

Can we hear black holes collide?

I. W. Harry

Submitted for the degree of Doctor of Philosophy
School of Physics and Astronomy
Cardiff University
March 2011

Declaration of authorship

- **DECLARATION:**

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed: (candidate) Date:

- **STATEMENT 1:**

This thesis is being submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (PhD).

Signed: (candidate) Date:

- **STATEMENT 2:**

This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references.

Signed: (candidate) Date:

- **STATEMENT 3**

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed: (candidate) Date:

Summary of thesis

The focus of this work is the attempt to detect gravitational waves emitted by compact binary coalescences (CBCs) using gravitational wave interferometers. We begin by reviewing the basic theory of gravitational waves and the methods for their detection, focusing on CBCs. We also briefly describe the laser interferometers that are being used to attempt to detect gravitational radiation.

We describe in detail the search pipeline that has been used to search for gravitational waves emitted from CBCs in data taken by the LIGO and Virgo detectors. We present the latest results of the all-sky, all-time search and electromagnetically triggered searches. We introduce a fully coherent, multi-detector analysis that can be used to search for CBC signals in coincidence with electromagnetically observed events. Using a number of signal consistency tests, including a coherent extension of the often used χ^2 test, we demonstrate that the coherent search offers an improvement in sensitivity when compared to the previous search method. Additionally we describe an extension of the coherent search that can be used to search for CBC signals where one of the components has spin. This method is well suited to searches for neutron star, black hole binaries.

We introduce a “stochastic” algorithm that can be used to create template banks in arbitrary parameter spaces of arbitrary dimension. We demonstrate this method in a search for super-massive CBCs in the mock LISA data challenge. Finally, we present the black hole hunter game, which has been widely used in outreach projects.

Contents

1	Introduction	1
2	Gravitational waves and interferometers – the basics	5
2.1	Gravitational radiation	5
2.1.1	Einstein’s equations	6
2.1.2	Linearized gravity	6
2.1.3	Gauge transformations	7
2.1.4	Gravitational waves in vacuum	8
2.1.5	The transverse traceless gauge	9
2.1.6	Gravitational wave interactions	10
2.1.7	Gravitational wave emission in linearized gravity	11
2.1.8	Power radiated as gravitational waves	15
2.2	Gravitational wave sources	16
2.2.1	Transient sources	16
2.2.2	Compact binary objects	16
2.2.3	Periodic sources	17
2.2.4	Stochastic background	18
2.3	Gravitational waves, the observational evidence	18
2.3.1	The Hulse Taylor binary (PSR B1913+16)	18
2.3.2	PSR J0737-3039	19
2.4	Gravitational wave detection with laser interferometers	20
2.4.1	A simple description of laser interferometers	20
2.4.2	Response of an interferometer to a gravitational wave	22
2.4.3	Sensitivity to different sky positions	24
2.4.4	Ground based interferometers	25
2.4.5	Interferometer sensitivities and noise sources	25
2.4.6	Non-stationary transient noise sources	29
2.4.7	Next generation ground based interferometers	29
2.4.8	Space based interferometers	29
3	Gravitational radiation emitted by a compact binary coalescence	31
3.1	Common definitions	31
3.2	Quadrupole radiation, to dominant order, from an inspiral	33
3.2.1	Time domain waveforms in the radiation frame	33
3.2.2	Energy loss in the system	34
3.2.3	Phase evolution of the system	35

3.2.4	Frequency domain waveforms	36
3.2.5	Detector response to a compact binary inspiral	38
3.2.6	Gravitational radiation from a binary neutron star merger	39
3.3	Higher order phase terms	40
3.4	Higher order amplitude terms	40
3.5	Numerical relativity	41
3.6	Waveform models for analyses	41
4	The coincidence search for compact binary coalescences in LIGO and Virgo data	43
4.1	Gravitational wave data analysis	44
4.1.1	The matched-filter	44
4.1.2	The likelihood	48
4.2	An all-sky, all-time search for compact binary coalescences	49
4.2.1	Pipeline overview	51
4.2.2	Search initialization	53
4.2.3	Generating a template bank	54
4.2.4	The matched-filter	56
4.2.5	Determining Coincidence	58
4.2.6	Second stage matched-filter	58
4.2.7	Detector characterization	61
4.2.8	Effective SNR and signal-based vetoes	63
4.2.9	False alarm rates	67
4.2.10	Simulated signals	69
4.2.11	Upper limits	72
4.2.12	Output pages	75
4.2.13	Results	75
4.3	The triggered search	77
4.3.1	The coincident triggered search pipeline	78
4.3.2	Results	79
5	A targeted coherent search for gravitational waves from compact binary coalescences	81
5.1	Coherent matched filtering	82
5.1.1	The binary coalescence waveform	83
5.1.2	Multi detector binary coalescence search	86
5.1.3	Comparison with coincident search	92
5.2	Signal consistency between detectors	93
5.2.1	Null stream consistency	94
5.2.2	Amplitude consistency	96
5.3	Coherent χ^2 tests	100
5.3.1	A general framework for χ^2 tests	100
5.3.2	The coherent bank χ^2 test	104
5.3.3	The coherent autocorrelation χ^2 test	105
5.3.4	The coherent χ^2 test	106
5.4	Implementation and performance of a coherent search	109

5.4.1	Implementation of a coherent triggered search for compact binary coalescences	109
5.4.2	Analysis of simulated data	112
5.4.3	Analysis of real data	113
5.5	Discussion	120
6	Searching for spinning compact binary coalescences	123
6.1	Spinning compact binary coalescence gravitational waveforms	125
6.2	Non precessing systems	127
6.3	The single spin PTF waveform	129
6.4	Spinning search using PTF waveforms	135
6.4.1	Formulating the PTF detection statistic	135
6.4.2	Physical freedom in the PTF SNR	139
6.4.3	PTF with no precession	140
6.4.4	SNR distribution in Gaussian noise	140
6.4.5	Remaining challenges for the coincidence PTF search	141
6.5	Triggered coherent PTF spinning search	142
6.5.1	SNR distribution in Gaussian noise	145
6.6	Identifying where the PTF search is most beneficial	146
6.7	Search method and example results	148
6.8	Discussion and generic spin systems	151
7	A stochastic template placement algorithm for gravitational wave data analysis	153
7.1	Stochastic template placement algorithm	155
7.1.1	Expected size of complete stochastic template banks	156
7.1.2	The convergence of a stochastic template bank	159
7.1.3	Computational cost of filtering templates	161
7.2	Testing the algorithm	163
7.2.1	Templates in flat spaces of different dimensions	163
7.2.2	Choice of coordinate system and convergence of template numbers	164
7.2.3	Templates on a sphere	167
7.3	Templates for gravitational wave chirps	168
7.3.1	Choice of coordinate system	169
7.3.2	Comparison of stochastic lattice with a square lattice	170
7.3.3	Efficiency of the Stochastic bank	171
7.4	Conclusions	173
8	A hierarchical search for gravitational waves from supermassive black hole binary mergers	175
8.1	Search method	177
8.1.1	Matched filtering with the F-statistic	178
8.1.2	Stochastically generated template bank	180
8.1.3	Hierarchical search technique	182
8.2	Results	183

8.3	Summary and future plans	185
9	Black hole hunter	186
9.1	“Can you hear black holes collide?”	186
9.2	The Black Hole Hunter game	188
9.3	Downloadable ringtones	190
9.4	Response to the Black Hole Hunter game	190
10	Concluding remarks	192

List of Figures

2.1	The deformation of a ring of particles due to the passage of a gravitational wave.	11
2.2	The angles that describe the relationship between the source and radiation frames.	15
2.3	The evidence for gravitational wave emission from the Hulse Taylor binary pulsar.	19
2.4	The principle components of a gravitational wave interferometer. . .	21
2.5	The angles that describe the relationship between the detector and radiation frames.	24
2.6	The various noise contributions to initial detector sensitivity.	27
2.7	Sensitivities of current detectors.	28
2.8	The predicted sensitivity of LISA.	30
3.1	An example CBC waveform in the time domain.	39
4.1	Inspiral horizon distances during S5 and VSR1.	50
4.2	A flowchart of the ihope pipeline.	52
4.3	SNR distribution at different stages of the pipeline.	59
4.4	A histogram of the ethinca distance.	60
4.5	The r^2 veto duration vs SNR.	64
4.6	The distribution of the fractional difference in effective distance. . .	65
4.7	The χ^2 test plotted against SNR.	66
4.8	Cumulative histograms of inverse false alarm rate.	70
4.9	Software injection recovery for BNS injections.	71
4.10	An example output html page.	76
4.11	The 90% rate upper limits as a function of mass.	77
4.12	The sky location error region of GRB 070201.	80
5.1	Distribution of the recovered inclination angle plotted against coherent SNR.	98
5.2	Distribution of single detector SNR plotted against coherent SNR. . .	99
5.3	The distribution of the bank χ^2 test in Gaussian data.	105
5.4	The single detector auto-correlation of a gravitational wave inspiral signal from a 1.4,1.4 solar mass binary neutron star.	106
5.5	The distribution of the auto χ^2 test in Gaussian noise.	107
5.6	The distribution of the standard χ^2 test in Gaussian noise.	108
5.7	SNR of triggers in the off-source region plotted against time in Gaussian noise.	112

5.8	SNR of triggers in the off-source region plotted against time in S4 data.	113
5.9	The distribution of the null SNR plotted against coherent SNR.	114
5.10	The distribution of single detector SNR for the more sensitive H1 and L1 detectors, plotted against coherent SNR.	115
5.11	The distribution of the bank χ^2 test in S4 data.	116
5.12	The distribution of bank and auto χ^2 test plotted against SNR.	116
5.13	The distribution of the standard χ^2 test plotted against SNR.	117
5.14	Efficiency of injection recovery.	119
6.1	The distributions of the single detector PTF SNR.	141
6.2	The distributions of the coherent PTF SNR.	146
6.3	The fraction of templates analysed by the full PTF statistic as a function of the masses.	148
6.4	The distribution of triggers found by the coherent physical template family (PTF) search.	150
7.1	The theoretical upper and lower bounds on normalized thickness of a stochastic template bank.	158
7.2	A histogram of the distribution of distances from a template to its nearest template.	158
7.3	The relationship between the covering fraction and the thickness of the bank.	161
7.4	The computational cost depends on the covering fraction.	162
7.5	As Figure 7.3 but with large boundary effect.	164
7.6	The distribution of trial points chosen uniformly in polar coordinates and the points that remain as templates after the application of the stochastic placement algorithm.	166
7.7	Template coverage when using spherical coordinates.	168
7.8	The distribution of templates placed by the hexagonal lattice algorithm.	171
7.9	The number of templates as a function of the number of random trial points.	172
7.10	The distribution of stochastically generated templates.	172
7.11	The areas of parameter space where the stochastic algorithm performs poorly.	173
9.1	The “Can you hear black holes collide” exhibition.	187
9.2	Screenshots of the Black Hole Hunter game.	188

List of Tables

4.1	Duration of LIGO and Virgo science times.	49
4.2	Marginalized upper limits and predicted rates for compact binary coalescence systems.	76
6.1	The SNR corresponding to a FAP of 10^{-10} for PTF and non spinning searches.	147
7.1	Number of templates n and fractional coverage f in a flat space. . .	165
7.2	Number of templates n and coverage f in a curved space.	169
8.1	Table showing the results of our analysis on a training dataset. . . .	184
8.2	Table showing the results of our analysis on the official challenge dataset.	184

Co-authored papers

Sections of this thesis include previously published collaborative work, as described below:

- Chapter 4 presents the method and results of searches for low mass compact binary coalescences in LIGO and Virgo data. These results were published by the LIGO and Virgo scientific collaborations in [1, 2, 3, 4, 5]. The author analysed one month of data and was one of the leading contributors to the results presented in [2].

DISCLAIMER: Chapter 4 presents results that were previously published by the LIGO and Virgo Scientific Collaborations. Some of the plots in this section, while generated by published analyses, have not previously been published. These plots have therefore not gone through the rigorous review process that the LIGO and Virgo Scientific Collaborations impose for results and plots published by those collaborations.

- Chapter 5 is taken from ‘*A targeted coherent search for gravitational waves from compact binary coalescences*’ [6]. This was a collaborative work between the author and Stephen Fairhurst.
- Much of chapter 6 is taken from ‘*A coherent triggered search for single spin compact binary coalescences in gravitational wave data*’ [7]. This was a collaborative work between the author and Stephen Fairhurst.
- Chapter 7 is taken from ‘*A stochastic template placement algorithm for gravitational wave data analysis*’ [8]. This was a collaborative work between the author, Bruce Allen and B.S. Sathyaprakash.
- Chapter 8 is taken from ‘*A hierarchical search for gravitational waves from supermassive black hole binary mergers*’ [9]. This was a collaborative work between the author, Stephen Fairhurst and B.S. Sathyaprakash.
- Chapter 9 is taken from ‘*Black Hole Hunter: the game that lets YOU search for gravitational waves*’. This is a paper that is in preparation and is being written by the Cardiff gravitational physics group of summer 2008; the author is one of the leading contributors.

Acknowledgements

I have to thank a number of people for their help and guidance on my course to writing this thesis. It has been amazing to work with so many knowledgeable and helpful collaborators!

First and foremost, I must thank Steve Fairhurst for tutoring me over the last 4 years. Steve has been an excellent and caring supervisor and friend and I wouldn't be in the situation I am now if it wasn't for him. Steve has always been capable of explaining just about any problem or concept in an understandable manner. I hope that supervising me has been as rewarding for him as it has been for me! I greatly appreciate the many hours we spent in his office while he explained to me why I was wrong. Never once was I right; if there is one thing I've learnt from the whole thesis process its that Steve is never ever wrong. I realise that I was very lucky to have had Steve supervise me through my PhD and very much hope that we can continue to work together in the future.

Next I must thank Sathyaprakash for supervising me twice as a summer student and through my masters project. While Sathya does spend 70% of his time jetting around the world, the 30% of the time when he is in the same country is extremely useful, Sathya's insight and experience in all manner of gravitational wave issues has been invaluable to me. Especially I must thank Sathya for his loyalty to me and all the members of the Cardiff group; Sathya has always ensured that we have received the credit that we deserve and created opportunities for all of us in the field of gravitational physics.

To Leonid Grishchuk, thank you for starting me on the path toward a career in gravy waves. Without your unique blend of encouragement and disparagement I would probably be on a very different path right now.

To everyone else who I have worked with in my years at Cardiff, thank you for your help, advice, beer and friendship over the years and a special thank you to everyone who shared an office with me and tolerated (to varying degrees) the pen throwing and constant fiddling. Thomas Adams, Deepak Baskaran, James Clark, Thomas Cokelear, Gerald Davies, Thomas Dent, Mark Edwards, Mark Hannam, Gareth Jones, Ioannis Kamaretsos, Duncan Macleod, Chris Messenger, Laura Nuttall, Devanka Pathak, Valeriu Predoi, Craig Robinson, Edmund Schluesel, Patricia Schmidt, Patrick Sutton, Chris Van Den Broeck, John Vietch, Jack Yu and Wen Zhao thank you all and I'm sorry that I couldn't be bothered to write an individual message to each of you. I hope I didn't forget anyone (sorry if I did)! A special thanks to Laura and Duncan for getting their priorities straight and visiting LIGO-Hanford instead of coming to my wedding!

I do have to single out Dave McKechnan (don't worry I hadn't forgotten you!)

my “partner in crime” for many years. There are people in the collaboration who still confuse the two of us! Thank you Dave for many years of support and friendship, now I have to find a new “chum” for the meetings!

Thank you to all the “real” astronomers who have worked in Cardiff during my time there. You are too numerous to list but I greatly value and appreciate our friendships. If nothing else your scepticism of the existence of gravity-waves has strengthened my debating skills.

During my PhD it has been my privilege to work with many outstanding people in the LIGO and Virgo collaborations. There are too many people to list individually here but a special thank you to Bruce Allen for collaborating with me on the stochastic placement algorithm. Also a big thank you to Marie-Anne Bizouard, Kipp Cannon, Collin Capano, Jessica Clayton, Diego Fazi, Nick Fotopolous, Chad Hanna, Andrew Lundgren, Drew Keppel, Frederique Marion, Lucia Santamaria, Ruslan Vaulin, Alan Weinstein and John Whelan for many interesting conversations, debates and debacles in various pubs and bars across the globe.

I must thank my thesis examination board, Alberto Vecchio, Mark Hannam and Ant Whitworth for the time and effort spent in reading through this thesis and in providing me with useful feedback and comments. Also thank you for giving me a harrowing 3 hours of viva!

I have to thank all of my family (including the in-laws), for their longstanding and continued support. Also thank you for feigning interest when talking about what I do, don’t worry your eyes don’t gloss over too much! Maybe now that I earn more than you I will buy my share of the beer in the pub! A special thank you, of course, has to go to my parents.

Finally, I have to thank Sarah to whom this thesis is dedicated. I used up my quota of sappy things to say at our wedding last year so it will simply suffice to say that without you I would be nothing and have nothing. I love you always.

Conventions

We will follow the standard convention of using latin indices to denote spatial vectors ($x=1,y=2,z=3$) and greek indices to denote space-time vectors ($t=0,x=1,y=2,z=3$) and follow the Einstein summation convention

$$a^\mu b_\mu = \sum_{\mu=0}^3 a^\mu b_\mu \quad (1)$$

. Unless stated otherwise we always use the assumption that $c = G = 1$. To simplify partial differentials we set:

$$\partial_\mu = \frac{\partial}{\partial x^\mu} \quad (2)$$

The Cristoffel symbol is given by:

$$\Gamma_{\mu\nu}^\rho = \frac{1}{2} g^{\rho\sigma} (\partial_\mu g_{\sigma\nu} + \partial_\nu g_{\sigma\mu} - \partial_\sigma g_{\mu\nu}) \quad (3)$$

The Riemann tensor is defined as:

$$R_{\nu\rho\sigma}^\mu = \partial_\rho \Gamma_{\nu\sigma}^\mu - \partial_\sigma \Gamma_{\nu\rho}^\mu + \Gamma_{\alpha\rho}^\mu \Gamma_{\nu\sigma}^\alpha - \Gamma_{\alpha\sigma}^\mu \Gamma_{\nu\rho}^\alpha \quad (4)$$

The Ricci tensor is:

$$R_{\mu\nu} = R_{\mu\alpha\nu}^\alpha \quad (5)$$

and the Ricci scalar is:

$$R = g^{\mu\nu} R_{\mu\nu}, \quad (6)$$

where $g^{\mu\nu}$ denotes the metric tensor of 4 dimensional spacetime.

Chapter 1

Introduction

At the beginning of the 20th century Albert Einstein published a series of papers, formulating his general theory of relativity, which revolutionized humanity's understanding of gravity. [10, 11, 12]. This theory suggests that gravity should not be considered a force, as in classical mechanics, but instead as an effect of the geometry of 4-dimensional space-time. It might appear that bodies are accelerated by a gravitational force but instead Einstein's theory suggests such bodies are travelling along straight line paths, geodesics, in space-time. In most situations, general relativistic gravity behaves identically to Newtonian gravity. However, some predictions of general relativity differ significantly.

One of the predictions of general relativity that differed from classical mechanics was that photons interact with gravity. This was also the first new prediction to be experimentally verified. Photons passing through a strong gravitational field might be deflected noticeably. This prediction was verified by Arthur Stanley Eddington in 1919 [13]. He observed the positions of stars in the Hyades cluster during a solar eclipse. These stars, which appeared very near to the Sun's position in the sky at the time of the eclipse, were observed to be shifted slightly from their normally observed positions. This shift was in agreement with the predictions of general relativity. General relativity has since passed every experimental test to which it has been subjected.

Another prediction of Einstein's theory is the existence of gravitational waves. This gravitational radiation is emitted by accelerating masses, in a similar manner to how accelerated charged particles emit electromagnetic waves. There is strong observational evidence of the emission of this gravitational radiation in binary systems that are losing energy, almost certainly through gravitational wave emission, as we describe in section 2.3.

Scientists have been trying to directly observe gravitational radiation for the

last 50 years. However, even the strongest sources of gravitational radiation would produce waves that, by the time they reach Earth, would induce strains¹ no bigger than 10^{-22} . Thus, it is not surprising that a direct observation of a gravitational wave has yet to be made. Nevertheless, there has been excellent progress towards gravitational wave astronomy over recent years. The first generation of large scale gravitational wave interferometers reached unprecedented sensitivities and have undertaken extended science runs. The U.S. Laser Interferometer Gravitational-wave Observatory (LIGO) [14], the French–Italian Virgo [15] and the German–British GEO600 [16] detectors now form a collaborative network of interferometers.

The data from LIGO’s fifth science run (S5) and Virgo’s first science run (VSR1) has been analyzed for gravitational waves from compact binary coalescence [3], stochastic background [17], unmodelled burst [18] and pulsar [19] sources. LIGO’s sixth science run (S6) and Virgo’s second and third science runs (VSR2 and VSR3) ended in October 2010 and yielded the most sensitive data yet taken; the analysis of this data is ongoing. In the meantime, the detectors are being upgraded to their advanced configurations [20, 21, 22], with the expectation of a ten-fold improvement in sensitivity. With these sensitivities, it is expected that gravitational waves will be observed regularly [23]. Furthermore, with a proposed advanced detector in Japan [24], a possible detector in Australia [25], and 3rd generation detectors on the horizon [26], the future of gravitational wave astronomy is promising. There is even an ambitious plan to launch an interferometer into space, called the Laser Interferometer Space Antenna (LISA) [27].

As the gravitational wave community matures it is essential that a relationship is built between gravitational wave (GW) and electromagnetic (EM) astronomers. The GW emission from a source is likely to provide complementary information to emission in various EM bands, and a joint observation is significantly more likely to answer outstanding astrophysical questions. Already this relationship is beginning to mature. A number of EM transients have already been followed up in GW data [5, 4, 28]. Additionally, infrastructure is also being put in place to allow for EM follow-up of GW observations [29].

One of the most promising sources of gravitational waves are compact binary coalescences (CBCs). As two compact objects orbit each other they will lose energy due to gravitational radiation. This loss in energy will cause the orbit to decay and the objects will eventually merge. As the objects near merger the energy

¹Strain is defined as a fractional change in length. For a gravitational wave detector with 4km arms, a strain of 10^{-22} would translate to an variation in length of 10^{-17} m.

release and thus rate of decay of the orbit also increases. The gravitational wave signal emitted by such a merger is often described as a “chirp”; the amplitude and frequency of the signal increases rapidly before the component bodies merge. Searches for CBC signals originating at any time, from any direction and with any orientation have been carried out in data from S5 and VSR1 [1, 2, 3].

CBCs are also an ideal candidate for joint GW-EM astronomy. CBCs containing at least one neutron star (NS) are expected to emit electromagnetically. Specifically, binary neutron stars (BNS) and neutron star–black hole binary (NSBH) mergers are the preferred progenitor model for the short gamma-ray burst (GRB) [30, 31]. It is also possible that these mergers will be observable electromagnetically as orphan afterglows [30], optical [32] or radio transients [33]. Since GRBs are well localized both in time and on the sky by EM observations, the corresponding GW search can be simplified by reducing the search volume relative to an all-sky, all-time search. Targeted searches for CBC waveforms associated to short GRBs were performed using data from S5 and VSR1 [5, 4].

The focus of this thesis is the analysis of data taken from gravitational wave interferometers to try to detect CBCs. We discuss a number of aspects of this effort in the various chapters. After giving some necessary background theory we describe the existing all-sky search and targeted search in detail. These searches use a “coincidence” technique in which data from each of the detectors is analysed separately before searching for any potential signal that is seen in coincidence between the detectors. We then describe an alternative technique, the coherent search method. We demonstrate that this method offers an improvement in sensitivity over coincidence and apply the technique to the targeted search.

Current search methods only utilize template waveforms that do not include spin effects. This can mean that such searches are not sensitive to certain systems where spin has a major effect on the dynamics. We describe how a search can be performed using template waveforms that include spin, but no precessional effects. We also discuss how to perform a search using single spin templates with precession and extend this search to use the coherence technique.

The current searches use a geometrical algorithm to create template banks to cover the desired range of masses. However, it is not clear how to create a template bank in higher dimensional parameter spaces, such as for the case of spinning searches where templates must cover the additional spin parameters as well as the masses. We present a “stochastic” template placement algorithm, which can be used to create template banks in arbitrary parameter spaces of arbitrary dimension. We apply this method to the case of searches for super massive black

hole binaries with LISA, which requires a 5 dimensional template placement.

The layout of this thesis is as follows. In chapter 2 we give a very brief introduction to gravitational wave theory and gravitational wave interferometers. In chapter 3 we discuss the nature of gravitational wave emission from CBCs and in chapter 4 we discuss the current, coincidence, searches for CBC signals. In chapter 5 we discuss a “coherent” search method, which could be used to search for CBC signals associated with GRBs. In chapter 6 we discuss methods for detecting CBC signals incorporating spin. In chapter 7 we discuss a “stochastic” template placement algorithm, applicable to any parameter space. This method is applied to searches for super-massive CBC signals in chapter 8. Finally, in section 9 we discuss the “black hole hunter” game and how it has been used in outreach projects around the world.

Chapter 2

Gravitational waves and interferometers – the basics

The main focus of this thesis is the effort to detect compact binary coalescences (CBCs) through gravitational wave observations. Before exploring this topic, it is useful to briefly discuss gravitational waves, their emission mechanism and the large interferometric detectors that are being used to attempt to detect them. In this chapter we give a broad overview of these subjects, but it is not intended to be a detailed description of the concepts discussed herein. If the reader wishes a comprehensive introduction to gravitational wave theory we would refer them to [34, 35, 36]. For a detailed description of how modern interferometric gravitational wave detectors operate we refer the reader to [37].

In this chapter we begin by introducing the basic concepts of gravitational radiation in section 2.1. In section 2.2 we briefly discuss astrophysical phenomena that may emit gravitational waves that are powerful enough to detect on the Earth. In section 2.3 we discuss the observational evidence for gravitational wave emission from binary systems. Finally, in section 2.4 we discuss gravitational wave interferometric detectors.

2.1 Gravitational radiation

In this section we will give a brief introduction to gravitational radiation. We will begin with Einstein's equations and describe how the existence of gravitational waves can be inferred. We will discuss how gravitational waves interact with matter and describe the mechanism by which gravitational waves are generated.

This overview of gravitational radiation is included to provide a foundation to many of the later chapters, but is far from a comprehensive description of the

subject. For readers seeking more information, we refer them to [34, 35, 36], from which we drew heavily when writing this chapter.

2.1.1 Einstein's equations

Einstein's equations are a set of 10 equations that underpin Einstein's theory of general relativity. They describe gravity in terms of the curvature of spacetime due to matter and energy. In the standard tensor notation they are given as

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi T_{\mu\nu} \quad (2.1)$$

where $R_{\mu\nu}$ is the Ricci tensor, R the Ricci scalar and $g_{\mu\nu}$ the metric tensor of 4 dimensional spacetime as defined in the opening conventions. $T_{\mu\nu}$ is the stress energy tensor, which describes the density and flux of matter (or energy) and momentum. The components of this tensor can be understood in the following way [35]:

- T_{00} is the energy (or relativistic mass) density.
- $T_{0i} = T_{i0}$ is the energy flux in the i th direction, or the density of i -momentum.
- T_{ij} is flux of i momentum in the j direction.

We note that Einstein's equations describe 10 equalities, rather than the 16 apparent. This is because $R_{\mu\nu}$, $g_{\mu\nu}$ and $T_{\mu\nu}$ are all symmetric.

2.1.2 Linearized gravity

To simplify Einstein's equations we can consider the case where there is only a *weak* gravitational field. The space time metric $g_{\mu\nu}$ can then be written as the sum of the Minkowski tensor $\eta_{\mu\nu}$ and some small perturbation due to the gravitational field, $h_{\mu\nu}$

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (2.2)$$

To describe the gravitational field as weak we demand that the magnitude of all components of the perturbation are much less than unity

$$|h_{\mu\nu}| \ll 1. \quad (2.3)$$

Einstein's equations can then be linearized by discarding all terms that are second order or higher in $h_{\mu\nu}$. It is straightforward to show that the left side of Einstein's

equations are simplified in the weak field limit according to

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{1}{2}(\partial^\sigma\partial_\mu\bar{h}_{\sigma\nu} - \partial^\sigma\partial_\sigma\bar{h}_{\mu\nu} + \partial_\nu\partial^\alpha\bar{h}_{\mu\alpha} - \eta_{\mu\nu}\partial^\alpha\partial^\beta\bar{h}_{\alpha\beta}) \quad (2.4)$$

where $\bar{h}_{\mu\nu}$ is defined as the “trace reverse” of $h_{\mu\nu}$

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h^\alpha{}_\alpha. \quad (2.5)$$

We note that the indices of $h_{\mu\nu}$ can be raised or lowered by a multiplication with $\eta_{\mu\nu}$

$$h^{\alpha\beta} = \eta^{\alpha\mu}\eta^{\beta\nu}h_{\mu\nu}, \quad (2.6)$$

this will be used in many of the derivations that follow.

2.1.3 Gauge transformations

At this point it is useful to consider the gauge freedoms present in the weak field Einstein’s equations. We can apply small coordinate translations to the perturbed spacetime while still keeping the metric perturbations small. Consider for example the coordinate transformation

$$x'^\mu = x^\mu + \xi^\mu(x). \quad (2.7)$$

The metric in the new coordinate system will transform according to

$$g'_{\mu\nu}(x') = \frac{\partial x^\rho}{\partial x'^\mu} \frac{\partial x^\sigma}{\partial x'^\nu} g_{\rho\sigma}(x) = \eta_{\mu\nu} + h'_{\mu\nu}. \quad (2.8)$$

Any coordinate transformation of this form is valid as long as the weak field theory still holds in the new coordinates

$$|h'_{\mu\nu}| \ll 1. \quad (2.9)$$

In addition to small coordinate translations, Lorentz rotations of the coordinate system are also permitted

$$x'^\mu = \Lambda^\mu{}_\nu x^\nu. \quad (2.10)$$

Where $\Lambda^\mu{}_\nu$ must satisfy

$$\Lambda^\mu{}_\rho \Lambda^\sigma{}_\nu \eta_{\rho\sigma} = \eta_{\mu\nu} \quad (2.11)$$

and is independent of x . Under this transformation the metric becomes

$$g'_{\mu\nu}(x) = \eta_{\mu\nu} + \Lambda_\mu{}^\rho \Lambda_\nu{}^\sigma h_{\rho\sigma}(x). \quad (2.12)$$

These gauge freedoms can be utilized to choose a coordinate system that will help to express the weak field equations in a simplified form. By examining equation (2.4) one can see that it would be simplified if we were to transform to a set of coordinates in which

$$\partial^\nu \bar{h}_{\mu\nu} = 0. \quad (2.13)$$

It can be shown [36] that one can always perform a coordinate transformation of the kind described in equation (2.7) to achieve this condition. In this *Lorentz gauge* Einstein's equations are simplified to

$$\square \bar{h}^{\mu\nu} = -16\pi T^{\mu\nu}. \quad (2.14)$$

2.1.4 Gravitational waves in vacuum

Let us now consider the weak-field equation far from any source of mass or energy, such that $T_{\mu\nu} = 0$. This simplifies equation (2.14) to

$$\square \bar{h}^{\mu\nu} = 0, \quad (2.15)$$

which is the well known wave equation. This is solved by the plane wave equation

$$\bar{h}^{\mu\nu} = A^{\mu\nu} e^{ik_\alpha x^\alpha}, \quad (2.16)$$

where $A^{\mu\nu}$ is a matrix with constant components and k_α is the wave vector $k_\alpha = (\omega, k_i)$, satisfying

$$k^\alpha k_\alpha = 0. \quad (2.17)$$

Thus, the propagation of gravitational radiation as plane waves is permissible in Einstein's general relativity.

From equation (2.17), it can be seen that $\omega^2 = |k_i|^2$. As we have generically $|k_i| = \omega/v$ we can see that gravitational radiation must propagate at a speed $v = 1$. Therefore, this implies that gravitational waves will travel at the speed of light. Additionally, applying the Lorentz gauge condition given in equation (2.13) to the wave equation gives $k_\mu A^{\mu\nu} = 0$. This implies that gravitational waves are transverse.

2.1.5 The transverse traceless gauge

We have made use of the Lorentz gauge to demonstrate that gravitational radiation will propagate in a vacuum as transverse plane waves at the speed of light. There are however, further gauge freedoms that can be used to further simplify the form of $h_{\mu\nu}$. These gauge freedoms can always be used to choose a coordinate system [36, 35] in which

$$\bar{h}^{0i} = 0 \tag{2.18a}$$

$$\bar{h}^{\mu}_{\mu} = 0. \tag{2.18b}$$

The first of these conditions, combined with the Lorentz gauge, will also imply that \bar{h}^{00} will have no time or spacial dependence. While equation 2.15 does allow there to be constant terms in $\bar{h}^{\mu\nu}$, for the case of gravitational waves these can be ignored and we can treat $\bar{h}^{00} = 0$. The second condition serves to set the trace of $\bar{h}^{\mu\nu}$ to 0. In this case $\bar{h}^{\mu\nu} = h^{\mu\nu}$. This is commonly referred to as the “transverse traceless” (TT) gauge.

Finally, consider a gravitational wave propagating in the z direction, $h^{\mu\nu}$ can be written in the TT gauge as

$$h^{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_{\times} & 0 \\ 0 & h_{\times} & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \tag{2.19}$$

Where we have defined

$$h_+ = A_+ \cos(\omega t - \omega z + \phi_0) \tag{2.20}$$

$$h_{\times} = A_{\times} \cos(\omega t - \omega z + \phi_0). \tag{2.21}$$

It can be seen that the metric perturbation $h_{\mu\nu}$ is now expressed in terms of only two independent components. A_+ and A_{\times} are the complex amplitudes of the two components. ϕ_0 is some constant phase offset. Note that here we consider only a monochromatic gravitational wave; however a gravitational wave with a time dependant frequency, or a number of coincident gravitational waves are also a solution.

Thus, despite the fact that Einstein’s equations consist of 10 equalities, we have demonstrated that gravitational radiation can be expressed in terms of only

two independent components. These components are referred to as the + (plus) and \times (cross) components of gravitational radiation.

2.1.6 Gravitational wave interactions

In equation (2.19) we demonstrated that, by fully utilizing the available gauge freedom to set the coordinate system, a gravitational wave can be expressed in terms of two independent components. Here we examine how a passing gravitational wave will distort spacetime.

We begin by considering the effect of a gravitational wave passage on a particle as observed in the transverse traceless coordinate system. It can be shown [36] that in the TT coordinate system a particle initially at rest will remain at rest during the passage of a gravitational wave. However, this is simply an effect of the TT gauge; we have chosen a coordinate system such that this will be the case. It would be more physically meaningful to calculate a coordinate invariant quantity such as how the proper distance changes during the passage of a gravitational wave. To do this let us first write down the equation defining the interval ds^2 in the TT gauge (with the direction of propagation of the wave the z direction).

$$ds^2 = -dt^2 + dz^2 + (1 + A_+ \cos(\omega t - \omega z + \phi_0)) dx^2 + (1 - A_+ \cos(\omega t - \omega z + \phi_0)) dy^2 + 2A_\times \cos(\omega t - \omega z + \phi_0) dx dy. \quad (2.22)$$

If we then consider two particles at positions $(x_1, y_1, 0)$ and $(x_2, y_2, 0)$ the proper distance between them at some arbitrary time t would be given by

$$ds^2 = (1 + A_+ \cos(\omega t - \omega z)) (x_1 - x_2)^2 + (1 - A_+ \cos(\omega t - \omega z)) (y_1 - y_2)^2 + 2A_\times \cos(\omega t - \omega z) (x_1 - x_2) (y_1 - y_2). \quad (2.23)$$

It is clear to see that this quantity is time dependent. Take for example the simplification that $y_2 - y_1 = 0$ this equation would become

$$ds^2 = (1 + A_+ \cos(\omega t - \omega z)) (x_1 - x_2)^2, \quad (2.24)$$

and the proper distance is

$$ds \approx \left(1 + \frac{1}{2} A_+ \cos(\omega t - \omega z)\right) (x_1 - x_2), \quad (2.25)$$

where we have used the fact that in the weak field limit $A_+ \ll 1$. To visualize

this, we demonstrate in Figure 2.1 how the proper distances between a ring of particles would deform in the presence of a gravitational wave containing only the $+$ and \times components respectively. We can see that the $+$ and \times components are interchangeable under a 45° rotation.

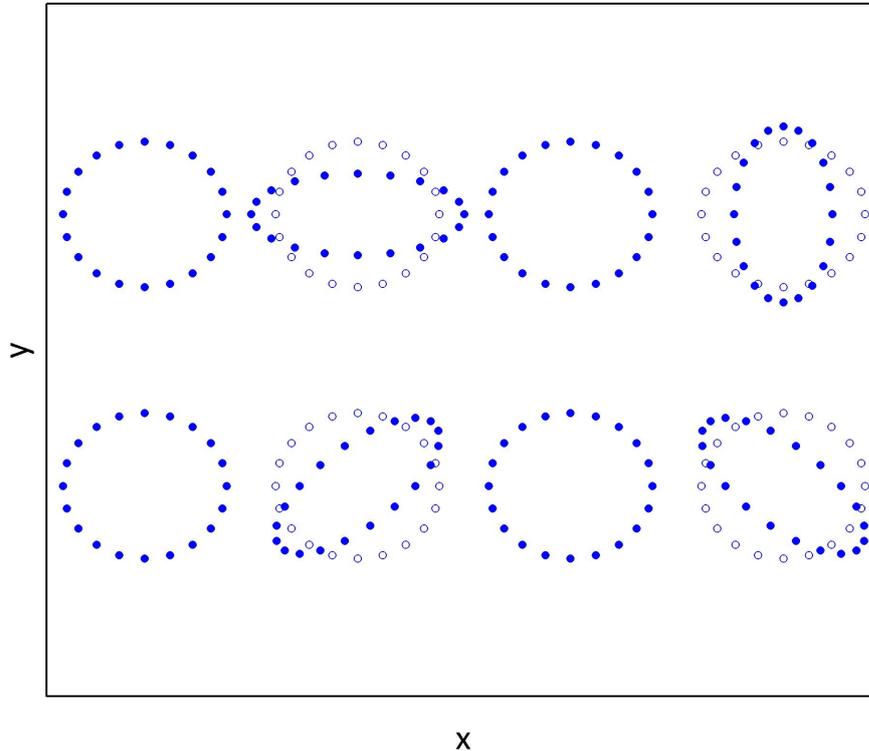


Figure 2.1: The deformation of a ring of particles due to the passage of a $+$ polarized gravitational wave (top) and a \times polarized wave (bottom), propagating in the z direction. From left to right we show the $0, \frac{\pi}{2}, \pi$ and $\frac{3\pi}{2}$ phases. The unfilled circles denote the positions of the particles at the 0 phase. This Figure was originally published in [38].

2.1.7 Gravitational wave emission in linearized gravity

In this section we investigate how gravitational waves would be emitted. We show that, to leading order, gravitational waves are a form of quadrupole radiation. Finally, we give the form of h_+ and h_\times , to dominant order, for a generic gravitational wave source. In this section we follow very closely the approach taken in [36], we refer the reader to this text for more detail, including a more detailed study of the

higher order terms.

We again begin with Einstein's equations in the weak-field limit,

$$\square \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}. \quad (2.26)$$

Consider a source at the origin of the coordinates and a distant observer displaced by \mathbf{n} from the origin. The distance to the observer D is given by $D = |\mathbf{n}|$. It is shown in [36] that we can find a solution to Einstein's equations by using the Green's function for the wave equation, this will give us

$$\bar{h}_{\mu\nu}(t, \mathbf{n}) = 4 \int d^3x \frac{1}{|\mathbf{n} - \mathbf{x}|} T_{\mu\nu}(t - |\mathbf{n} - \mathbf{x}|, \mathbf{x}). \quad (2.27)$$

Where \mathbf{x} describes the spatial coordinates of the mass elements generating the gravitational radiation.

We wish to simplify equation (2.27) to understand under what conditions gravitational waves would be emitted. As long as an observer is distant from the source such that $|\mathbf{n}| \gg |\mathbf{x}|$, we can expand the quantity $|\mathbf{n} - \mathbf{x}|$ as

$$|\mathbf{n} - \mathbf{x}| = D - \mathbf{x} \cdot \hat{\mathbf{n}} + O\left(\frac{r^2}{D}\right). \quad (2.28)$$

Where r is the typical size of the source. Thus we rewrite equation (2.27) as

$$\bar{h}_{\mu\nu}(t, \mathbf{n}) = \frac{4}{D} \int d^3x T_{\mu\nu}(t - D, \mathbf{x}). \quad (2.29)$$

Here we have kept only dominant order terms in $|\mathbf{n} - \mathbf{x}|$. This can be simplified further by using the TT gauge. In this gauge $h_{0\mu} = 0$ and h_{00} is constant. We can therefore drop the timelike components and restrict ourselves to

$$h_{ij}^{TT}(t, \mathbf{n}) = \frac{4}{D} S_{ij}(t - D). \quad (2.30)$$

Where the first moment of the stress tensor T^{ij} , is defined as

$$S^{ij}(t) = \int d^3x T^{ij}(t, \mathbf{x}). \quad (2.31)$$

To gain insight into the physical meaning of S_{ij} it is useful to re-express it in terms of T^{00} (the energy density) and T^{0i} (the momentum density) instead of T^{ij} . We

therefore define

$$M(t) = \int d^3x T^{00}(t, \mathbf{x}), \quad (2.32a)$$

$$M^i(t) = \int d^3x T^{00}(t, \mathbf{x}) x^i, \quad (2.32b)$$

$$M^{ij}(t) = \int d^3x T^{00}(t, \mathbf{x}) x^i x^j, \quad (2.32c)$$

where these three quantities are respectively, the mass monopole, mass dipole, and mass quadrupole moments. There are of course higher order terms but we don't consider them here. We note that in linearized gravity the mass monopole is interpreted as the total mass of the source. We can also write the momentum monopole and dipole as

$$P^i = \int d^3x T^{0i}(t, \mathbf{x}), \quad (2.33a)$$

$$P^{ij} = \int d^3x T^{0i}(t, \mathbf{x}) x^j. \quad (2.33b)$$

Again we note here that in linearized gravity P^i is interpreted as the total momentum of the system. To express S^{ij} in terms of the mass multipole moments we make use of a consequence of the Lorentz gauge condition, $\partial_\mu T^{\mu\nu} = 0$, and the divergence theorem to obtain the following identities [36]

$$\dot{M} = 0, \quad (2.34a)$$

$$\dot{M}^i = P^i, \quad (2.34b)$$

$$\dot{M}^{ij} = P^{i,j} + P^{j,i}, \quad (2.34c)$$

$$\dot{P}^i = 0, \quad (2.34d)$$

$$\dot{P}^{i,j} = S^{ij}, \quad (2.34e)$$

$$(2.34f)$$

From these equations, as well as the fact that $S^{ij} = S^{ji}$ we obtain the identity

$$\ddot{M}^{ij} = 2S^{ij}. \quad (2.35)$$

This shows that, as S^{ij} is the leading term in the expansion of h_{ij} , gravitational wave radiation has no monopole or dipole component and the leading order term is the quadrupole moment.

If we were to evaluate h_{ij} to higher order in $|\mathbf{n} - \mathbf{x}|$ we would discover that

gravitational waves have higher order components coming from the mass octopole and higher order moments. We refer the reader to [35, 36] for a discussion of these higher order terms. We will, in this work, restrict ourselves to considering only the dominant quadrupole term. The equation for h_{ij} is then written as

$$h_{ij}^{TT}(t) = \frac{2}{D} \ddot{M}_{ij}(t - D). \quad (2.36)$$

Thus, the h_+ and h_\times components are given by

$$h_+ = \frac{1}{D} (\ddot{M}_{11} - \ddot{M}_{22}) \quad (2.37)$$

$$h_\times = \frac{2}{D} \ddot{M}_{12}. \quad (2.38)$$

To obtain these equations we have assumed that \ddot{M}_{ij} is evaluated in the transverse-traceless frame, with the gravitational radiation propagating along the z direction. This is commonly called the “radiation frame” whose basis vectors we denote as \hat{x}_i . However, it is often easier to evaluate \ddot{M}_{ij} in a different frame, for example when considering a binary system it is useful to define a “source frame” in which the binaries orbit in the x' - y' plane. We denote the basis vectors in the source frame as \hat{x}'_i .

To understand the relationship between the source and radiation frames let us consider this geometrically. Two angles are needed to rotate the source frame into a frame in which the radiation is travelling in the z direction. This is illustrated in Figure 2.2. We note that normally there would be three rotation angles involved in such a transformation. Here the third rotation would correspond to the freedom to rotate the x - y plane of the radiation frame. For now, we arbitrarily set this angle to 0 for simplicity. This condition ensures that the x direction in the radiation frame points in the plane defined by the x' and y' directions in the source frame. Explicitly, the radiation frame is related to the source frame as

$$\begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} = \begin{pmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \iota & -\sin \iota \\ 0 & \sin \iota & \cos \iota \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad (2.39)$$

$$\begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} = R \begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad (2.40)$$

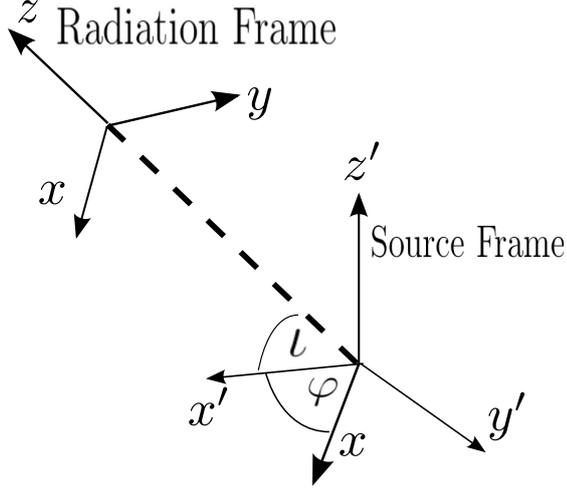


Figure 2.2: An illustration of the angles that describe the relationship between the source and radiation frames. Figure originally published in [38].

where the angles, (ι, φ) are defined in Figure 2.2. M'_{ij} is then related to M_{ij} by

$$M'_{ij} = R_{ki} M_{kl} R_{lj}. \quad (2.41)$$

Thus, in the radiation frame we can write the h_+ and h_\times components of the gravitational radiation as

$$\begin{aligned} h_+ = \frac{1}{D} & \left[\ddot{M}'_{11} (\cos^2 \varphi - \sin^2 \varphi \cos^2 \iota) \right. \\ & + \ddot{M}'_{22} (\sin^2 \varphi - \cos^2 \varphi \cos^2 \iota) \\ & - \ddot{M}'_{33} \sin^2 \iota - \ddot{M}'_{12} \sin 2\varphi (1 + \cos^2 \theta) \\ & \left. + \ddot{M}'_{13} \sin \varphi \sin 2\iota + \ddot{M}'_{23} \cos \varphi \sin 2\iota \right] \end{aligned} \quad (2.42a)$$

$$\begin{aligned} h_\times = \frac{1}{D} & \left[(\ddot{M}'_{11} - \ddot{M}'_{22}) \sin 2\varphi \cos \iota + 2\ddot{M}'_{12} \cos 2\varphi \cos \iota \right. \\ & \left. - 2\ddot{M}'_{13} \cos \varphi \sin \iota + 2\ddot{M}'_{23} \sin \varphi \sin \iota \right]. \end{aligned} \quad (2.42b)$$

2.1.8 Power radiated as gravitational waves

Finally, we wish to examine the total power that would be radiated by a source emitting gravitational waves. This can be calculated by determining an effective stress energy tensor due to the gravitational radiation itself. The energy flux in a given direction can then be determined. This flux can be integrated over all sky angles to give the total energy radiated. The power radiated by a gravitational

wave is then given as [36]

$$\frac{dE}{dt} = \frac{1}{16\pi} \int d\Omega \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle. \quad (2.43)$$

The angle brackets represent that this quantity is averaged over time. In [36] it is demonstrated that this integral can be evaluated to give the following equation for the emitted gravitational wave power

$$\frac{dE}{dt} = \frac{1}{5} \langle \ddot{M}'_{ij} \ddot{M}'^{ij} - \frac{1}{3} (\delta^{kl} \ddot{M}'_{kl})^2 \rangle. \quad (2.44)$$

2.2 Gravitational wave sources

In this section we briefly discuss the different classes of physical sources that are likely to generate gravitational waves of sufficient amplitude to be detectable by current or theorized gravitational wave detectors. As this work focuses on gravitational wave data analysis from compact binary objects we discuss this source more comprehensively in chapter 3.

2.2.1 Transient sources

A transient or *burst* source is the name given to an event that releases a large amount of gravitational energy over a very short period of time, typically less than a few seconds. Astrophysical events that are believed to result in a burst of gravitational waves include gamma ray bursts and supernovae explosion as well as the final stages of a coalescing binary [39, 40, 41, 42].

When performing a search for burst sources nothing is assumed about the form the gravitational radiation will take, only that the gravitational signal will be of short duration. Therefore, transient gravitational wave sources could be detectable even if the source is an object we haven't discovered or predicted as of yet.

For details of recent searches for burst sources see [28, 43, 44, 45].

2.2.2 Compact binary objects

We will demonstrate in chapter 3 that two compact binary objects orbiting around each other will emit energy in gravitational waves. This emission of energy will cause the orbital radius of the system to decay until the two objects eventually coalesce. When the system nears coalescence the power and frequency of the emitted gravitational energy will increase, producing a distinctive chirp-like signal.

A binary neutron star system at a distance of 10Mpc would produce an observed gravitational wave amplitude, near merger, of order 10^{-21} , as we derive in section 3.2.6.

Searches for compact binary coalescences are similar to unmodelled burst searches. However, the fact that the form of the gravitational radiation can be predicted allows a more sensitive search to be performed. As we will explore in later chapters, knowing the form of the signal that is being searched for allows powerful matched filtering techniques and signal consistency tests to be used in the attempt to detect such signals.

The majority of this thesis is devoted to the quest to detect compact binary coalescence (CBC) signals.

2.2.3 Periodic sources

A periodic source is a source that emits a continuous, almost monochromatic gravitational wave. These sources should be present throughout the operational lifetime of a detector, so the greater the observation time, the better the sensitivity to periodic sources becomes. The periodic source of most interest to gravitational wave astronomers is that of a rapidly spinning, slightly spherically asymmetric neutron star

Such spinning neutron stars will lose energy and spin down over time [46]. This energy loss is due to a number of different mechanisms, including emission of gravitational radiation [47], though the fraction of energy emitted as gravitational waves is unclear. To motivate gravitational wave searches for spinning down pulsars we can consider the Crab pulsar. This pulsar provides the best opportunity to detect continuous gravitational waves with current gravitational wave interferometers [47]. In the extremal case where all the spin down energy is radiated as gravitational radiation the Crab pulsar would produce a gravitational strain amplitude of 1.4×10^{-24} [47]. Although this is below the detector noise floor, the source can be observed over a timescale of months or years, therefore it is within the observational limits of the detectors.

Indeed, in [47] an extended period of observation with the LIGO detector placed upper limits on gravitational radiation from the Crab pulsar, which imply that less than 6% of the spin down energy of the Crab is emitted as gravitational waves. For more details on periodic gravitational wave sources, and the current results of gravitational wave searches for these sources we refer the reader to [46, 19, 48, 49] and references therein.

2.2.4 Stochastic background

As the universe has background electromagnetic radiation, the cosmic microwave background, so it is believed that at the very earliest moments of the universe a stochastic background of gravitational waves was generated [50]. As well as the cosmological gravitational wave background, there is expected to be stochastic radiation emitted from indistinguishable gravitational wave sources, such as the large number of white dwarf binaries in our galaxy. It is difficult however to estimate how strong this *stochastic* radiation is likely to be and whether or not it will be observable. Nevertheless, gravitational wave astronomers perform searches in their data for stochastic gravitational wave radiation.

For more details on stochastic gravitational wave sources and searches we refer the reader to [50, 17, 51, 52] and references therein.

2.3 Gravitational waves, the observational evidence

Although it was mentioned in the opening paragraphs that we have not yet directly observed a gravitational wave, we do have strong observational evidence for their existence. In this section we will briefly examine this evidence.

2.3.1 The Hulse Taylor binary (PSR B1913+16)

The Hulse Taylor binary or PSR B1913+16 is a binary neutron star system, which is named after its discoverers Joseph Taylor and Russell Hulse. It is the most famous observational evidence for the existence of gravitational waves and won Hulse and Taylor the 1993 Nobel prize for physics. One of the objects is directly observed as a pulsar, with pulse period of 59ms and, from measurements of the variation of the arrival time of the pulses, it is inferred that the other object must also be a neutron star [53].

From conservation of energy, it follows that if a system is emitting energy in the form of gravitational waves, then the system must be losing energy itself. In the case of a binary system we would expect the system to slowly lose rotational energy as it emits gravitational waves. This loss of energy can be observed by measuring the time of periastron of the system over a long period of time. If the system is losing energy, the orbital radius is expected to decrease and thus the orbital period will also decrease in agreement with Kepler's third law. Such

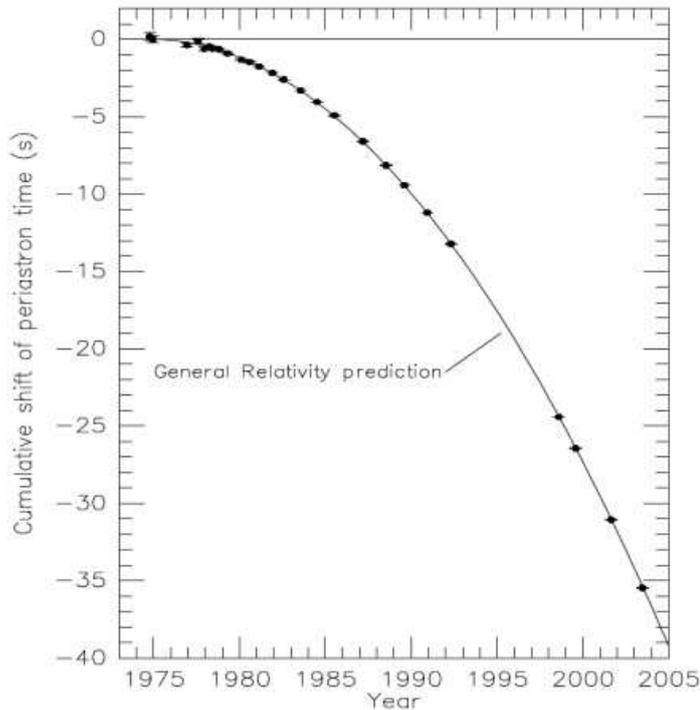


Figure 2.3: The cumulative shift of periastron time of the Hulse-Taylor binary as a function of time. As can be seen the observations fit exactly the predictions made by general relativity. The horizontal bar at $y = 0$ shows the prediction from Newtonian gravity. This Figure is taken from [54].

observations of this system have been taken between 1975 and the present, the results of which, shown in figure 2.3, clearly show agreement with the predictions of general relativity and the emission of gravitational waves [54]. The difference between the predicted and observed energy loss is quoted as $(0.13 \pm 0.21)\%$ in [54].

2.3.2 PSR J0737-3039

The observational evidence from the Hulse-Taylor binary strongly indicates that this system is emitting gravitational waves. Since then a handful of other similar systems have also been observed to show the same behaviour, all agreeing with Einstein’s theory of general relativity [55, 56]. The best evidence we have to date comes from a system called PSR J0737-3039 [56].

PSR J0737-3039 is a system consisting of two neutron stars, but here, uniquely, both objects are observed directly as pulsars and their orbit is oriented almost face on. This allows the dynamics of this system to be observed to a greater degree of accuracy than is possible with other binary systems and general relativity to be tested in a number of ways. Measurements of this system have allowed astronomers

to accurately determine the gravitational redshift, time dilation parameters and Shapiro-time delay parameters as well as the rate of change of periastron. All observations have been fully consistent with the predictions of general relativity [56]. In [56] it is quoted that observations of this system agree with general relativity with a error of only 0.05%.

2.4 Gravitational wave detection with laser interferometers

The history of gravitational wave detectors stretches back almost 50 years to the 1960's. The first gravitational wave detector was a *resonant bar detector* built by Joseph Weber. In the following years the sensitivity of these bar detectors has increased, but still has not reached a level where they will be able to detect a source outside of the galaxy and its nearby surroundings. For a detailed explanation of resonant bar detectors we refer the reader to [36].

The last ten to fifteen years have seen the development and operation of a different kind of gravitational wave detector, the large scale laser interferometer. These massive instruments, requiring large collaborations of people to operate, are beginning to reach sensitivity levels in which detection of gravitational waves becomes possible.

In section 2.4.1 we begin by discussing the basic operation of a laser interferometer, in sections 2.4.2 and 2.4.3 we describe how an interferometer would react to the passage of a gravitational wave. In section 2.4.4 we discuss current interferometers and in sections 2.4.5 and 2.4.6 the noise sources that limit their sensitivities. Finally, we discuss next generation interferometers in section 2.4.7 and a proposed space-based detector in section 2.4.8.

2.4.1 A simple description of laser interferometers

The Michelson interferometer dates back to 1887, when it was famously used to show the non-existence of the aether. A Michelson laser interferometer consists of two equal length arms (ideally) oriented at a 90° angle with lasers running along the length of the arms. The principle of the detector is as follows.

- Laser light is emitted at the centre of the “L” shape and split using a beam splitter.

- The light then travels along each of the arms and is reflected by mirrors at the end of each arm.
- The light passes back down the arms and recombines at the beam splitter.
- Some of this recombined light will then be directed toward a photodiode (while the rest will travel back toward the laser).

The idea is that if gravitational radiation causes the relative light travel time up and down the arms to vary, the laser light being recombined will have a time varying phase offset. Thus by observing this phase offset the change in path length between the two arms can be measured. This is illustrated in figure 2.4.

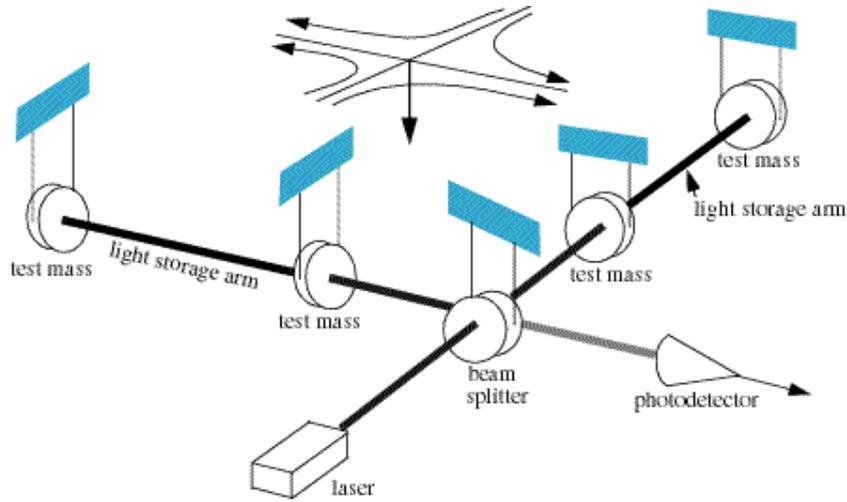


Figure 2.4: An illustration showing the principle components of a gravitational wave interferometer.

To try to demonstrate how this might work in practice let us consider a simple example given in [37]. We begin by considering the electric field of the input laser light as

$$E_{\text{input}} = E_0 e^{i\omega t - i\mathbf{k}\cdot\mathbf{x}}, \quad (2.45)$$

where ω is the angular frequency of the light and \mathbf{k} is the wave vector. E_0 denotes the amplitude of the laser light. Consider the beam splitter to be at $\mathbf{x} = 0$. As the beam splitter will send equal power up each arm, we can denote the light, transmitted by the beam splitter and travelling down the x arm as having an electric field given by

$$E_x = \frac{E_0}{\sqrt{2}} e^{i\omega t - i|\mathbf{k}|x}. \quad (2.46)$$

The electric field reflected by the beam splitter and travelling down the y arm will have an electric field given by

$$E_y = i \frac{E_0}{\sqrt{2}} e^{i\omega t - i|\mathbf{k}|y}, \quad (2.47)$$

where $1/\sqrt{2}$ is the transmission coefficient and $i/\sqrt{2}$ the reflection coefficient.

At the end of each arm the light is reflected. Thus the x and y light returning to the beam splitter is given by

$$E_x = \frac{E_0}{\sqrt{2}} e^{i\omega_L t - i2kL_x} \quad (2.48)$$

$$E_y = i \frac{E_0}{\sqrt{2}} e^{i\omega_L t - i2kL_y}, \quad (2.49)$$

where L_x and L_y denote the path length along the x and y arms respectively. The light combining at the beam splitter will again be reflected or transmitted. Thus, the light exiting the beam splitter toward the photodetector is given by

$$E_{\text{out}} = i \frac{E_0}{2} (e^{i\omega_L t - 2kL_x} + e^{i\omega_L t - 2ikL_y}). \quad (2.50)$$

With some manipulation this can be written as

$$E_{\text{out}} = iE_0 e^{i\omega_L t - k(L_x + L_y)} \cos(k(L_x - L_y)). \quad (2.51)$$

Since the power of a beam of light is proportional to E^2 we can see that

$$P_{\text{out}} \propto 1 + \cos(2k(L_x - L_y)). \quad (2.52)$$

Thus any variation of the relative path length of the arms would cause a variation in the power incident on the photodiode. Or, if the power incident on the photodiode varies with time, it implies that the relative path length of the two arms is varying.

In this section we have only given a very simplistic description of how an interferometer works. For a more comprehensive description of the operation of modern interferometers and the methods used to increase sensitivity see [37, 36].

2.4.2 Response of an interferometer to a gravitational wave

In this subsection we will consider how an interferometer reacts to the passage of a gravitational wave. Consider a gravitational wave detector with two equal length

arms pointing along the x and y directions interacting with a gravitational wave propagating in the z direction. We can calculate the light travel time up each of the arms by using the fact that for light $ds^2 = 0$. Thus for the x axis

$$0 = -dt^2 + (1 + h_+)dx^2 \quad (2.53)$$

From integration, the light travel time up the arm is given by,

$$t_x = L_0 + \int_0^{L_0} h_+ dx^2, \quad (2.54)$$

where L_0 is the length of the arm when no gravitational wave is incident on the detector and we have ignored terms that are second order in h_+ . The integral on the right of this equation is easily evaluated if we assume that h_+ does not vary significantly during travel up and down the arm. This is a fair assumption for the realistic case of an interferometer with 4km arms and a gravitational wave with 100Hz frequency. The light travel time to travel up the x arm is then given by

$$t_x = L_0(1 + h_+). \quad (2.55)$$

Similarly the light travel time down the arm is evaluated in the same way and has the same value. However, the light travel time up the y arm is given by

$$t_y = L_0(1 - h_+). \quad (2.56)$$

The difference in the light travel time to go up each arm and back to the beam splitter is then given by

$$\Delta t = 4L_0h_+. \quad (2.57)$$

Therefore as h_+ varies, the relative light travel time up the two arms will also vary. This will then cause variations in the power observed at the photodetector, allowing us to directly observe the variation in the gravitational field. We note that in the case described, the change in light travel time has no dependence on h_\times . Thus, a single gravitational wave detector is only sensitive to one polarization of gravitational waves.

There is, however, a large gap between this simple explanation and a gravitational wave interferometer capable of detecting gravitational radiation with amplitude $|h| \sim 10^{-22}$. For example consider a gravitational wave detector with 4km arms and a 500nm laser. To be able to achieve sensitivities to $|h| = 10^{-22}$ the detector would need to be sensitive to phase changes of order 10^{-12} .

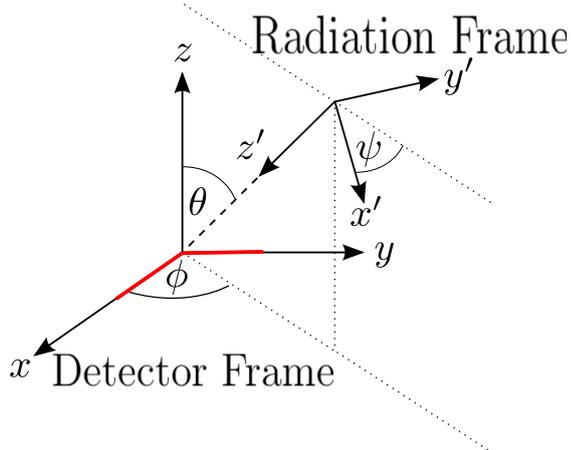


Figure 2.5: An illustration of the angles that describe the relationship between the detector and radiation frames. Figure originally published in [38].

To achieve this level of sensitivity the operation of gravitational wave interferometers is much more complex than the description here. For example, Fabry-Perot cavities are used to increase the effective path length of the laser light. Power recycling techniques and high powered lasers are used to maximize the power coming out of the beam splitter. A much more comprehensive description of the operation of gravitational wave interferometers and the difficulties they have to overcome can be found in [37, 36].

2.4.3 Sensitivity to different sky positions

Let us now consider a more generic case, where the detector is not aligned with the radiation frame. To deal with this we simply perform a series of rotations on the radiation frame to transform it into a frame where the x and y directions point along the arms of the detector. We will refer to this frame as the detector frame. The angles relating the detector frame to the radiation frame are shown in Figure 2.5. In words, the angles (θ, ϕ) give the sky location of the source relative to the detector frame. These two angles translate us to a frame in which z points from the source to the detector. A further angle, the polarization phase, ψ , is needed to rotate the x and y axes of this frame into the radiation frame described in section 2.1.7¹. Equation (2.57) then becomes

$$\Delta t = 4L_0 h(t). \quad (2.58)$$

¹We note that an alternative definition of the radiation frame uses the polarization phase in the translation between source and radiation frames instead of between detector and radiation frames.

where

$$h(t) = F_+ h_+ + F_\times h_\times. \quad (2.59)$$

In this formulation F_+ and F_\times give the detector response to the h_+ and h_\times components respectively of the gravitational wave in the radiation frame. Explicitly these are given by [57]

$$F_+(\theta, \phi, \psi) = -\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi \quad (2.60a)$$

$$F_\times(\theta, \phi, \psi) = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \sin 2\psi - \cos \theta \sin 2\phi \cos 2\psi \quad (2.60b)$$

2.4.4 Ground based interferometers

The last decade has seen the emergence of gravitational wave astronomy. The Laser Interferometer Gravitational-wave Observatory (LIGO) scientific collaboration has conducted two extended science runs, S5 and S6, taking gravitational wave data at their design sensitivities. The LIGO detectors are situated in the United States. One of these in Livingston, LA with a 4km armlength. The other two detectors are housed in the same beam tubes in Hanford, WA; one with a 4km armlength and one with a 2km armlength. The 2km detector was not active during S6. These “initial LIGO” detectors are now being decommissioned in preparation for the “advanced” detectors as discussed in subsequent sections.

The Virgo scientific collaboration has conducted three science runs in coincidence with the LIGO detectors using a 3km detector near Pisa in Italy. Virgo’s first science run (VSR1) overlapped the end of LIGO’s S5 while Virgo’s second and third science runs (VSR2 and VSR3) were conducted during S6. The Virgo collaboration is currently planning a 4th science run before beginning construction of their “advanced Virgo” detector. The LIGO and Virgo collaborations have been working together to jointly analyse their data since Virgo began VSR1.

In addition to LIGO and Virgo there is a 600m detector near Hannover in Germany called GEO600, the operation and data analysis from this detector is carried out in consortium with the LIGO detectors by the LIGO Scientific Collaboration (LSC). This detector, while not as sensitive as the others, is often used as a prototype for new interferometer technology.

2.4.5 Interferometer sensitivities and noise sources

To be able to compare the sensitivities between these detectors and determine what sources these detectors would be able to detect we have to consider sources of

noise. There are many sources of noise which effect a gravitational wave detector, but the ones that, in principle, dominate the limits of our sensitivity are

- Thermal noise: The various components of the interferometer, including the mirrors will vibrate due to their thermal energy. This noise can be reduced by employing mirrors which do not have resonance frequencies in the most sensitive region. Alternatively the detector could be cryogenically cooled as proposed in the LCGT project [24].
- Seismic noise: Vibrations of the ground will cause vibrations in the mirrors of the interferometer, which will cause changes in the light travel time up the arms. These vibrations can be caused by many things, including earthquakes, human activity and ocean waves. This noise can be limited by trying to isolate the mirrors from the ground as much as possible using advanced suspension systems.
- Shot noise: The number of photons seen at the photodetector in any interval of time will be poisson distributed. This error in the photon count at any time places a limit on the sensitivity. This noise source can be reduced simply by increasing the power of the laser.
- Radiation pressure: Photons hitting the end mirrors exerts a pressure on the mirrors. This noise source can be reduced by decreasing the power of the laser, thus reducing the pressure. However, this will increase the shot noise. A balance must therefore be reached between the radiation pressure and the shot noise.

An illustration of how the different noise sources contribute to the sensitivity of the detectors can be seen in Figure 2.6. Here the y axis shows the predicted contribution of the various noise sources to the overall noise power spectral density (PSD).

For real detectors the noise PSD is somewhat more complicated than the theoretical case described above. Many more noise sources will determine the PSD. For a more detailed discussion of such noise sources we refer the reader to [37, 36]. Figure 2.7 compares the actual sensitivity of the LIGO, Virgo and GEO detectors in 2006-2007.

We also refer the reader to the “Space Time Quest” game, which illustrates how these different sources of noise will contribute to a detector’s sensitivity and how they can be balanced, within a set budget [60]!

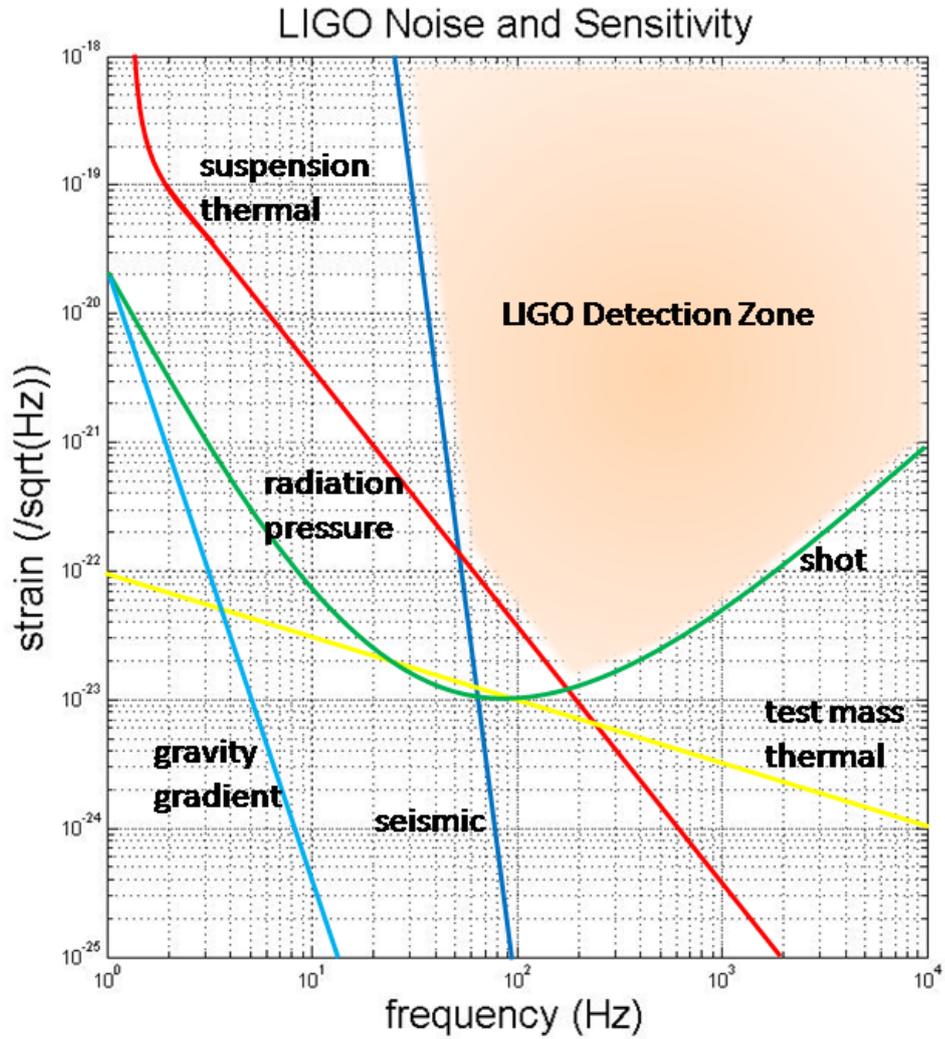


Figure 2.6: An illustration of the sensitivity of current gravitational wave interferometers, highlighting how different noise sources will determine the strain sensitivity of the instrument. Figure originally published in [58].

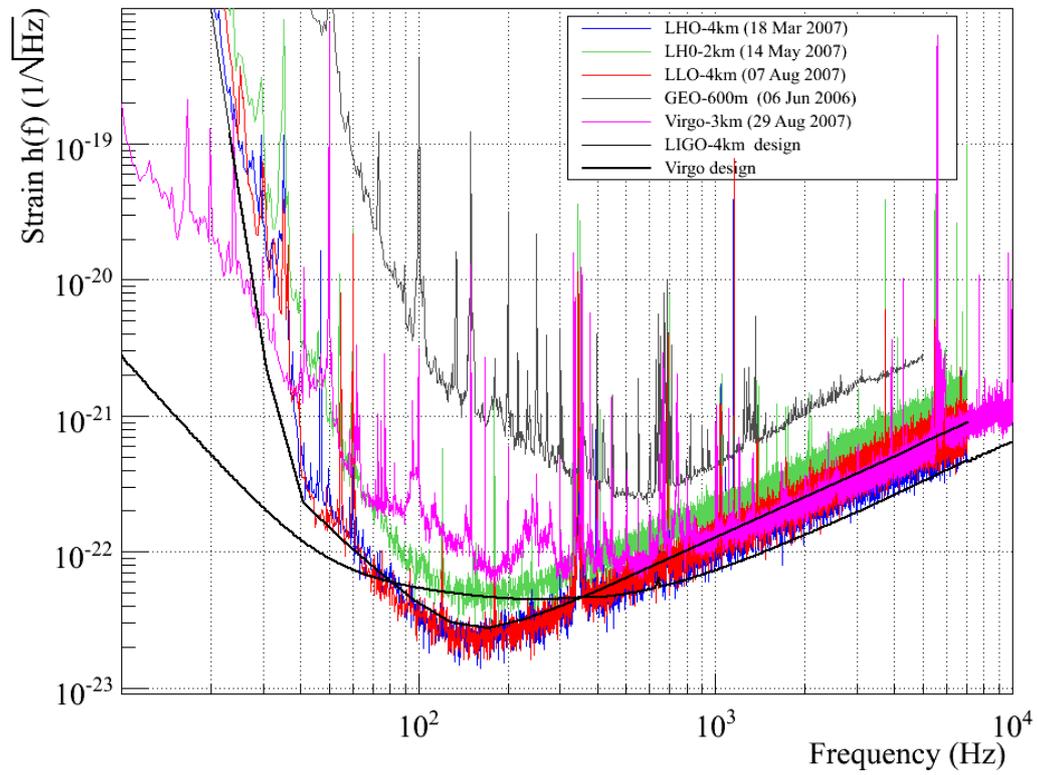


Figure 2.7: The various sensitivities of ground based detectors. The design sensitivities for the LIGO and Virgo instruments are also shown. Figure originally published in [59].

2.4.6 Non-stationary transient noise sources

In the section above we described the stationary noise sources that describe the sensitivity curve of gravitational wave detectors. However, the ability to detect gravitational waves is often dominated by non-stationary transient noise sources. Such transient noise sources can occur for a variety of reasons. Some of these are well understood, for example it is known that a truck driving near the instrument will generate ground motion, which couples to mirror motion. However, the reason for a lot of non-stationary noise in gravitational wave interferometers is not well understood. A large number of auxiliary monitors, such as seismometers, operate around the detectors to attempt to determine the origins of such noise transients and mitigate their effects.

Non-stationary transients or “glitches” are often mistakenly picked up by the data analysis software as interesting events. We will spend a lot of time in the subsequent chapters discussing methods for separating these glitches from gravitational wave signals in the context of searches for CBC signals.

2.4.7 Next generation ground based interferometers

The LIGO detectors are currently being rebuilt to become the “advanced LIGO” detectors [21]. This upgrade, which aims to reduce the effect of all sources of noise and increase sensitivity by a factor of 10 over existing instruments, is planned to be complete around 2014. Virgo will also soon be rebuilt and reconstructed as Advanced Virgo with a similar increase in sensitivity [22]. In addition to this the GEO detector is planned to be rebuilt as a detector with high sensitivity in the kHz range [61].

As well as the planned upgrades to the LSC-Virgo detectors the Japanese plan to build a cryogenically cooled detector LCGT in Japan [24]. An Australian collaboration ACIGA has plans to build a large scale interferometer [25]. Finally, a third generation of detectors designed to replace Advanced LIGO and Advanced Virgo are already being considered, such as the Einstein Telescope (ET) [62].

2.4.8 Space based interferometers

The most ambitious gravitational wave detector currently under serious development is a space based detector called Laser Interferometer Space Antenna (LISA). LISA is planned to consist of three separate spacecraft orbiting the Sun in such a way that the three spacecraft will always form an equilateral triangle. This detector would have an armlength of 5 million kilometers. The LISA constellation

is planned to orbit the Sun at a distance of 1AU following the Earth’s orbit, trailing it by 20° . Because of the lack of seismic noise and the greatly increased path length of the lasers LISA will be able to achieve a sensitivity to gravitational waves of a much lower frequency than those that ground based detectors could observe. LISA’s sensitivity range would span from $100\mu\text{Hz}$ to 100mHz (see figure 2.8). This would allow this instrument to detect a different range of sources to those that ground based detectors would observe. We discuss in more detail searching for super massive black hole systems with LISA in chapter 8.

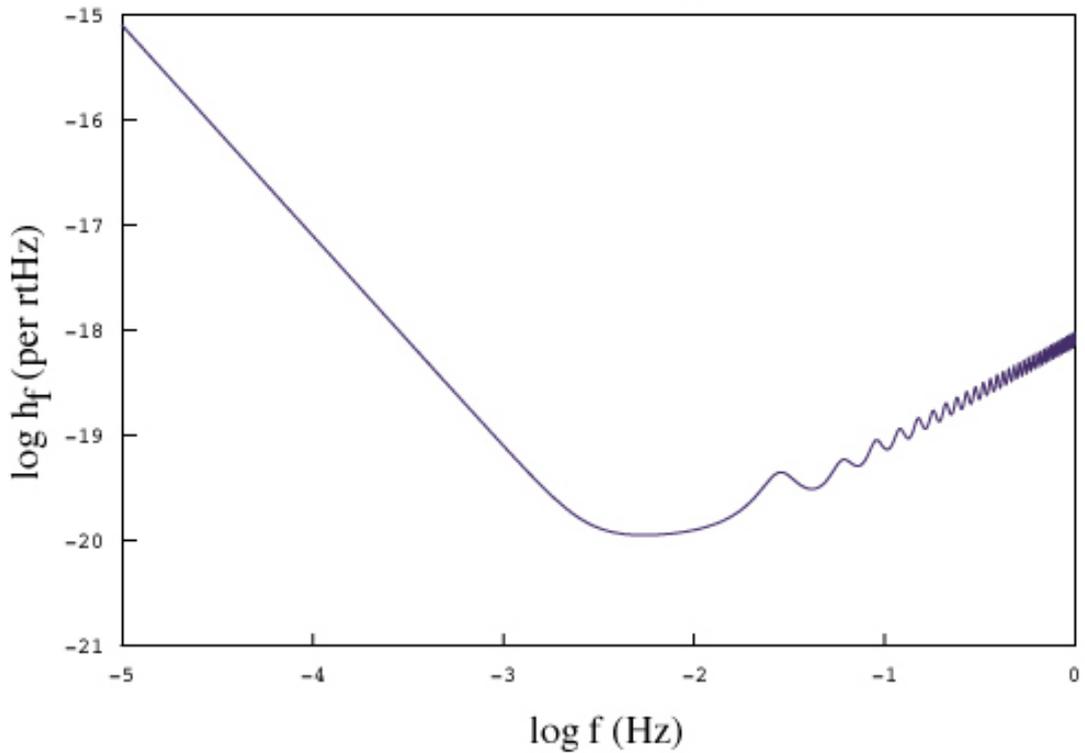


Figure 2.8: A diagram showing the predicted sensitivity curve of the proposed space-based LISA detector. Figure originally published in [63].

Chapter 3

Gravitational radiation emitted by a compact binary coalescence

In this thesis we describe a variety of data analysis techniques that have been used or can be used to search for compact binary coalescences (CBCs) using gravitational wave interferometers. All of these techniques utilise matched filtering, which we will describe in section 4.1.1. To achieve optimum sensitivity to CBC signals when using matched filtering, the form of the signal, as seen by the gravitational wave detectors, must be predicted as accurately as possible.

In this chapter we will derive the form of gravitational radiation emitted from two inspiralling compact objects, to leading order in sections 3.2.1 to 3.2.4. We will also derive the response of a gravitational interferometer to this radiation in section 3.2.5. We discuss the Post-Newtonian (PN) approximation, which is used to predict the phase evolution of inspiral systems beyond the dominant term in section 3.3. We also discuss the effect of higher order amplitude terms in section 3.4. We will then mention the efforts in the numerical relativity community to predict the gravitational radiation emitted during the merger and ringdown phases of a CBC in section 3.5. Finally, we discuss waveform families that are used in the analyses described in later chapters in section 3.6.

3.1 Common definitions

It will be useful to begin by describing the various quantities that are used in defining a CBC system. These definitions will be used heavily in this and all the following chapters. A CBC can be completely described by 17 physical parameters, these are

- The two masses, (m_1, m_2)
- The coalescence time of the signal at the detector, t_c
- The sky location of the source, (θ, ϕ)
- The distance to the source, D
- The inclination angle, ι
- The coalescence phase, Ψ_c
- The polarization phase, ψ
- The spin of the two components, $(\mathbf{S}_1, \mathbf{S}_2)$
- The ellipticity of the system and the orbital phase of periastron at t_c .

The masses of the system are often combined in a number of different ways, these are given by

- The total mass, $M = m_1 + m_2$
- The chirp mass, $\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}$
- The symmetric mass ratio, $\eta = \frac{m_1 m_2}{(m_1 + m_2)^2}$
- The reduced mass, $\mu = \frac{m_1 m_2}{m_1 + m_2}$.

We also use the following definitions

- The *orbital* phase of the system, Ψ
- The phase *of the dominant mode of the emitted signal*, $\Phi = 2\Psi$
- The *orbital* angular frequency, ω
- The frequency *of the emitted signal*, f
- The orbital radius of the system, r
- The vector describing the orbital angular momentum, \mathbf{L}_N .

Note that the frequency of the dominant component of the emitted signal will be twice the orbital frequency, as we will show in section 3.2, thus $\omega = \pi f$.

3.2 Quadrupole radiation, to dominant order, from an inspiral

In section 2.1.7 we discussed the nature of gravitational wave emission in a generic context. In section 2.4.2 we discussed how this generic gravitational radiation would be observed in an interferometric gravitational wave detector on the Earth. In this section we will apply these generic equations to the case of gravitational wave emission from a compact binary inspiral.

3.2.1 Time domain waveforms in the radiation frame

We begin by modelling the compact binary inspiral, assuming that there is no eccentricity and that the component bodies have no spin. Apart from CBCs with extreme mass ratios, it is believed that as the system nears merger any ellipticity that would have any effect on the system would have been radiated away [64, 65, 66]. We will explore how the presence of spin will effect the dynamics of a CBC in chapter 6. We also ignore the effects of, for example tidal effects in BNS systems, by modelling the objects as two point masses.

Therefore we model the inspiral as two point masses orbiting in circular orbits with no change in the direction of orbital angular momentum over time. We set our coordinates such that there is no motion in the z direction and the x direction is aligned with the orbital phase at $t = 0$

$$x(t) = r(t) \cos \left(\int_0^t \omega(t') dt' \right) \quad (3.1a)$$

$$y(t) = r(t) \sin \left(\int_0^t \omega(t') dt' \right) \quad (3.1b)$$

$$z(t) = 0. \quad (3.1c)$$

The radius and frequency are time dependent due to energy loss to gravitational waves in the system. To fit this into the generic form given in equation (2.42) we need to calculate the second differentials of the second mass moment. The second mass moment for two objects in orbit is given by [36]

$$M^{ij} = \mu x^i(t) x^j(t). \quad (3.2)$$

If we differentiate this with respect to time twice we obtain

$$\ddot{M}_{11} = -\ddot{M}_{22} = 2\mu r^2 \omega^2 \cos\left(2 \int_0^t \omega(t') dt'\right) \quad (3.3a)$$

$$\ddot{M}_{12} = -2\mu r^2 \omega^2 \sin\left(2 \int_0^t \omega(t') dt'\right) \quad (3.3b)$$

$$\ddot{M}_{13} = \ddot{M}_{23} = \ddot{M}_{33} = 0, \quad (3.3c)$$

where we assume that $r\omega \gg \dot{r}$. When this is not true the system is not really in circular motion. Thus, by combining these expressions with equations (2.42a) and (2.42b) we obtain the following expressions for h_+ and h_\times

$$h_+ = \frac{2\mu\omega^2 r^2}{D} (1 + \cos^2 \iota) \cos(\Phi(t) + 2\varphi) \quad (3.4a)$$

$$h_\times = -\frac{2\mu\omega^2 r^2}{D} 2 \cos \iota \sin(\Phi(t) + 2\varphi), \quad (3.4b)$$

where the gravitational wave phase is defined as

$$\Phi(t) = 2 \int_0^t \omega(t') dt'. \quad (3.5)$$

We now wish to simplify the equation by expressing r in terms of ω . To dominant order, this is given by Kepler's law

$$\omega^2 = \frac{m_1 + m_2}{r^3}. \quad (3.6)$$

We can then write h_+ and h_\times as

$$h_+(t) = \frac{2}{D} \mathcal{M}^{5/3} \omega(t)^{2/3} (1 + \cos^2 \iota) \cos(\Phi(t) + 2\varphi) \quad (3.7a)$$

$$h_\times(t) = -\frac{2}{D} \mathcal{M}^{5/3} \omega(t)^{2/3} (2 \cos \iota) \sin(\Phi(t) + 2\varphi). \quad (3.7b)$$

3.2.2 Energy loss in the system

Now we wish to investigate how ω varies with time, due to emission of gravitational radiation causing energy loss in the system. Firstly, we must quantify how much energy is being lost to gravitational waves. Equation (2.44) gives us a general formula for the energy loss, to dominant order

$$\frac{dE}{dt} = \frac{1}{5} \langle \ddot{M}'_{ij} \ddot{M}'^{ij} - \frac{1}{3} (\delta^{kl} \ddot{M}'_{kl})^2 \rangle. \quad (3.8)$$

We then perform a further differentiation on equation (3.3) to obtain the third differentials of the second mass moment

$$\ddot{M}_{11} = -\ddot{M}_{22} = -4\mu r^2 \omega^3 \sin\left(2 \int_0^t \omega(t') dt'\right) \quad (3.9a)$$

$$\ddot{M}_{12} = -4\mu r^2 \omega^3 \cos\left(2 \int_0^t \omega(t') dt'\right) \quad (3.9b)$$

$$\ddot{M}_{13} = \ddot{M}_{23} = \ddot{M}_{33} = 0. \quad (3.9c)$$

The power radiated by the system is then calculated by inserting equation (3.9) into (2.44). Using the assumption that for circular orbits $\langle \sin^2 \Phi(t) \rangle = \langle \cos^2 \Phi(t) \rangle = \frac{1}{2}$ we get

$$P = -\frac{dE_{orbit}}{dt} = \frac{32}{5} (\mathcal{M}\omega)^{10/3}, \quad (3.10)$$

which is the total radiated power of the system, to dominant order.

3.2.3 Phase evolution of the system

Now we have an expression for the radiated energy we can calculate how the frequency will vary with time. Consider the total energy of the system

$$\begin{aligned} E_{orbit} &= -\frac{m_1 m_2}{r} + \frac{m_1 m_2}{2r} = -\frac{m_1 m_2}{2r} \\ &= -(\mathcal{M}^5 \omega^2 / 8)^{1/3}, \end{aligned} \quad (3.11)$$

differentiate this with respect to time and insert it into equation (3.10) to get

$$\frac{32}{5} (\mathcal{M}\omega)^{10/3} = \mathcal{M}^{5/3} \frac{1}{3} \omega^{-1/3} \dot{\omega}. \quad (3.12)$$

This formula can be rearranged to give the change in orbital frequency with respect to time, $\dot{\omega}$, as

$$\dot{\omega} = \frac{96}{5} \mathcal{M}^{5/3} \omega^{11/3}. \quad (3.13)$$

The orbital angular frequency can then be expressed as a function of time, by integrating this equation from a fiducial time t_0 with fiducial frequency ω_0

$$\omega(t) = \left(-\frac{256}{5} \mathcal{M}^{5/3} (t - t_0) + \omega_0^{-8/3} \right)^{-3/8}. \quad (3.14)$$

From examination we can see that there will always be a time t when ω will become infinite. For simplification let us take this as our fiducial time, which we

will call t_c , this can be thought of as the time of coalescence. This choice serves to set $\omega_0^{-8/3}$ to zero. If we then define

$$\tau = t_c - t. \quad (3.15)$$

ω can be re-expressed in terms of τ as

$$\omega(\tau) = \frac{1}{8} \left(\frac{\tau}{5} \right)^{-3/8} \mathcal{M}^{-5/8}, \quad (3.16)$$

Now $\Phi(t)$ can be evaluated by combining equations (3.5) and (3.16)

$$\Phi(\tau) = 2 \int_{t_c}^{\tau} \frac{1}{8} \left(\frac{\tau'}{5} \right)^{-3/8} \mathcal{M}^{-5/8} d\tau' \quad (3.17a)$$

$$\Phi(\tau) = 2 \left(\frac{\tau}{5\mathcal{M}} \right)^{5/8} - 2 \left(\frac{t_c}{5\mathcal{M}} \right)^{5/8}. \quad (3.17b)$$

Finally, this allows h_+ and h_\times to be evaluated in the time domain by combining equations (3.7), (3.16) and (3.17b).

3.2.4 Frequency domain waveforms

It is often useful in gravitational wave searches to express h_+ and h_\times in the frequency domain. This can be done by performing a Fourier transform on h_+ and h_\times .

Firstly, to express τ as a function of frequency, equation (3.16) is rearranged, remembering that $\omega = \pi f$, to get

$$\tau(f) = \frac{5}{256} (\pi f)^{-8/3} \mathcal{M}^{-5/3}. \quad (3.18)$$

Inserting this into equation (3.17b) gives

$$\Phi(f) = \frac{1}{16} f^{-5/3} \mathcal{M}^{-5/3} - 2 \left(\frac{t_c}{5\mathcal{M}} \right)^{5/8} \quad (3.19)$$

and the time domain waveforms can be written in terms of the frequency

$$h_+(\tau) = \frac{2}{D} \mathcal{M}^{5/3} (\pi f(\tau))^{2/3} (1 + \cos^2 \iota) \cos(\Phi(f(\tau)) + 2\varphi) \quad (3.20a)$$

$$h_\times(\tau) = -\frac{2}{D} \mathcal{M}^{5/3} (\pi f(\tau))^{2/3} (2 \cos \iota) \sin(\Phi(f(\tau)) + 2\varphi). \quad (3.20b)$$

To convert this into the Fourier domain, \tilde{h}_+ and \tilde{h}_\times , waveforms a Fourier transform

could simply be performed on the time domain waveforms. However, performing a Fourier transform numerically can often be computationally expensive and it is therefore useful to have an analytical formula for the frequency domain waveforms.

This can be obtained by using the stationary phase approximation. The stationary phase approximation is defined in the following manner [67]. Given some function

$$h(t) = A(t) \cos(\phi(t)), \quad (3.21)$$

where

$$\frac{1}{A} \frac{dA}{dt} \ll \frac{d\phi}{dt} \quad (3.22)$$

at all t , The Fourier transform of $h(t)$ can be approximated as

$$\tilde{h}(f) \approx \frac{1}{2} A(f) \left(\frac{df}{dt} \right)^{1/2} \exp \left[-i \left(2\pi f t' - \phi(f) - \frac{\pi}{4} \right) \right], \quad (3.23)$$

where t' is defined as the time at which

$$\frac{d\phi(t)}{dt} = \pi f. \quad (3.24)$$

The stationary phase approximation can be applied to h_+ and h_\times to give analytical formulae for the frequency domain waveforms. To evaluate these frequency domain waveforms we first need to evaluate $\frac{df}{dt}$. Using equation (3.16) and $\omega = \pi f$ we get

$$\frac{df}{dt} = -\frac{df}{d\tau} = \frac{3}{320\pi} \left(\frac{\tau}{5} \right)^{-11/8} \mathcal{M}^{-5/8}, \quad (3.25)$$

this can be expressed in terms of frequency by substituting equation (3.18)

$$\frac{df}{dt} = \frac{96}{5} \pi^{8/3} f^{11/3} \mathcal{M}^{5/3}. \quad (3.26)$$

The stationary phase frequency domain waveforms can then be written, to dominant order, as

$$\tilde{h}_+(f) = \frac{1}{D} \left(\frac{5}{96} \right)^{1/2} \mathcal{M}^{5/6} \pi^{-2/3} f^{-7/6} (1 + \cos^2 \iota) \exp \left[i \left(2\pi f t' - \Phi(f) - \frac{\pi}{4} - 2\varphi \right) \right] \quad (3.27a)$$

$$\tilde{h}_\times(f) = -\frac{1}{D} \left(\frac{5}{96} \right)^{1/2} \mathcal{M}^{5/6} \pi^{-2/3} f^{-7/6} (2 \cos \iota) \exp \left[i \left(2\pi f t' - \Phi(f) + \frac{\pi}{4} - 2\varphi \right) \right]. \quad (3.27b)$$

We have now expressed the h_+ and h_\times components of the gravitational wave produced by a compact inspiral system in both the time and the frequency domain, to first order.

3.2.5 Detector response to a compact binary inspiral

Now we wish to calculate the strain that would be observed at a gravitational wave detector due to the passage of a gravitational wave emitted by a CBC. To do this the expressions for h_+ and h_\times can be combined with equations (2.59) and (2.60)

$$\begin{aligned}
h(\tau) &= F_+ h_+(\tau) + F_\times h_\times(\tau) \tag{3.28a} \\
&= \left(-\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi \right) \\
&\quad \left(\frac{2}{D} \mathcal{M}^{5/3} (\pi f(\tau))^{2/3} (1 + \cos^2 \iota) \cos (\Phi(\tau) + 2\varphi) \right) \\
&\quad + \left(\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \sin 2\psi - \cos \theta \sin 2\phi \cos 2\psi \right) \\
&\quad \left(-\frac{2}{D} \mathcal{M}^{5/3} (\pi f(\tau))^{2/3} (2 \cos \iota) \sin (\Phi(\tau) + 2\varphi) \right).
\end{aligned}$$

With some manipulation it is possible to write this as

$$h(\tau) = A(D, \iota, \theta, \psi, \phi) \mathcal{M}^{5/3} (f(\tau))^{2/3} \cos (\Phi(\mathcal{M}, \tau) + \Phi_0(\iota, \varphi, \theta, \psi, \phi)), \tag{3.29}$$

where A is a constant amplitude term and Φ_0 a constant phase offset. Or equivalently in the frequency domain as

$$\tilde{h}(f) = \tilde{A}(D, \iota, \theta, \psi, \phi) \mathcal{M}^{5/6} f^{-7/6} \exp \left[i \left(\Phi(\mathcal{M}, f) + \tilde{\Phi}_0(\iota, \varphi, \theta, \psi, \phi, t_c) \right) \right], \tag{3.30}$$

This implies that a single gravitational wave detector could not disentangle the various orientation angles and distance. With a single detector one could only recover an amplitude term and a phase offset. Nevertheless, being able to write $h(t)$ in this form does make the task of data analysis easier, as we will explore in the next chapter. We note that if the phase evolution is evaluated to higher order it will depend on η as well as \mathcal{M} and time/frequency.

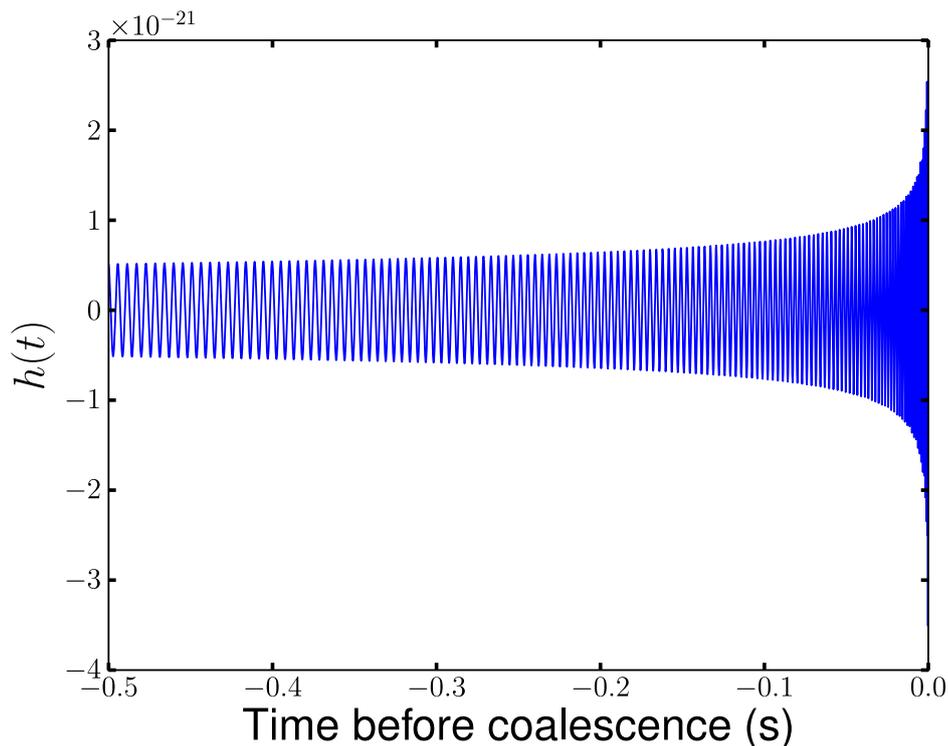


Figure 3.1: The gravitational waveform produced by the system described in section 3.2.6 plotted in the time domain.

3.2.6 Gravitational radiation from a binary neutron star merger

In this section we have derived, to dominant order, the form of the gravitational radiation emitted by a compact binary coalescence. To try to put this into perspective and understand what such a waveform would look like, let us consider an example system.

Consider a compact binary neutron star inspiral, with both components having masses $1.4M_{\odot}$. Let this system be located at a distance of 10Mpc and directly over a detector $\theta = \phi = 0$. Additionally, let it be oriented such that $\iota = \varphi = \psi = 0$. The gravitational waveform that this system would produce in the detector is shown, in the time domain, in Figure 3.1.

One must remember that this waveform has only been generated to dominant order. For a real signal, higher order terms will be important close to merger, as we will describe in the next sections. Additionally, the assumption of point masses will break down at merger, as effects due to the size of the objects, such as the tidal interaction between 2 neutron stars, can become noticeable.

3.3 Higher order phase terms

The gravitational waveform derived in equation (3.20) is only accurate to the dominant term both in amplitude and frequency because of various assumptions we made when deriving it. With ground based interferometers it is hoped that gravitational wave emission from CBC systems will be detected up to and including the merger. It is therefore vital that the phase evolution of the system is modelled as accurately as possible.

The Post-Newtonian (PN) formalism allows the phase evolution of the system to be predicted with much higher accuracy than the derivation given in the previous section [68, 69]. The PN expansion uses perturbative techniques to expand the phase of the system to higher order. Generally the expansion is performed around $(\pi M f)^{1/3}$.

As we have demonstrated, the dominant term in the time domain phase evolution is a multiple of $f^{-5/3}$. The next term, the “1 PN” term enters at $f^{-3/3}$, the “1.5 PN” term enters at $f^{-2/3}$ and so on (there is no 0.5 PN term proportional to $f^{-4/3}$). Current non-spinning PN expansions generally include all terms up to 3.5 PN order [68].

There are a number of different methods of constructing the PN expansion and these produce waveforms which are not identical when expressed to the same PN order. A detailed explanation and comparison of the various PN approximants can be found in [68].

3.4 Higher order amplitude terms

As well as higher order terms in the phase evolution there are also higher order amplitude terms. These higher order amplitude terms arise from the octopole (and higher) moments and therefore the phase of these terms is not necessarily twice the orbital phase. A study of these higher order amplitude terms and how they might affect our ability to detect CBC systems can be found in [38], but the importance of these higher order amplitude terms is much less than the higher order frequency terms.

For the current generation of gravitational wave detectors it is not vital to have waveforms including higher order amplitude terms for detection of CBC systems [38]. However, when the next generation of gravitational wave detectors comes online, it will be increasingly important for the signal models to be as accurate as possible, this will involve including higher order amplitude terms. In this thesis

we will only consider the dominant amplitude term.

3.5 Numerical relativity

The PN theory has allowed us to model the phase evolution of the inspiral part of the signal accurately. Perturbation theory also offers us a way to model the ringdown part of the signal. However, there is no theoretical model for the merger component of a CBC. Here we must rely on numerical modelling to provide us with waveforms that contain a merger component.

Recent advances in numerical relativity have allowed for the production of many waveforms containing the very late inspiral, merger and ringdown phases. These waveforms can be combined with inspiral waveforms generated using the Post-Newtonian approximation to produce “complete” inspiral, merger and ringdown (IMR) waveforms.

A recent review of numerical relativity and a list of waveforms that are currently available can be found in [70, 71].

3.6 Waveform models for analyses

There are a number of options for waveform models that can be used in data analysis. In [68] it is demonstrated that PN waveform models containing only the inspiral are sufficient for searching for CBC systems with a total mass below $12M_{\odot}$, with current detectors. The authors of [68] recommend using the “Taylor F2” frequency domain waveform family in such cases because it is the fastest waveform to compute. Taylor F2 waveforms are computed from a PN expansion of the stationary phase waveforms. This is in contrast to the Taylor T1-4 waveforms, which perform various PN expansions on the time domain waveform [68].

Above a total mass of about $12M_{\odot}$, the merger and ringdown components of the signal become important, inspiral only waveforms would be expected to pick up noticeably less signal-to-noise ratio (SNR) than complete IMR waveforms. Here the “EOBNR” waveform model, which attaches a numerically produced merger and ringdown to the inspiral signal produced by the effective one-body (EOB) model, is generally used [72].

In chapter 6 we describe the physical template family, used to model single spin systems. This is an extension of the Taylor T4 model [68]. The coherent search described in chapter 5 uses the same architecture as the single spin search

discussed in chapter 6, it therefore uses the Taylor T4 model to produce inspiral waveforms.

Chapter 4

The coincidence search for compact binary coalescences in LIGO and Virgo data

The LIGO and Virgo collaborations have performed many “all-sky” searches for compact binary coalescence (CBC) systems in the data taken by their interferometers [73, 74, 75, 76, 77, 1, 2, 3]. These all-sky searches attempt to detect CBC systems that have originated from any sky location, with any orientation and at any time when data is taken. While these all-sky searches have not yet made a direct detection of gravitational waves, they have allowed for upper limits to be placed on the merger rates of compact objects [3].

Input from electromagnetic (EM) observations has also allowed “triggered” searches for CBC signals to be performed [5, 4]. This is where an EM observation is made of some event whose progenitor might have been a CBC. The time and sky location of the event is sent to gravitational wave astronomers who then search their data for any coincident signal. Short GRBs are especially interesting for CBC triggered searches because the favoured progenitor model for a short GRB is a BNS or NSBH merger [30, 31]. Triggered searches have not yet made a direct detection, but are able to place lower distance limits on the progenitor of these events, assuming that it was a CBC.

The layout of this chapter is as follows. We will begin by discussing the matched-filtering techniques that make these searches possible in section 4.1. We then, in section 4.2, discuss in detail the all-sky search that has been used to produce the most recent results [1, 2, 4]. Finally, in section 4.3 we discuss the triggered search used to search for CBCs in coincidence with EM observations.

It is worth noting at this point that current searches use the “coincidence”

method where data from each of the detectors are analysed separately before searching for coincident events between them. In chapter 5 we will discuss an alternative coherent method, which simultaneously analyses data from all detectors; we then apply this to the triggered search.

Additionally, these searches use template waveforms that do not incorporate spin effects. While it is expected that there will be good sensitivity to most spinning systems, an improvement in sensitivity can be gained, in certain regions of parameter space, by using templates that include spin effects. We explore this in detail in chapter 6.

4.1 Gravitational wave data analysis

Gravitational wave signals in the data taken by interferometric detectors are not likely to be easily distinguishable from the noise. Therefore, a lot of the work currently being carried out by gravitational wave astronomers is aimed at being able to detect a weak signal buried in the noise produced from the detectors.

The problem of extracting signals from a noisy data stream is one that has been faced in other contexts for a long time [78]. For example, radio signals are often distorted by background noise and recovering the original signal can be difficult. The matched-filtering techniques discussed in this chapter were developed in the context of radio signal extraction and have been adapted to the detection of gravitational wave signals [78].

4.1.1 The matched-filter

The purpose of matched-filtering is to determine whether a known signal is present in a noisy data stream. We remind the reader that the form of a CBC signal in a gravitational wave detector *is* known, as we discussed in chapter 3. In this section our derivation follows closely those given in [78, 36].

Let us begin by considering the data output by a gravitational detector $s(t)$. This data is given by

$$s(t) = h(t) + n(t), \quad (4.1)$$

where $n(t)$ is the detector noise, which we will assume to be Gaussian with zero mean and stationary. $h(t)$ is a gravitational wave signal that we wish to extract from the noise. We wish to determine how likely it is that $h(t)$ is present in the data.

Before we describe the optimal matched-filter let us consider a naive approach.

One might assume that if we wanted to determine if $h(t)$ is present, we could simply multiply $s(t)$ by $h(t)$ and integrate over all time. If no signal is present, the expected value of this is zero, with some error. If, however, there is a signal present the expected value would be the integral of $h^2(t)$ over all time. If a value is obtained that is substantially larger than the expected deviation due to noise alone, then it could be inferred that the signal $h(t)$ is likely to be present in the data.

It is not clear, however, that the process of multiplying by $h(t)$ is the “optimal” method of extracting the signal from the data. Let us now consider a more general filter. Let us impose some linear filter $K(t)$ on $s(t)$. We want to determine the filter $K(t)$ such that the SNR is maximized for the signal $h(t)$. We can write this filter as

$$m = \int_{-\infty}^{\infty} K(t)s(t)dt = \int_{-\infty}^{\infty} K(t)h(t)dt + \int_{-\infty}^{\infty} K(t)n(t)dt. \quad (4.2)$$

We note that this is equivalent to the frequency domain representation

$$\int_{-\infty}^{\infty} K(t)s(t)dt = \int_{-\infty}^{\infty} \tilde{K}^*(f)\tilde{s}(f)df, \quad (4.3)$$

where the tilde represents that the quantity has been Fourier transformed. We note that for real data one would have a finite length of data with a finite sample rate, so the integral would become a summation over all discrete time points, or discrete frequency points. For this derivation however we will assume that the data is continuous.

To evaluate the optimal form of $K(t)$ we want to find the filter that would maximize the signal-to-noise ratio (SNR). This optimal SNR is defined as $\rho_{\text{opt}} = S/N$, where S is the expected value of m when the signal $h(t)$ is present and N is the rms value of m when no signal is present [78, 36].

Let us first consider N . Recalling that we imposed the condition that the noise is Gaussian and stationary we can write

$$\langle \tilde{n}(f)\tilde{n}(f') \rangle = \delta(f - f')S_n^{\text{d}}(f), \quad (4.4)$$

where we have defined S_n^{d} as the two-sided noise power spectral density (PSD).

Therefore we can express N as

$$\begin{aligned}
N^2 &= \langle m^2 \rangle |_{h=0} \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{K}^*(f) \tilde{K}(f') \langle \tilde{n}^*(f) \tilde{n}(f') \rangle df df' \\
&= \int_{-\infty}^{\infty} |\tilde{K}(f)|^2 S_n^d(f) df.
\end{aligned} \tag{4.5}$$

To evaluate S , the expected value of m with a signal present, we use the fact that the average value of the noise, at a given frequency is zero, $\langle \tilde{n}(f) \rangle = 0$ to obtain

$$S = \int_{-\infty}^{\infty} K(t)h(t)dt = \int_{-\infty}^{\infty} \tilde{K}^*(f)\tilde{h}(f)df. \tag{4.6}$$

Thus, we can express the SNR as

$$\rho_{\text{opt}} = \frac{\int_{-\infty}^{\infty} df \tilde{h}(f) \tilde{K}(f)}{\left(\int_{-\infty}^{\infty} df' S_n^d(f') |\tilde{K}(f')|^2 \right)^{1/2}}. \tag{4.7}$$

This equation can be simplified if we define the inner product between two real time series $a(t)$ and $b(t)$ to be

$$\begin{aligned}
(a|b) &= \int_{-\infty}^{\infty} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n^d(f)} df \\
&= 4 \operatorname{Re} \int_0^{\infty} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} df,
\end{aligned} \tag{4.8}$$

where the second line of this equation comes about because $S_n^d(f) = S_n^d(-f)$ and we demand that $a(t)$ and $b(t)$ are real functions, such that $\tilde{a}(f) = \tilde{a}^*(-f)$. $S_n(f)$ denotes the one-sided noise PSD, which will be used in the rest of this work, it is related to the two-sided noise PSD by $S_n(f) = 2S_n^d(f)$ for positive f . We can then re-express the SNR as

$$\rho_{\text{opt}} = \frac{(u|h)}{(u|u)^{1/2}}, \tag{4.9}$$

where \tilde{u} is given as

$$\tilde{u}(f) = \frac{1}{2} \tilde{K}(f) S_n(f). \tag{4.10}$$

It is then clear that ρ will be maximized when u is proportional to h . Thus the

maximum value of the optimal filter $\tilde{K}(f)$ is given by

$$\tilde{K}(f) = A \frac{\tilde{h}(f)}{S_n(f)}, \quad (4.11)$$

where A is an arbitrary constant. We have now calculated the optimal filter for a given signal in Gaussian, stationary noise. The optimal SNR for a signal h is given by

$$\rho_{\text{opt}} = (h|h)^{1/2}. \quad (4.12)$$

To filter the data to determine if the given signal is present one would calculate the matched-filter SNR as

$$\rho_{\text{mf}} = \frac{(s|h)}{(h|h)^{1/2}}. \quad (4.13)$$

Note that the constant A will cancel between numerator and denominator. If a signal is present we would expect ρ_{mf}^2 to follow a non-central χ^2 distribution with 1 degree of freedom and non-centrality parameter ρ_{opt}^2 . If no signal is present the distribution of ρ_{mf}^2 simply becomes a χ^2 distribution with 1 degree of freedom. We are keen to emphasize that the matched-filter SNR, ρ_{mf} and the optimal SNR, ρ_{opt} are not equivalent. There can be confusion between these two quantities as both are commonly referred to as simply SNR. In this work we will clearly state which of the quantities we are referring to when we introduce a SNR.

At the end of this derivation it is worth reminding the reader that we have assumed that we know the exact form of the signal, up to an overall amplitude factor. However, a CBC signal is characterised by a large number of physical parameters. It is thus necessary to filter the data against a set of templates, called a template bank [79]. If one were to naively create this template bank to cover all the physical parameters, the size of the bank would be far too large to filter with the current computing infrastructure. However, it is possible to analytically maximize over a number of the parameters and greatly reduce the size of the template bank.

For the non-spinning coincidence search, we will discuss template bank placement and the maximization technique in sections 4.2.3 and 4.2.4. We also describe matched-filtering techniques in the context of a coherent search in chapter 5 and spinning searches in chapter 6

4.1.2 The likelihood

Alternatively, this can be considered from a Bayesian perspective and the likelihood ratio can be calculated. This is the ratio of the likelihood of obtaining the given data assuming a signal h is present in the data and the likelihood of obtaining the given data assuming no signal is present.

As before we assume that the noise is Gaussian and stationary. The probability of a given noise realization n_0 occurring is given by [36]

$$p(n_0) = N \exp \{ - (n_0 | n_0) / 2 \}, \quad (4.14)$$

where N is a normalization constant.

We can then estimate the probability of a given realisation of data if we make the hypothesis that a signal is present with parameters given by ξ_i by taking $n_0 = s - h(\xi_i)$ and inserting this in the above equation to give us the conditional probability

$$p(s|h(\xi_i)) = N \exp \{ - (s - h(\xi_i) | s - h(\xi_i)) / 2 \} \quad (4.15)$$

$$= N \exp \left\{ (h|s) - \frac{1}{2} (h|h) - \frac{1}{2} (s|s) \right\}. \quad (4.16)$$

Similarly the probability of obtaining the given realization of data if no signal is present is obtained by setting $n_0 = s$ to give

$$p(s|0) = N \exp \{ - (s|s) / 2 \}. \quad (4.17)$$

We then define the likelihood ratio

$$\Lambda(h(\xi_i)) = \frac{p(s|h(\xi_i))}{p(s|0)} = \exp \left((h|s) - \frac{1}{2} (h|h) \right). \quad (4.18)$$

As we will see later on, this is often used as the log likelihood ratio

$$\lambda := \log \Lambda = (h|s) - \frac{1}{2} (h|h). \quad (4.19)$$

It would now be useful to try to relate this back to the matched-filter SNR. To do this we remember that the matched-filter SNR was maximized over an overall amplitude, whereas the likelihood formula is not. If we were to write $h = Ah_0$ and

Science run	Start time	End time
LIGO – S5	4 November 2005	30 September 2007
Virgo – VSR1	18 May 2007	30 September 2007
LIGO – S6	July 7 2009	October 20 2010
Virgo – VSR2	July 7 2009	January 11 2010
Virgo – VSR3	August 11 2010	October 20 2010

Table 4.1: The duration of LIGO and Virgo science times.

extract A from the log likelihood we obtain

$$\log \Lambda = A(h|s) - \frac{A^2}{2}(h|h). \quad (4.20)$$

This can then be maximized over A to get

$$\lambda|_{\text{Max},A} = \frac{1}{2} \frac{(s|h_0)^2}{(h_0|h_0)} = \rho^2/2. \quad (4.21)$$

Thus the log likelihood, maximized over amplitude, is easily related to the matched-filter SNR.

4.2 An all-sky, all-time search for compact binary coalescences

The LIGO and Virgo detectors have recently completed two joint science runs. The durations of these science runs are given in Table 4.1 and the sensitivities of the detectors to BNS inspirals during these science runs can be seen in Figure 4.1. With science runs spanning many years and the ability to detect BNS systems up to 40Mpc and NSBH and binary black holes (BBH) systems up to greater distances, the challenge is now to detect CBC signals buried in the detector noise.

The “ihope” search pipeline has been used to search for CBC signals originating at any time, from any direction and with any orientation in the data taken in LIGO’s fifth science run (S5) and Virgo’s first science run (VSR1) [1, 2, 3]. This technique is also being used to perform the same search on data from LIGO’s sixth science run (S6) and Virgo’s second and third science runs (VSR2 and VSR3). The search utilizes the “coincidence” technique in which data from all of the detectors is analysed separately, before looking for events which are coincident between detectors.

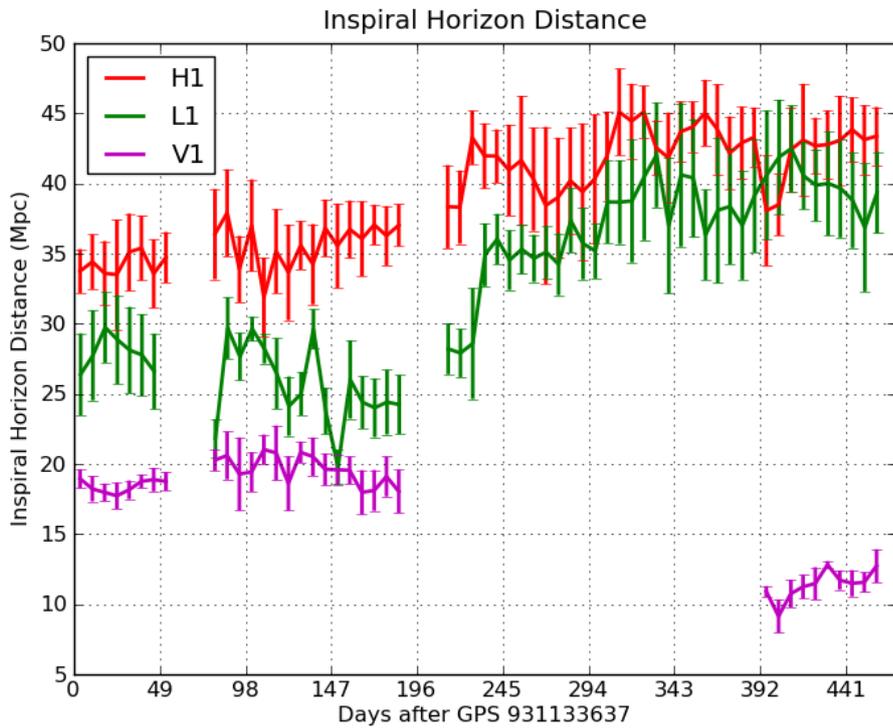
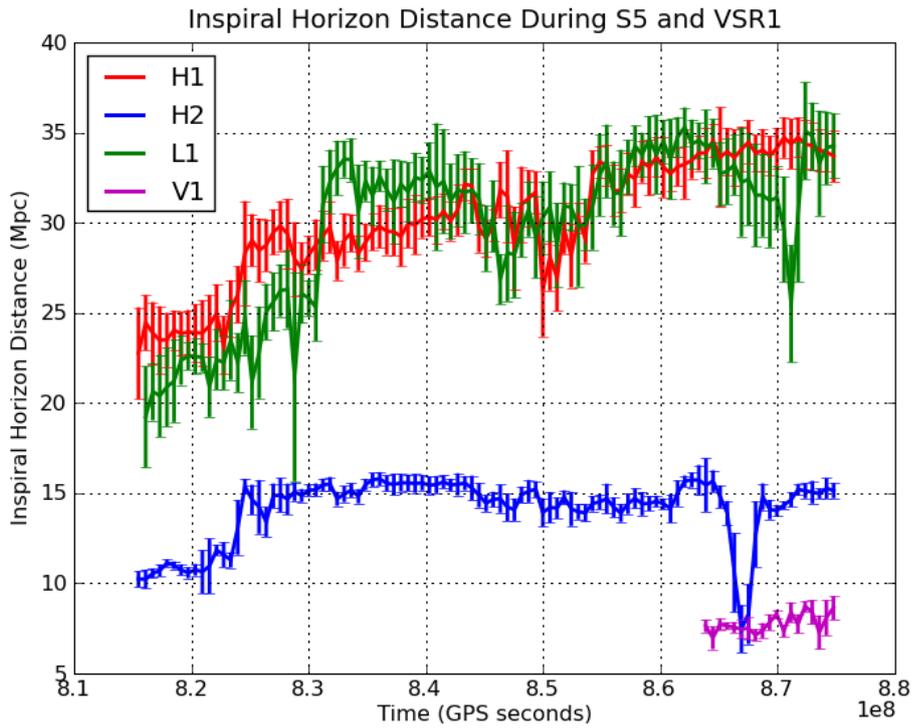


Figure 4.1: The inspiral horizon distance plotted against time for S5 and VSR1 (top) and for S6, VSR2 and VSR3 (bottom). The inspiral horizon distance is defined as the distance at which an optimally positioned, optimally oriented BNS merger would obtain an SNR of 8 in the respective detectors.

The ihope implementation described here is used to search for CBC systems whose total masses lie between 2 and $35 M_{\odot}$ with a minimum component mass of $1 M_{\odot}$. This “low mass” search is very similar to the “high mass” search described in [72]. The reason for the split is that the “high mass” search is used to search for CBC systems that merge within the sensitivity band of the instrument, therefore template waveforms incorporating the merger and ringdown must be used.

In this section we begin by giving an overview of the ihope pipeline in section 4.2.1. We will then describe the stages of the pipeline in detail in sections 4.2.2 to 4.2.12. Finally we will present the latest results of the search in section 4.2.13.

In Figure 4.1 and in subsequent text, we will refer to the various detectors using the following acronyms

- H1: The 4km LIGO instrument at Hanford. Operational in all LIGO science runs.
- H2: The 2km LIGO instrument at Hanford. Operational in LIGO’s S1 - S5.
- L1: The 4km LIGO instrument at Livingston. Operational in all LIGO science runs.
- V1: The 3km Virgo instrument at Cascina. Operational in all Virgo science runs.

4.2.1 Pipeline overview

The steps of the pipeline are described in detail in the following sections. What follows is a brief overview of these steps. An illustration of the various steps in the pipeline is given in Figure 4.2.

As this is a coincidence search, the first stage of the pipeline is to determine if there is any loud SNR event in any of the detectors. For each detector the process is to

- Create a template bank to cover the full range of masses.
- Filter the data against every template in the bank for each detector.
- Retain a “trigger” whenever a loud SNR is observed.

This results in a list of single detector triggers for each detector. The lists are then examined for any triggers that are coincident between detectors. A trigger is discarded if it is not seen in more than one detector. Coincidence is determined using the masses of the templates as well as the observed time.

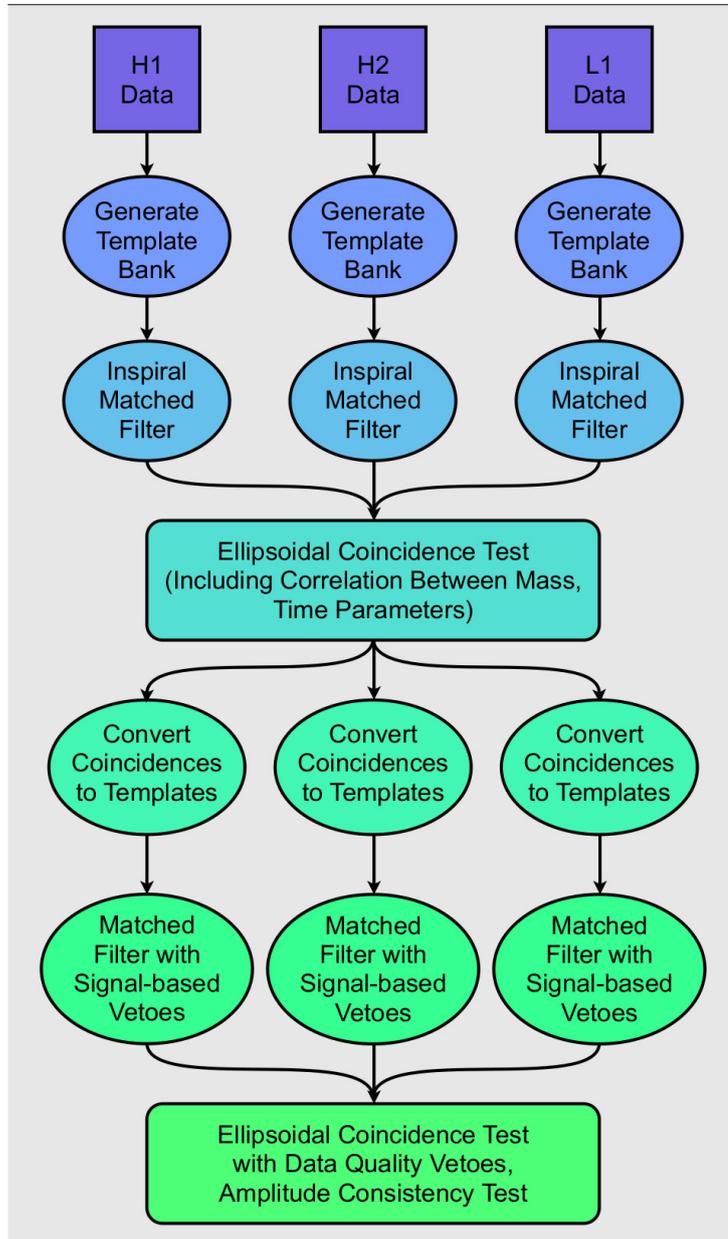


Figure 4.2: A flowchart illustrating the various steps in the ihope pipeline for the case of a search of H1, H2 and L1 data.

If the data taken by gravitational wave interferometers was Gaussian and stationary then the task of data analysis would stop here. A significance would be attached to all of the triggers and the likelihood of a CBC signal being present would be determined. However, the data taken is neither Gaussian nor stationary. Non-Gaussian transients or “glitches” often match well with the templates and are seen with a high SNR. It is vital to be able to distinguish such noise transients from real CBC signals and discard or downweight triggers that are due to noise transients.

A two-stage pipeline is therefore utilized. A second matched-filter is performed on any trigger that is seen in coincidence. However, during the second matched-filter a number of signal consistency tests are calculated, such as the powerful, but computationally expensive χ^2 test [80]. These tests are then used to separate glitches from gravitational wave signals. In addition to these signal consistency tests, triggers are also discarded that come from times when the detectors are known to be likely to produce glitches.

To determine if a signal is present in the data, the background rates of different types of triggers are estimated. These are used to associate a false alarm rate (FAR) to any trigger that survives the signal consistency tests. If no significant triggers are observed then upper limits on merger rates can be computed.

The figures that are shown in this section, unless indicated otherwise, were generated using results from this pipeline run on one month of data in S5. During this month all three LIGO detectors were operational. Virgo was not operating at this time. This month was the third month analysed in the search reported in [2]. In a number of the plots we have made use of simulated signals in the data to demonstrate the tuning of the pipeline. The mechanism by which such simulated signals are made into the data is described in section 4.2.10. For these plots the simulated signals were taken to be a set of BNS inspiral signals, uniformly distributed in distance, coalescence time, sky position and orientation angles.

4.2.2 Search initialization

The analysis pipeline begins with the calibrated $h(t)$ strain data taken from each of the instruments. Obtaining this calibrated data from the detector is not a trivial task, but we will not discuss these issues in this work. For details on the calibration of the detectors we refer the reader to [81, 82].

Given the calibrated data the coincidence analysis can be initialized. The first

step is to determine the available “science” data for each instrument¹. The data is then split up into blocks of 2048s duration, which are analysed separately. This value is chosen because of memory limitations and to ensure that as many CPU cores as possible are utilized. One noise PSD is calculated for each block and is used to analyse all the data in that block. The calculation of this noise PSD is described in [67].

These blocks of data are then further separated into 16 overlapping 256s segments. The first segment begins at the start of the block, the next 128s later and so on. Each of these segments is matched-filtered separately. The reason for overlapping the segments is that the finite length of filters means that the process of Fourier transforming the data will corrupt times at the beginning and end of each segment [67], which cannot be analysed. Additionally, the length of a template should not exceed the length of available data preceeding the point being analysed. As an example, consider a matched-filter that is performed with a template of 10s duration; one could not search for a signal that coalesced 5s from the start of a segment. To be conservative and ensure that these boundary issues are avoided, 64s on either side of each segment is discarded. Thus, only the middle 128s of any segment is analysed. This does mean that the first and last 64s of any block is not analysed and therefore the blocks must overlap to ensure that all available data is analysed. However, the first and last 64s of any stretch of continuous science data is never analysed. Additionally, a stretch of data that is shorter than 2048s in length is not analysed by the pipeline as there is not sufficient data to reliably calculate a PSD.

As a final note, because the different detectors will not be taking data at the same times it is not required that the times of the blocks are the same in all detectors.

4.2.3 Generating a template bank

To perform a matched-filtered search that would recover any CBC system with minimal loss in SNR over a given range of parameters one must filter the data against a set of waveforms or “template bank”. The computational cost of any gravitational wave search is directly proportional to the number of templates used. It is therefore vital to have a method that enables one to place a template bank using as few templates as possible.

For the case of matched-filter searches for non-spinning CBC signals this prob-

¹ We note that “category 1” data is not included as we will explain in section 4.2.7

lem is well explored [83, 84, 85, 86, 87, 88, 79]. Here, it is only necessary to place templates to cover the two dimensional parameter space of the masses. Other parameters enter only as amplitude or phase corrections, which are maximized over analytically when matched-filtering, as we show in section 4.2.4.

The basic idea of current template placement algorithms is to place a grid of templates using a hexagonal geometrical placement algorithm [83]. The spacing between the templates is determined from calculating the *mismatch* between templates, where the mismatch MM between two templates described by parameters $\boldsymbol{\xi}^1, \boldsymbol{\xi}^2$ is given by [84]

$$MM = 1 - (h(\boldsymbol{\xi}^1)|h(\boldsymbol{\xi}^2)), \quad (4.22)$$

where the templates are normalized such that for any set of parameters $(h(\boldsymbol{\xi}^n)|h(\boldsymbol{\xi}^n)) = 1$. This mismatch can be interpreted as the loss in optimal SNR that would be incurred by searching for $h(\boldsymbol{\xi}^1)$ with $h(\boldsymbol{\xi}^2)$. The template banks are created such that no point in the parameter space would have a mismatch of larger than 0.03 with the closest template in the bank. This value is chosen such that signals in the data will be recovered with close to the optimal SNR and the number of templates will not become excessive.

Explicitly calculating the mismatch between various points in parameter space can be computationally expensive. Therefore, the parameter space metric is normally used to approximate this. This metric is given by [84]

$$g^{ij}(\boldsymbol{\xi}) = -\frac{1}{2} \frac{\partial^2 (h(\boldsymbol{\xi})|h(\boldsymbol{\xi}))}{\partial \xi_i \partial \xi_j}, \quad (4.23)$$

which describes the distance between two templates infinitesimally separated in parameter space

$$1 - (h(\boldsymbol{\xi})|h(\boldsymbol{\xi} - \delta\boldsymbol{\xi})) = \sum_{ij} g^{ij}(\boldsymbol{\xi}) \delta\xi_i \delta\xi_j. \quad (4.24)$$

This can be used as an approximation of the mismatch between two templates with non-infinitesimal separation. The approximation holds as long as the metric is roughly constant in the parameter space between the templates. For mismatches of 0.03 the metric is a good approximation [84] and is used when placing the template banks.

While the problem of placing templates for a non-spinning CBC search is solved by the hexagonal lattice method, it is not yet clear how to place template banks in higher dimensional parameter spaces, such as might be needed for searches for CBC signals including spin. We will discuss the difficulties of placing template in

higher dimension arbitrary parameter spaces in chapter 7, where we discuss the possibility of using a stochastic placement algorithm to place template banks in generic parameter spaces.

4.2.4 The matched-filter

The next stage in the pipeline is to matched-filter every template in the template bank against all the data that has been taken by the detectors. We begin by deriving the maximized matched-filter statistic that is used in coincidence modelled searches for CBC systems.

We demonstrated in equation (3.29) that a CBC waveform observed by a gravitational wave detector can be expressed as

$$h(\tau) = A(D, \iota, \theta, \psi, \phi) \mathcal{M}^{5/3} (f(\tau))^{2/3} \cos(\Phi(\mathcal{M}, \eta, \tau) + \Phi_0(\iota, \varphi, \theta, \psi, \phi)), \quad (4.25)$$

where τ is defined as the time until the binary system will merge, the other parameters are defined as in section 3.1. We remind the reader that the frequency evolution will depend only on the masses and coalescence time of the system. All other parameters enter the waveform as an amplitude or phase offset.

In equation (4.21), we showed that the log likelihood ratio maximized over amplitude can be written as

$$\lambda|_{\text{Max},A} = \frac{1}{2} \frac{(s|h)^2}{(h|h)}. \quad (4.26)$$

Further to this we wish to maximize over the phase offset. To do this we rewrite h in terms of two components

$$h(\tau) = h_0(\tau) \cos \Phi_0 + h_{\pi/2}(\tau) \sin \Phi_0, \quad (4.27)$$

where $h_0(\tau)$ and $h_{\pi/2}(\tau)$ are given explicitly by

$$\begin{aligned} h_0(\tau) &= A(\mathcal{M}, D, \iota, \varphi, \theta, \psi, \phi) (f(\tau))^{2/3} \cos(\Phi(\tau)) \\ -h_{\pi/2}(\tau) &= A(\mathcal{M}, D, \iota, \varphi, \theta, \psi, \phi) (f(\tau))^{2/3} \sin(\Phi(\tau)). \end{aligned} \quad (4.28)$$

The log likelihood can then be written in terms of $h_0(t)$ and $h_{\pi/2}(t)$ as

$$\lambda|_{\text{Max},A} = \frac{1}{2} \frac{[(s|h_0) \cos \Phi_0 + (s|h_{\pi/2}) \sin \Phi_0]^2}{(h_0|h_0)}, \quad (4.29)$$

where we have made use of $\tilde{h}_0 = i\tilde{h}_{\pi/2}$. This is identically true when using waveforms generated in the frequency domain using the stationary phase approximation. In this form it is possible to maximize the log likelihood over the phase offset, Φ_0 . Thus, we obtain the final maximized form of the log likelihood ratio for a single detector matched-filter CBC search

$$\lambda|_{\text{Max}(A, \Phi_0)} = \frac{\rho^2}{2} = \frac{[(s|h_0)^2 + (s|h_{\pi/2})^2]}{2(h_0|h_0)}, \quad (4.30)$$

where we have defined ρ to be the maximized, matched-filter, single detector SNR. We note that as A and Φ_0 have been maximized over, the statistic will only depend on the masses and coalescence time. It will have no dependence on the other ‘‘extrinsic’’ parameters.

To calculate ρ at all times an inverse fourier transform on the matched-filter is utilized [78]

$$(s|h)(t_c) = \int_{-\infty}^{\infty} \frac{\tilde{s}(f)[\tilde{h}_0(f)]^*}{S_h(f)} e^{-2\pi i f t_c} df, \quad (4.31)$$

where t_c is the coalescence time of the signal. This quantity will be complex; if h_0 is used as the template waveform then the real component will give $(s|h_0)(t_c)$, the imaginary component will give $(s|h_{\pi/2})(t_c)$. Fourier transforms can be computed efficiently using the FFTW algorithm [89]. For the low mass ihope searches \tilde{h}_0 is calculated directly in the frequency domain using the ‘‘Taylor F2’’ waveforms as described in section 3.6.

Equation (4.30) gives the maximized SNR that can be calculated at all times in the segment being analysed using equation (4.31). It is easy to show that the expected distribution of ρ^2 in Gaussian noise will follow a χ^2 distribution with two degrees of freedom. Triggers are retained only where the SNR at that point in time is larger than 5.5 and is the largest SNR within a small time interval. This value is chosen empirically such that the number of triggers in the pipeline is limited to a manageable level. The value of 5.5 may initially seem rather large; the probability for a χ^2 distributed statistic with 2 degrees of freedom to give a value larger than $(5.5)^2$ is 2.7×10^{-7} . However, consider that there may be approximately 1000 independent trials at any point in time due to the range of masses being covered and an independent trial at least every 0.1 seconds in time. We would therefore expect to see such events in even a few hundred seconds of data. It is not at all surprising to see a large number of $\rho > 5.5$ events in a year of observing time in Gaussian noise.

The distribution of triggers obtained by filtering 1 month of S5 data in the H1,

H2 and L1 detectors is shown in Figure 4.3. As can be seen on the figure there are a huge number of triggers, including a large number which have been found with an SNR of at least 100. Clearly this is not the distribution of triggers that would be expected from Gaussian noise. This illustrates that a single detector search, using only SNR to rank triggers, would have no realistic hope of detecting a gravitational wave. Much of what follows is focused on separating these noise transient triggers from gravitational wave signals.

4.2.5 Determining Coincidence

The next stage of the pipeline is to perform coincidence tests between the triggers that are produced for each of the detectors. Any trigger is discarded that is not seen in “coincidence” in at least two detectors. To be coincident, triggers must have occurred at the same time and with similar masses.

The exact definition of coincidence, which is used in the most recent CBC searches in LIGO and Virgo data, is given in [90]. This uses a similar approach to that of template placement to define a “distance” between two triggers in two different detectors. This distance, calculated using the parameter space metric given in equation (4.23) is known as the “ethinca distance”. Triggers are only considered to be coincident if their ethinca distance is less than a preset threshold. The threshold is empirically set at 0.5 by investigating the distribution of the ethinca distance that simulated signals are recovered with; as we demonstrate in Figure 4.4.

In Figure 4.3 we show the distribution of single detector events after this coincidence cut is applied. We can see from comparing the first two rows of Figure 4.3 that the coincidence cut has reduced the number of triggers, but that some of the loudest triggers are still present after the coincidence test has been applied. It is clear that a large number of noise triggers pass the coincidence test.

4.2.6 Second stage matched-filter

The “standard” χ^2 test (often just χ^2 test) [80], which we describe in detail in section 5.3, tests whether a potential trigger has the expected power in a number of different frequency bins. It is one of the most effective methods for separating non-Gaussian noise transients from gravitational wave signals. Unfortunately it is also very expensive to calculate; a matched-filter must be calculated for every frequency bin used in the χ^2 test. In the low mass search this is 16 frequency bins. Thus, it is desirable to only calculate the χ^2 test when necessary. A two stage

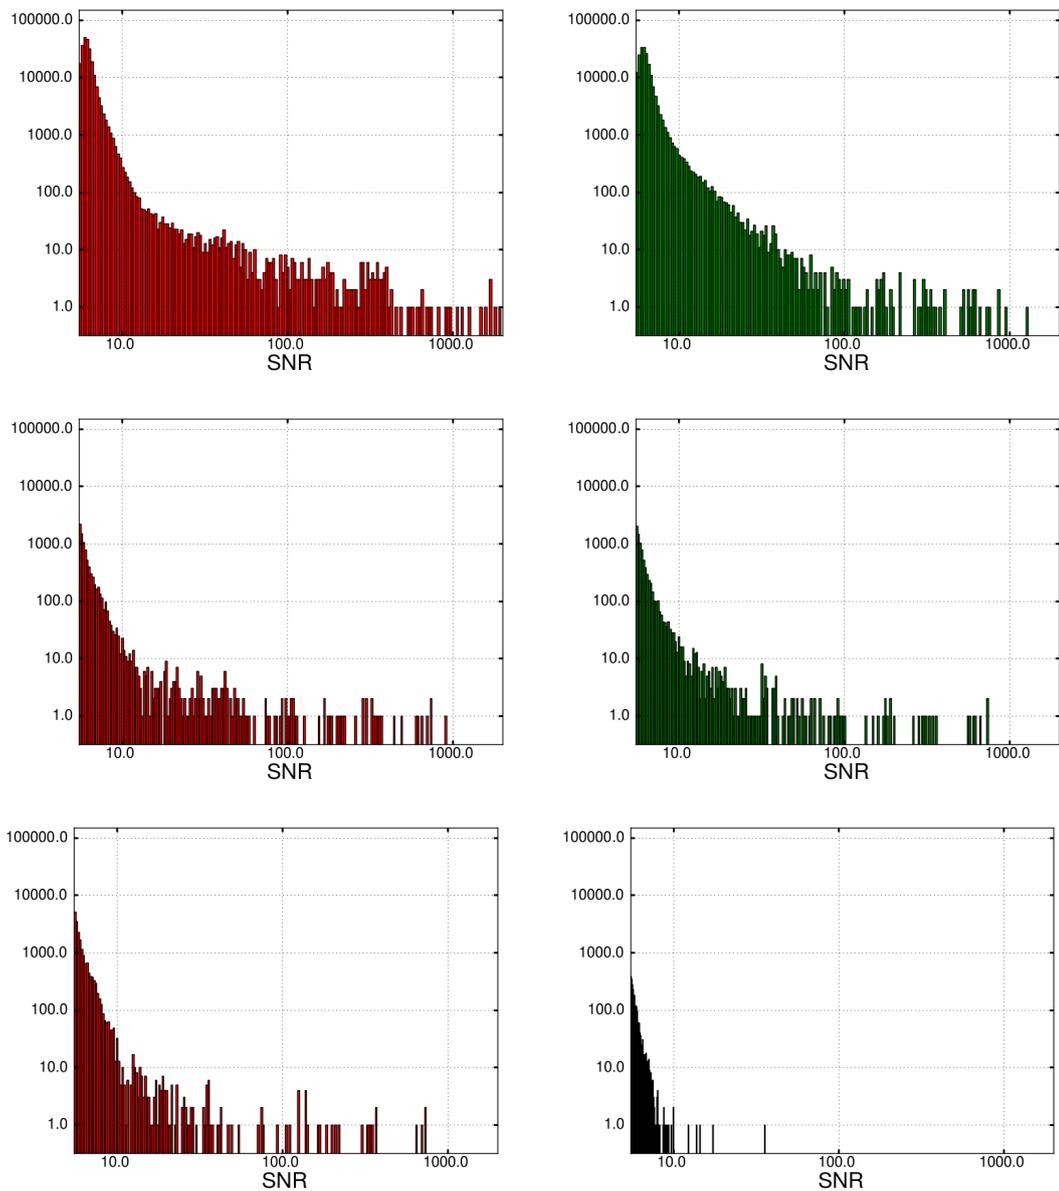


Figure 4.3: The number of triggers vs SNR for triggers obtained in the H1 (left) and L1 (right) detectors at a variety of stages in the pipeline. The first row shows triggers before coincidence is applied. The second row shows triggers after coincidence is applied. The third row shows triggers after coincidence and after signal based vetoes have been applied and after triggers from times marked as category 2 or category 3 have been omitted. This data is taken from the third month of the analysis described in [2].

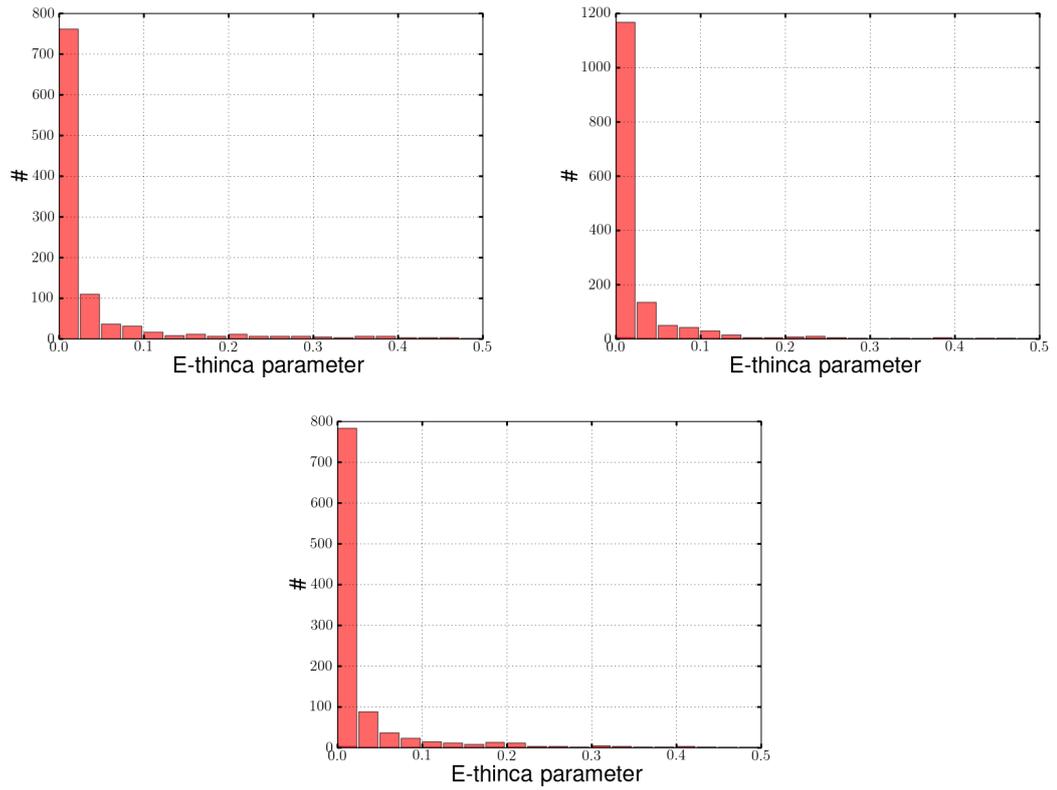


Figure 4.4: A histogram of the ethinca distance between triggers associated with simulated signals. Shown for triggers between H1 and H2 (top left), H1 and L1 (top right) and H2 and L1 (bottom). The ethinca cut was set to a value of 0.5. The distribution of background triggers is uniform in the ethinca distance.

pipeline is therefore utilized. In the first stage, matched-filtering is performed and coincidences determined without calculating the χ^2 test. A reduced template bank or “triggered bank” is then created for each block of data, containing only the templates that formed a coincidence within that block. The matched-filtering and coincidence is then performed a second time using the triggered bank. However this time the χ^2 test is calculated for every trigger at this second stage. We discuss how the χ^2 test is used to separate noise events from real signals in section 4.2.8.

4.2.7 Detector characterization

To reduce the number of triggers due to non-Gaussian noise in the analyses, it is useful to try to identify times during which noise transients are likely to occur. These glitches are normally caused by one of two reasons. Firstly there might be a problem with the operation of the detector that produces noise transients. For example the H1 and H2 detectors share the same beam tube. If one of these detectors is “down” and not taking science data the process of returning it to taking science data can affect the operation of the second detector. The second reason is that there might be some environmental feature that produces glitches in the detector data. For example seismic activity is known to couple to transverse motion in the mirrors and can produce transient noise events.

For these reasons the detectors are continually monitored by a host of sensors, which monitor the internal and external conditions. A lot of effort within the LIGO and Virgo collaborations is spent in trying to identify data that is likely to contain glitches and identifying, and if possible removing, the internal or external cause. For more details of these activities see [91, 92, 93, 94, 95, 96].

To consider this from the point of view of the data analyst it is sufficient to know whether the data should be analysed or not. The end product of the detector characterisation process is to assign all data a data quality category. Analyses for CBC signals treat these data quality categories in the following way²

- Category 1: Data marked as category 1 indicates that the detector is not operating correctly. This data are not used for any part of the analysis. The analysis blocks discussed earlier contain no category 1 data as it would corrupt any estimate of the noise PSD. An example of a category 1 flag might be if the data is unable to be calibrated and $h(t)$ is not available.

² Note that there is also a category 4, but for CBC searches this category is only used when following up interesting triggers. In the ihope pipeline category 4 data is treated equally to data with no category.

- Category 2: Data marked as category 2 indicates that some mechanism known to have a strong correlation with glitches was active at the time this data was taken. The analysis blocks *can* contain data marked as category 2, as it can be used to estimate a PSD. However, any trigger occurring during time marked as category 2 is discarded. This is referred to as “vetoing” a trigger³. An example of a category 2 flag is when a hardware injection is being made into the data.
- Category 3: Data marked as category 3 indicates that some mechanism known to have some correlation with noise transients was active at the time this data was taken. Category 3 data *is* analysed and false alarm rates are calculated for coincident triggers occurring during category 3 data. Category 3 data is, however, discarded when calculating upper limits. An example of a category 3 flag might be that there is elevated seismic noise at the time the data is taken..
- No category: Data that does not have a category assigned to it indicates that there is no understood environmental or internal cause of noise transients active at that time. This data is analysed normally.

Results obtained from category 3 data and “good” data are treated separately. We will discuss how these separate results are dealt with in [4.2.9](#).

To ensure that the detector characterization efforts would not lead to gravitational wave signals being vetoed more than would be expected by random chance we use hardware injections (see section [4.2.10](#) for details). These hardware injections simulate the response of the detector to a gravitational wave signal. The “safety” of detector characterization is assessed by ensuring that the percentage of these hardware injections vetoed by data quality is consistent with the percentage of the total data vetoed.

As an example of the amount of analysis time that is lost to data quality vetoes, we can consider the case of the third month in the analysis in [\[2\]](#). For this month, there is approximately 1.2 million seconds of data that is coincident in at least two detectors after category 1 times have been removed. After category 2 and category 3 times are removed the figure is reduced to approximately 1 million seconds, though the majority of the lost time is category 3 time and is still analysed.

³Triggers can also be vetoed by signal consistency tests as we explore in section [4.2.8](#).

4.2.8 Effective SNR and signal-based vetoes

The next stage in the pipeline is to apply signal consistency tests. Any trigger that does not have the expected characteristics of a CBC signal should be discarded. Currently the most powerful test for doing this is the χ^2 test, which we derive and explore in section 5.3. An SNR consistency test can also be utilized if a signal is seen in coincidence between the H1 and H2 detectors.

In this subsection we begin by describing how the χ^2 and SNR consistency tests are used to discard triggers that do not have the characteristics of a CBC signal. We then describe how triggers are ranked using a combination of SNR and the χ^2 and explain how this “effective SNR” is defined.

Vetoing triggers with χ^2

The value of the χ^2 combined with the SNR is used to rank triggers, thus down-weighting high SNR triggers with poor χ^2 , as we will describe. Additionally, triggers are vetoed based on the χ^2 values around the time of the trigger. The “ r^2 veto duration” is the duration in which the χ^2 is above 15 per degree of freedom within 6s of the trigger. The value of 6s for the time window of has been chosen empirically [97].

The cut on this is illustrated in Figure 4.5. Below an SNR of 12, triggers are discarded if they have a r^2 veto duration longer than 0.002s; above this the cut is dependent on SNR as demonstrated in the figure. There is a big discontinuity in this cut at SNR 12, but no triggers that are vetoed would be potential detection candidates. This values are also chosen empirically, as described in [97].

The H1, H2 SNR consistency cut

For events seen in coincidence between the H1 and H2 detectors a SNR consistency test is applied. These two detectors have aligned arms, as they share the same vacuum tubes, therefore the ratio of SNRs seen in H1 and H2 should be equal to the relative sensitivity of the detectors. As can be seen in Figure 4.1, the H1 detector is normally twice as sensitive as H2 so one would normally expect a signal to be observed in H1 with twice the SNR as the same signal seen in H2.

Formally the SNR consistency cut can be expressed in terms of “effective distance”. For a single detector trigger the amplitude and phase are maximized over. The orientation angles, source sky location and the distance which form the amplitude parameter cannot be separated. Instead an “effective distance” can be determined, this is defined as the distance at which an optimally oriented and lo-

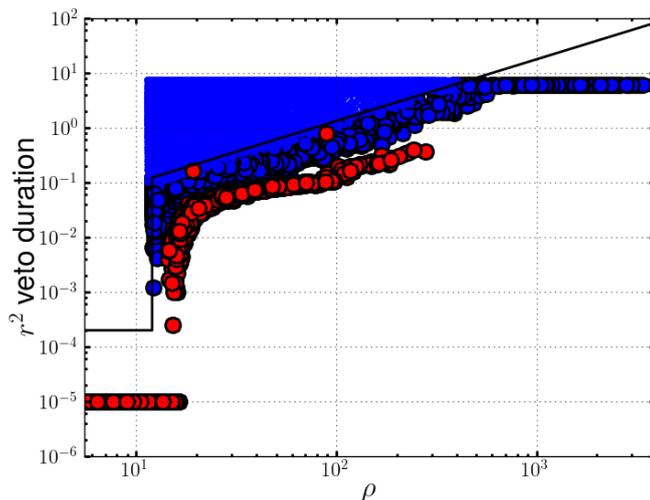


Figure 4.5: The r^2 veto duration plotted against SNR for all second stage triggers seen in the H1 detector. The solid mass of blue is comprised of blue crosses marking the position of triggers failing the r^2 cut, blue circles indicate triggers which do not fail the r^2 cut. The red circles are triggers produced by simulated signals, we note that none of these are vetoed by the r^2 cut. Triggers that are vetoed below an SNR of 12 are not shown in this figure.

cated source would give the observed SNR. For H1 and H2 the recovered effective distances for a CBC signal should be equal. In contrast, a glitch might occur in only one of the detectors, thus producing triggers with different effective distances. Triggers that have a fractional difference in effective distance greater than 0.6 are discarded; this allows for some spread due to measurement uncertainties. This cut is illustrated in Figure 4.6.

The effective distance cut was not used between other pairs of detectors because in principle any ratio of effective distance is possible for a real signal seen in two unaligned detectors. However, it is rather unlikely that, for a given system, the fractional difference in effective distances in two detectors will be large. Therefore, it should be possible to use amplitude consistency information to separate glitches from CBC signals in any detector combination. We investigate the possibility of using an effective amplitude cut between any two detectors when discussing coherent search methods in chapter 5.

Effective SNR and New SNR

We have now described the various instrumental and signal based vetoes that are applied to triggers passing through the pipeline. Unfortunately, as can be seen in Figure 4.3 while the majority of triggers have not passed the various checks, there

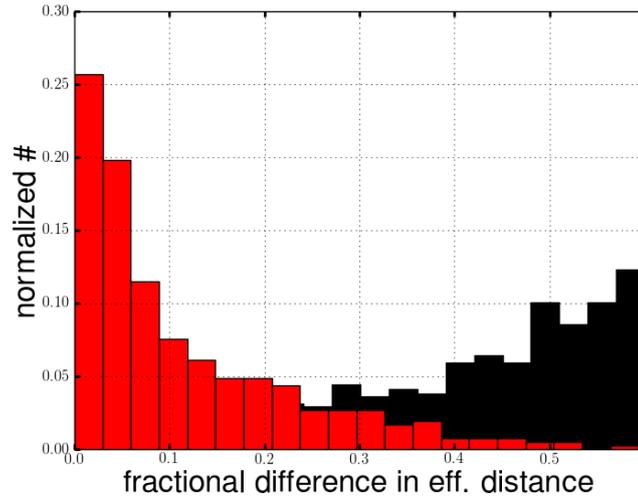


Figure 4.6: A histogram of the fractional difference in effective distance. The distributions for simulated signals (red) and background triggers (black) are shown. Here triggers with a fractional difference bigger than 0.6 have been removed.

are still many loud events surviving after the data quality and signal based vetoes have been applied. Figure 4.7 shows the distribution of the χ^2 statistic plotted against SNR for triggers that have survived all vetoes. A clear separation between noise background (see section 4.2.9) and simulated signals (see section 4.2.10) can easily be observed.

Triggers are therefore ranked using a combination of SNR and χ^2 . In S5 and VSR1 a combination called “effective SNR” was used, this is defined as

$$\rho_{\text{eff}}^2 = \frac{\rho^2}{\sqrt{\left(\frac{\chi^2}{n_{\text{dof}}}\right) \left(1 + \frac{\rho^2}{250}\right)}}, \quad (4.32)$$

where n_{dof} is used to denote the number of degrees of freedom in the χ^2 test (see section 5.3). The factor of 250 in the denominator was chosen to provide a suitable separation between background triggers and simulated signals. We demonstrate contours of effective SNR in Figure 4.7. It can be seen that these contours provide a much better separation between simulated signals and noise background triggers than SNR.

In S6 and VSR2 and VSR3 a different statistic is being used called *new SNR* [98]. This *new SNR* is designed to improve the contours to allow for better separation between the background and real events. The definition of *new SNR* is

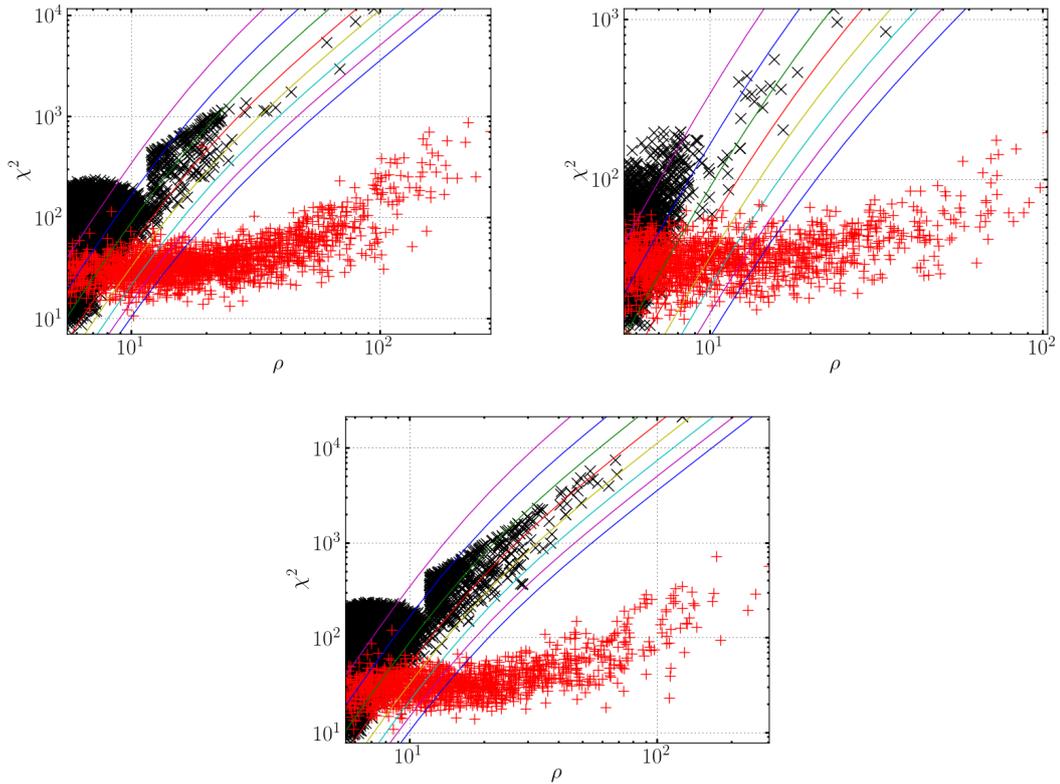


Figure 4.7: The χ^2 test plotted against SNR for triggers after all cuts have been applied and category 2 and 3 data removed. This is shown for triggers in the H1 detector (top left), the H2 detector (top right) and the L1 detector (bottom). The solid, coloured lines on the plots indicate lines of constant effective SNR, larger values of effective SNR are at the bottom and right end of the plots. The clearly visible notch in the H1 and L1 plots is caused by the discontinuity in the r^2 cut at an SNR of 12 (section 4.2.8). Here background triggers are represented by black crosses and injections by red plusses.

explained and given in [98] as

$$\rho_{\text{new}} = \begin{cases} \rho, & \chi^2 \leq n_{\text{dof}} \\ \frac{\rho}{\left[\left(1 + \left(\frac{\chi^2}{n_{\text{dof}}} \right)^{4/3} \right) / 2 \right]^{1/4}}, & \chi^2 > n_{\text{dof}} \end{cases} \quad (4.33)$$

and has been found to provide a better separation, especially for low mass non spinning signals and templates [98]. As well as a different form of the contours, it does not allow the significance of a trigger to be increased if the χ^2 value is less than 1 per degree of freedom.

The effective SNR (or *new SNR*) for triggers in coincidence are added together in quadrature to give a “combined effective SNR”. This is then used to rank triggers, as we describe in the next subsection.

4.2.9 False alarm rates

The signal consistency tests, data quality conditions and ranking statistic serve to remove or downweight loud non-Gaussian transient events that occur in the data. The next step is to calculate the significance of surviving triggers. To do this the rate at which different types of triggers appear due to the noise background must be calculated. This is then used to calculate a false alarm rate (FAR) for every trigger surviving the various vetoes. False alarm rates for data marked as category 3 are calculated separately from data that has no data quality flag. Both are used to search for detection candidates. To be a detection candidate a trigger must have a FAR that is significantly smaller than 1 per analysis time.

In this subsection we begin by describing how the background rate of triggers can be evaluated through a process of time shifting the data. We then discuss how triggers are separated into different categories, how FARs are calculated for each of these categories and how the values are combined to produce a final combined FAR.

Time shifted background

The background rate of triggers can be calculated by repeating the analysis with the triggers from different detectors shifted in time relative to each other. Time shifts are done in multiples of 5s, which ensures that the shifts are significantly larger than the light travel time between the detectors and the autocorrelation of the templates. Therefore, any coincidences seen in the time shifts must be due to

noise background events and not gravitational wave signals.

The normal practice is to perform 100 time shifts to provide a good estimate of the background. If any triggers are found in the non time shifted data that are louder than all the triggers in the background, additional time shifts can be performed to provide a better estimate of the background rates. To distinguish from the time slide analysis the non-time slide analysis is often referred to as the “zero lag” analysis.

This time shifting technique does not require a great deal of extra computational time. The single detector triggers already exist so no additional matched-filtering needs to be performed, only the coincidence stage must be recalculated for each time shift. We should note however, that in the two stage pipeline, the triggered bank (section 4.2.6) contains all templates that produced coincident triggers in the first stage both in the zero lag and in the time slides. Therefore additional matched-filtering is often necessary at the second stage.

The H1 and H2 detectors share the same beam tube and therefore background noise triggers in the two detectors have been observed to be correlated. Therefore, the time sliding technique does not work for estimating the background level of triggers detected by only these two detectors. While the loudest triggers seen in only H1 and H2 are briefly examined, no FARs are calculated. Because of this inability to calculate a false alarm rate, times when only H1 and H2 are operating are not used in the calculation of upper limits [1].

Trigger categories

The background rate of triggers is not uniform over the parameter space and different detectors can exhibit different levels of “glitchiness”. It is known, for example, that the background rate of triggers is dependent on the masses of the templates being filtered. In general, glitches match better to the shorter, higher mass templates than to the low mass ones and the signal consistency tests are less effective for such templates [1]. It is also known that triple (or quadruple) coincident triggers are considerably less likely to occur in background than double coincident triggers.

It is therefore necessary to split up the triggers into different categories and calculate FAR measurements relative to background in that category. The separation of triggers occurs as follows

- Triggers are separated according to the chirp mass of the trigger, there are three mass bins

- $\mathcal{M} < 3.48M_{\odot}$
- $3.48M_{\odot} < \mathcal{M} < 7.4M_{\odot}$
- $7.4M_{\odot} < \mathcal{M}$.

- Triggers are separated according to the detectors that observed the trigger.
- Triggers are separated according to the detectors that were active when the trigger was observed.

With H1 - H2 coincident triggers removed, there are 15 different categories before Virgo began VSR1 and 87 different categories after. A FAR can be calculated for a trigger in any of these categories by dividing the number of events, in the background, in the same category as the candidate trigger that are more or equally significant than the trigger, by the total analysed background time.

Combined false alarm rates

For triggers that occur when a given set of detectors were operating, a “combined FAR” can be calculated by simply dividing the uncombined FAR by the number of categories. The joint LIGO and Virgo search described in [3] took this further and developed a likelihood based method to combine the various categories based on how likely a real signal was to have been found in that category. A combined FAR was then calculated from this likelihood ranking.

The inverse of the combined FAR is often quoted as the ranking statistic. This is so that larger values of “IFAR” indicate that the event is more significant. Some examples of inverse FAR result plots are shown in Figure 4.8. There are no detection candidates in these plots. A potential detection candidate would appear as a blue triangle in the bottom right corner of these plots, with a very large IFAR.

4.2.10 Simulated signals

It is useful to test the sensitivity of the pipeline and demonstrate that it is capable of detecting CBC signals. To do this, simulated signals can be added to the data before filtering. The pipeline analyses these signals as it would any other stretch of data and the FAR of any trigger associated with these simulations can be determined.

These simulations can be performed in two ways. *Software injections* are performed by adding a simulated waveform to the data after it has been read into the

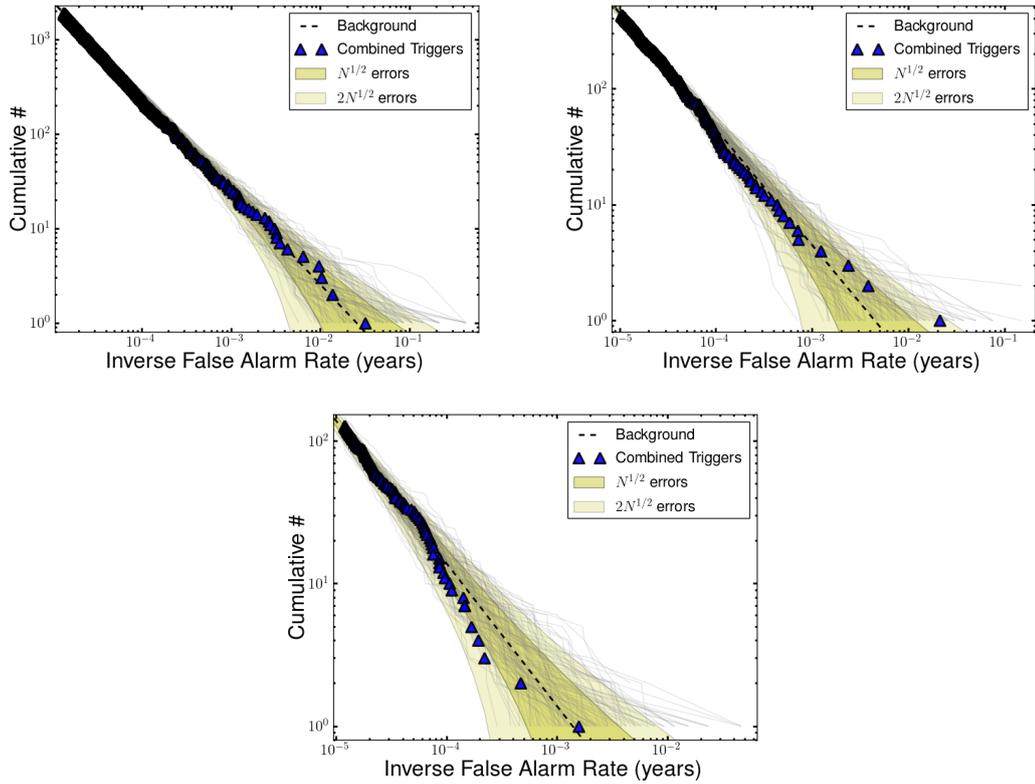


Figure 4.8: Cumulative histograms of triggers vs inverse false alarm rate. This is shown for triggers when H1, H2 and L1 were operating (top left), when H1 and L1 were operating (top right) and when H2 and L1 were operating (bottom). The black dashed line indicates the expected trigger level from the background, the shaded regions indicate one and two σ error regions. The thin grey lines give the cumulative number of triggers for each of the time slides (100 in total), these can give an additional indication of the expected deviation from the mean. In these plots the triggers are consistent with the background and there are no detection candidates.

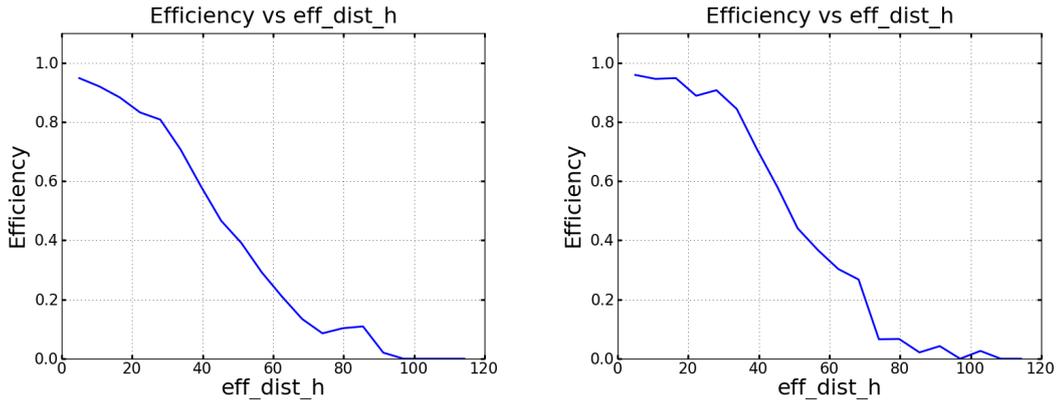


Figure 4.9: Efficiency of software injection recovery plotted against H1 effective distance for BNS injections. The left figure shows results using non-spinning injections, the right figure shows results using injections where spin is present and uniformly distributed. Injections are “found” if they are recovered with a value of effective SNR that is larger than the loudest event in the zero lag.

pipeline’s analysis codes. In contrast *hardware injections* are performed by actuating the mirrors at the end of the interferometers’ arms to simulate the response of a gravitational wave passage.

Hardware injections offer the better test of the sensitivity of gravitational wave instruments to CBC systems. A hardware injection allows for the whole process of gravitational wave astronomy to be tested, from the gravitational wave being incident on the detector through to making a detection. However, as hardware injections cannot be removed from the data, times when hardware injections are made cannot be analysed for any other signals that might be in the data. Such times are marked as category 2. Therefore, only a limited number of hardware injections can be made in a given stretch of data.

Software injections do not offer as comprehensive a test of the infrastructure as hardware injections. However, there is no limit to the number of software injections that can be performed. The analysis of the same data can simply be repeated with different software injections present as often as is desired.

For the CBC low mass analysis, the process is to use software injections to tune the pipeline, as we have shown in a number of plots already in this section. They are also used when calculating upper limits, as we will describe in section 4.2.11. Hardware injections are used to verify the results obtained from software injections. In Figure 4.9 we demonstrate the ability to recover BNS software injections with and without spin.

4.2.11 Upper limits

The final stage of the pipeline is to determine rate limits on the mergers of various compact objects based on the sensitivity of the pipeline. In the current era, where no CBC signals have yet been detected, this is limited to placing upper limits on merger rates. The upper limit quoted for the ihope pipeline is calculated based only on the loudest event. From the observation that no event was observed in the data louder than the loudest event, and given the sensitivity of the detectors, a 90% confidence upper limit can be placed on the rate of mergers. This can be calculated for BNS, NSBH and BBH signals. In this subsection we will present the method used to calculate the upper limits quoted in [1, 2, 4]. We note that category 1, 2 and 3 data are excluded when calculating upper limits. The upper limit calculation for CBCs was derived previously in [99, 100, 101]. We try to follow the notation of these earlier works as much as possible.

Derivation of rate upper limit

Here we will derive the rate upper limit calculation used in CBC searches. We begin by evaluating the probability that no gravitational wave signal will be recovered with an IFAR equal to or greater than α during the search

$$P(\alpha) = e^{-\mu(\alpha)T}. \quad (4.34)$$

Where T is the duration of the search and μ is the rate of signals which would be recovered louder than α . It is believed that the rate of these mergers is proportional to the rate of star formation, which is itself proportional to blue light luminosity [102]. Therefore, this can be written as

$$\mu = C_L(\alpha)R, \quad (4.35)$$

where R is the rate of mergers per unit luminosity per year and C_L is the blue light luminosity in which the search would expect to recover a signal louder than α . We will discuss the derivation of C_L later, for now we assume it is known. These quantities normally use L_{10} as the unit of luminosity, L_{10} is defined as 10^{10} of the blue light luminosity of the sun. To put this in context, the Milky Way galaxy contains $1.7L_{10}$. Over a large enough volume, where the distribution of galaxies can be considered isotropic, it is sufficient to use volume instead of luminosity in this calculation.

We also wish to calculate the probability that no noise signal will be recovered

with an IFAR greater than or equal to α , $P_B(\alpha)$. This is simply given as

$$P_B(\alpha) = e^{-1/\alpha}. \quad (4.36)$$

We can then calculate the probability that the zero lag search will produce no triggers with IFAR greater than or equal to α , either from noise or CBC signals, as

$$P(\alpha|R, T) = P_B(\alpha)P(\alpha) = P_B(\alpha)e^{-RC_L(\alpha)T}. \quad (4.37)$$

This is then differentiated with respect to α to obtain the probability per unit FAR interval

$$p(\alpha|R, T) = \frac{d}{d\alpha} (P_B(\alpha)e^{-RC_L(\alpha)T}) \quad (4.38a)$$

$$= P'_B(\alpha)e^{-RC_L(\alpha)T} \left(1 + \frac{P_B(\alpha)}{P'_B(\alpha)}(-RT)C'_L(\alpha) \right), \quad (4.38b)$$

where a dash indicates that this is the derivative of the quantity with respect to IFAR.

$$P'_B(\alpha) = \frac{\partial P_B}{\partial \alpha} \quad C'_L(\alpha) = \frac{\partial C_L}{\partial \alpha} \quad (4.39)$$

Using Bayes Theorem this can be re-expressed as the probability density function (pdf) of the rate of mergers given that no event was seen louder than the loudest event, α_m ,

$$p(R|\alpha_m, T) \propto p(R)e^{-RC_L(\alpha_m)T} (1 + \Lambda(\alpha_m)RTC_L(\alpha_m)), \quad (4.40)$$

where we have defined $p(R)$ as the prior distribution of R ; usually taken as the result of the previous search. If there was no previous search a uniform prior is used. Λ is defined as

$$\Lambda = \frac{|C'_L(\alpha_m)| P_B(\alpha_m)}{C_L(\alpha_m) P'_B(\alpha_m)} \quad (4.41)$$

and we can interpret Λ as the relative probability that the loudest event was due to a gravitational wave against it being due to background. If there is no chance it was a gravitational wave Λ will tend to zero. Whereas if it was definitely a gravitational wave Λ will tend to negative infinity.

If we were to normalize p we can then obtain an upper limit of the rate (R_*) for a given confidence level (γ) by evaluating

$$\gamma = \int_0^{R_*} p(R|\alpha_m, T) dR. \quad (4.42)$$

The upper limits calculated for CBC searches are calculated at 90% confidence.

Calculating the total luminosity sensitivity of a search

To calculate the upper limits described above one must determine the total luminosity a search is sensitive to, C_L , and the derivative of that with respect to IFAR. To calculate this one needs to evaluate, as a function of distance and mass,

- L_B , the total luminosity. This is evaluated using a galaxy catalogue, completed with an isotropic distribution of galaxies where the catalogue is incomplete [103].
- ϵ , the probability of being able to detect a system with a IFAR louder than or equal to the loudest event. This is evaluated by adding a large number of software injections into the data and seeing whether or not the pipeline detects them (and with what IFAR). The derivative of the efficiency with respect to IFAR can also be calculated numerically.

The total sensitive blue light luminosity is then given by multiplying these quantities and integrating over all distances and masses

$$C_L(\alpha_m) = \int \epsilon(\alpha_m, D, \mathcal{M}) L_B(D, \mathcal{M}) dD d\mathcal{M}. \quad (4.43)$$

As only the efficiency depends on FAR the derivative of C_L with respect to FAR can be given as

$$C'_L(\alpha_m) = \int \epsilon'(\alpha_m, D, \mathcal{M}) L_B(D, \mathcal{M}) dD d\mathcal{M}. \quad (4.44)$$

To evaluate the integrals numerically the efficiency and total luminosity must be evaluated in a number of distance and mass bins. For simplicity the mass is binned only in the chirp mass as it is assumed that the efficiency has little dependence on the mass ratio. The distance is binned in terms of effective distance, weighted by the chirp mass. This is to try to ensure that the efficiency can be reliably calculated in every bin, while limiting as much as possible the number of software injections that need to be performed.

Marginalizing over errors

In the calculation above we have described how upper limits are calculated for the CBC searches ignoring any errors in the calculation. However, there are various

errors that appear in this calculation and they must be considered when quoting upper limits. These errors include

- Imprecise knowledge of the astrophysical distributions of the mass and distance of binaries.
- Differences between the physical signal and the non-spinning, restricted post-Newtonian waveforms.
- Statistical fluctuations in the measured efficiency.
- Uncertainties in instrumental calibration.
- Errors in the calculated value of Λ , arising from the above uncertainties and errors in the background estimation.

The effect of these errors are incorporated by marginalizing the rate upper limit over them. The method of this marginalization is given explicitly in [101].

4.2.12 Output pages

When the analysis is completed and the results are ready to be examined, it is useful to have a tool that makes it easy to quickly examine the results and determine whether or not a statistically significant gravitational wave signal is present in this signal.

To do this a tool has been developed that parses the output created by the pipeline and creates a html page with all the relevant information easily accessible. An example of such an output page can be seen in Figure 4.10. This is now a standard tool and has greatly sped up the analysis of the output of the ihope pipeline.

4.2.13 Results

The search described above has been run to search for low mass CBCs in data from S5 and VSR1. The search was published in three parts, the first year of S5 data [1], the second year of S5 data up until Virgo became operational [2] and then a final results paper containing the joint S5-VSR1 analysis [3]. Unfortunately none of these searches found any gravitational wave event.

Upper limits on the merger rates for BBH, BNS and NSBH objects were produced by these searches, where the black hole mass is defined to be distributed as $(5 \pm 1)M_{\odot}$ and the neutron star mass to be distributed as $(1.35 \pm 0.04)M_{\odot}$. The

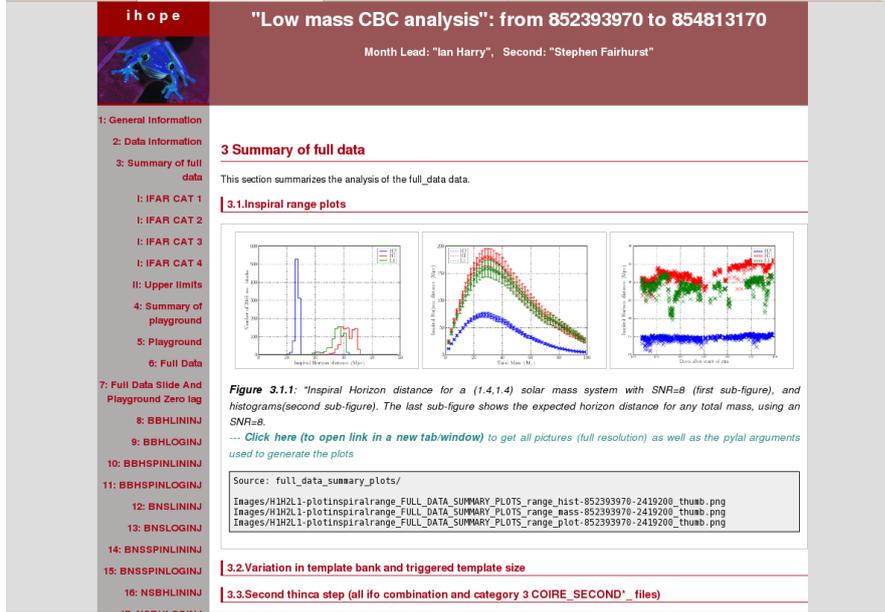


Figure 4.10: An example of the output html pages that are used to easily view ihope results.

	BNS	BHNS	BBH
Component Masses (M_{\odot})	1.35/1.35	5.0/1.35	5.0/5.0
Non-spinning Upper Limit ($\text{yr}^{-1}L_{10}^{-1}$)	8.7×10^{-3}	2.2×10^{-3}	4.4×10^{-4}
Spinning Upper Limit ($\text{yr}^{-1}L_{10}^{-1}$)	...	2.7×10^{-3}	5.3×10^{-4}
Predicted “optimistic” rates	6×10^{-4}	4.7×10^{-5}	1.1×10^{-5}
Predicted “realistic” rates	6×10^{-5}	1.6×10^{-6}	1.1×10^{-7}
Predicted “pessimistic” rates	6×10^{-7}	2.3×10^{-8}	3.4×10^{-9}

Table 4.2: Marginalized upper limits for various CBC systems. The effects of spin on BNS systems are negligible and not reported here. Also shown is the “predicted”, “realistic” and “pessimistic” rate predictions as quoted in [23]. Note that as [23] gives results where the black hole has a mass of $10M_{\odot}$, these numbers have been rescaled to correspond to a black hole mass of $5M_{\odot}$.

values for these upper limits are shown in Table 4.2. Upper limits were also calculated as a function of total mass and as a function of the black hole mass for NSBH systems, these upper limits are shown in Figure 4.11. The quoted upper limits are considerably larger than the rate predictions in [23], so it is not surprising that a signal was not seen in these searches.

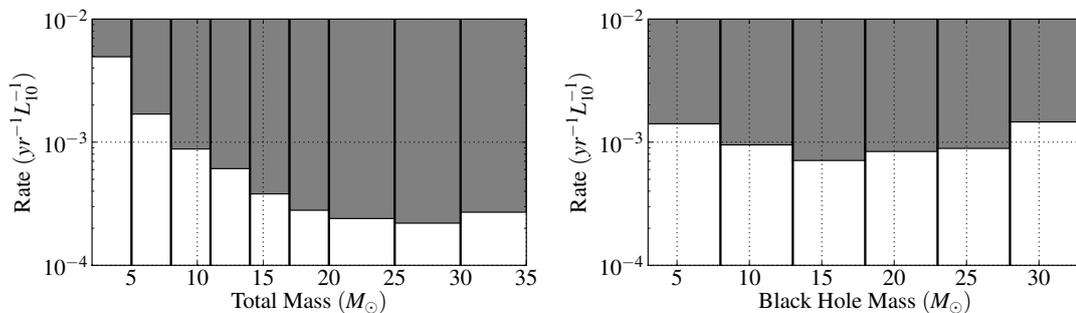


Figure 4.11: The 90% rate upper limits as a function of mass. The left figure gives the upper limit on the rate of coalescence from BBH system as a function of the total mass of the system. The right figure gives the BHNS upper limit as a function of black hole mass, assuming a fixed neutron star mass of $1.35 M_{\odot}$.

The analysis of S6 and VSR2 and VSR3 data is progressing, but as sensitivity has not significantly improved, the probability of making a detection is not significantly increased. However, the detectors are now undergoing upgrades for the advanced era. If these advanced detectors provide a 10-fold sensitivity increase as predicted [21, 22], which would increase the sensitivity in volume by a factor of 1000, then gravitational wave astronomers in the next decade should be regularly making detections!

4.3 The triggered search

In this section we discuss the coincidence “triggered” search for CBCs. These searches are so named because they are triggered by an electromagnetic (EM) observation. If an EM telescope has observed a source, which might have also produced gravitational waves, the information is sent to gravitational wave astronomers. A targeted search is then carried out for a gravitational wave signal coincident with the time and sky location of the electromagnetic observation.

Such triggered searches offer an improvement in sensitivity over all-sky, all-time searches because the time over which a search is performed is greatly reduced. Additionally, knowing the sky location gives the delay in time of arrival between detectors. Both of these things serve to reduce the number of background noise

triggers and thus reduce the SNR a trigger needs to be a detection candidate. Additionally, if an event is simultaneously observed with both electromagnetic and gravitational wave telescopes the science gain can be significantly greater than an observation with only one of these. For example, with joint observations, both the redshift and luminosity distance of a source can be observed, thus offering precision measurements of Hubble’s constant [104].

For the case of CBC searches the most interesting electromagnetically observable source is a gamma-ray burst (GRB). GRBs emit short duration, but extremely high energy gamma rays [105]. There are two types of GRB, believed to be due to different progenitors [106]. The “long GRB” has a longer duration of emission but emits lower energy photons compared to the “short GRB”. For CBC sources short GRBs are especially interesting as the favoured progenitor model is a compact binary coalescence [30, 31].

Unfortunately, the majority of short GRBs with measured distances are believed to originate at distances that are far too large for current gravitational wave detectors to be able to observe [105]. Triggered searches are, however, performed at the time of all short GRBs in case these distance estimates are wrong and to be able to set lower limits on the distance at which they originated. Nevertheless a number of short GRBs do not have measured distances and there is also the occasional short GRB that may have originated within range of our detectors, such as GRB 070201 as we describe in section 4.3.2.

During S5 and VSR1, searches were carried out for CBC signals in coincidence with observed GRBs using a coincidence method very similar to the all sky search [4, 5]. We introduce the coincidence triggered search in this section to put into context the coherent triggered search that we will describe in the next chapter. It is our hope that the coherent search, which we will demonstrate is more sensitive than the coincidence search, will be used to search for CBC signals in coincidence with short GRBs during S6 and VSR2 and VSR3.

In this section we begin, in section 4.3.1, by describing how the triggered coincidence search differs from the all sky search described in section 4.2. We then discuss the results of this search in section 4.3.2 focusing on the special case of GRB 070201.

4.3.1 The coincident triggered search pipeline

To search for gravitational waves in coincidence with a short GRB observation we use a very similar procedure to that described in section 4.2. The same vetoes and cuts are applied and triggers are still ranked by effective SNR (for runs in S5

and VSR1 data). The differences between the two searches are as follows

- The SNR thresholds in each detector are lowered from 5.5 to 4.
- In accordance with astrophysical predictions, we only analyse $[-5, +1)$ seconds around the peak emission time of the short GRB [30, 31, 4].
- Coincidence windows are tightened to account for the sky location, and thus difference in arrival times, being known.
- The efficiency of the search and its tuning is carried out using software injections placed at the same sky location as the GRB.
- The background rates are estimated differently (as described below).

To be able to estimate background rates for any triggers occurring within this 6s “onsource” time triggers in nearby coincident times are used. One block (2048s) of data is analysed around the GRB. The background is then estimated from the 304 available six second long, non-overlapping “offsource” trials within this block, each of which should be separated from the onsource by at least 48s to avoid contamination. This means that a FAR can be calculated without having to time slide the data (though time sliding could be performed on the offsource if a better background estimate was required).

4.3.2 Results

GRB 070201 was a short GRB observed by 4 satellites on the 2nd February, 2007. The sky location error box for this GRB was found to overlap the Andromeda galaxy, see Figure 4.12, taken from [107]. At 770kpc, the Andromeda galaxy was well within the range of detection for the H1 and L1 detectors, which were both taking good quality science data at the time.

A search for any CBC or unmodelled gravitational wave signal at the time and sky location of this GRB was carried out [5], but no gravitational wave signal was observed. Thus, either this GRB was caused by a CBC at a considerably greater distance than the Andromeda galaxy, or this GRB’s progenitor was something else. The favoured theory is that this was a soft gamma-ray repeater originating in the Andromeda galaxy [108, 107].

Apart from GRB 070201, searches have been carried out for CBC gravitational wave signals in coincidence with 21 GRBs in S5 and VSR1 [4]. Again, no gravitational wave signal was detected. Exclusion distances were set for BNS and NSBH progenitor models for each of these GRBs and these can be seen in [4].

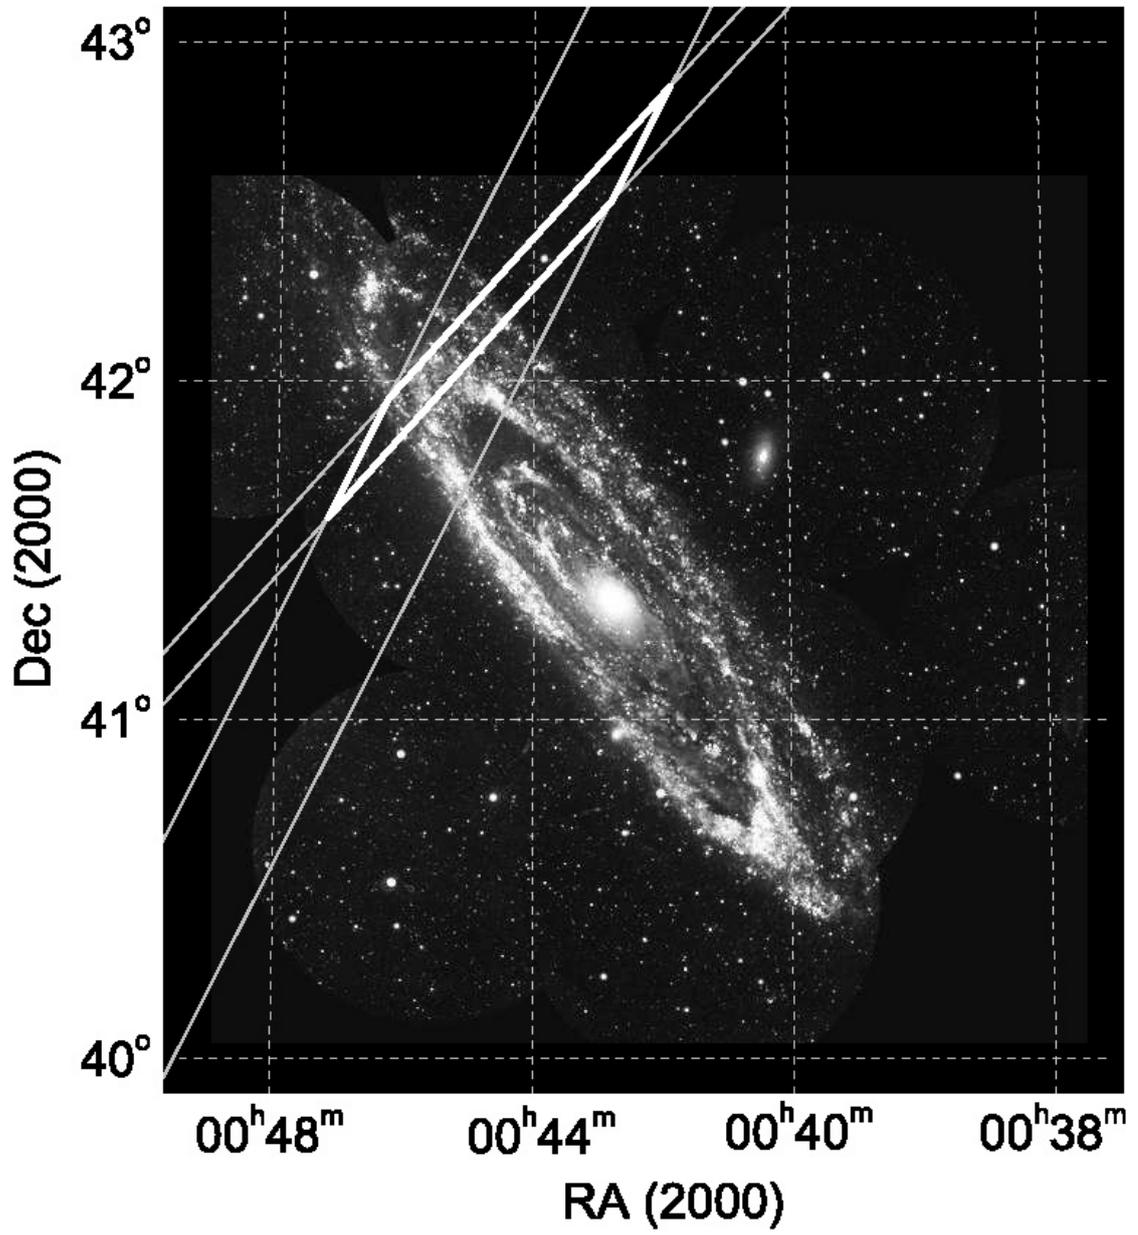


Figure 4.12: The sky location error region of GRB 070201 [107].

Chapter 5

A targeted coherent search for gravitational waves from compact binary coalescences

In this chapter, we introduce a targeted coherent analysis for CBC signals, appropriate for searching for GWs from EM transients such as GRBs. Coherent search methods have been introduced by numerous authors [109, 110, 111]. In coherent searches, data from all operational detectors are combined in a coherent manner before searching for a signal. Additionally, coherent analyses naturally impose the restriction that gravitational waves have only two independent polarizations. However, primarily for computational reasons, many searches have instead made use of a coincidence requirement — namely that a signal with consistent parameters is observed in two or more detectors in the network. Indeed, all previous LIGO and Virgo CBC search result papers have used a coincidence search [5, 4]. Coincidence searches can approach the sensitivity of a fully coherent analysis, but will generally not achieve the same sensitivity.

We begin by deriving the coherent analysis for a templated CBC search. This has been presented in the literature previously [109]. We present an alternative derivation based on the \mathcal{F} -statistic formalism [46] introduced for continuous wave analyses and now widely used. This allows for a more straightforward derivation of the coherent detection statistic. It also allows for a simple comparison to the coincident search and a straightforward derivation of the null stream [112], which by definition contains no gravitational wave signal.

The data output by gravitational wave interferometers are neither stationary nor Gaussian, but are contaminated by noise transients of instrumental and environmental origin. This makes the task of analyzing the data a complex one,

and matched filtering alone is not sufficient to distinguish signal from noise — the most significant events by SNR would typically be due to non-Gaussian transients, or “glitches”, in the data. A significant effort goes into understanding the cause of these glitches [93] and removing times of poor data quality from the analysis. While these efforts greatly reduce the number of glitches, they cannot remove them entirely. Therefore the analysis must also employ methods to distinguish signal from noise transients. In previous CBC searches, signal consistency tests [80, 113] have proved very effective at removing the non-Gaussian background. We provide an overview of the formalism for these χ^2 consistency tests, and then extend these tests to the coherent analysis. We demonstrate that these tests continue to be effective in separating signal from non-stationary noise in the coherent analysis. In addition, coherent analyses naturally lend themselves to multi-detector consistency tests, such as the null stream. We describe a number of consistency tests for the coherent CBC search and again demonstrate their efficacy.

This chapter presents the first implementation of a modelled coherent detection search for CBC signals, which has been run on real LIGO data. We present results from running this analysis on both simulated Gaussian data and real detector data taken from LIGO’s fourth science run (S4). We are able to show that the signal consistency tests we have implemented are sufficient to remove the majority of non-Gaussian transients and render the search almost as sensitive as if the data were Gaussian and stationary. This search is available to be used to search for CBCs associated with GRBs in S6 and VSR2 and VSR3.

The layout of this chapter is as follows. In section 5.1, we describe the formulation of a targeted coherent search for CBC signals. In section 5.2 we discuss an implementation of the null stream formalism and other multi-detector consistency tests. In section 5.3 we describe a number of χ^2 tests that can be applied in a coherent search to try to separate and veto glitches. Finally, in 5.4 we outline an implementation of a targeted, coherent search for CBCs and present results on both simulated and real data.

5.1 Coherent matched filtering

In this section, we describe the coherent matched-filtering search for a gravitational wave signal from a coalescing binary in data from a network of detectors. We restrict attention to binaries where the component spins can be neglected. The description is primarily tailored towards searches where the sky location of the gravitational wave event is known a priori, as is appropriate when performing a

followup of an EM transient such as a GRB [30, 4]. Finally, since all previously published CBC search results [3, 2, 1] have used a coincidence search between detectors, we compare the coherent analysis with the multi-detector coincident analysis.

The coherent analysis for coalescing binary systems has been derived previously using a similar method in [109, 114, 115, 116]. Our presentation makes use of the \mathcal{F} -statistic formalism, introduced in [46]. This was originally defined as a method for performing searches for continuous wave searches and has been regularly used for this task (see for example [48]). It was noted in [117] that the \mathcal{F} -statistic and the multiple detector inspiral statistic derived in [114] are similar and the \mathcal{F} -statistic was adapted to searches for CBC signals in [118].

5.1.1 The binary coalescence waveform

The generic binary coalescence waveform depends upon as many as seventeen parameters. However, we restrict attention to binaries on circular orbits with non-spinning components. This reduces the parameter space to nine dimensions: the two component masses M_1 and M_2 ; the sky location of the signal (θ, ϕ) ; the distance, D , to the signal; the coalescence time of the signal, t_o ; the orientation of the binary, given by the inclination ι , the polarization angle ψ and the coalescence phase ϕ_o . We also assume that the sky location (θ, ϕ) of the signal is known, thereby reducing the number of unknown parameters to seven.

In the radiation frame, where the gravitational wave propagates in the \mathbf{e}_z^R -direction, the gravitational waveform is given by

$$\mathbf{h} = h_+ \mathbf{e}_+ + h_\times \mathbf{e}_\times \quad (5.1)$$

where

$$\begin{aligned} \mathbf{e}_+ &= \mathbf{e}_x^R \otimes \mathbf{e}_x^R - \mathbf{e}_y^R \otimes \mathbf{e}_y^R, \\ \mathbf{e}_\times &= \mathbf{e}_x^R \otimes \mathbf{e}_y^R + \mathbf{e}_y^R \otimes \mathbf{e}_x^R, \end{aligned} \quad (5.2)$$

and the waveforms $h_{+,\times}$ depend upon seven parameters $(M_1, M_2, t_o, D, \iota, \psi, \phi_o)$. The three angles (ι, ψ, ϕ_o) give the relationship between the radiation frame and the source frame (in which \mathbf{e}_z^S lies in the direction of the binary's angular momentum and \mathbf{e}_x^S along the separation between the binary components at t_o). Even for a known sky location, it is necessary to search a seven dimensional parameter space of signals. Naively covering this space with a grid of templates would be pro-

hibitively costly [85]. However, the analysis is greatly simplified by the observation that the last four parameters enter only as amplitude parameters which can be analytically maximized over at minimal cost.¹ Specifically, the two polarizations of the waveform can be expressed as

$$\begin{aligned} h_+(t) &= \mathcal{A}^1 h_0(t) + \mathcal{A}^3 h_{\frac{\pi}{2}}(t) \\ h_\times(t) &= \mathcal{A}^2 h_0(t) + \mathcal{A}^4 h_{\frac{\pi}{2}}(t). \end{aligned} \quad (5.3)$$

The two phases of the waveform are written as h_0 and $h_{\frac{\pi}{2}}$. These depend upon the physical parameters of the system (in this case just the masses) as well as the coalescence time t_o .² \mathcal{A}^μ are constant amplitude terms and are given explicitly as [118, 116]

$$\begin{aligned} \mathcal{A}^1 &= A_+ \cos 2\phi_o \cos 2\psi - A_\times \sin 2\phi_o \sin 2\psi \\ \mathcal{A}^2 &= A_+ \cos 2\phi_o \sin 2\psi + A_\times \sin 2\phi_o \cos 2\psi \\ \mathcal{A}^3 &= -A_+ \sin 2\phi_o \cos 2\psi - A_\times \cos 2\phi_o \sin 2\psi \\ \mathcal{A}^4 &= -A_+ \sin 2\phi_o \sin 2\psi + A_\times \cos 2\phi_o \cos 2\psi, \end{aligned} \quad (5.4)$$

where

$$\begin{aligned} A_+ &= \frac{D_o}{D} \frac{(1 + \cos^2 \iota)}{2} \\ A_\times &= \frac{D_o}{D} \cos \iota, \end{aligned} \quad (5.5)$$

and D_o is a fiducial distance which is used to scale the amplitudes \mathcal{A}^μ and waveforms $h_{0, \frac{\pi}{2}}$. Thus, the amplitudes \mathcal{A}^μ depend upon the distance to the source and the binary orientation as encoded in the three angles (ι, ψ, ϕ_0) . For *any* set of values \mathcal{A}^μ , the expressions (5.4) can be inverted to obtain the physical parameters, unique up to reflection symmetry of the system [118].

The gravitational waveform observed in a detector X is

$$h^X = h^{ij} D_{ij}^X \quad (5.6)$$

where D_{ij}^X denotes the detector response tensor. For an interferometric detector,

¹This was observed for the inspiral signal in [114] and independently for continuous wave signals in [46].

²This decomposition is actually valid for all binaries in which the plane of the orbit does not precess. Thus, binary coalescence waveforms in which the spins are aligned with the orbital angular momentum can also be expressed in this form. However, for generic spin configurations, the orbit will precess and this simple decomposition is no longer applicable.

in the long wavelength limit, the response tensor is given by

$$\mathbf{D}^X = (\mathbf{e}_x^X \otimes \mathbf{e}_x^X - \mathbf{e}_y^X \otimes \mathbf{e}_y^X) \quad (5.7)$$

where the basis vectors \mathbf{e}_x^X and \mathbf{e}_y^X point in the directions of the arms of the detector. It is often convenient to re-express the gravitational wave signal observed in a given detector as

$$h^X(t) = F_+(\theta^X, \phi^X, \chi^X)h_+(t) + F_\times(\theta^X, \phi^X, \chi^X)h_\times(t), \quad (5.8)$$

where the detector response to the two polarizations of the gravitational wave is encoded in the functions

$$F_+(\theta, \phi, \chi) = -\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \cos 2\chi - \cos \theta \sin 2\phi \sin 2\chi, \quad (5.9)$$

$$F_\times(\theta, \phi, \chi) = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \sin 2\chi - \cos \theta \sin 2\phi \cos 2\chi. \quad (5.10)$$

These response functions depend upon the three angles $(\theta^X, \phi^X, \chi^X)$ which relate the detector frame to the radiation frame: θ^X and ϕ^X give the sky location relative to the detector, while χ^X is the polarization angle between the detector and the radiation frames. We have, somewhat unconventionally, allowed for a polarization angle in transforming from source to radiation *and* radiation to detector coordinates. In what follows, we will often find it convenient to fix the angle χ^X by explicitly tying it to the detector (or equatorial) frame; for example, by maximizing the detector (or network) sensitivity to the + polarization. The angle ψ then describes the orientation of the source with respect to this preferred radiation frame.

Since we are considering CBC observed in ground-based detectors, the time that a potential signal would spend in the sensitivity band of any detector will be short (less than 60s for the initial detectors). Thus the change in the source's sky location over the observation time may be neglected, and the angles $(\theta^X, \phi^X, \chi^X)$ can be treated as constants. When working with a network of detectors, it is often useful to work in the equatorial frame. The location of the source (θ, ϕ, χ) is measured relative to this frame, and coalescence time is measured at the Earth's centre. In this case, the location and orientation of the detector X are specified by three angles, which we denote $\vec{\alpha}^X$, and the detector response will depend upon

six angles $(\vec{\alpha}^X, \theta, \phi, \chi)$. Then, the observed signal in a given detector is³

$$h^X(t) = F_+(\vec{\alpha}^X, \theta, \phi, \chi)h_+(t^X) + F_\times(\vec{\alpha}^X, \theta, \phi, \chi)h_\times(t^X), \quad (5.11)$$

where t^X is the time of arrival of the signal at detector X,

$$t^X = t - dt(\vec{\alpha}^X, \theta, \phi, \chi) \quad (5.12)$$

and dt gives the difference in arrival time of the signal between the geocenter and detector, for the given sky position.

Combining the final expressions for the binary coalescence waveform (5.3) and the detector response (5.11), we can express the gravitational waveform observed in a given detector as

$$h^X(t) = \mathcal{A}^\mu(D, \psi, \phi_o, \iota)h_\mu^X(t) \quad (5.13)$$

where the \mathcal{A}^μ are defined in (5.4), h_μ^X are given by

$$\begin{aligned} h_1^X(t) &= F_+^X h_0(t^X) \\ h_2^X(t) &= F_\times^X h_0(t^X) \\ h_3^X(t) &= F_+^X h_{\frac{\pi}{2}}(t^X) \\ h_4^X(t) &= F_\times^X h_{\frac{\pi}{2}}(t^X), \end{aligned} \quad (5.14)$$

and we use the standard summation convention over the repeated index μ .

5.1.2 Multi detector binary coalescence search

Matched filtering theory [78] provides a method for determining whether the signal $h(t, \xi)$, parametrized by the time and other parameters ξ , is present in a noisy data stream. The data output by a detector is

$$s^X(t) = n^X(t) + h^X(t, \xi) \quad (5.15)$$

where $n^X(t)$ is the noise, taken to be Gaussian and stationary. The noise $n^X(t)$ of the detectors is characterized by the noise PSD $S_h^X(f)$ as

$$\langle \tilde{n}^X(f)[\tilde{n}^X(f')]^* \rangle = \delta(f - f')S_h^X(f). \quad (5.16)$$

³We do not give the explicit formula for the response function dependent on the six angles $(\vec{\alpha}^X, \theta, \phi, \chi)$, as the expression is somewhat lengthy. It can be obtained by performing six successive rotations to the detector response tensor to transform from the detector frame, via the equatorial frame, to the radiation frame. The calculation is detailed in [119].

With this, we define the single detector inner product between two time series $a(t)$ and $b(t)$ as

$$(a^X|b^X) = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}^X(f)[\tilde{b}^X(f)]^*}{S_h^X(f)}. \quad (5.17)$$

Then, the likelihood ratio of there being a signal h present in the data is given by:

$$\Lambda(h) = \frac{P(s|h)}{P(s|0)} = \frac{e^{-(s^X - h^X|s^X - h^X)/2}}{e^{-(s^X|s^X)/2}}. \quad (5.18)$$

For a known signal h , with no unknown parameters, in Gaussian noise, the likelihood ratio is optimal in the Neyman-Pearson sense. However, this statistic is not optimal in the presence of non-Gaussian noise, as is discussed in greater detail in sections 5.2 and 5.3. It is often more convenient to work with the log-likelihood,

$$\log \Lambda = (s|h) - \frac{1}{2}(h|h), \quad (5.19)$$

and we will do so.

The likelihood ratio for multiple detectors is a straightforward generalization of the single detector expression (5.18). Assuming that the noise in different detectors are independent, in the sense that

$$\langle \tilde{n}^X(f)[\tilde{n}^Y(f')^*] \rangle = \delta^{XY} \delta(f - f') S_h^X(f), \quad (5.20)$$

the multi-detector inner product is simply given by the sum of the single detector contributions⁴

$$(\mathbf{a}|\mathbf{b}) := \sum_X (a^X|b^X). \quad (5.21)$$

The multi-detector log-likelihood is given by

$$\ln \Lambda = (\mathbf{s}|\mathbf{h}) - \frac{1}{2}(\mathbf{h}|\mathbf{h}). \quad (5.22)$$

Specializing to the case of binary coalescence, we can substitute the known waveform parametrization (5.13) into the general matched filter likelihood ratio (5.22). The multi-detector likelihood ratio becomes

$$\ln \Lambda = \left[\mathcal{A}^\mu(\mathbf{s}|\mathbf{h}_\mu) - \frac{1}{2} \mathcal{A}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}^\nu \right] \quad (5.23)$$

⁴Note that we will explicitly write out the summation over detectors X , and *do not* use implicit summation over these indices.

where the matrix $\mathcal{M}_{\mu\nu}$ is defined as

$$\mathcal{M}_{\mu\nu} := (\mathbf{h}_\mu | \mathbf{h}_\nu). \quad (5.24)$$

The derivative of (5.23) with respect to \mathcal{A}^μ provides the values of \mathcal{A}^μ which maximize the likelihood ratio as

$$\hat{A}^\mu = \left[\mathcal{M}^{\mu\nu}(\mathbf{s} | \mathbf{h}_\nu) \right], \quad (5.25)$$

where, following [120], we take $\mathcal{M}^{\mu\nu}$ to be the inverse of $\mathcal{M}_{\mu\nu}$. We then define the maximized ‘‘coherent SNR’’ via the maximum likelihood ratio as

$$\rho_{\text{coh}}^2 := 2 \ln \Lambda |_{\text{max}} = \left[(\mathbf{s} | \mathbf{h}_\mu) \mathcal{M}^{\mu\nu}(\mathbf{s} | \mathbf{h}_\nu) \right]. \quad (5.26)$$

An identical maximization is performed in the derivation of the \mathcal{F} -statistic, used in searches for gravitational waves from asymmetric neutron stars [46]. In that context, the quantity obtained in (5.26) is typically denoted $2\mathcal{F}$. However, to make a closer connection to previous CBC search methods, we denote this quantity ρ_{coh}^2 .

It is not difficult to show that ρ_{coh}^2 follows a χ^2 distribution with four degrees of freedom in the absence of a signal, and a non-central χ^2 distribution (again with 4 degrees of freedom) when a signal is present. See, for example, [120] for more details. Furthermore, ρ_{coh}^2 is now a function of only the waveform components h_μ and no longer the \mathcal{A}^μ parameters. Thus four of the original seven waveform parameters have been analytically maximized, leaving three to be searched over.

Calculating the maximized likelihood ratio, as well as estimating the parameters \hat{A}^μ requires an inversion of the matrix $\mathcal{M}_{\mu\nu}$. CBC signals will spend a large number of cycles in the sensitive band of the detector and consequently the 0 and $\frac{\pi}{2}$ phases will be (close to) orthogonal. Since the frequency evolves slowly, the amplitudes of the two phases will be close to equal,⁵ i.e.

$$\begin{aligned} (h_0^X | h_{\frac{\pi}{2}}^X) &\approx 0 \\ (h_{\frac{\pi}{2}}^X | h_{\frac{\pi}{2}}^X) &\approx (h_0^X | h_0^X) =: (\sigma^X)^2. \end{aligned} \quad (5.27)$$

⁵Indeed, several CBC waveforms are generated directly in the frequency domain [121], making these equalities exact.

Therefore, the matrix \mathcal{M} simplifies to

$$\mathcal{M}_{\mu\nu} = \begin{pmatrix} A & C & 0 & 0 \\ C & B & 0 & 0 \\ 0 & 0 & A & C \\ 0 & 0 & C & B \end{pmatrix} \quad (5.28)$$

where

$$\begin{aligned} A &= \sum_X (\sigma^X F_+^X)^2 \\ B &= \sum_X (\sigma^X F_\times^X)^2 \\ C &= \sum_X (\sigma^X F_+^X)(\sigma^X F_\times^X). \end{aligned} \quad (5.29)$$

Dominant polarization

Since we have included a polarization angle in both the transformation between equatorial and radiation frame (χ) and between radiation and source frame (ψ), we have the freedom to specify one of these without placing any physical restriction on the signal. The coherent SNR is further simplified by introducing a dominant polarization frame which renders $\mathcal{M}_{\mu\nu}$ diagonal.

Under a rotation of the radiation frame by an angle χ^{DP} , the detector response functions transform as

$$\begin{aligned} F_+^X &\rightarrow F_+^{\text{DP},X} = F_+^X \cos 2\chi^{\text{DP}} + F_\times^X \sin 2\chi^{\text{DP}} \\ F_\times^X &\rightarrow F_\times^{\text{DP},X} = -F_+^X \sin 2\chi^{\text{DP}} + F_\times^X \cos 2\chi^{\text{DP}}. \end{aligned} \quad (5.30)$$

The rotation through χ^{DP} will have an identical effect on all detectors. Thus, there exists a polarization angle χ^{DP} which satisfies

$$C^{\text{DP}} = \sum_X (\sigma^X F_+^{\text{DP},X})(\sigma^X F_\times^{\text{DP},X}) = 0. \quad (5.31)$$

This can be solved to give χ^{DP} as

$$\tan 4\chi^{\text{DP}} = \frac{2 \sum_X (\sigma^X F_+^X)(\sigma^X F_\times^X)}{\sum_X [(\sigma^X F_+^X)^2 - (\sigma^X F_\times^X)^2]}. \quad (5.32)$$

This choice serves to diagonalize the matrix \mathcal{M} . To uniquely determine χ^{DP} , we impose an additional requirement that the network be more sensitive to the + polarization than to the \times polarization. The value of χ^{DP} is a function of the detector network, the source location and waveform; in particular it depends upon $F_{+,\times}^X$ and σ^X . From now on, we assume that we are working in the dominant

polarization frame and drop the DP superscript from our expressions.

The concept of the dominant polarization frame has been introduced previously in un-modelled burst searches [122, 123, 111]. While the idea is very similar, the actual implementation is somewhat different.

In the case that both A and B are non-zero, i.e. that the detector has some sensitivity to both polarizations, the coherent SNR can be written, in the dominant polarization, as

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s}|\mathbf{F}_+\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\frac{\pi}{2}})^2}{(\mathbf{F}_+\mathbf{h}_0|\mathbf{F}_+\mathbf{h}_0)} + \frac{(\mathbf{s}|\mathbf{F}_\times\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_{\frac{\pi}{2}})^2}{(\mathbf{F}_\times\mathbf{h}_0|\mathbf{F}_\times\mathbf{h}_0)}. \quad (5.33)$$

The coherent SNR can then be seen to arise as the quadrature sum of the power in the two phases of the waveform (0 and $\frac{\pi}{2}$) in the two gravitational wave polarizations (+ and \times).

Synthetic + and \times detectors

In the dominant polarization the coherent SNR is comprised of separate + and \times components, with no cross terms. We can go one step further and interpret the coherent SNR as arising from two synthetic detectors, one sensitive to only the + polarization and one sensitive to only the \times polarization. These synthetic detectors are most easily formed by combining the “over-whitened” data streams o^X from the various detectors, where

$$o^X(f) = \frac{s^X(f)}{S_h^X(f)}. \quad (5.34)$$

The over-whitened synthetic data streams are simply

$$o_{+,\times}(f) = \sum_X F_{+,\times}^X o^X(f); \quad (5.35)$$

and the power spectra for these over-whitened data streams are

$$S_{+,\times} = \left(\sum_X \frac{(F_{+,\times}^X)^2}{S_h^X(f)} \right)^{-1}. \quad (5.36)$$

Using this, the un-whitened synthetic data streams are given as⁶

$$s_{+, \times}(f) = \sum_X \frac{F_{+, \times}^X s^X(f)}{S_h^X(f)} \left(\sum_Y \frac{(F_{+, \times}^Y)^2}{S_h^Y(f)} \right)^{-1}. \quad (5.37)$$

In terms of these synthetic detectors the detection statistic becomes

$$\rho_{\text{coh}}^2 = \frac{(s_+ | h_0)_+^2 + (s_+ | h_{\frac{\pi}{2}})_+^2}{(h_0 | h_0)_+} + \frac{(s_{\times} | h_0)_{\times}^2 + (s_{\times} | h_{\frac{\pi}{2}})_{\times}^2}{(h_0 | h_0)_{\times}}, \quad (5.38)$$

where the subscripts $+$, \times on the inner products denote the fact that the power spectrum of the synthetic detectors is used in their evaluation.

Network degeneracy

In many cases, a detector network is much more sensitive to one gravitational wave polarization than the other. In the extreme limit (e.g. co-located and co-aligned detectors such as those at the Hanford site) the network is entirely insensitive to the second polarization. In the dominant polarization frame, the network becomes degenerate as $B \rightarrow 0$ or equivalently

$$\sum_X (\sigma^X F_{\times}^X)^2 \rightarrow 0. \quad (5.39)$$

Thus the network will only be degenerate if $F_{\times}^X = 0$ for all detectors X . If the network is degenerate then it is easy to see that the detection statistic will be degenerate as well. In this case it is logical to remove the \times terms from the detection statistic reducing it to

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0)^2 + (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_{\frac{\pi}{2}})^2}{(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0)}, \quad (5.40)$$

which is χ^2 distributed with two degrees of freedom.

In this formalism the coherent SNR changes abruptly from (5.40) to (5.33). If there is any sensitivity, no matter how small, to the \times polarization, there is an entirely different detection statistic. This arises due to maximization over the \mathcal{A}^μ parameters, allowing them to take any value. Thus, even though a network may have very little sensitivity to the \times polarization, and consequently there will be

⁶There is some ambiguity in fixing the overall normalization of the synthetic detectors. We require that our synthetic detectors have the same sensitivity to the two polarizations as the original network did by requiring $(h_0^{+, \times} | h_0^{+, \times})_{+, \times} = \sum_X (F_{+, \times}^X \sigma^X)^2$.

little chance of observing the waveform in the \times polarization, this is not taken into account in the derivation. A possible modification is to place an astrophysical prior on the parameters (D, ι, ψ, ϕ_0) and propagate this to the distribution of the \mathcal{A}^μ [124]. This would provide a smooth transition from the degenerate to non-degenerate search.

5.1.3 Comparison with coincident search

The single detector search is a special case of the degenerate network (5.40) and can be written as

$$\rho_X^2 = \frac{(s^X|h_0)^2 + (s^X|h_{\frac{\pi}{2}})^2}{(\sigma^X)^2}. \quad (5.41)$$

A coincidence search requires a signal to be observed in two or more detectors, without requiring consistency of the measured waveform amplitudes in the different detectors. In many cases, coincidence searches have made use of different template banks in the different detectors [4, 3, 2] and required coincidence between the recovered mass parameters [90]. A comparison with the coherent analysis discussed above is facilitated if we consider a coincident search where an identical template is used in all detectors, as was done in an analysis of early LIGO data [74]. In this case, the multi-detector coincident SNR is given by

$$\rho_{\text{coinc}}^2 = \sum_X \rho_X^2 = \sum_X \frac{(s^X|h_0)^2 + (s^X|h_{\frac{\pi}{2}})^2}{(\sigma^X)^2}. \quad (5.42)$$

This is not immediately comparable to the coherent SNR given in (5.33). However, both can be re-cast into similar forms by writing the coincident SNR as

$$\rho_{\text{coinc}}^2 = \sum_{X,Y} \sum_{i=0, \frac{\pi}{2}} \left(s^X \left| \frac{h_i}{\sigma^X} \right. \right) [\delta^{XY}] \left(s^Y \left| \frac{h_i}{\sigma^Y} \right. \right), \quad (5.43)$$

and the coherent SNR can be written as

$$\rho_{\text{coh}}^2 = \sum_{X,Y} \sum_{i=0, \frac{\pi}{2}} \left(s^X \left| \frac{h_i}{\sigma^X} \right. \right) [f_+^X f_+^Y + f_\times^X f_\times^Y] \left(s^Y \left| \frac{h_i}{\sigma^Y} \right. \right), \quad (5.44)$$

where we have defined the orthogonal unit vectors (in detector space) f_+^X, f_\times^X as

$$f_{+, \times}^X = \frac{\sigma^X F_{+, \times}^X}{\sqrt{\sum_Y (\sigma^Y F_{+, \times}^Y)^2}}. \quad (5.45)$$

The SNR of the coincident search (5.43) is simply the sum of all power consistent with the template waveform in each detector. The coherent SNR (5.44) makes use of the fact that gravitational waves have only two polarizations to restrict the accumulated SNR to the physical subspace spanned by f^+ and f^\times . For a signal, the power will lie entirely in this subspace, while noise in the detectors will contribute to all components of the coincident SNR. Thus, the coherent analysis obtains precisely the same signal SNR but reduces the noise background. Specifically, the coherent SNR acquires contributions from four noise degrees of freedom, while the coincident SNR has $2N$ noise degrees of freedom, where N indicates the number of active detectors. For a non-degenerate two detector network, the coincident and coherent SNRs are equal as in this case $f_+^X f_+^Y + f_\times^X f_\times^Y = \delta^{XY}$.

In the case where a network is sensitive to only one polarization, the coherent SNR is constructed solely from the f_+^X direction and coherent SNR is χ^2 distributed with 2 degrees of freedom.

Finally, we note that restricting to the coherent SNR can help to separate transients from gravitational wave signals as those transients which do not contribute power to the signal space will be ignored. However, many noise transients will contribute to the coherent SNR and more active methods of removing them are required. These methods are the focus of sections 5.2 and 5.3.

5.2 Signal consistency between detectors

As discussed in the introduction, due to the presence of non-Gaussian noise transients, it is essential to make use of signal consistency requirements within search algorithms to distinguish glitches from gravitational wave signals.

Multi-detector analyses have made good use of signal consistency between detectors (see e.g. [111]). A particularly powerful test is the use of a “null stream” [112] which, by construction, contains no gravitational wave signal. Many noise transients will contribute power to the null stream and can therefore be eliminated as candidate events. In addition, requiring that the gravitational wave signal is recovered consistently between detectors can eliminate other noise transients, in our case this is equivalent to imposing restrictions on the recovered values of the parameters, \hat{A}^μ . These two methods will be considered in turn. For matched filtering searches, requiring consistency between the observed signal and template waveform has also proven very powerful [80]. A full description of waveform consistency tests is presented in the next section.

5.2.1 Null stream consistency

The gravitational waveform consists of two polarizations. Thus for networks comprising three or more detectors it is possible to construct one or more null data streams which contain no gravitational wave signal [112]. In the context of a coherent search for CBC signals, the null consistency tests arise quite naturally. In section 5.1.3, we noted that the coherent SNR can be thought of as a projection of the coincident multi-detector SNR onto a four dimensional signal subspace. The remaining dimensions in the coincident search do not contain any gravitational wave signal, but will be subject to both Gaussian and non-Gaussian noise. Thus, we can define the null SNR as

$$\begin{aligned}\rho_N^2 &= \rho_{\text{coinc}}^2 - \rho_{\text{coh}}^2 \\ &= \sum_{X,Y} \sum_{i=0, \frac{\pi}{2}} \left(s^X \middle| \frac{h_i}{\sigma^X} \right) [N^{XY}] \left(s^Y \middle| \frac{h_i}{\sigma^Y} \right),\end{aligned}\tag{5.46}$$

where

$$N^{XY} = \delta^{XY} - f_+^X f_+^Y - f_\times^X f_\times^Y.\tag{5.47}$$

A gravitational wave signal matching the template h will provide no contribution to the null SNR, so we expect that, for signals, this statistic will be χ^2 distributed with $(2N - 4)$ degrees of freedom. A noise transient that is incoherent across the data streams, may give a large coherent SNR, but it is likely to also give a large null SNR. Thus requiring a *small* null SNR will prove effective at distinguishing incoherent noise transients from real gravitational wave signals. Since the definition of the null SNR (5.46) makes use of the template waveform, gravitational waveforms which do not match the template h can contribute to the null SNR.

We can go one step further and introduce synthetic null detectors in analogy with the synthetic $+$ and \times detectors. For concreteness, we describe the three detector case, but this can be extended in a straightforward manner to larger networks. When working with a three detector network, there is a single null direction,

$$n^X = \sum_{Y,Z} \epsilon_{XYZ} f_+^Y f_\times^Z,\tag{5.48}$$

where ϵ_{XYZ} denotes the Levi-Cevita symbol, and the projection onto the null space is given by $N^{XY} = n^X n^Y$. Then, the over-whitened synthetic null detector

is

$$o_N(f) = \sum_X \frac{n^X}{\sigma^X} o^X(f). \quad (5.49)$$

The power spectrum of the null stream is⁷

$$S_N(f) = \left(\sum_X \frac{(n^X)^2}{(\sigma^X)^2 S_h^X(f)} \right)^{-1} \quad (5.50)$$

and the un-whitened null stream is

$$s_N(f) = \left(\sum_X \frac{n^X s^X(f)}{\sigma^X S_h^X(f)} \right) \cdot S_N(f). \quad (5.51)$$

Finally, the null SNR can be written as

$$\rho_N^2 = \frac{(s_N|h_0)_N^2 + (s_N|h_{\frac{\pi}{2}})_N^2}{(h_0|h_0)_N}. \quad (5.52)$$

The null SNR described above differs from the multi-detector null stream formalism introduced in [112] and used by several other authors. A null stream is constructed to be a data stream which contains no contribution from the h_+ and h_\times gravitational waveforms, regardless of the details of the waveform. To provide a concrete comparison between the null stream and null SNR, we again restrict attention to a three detector network. The null stream is explicitly constructed as

$$s_{\text{Null}}(f) = \sum_{X,Y,Z} \epsilon_{XYZ} F_+^X F_\times^Y s^Z(f). \quad (5.53)$$

By comparing the null stream in (5.53) with the synthetic null detector (5.51), it is clear that these will generically differ. To get an insight into the differences, consider a network with two co-located detectors A and B , with power spectra $S_h^A(f)$ and $S_h^B(f)$ respectively, and a third detector C which is sensitive to the other polarization of gravitational waves. For this network, the null stream will be a combination of only the A and B detector data. The power spectrum of the null stream is

$$S_{\text{Null}}(f) = S^A(f) + S^B(f). \quad (5.54)$$

⁷In this case, there is a normalization ambiguity. For the synthetic plus and cross streams, it was natural to require that the synthetic detectors have the same sensitivities as the original network. For the null stream this is not feasible as the network has zero sensitivity to a signal in the null stream, so we normalize such that $(h_0|h_0)_N = 1$.

while for the synthetic null detector, it is

$$\frac{1}{S_{\text{N}}(f)} = \frac{1}{(\sigma^A)^2 S^A(f)} + \frac{1}{(\sigma^B)^2 S^B(f)}. \quad (5.55)$$

Thus, if the power spectra of detectors A and B are identical, then the two null streams are also identical. In the extreme case that the sensitivity bands of the two detectors do not overlap at all, then there is no null stream ($S_{\text{Null}} \rightarrow \infty$). However, the null SNR need not vanish and is similar to a two bin version of the χ^2 test described in section 5.3.4. Thus, it is possible to construct scenarios in which these two null stream formulations differ significantly.

For the most part, the power spectra of the ground based detectors are rather comparable. So, in general there will not be a significant difference between these two forms. There are advantages to both methods. The null stream is designed to cancel all gravitational wave signals from the data, thus making it more robust when the signal is not well known. However, by making use of the template signal, there are instances in which the null SNR provides a more powerful consistency test. Furthermore, it has a computational benefit in that it does not require the production of a null stream — all manipulations are performed on the single detector SNR data streams which are subsequently separated into coherent and null components. In practice we have found very little difference in performance, and choose to compute the null SNR (5.46) for computational simplicity.

Finally, we note that both null stream formalisms will perform optimally only if the three detectors have similar sensitivities. In the case where one detector is significantly less sensitive than the others, the null stream will generally tend to the data of that least sensitive detector. Also, the null formalisms described here will only completely cancel a gravitational wave signal provided that the calibration of the data streams is accurate, any error in calibration will lead to a signal surviving in the null stream.

5.2.2 Amplitude consistency

The four amplitude parameters \mathcal{A}^μ , encoding the distance to and orientation of the binary system, can take any values. Indeed, any set of \mathcal{A}^μ corresponds to unique values of the distance, inclination angle, coalescence phase and polarization angle, up to symmetries of the system. However, some of these values will be significantly more likely to occur astrophysically than others. For example, the number of binary coalescence events is expected to be approximately proportional to star formation rate [23] and consequently should be roughly uniform in volume. Thus,

events are more likely to occur at a greater distance. Similarly, the gravitational wave amplitude, at a fixed distance, is greater for face on signals than edge on ones, as is clear from (5.5). Therefore, we are more likely to detect face on signals at a large distance than nearby, edge on ones. Consequently, certain values of \mathcal{A}^μ are astrophysically more likely than others. This can be taken into account in a Bayesian manner by marginalizing over an appropriate distribution for the \mathcal{A}^μ rather than performing a maximization [124].

The distribution of noise events will follow its own characteristic distribution. For Gaussian noise, the expected distributions of the \mathcal{A}^μ can be explicitly calculated. Non-stationarities in the data will again produce a different distribution of amplitudes which cannot be analytically modelled in a straightforward manner. Specifically, the majority of transients are caused by a disturbance or glitch in a single detector with little or no signal in the other detectors. For networks with three or more detectors, this will typically be inconsistent with a coherent signal across the network, leading to a large value of the null SNR. In certain scenarios, most notably for two detector networks, there will be a consistent set of values for the \mathcal{A}^μ . However, these values carry the characteristic signature of a glitch. Specifically, the SNR contributions will typically be consistent with a nearby, close to edge on system ($A_\times \approx 0$), with a very specific orientation to provide essentially no response in all but one detector. Thus, the glitch distribution of the \mathcal{A}^μ parameters, will be significantly different from the distribution expected for gravitational wave signals. In the remainder of this section, we explore the possibility of making use of the extracted \mathcal{A}^μ parameters to distinguish between glitches and signals. Unlike the null stream, amplitude consistency tests are available for two detector networks. They should be especially useful in the case of the two 4km LIGO instruments, which have similar sensitivities to the majority of points on the sky.

We have argued that the majority of gravitational wave signals will originate from (close to) face on binaries while the majority of noise transients will mimic (close to) edge on binaries. The recovered value of the inclination angle ι should then serve to separate signals from noise. To investigate this, we simulated a large number of CBC signals and a large number of noise glitches; added Gaussian noise and plotted the recovered inclination angle in Figure 5.1. The glitches were generated as events with a large SNR in one detector coincident with Gaussian noise in a second detector. The signals were separated into two groups: the first with only face on binaries ($|\cos \iota| = 1$) and the second a uniform distribution over the two sphere (uniform in $\cos \iota$ and ψ) of the binary orientation. In both cases, they were distributed uniformly in volume and orbital phase. We also consider two

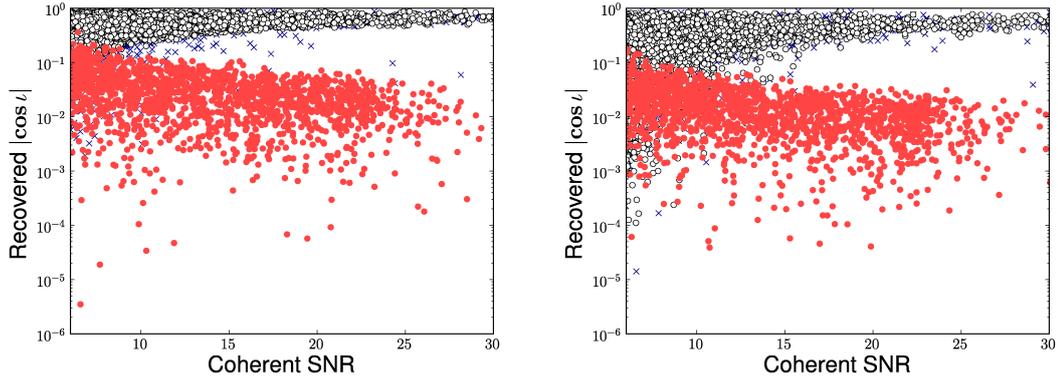


Figure 5.1: The distribution of the recovered inclination angle plotted against coherent SNR for optimally oriented signals (unfilled circles), uniformly distributed orientations (blue crosses) and simulated glitches (red circles). The left figure shows a network configuration where we are equally sensitive to both gravitational wave polarizations. The right figure shows a configuration where we are 5 times more sensitive to the $+$ polarization than to the \times .

different network configurations, both containing two equally sensitive detectors. In the first case one detector is sensitive to $+$ and the other to \times polarization; in the second case both detectors have strong and equal sensitivity to the $+$ polarization and weak but opposite sensitivity to the \times polarization — rather typical for the Hanford, Livingston network. For both sets of signals and choices of network, there is a clear distinction between signal and glitch distribution. However, there is a clear downwards bias on the recovered values of ι . This can be understood by looking at the expressions for A_+ and A_\times . For face on binaries, these will be equal but, in the presence of noise, A_\times will be reconstructed to be somewhat smaller than A_+ . A relative difference of only 5% leads to a recovered inclination of 45° , so even for loud signals there can be large discrepancy between the actual and recovered inclination angle.

Despite the difference in distribution between signal and noise, there is also a significant overlap of the populations at low SNRs. Consequently, any threshold imposed on the recovered inclination angle is liable to either reject a fraction of signals or pass a fraction of glitches. It is, however, quite possible that knowledge of these expected distributions could be folded into the detection statistic in a Bayesian manner.

We have found that using the observed SNR in the individual instruments to be a more effective discriminator of signal and noise. To demonstrate the efficacy of such an approach, in Figure 5.2 we plot the single detector SNR as a function of the coherent SNR for the same population of glitches and the two classes of signals

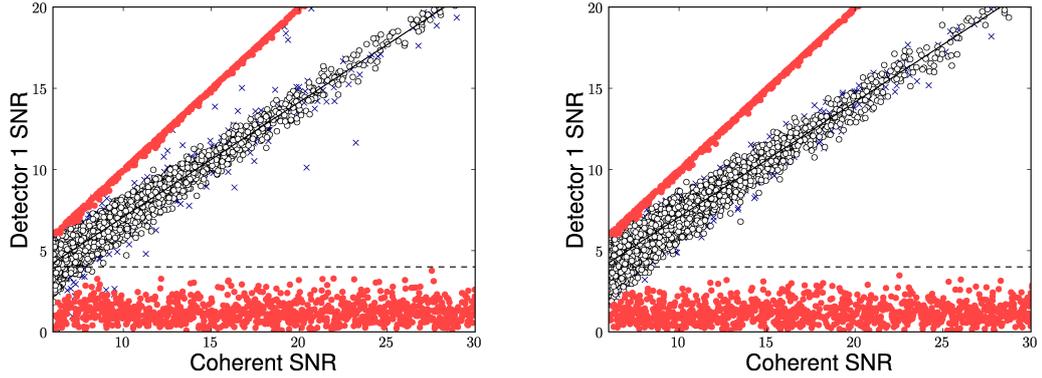


Figure 5.2: The distribution of single detector SNR plotted against coherent SNR for optimally oriented signals (unfilled circles), uniformly oriented signals (blue crosses) and glitches (red circles). The left figure shows a network of two equally sensitive detectors, where one detector sees only the $+$ polarization and the second detector sees only the \times polarization. The right figure shows a similar network where both detectors have strong and equal sensitivity to the $+$ polarization and weak but opposite sensitivity to the \times polarization. This is typical for the Hanford, Livingston network. The diagonal solid black line shows the expected SNR for the optimally oriented signals. The horizontal dashed black line indicates an SNR of 4.

(face on and uniformly distributed orientation) described above. The glitches fall into two groups depending upon which detector suffered the glitch. Since our model detectors are equally sensitive, then on average one expects each detector to accrue $1/\sqrt{2}$ of the coherent SNR. Even allowing for non-optimally oriented signals and the addition of Gaussian noise, the signals follow this expectation. Only a small number of signals are found with SNRs inconsistent with the expected values, these are ones that have very specific orientations. Overall, the signal and glitch populations are very well separated, at least until the coherent SNR becomes rather small.

The most effective strategy we have found is to require that all events have an SNR above 4 in the two most sensitive detectors in the network. The cut is illustrated in Figure 5.2. This strategy removes the majority of glitch signals while having a negligible effect on the signal population at large SNR. For lower SNR the signals which are lost due to this cut would be unlikely to be detection candidates as Gaussian noise alone produces similar events.

5.3 Coherent χ^2 tests

Data from gravitational wave detectors contain numerous non-stationarities due to both instrumental and environmental causes. These non-stationarities, or glitches, typically do not match well with the CBC waveform. However, they often contain enough power that, even though the match with the template is poor, a large SNR is observed. In the previous section, we have seen how the use of various coherent consistency tests can mitigate this problem. Additionally, a number of other signal consistency tests have been implemented [80, 113] and used in searches for CBC signals [3, 2, 1]. These tests are all designed to eliminate glitches which have a different signal morphology than the template waveform. This is essentially done by testing whether the detector data orthogonal to the signal is well described as Gaussian and stationary — for a glitch, there will be residual power which does not match the template waveform. These tests are commonly known as “ χ^2 tests” as they construct a statistic which is χ^2 distributed in the presence of Gaussian noise plus a signal matching the template waveform. If the data contains a glitch, the χ^2 statistic will generally have a large value, thereby allowing for differentiation of signal from non-stationary noise. In this section, we briefly review the general formulation of χ^2 tests before presenting a detailed description of three such tests which have been implemented for the coherent search described in section 5.1.

5.3.1 A general framework for χ^2 tests

Consider the data from a gravitational wave detector at a time t which has produced a large SNR when filtered against a template $h(t)$. Generically, the data $s(t)$ can be decomposed as

$$s(t) = n(t) + Ah(t) + Bg(t) \quad (5.56)$$

where $n(t)$ represents a Gaussian noise component, $h(t)$ is the template waveform, $g(t)$ is an additional non-Gaussian noise contribution to the data stream and A and B are amplitude factors. The glitch contribution $g(t)$ is taken to be the power orthogonal to $h(t)$ and both $g(t)$ and $h(t)$ are normalized, so that

$$(g|g) = 1, (h|h) = 1, (g|h) = 0. \quad (5.57)$$

In order to construct a χ^2 test, we must introduce an additional set of waveforms T^i . These waveforms are required to be orthonormal and orthogonal to

h ,

$$(h|T^i) = 0, (T^i|T^j) = \delta^{ij}. \quad (5.58)$$

Furthermore, for the χ^2 test to be effective, the T^i must have a good overlap with the glitch waveform $g(t)$.

The χ^2 discriminator is constructed as

$$\chi^2 = \sum_{i=1}^N (T^i|s)^2. \quad (5.59)$$

When the data comprises only signal plus Gaussian noise, i.e. $B = 0$ in equation (5.56),

$$\chi^2 = \sum_{i=1}^N (T^i|n)^2 \quad (5.60)$$

and the statistic is the sum of squares of independent Gaussian variables with zero mean and unit variance. Thus the test is χ^2 distributed with N degrees of freedom, with a mean and variance of

$$\langle \chi^2 \rangle = N, \quad \text{Var}(\chi^2) = 2N. \quad (5.61)$$

This is true for *any* set of waveforms T^i given the above assumptions.

In the case where the data are not an exact match to the signal, we take both A and B non-zero, i.e. any signal or glitch can be decomposed into a part $Ah(t)$ proportional to the template under consideration plus a second orthogonal contribution $Bg(t)$. Clearly, for different glitches, the waveform $g(t)$ as well as the amplitude factor B will be different. In this case the χ^2 test takes the form

$$\chi^2 = \sum_{i=1}^N [(T^i|n)^2 + 2B(T^i|n)(T^i|g) + B^2(T^i|g)^2]. \quad (5.62)$$

This has a mean

$$\langle \chi^2 \rangle = N + B^2 \sum_i (T^i|g)^2 \quad (5.63)$$

and a variance

$$\text{Var}(\chi^2) = 2N + 4B^2 \sum_i (T^i|g)^2. \quad (5.64)$$

The χ^2 test is distributed as a non-central χ^2 distribution with N degrees of

freedom and a non-centrality parameter [80]

$$\lambda = B^2 \sum_{i=1}^N (T^i | g)^2. \quad (5.65)$$

The challenge in constructing a χ^2 test is to select the basis waveforms T^i such they have large overlaps with the observed glitches in the data. If this is done successfully, then any glitch producing a large SNR will also give a large value of χ^2 , inconsistent with a signal in Gaussian noise.

In many cases, there is some uncertainty in the template waveform. For example, the post-Newtonian (PN) expansion used in generating CBC waveforms is truncated at a finite (typically 3 or 3.5 PN [69]) order and there will be differences between this analytically calculated waveform and the one provided by nature. There are similar uncertainties in waveforms obtained from numerical relativity simulations [125]. Additionally, to search the full parameter space of coalescing binaries, a discrete template bank is used which allows for some mismatch between the templates and any potential signal within the parameter space [85]. Normally the template bank is created so that the mismatch is no larger than 3% at any point in the parameter space. Finally, there are uncertainties in instrumental calibration [81] which will affect the match between signal and template.

We model these effects by parametrizing the signal as

$$H(t) = A(\sqrt{1 - \epsilon^2} h(t) + \epsilon m(t)), \quad (5.66)$$

where $m(t)$ is the component of H that is orthogonal to h [$(m|h) = 0$] and ϵ encodes the mismatch between signal and template in the sense that

$$1 - \frac{(H|h)}{\sqrt{(H|H)(h|h)}} = 1 - \sqrt{1 - \epsilon^2} \approx \epsilon. \quad (5.67)$$

In most cases, it is reasonable to assume a mismatch of less than 5%. The obvious counter-example is when searching for highly spinning systems using non-spinning waveforms, see e.g. [126, 127].

Since (5.66) is a special case of (5.56) it follows directly that the mean and

variance of the χ^2 test in the presence of a mis-matched signal are

$$\begin{aligned}\langle \chi^2 \rangle &= N + A^2 \epsilon^2 \sum_{i=1}^N (T^i | m)^2 \\ \text{Var}(\chi^2) &= 2N + 4A^2 \epsilon^2 \sum_{i=1}^N (T^i | m)^2.\end{aligned}\tag{5.68}$$

As the SNR of the signal is proportional to A , the expected χ^2 value for a mis-matched signal increases with the strength of the signal. However, for mismatched signals $\chi^2 \propto \epsilon^2 A^2$ while for glitches $\chi^2 \propto B^2$ and provided $\epsilon A \ll B$ the two can be separated. See [80] for a more detailed discussion.

When introducing the χ^2 test, we assumed that the T^i were orthonormal and orthogonal to the template waveform h . In practice, this can be difficult to guarantee. The signal consistency tests discussed in the remainder of this section are constructed from gravitational waveforms. If one picks a set of gravitational waveforms, t^i , there is no guarantee that they will be either orthonormal or orthogonal to h . We can, at least, construct waveforms which are orthogonal to h by introducing

$$T^i = \frac{t^i - (t^i | h)h}{\sqrt{1 - (t^i | h)^2}}.\tag{5.69}$$

While this ensures $(h | T^i) = 0$ it does not guarantee orthonormality of the T^i , $(T^i | T^j) = \delta^{ij}$. Thus this method will not produce a χ^2 distribution, and will instead form a generalized χ^2 distribution. The mean of the distribution remains N but the variance is increased,

$$\text{Var}(\chi^2) = 2N + 2 \sum_{i \neq j} (T^i | T^j)^2.\tag{5.70}$$

It has been found, however, that this does not present a significant obstacle to using these tests, especially as the thresholds are tuned empirically [128].

Multi-detector χ^2 tests

In section 5.1, we derived a coherent multi-detector search for coalescing binaries. The search involves filtering four waveform components h_μ against the multi-detector data stream. Our initial discussion of χ^2 tests was limited to the description of a single phase template waveform h and test waveforms T^i . The extension to a two phase waveform has been described previously [80] and here we extend that to a four component waveform across multiple detectors, as is appropriate

for this search. We begin by noting that the four waveform components h_μ are orthogonal in the dominant polarization basis. They are, however, not generally normalized, as

$$(\mathbf{h}_\mu | \mathbf{h}_\nu) = \mathcal{M}_{\mu\nu} = \text{diag}(A, B, A, B), \quad (5.71)$$

where A and B are defined in (5.29). Thus, we first normalize so that

$$(\hat{\mathbf{h}}_\mu | \hat{\mathbf{h}}_\nu) = \delta_{\mu\nu}. \quad (5.72)$$

To construct a network χ^2 test, we require a set of (4-component) normalized, test waveforms \hat{t}_μ^i . The components

$$T_\mu^i = \frac{\hat{t}_\mu^i - \sum_\nu (\hat{\mathbf{t}}_\mu^i | \hat{\mathbf{h}}_\nu) \hat{h}_\nu}{\sqrt{1 - \sum_\sigma (\hat{\mathbf{t}}_\mu^i | \hat{\mathbf{h}}_\sigma)^2}}, \quad (5.73)$$

constructed to be orthogonal to h_μ , are used in the χ^2 test. Thus, the coherent, multi-detector χ^2 test is

$$\chi^2 = \sum_{\mu=1}^4 \sum_{i=1}^N (\mathbf{T}_\mu^i | \mathbf{s})^2. \quad (5.74)$$

Provided the test waveforms are orthonormal, in the sense that

$$(\mathbf{T}_\mu^i | \mathbf{T}_\nu^j) = \delta^{ij} \delta_{\mu\nu}, \quad (5.75)$$

the distribution for a signal matching h_μ plus Gaussian noise will be χ^2 distributed with $4N$ degrees of freedom. As for the single phase filter, we cannot always guarantee (5.75) is satisfied, although it is relatively simple to ensure the four components of a given template *are* orthogonal. This means that the statistic will not, in general, be χ^2 distributed: The mean remains $4N$ but the variance increases to

$$\text{Var}(\chi^2) = 8N + 2 \sum_{i,j=1}^N \sum_{\mu,\nu=1}^4 [(T_\mu^i | T_\nu^j)^2 - \delta^{ij} \delta_{\mu\nu}]. \quad (5.76)$$

5.3.2 The coherent bank χ^2 test

The bank χ^2 test was designed to test the consistency of the observed SNR across different templates in the bank at the time of a candidate signal. It was first described in [113] for the case of a single detector. A glitch will typically cause a high SNR in many templates across the bank, while a real signal will give a well prescribed distribution of SNR across the template bank.

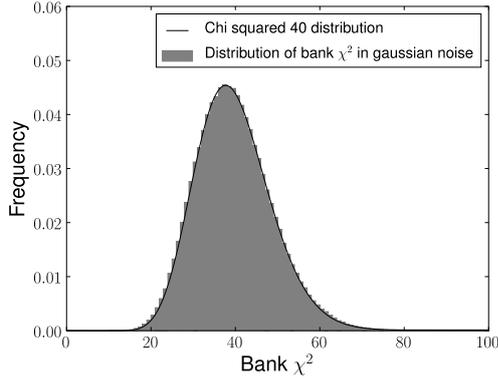


Figure 5.3: The distribution of the bank χ^2 test for a single template h , with a bank of size 10. The plot shows the distribution of the bank veto calculated for every time sample in 128s of simulated Gaussian data (with no signal present). In the case that the ten bank templates are orthogonal, the expected distribution is χ^2 with 40 degrees of freedom (shown as the solid black line). As can be seen, the actual distribution follows the expected one closely.

The bank χ^2 makes use of other CBC templates as the waveforms t^i to construct the χ^2 test. These N templates are taken from different points across the mass space. In implementing the bank χ^2 , we choose a *fixed* set of template waveforms t^i which remain the same for every template h in the search template bank. The bank χ^2 statistic is then constructed following (5.73) and (5.74). The test is most effective when the set of T_μ^i is close to orthogonal [113] so we select templates which are well distributed across the mass space, ensuring the overlaps $(T_\mu^i | T_\nu^j)$ are small for $i \neq j$. Figure 5.3 shows the distribution of the bank χ^2 for a single template filtered against Gaussian noise. The set of fixed bank waveforms consisted of ten waveforms distributed over the full mass parameter space. Using these waveforms, the deviation from a χ^2 distribution is negligible.

For the bank χ^2 to be effective, glitches in the data must have a good overlap with a reasonable fraction of the templates t^i . While, in general, it is difficult to predict the composition of glitches in the data, it seems reasonable to assume that glitches which produce a large SNR for the template h will also have a good overlap with other waveforms in the template space. Thus, the set of templates which is spread across the parameter space is suitable.

5.3.3 The coherent autocorrelation χ^2 test

Filtering a gravitational wave template against data containing a matching signal produces a peak in the SNR at the time of the signal. Furthermore, there is a

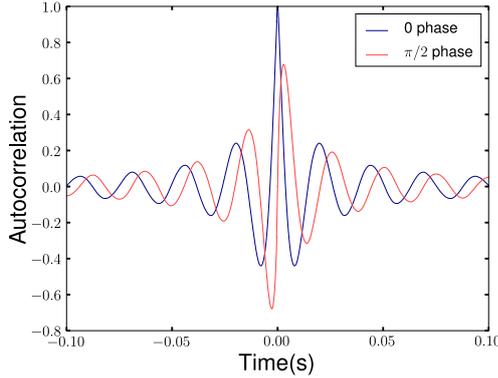


Figure 5.4: The single detector auto-correlation of a gravitational wave inspiral signal from a 1.4,1.4 solar mass binary neutron star. Both phases of the waveform are shown.

characteristic shape of this peak which depends upon the template waveform and also the noise power spectrum of the data. An example of this “autocorrelation” for a BNS template is shown in Figure 5.4. A noise transient in the data will produce a peak in the SNR but it will typically lack the characteristic shape produced by a genuine CBC signal.

The “auto” χ^2 test was designed to test the consistency of the SNR peak [113]. It is a similar test to the bank χ^2 , but where the bank χ^2 investigates consistency in SNR across the mass space, the auto χ^2 tests for consistency of the SNR time series. The set of templates t^i are chosen to be the original template h with time shifts δt^i applied. The values of δt^i are all unique and chosen to be of the same time-scale as the auto-correlation of the template waveform (typically 0.1s or less) and the duration of non-stationarities in the data, which is similar.

In Figure 5.5, we show the distribution of the auto χ^2 for a single template waveform filtered in Gaussian data. For this result, forty waveforms t^i were used, equally spaced with a 1 ms spacing, and all with coalescence times prior to that of h . Thus, the auto χ^2 is testing the consistency of the SNR time series for 0.04 seconds prior to the SNR peak. The overlap ($t^i|t^j$) depends only upon the difference $\delta t^i - \delta t^j$ and Figure 5.4 shows clearly that a significant fraction of the overlaps are far from zero. Consequently, the auto χ^2 test has a distribution with a large deviation from a χ^2 distribution with $4N$ degrees of freedom.

5.3.4 The coherent χ^2 test

The “standard” χ^2 test originally proposed in [80] has been used as a discriminator in many gravitational wave searches for CBCs. Given the template waveforms

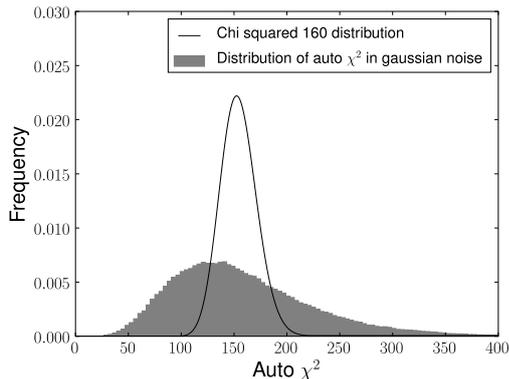


Figure 5.5: The distribution of the auto χ^2 test for a single template h , generated with 40 time shifted templates, with shifts between 0.001 and 0.04 seconds. The plot shows the distribution of the auto veto calculated for every time sample in 128s of simulated Gaussian data (with no signal present). In the case that the forty time shifted templates are orthogonal, the expected distribution is χ^2 with 160 degrees of freedom (shown in black). As can be seen, the actual distribution differs significantly from this due to the non-orthogonality of the t^i waveforms.

and the detector sensitivity, it is possible to predict the accumulation of SNR as a function of frequency. By calculating the observed SNR contribution from a number of frequency bins, and comparing to the predicted value, one can construct a χ^2 consistency test.

Formally, given a template h which produced a candidate signal with an SNR of ρ , calculate N non-overlapping frequency windows such that the expected SNR is ρ/N in each. Then, calculate the actual SNR ρ^i in each of these frequency bins and compare with the expected value by calculating

$$\chi^2 = N \sum_i^N (\rho^i - \rho/N)^2. \quad (5.77)$$

For a gravitational wave signal matching the template h plus Gaussian noise, this statistic will be χ^2 distributed with $N - 1$ degrees of freedom. Written in the form (5.77) it appears different from the general case we discussed earlier. In [80] it was shown that it can be re-expressed in the form (5.59).

This χ^2 test can be extended to coherent, multi-detector searches. Indeed, in [129], the construction was applied to a coherent search for continuous gravitational waves. Here, we present the extension to a coherent CBC search. First, define

$$\rho_\mu^i = \frac{(\mathbf{s}|\mathbf{h}_\mu^i)}{\sqrt{(\mathbf{h}_\mu|\mathbf{h}_\mu)}} \quad (5.78)$$

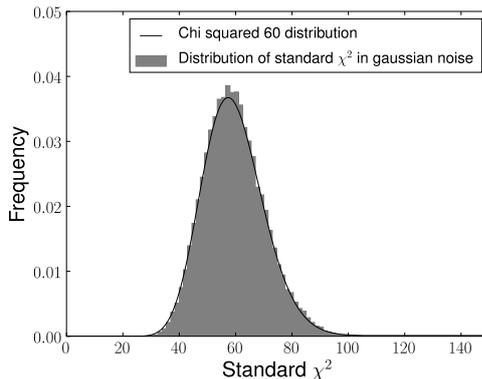


Figure 5.6: The distribution of the standard χ^2 test for a single template h , split into 16 non-overlapping frequency bins. The plot shows the distribution of the χ^2 test calculated for every time sample in 128s of simulated Gaussian data (with no signal present). The observed distribution of values (shown in grey) is an excellent match with the expected χ^2 distribution with sixty degrees of freedom (shown in black).

to be the SNR contribution in the i th frequency bin to the SNR.⁸ The coherent χ^2 statistic is then constructed as

$$\chi^2 = N \sum_{i=1}^N \sum_{\mu=1}^4 (\rho_{\mu}^i - \rho_{\mu}/N)^2. \quad (5.79)$$

As all the components are orthogonal it is easy to see that this statistic will be exactly χ^2 distributed with $4(N - 1)$ degrees of freedom. One can interpret this as the sum of the single detector χ^2 values for the h_0 and $h_{\pi/2}$ waveforms in the synthetic + and \times detectors. Figure 5.6 shows the distribution of the “standard” χ^2 , using 16 frequency bins. The distribution matches the expected χ^2 with 60 degrees of freedom.

An alternative approach to applying the χ^2 test to a coherent search was proposed in [109]. This approach involves calculating the χ^2 values for each of the active detectors and using these values to veto glitches.

⁸Strictly speaking the frequency bins for the F_+ and F_{\times} components will be different because, as we have noted in equation (5.36), the PSDs for the synthetic + and \times detectors are not equal. One can avoid this problem by using separate frequency bins for the + and \times detectors, however this doubles the number of filters. Alternatively, in [80] a method was presented for calculating the standard χ^2 test using unequal frequency bins, that method could easily be incorporated into a coherent search.

5.4 Implementation and performance of a coherent search

In this section, we describe an implementation of the targeted, coherent search for gravitational waves from CBCs. The search calculates the coherent SNR described in 5.1 and the various strategies for discriminating between signal and noise in non-Gaussian, non-stationary data introduced in sections 5.2 and 5.3. We demonstrate the efficacy of the search by performing test analyses of simulated data and data from LIGO’s S4 run.

5.4.1 Implementation of a coherent triggered search for compact binary coalescences

Here, we describe the main steps by which the algorithms described in sections 5.1, 5.2 and 5.3 have been implemented. The analysis is available in the LIGO Scientific Collaboration Applications Library (LAL) suite [130], and makes use of a large number of tools and methods previously implemented in that library.

Analysis setup

A targeted, coincident search for gravitational waves from CBCs associated to GRBs has been implemented, and used in a search of S5 and VSR1 data [4]. The coherent search pipeline uses many of the same definitions, and much of the same architecture as the coincidence search pipeline to determine the analysis details. Specifically, “on-source” time is $[-5,+1)$ seconds around the reported time of the GRB; this is when a gravitational wave signal would be expected [31, 131] and is the time over which we perform our search. The noise background is estimated using 1,944 seconds of “off-source” data split into 324 trials of 6 second length each. These are used to calculate the significance of any event occurring in the on-source. As in Ref. [4] we impose a 48s “buffer zone” between the on-source and off-source regions, which is not used in the analysis. To obtain an accurate estimate of the detectors’ power spectra, we only analyse data from a detector if it has taken at least 2190s of continuous data around the time of the GRB. Modulo this restriction, the coherent analysis is designed to make use of data from all detectors that were operating at the time of the GRB.

Template bank generation

The problem of placing a non-spinning template bank in (m_1, m_2) space for a single detector has been extensively studied [86, 85, 84, 83]. However, less thought has been given to the problem of placing an appropriate bank for a coherent analysis, targeted or otherwise. For the targeted, coherent SNR statistic, described in section 5.1, the maximization procedure ensures that we need only place templates in (m_1, m_2) space. In the single detector search, the parameter space metric is independent of the amplitude and phase parameters that are maximized over. However, for the coherent analysis, this is generically *not* the case, and the metric depends upon the \mathcal{A}^μ parameters [120].

At present, we have not implemented an optimal template bank for a coherent analysis. Instead, we simply make use of an over-dense template bank generated for one of the detectors in the network. In the results presented later, we have used a bank generated with the initial LIGO design spectrum, with a maximum total mass of $40M_\odot$ and a minimum component mass of $1M_\odot$, which are the same values as used in previous searches for GRBs [4]. While this method of template placement enables us to demonstrate the efficiency of the coherent search, it is clearly not the optimal solution. A simple improvement would involve placing a template bank appropriate for the (maximally sensitive) synthetic + detector defined in equation (5.36). For the relatively common situation where the network is significantly more sensitive to the + polarization than the \times , this will be close to the optimal solution. However, for network and source configurations which provide good sensitivity to both polarizations, a more detailed investigation along the lines of [120] is warranted.

Coherent SNR and null streams

The data are read in and conditioned using the methods and algorithms developed for the S4 search for post-merger ringdowns from CBCs [132, 133]. The data are then downsampled to a frequency of 4096 Hz and split into overlapping 256 second segments for analysis. The noise PSDs are calculated using the same method as in [132].

Each template in the bank is filtered against the data from each detector to generate the single detector filters ($s^X | h_{0, \frac{\pi}{2}}^X$) and sensitivities, σ^X (defined in equation (5.27)). The algorithms used are taken from the LAL FindChirp library [57], specifically those written to perform a search for spinning waveforms [127] using

the physical template family (PTF) waveforms [134].⁹ The waveform templates are generated using the TaylorT4 post-Newtonian approximant [121]. The single detector filter outputs are shifted in time to account for the relative delays from the given GRB sky location. They are then combined to form the coherent and the null SNRs as described by (5.44) and (5.46). A “trigger” is recorded at any time the coherent SNR is greater than 6, and no louder event occurred in any template in the bank within 0.1 seconds.

Calculating the χ^2 tests

The analysis calculates signal based vetoes in the same manner that it does the coherent SNR: The necessary single detector filters are constructed and then these are combined together to create the χ^2 tests as described in section 5.3.

Calculating the “standard” χ^2 test is computationally expensive. Therefore this veto is only calculated for a segment if there is at least one event with SNR above threshold and values of bank χ^2 , auto χ^2 and null SNR that do not immediately lead to it being dismissed as a glitch.

Event significance

After the analysis, we have a set of triggers with associated masses and coalescence times. Each trigger is characterized by the coherent SNR as well as a number of other quantities: the null SNR, single detector SNRs and values of the χ^2 tests. In order to identify candidate gravitational wave events, these triggers are compared to those obtained from an identical analysis performed over the off-source times, and over the off-source data with simulated gravitational wave signals added. A simulation is deemed to have been recovered by the analysis if there is any trigger within 0.1 seconds of the signal time; no attempt is made to guarantee a good match between simulated and recovered parameters. Although this means that signals may be found due to a nearby glitch, in practice we find that the effect is minimal, particularly when considering detection candidates, which are louder than all background.

By comparing the observed triggers from the off-source and simulations, we construct a detection statistic designed to best separate signal from noise. At present, this is done by performing a number of simple cuts in various parameters, as described in detail in the remainder of the section. In the future, we plan to investigate the use of multi-dimensional classifiers to improve the efficacy of

⁹This choice stems from the desire to extend this search to incorporate a single spin. This is particularly appropriate for NSBH binaries where the spin of the NS can be safely neglected.

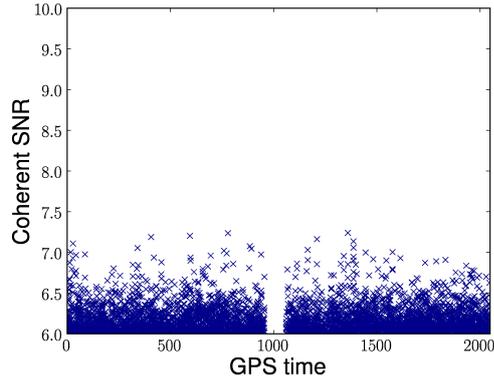


Figure 5.7: The SNR of triggers in the off-source region plotted against time for an analysis of simulated Gaussian noise in the initial LIGO (H1, H2, L1) network. The gap in the middle of the plot contains the on-source region and the buffer, this is not used in estimating noise background rates. The loudest trigger occurs with an SNR of 7.24.

the analysis. The detection statistic is used to calculate a false alarm probability for a given on-source observation, by comparing with the off-source results. We use three hundred background trials, allowing for a false alarm probability as low as 3×10^{-3} to be assigned to an event. This suffices to identify GRBs for which an interesting gravitational wave trigger has been observed. Realistically a false alarm probability closer to 10^{-4} or 10^{-5} would be required for a detection candidate. This could be achieved by performing additional background trials on time-shifted data from the detectors. This has not yet been implemented.

5.4.2 Analysis of simulated data

The analysis was first run on simulated data for the initial LIGO network, comprising a 4km detector at the Livingston site (L1) and 4km and 2km detectors at the Hanford site (H1 and H2 respectively). The coherent analysis pipeline analyzed a 2048 second stretch of stationary, Gaussian data as if a GRB had occurred in the middle of the data stretch.

Figure 5.7 shows the triggers produced by the pipeline in the off-source time. The loudest event in the approximately 2000 seconds of off-source data has an SNR of 7.24. In Gaussian noise, the signal consistency tests have no power as they are designed specifically to reduce the effect of non-Gaussian transients in the data.

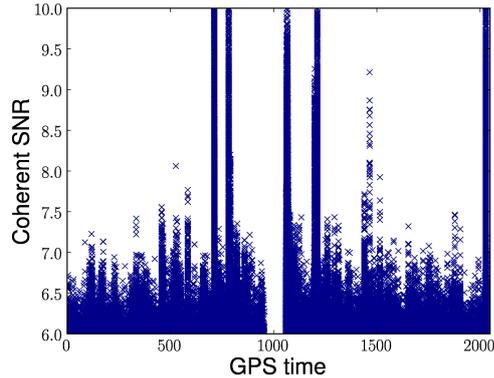


Figure 5.8: The SNR of triggers in the off-source region plotted against time for an analysis of a mock S4 GRB. The axes on the plot are chosen to be identical to those for Figure 5.7 to make the plots easier to compare. The S4 data has a large number of non-Gaussian features. The largest of these peaks extends to a coherent SNR of 40, although non-Gaussian structure is visible at SNRs as low as 7.

5.4.3 Analysis of real data

A test analysis was also performed on real data taken from S4. We analyzed an arbitrary stretch of 2048 seconds of data for which all three of the LIGO detectors were operating and ran the analysis as if a GRB had occurred during this time. The simulated sky location of the GRB was $(184.6^\circ, 42.34^\circ)$ in right ascension and declination respectively. For this chosen time and sky location the sensitivity of the H1 and L1 detectors were roughly equal and the H2 detector was half as sensitive as the other two.

The data from the detectors is neither Gaussian nor stationary. Thus, the goal of the analysis is to reduce the background from non-stationary data using the signal consistency tests described in sections 5.2 and 5.3. We proceed by investigating the various signal consistency tests one by one before combining these cuts into a detection statistic.

The final goal is to obtain a search sensitivity as close as possible to that obtained in Gaussian data. The sensitivity is assessed by evaluating the efficiency of observing simulated signals in the data. We make use of a set of simulated BNS signals (component masses limited to be between 1 and $3M_\odot$) all oriented face on to the detectors. We evaluate the efficiency of the search as the fraction of simulated signals observed with a detection statistic greater than any event in the off-source data.

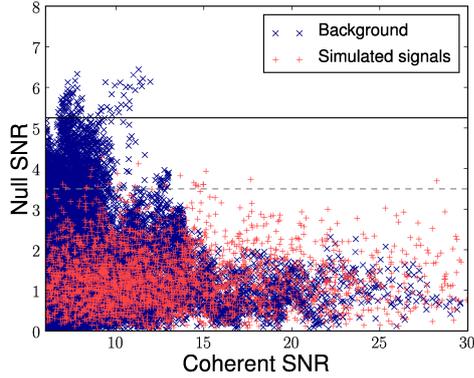


Figure 5.9: The distribution of the null SNR plotted against coherent SNR. The solid line at null SNR of 5.25 is the line above which triggers are vetoed. The dashed line at 3.5 is the line above which triggers are downweighted (see section 5.4.3).

Coherent SNR

Figure 5.8 shows the coherent SNR of triggers produced during the analysis of the S4 data. It clearly demonstrates that the data are not well characterized by Gaussian noise alone. A number of loud transients are present in the data and show up as short duration peaks of large SNR. The largest of these has an SNR of almost 40. If events were ranked by coherent SNR, a signal would have to be very loud to stand out above this non-Gaussian background. In addition to the loud peaks there are also a large number of smaller non-Gaussian peaks that occur rather frequently in the data. All of these affect the sensitivity of the search.

Null SNR

Figure 5.9 shows the performance of the null stream for both simulated signals and background noise. The ability of the null SNR to distinguish signal from noise is relatively poor in this example. The mock GRB analysis uses data from the two Hanford detectors and the detector at Livingston. As the two Hanford detectors are aligned, the null stream is derived from a combination of these detectors; the Livingston detector does not contribute. The loudest glitches during the time of this analysis originated in L1, and therefore do not contribute to the null SNR. However, quieter glitches in the Hanford detectors at an SNR around 10 do produce a large null SNR. Any trigger with a null SNR greater than 5.25 is eliminated from the analysis. In this example analysis, this removes a small fraction of the background and none of the simulated signals.

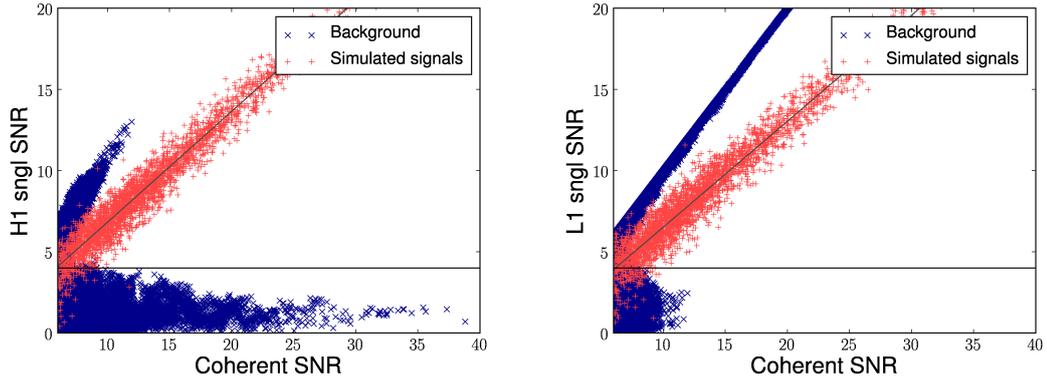


Figure 5.10: The distribution of single detector SNR for the more sensitive H1 and L1 detectors, plotted against coherent SNR. The left figure shows the H1 SNR, the right figure shows the L1 SNR. The horizontal line indicates SNR=4. Below this line triggers will be vetoed. The inclined dark gray line indicates the expected SNR of these face on simulated signals.

Single detector SNR

The most straightforward, and most effective, amplitude consistency test we have found is the requirement of an SNR greater than 4 in the two most sensitive detectors; in this analysis, the L1 and H1 detectors. Figure 5.10 demonstrates that this is a particularly effective strategy for removing noise glitches. Triggers arising due to glitches in the L1 detector have large coherent SNR but a negligible contribution from H1 and are consequently discarded. The single detector SNR threshold is very effective at removing background triggers — particularly those associated to glitches in L1. There is some loss of simulated signals, but these generally have a small enough SNR as to be indistinguishable from the background, even in Gaussian noise.

χ^2 tests

In section 5.3 we introduced three χ^2 tests designed to separate signals from noise glitches in the data. Figure 5.11 shows the distribution of the bank χ^2 for every time sample for a single template. This is directly comparable to Figure 5.3 which shows the same for Gaussian data. The deviation from the predicted χ^2 distribution is due to the non-Gaussianity of the data.

The distribution of bank and auto χ^2 for both simulated signals and noise triggers is shown in Figure 5.12. Both of these tests are effective at separating the simulated signals from noise transients. In order to quantify this, we make use of the *newSNR* formalism that was developed for the latest coincident searches for

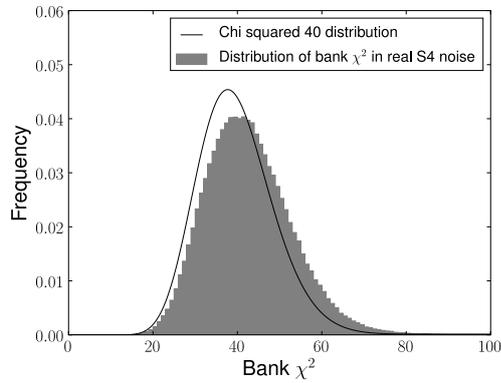


Figure 5.11: The distribution of the bank χ^2 test for a single template h , with a bank of size 10. The plot shows the distribution of the bank veto calculated for every time sample in 128s of data. The observed distribution is inconsistent with the expected result in Gaussian noise (the black curve).

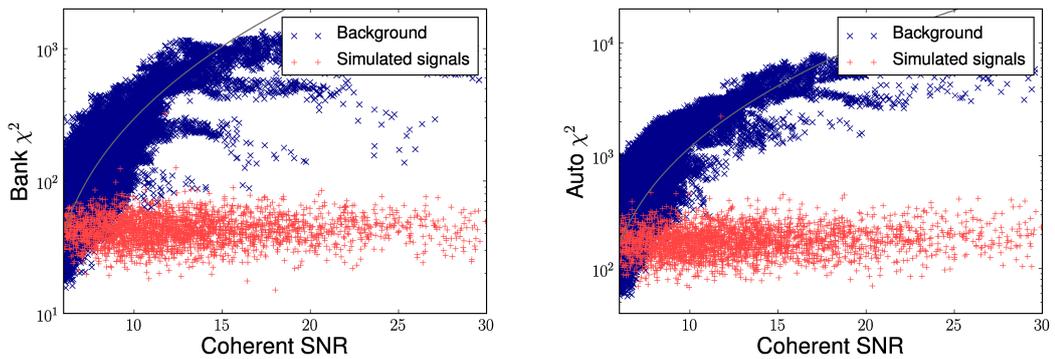


Figure 5.12: The distribution of bank and auto χ^2 test plotted against SNR. There is a clear separation between simulated signal and background at coherent SNRs above 10. The solid line shows the line of $newSNR = 6$. Triggers with $newSNR < 6$ (above and to the left of the line) are discarded.

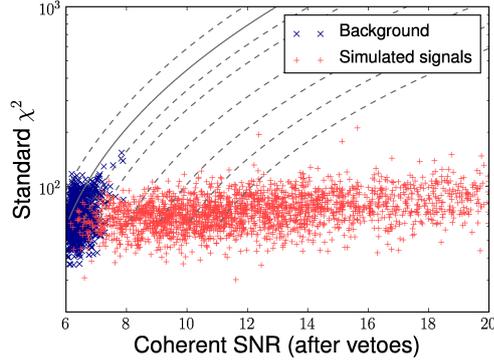


Figure 5.13: The distribution of the standard χ^2 test plotted against SNR. The majority of the background has already been removed by utilizing other signal consistency tests. The dashed lines show contours of *newSNR*, with the value of *newSNR* increasing towards the bottom right of the plot. The solid line shows the line of *newSNR* = 6; triggers above this line are discarded.

CBCs [98, 135]. For a signal matching the template waveform, the expected value of χ^2 is one per degree of freedom, while for noise transients this will be larger. The idea is to down-weight the significance of noise triggers with large χ^2 values relative to signals. This is achieved by introducing the “*newSNR*”:

$$\rho_{\text{new}} = \begin{cases} \rho, & \chi^2 \leq n_{\text{dof}} \\ \frac{\rho}{\left[\left(1 + \left(\frac{\chi^2}{n_{\text{dof}}} \right)^{4/3} \right) / 2 \right]^{1/4}}, & \chi^2 > n_{\text{dof}} \end{cases} \quad (5.80)$$

where n_{dof} is the number of degrees of freedom of the χ^2 test. For signals, *newSNR* will be similar to the SNR, while noise transients with a large χ^2 value are significantly down-weighted.

The *newSNR* can be calculated with all of the χ^2 tests introduced in section 5.3. Any trigger with *either* an auto or bank *newSNR* less than 6 is discarded. The curves on Figure 5.12 show this *newSNR* threshold for the two χ^2 tests. Finally, we turn to the standard χ^2 test. As this is rather costly to compute, we only do so for triggers which have passed all of the previously described thresholds (on coherent, null and single detector SNR, bank and auto *newSNR*). Figure 5.13 shows the distribution of the standard χ^2 test for simulated signals and noise. The preceding tests have succeeded in removing the vast majority of non-Gaussian triggers from the data. A threshold of 6 on *newSNR* serves to eliminate a few more.

We have found that the standard χ^2 is the most effective of our χ^2 tests at

separating signal from background, so we also make use of it in the final ranking of events. Figure 5.13, shows contours of constant $newSNR$ which will be used in the final ranking.

Detection statistic

In the preceding discussion, we have imposed a number of cuts on the initial candidate events produced by the analysis pipeline. Let us briefly recap those cuts:

1. Coherent SNR: Generate a trigger at any time for which $\rho > 6$. Only keep the loudest trigger in each 0.1 seconds.
2. Null SNR: Discard any triggers with $\rho_N > 5.25$.
3. Single detector SNR: Discard any triggers for which $\rho_{H1} < 4$ or $\rho_{L1} < 4$.
4. χ^2 tests: Discard any triggers for which $\rho_{new} < 6$ for the bank or auto χ^2 .

Finally, we rank the remaining triggers based upon the $newSNR$ calculated using the standard χ^2 and the null SNR:

5. Detection statistic: Rank remaining triggers using a detection statistic ρ_{det} given by

$$\rho_{det} = \begin{cases} \rho_{new}, & \rho_N \leq 3.5 \\ \frac{\rho_{new}}{\rho_N - 2.5}, & 3.5 < \rho_N < 5.25. \end{cases} \quad (5.81)$$

The length of time a CBC spends in the sensitive band of the detector varies greatly with the mass, and it has been found that the shorter, high mass templates are more susceptible to occurring with large SNR at the time of glitches [1]. Also, the various signal consistency tests are less effective for these short templates. Therefore, we follow Ref. [1] and split the template bank into three regions based on the chirp mass of the template, and calculate the false alarm probability for each mass bin separately:

6. False alarm probability: For each trigger, calculate the false alarm probability by counting the fraction of off-source trials with a louder trigger in the same mass bin.

It is this false alarm probability which allows us to assess the significance of any events in the on-source data.

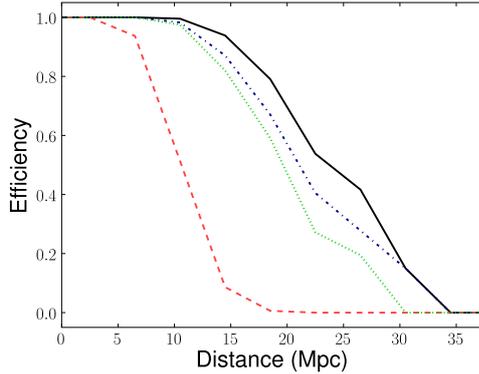


Figure 5.14: Efficiency of injection recovery is shown against distance at a false alarm probability of less than one in 324. This is calculated from a set of simulated signals in real data. The efficiency is shown for four different cases: i) signals found above an SNR of 7.24, the loudest background trigger in Gaussian noise (black solid line); ii) signals found above an SNR of 12.88, the loudest background trigger in real noise (red dashed line); iii) signals found with a value of the detection statistic (given by equation 5.81) above 7.41, the loudest background event (blue dot-dashed line); iv) signals found louder than all background in an H1-L1 search (green dotted line).

Performance of search

To quantitatively assess the sensitivity of the analysis we evaluate the efficiency of recovering simulated signals at a fixed false alarm probability. Any simulated signal which is found louder than the loudest off-source event is considered as found by the analysis — this ensures that the false alarm probability is less than one in 324 (the number of off-source trials). We use a population of optimally oriented BNS signals from the location of the fake GRB. Since all BNS injections are recovered with triggers in the low mass bin (with chirp mass less than $3.5M_{\odot}$), we compare to off-source triggers in that bin.

Figure 5.14 shows the efficiency with which simulated signals are recovered as a function of distance for a variety of search methods. We begin by considering the sensitivity of a search in Gaussian noise; this will provide a benchmark against which to compare the searches in real data. The largest SNR recorded in the analysis of simulated, Gaussian data was 7.24 (see Figure 5.7). Thus, the efficiency of a search in Gaussian data is given by the fraction of simulations found with an SNR greater than this. We have seen that the real detector data is *not* Gaussian, so only if the signal consistency tests are perfectly able to separate signals from noise transients can we hope to achieve a similar sensitivity in real data.

If we were to use *only* coherent SNR to rank events, and ignore all signal

consistency tests, the sensitivity of the search is about a factor of two worse than in Gaussian data. This is expected as the loudest low mass trigger had a coherent SNR of 12.88, almost double the value observed in Gaussian noise. Once all of the signal consistency tests are taken into account, the loudest off-source event has a detection statistic of 7.41. Any simulated signal which produces a trigger which passes all signal consistency cuts and has $\rho_{\text{det}} > 7.41$ is considered found. The use of signal consistency tests and a new detection statistic greatly improves the sensitivity of the search. The distance at which 50% efficiency is achieved is only about 10% less than in Gaussian noise. The 10% loss in sensitivity can be attributed to a slightly louder off-source event (7.41 rather than 7.24) and a *small* number of the simulated signals being vetoed by the signal consistency cuts.

Finally, we would like to illustrate the benefits of a coherent search over a coincidence search. A coincidence search filters each detector independently and records single detector triggers before searching for coincidence between the triggers in different detectors. For the initial LIGO network, signals close to detection threshold would be unlikely to be seen in the less sensitive H2 detector. Specifically, for a signal with a coherent SNR of 7.5, the expected SNR in H2 would be around 2.5 which would be insufficient to generate a trigger. Therefore the sensitivity of a coincidence search would be limited by events observable in the two detector H1-L1 network. For a two detector search, the coherent and coincident SNRs are equal, and the null stream test is not available. Consequently, the performance of two detector coherent and coincident searches should be comparable. The efficiency of this two detector search is also shown in Figure 5.14; it is about 10% less sensitive than the three detector coherent search. This demonstrates that the coherent analysis, which incorporates the H2 data, can increase the distance reach of the search by 10%. For a network of three approximately equally sensitive detectors, we would expect an even greater sensitivity improvement from employing coherent techniques.

5.5 Discussion

We have presented a formulation of a targeted coherent search for compact binary coalescences. For Gaussian noise, the coherent SNR would be ideal for separating signals from the noise background. However, since data from gravitational wave interferometers is neither Gaussian nor stationary, we have also discussed a number of methods of separating the non-stationary noise background from the signal population. These tests include various χ^2 tests, which were originally designed

for use in single detectors. We have extended them to the network analysis and demonstrated their continued efficacy. Additionally, the coherent analysis naturally allows for tests which are not readily available in the coincidence case. The most significant of these is the null SNR which can be used to reject events which are not consistent with two gravitational wave polarizations. We also explored consistency tests between the recovered amplitudes of the gravitational wave and found that a simple SNR threshold on the two most sensitive detectors gave excellent results. There are various other glitch rejection techniques, which have been recently discussed in the literature [136, 137, 138], it might be possible to utilize such methods to improve the coherent search described here.

The analysis described in this chapter has been implemented and in the final section we showed results of a test run. This made use of the S4 data from the LIGO detectors. Although the data was far from Gaussian, after the application of all of the signal consistency tests, the results were remarkably close to what would be expected in Gaussian noise. This analysis is available to be used in searches for GW inspiral signals associated with GRBs in more recent LIGO and Virgo data, such as S6 and VSR2 and VSR3.

There are a number of ways in which this analysis could be enhanced to broaden its use and increase its sensitivity. First, a number of GRBs, particularly those observed by Fermi [139] and IPN [140] are not localized sufficiently accurately that the error box can be treated as a point on the sky. Thus, it would be nice to extend this analysis to cover a region of the sky. This would require looping over the relevant sky points; incorporating the correct detector sensitivities $F_{+,x}$ and time delays. In principle, this would not greatly slow down the analysis as the majority of time is taken in performing the single detector filters and these would *not* need to be re-calculated. As well as looking at a patch on the sky, the analysis could be extended to cover the whole sky, as appropriate for an un-triggered search. This brings in a host of new complications which have been met and dealt with by other coherent search methods [111, 110]. In order to obtain a good estimate of the background for an all sky, un-triggered search we would need to implement background estimation and time shifting the data would likely be the best way to do this.

The \mathcal{F} -statistic technique described in section 5.1.2 is formulated by maximizing over the extrinsic parameters of the system. From a Bayesian perspective this would imply that we have placed a uniform prior on the distributions of the \mathcal{A}^μ . If instead we were to place an astrophysical prior on these amplitudes based on the expected distributions of (D, ι, ψ, ϕ_0) we would expect to increase the efficiency

of the search [124]. Alternatively, since GRBs are thought to be rather tightly beamed, it is reasonable to take them as being face on, or close to. In this case, the gravitational waves are circularly polarized and there is, in effect, only a single polarization. This opens the possibility of limiting the signal space to just this one polarization and adding an extra “null” test.

The progenitors of short GRBs are thought to be BNS or NSBH. The search we have described is ideal for the BNS case as the spins of the neutron stars are unlikely to have a significant effect on the waveform. However, when one of the components of the binary is a black hole, the spin could be large. Furthermore, the mass ratio is likely to be relatively large. In this case, the spin of the black hole can have a significant effect on the observed waveform [134]. Consequently, we would like to extend this search to incorporate spin effects. We discuss this implementation in the next chapter.

Chapter 6

Searching for spinning compact binary coalescences

In chapters 4 and 5 we have described coincidence and coherent methods for detecting CBC signals in data taken by gravitational wave interferometers. These searches have utilized template waveforms where the spins of the individual components are neglected. In some areas of the parameter space, spin can have a significant effect on the evolution of the system, and consequently the emitted gravitational waveform [141, 142], leading to a poor match with the non-spinning templates. In these regions of parameter space the use of templates incorporating spin will provide an increase in search sensitivity.

Incorporating spin into template waveforms in a gravitational wave search is a complex problem. In a search for CBCs with non spinning components in circular orbits, a source is described by nine physical parameters [143]. The majority of these do not affect the signal morphology, but serve to change the overall amplitude, phase or coalescence time of the signal and are easily maximized over, as we demonstrated in chapter 4. Therefore, template placement can be restricted to the two dimensional space of component masses [84]. A CBC with spinning components in a circular orbit, however, is described by 15 physical parameters [143]. The challenge is to formulate a method to detect any manner of spinning system while limiting the number of templates, such that an analysis can be run in a reasonable amount of time.

The problem is simplified if the spins are aligned with the orbital angular momentum. In this case the system will have no precession and is described by just two extra parameters — the spins of both bodies in the direction of the angular momentum. Furthermore, the dominant spin term in the PN expansion is well described by a single spin parameter [144], and it is therefore feasible to

search for non-precessing waveforms using a three dimensional template bank.

At the time of writing only one detection search for CBCs using spinning templates with precession in LIGO/Virgo data has been published [145]. This search utilized a phenomenological waveform family designed to capture precessional effects [146]. It was, however, later abandoned because it was not found to increase efficiency relative to the non-spinning search [1, 126]. This was due to the ability of the phenomenological templates to match non-stationarities in the data and the lack of an effective signal consistency test to veto them such as the χ^2 test used in the non spinning search [80].

The PTF waveforms proposed in [146] and further explored in [134, 147, 127] give a different method for searching for spinning binaries with precession. This method uses *single-spin* precessing waveforms as templates. Making clever use of maximization, it was shown [134] that a PTF search could be performed with a four dimensional template bank: the two masses, the magnitude of the spin and the angle between the spin and the orbital plane. This method is especially useful for detecting neutron star–black hole binary (NSBH) systems, where the spin of the neutron star would have a negligible effect on the dynamics of the system [134]. A coincidence search utilizing the PTF waveforms has been developed [127]. However, while coincidence for non-spinning searches is well defined as we discussed in section 4.2, it is less clear how to define coincidence when additional spin parameters are present.

Alternatively the PTF search could be conducted using the coherent technique, as described in chapter 5. This would remove the need to define coincidence. The coherent technique is especially useful when the sky position is known, such as when searching for gravitational waves in coincidence with an electromagnetic transient, such as a GRB [28]. Since NSBH and BNS mergers are the preferred progenitor model [30] for short GRB, a coherent single-spin search is ideally suited to this source.

In this chapter we describe the implementation of a coherent search for single spin binaries with known sky location, using the PTF waveforms. We briefly review the PTF formalism before deriving the coherent PTF SNR. Due to the increased complexity of the spinning waveforms, the coherent SNR has a different distribution than its non-spinning counterpart. In particular, there is a greater chance of obtaining a large value of the spinning SNR, even in Gaussian noise. Thus there is a trade-off between the improved spinning signal model and the increased false alarm rate at a fixed SNR. We explore the single-spin CBC parameter space to identify regions where spin (and precession) effects are significant

enough to make the spinning search worthwhile. We will also briefly discuss some possibilities for vetoing background non-Gaussian transients in the data when using the PTF search and present results of this search run on a short stretch of S4 data.

The layout of this chapter is as follows: In section 6.1 we explore how the evolution of a CBC is effected by the presence of spin and how this will change the gravitational wave signal that we hope to detect. In section 6.2 we discuss how non-precessing templates could be used in a gravitational wave search. In section 6.3 we formulate the PTF waveform before describing the single detector PTF search in section 6.4 and investigating the distribution of the spinning SNR in Gaussian noise. In section 6.5 we introduce the coherent PTF search and investigate the distribution of the coherent spinning SNR. In section 6.6 we identify regions of the parameter space where the PTF search offers increased sensitivity over the non-precessing search. Section 6.7 briefly describes our search pipeline and the results of these methods applied to a stretch of data from S4. Finally, in section 6.8 we briefly discuss how these techniques could be used to detect generic spin systems.

6.1 Spinning compact binary coalescence gravitational waveforms

To be able to discuss different strategies for searching for CBC systems incorporating spin, we must first understand how spinning CBC waveforms will differ from the non-spinning ones. We will begin by discussing how the inspiral part of the waveform is effected by spin, something that is well explored in the literature [141, 148, 149, 146]. Then we will describe how recent advances in numerical relativity have enabled the production of spinning signals incorporating inspiral, merger and ringdown components.

The presence of spin in a compact binary inspiral effects the evolution of the inspiral in two different ways. Firstly, spin will modify the frequency evolution of the inspiralling system. Secondly, if the spins of the components are not (anti-) aligned with the orbital angular momentum then the system will precess. This precession will further complicate the frequency evolution of the system and will cause the orientation of the system to be time dependent.

The frequency evolution of an inspiral up to 3.5PN order with spin effects up

to 2PN order is given in Eq (1) of [146] as

$$\begin{aligned}
\frac{\dot{\omega}}{\omega^2} = & \frac{96}{5} \eta (M\omega)^{5/3} \left\{ 1 - \frac{743 + 924 \eta}{336} (M\omega)^{2/3} \right. \\
& - \left. \left(\frac{1}{12} \sum_{i=1,2} \left[\chi_i (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_i) \left(113 \frac{m_i^2}{M^2} + 75\eta \right) \right] - 4\pi \right) (M\omega) \right. \\
& + \left. \left(\frac{34\,103}{18\,144} + \frac{13\,661}{2\,016} \eta + \frac{59}{18} \eta^2 \right) (M\omega)^{4/3} \right. \\
& - \frac{1}{48} \eta \chi_1 \chi_2 \left[247 (\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{S}}_2) - 721 (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_1) (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_2) \right] (M\omega)^{4/3} \\
& - \frac{1}{672} (4\,159 + 15\,876 \eta) \pi (M\omega)^{5/3} \\
& + \left[\left(\frac{16\,447\,322\,263}{139\,708\,800} - \frac{1\,712}{105} \gamma_E + \frac{16}{3} \pi^2 \right) \right. \\
& + \left. \left(-\frac{273\,811\,877}{1\,088\,640} + \frac{451}{48} \pi^2 - \frac{88\,1039}{3\,4620} \right) \eta \right. \\
& + \left. \frac{541}{896} \eta^2 - \frac{5\,605}{2\,592} \eta^3 - \frac{856}{105} \log [16(M\omega)^{2/3}] \right] (M\omega)^2 \\
& + \left. \left(-\frac{4\,415}{4\,032} + \frac{358\,675}{6\,048} \eta + \frac{91\,495}{1\,512} \eta^2 \right) \pi (M\omega)^{7/3} \right\}, \tag{6.1}
\end{aligned}$$

where γ_E is Euler's constant. The presence of spin causes an additional term to be present at 1.5PN and another at 2PN order. The 1.5PN term is the dominant spin-orbit coupling term and is given by [146]

$$\frac{1}{12} \sum_{i=1,2} \left[\chi_i (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_i) \left(113 \frac{m_i^2}{M^2} + 75\eta \right) \right]. \tag{6.2}$$

We remind the reader that we are still using the definitions given in section 3.1 to describe a CBC system. Specifically $\hat{\mathbf{S}}_i$ gives the spin unit vector for both components, χ_i is the magnitude of that spin ($0 \leq \chi \leq 1$). $\hat{\mathbf{L}}_N$ gives the unit vector describing direction of the orbital angular momentum. The 2PN term is the dominant spin-spin coupling term and is given by

$$\frac{1}{48} \eta \chi_1 \chi_2 \left[247 (\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{S}}_2) - 721 (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_1) (\hat{\mathbf{L}}_N \cdot \hat{\mathbf{S}}_2) \right]. \tag{6.3}$$

There are additional, higher order spin-related corrections to the frequency evolution equation [150, 151] these higher order terms should also be used when creating waveforms incorporating spin for gravitational wave data analysis.

The evolution of the spin directions due to precession is given in [141, 149, 146] to dominant order as

$$\dot{\mathbf{S}}_{1,2} = \frac{(M\omega)^2}{2M} \left\{ \eta(M\omega)^{-1/3} \left(4 + 3\frac{m_2}{m_1} \right) \hat{\mathbf{L}}_N + \frac{1}{M^2} \left[\mathbf{S}_{2,1} - 3 \left(\mathbf{S}_{2,1} \cdot \hat{\mathbf{L}}_N \right) \hat{\mathbf{L}}_N \right] \right\} \times \mathbf{S}_{1,2}. \quad (6.4)$$

The evolution of the orbital angular momentum is given as [146]

$$\dot{\hat{\mathbf{L}}}_N = \frac{\omega^2}{2M} \left[\mathbf{S}_{\text{eff}} \times \hat{\mathbf{L}}_N - \frac{3\omega^{1/3}}{\eta M^{5/3}} \left(\left(\mathbf{S}_2 \cdot \hat{\mathbf{L}}_N \right) \mathbf{S}_1 + \left(\mathbf{S}_1 \cdot \hat{\mathbf{L}}_N \right) \mathbf{S}_2 \right) \times \hat{\mathbf{L}}_N \right], \quad (6.5)$$

where

$$\mathbf{S}_{\text{eff}} \equiv \left(1 + \frac{3m_2}{4m_1} \right) \mathbf{S}_1 + \left(1 + \frac{3m_1}{4m_2} \right) \mathbf{S}_2. \quad (6.6)$$

\mathbf{S}_{eff} can be thought of as the effective spin of the system. These equations can then be used to carry out a time-evolution to model the dynamics of the inspiral. These can be converted into the $h(t)$ response at a gravitational wave detector in the same manner as described in chapter 3 but one must remember that the orientation of the system will be time dependent due to precession of the orbital angular momentum.

It is worth noting at this point that a special case exists when \mathbf{S}_1 , \mathbf{S}_2 and $\hat{\mathbf{L}}_N$ are all aligned or anti-aligned. In this case these vectors will be constant and the system will experience no precession, though the phase evolution is still effected by the presence of spin. We call this special case the ‘‘aligned spin’’ case or non-precessing case.

The Post-Newtonian expansion can only model the inspiral part of a spinning waveform. Numerical relativity must be used to produce waveforms that contain the merger phase and to connect the inspiral, merger and ringdown phases together. The presence of spin and precession presents additional complications to modelling such systems numerically. However, a lot of work in the numerical relativity community has been focused on generating spinning waveforms and a number of numerical relativity waveforms incorporating spin have been produced. A review of numerical relativity results and a table of currently available waveforms is given in [70, 71].

6.2 Non precessing systems

The task of performing gravitational wave data analysis using generic spinning waveforms as templates is a complicated one. It is therefore useful to consider

simplified cases first. In this section we will describe how gravitational wave searches can be performed using waveforms that have spin, but no precessional effects.

In equation (6.5) we demonstrated how the orbital angular momentum will precess in the presence of spin. In the case that both spins are (anti-)aligned with the orbital angular momentum, there will be no precession. With no precessional effects, modelling the waveform is easier than the generic case, as the spin terms are non-variant. One simply has to add the constant spin terms to the PN expansion. Waveform families also exist to produce such aligned-spin systems with a merger and ringdown component, both for the phenomenological and effective one body models [144, 152].

From the perspective of data analysis, performing a search with non-precessing template waveforms is easier than with generic spin templates because $h(t)$ can still be expressed in the form derived for non-spinning templates in equation (3.29)

$$h(\tau) = A(\mathcal{M}, D, \iota, \theta, \psi, \phi) (f(\mathcal{M}, \eta, \chi_1, \chi_2, \tau))^{2/3} \cos(\Phi(\mathcal{M}, \eta, \chi_1, \chi_2, \tau) + \Phi_0(\iota, \varphi, \theta, \psi, \phi)). \quad (6.7)$$

This is not true of a generic signal incorporating spin effects where A is time dependent. The only difference between here and the non-spinning case is that the frequency evolution now depends on two additional parameters, the magnitudes of both spins.

The data analysis strategies that we described in chapters 4 and 5 to search for CBCs using non-spinning templates can therefore be applied directly. Thus, the non-spinning search method could accurately be described as a non-precessing search method. The only complications that will arise from this are that the template bank must be placed to cover the two additional spin parameters and, if a coincidence search is used, the definition of coincidence must also be extended for these extra parameters. In [144] the authors show that these two spin parameters can be described by only one “effective spin” parameter, simplifying both of these problems. Thus, the precessing search is just the non-spinning search with templates covering the two masses *and* the effective spin.

The problem of template bank placement in 3 dimensions could be solved by using the stochastic placement algorithm, which is presented in chapter 7. This method would, however, require a parameter space metric to be calculated on the non-precessing parameter space. To define coincidence, it would be possible in theory to extend the definition of coincidence ellipsoids in [90] to 4 dimensions. Remember that time is also used in defining coincidence. Again, however, the

parameter space metric is required to do this. Alternatively the problem of coincidence could be avoided by using the coherent search method described in chapter 5.

At the time of writing this thesis, no gravitational wave search using non-precessing templates has been run on data from LIGO and Virgo. I hope to conduct such a search in the near future. A search using non-precessing templates would not only be useful in detecting spin-aligned systems it could be useful in detecting precessing systems as well. We showed in section 4.2.10 that the non-spinning search was capable of detecting spinning systems in large portions of the full spinning parameter space. The non-precessing search would only extend our detection reach in this parameter space. It would be informative to carry out a study to determine how the non-precessing search could improve detection efficiency as a function of the spin and mass of the system.

6.3 The single spin PTF waveform

To be able to detect CBC systems that have a significant amount of precession, it is useful to search with templates that incorporate precessional effects. A spinning system is described by 15 physical parameters. An efficient detection search over this whole parameter space has not yet been devised. However, one can consider systems where only one of the bodies has spin. This reduces our problem to a 12 dimensional parameter space, which, as we will show in the rest of this section, can be solved. Searching with single spin templates can be physically motivated. The effect of the spin of the neutron star in neutron star–black hole binary (NSBH) systems will be considerably less than the effect of the spin of the black hole. This can be seen from the equations in section 6.1. Therefore single spin templates should be able to detect any manner of physical NSBH coalescence.

The physical template family (PTF) waveforms proposed in [146] and further explored in [134, 147, 127] give a method for searching for spinning binaries with precession using single-spin precessing waveforms as templates. Through a maximization technique, which we will explore fully in section 6.4.1, we can conduct a search, using the PTF waveforms as templates, where we only need to place templates to cover 4 of the parameters. Those parameters are

- The two masses, (M_1, M_2) ,
- The magnitude of the spin, χ ,

- The cosine of the angle between the spin direction and the initial orbital angular momentum, κ .

The PTF waveforms are constructed by demanding that $\mathbf{S}_2 = 0$.

We will begin by describing the form of the PTF waveforms. In the non-spinning search we always defined the source frame such that the z direction is aligned with the orbital angular momentum. For precessing systems the orbital angular momentum is time dependant. Therefore we can either use a source frame that is time dependant where the z direction tracks the orbital angular momentum. Alternatively we could define a stationary source frame that does not generically have the z direction aligned with the orbital angular momentum.

To define the PTF waveform we will make use of both time dependant and stationary source frames to express the waveform in the stationary frame. We follow the work of [134] to define the coordinate system and calculate the form of the signal.

Let us begin by giving the definition of the stationary source frame. This coordinate system is set using the following conditions:

- The z axis is chosen such that it lies in the direction of the orbital angular momentum at some initial time $\hat{\mathbf{L}}_N(0)$.
- The x axis is chosen such that the initial direction of spin $\mathbf{S}_1(0)$ will lie in the x, z plane.

where the initial time corresponds to the start of the template waveform.

To calculate the gravitational waveform we need to understand how this frame is related to the precessing source frame. To do this we need to be careful to distinguish between the basis vectors in both frames. We use the following convention¹

$$\mathbf{e}_{x,y,z}^P = \text{Basis vectors of precessing source frame} \quad (6.8a)$$

$$\mathbf{e}_{x,y,z}^S = \text{Basis vectors of stationary source frame} \quad (6.8b)$$

$$\mathbf{e}_{x,y,z}^R = \text{Basis vectors of radiation frame.} \quad (6.8c)$$

At the initial time, the coordinates of both source frames will be identical. From this time the precessing source frame evolves in the stationary frame according to [134]

$$\left(\dot{\mathbf{e}}_{x,y,z}^P\right)^S(t) = \boldsymbol{\Omega}_e(t) \times \left(\mathbf{e}_{x,y,z}^P\right)^S(t), \quad (6.9)$$

¹ The stationary radiation frame, which we will use later, is defined in the same way as in chapter 4 for the non spinning search. e_z^R always points toward the Earth.

$$\boldsymbol{\Omega}_e(t) = \frac{\omega^2}{2M} \left(4 + 3\frac{m_2}{m_1}\right) \mathbf{S}_1 - \left(\frac{\omega^2}{2M} \left(4 + 3\frac{m_2}{m_1}\right) \mathbf{S}_1 \cdot \hat{\mathbf{L}}_N\right) \hat{\mathbf{L}}_N. \quad (6.10)$$

The radiation frame is then defined from the stationary source frame in the same manner as for the non-spinning search using the rotation angles (ι, φ) , see section 2.1.7. The initial phase of the system, Φ_0 is not used when defining the source coordinate system, so needs to be included in the phase evolution of the system.

Now we can construct how a ground-based gravitational wave interferometer would respond to a single spin gravitational waveform. Explicitly the response is given by [134]

$$h(t) = -\frac{2}{D} \mathcal{M}^{5/3} \omega^{2/3} ([\mathbf{e}_+]^{ij} \cos 2(\Phi + \Phi_0) + [\mathbf{e}_\times]^{ij} \sin 2(\Phi + \Phi_0)) ([\mathbf{T}_+]_{ij} F_+ + [\mathbf{T}_\times]_{ij} F_\times), \quad (6.11)$$

where $\mathbf{T}_{+,\times}$ is the detector polarization tensor and $\mathbf{e}_{+,\times}$ is the gravitational wave polarization tensor

$$\mathbf{T}_+ \equiv \mathbf{e}_x^R \otimes \mathbf{e}_x^R - \mathbf{e}_y^R \otimes \mathbf{e}_y^R \quad \mathbf{T}_\times \equiv \mathbf{e}_x^R \otimes \mathbf{e}_y^R + \mathbf{e}_y^R \otimes \mathbf{e}_x^R \quad (6.12a)$$

$$\mathbf{e}_+ \equiv \mathbf{e}_x^P \otimes \mathbf{e}_x^P - \mathbf{e}_y^P \otimes \mathbf{e}_y^P \quad \mathbf{e}_\times \equiv \mathbf{e}_x^P \otimes \mathbf{e}_y^P + \mathbf{e}_y^P \otimes \mathbf{e}_x^P. \quad (6.12b)$$

Clearly $[\mathbf{e}_{+,\times}]^{ij}$ is most easily expressed in the precessing source frame and $[\mathbf{T}_{+,\times}]_{ij}$ in the radiation frame. However, it is convenient to evaluate $[\mathbf{e}_{+,\times}]^{ij} [\mathbf{T}_{+,\times}]_{ij}$ in the stationary source frame

$$[\mathbf{T}_+]_{ij}^S = \begin{pmatrix} \sin^2 \varphi - \cos^2 \Theta \cos^2 \varphi & -(\cos^2 \Theta + 1) \sin \varphi \cos \varphi & \cos \Theta \sin \Theta \cos \varphi \\ -(\cos^2 \Theta + 1) \sin \varphi \cos \varphi & \cos^2 \varphi - \cos^2 \Theta \sin^2 \varphi & \cos \Theta \sin \Theta \sin \varphi \\ \cos \Theta \sin \Theta \cos \varphi & \cos \Theta \sin \Theta \sin \varphi & -\sin^2 \Theta \end{pmatrix} \quad (6.13a)$$

$$[\mathbf{T}_\times]_{ij}^S = \begin{pmatrix} 2 \cos \Theta \sin \varphi \cos \varphi & \cos \Theta (\sin^2 \varphi - \cos^2 \varphi) & -\sin \Theta \sin \varphi \\ \cos \Theta (\sin^2 \varphi - \cos^2 \varphi) & -2 \cos \Theta \sin \varphi \cos \varphi & \sin \Theta \cos \varphi \\ -\sin \Theta \sin \varphi & \sin \Theta \cos \varphi & 0 \end{pmatrix} \quad (6.13b)$$

$$[\mathbf{e}_+]_{ij}^S = \begin{pmatrix} ((\mathbf{e}_x^P)^S_x)^2 - ((\mathbf{e}_y^P)^S_x)^2 & (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_y - (\mathbf{e}_y^P)^S_x (\mathbf{e}_x^P)^S_y & (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_z - (\mathbf{e}_y^P)^S_x (\mathbf{e}_x^P)^S_z \\ (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_y - (\mathbf{e}_y^P)^S_x (\mathbf{e}_x^P)^S_y & ((\mathbf{e}_x^P)^S_y)^2 - ((\mathbf{e}_y^P)^S_y)^2 & (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_z - (\mathbf{e}_y^P)^S_y (\mathbf{e}_x^P)^S_z \\ (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_z - (\mathbf{e}_y^P)^S_x (\mathbf{e}_x^P)^S_z & (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_z - (\mathbf{e}_y^P)^S_y (\mathbf{e}_x^P)^S_z & ((\mathbf{e}_x^P)^S_z)^2 - ((\mathbf{e}_y^P)^S_z)^2 \end{pmatrix} \quad (6.13c)$$

$$[\mathbf{e}_\times]_{ij}^S = \begin{pmatrix} 2(\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_x & (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_y + (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_x & (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_z + (\mathbf{e}_x^P)^S_z (\mathbf{e}_y^P)^S_x \\ (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_x + (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_y & 2(\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_y & (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_z + (\mathbf{e}_x^P)^S_z (\mathbf{e}_y^P)^S_y \\ (\mathbf{e}_x^P)^S_z (\mathbf{e}_y^P)^S_x + (\mathbf{e}_x^P)^S_x (\mathbf{e}_y^P)^S_z & (\mathbf{e}_x^P)^S_z (\mathbf{e}_y^P)^S_y + (\mathbf{e}_x^P)^S_y (\mathbf{e}_y^P)^S_z & 2(\mathbf{e}_x^P)^S_z (\mathbf{e}_y^P)^S_z \end{pmatrix}. \quad (6.13d)$$

This is similar to the definition of $h(t)$ for the non-spinning waveforms, however

here \mathbf{e}_+ and \mathbf{e}_\times will be time dependent with no components that are explicitly 0. Therefore the equation cannot be simplified as it was for the non-spinning case. However we note that both $\mathbf{e}_{+,\times}$ and $\mathbf{T}_{+,\times}$ are transverse and traceless. Any transverse, traceless (TT) 3x3 matrix can be expressed in terms of 5 base matrices defined from the $2m$ spin-weighted spherical harmonics

$$A_{ij} = A^I (M^I)_{ij} \quad (6.14)$$

where the M^I satisfy

$$(M^I)_{ij} (M^J)^{ij} = \delta^{IJ}. \quad (6.15)$$

For the PTF definition, the authors of [134] found it convenient to define M^I as a combination of the spin weighted spherical harmonics given by

$$(M^1)_{ij} = \left(\frac{4\pi}{15}\right)^{1/2} (\mathcal{Y}_{ij}^{22} + \mathcal{Y}_{ij}^{2-2}) = \begin{bmatrix} (2^{-1/2}) & 0 & 0 \\ 0 & -(2^{-1/2}) & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (6.16)$$

$$(M^2)_{ij} = -i \left(\frac{4\pi}{15}\right)^{1/2} (\mathcal{Y}_{ij}^{22} - \mathcal{Y}_{ij}^{2-2}) = \begin{bmatrix} 0 & (2^{-1/2}) & 0 \\ (2^{-1/2}) & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (6.17)$$

$$(M^3)_{ij} = -\left(\frac{4\pi}{15}\right)^{1/2} (\mathcal{Y}_{ij}^{21} - \mathcal{Y}_{ij}^{2-1}) = \begin{bmatrix} 0 & 0 & (2^{-1/2}) \\ 0 & 0 & 0 \\ (2^{-1/2}) & 0 & 0 \end{bmatrix} \quad (6.18)$$

$$(M^4)_{ij} = i \left(\frac{4\pi}{15}\right)^{1/2} (\mathcal{Y}_{ij}^{21} + \mathcal{Y}_{ij}^{2-1}) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & (2^{-1/2}) \\ 0 & (2^{-1/2}) & 0 \end{bmatrix} \quad (6.19)$$

$$(M^5)_{ij} = -\left(\frac{8\pi}{15}\right)^{1/2} \mathcal{Y}_{ij}^{20} = \begin{bmatrix} 6^{-1/2} & 0 & 0 \\ 0 & 6^{-1/2} & 0 \\ 0 & 0 & -2(6^{-1/2}) \end{bmatrix} \quad (6.20)$$

where \mathcal{Y}_{ij}^{lm} give the spin-weighted spherical harmonics according to

$$\mathcal{Y}_{ij}^{22} = \frac{1}{4} \left(\frac{15}{2\pi}\right)^{1/2} \begin{bmatrix} 1 & i & 0 \\ i & -1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (6.21)$$

$$\mathcal{Y}_{ij}^{2-2} = \frac{1}{4} \left(\frac{15}{2\pi} \right)^{1/2} \begin{bmatrix} 1 & -i & 0 \\ -i & -1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (6.22)$$

$$\mathcal{Y}_{ij}^{21} = \frac{1}{4} \left(\frac{15}{2\pi} \right)^{1/2} \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & -i \\ -1 & -i & 0 \end{bmatrix} \quad (6.23)$$

$$\mathcal{Y}_{ij}^{2-1} = \frac{1}{4} \left(\frac{15}{2\pi} \right)^{1/2} \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -i \\ 1 & -i & 0 \end{bmatrix} \quad (6.24)$$

$$\mathcal{Y}_{ij}^{20} = \frac{1}{4} \left(\frac{5}{4\pi} \right)^{1/2} \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 2 \end{bmatrix}. \quad (6.25)$$

Therefore, it is possible to rewrite $h(t)$ as

$$h(t) = \sum_{I=1}^5 P_I(D, \theta, \phi, \psi, \iota, \varphi) Q^I(M_1, M_2, \chi, \kappa, \Phi_0, t_c), \quad (6.26)$$

where

$$Q^I = \mathcal{M}^{5/3} \omega^{2/3} \sqrt{2} ([\mathbf{e}_+]^{ij} \cos 2(\Phi + \Phi_0) + [\mathbf{e}_\times]^{ij} \sin 2(\Phi + \Phi_0)) (M^I)_{ij} \quad (6.27a)$$

$$P_I = -\frac{\sqrt{2}}{D} ([\mathbf{T}_+]_{ij} F_+ + [\mathbf{T}_\times]_{ij} F_\times) (M_I)^{ij}. \quad (6.27b)$$

In this form $h(t)$ is written as the sum of 5 components. Each component consists of an amplitude term, P_I , and a time dependent waveform, Q^I . As we will demonstrate in the next section it will be possible to maximize over the P_I components when performing the matched-filter, thus reducing the number of parameters that

must be templated over. The P_I amplitude terms are given explicitly by

$$\begin{aligned}
P_1 &= -\frac{1}{D} \left([\mathbf{T}_+]_{11} F_+ - [\mathbf{T}_+]_{22} F_+ + [\mathbf{T}_\times]_{11} F_\times - [\mathbf{T}_\times]_{22} F_\times \right) \\
&= -\frac{2}{D} \left[\frac{-1}{4} (3 + \cos(2\Theta)) \cos(2\varphi) \right] F_+ + -\frac{2}{D} \left[\cos \Theta \sin(2\varphi) \right] F_\times \quad (6.28a)
\end{aligned}$$

$$\begin{aligned}
P_2 &= -\frac{2}{D} \left([\mathbf{T}_+]_{12} F_+ + [\mathbf{T}_\times]_{12} F_\times \right) \\
&= -\frac{2}{D} \left[\frac{-1}{4} (3 + \cos(2\Theta)) \sin(2\varphi) \right] F_+ - \frac{2}{D} \left[\cos \Theta \cos(2\varphi) \right] F_\times \quad (6.28b)
\end{aligned}$$

$$\begin{aligned}
P_3 &= -\frac{2}{D} \left([\mathbf{T}_+]_{13} F_+ + [\mathbf{T}_\times]_{13} F_\times \right) \\
&= -\frac{2}{D} \left[\frac{1}{2} \sin(2\Theta) \cos(\varphi) \right] F_+ - \frac{2}{D} \left[-\sin(\Theta) \sin(\varphi) \right] F_\times \quad (6.28c)
\end{aligned}$$

$$\begin{aligned}
P_4 &= -\frac{2}{D} \left([\mathbf{T}_+]_{23} F_+ + [\mathbf{T}_\times]_{23} F_\times \right) \\
&= -\frac{2}{D} \left[\frac{1}{2} \sin(2\Theta) \sin(\varphi) \right] F_+ - \frac{2}{D} \left[\sin(\Theta) \cos(\varphi) \right] F_\times \quad (6.28d)
\end{aligned}$$

$$\begin{aligned}
P_5 &= -\frac{2}{\sqrt{12}D} \left(([\mathbf{T}_+]_{11} + [\mathbf{T}_+]_{22} - 2[\mathbf{T}_+]_{33}) F_+ + ([\mathbf{T}_+]_{11} + [\mathbf{T}_+]_{22} - 2[\mathbf{T}_+]_{33}) F_\times \right) \\
&= \frac{-\sqrt{3}}{D} \left([\mathbf{T}_+]_{33} F_+ + [\mathbf{T}_\times]_{33} F_\times \right) \\
&= -\frac{2}{D} \left[\frac{\sqrt{3}}{4} (1 - \cos(2\Theta)) \right] F_+. \quad (6.28e)
\end{aligned}$$

Note that the trace free condition is used in the formulation of P_5 . The Q^I terms

describe the time dependent waveforms and are given by

$$\begin{aligned}
Q^1 &= \mathcal{M}^{5/3} \omega^{2/3} \left(\cos 2(\Phi + \Phi_0) \left([\mathbf{e}_+]^{11} - [\mathbf{e}_+]^{22} \right) + \sin 2(\Phi + \Phi_0) \left([\mathbf{e}_\times]^{11} - [\mathbf{e}_\times]^{22} \right) \right) \\
&= \mathcal{M}^{5/3} \omega^{2/3} \left[\cos(2\Phi_0) \left([\mathbf{e}_+]^{11} - [\mathbf{e}_+]^{22} \right) + \sin(2\Phi_0) \left([\mathbf{e}_\times]^{11} - [\mathbf{e}_\times]^{22} \right) \right] \cos(2\Phi) \\
&\quad + \mathcal{M}^{5/3} \omega^{2/3} \left[-\sin(2\Phi_0) \left([\mathbf{e}_+]^{11} - [\mathbf{e}_+]^{22} \right) + \cos(2\Phi_0) \left([\mathbf{e}_\times]^{11} - [\mathbf{e}_\times]^{22} \right) \right] \sin(2\Phi)
\end{aligned} \tag{6.29a}$$

$$\begin{aligned}
Q^2 &= 2\mathcal{M}^{5/3} \omega^{2/3} \left(\cos 2(\Phi + \Phi_0) [\mathbf{e}_+]^{12} + \sin 2(\Phi + \Phi_0) [\mathbf{e}_\times]^{12} \right) \\
&= 2\mathcal{M}^{5/3} \omega^{2/3} \left[\cos(2\Phi_0) [\mathbf{e}_+]^{12} + \sin(2\Phi_0) [\mathbf{e}_\times]^{12} \right] \cos(2\Phi) \\
&\quad + 2\mathcal{M}^{5/3} \omega^{2/3} \left[-\sin(2\Phi_0) [\mathbf{e}_+]^{12} + \cos(2\Phi_0) [\mathbf{e}_\times]^{12} \right] \sin(2\Phi)
\end{aligned} \tag{6.29b}$$

$$\begin{aligned}
Q^3 &= 2\mathcal{M}^{5/3} \omega^{2/3} \left(\cos 2(\Phi + \Phi_0) [\mathbf{e}_+]^{13} + \sin 2(\Phi + \Phi_0) [\mathbf{e}_\times]^{13} \right) \\
&= 2\mathcal{M}^{5/3} \omega^{2/3} \left[\cos(2\Phi_0) [\mathbf{e}_+]^{13} + \sin(2\Phi_0) [\mathbf{e}_\times]^{13} \right] \cos(2\Phi) \\
&\quad + 2\mathcal{M}^{5/3} \omega^{2/3} \left[-\sin(2\Phi_0) [\mathbf{e}_+]^{13} + \cos(2\Phi_0) [\mathbf{e}_\times]^{13} \right] \sin(2\Phi)
\end{aligned} \tag{6.29c}$$

$$\begin{aligned}
Q^4 &= 2\mathcal{M}^{5/3} \omega^{2/3} \left(\cos 2(\Phi + \Phi_0) [\mathbf{e}_+]^{23} + \sin 2(\Phi + \Phi_0) [\mathbf{e}_\times]^{23} \right) \\
&= 2\mathcal{M}^{5/3} \omega^{2/3} \left[\cos(2\Phi_0) [\mathbf{e}_+]^{23} + \sin(2\Phi_0) [\mathbf{e}_\times]^{23} \right] \cos(2\Phi) \\
&\quad + 2\mathcal{M}^{5/3} \omega^{2/3} \left[-\sin(2\Phi_0) [\mathbf{e}_+]^{23} + \cos(2\Phi_0) [\mathbf{e}_\times]^{23} \right] \sin(2\Phi)
\end{aligned} \tag{6.29d}$$

$$\begin{aligned}
Q^5 &= -\sqrt{3} \mathcal{M}^{5/3} \omega^{2/3} \left(\cos 2(\Phi + \Phi_0) [\mathbf{e}_+]^{33} + \sin 2(\Phi + \Phi_0) [\mathbf{e}_\times]^{33} \right) \\
&= -\sqrt{3} \mathcal{M}^{5/3} \omega^{2/3} \left[\cos(2\Phi_0) [\mathbf{e}_+]^{33} + \sin(2\Phi_0) [\mathbf{e}_\times]^{33} \right] \cos(2\Phi) \\
&\quad - \sqrt{3} \mathcal{M}^{5/3} \omega^{2/3} \left[-\sin(2\Phi_0) [\mathbf{e}_+]^{33} + \cos(2\Phi_0) [\mathbf{e}_\times]^{33} \right] \sin(2\Phi).
\end{aligned} \tag{6.29e}$$

Note that we have written these Q^I in terms of $\cos(2\Phi_0)$ and $\sin(2\Phi_0)$. This will be necessary when performing a maximization over Φ_0 in the next section.

6.4 Spinning search using PTF waveforms

In this section we describe the single detector PTF search and its implementation, which was first derived in [134, 127]. We will follow the conventions of these earlier publications as much as possible.

6.4.1 Formulating the PTF detection statistic

Here we will derive the PTF detection statistic. We note again that this derivation has been done previously in [134, 127]. For completeness of this chapter we will describe this derivation in detail here.

As we derived in section 4.1.2 the likelihood ratio of there being a signal h present in the data s for a single detector is given by

$$\Lambda(h) = \frac{P(s|h)}{P(s|0)}. \quad (6.30)$$

Assuming the noise is Gaussian, the log likelihood can be written as

$$\log \Lambda = (s|h) - \frac{1}{2}(h|h). \quad (6.31)$$

Where again we are using the single detector inner product

$$(a|b) = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}(f)[\tilde{b}(f)]^*}{S_h(f)}. \quad (6.32)$$

$S_h(f)$ is the PSD of the detector, as described in section 4.1.1. From this starting point, h must be re-expressed in such a way that it is possible to maximize over the majority of the parameters, leaving us with only a small number of dimensions over which to carry out a templated search.

In the previous section we demonstrated that the gravitational waveform for a single spin binary can be expressed as

$$h(t) = \sum_{I=1}^5 P_I(D, \theta, \phi, \psi, \iota, \varphi) Q^I(M_1, M_2, \chi, \kappa, \Phi_0, t_c). \quad (6.33)$$

where P_I are five amplitudes and Q^I describe five waveform components. If we insert this into the log likelihood we obtain

$$\ln \Lambda = (P_I Q^I, s) - \frac{1}{2}(P_I Q^I, P_J Q^J). \quad (6.34)$$

We now wish to maximize over as many parameters in this PTF waveform model as possible, to reduce the number of dimensions that we need to perform a templated search over. We can initially carry out a maximization over phase and overall amplitude using a similar method as in the non-spinning search. Firstly we express the Q_I components in terms of its 0 and $\frac{\pi}{2}$ phases, as is illustrated in equation (6.29),

$$Q^I = Q_0^I \cos 2\Phi_0 + Q_{\frac{\pi}{2}}^I \sin 2\Phi_0. \quad (6.35)$$

If we then make the standard assumption that $Q_0^I = iQ_{\frac{\pi}{2}}^I$ we can express the log

likelihood equation as

$$\ln \Lambda = P_I(A^I \cos 2\Phi_0 + B^I \sin 2\Phi_0) - \frac{1}{2} P_I P_J M^{IJ}. \quad (6.36)$$

Where we have defined

$$A^I = (s, Q_0^I) \quad B^I = (s, Q_{\pi/2}^I) \quad (6.37)$$

as the inner products of the 0 and $\frac{\pi}{2}$ components of the waveforms with the data.

$$M^{IJ} = (Q_0^I, Q_0^J) = (Q_{\pi/2}^I, Q_{\pi/2}^J) \quad (6.38)$$

is the inner product of the waveform components with themselves. We can then pull out an overall common amplitude A from the P_I terms and reexpress the log likelihood as

$$\ln \Lambda = A \hat{P}_I (A^I \cos 2\Phi_0 + B^I \sin 2\Phi_0) - \frac{A^2}{2} \hat{P}_I \hat{P}_J M^{IJ}. \quad (6.39)$$

Now, we can maximize over the phase and the distance by defining

$$\gamma = A \cos 2\Phi_0 \quad (6.40)$$

$$\beta = A \sin 2\Phi_0. \quad (6.41)$$

The log likelihood expressed in terms of γ and β is

$$\ln \Lambda = \hat{P}_I (A^I \gamma + B^I \beta) - \frac{1}{2} (\gamma^2 + \beta^2) \hat{P}_I \hat{P}_J M^{IJ}. \quad (6.42)$$

We then maximize $\ln \Lambda$ over γ and β . This will give us

$$\gamma_{\max} = \frac{\hat{P}_I A^I}{\hat{P}_J \hat{P}_K M^{JK}} \quad (6.43)$$

$$\beta_{\max} = \frac{\hat{P}_I B^I}{\hat{P}_J \hat{P}_K M^{JK}}. \quad (6.44)$$

Thus our detection statistic maximized over distance and phase is given by

$$\ln \Lambda|_{\max(D, \Phi_0)} = \frac{1}{2} \frac{\hat{P}_I \hat{P}_J (A^I A^J + B^I B^J)}{\hat{P}_I \hat{P}_J M^{IJ}}. \quad (6.45)$$

We now wish to maximize over the P_I amplitude terms, before we do this we can

simplify the form of the log likelihood. We can perform a transformation such that both Q_0^I and $Q_{\frac{\pi}{2}}^I$ are orthonormal. First, perform a rotation on the Q_0^I to make M^{IJ} diagonal, then normalize the basis vectors. We denote the orthonormal basis \tilde{Q}_0^I . This transformation will also orthonormalize $\tilde{Q}_{\frac{\pi}{2}}^I$ and render \tilde{M}^{IJ} the identity matrix. We note that this orthonormalization is not carried out in [134, 127] but will simplify the final form of the PTF SNR.

Furthermore, we are free to rescale the P_I terms in both the numerator and denominator without changing the log likelihood. Therefore, we are free to scale the P_I terms such that $\tilde{P}_I \tilde{P}_J \delta^{IJ} = 1$. The log likelihood is then simplified to

$$\ln \Lambda|_{\max(\mathcal{D}, \Phi_0)} = \frac{1}{2} \tilde{P}_I \tilde{P}_J (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J). \quad (6.46)$$

We can then maximize this with the constraint that $\tilde{P}_I \tilde{P}_J \delta^{IJ} = 1$ using the method of Lagrangian multipliers as described in [134, 127]

$$\ln \Lambda = \frac{1}{2} \tilde{P}_I \tilde{P}_J (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J) - \lambda (\tilde{P}_I \tilde{P}_J \delta^{IJ} - 1). \quad (6.47)$$

Differentiating this with respect to \tilde{P}_I and setting to 0 gives

$$\frac{1}{2} (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J) \tilde{P}_J = \lambda \delta^{IJ} \tilde{P}_J. \quad (6.48)$$

This equation is only solved when λ is an eigenvalue of $\frac{1}{2} (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J)$ and \tilde{P}_I is the corresponding eigenvector. Further, if we multiply (6.48) by \tilde{P}_J we can show that

$$\lambda = \frac{1}{2} \tilde{P}_I \tilde{P}_J (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J). \quad (6.49)$$

Thus, the maximum log likelihood is given by the maximum eigenvalue of $\frac{1}{2} (\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J)$. To calculate the value of this maximum eigenvalue it is useful to study the form of $(\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J)$. Firstly we note that any vector that is orthogonal to both A^I and B^I , will also be orthogonal to $(\tilde{A}^I \tilde{A}^J + \tilde{B}^I \tilde{B}^J)$. Therefore there is a 3-dimensional sub-space of vectors that is orthogonal to this matrix. Equivalently, this matrix has a three dimensional *kernel* and thus can only have 2 non zero eigenvalues. These eigenvalues must lie in the plane described by \tilde{A}^I and \tilde{B}^I . We can therefore write the eigenvectors corresponding to the non-zero eigenvalues as

$$\tilde{P}^I = \alpha \tilde{A}^I + \beta \tilde{B}^I. \quad (6.50)$$

If we then insert equation (6.50) into (6.48) we can form two simultaneous equa-

tions

$$\alpha\lambda = \frac{\alpha}{2}\tilde{A}^I\tilde{A}_I + \frac{\beta}{2}\tilde{A}^I\tilde{B}_I \quad (6.51a)$$

$$\beta\lambda = \frac{\alpha}{2}\tilde{A}^I\tilde{B}_I + \frac{\beta}{2}\tilde{B}^I\tilde{B}_I, \quad (6.51b)$$

which can be solved to find the values of λ . Doing this gives

$$\lambda = \frac{1}{2} \left[\tilde{A} \cdot \tilde{A} + \tilde{B} \cdot \tilde{B} \pm \sqrt{(\tilde{A} \cdot \tilde{A} - \tilde{B} \cdot \tilde{B})^2 + 4(\tilde{A} \cdot \tilde{B})^2} \right]. \quad (6.52)$$

The maximum eigenvalue is obtained by taking the positive form of the square root. Thus, the maximized PTF detection statistic is given by

$$\frac{\rho^2}{2} = \log \Lambda|_{\max(D, \Phi_0, P_I)} = \frac{1}{2} (\tilde{A} \cdot \tilde{A} + \tilde{B} \cdot \tilde{B}) + \frac{1}{2} \sqrt{(\tilde{A} \cdot \tilde{A} - \tilde{B} \cdot \tilde{B})^2 + (2\tilde{A} \cdot \tilde{B})^2}. \quad (6.53)$$

6.4.2 Physical freedom in the PTF SNR

We have performed a free maximization over the five P_I amplitudes. In principle, these depend upon *six* physical parameters. However, these parameters only enter in four different combinations as

- an amplitude parameter, dependent on (D, θ, ψ, ϕ)
- the relative sensitivity of the instrument to the $+$ and \times polarizations, dependent on (θ, ψ, ϕ)
- the inclination angle, ι
- the spin orientation, φ .

Therefore performing a free maximization over the five P_I components means that the maximized P_I values may not correspond to a physical set of parameters. This is discussed in [134] and various methods for projecting onto the physical sub-space have been proposed. For the case of an externally triggered search, as we describe in section 4.3, where the sky location is known, the situation is unchanged as the P_I are still described by the same four unknown parameters.

Additionally, in the case of a coincidence search, these parameters are freely maximized over *for each detector*. This can mean that, even if the values are projected onto the physical sub-space it is possible to have rather different maximized values of the parameters across the different detectors.

6.4.3 PTF with no precession

When the orbital plane of the system does not precess, there is gravitational wave emission in only two of the harmonics, Q^1 and Q^2 . The other components vanish identically. Furthermore, these two harmonics are related by a phase shift: $Q^1 = iQ^2$. Thus, the matrix M is degenerate and the PTF maximization breaks down. It is, however, straightforward to maximize over the two remaining amplitudes, and obtain the SNR as

$$\rho^2 = \frac{(s|Q_0^1|^2 + (s|Q_{\frac{\pi}{2}}^1|^2)}{(Q_0^1|Q_0^1)}. \quad (6.54)$$

This is identical to the well known SNR for the non-spinning search [57], and the two phases of Q^1 correspond to the 0 and $\pi/2$ phases of the non-precessing template.

6.4.4 SNR distribution in Gaussian noise

The PTF template waveform will provide a better match than a non-spinning template to a gravitational wave signal from a spinning binary. However, we pay a price since we must filter the data against more waveform components, Q^I , thereby increasing the chance of a spurious match with the noise. Additionally, the spinning SNR takes a more complex form (6.53) than the simple quadratic expression (6.54) when there is no precession. Here, we will investigate the SNR distributions in Gaussian noise for these two cases. In section 6.6 we use this to identify regions of parameter space with sufficient spin effects to warrant the use of the PTF search.

Ten filters are used in the calculation of the PTF detection statistic: (\tilde{Q}_0^I, s) and $(\tilde{Q}_{\frac{\pi}{2}}^I, s)$. As both \tilde{Q}_0^I and $\tilde{Q}_{\frac{\pi}{2}}^I$ are orthonormal, the only remaining freedom is the relation between the \tilde{Q}_0^I and $\tilde{Q}_{\frac{\pi}{2}}^I$ terms,

$$\tilde{N}^{IJ} = (\tilde{Q}_0^I, \tilde{Q}_{\frac{\pi}{2}}^J). \quad (6.55)$$

This \tilde{N}^{IJ} is a 5×5 antisymmetric matrix which can therefore have 4 non-zero eigenvalues: $\pm\lambda_1, \pm\lambda_2$. The values of these eigenvalues determine the distribution of the PTF detection statistic and these values will differ for each template.

For every NSBH waveform we have tested using the initial LIGO sensitivity curve, the magnitudes of λ_1 and λ_2 have been very close to 1. Thus, although there are ten different waveform components, we find that, in effect, only six of these are independent — the others are linear combinations of these six. There

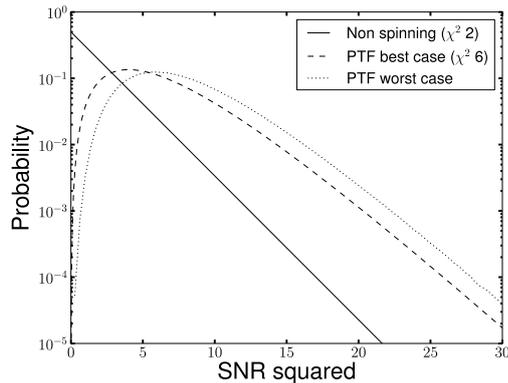


Figure 6.1: The “best” and “worst” possible distributions of the single detector PTF SNR squared, this is compared with the non precessing SNR squared.

are then only six independent filters. In this case it is not difficult to show that the spinning SNR (6.53) collapses to a quadratic form which is χ^2 distributed with six degrees of freedom in Gaussian noise. This is the “best” case for the detection statistic. The “worst” case occurs when both λ_1 and λ_2 are zero and all ten of the filters are independent. In this case, the SNR expression cannot be simplified and its distribution does *not* correspond to a χ^2 distribution with 10 degrees of freedom as might be expected; the real distribution is somewhat more complex. Both best and worst cases are illustrated in Figure 6.1.

The SNR (6.54) for a non-precessing template follows a χ^2 distribution with two degrees of freedom in Gaussian noise. This is also plotted on Figure 6.1. By comparing the distributions of the PTF and non-precessing SNRs, it is clear that the background triggers produced by the PTF search will have, on the average, considerably larger SNR than those produced by the non-precessing search. We explore the effect that this has on a search further in section 6.6.

6.4.5 Remaining challenges for the coincidence PTF search

The PTF search allows one to perform a coincidence search using single spin waveforms in a reasonable amount of time [127]. However, it remains a challenge to derive a metric on the four dimensional mass and spin space that could be used in generating a template bank and in defining coincidence requirements. Furthermore, a strategy for vetoing non-transient glitches, such as we described for the non-spinning search in chapter 4, would be required for running this search. Such a strategy has been suggested in [127], but more investigation is needed to show how this would increase detection confidence over a non-spinning search.

6.5 Triggered coherent PTF spinning search

In this section, we introduce a multi-detector, coherent formulation of the PTF search. As in chapter 5, this is applied in terms of a triggered search where we will restrict attention to the case where the sky location is known. This simplifies a coherent search as the sensitivity of the detectors to the two GW polarizations and the relative time delays between detectors are known. Astrophysically, this is of interest when searching for gravitational waves associated to electromagnetic transients such as GRBs [30, 31].

The data from various detectors are combined together coherently to form two coherent data streams, with one stream containing the + polarization of any gravitational wave signal present in the network and the second containing the \times polarization. The coherent method will offer an improvement in sensitivity over the coincidence method when more than two detectors are used, as only two data streams are searched. For networks with greater than two detectors, it is also possible to construct null streams which will contain no gravitational wave signal, and can be used as a consistency test as described in [112] and chapter 5.

We begin by formulating the coherent SNR for the spinning PTF search and go on to explore how this will be distributed in Gaussian noise. To formulate a coherent detection statistic for the PTF templates we draw on many of the methods and techniques that were used in deriving the single detector statistic described in section 6.4 and in [134, 127]. We follow the conventions of chapter 5 in extending this to a coherent, multi-detector search. As in chapter 5, assuming that the noise in different detectors is independent, the multi-detector likelihood is given by

$$\ln \Lambda = (\mathbf{h}|\mathbf{s}) - \frac{1}{2}(\mathbf{h}|\mathbf{h}). \quad (6.56)$$

We again use a multiple detector matched filter

$$(\mathbf{a}|\mathbf{b}) = \sum_X (a^X | b^X) \quad (6.57)$$

where the index X runs over the detectors in the network.

As before, we want to maximize over as many of the parameters as possible to minimize the dimension of the required template bank. We start by maximizing

this over the distance, D , and initial orbital phase, Φ_0 , to obtain²

$$\ln \Lambda|_{\max(D, \Phi_0)} = \frac{1}{2} \frac{\left[\sum_{X,I} (P_I^X A_I^X) \right]^2 + \left[\sum_{X,I} (P_I^X B_I^X) \right]^2}{\sum_{Y,J,K} [P_J^Y P_K^Y (Q_0^J | Q_0^K)_Y]}, \quad (6.58)$$

where P_I are the amplitudes of the various waveform components Q^I , and A_X^I, B_X^I are the matched-filters of the Q components with the data as defined in (6.37). Although the P_I depend upon an overall amplitude the maximized likelihood is independent of it as scaling the amplitude has an identical effect on both the numerator and denominator of eq. (6.58).

As in the single detector case, we would like to maximize over the P_I to eliminate them. However, these are detector dependent since the sensitivity of the detectors to the $+$ and \times gravitational wave polarizations will differ. These sensitivities are encoded in the detector response functions, F_+ and F_\times , which depend on the sky location of the source in the detector frame. As we are focusing on an externally triggered search, where the sky location is known, these values will be known for each detector. We can then factor the detector dependent terms out of the P_I^X if we remember that

$$F_+ = -a(\theta, \phi) \cos(2\psi) - b(\theta, \phi) \sin(2\psi) \quad (6.59a)$$

$$F_\times = a(\theta, \phi) \sin(2\psi) - b(\theta, \phi) \cos(2\psi), \quad (6.59b)$$

where a and b are given by

$$a(\theta, \phi) = 0.5 (1 + \cos^2 \theta) \cos(2\phi) \quad (6.60a)$$

$$b(\theta, \phi) = \cos \theta \sin(2\phi). \quad (6.60b)$$

These expressions were derived in section 2.4.2. We can then express the P_I terms as

$$P_I^X = a^X(\theta, \phi) S_I(D, \iota, \psi, \varphi_0) + b^X(\theta, \phi) T_I(D, \iota, \psi, \varphi_0) \quad (6.61)$$

where S_I and T_I can be thought of as the amplitude of the $+$ and \times components

²The Y subscript in the inner product in the denominator denotes the fact that the PSD of detector Y is used in evaluating the inner product. We do not require the noise PSDs of the different detectors to be the same.

respectively of the 5 Q_I . S_I and T_I are given as

$$S_I = -R_{I1} \cos(2\psi) + R_{I2} \sin(2\psi) \quad (6.62a)$$

$$T_I = -R_{I2} \cos(2\psi) - R_{I1} \sin(2\psi) \quad (6.62b)$$

where we have defined

$$\begin{aligned} R_{11} &= -0.25(3 + \cos(2\iota)) \cos(2\varphi) & R_{12} &= \cos(\iota) \sin(2\varphi) \\ R_{21} &= -0.25(3 + \cos(2\iota)) \sin(2\varphi) & R_{22} &= -\cos(\iota) \cos(2\varphi) \\ R_{31} &= 0.5 \sin(2\iota) \cos(\varphi) & R_{32} &= -\sin(\iota) \sin(\varphi) \\ R_{41} &= 0.5 \sin(2\iota) \sin(\varphi) & R_{42} &= \sin(\iota) \cos(\varphi) \\ R_{51} &= \frac{\sqrt{3}}{4}(1 - \cos(2\Theta)) & R_{52} &= 0. \end{aligned} \quad (6.63)$$

We can re-cast the log-likelihood into a form which more closely resembles the single detector case by introducing ten-dimensional analogues of the P_I and Q^I by defining

$$\begin{aligned} \mathcal{P}_\alpha &:= [S_1, S_2, S_3, S_4, S_5, T_1, T_2, T_3, T_4, T_5] \\ \mathcal{Q}_{0, \frac{\pi}{2}}^\alpha &:= \left[F_+ Q_{0, \frac{\pi}{2}}^1; \dots; F_+ Q_{0, \frac{\pi}{2}}^5; F_\times Q_{0, \frac{\pi}{2}}^1; \dots; F_\times Q_{0, \frac{\pi}{2}}^5 \right]. \end{aligned} \quad (6.64)$$

The change to ten dimensions naturally arises because a multiple detector coherent network is sensitive to both the $+$ and \times components, whereas a single detector network is only sensitive to one polarization. We also define the multi-detector inner products between signal and waveform components

$$\mathcal{A}^\alpha = (\mathbf{s} | \mathcal{Q}_0^\alpha) \quad \text{and} \quad \mathcal{B}^\alpha = (\mathbf{s} | \mathcal{Q}_{\frac{\pi}{2}}^\alpha) \quad (6.65a)$$

$$\mathcal{M}^{\alpha\beta} = (\mathcal{Q}_0^\alpha | \mathcal{Q}_0^\beta) = (\mathcal{Q}_{\frac{\pi}{2}}^\alpha | \mathcal{Q}_{\frac{\pi}{2}}^\beta). \quad (6.65b)$$

The log likelihood equation can then be written as

$$\ln \Lambda|_{\max(\mathcal{D}, \Phi_0)} = \frac{1}{2} \frac{\mathcal{P}_\alpha \mathcal{P}_\beta (\mathcal{A}^\alpha \mathcal{A}^\beta + \mathcal{B}^\alpha \mathcal{B}^\beta)}{\mathcal{P}_\alpha \mathcal{P}_\beta \mathcal{M}^{\alpha\beta}}. \quad (6.66)$$

We proceed, as before, by transforming to an orthonormal basis $\tilde{\mathcal{Q}}_0^\alpha, \tilde{\mathcal{Q}}_{\frac{\pi}{2}}^\alpha$ for the waveform components. Then, we can proceed in an identical manner as section

6.4.1 to maximize freely over \mathcal{P}_α , thus yielding the coherent PTF SNR

$$\rho_{\text{coh}}^2 = [\tilde{\mathcal{A}} \cdot \tilde{\mathcal{A}} + \tilde{\mathcal{B}} \cdot \tilde{\mathcal{B}}] + \sqrt{(\tilde{\mathcal{A}} \cdot \tilde{\mathcal{A}} - \tilde{\mathcal{B}} \cdot \tilde{\mathcal{B}})^2 + (2\tilde{\mathcal{A}} \cdot \tilde{\mathcal{B}})^2}, \quad (6.67)$$

where, as before, the tilde denotes that we are in the orthonormal basis.

When the network is only sensitive to one polarization, the matrix $\mathcal{M}^{\alpha\beta}$ becomes degenerate and the maximization procedure must be re-visited. Here it is natural to remove all terms corresponding to the second polarization and reduce to 5 dimensions, as in the single detection search. Additionally, in section 6.4 we noted that when the template has no precession the single detection PTF SNR collapses to the familiar SNR formalism used in the non-spinning search. Similarly, in the coherent PTF search, when the template has no precession, the coherent SNR will collapse to the non-spinning coherent SNR given in equation (5.33).

The coherent SNR of equation (6.67) can be used as a detection statistic in performing a coherent search using PTF templates, as we explore in section 6.7. In the single detector search, we maximized freely over five P_I which were dependent upon four physical parameters. Here, the \mathcal{P}_α still depend on only 4 parameters but we are now maximizing over ten amplitudes. This clearly introduces a lot of unnecessary freedom. However, we should note that the coincidence search allows for a similar freedom as the P_I are maximized independently for each detector. Consequently, for a network with three or more detectors, the coherent search provides a sensitivity improvement.

For either search it would be useful to restrict the statistic such that the maximized values take physical values. We are investigating methods of doing this. Additionally, it would be useful to use marginalization techniques to downweight physical values that are “unlikely”, such as very small distances with poor orientation. Such techniques would serve to increase the overall efficiency of a PTF search.

6.5.1 SNR distribution in Gaussian noise

In section 6.4.4 we explored how the single detector PTF statistic is distributed in Gaussian noise. For the coherent PTF search we can use a similar strategy to investigate the distribution of the coherent SNR. In the coherent case there are twenty filters $\tilde{\mathcal{A}}^\alpha$ and $\tilde{\mathcal{B}}^\alpha$ and we have constructed the detection statistic such that $\tilde{\mathcal{Q}}_0^\alpha$ and $\tilde{\mathcal{Q}}_{\frac{\pi}{2}}^\alpha$ are orthonormal. As before, the only freedom is the relationship

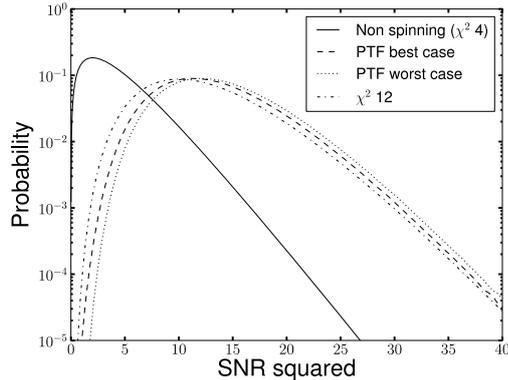


Figure 6.2: The “best” and “worst” possible distributions of the coherent PTF SNR squared as well as the distribution of the non spinning SNR squared.

between the 0 and $\frac{\pi}{2}$ terms encoded in

$$\tilde{\mathcal{N}}^{\alpha\beta} = (\tilde{\mathcal{Q}}_0^\alpha, \tilde{\mathcal{Q}}_{\frac{\pi}{2}}^\beta). \quad (6.68)$$

This is a 10x10 antisymmetric matrix comprised of four 5x5 blocks, each of which is antisymmetric. Therefore this matrix can have 8 non-zero eigenvalues: $\pm\lambda_{1,2,3,4}$. These eigenvalues determine the distribution of coherent SNR in Gaussian noise — for smaller eigenvalues, the large SNR tail of the distribution becomes more significant. In the tests that we have performed using the initial LIGO sensitivity curve and NSBH precessing templates, all four eigenvalues give values close to unity, the “best” case in which there are 12 independent waveform components. However, in this case it is not possible to separate these 12 components in the detection statistic and the distribution does not collapse to a χ^2 distribution. In Figure 6.2, we demonstrate that this gives a distribution *similar* to a χ^2 distribution with 12 degrees of freedom. In the “worst” case, where all of the eigenvalues are equal to 0, there are 20 independent waveform components and this distribution is also shown in Figure 6.2.

6.6 Identifying where the PTF search is most beneficial

In sections 6.4 and 6.5 we have derived the spinning SNR that can be used to perform a gravitational wave search using single spin inspiral waveforms as templates. We have demonstrated that, on the average, background triggers will have larger values of SNR in the PTF search than in the non-precessing search. At

	Non Precessing	PTF
Single detector	6.79	7.63
Coherent	7.26	8.53

Table 6.1: The SNR corresponding to a FAP of 10^{-10} for the non precessing and the PTF search, for both coherent and single detector cases. Here, for the PTF case the single detector and coherent detection statistics are assumed to be χ^2 distributions with 6 and 12 degrees of freedom respectively.

the same time, precessing PTF waveforms will be a better match to any spinning, precessing signals in the data. This begs the question as to whether it is preferable to use a search with non-precessing waveforms or single spin PTF waveforms to detect precessing systems. The PTF triggers will match the waveform better but this comes at the cost of searching a larger parameter space.

To quantify this, in Table 6.1 we give the SNR that corresponds to a false alarm probability (FAP) of 10^{-10} in Gaussian noise for the various searches. We chose this value because it roughly corresponds to the loudest background events we observe when running the search on 2000 seconds of Gaussian noise, as is appropriate for a GRB search. The figures in the table show that the PTF search must obtain 26% more signal power (SNR squared) to be more efficient in the single detector case at this FAP and 38% more signal power for the coherent case.

There are large areas of the parameter space where precession will not significantly effect the evolution of the binary and thus a non-precessing template will pick up the majority of the power in a precessing signal. In these areas it would be better to search for the spinning signal with a non-precessing template, achieving a lower FAP than for the PTF search using an exactly matching template. Equivalently, when a system has little precession, the majority of power is contained in the Q^1 and Q^2 components of the PTF waveform and these two components are very similar, up to an overall phase shift. We can then consider performing a “restricted PTF” search, where we filter only these two components of the waveform against the data. This serves to reduce the FAP at a fixed SNR while losing only a small amount of the power in the signal.

To do this, we test every template waveform, before filtering, to determine whether the template would be more likely to detect a matching signal below a false alarm probability of 10^{-10} using the restricted or full PTF search. This can be calculated by simulating a large number of gravitational wave signals, with masses and spin matching those of the template, but uniformly distributed in volume and orientation. Then, simply count the number of simulated signals expected to give

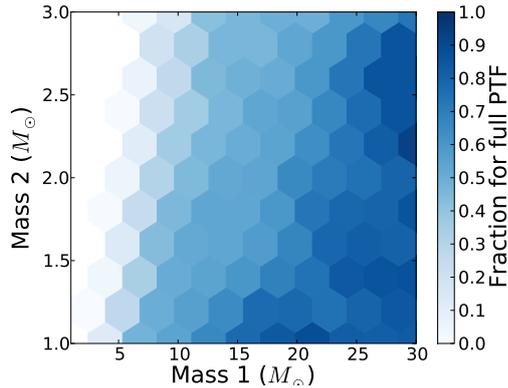


Figure 6.3: The fraction of templates analysed by the full PTF statistic as a function of the masses in the NSBH region of the parameter space.

an SNR greater than the value corresponding to a FAP of 10^{-10} (given in Table 6.1) for both methods. Whichever of the PTF or restricted methods is expected to perform better is then used when filtering the data with that template. Using this method, we are able search the full parameter space of NSBH binaries in a single search, including non-spinning, non-precessing, marginally precessing and fully precessing configurations. This method works equally well for the single detector or the coherent search.

In Figure 6.3 we illustrate the fraction of templates analysed by the full PTF statistic, as a function of the masses, for the coherent search. The splitting of the templates into full and restricted does not require filtering against the data, but it does make use of the PSDs of the detectors. For this study, we use data from the three LIGO detectors during the S4 run.

A template bank was generated by taking a standard non-spinning template bank [83] in the mass space and, for each value of the masses, creating 15 templates with identical masses but spin parameters gridded over the two dimensional spin space, as described in [127]. The precessing single spin templates are most needed in the high mass ratio region of the parameter space. For this template bank, there are 35395 templates to be analysed with the restricted method and 14660 templates to be analysed with the full PTF method.

6.7 Search method and example results

In section 6.5 we derived a detection statistic appropriate for a coherent search using the PTF waveforms as templates. In section 6.6 we described a method through which one can identify where the PTF search is most needed and to split

a template bank into those templates that should be analysed with the full PTF statistic and those that should be analysed with the restricted PTF statistic. We have combined these two methods together to create a search pipeline that can be used to coherently analyse gravitational wave data to search for precessing NSBH signals associated to short GRBs. We will briefly describe the analysis procedure before presenting an example result.

The search uses much of the same architecture as that described in chapter 5 and [4]. Namely we search for gravitational wave signals in the “on-source” time, defined to be $[-5,+1)$ seconds around the reported time of the GRB. Background is estimated from performing 324, 6 second trials around the GRB time, but at least 48s away from the on-source time. The coherent PTF search makes use of the same infrastructure as the coherent non-spinning search described in chapter 5. In particular, the data handling, PSD estimation and matched filtering routines are the same. Of course, the coherent PTF search makes use of spinning, precessing waveforms in the filtering and computes the SNR given in equation (6.67).

To demonstrate the performance of the coherent PTF, we ran it over a stretch of data from S4. The data was chosen randomly, subject to the condition that all three of the LIGO detectors were operational at the time. This is the same data as was used to illustrate the template bank splitting in section 6.6, and the same bank with 15,000 full PTF and 35,000 restricted PTF templates was used.

In Figure 6.4 we show the distribution of the SNR of the triggers produced using both the full and restricted statistic. This is shown for the stretch of real S4 data and for a stretch of simulated Gaussian data. As expected, the SNRs of triggers in Gaussian noise are larger for the precessing templates than the restricted ones, even though significantly fewer templates were analysed with the full statistic. The results from real data are badly affected by non-Gaussianities in the data. A number of loud transients are clearly visible as short duration peaks of large SNR, while there are an even greater number of quieter peaks throughout the analyzed time. This has a similar effect on both the full and restricted waveforms.

In chapter 5, we described and developed a number of tools which can be used to effectively remove the majority of the non-Gaussian features from a non-spinning, coherent analysis. These include null stream consistency [112], amplitude consistency and χ^2 signal consistency tests [80, 113]. All of these can be applied without modification to the restricted PTF search, and it seems reasonable to expect they would be similarly effective in reducing the effect of non-Gaussianities in the data. However, we currently have no such tools which can be used for the full PTF waveforms. Before using this search on real data we will need to implement

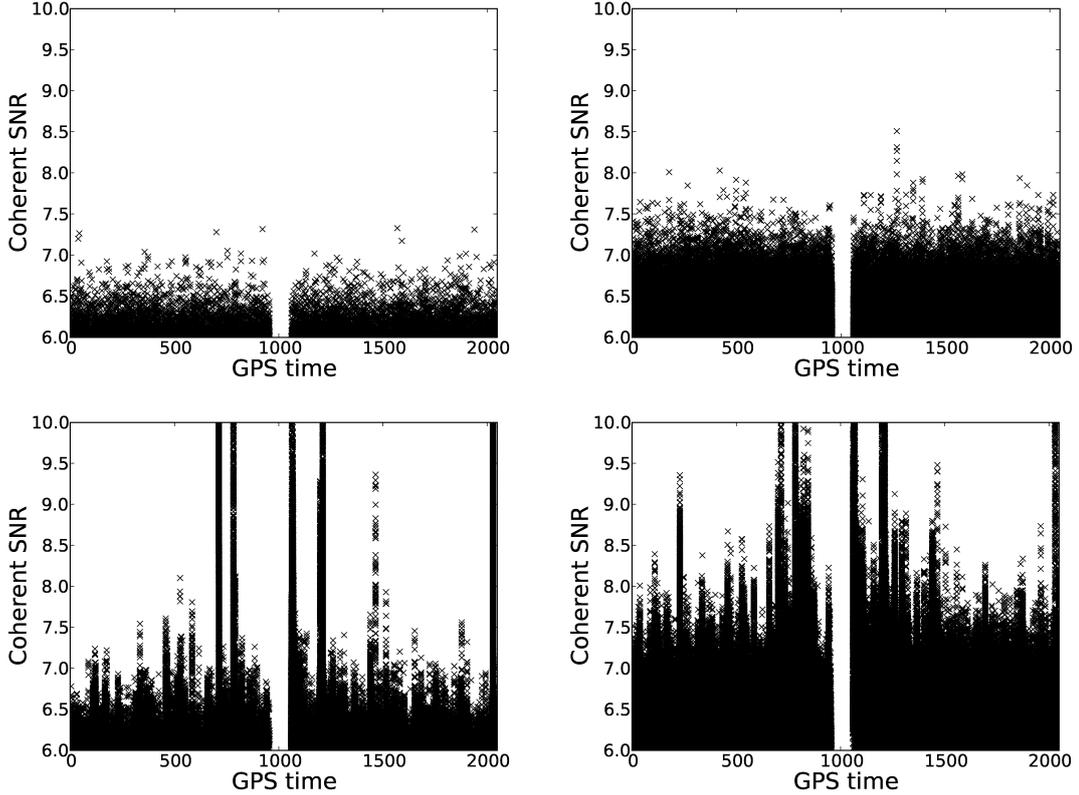


Figure 6.4: The distribution of triggers found by the coherent PTF search. The left panels show the distribution of triggers from templates that were analysed using the “restricted” coherent PTF search, the right panels show the distribution from the templates that were analysed using the “full” coherent PTF search. The top panels were created from analysing Gaussian noise. The bottom panels were created from analysing a stretch of real data from S4. All these plots have been rescaled to use the same y-axis. For the two cases using real data the non-Gaussian spikes extend much higher than is shown, the loudest trigger has an SNR of 39 in the restricted case and 45 in the full case.

a set of tests that can discriminate glitches from real signals for the full statistic. It should be relatively straightforward to implement the null stream consistency test. Unfortunately, as discussed in chapter 5, the null stream for the LIGO S4 detectors is constructed only from the two instruments in Hanford. In this stretch of data the loudest background triggers are caused by non-stationarities in the Livingston detector and thus the null stream is ineffective. Alternatively a χ^2 test such as the ones described in [80, 113] could be adapted to this search, [127] presents a possible way of doing this for single detectors. We are working on developing an alternative version of this χ^2 test, which would test the consistency of the six independent components of a single detector PTF waveform, and then extending this to the fully coherent analysis.

6.8 Discussion and generic spin systems

In this section we have explored how the presence of spin on the components affects the dynamics of a compact binary coalescence. We have discussed how one could conduct a search using templates that include spin, but not precessional effects, and presented a method for performing a coincidence and coherent search for precessing, single spin black hole–neutron star coalescences using PTF templates. We have compared the performance of PTF to searches using non-precessing waveforms and have identified regions of the parameter space where the PTF search offers increased sensitivity. We have presented a method by which these areas could be identified and demonstrated these techniques on a short stretch of S4 data.

The PTF method should allow for the detection of highly precessing NSBH systems with greater efficiency than the current non-spinning searches. However, more work is required before this search is ready to be used. The main need is for the development of effective methods of separating glitches from real events in the full PTF search, whether performing a coincident or coherent search. It should be possible to adapt a lot of the methods that have proven effective in non-spinning searches but this is a non-trivial task.

We also demonstrated that the coincidence and coherent PTF statistics allow for a large degree of unphysical freedom when maximizing over the extrinsic parameters. The efficiency of a PTF search could be increased if this unphysical freedom is reduced or removed. Additionally the efficiency of the PTF search could be increased if a marginalization technique was used to downweight physically unlikely values of the parameters. We are investigating methods of doing both of these things.

Real CBC systems will likely have some spin associated to *both* compact bodies and it is important that we are able to detect any manner of physically possible CBC system. Unfortunately, at the moment there is no reliable search method that uses double spin systems as templates. However, the aligned spin and single spin templates are capable of covering a good deal of the parameter space. It is believed that the spin of the components of a BNS system will be very small when the system becomes sensitive to ground based interferometers [141]. For these cases non-spinning templates have been shown to have a good sensitivity to most physically possible BNS systems, at least given initial detector sensitivities [1]. For NSBH systems, any spin on the neutron star will have a negligible effect on the overall dynamics of the system. Therefore the PTF search should be able

to detect any manner of NSBH system.

The only potential problem arises with BBH systems. However, even here it is hoped that aligned spin templates and single spin templates (including single spin templates with overmaximal spin, $\chi > 1$), have a good ability to detect such systems. For initial LIGO, the combination of non spinning, aligned spin and single spin searches should be able to detect almost any manner of CBC signal that might be in the data, though more investigation is needed to verify this. It is important however, that the aligned spin and PTF searches are run on the current LIGO and Virgo data. It is unlikely, but there could be a highly precessing NSBH signal buried in the current data that the non-spinning pipeline has failed to detect.

As the second and third generation of gravitational wave detectors are developed, and the sensitive frequency range of the instruments increases, it will become increasingly important to develop a search using double spin systems as templates. Many of the techniques we have discussed for the single spin PTF search would be equally applicable to the dominant harmonic of *any* family of precessing waveforms. In particular, the method of maximizing freely over the amplitudes of the five components of the $l = 2$ spin weighted-spherical harmonic is directly applicable to other waveform families. As the catalogue of numerical simulations of precessing binaries grows, these methods may well find applications in searches such as ones using numerical relativity inspired inspiral-merger-ringdown waveforms.

Chapter 7

A stochastic template placement algorithm for gravitational wave data analysis

The past two decades have seen the development of methods [153, 79, 88, 154] for setting up template banks which minimize the computational cost in a search without reducing the detectability of signals. For instance, a geometric framework was developed [155, 87, 86] in the 1990’s to address the problem of template placement. This works quite well when the parameter space is of a small dimension (2, 3, or 4 at most) [85, 156, 84, 83]. The most important tool in this geometric framework is a positive-definite *metric* which measures the fractional loss in (squared) signal-to-noise ratio of a putative signal (at one point in the parameter space) filtered through the optimal filter corresponding to a nearby point in the parameter space. The metric gives the parameter space the geometric structure of a (possibly curved) Riemannian manifold, which is often called the *signal manifold* (in this chapter we continue to refer to it as the parameter space).

When the dimension of the parameter space becomes large there are problems with existing methods. First, even for flat parameter spaces, there are no known optimal placement algorithms for dimensions greater than 5 (the analogue of the two-dimensional hexagonal lattice) [157] (and references therein). Second, it is not clear how to place templates in a curved parameter space. For example, one cannot set up an optimal (equally-spaced) lattice on a two-sphere unless the number of points is very small (for example, 12). This issue becomes increasingly important in parameter spaces with dimension greater than 2. Third, if the parameter space includes irregular boundaries, or is formed of regions with differing dimensions, it is extremely difficult to “step around” the parameter space in a deterministic way

that covers the parameter space completely but does not significantly over-cover it.

This chapter gives a template placement algorithm that works for any parametrized signal model in any number of dimensions, provided that one can determine if two points in the parameter space are a large metric distance apart, and, if they are not, accurately calculate the metric distance between them. The idea is simple. Pick points at random in parameter space, rejecting any points that are too close to those previously retained. Continue this process until no new points are added, because any newly selected random points are close to previously retained points. We call this a *stochastic template placement* algorithm, and the resulting grid a *stochastic template grid* or *stochastic template bank*.

By construction, the stochastic template bank does not over-populate the parameter space. But does it properly populate all regions? The answer depends upon the properties of the signal manifold and its metric. It is very similar to the question of whether the Monte-Carlo approximation to an integral converges to the correct value. And in the same way as with Monte-Carlo integration, these stochastic template banks appear to perform very well in real-world applications.

This method is closely related to another way of creating random template banks, [158], in which the filtering stage is not carried out, but has certain advantages. In particular, fewer templates are needed to obtain a given degree of coverage of the parameter space. However, the filtering stage can become computationally expensive.

Some practical issues remain. The most convenient way to generate a random template bank is to use computer-generated uniformly-distributed random numbers as random coordinate values in parameter space. However, the distribution of the resulting points then depends strongly upon the choice of the coordinate system. If global coordinates can be found in which the determinant of the metric is constant (or nearly constant) then choosing uniformly distributed random numbers for the coordinate values will result in a uniform density of points. This is optimal. If not, the random points should ideally be generated with a probability density in coordinate space proportional to the square root of the determinant of the metric in those coordinates. (One can also pick a small number of points in the space, and at each point define a local coordinate system in which the metric is proportional to δ_{ab} , then place many points uniformly in those coordinates.) In practice, this is not necessary: this chapter shows that a stochastic template bank can still be effectively generated by choosing uniform probability distributions for the coordinate values, even if the determinant of the metric is *not* constant

on those coordinates. The only downside is additional computational cost as we discuss in Sec. 7.2.

Later in this chapter, two examples are shown to illustrate this: the placement of templates in a D -dimension cube, and the placement of templates on the signal manifold of gravitational wave chirps from inspiralling compact binaries calculated in the first post-Newtonian approximation. In both cases, one can create stochastic template banks using coordinates (polar, and masses (m_1, m_2) , respectively) in which the determinant of the metric is not constant. This incurs unnecessary computation cost, but it works. Alternatively, one can create a stochastic template bank using coordinates (Cartesian, and chirp-time coordinates (τ_1, τ_3) , respectively) in which the metric (and hence its determinant) is approximately constant [154]. This works better, since it is computationally more efficient, but the end result is the same.

The chapter is organized as follows. Sec. 7.1 presents the stochastic template placement algorithm. An implementation and results of testing are presented in Sec. 7.2 for some simple cases where the number of templates is known analytically. Sec. 7.3 is devoted to the application of the algorithm to the case of gravitational wave chirps from inspiralling compact binaries where the performance of the stochastic template placement method is compared with existing geometrical template placement algorithms.

7.1 Stochastic template placement algorithm

Let \mathcal{M} denote a signal manifold of dimension D , with $d(x, y)$ being a positive-definite distance function. Here $x, y \in \mathcal{M}$ are points in the manifold. Note that the signal manifold \mathcal{M} might cover only part of the space of possible signals of a particular type, for example one might only want to lay a bank to search for binary inspiral signals within a specific range of masses.

A template bank T is a set of n points taken from \mathcal{M} : $T = \{x_1, \dots, x_n; x_i \in \mathcal{M}\}$. A template bank is said to cover the signal manifold with radius Δ (or to be complete) if every point in \mathcal{M} lies within distance Δ of at least one of the n points: $\forall y \in \mathcal{M}, d(y, x_i) \leq \Delta$ for at least one $i \in 1, \dots, n$.

An optimal template bank of radius Δ would fulfill two conditions. First, it would cover the signal manifold with radius Δ . Second, it would contain the minimum number of points. However, it is difficult to achieve this in practice!

The method proposed in this chapter creates a template bank according to the following algorithm:

1. Let T be a list of n points from \mathcal{M} . Initially $n = 0$ and the list is empty. As points get added to this list, they will be denoted by x_1, \dots, x_n .
2. Pick a point z at random from \mathcal{M} . If $d(z, x_i) > \Delta$ for all points in the list T , then add z to T and increment n by one. Else discard the point z .
3. Repeat the previous step, until the list T stops changing in length, or some other stopping criterion is met.

7.1.1 Expected size of complete stochastic template banks

An important question to ask is at what point will this iterative process terminate? This is determined by the number of templates needed to completely cover the space. To understand this, it is useful to first ask the more general question, how large does a complete template bank (not necessarily one generated by the algorithm above) need to be? To try to understand these questions this sub-section begins by discussing upper and lower bounds on the size of the stochastic bank. Two commonly used lattice algorithms are then discussed and the performance of the stochastic bank, at low dimension, is compared to these quantities.

In this discussion we follow [159, 158] and use *thickness* (Θ) and *normalized thickness* (θ) to assess the efficiency of a specific template covering. Thickness is defined [159] as the average number of templates covering any point in the parameter space while normalized thickness is defined as the number of templates per unit volume in the case where the radius of the templates is unity. They are related by [159]

$$\theta = \Theta/V_S, \tag{7.1}$$

where V_S is the volume enclosed by a D -dimensional sphere of unit radius

$$V_S = \frac{2\pi^{D/2}}{D \Gamma(D/2)}. \tag{7.2}$$

The advantage of using these quantities is that they are independent of the size of the parameter space and independent of the template radius. These quantities are also directly related to the number of templates that will be required [158], by

$$\theta = \frac{n\Delta^D}{V} \tag{7.3}$$

where V is the proper D -volume of the parameter space,

$$V = \int_{\mathcal{M}} \sqrt{g} d^D x \quad (7.4)$$

and g is the determinant of the metric g_{ij} on the manifold \mathcal{M} .

We also assume, in this section, that “boundary effects” can be ignored. Except in pathological cases, this is true if the total volume within distance Δ of $\partial\mathcal{M}$ is small compared with the total volume of \mathcal{M} .

A simple theoretical lower bound on the number of templates needed in any complete template bank is the ratio of the volume of the parameter space to the volume of a single template. The volume of a single template is the D -volume contained in a ball of radius Δ is given by

$$V_{\text{template}} = B(\Delta) = V_S \Delta^D. \quad (7.5)$$

Hence the number of required templates is bounded below by V/V_{template} . Alternatively we can say that the thickness of a complete template bank must be greater than unity or that the normalized thickness must be greater than $1/V_S$. For the case of flat spaces a great deal of work has been carried out in trying to obtain better estimates of the minimum possible thickness for a complete template bank; it is clear, for example, that even in the 2 dimensional case a complete template bank cannot have a thickness of 1, there must be some overlap between the templates. In [159] the best currently known theoretical bounds on thickness are given and these are the values that are shown as the lower bound in Figure 7.1.

To try to obtain an upper bound on the thickness of the stochastic template bank one can consider the sphere-packing problem, this is the question of how many non-overlapping spheres can be packed into a certain volume. Consider the packing problem with hard spheres of radius $\Delta/2$. Since the centers of any of these spheres are distances of Δ or more apart, they are suitable locations for a stochastic template bank. In Ref. [159] a bound is given on the number of hard spheres of radius $\Delta/2$ that can be placed into a volume V . This can be considered as an upper bound on the number of templates that the stochastic algorithm can place. Figure 7.2 also suggests that this bound may be a reasonable estimate of the thickness of a complete stochastic bank, it can be seen that the average minimum distance between any template and the rest of the bank is close to Δ , as it would be in the sphere packing problem. However, at least for low dimension ($D < 4$), Figure 7.1 shows that a complete stochastic template bank requires considerably fewer templates than this sphere packing upper bound.

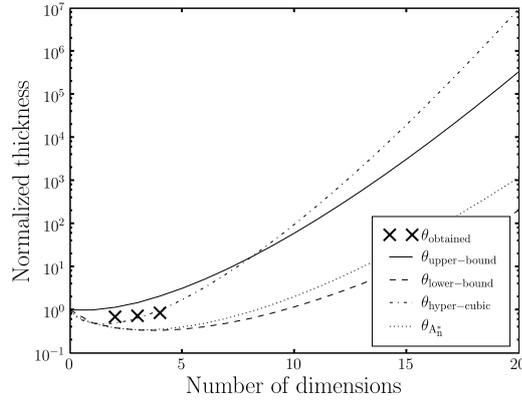


Figure 7.1: The theoretical upper and lower bounds on normalized thickness of a stochastic template bank [159] and the normalized thickness of known lattice algorithms as a function of dimension as defined by equations (7.7) and (7.8). Also the obtained thickness of stochastic banks at dimensions less than 5.

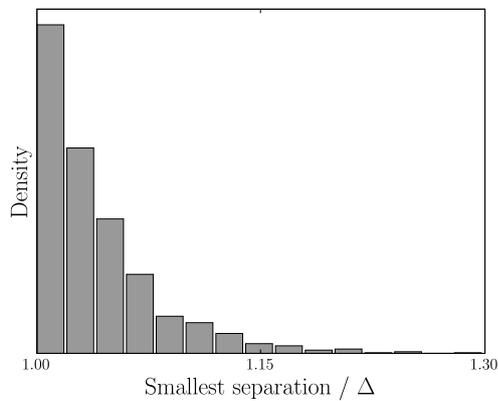


Figure 7.2: A histogram of the distribution of distances from a template to the nearest template, in units of the closest possible spacing Δ , for a simple three-dimensional example. The distances are clustered close to the minimum possible spacing Δ , showing that the covering locations found by the stochastic template placement method are close to the positions found by packing spheres of radius $\Delta/2$.

It is also useful to compare this with the performance of known lattice algorithms. In this work two different lattice algorithms are considered. The first is the hyper-cubical lattice, where the hyper-cubes of the lattice are just small enough to fit entirely inside a single ball of radius Δ . The side length δ of such a cube is given by

$$\delta = 2\Delta (1/D)^{0.5}, \quad (7.6)$$

since the longest diagonal of this D -cube is then 2Δ long. Thus

$$\theta_{\text{hyper-cubic}} = \frac{D^{D/2}}{2^D} \quad (7.7)$$

describes the normalized thickness of a template bank in a hyper-cubic arrangement.

The second lattice algorithm considered in this work is the A_n^* lattice [159, 157]. The 2 dimensional A_n^* lattice is the well known hexagonal lattice. For this algorithm the normalized thickness is given by [158]

$$\theta_{A_n^*} = \sqrt{D+1} \left[\frac{D(D+2)}{12(D+1)} \right]^{D/2}. \quad (7.8)$$

From Figure 7.1 it can be seen that the A_n^* lattice requires less templates than the hyper-cubic lattice in all dimensions (except the trivial one-dimensional case). It is also the most efficient lattice known in dimensions up to 20 [159]. This figure also shows that the number of templates required to create a complete stochastic bank is less than the hyper-cubic lattice, but only when the dimension D is greater than 3. A stochastic template bank with full coverage, however, will require more templates than the A_n^* lattice at least up to four dimensions. We have no reason to believe that a complete stochastic bank will be more efficient than the A_n^* lattice in any dimension.

One must consider however that these lattice algorithms are only defined in the case of flat parameter spaces. The stochastic algorithm on the other hand can be used in any parameter space and it is in the cases where the parameter space is not flat that we believe the stochastic bank would be the most useful.

7.1.2 The convergence of a stochastic template bank

In real world applications it may not be necessary for the template bank to be complete. It is therefore useful to be able to understand the convergence of the iteration that creates a stochastic template bank. This sub-section is devoted to

trying to understand this convergence and comparing it to the method describe in [158].

To begin to understand how a stochastic bank converges it is necessary to define a *covering fraction* $f \in [0, 1]$. The covering fraction is the ratio of the volume of the subset of \mathcal{M} that lies within a distance Δ of the points in the template bank, to the total volume of \mathcal{M} . The expected number of trials required to add a new template to the list is given by $1/(1-f)$, as can be seen by considering the template placement process as a form of Monte-Carlo integration.

At the beginning of the iterative process, the template bank is empty, and $f = 0$. After the first template is added (and assuming that boundary effects can be ignored!) the covering fraction is $f = \epsilon$, where $\epsilon = V_{\text{template}}/V$, which is the fraction of the entire volume covered by a single template. During the first iterative steps, while the number of templates n in the bank is small, $n \ll 1/\epsilon$, the covering fraction increases linearly with the template number according to $f = \epsilon n = nV_{\text{template}}/V$.

How does the covering fraction increase when n becomes larger? To understand this, it is helpful to first consider the behavior that the covering fraction would have in the case where the n points in the template bank were simply selected at random from \mathcal{M} , without any consideration of whether or not they were closer together than Δ . This case is considered in some detail in a recent paper on **random** template banks [158]. (In contrast, we use the name *stochastic* template bank.) In that case, since on average each additional template removes a fraction ϵ of the volume that is not already covered, one obtains

$$E(f(n)) = 1 - \exp(-\epsilon n) \quad (7.9)$$

or

$$E(f(\Theta)) = 1 - \exp(-\Theta) \quad (7.10)$$

for the expectation value of the coverage. For small n , this gives a linear increase in the covering fraction, which also describes the stochastic template bank.

Compared to the random template bank, on the average, a stochastic template bank gives higher coverage for a given number of templates. This is illustrated in Figure 7.3, which shows the covering fraction as a function of thickness, where the signal manifold \mathcal{M} is a unit box in 2, 3 and 4 dimensions. Thus, if it is desirable to minimize the computation cost because a single template bank is going to be used and re-used many times, the stochastic banks could offer a significant improvement compared with the random ones. The graph does seem to indicate, however, that

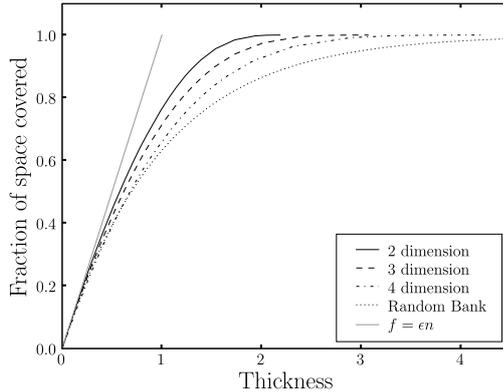


Figure 7.3: The relationship between the covering fraction and the thickness of the bank in 2, 3 and 4 dimensions. This is also compared to what one would expect in the case of the random template bank [158], as well as the case where no templates overlap each other.

the stochastic bank converges toward the random case as dimension increases. Further investigation is needed to demonstrate what level of improvement the stochastic bank would have over the random bank at high dimension.

7.1.3 Computational cost of filtering templates

While the stochastic bank will provide a better coverage for the same number of templates, one must incur an extra computational cost to carry out the filtering stage of the stochastic placement algorithm. This sub-section investigates what this computational cost would be as a function of number of templates and covering fraction.

If every random point was accepted as a template, because it was farther than Δ from all previous templates, then the computational cost would be

$$C = \alpha n(n - 1)/2, \quad (7.11)$$

where α is the cost of computing the distance between two points. This follows because the distance must be calculated between all possible pairs of templates, and there are $n(n - 1)/2$ such pairs. This also correctly describes the cost of stochastic template bank creation when the covering fraction is substantially less than one, and few potential templates are rejected. But when the covering fraction approaches one, the computational cost explodes, because the dominant computational cost is the cost of rejecting templates. This is shown in Figure 7.4.

This also allows us to provide an estimate of the computational cost. In prac-

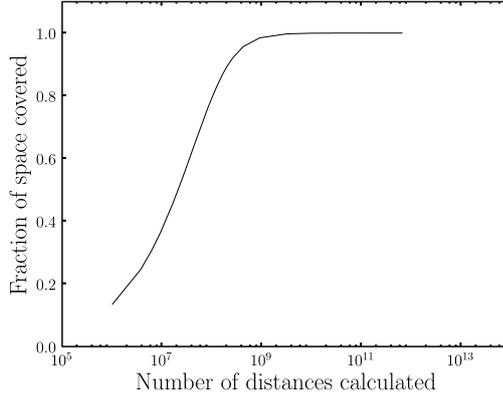


Figure 7.4: The computational cost (number of distance calculations) depends on the covering fraction.

tice, 100% coverage is not necessary or desired. For a typical binary inspiral search one might be happy with a coverage $f \in [0.9, 0.99]$. For such coverages the computation cost is bounded above by

$$C = \frac{\alpha}{1-f} \frac{n_{\text{estimated}}^2}{2}, \quad (7.12)$$

which is obtained by assuming that the cost of adding the last template is the same as the cost of adding every template. This is an upper bound because the factor of $1/2$ is larger than $n(f = 1/2)/n_{\text{estimated}}$, and because the computational cost of adding the earlier templates is smaller than that of adding the final template.

The computational cost of this method grows faster than the square of the number of templates. However, there is a modified version of this algorithm in development that has a cost proportional to $n \log n$.

This works by (conceptually) dividing the space into a set of hypercubic cells whose linear size is 2Δ . Each template is assigned to a particular cell. The algorithm maintains an internal list of hyper-cubic cells, which contain points separated by distances of less than 2Δ . When a new random template is considered, its distance only needs to be compared to the points in the same cell, and the $3^D - 1$ neighboring cells. The process of looking to see if there are neighboring cells requires a binary search in an index list, and accounts for the additional $\log n$ factor.

It is this prohibitive computational cost that has prevented us from being able to test the stochastic template bank in dimensions higher than 4 without boundary effects becoming rather pronounced. With this improved version of the algorithm it is hoped that a test of the stochastic bank in higher dimensions can

be performed.

7.2 Testing the algorithm

This section investigates how the stochastic template placement performs in dimensions less than 5 and how this compares with geometrical placement algorithms. An investigation of how the algorithm performs when the distribution of initial seed points is not proportional to the determinant of the metric is carried out as well as a demonstration that the stochastic algorithm will perform well in intrinsically curved parameter spaces.

7.2.1 Templates in flat spaces of different dimensions

First consider a flat unit hyper-cube in D -dimensions, with Cartesian coordinates and the metric $g_{ij} = \delta_{ij}$. Each coordinate lies in the range $[0, 1]$. Δ is chosen so that $1/\epsilon$ is equal to 10000.

Figures 7.3 and 7.4 show the coverage, in 2, 3 and 4 dimension, and computational cost as a function of the number of templates in the bank. The coverage in Fig. 7.3 was computed using Monte-Carlo integration with 20,000 sample points. The coverage is the fraction of these points that are less than Δ from a template in the bank. To generate Fig. 7.3 as well as Figures 7.5, 7.4, 7.9 and Tables 7.1 and 7.2 this process was carried out 100 times and the mean of the values obtained was used.

Figure 7.1 compares the number of templates being converged upon by the stochastic bank with the estimates and the lattice algorithms as described in section 7.1. It can be seen from this table that the stochastic template banks perform better than the naive hyper-cubic lattices as the dimension D of the parameter space increases. This is what was predicted in the previous section: the stochastic template banks converge to “complete” coverage with fewer templates than would be needed in a cubic lattice. Also, as predicted, the A_n^* lattice is more efficient than the stochastic bank when the stochastic bank has reached complete coverage.

An interesting feature, which is more noticeable when the number of templates in the banks is reduced, is that they show effects due to the boundaries, especially noticeable in the higher dimensions. This is easy to understand. Any template located closer than distance Δ to the boundary of the unit D -cube will have part of its coverage region lying outside the cube. Consequently, if Δ is too large, then

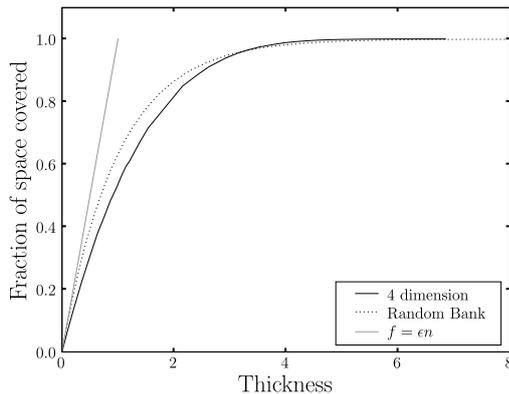


Figure 7.5: As Figure 7.3 but setting the value of $N_{\text{lower-bound}}$ to be 50. By comparing the two figures one can see how boundary effects manifest themselves both by decreasing the initial slope df/dn and by requiring a much larger number of templates than the estimate.

many of the templates will fail to produce the amount of coverage that would arise if no boundaries were present. Thus, a sign that boundary effects are appearing is that the initial coverage grows more slowly with template number than expected.

This effect can be seen in Figure 7.5 where the initial slope $df/d\Theta$ at $\Theta = 0$ is smaller than unity and also the final thickness is much larger than the estimate, which was not seen in Figure 7.3. At what template radius Δ do boundary effects become significant? This can be easily understood by estimating the volume that lies within distance Δ of the boundary of the D -cube. This is $V_{\Delta\text{-boundary}} = V - (V^{1/D} - 2\Delta)^D \approx 2D \Delta V^{1-1/D}$. Hence the initial coverage, when $\Theta \ll 1$, is

$$\left(\frac{df}{d\Theta}\right)_{\Theta=0} = 1 - \beta \frac{V_{\Delta\text{-boundary}}}{V}. \quad (7.13)$$

Here β is a numerical factor, of order $1/6$ in three dimensions, which measures the average fraction of volume of a template that lies outside the cube, as the center of the template moves through all positions in the Δ -boundary.

7.2.2 Choice of coordinate system and convergence of template numbers

How does the convergence of the stochastic template bank generation depend upon the distribution of the random template candidates in the underlying parameter space? This question is of practical interest, because the optimal distribution of the random points has a probability proportional to the volume element

N	Cartesian				Polar			
	n	σ_n	f	σ_f	n	σ_n	f	σ_f
1500	1397.1	1.0	0.1353	0.0003	1313.4	1.3	0.1246	0.0003
5000	3994.4	2.6	0.3632	0.0004	3613.9	2.6	0.3237	0.0004
15000	8494.5	3.9	0.6818	0.0004	7605.3	3.9	0.6140	0.0004
50000	13961.5	4.1	0.9221	0.0002	12979.6	4.0	0.8825	0.0003
150000	17307.8	4.0	0.9847	0.0001	16676.6	3.6	0.9747	0.0001
500000	19365.5	3.1	0.99746	0.00004	19025.4	3.5	0.99582	0.00005
1500000	20439.3	3.0	0.99949	0.00001	20241.3	3.6	0.99917	0.00002
5000000	21141.9	3.5	0.99990	0.00001	21023.4	3.2	0.99987	0.00001
10000000	21401.3	3.1	0.999971	0.000003	21305.4	3.3	0.999948	0.000005

Table 7.1: Number of templates n and fractional coverage f with associated standard deviations as a function of the cumulative number of trials N in the case of Cartesian (n_C) and polar (n_P) coordinates.

$\sqrt{\det(g_{ab})}d^Dx$. However, it can be difficult in practice to generate such a distribution, whereas it is simple to generate random points that have a uniform distribution in the coordinates.

For example, in two-dimensional flat space, one could choose trial points with uniform probability distributions in polar coordinates. This means that too many random templates are tested from the region near the origin, and then rejected. However, they are soon rejected, as being too close to points already in the template bank, and in the end, the template points that survive have the correct probability distribution proportional to $dx dy = r d\theta dr$. This is shown in Figure 7.6.

The only disadvantage of using a non-optimal distribution of random points is that the computational cost of generating the stochastic bank will increase. The added computational cost can be easily estimated: it is proportional to the square of the number of grid points which are un-necessarily added and then subsequently removed. This is bounded above by the ratio $\sqrt{g_{\max}}/\sqrt{g_{\min}}$ where the maximum and minimum are taken over the entire parameter space. More precisely, the factor is proportional to the square of the average over the space of $\langle \sqrt{g} \rangle / \sqrt{g}$ where the angle brackets mean “volume average”.

Table 7.1 shows the number of templates n as a function of the number of trial points N for random template candidates distributed uniformly in Cartesian and polar coordinates. Also shown is the coverage of the template bank, calculated using Monte-Carlo integration as described in the previous section.

This is very useful because in many cases one does not know coordinates in which the determinant of the metric is constant. Of course one could simply distribute points with a probability density proportional to the volume measure!

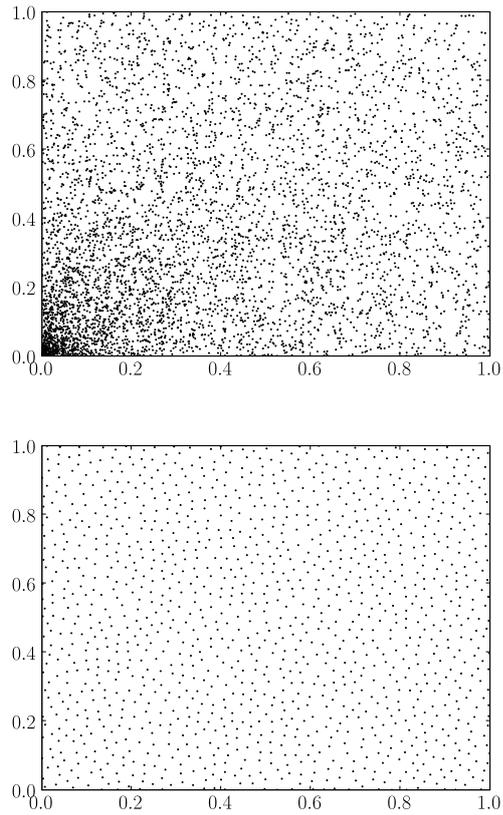


Figure 7.6: The distribution of trial points chosen uniformly in polar coordinates (left panel) and the points that remain as templates after the application of the stochastic placement algorithm (right panel).

7.2.3 Templates on a sphere

So far in this chapter we have considered templates in a flat signal space. However one can consider examples where the signal manifold is not flat, but is curved. This introduces two new issues.

First, the distance between widely separated points can no longer be easily computed. However, the only important case is the one in which the points are nearby. In this case, one can use the metric to approximate the distance at small separations:

$$dl^2 = g_{ik}(x_A^j) (x_A^i - x_B^i) (x_A^k - x_B^k). \quad (7.14)$$

Since the components of the metric can be expensive to calculate, an efficient approach is to calculate and store those components only for points that are included in the template bank. Those metric components are then used for the distance comparisons with potential new (randomly chosen) template candidates.

Second, depending upon the choice of the coordinate system, the determinant of the metric may be non-constant. In this case, an efficient approach would be to generate random points with a probability distribution proportional to the volume element $\sqrt{\det(g_{ab})}d^Dx$. However, in practice one can generate points with *any* distribution in the coordinates: the stochastic template placement algorithm simply rejects those points that are not needed, and produces a distribution with the correct density proportional to $\sqrt{\det(g_{ab})}d^Dx$.

To demonstrate the performance of the stochastic template placement algorithm on a curved manifold, consider a unit-radius two-sphere S^2 with standard spherical polar coordinates (θ, φ) . The metric is

$$dl^2 = d\theta^2 + \sin^2 \theta d\varphi^2. \quad (7.15)$$

Table 7.2 shows the number of templates n as a function of the number of trial points N . The size of the templates has been chosen so that the ratio $\epsilon = V_{\text{template}}/V$ is the same as for the unit cube examples given in the previous section. In this case the stochastic algorithm converges for a smaller number of templates than for the unit cube. This is for the reasons described above: since the unit sphere has no boundary, no templates lie partly outside the space, so every template provides the largest possible coverage.

Figure 7.7 shows the distribution of 5000 candidate points, chosen uniformly in spherical polar coordinates (θ, φ) (top panel). The points that survive and remain in the stochastic template bank are shown in the middle panel. A histogram of the distribution of the templates as a function of θ is also shown (right panel). As ex-

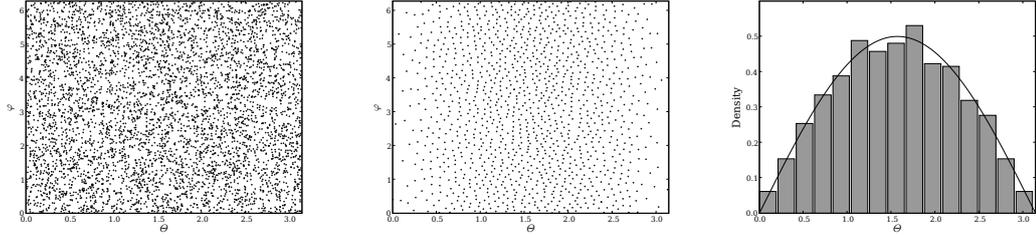


Figure 7.7: Trial points chosen uniformly in the (θ, φ) coordinates (left panel) and the templates on the surface of a sphere of unit area that remain after the application of the stochastic template placement algorithm (middle panel). Also the distribution of these templates in θ -coordinate (right panel) where the solid line shows the expected distribution.

pected, the density of templates is proportional to the volume element $\sin \theta d\theta d\phi$: it is smallest at the poles and the greatest at the equator.

7.3 Templates for gravitational wave chirps

As we discussed in chapter 3, binary systems of compact objects (i.e., black holes and/or neutron stars) evolve by emitting gravitational radiation. The loss of energy and angular momentum into gravitational waves causes the two bodies to spiral in toward each other, emitting a burst of radiation just before they merge. Although there is no exact solution to the two-body problem in general relativity, an approximation method called the post-Newtonian formalism has been used to compute the amplitude and phase of the waves emitted in the adiabatic inspiral phase to a very high accuracy [160, 39, 161] (for a recent review see [69]). Moreover, recent progress in numerical relativity has provided a good knowledge of the waveform even in the strong gravity regime of the merger dynamics [162]. Thus, one can use matched filtering to dig out astrophysical signals from the noise background of an interferometric detector.

In general, the radiation from a binary is characterized by as many as seventeen parameters. However, some of these parameters (the distance to the binary, the inclination of the orbit relative to the detector, etc.) only affect the amplitude of the waveform, which does not modify the search template. Therefore, one would only need to place templates in a lower-dimensional parameter space. State-of-the-art template placement algorithms deal only with a binary composed of non-spinning objects; in which case templates are only needed in the two-dimensional parameter space of the masses of the component objects. We discussed some simple techniques to create 3 or 4 dimensional template banks in chapter 6 but

N	n	σ_n	f	σ_f
1500.0	1330.3	1.2	0.1989	0.0004
5000.0	3688.3	2.8	0.4253	0.0004
15000.0	7807.1	3.7	0.7027	0.0004
50000.0	13198.7	4.1	0.9192	0.0002
150000.0	16779.5	4.4	0.9836	0.0001
500000.0	19007.0	3.6	0.99735	0.00003
1500000.0	20171.6	3.6	0.99947	0.00002
5000000.0	20906.6	3.8	0.99991	0.00001
10000000.0	21182.9	4.0	0.999967	0.000004

Table 7.2: Number of templates n and coverage f as a function of the number of trials N on a sphere of unit radius.

these methods were far from optimal and did not guarantee that a template bank would be complete. The stochastic placement algorithm would be preferable in such cases.

The goal of the stochastic template placement algorithm is to address the problem of choosing templates on a manifold of arbitrary dimensions. However, in this chapter, the algorithm is only applied to the case of a binary consisting of non-spinning bodies where the results are well known, thus facilitating a straightforward comparison with established results. This algorithm has also been applied in a recent search for spinning binaries in the first year of LIGO’s fifth science run using templates placed in a three-dimensional parameter space [126] as well as a five-dimensional search for super-massive black holes in a mock data challenge, see chapter 8. A similar, but independently developed, algorithm was also used in this mock data challenge [163]. This algorithm was effectively the same as the one described in this work but the author calculates the overlap between points explicitly, instead of using the metric approximation as in this work. While this will more accurately determine the overlap, especially for overlaps not close to unity, it will come at considerable additional computational cost.

7.3.1 Choice of coordinate system

Begin by choosing a suitable coordinate system on the signal manifold. The masses m_1 and m_2 of the component stars are the most obvious coordinates on the manifold. However, when one uses masses as the coordinate system the determinant of the metric will vary significantly over this parameter space [85]. Because of this a much higher density of templates is needed in the low mass region than in the high mass region.

A better coordinate system is *chirp times* [154, 84], defined by

$$\tau_0 = \frac{5}{128\pi f_L \eta} (\pi M f_L)^{-5/3}, \quad \tau_3 = \frac{1}{8f_L \eta} (\pi M f_L)^{-2/3} \quad (7.16)$$

Using this coordinate system the determinant of the metric does not vary much over this parameter space. This can be illustrated by looking at the distribution of templates in both coordinate-systems as shown in Figure 7.8. The algorithm used in this case [83] places templates first along the $m_1 = m_2$ curve (the lower right boundary in the left panel and upper left boundary in the right panel). This is dictated by the fact that the region below the equal-masses curve (in τ_0, τ_3 coordinates) corresponds to binaries with imaginary component masses¹. The algorithm uses a hexagonal placement over the rest of the parameter space.

7.3.2 Comparison of stochastic lattice with a square lattice

Let us now compare the stochastically generated template bank with a hexagonal lattice and with a square lattice. In this case the template banks are created to cover binary compact objects whose components have a mass range of 1 to 10 solar masses such that any real signal within this range of masses would have an “overlap” greater than 0.96 with at least one of the templates in the bank. This overlap, defined by $1 - \Delta^2$, is calculated using the assumption made in equation (7.14) and the metric defined in [84]. For the stochastic algorithm the trial points are placed uniformly in (τ_0, τ_3) coordinates (and limited by the restrictions on the masses). This is also compared to trial points placed uniformly in (m_1, m_2) coordinates.

For this choice of parameters and for trial points placed uniformly in both coordinate systems, the number of templates is plotted as a function of the number of trial points in Fig. 7.9. Fig. 7.10 shows the distribution of resultant templates for both initial trial point distributions.

In this two-dimensional example, the stochastic algorithm, in both cases, converges at about 7500 templates. For comparison, with the same range of masses a hexagonal lattice has 5914 templates and a square lattice has 8353 templates. This may seem to be in conflict with the statement in section 7.1 that the stochastic algorithm performs worse than hyper-cubic lattices in two dimensions. However, one must remember that the geometrical algorithm used here begins by

¹Although the waveform, which depends only on the total mass M and symmetric mass ratio η , which are real in that region, can be generated in this part of the parameter space, it is unphysical and, therefore, not of any interest.

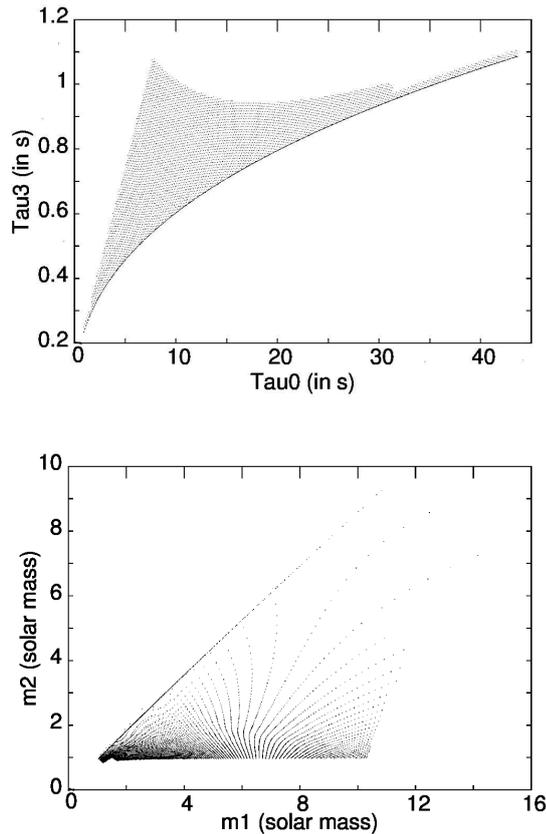


Figure 7.8: The distribution of templates placed by the hexagonal lattice algorithm in (τ_0, τ_3) coordinates (top panel) and the same templates in (m_1, m_2) coordinates (bottom panel). Clearly, the distribution is highly skewed in the latter coordinates.

placing templates along the boundaries, which is quite inefficient. One also must remember that though this parameter space is close to flat it is not flat.

7.3.3 Efficiency of the Stochastic bank

The quality or performance of a template bank can be assessed by measuring the overlap between randomly simulated compact binary signals in the relevant range of parameters of the template bank in question. To test the performance of the stochastic template banks, a set of 20,000 signals (standard post-Newtonian waveforms of type TaylorT3 [164]) was generated and the maximum overlap of each over the entire template bank was calculated. In this case templates with masses between 3 and 30 solar masses were used (a different mass range was used here to produce less templates, thus making it easier to show the results graphically) and an overlap of 0.95 was used, equivalently, $\Delta^2 = 0.05$. The template bank was

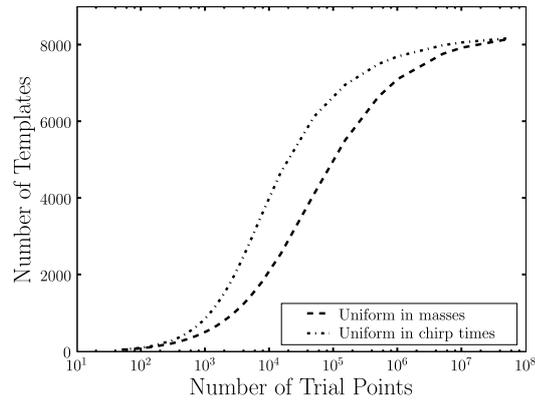


Figure 7.9: The number of templates as a function of the number of random trial points is shown when the trial points are assigned uniformly in (τ_0, τ_3) coordinates as well as uniformly in (m_1, m_2) coordinates.

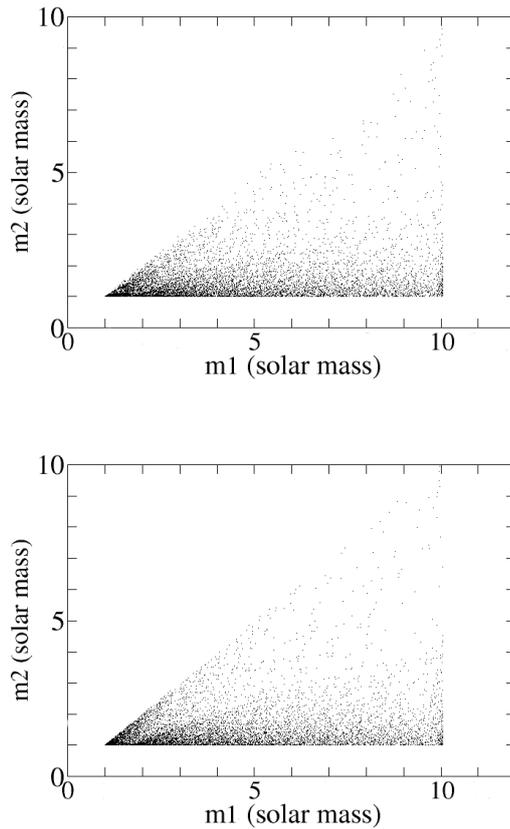


Figure 7.10: The distribution of stochastically generated templates in (τ_0, τ_3) coordinates (top panel) and in (m_1, m_2) coordinates (bottom panel).

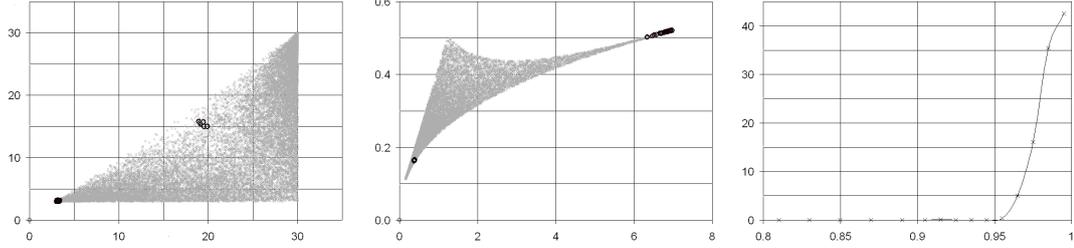


Figure 7.11: The signals which had an overlap larger than 0.95 (gray crosses) as well as signals with an overlap less than 0.95 (black circles) in (m_1, m_2) coordinates (left) and (τ_0, τ_3) coordinates (middle). Also a histogram of the overlaps of all injections with the stochastic template bank (right).

generated from 60,000 trial points placed uniformly in (τ_0, τ_3) coordinates.

The result of the test is shown in Fig. 7.11. One can see that the stochastic placement algorithm struggles to cover certain areas of this parameter space. If a larger number of trial points had been used, the coverage would have been better.

The areas of the parameter space with poor coverage from the stochastic template bank are in regions of the parameter space that are very thin, almost one-dimensional and therefore have very little volume. The hexagonal and square lattices overcome this problem by placing templates along the boundary of the space, especially along the $m_1 = m_2$ curve. A stochastic placement algorithm can overcome this problem in the same manner, or by increasing the mass range of allowed templates. But both solutions come with the cost of additional templates in the bank.

7.4 Conclusions

This chapter presents a method for stochastically generating template banks in parameter spaces of arbitrary dimension and with arbitrary metrics. The relationship between coverage and the number of templates required to reach that coverage has been investigated for dimensions up to 4. The performance of the stochastic placement algorithm has been compared to lattice placement algorithms in flat spaces and was found to only be marginally less effective at dimension less than 4. The area where we believe this algorithm would be of most use is in signal manifolds that have a large intrinsic curvature, where lattice placement algorithms can not easily be applied. Stochastic banks which cover less than 100% of the signal manifold may be useful for large dimensional manifolds, though further investigation is needed to show that this is the case.

For cases where the number of required templates is very high, the algorithm will become very computationally expensive. In these cases other “random template banks”, which do not use our filtering stage might become more practical [158]. Nevertheless the stochastic template bank will provide better coverage for a given number of templates. The construction of stochastic template banks can be made less expensive, however, by utilizing the fact that it is not necessary to compute the distance between a trial point and *every* template in the bank. This is a topic of ongoing investigation.

Chapter 8

A hierarchical search for gravitational waves from supermassive black hole binary mergers

There is growing evidence that some fraction of quasars [165], and X-ray and infrared sources [166] host supermassive binary black holes (SMBBH) that are potential sources of gravitational radiation. The late time evolution of such systems is dominated by the emission of gravitational waves, the radiation back reaction torque driving the system to coalesce. The proposed space-based Laser Interferometer Space Antenna (LISA) [167] targets gravitational waves from these systems in the frequency range of [few $\times 10^{-5}$, 0.1] Hz which corresponds to SMBBH of masses in the range [10^4 , few $\times 10^7$] M_{\odot} . Within a redshift of $z \sim 10$, SMBBH coalescence rates could be as high as several tens per year but depending on the way galaxies and black holes at their cores formed the rates could be several hundreds per year [168, 169].

At a redshift of $z \simeq 1$ SMBBHs would be expected to produce a signal to noise ratio (SNR) of several thousands. This implies that even at redshifts of $z \simeq 10$ LISA could detect SMBBH systems with an SNR of $z \sim 10$. LISA is, therefore, an excellent probe of the seed black holes that are believed to be responsible for the formation and evolution of galaxies [168, 169, 170] and the large-scale structure in the Universe.

SMBBH mergers from redshifts up to about $z \sim 3$ are expected to produce a response in LISA loud enough that they can be detected without any sophisticated data analysis; Although accurate models of the merger dynamics would be needed

for parameter extraction. Indeed, these sources will be so bright that one has to worry about systematics due to our limited theoretical understanding of their dynamics [171]. At larger red-shifts, however, it would be necessary to employ data analysis techniques that are sensitive to weaker signals, such as those discussed in earlier chapters. This is an important goal for LISA as there is significant uncertainty in when the first seed black holes and galaxies might have formed and it would be good to be able to probe as far back as a red-shift of $z \sim 10-15$. At a red-shift of z , an SMBBH of intrinsic total mass M would appear in LISA to have a red-shifted total mass of $(1+z)M$. Thus, at $z = 10$ LISA would probe masses that are intrinsically 11 times smaller than at $z = 0$. Therefore, searching for SMBBH at higher red-shifts would probe smaller masses too.

In addition to SMBBH mergers LISA will observe a host of other sources (see [172] and references therein). These include binary white dwarfs in the Milky-way (both a stochastic signal from an unresolved background population and continuous signals from resolved foreground sources), inspirals of small black holes into supermassive black holes (again a stochastic background from overlapping sources and a foreground of individual sources), etc. Analysing LISA data and resolving tens of thousands of signals belonging to different classes is unprecedented and likely to be a daunting task. As we have demonstrated in earlier chapters, matched filtering is a very powerful approach that has been extensively used in searching for gravitational wave signals in ground based detectors. In this paper we report the results from a hierarchical matched filtering algorithm to search for SMBBH mergers.

From a computational point of view, matched filtering is very expensive, computational cost increasing as a power-law of the number of search parameters. While the parameter space of stellar mass binaries consisting of non-spinning black holes is only two-dimensional, the number of parameters in the case of SMBBH, even while neglecting spins, is quite large. This is because the source's position relative to LISA changes during the course of observation, causing a modulation in the signal's amplitude and phase that must be taken into account in the search templates as well as the waveform's polarization angle. Thus, the computational cost of a naive implementation of a matched filtered search would be formidable.

In an attempt to test our ability to detect signals observed in LISA, the LISA International Science Team has put together a task force to develop a set of data analysis challenges [173] of ever increasing complexity [174] to encourage data analysts to explore and test their search algorithms on simulated data. These are referred to as the Mock LISA Data Challenges (MLDC). We participated in

Challenge 1B, a rerun of the Challenge 1, consisting of the simplest possible data sets with only one inherent signal.

In the MLDC 1B one of the challenges was to detect and characterize one SMBBH coalescence buried in LISA instrumental noise only. Two datasets were released, one where the coalescence was in the middle of the observation period and a second where the coalescence was two months after the observation period ended. We only took part in the Challenge where the binary coalesced during the observation time.

We developed a hierarchical approach to this challenge in which the goal was to zoom-in onto an interesting region of the parameter space in several steps, each of which uses a progressively greater density of templates. We tested our algorithm on the *training* and *challenge* data sets from Challenge1B of the MLDC.

For our templates we used post-Newtonian waveforms at the second post-Newtonian order. We tapered the end point of our templates to prevent the bleeding of spurious power in the Fourier domain arising from the step function that is implicit if the waveform were to be terminated abruptly. The signal is characterized by nine independent parameters. We separate these into the ‘intrinsic’ parameters consisting of the two component masses, the binary’s position on the sky and its epoch of coalescence and ‘extrinsic’ parameters comprising the inclination angle, the polarization phase, the coalescence phase and the distance to the binary. We devised a search that was capable of determining all these parameters in an efficient manner, albeit not to accuracies that are theoretically possible. Let us note, however, that the goal of this exercise is not to measure the parameters accurately but to efficiently detect the signal and constrain the parameter space well enough so that other techniques, such as the Markov Chain Monte Carlo, can be deployed in a follow-up study to determine the parameters more accurately.

Other groups have, of course, participated in the search for SMBBHs in the Mock LISA Data Challenges, and their methods differ from ours [173, 175, 176, 177, 178, 179, 118]. The majority of these searches involve a variety of methods to detect the source and constrain somewhat its parameters followed by a Markov Chain Monte Carlo followup to determine accurately all the parameters.

8.1 Search method

A search for supermassive binary black hole signals in the LISA data requires, in general, the determination of seventeen parameters. In this paper, and in Challenge 1B of the MLDC, spins of the component black holes are ignored, restricting

to non-spinning components. In addition, the orbit is assumed to be circularized sufficiently by the time it enters LISA’s sensitive band that eccentricity can be ignored. This allows us to neglect eight parameters leaving the parameters of interest to be:

- The masses of the two components of the binary, M_1 and M_2 . It is often convenient to express the mass parameters in terms of the chirp mass \mathcal{M} and reduced mass μ , defined in section 3.1.
- The time that the binary coalescences, t_c , which is assumed to be within the LISA data.
- The sky location of the binary, determined by its ecliptic latitude, θ , and longitude, ϕ .
- The orientation of the binary system, given by the inclination angle, ι , and polarization angle, ψ .
- The initial phase of the binary, Ψ_0 .
- The luminosity distance to the binary, D_L .

In this search, similarly to the coherent search described in chapter 5, we make use of the ‘ \mathcal{F} -statistic’ [46]. This lets us analytically maximize over four of the parameters introduced above (the ‘extrinsic parameters’): the distance to the binary, and its inclination, polarization and initial phase. This procedure is discussed in Section 8.1.1. We note that, unlike the earlier ground-based coherent search, the sky location cannot be treated as constant in the LISA case and thus cannot be easily maximized over.

The remaining five parameters are determined by searching over stochastically generated template banks as described in chapter 7. Since some parameters, in particular the chirp mass and coalescence time, are more easily determined we employ a hierarchical search whereby we obtain good estimates of these parameters before refining our search to determine the full parameter set. The stochastic bank used in this search is described in Section 8.1.2 and the hierarchical search method is discussed in 8.1.3. Results of this search applied to the MLDC 1B data are presented in Section 8.2.

8.1.1 Matched filtering with the F-statistic

When the signal waveform is well known, the technique of matched-filtering is typically used to search for the signal (see, for example [78, 180]). The \mathcal{F} -statistic

is an elegant way to maximize over the extrinsic parameters, and thereby simplify the search. As well as the coherent search discussed earlier, it is used in a number of other searches by other groups for the MLDC and in the pulsar search with ground-based detectors (for example [48, 181]). We quickly recap how it is used in this section. We follow closely Ref. [46] when describing this method. Given a template waveform h and the data s , we calculate the *likelihood function*, defined as:

$$\ln \Lambda = (s|h) - \frac{1}{2} (h|h) , \quad (8.1)$$

where the inner product between the signal and template is given by

$$(s|h) = 4 \operatorname{Re} \left[\int_0^\infty \frac{\tilde{h}^*(f) \tilde{s}(f)}{S_n(f)} df \right] , \quad (8.2)$$

and $S_h(f)$ is the one-sided power spectral density of the LISA detector.

For a signal present in the LISA data, it can be shown that the gravitational wave signal can be decomposed as [182]

$$h(t) = \sum_{i=1}^4 A_i(D_L, \Psi_0, \iota, \psi) \cdot h_i(t; t_c, M_1, M_2, \theta, \phi) . \quad (8.3)$$

The amplitudes A_i are functions only of the extrinsic parameters: D_L, Ψ_0, ι and ψ . The $h_i(t)$ are functions of the remaining, intrinsic, parameters only. In contrast to searches with ground based detectors, the sky location cannot be considered during the observation of a signal and must enter as an intrinsic parameter. The benefit of expressing the waveform in this manner is that it is straightforward to maximize the likelihood parameters over these A_i by requiring

$$\frac{\partial \ln \Lambda}{\partial A_i} = 0 . \quad (8.4)$$

As discussed in chapter 5 this can be shown to be equivalent to

$$A_i = \sum_{j=1}^4 \mathcal{M}_{ij}^{-1} (s|h_j) \quad \text{where} \quad \mathcal{M}_{ij} = (h_i|h_j) . \quad (8.5)$$

Therefore, the A_i can be determined from the $(h_i|h_j)$ and $(s|h_i)$. Furthermore, for each possible set of values for A_i we obtain a unique value for the four extrinsic parameters: distance D_L , initial phase φ_0 , inclination angle ι and polarization angle ψ . However, there remain implicit degeneracies in these values. Specifically,

as we use only the dominant, 2Φ harmonic in the waveform, there is a degeneracy in the initial phase corresponding to $\Psi_0 \rightarrow \Psi_0 + \pi$. The same degeneracy exists for the polarization angle. Additionally, a system with polarization ψ and phase Ψ_0 is indistinguishable from one with values $\psi + \frac{\pi}{2}$ and phase $\Psi_0 + \frac{\pi}{2}$. Finally, by substituting the expression for A_i from (8.5), the likelihood expression becomes

$$\ln \Lambda = \frac{1}{2} \sum_{i,j=1}^4 (s|h_i) \mathcal{M}_{ij}^{-1} (s|h_j) . \quad (8.6)$$

In the above discussion, we have used the gravitational wave strain $h(t)$ in discussing the \mathcal{F} -statistic. In the MLDC, the signals were released in the form of time delay interferometry (TDI) variables X , Y and Z [183, 184]. These TDI variables are used as a way of cancelling the laser phase noise in the output of LISA.¹ The \mathcal{F} -statistic method is equally applicable to the TDI variables. To maximize the efficiency of our search method we simultaneously utilize two of the TDI outputs, X and Y , to conduct our search. We do not use the Z output since the gravitational wave content in it can be constructed from the other two and is therefore not independent.

It is a trivial matter to convert the one-detector search outlined above to a two-detector TDI search. We simply rewrite our likelihood function as:

$$\ln \Lambda = (s_X|h_X) + (s_Y|h_Y) - \frac{1}{2} (h_X|h_X) - \frac{1}{2} (h_Y|h_Y) , \quad (8.7)$$

where the subscripts X and Y denote the data or template appropriate for either the X or Y TDI data stream². The \mathcal{F} -statistic maximization can similarly be extended to the two detector search. In this case, the expressions in (8.5) and (8.6) generalize to include a summation over detector. For example:

$$\mathcal{M}_{ij} = (h_{i,X}|h_{j,Y}) + (h_{i,Y}|h_{j,Y}) . \quad (8.8)$$

8.1.2 Stochastically generated template bank

Even after maximizing over the ‘extrinsic’ variables, there are still five remaining, ‘intrinsic’ parameters that we would like to determine. We utilize a template bank

¹It is assumed for the purposes of this Challenge that this time delay interferometry process will completely cancel all of the laser phase noise.

²Strictly speaking this expression is incorrect for the X and Y channels as the noise in them is correlated. It is, of course, preferable to use the synthetic A and E variables which are generated from X , Y , Z and are independent. Due to time constraints, for this challenge we did not get around to moving the code over to A E and T . This has since been implemented.

to search over this five dimensional parameter space [86, 85]. The signal manifold for the SMBBH search in LISA data is 5 dimensional and is not flat. Therefore, we use the method outlined in chapter 5 and create *stochastically* generated template banks.

While the stochastic bank generation is generic, there are certain subtleties which arise in employing it for the \mathcal{F} -statistic search for SMBBH described in Section 8.1.1. First note that, in contrast to searches for binaries in ground-based detectors, we must include the coalescence time when generating the template bank. Binary coalescence signals in ground-based detectors last at most $\sim 1,000$ seconds, during which the motion of the Earth, and the detector, can be neglected. Hence, the waveforms of binaries with different coalescence times differ only by a time-shift and amplitude rescaling. However, SMBBH signals spend several months in LISA’s sensitive band, during which LISA completes a significant fraction of an orbit around the Sun. Consequently, the template-shape depends on the coalescence time, and this parameter must be included in the template bank.

Next, we consider the effect of maximization over the four extrinsic parameters in the \mathcal{F} -statistic. This is dealt with by generating a metric on the full parameter space and projecting down to the five-dimensional subspace (see [84] for details). A complication arises in that the projected metric depends upon the value of three of the extrinsic parameters ι , ψ and Ψ_0 .³ This is a well known issue, see for example [120]. To proceed, we simply choose a fiducial value of 0.5 radians for these angles. The value was chosen arbitrarily, ensuring that none of the four A_i values was zero and they would all contain contributions from both gravitational-wave polarizations.

To generate the metric, we calculated the inner product (8.2) for the X -detector using gravitational-wave strain $h(t)$ rather than the TDI variables. This introduces two additional approximations. First, by using the strain, rather than the TDI variables, we are neglecting the directional dependence of the detector’s response function and implicitly working at the long-wavelength approximation. Second, we have performed the search using both the X and Y data streams while only the metric for X was used to generate the template bank. The above simplifications will mean that the stated minimal match of the metric would not have been achieved. However, in performing the search, as described in Section 8.1.3, we continually refined the template bank to determine the correct parameter values and did not rely on the minimal match to decide stopping conditions.

³The distance, D_L , will always enter only as an amplitude scaling and therefore will not affect the metric at all and can safely be neglected.

Clearly these issues would need to be resolved if this search were to be used in future MLDCs or in a future search of real LISA data.

8.1.3 Hierarchical search technique

Populating even the reduced, 5-dimensional parameter space with sufficient templates to determine the binary's parameters to the required accuracy would necessitate far more templates than could feasibly be filtered. Thus we must employ a hierarchical method to search for the parameters.

We began by match-filtering the data against a bank comprised of templates that are sparsely spaced and placed in only the two-dimensional space of mass parameters. This bank enabled us to make an initial estimate of the binary's masses and coalescence time with 1,000 templates in the allowed range of masses, setting the sky location arbitrarily to $\lambda = 0.5$ and $\beta = 0.5$ for all templates and fixing the coalescence time to be the value at the beginning of the allowed range. This enabled us to estimate the chirp mass and reduced mass to within 30% accuracy and coalescence time to within 10,000 seconds.

We then placed a second bank of 1,000 templates within a reduced range of the parameter space, using the best estimate of the coalescence time, sky locations again set to $\lambda = \beta = 0.5$ and repeat the process. By this method we could estimate the chirp mass to at least $\pm 5\%$, the reduced mass to at least $\pm 10\%$ and the coalescence time to within 10,000s. Using these initial estimates we were then able to place a template bank with restricted parameter ranges to determine all five of the 'intrinsic' parameters.

The final step in determining the parameters could be performed by two different methods. The first method involves placing a template bank over the full five-dimensional parameter space and using a hierarchical procedure to 'zoom in' on the true values of the binary's parameters. While this is the preferred search method, a large number of templates are still required to fill this reduced five-dimensional template space, to do this in one step would require 10^{13} templates. We would thus have to use a hierarchical procedure to construct a series of five dimensional banks, but, this search can still become computationally costly. An alternative technique is to alternate between placing two-dimensional template banks in the mass space, using the best current estimates of coalescence time and sky location, and placing three-dimensional banks in sky location and coalescence time, using the best current estimates for the masses. This method is computationally quicker as we limit the template bank size to under 1000 templates for the two dimensional case and under 10,000 templates for the three dimensional

banks. However, much more than when using 5 dimensional banks, care must be taken to avoid ‘zooming in’ on secondary maxima. For example, LISA has similar sensitivity to binary systems on opposite sides of the sky, so restricting the range of sky locations used in our template bank searches is not trivial.

The figures quoted above for template bank size and parameter accuracy are those for the binary systems in the MLDC datasets where the SNR is very large (approx. 500). For SMBBH systems where the SNR is significantly lower the main issue would be whether any templates at the initial stage were similar enough to the signal to pick it up. If so, the parameter accuracies at this stage would be similar as they are limited by the template spacing. Further investigation is warranted to determine what strength of signals can be detected by this method, how many more templates are needed at initial stage to detect weaker signals and how final parameter accuracy depends on SNR.

In future searches using this method it would be desirable to automate the hierarchical technique. To do so, we would need to quantify how many iterations are needed to adequately determine the parameters and how much each iteration reduces the possible range of values for each parameter. Although this method is still under development, it is interesting to note that it uses a comparable number of templates as the MCMC search implemented in [118]. It is also worth noting that in a template bank based search it is straightforward to parallelize the search over numerous computers.

8.2 Results

The MLDC Challenge 1B data set for SMBBH consists of one year of simulated LISA data with a single supermassive binary black hole coalescence occurring during the year. In addition, a “training” data set was released for which the binary’s parameters were also made public. Due to unforeseen technical issues we were unable to run as full an analysis as we would have liked on the challenge dataset, and our results reflect this. Therefore, we have also included the results from the training run, as they provide a more accurate reflection of the sensitivity of our current search technique. The released training data parameters were not used in running the search, as it was treated as a warm up to the challenge. For both training and challenge results we have taken into account the parameter degeneracies discussed in Section 8.1.1 by choosing the values of polarization and initial phase that are closest to the true values.

The results from the training data set are presented in Table 8.1, while the

Parameter	True Value	Our Value	Error	Fract. Error
Chirp Mass, \mathcal{M} (M_\odot)	1.3769×10^6	1.3772×10^6	360	2.6×10^{-4}
Symmetric Mass Ratio, η	0.1959	0.1972	0.0013	–
Ecliptic Latitude, β	1.028	1.072	0.044	–
Ecliptic Longitude, λ	5.050	5.037	0.013	–
Coalescence Time, t_c (s)	17523096.4	17523090	6.4	–
Polarization Angle, ψ	0.826	0.668	0.158	–
Inclination Angle, ι	2.846	2.313	0.533	–
Initial Phase, φ_o	1.844	1.836	0.048	–
Luminosity Distance, D_L (Gpc)	36.3	26.6	9.6	0.27

Table 8.1: Table showing the results of our analysis on a training dataset.

Parameter	True Value	Our Value	Error	Fract. Error
Chirp Mass, \mathcal{M} (M_\odot)	2.6832×10^6	2.6904×10^6	7178.8	2.68×10^{-3}
Symmetric Mass Ratio, η	0.2159	0.2316	0.0158	–
Ecliptic Latitude, β	1.139	-0.235	1.374	–
Ecliptic Longitude, λ	3.931	3.382	0.549	–
Coalescence Time, t_c (s)	15045887.8	15046429.6	541.2	–
Polarization Angle, ψ	6.063	5.941	0.123	–
Inclination Angle, ι	1.939	1.252	0.687	–
Initial Phase, φ_o	0.213	1.031	0.818	–
Luminosity Distance, D_L (Gpc)	10.7	26.0	15.3	1.43

Table 8.2: Table showing the results of our analysis on the official challenge dataset.

Challenge results are shown in Table 8.2. It is interesting to compare our results to those obtained by other groups applying different methods to search for SMBBH coalescences in the Mock LISA Data Challenge [173, 175, 176, 177, 178, 179, 118]. It is clear that our Challenge results are substantially less accurate, for reasons described above. However, our results from the training data set are comparable to those obtained using other methods. In particular, it is gratifying to see that we were able to obtain the correct sky location. Furthermore, the sky location is recovered to within a few square degrees, which is the accuracy required to make an optical followup feasible (see for example [185]).

8.3 Summary and future plans

We have presented a hierarchical, template-based search method for SMBBH in LISA data. This method makes use of the \mathcal{F} -statistic to reduce the parameter space for non-spinning black holes from nine to five dimensions, and then employs a stochastically generated template bank to search over the remaining parameter space. This method has been applied to perform a search on the data from Challenge 1B of the MLDC. We were able to successfully locate the signal and, in the case of the training data, recover its parameters with good accuracy.

In the future, we hope to continue our participation in the mock LISA data challenges. Challenges 3 and 4 have already been started and include SMBBH data sets where spin effects have been included in the waveform [176]. In order to participate, we must develop an analysis technique to be able to search for inspiralling supermassive black holes with spin. Initially, we want to investigate how effectively we are able to search for spinning binaries with non-spinning templates and see if this approach might enable us to get a good estimate of the masses and coalescence time of the binaries. However, to obtain good parameter estimates, we will need to incorporate the effects of spin into our signal model. Unfortunately the \mathcal{F} -statistic is not directly applicable due to the added complications spinning binaries bring. We will either have to develop a new technique to analytically maximize over some of the parameters, adapt some of the techniques discussed in chapter 6, or be forced to place templates in a much higher dimensional signal manifold.

Chapter 9

Black hole hunter

The fast developing field of gravitational wave astronomy provides an ideal opportunity to inspire public interest and excitement in science. There is a large outreach effort in the gravitational wave community, including public education centres [186], teaching projects in schools, and a travelling gravitational waves exhibit [187]. In 2008, a collaboration of gravitational wave physics groups from across the U.K. and Germany presented “Can you hear black holes collide?” [188] at the Royal Society Summer Exhibition [189]. As part of the exhibit, we developed Black Hole Hunter to provide insight into how searches for gravitational waves are performed. The game was enjoyed by the attendees at the exhibition and has subsequently been incorporated in many further outreach projects.

9.1 “Can you hear black holes collide?”

The Royal Society annually hosts a summer science exhibition at its offices in central London. This exhibition, which is open to the general public, aims to inform visitors of the latest developments and discoveries in all fields of science and inspire young people’s interest and involvement in science. The Royal Society Summer Exhibition 2008 [189] consisted of twenty-three science, two art and two history of science exhibits. These exhibits covered a vast range of scientific fields from bioscience to astrophysics, and the exhibition was attended by several thousand visitors over four days.

Among the exhibits selected for the summer exhibition in 2008 was “Can you hear black holes collide?” presented by a collaboration of British and German gravitational wave researchers. Detectors such as LIGO and GEO, are sensitive to gravitational waves with frequencies between approximately 50Hz and a few thousand Hertz. This range is comparable to the frequency range of the human



Figure 9.1: People enjoying the Black Hole Hunter game and the “Can you hear black holes collide” exhibition, London 2008.

ear, motivating the choice of title.

The goal of this exhibit was two-fold: to give the public an idea of what gravitational waves are; and how we go about searching for them. The exhibition featured a looping 5-minute long video to attract visitors. A “rubber sheet universe” was used to illustrate Einstein’s concept of space-time and curvature and to demonstrate heuristically how orbiting bodies might emit gravitational radiation. A fully-functional table-top interferometer was used to explain and demonstrate to visitors the basic principles of laser interferometric detectors. In order to illustrate the methods and challenges involved in searching for gravitational waves, the “Black Hole Hunter” game was available to play on multiple computers.

Additionally a group of researchers actively involved in gravitational wave science were stationed at the exhibit to talk to visitors and to answer their questions. A variety of handouts were distributed which provided visitors with some additional information about gravitational wave astronomy and links to relevant web pages. Several images of the exhibition and the Black Hole Hunter setup are shown in Figure 9.1.

9.2 The Black Hole Hunter game

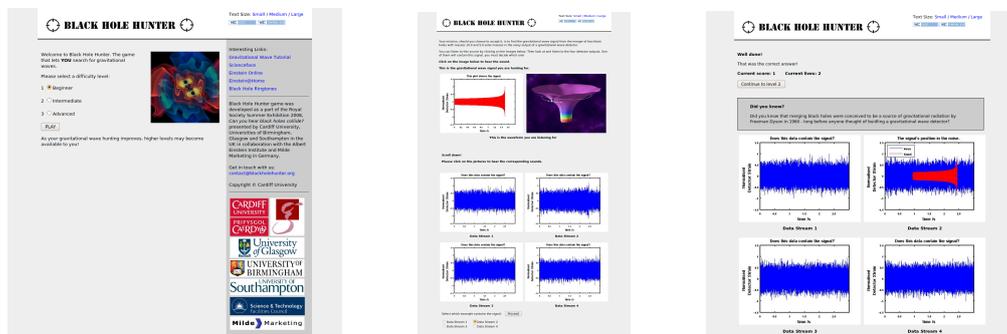


Figure 9.2: Screenshots of the Black Hole Hunter game. The first screenshot shows the welcome screen. The second shows the gravitational wave signal at the top and the four data streams below, one of which has the signal embedded in it. The third screenshot displays the four data streams again, this time with the gravitational wave signal clearly visible. The third page also has a “Did you know” factoid about gravitational waves and general relativity.

The aim of the Black Hole Hunter game is to give the player insight into the various techniques used, and challenges faced, in the search for gravitational waves. There are many potential sources of gravitational waves, and for the game we chose to focus on those emitted during the merger of two black holes or neutron stars.

These systems produce a characteristic “chirp” waveform which sweeps upwards in both frequency and amplitude as the stars draw closer to merger.

The game begins by showing the player a gravitational wave signal from a binary merger, as a time-domain waveform, and also what it sounds like, by allowing the player to play a short sound clip.¹ The player is then told that they must “detect” this gravitational wave signal. Once the player has listened to the signal he/she is presented with four graphs, and corresponding audio clips, of simulated data output from a gravitational wave detector, one of which contains the signal. The signal to noise ratio (SNR), which determines the relative amplitudes of the signal and the simulated detector noise, varies depending on the difficulty level. The idea is that the player must work in a similar way to real search algorithms and match the gravitational wave signal to what they can see or hear in the noisy data. Interestingly, it is much easier to pick out a signal by listening than by looking at the plots.

Once the player has decided which of the four data streams they think contains the signal he/she gives an answer and the game reveals whether they are correct by showing which of the data streams contained the signal and the position of the signal in the noise. If the chosen answer is correct the player will proceed on to the next level where the SNR will be lower, and thus harder to find. If the wrong answer is given the player will be able to try again with a different signal at the same SNR. This repeats until the player runs out of “lives” or reaches the furthest level. The player can choose between *beginner*, *intermediate* or *advanced* at the start of the game, which adjusts the SNR of the first and hardest levels accordingly.

To demonstrate some of the problems faced in real gravitational wave data analysis (and to make the game more fun), the hardest levels also contain “glitches” in some of the simulated data. The glitches are designed to confuse the player. They are either short sine waves of random frequency with Gaussian envelopes or *other* simulated gravitational waves that are similar to the signal, but shorter in duration. The hardest levels contain simulated data with several glitches of both kinds!

As well as giving a basic demonstration of the problems data analysts face in searching for gravitational wave sources, the Black Hole Hunter game aims to teach the player more about gravitational physics in general. This is achieved in two ways during the game. Firstly, the home page and the “Game Over” pages of Black

¹Although the signal frequencies are within human hearing range they were in fact shifted to higher pitches because typical laptop speakers and headphones were not deemed adequate at low frequencies.

Hole Hunter both have an information bar on the right hand side, which contains links to a variety of pages where the player can find out more about gravitational wave physics, and even actively participate in real gravitational wave research through the Einstein@Home project [190]. Secondly, when the player has given their answer he/she is presented with a prominent “Did you know?” box. The box contains a snippet of information about gravitational physics and an associated internet link leading to more information. There is a set of nearly 100 of these snippets, so it is unlikely that a player will encounter the same information twice in one game. Figure 9.2 shows screenshots of various parts of the game.

In addition to the website the Black Hole Hunter game has been modified to run on a local machine without requiring access to the internet. This version is available in German as well as English.

9.3 Downloadable ringtones

In addition to the game itself, the Black Hole Hunter website also gives players the opportunity to download gravitational wave ringtones. These consist of short snippets of sound or music in WAV and MP3 format which are suitable for use as a ringtone on a mobile phone. The ringtones themselves were produced by manipulating sound files generated from the expected gravitational wave signals of a variety of sources. The manipulations included significant editing, pitch shifting, layering signals on top of each other, and applying a number of audio effects. These processes were performed using audio editing software such as Cubase[191], LMMS [192] and Audacity [193].

9.4 Response to the Black Hole Hunter game

Following its launch at the 2008 exhibition, Black Hole Hunter was featured in a New Scientist blog [194] and linked from the Einstein@Home web site [190]. With this publicity, in the first month the site received 3123 unique visitors (IP addresses) from at least 66 different countries. Since then, popularity has stabilised at about 900 unique visitors per month in 2009 and around 800 unique visitors per month in 2010.

Black Hole Hunter has been used in exhibitions in the U.K. and Germany, as a teaching aid in Australia and is forming a major part of a travelling exhibition in the U.S.A. [187]. Visitors to these exhibitions typically include school teachers, school children and their parents. By appealing to their existing interests in

computer games and mobile phones and presenting science in a non-traditional manner both the game and ringtones have been commended as progressive methods of promoting science to schoolchildren.

Chapter 10

Concluding remarks

In this work we have described many aspects of the search for compact binary coalescences using interferometric detectors. We have discussed the current non-spinning LIGO-Virgo all-sky, all-time search, which has been used to search for CBC signals in S5 and VSR1 and is being used to search for CBC signals in S6 and VSR2 and VSR3. We have also described the coincident triggered search, that has been used to search for CBC signals in coincidence with electromagnetic observations in S5 and VSR1. We have presented the latest results from these searches, which have, as of yet, detected no gravitational waves. Nevertheless, in the next five years LIGO and Virgo will begin taking data from the “advanced” detectors. These are expected to increase sensitivity by a factor of 10 over the current instruments. With this increase in sensitivity these searches will be capable of regularly observing CBC systems.

We have presented the coherent search, which is an alternative method to the coincidence search in which the data from all detectors is combined coherently and events are searched for in this coherent data. We have demonstrated that the coherent search will offer an improvement in sensitivity when compared to the coincidence method. We have applied the coherent search to searches for CBC signals triggered by electromagnetic observations and demonstrated its performance on a stretch of LIGO S4 data. We hope that this technique will be used to search for CBC signals coincident with short GRBs during S6 and VSR2 and VSR3. Additionally, it would require minor additions to this triggered coherent search method to allow it to be used as an all-sky, all-time coherent search. The ability to search over multiple sky points and the ability to do time sliding must be added. We plan to implement these changes in the near future, however it is not clear how much of a computational burden this will become as SNR must be calculated for every time point, sky location and time slide.

We have discussed how we might search for CBC signals using templates that incorporate spin effects. We have described a method which would allow non-precessing templates to be used in the current search implementation. To do this the problem of template placement over the additional spin parameters must be solved and the definition of coincidence must be extended to cover the additional parameters. We have discussed the single spin PTF search method and described how this is implemented as a coincidence search. We have also presented an implementation of a coherent PTF search method, and demonstrated that it would provide an improvement in sensitivity over the coincident PTF search when more than two detectors are operating. As no searches for CBC signals using spinning templates has yet been performed on data more recent than S4 we feel it is vital that such a search is run on the most recent LIGO and Virgo data.

Finally we have presented a “stochastic” algorithm to place efficient template banks in arbitrary parameter spaces of arbitrary dimension. We have demonstrated the performance of this algorithm in some simple examples and shown that its performance when placing a two dimensional template bank for the non spinning CBC search is not significantly worse than the current geometrical algorithm. However, this algorithm would be most useful in placing template banks in higher dimension parameter spaces, such as that of searches that incorporate spin effects.

It is now a very exciting time for everyone involved in gravitational wave astronomy as we wait for the next generation of detectors to be built and reach the predicted sensitivity levels. It is (hopefully!) only a matter of time now before the first direct detection of gravitational waves is made and we truly become gravitational wave astronomers.

Bibliography

- [1] B. P. Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. *Phys. Rev.*, D79:122001, 2009.
- [2] 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.
- [3] J. Abadie et al. Search for Gravitational Waves from Compact Binary Coalescence in LIGO and Virgo Data from S5 and VSR1. *Phys. Rev.*, D82:102001, 2010.
- [4] J. Abadie et al. Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO's fifth and Virgo's first science run. *Astrophys. J.*, 715:1453–1461, 2010.
- [5] B. Abbott et al. Implications for the Origin of GRB 070201 from LIGO Observations. *Astrophys. J.*, 681:1419–1428, 2008.
- [6] Ian W. Harry and Stephen Fairhurst. A targeted coherent search for gravitational waves from compact binary coalescences. *Phys. Rev.*, D83:084002, 2011.
- [7] Ian Harry and Stephen Fairhurst. A coherent triggered search for single spin compact binary coalescences in gravitational wave data. 2011.
- [8] Ian W. Harry, Bruce Allen, and B. S. Sathyaprakash. A stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev.*, D80:104014, 2009.
- [9] I. W. Harry, S. Fairhurst, and B. S. Sathyaprakash. A hierarchical search for gravitational waves from supermassive black hole binary mergers. *Class. Quant. Grav.*, 25:184027, 2008.

- [10] Albert Einstein. The Formal Foundation of the General Theory of Relativity. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1914:1030–1085, 1914.
- [11] Albert Einstein. On the General Theory of Relativity. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1915:778–786, 1915.
- [12] Albert Einstein. The Field Equations of Gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1915:844–847, 1915.
- [13] Peter Coles. Einstein, Eddington and the 1919 Eclipse. 2001.
- [14] B. Abbott et al. Ligo: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, 72(7):076901, 2009.
- [15] F. Acernese et al. The virgo status. *Class. Quant. Grav.*, 23:S635–S642, 2006.
- [16] Benno Willke. GEO600: Status and plans. *Class. Quant. Grav.*, 24:S389–S397, 2007.
- [17] B. P. Abbott et al. An Upper Limit on the Stochastic Gravitational-Wave Background of Cosmological Origin. *Nature*, 460:990, 2009.
- [18] J. Abadie et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys. Rev.*, D81:102001, 2010.
- [19] B. Abbott et al. Searches for gravitational waves from known pulsars with S5 LIGO data. *Astrophys. J.*, 713:671–685, 2010.
- [20] Gregory M. Harry. Advanced LIGO: The next generation of gravitational wave detectors. *Class. Quant. Grav.*, 27:084006, 2010.
- [21] Advanced LIGO. <http://www.ligo.caltech.edu/advLIGO/>.
- [22] Advanced Virgo. <http://wwwcascina.virgo.infn.it/advirgo/>.
- [23] J. Abadie et al. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *Class. Quant. Grav.*, 27:173001, 2010.
- [24] LCGT: Large Scale Cryogenic Gravitational Wave Telescope. <http://gw.icrr.u-tokyo.ac.jp/lcgt/>.

- [25] AIGO: Australian International Gravitational Observatory. <http://www.aigo.org.au/>.
- [26] Stefan Hild, Simon Chelkowski, and Andreas Freise. Pushing towards the ET sensitivity using 'conventional' technology. 2008.
- [27] K Danzmann. LISA - an ESA cornerstone mission for a gravitational wave observatory. *Class. Quantum Grav.*, 14:1399, 1997.
- [28] B. P. Abbott et al. Search for gravitational-wave bursts associated with gamma-ray bursts using data from LIGO Science Run 5 and Virgo Science Run 1. *Astrophys. J.*, 715:1438–1452, 2010.
- [29] Jonah Kanner et al. LOOC UP: Locating and observing optical counterparts to gravitational wave bursts. *Class. Quant. Grav.*, 25:184034, 2008.
- [30] Ehud Nakar. Short-hard gamma-ray bursts. *Phys. Rept.*, 442:166–236, 2007.
- [31] Masaru Shibata and Keisuke Taniguchi. Merger of black hole and neutron star in general relativity: Tidal disruption, torus mass, and gravitational waves. *Phys. Rev.*, D77:084015, 2008.
- [32] B. D. Metzger et al. Electromagnetic Counterparts of Compact Object Mergers Powered by the Radioactive Decay of R-process Nuclei. 2010.
- [33] V. Predoi et al. Prospects for joint radio telescope and gravitational wave searches for astrophysical transients. *Class. Quant. Grav.*, 27:084018, 2010.
- [34] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Freeman, San Francisco, 1973.
- [35] Bernard F. Schutz. *A first course in general relativity*. Cambridge university press, 1985.
- [36] Michele Maggiore. *Gravitational Waves*. Oxford university press, 2008.
- [37] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, Singapore, 1994.
- [38] Dave McKechnan. *On the use of higher order waveforms in the search for gravitational waves emitted by compact binary coalescences*. PhD thesis, Cardiff University, 2010.

- [39] Luc Blanchet, Thibault Damour, Gilles Esposito-Farese, and Bala R. Iyer. Gravitational radiation from inspiralling compact binaries completed at the third post-Newtonian order. *Phys. Rev. Lett.*, 93:091101, 2004.
- [40] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Binary black hole merger dynamics and waveforms. *Phys. Rev.*, D73:104002, 2006.
- [41] Christian D. Ott, Adam Burrows, Luc Dessart, and Eli Livne. A new mechanism for gravitational wave emission in core-collapse supernovae. *Phys. Rev. Lett.*, 96:201102, 2006.
- [42] Adam Burrows, Eli Livne, Luc Dessart, Christian Ott, and Jeremiah Murphy. A New Mechanism for Core-Collapse Supernova Explosions. *Astrophys. J.*, 640:878–890, 2006.
- [43] B. P. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev.*, D80:102001, 2009.
- [44] B. P. Abbott et al. Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm. *Astrophys. J.*, 701:L68–L74, 2009.
- [45] B. P. Abbott et al. Search for High Frequency Gravitational Wave Bursts in the First Calendar Year of LIGO’s Fifth Science Run. *Phys. Rev.*, D80:102002, 2009.
- [46] Piotr Jaranowski, Andrzej Krolak, and Bernard F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars. I: The signal and its detection. *Phys. Rev.*, D58:063001, 1998.
- [47] B. Abbott et al. Beating the spin-down limit on gravitational wave emission from the Crab pulsar. *Astrophys. J.*, 683:L45–L50, 2008.
- [48] B. P. Abbott et al. Einstein@Home search for periodic gravitational waves in early S5 LIGO data. *Phys. Rev.*, D80:042003, 2009.
- [49] B. Abbott et al. All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data. *Phys. Rev. Lett.*, 102:111102, 2009.
- [50] Bruce Allen. The stochastic gravity-wave background: Sources and detection. *gr-qc/9604033*, 1996.

- [51] B. Abbott et al. Searching for a stochastic background of gravitational waves with LIGO. *Astrophys. J.*, 659:918–930, 2007.
- [52] B. Abbott et al. Analysis of first LIGO science data for stochastic gravitational waves. *Phys. Rev.*, D69:122004, 2004.
- [53] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophys. J.*, 195:L51–L53, 1975.
- [54] Joel M. Weisberg and Joseph H. Taylor. Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis. 2004.
- [55] I. H. Stairs, S. E. Thorsett, J. H. Taylor, and A. Wolszczan. Studies of the Relativistic Binary Pulsar PSR B1534+12: I. Timing Analysis. *Astrophys. J.*, 581:501–508, 2002.
- [56] M. Kramer et al. Tests of general relativity from timing the double pulsar. *Science*, 314:97–102, 2006.
- [57] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. 2005.
- [58] <http://commons.wikimedia.org/wiki/File:LIGOsensitivity.png>.
- [59] https://www.cascina.virgo.infn.it/DataAnalysis/Calibration/Sensitivity/LigoVirgo/LIGO_VIRGO_VSR1.gif.
- [60] Space time quest game. http://www.gwoptics.org/processing/space_time_quest/.
- [61] GEO website. <http://www.geo600.org>.
- [62] Einstein Telescope. <http://www.et-gw.eu>.
- [63] <http://www.srl.caltech.edu/~shane/sensitivity/IMG/LISApositionPSD.jpg>.
- [64] P. C. Peters and J. Mathews. Gravitational radiation from point masses in a Keplerian orbit. *Phys. Rev.*, 131:435–439, 1963.
- [65] Duncan A. Brown and Peter J. Zimmerman. The Effect of Eccentricity on Searches for Gravitational- Waves from Coalescing Compact Binaries in Ground-based Detectors. *Phys. Rev.*, D81:024007, 2010.

- [66] T. Cokelaer and D. Pathak. Searching for gravitational-wave signals emitted by eccentric compact binaries using a non-eccentric template bank: implications for ground-based detectors. *Class. Quant. Grav.*, 26:045013, 2009.
- [67] Duncan A. Brown. *Searching for gravitational radiation from black hole machos in the galactic halo*. PhD thesis, University of Wisconsin-Milwaukee, December 2004.
- [68] Alessandra Buonanno, Bala Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev.*, D80:084043, 2009.
- [69] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 5:3, 2002.
- [70] Joan M. Centrella, John G. Baker, Bernard J. Kelly, and James R. van Meter. Black-hole binaries, gravitational waves, and numerical relativity. *Rev. Mod. Phys.*, 82, 2010.
- [71] Mark Hannam. Status of black-hole-binary simulations for gravitational-wave detection. *Class. Quant. Grav.*, 26:114001, 2009.
- [72] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. 2011.
- [73] B. Abbott et al. Analysis of LIGO data for gravitational waves from binary neutron stars. *Phys. Rev.*, D69:122001, 2004.
- [74] B. Abbott et al. Search for gravitational waves from galactic and extragalactic binary neutron stars. *Phys. Rev.*, D72:082001, 2005.
- [75] B. Abbott et al. Search for gravitational waves from primordial black hole binary coalescences in the galactic halo. *Phys. Rev.*, D72:082002, 2005.
- [76] B. Abbott et al. Search for gravitational waves from binary black hole inspirals in ligo data. *Phys. Rev.*, D73:062001, 2006.
- [77] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Phys. Rev.*, D77:062002, 2008.
- [78] L A Wainstein and V D Zubakov. *Extraction of Signals from Noise*. Prentice-Hall, Englewood Cliffs, 1962.

- [79] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev.*, D44:3819–3834, 1991.
- [80] Bruce Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev.*, D71:062001, 2005.
- [81] J. Abadie et al. Calibration of the LIGO Gravitational Wave Detectors in the Fifth Science Run. *Nucl. Instrum. Meth.*, A624:223–240, 2010.
- [82] T. Accadia et al. Calibration and sensitivity of the Virgo detector during its second science run. *Class. Quant. Grav.*, 28:025005, 2011.
- [83] Thomas Cokelaer. Gravitational waves from inspiralling compact binaries: hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev.*, D76:102004, 2007.
- [84] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash. A template bank to search for gravitational waves from inspiralling compact binaries. I: Physical models. *Class. Quant. Grav.*, 23:5477–5504, 2006.
- [85] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev.*, D60:022002, 1999.
- [86] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev.*, D53:6749–6761, 1996.
- [87] R Balasubramanian, B S Sathyaprakash, and S V Dhurandhar. Gravitational waves from coalescing binaries: detection strategies and monte carlo estimation of parameters. *Phys. Rev. D*, 53:3033, 1996. Erratum-ibid. D **54**, 1860 (1996).
- [88] S. V. Dhurandhar and B. S. Sathyaprakash. Choice of filters for the detection of gravitational waves from coalescing binaries. 2. detection in colored noise. *Phys. Rev.*, D49:1707–1722, 1994.
- [89] FFTW - Fastest Fourier Transform in the West. <http://www.fftw.org/>.
- [90] C. A. K. Robinson, B. S. Sathyaprakash, and Anand S. Sengupta. A geometric algorithm for efficient coincident detection of gravitational waves. *Phys. Rev.*, D78:062002, 2008.

- [91] Nelson Christensen, Peter Shawhan, and Gabriela Gonzalez. Vetoes for inspiral triggers in LIGO data. *Class. Quant. Grav.*, 21:S1747–S1756, 2004.
- [92] Nelson Christensen. Veto Studies for LIGO Inspiral Triggers. *Class. Quant. Grav.*, 22:S1059–S1068, 2005.
- [93] L Blackburn et al. The LSC Glitch Group : Monitoring Noise Transients during the fifth LIGO Science Run. *Class. Quant. Grav.*, 25:184004, 2008.
- [94] J. Slutsky et al. Methods for Reducing False Alarms in Searches for Compact Binary Coalescences in LIGO Data. *Class. Quant. Grav.*, 27:165023, 2010.
- [95] T. Accadia et al. Noise from scattered light in Virgo’s second science run data. *Class. Quant. Grav.*, 27:194011, 2010.
- [96] T. Accadia et al. Tools for noise characterization in Virgo. *J. Phys. Conf. Ser.*, 243:012004, 2010.
- [97] Andres Rodríguez. Reducing false alarms in searches for gravitational waves from coalescing binary systems. Master’s thesis, Louisiana State University, 2007.
- [98] J. Abadie et al. An all sky search for Compact Binary Coalescences in LIGO’s 6th and Virgo’s 2nd and 3rd Science runs. *In preparation*, 2011.
- [99] Stephen Fairhurst and Patrick Brady. Interpreting the results of searches for gravitational waves from coalescing binaries. *Class. Quant. Grav.*, 25:105002, 2008.
- [100] Rahul Biswas, Patrick R. Brady, Jolien D. E. Creighton, and Stephen Fairhurst. The Loudest Event Statistic: General Formulation, Properties and Applications. *Class. Quant. Grav.*, 26:175009, 2009.
- [101] Drew Keppel. *Signatures and Dynamics of Compact Binary Coalescences and a Search in LIGO S5 Data*. PhD thesis, California Institute of Technology, 2009.
- [102] E. S. Phinney. The rate of neutron star binary mergers in the universe: Minimal predictions for gravity wave detector. *Astrophys. J.*, 380:L17, 1991.
- [103] Ravi Kumar Kopparapu et al. Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events. 2007.

- [104] Joshua S. Bloom et al. Astro2010 Decadal Survey Whitepaper: Coordinated Science in the Gravitational and Electromagnetic Skies. 2009.
- [105] Peter Meszaros. Gamma-Ray Bursts. *Rept. Prog. Phys.*, 69:2259–2322, 2006.
- [106] Neil Gehrels et al. Swift detects a remarkable gamma-ray burst, GRB 060614, that introduces a new classification scheme. *Nature*, 444:1044, 2006.
- [107] E. P. Mazets et al. A Giant Flare from a Soft Gamma Repeater in the Andromeda Galaxy, M31. 2007.
- [108] E. O. Ofek et al. GRB 070201: A possible Soft Gamma Ray Repeater in M31. 2007.
- [109] Archana Pai, Sanjeev Dhurandhar, and Sukanta Bose. A data-analysis strategy for detecting gravitational-wave signals from inspiraling compact binaries with a network of laser-interferometric detectors. *Phys. Rev.*, D64:042004, 2001.
- [110] S. Klimenko, I. Yakushin, A. Mercer, and Guenakh Mitselmakher. Coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 25:114029, 2008.
- [111] Patrick J. Sutton et al. X-Pipeline: An analysis package for autonomous gravitational-wave burst searches. *New J. Phys.*, 12:053034, 2010.
- [112] Y. Guersel and M. Tinto. Near optimal solution to the inverse problem for gravitational wave bursts. *Phys. Rev.*, D40:3884–3938, 1989.
- [113] Chad. Hanna. *Searching for gravitational waves from binary systems in non-stationary data*. PhD thesis, Louisiana State University, 2008.
- [114] Sukanta Bose, Archana Pai, and Sanjeev V. Dhurandhar. Detection of gravitational waves from inspiraling compact binaries using a network of interferometric detectors. *Int. J. Mod. Phys.*, D9:325–329, 2000.
- [115] Sukanta Bose, Sanjeev V. Dhurandhar, and Archana Pai. Detection of gravitational waves using a network of detectors. *Pramana*, 53:1125–1136, 1999.
- [116] Sukanta Bose, Thilina Dayanga, Shaon Ghosh, and Dipongkar Talukder. A blind heirarchical coherent search for gravitational-wave signals from coalescing compact binaries in a network of interferometric detectors. *In preparation*, 2011.

- [117] Curt Cutler and Bernard F. Schutz. The generalized F-statistic: Multiple detectors and multiple GW pulsars. *Phys. Rev.*, D72:063006, 2005.
- [118] Neil J. Cornish and Edward K. Porter. The search for supermassive black hole binaries with LISA. *Class. Quant. Grav.*, 24:5729–5755, 2007.
- [119] Jeffrey C. Livas. *Upper limits for gravitational radiation from some astrophysical sources*. PhD thesis, Massachusetts Institute of Technology, 1987.
- [120] Reinhard Prix. Search for continuous gravitational waves: Metric of the multi-detector F-statistic. *Phys. Rev.*, D75:023004, 2007.
- [121] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. A comparison of search templates for gravitational waves from binary inspiral. *Phys. Rev.*, D63:044023, 2001.
- [122] S. Klimenko, S. Mohanty, Malik Rakhmanov, and Guenakh Mitselmakher. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev.*, D72:122002, 2005.
- [123] S. Klimenko, S. Mohanty, Malik Rakhmanov, and Guenakh Mitselmakher. Constraint likelihood method: Generalization for colored noise. *J. Phys. Conf. Ser.*, 32:12–17, 2006.
- [124] Reinhard Prix and Badri Krishnan. Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics. *Class. Quant. Grav.*, 26:204013, 2009.
- [125] Mark Hannam et al. The Samurai Project: verifying the consistency of black-hole-binary waveforms for gravitational-wave detection. *Phys. Rev.*, D79:084025, 2009.
- [126] 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.
- [127] Diego Fazi. *Development of a physical-template search for gravitational waves from spinning compact-object binaries with LIGO*. PhD thesis, Università di Bologna, 2009.
- [128] B. Abbott et al. Tuning matched filter searches for compact binary coalescence. Technical Report LIGO-T070109-01, 2007.

- [129] Y. Itoh, M. A. Papa, B. Krishnan, and X. Siemens. Chi-square test on candidate events from CW signal coherent searches. *Class. Quant. Grav.*, 21:S1667–S1678, 2004.
- [130] LSC Algorithm Library Suite. <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [131] Melvyn B. Davies, Andrew J. Levan, and Andrew R. King. The ultimate outcome of black hole - neutron star mergers. *Mon. Not. Roy. Astron. Soc.*, 356:54–58, 2005.
- [132] B. P. Abbott et al. Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data. *Phys. Rev.*, D80:062001, 2009.
- [133] Lisa Maria Goggin. A Search For Gravitational Waves from Perturbed Black Hole Ringdowns in LIGO Data. 2008.
- [134] 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.*, D69:104017, 2004. Erratum-ibid. **D74**, 029905(E) (2006).
- [135] J. Abadie et al. A triggered search for Gamma Ray Burst gravitational wave counterparts in LIGO’s 6th and Virgo’s 2nd and 3rd Science runs. *In preparation*, 2011.
- [136] James Clark, Ik Siong Heng, Matthew Pitkin, and Graham Woan. An evidence based search for gravitational waves from neutron star ring-downs. *Phys. Rev.*, D76:043003, 2007.
- [137] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys. Rev.*, D81:062003, 2010.
- [138] Tyson B. Littenberg and Neil J. Cornish. Separating Gravitational Wave Signals from Instrument Artifacts. *Phys. Rev.*, D82:103007, 2010.
- [139] Fermi Gamma-ray space telescope. http://www.nasa.gov/mission_pages/GLAST/science/index.html.
- [140] The Interplanetary Network Progress Report. <http://ipnpr.jpl.nasa.gov/index.cfm>.

- [141] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev.*, D49:6274–6297, 1994.
- [142] Philippe Grandclement, Vassiliki Kalogera, and Alberto Vecchio. Searching for gravitational waves from the inspiral of precessing binary systems. I: Reduction of detection efficiency. *Phys. Rev.*, D67:042003, 2003.
- [143] Kip S. Thorne. Gravitational radiation. In S.W. Hawking and W. Israel, editors, *Three hundred years of gravitation*, pages 330–458. Cambridge University Press, 1987.
- [144] P. Ajith et al. Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins. 2009.
- [145] 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.
- [146] A. Buonanno, Y. Chen, and M. Vallisneri. Detection template families for precessing binaries of spinning compact binaries: Adiabatic limit. *Phys. Rev. D*, 67:104025, 2003. Erratum-ibid. **D** 74, 029904(E) (2006).
- [147] Alessandra Buonanno, Yan-bei Chen, Yi Pan, and Michele Vallisneri. A quasi-physical family of gravity-wave templates for precessing binaries of spinning compact objects: II. Application to double-spin precessing binaries. *Phys. Rev.*, D70:104003, 2004.
- [148] Lawrence E. Kidder, Clifford M. Will, and Alan G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. *Phys. Rev.*, D47:4183, 1993.
- [149] Lawrence E. Kidder. Coalescing binary systems of compact objects to post-Newtonian 5/2 order. 5. Spin effects. *Phys. Rev.*, D52:821–847, 1995.
- [150] Guillaume Faye, Luc Blanchet, and Alessandra Buonanno. Higher-order spin effects in the dynamics of compact binaries. i: Equations of motion. *Phys. Rev.*, D **74**:104033, 2006.
- [151] K. G. Arun, Alessandra Buonanno, Guillaume Faye, and Evan Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev.*, D79:104023, 2009.

- [152] Yi Pan et al. Effective-one-body waveforms calibrated to numerical relativity simulations: coalescence of non-precessing, spinning, equal-mass black holes. *Phys. Rev.*, D81:084041, 2010.
- [153] B.F. Schutz. Data processing, analysis, and storage for interferometric antennas. In D Blair, editor, *The detection of gravitational waves*, pages 406–452, England, 1989. Cambridge University Press.
- [154] B. S. Sathyaprakash. Filtering postnewtonian gravitational waves from coalescing binaries. *Phys. Rev.*, D50:7111–7115, 1994.
- [155] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Estimation of parameters of gravitational waves from coalescing binaries. *Pramana*, 45:L463, 1995.
- [156] Patrick R. Brady, Teviet Creighton, Curt Cutler, and Bernard F. Schutz. Searching for periodic sources with ligo. *Phys. Rev. D*, 57:2101–2116, 1998.
- [157] Reinhard Prix. Template-based searches for gravitational waves: efficient lattice covering of flat parameter spaces. *Classical and Quantum Gravity*, 24(19):S481–S490, 2007.
- [158] C. Messenger, R. Prix, and M. A. Papa. Random template banks and relaxed lattice coverings. *Phys. Rev.*, D79:104017, 2009.
- [159] J.H. Conway and N.J.A. Sloane. *Sphere Packings, Lattices and Groups*. Springer-Verlag, New York, 2nd edition edition, 1993.
- [160] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet. Gravitational-wave inspiral of compact binary systems to 7/2 post-newtonian order. *Phys. Rev. D*, 65:061501(R), 2002. Erratum-ibid **71**, 129902(E) (2005).
- [161] K G Arun, Luc Blanchet, Bala R. Iyer, and Moh'd S S. Qusailah. The 2.5pn gravitational wave polarisations from inspiralling compact binaries in circular orbits. *Class. Quantum Grav.*, 21:3771, 2004. Erratum-ibid. **22**, 3115 (2005).
- [162] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.
- [163] Stanislav Babak. Building a stochastic template bank for detecting massive black hole binaries. *Class. Quant. Grav.*, 25:195011, 2008.

- [164] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. A comparison of search templates for gravitational waves from binary inspiral. *Phys. Rev. D*, 63:044023, 2001. Erratum-ibid. **D** 72 (2005) 029902.
- [165] M. J. Valtonen and Others. A massive binary black-hole system in OJ 287 and a test of general relativity. *Nature*, 452:851–853, 2008.
- [166] S. Komossa et al. Discovery of a binary agn in the ultraluminous infrared galaxy ngc 6240 using chandra. *Astrophys. J.* , 582:L15–L20, 2003.
- [167] P. L. Bender. Lisa: Laser interferometer space antenna for the detection and observation of gravitational waves: Pre-phase: A report, 1995. Unpublished.
- [168] Martin J. Rees and Marta Volonteri. Massive Black Holes: formation and evolution. *Proceedings of the International Astronomical Union*, 2:51–55, 2007.
- [169] Alberto Sesana, Marta Volonteri, and Francesco Haardt. The imprint of massive black hole formation models on the LISA data stream. *Mon. Not. Roy. Astron. Soc.*, 377:1711–1716, 2007.
- [170] Marta Volonteri, Francesco Haardt, and Piero Madau. The Assembly and Merging History of Supermassive Black Holes in Hierarchical Models of Galaxy Formation. *Astrophys. J.*, 582:559–573, 2003.
- [171] Curt Cutler and Michele Vallisneri. LISA detections of massive black hole inspirals: parameter extraction errors due to inaccurate template waveforms. *Phys. Rev.*, D76:104018, 2007.
- [172] LISA Mission Science Office. LISA: probing the Universe with gravitational waves. www.lisascience.org/resources/talks-articles/science, 2007.
- [173] K. A. Arnaud et al. Report on the first round of the Mock LISA Data Challenges. *Class. Quant. Grav.*, 24:S529–S540, 2007.
- [174] K. A. Arnaud and Others. An overview of the second round of the Mock LISA Data Challenges. *Class. Quant. Grav.*, 24:S551–S564, 2007.
- [175] Stanislav Babak and Others. Report on the second Mock LISA Data Challenge. *Class. Quant. Grav.*, 25:114037, 2007.
- [176] Stanislav Babak et al. The Mock LISA Data Challenges: from Challenge 1B to Challenge 3. *Class. Quant. Grav.*, 25:184026, 2008.

- [177] Neil J. Cornish and Edward K. Porter. Searching for massive black hole binaries in the first Mock LISA Data Challenge. *Class. Quant. Grav.*, 24:S501–S512, 2007.
- [178] Duncan A. Brown, Jeff Crowder, Curt Cutler, Ilya Mandel, and Michele Vallisneri. A Three-Stage Search for Supermassive Black Hole Binaries in LISA Data. *Class. Quant. Grav.*, 24:S595–S606, 2007.
- [179] Jordan B. Camp, John K. Cannizzo, and Kenji Numata. Application of the Hilbert-Huang transform to the search for gravitational waves. *Phys. Rev.*, D75:061101, 2007.
- [180] L.S. Finn and D.F. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. *Phys. Rev. D*, 47:2198–2219, 1993.
- [181] John T. Whelan, Reinhard Prix, and Deepak Khurana. Searching for Galactic White Dwarf Binaries in Mock LISA Data using an F-Statistic Template Bank. *Class. Quant. Grav.*, 27:055010, 2010.
- [182] Andrzej Krolak, Massimo Tinto, and Michele Vallisneri. Optimal filtering of the LISA data. *Phys. Rev.*, D70:022003, 2004.
- [183] Massimo Tinto, F. B. Estabrook, and J. W. Armstrong. Time-delay interferometry for LISA. *Phys. Rev.*, D65:082003, 2002.
- [184] Michele Vallisneri. Geometric time delay interferometry. *Phys. Rev.*, D72:042003, 2005.
- [185] J. Anthony Tyson. Large Synoptic Survey Telescope: Overview. *Proc. SPIE Int. Soc. Opt. Eng.*, 4836:10–20, 2002.
- [186] M. Cavaglia, M. Hendry, D. Ingram, S. Milde, S. R. Pandian, D. Reitze, K. Riles, B. Schutz, A. L. Stuver, T. Summerscales, D. Ugolini, J. Thacker, M. Vallisneri, and A. Zerneno. Gravitational Wave Astronomy: Opening a New Window on the Universe for Students, Educators and the Public. In *Bulletin of the American Astronomical Society*, volume 40 of *Bulletin of the American Astronomical Society*, page 231, May 2008.
- [187] M. Cavaglia, M. Hendry, S. Marka, D. Reitze, and K. Riles. Astronomy’s New Messengers: A traveling exhibit to reach out to a young adult audience.
- [188] Can you hear black holes collide web page. <http://www.soton.ac.uk/math/blackholes.html>.

- [189] Royal society summer science exhibition website. <http://www.summerscience.org.uk/>.
- [190] Einstein@home website. <http://einstein.phys.uwm.edu>.
- [191] Cubase - advanced music production system. <http://www.steinberg.net/en/products/musicproduction/cubase5>.
- [192] Linux multimedia studio. <http://lmms.sourceforge.net>.
- [193] Audacity. <http://audacity.sourceforge.net/download/linux>.
- [194] Hazel Muir. Online game makes anyone a 'black hole hunter'. *New Scientist Blog*, 2008. http://www.newscientist.com/blog/space/2008/06/online-game-makes-anyone-black-hole_30.html.