaussian noise with known signals, it is decidedly non-optimal when the background in non-Gaussian or the actual signal does not match the template. The latter observation provides even more motivation for the move to including more physical effects in the templates. We want not only to maximize the recovered signal-to-noise, but also to minimize the χ^2 for signal.

3.3.4 Data quality metrics

Another approach we take against non-Gaussian behavior is to monitor the state of the interferometer and its environment and look for disturbances which are likely to propagate into the data stream. In severe cases, one might flag the data as unanalyzable; in less severe cases, one might flag the data as suspicious but analyzable. In post-processing of the triggers, these *data quality* (DQ) flags help to search for problems in the data which might have appeared to be a real event. Vetoes are assigned to categories based on the severity of instrumental problems and on how well the couplings between the GW and auxiliary channels are understood. Correspondingly, CBC searches assign data to one of four DQ categories, as described in Tbl. 3.2.

All LIGO detectors are equipped with environmental and instrumental monitors; their output is

Table 3.2: Data quality metrics for gravitational wave searches. Searches for gravitational waves from compact binary coalescence rely crucially on having data that is as clean as possible. A multitude of channels monitor the state of the interferometer and its environment. These channels are used to partition the data quality into four categories.

Category 1	Seriously compromised or missing data. The data are entirely unusable, to the extent that they would corrupt noise PSD estimates. These times are excluded from the analysis as if the detector was not operating.		
Category 2	2 Instrumental problems with known couplings to the GW channel. Althoug the data are compromised, these times can still be used for PSD estimation Data flagged as Category 2 are analyzed in the pipeline, but any trigger occurring during these times are discarded. This reduces the fragmentation of data, maximizing the amount of data that can be analyzed.		
Category 3	Likely instrumental problems, casting doubt on triggers found during the times. Data flagged as category-3 are analyzed and triggers are processe However, the excess noise in such times may obscure signals in clean dat Consequently, the analysis is also performed excluding time vetoed at Categor 3, allowing weaker signals in clean data to be extracted. Data vetoed Category 3 are excluded from the estimation of upper limits on GW ever rates.		
Category 4	The data remaining after the removal of Category 1–3 data. Category 4 data are considered good data and are analyzed in full.		

recorded in the detectors auxiliary channels. Periods of heightened activity in these channels (e.g., as caused by elevated seismic noise) are automatically marked with DQ flags. Data quality flags can also be added manually if the detector operators observe poor instrumental behavior. If a DQ flag is found to be strongly correlated with CBC triggers, then it can be used as a DQ veto. Data quality vetoes must be *safe*, i.e., not triggered by real GWs. Veto safety is assessed by comparing the fraction of *hardware injections* that are vetoed with the total fraction of data that is vetoed. Hardware injections are signals injected *directly* into the detector by actuating on the test masses. The hardware injections are meant to simulate exactly the effects of a real gravitational wave passing through the instrument. In particular, hardware injections may couple in some unknown way to channels other than the gravitational wave channel. Such channels must not be used for identifying bad data.

3.4 The IHOPE Pipeline

We conclude this chapter with an overview of one specific implementation of the techniques discussed here, the IHOPE pipeline [5, 63], which was used in the initial LIGO compact binary searches described in the next chapter [3, 4]. Here, we discuss implementation features which are specific



Figure 3.5: Architecture of the IHOPE analysis pipeline.

to the IHOPE pipeline. Later we will meet the GSTLAL pipeline, which implements some of these signal processing techniques in different forms.

Fig. 3.5 gives a schematic overview of the IHOPE pipeline. Before the analysis, we generate a collection of *science segments*, which are those times in which the detector was operating and collecting Category 2 or better data. The individual instrument observing times are broken up into 2048 second segments; segments shorter than 2048 seconds are discarded. In each segment, we compute the PSD, generate template banks, and filter the data. Events exceeding some predetermined SNR threshold are recorded. Since a single event can produce more than one trigger, we *cluster* triggers, taking triggers which are nearby in time and keeping only the loudest one. After clustering, we look for *coincidence* between the triggers in separate detectors. The triggers in each detector must be nearby in both time and mass to be considered further [74]. Only coincident triggers are considered for further analysis. After the coincidence stage, we perform a second pass through the data using only those templates that generated coincident triggers. For the triggers that are generated in this second pass, we compute the computationally expensive χ^2 statistic; performing this calculation in the second stage allows us to avoid computing this statistic for triggers that can be discarded by other means (such as coincidence). At the end of the second pass through the pipeline, including inter-detector coincidence, we obtain a list of triggers ranked by a detection statistic which depends on the search but generically takes the form of the signal-to-noise ratio re-weighted by χ^2 .

To understand the significance of the foreground triggers, we measure the background rate of coincident triggers as a function of combined re-weighted SNR by performing numerous time-shift analyses. In each one we artificially introduce different relative time shifts in the data from each detector. The time shifts that are introduced must be large enough such that each time-shift analysis is statistically independent. To perform the time-shift analysis in practice, we simply shift the triggers generated at the first matched- filtering stage of the analysis, and repeat all subsequent stages from multi-detector coincidence onwards. Shifts are performed on a ring: for each observational coincidence period (i.e., data segment where a certain set of detectors is operational), triggers that are shifted past the end are re-inserted at the beginning. To ensure statistical independence, time shifts are performed in multiples of 5s; this ensures that they are significantly larger than the light travel time between the detectors, the autocorrelation time of the templates, and the duration of most non-transient glitches seen in the data. Therefore, any coincidences seen in the time shifts cannot be due to an astrophysical source, and are most likely due to noise-background triggers. It is possible, however, for a GW-induced trigger in one detector to arise in time-shift coincidence with

noise in another detector. Indeed, this issue arose in Ref. [3], where a blind injection was added to the data to test the analysis procedure.

In the following chapter, we report on two searches performed using the IHOPE pipeline. There we will see what the output of the pipeline looks like and the methods used to interpret the results, which is the final stage of the IHOPE pipeline.

Chapter 4

Two Searches for CBCs in Data from LIGO's Sixth and Virgo's Second and Third Science Runs

It is the weight, not the number of experiments, that is to be regarded.

Isaac Newton

LIGO's fifth science run concluded in 2007, with the Hanford and Livingston detectors operating at the initial LIGO design sensitivity for frequencies above 60Hz. LIGO's sixth science run took place from July 7, 2009 to October 20, 2010, featuring improvements upon the initial design [55,56] which allowed for a stable high-power laser configuration. The upgrades led to detectors with sensitivities that exceeded the initial design sensitivity above 60Hz by up to a factor of two in strain, as shown in Fig. 4.1. Thus, LIGO's sixth science run went beyond the goals of initial LIGO, demonstrating the most sensitive instrument ever constructed for the detection of gravitational waves. The Virgo observatory also collected data during 2009 and 2010 in what were its second and third science runs.

We performed two searches [3,4] for gravitational waves signals from compact binary coalescence in these data. The templates used in our searches covered the mass range $2M_{\odot} \leq M_{\text{total}} \leq 100M_{\odot}$ and component masses $1M_{\odot} \leq m_i \leq 99M_{\odot}$, but neglected component spins. This parameter space includes each of the three compact binary sources: binary neutron stars, neutron star-black hole binaries, and binary black holes. For our searches, we break down the parameter space into two subregions, "low mass" and "high mass", depending on the filter templates used for the search. Lower mass systems ($M_{\text{total}} \leq 25M_{\odot}$) merge out of band and for these systems a post-Newtonian approximation to the gravitational wave signal is sufficient for detection¹. Higher mass systems merge in LIGO's most sensitive band 100–300 Hz. For these systems, we use templates that include each of the inspiral, merger, and ringdown phases of the coalescence.

All candidate gravitational wave events identified in both searches were consistent with background. In the absence of a confident detection, we place upper limits on the rates of CBCs in the nearby Universe. The inferred upper limits on event rates depend on the assumed distributions of compact binary parameters, which presently are not well-constrained. We therefore considered several compact binary mass and spin distributions in our calculations, including non-spinning binaries, aligned-spin binaries, precessing-spin binaries (where possible), and a focused study of BNS, NSBH, BBH systems with masses Gaussian-distributed about nominal values for the neutron star and black hole masses ($m_{\rm NS} = 1.34 \pm 0.05 M_{\odot}$ and $m_{\rm BH} = 5 \pm 1 M_{\odot}$). For each population distribution, we measured the mean detectable volume of the search to a simulated population derived from that distribution . The absence of any event above the loudest event in the observed volume then implies an upper limit on the coalescence rate density [77, 78]. The upper limits for CBC rates computed from our gravitational wave observations do not yet constrain even the most optimistic estimates of event rates obtained through other approaches, which were discussed in Sec. 1.4.

In the final section of this chapter, we take a look at the results of the searches in terms of sensitivity to spinning binaries. We note that the searches performed quite well in this regard, matching closely with back-of-the-envelope calculations for the expected sensitive volume. Still there is room for improvement, which we will quantify in Chap. 6. The remainder of this thesis will then focus on demonstrating that some of this sensitive volume can in fact be recovered with the use of spinning templates.

4.1 LIGO's sixth and Virgo's second and third science runs

During LIGO's sixth science run (S6), each of the two LIGO sites operated a single four kilometer laser interferometer (H1 and L1). The two kilometer H2 instrument which operated at the Hanford site in earlier science runs was not operational in S6. Following LIGO's fifth science run (S5), several hardware changes were made to the LIGO detectors so that prototypes of aLIGO technology could be installed and tested [55, 56]. These upgrades included the installation of a higher power laser, a thermal compensation system, and the implementation of a DC readout system that included

¹The searches were divided at $M_{\text{total}} = 25 M_{\odot}$, but this was probably a sub-optimal choice. See Fig. 4.3.



Figure 4.1: Detector sensitivities in S6 and VSR2/3. noise for the LIGO and Virgo detectors in the S6 and Virgo VSR2/3 runs [2]. To the right, we show the horizon distance computed from these spectra as a function of M_{total} for equal-mass non-spinning binaries in S6/VSR2-3 (solid lines). Using the spectra shown in Fig. 2.6, we also compute the horizon distance for the LIGO and Virgo detectors in S5 and VSR1 (dashed lines). The horizon distance is computed with the IMRPHENOMB waveform approximation. We note that while the strain sensitivity improved between S5 and S6 by up to a factor of two, the horizon distance to CBCs increased only ~ 15% because most of the sensitivity improvement occurred at high frequencies, where the CBC signal spends very little time. On the other hand, the Virgo detector achieved more than a factor of two improvement in horizon distance between VSR1 and VSR2/3, since its greatest improvements were at low frequencies.

a new output mode cleaner on an aLIGO seismic isolation table [79]. In addition, the hydraulic seismic isolation system was improved by fine tuning its feed-forward path. In the period between Virgo's first science run (VSR1) and Virgo's second science run (VSR2), several enhancements were made to the Virgo detector. Specifically, a more powerful laser was installed, along with a thermal compensation system and improved scattered light mitigation. During early 2010, monolithic suspensions were installed, which involved replacing Virgo's test masses with new mirrors hung from fused-silica fibers [80]. Virgo's third science run (VSR3) followed this upgrade. In Fig. 4.1a, we show "representative" (see Ref. [2]) strain spectral densities for each of the three detectors operating during the S6 and VSR2/3.

A convenient measure of the sensitivity of a detector to CBC signals is the *horizon distance*, which is the distance at which an optimally oriented and optimally sky-positioned compact binary would produce a signal with expected SNR of $\langle \rho \rangle = 8$ in that detector. The horizon distance combines the detector strain sensitivity, encoded in its power spectrum, with the expected form of the CBC signals to produce a single quantity that summarizes the sensitivity of the detector to those signals at a given time. The horizon distance also helps establish a benchmark for the performance of a CBC search pipeline. If we assume an isotropic distribution of sky position and binary orientation angles, then the mean sensitive distance at $\langle \rho \rangle = 8$ is given in terms of the optimal horizon distance by

$$D_{\rm SM} = \frac{D_{\rm horiz}}{2.26}.\tag{4.1}$$

The angle-averaged sensitive distance $D_{\rm SM}$ is called the *SenseMon* distance. From the SenseMon distance, we can compute a naive expectation for the sensitive volume

$$V_{\rm SM} = \frac{4\pi}{3} D_{\rm SM}^3.$$
(4.2)

In actual searches, we *measure* the sensitive volume by performing the analysis with simulated signals added to the data. Comparing the measured sensitive volume to the SenseMon sensitive volume gives us a sense of whether the pipeline is performing as expected.

The expected SNR for a signal h in a detector with spectral density $S_n(f)$ is given by

$$\langle \rho \rangle = \sqrt{\langle h | h \rangle},\tag{4.3}$$

where the inner product $\langle | \rangle$ is defined by Eqn. 3.3. If $h_{1 \text{Mpc}}$ denotes the signal from the same binary system located at 1 Mpc, and D is the distance to the source in Mpc, then $h = h_{1 \text{Mpc}}/D$. Setting $\langle \rho \rangle = 8$ in Eqn. 4.3, we obtain the expression

$$D_{\text{horiz}} = \frac{\sqrt{\langle h_{1\text{Mpc}} | h_{1\text{Mpc}} \rangle}}{8} \text{ Mpc.}$$
(4.4)

In Fig. 4.1b, we show the horizon distance for equal mass non-spinning binaries computed from the spectra in Fig. 4.1a. Our calculation uses the IMRPHENOMB approximation, so as to include the effects of merger and ringdown. We show the horizon distance only up to $M_{\text{total}} = 100 M_{\odot}$, since this value is the upper mass cutoff in the searches described below, but the detectors have sensitivity to systems up to roughly

$$M_{\rm max} \approx \frac{c^3}{6\sqrt{6}Gf_{\rm low}} \approx 350 M_{\odot},\tag{4.5}$$

where we have taken $f_{\text{low}} = 40$ Hz. At masses beyond $350 M_{\odot}$, the entire CBC signal is out of the LIGO band².

Since the noise in the LIGO and Virgo detectors is not stationary, the horizon distances shown

²Actually, the GW signal consists of a superposition of waves with frequencies that are integer multiples of the orbital frequency. The dominant frequency is $f_{\rm GW} = 2f_{\rm orb}$, but other multiples also contribute to the signal, bringing part of the signal for masses beyond $M_{\rm total} = 350 M_{\odot}$ back into band.



Figure 4.2: Variation of sensitivity over the course of the S6/VSR2-3 science runs. We show (left) the distribution of horizon distances for a $m_1 = m_2 = 1.4 M_{\odot}$ binary neutron star signal, computed from each of the 2048 s blocks of data analyzed in the S6/VSR2-3 low mass CBC search. Each point in the time series (right) indicates the average sensitive distance over an interval of seven days and the error bar indicates the variance in the sensitivity during that week. The detector sensitivity is a strong function of time, changing in response to local conditions and periodic maintenance and commissioning of the detectors.

in Fig. 4.1b conceal important information about the variability of the detector sensitivities over the course of S6/VSR2-3. In the S6/VSR2-3 CBC searches, the spectral density was computed on 2048-second blocks of contiguous data [81], as explained in Sec. 3.4. We account for non-stationary detector behavior by recomputing the spectral density for every 2048 second block of data. In Fig. 4.2, we show the distribution of the horizon distance towards a canonical $m_1 = m_2 = 1.4 M_{\odot}$ binary neutron star source for each of the 2048 second blocks of data analyzed in the low mass CBC search We note the wide variation in sensitivity over the course of such a long run. The Hanford site finished S6 operating at a sensitivity with horizon distance of 45 Mpc. At the beginning of the S6 observational run, Hanford had an observable distance of only 30 Mpc. This improvement corresponds to an increase is detection rate by a factor of $(45/30)^3$, which is slightly more than a factor of three. The decrease in horizon distance for the Virgo detector in VSR3 is due to a mirror with an incorrect radius of curvature being installed during the conversion to monolithic suspensions.

4.2 Search parameters

Together our two CBC searches covered the mass space $2M_{\odot} \leq M_{\text{total}} \leq 100M_{\odot}$ and $1M_{\odot} \leq m_i \leq 99M_{\odot}$. We split this parameter space into two subregions depending on whether we use inspiralonly post-Newtownian filters or templates which also include merger and ringdown. The transition

Table 4.1: Observational times for the $S6/VSR_2$ -3 high mass search. We list the nine observational epochs by their GPS start and end times, together with the total amount of analyzed time in each coincident detector combination for the S6/VSR2-3 high mass search. These analysis times correspond to the data that are used for computing upper limits. Therefore, we have removed playground data and applied Category 1–3 vetoes (see Sec. 3.3.4). Due to search-specific vetoes, the analyzed times for the low mass search are not precisely the same; however, the difference is on the order of 1%.

Enceh	Analyzed time in seconds (days)			
Бросп	H1L1V1	H1L1	H1V1	L1V1
931035296-935798487	555544(6.43)	98191(1.14)	835797 (9.67)	784626 (9.08)
937800015 - 944587815	372986 (4.32)	$203292 \ (2.35)$	$1396983 \ (16.2)$	$396321 \ (4.59)$
944587815 – 947260815	451084 (5.22)	$350747 \ (4.06)$	210478(2.44)	344620 (3.99)
949449543 - 953078487	_	1087283 (12.6)	_	_
953078343 - 957312087	_	1090657 (12.6)	_	_
957311943 - 961545687	_	895439(10.4)	—	_
961545543 - 965174487	_	$953236\ (11.0)$	_	_
965174343 - 968544087	671531(7.77)	442106(5.12)	197317 (2.28)	$354660 \ (4.10)$
968543943 - 971622087	625788(7.24)	$221240 \ (2.56)$	499116(5.78)	299434(3.47))
Total	2676933 (31.0)	$5342191 \ (61.8)$	$3139691 \ (36.3)$	$2179661 \ (25.2)$

between the two regimes for an iLIGO design sensitivity is depicted in Fig. 4.3. This figure shows that for binary systems with total masses exceeding ~ 25 M_{\odot} , the neglect of the merger and ringdown phases of coalescence leads to a loss of more than 10% in SNR. Consequently, the upper mass limit for the templates used in the "low mass" search was set to $M_{\text{total}} = 25M_{\odot}$. The "high mass" search covered the remaining parameter space $25M_{\odot} \leq M_{\text{total}} \leq 100M_{\odot}$. For even higher masses, the expected signals spend so little time in the iLIGO band that different detection methods are required [82]. We note that while the templates used for these searches neglect the spin of the binary components, the search is still sensitive to spinning signals, as we will demonstrate below. However, in later chapters, we will also demonstrate that the neglect of spin effects leads to a suboptimality in the search, similar to the way that neglecting merger and ringdown effects above $M_{\text{total}} = 25M_{\odot}$ would also lead to a suboptimal search.

For each of the two searches, we divided the analysis of the S6/VSR2-3 data into 24 disjoint sub-periods (see Tbl. 4.1) and analyzed each sub-period independently in order to crudely capture the variability of the background across different sets of observing instruments and variability over time for any given instrument. We divided the data into nine different analysis time periods (each roughly 2 months long) with the boundaries corresponding to times of significant commissioning of one or more of the instruments. In five of these periods, all three instruments operated, and there are therefore four different coincident observation modes: H1L1, H1V1, L1V1, and H1L1V1. In the other four, only the H1 and L1 instruments operated, and only the H1L1 coincident observation mode is possible. Thus, the S6/VSR2-3 searches, nominally just a single search, actually consisted



Figure 4.3: Defining the inspiral-merger-ringdown boundary. We compute the fitting factor of IMRPHE-NOMB signals (as used in the "high mass" analysis) against a bank of post-Newtonian inspiral-only waveforms assuming zero spin (as used for the "low mass" search). The solid lines indicate the contours of constant fitting factor. We see that the effects of merger and ringdown in the template waveform become non-negligible (F < 0.97) at total masses above $M_{\text{total}} = 15 M_{\odot}$. The IMRPHENOMB waveforms used for creating this figure were not fully trusted at the beginning of the science run, and it was thought that merger effects become important only at higher masses; hence the low mass search used $M_{\text{total}} = 25M_{\odot}$ for the upper mass limit.

formally of 5 * 4 + 4 = 24 independent sub-searches. We will show below how we combine the results of independent searches.

4.3 Evaluating the search sensitivity

It is crucial to quantify the performance of the search pipeline in terms of its ability to detect its target signals. For one, doing so allows us to identify sub-optimalties in the search techniques. In Sec. 4.1, we developed the notion of the horizon distance in order to characterize the performance of the *detectors* from the perspective of observing CBC signals. In that discussion, we rather arbitrarily chose a nominal value of $\langle \rho \rangle = 8$ to correspond to a detectable signal, but the actual threshold for detection depends on the search methods and implementation. We now aim to understand the extent to which our search pipeline exploited the detectors' sensitivities. In particular, how does the *pipeline* performance on actual data compare to these idealized expectations for search sensitivity based only the detector behavior and signal morphology? Understanding the pipeline sensitivity is also fundamental to the interpretation of the search results in terms of coalescence event rates.

We estimate the sensitive search volume by reanalyzing the data with the addition of a large number of simulated signals ("software injections") distributed according to some distribution $f(\vec{\theta})$ on the source parameters $\vec{\theta}$ (see Tbl. 3.1 for a description of these parameters). We treat the triggers arising from the simulation analyses in exactly the same way we treat triggers in the actual analysis of the LIGO data. We apply the same filtering, the same clustering methods, the same χ^2 statistic, the same template banks, etc., to ensure that the triggers resulting from the simulated signals are exactly as they would have appeared in the pipeline in the actual analysis had that signal been present. After performing the simulation analysis, we then look for associations between triggers and the performed injections, which are usually constrained to lie within some small time window around the injected time (e.g. one second).

We then compute the sensitive volume by specifying a fixed false alarm probability (FAP), which is used as a threshold to cut between "found" and "missed" injections. An injection is considered "found" if there is a trigger within its coincidence time window whose FAP is smaller than the given threshold. All other injections are considered "missed". We construct N distance bins d_i and compute the efficiency in each distance bin as

$$\varepsilon_i = \frac{N_{\text{found}}}{N_i},\tag{4.6}$$

where N_i is the total number of simulated events in that distance bin. We compute the volume from the efficiency with the numerical integral

$$V = \sum_{i=1}^{N} 4\pi \varepsilon_i d_i^2 \Delta d_i, \qquad (4.7)$$

where Δd_i is the distance bin width. The distances d_i are bin representatives (e.g. left endpoint, right endpoint, midpoint). In forming this expression, we have implicitly assumed a uniform spatial distribution of sources.

The observed volume clearly depends on the input distribution $f(\vec{\theta})$. In cases where the correct astrophysical distribution is unknown, such as with the binary masses, it is easier to interpret simulations in which these parameters are fixed to some small range and only the parameters with known distributions are varied. Hence, below we quote sensitive volumes and rate estimates in narrow mass bins, for example total mass or m_1 - m_2 bins, which mitigates the distribution-dependent effects of the sensitivity calculation. In some cases, there is a natural choice for the distribution. For instance, Eqn. 4.7 already implicitly assumes a uniform spatial distribution of sources, a reasonable approximation for searches sensitive to distances beyond ~ 20 Mpc. We can see from Fig. 4.2a that the H1 and L1 were in this regime for most of S6; V1 on the other hand sometimes was not (note the significant difference between VSR2 and VSR3 in Fig. 4.2b). Remember, however, that Fig. 4.2 shows the horizon distance to binary neutron stars, one of our weakest sources.

Even with a well-known distribution, the sensitivity calculation suffers from a number of other uncertainties. One such uncertainty is the binomial errors

$$\Delta \varepsilon_i = \sqrt{\frac{\varepsilon_i (1 - \varepsilon_i)}{N_i}} \tag{4.8}$$

in the computed efficiency due to the finite number of injections. Uncertainties in the calculation also arise from imperfect knowledge of the detector calibration (conservatively estimated to be 14% in amplitude combined over all three detectors and over the entire observational period) and imperfect waveform models (taken to be a one-sided 10% bias towards lower sensitive distance). We marginalize over the uncertainty in volume to obtain an upper limit which takes into account these uncertainties [77].

4.4 Estimating binary coalescence rates

Here we review the loudest event statistic, described in Refs. [77,78,83], which we will use in Sec. 4.6 to compute Bayesian 90% confidence level upper limits on the coalescence rate from our S6/VSR2-3 observations. One of the advantages to using a Bayesian formalism for estimating event rates is that it provides a natural prescription for combining the results of independent experiments. As we mentioned above, these searches each consisted of 24 independent sub-experiments. Additionally, we have upper limit estimates from the S5/VSR1 searches [47, 84, 85], which we would like to fold into our present observations. After describing the loudest event statistic, we show how we combine results from multiple independent experiments, marginalize over the various sources of uncertainty, and compute upper limits.

We formalize the problem by considering a set $\vec{\mu}$ of model parameters we wish to constrain and a set \vec{x} of outcomes for N independent experiments. In our case, the model parameters $\vec{\mu}$ will consist of the rate density μ plus a number or "error" parameters which we will marginalize over. The outcomes \vec{x} of the experiments are the loudest event observed in each experiment. To make inferences on the parameters $\vec{\mu}$ after having observed the outcomes \vec{x} , we apply Bayes' theorem, which states

$$p(\vec{\mu}|\vec{x}) = p(\vec{x}|\vec{\mu}) \frac{p(\vec{\mu})}{p(\vec{x})}.$$
(4.9)

If the observations \vec{x} are independent, as for the outcomes of our 24 sub-periods, then the expression above factorizes as

$$p(\vec{\mu}|\vec{x}) = \prod_{i} p(x_i|\vec{\mu}) \frac{p(\vec{\mu})}{\prod_{i} p(x_i)}.$$
(4.10)

The conditional probability density $p(\vec{x}|\vec{\mu})$ is called the *likelihood* of the outcome \vec{x} given the parameter $\vec{\mu}$, $p(\mu)$ is the prior distribution on the parameters $\vec{\mu}$, and $p(\vec{x})$ is the prior distribution on the outcome of the experiments.

Before delving into the loudest event statistic, we take a step back and consider the simple and familiar problem of estimating event rates when the number of observed events is certain. One can view the loudest event statistic as essentially performing the same calculation but marginalizing over the unknown number of observed events. Suppose that in performing a search for gravitational waves, we observe exactly n CBC events. If the search is sensitive to a volume V and lasts for a duration T then one can show by applying Bayes theorem that the posterior distribution on the rate is

$$p(\mu|n) = \frac{p(\mu)(\mu VT)^n e^{-\mu VT}}{\int_0^\infty p(\mu)(\mu VT)^n e^{-\mu VT} d\mu}.$$
(4.11)

In deriving this result, we treat the foreground process as a Poisson process with mean number of events μVT .

For the special case in which the prior is uniform in the rate, the posterior becomes

$$p(\mu|n) = \frac{(\mu VT)^n e^{-\mu VT}}{\int (\mu VT)^n e^{-\mu VT} d\mu}.$$
(4.12)

A simple way to understand this result is to compute the peak of the distribution, which turns out to be

$$\mu_{max} = \frac{n}{VT}.\tag{4.13}$$

We see that given the observation of n events and no prior information about the rate (uniform prior), the most likely rate is just the number of observed events divided by the observed volume.

Another important special case are priors of the form $p(\mu) \propto \mu^n e^{-\mu V_1 T_1}$ with *n* an integer. This prior would arise naturally as a posterior from an initial search (using a flat prior) which identified

\overline{n}	5% lower limit	90% upper limit	95% upper limit
0	-	2.302	2.996
1	0.355	3.890	4.744
2	0.818	5.322	6.296
3	1.366	6.681	7.754
4	1.970	7.994	9.154
5	2.613	9.275	10.513

Table 4.2: Event rate confidence intervals for a simple counting experiment.

n events as in Eqn. 4.12. If in a subsequent search with sensitivity V_2T_2 , we identify *m* more events and wish to use the posterior from the previous search as a prior to this search, we find that the combined posterior is

$$p(\mu|n,m) \propto \mu^{n+m} e^{-\mu(V_1 T_1 + V_2 T_2)}.$$
 (4.14)

This result illustrates an important property of this approach to estimating the rate: the posterior distribution one obtains starting from a flat prior *depends only on the total number of observed events and the total observed volume*. The loudest event statistic does not have this property – the posterior is sensitive to the way in which the data are divided.

We can use the posterior in Eqn. 4.12 to compute upper limits or confidence intervals. In the case of an experiment using a uniform prior which makes no detections, the upper limit at 90% confidence is

$$\hat{\mu} = \frac{2.303}{VT}.$$
(4.15)

For n = 1, the 90% upper limit is

$$\hat{\mu} = \frac{3.890}{VT}.$$
(4.16)

We will see below that in loudest event statistic formalism, these two results are limiting cases, and a parameter Λ interpolates between them to take into account uncertainty in possibly having made a false dismisal. We present upper and lower limits for other values of n in Tbl. 4.2. This table is useful for interpreting the upper limits obtained by the loudest event method.

We now generalize these results to the case in which the identity of the foreground events (signal or noise) is unknown. The loudest event statistic provides a method to compute the consistency of the loudest event with the background, and in doing so the statistic effectively interpolates between the two extreme cases: (i) where the loudest event is definitely due to signal and (ii) where the loudest event is definitely due to noise. Beginning with the trivial observation that there were no events observed above the loudest event, one can show that Bayes' theorem implies that the rate posterior for a single search with loudest event x, is given by

$$p(\mu|x) = \frac{p(\mu)(1+\mu V(x)T\Lambda(x))e^{-\mu V(x)T}}{\int p(\mu)(1+\mu V(x)T\Lambda(x))e^{-\mu V(x)T}d\mu},$$
(4.17)

where V(x) is the average volume of the region in which events would register in the pipeline louder than x, and $\Lambda(x)$ is a measure of the relative likelihood of the loudest event coming from the foreground versus coming from the background. Compare Eqn. 4.17 to Eqn. 4.12 with n = 0 and n = 1. We can think of the parameter Λ as interpolating between these two cases. We define Λ by

$$\Lambda(x) = -\frac{d\log n_f(x)/dx}{d\log P_b(x)/dx},\tag{4.18}$$

where $n_f(x)$ is the expected number of signal events above the loudest event and $P_b(x)$ is the probability of having a background event above the loudest event. For $\Lambda < 1$, the loudest event is more consistent with background, while $\Lambda > 1$ indicates the loudest event is inconsistent with the background. For our searches, the detection statistic x is false alarm number x = RT, where R is the false alarm rate, and therefore $P_b(x) = e^{-x}$. The mean number of foreground events is $n_f = \mu VT$. Putting these into Eqn. 4.18, we obtain the expression

$$\Lambda(x) = \frac{1}{V} \frac{dV}{dx},\tag{4.19}$$

which is how $\Lambda(x)$ is actually computed in the IHOPE pipeline.

To combine multiple experiments, we use the posterior from one analysis as the prior to the next. We ignore marginalization over volume uncertainties for now, treating that issue in the next section. For two experiments with sensitive volumes V_1, V_2 and likelihoods Λ_1, Λ_2 , the rate posterior is given by

$$p(\mu|V_1T_1, \Lambda_1, V_2, \Lambda_2) \propto (1 + \mu V_1T_1\Lambda_1)(1 + \mu V_2T_2\Lambda_2)e^{-\mu(V_1T_1 + V_2T_2)}.$$
(4.20)

We note three special corner cases (see Fig. 4.4) that give a good summary of the overall behavior of this posterior. In each of these cases, the posterior in Eqn. 4.20 simply reduces to a posterior of the form in Eqn. 4.14 for a certain value of n. We have (i) $\Lambda_1 = \Lambda_2 = 0$, corresponding to Eqn. 4.14 with n = 0, (ii) $\Lambda_1 = 0, \Lambda_2 = \infty$, or vice versa, corresponding to Eqn. 4.14 with n = 1 and (iii) $\Lambda_1 = \Lambda_2 = \infty$, corresponding to Eqn. 4.14 with n = 2. Thus, for two experiments the upper limits satisfy $2.30 \leq \mu_{90} \sum V_i T_i \leq 5.32$. By similar reasoning, one can argue that for three experiments $2.30 \leq \mu_{90} \sum V_i T_i \leq 6.68$. Where in that range the upper limit falls depends on the computed



Figure 4.4: Combining independent experiments with the loudest event statistic. When you combine two experiments using the loudest event statistic, the resulting 90% upper limit (solid vertical lines) is bounded by 2.302VT to 5.322VT, corresponding to the two extreme values for the counting experiment with n = 0 or n = 2 observed events, respectively. The outcome depends on whether there are large values of Λ in one or both of the experiments.

values for Λ_i . For an arbitrary number of independent experiments the posterior distribution on the combined analysis is given by

$$p(\mu|\vec{V},\vec{\Lambda}) \propto p(\mu) \exp(-\mu \sum_{k} V_k T) \prod_{k} (1 + \mu \Lambda_k V_k T), \qquad (4.21)$$

where V_k and Λ_k are the measured volumes and Λ values for the individual experiments.

4.5 Marginalizing over uncertainties

The upper limit calculation suffers from a number of significant uncertainties. We propagate three sources of uncertainty into the upper limit: Monte Carlo errors, waveform errors, and (amplitude) calibration errors. The Monte Carlo errors arise from the finite number of simulations use to assess the sensitivity. The calibration uncertainty arises from an imperfect measurement of the response function R(f), which converts ADC counts into strain [86,87]. We translate the overall amplitude uncertainties in R(f) for each detector into an equivalent volume uncertainty and propagate this uncertainty to the measured sensitive volume. Similarly, uncertainties in the injection waveforms, which are known to differ from precise numerical relativity calculations, directly translate into an uncertainty in the volume, and we apply corrections to account for this uncertainty when these waveform errors can be quantified. We marginalize over calibration errors coherently over the entire run, while Monte Carlo errors are marginalized over independently in each sub-search.

We compute the volume integral in Eqn. 4.7 as

$$V = \sum_{i=1}^{N} \frac{N_{\text{found}}}{N_i} 4\pi d_i^2 \Delta d_i, \qquad (4.22)$$

where d_i is a representative for the *i*th distance bin (e.g. the center of the bin) and Δd_i is the bin width. There are statistical uncertainties in the estimated efficiency in each distance bin due to the finite number of injections. The statistical uncertainties in each distance bin are independent, and therefore they propagate to the volume according to

$$\Delta V = \sqrt{\sum_{i=1}^{N} \frac{N_{\text{found}} (1 - N_{\text{found}})}{N_i} (4\pi d_i^2 \Delta d_i)^2},$$
(4.23)

assuming that the error in the efficiency calculation is binomially distributed.

The Monte Carlo errors associated with independent experiments are of course also independent. To marginalize over the statistical uncertainty, we input an assumed prior distribution on the statistical volume fluctuations and marginalize over this distribution. Consider a prior $p(\mu, \vec{\delta}) = p(\mu)p(\vec{\delta})$, where $\vec{\delta}$ is a vector of the statistical errors ΔV_i in the measured volume from each experiment as given in Eqn. 4.23. Using Bayes' theorem in the form of Eqn. 4.10, we marginalize over these Monte Carlo errors by

$$p(\mu|\vec{x}) = \int \prod_{i} p(x_i|\mu, \vec{\delta}) \frac{p(\mu)p(\vec{\delta})}{\prod_{i} p(x_i)} d\vec{\delta}.$$
(4.24)

Since the errors are independent, $p(\vec{\delta})$ factorizes and the equation above is equivalent to

$$p(\mu|\vec{x}) \propto p(\mu) \prod_{i} \int p(x_i|\mu, \delta_i) p(\delta_i) d\delta_i.$$
(4.25)

That is, when the uncertainties on the individual experiments are independent, then it is acceptable to marginalize over each experiment separately and multiply the marginalized likelihoods together at the end. Now consider where δ represents the characteristic magnitude of calibration uncertainty. Since the exact calibration error is unknown, we wish to marginalize over this parameter. We again assume a prior distribution $p(\mu, \delta) = p(\mu)p(\delta)$ for volume uncertainties arising from the calibration. Applying Bayes' theorem, we find that

$$p(\mu|\vec{x}) = \int \prod_{i} p(x_i|\mu, \delta) \frac{p(\mu)}{\prod_{i} p(x_i)} p(\delta) d\delta.$$
(4.26)

Thus, to marginalize over the calibration uncertainty, we must first compute the posterior first each individual experiment – marginalized over the statistical uncertainties – and then marginalize the whole result over δ . For our calculations, we assume that δ couples to the volume according to³

$$V(x|\delta) = V(x)\delta, \tag{4.27}$$

and that δ is log-normal distributed with unit mean. Note that since Λ is the logarithmic derivative of V, Λ is actually invariant under the transformation in Eqn. 4.27.

4.6 Search results

We now present the results of our two CBC searches. In Fig. 4.5, we show the achieved mean sensitive distance for these two search as a function of total mass for EOBNR (gray bars) and IMRPHENOMB (colored bars) signals. On top of the bars, we plot the mean sensitive distance predicted by the SNR=8 detection criteria for IMRPHENOMB signals with $\chi = -0.5$, $\chi = 0$ and $\chi = 0.5$. These curves are derived from those in Fig. 4.1b by (i) converting the horizon distance into an angle-averaged mean distance $D_{\text{avg}} = D_{\text{horiz}}/2.26$ for each detector and (ii) considering all possible coincident observational modes (H1L1, H1V1, L1V1, H1L1V1), taking the horizon distance for the second most sensitive operating detector (a crude approximation to a trigger coincidence criterion), and finally (iii) averaging these values weighted by the coincident analyzed time for the given detector combination (see Tbl. 4.1). The fantastic agreement between these two calculations is a compliment to the search methods, delivering on what the detectors' sensitivities had made available. In the final section of this chapter, we will discuss in a bit more detail the results of the simulations of IMRPHENOMB aligned spin signals (also depicted in Fig. 4.5).

Neither search revealed a plausible GW signal from compact binary coalescence. In the absence

³Thus, δ is the cube of the uncertainty in the calibration amplitude.



Figure 4.5: Comparison of achieved mean sensitive distance to SenseMon expectations. We show (bars) the mean sensitive distance above the loudest event for the low and high mass searches as a function of the binary total mass, as measured by the full pipeline simulations. We see that the pipeline sensitivity to non-spinning signals are different for the EOBNR and IMRPHENOMB models; this gives an indication of the level of disagreement between the them. The pipeline is most sensitive to positively aligned spin systems, as expected, as these systems are intrinsically louder than negatively aligned spin systems. In all cases, the measured pipeline sensitivity agrees quite well with the naive SenseMon estimate. Here we have combined the various detector sensitivities shown in Fig. 4.1b by taking an average, weighted by analyzed time, of the sensitive distance of the second most sensitive detector in each of the four possible coincident detector combinations.

of such an event, we use the methods just described to constrain the event rate. The upper limits derived here will turn out to be consistent with even the most optimistic predictions for the binary coalescence rate. Previous searches [47] presented upper limits in terms of blue light luminosity, using units of L_{10}^{-1} yr⁻¹, where one L_{10} is 10^{10} times the solar blue light luminosity. However, during the current analysis, the sensitivity of the detectors to the systems of interest (see Fig. 4.1b) was sufficiently large that we may approximate our signals as uniformly distributed in volume. Hence, we quote our upper limits in terms of a volume density Mpc⁻³ yr⁻¹ (see Sec. 1.4 for more discussion).

4.6.1 Low mass

The low mass search used non-spinning post-Newtonian inspiral templates computed to 0PN in amplitude and 3.5PN in phase to search for compact binary coalescence in the range $2M_{\odot} \leq M_{\text{total}} \leq$
Table 4.3: Rate upper limits of BNS, BBH, and NSBH coalescence. We give the computed rate upper limits for our three canonical sources, assuming Gaussian mass distributions around the central values $m_{\rm NS} =$ $1.35\pm0.04M_{\odot}$ and $m_{\rm BH} = 5.0\pm1.0M_{\odot}$. The effective mean sensitive distance is given by $D_{\rm avg} = (3V/4\pi)^{1/3}$; the horizon distance is $D_{\rm horiz} = 2.26D_{\rm avg}$ (see Sec. 4.1). For the spinning upper limits, we assume a black hole spin distribution that is uniform in magnitude between zero and the maximal value of Gm^2/c and spin angles that are uniformly distributed on the sphere. We neglect neutron star spin.

System	BNS	NSBH	BBH
Component masses	1.35 / 1.35	1.35 / 5.0	5.0 / 5.0
(M_{\odot})			
$D_{ m horizon} (m Mpc)$	40	80	90
Non-spinning upper	1.3×10^{-4}	3.1×10^{-5}	6.4×10^{-6}
limit $(Mpc^{-3}yr^{-1})$			
Spinning upper limit		$3.6 imes 10^{-5}$	$7.4 imes 10^{-6}$
$(Mpc^{-3}yr^{-1})$			

 $25M_{\odot}$. Single detector triggers in the search were ranked according to the statistic

$$\hat{\rho} = \begin{cases} \frac{\rho}{[(1+\chi_{\rm red}^6)/2]^{1/6}} & \text{if } \chi_{\rm red}^2 > 1\\ \rho & \text{if } \chi_{\rm red}^2 \le 1, \end{cases}$$
(4.28)

where ρ is the usual signal-to-noise ratio and χ^2_{red} is the traditional χ^2 per degree of freedom (see Sec. 3.3.3). The (effective) signal-to-noise for a coincident event is given by

$$\rho_c = \sqrt{\sum_{\text{detectors } i} \hat{\rho}_i^2}.$$
(4.29)

The background is a strong function of mass. Therefore, background triggers were split into three chirp mass bins, and the significance of each event in the actual search was computed relative the background in that event's chirp mass bin. In a later step, the triggers from different mass bins are combined into one set to account for the trials factor associated with having three separate mass bins. The most significant event was an L1V1 coincidence in L1V1 time with a combined false alarm rate of 1.2 yr^{-1} . The second and third most significant events had combined false alarm rates of 2.2 yr^{-1} and 5.6 yr^{-1} respectively. All of these events were consistent with background. Having analyzed approximately half a year of data, we expect the loudest event to have a false alarm rate of $\sim 2 \text{ yr}^{-1}$.

In the absence of any plausible gravitational wave signal, we place upper limits on the event rate using the loudest event statistic just described. We considered several different source populations for computing the upper limits. In Tbl. 4.3 we present the marginalized upper limits at the 90% confidence level, assuming canonical mass distributions for BNS ($m_1 = m_2 = 1.35 \pm 0.04 M_{\odot}$), BBH



Figure 4.6: Upper limits on the binary coalescence rate as a function of mass. We show the upper limits on coalescence rate derived from the measured search depth for the two CBC searches as a function of M_{total} . The vertical line indicates the boundary between the low and high mass searches. For the low mass search, the injections were distributed uniformly M_{total} and in m_1 for a given M_{total} . For the high mass search, we used injections distributed uniformly over m_1 - m_2 ; the total mass bins above correspond the equal mass bins in Tbl. 4.4. The dark bars indicate upper limits from previous searches. The light bars indicate the combined upper limits, including the results of this search.

 $(m_1 = m_2 = 5 \pm 1 M_{\odot})$, and NSBH $(m_1 = 1.35 \pm 0.04 M_{\odot}, m_2 = 5 \pm 1 M_{\odot})$ systems. In our simulations, the non-spinning BNS injections were computed in the post-Newtonian approximation; all other systems were modeled with the EOBNR waveform family. Spinning injections were all performed using the SPINTAYLORT4 model⁴. We considered (i) a population of non-spinning compact binaries and (ii) a population of binaries in which the black hole spin is uniformly distributed in magnitude between 0 and the maximal spin Gm^2/c and spin angles uniformly distributed on the sphere. We neglect neutron star spin, as we do not expect neutron stars to have large spins. From the table, we see that the search was less sensitive to the population with spin compared to the population without spin.

We also compute upper limits, shown in Fig. 4.6, as a function of M_{total} , using an injection population distributed uniformly over M and uniformly over m_1 for a given M. For NSBH systems we present the upper limit as a function of black hole mass, keeping the neutron star mass fixed in

 $^{^{4}}$ Note that the use of different waveform models for the spinning and non-spinning injections was accidental, but conservatively we can estimate the waveform differences to be about 10% in amplitude (introducing a 30% error in volume).



Figure 4.7: Coalescence rate upper limit rates for BNS, NSBH, and BBH systems. In the left figure, we show the upper limits for NSBH as a function of the black hole mass, with the neutron star mass restricted to the range $1 - 3 M_{\odot}$. To the right, we compare the upper limits for BNS, NSBH, and BBH given in Tbl. 4.3 to the predicted rates given in Ref. [31]. The light gray regions display the upper limits obtained in the S5/VSR1 analysis; dark gray regions show the upper limits obtained in this analysis, using the S5/VSR1 limits as priors. The new limits are up to a factor of 1.4 improvement over the previous results. The lower (blue) regions show the spread in the astrophysically predicted rates, with the dashed-black lines showing the "realistic" estimates. Note: In Ref. [31], NSBH and BBH rates were quoted using a black-hole mass of $10 M_{\odot}$. We have therefore rescaled the S5 and S6 NSBH and BBH upper limits in this plot by a factor of $(\mathcal{M}_5/\mathcal{M}_{10})^{5/2}$, where \mathcal{M}_{10} is the chirp mass of a binary in which the black hole mass is $10 M_{\odot}$, and \mathcal{M}_5 is the chirp mass of a binary in which the black hole mass is $5 M_{\odot}$.

the range $1-3 M_{\odot}$ (Fig. 4.7a). In these calculations, the injected signals were non-spinning inspiralmerger-ringdown signals computed in the EOBNR approximation. Fig. 4.7b compares the upper limits obtained in this analysis (dark gray regions) to limits obtained in our previous searches up to S5/VSR1 [47] (light gray region) and to predicted rates (blue regions) for BNS, NSBH, and BBH systems. The improvement over the previous limits is up to a factor of 1.4, depending on binary mass; this reflects the additional observation time and improved sensitivity of the S6/VSR2-3 data with respect to all previous observations.

While the rates presented here represent an improvement over the previously published results from earlier LIGO and Virgo science runs, they are still above the astrophysically predicted rates of binary coalescence. There are numerous uncertainties involved in estimating astrophysical rates, including limited numbers of observations and unknown model parameters; consequently the rate estimates are rather uncertain. For BNS systems the estimated rates vary between 1×10^{-8} and 1×10^{-5} Mpc⁻³yr⁻¹, with a "realistic" estimate of 1×10^{-6} Mpc⁻³yr⁻¹. For BBH and NSBH, realistic estimates of the rate are 5×10^{-9} Mpc⁻³yr⁻¹ and 3×10^{-8} Mpc⁻³yr⁻¹ with at least an order of magnitude uncertainty in either direction [31]. In all cases, the upper limits derived here are two



Figure 4.8: Loudest events in the high mass search. We show the cumulative distributions of coincident events in the long (left) and short (right) duration categories for all of the S6/VSR2-3 high mass search. Grey shaded bands indicate 1σ -5 σ consistency with the estimated background distribution.

to three orders of magnitude above the "realistic" estimated rates, and about a factor of ten above the most optimistic predictions.

4.6.2 High mass

Our high mass search used non-spinning EOBNR (v1) templates, which include the effects of merger and ringdown, to search for coalescing binaries in the range $25M_{\odot} \leq M_{\text{total}} \leq 100M_{\odot}$. In this regime, the template waveforms can become incredibly short for the highest mass systems and these short duration templates tend to be very efficient at picking up glitches. Thus, as in the low mass search, we divided our background into categories, here based on template duration. For coincident triggers with all single detector templates having duration greater than 0.2 s, we assign the single detector triggers the same statistic used for the low mass search, given in Eqn. 4.28. However, for triggers which involve at least one template with duration smaller than 0.2 s, we compute instead

$$\hat{\rho} = \frac{\rho}{[\chi_{\rm red}^2 (1 + \rho^2 / 50)]^{1/4}}.$$
(4.30)

The combined (effective) signal-to-noise is given by Eqn. 4.29. As before, we assess the significance of an actual event relative to the background in its duration bin and later re-rank the candidate events in each category to account for the trials factor. The loudest events for the high mass search for both long and short duration events are shown in Fig. 4.8. All events in both categories are consistent with background.

Table 4.4: *High mass search sensitivity*. We give the mean sensitive distances and coalescence rate upper limits, quoted over $9 M_{\odot}$ -wide component mass bins labeled by their central values. We also quote the

limits, quoted over $9 M_{\odot}$ -wide component mass bins labeled by their central values. We also quote the chirp mass \mathcal{M} at the center of each bin. The sensitive distance in Mpc (averaged over the observation time and over source sky location and orientation) is given for EOBNR waveforms in S5 data rescaled for consistency with NR results [48], and for EOBNR (v2), IMRPHENOMB non-spinning ("PhenB nonspin"), and IMRPHENOMB spinning ("PhenB spin") waveforms in the S6-VSR2/3 data. The last two columns report 90%-confidence rate upper limits in units of 10^{-7} Mpc⁻³yr⁻¹, for bins with component mass ratios $1 \leq m_1/m_2 \leq 4$, for S5 data [6] and the cumulative upper limits over S5 and S6-VSR2/3 data.

Wa	avefor	ms		EOBNR	PhenB nonspin	PhenB spin	EOBNR	EOBNR
Sea	arch d	ata	S5	S6/VSR2-3	S6/VSR2-3	S6/VSR2-3	S5	S5-6/VSR1-3
m_1	m_2	\mathcal{M}	Distance		Distance		Upper Limit	
14	14	13	81	102	105	106	18	8.7
23	14	16	95	116	126	126	12	5.9
32	14	18	102	140	132	135	8.8	4.2
41	14	21	107	139	141	145	7.8	4.1
50	14	22	107	131	137	149	8.2	4.3
23	23	20	116	152	148	149	7.4	3.3
32	23	24	133	172	172	179	4.9	2.4
41	23	27	143	181	178	183	4.3	2.2
50	23	29	145	187	188	198	3.4	1.7
59	23	32	143	189	188	192	3.2	1.5
68	23	34	140	177	180	191	3.7	1.8
77	23	36	119	156	176	170	5.6	3.8
32	32	28	148	194	190	197	3.4	1.7
41	32	32	164	210	219	220	2.5	1.4
50	32	35	177	224	221	214	1.9	1.0
59	32	38	174	223	221	214	2.0	1.0
68	32	40	162	201	199	210	2.4	1.3
41	41	36	183	230	222	224	1.6	0.9
50	41	39	191	253	253	258	1.4	0.7
59	41	43	194	224	239	236	1.4	0.8
50	50	44	192	257	218	217	1.4	0.7

In order to evaluate the search sensitivity, we used an implementation of the recently developed EOBNR (v2) waveform family [88] as simulated signals. The injections were distributed approximately uniformly over the component masses m_1 and m_2 , within the ranges $1 M_{\odot} \leq m_i \leq 99 M_{\odot}$ and $20 M_{\odot} \leq M \leq 109 M_{\odot}$. The resulting 90% confidence upper limits on non-spinning coalescence rates are displayed in Tbl 4.4. For binaries with both component masses lying between 19 and 28 M_{\odot} we find a 90% limit of 3.3×10^{-7} mergers Mpc⁻³yr⁻¹. We use the results from Ref. [6], revised from Ref. [48], as priors for these calculations. Since the injected waveforms are phenomenological models, our upper limits will also be systematically affected to the extent that the true IMR waveforms differ from these models. We therefore also evaluated the mean sensitive distance to injections performed with the IMRPHENOMB approximation, which we discuss in more detail in the next section.

100

4.7 Comments on spin in these searches

Although each of these searches considered the pipeline sensitivity to spinning signals, neither analysis used spinning templates. As discussed in Sec. 1.3, most predictions based on the formation of compact binaries indicate that the components will have significant spin. Furthermore, as we pointed out in Secs. 1.4 and 1.5, the spins of the component black holes are known to have a potentially large effect on the emitted waveform. Component spin is expected to have several effects on our searches, compared to its performance for non-spinning systems. First, the non-spinning templates used in our searches have reduced overlap with the spinning signals, leading to a loss of sensitivity. Second, the signal-based χ^2 test values are expected to be higher than if exactly matched spinning templates were used, due to "unmatched" excess power in the signals; this would further reduce the search sensitivity.

In the low mass search, we considered the effects of spin *only* for NSBH and BBH systems, as described in Tbl. 4.3. Our simulations were performed using inspiral-only post-Newtonian waveform models, in which the neutron spin was assumed to be zero and the black hole spin was arbitrarily oriented with a uniform distribution in spin magnitude. We found through these simulations that the analysis was less sensitive to this population of spinning signals compared to an otherwise equivalent population without spin by about ~15–20%. These simulations answer the question of how far the search was sensitive to generically spinning systems in this mass range (given that the search was conducted with non-spinning templates) but it gives no indication of how well the search could have done with templates that include spin effects⁵. In Sec. 6.3, we'll take a much closer look at how much signal-to-noise may be at stake if we ignore spin effects in aLIGO low mass searches.

In the high mass regime, we use the IMRPHENOMB waveform family [19], discussed in Sec. 1.5.2, which models IMR signals from BBH with aligned/anti-aligned spins. This waveform family is parametrized by two mass parameters and a single "effective spin" parameter χ_{eff} , defined in Eqn. 1.70. We performed two sets of IMRPHENOMB injections, a non-spinning set and a spinning set. Both were uniformly distributed in total mass between 25 and 100 M_{\odot} , and uniformly distributed in $q/(q + 1) \equiv m_1/M$ for a given M, between the limits $1 \leq q < 4$. In addition, the spinning injections were assigned aligned spin components χ_i uniformly distributed between -0.85and 0.85. To illustrate the effect of aligned spin on the search sensitivity, we plot in Fig. 4.5 the

⁵In particular, one should take note that no spinning IMR injections were performed in the low mass search, and in particular, none in the range $M_{\text{total}} \in [12, 25] M_{\odot}$, where merger effects are known to be significant. This mass regime may be very relevant for spin, and it would be a worthwhile to measure the sensitivity of the search pipeline to aligned spin signals in this mass range.

average sensitive distance over the S6/VSR2-3 observation time, in bins of total mass M, for both non-spinning simulated signals and for injections with $\chi < 0$ and $\chi > 0$, respectively.

Fig. 4.5 indicates higher sensitivity to positive- χ signals even with the current non-spinning templates, but also shows that the search is significantly less sensitive to negative- χ signals at higher values of total mass M. This result is expected because the intrinsic luminosity of a compact binary gravitational wave signal increases as the spins of the components increase. Interestingly, the search is *more* sensitive to non-spinning IMRPHENOMB signals than it is to non-spinning EOBNR signal, a rough indication of the level of disagreement between to the waveform models. As with the low mass case, these simulations only show how the actual search performed with respect to detecting spinning signals; it gives no hint as to whether the inclusion of spin effects in the templates could improve the sensitivity further (except inasmuch as you place weight on the SenseMon calculation, which seems to indicate that the search is already quite nearly optimal). In the following chapters, we will see that this is not the case. There is plenty of signal-to-noise available from the use of spinning templates, enough to offset the increased background incurred by using larger template banks.

Chapter 5

Towards Advanced LIGO Searches

Before we acquire great power, we must acquire wisdom to use it well.

Ralph Waldo Emerson

In the previous chapter, we reported on the results of two recent searches for gravitational waves from compact binary coalescence in the final initial LIGO (iLIGO) observational data [3, 4]. For the remainder of this work, we will focus on the projection of such analyses to aLIGO searches. We conduct two basic lines of inquiry: (i) how do the data analysis requirements scale from iLIGO to aLIGO and (ii) can the sensitivity of such searches benefit from the inclusion of additional physical effects. The first question is the subject of the present chapter.

Advanced LIGO instruments are expected to be sensitive to gravitational wave frequencies as low as $f \sim 10$ Hz. In order to recover CBC signals from such data with maximal significance – and in particular to recover otherwise sub-threshold events – the filter templates must extend down to these low frequencies. However, the full exploitation of the low frequency sensitivity leads to a number of technical problems which are not fully solved in the IHOPE implementation [5]. For example, the duration of the templates (which sets the analysis cost per filter) and the required number of templates are a very strong function of the low frequency cutoff. Furthermore, the increased length of the signal in the LIGO band presents the problem of how to detect a signal which may overlap with more than one stable lock stretch and how to estimate the background power spectral density (PSD) in data which are fully "contaminated" by signal.

One promising solution to these computational problems is to exploit the relatively simple frequency evolution of the CBC signal and the large redundancy between different filters in a template bank. While the signals we are looking for may be as long as an hour and extend up to very high frequencies, the signal spends most of its time within a very small bandwidth, allowing for the downsampling of the filters without significant loss of SNR. Additionally, template banks are by design highly redundant. As a result, an orthogonal decomposition of the template bank can greatly reduce the required number of templates. We filter with the non-physical orthogonal waveforms instead of the redundant physical templates and only reconstruct the SNR when an excess power is found within the orthogonal projection.

We have implemented these techniques in a new pipeline, known as GSTLAL [89,90], which is built upon a stream-based infrastructure known as GSTREAMER [91] and the same LAL library [92] upon which IHOPE is built. The stream-based infrastructure provides an elegant solution to a number of other computational problems, such as the ability to filter over gaps in the data, a crucial requirement for signals that exceed the length of a typical stable lock stretch. Furthermore, the use of streambased technology allows for a tunable latency of trigger generation, ideal for electromagnetic followup of mergers involving neutron stars, for which there may be bright electromagnetic counterparts [93].

In the following chapters, we use the GSTLAL pipeline described here to address the problem of extending the search parameter space to include spin.

5.1 Scaling CBC Searches from Initial LIGO to Advanced LIGO

As shown in Fig. 5.1, the aLIGO detector designs are predicted to have low frequency sensitivity well beyond both the iLIGO design and the actual achieved sensitivity of the detectors in S6. This larger bandwidth in turn implies that CBC signals spend more time in the LIGO sensitive band. Recall the expression we obtained in Eqn. 1.48 for the duration of a binary inspiral beginning from some initial separation, which we rewrite here in terms of the initial gravitational wave frequency:

$$t_{\rm coal} = \frac{5}{256} \frac{c^5}{G^2} \frac{(m_1 + m_2)^{1/3}}{m_1 m_2} (\pi f_{\rm low})^{-8/3}.$$
 (5.1)

We note that the coalescence time is a very strong function of the masses and initial frequency. Initial LIGO detectors had negligible sensitivity below 40 Hz. In this case, as seen in Fig. 5.2, even the longest filter down to the lowest masses $m_i \sim 1M_{\odot}$ would have been shorter than one minute.



Figure 5.1: Comparison of initial and advanced LIGO spectra. The bandwidth of aLIGO will extend beyond that of iLIGO in the low frequency regime. With this added bandwidth comes an increased computational scale for CBC searches in aLIGO. Even the early-stage detector design (operating at low power without a signal recycling mirror) indicates a low frequency sensitivity well beyond that of iLIGO and is in fact quite comparable to the late-stage detector design.

Early and late stage aLIGO detectors may have non-negligible sensitivity down to frequencies as low as $f \sim 10$ Hz. Extending the templates down to this frequency requires at the lowest masses filters nearly half an hour long.

Recall that to account for the non-stationary behavior of the detectors, we break the data into small segments which can be analyzed independently and during which the background can be considered approximately stationary. Estimating the background PSD requires using segments which are factors of ~ 15 longer than the signal itself; otherwise, the signal corrupts the PSD estimation. In our iLIGO analyses, we used a segment length of 2048 seconds and measured the PSD in sixteen partially overlapping sub-segments over the course of the larger 2048 second segment [63]. Given that the longest template for $f_{\rm low} = 40$ Hz is shorter than 60 seconds, any signal existing in the data would span no more than two of these measurements. We take the median of the measured PSDs to minimize the impact of the signal on measuring the background. If we apply this method of PSD estimation directly to an aLIGO search, we will need to have contiguous coincident segments that are nearly half a day long to accommodate waveforms as long as half an hour. Restricting an analysis only to segments of such a long duration would be disastrous for the search.

A natural question to ask while examining Fig. 5.2 is whether there are small sacrifices we can



Figure 5.2: Dependence of coalescence signal duration on the masses of the system. The duration of an inspiral template is a strong function the low frequency cutoff chosen for filtering. In the plots above, we see that templates exceeding several minutes and as long as half an hour will be required to filter LIGO data down to 10 Hz. We show the dependence of duration on mass for mass ratios q = 1 (left) and q = 4 (right). For fixed total mass, high mass ratio systems tend to be longer in duration compared to nearly equal mass systems. Spin effects can further increase the duration of the signal.



Figure 5.3: Signal-to-noise available in the low frequency band of aLIGO. For detectors with aLIGO sensitivities, a non-negligible amount of the SNR is contained in the band 10–40 Hz, a frequency range which was inaccessible to iLIGO observations. Above we show the fraction of the recovered SNR as a function of the low frequency cutoff of integration for three different noise curves. Each curve uses the same inspiral-only approximation for a binary neutron star signal with $m_1 = m_2 = 1.4M_{\odot}$. We see that for such a signal, as much as 50% of the available SNR may be contained in the low frequency band 10–40 Hz.



Figure 5.4: Number of templates required to cover a space as a function of low frequency cutoff. We show the number of templates required to fully cover the parameter space $1M_{\odot} \leq m_i \leq 24M_{\odot}$ and $2M_{\odot} \leq M_{\text{total}} \leq 25M_{\odot}$ at a minimal match of M = 0.97 for a late aLIGO sensitivity as we vary the low frequency cutoff. We see clearly that the required number of templates increases precipitously as the low frequency cutoff approaches $f_{\text{low}} = 10$ Hz. We have furthermore illustrated the concentration of templates as a function of total mass as we vary f_{low} , which shows that the plurality of templates is in the low mass range.

make in order to avoid using the long filters that go all the way down to $f_{\text{low}} = 10$ Hz. The figure shows that the long duration templates span only a small subset of the much larger parameter space. One should consider whether the loss of SNR from neglecting the low frequency content of the signal is sufficiently small to not warrant having to develop new data analysis solutions. To quantitatively understand the impact of the low frequency sensitivity for signal recovery, we must also fold in information about the detectors. In Fig. 5.3, we show the fractional accumulated SNR as a function of the low frequency cutoff and detector noise model for a BNS system with $m_i = 1.4M_{\odot}$. We see that for an aLIGO sensitivity as much as 20% of the available SNR is accumulated between 10 Hz and 20 Hz for both the early and late aLIGO models. In terms of detection rates, including low frequency content for the filter could yield up to a factor $(1/0.8)^3 \approx 2$ increase. We simply cannot afford to neglect the low frequency content of the signal. We must be able to handle filters that span up to half an hour in duration.

Long templates increase the computational cost of the search *per filter*. To compound on this situation, longer filters require many more templates to cover a given parameter space. As the waveforms get longer and longer, there are more cycles in band and a longer time over which two similar signals can dephase. In Fig. 5.4, we show the number of templates generated by lalapps_tmpltbank [92] for the same mass range covered by the low mass search presented in the last chapter as a function

106

of the low frequency cutoff (namely components in $m_i \in [1, 24] M_{\odot}$ and $M_{\text{total}} \in [2, 25] M_{\odot}$). We use for this figure a minimal match of M = 0.97. We see that moving from $f_{\text{low}} = 20$ Hz to $f_{\text{low}} = 10$ Hz increases the number of required templates by more than a factor of two. Relative to $f_{\text{low}} = 40$ Hz, the required number of templates increases by a factor of ten in going to $f_{\text{low}} = 10$ Hz.

It is worth pointing out the historical difference between detector design and actual performance. Note the curious feature in Fig. 5.3, in which there appears to be roughly 20% of the available SNR for an analysis with an iLIGO sensitivity in the region between 30 Hz and 40 Hz. However, our most recent searches only used filters down to 40 Hz. The reason for this difference is that the initial detectors never reached design sensitivity below 40 Hz, as one can see in Fig. 5.1. Thus, in the actual pipeline analyses discussed in the previous chapter, there was very little sensitivity lost by choosing a 40 Hz low frequency cutoff instead of a 30 Hz cutoff. It is possible that it will take some time for aLIGO detectors to reach their design sensitivity at the lowest frequencies, and that some of the problems we consider here will be less relevant.

All these considerations pertain only to the problem of developing an analysis strategy that does essentially the same kind of analysis as done for iLIGO. We have already incurred two factors of ten; one from the increased cost of filtering longer templates, and one from having to use more filters to cover the parameter space. However, we also want to take the pipeline beyond what we have done in the past. In particular, we want to see if we can improve the sensitivity of the pipeline to spinning signals with the use of spinning templates. As we have discussed, spin in compact binary systems is expected to be significant, but our iLIGO analyses used non-spinning templates, as no known methods were available to profitably extend the parameter space to include spin. Including more physical parameters in our search templates will also increase the computational scale of the analysis. Even if we consider only aligned spins, the inclusion of spin in the template banks can further increase the required number of templates by another factor of ten (see Chap. 6).

5.2 Multibanding and down-sampling

We have just seen that the full exploitation of the aLIGO detector sensitivity for CBC searches requires a computational scale which is several orders of magnitude greater than our previous searches on iLIGO data. Fortunately, the morphology of inspiral signals can be leveraged to significantly offset some of the computational cost of filtering. Our first trick is to realize that compact binaries spend most of their lives at low frequencies. The Nyquist theorem states that a signal which is band limited to $0 \leq f \leq f_{high}$ can be fully reconstructed, provided that the signal is sampled at the critical frequency $f_s = 2f_{high}$ or higher. On the other hand, the merger frequency for low mass compact binaries, which traditionally sets the analysis sampling rate, is usually in the kHz range. While treating the signal as a whole, we therefore require sampling rates of the order a few kHz, but we are actually over-sampling the signal throughout most of its duration. Thus, one approach to mitigating the computational cost for filtering is to break the signal into contiguous band-limited time intervals and process each interval separately at possibly lower sample rates, as depicted in Fig. 5.5. By the Nyquist theorem, one can fully reconstruct the whole SNR time series that one would have obtained by filtering the whole signal at a single, higher sampling rate.

To make the above statements more quantitative, we return to Eqn. 5.1, which gives the timefrequency relation for a compact binary inspiral at Newtonian order. Higher order post-Newtonian terms, including those from spin and merger effects, will significantly alter the time-frequency relation, which should be accounted for in an actual analysis, but here we are simply trying to set the scale. We see from Eqn. 5.1 that an inspiralling binary spends a time

$$\frac{\Delta t}{t_1} = 1 - \left(\frac{f_1}{f_2}\right)^{8/3},\tag{5.2}$$

emitting gravitational waves in the frequency band $[f_1, f_2]$. Here t_1 is the time to coalescence from initial frequency f_1 and Δt is the time it takes for the binary to evolve from f_1 to f_2 . Of common practical interest is the case $f_2 = 2f_1$; since Fourier transforms are most efficiently executed on arrays having power-of-two lengths, we typically only ever change the sampling rate by factors of two. Putting this numerical value into Eqn. 5.2, we find

$$\frac{\Delta t}{t_1} \approx 0.84. \tag{5.3}$$

This result says that the signal spends about 84% of its *remaining* lifetime in the band $[f_1, 2f_1]$. In particular, a signal which lasts for 30 minutes starting from 10 Hz will spend only five minutes above 20 Hz, and only 45 seconds above 40 Hz. To filter such a signal, we can safely sample all but the last minute at $f_s \sim 100$ Hz, typically at least a factor of ten smaller than the usual sampling rate.

Our strategy is to reduce the filtering cost of a large fraction of the waveform by computing



Figure 5.5: Computational benefits of multibanding and the singular value decomposition. We depict (left) the decomposition of an inspiralling compact binary signal into separate frequency bands in which the signal can be sampled at lower rates than would be required for the signal as a whole. This particular signal lasts ~1100 s, reaching at coalescence gravitational wave frequencies of $f \sim 1$ kHz. However, all but the last 12.5 seconds of the signal can be sampled at $f_s \leq 128$ Hz without aliasing. Decomposing a template bank of 1314 templates sorted by chirp mass, including the example signal, we indicate (right) the number of basis vectors required to achieve a specified level of SVD tolerance. Even for a SVD tolerance of 10^{-6} , corresponding to the right most column, the SVD decomposition reduces the number of filter templates by more than an order of magnitude.

part of the convolution at a lower sample rate. Similar techniques were successfully implemented in iLIGO for the purpose of achieving low latency trigger generation and exercising the procedures for electromagnetic followup [94]. One example is MBTA [59,95], which was deployed in S6/VSR3. MBTA consists of multiple, usually two, template banks for different frequency bands, one which is matched to the early inspiral and the other which is matched to the late inspiral. An excursion in the output of any filter bank triggers coherent reconstruction of the full matched filtered output. Final triggers are built from the reconstructed matched filter output.

In the GSTLAL implementation, we divide the templates into time slices in a time-domain analog to the frequency-domain decomposition employed by MBTA. We consider a bank of filters $B = \{h_i\}$ expressed in the time domain and sampled initially at sufficiently high frequency f_0 such that there is no aliasing in any template in the bank. The templates in the bank are all zero-padded as necessary so that they have the same number of samples N. We trivially decompose each template h_i into a sum of S non-overlapping templates h_i^s such that

$$h_i[k] = \sum_{s=0}^{S-1} \begin{cases} h_i^s[k] & \text{if } t_s \leqslant k/f_0 < t_{s+1} \\ 0 & \text{otherwise} \end{cases}$$
(5.4)

for S integers $\{f_0t_s\}$ such that $0 = f_0t_0 < f_0t_1 < \cdots < f_0t_S = N$. We filter the data separately with

each time-sliced filter, and the outputs of these new time-sliced filters form an ensemble of partial SNR streams. By linearity of the filtering process, these partial SNR streams can be summed to reproduce the SNR of the full template.

Since waveforms with neighboring intrinsic source parameters have similar time-frequency evolution, it is possible to design computationally efficient time slices for an extended region of parameter space rather than having to design different time slices for each template. We construct from our time-sliced filters S sub-banks $B_s = \{h_i^s\}_i$, and within each sub-bank we choose time slice boundaries with the smallest power-of-two sample rates that sub-critically sample all time-sliced templates in that bank. The time slices consist of the S intervals $[t_0, t_1)$, $[t_1, t_2)$, ..., $[t_{S-1}, t_S)$, sampled at frequencies $f_0, f_1, \ldots, f_{S-1}$, where f_s is at least twice the highest nonzero frequency component of any filter in the bank B_s for the *s*th time slice. The time-sliced templates can then be downsampled in each interval without aliasing, so we define them as

$$h_i^s[k] \equiv \begin{cases} h_i \left[k \frac{f}{f_s} \right] & \text{if } t_s \leqslant k/f_s < t_{s+1} \\ 0 & \text{otherwise.} \end{cases}$$
(5.5)

We note that the time slice decomposition in Eqn. 5.4 is manifestly orthogonal since the time slices are disjoint in time. In the next section, we examine how to reduce the number of filters within each time slice via the singular value decomposition (SVD) of the time-sliced templates.

In the case where the time-frequency relationship is not known precisely, as for example during merger, we can still apply this multibanding prescription. We treat these cases as we have done here, except that we err on the side of over sampling.

5.3 Orthogonal template banks

Templates banks used for CBC searches are by design highly correlated. The redundancy is required to avoid missing signals "in between" any two templates when covering a continuous parameter space with a finite sized bank. Given this fact, it is natural to consider whether the computational cost for this filtering strategy may be improved by choosing a different set of filters with reduced overlaps. In Ref. [96], they showed that applying a SVD to inspiral template banks greatly reduces the number of filters required to achieve a particular minimal match. A similar technique can be applied to the time-sliced templates defined in Eqn. 5.5 above. Any $n \times m$ rectangular matrix M can be uniquely decomposed into a unitary $n \times n$ mapping U, a unitary $m \times m$ mapping V, and an $n \times m$ diagonal scaling map Σ such that

$$M = U\Sigma V^{\dagger},\tag{5.6}$$

where [†] denotes the conjugate transpose and Σ has all non-negative values. This decomposition is known as the *singular value decomposition*. An equivalent expression for the SVD is given by

$$M = \sum_{i=1}^{n} \sigma_i |u_i\rangle \langle v_i|, \qquad (5.7)$$

where $|\rangle\langle|$ is the vector outer product, σ_i are called the singular values, and u_i and v_i are the normalized (left and right) eigenvectors of U and V, respectively. The first representation is most convenient for the actual software implementation of the SVD, whereas the second is useful for understanding the theoretical properties of the expansion and why it works for GW data analysis. If the singular values are ordered so that $\sigma_i \geq \sigma_{i+1}$, then this decomposition is unique.

Now we apply the SVD to a template bank. Let M be a matrix whose rows correspond to the whitened, normalized time domain representation of templates in a bank. As usual, the SNR for a normalized template \hat{h} in the frequency domain is given by

$$\rho = 4 \operatorname{Re} \int_0^\infty \frac{\hat{h}_i^*(f)\tilde{s}(f)}{S_n(f)} df.$$
(5.8)

Note that with this expression for the SNR, we have not maximized over coalescence time or phase. For any time series s(t) and power spectrum $S_n(f)$, we define the whitehed time series $s_W(t)$ by

$$s_W(t) = \int_0^\infty \frac{\tilde{s}(f)}{\sqrt{S_n(f)}} e^{2\pi i f t} df.$$
(5.9)

The whitened time series is the inverse Fourier transform of the signal Fourier transform, after the latter has been rescaled by the noise amplitude spectral density. Note that for white noise, $S_n(f)$ is a constant and the whitening process has no effect, as one might expect.

We now use Parseval's theorem to move back to the time domain representations of the filters. According to our definitions

$$\rho = 4\text{Re } \int_0^\infty \hat{h}_W^*(f)\tilde{s}_W(f)df, \qquad (5.10)$$

which by Parseval's theorem implies

$$\rho = 4\text{Re } \int_0^\infty \hat{h}_W(t) s_W(t) dt.$$
(5.11)

Since the matrix M has as rows the time-domain whitehed templates, the SNR for all templates may be computed with the single matrix operation

$$\rho = 4Ms_W,\tag{5.12}$$

noting that both the whitened data s_W and the matrix elements of M are real. Now we see the power of the SVD. We carry out the SVD for M, and then truncate the expansion in Eqn. 5.7 to obtain an *approximate* expression for the SNR,

$$\rho_L \approx \sum_{i=1}^L \sigma_i |u_i\rangle \langle v_i | s \rangle, \tag{5.13}$$

where $L \leq n$ is the truncation order. The quantity ρ_L is a column vector, whose rows correspond to the signal-to-noise for each template in the bank before maximization over coalescence time and phase.

By filtering with the orthogonal templates and allowing for a small loss in SNR, we significantly reduce the computational cost. The problem is to achieve $L \ll n$ while still having $\rho_L \approx \rho$. The crucial question therefore is how much SNR is lost in making the approximation in Eqn. 5.13. A straightforward calculation shows that the expected loss of SNR due to truncation of the SVD expansion is given by

$$(\Delta \rho)^2 = \left\langle \left| \rho - \rho_L \right|^2 \right\rangle = \sum_{i=L+1}^n \sigma_i^2.$$
(5.14)

The normalized SNR loss $\Delta \rho / \rho$ is referred to as the *SVD tolerance* and is clearly determined by the number of basis templates that are kept in the approximation. We refer to the inner products

$$\rho_{\perp,i} = \sigma_i \langle v_i | s \rangle \tag{5.15}$$

as the *orthogonal* or *partial* SNRs. We note in particular that these SNRs are non-physical. In order to detect real signals, we need to be able to reconstruct the original SNR. In terms of the orthogonal



Figure 5.6: Illustration of the transformation from correlated physical templates to orthogonal unphysical templates. Template banks for the detection of GW signals from CBCs are highly redundant by design. By applying a singular value decomposition to the template waveforms, one can pick out the dominant contributing pieces and filter only with these. In doing so, we change from filtering with the physical templates (left) and unphysical templates (right). With this method, the number of required filters can be reduced by an order of magnitude or more with a negligible loss of SNR.

SNRs, the reconstructed SNR is given by

$$\rho_L = \sum_{i=1}^L \rho_{\perp,i} |u_i\rangle, \tag{5.16}$$

which follows from Eqn. 5.13. Since we have not maximized over phase, we must include in M two filters for each set of intrinsic parameters: one filter with $\varphi_0 = 0$ and one with $\varphi_0 = \pi/4$. Then combining the SNRs for these two filters in quadrature will yield the maximization over coalescence phase.

As shown in Fig. 5.5 (right), the SVD can yield greater than an order of magnitude reduction in the required number of filters, even at the strictest of SVD tolerance. The key to the efficacy of this technique is that we fill each matrix M with templates that are very nearby in parameter space, maximizing the redundancy among them. A graphical illustration of the singular value decomposition applied to inspiral signals is shown in Fig. 5.6. The orthogonal templates shown in Fig. 5.6b are sorted in order of decreasing singular values (top to bottom). We see clearly that the orthogonal templates with the smallest singular values bear almost no resemblance to the original templates, an indication of the negligible contribution to the signal-to-noise from this part of the



Figure 5.7: Architecture of the GSTLAL analysis pipeline. The above flow chart depicts the workflow of the GSTLAL pipeline for CBC searches. First one constructs the template bank and performs the time-slicing decomposition discussed above, choosing the optimal sampling frequency for each slice. The data is read from disc, whitened and split into a number of parallel streams. Each stream corresponds to one of the time slices, and the data in the stream are down-sampled to match the filters. Filtering the data with the orthogonal SVD templates, we create an equal number of SNR streams. If the orthogonal filters show a large excursion, then the physical SNR is reconstructed from the partial streams. GSTREAMER handles the synchronization and up-sampling seamlessly such that the SNR may be added together across time slices to obtain the full SNR time series. Finally, the reconstructed SNR time series is further analyzed (clustered, checked for coincidence, etc.) to obtain a trigger.

signal space.

5.4 gstlal: A stream based approach to gravitational wave searches

We have described two methods of reducing the computational scale of aLIGO analyses for CBC signals. We still require a solution for the problem of having signals which may span more than one lock stretch. For this, we implement the techniques described above around a stream-based technology known as GSTREAMER [91], a widely-used multimedia library which is highly modular and naturally suited for processing signals over gaps, such as would be necessary for viewing multimedia over a flaky internet connection. By filtering over gaps, the pipeline by design does not require

segments of any particular length, allowing for a continuous measurement of the PSD. Computing a running average of the PSD over the last several hours, the GSTLAL pipeline also manages to avoid contamination of the PSD estimate by long duration signals. Additionally, GSTREAMER is a multi-threaded framework, and easily allows for the parallelization of independent computations. The multi-threaded aspect of GSTREAMER requires robust data stream synchronization, so that one stream can be buffered while waiting for another data stream to be analyzed.

A GSTREAMER pipeline consists of a number of "elements" connected together by "pads". A pad can either be a *source*, which emits data, or a *sink*, which accepts data. All elements are set up with a single clock which can be used to synchronize data streams, which otherwise may process independently. In Fig. 5.7, we draw a schematic of the basic elements of the GSTLAL pipeline, the connections between them, and where in the pipeline the multibanding and SVD techniques are applied. The GSTLAL pipeline starts with a frame reading element, which finds the data on disk (or from a real-time broadcasting source) and stores it in memory. The data from this element are then passed a whitening element, which computes the power spectrum and whitens the data by that spectrum. From the whitener, the data is split into parallel streams to be processed separately. Each data stream is down sampled according to the portion of the bank that will be used to filter the bank. The whitened data are then passed to the filtering element, which uses the orthogonal filters obtained from the SVD (which has been computed beforehand along with the template bank), outputting a time series of partial SNR. Since the filter templates are non-physical, we must reconstruct the physical SNRs to interpret the results, but to maintain the reduced computational cost earned with the SVD we must also only do this if we think there is a good chance for a potential signal in the data. This is achieved with a conditional statistic, which applies to the orthogonal SNRs and triggers reconstruction only when necessary. If reconstruction is triggered, then the partial SNR time series is passed to upsampling elements, added together with appropriate time delays, and written to disk (and possibly sent as an alert to other astronomers for electromagnetic followup).

5.5 Low latency trigger generation

As mentioned briefly in Sec. 1.4, compact binary coalescence is a plausible progenitor for most short γ -ray bursts [97, 98] (GRBs). However, the association is not iron-clad [99]. It is thought that following the merger of a neutron star with another compact object tidally disrupted neutron star material falls onto the newly formed, rapidly spinning compact object. The material is accelerated



Figure 5.8: Possible electromagnetic counterparts to compact binary mergers involving matter. Electromagnetic counterparts to CBC events come in many radiation bands on many timescales. Being able to analyze the data and generate candidate events in low-latency will be key to making the connection between the gravitational wave and electromagnetic spectrum and provide fundamental insight into the dynamics of the post-merger. Figure credit [93].

in jets along the spin axis with a timescale of 0.1–1 s after the merger [100], matching the short GRB duration distribution well. Prompt electromagnetic emission, including the γ -ray burst itself, can arise as fast outflowing matter collides with slower matter ejected earlier in inner shocks. The same inner shocks, or potentially reverse shocks, can produce an accompanying optical flash [101]. These and other potential electromagnetic counterparts to a compact binary merger are depicted in Fig. 5.8.

Because the gstlal pipeline framework naturally handles the added complexities associated with longer waveforms in the wider aLIGO band, the GSTLAL framework is a promising solution for low latency trigger generation in aLIGO. By breaking the templates up into time slices, the pipeline necessarily computes the signal-to-noise for a template in pieces. Thus, if enough signal-to-noise can be accumulated in the early phases of the inspiral, GSTLAL can trigger a followup before analyzing the entire template. Given a loud enough signal, the pipeline could be used to point telescopes *before* the merger even happens and catch the prompt electromagnetic emission [89].

The prompt emission is a probe into the extreme initial conditions of the outflow in contrast with afterglows, which arise in the external shock with the local medium and are relatively insensitive to initial conditions. Optical flashes have been observed for a handful of long GRBs [102] by telescopes with extremely rapid response or, in the case of GRB 080319b, by pure serendipity, where several telescopes were already observing the afterglow of another GRB in the same field of view [103]. The observed optical flashes peaked within tens of seconds and decayed quickly. For short GRB energy balance and plasma density, however, the reverse shock model predicts a peak flux in radio, approximately 20 minutes after the GRB, but also a relatively faint optical flash [98]; for a onceper-year Advanced LIGO event at 130 Mpc, the radio flux will peak around 9 GHz at ~5 mJy, with emission in the *R*-band at ~19 mag. Interestingly, roughly a quarter to half of the observed short GRBs also exhibit extended X-ray emission of 30–100 s in duration beginning ~10 s after the GRBs and carrying comparable fluence to the initial outburst. This can be explained if the merger results in the formation of a proto-magnetar that interacts with ejecta [104]. Rapid GW alerts would enable joint EM and GW observations to confirm the short GRB-CBC link and allow the early EM observation of exceptionally nearby and thus bright events.

In S6/VSR3, several all-sky detection pipelines operated in a low-latency configuration to send astronomical alerts [105–108], which were sent only when all three detectors (H1, L1, and V1) were operating. Among the pipelines, the MBTA pipeline (mentioned briefly in Sec. 5.2) achieved the best latency, delivering candidate gravitational wave events within five minutes of the event. Alerts were sent with latencies of 30–60 minutes, dominated by human vetting. Candidates were sent for EM follow-up to several telescopes; *Swift*, LOFAR, ROTSE, TAROT, QUEST, SkyMapper, Liverpool Telescope, Pi of the Sky, Zadko, and Palomar Transient Factory [105, 109] imaged some of the most likely sky locations.

There were a number of sources of latency associated with the search for CBC signals in S6/VSR3 [105], listed here.

Data acquisition and aggregation ($\gtrsim 100 \text{ ms}$) The LIGO data acquisition system collects data from detector subsystems 16 times a second [110]. Data are also copied from all of the GW observatories to the analysis clusters over the Internet, which is capable of high bandwidth but only modest latency. Together, these introduce a latency of $\gtrsim 100 \text{ ms}$. These technical sources of latency could be reduced with significant engineering and capital investments, but they are minor compared to any of the other sources of latency.

Data conditioning ($\sim 1 \text{ min}$) Science data must be calibrated using the detector's frequency response to gravitational radiation. Currently, data are calibrated in blocks of 16 s. Within $\sim 1 \text{ minute}$, data quality is assessed in order to create veto flags. These are both technical sources of latency that might be addressed with improved calibration and data quality software for advanced detectors.

Trigger generation (2–5 min) Low-latency data analysis pipelines deployed in S6/VSR3 achieved an impressive latency of minutes. However, second to the human vetting process, this dominated the latency of the entire EM follow-up process. Even if no other sources of latency existed, this trigger generation latency is too long to catch prompt or even extended emission. Low-latency trigger generation will become more challenging with advanced detectors, because inspiral signals will stay in band up to ten times longer.

Alert generation (2–3 min) S6/VSR3 saw the introduction of low-latency astronomical alerts, which required gathering event parameters and sky localization from the various online analyses, downselecting the events, and calculating telescope pointings. If other sources of latency improve, the technical latency associated with this infrastructure could dominate, so work should be done to improve it.

Human validation (10–20 min) Because the new alert system was commissioned during S6/VSR3, all alerts were subjected to quality control checks by human operators before they were disseminated. This was by far the largest source of latency during S6/VSR3. Hopefully, confidence in the system will grow to the point where no human intervention is necessary before alerts are sent.

The remainder of our work here deals primarily with binary black holes mergers, for which there is no well-understood mechanism for producing an electromagnetic counterpart. However, we will discuss the importance of spin for binary neutron star and neutron star-black hole mergers in aLIGO in Sec. 6.3; the conclusions reached there could have implications for the efficiency of searches for these systems in aLIGO. We mention the low latency capabilities of the GSTLAL pipeline here as a reminder of this potential application of GSTLAL and its importance for constructing a detailed understanding of our sources.

Chapter 6

The Importance of Spin for the Detection of Gravitational Waves from Compact Binary Coalescences

It is reasonable to hope that in the not too distance future we shall be competent [enough] to understand so simple a thing as a star.

> Arthur Eddington in *The Internal Constitution* of Stars (1926)

As we have emphasized throughout, the most recent searches for CBCs in iLIGO data [3, 4] have used non-spinning templates. Phenomenological models for generically spinning binaries in the inspiral phase have been constructed and implemented as search templates in an analysis of data from LIGO's third science run [49], but later studies showed that the spinning templates pulled as much noise as signal, leading to no net gain in search sensitivity relative to a non-spinning search [50]. At the time of S5, complete inspiral-merger-ringdown (IMR) waveforms suitable for CBC searches were just coming onto the scene [111], and it was already a huge advance in our search methods to use these waveforms for the high mass search [48]. By LIGO's sixth science run (S6), IMR waveforms with leading order aligned spin effects began to appear [19, 112], but the tools were not available to use these models as search templates. Experience had shown that the inclusion of spin effects in search filters will not work without techniques for managing the elevated background. Consequently, these spinning models were not used as search templates in the analysis of S6 data.

In Sec. 1.5, we met two models for gravitational waves from compact binaries with aligned spins. These models are promising for data analysis applications in part because they reduce the large mass and spin parameter space to just three parameters. We expect that these models will capture some, but not all, of the spin effects in a generically spinning system. Furthermore, depending on the formation mechanism ones favors, most of our sources may have spins that are very nearly aligned with the orbital angular momentum, and we expect these waveforms to cover this class of signals very well. With these reduced dimensionality models in hand, we reassess here the importance of spin effects for the detectability of compact binary coalescence signals and the viability of using these templates to improve the search sensitivity.

In this chapter, we focus on the *signal-to-noise* that may be gained from including spin effects. In practice, in quantifying the performance of a search pipeline, we must take into account the production of false triggers due to random overlap between the templates and detector noise. The inclusion of extra physical parameters in the template waveforms typically requires one to analyze the data with many more templates. Each additional template picks up a different portion of the background and increases the chance for a false trigger. Thus, the gain in SNR achieved by the use of more accurate templates may be offset by the elevation in the background rate. This interplay between SNR and false alarm rate (FAR) is precisely the reason that precessing templates were previously found to be ineffective. The final element of this thesis will be the demonstration of an analysis pipeline which effectively uses additional physical effects (spin) in the search templates to improve the search sensitively, even with consideration of the elevated background.

6.1 Astrophysical expectations for spin in compact binaries

If compact binaries have significant spins, then the inclusion the effects of spin in the template waveforms could significantly impact the detection rate of such systems. Conversely, by determining the spin distribution of black holes in binaries, we can better understand the formation mechanisms which give rise to these systems. Indeed, the spin of components of the binaries is strongly coupled to the formation history of the binary, as discussed in Sec. 1.3. Regardless of the formation mechanism, however, most theoretical models predict some amount of spin in the resulting components. In this section, we review what is known or expected about the spins in compact binaries.

We consider first an order of magnitude estimate for the spin of neutron stars based on observational considerations. Recall that the dimensionless spin χ is related to the spin angular momentum J by

$$\chi = \frac{J}{Gm^2/c},\tag{6.1}$$

where m is the mass of the spinning body. The moment of inertia for a uniform density spherical body is given by $I = \frac{2}{5}mR^2$, where R is the radius of the body. For such a body, rotating with frequency f, its dimensionless spin is given by

$$\chi = \frac{I\omega}{Gm^2/c} = \frac{2}{5} \frac{c}{G} \frac{R^2}{m} 2\pi f.$$
(6.2)

This expression allows us to set the scale for neutron star spins. Consider a millisecond pulsar rotating with frequency f = 50Hz¹. Taking $m_{\rm NS} = 1.5 M_{\odot}$ and $R_{\rm NS} = 10$ km, we obtain

$$\chi_{NS} \approx 0.02 \tag{6.3}$$

as a baseline estimate for the dimensionless spin of a neutron star. Even such a small spin magnitude, as we shall see, will be non-negligible for CBC searches in aLIGO. The spin of the neutron may exceed this value, but arguments based on the tidal disruption of a neutron star generally agree that the spin cannot exceed $\chi \sim 0.7$. The exact value for this upper limit depends on the neutron equation of state which is currently not well-constrained. Softer equations of state support larger spins, since the neutron star is more compact.

The distribution of spins in black hole binaries in not known, but measurements accreting supermassive black holes and black holes in binaries with luminous, massive stars tend to indicate that the (dimensionless) spin magnitudes of stellar-mass black holes are believed to be close to unity. Some observational determinations yield values around or exceeding 0.9 [113–115]; however, see [116–118], where smaller values were obtained. Note that such measurements apply to black holes in X-ray binaries, which may not be representative of spins in BBH systems. X-ray observations of the spins of accreting black holes in binary systems, while technically challenging, indicate a fairly uniform distribution of component spins over the entire range $0 \le a \equiv S/m^2 \le 1$ [114, 115, 119–123]. Note that such measurements apply to black holes in X-ray binaries, which may not be representative of spins in BBH systems. Indications that spin-orbit misalignment in field binaries may be small come from observations of the microquasar XTE J1550-564 [124], and from population synthesis models of

¹Note that the pulsar PSR J0737-3039A, one of the two observed pulsars in the associated relativistic binary, has a spin frequency of $f \approx 50$ Hz, which motivates this particular choice. The other pulsar, J0737-3039B, has a spin frequency of only $f \approx 0.4$ Hz. This binary system is the *only* relativistic compact binary with spin measurements of both components.

Fragos et al. [125]. For dynamically formed binaries, however, the component spins may be largely independent of each other and of the orbital parameters.

6.2 SBank: a generic infrastructure for template bank construction

The first problem we address is that of constructing template banks with efficiently spaced templates that minimize the loss of SNR arising from the discreteness of the template bank. Recent LIGO and Virgo compact binary searches, which used non-spinning templates, have relied on the lattice placement technique described in Sec. 3.2.3, which is highly computationally efficient and nearly optimal for non-spinning inspiral signals. However, this technique relies on having knowledge of certain special parameters in which to construct the lattice. In larger parameter spaces, or when using waveforms which include effects beyond the inspiral portion of the coalescence, we often do not have any known special parameters in which to form the lattice and other approaches to template placement become necessary.

On the other hand, stochastic placement techniques (described in Sec. 3.2.4) are applicable to a wide variety of waveform approximants, requiring no prior knowledge of special parameters and extending straight-forwardly to higher dimensions. We have implemented the random placement algorithm described in Sec. 3.2.4 in a completely generic infrastructure called SBANK, which is independent of the particular details of the waveform family one is using. Waveforms are constantly under development for compact binary searches, and since SBANK is designed to be ignorant of the waveform details, one can immediately begin constructing banks with SBANK once your model is coded up. Here we will see applications of this code base to many different waveform approximations, some in the time domain – some in the frequency domain, some with metrics and some without. Our templates will be described by three parameters, but the infrastructure supports an arbitrary number of parameters (and in fact our bank simulations make use of this feature).

For practical applications, we applied a few heuristics to improve the speed of SBANK. First, we test each proposal against not the whole bank, but only its neighborhood of templates, defined by some fractional difference in τ_0 , the coordinate that is best fractionally measured; we assume that any template far away from the proposal in τ_0 cannot have a high match. Furthermore, so that we find the high matches even sooner, we evaluate matches in the order of increasing τ_0 difference Table 6.1: Reduced spin template bank parameters. We list the parameters used for generating the reduced spin template bank described in the text. The non-spinning template bank is generated with the same mass parameters, but with all spins set to zero. The spin limits for black holes and neutron stars are different, corresponding to the different astrophysical expectations for the spins of these bodies. In constructing the aligned spin template bank, we place cuts on the individual component spins, which are astrophysically motivated, rather than placing cuts on the reduced spin. We consider a neutron star to be a body with mass $m_i \leq 2 M_{\odot}$ and a black hole to be a body with mass $m_i > 2 M_{\odot}$.

Bank parameter	Value
Template waveform	TAYLORF2REDSPIN
Noise PSD model	ALIGOZERODETHIGHPOWER
Low-frequency cutoff: f_{low}	$20\mathrm{Hz}$
Component mass: m_1, m_2	$[1,20]M_{\odot}$
Total mass: m	$[2,21]M_{\odot}$
NS spin: χ_i	[-0.4, 0.4]
BH spin χ_i	[-0.98, 0.98]
Minimum match: \mathcal{M}_{\min}	0.95

between the proposal and the bank seed. We stop computing further matches immediately upon finding a match greater than the target minimum match. The sorting operation is costly, but this is outweighed by the reduction in the required number of match calculations.

Another technique we use is to draw proposals uniformly in (τ_0, τ_3) space. As the true template density is proportional to $\sqrt{|\mathbf{g}|}$ and \mathbf{g} is slowly varying in these coordinates, this reduces the number of proposals thrown at already over-tested regions of parameter space and puts them in undertested regions. More generally, it is straightforward to drop and replace custom proposal functions in SBANK, should we find a better set of coordinates in which to draw the proposals. Finally, the availability of the metric gives an analytic approximation to the mismatch, which significantly speeds up each iteration of this algorithm, but it is not strictly necessary. SBANK is available for use in the LALSUITE data analysis package [92]. SBANK is implemented primarily in the Python programming language with speed-critical components in C.

6.3 An aligned spin template bank for the detection of generically spinning compact binary inspirals in advanced LIGO

Here we demonstrate the construction of a three-dimensional aligned spin template bank suitable for the search of gravitational waves from inspiralling compact binaries in aLIGO. We show that although the templates model only leading order *aligned* spin effects, the template bank is highly effective for detecting signals from *generically* spinning binaries. The high effectiveness of the aligned

Table 6.2: Precessing binary population simulation parameters. We indicate the source parameter distributions chosen for the efficiency studies of our TAYLORF2REDSPIN template banks described in the text. A binary component is deemed a neutron star (NS) if its mass is $\leq 2 M_{\odot}$, and is deemed a black hole (BH) if its mass is $> 2 M_{\odot}$. We note that all simulated signals are placed at the same luminosity distance to take into account inherent selection effects in the source population.

Simulation parameter	Value
Waveform approximant	SpinTaylorT5
BH spin magnitudes: $ \chi_i $	uniform(0, 0.98)
NS spin magnitudes: $ \chi_i $	uniform(0, 0.4)
Cosine of spin orientations: $\mathbf{L}_{\mathrm{N}}^{\mathrm{ini}} \cdot \chi_{i}^{\mathrm{ini}}$	uniform(-1, 1)
Cosine of sky location (polar): $\cos \theta$	uniform $(-1, 1)$
Sky location (azimuth): ϕ	uniform $(0, 2\pi)$
Cosine of inclination angle: $\cos \iota$	uniform(0, 1)
Polarization angle : ψ	uniform $(0, 2\pi)$
Luminosity distance: d_L	$1 { m Mpc}$
Noise PSD model	ALIGOZERODETHIGHPOWER
Low-frequency cutoff: f_{low}	$20\mathrm{Hz}$

spin bank arises in part from selection effects inherent to our sources by which the systems with the greatest precession effects are also intrinsically dim.

The templates in our bank are the aligned spin TAYLORF2REDSPIN inspiral waveforms described in Sec. 1.5.1. The template parameter space consists of two dimensions describing the binary masses and a single reduced-spin parameter

$$\chi_{\rm red} \equiv \chi_{\rm s} + \delta \chi_{\rm a} - \frac{76\eta}{113} \chi_{\rm s}, \tag{6.4}$$

where

$$\chi_{\rm s} \equiv \frac{1}{2} \left(\vec{\chi}_1 + \vec{\chi}_2 \right) \cdot \hat{L} \tag{6.5}$$

$$\chi_{\rm a} \equiv \frac{1}{2} \left(\vec{\chi}_1 - \vec{\chi}_2 \right) \cdot \hat{L} \tag{6.6}$$

are the symmetric and anti-symmetric combinations of the component spins and $\delta \equiv (m_1 - m_2)/m_1$ is the asymmetric mass ratio. The range of parameters used for the construction of our template banks is listed in Tbl. 6.1. We construct both aligned spin and non-spinning template banks with SBANK from the TAYLORF2REDSPIN waveform model, making use of a fast semi-analytic approximation to the mismatch function. To quantify the relative signal recovery by these banks, we simulate a population of generically spinning compact binaries and compute the fitting factors (defined in Sec. 3.2.2) of the two template banks to these sources.

Though the templates in this bank model only the leading order effects of aligned spins, we



Figure 6.1: *Effective fitting factors towards a simulated population of precessing binaries.* We show the achieved effective fitting factors (color bar) for our two banks constructed here towards signals from generically spinning systems computing in the SPINTAYLORT5 approximation. The left and right panels correspond to the non-spinning and reduced spin template banks, respectively.



Figure 6.2: Increase in detection volume at fixed signal-to-noise. We show the expected relative increase in sensitive volume of a search employing the reduced-spin template bank as compared to one employing a non-spinning template bank (corresponding to a fixed SNR threshold). The reduced spin template bank is expected to bring about a $\sim 20-52\%$ increase in the average detection volume, assuming that the maximum spin of neutron stars is 0.4.

find that the template bank is efficient for the detection of generically spinning binaries due (at least partly) to inherent population selection effects. The intrinsic luminosity of the target binary, as well as the fitting factor of the templates, depends not only on the masses and spins, but also on the parameters describing the location and orientation of the target binary. For example, the modulational effects of precession are the highest for binaries highly inclined with respect to the detector, while the intrinsic luminosity of such binaries is lower (as compared to binaries which are nearly "face on"). Thus, highly inclined binaries (which show the largest modulational effects of precession) are intrinsically less likely to be observed as compared to binaries that are face-on.

In order to take into account such selection effects in evaluating the coverage of the template bank, we average the fitting factor in our simulations over the population. The average is weighted by the intrinsic luminosity of the source at fixed distance and called the *effective fitting factor* FF_{eff}

$$FF_{eff} = \left(\frac{\overline{\rho_{bank}^3}}{\overline{\rho^3}}\right)^{1/3},\tag{6.7}$$

where $\rho \equiv \sqrt{\langle h^{\text{targ}} | h^{\text{targ}} \rangle}$ is the expected SNR for the binary signal in the detector. Our template bank recovers only a portion FF of this SNR, where FF is the signal's fitting factor against the bank; therefore, by definition of fitting factor, we have $\rho_{\text{bank}} \equiv \rho$ FF. The bars in Eqn. 6.7 indicate ensemble averages over the full parameter space (while keeping the component masses fixed) with the parameter distributions given in Tbl. 6.2. The effective fitting factor FF_{eff} describes average detection range by a suboptimal template bank as a fraction of the detection range using an optimal template bank. The corresponding fractional increase in detection rate due to the inclusion of spinning waveforms in the template bank is then given by

$$\frac{V_{\rm spin}}{V_{\rm nonspin}} = \left(\frac{\rm FF_{\rm spin}}{\rm FF_{\rm nonspin}}\right)^3,\tag{6.8}$$

where $FF_{nonspin}$ and FF_{spin} are the effective fitting factors for the non-spinning and spinning template banks, respectively, to the target signal. These statements regarding the increase in detection rate assume that the only gain in sensitivity comes from the SNR. In practice, our data also contain non-Gaussian artifacts and SNR is not an optimal detection statistic, as mentioned above.

The parameters chosen for our simulation are listed in Tbl. 6.2. The waveforms are generated by solving the ordinary differential equations given by Eqns. 1.60-1.62 in the SPINTAYLORT5 approximation (see Sec. III of Ref. [25] for the full description)². The target binaries (for a set of fixed values of component masses) are uniformly distributed in volume throughout the local universe. Spin magnitudes are distributed uniformly between zero and a maximum value (see Tbl. 6.2) and the spin angles are isotropically distributed. Bodies under two solar masses are assumed to be neutron stars; their spin magnitudes are limited to less that 0.4. All others are assumed to be black holes,

 $^{^{2}}$ This particular approximant is chosen so as to disentangle the effects of precession from the effects of different PN approximants; see below for a discussion.



Figure 6.3: Dependence of precession effects on mass ratio. We show the fitting factor (indicated by the color of the dots) of the reduced-spin template bank in detecting generic spinning binaries with component masses $(6M_{\odot}, 6M_{\odot})$ (left) and $(10M_{\odot}, 1.4M_{\odot})$ (right). The x-axis corresponds to the spin magnitude of the more massive compact object, while the y-axis corresponds to the cosine of the angle between the spin and initial Newtonian orbital angular momentum. In the left plot (equal-mass binary) fitting factors are ~ 1 irrespective of the magnitude and orientation of the spin vector, while in the right plot (highly unequal-mass binary) fitting factors can be as low as ~ 0.7 for binaries with large, misaligned spins.

with spin magnitudes limited to below 0.98. Cosine of the angle ι describing the relative orientation of the initial total angular momentum of the binary with respect to the line of sight is uniformly distributed in the interval (0, 1), while the polarization angle ψ is uniformly distributed in $(0, \pi)$.

The computed effective fitting factor FF_{eff} of the *reduced spin* template bank is shown in the top panel of Fig. 6.1. For each point, we have injected 5,000 waveforms with fixed mass but arbitrary spins and binary orientations. The figure indicate that the bank is effective for the detection of generic spinning binaries over almost all the relevant regions in the "low-mass" parameter space $(m_1 + m_2 < 12 M_{\odot})^3$. Even in the worst case of the $(4.5 M_{\odot}, 2.95 M_{\odot})$ binary, the effective fitting factor is ~ 0.93, and in almost all regions in the parameter space the fitting factor is greater than 0.95 (note that the minimum match requirement \mathcal{M}_{min} on the template bank was chosen to be 0.95). Note that the region above the gray line in the figure is the region where the contribution from the post-inspiral stages are expected to be significant, and the inspiral template bank needs to be replaced by an inspiral-merger-ringdown bank.

The effective fitting factors for the *non-spinning* template bank (covering the same mass range)

³In conflict with our previously terminology used in the discussion of S6/VSR2-3 searches, we redefine "low-mass" to the region of the parameter space with roughly $M_{\text{total}} \leq 12M_{\odot}$, based on recent studies which showed that it is essential to include the effects of post-inspiral stages in the waveform for binaries with total mass $\gtrsim 12M_{\odot}$ [126] (see also Fig. 4.3).



Figure 6.4: Selection bias against precessing systems. We plot the normalized SNR (such that the maximum SNR is 1) of generic spinning binaries against the fitting factor (FF) of the reduced-spin template bank in detecting them. It can be seen that fitting factors are high towards binaries with large SNR. The color of the dots correspond to the sine of the inclination of the total angular momentum vector with respect to the line of sight (darker shades correspond to binaries whose total angular momentum is along the line of sight). The left plot corresponds to binaries with component masses $(6M_{\odot}, 6M_{\odot})$ and the right plot to binaries with component masses $(10M_{\odot}, 1.4M_{\odot})$.

is shown in the right panel of Fig. 6.1. The fitting factor of the non-spinning bank is 0.84-0.89 over the same parameter space. The average increase in the detection volume provided by a search employing the reduced-spin template bank (as compared against the corresponding non-spinning template bank) is shown in Fig. 6.2. The figure suggests that we can expect an increase of $\sim 19-58\%$ in the average detection volume at a *fixed SNR threshold*. Note that the real figure of merit of the improvement would be the increase in the detection volume for a *fixed false-alarm rate*. Calculation of this requires the calculation of the increase in the false-alarm rate due to the increased number of templates in the bank. This work is currently underway.

The high signal recovery by the reduced-spin template bank from generically spinning binaries can be attributed to two reasons. Firstly, for binaries with comparable masses $(m_1 \sim m_2)$ the total angular momentum is dominated by the orbital angular momentum, and hence the modulational effects of spin precession on the orbit, and hence on the observed signal, is small. In this regime, nonprecessing waveforms provide a good approximation to the observed signal. However, as the mass ratio increases, spin angular momentum becomes comparable to the orbital angular momentum and the modulational effects of precession become appreciable. Effectualness of non-precessing templates thus decrease with increasing mass ratio (see Fig. 6.3).



Figure 6.5: Difference between the T4 and T5 post-Newtonian approximations. Same as Fig. 6.1, except that in this plot, the target waveforms are generated using the SPINTAYLORT4 approximation. The difference in the effectualness between Fig. 6.1 and this figure is due to the difference between the two different PN approximants, and is a reflection of the current uncertainty in the PN waveforms. This could be improved by computing the higher order (spin-dependent) PN terms.

Secondly, there is an intrinsic selection bias towards binaries that are nearly "face-on" with the detector (where the modulational effects of precession are weak while the signal is strong) as opposed to binaries that are nearly "edge-on" (where the modulational effects are strong while the signal is weak). Thus the fitting factors are high towards binaries with large SNR. This effect is illustrated in Fig. 6.4 for the case of an equal-mass binary (left) and for the case of a highly unequal-mass binary (right). This helps the reduced-spin template bank to have reasonably high effective fitting factor towards a population of generic spinning binaries.

The reduction in the fitting factor of the reduced-spin template bank in the high-mass and highmass-ratio regimes is due to multiple reasons. The modulational effects of precession increase with increasing mass ratio, and are not modeled by our templates. There are additional factors causing the loss: The difference between different PN approximants become considerable at the high-mass-, high-mass-ratio regime (reflecting the lack of knowledge of the higher order spin-dependent PN terms), causing appreciable mismatch between the target waveforms and the template waveforms even in regions where they should agree [127] (e.g., in the limit of non-precessing spins). Hence, it is likely that the fitting factor can be further improved by including the higher order PN terms, assuming that these higher order terms will reduce the difference between different PN approximants.

The effective fitting factors for our template banks are computed using a population of generic spinning binaries, assuming that the target signals are given by the SPINTAYLORT5 post-Newtonian approximation. The SPINTAYLORT5 model is one of the many approximations that can be used

to compute PN waveforms from inspiralling compact binaries, and these different approximations can produce somewhat different results (see Ref. [18] for an overview of different approximations). We would like to disentangle the loss of effectualness due to this effect from the loss due to the effect of precession. Thus, as the target waveform we need to use an approximant that is closest to the template in regions of parameter space where the target and template are expected to agree very well (e.g., in the limit of non-precessing spins). This is the motivation for choosing SPINTAY-LORT5 approximation as the target waveform.

We do not expect *a priori* one approximant to be closer to the signals given by nature than any other approximants. This was further confirmed by comparisons of PN approximants with numerical-relativity simulations [128, 129]. Thus, in order to get a conservative estimate of the effectualness of the template banks, we compute their effective fitting factors towards signals from generic spinning binaries computed in the SPINTAYLORT4 approximation (Fig. 6.5). Note that the fitting factors at high mass ratios are slightly lower than what we see in Fig 6.1. This difference arises from the fact that the waveforms computed using different approximants can be somewhat different, reflecting the current uncertainty in the PN waveforms (see also [130] for a detailed discussion). It is likely that this uncertainty will decrease with the knowledge of higher PN terms (note that currently the spin-dependent terms are known only up to 2.5PN).

We argue that one of the main reasons for the lower effectualness of the reduced-spin template bank towards SPINTAYLORT4 waveforms at high mass ratios is, apart from the modulational effects of precession, the difference between PN approximants. In order to demonstrate this, we compute the effective fitting factor of the reduced-spin template bank towards SPINTAYLORT4 waveforms *with non-precessing* spins (Fig. 6.5). It can be seen that the mismatch of the template bank at high mass ratios (in the low-mass regime) can be as large as 5%. This cannot be attributed to the effects of precession. These results greatly motivate the need of computing higher order spin terms in the post-Newtonian approximation.

As discussed above, all the currently observed neutron stars in binaries have spin periods, which correspond to spin magnitudes of $|\vec{\chi}_i| \leq 0.05$. While this is not necessarily an upper limit on neutron-star spins, this could be indicative of the typical spins. We have repeated the simulations by restricting the spin range of neutron stars in the target binaries to the interval (0, 0.05). This was found to make an appreciable difference only in the binary-neutron-star $(m_{1,2} \leq 2M_{\odot})$ region of the parameter space. In this region, the effective fitting factor of the non-spinning template bank


Figure 6.6: Capturing aligned spin effects in high mass template banks. We show the expected fractional signal-to-noise recovery for a population of aligned-spin binary black holes using a bank of IMRPHE-NOMB waveforms with (a) $\chi_{\text{eff}} = 0$ and (b) $\chi_{\text{eff}} \ge 0$. The solid lines indicate the approximate fitting factor contours in the $M_{\text{total}} - \chi_{\text{eff}}$ plane, averaging over the mass ratio dimension with $1 \le m_1/m_2 \le 4$. The template banks are both constructed with SBANK assuming the design iLIGO sensitivity [11] with $f_{\text{low}} = 40$ Hz. We find that with this sensitivity, a template bank that neglects spin achieves fitting factors exceeding the nominal $FF_{\text{min}} = 0.97$ from aligned-spin systems over a wide region of parameter space, spanning roughly $-0.25 \le \chi_{\text{eff}} \le 0.2$ over the entire mass range. As the mass of the system increases, the loss of signal-to-noise incurred from neglecting spin becomes small, and we therefore do not consider systems with total masses exceeding $M_{\text{total}} = 35 M_{\odot}$. The $\chi_{\text{eff}} = 0$ bank has ~ 700 templates, whereas the $\chi_{\text{eff}} \ge 0$ bank has ~ 3000 templates.

was increased to 0.98. Thus, under this assumption, the non-spinning template bank appears to be adequate for the detection of GWs from binary neutron stars.

6.4 Aligned spin template banks for the detection of merging binary black holes

We now consider higher mass systems for which each of inspiral, merger, and ringdown are important for signal recovery. Here we study template banks generated again with SBANK, but now using the IMRPHENOMB approximation as described in Sec. 1.5.2. This waveform models aligned spin effects again with a single additional parameter, called the *effective spin* and is defined as

$$\chi_{\text{eff}} = \frac{m_1 \chi_1 + m_2 \chi_2}{m_1 + m_2}.$$
(6.9)

Unfortunately, in the IMR regime we do not yet have precessing waveform models available to study the recovery of precessing binary black holes with aligned spin template banks. Instead, we focus on the improvement in signal recovery for aligned spin systems only.

The IMRPHENOMB waveform model consists of a parametrized phenomenological fit to hybrid waveforms constructed from numerical relativity simulations of the late-inspiral [that is, $f_{\rm GW} \gtrsim$ $10^{-3}/(GM_{\rm total}/c^3)$, where $f_{\rm GW}$ is the dominant mode gravitational wave frequency], merger and ringdown of binary black holes matched to a post-Newtonian approximation describing the early inspiral. As such, the validity of these waveforms have restrictions on the mass ratio and spins based on the availability of numerical simulations with which to fit. Specifically, the IMRPHENOMB family is expected to be accurate only for low to moderate mass ratios and spins. Hence, in these studies, we consider only binaries for which $1 \le m_1/m_2 \le 4$ and $-0.5 \le \chi_{\rm eff} < 0.85$.

We choose to further focus only on the regions in the parameter space where the merger and ringdown stages are important for detection. For both an initial and advanced LIGO design sensitivity, the effects of merger and ringdown begin to contribute significantly to the SNR when the total mass of the binary exceeds $M_{\text{total}} \approx 12 \ M_{\odot}$ [18]. For lower mass systems, accurate and generically spinning post-Newtonian waveforms are available [92], and can be used to give a more detailed understanding the effects of spin on the search. We therefore consider only systems with $M_{\text{total}} \geq 10 \ M_{\odot}$, giving a small safety factor between the transitional region and considering the degeneracy between the mass and spin parameters.

Since the finite size of neutron stars can have a significant impact on the gravitational waveform observed in the merger phase of coalescence, we restrict our attention to binary black holes and take $m_i = 3 M_{\odot}$ as the minimal component mass. We note that from astrophysical considerations, neutron stars in coalescing compact binaries are not expected to have large spins. Further, from physical considerations of the possible neutron star equations of state, the dimensionless spin for a neutron cannot exceed ~ 0.7 without undergoing tidal disruption.

We consider two regimes separately. First, we consider the available gains in SNR by the use of aligned spin templates to analyze data with iLIGO sensitivity. This study is intended to presage our development of a pipeline, which includes aligned spin templates to recover aligned spin signals in *real* LIGO noise. Then we turn our attention to signal recovery in data obtained at aLIGO sensitivities. Here we immediately bump up against questions of the validity of the phenomenological models. Nonetheless, as we discuss in this and the next chapter, preliminary studies pushing these waveforms to their maximum indicate that spin effects for aLIGO binary black hole searches will be *extremely* important. The development and validation of IMR waveforms with spin effects must be a top priority for the success of these searches.

6.4.0.1 Initial LIGO sensitivity

As a proof of principle, we are going to demonstrate our analysis pipeline with spinning templates on initial LIGO data, where the computational scale is not a limiting factor. Here we investigate template banks covering an initial LIGO design sensitivity to determine the regions in parameter space which have the greatest potential for improvement in SNR recovery.

We used SBANK to construct banks of IMRPHENOMB templates with $\chi_{\text{eff}} = 0$ using the above mass parameter restrictions and $f_{low} = 40$ Hz. We then computed the fittings factors of this template bank towards aligned spin signals in the same $m_1 - m_2$ parameter space. In Fig. 6.6a, we show that a template bank with $\chi_{\text{eff}} = 0$ already captures greater than 97% of the possible SNR over a wide mass and spin range. In particular, we note that the $\chi_{eff} = 0$ bank covers signals with $\chi_{eff} < 0$ down to roughly $\chi_{\rm eff} \sim -0.25$ over the entire mass range. From astrophysical considerations of binary evolution, spins positively aligned with the orbital angular momentum are considered the more likely scenario for binary black holes [26]. Given these factors, along with the potential for artifacts in the waveforms at large negative χ_{eff} , we develop our search using only $\chi_{\text{eff}} \geq 0$ templates. Note that since χ_{eff} is a mass-weighted sum of the two component spins, this restriction does not necessarily exclude the possibility that one of the black holes has an anti-aligned spin. We also see that as the total mass of the target system increases, the fractional loss of SNR incurred from neglecting spin decreases. This effect is due to the fact that higher mass systems merge at lower frequencies and have fewer cycles in the LIGO sensitive band, and consequently the matched filtering is more tolerant of imperfect templates. We thus expect that for systems with total masses exceeding $M_{\text{total}} = 35 M_{\odot}$, the benefits of including spin effects will be small.

In Fig. 6.6b, we demonstrate the coverage of the parameter space obtained by including only waveforms for non-negative aligned spins ($\chi_{\text{eff}} \geq 0$) in the template bank. The improvement in SNR recovery obtained by using such a bank comes at the cost of having more than three times as many templates in the bank. For the non-spinning case, we constructed a bank with ~ 700 templates, while to cover the positively aligned signals, we require ~ 3000 templates. The increase in the number of templates will increase the number of background triggers, and detecting a signal at a given false probability requires raising the SNR threshold used for detection.

The characteristics of the background can change in complicated ways when new template wave-



Figure 6.7: Importance of spin for aLIGO binary black hole searches. The two plots above show the coverage of a non-spinning template bank to aligned spin target signals, as in Fig. 6.6a, but for a late aLIGO design spectrum. To the left, we show the signal recovery by a non-spinning template bank integrating from 25 Hz. To the right, we show the same for higher mass systems integrating from 10 Hz. Note the regime of validity for the IMRPHENOMB waveforms requires one to make sacrifices between the low frequency cutoff for integration and the mass range covered by the template bank. For maximal signal recovery in aLIGO, we must have valid IMR waveforms that extend down to 10 Hz in the entire BBH mass range.

forms are introduced to a search. The results presented in Fig. 6.6 do not reflect the impact of non-Gaussianity in the data, nor do they capture the effects of multi-detector coincidence requirements, the use of χ^2 statistics, increased false alarms due to larger template banks, or other effects which are important in realistic search pipelines. In the following chapter, we describe the implementation of these spinning template banks in a search pipeline. We show that even in non-Gaussian data, we are able to sufficiently suppress the extra background to achieve a net gain in the search sensitivity.

6.4.0.2 Advanced LIGO sensitivity

In Fig. 6.7, we show how the results presented for iLIGO extend to an aLIGO sensitivity. Here we note that we are limited by the range of validity of the IMRPHENOMB waveform family. Thus, we show two cases: (i) low mass IMR aligned spin with 25 Hz lower frequency cutoff and (ii) high mass aligned spin with 10 Hz lower frequency cutoff. The simulation results indicate, on the basis of SNR, that spin effects are very important for binary black holes in aLIGO and up to even higher masses than in iLIGO. In the next chapter, we demonstrate an analysis of simulated aLIGO noise and show that in fact these expectations hold even after considering false alarm rates.

Chapter 7

A Sensitive Search Pipeline for Binary Black Holes with Aligned Spin

When it is obvious that the goals cannot be reached, don't adjust the goals, adjust the action steps.

Confucius

In the previous chapter, we showed that the inclusion of spin effects in templates for compact binary searches may significantly increase the detection rate if such spinning systems are common in astrophysical populations of binary black holes. Our discussion there was based on the expected improvement in signal-to-noise in idealized pipelines running on data with Gaussian background. In reality, pipeline implementations, even in Gaussian noise, require compromises between optimality and computational feasibility. Furthermore, LIGO data contain non-Gaussian artifacts which can mimic the gravitational wave signals we are trying to detect. This latter problem becomes considerably worse the shorter the template waveforms, as is the case for binary black holes.

In this chapter, we demonstrate the implementation of a sensitive search pipeline in GSTLAL for binary black holes whose components have significant spin. We demonstrate an analysis of simulated binary black holes signal in the range $M_{\text{total}} \in [15, 25] M_{\odot}$ added to iLIGO detector noise, showing that the pipeline recovers 45% more volume from highly spinning binaries than the equivalent search using non-spinning templates, even in the presence of realistic non-Gaussian noise. After demonstrating this analysis in real iLIGO data, we extend our results to an analysis of simulated Gaussian data with a late aLIGO spectra. In this case, our simulated signals from $M_{\text{total}} \in [60, 100] M_{\odot}$ and we find more than a factor of two increase in sensitive volume for the highest spinning systems. As in the previous chapter, we will see that the improvement in recoverable search volume obtained by the inclusion of spin effects is dramatic when considering aLIGO sensitivities, but that covering the whole inspiral-merger-ringdown parameter space in this regime is limited by the availability of accurate waveform models.

7.1 Pipeline implementation challenges

Given that BBH systems probably have significant spins and that the search sensitivity depends strongly on having accurate waveform models, the inclusion of spin effects in search templates has for good reason been a long-standing goal in the field [8,49,50,69–71,131–134]. Yet, as discussed in the introduction to Chap. 6, none of the previous attempts to include spin effects in templates have resulted in improved search sensitivity.

The null result in Ref. [49] was attributed to the elevation in the rate of background events owing to the large number of parameters required to describe the spinning template waveforms. It was immediately realized that to make the use of spinning templates in an analysis beneficial, better signal consistency tests would have to be developed and implemented to suppress the increased background event rates. In particular, the highly effective χ^2 test used in the contemporary LIGO search with non-spinning templates [135] was never integrated into the spinning search pipeline, probably at a severe cost for the results of the sensitivity analysis. Following this work, two other pipelines were developed [133, 134] which included spinning templates based on a modified phenomenological model for single-spin binaries [132], but neither analysis demonstrated conclusively that the techniques would improve the search sensitivity. Neither of these investigations examined the use of signal consistency tests for suppressing the background. In our implementation below, we use an autocorrelation χ^2 test (described in Sec. 3.3.3), which unlike the traditional χ^2 requires no knowledge of the frequency evolution of the waveform.

Two other implementation issues arose from these early investigations on spin, ultimately concerning the question of having a measure for the "distance" between nearby templates. The first was the problem of defining coincidence between triggers in different detectors. In the two-dimensional mass parameter space, elegant and rigorous techniques exist for defining when the parameters of two triggers are close enough to be considered the same [74]. However, it is not clear how to extend this method to define a robust coincidence criterion in higher-dimensional parameter spaces. The analyses in Refs. [49] and [134] defined coincidence between triggers in terms of the standard mass coincidence criterion proposed in Ref. [74] together with a simple interval cut on the remaining parameters. The study in Ref. [133] was for a coherent analysis, for which the question of coincidence is irrelevant.

Furthermore, as we have discussed extensively, first in Sec. 3.2.2 and again in Sec. 6.2, another problem that comes with having a larger template parameter space is that of efficiently placing templates to minimize the loss of SNR arising from the discreteness of the template bank. Recent LIGO and Virgo compact binary searches, which used non-spinning templates, have relied on a lattice placement technique (discussed in Sec. 3.2.3) known to select the fewest number of templates for a given tolerance of SNR loss [66,67,72,73]. In Ref. [49], the authors explored the use of a stochastically generated bank and found it to give significantly fewer templates than the stacking approach. More recently, several groups have conducted thorough and systematic studies of the stochastic template placing techniques and the outlook is quite promising [8,64,65]. Stochastic placement techniques (discussed in Sec. 3.2.4) are applicable to a wide variety of waveform approximants, requiring no prior knowledge of special parameters and extending straight-forwardly to higher dimensions. In Sec. 6.2, we described a generic infrastructure called SBANK for creating stochastic template banks, which we then used to create template banks of aligned spin templates.

In following sections, we integrate template banks constructed in the previous chapter into a real analysis pipeline and measure the sensitivity of the pipeline in real detector noise. Our analysis uses the three-parameter IMRPHENOMB waveform family, introduced in Sec. 1.5.2, which models the inspiral, merger, and ringdown of binary black hole system with aligned spins. This waveform family captures the dominant effects of aligned spin with a single *effective spin* parameter. As before, we construct template banks from these waveforms using our generic, extensible stochastic placement infrastructure SBANK [8] implemented in the LAL gravitational wave data analysis library [92], and incorporate the SBANK infrastructure into the GSTLAL pipeline described in Chap. 5 [136].

7.2 A sensitive search pipeline for gravitational waves from merging binary black holes with aligned spin

We now demonstrate the use of the template banks constructed in Sec. 6.4 as filters in the GST-LAL search pipeline [136]. Using these template banks, we measured the sensitivity of the pipeline to a simulated population of more than 200,000 binary black hole mergers. Here we compare the mean sensitive distance of the pipeline analysis when using a bank of aligned-spin templates against that of an otherwise identical analysis which uses non-spinning templates. Our simulated binary black hole systems were populated with a uniform distribution in mass ratio, total mass, and effective spin with $m_1/m_2 \in [1, 4]$, $M_{\text{total}} \in [15, 25] M_{\odot}$ and $\chi_{\text{eff}} \in [0, 0.85]$. As with the templates, the simulated waveforms were computed using the IMRPHENOMB approximation. We conducted our study on 25.9 days of coincident detector strain data obtained from observations of the Hanford and Livingston detectors during LIGO's fifth science run. Typical strain sensitivities for these two detectors during this science run are shown in Fig. 2.6a.

In Fig. 7.1, we show the measured sensitivities of our two analyses in terms of the mean distance accessible to each search as a function of the false alarm rate threshold. We show our results only for systems with total masses in the range $M_{\text{total}} \in [15, 25] M_{\odot}$ to avoid complications associated with the boundaries of the template banks, which covered the range $M_{\text{total}} \in [10, 35] M_{\odot}$. As expected and demonstrated in Fig. 7.1c, we find that the greatest improvement in sensitivity is for target systems with high effective spins. As seen in Fig. 7.1d, the volume improvement, and therefore the increase in detection rate, can be as high as 45% for these highly spinning systems. We emphasize the nontrivial result shown in Fig. 7.1a that for weakly-spinning target systems ($\chi_{\text{eff}} \leq 0.2$), the analysis with spinning templates and the analysis with non-spinning templates have comparable sensitivities. with the aligned-spin template analysis achieving at worst 95% of the sensitive volume of the nonspin template analysis. The apparent loss of detection rate in the small effective spin regime is only applicable if we are wrong in our expectations that black holes have significant spins. Otherwise, we expect this search method to increase the overall detection rate of spinning BBH systems, provided that these spins are aligned. These results demonstrate for the first time an analysis of real detector data, which is made more sensitive to spinning signals by the use of spinning templates compared to the same analysis performed with non-spinning templates.

The analysis performed here differs from previous attempts towards the inclusion of spin effects



Figure 7.1: Comparison of search sensitivities as a function of false alarm rate threshold. We compare the sensitivities to aligned spin systems with $M_{\text{total}} \in [15, 25] M_{\odot}$ for an analysis which used templates with $\chi_{\text{eff}} \geq 0$ and an analysis which used templates with $\chi_{\text{eff}} = 0$. The template banks each covered the mass range $M_{\text{total}} \in [10, 35] M_{\odot}$. In (a-c), we show the absolute sensitivities for these analyses in terms of the average distance to which the analyses identify an injection with a trigger above a given false alarm rate threshold. In (d), we show the ratios of the sensitive volumes for each of the three spin bins. We find that for injections with $\chi_{\text{eff}} \geq 0.2$, the spinning search observes a larger sensitive volume than the non-spinning search for all false alarm rates by as much as 45%. For injections with $0 \leq \chi_{\text{eff}} \leq 0.2$, we observe a small but statistically significant decrease in sensitive volume on the order of 5% incurred by the use of spinning templates.

in search templates in a number of ways. Firstly, this analysis makes use of a template family that captures the effect of non-precessing spins by using a small number of *physical* parameters, which allows us to construct a simple three-dimensional template bank. Recent studies have suggested that such template banks are effectual for a significant fraction or precessing binaries as well [8,25]. This is in sharp contrast with the earlier work, which either used phenomenological parameters to capture spin effects [49,131] or methods to maximize the SNR over a number of extrinsic parameters that produced elevated background [132].



14

12

Recovered M_{chip} (M_o. Recovered $M_{\rm chip}$ 2 10 $^{6}_{\text{Injected }M_{\text{chirp}}} M_{\odot}$ 0.4 0.5 0.6 0.7 0.8 Injected M_{chirp} (M_{\odot}) Injected χ (a) $\chi_{\rm eff} = 0$ Template Bank (b) $\chi_{\text{eff}} \geq 0$ Template Bank (c) $\chi_{\text{eff}} \geq 0$ Template Bank

Figure 7.2: Improvement in parameter recovery with aligned spin templates. In addition to improving the sensitive search volume, the inclusion of spin effects helps to curtail systematic biases in the recovery of template parameters. Here we demonstrate that the improvement in chirp mass and effective spin parameter recovery when spin is included in the templates. The left panel corresponds to the analysis with non-spinning templates, while the middle and right panels corresponds to the analysis with aligned-spin templates. In the right panel, we see that the aligned spin template bank allows us to recover the effective spin parameter of the simulated signal with fair accuracy.

This analysis also used an autocorrelation χ^2 statistic, analogous to the traditional time-frequency statistic χ^2 used in recent LIGO and Virgo compact binary searches (these statistics are described in Sec. 3.3.3). The autocorrelation statistic is based on the principle that the SNR time series obtained from filtering data which contain a signal against a template that closely matches the signal is approximately equal to the autocorrelation function of the template plus noise. Subtracting the template autocorrelation from the SNR time series and computing the residual noise power gives a measure of the consistency of that data with the signal model.

As discussed in the previous section, early studies on the inclusion of spin effects in template waveforms suffered in part due to the lack of a sufficiently strong signal consistency tests to reject triggers occurring due to non-Gaussian artifacts in the data. We suggest that the autocorrelation test used here was instrumental towards achieving our results, and encourage the development and implementation of other signal-based consistency tests which could be added to this analaysis to improve upon these results (one such consistency test, known as the bank veto [75], is currently being tested within the gstlal pipeline). We also point out that the autocorrelation consistency test is appealing from a computational point of view since once a trigger has been produced by the pipeline, the needed SNR time series is already available in memory, and the calculation comes at nearly no extra cost.

We have also taken a simplistic, but seemingly quite powerful, approach to the matter of defining the coincidence of triggers between detectors. For coincidence in time, we follow the standard



Figure 7.3: Increase in false alarm rate from additional templates. We show the number of coincident background triggers observed in our two analyses of iLIGO data. The log-likelihood (x-axis) is a measure of the trigger significance and is one-to-one with the false alarm rate. We that the inclusion of spin in the search templates leads to an elevation in the background of a factor of a few, comparable to the relative sizes of the aligned spin and non-spin banks.

interval approach, requiring that triggers occur within 3 ms (plus the light travel time between detectors) of each other. For mass and spin coincidence, however, we require that triggers in each detector have identical parameters. This choice is possible in the GSTLAL framework since the same template bank is used for all detectors and all times in the analysis. Previous LIGO and Virgo searches for compact binary coalescence have used template banks whose parameters depend on the local power spectral density of a detector, resulting in template banks which are different in each detector and at different times. In the latter implementation, the coincidence criterion must allow for some small mismatch in the trigger parameters from different detectors. Recent searches using two-parameter non-spinning template banks have achieved this tolerance using estimates of the expected uncertainty in parameter recovery to define a small error region [74], but the generalization of this technique to higher dimensional parameter spaces is not straightforward, and the exact size of the error region typically requires careful tuning in order to be effective. On the other hand, the exact parameter coincidence feature of the GSTLAL search pipeline generalizes trivially to higher dimensional parameter spaces and requires no tuning. The results here suggest the exact coincidence criterion is a strong discriminator between background and signal, but we do not systematically examine the relative merits of these two approaches.

Ultimately, the key to improving the sensitivity of a search pipeline by the inclusion of more

physical effects in the search templates is the ability to manage the background trigger rates while exploiting the elevation of the signal. The methods described here have proven successful in mitigating the background elevation relative to the signal to obtain a net gain in sensitivity. We point out, however, that the inclusion of spin effects in the templates does increase the background levels in proportion to the increase in the size of the bank, as shown in Fig. 7.3. We have highlighted in this section the use of exact template parameter coincidence and the autocorrelation χ^2 . These are just two features of the GSTLAL pipeline which are manifestly different from other studies, and lie at the core of the background rejection techniques currently implemented in the pipeline. Given that the GSTLAL pipeline has not previously been used for an analysis of this type, there are of course many other differences between this work and previous studies, but isolating the particular features which made these results possible is a difficult task.

7.3 Prospects for advanced LIGO

We have extended the IMR study described in Sec. 7.2 to simulated Gaussian noise with a late advanced LIGO spectrum. The very first problem we encounter in extending the study performed in iLIGO data is the validity of the IMRPHENOMB waveform. Advanced LIGO detectors may have sensitivity all the way down to 10 Hz, but the IMRPHENOMB model becomes inconsistent with post-Newtonian predictions if extended to have too many cycles. Specifically, to extend the waveforms down to 10 Hz with confidence, it is recommended in Ref. [19] to restrict the systems we consider to $M_{\text{total}} \geq 40 M_{\odot}$. This situation is clearly not ideal, since most astrophysical priors suggest that most binary black holes are smaller in mass. However, we wanted to conduct a study in which the *full* advanced LIGO sensitivity is exploited and we therefore generated our template banks to obey this restriction. As before, we generate two template banks, one with spin and one without, and examine the pipeline sensitivity to spinning signals.

The results of our analysis are displayed in Fig. 7.4. Our template banks covered the mass space $M_{\text{total}} \in [40, 125] M_{\odot}$, in accordance with the limits on the waveform validity. We find that the inclusion of spin effects in the template waveforms can increase the observed volume by more than a factor of two for the highest spinning signals with masses $M_{\text{total}} \in [60, 80] M_{\odot}$. Note that the injection set lies totally *within* the template bank. We do this to avoid spurious adverse effects to the sensitivity in which the best matching template is slightly outside the template bank. In that case, another search, which included templates in that mass range, would have picked up the signal.



Figure 7.4: Improved sensitivity to spinning signals in simulated aLIGO noise. We compare the sensitive distances measured in the GSTLAL pipeline with (color) and without (black) spin effects in the templates. The template bank covers the space $M_{\text{total}} \in [40, 125] M_{\odot}$ and the simulated signals systems have $M_{\text{total}} \in [60, 80] M_{\odot}$, each restricted to mass ratios $m_1/m_2 \leq 4$. The template bank extends beyond the injection space in order to avoid template bank edge effects. We see that the inclusion of aligned spin effects in the template waveforms can increase the observable volume by more than a factor of two in this mass range for systems with the largest spins.

However, it is worth noting that if IMRPHENOMB were the only approximation we had available for aligned spin systems, this would be a *real* effect – we don't have IMR aligned spin templates below $M_{\text{total}} = 40 M_{\odot}$ that are valid down to 10 Hz. Fortunately, more accurate models with a wider range of validity are currently being developed. As they are developed, we plan to insert these models into the SBANK infrastructure to generate template banks with them and extend the study presented here down to even lower masses, where the effects of spin only become greater.

In Fig. 7.5, we illustrate the effect of waveform mismatch on the χ^2 . Recall that we have conducted our study in Gaussian noise. As a result, the χ^2 is in some sense superfluous; theory indicates that SNR is an optimal statistic in Gaussian noise. Nonetheless, even in Gaussian noise,



Figure 7.5: The effect of waveform mismatch on the χ^2 statistic. We compare the values of the autocorrelation χ^2 statistic for an analysis of simulated spinning signals with and without aligned spin templates for our study in simulated aLIGO noise. The injections here are the same as those in Fig. 7.4 except that the upper limit on total mass is 100 M_{\odot} . This figure illustrates the sensitivity of the χ^2 statistic to mismatch between the signal and the triggered template. In the non-spinning template case, the χ^2 statistic can have values that rival that of non-Gaussian background triggers (cf. Fig. 3.4), although such triggers are not present in this study on Gaussian noise. Hence when recovering spinning signals with non-spinning templates, the loss of SNR is compounded by the waveform mismatch which appears in the χ^2 .

the χ^2 statistic reflects the consistency of the template with the signal in the data. We see that when we use non-spinning templates to search for aligned-spin systems, the χ^2 values lie dangerously near where we would expect non-Gaussian triggers to appear (compare to Fig. 3.4, although this figure uses the traditional χ^2 in the IHOPE pipeline).

We have much further to go. We have completely neglected the mass range $M_{\text{total}} \in [15, 40] M_{\odot}$, relevant for IMR templates. We have at present no waveform which models both inspiral-mergerringdown and precession. Without these, we are unable to quantify how well these aligned spin filters work for detecting *generically* (not necessarily aligned) spinning binary black holes. The inclusion of spin effects for binary black hole searches could very well prove to be vital to their successful detection, but we need the waveforms.

Chapter 8

Conclusions

I don't understand why no adult knows that Sirius is the brightest star in the sky. I know Uncle Stephen would know.

My nephew Nicholas (age 5)

The LIGO experiment is one of the most important and exciting efforts in contemporary physics. Discovering gravitational waves from coalescing compact binaries would open up a whole new way of studying the sky. What we have demonstrated is one small piece of a much larger puzzle. What is the impact of the results presented here? How will the answers to the exam be different next year? We close our discussion in this thesis by zooming out and reminding the reader where this piece fits in the big picture.

8.1 Summary of results

In this thesis, we have focused principally on the problem of improving the detection rate of gravitational waves from compact binary coalescence in next generation searches by the inclusion of spin effects in the template waveforms. Along the way, we have covered a lot of ground, summarized below with references to the relevant sections.

We began our study with an analysis of data collected during LIGO's sixth and Virgo's second and third observational runs (Chap. 4). We showed that by the end of LIGO's fifth science run, the Hanford and Livingston LIGO detectors had achieved (above 60 Hz) the iLIGO design sensitivity written down decades earlier before any construction had taken place (Sec. 2.5). In the latest run, with improvements upon the initial design, these detectors exceeded the design sensitivity by up to a factor of two in strain – a factor of eight in observable volume (Secs. 2.5). We used the IHOPE pipeline (Sec. 3.4) to analyze these data for compact binary coalescence. No gravitational wave events significantly above the background were identified in these searches. Based on these observations, we computed upper limits on the coalescence rate as a function of mass over the searched parameter space, which spanned $M_{\text{total}} \in [2, 100] M_{\odot}$ (Sec. 4.6). Unfortunately, the derived limits do not yet constrain even the most optimistic rate predictions. However, we found that the naive SenseMon criterion (Sec. 4.1) for the detectability of a compact binary coalescence signal has impressive agreement with the achieved sensitivity of the analysis pipeline (Sec. 4.6), an incredible testament to the search methods.

We then turned our attention to scaling these analyses to aLIGO (Chap. 5) and considering improvements that could be obtained through the use of aligned spin templates (Chap. 6). In looking towards aLIGO searches for compact binary coalescence, we uncovered a number of potential scalability problems associated with (i) going to lower frequencies and (ii) expanding the template parameter space to include spin (Sec. 5.1). We discussed multi-banding and the singular value decomposition as techniques for reducing the computational scale of aLIGO analyses (Secs. 5.2 and 5.3). These techniques are implemented in a new pipeline infrastructure GSTLAL (Sec. 5.4) which is capable of producing event candidates with very low latency (Sec. 5.5). We used the GSTLAL pipeline to demonstrate an analysis of actual (initial) LIGO detector noise in which the pipeline sensitivity to spinning signals *improves* with the use of aligned spin templates (Sec. 7.2).

In the process of developing this pipeline, we implemented a generic random template placement infrastructure SBANK (Sec. 6.2), and applied this infrastructure to generate template banks of TAYLORF2REDSPIN (Sec. 6.3) and IMRPHENOMB aligned spin signals (Sec. 6.4). Using a late aLIGO spectrum, we found that while spin effects are important for the detection of gravitational waves from compact binaries, Nature selects against complicated signals. Aligned spin templates may be sufficient to detect a majority of fully-precessing binaries in aLIGO, provided their signals are dominated by the inspiral (Sec. 6.3).

Finally, we considered the problem of demonstrating the aligned spin pipeline analysis (performed on initial LIGO noise) to searches in aLIGO. We found that we do not presently have implemented the waveforms that will allow us to take full advantage of the aLIGO detector sensitivity, at least in the mass range $M_{\text{total}} \in [15, 40] M_{\odot}$ (Sec. 6.4). The development of accurate waveform models which include the effects of merger, ringdown, and spin is crucial to extracting the most of LIGO and Virgo observations. Nonetheless, pushing the waveforms we do have to the limit of their applicability, we showed in simulated Gaussian noise with a late aLIGO spectrum that gains of a factor of two or more in volume are obtainable by the inclusion of spin in the templates for systems with masses in the range $M_{\text{total}} \in [60, 100] M_{\odot}$ (Sec. 7.3).

8.2 Future work

A number of followup projects are already underway. We list them here in the form of questions they aim to answer.

How can we cover the full inspiral-merger-ringdown (mass and spin) parameter space with templates that extend down to 10 Hz?

At the end of Sec. 6.4, we noted that the IMRPHENOMB has limitations in its regime of validity that make it unsuitable for use as a template in aLIGO searches, particularly in the regime $M_{\text{total}} \in$ [15, 40] M_{\odot} . In this mass range, merger and ringdown effects are still important for detection; however, IMRPHENOMB disagrees with the post-Newtonian expansion when extended down to 10 Hz in this mass range. Furthermore, the IMRPHENOMB model also breaks down at high mass ratios and large effective spins. Another phenomenological model IMRPHENOMC [23] has recently been implemented in the LALSUITE code base and is currently being investigated to understand its regime of validity (or "believability", since we do not actually have the numerical relativity waveforms to compare to). In particular, IMRPHENOMC is expected to be valid up to mass ratios as high as $m_1/m_2 = 20$, possibly the full effective spin range [-1, 1], and is faithful to the post-Newtonian expansion for many more cycles than IMRPHENOMB . Should these expectations hold up to scrutiny, the IMRPHENOMC model will be a strong candidate for use as search templates in next generation binary black hole searches. Furthermore, it is trivial to drop IMRPHENOMC into SBANK to create template banks and analyze simulated signals with these template banks in GSTLAL . This work is already underway.

How well do aligned spin templates capture precession?

This question has to be separated into two regimes, based on the availability of waveforms. For systems with $M_{\text{total}} \in [6, 15] M_{\odot}$, post-Newtonian templates are sufficient for near-optimal detection and we are still in a mass regime likely to contain binary black holes, for which precession may be important. We have already fully precessing waveforms in the post-Newtonian approximation. Indeed, in Sec. 6.3 we studied the ability for an aligned spin template bank to capture the signal-tonoise from fully precessing signals and we found that – from the perspective of SNR recovery – that the aligned spin template bank does quite well. Now we want to address the question of how this template bank performs in actual data. Thus, our goal is to study the precessing and aligned spin signal recovery of this bank in the GSTLAL pipeline on simulated aLIGO noise. This work is already underway, using fake Gaussian aLIGO noise, and is showing great promise. The hard problem of course is how the template bank responds to the presence of glitches. Without actual aLIGO data, this is a difficult question to address. Our plan is to use recolored S6 data as a proxy to aLIGO noise. We expect that longer filters will be less sensitive to glitches and, given the success of the IMRPHENOMB pipeline studies reported in Sec. 7.2, we are optimistic about the fruits this work will bear.

In the IMR regime, our ability to address the question of how well aligned spin templates capture the effects of precession is severely limited by the availability of (i) phenomenological models for precessing IMR signals and (ii) numerical simulations of such systems, upon which the development of reliable precessing IMR models depend. Nonetheless, recently a new phenomenological IMR waveform model IMRPHENOMP with precessional effects has been developed [24], although its implementation in LALSUITE is still in its early stages. Once the implementation of this model is fully trusted, we can use it to make the first systematic study of the performance of aligned spin IMR templates against precessing IMR signals in an actual pipeline analysis. Work on this project will begin in the summer.

8.3 Long term prospects

Finally, we remind the reader of the long term goals for the work presented here. These are very big picture items, for which the work described here plays only a partial, but hopefully significant, role.

How do binary stellar systems evolve?

Although we have a fairly good understanding of how isolated stellar systems evolve, there are still a lot of unknowns on how interacting stellar systems evolve. What is the result of the common envelope stage? What is the distribution for the magnitude and direction of natal supernova kicks in a binary? Are Kozai resonances capable of forming short period binary black holes? These are hard questions to answer with electromagnetic observations. The common envelope phase is incredibly short compared to the lifetime of the binary, making its observation by light comparatively rare; but its outcome is imprinted on the properties of the compact binary it leads to, if it does at all. Furthermore, binary black holes are of course black. If these systems exist in globular clusters, gravitational waves are the only to directly observe them. Many other interesting astrophysical conclusions may come from gravitational wave observations of compact binary coalescence, for instance, on the origin (or existence) of a mass gap between neutron stars and black holes, or on the origin (or existence) of a mass gap between stellar mass black holes ($\sim 10-100M_{\odot}$) and supermassive black holes ($\sim 10^8M_{\odot}$) observed in the center of many galaxies.

What are the origins of short hard γ -ray bursts?

We recall that short hard γ -ray burst are thought to be due to fall back of tidally disrupted neutron stars onto a newly formed black hole resulting from a binary merger. In order to confirm this hypothesis, it is important to be able to generate candidate events in very low latency so that we can immediately perform electromagnetic followup observations of the event. The GSTLAL pipeline used here has demonstrated its capacity to deliver low latency triggers. Although we use the pipeline to develop improved search methods for binary black holes, which are thought to not have any electromagnetic counterpart, our methods can in principle be applied to the binary neutron star and neutron star-black hole regime, although the computational scale for studies in this mass (BNS) and mass ratio (NSBH) regime quickly becomes challenging to making progress, particularly because the inclusion of spin requires many factors more templates than are already required.

Is General Relativity an accurate description of gravitational dynamics?

We have not directly addressed this question at all in this thesis. Indeed, the question of whether aLIGO observations have the strength to answer this question forms a vast body of active research in itself. However, if LIGO observations can make such determinations, careful examination of the gravitational wave signals from compact binary coalescences are one of the most promising prospects for doing so. Of the currently understood potential sources of gravitational waves, binary coalescences are the cleanest, the easiest to model, and the systems for which we can make the most precise predictions for the observed waveforms. Furthermore, among compact binaries, binary black holes offer the unique opportunity to test General Relatively in an extraordinarily relativistic setting. Binary black holes are also potentially cleaner systems than binaries involving neutron stars, where matter effects probably become important in the merger regime. What we have shown here is that if black holes in binaries have large spins, then the inclusion of spin effects in the search templates will significantly increase the detection rate. We hope that black holes spin very fast, and that the methods presented here will help to maximize the number of detected events in aLIGO.

We look forward to collecting the dividends suggested above – and others that we have failed to anticipate – in the next generation of LIGO and Virgo observations. But there is no time to spare and we must be leaving you now to prepare!!

Bibliography

- J. Abadie et al. Sensitivity to Gravitational Waves from Compact Binary Coalescences Achieved during LIGO's Fifth and Virgo's First Science Run. 2010.
- [2] J. Abadie et al. Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO's Sixth and Virgo's Second and Third Science Runs. 2012.
- [3] The LIGO and Virgo Collaborations. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3. *Phys. Rev. D*, 85:082002, 2012.
- [4] The LIGO and Virgo Collaborations. Search for Gravitational Waves from Binary Black Hole Inspiral, Merger and Ringdown in LIGO-Virgo Data from 2009-2010. *Phys. Rev. D*, 87:022002, 2013.
- [5] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T Cokelaer, J. D. E. Creighton, T. Dent, A. Dietz, S. Fairhurst, N. Fotopoulos, G. Gonzalez, C. Hanna, I. W. Harry, G. Jones, D. Keppel, D. J. A. McKechan, L. Pekowsky, S. Privitera, C. Robinson, A. C. Rodriguez, B. S. Sathyaprakash, A. S. Sengupta, M. Vallisneri, R. Vaulin, and A.J. Weinstein. Searching for gravitational waves from binary coalescence. *Phys. Rev. D*, 87(2):024033, 2013.
- [6] J. Abadie et al. Erratum: Search for gravitational waves from binary black hole inspiral, merger and ringdown. 2012. To be published, Phys. Rev. D.
- [7] R. Biswas, P.R. Brady, J. D. E. Creighton, S. Fairhurst, G. Mendell, and S. Privitera. Corrigendum: The loudest event statistic: general formulation, properties and applications. 30:079502, 2013.

- [8] Ajith, P. and Fotopoulos, N. and Privitera, S. and Mazumdar, N. and Weinstein, A.J. An effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Phys. Rev. D*, 89:084041, 2014.
- [9] Privitera, S. and Mohapatra, S.R.P. and Fotopoulos, F. and Frei, M.A. and Hanna, C. and Weinstein, A.J. and Whelan, J.T. Improving the sensitivity of a search for coalescing binary black holes with non-precessing spins in gravitational wave data. *Phys. Rev. D*, 89:024003, 2014.
- [10] A. Einstein. Nherungsweise Integration der Feldgleichungen der Gravitation. Preussische Akademie der Wissenschaften, pages 688–696, 1916.
- B. Abbott et al. LIGO: The Laser Interferometer Gravitational Wave Observatory. *Rep. Prog. Phys.*, 72:076901, 2009.
- [12] T. Accadia et al. Virgo: A Laser Interferometer to Detect Gravitational Waves. JINST, 7:P03012, 2012.
- [13] H. Grote. The Status of GEO 600. Class. Quantum Grav., 25:114043, 2008.
- [14] K. Somiya. Detector Configuration of KAGRA: the Japanese Cryogenic Gravitational Wave Detector. Class. Quantum Grav., 29:124007, 2012.
- [15] G.M. Harry for the LIGO Scientific Collaboration. Advanced LIGO: The Next Generation of Gravitational Wave Detectors. *Class. Quantum Grav.*, 27(8):084006, 2010.
- [16] Advanced Virgo Baseline Design, 2009. VIR027A09.
- [17] Sean M. Carroll. Lecture Notes on General Relativity. http://arxiv.org/abs/gr-qc/9712019, 1997.
- [18] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash. Comparison of post-Newtonian Templates for Compact Binary Inspiral Signals in Gravitational Wave Detectors. *Phys. Rev. D*, 80(8):084043, 2009.
- [19] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Bruegmann, N. Dorband, D. Mueller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaria, and J. Seiler. Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins. *Phys. Rev. Lett.*, 106:241101, 2011.
- [20] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, 1999.

- [21] A. Buonanno, Y. Pan, H.P. Pfeiffer, M.A. Scheel, L.T. Buchman, and L.E. Kidder. Effectiveone-body waveforms calibrated to numerical relativity simulations: Coalescence of nonspinning, equal-mass black holes. *Phys. Rev. D*, 79:124028, 2009.
- [22] F. Pretorius. Evolution of Binary Black Hole Spacetimes. Phys. Rev. Lett., 95:121101, 2005.
- [23] L. Santamaria et al. Matching post-newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries. *Phys. Rev.* D, 82:064016, 2010.
- [24] M. Hannam et al. Twist and shout: A simple model of complete precessing black-hole-binary gravitational waveforms. 2013.
- [25] P. Ajith. Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins. *Phys. Rev. D*, 84(8):084037, 2011.
- [26] K. Postnov and L. Yungelson. The Evolution of Compact Binary Star Systems. Living Rev. Rel., 6, 2007.
- [27] M.J. Benacquista and J.M.B. Downing. Relativistic Binaries in Globular Clusters. Living Reviews in Relativity, 16(4), 2013.
- [28] Belczynski, K. and others. The Formation and Gravitational Wave Detection of Massive Stellar Black Hole Binaries. Ap. J., 2014.
- [29] Anderson, SB and Gorham, PW and Kulkarni, SR and Prince, TA and Wolszczan, A. Discovery of two radio pulsars in the globular cluster M15. *Nature*, 346:42–44, 1990.
- [30] V. Kalogera. Binary compact object inspiral: Detection expectations. Pramana, 63:673–683, 2004. 10.1007/BF02705191.
- [31] J. Abadie et al. Predictions for the rates of compact binary coalescences observable by groundbased gravitational-wave detectors. 27:173001, March 2010.
- [32] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. Astroph. J. Lett., 195:L51–L53, January 1975.
- [33] J. M. Weisberg and J. H. Taylor. Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis. 2004.

- [34] E. S. Phinney. The rate of neutron star binary mergers in the universe: Minimal predictions for gravity wave detector. 380:L17, 1991.
- [35] Ehud Nakar, Avishay Gal-Yam, and Derek B. Fox. The local rate and the progenitor lifetimes of short-hard gamma-ray bursts: Synthesis and predictions for ligo. 650:281, 2006.
- [36] Kopparapu, R.K. and Hanna, C. and Kalogera, V. and O'Shaughnessy, R. and Gonzalez, G. and Brady, P.R. and Fairhurst, S. Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events. 675(2):1459–1467, 2008.
- [37] Silverman, J.M. and Filippenko, A.V. On IC 10 X-1, the Most Massive Known Stellar-Mass Black Hole. Ap. J., 678(17), 2008.
- [38] T. Bulik, K. Belczynski, and A. Prestwich. IC10 X-1/NGC300 X-1: the very immediate progenitors of BH- BH binaries. 730:140, 2011.
- [39] Belczynski, K. and Bulik, T. and Mandel, I. and Sathyaprakash, B.S. and Zdziarski, A. and Mikolajewska, J. Cyg X-3: a Galactic double black hole or black hole-neutron star progenitor. *Ap. J.*, 764(96), 2013.
- [40] J.E. McClintock and R.A. Remillard. Black hole binaries. In Walter Lewin and Michiel van der Klis, editors, *Compact Stellar X-Ray Sources*. Cambridge University Press, Cambridge, UK, 2006.
- [41] Belczynski, K. and Bulik, T. and Fryer, C.L. High Mass X-ray Binaries: Future Evolution and Fate. 2012.
- [42] P. A. Crowther, R. Barnard, S. Carpano, J. S. Clark, V. S. Dhillon, and A. M. T. Pollock. NGC 300 X-1 is a Wolf-Rayet/black hole binary. *Mon. Not. Roy. Astron. Soc.*, pages L11+, January 2010.
- [43] V. Kalogera, C. Kim, D. R. Lorimer, M. Burgay, N. D'Amico, A. Possenti, R. N. Manchester, A. G. Lyne, B. C. Joshi, M. A. McLaughlin, M. Kramer, J. M. Sarkissian, and F. Camilo. The cosmic coalescence rates for double neutron star binaries. 601:L179–L182, 2004. Erratum-ibid. 614 (2004) L137.
- [44] C. Kim, V. Kalogera, and D. R. Lorimer. Effect of PSR J0737-3039 on the DNS Merger Rate and Implications for GW Detection. In A life with stars, 2006. astro-ph/0608280.

- [45] R. O'Shaughnessy, C. Kim, V. Kalogera, and K. Belczynski. Constraining Population Synthesis Models via Empirical Binary Compact Object Merger and Supernova Rates. 672:479–488, January 2008.
- [46] V. Kalogera, K. Belczynski, C. Kim, R. O'Shaughnessy, and B. Willems. Formation of double compact objects. *Physics Reports*, 442:75–108, April 2007.
- [47] J. Abadie et al. Search for compact binary coalescences in ligo and virgo data from s5 and vsr1. Phys. Rev. D, 82:102001, 2010.
- [48] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. Phys. Rev. D, 83:122005, 2011.
- [49] B. Abbott et al. Search of S3 LIGO data for gravitational wave signals from spinning black hole and neutron star binary inspirals. *Phys. Rev.*, D78:042002, 2008.
- [50] Chris Van Den Broeck et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Phys.Rev.*, D80:024009, 2009.
- [51] P. Ajith et al. The NINJA-2 catalog of hybrid post-Newtonian/numerical-relativity waveforms for non-precessing black-hole binaries. *Class. Quantum Grav.*, 29:124001, 2012.
- [52] Peter R. Saulson. Fundamentals of Interferometric Gravitational Wave Detectors. World Scientific, Singapore, 1994.
- [53] M. Rakhmanov, J.D. Romano, and J.T. Whelan. High-frequency corrections to the detector response and their eect on searches for gravitational waves. *Class. Quantum Grav.*, 25:184017, 2008.
- [54] The LIGO Scientific Collaboration and the Virgo Collaboration. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics*, 7:613– 619, 2013.
- [55] Adhikari, R. and Fritschel, P. and Waldman, S. Enhanced LIGO. Technical Report LIGO-T060156, 2006.
- [56] J. Smith for the LIGO Scientific Collaboration. The Path to the Enhanced and Advanced LIGO Gravitational-Wave Detectors. *Class. Quant. Grav*, 26(11):114013, June 2009.
- [57] B.F. Schutz. Networks of gravitational wave detectors and three figures of merit. Class. Quantum Grav., 28:125023, 2011.

- [58] L. Singer, L. Price, and A. Speranza. Optimizing optical follow-up of gravitational-wave candidates. 2012.
- [59] F. Marion and the Virgo Collaboration. Multi-band search of coalescing binaries applied to CITF data. In Jacques Dumarchez, editor, Proc. Rencontres de Moriond 2003, Les Arcs, 2003.
- [60] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. *accepted to Phys. Rev. D*, 2013.
- [61] L A Wainstein and V D Zubakov. Extraction of signals from noise. Prentice-Hall, Englewood Cliffs, NJ, 1962.
- [62] P. Jaranowski and A. Królak. Gravitational-Wave Data Analysis. Formalism and Sample Applications: The Gaussian Case. *Living Reviews in Relativity*, 15, 2012.
- [63] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85(12):122006, June 2012.
- [64] I. W. Harry, B. Allen, and B. S. Sathyaprakash. Stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D*, 80:104014, Nov 2009.
- [65] Gian Mario Manca and Michele Vallisneri. Cover art: Issues in the metric-guided and metricless placement of random and stochastic template banks. *Phys. Rev. D*, 81:024004, Jan 2010.
- [66] Thomas Cokelaer. Gravitational wave from inspiralling compact binaries: hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D*, 76:102004, 2007.
- [67] B. J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, 1996.
- [68] T. A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. Phys. Rev. D, 52:605, 1995.
- [69] D.A. Brown, A. Lundgren, and R. O'Shaughnessy. Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation. *Phys. Rev. D*, 86:064020, 2012.
- [70] D.A. Brown, I.W. Harry, A. Lundgren, and A.H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Phys. Rev. D*, 86:084017, 2012.

- [71] I.W. Harry, A.H. Nitz, D.A. Brown, A. Lundgren, E. Ochsner, and D. Keppel. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. *ArXiv e-prints*, 2013.
- [72] B.J. Owen and B.S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, 1999.
- [73] B.S. Sathyaprakash. Filtering post-Newtonian gravitational waves from coalescing binaries. *Phys. Rev. D*, 50:7111–7115, 1994.
- [74] C.A.K Robinson et al. A geometric algorithm for efficient coincident detection of gravitational waves. *Phys. Rev. D*, 78:062002, 2008.
- [75] C. Hanna. Searching for gravitational waves from binary systems in non-stationary data. PhD thesis, Louisiana State University, 2008.
- [76] B. Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005.
- [77] Patrick R. Brady and Stephen Fairhurst. Interpreting the results of searches for gravitational waves from coalescing binaries. 25(10):105002, 2008.
- [78] R. Biswas, P.R. Brady, J. D. E. Creighton, and S. Fairhurst. The Loudest Event Statistic: General Formulation, Properties and Applications. 26:175009, 2009.
- [79] T. Fricke et al. DC Readout Experiment in Enhanced LIGO. 2011.
- [80] M. Lorenzini (for the Virgo Collaboration). The monolithic suspension for the virgo interferometer. 27(8):084021, 2010.
- [81] Duncan A Brown. Search for gravitational radiation from black hole MACHOs in the Galactic halo. PhD thesis, University of Wisconsin–Milwaukee, 2004.
- [82] The LIGO and Virgo Collaborations. Search for gravitational waves from intermediate mass binary black holes. *Phys. Rev. D*, 85:102004, 2012.
- [83] P. R. Brady, J. D. E. Creighton, and A. G. Wiseman. Upper limits on gravitational-wave signals based on loudest events. 21:S1775–S1781, 2004.

- [84] B.P. Abbott et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run. *Phys.Rev.*, D80:047101, 2009.
- [85] B. Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. *Phys. Rev. D*, 79:122001, 2009.
- [86] I. Bartos. Frequency domain calibration error budget of LIGO instruments in S6. Technical Report LIGO-T1100071-v9, 2011.
- [87] The Virgo Collaboration. Calibration and sensitivity of the Virgo detector during its second science run. Technical report, 2011.
- [88] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel. Inspiral-merger-ringdown multipolar waveforms of nonspinning black-hole binaries using the effective-one-body formalism. *Phys. Rev. D*, 84(12):124052, December 2011.
- [89] Cannon, K. and Cariou, R. and Chapman, A. and Crispin-Ortuzar, M. and Fotopoulos, N. and Frei, M.A. and Hanna, C. and Kara, E. and Keppel, D. and Liao, L. and Privitera, S. and Searle, A. and Singer, L. and Weinstein, A.J. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. Ap. J., 748:136, 2012.
- [90] https://www.lsc-group.phys.uwm.edu/daswg/projects/gstlal.html.
- [91] http://gstreamer.freedesktop.org.
- [92] http://www.lsc-group.phys.uwm.edu/lal.
- [93] B.D. Metzger and E. Berger. What is the Most Promising Electromagnetic Counterpart of a Neutron Star Binary Merger? 746:48, 2012.
- [94] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, and et al. First low-latency LIGO+Virgo search for binary inspirals and their electromagnetic counterparts. 541:A155, May 2012.
- [95] D Buskulic, the LIGO Scientific Collaboration, and the Virgo Collaboration. Very low latency search pipeline for low mass compact binary coalescences in the LIGO S6 and Virgo VSR2 data. Class. Quantum Grav., 27(19):194013, 2010.
- [96] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein. Singular value decomposition applied to compact binary coalescence gravitational-wave signals. *Phys. Rev. D*, 82(4):044025, August 2010.

- [97] William H. Lee, Enrico Ramirez-Ruiz, and Jonathan Granot. A compact binary merger model for the short, hard grb 050509b. Ap. J., 630:L165–L168, 2005.
- [98] Ehud Nakar. Short-hard gamma-ray bursts. Phys. Rept., 442:166–236, 2007.
- [99] F. J. Virgili, B. Zhang, P. O'Brien, and E. Troja. Are All Short-hard Gamma-ray Bursts Produced from Mergers of Compact Stellar Objects? Ap. J., 727:109-+, February 2011.
- [100] H. Thomas Janka, Thomas Eberl, Maximilian Ruffert, and Chris L. Fryer. Black hole-neutron star mergers as central engines of gamma-ray bursts. Ap. J., 527:L39–L42, December 1999.
- [101] Re'em Sari and Tsvi Piran. Predictions for the very early afterglow and the optical flash. Ap. J., 520(2):641, 1999.
- [102] J.-L. Atteia and M. Boër. Observing the prompt emission of GRBs. C. R. Phys., 12:255–266, April 2011.
- [103] J. L. Racusin et al. Broadband observations of the naked-eye γ -ray burst GRB080319B. Nature, 455:183–188, September 2008.
- [104] N. Bucciantini, B. D. Metzger, T. A. Thompson, and E. Quataert. Short gamma-ray bursts with extended emission from magnetar birth: jet formation and collimation. MNRAS, 419(2):1537–1545, 2012.
- [105] Brennan Hughey. Electromagnetic follow-ups of candidate gravitational wave triggers in the recent ligo and virgo science runs. In *Gravitational Wave Physics and Astronomy Workshop* (GWPAW), 2011.
- [106] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, and et al. Implementation and testing of the first prompt search for electromagnetic counterparts to gravitational wave transients. arXiv:1109.3498v1, 2011.
- [107] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, and et al. The low-latency search for binary inspirals and their electromagnetic counterparts in ligo s6 and virgo vsr3. arXiv:1112.6005v4, 2012.
- [108] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, and et al. Swift follow up observations of candidate gravitational-wave transient events. Ap. J. (in press), 2011.

- [109] Jonah Kanner, Tracy L Huard, Szabolcs Márka, David C Murphy, Jennifer Piscionere, Molly Reed, and Peter Shawhan. Looc up: locating and observing optical counterparts to gravitational wave bursts. *Class. Quantum Grav.*, 25(18):184034, 2008.
- [110] R. Bork, R. Abbott, D. Barker, and J. Heefner. An overview of the ligo control and data acquisition system. In Hamid Shoaee, editor, Proc. 8th Int. Conf. on Accelerator and Large Experimental Physics Control Systems, pages 19–23, Menlo Park, CA, 2001. Stanford Linear Accelerator Center Technical Publications.
- [111] Alessandra Buonanno et al. Toward faithful templates for non-spinning binary black holes using the effective-one-body approach. *Phys. Rev. D*, 76:104049, 2007.
- [112] Taracchini, A. and Pan, Yi and Buonanno, A. and Barausse, E. and Boyle, M. and Chu, T., and Lovelace, G. and Pfeiffer, H.P. and Scheel, M.A. Prototype effective-one-body model for nonprecessing spinning inspiral-merger-ringdown waveforms. *Phys. Rev. D*, 86:024011, 2012.
- [113] Lijun Gou et al. The Extreme Spin of the Black Hole in Cygnus X-1. 2011.
- [114] Lijun Gou et al. A Determination of the Spin of the Black Hole Primary in LMC X-1. 701:1076–1090, 2009.
- [115] Jeffrey E. McClintock et al. The Spin of the Near-Extreme Kerr Black Hole GRS 1915+105. 652:518–539, 2006.
- [116] James F. Steiner, Jeffrey E. McClintock, and Mark J. Reid. The Distance, Inclination, and Spin of the Black Hole Microquasar H1743-322. 2011.
- [117] James F. Steiner et al. The Spin of the Black Hole Microquasar XTE J1550-564 via the Continuum-Fitting and Fe-Line Methods. Mon. Not. Roy. Astron. Soc., 416:941–958, 2011.
- [118] Lijun Gou et al. The Spin of the Black Hole in the Soft X-ray Transient A0620–00. 2010.
- [119] J.M. Miller, C.S. Reynolds, A.C. Fabian, G. Miniutti, and L.C. Gallo. Stellar-mass Black Hole Spin Constraints from Disk Reflection and Continuum Modeling. 697:900–912, 2009.
- [120] R. Shafee, J. E. McClintock, R. Narayan, S. W. Davis, L.-X. Li, and R. A. Remillard. Estimating the Spin of Stellar-Mass Black Holes by Spectral Fitting of the X-Ray Continuum. 636:L113–L116, January 2006.

- [121] J. Liu, J. E. McClintock, R. Narayan, S. W. Davis, and J. A. Orosz. Erratum: "Precise Measurement of the Spin Parameter of the Stellar-mass Black Hole M33 X-7". 719:L109, August 2010.
- [122] S. W. Davis, C. Done, and O. M. Blaes. Testing Accretion Disk Theory in Black Hole X-Ray Binaries. 647:525–538, August 2006.
- [123] L.-X. Li, E. R. Zimmerman, R. Narayan, and J. E. McClintock. Multitemperature Blackbody Spectrum of a Thin Accretion Disk around a Kerr Black Hole: Model Computations and Comparison with Observations. Astrophys. J.s, 157:335–370, April 2005.
- [124] James F. Steiner and Jeffrey E. McClintock. Modeling the Jet Kinematics of the Black Hole Microquasar XTE J1550-564: A Constraint on Spin-Orbit Alignment. 745:136, 2012.
- [125] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski. Black Hole Spin-Orbit Misalignment in Galactic X-ray Binaries. 719:L79–L83, August 2010.
- [126] P. Ajith. Gravitational-wave data analysis using binary black-hole waveforms. 25:114033, 2008.
- [127] P. Ajith and M. Favata. Comparison of post-newtonian templates for gravitational waveforms from spinning compact binaries. In preparation., 2012.
- [128] Ilana MacDonald, Abdul H. Mroue, Harald P. Pfeiffer, Michael Boyle, Lawrence E. Kidder, et al. Suitability of hybrid gravitational waveforms for unequal-mass binaries. 2012.
- [129] Mark Hannam, Sascha Husa, Frank Ohme, Doreen Muller, and Bernd Bruegmann. Simulations of black-hole binaries with unequal masses or nonprecessing spins: Accuracy, physical properties, and comparison with post-Newtonian results. *Phys. Rev.*, D82:124008, 2010.
- [130] Alexander H. Nitz, Andrew Lundgren, Duncan A. Brown, Evan Ochsner, Drew Keppel, et al. Accuracy of gravitational waveform models for observing neutron-star-black-hole binaries in Advanced LIGO. 2013.
- [131] Alessandra Buonanno, Yanbei Chen, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Phys. Rev. D*, 67:104025, 2003. Erratum-ibid. 74 (2006) 029904(E).

- [132] Yi Pan, Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. A physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single-spin binaries. *Phys. Rev. D*, 69:104017, 2004. Erratum-ibid. 74 (2006) 029905(E).
- [133] I. W. Harry and S. Fairhurst. A coherent triggered search for single spin compact binary coalescences in gravitational wave data. *Class. Quant. Grav*, 28:134008, 2011.
- [134] D. Fazi. Development of a physical template search for gravitational waves from spinning compact object binaries with LIGO. PhD thesis, Universita di Bologna, 2009.
- [135] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Phys. Rev.*, D77:062002, 2008.
- [136] K Cannon et al. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. Ap.J., 748, 2012.