

SEARCHING FOR GRAVITATIONAL WAVES FROM COMPACT BINARY COALESCENCES

Craig A. K. Robinson

School of Physics and Astronomy, Cardiff University,
Queens Buildings, The Parade, Cardiff CF24 3AA

A thesis submitted to
Cardiff University
for the degree of
Doctor of Philosophy

December 2007

Supervisor: Prof. B.S. Sathyaprakash

UMI Number: U585054

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U585054

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Acknowledgements

I would like to thank PPARC for funding this study. I would also like to thank Prof. B.S. Sathyaprakash for providing additional funding for the last few months of the study.

I would like to thank my supervisor, B.S. Sathyaprakash, for the support and guidance over the course of the study. It has been an inspiration to work with someone who has an obvious love of the subject, and who has much insight into the problems which arise. I would also like to thank Dr. R. Philp and Dr. I.J. Taylor for the useful discussions, particularly in the early stages of the study.

I would also like to thank my collaborators, in no particular order: B.S. Sathyaprakash, A.S. Sengupta, B.R. Iyer, P.Ajith, the Triana team, and the Compact Binary Coalescence group of the LIGO Scientific Collaboration. In addition, I would like to thank the members of the Cardiff Gravitational Physics Group: Gareth, Thomas, Chris, Steve, Dave, Ian, Leonid, Alex, Edmund, Rhiannon, Jack, Gerald, Deepak and Patrick for many interesting discussions over the years, not just about science. I am also grateful to the staff in the Physics office - in particular, Louise Winter and Philip Treadgold, for much support on the administrative side of things.

On the personal side, I would like to thank bandmates past and present: Grant, Nick, Jeff, Ben, Luke, Matt, Rich, Lee and Kieron for accompanying me in various musical endeavours over the years, and probably helping me remain (very nearly!) sane throughout this study.

And finally, special thanks to my parents for the help and support through the years.

Many apologies to anybody I've forgotten!

Disclaimer

This thesis contains some results concerning the playground data from the LIGO S5 run. These results do not reflect the opinion of the LIGO Scientific Collaboration, and have not been reviewed by the Collaboration.

Summary of thesis

Firstly, we give a general introduction to gravitational waves, the instruments used to detect them, potentially interesting sources, and the basics of gravitational wave data analysis with respect to the compact binary search.

Following this, we look at a new class of approximants for inspiral waveforms. In these complete approximants, instead of truncating the binding energy and flux functions at the same post-Newtonian order, we instead keep terms such that the approximant corresponds in spirit to the dynamics of the system, with no missing terms in the acceleration. We compare the overlaps with an exact signal (in the adiabatic approximation) for a test mass orbiting a Schwarzschild black hole, for standard and complete approximants in the adiabatic approximation, and beyond the adiabatic approximation using Lagrangian models. A limited extension to the comparable mass case is also given.

We then investigate two approaches to performing inspiral searches in a time-critical manner. Both involve splitting the search parameter space across several compute nodes. The first attempts to split the parameter space in an efficient manner by using information from previous runs. The second is balanced dynamically, with slave nodes requesting work off a master node.

We then develop a new method for coincidence analysis. In this method, each trigger has associated with it an ellipsoidal region of the parameter space defined by the covariance matrix. Triggers from different detectors are deemed coincident if their ellipsoids are found to overlap. Compared to an algorithm which uses uncorrelated windows separately for each parameter, the method significantly reduces the background rate for comparable detection efficiency.

We then give a summary of the current status of the ongoing search for high mass compact binary coalescences in the first calendar year of LIGO's fifth science run.

Contents

1	Introduction	1
1.1	Gravitational waves	3
1.1.1	Gravitational waves in linearised theory	3
1.1.2	Radiation of gravitational waves	7
1.2	Detecting gravitational waves	8
1.2.1	Indirect detection: the Hulse-Taylor pulsar	9
1.2.2	Direct detection: the laser interferometer	9
1.3	Sources of gravitational radiation	12
1.4	Compact binary coalescences	15
1.5	Data analysis	17
1.5.1	Searching for signals of known form	18
2	A new class of post-Newtonian approximants for gravitational wave data analysis	27
2.1	Introduction	27
2.1.1	Standard approach to phasing formula	28
2.1.2	Complete phasing of the adiabatic inspiral: An alternative	30
2.1.3	Non-adiabatic inspiral	32
2.1.4	What this study is about	33
2.1.5	Effectualness and Faithfulness	34
2.1.6	Noise spectra of the interferometers	36
2.2	Test mass waveforms in the adiabatic approximation	36
2.2.1	The energy function	38
2.2.2	The flux function	39
2.2.3	Comparison of standard and complete adiabatic approximants	39
2.3	Non-adiabatic models	51
2.3.1	Standard and complete non-adiabatic approximants in the test mass case	55
2.4	Comparable Mass Waveforms	58
2.4.1	The energy function	58
2.4.2	The flux function	60
2.4.3	Comparable mass results in the adiabatic approximation	61
2.4.4	Comparable mass results beyond the adiabatic approximation	64
2.5	Summary and Conclusion	64
3	Developing a low-latency search for inspiralling compact binaries	68
3.1	Step-wise load balancing algorithm	72
3.1.1	How the algorithm works	72
3.1.2	Marching ahead: how we deal with failures	73
3.1.3	Benefits of using the algorithm	75
3.2	Dynamic load-balancing algorithm	75

3.2.1	Benefits of using the algorithm	77
3.3	Conclusions	78
4	A geometric algorithm for efficient coincident detection of gravitational waves	80
4.1	Motivation	80
4.1.1	The problem of coincident detection	81
4.1.2	A geometric approach to choosing coincident windows	82
4.1.3	Organization of the Chapter	83
4.2	A geometric coincidence algorithm	83
4.2.1	Scalar Product, Signal Manifold and Metric	84
4.2.2	Coincidence windows	86
4.3	Overlap of ellipsoids	91
4.4	Expected reduction in false alarm rate	95
4.5	Application to coalescing binaries	97
4.6	Summary and Conclusions	101
5	Searching for high-mass binary coalescences in LIGO S5 data	102
5.1	The S5 first calendar year	106
5.1.1	Details of the data analysed	106
5.2	Structure of the pipeline	107
5.2.1	Filtering	108
5.2.2	Coincidence analysis	109
5.2.3	Signal-based vetoes	110
5.3	Monte-Carlo simulations and background	113
5.3.1	Coincidence step	113
5.3.2	The χ^2 veto	115
5.3.3	The r^2 veto	116
5.4	Estimation of parameters	118
5.5	The S5 first-calendar year playground results	123
5.5.1	Efficiency and projected upper limit	127
5.6	Status and future plans	129
6	Conclusion	131
	Bibliography	135

Chapter 1

Introduction

In 1919, Eddington announced the results of an expedition to Principe, an island off the west coast of Africa, to view the total eclipse which took place there[1, 2]. The results had confirmed the prediction of general relativity of the bending of light rays from distant stars by the Sun. The theory had already been shown to explain the precession of the perihelion of Mercury[3], and as a consequence, Albert Einstein was propelled to international stardom. Since then, every experiment which has been carried out has confirmed the predictions of general relativity with remarkable accuracy.

One of the predictions of general relativity is the existence of gravitational radiation. In fact, gravitational radiation of some sort would be a prediction of any relativistic theory of gravity, due to the finite speed of propagation of interactions. In general relativity, gravitational waves manifest themselves as fluctuations in the curvature of spacetime. Such waves will be emitted by any object, or objects, which experience acceleration, provided the motion is not axisymmetric. Strong sources of gravitational waves include binary systems of compact objects, such as the Hulse-Taylor binary pulsar, which has been shown to behave as predicted by general relativity to a great degree of accuracy.

Many attempts have been made to detect gravitational waves over the years. However, despite the often dramatic nature of potential sources, the effects of gravitational radiation incident on the Earth from these sources is expected to be very weak. This means that the signals due to gravitational radiation will be obscured by noise, making the detection of the gravitational radiation very difficult. So far, there has been no satisfactory claim of the direct detection of

gravitational radiation. Nonetheless, detectors such as LIGO are now taking data at unprecedented sensitivities, and there is optimism that the first sources will be detected over the next few years.

This thesis documents some aspects of the ongoing search for gravitational waves; in particular, the search for gravitational wave signals from inspiralling compact binary systems, consisting of neutron stars and/or black holes. Systems of this type are one of the most promising sources for detection, with next-generation detectors such as Advanced LIGO expected to see several such events per year. The search for such signals requires much effort, with work required in theoretical aspects, computational aspects, and optimisation of the search procedure. All of these aspects will be touched upon in this thesis.

The structure of the thesis is as follows: in this Chapter, we have an introduction to the basics underlying the search for gravitational waves, including the basics of gravitational wave physics, and how we attempt to detect them; a brief review of promising gravitational wave sources, with particular emphasis on compact binary coalescences (CBCs); and the basics of the methods used to extract the signals from noisy data. The next Chapter consists of a study of a new class of approximants for binary inspiral waveforms. This work was done in collaboration with P. Ajith, B. Iyer, and B.S. Sathyaprakash, and was published in [4] and [5]. My contribution to this work mainly consisted of the studies of non-adiabatic waveforms. The next Chapter consists of an investigation into possible methods for performing compact binary searches in a time-critical manner. This is likely to be a serious issue in the future, particularly when spinning sources are considered. The next Chapter gives details of a new method for performing coincidence analysis, developed by myself in collaboration with A.S. Sengupta and B.S. Sathyaprakash, and being used for all ongoing searches for compact binary inspirals in LIGO data. The final Chapter details such an ongoing search - namely the search for inspiral compact binaries with masses ($25M_{\odot} < M < 100M_{\odot}$). This search is being carried out as part of the efforts of the Compact Binary Coalescence group of the LIGO Scientific Collaboration (LSC).

1.1 Gravitational waves

1.1.1 Gravitational waves in linearised theory

We consider a weak gravitational field in a vacuum. In this situation, we can consider the spacetime metric g_{ik} to be a weak perturbation of the flat spacetime metric η_{ik} [6], i.e.

$$g_{ik} = \eta_{ik} + h_{ik}, \quad \eta_{ik} = \text{diag}(1, -1, -1, -1). \quad (1.1)$$

To terms linear in the perturbation h_{ik} , we find

$$g^{ik} = \eta^{ik} - h^{ik} \quad (1.2)$$

and

$$g = (1 + h)\eta \quad (1.3)$$

where $h = h^i_i$. We raise and lower indices using the unperturbed metric η .

In this case, the Riemann curvature tensor becomes

$$R_{iklm} = \frac{1}{2}(h_{im,kl} + h_{kl,im} - h_{il,km} - h_{km,il}) \quad (1.4)$$

This in no way fixes a particular choice of reference frame. In particular, under a change of co-ordinates $x'^i = x^i + \xi^i$, where the ξ^i are small quantities, h_{ik} will transform as

$$h'_{ik} = h_{ik} - 2\xi_{(i,k)}. \quad (1.5)$$

For our purposes, it proves useful to introduce the trace-reversed metric perturbation

$$\bar{h}_{ik} = h_{ik} - \frac{1}{2}\eta_{ik}h. \quad (1.6)$$

We use the gauge freedom given in Eq. 1.5 to impose the condition

$$\bar{h}^{,k}_{ik} = 0. \quad (1.7)$$

This choice is known as the *harmonic* or *de Donder* gauge ¹.

The Einstein field equations are given by

$$G_{ik} \equiv R_{ik} - \frac{1}{2}Rg_{ik} = \frac{8\pi k}{c^4}T_{ik}, \quad (1.8)$$

where G_{ik} is the Einstein tensor, R_{ik} is the Ricci tensor, k is the gravitational constant, and T_{ik} is the energy-momentum tensor of the matter present. If we impose the conditions above in a vacuum, Eq. 1.8 reduces to

$$\square \bar{h}_{ik} = 0. \quad (1.9)$$

This, taken with Eq. 1.7 are linear equations on a homogeneous Minkowski background. It is sufficient to consider a single Fourier mode

$$\bar{h}_{ik} = \hat{h}_{ik}(n_a)e^{in_a x^a} \quad (1.10)$$

Plugging this into the field equation 1.9 gives

$$(n_a n^a) \hat{h}_{ik} = 0, \quad n^k \hat{h}_{ik} = 0. \quad (1.11)$$

These equations tell us that n^i is a null vector, and that \hat{h}_{ik} is orthogonal to n^i . As an illustration, we shall take the specific case of a plane wave travelling in the z -direction. In this case, the spacetime dependence of the Fourier mode is of the form $e^{i\omega(t-z/c)}$. Using the residual gauge freedom we have, we impose the following further conditions:

$$\hat{h}_{0k} = 0, \quad \hat{h}_m{}^m = 0. \quad (1.12)$$

When these are applied, $\hat{h}_{ik} = h_{ik}$, and we have

$$h_{ik} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.13)$$

¹Note that this does not completely fix the gauge, since we can still add $\tilde{\xi}^i$ to the ξ^i , provided they satisfy $\square \tilde{\xi}^i = 0$.

From this, we could conclude that gravitational effects propagate as waves travelling at the speed of light. However, we have chosen a specific gauge, so the existence of non-zero h_{ik} in this gauge is not an invariant indication of a gravitational field. A stronger argument for the existence of the field follows from combining Eq. 1.4 with Eq. 1.9 to get

$$\square R_{iklm} = 0, \quad (1.14)$$

i.e. the curvature tensor itself obeys the wave equation[7].

Polarization of gravitational waves

For the case described above, the line element becomes

$$ds^2 = dt^2 - (1 - h_+)dx^2 - (1 + h_+)dy^2 + 2h_\times dx dy - dz^2 \quad (1.15)$$

from which we can see that gravitational radiation has two different polarization states.

Let us first consider the case where $h_\times = 0$ (i.e. $\hat{h}_\times = 0$), and investigate the effect on a system of test particles. If we consider two neighbouring particles in the (x, y) plane, which initially have co-ordinates (x_0, y_0) and $(x_0 + dx, y_0)$, the proper distance between them will be given by

$$dl^2 = (1 - h_+)dx^2. \quad (1.16)$$

If we now consider two particles located at co-ordinates (x_0, y_0) and $(x_0, y_0 + dy)$, the proper distance will be given by

$$dl^2 = (1 + h_+)dy^2. \quad (1.17)$$

The proper distance is a co-ordinate invariant quantity, so it can be seen that as the value of h_+ oscillates, the particles will move closer together in the x -direction while moving further apart in the y -direction, and vice versa. Figure 1.1 shows the effect of such a wave on a ring of particles. This is known as the $+$ *polarization*.

We now consider the case when $h_+ = 0$. In this case, the line element becomes

$$ds^2 = dt^2 - dx^2 - dy^2 + 2h_\times dx dy - dz^2. \quad (1.18)$$

Figure 1.1: The effect of a $+$ -polarized gravitational wave on a system of particles.





Figure 1.2: The effect of a \times -polarized gravitational wave on a system of particles.



If we change the co-ordinates by performing a rotation of 45° about the z -axis,

$$x' = \frac{1}{\sqrt{2}}(x+y), \quad y' = \frac{1}{\sqrt{2}}(y-x), \quad (1.19)$$

we see that the line element becomes

$$ds^2 = dt^2 - (1 - h_\times)dx'^2 - (1 + h_\times)dy'^2 - dz^2. \quad (1.20)$$

Thus, we can see that the effect of such a wave is the same as a h_+ wave, only rotated through 45° , as illustrated in Figure 1.2. This is known as \times *polarization*.

Note that equations 1.16 and 1.17 illustrate the fact that the change in separation between particles due to a passing gravitational wave is proportional to their initial separation. This is important when considering detection of gravitational waves.

1.1.2 Radiation of gravitational waves

So far, we have seen that the solutions to the linearised Einstein equations in a vacuum allow wave-like solutions. However, we have not yet seen how these waves are generated. To investigate this, let us consider a weak gravitational field produced by bodies moving with velocities much smaller than the speed of light. In this case, we end up with the following equation[6]

$$\frac{1}{2}\Box\bar{h}_{ik} = \frac{8\pi k}{c^4}\tau_{ik}. \quad (1.21)$$

Due to Eq. 1.7, we can see that

$$\tau_{ik}{}^{,k} = 0. \quad (1.22)$$

Eq. 1.21 has a standard ‘retarded potential’ solution of the following form:

$$\bar{h}_{ik} = -\frac{4k}{c^4} \int \frac{\tau_{ik}(t - |\mathbf{x} - \mathbf{x}'|/c, \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d^3x' \quad (1.23)$$

If we are far away from a compact source, i.e. $|\mathbf{x}'| \ll |\mathbf{x}| = r$, then we have, to leading order

$$\bar{h}_{ik} = -\frac{4k}{c^4 r} \int \tau_{ik}(t - r/c, \mathbf{x}') d^3x' \quad (1.24)$$

To solve this equation, we use the result that

$$\int \tau_{\alpha\beta} d^3x = \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \int \tau_{00} x^\alpha x^\beta d^3x. \quad (1.25)$$

In the approximation we are using, the value of τ_{00} can be taken to be μc^2 , where μ is the mass density of the source. Therefore, we obtain

$$\bar{h}_{\alpha\beta} = -\frac{2k}{c^4 r} \frac{\partial^2}{\partial t^2} \int \mu x^\alpha x^\beta d^3x, \quad (1.26)$$

or

$$h_{\alpha\beta} = -\frac{2k}{3c^4 r} \ddot{Q}_{\alpha\beta}, \quad Q_{\alpha\beta} = \int \mu (3x^\alpha x^\beta - \delta_{\alpha\beta} x^\mu x_\mu) d^3x. \quad (1.27)$$

Q is the mass quadrupole tensor of the system. Note that this illustrates that there is no monopole or dipole gravitational radiation. This is a consequence of the conservation of 4-momentum

in general relativity, and is analogous to the fact that there is no monopole electromagnetic radiation due to the conservation of charge.

1.2 Detecting gravitational waves

The attempt to detect gravitational waves has a long history[8, 9]. The field began in the 1960's when J. Weber developed the first bar detectors. Following this, many other groups began to develop similar detectors. Inspired by the work done by these groups, Gertsenshtein and Pustovoit proposed that gravitational waves could be detected by looking for small shifts in the interference pattern of a Michelson interferometer formed between freely hung mirrors. The first interferometer of this type was created by Robert Forward in the late 1960's. From the 1970's onward, Weiss and Drever, amongst others, continued the development of the fundamental designs of laser interferometric detectors, and from that time through to the 1990's, major collaborations such as GEO[10], LIGO[11], Virgo[12] and TAMA[13] came together to build large-scale detectors. Since 2002, the LIGO detectors, consisting of two 4km detectors located at Hanford, Washington, and Livingston, Louisiana (called H1 and L1 respectively), plus a 2 km detector also located at Hanford (called H2); along with GEO600, a 600m detector located in Hanover, Germany, have undertaken several science runs searching for gravitational waves, with the LIGO detectors recently reaching their design sensitivity. The Virgo detector, a 3 km detector located in Cascina, Italy, began its first full scale science run in May 2007, and ran in coincidence with the LIGO detectors until the end of their fifth science run (known as S5).

Despite these many years of efforts, and although there is optimism for the near future, at present, there has been no direct detection of gravitational waves. However, strong indirect evidence of gravitational waves has been observed. Observations of the Hulse-Taylor binary pulsar showed that the decrease in its orbital period agrees remarkably well with the prediction made by general relativity. For these observations, Hulse and Taylor were awarded the 1993 Nobel Prize for Physics.

In this Section, we give a brief description of the indirect detection of gravitational waves, followed by a basic explanation of how laser interferometers such as LIGO can be used for their direct detection.

1.2.1 Indirect detection: the Hulse-Taylor pulsar

The Hulse-Taylor binary pulsar, or PSR B1913+16, was discovered in 1974[14]. It is a system containing two neutron stars, which are thought to have roughly equal masses $m_1, m_2 \approx 1.4M_\odot$, with a separation of $\approx 10^9$ m and an orbital period of ≈ 7.75 Hr. Immediately, this was seen as an ideal system to use as a test of general relativity. Such a system would be expected to emit gravitational waves, and as a result of balance of energy and angular momentum, the angular frequency of the system would be expected to change, to the lowest order, according to the following relation²:

$$\dot{\omega} = \frac{96}{5} \eta M^{5/3} \omega^{11/3}, \quad (1.28)$$

where M is the total mass of the system, and $\eta \equiv m_1 m_2 / M^2$ is the symmetric mass ratio. The system has been observed throughout the years since its discovery, and the change in its orbital angular frequency has been found to agree with general relativity to within 0.3%[15].

1.2.2 Direct detection: the laser interferometer

There are now a number of efforts underway to directly detect gravitational waves using laser interferometric detectors. For the purposes of this thesis, we are mainly concerned with the three LIGO detectors. The LIGO detectors are power recycled Fabry-Perot-Michelson interferometers[16]. The mirrors in the interferometer play the role of the test masses.

To illustrate the basic principles involved, let us consider a simple Michelson interferometer with its arms aligned with the x and y axes, in the presence of a plane, sinusoidal, $+$ -polarized gravitational wave. In this case, the line element Eq. 1.15 becomes


$$ds^2 = dt^2 - (1 - h_+)dx^2 - (1 + h_+)dy^2 - dz^2. \quad (1.29)$$

We can assess the response of the detector to the gravitational wave by looking at the phase shift of the light in the arms. The phase acquired by the light during a round trip along the arm aligned with the x -axis is given by

$$\phi_x = \int_0^{\tau_{RT}} 2\pi f_l dt, \quad (1.30)$$

²Hereafter, we use $k = c = 1$.

Figure 1.3: The basic construction of a Michelson interferometer.



where f_l is the frequency of the laser, and τ_{RT} is the round trip time. Using the fact that $ds^2 = 0$ for light rays, to first order, we end up with [16, 17]

$$\begin{aligned} \phi_x &= \int_0^L 2\pi f_l (1 - h_+)^{1/2} dx - \int_L^0 2\pi f_l (1 - h_+)^{1/2} dx \\ &\approx 4\pi f_l L \left(1 - \frac{h_+}{2}\right), \end{aligned} \quad (1.31)$$

where L is the length of the arm. This gives a phase shift due to the gravitational wave of density, which is the space-time metric tensor, $h_{\mu\nu}$, which is a function of time and space.

$$\delta\phi_x = -2\pi f_l L h_+. \quad (1.32)$$

A similar calculation for the arm aligned with the y-axis gives

$$\delta\phi_y = 2\pi f_l L h_+. \quad (1.33)$$

which gives a total phase shift between the arms of

$$\Delta\phi = 4\pi f_l L h_+. \quad (1.34)$$

This phase shift can be detected by looking for changes in the interferometers interference pat-

tern. The phase shift can be related to the change in length of the arms by using

$$\Delta L \sim \frac{\Delta \phi}{f_l}. \quad (1.35)$$

It is useful to define the *gravitational wave strain* h , as

$$h = \frac{\Delta L}{2L}. \quad (1.36)$$

To be able to detect the gravitational waves from expected sources, we need to be able to measure $h \sim 10^{-22}$. For a 4 km detector this means it is necessary to measure changes in length

$$\Delta L \sim 10^{-22} \times 10^4 = 10^{-18} m, \quad (1.37)$$

To increase the signal response of the detector, further mirrors are added near the beam splitter of the interferometer, creating a Fabry-Perot cavity in each arm, in which the light is stored for approximately 200 bounces. This increases the phase shift $\Delta \phi$ by 3 orders of magnitude.

Noise sources in interferometric detectors

There are many sources of noise which can mask the signal produced by a gravitational wave. The job of the experimenter is to minimise the noise sources, so as to maximise the chance of detecting the signal. The noise in the detector is usually measured as the amplitude spectral density, which is the square root of the power spectral density of the detector strain. A complete overview of the noise searches present in the LIGO interferometers can be found in [18]. Here, we give a brief overview of the most important sources[8].

Seismic noise. Mechanical noise caused by seismic motion of the Earth is an important source of noise, being most dominant at frequencies $f < 40\text{Hz}$. To reduce this noise, it is necessary to isolate the mirrors from the motion of the ground. This is typically done by using mirror suspension systems consisting of coupled pendulums.

Thermal noise. Thermal vibrations of the mirrors and suspensions is the dominant source of noise for frequencies in the range $40 \lesssim f \lesssim 200\text{Hz}$. To mitigate the effects of this noise, systems are designed so that their resonant frequencies are far from the measurement

frequencies. In addition, materials are chosen so that they have high Q , which means their resonant peaks are very sharp, and there will be minimal leakage into the frequency bands of interest.

Shot noise. Shot noise is the dominant source of noise for frequencies $\gtrsim 200\text{ Hz}$. It arises due to the quantum nature of light. We know that the *actual* number of photons collected by the photodetector per unit time is a Poisson process. Therefore, the error will improve as $\sqrt{\bar{N}}$, where \bar{N} is the average number of photons per unit time. It can be calculated that, to be able to detect changes in phase for a gravitational wave with a frequency of $\sim 100\text{ Hz}$, the laser intensity would need to be $\sim 10^2\text{ W}$, which is currently beyond the capability of any continuous laser. To get around this limitation, it is necessary to introduce a power recycling mirror into the interferometer design. This reflects wasted light back into the interferometer, and allows power to build up in the cavity. Using such a design, it is possible to obtain the desired sensitivity using the 10 W lasers typically used in the first generation interferometric detectors.

Other quantum effects. Although shot noise can be reduced by increasing the laser power, this in turn will increase the size of the fluctuations of the laser intensity, and therefore will increase the fluctuations of the light pressure on the mirrors. At high powers, these fluctuations will become a limiting factor of the noise, known as the quantum limit. For future interferometers, investigations are underway into ways to beat this quantum limit, including signal recycling, and the use of squeezed light.

1.3 Sources of gravitational radiation

There are many potential sources of gravitational radiation, each of which has its own signature, and poses its own detection challenges. In the case of ground-based detectors, due to the effects of seismic noise, they are confined to search for sources with frequencies greater than about 10 Hz ³. This means that we are restricted to searching for sources which are highly relativistic and not too massive. Future space-based detectors, such as LISA, will be able to search for more massive systems, and those in a less relativistic regime. In this Section, we give a brief

³For the case of current searches on initial LIGO data, the low-frequency cut-off is taken to be 40 Hz .

overview of some of the sources being searched for with the current generation of ground-based detectors [8, 19].

Spinning neutron stars. Slightly deformed, spinning neutron stars will be a source of quasi-periodic gravitational radiation. A particularly interesting class of objects of this type are low-mass X-ray binaries (LMXBs). In these systems, the neutron stars are spun up by accretion from a companion object; however, the systems seem to be locked into a spin frequencies of $\sim 300 - 600\text{Hz}$. A plausible mechanism for this locking is that the torque due to accretion is being counteracted by an equal and opposite torque due to gravitational radiation. There are a number of potential mechanisms for the emission of gravitational waves by spinning neutron stars. These include density inhomogeneities, the rotation axis not being the the axis of symmetry of the star, causing it to ‘wobble’, or instabilities in the star driven by gravitational radiation reaction. The GW signal emitted by such signals will be weak, but due to their continuous nature, it is possible to use long integration times to increase the signal to noise ratio (SNR). However, for the case of LMXBs, due to the wandering of the spin frequency, care must be taking when analysing more than ~ 20 days of data.

Stochastic background. It is possible to search for a background of gravitational radiation. Such a background could have been created in the early universe, and permeate space in a manner analogous to the Cosmic Microwave Background (CMB). Additionally, a background could exist due to unresolved astrophysical sources, such as supernovae with asymmetric core collapse, binary black hole mergers, or neutron stars [20]. Searches for stochastic backgrounds typically involve cross-correlating the data between different detectors. Directional information can be obtained by introducing suitable time delays between the data streams from each detector. When considering stochastic backgrounds, it is usual to consider the energy density of the gravitational waves, $\Omega_{gw}(f)$, which is usually measured in units of the critical energy density of the Universe, ρ_c . ρ_c is defined as the energy density for which the Universe would just stop expanding at infinite time. The value of ρ_c can be calculated from the Hubble parameter, H , in the following way:

$$\rho_c = \frac{3H^2}{8\pi G}. \quad (1.38)$$

Given this, $\Omega_{gw}(f)$ is defined as

$$\Omega_{gw}(f) = \frac{1}{\rho_c} f \frac{dE_{gw}}{df}. \quad (1.39)$$

Initial detectors should be able to place an upper limit on Ω_{gw} of $\sim 10^{-5}$. The next generation of detectors, such as Advanced LIGO, will be able to measure Ω_{gw} down to $\sim 10^{-9}$. Such a sensitivity will be able to place meaningful constraints on theories of the early universe. However, this limit is orders of magnitude higher than the level predicted by standard inflationary models, so an actual detection would be of profound significance.

Burst sources. There are many potential sources of gravitational radiation which would be expected to be signals of short duration, and be spread over a wide spectral bandwidth. Moreover, the precise form of these signals may presently be poorly understood, thus making optimal searches using matched filtering (see Section 1.5.1) impossible. Examples of such sources are mergers of binary compact objects, and core collapse supernovae. Searching for such signals typically involves looking for excess power in certain frequency bands at certain times. Various consistency checks are then used, including looking at the cross-correlation between data streams from different detectors, and requiring consistency in the relative amplitudes of signals in co-located detectors. While burst searches are less sensitive than matched filtering searches, they have the advantage of not assuming any form for the signal, thus allowing for the prospect of detecting previously unexpected types of signal which may turn out to be present.

Inspiralling compact binaries. A major search effort is underway to detect gravitational waves from inspiralling compact binaries, consisting of neutron stars and/or black holes. As the objects orbit each other, they emit gravitational radiation at twice the orbital frequency, causing the system to lose energy. After a period of adiabatic inspiral, the objects will plunge towards each other, and merge to form a single black hole, which will then radiate gravitational waves as it approaches a stable state. The radiation produced by adiabatic inspiral, and the final black hole's 'ring-down' are reasonably well modelled. However, the standard analytic treatments become less reliable during the plunge and merger phase. Recent advances in numerical relativity have begun to shed some light on these phases of

the evolution. For a more detailed look at compact binary inspirals, see the next Section.

1.4 Compact binary coalescences

Inspiralling compact binaries, consisting of neutron stars and/or black holes are one of the most promising candidates for detection. Estimates of the rates of binary neutron star coalescences per unit L_{10} , where L_{10} is 10^{10} solar blue luminosity, can be inferred from the fact that we know of four systems which will merge within a Hubble time. Current estimates for the merger rate of BNS suggest a rate of $10 - 170 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ at 95% confidence [21]. The rates for black hole-neutron star (BH-NS) systems, and binary black hole systems are highly uncertain, and the systematic uncertainties are poorly understood. However, by exploring the parameter space of the population synthesis models, while ensuring the BNS merger rate inferred from observations is preserved, one can obtain rate estimates for field binaries at the 95% confidence level of $0.1 - 15 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ for BH-NS, and $0.15 - 10 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ for BBH [22]. Other discussions have suggested that certain fractions of dense clusters may form many BBHs. In addition, there is evidence to suggest that short, hard gamma-ray bursts may be associated with the merger of a BNS or BH-NS system. Such discussions suggest that is a real chance of observing systems even with Initial LIGO.

The waveform for a compact binary coalescence consists of three phases: the inspiral phase, the plunge and merger phase, and the ringdown phase. In this thesis, we are mainly concerned with the inspiral phase of the waveform. The inspiral phase is well modelled as a post-Newtonian (PN) expansion in the adiabatic approximation[23]. In this approximation, the dynamics of the system are considered to be a series of quasi-circular orbits, in which the change in orbital frequency is negligible over an orbital period, i.e. $\dot{\omega}/\omega^2 \ll 1$. In this regime, the phasing formula can be obtained by using the energy balance equation, $dE_{tot}/dt = -\mathcal{F}$, where $E_{tot}(x \equiv v^2) = M(1 + E(x))$. Here, M is the total mass of the binary, and $E(x)$ is the PN expansion of the energy function. $\mathcal{F}(v)$ is the GW luminosity function. The invariantly defined velocity v is related to the GW frequency by the following:

$$v = (\pi M f)^{1/3}. \quad (1.40)$$

The GW frequency is twice the orbital frequency. From this, it can be shown that the phasing can be obtained by solving the following system of ODEs:

$$\frac{d\phi}{dt} = \frac{2v^3}{M}, \quad \frac{dv}{dt} = \frac{-\mathcal{F}(v)}{ME'(v)}, \quad (1.41)$$

where $E'(v) = dE/dv$. We integrate these formulae up to v_{lso} , which is the value of v for which E is a minimum, i.e. $E'(v_{lso}) = 0$.

It is known that the PN series is poorly convergent, and so it may be the case that the standard PN expansions are inadequate for detection and parameter estimation of inspirals, particularly for high mass systems. To counteract this, it is possible to use resummation techniques such as Padé resummation to improve convergence. Further information about the use of such techniques can be found in [23].

For systems $\gtrsim 25M_\odot$, the gravitational wave frequency at v_{lso} is close to the peak detector sensitivity. We know that at this point, the adiabatic approximation starts to break down. Moreover, it is found that different ways of re-summing the PN expansions can give significantly different results for the phasing in this regime. Therefore, for systems of this mass and greater, it is necessary to go beyond the adiabatic approximation, and also include the plunge waveform. This can be done using the effective one-body approach [24]. This approach works by mapping the conservative part of the dynamics from a two body problem to an effective one-body problem. With no radiation, the effective metric is a spherically symmetric deformation of the Schwarzschild metric, with the symmetric mass ratio $\eta \equiv m_1 m_2 / M^2$ acting as the deformation parameter. The damping force is re-summed using Padé resummation. The 2PN waveform obtained from such an approach is currently being used as the signal template for the S5 LIGO search for high-mass compact binary coalescences, the current status of which will be described in Section 5. To generate this waveform, we start from the reduced Hamiltonian, given by:

$$\hat{H}(r, p_r, p_\phi) = \frac{1}{\eta} \sqrt{1 + 2\eta \left[\sqrt{A(r) \left(1 + \frac{p_r^2}{B(r)} + \frac{p_\phi^2}{r^2} \right)} - 1 \right]}, \quad (1.42)$$

where the functions $A(r)$ and $B(r)$ are given by

$$A(r) \equiv 1 - \frac{2}{r} + 2\eta \left(\frac{1}{r} \right)^3, \quad B(r) \equiv \frac{1}{A(r)} \left[1 - \frac{6\eta}{r^2} \right]. \quad (1.43)$$

We can then calculate the orbital phase φ by integrating the following equations (here, $\hat{t} \equiv t/M$):

$$\begin{aligned}\frac{dr}{d\hat{t}} &= \frac{\partial \hat{H}}{\partial p_r}, \\ \frac{d\varphi}{d\hat{t}} &= \frac{\partial \hat{H}}{\partial p_\varphi}, \\ \frac{dp_r}{d\hat{t}} &= -\frac{\partial \hat{H}}{\partial r}, \\ \frac{dp_\varphi}{d\hat{t}} &= \mathcal{F}_\varphi.\end{aligned}\tag{1.44}$$

From these equations, the time-domain waveform is then given by

$$h(\hat{t}) = C v_\omega(\hat{t}) \cos[\phi_{GW}(\hat{t})], \quad v_\omega = \left(\frac{d\varphi}{d\hat{t}}\right)^{1/3}, \quad \phi_{GW} = 2\varphi.\tag{1.45}$$

where $\mathcal{F}_\varphi \equiv \frac{1}{\eta v_\omega^3} \mathcal{F}_{P_n}(v_\omega)$. $\mathcal{F}_{P_n}(v_\omega)$ is the flux function used for Padé approximants, which can be found in [23].

The EOB equations of motion are integrated until the termination of the plunge phase at the light ring. At this point, the radial co-ordinate has the value r_{lr} , which is given by solving the following equation:

$$r_{lr}^3 - 3r_{lr}^2 + 5\eta = 0.\tag{1.46}$$

It is suggested in [23] that the best potential candidate sources for detection are stellar mass binaries with a total mass in the range $30 - 90M_\odot$. Including the plunge part by using the EOB waveform increases the SNR significantly compared to the standard post-Newtonian templates for total mass $\gtrsim 35M_\odot$. This highlights the importance of carrying out the high stellar mass compact binary search.⁴

1.5 Data analysis

We shall now discuss the means by which signals are extracted from noise, in particular, for the case of compact binaries. Since, as we have seen, the form of the signals is known in this case, we will first derive the optimal filter for searching for signals of known form, the *matched filter*. We shall then look in more detail at the methods used in data analysis for compact binary

⁴Systems of still higher mass can be searched for using the ringdown signal of the final black-hole. The ringdown in itself is not expected to provide a significant contribution to SNR for systems $\lesssim 100M_\odot$.

searches.

1.5.1 Searching for signals of known form

The optimal filter

To begin with, we will start by deriving the form of the optimal linear filter for a signal whose form is completely known. Let us consider the detector output $x(t)$, of the form

$$x(t) = h(t) + n(t) \quad (1.47)$$

where $h(t)$ is a known signal, and $n(t)$ is the noise present in the detector. We assume the noise to be stationary and Gaussian, as defined by the following equations:

$$\overline{n(t)} = 0, \quad \overline{n^2(t)} = \frac{1}{2\pi} \int_{-\infty}^{\infty} S_n(\omega) d\omega \quad (1.48)$$

We filter the data using a linear filter \mathcal{K} with transfer function $K(\omega)$. It is clear that the output of such a filter will be of the form [25]

$$\varphi(t) = \mu(t) + v(t), \quad (1.49)$$

where $\mu(t)$ is the output due to filtering the signal $h(t)$, and $v(t)$ is the output due to filtering the noise. It can clearly be seen that these functions are given by:

$$\mu(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{i\omega t} K(\omega) \tilde{h}(\omega) d\omega, \quad \overline{v^2(t)} = \frac{1}{2\pi} \int_{-\infty}^{\infty} |K(\omega)|^2 S_n(\omega) d\omega, \quad (1.50)$$

where $\tilde{h}(\omega)$ is the Fourier representation of the signal.

The optimum filter for a known signal is that which maximises the *signal-to-noise ratio* (SNR), ρ , which is defined by

$$\rho^2 = \frac{\mu^2(t_0)}{v^2}, \quad (1.51)$$

where $\mu(t_0)$ is just the output due to filtering the signal at some reference time t_0 . It can be seen

from equations 1.50 that ρ^2 is given by

$$\rho^2 = \frac{1}{2\pi} \frac{|\int_{-\infty}^{\infty} e^{i\omega t_0} K(\omega) \tilde{h}(\omega) d\omega|^2}{\int_{-\infty}^{\infty} |K(\omega)|^2 S_n(\omega) d\omega}. \quad (1.52)$$

Using the Schwarz inequality, it can be proved that ρ^2 obeys the inequality

$$\rho^2 \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{|\tilde{h}(\omega)|^2}{S_n(\omega)} d\omega. \quad (1.53)$$

This maximum value for ρ^2 is obtained if the transfer function $K(\omega)$ takes the following form:

$$K(\omega) = c \frac{e^{-i\omega t_0} \tilde{h}^*(\omega)}{S_n(\omega)}, \quad (1.54)$$

where c is an arbitrary constant. This defines the optimal filter for a known signal.

The same filter can be used for signals with unknown time of arrival τ , and signal amplitude G , i.e. signals of the form

$$h'(t) = Gh(t - \tau) \quad (1.55)$$

In this case, the time of arrival can be found by maximising $\mu(t)$ over time. The SNR at this maximum is $\rho = G\rho_0$, where ρ_0 is the SNR when $G = 1$, given in Eq. 1.53.

The statistical nature of detection

In practice, we do not know the particular parameters of a signal ahead of time; or even if there is a signal present. Due to the random nature of the noise present, detection of signals, and estimation of the parameters, cannot be achieved with any certainty. It is possible that the noise can conspire to give the impression of the presence of a signal, when there is none. Moreover, if a signal is detected, the random contamination of the signal by the noise will introduce errors in the measurement of the parameters. Due to these reasons, it is necessary to treat detection of signals in a statistical manner [25, 26].

We will assume that we are searching for a signal $h(t; \mu)$, where μ is a vector of the unknown parameters characterising the system. The data stream from the detector will be of the form

$$x(t) = h(t; \mu) + n(t) \quad (1.56)$$

if there is a signal present, or

$$x(t) = n(t) \quad (1.57)$$

if there is no signal present.

A useful point to start the analysis of this problem is Bayes' theorem, which is given by

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)}, \quad (1.58)$$

where $P(A|B)$ posterior probability of event A occurring given event B has occurred, $P(A)$ is the prior probability of event A occurring, $P(B|A)$ is the probability of event B occurring given event A has occurred (known as the likelihood), and $P(B)$ is the prior probability event B occurring. In our case, we are interested in the probability $P(h|x)$ of a signal h being present given we have received data x . The posterior probability of there being a signal present is then given by

$$P(h|x) = \frac{P(h)P(x|h)}{P(x)}. \quad (1.59)$$

The probability $P(x)$, taking into account the unknown parameters μ of the signal, is given by

$$P(x) = \int p[h(\mu)]P[x|h(\mu)]d^n\mu + P(0)P(x|0) \quad (1.60)$$

where $P(0)$ is the prior probability of there being no signal present, and $P(x|0)$ is the likelihood of the data given there is no signal. Using this in Eq. 1.59, and making a simple transformation, we get the following

$$P(h|x) = \frac{\Lambda}{\Lambda + [P(0)/P(h)]}, \quad (1.61)$$

where Λ is the *likelihood ratio*, defined as

$$\Lambda = \int \Lambda(\mu)d^n\mu = \frac{P(x|h)}{P(x|0)} = \int \frac{p(\mu|h)P[x|h(\mu)]}{P(x|0)}d^n\mu. \quad (1.62)$$

Similarly, the posterior probability of there being no signal present is given by

$$P(0|x) = 1 - P(h|x) = \frac{P(0)/P(h)}{\Lambda + [P(0)/P(h)]}. \quad (1.63)$$

It can be seen that the posterior probability of a signal being present given the data received

depends on the prior probability of a signal being present. In practice, this prior probability is unknown, so the decision as to whether a signal is present is taken on the basis of the likelihood ratio. The decision rule takes the following form:

If $\Lambda \geq \Lambda_*$, we decide the signal is present.

If $\Lambda < \Lambda_*$, we decide the signal is absent.

where Λ_* is the threshold value of the likelihood ratio.

If we consider this decision rule, it is clear that there are two types of correct decisions which can be made: correctly asserting that a signal is detected; and correctly asserting that a signal is not detected. Similarly, there are two types of error which can occur: asserting that a signal has been detected when there was no signal present, known as a *false alarm*; and asserting that no signal was detected when a signal was present, known as a *false dismissal*. We denote the probability of a correct detection as D , and the probability of a correct non-detection as F_0 . From this, it is clear that

$D_0 = 1 - D$ is the probability of a false dismissal.

$F = 1 - F_0$ is the probability of a false alarm.

The total probability of a correct detection is clearly given by

$$P(\text{correct}) = P(h)D + p(0)(1 - F). \quad (1.64)$$

The observer who, on the basis of the data obtained, maximizes the detection probability for a fixed false alarm probability, is known as a *Neyman-Pearson observer*. The decision rule given above is the optimal rule in this sense.

Application to binary inspirals

We will now consider the construction of the likelihood ratio for the case of a signal $h(t)$, and a data stream $x(t)$. We assume the noise to be stationary and Gaussian, with zero mean, and a one-sided power spectral density $S_n(|f|)$ defined by

$$\overline{n(t)} = 0, \quad \overline{n(f)n(f')} = \frac{1}{2}S_n(|f|)\delta(f - f'). \quad (1.65)$$

It should be noted that the probabilities $P(x|h)$ and $P(x|0)$ which come into the calculation of Λ , are usually zero. To remove this issue, we write

$$\Lambda = \frac{P(x|h)}{P(x|0)} = \frac{p(x|h)dx}{p(x|0)dx} = \frac{p(x|h)}{p(x|0)}, \quad (1.66)$$

i.e. instead of using the actual probabilities to calculate Λ , we use the probability densities.

We introduce an inner product, defined as

$$\langle a, b \rangle = 2 \int_{-\infty}^{\infty} df \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} \quad (1.67)$$

It can be shown that the probability density of obtaining a given realisation of detector noise is given by

$$p(n) = \mathcal{K} e^{-\frac{1}{2}\langle n, n \rangle}. \quad (1.68)$$

Therefore, the probability density of obtaining data $x(t)$ given there is no signal present is given by

$$p(x|0) = \mathcal{K} e^{-\frac{1}{2}\langle x, x \rangle}. \quad (1.69)$$

Similarly, the probability density of obtaining data $x(t)$ given that a signal $h(t)$ is present is given by

$$p(x|h) = \mathcal{K} e^{-\frac{1}{2}\langle x-h, x-h \rangle}. \quad (1.70)$$

Thus, the likelihood ratio Λ is given by

$$\Lambda = \frac{p(x|h)}{p(x|0)} = \exp[\langle x, h \rangle - \frac{1}{2}\langle h, h \rangle]. \quad (1.71)$$

It is clear in Eq. 1.71 that the only dependence of Λ on the data received is through $\langle x, h \rangle$. Moreover, Λ is a monotonically increasing function of $\langle x, h \rangle$. Therefore, instead of considering the likelihood ratio for making the decision regarding the presence of a signal, we can instead use $\langle x, h \rangle \equiv \varphi$:

If $\varphi \geq \varphi_*$, we decide the signal is present.

If $\varphi < \varphi_*$, we decide the signal is absent. It is worth noting that, for completely known signal h , this is the same as using the optimal filter given in Eq. 1.54, and making

the decision based on its output.

It is clear that φ is a random variable with a Gaussian distribution. If only noise is received, the mean value of φ will be zero. The variance will be given by

$$\sigma^2 \equiv \overline{\varphi^2} = \langle h, h \rangle. \quad (1.72)$$

In the presence of a signal h , the mean and variance will be given by

$$\bar{\varphi} = \langle h, h \rangle, \quad \sigma^2 \equiv \overline{(\varphi - \bar{\varphi})^2} = \langle h, h \rangle \quad (1.73)$$

Given the distributions of φ in the presence and absence of a signal, the detection and false alarm probabilities can be calculated, and the threshold value φ_* can be chosen.

If the signal is of an unknown amplitude, i.e. $h'(t; G) = Gh(t)$, the likelihood ratio changes to the form

$$\Lambda = \exp[G \langle x, h \rangle - \frac{1}{2} G^2 \langle h, h \rangle]. \quad (1.74)$$

As before, this function is monotonic in $\langle x, h \rangle$, so the same decision rule can be applied as before. Note that this implies that one need only filter the data with templates h of unit norm. In this case, the mean value of φ in the presence of a signal $h'(t; G)$ is $G \langle h, h \rangle$.

We now consider the case of a signal with unknown phase. In this case, the likelihood ratio becomes

$$\Lambda(\theta) = p(\theta|h) \exp[\langle x, h(\theta) \rangle - \frac{1}{2} \langle h, h \rangle]. \quad (1.75)$$

We write the signal in the form

$$h(t, \theta) = A(t) \cos[2\phi(t) - \theta]. \quad (1.76)$$

Substituting Eq. 1.76 in Eq. 1.75, and using simple trigonometric identities, we end up with

$$\begin{aligned} \varphi(\theta) &= \cos \theta \langle x, A(t) \cos 2\phi(t) \rangle + \sin \theta \langle x, A(t) \sin 2\phi(t) \rangle \\ &= a \cos \theta + b \sin \theta \\ &= \mathcal{E} \cos(\Phi - \theta). \end{aligned} \quad (1.77)$$

In calculating the likelihood ratio, we assume that the value of θ is equally likely to be any value between 0 and 2π . This means that

$$p(\theta|h) = \frac{1}{2\pi}. \quad (1.78)$$

The likelihood ratio is therefore given by

$$\begin{aligned} \Lambda &= \int_0^{2\pi} \Lambda(\theta) d\theta = \frac{1}{2\pi} \int_0^{2\pi} e^{\mathcal{E} \cos(\Phi - \theta) - \frac{1}{2} \langle h, h \rangle} d\theta \\ &= e^{-\frac{1}{2} \langle h, h \rangle} I_0(\mathcal{E}), \end{aligned} \quad (1.79)$$

where I_0 is the modified Bessel function of the first kind of order zero.

It can be seen that, similarly to the previous cases, the data received only enters into the likelihood ratio via \mathcal{E} , and the likelihood ratio is a monotonically increasing function of \mathcal{E} . Therefore, we can make the decision whether a signal is present based on the value of \mathcal{E} .

If we look more closely at the value of \mathcal{E} , we see that it has the following form

$$\mathcal{E} = \sqrt{a^2 + b^2} = \sqrt{\langle x, h_c \rangle^2 + \langle x, h_s \rangle^2}, \quad (1.80)$$

where

$$h_c = A(t) \cos 2\phi(t), \quad h_s = A(t) \sin 2\phi(t), \quad (1.81)$$

i.e. the value of \mathcal{E} is given by the sum of the squares of two optimal linear filters in quadrature.

It is convenient in practice to work with a normalised signal-to-noise ratio, defined by

$$\rho^2 = \frac{\mathcal{E}^2}{\sigma^2}, \quad \sigma^2 = \langle h_c, h_c \rangle. \quad (1.82)$$

It is clear that, in the presence of stationary Gaussian noise only, ρ^2 will obey a χ^2 distribution with two degrees of freedom.

As well as an unknown phase, it is clear that the time of the signal will also be unknown. In searching for coalescing binary systems, it is typical to use the time of coalescence, t_c , to locate the signal in time. This is defined as the time at which the inspiral waveforms used in the

search are terminated, e.g. $t(v_{lso})$ in the case of the TaylorT1 approximant defined in Eq. 1.41. We can consider the case of searching for the signal $h(t - t_C)$. In this case, due to the properties of Fourier transforms, the signal in the Fourier domain becomes

$$\tilde{h}_{t_C}(f) = e^{-2\pi i f t_C} \tilde{h}(f), \quad (1.83)$$

where $\tilde{h}_{t_C}(f)$ is the Fourier transform of $h(t - t_C)$, and $\tilde{h}(f)$ is the Fourier transform of $h(t)$. If we then consider the inner products which come into the construction of the statistic ρ^2 for signals ending at t_C , we see that

$$\begin{aligned} \langle x, h_c(t_C) \rangle &= 2 \int_{-\infty}^{\infty} e^{2\pi i f t_C} \frac{\tilde{x}(f) \tilde{h}_c^*(f)}{S_n(f)} df, \\ \langle x, h_s(t_C) \rangle &= 2 \int_{-\infty}^{\infty} e^{2\pi i f t_C} \frac{\tilde{x}(f) \tilde{h}_s^*(f)}{S_n(f)} df. \end{aligned} \quad (1.84)$$

It can be seen that the inner products are inverse Fourier transforms, thus making ρ^2 a function of t_C . To perform the measurement, we state that the most likely value for t_C is that which maximizes the likelihood ratio. In performing the search, we therefore maximise $\rho(t)$ over time. If this maximum is greater than the threshold value ρ^* , we record an event as having occurred at time t . Since there is the possibility of more than one signal occurring in a given stretch of data, we do not maximise over all time; we instead maximise over a certain time interval δt .

In searching for binary inspiral signals, there are many other parameters which will be unknown. For detection in a single detector, certain parameters, e.g. the sky location and the inclination of the binary, can be folded into the amplitude of the signal. This is done by introducing an *effective distance*, \mathcal{D} , which is related to the amplitude G in the following way [27, 28]:

$$\frac{1}{G} \propto \mathcal{D} \equiv \frac{r}{\sqrt{F_+^2(1 + \cos^2 i)^2/4 + F_\times^2 \cos^2 i}}, \quad (1.85)$$

where i is the inclination angle of the binary, and F_+ and F_\times are the antenna pattern functions of the detector.

For the other parameters, which, in the cases of non-spinning inspiralling compact binaries, correspond to the component masses of the system, (m_1, m_2) , it is not feasible to search over them in a continuous manner. Therefore, a bank of *templates* is defined with a certain granularity

on the parameter space. The spacing is chosen so that the *match* between any signal within the search space and at least one template will be more than a given minimum value. A summary of the geometrical formalism used for the construction of a template bank, can be found in Ch. 4. The details of how this formalism is used for template placement can be found in [29].

In measuring the parameters m_1 and m_2 , as with measuring t_C , we choose the parameters which maximise the value of the likelihood ratio. As is the case with t_C , to allow for the possibility of more than one signal, each having different parameters, we do not maximise over the entire parameter space, but maximise over a certain neighbourhood within this space. A description of an algorithm which performs such clustering can be found in [30].

The case of searching real data

In the case of real data, as produced by the current generation of interferometric detectors, the noise present in the data stream is neither stationary nor Gaussian. As a result of this, there is likely to be a large number of false alarms with a high signal-to-noise ratio when looking at a single detector. To reduce this false alarm rate, and to increase the confidence of a detection, we look for coincident events appearing in different detectors. For events in different detectors to be considered coincident, we require the parameters of the events to agree to within a certain tolerance. A new method for checking for coincidences can be found in Ch. 4.

Chapter 2

A new class of post-Newtonian approximants for gravitational wave data analysis

2.1 Introduction

The late-time dynamics of astronomical binaries consisting of neutron stars and/or black holes is dominated by relativistic motion and non-linear general relativistic effects. The component bodies would be accelerated to velocities close to half the speed of light before they plunge towards each other, resulting in a violent event during which the source would be most luminous in the gravitational window. Such events are prime targets of interferometric gravitational wave (GW) detectors that are currently taking data at unprecedented sensitivity levels and bandwidths [11, 12, 10, 13].

Binary coalescences are the end state of a long period of adiabatic dynamics in which the orbital frequency of the system changes as a result of gravitational radiation backreaction but the change in frequency per orbit is negligible compared to the orbital frequency itself. Indeed, the adiabatic inspiral phase is well-modelled by the post-Newtonian (PN) approximation to Einstein's equations but this approximation becomes less accurate close to the merger phase. Additionally, there are different ways of casting the gravitational wave phasing formula – the formula that gives the phase of the emitted gravitational wave as a function of time and the pa-

rameters of the system. These different approaches make use of the post-Newtonian expansions of the binding energy and gravitational wave luminosity of the system¹.

2.1.1 Standard approach to phasing formula

The standard approach in deriving the phasing formula uses the *specific* gravitational binding energy $E(v)$ (i.e. the binding energy per unit mass) of the system and its luminosity $\mathcal{F}(v)$, both to the same relative accuracy [31]. Including the radiation reaction at dominant order, however, is not a first order correction to the dynamics of the system, rather it is a correction that arises at $O[(v/c)^5]$, where v is the post-Newtonian expansion parameter describing the velocity in the system and c is the speed of light². Thus, the phasing of the waves when translated to the relative motion of the bodies implies that the dynamics are described by the dominant Newtonian force and a correction at an order $(v/c)^5$, but neglecting conservative force terms that occur at orders $(v/c)^2$ and $(v/c)^4$. Such considerations have led to an approximation scheme in which one constructs the phasing of gravitational waves using the following ordinary, coupled differential equations:

$$\frac{d\phi}{dt} = \frac{2v^3}{m}, \quad \frac{dv}{dt} = -\frac{\mathcal{F}(v)}{mE'(v)}, \quad (2.1)$$

where $E'(v) = dE(v)/dv$ and $m = m_1 + m_2$ is the total mass of the binary. The phasing obtained by numerically solving the above set of differential equations is called the *TaylorT1* approximant [23]. If the detector's motion can be neglected during the period when the wave passes through its bandwidth then the response of the interferometer to arbitrarily polarized waves from an inspiralling binary is given by

$$h(t) = \frac{4A\eta M}{D} v^2(t) \cos[\phi(t) + \phi_C], \quad (2.2)$$

where $\phi(t)$ is defined so that it is zero when the binary coalesces at time $t = t_C$, ϕ_C is the phase of the signal at t_C , $\eta = m_1 m_2 / M^2$ is the symmetric mass ratio, D is the distance to the source and A is a numerical constant whose value depends on the relative orientations of the interferometer and the binary orbit. It suffices to say for the present purpose that for an optimally oriented

¹In the case of binaries consisting of spinning bodies in eccentric orbit one additionally requires equations describing the evolution of the individual spins and the orbital angular momentum, but this complication is unimportant for our purposes.

²Throughout this Chapter we use units in which $G = c = 1$.

source $A = 1$.

One can compute the Fourier transform $H(f)$ of the waveform given in Eq. (2.2) using the stationary phase approximation:

$$H(f) = \frac{4Am^2}{D} \sqrt{\frac{5\pi\eta}{384}} v_f^{-7/2} e^{i[2\pi f t_c - \phi_c + \psi(f) - \pi/4]}, \quad (2.3)$$

where the phase of the Fourier transform obeys a differential equation given by

$$\frac{d\psi}{df} = 2\pi t, \quad \frac{dt}{df} = -\frac{\pi m^2}{3v_f^2} \frac{E'(v_f)}{\mathcal{F}(v_f)}. \quad (2.4)$$

The above expressions contain the post-Newtonian expansions of the energy and flux functions, as a function of the velocity parameter v_f . In terms of the frequency, the parameter is given by $v_f = (\pi m f)^{1/3}$. The waveform Eq.(2.3) computed by numerically solving the differential equations Eq.(2.4) is called *TaylorF1* [23] approximant.

Before we proceed further, let us recall the notation used in post-Newtonian theory to identify different orders in the expansion. In the conservative dynamics of the binary, wherein there is no dissipation, the energy is expressed as a post-Newtonian expansion in $\varepsilon = (v/c)^2$, with the dominant order termed Newtonian or 0PN and a correction at order $\varepsilon^n = (v/c)^{2n}$, $n = 1, 2, \dots$, called n PN, with the dynamics involving only even powers of $\sqrt{\varepsilon} = (v/c)$. When dissipation is added to the dynamics, then the equation of motion will have terms of both odd and even powers of v/c . Thus, a correction of order $(v/c)^m$ is termed as $(m/2)$ PN.

In the case of a test-particle orbiting a Schwarzschild black hole, the energy function $E(v)$ is exactly calculable analytically, while the flux function $\mathcal{F}(v)$ is exactly calculable numerically [32, 33, 34, 35]. In addition, $\mathcal{F}(v)$ has been calculated analytically to 5.5PN order [36] by black hole perturbation theory [37]. In contrast, in the case of a general binary including bodies of comparable masses, the energy function $E(v)$, has been calculated recently to 3PN accuracy by a variety of methods [38, 39, 40, 41, 42, 43, 44, 45, 46, 47]. The flux function $\mathcal{F}(v)$, on the other hand, has been calculated to 3.5PN accuracy [48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58] up to now only by the multipolar-post-Minkowskian method and matching to a post-Newtonian source [59].

2.1.2 Complete phasing of the adiabatic inspiral: An alternative

The gravitational wave flux arising from the lowest order quadrupole formula, that is the 0PN order flux, leads to an acceleration of order 2.5PN in the equations-of-motion. This far-zone computation of the flux requires a control of the dynamics, or acceleration, to only *Newtonian* accuracy. The lowest order GW phasing in the adiabatic approximation uses only the leading terms in the energy (Newtonian) and flux (quadrupolar) functions. For higher order phasing, the energy and flux functions are retained to the same relative PN orders. For example, at 3PN phasing, both the energy and flux functions are given to the same *relative* 3PN order beyond the leading Newtonian order. We refer to this usual physical treatment of the phasing of GWs computed in the adiabatic approximation, and used in the current LIGO/Virgo/GEO/TAMA searches for the radiation from inspiralling compact binaries, as the *standard adiabatic* approximation. We will denote the $n\text{PN}$ *standard adiabatic* approximant as $T(E_{[n]}, \mathcal{F}_n)$, where $[p]$ denotes the integer part of p .

In the adiabatic approximation, we assume the system evolves through a series of quasicircular orbits, for which the change in the orbital frequency is small over the orbital period. In this regime, we are justified in determining the evolution of the inspiral by using the energy balance equation, $\dot{E} = -\mathcal{F}$. As a prelude to go beyond the standard adiabatic approximation, let us consider the phasing of the waves in terms of the equations of motion of the system. To this end, it is natural to order the PN approximation in terms of its dynamics or acceleration. From the view point of the dynamics, the leading order standard adiabatic approximation is equivalent to using the 0PN (corresponding to 0PN conserved energy) and 2.5PN (corresponding to the Newtonian 0PN flux) terms in the acceleration ignoring the intervening 1PN and 2PN terms. A complete, mathematically consistent treatment of the acceleration, however, should include *all* PN terms in the acceleration up to 2.5PN, *without any gaps*. We refer to the dynamics of the binary, and the resulting waveform, arising from the latter as the 0PN *complete non-adiabatic* approximation. In contrast, the waveform arising from the former choice, with gaps in the acceleration at 1PN and 2PN, is referred to as the 0PN *standard non-adiabatic* approximation. Extension to higher-order phasing is obvious. At 1PN the standard non-adiabatic approximation would involve acceleration terms at orders 0PN, 1PN, 2.5PN and 3.5PN, whereas the complete non-adiabatic would additionally involve the 2PN and 3PN acceleration terms.

Finally, we propose a simple extension of the above construction to generate a new class of approximants in the adiabatic regime. To understand the construction let us examine the lowest order case. Given the 0PN flux (leading to an acceleration at 2.5PN), one can choose the energy function at 2PN (equivalent to 2PN dynamics) instead of the standard choice 0PN (equivalent to 0PN or Newtonian dynamics). This is the adiabatic analogue of the complete non-adiabatic approximant³. Extension to higher PN orders follows naturally. For instance, corresponding to the flux function at 1PN (1.5PN), the dissipative force is at order 3.5PN (4PN), and, therefore, the conservative dynamics, and the associated energy function, should be specified up to order 3PN (4PN). In general, given the flux at n PN-order, a corresponding complete adiabatic approximant is constructed by choosing the energy function at order $[n + 2.5]$ PN, where as mentioned before, $[p]$ denotes the integer part of p . We refer to the dynamics of the binary and the resulting waveform arising from such considerations, as the *complete adiabatic* approximation. We will denote the n PN *complete adiabatic* approximant as $T(E_{[n+2.5]}, \mathcal{F}_n)$.

Before moving ahead the following point is worth emphasizing: The standard adiabatic phasing is, by construction, consistent in the *relative PN* order of its constituent energy and flux functions, and thus *unique* in its ordering of the PN terms. Consequently, one can construct different inequivalent, but consistent, approximations as discussed in Ref. [23] by choosing to retain the involved functions or re-expand them. The complete adiabatic phasing, on the other hand, is constructed so that it is consistent in spirit with the underlying dynamics, or acceleration, rather than with the relative PN orders of the energy and flux functions. Consequently, it has a unique meaning only when the associated energy and flux functions are used *without any further re-expansions* when working out the phasing formula. As a result, though the complete non-adiabatic approximant is more consistent than the standard non-adiabatic approximant in treating the PN accelerations, in the adiabatic case there is no rigorous sense in which one can claim that either of the approximants is more consistent than the other. The important point, as we shall see is that, not only are the two approximants *not* the same but the new complete adiabatic approximants are closer to the exact solution than the standard adiabatic approximants.

In our view, these new approximants should be of some interest. They are simple generalizations of the *standard adiabatic* approximants coding information of the PN dynamics beyond

³In this case one may also choose the energy function to 3PN accuracy and construct a complete approximant leading to 3PN acceleration

the standard approximation without the need for numerical integration of the equations of motion. They should be appropriate approximants to focus on when one goes beyond the adiabatic picture and investigates the differences stemming from the use of more complete equations of motion (see Section 2.3).

In the case of comparable mass binaries, the energy function is currently known up to 3PN order and hence it would be possible to compute the complete adiabatic phasing of the waves to only 1PN order. One is thus obliged in practice to follow the standard adiabatic approximation to obtain the phasing up to 3.5PN order. Consequently, it is a relevant question to ask how ‘close’ are the complete and standard adiabatic approximants. The standard adiabatic approximation would be justified if we can verify that it produces in most cases a good lower bound to the mathematically consistent, but computationally more demanding, complete adiabatic approximation. In this Chapter we compare the standard and complete models by explicitly studying their overlaps with the exact waveform which can be computed in the adiabatic approximation of a test mass motion in a Schwarzschild spacetime. The availability, in this case, of the exact (numerical) and approximate (analytical) waveforms to as high a PN order as $(v/c)^{11}$, allows one to investigate the issue exhaustively, and provides the main motivation for the present analysis. Assuming that the comparable mass case is qualitatively similar and a simple η -distortion of the test mass case would then provide a plausible justification for the standard adiabatic treatment of the GW phasing employed in the literature⁴.

2.1.3 Non-adiabatic inspiral

The phasing formulas derived under the various adiabatic approximation schemes assume that the orbital frequency changes slowly over each orbital period. In other words, the change in frequency Δf over one orbital period P is assumed to be much smaller than the orbital frequency $f \equiv P^{-1}$. Denoting by \dot{f} the time-derivative of the frequency, the adiabatic approximation is equivalent to the assumption that $\Delta f = \dot{f}P \ll f$ or $\dot{f}/f^2 \ll 1$. This assumption becomes somewhat weaker, and it is unjustified to use the approximation $\dot{E} = -\mathcal{F}$, when the two bodies are quite close to each other. Buonanno and Damour [61, 24] introduced a non-adiabatic approach

⁴Note, however, that the view that the comparable mass case is just a η -distortion of the test mass approximation is not universal. In particular, Blanchet [60] has argued that the dynamics of a binary consisting of two bodies of comparable masses is very different from, and possibly more accurately described by post-Newtonian expansion than, the test mass case.




Figure 2.1: Various T-approximants of Newton-normalized (v -derivative of) energy function $E'_T(v)/E'_N(v)$ (left), and flux function $\mathcal{F}_T(v)/\mathcal{F}_N(v)$ (right) in the test mass limit along with the exact functions (denoted by X). Also plotted is the amplitude spectral density (per $\sqrt{\text{Hz}}$) of Initial LIGO noise in arbitrary units.

to the two-body problem called the *effective one-body* (EOB) approximation. In this approximation one solves for the relative motion of the two bodies using an effective Hamiltonian with a dissipative force put-in by hand. EOB allows to extend the dynamics beyond the adiabatic regime, and the last stable orbit, into the plunge phase of the coalescence of the two bodies [24, 62, 63, 64].

Recently, Buonanno, Chen and Vallisneri [65] have studied a variant of the non-adiabatic model but using the effective Lagrangian constructed in the post-Newtonian approximation. We shall use both the *standard* and *complete* non-adiabatic Lagrangian models in this study and see how they converge to the exact waveform defined using the adiabatic approximation⁵.

2.1.4 What this study is about

In our study we will use the *effectualness* and *faithfulness* (see below) to quantify how good the various approximation schemes are. There are at least three different contexts in which one can examine the performance of an approximate template family relative to an exact one. Firstly, one can think of a mathematical family of approximants and examine its convergence towards some exact limit. Secondly, one can ask whether this mathematical family of approximants correctly represents the GWs from some physical system. Thirdly, how does this family of approximate templates converge to the exact solution in the sensitive bandwidth of a particular GW detector. In the context of GW data analysis, the third context will be relevant and studied in this Chapter. Although there is no direct application to GW data analysis, equally interesting is the

⁵See Sec. 2.3 for a caveat in this approach.

mathematical question concerning the behavior of different approximations, and the waveforms they predict, in the strongly non-linear regime of the dynamics of the binary, which is also studied in this Chapter. The latter obviously does not require the details of the detector-sensitivity and it is enough to study the problem assuming a flat power spectral density (i.e. a *white-noise* background) for the detector noise.

To summarize, our approach towards the problem will be two-pronged. First, we will study the problem as a general mathematical question concerning the nature of templates defined using PN approximation methods. We shall deal with four families of PN templates – the standard adiabatic, complete adiabatic, standard non-adiabatic and complete non-adiabatic (in particular, Lagrangian-based) approximants – and examine their closeness, defined by using effectualness and faithfulness, to the exact waveform defined in the adiabatic approximation. Since this issue is a general question independent of the characteristics of a particular GW detector, we first study the problem assuming the white-noise case. Having these results, we then proceed to see how and which of these results are applicable when applied to specific detectors. To this end, we study the case of the Initial LIGO, Advanced LIGO, and Virgo detectors. During the course of this study, we also attempt to assess the relative importance of improving the accuracy of the energy and flux functions by studying the overlaps of the PN templates constructed from different orders of energy and flux functions with the exact waveform. It should be kept in mind that this work is solely restricted to the inspiral part of the signal and neglects the plunge and quasi-normal mode ringing phases of the binary [66, 67, 24, 23, 65, 68, 69].

2.1.5 Effectualness and Faithfulness

In order to measure the accuracy of our approximate templates we shall compute their overlap with a fiducial exact signal. We shall consider two types of overlaps [70, 71, 23, 72]. The first one is the *faithfulness* which is the overlap of the approximate template with the exact signal computed by keeping the intrinsic parameters (e.g. the masses of the two bodies) of both the template and the signal to be the same but maximizing over the extrinsic (e.g. the time-of-arrival and the phase at that time) parameters. The second one is the *effectualness* which is the overlap of the approximate template with the exact signal computed by maximizing the overlap over both the intrinsic and extrinsic parameters. Faithfulness is a measure of how good

is the template waveform in both detecting a signal and measuring its parameters. However, effectualness is aimed at finding whether or not an approximate template model is good enough in detecting a signal without reference to its use in estimating the parameters. As in previous studies, we take overlaps greater than 96.5% to be indicative of a good approximation.

In the next section we study the test-mass waveforms in the adiabatic approximation. We discuss the construction of the exact energy and flux functions as well as the respective T-approximants. The overlaps of various standard adiabatic and complete adiabatic approximants are also compared in this Section. Section III deals with the non-adiabatic approximation. Section IV explores the extension of the results in the comparable mass case. It presents the energy and flux functions which are the crucial inputs for the construction of the fiducial exact waveform as well as the approximate waveforms followed by a discussion of the results. In the last section we summarize our main conclusions.

One of the main conclusions of this Chapter is that the effectualness of the test-mass approximants significantly improves in the complete adiabatic approximation at PN orders below 3PN. However, standard adiabatic approximants of order ≥ 3 PN are nearly as good as the complete adiabatic approximants for the construction of effectual templates. In the comparable mass case the problem can be only studied at the lowest two PN orders. No strong conclusions can be drawn as in the test mass case. Still, the trends indicate that the standard adiabatic approximation provides a good lower bound to the complete adiabatic approximation for the construction of both effectual and faithful templates at PN orders ≥ 1.5 PN. From the detailed study of test-mass templates we also conclude that, provided the comparable mass case is qualitatively similar to the test mass case, neither the improvement of the accuracy of energy function from 3PN to 4PN nor the improvement of the accuracy of flux function from 3.5PN to 4PN will result in a significant improvement in effectualness in the comparable mass case. As far as faithfulness is concerned, it is hard to reach any conclusion. To achieve the target sensitivity of 0.965 in effectualness corresponding to a 10% loss in the event-rate, standard adiabatic approximants of order 2PN and 3PN are required for the $(10M_{\odot}, 10M_{\odot})$ and $(1.4M_{\odot}, 1.4M_{\odot})$ binaries, respectively, when restricting to only the inspiral phase. (Be warned that this is not a good approximation in the BH-BH case since the approach to the plunge and merger lead to significant differences.)

2.1.6 Noise spectra of the interferometers

The noise power spectral density Initial LIGO, Virgo and Advanced LIGO are given in terms of a dimensionless frequency $x = f/f_0$ by

- Initial LIGO [23]

$$S_n(f) = 9 \times 10^{-46} [(4.49x)^{-56} + 0.16x^{-4.52} + 0.52 + 0.32x^2] \quad (2.5)$$

where $f_0 = 150$ Hz.

- Virgo [23]

$$S_n(f) = 3.24 \times 10^{-46} [(6.23x)^{-5} + 2x^{-1} + 1 + x^2] \quad (2.6)$$

where $f_0 = 500$ Hz.

- Advanced LIGO [19, 104]

$$S_n(f) = 10^{-49} \left[x^{-4.14} - 5x^{-2} + 111 \frac{1 - x^2 + x^4/2}{1 + x^2/2} \right] \quad (2.7)$$

where $f_0 = 215$ Hz.

2.2 Test mass waveforms in the adiabatic approximation

Our objective is to compare the *effectualness* (i.e. larger overlaps with an exact signal) and *faithfulness* (i.e. smaller bias in the estimation of parameters) of the standard adiabatic $T(E_{[n]}, \mathcal{F}_n)$ and complete adiabatic $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants. As a by-product of this study, we would also like to have an understanding of the relative importance of improving the accuracy of the energy function and flux function. Thus, what we will do is to take all possible combinations of T-approximants⁶ of energy and flux functions, construct PN templates and calculate the overlap of these templates with the exact waveform. In all cases, the exact waveform is constructed by numerically integrating the phasing formula in the time-domain [*TaylorT1* approximant, cf. Eqs. (2.1) and (2.2)]. The waveforms (both the exact and approximate) are all terminated at $v_{\text{iso}} = 1/\sqrt{6}$, which corresponds to $F_{\text{iso}} \simeq 43$ Hz for the $(1M_\odot, 100M_\odot)$ binary, $F_{\text{iso}} \simeq 86$ Hz for

⁶We follow [23] in denoting the precise scheme used for constructing the approximant.

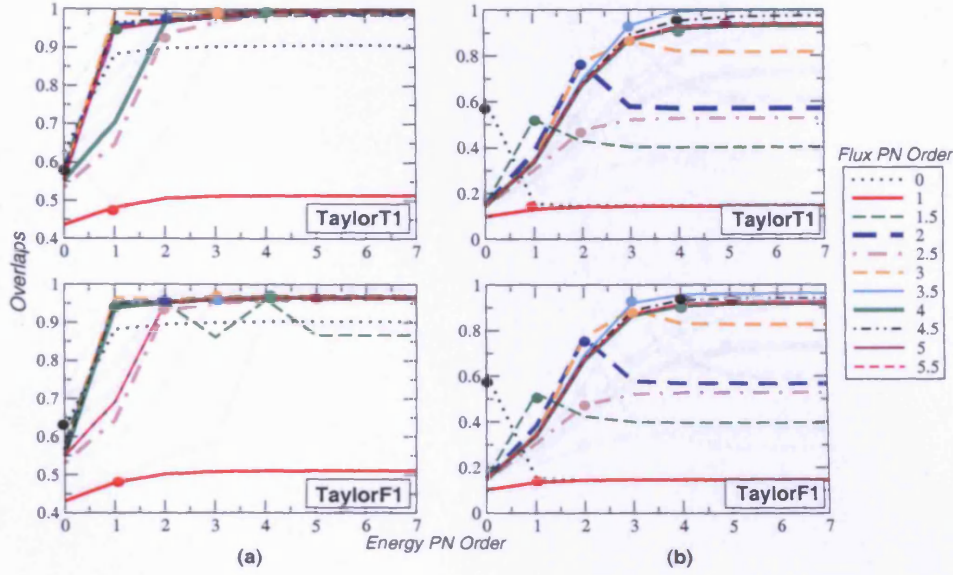


Figure 2.2: Effectualness (left panels) and faithfulness (right panels) of various test mass *TaylorT1* and *TaylorF1* templates in detecting a signal from a $(1M_{\odot}, 10M_{\odot})$ binary in white-noise. Different lines in the panels correspond to different orders of the flux function. Each line shows how the overlaps are evolving as a function of the accuracy of the energy function. Standard adiabatic approximants $T(E_{[n]}, \mathcal{F}_n)$ are marked with thick dots. All values are max-max overlaps.

the $(1M_{\odot}, 50M_{\odot})$ binary and $F_{\text{iso}} \simeq 399$ Hz for the $(1M_{\odot}, 10M_{\odot})$ binary⁷. Also, we take the noise power spectral density (PSD) of the detector to be infinite below a lower frequency cut-off $F_{\text{low}} = 20$ Hz.

As mentioned in Section 2.1.2, it is possible to create different approximants by re-expanding the functions involved in the phasing formula. However, the complete adiabatic approximant is motivated by the desire to be consistent with the underlying dynamics of the system. This consistency is achieved by identifying terms in the PN expansions of the energy and flux functions, with terms in the PN expansion of the acceleration of the system. This only has a unique meaning if the rational function E'/\mathcal{F} in the phasing formula is used without any further re-expansion. Therefore, in this study, to avoid any ambiguities arising from the use of unequal orders of the energy and flux functions in the phasing formula, we only look at the aforementioned *TaylorT1* and *TaylorF1* approximants.

⁷Here, v_{iso} is the velocity at the last stable circular orbit of Schwarzschild geometry having the same mass as the total mass $m_1 + m_2$ of the binary (we adopt units in which $c = G = 1$) and F_{iso} is the GW frequency corresponding to it.

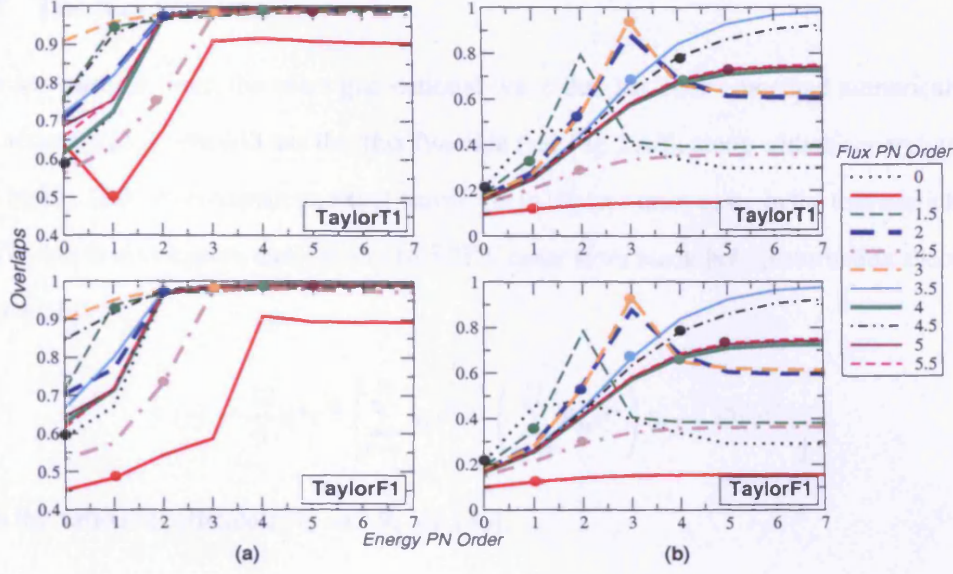


Figure 2.3: As in Fig. 2.2 except that the signal corresponds to a $(1M_{\odot}, 50M_{\odot})$ binary.

2.2.1 The energy function

In the case of a test-particle m_2 moving in circular orbit in the background of a Schwarzschild black hole of mass m_1 , where $m_2/m_1 \rightarrow 0$, the energy function $E(x)$ in terms of the invariant argument $x \equiv v^2$ is given by

$$E_{\text{exact}}(x) = \eta \frac{1-2x}{\sqrt{1-3x}}, \quad (2.8)$$

The associated v -derivative entering the phasing formula is

$$E'_{\text{exact}}(v) = 2v \left. \frac{dE(x)}{dx} \right|_{x=v^2} = -\eta v \frac{(1-6v^2)}{(1-3v^2)^{3/2}}. \quad (2.9)$$

We use the above exact $E'(v)$ to construct the exact waveform in the test-mass case. In order to construct various approximate PN templates, we Taylor-expand $E'_{\text{exact}}(v)$ and truncate it at the necessary orders.

$$E'_{7PN}(v) = -\eta v \left[1 - \frac{3v^2}{2} - \frac{81v^4}{8} - \frac{675v^6}{16} - \frac{19845v^8}{128} - \frac{137781v^{10}}{256} - \frac{1852389v^{12}}{1024} - \frac{12196899v^{14}}{2048} + O(v^{16}) \right]. \quad (2.10)$$

Different T-approximants of the energy function $E'_T(v)$ along with $E'_{\text{exact}}(v)$ are plotted in Fig. 2.1a.

2.2.2 The flux function

In the test-particle limit, the exact gravitational-wave flux has been computed numerically with good accuracy [33]. We will use this flux function (see Fig. 2.1*b*), along with the energy function given by Eq. (2.9), to construct an exact waveform in the test-mass case. In the test-particle limit, the GW flux is also known analytically to 5.5PN order from black hole perturbation theory [36] and given by

$$\mathcal{F}(v) = \frac{32}{5} \eta^2 v^{10} \left[\sum_{k=0}^{11} A_k v^k + \left(\sum_{m=6}^{11} B_m v^m \right) \ln v + O(v^{12}) \right], \quad (2.11)$$

where the various coefficients A_k and B_k are [36],

$$\begin{aligned} A_0 &= 1, \quad A_1 = 0, \quad A_2 = -\frac{1247}{336}, \quad A_3 = 4\pi, \quad A_4 = -\frac{44711}{9072}, \quad A_5 = -\frac{8191\pi}{672}, \\ A_6 &= \frac{6643739519}{69854400} + \frac{16\pi^2}{3} - \frac{1712\gamma}{105} - \frac{1712\ln 4}{105}, \quad A_7 = -\frac{16285\pi}{504}, \\ A_8 &= -\frac{323105549467}{3178375200} + \frac{232597\gamma}{4410} - \frac{1369\pi^2}{126} + \frac{39931\ln 2}{294} - \frac{47385\ln 3}{1568}, \\ A_9 &= \frac{265978667519\pi}{745113600} - \frac{6848\gamma\pi}{105} - \frac{13696\pi\ln 2}{105}, \\ A_{10} &= -\frac{2500861660823683}{2831932303200} + \frac{916628467\gamma}{7858620} - \frac{424223\pi^2}{6804} - \frac{83217611\ln 2}{1122660} + \frac{47385\ln 3}{196}, \\ A_{11} &= \frac{8399309750401\pi}{101708006400} + \frac{177293\gamma\pi}{1176} + \frac{8521283\pi\ln 2}{17640} - \frac{142155\pi\ln 3}{784}, \quad B_6 = -\frac{1712}{105}, \\ B_7 &= 0, \quad B_8 = \frac{232597}{4410}, \quad B_9 = -\frac{6848\pi}{105}, \quad B_{10} = \frac{916628467}{7858620}, \quad B_{11} = \frac{177293\pi}{1176}. \end{aligned} \quad (2.12)$$

We will use the energy and flux functions given by Eq. (2.10) – Eq. (2.12) to construct various approximate templates by truncating the expansions at the necessary order. The different T-approximants of the flux function $\mathcal{F}_T(v)$ along with the (numerical) exact flux $\mathcal{F}_{\text{exact}}(v)$ are plotted in Fig. 2.1*b*.

2.2.3 Comparison of standard and complete adiabatic approximants

We present the results of our study in the test mass limit in four parts. In the first part we discuss our conclusions on the mathematical problem of the closeness of the standard adiabatic and complete adiabatic template families with the family of exact waveforms in the adiabatic approximation. In the next part we exhibit our results in the case of the Initial LIGO, Advanced

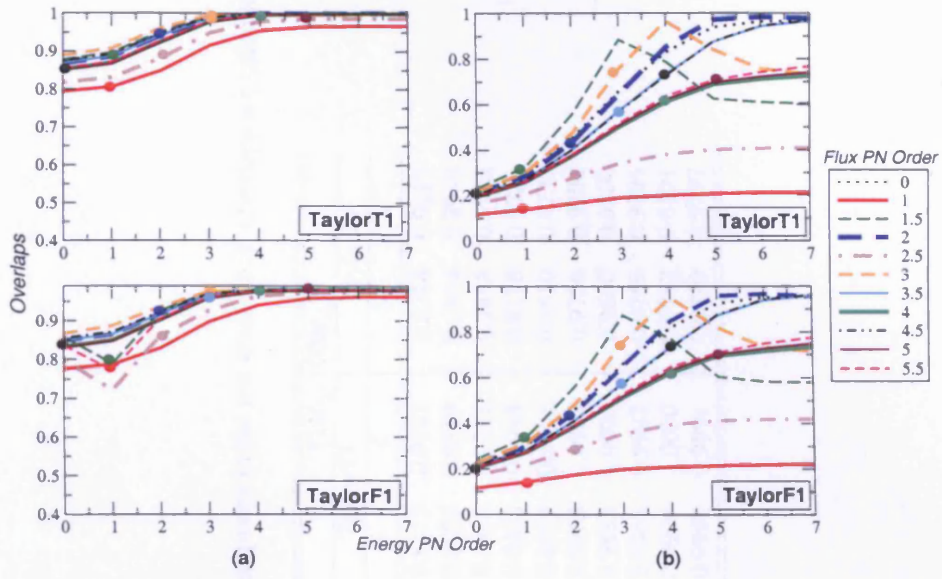


Figure 2.4: As in Fig. 2.2 except that the signal corresponds to a $(1M_{\odot}, 100M_{\odot})$ binary.

LIGO and Virgo detectors. In the third part we compare the relative importance of improving the accuracy of the energy and flux functions. Finally, in the fourth part we compare the total number of GW cycles and the number of useful cycles accumulated by various standard adiabatic and complete adiabatic approximants.

Table 2.1: Effectualness of *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ templates in the test mass limit. Overlaps are calculated assuming a flat spectrum for the detector noise (white-noise).

Order (n)	$(1M_{\odot}, 10M_{\odot})$				$(1M_{\odot}, 50M_{\odot})$				$(1M_{\odot}, 100M_{\odot})$			
	<i>TaylorTl</i>		<i>TaylorFl</i>		<i>TaylorTl</i>		<i>TaylorFl</i>		<i>TaylorTl</i>		<i>TaylorFl</i>	
	S	C	S	C	S	C	S	C	S	C	S	C
0PN	0.6250	0.8980	0.6212	0.8949	0.5809	0.9726	0.5917	0.9644	0.8515	0.9231	0.8318	0.9017
1PN	0.4816	0.5119	0.4801	0.5086	0.4913	0.9107	0.4841	0.5871	0.8059	0.9169	0.7874	0.8980
1.5PN	0.9562	0.9826	0.9448	0.9592	0.9466	0.9832	0.9370	0.9785	0.8963	0.9981	0.7888	0.9788
2PN	0.9685	0.9901	0.9514	0.9624	0.9784	0.9917	0.9719	0.9872	0.9420	0.9993	0.9178	0.9785
2.5PN	0.9362	0.9924	0.9298	0.9602	0.7684	0.9833	0.7326	0.9772	0.8819	0.9858	0.8610	0.9730
3PN	0.9971	0.9991	0.9677	0.9713	0.9861	0.9946	0.9821	0.9886	0.9965	0.9959	0.9756	0.9792
3.5PN	0.9913	0.9996	0.9636	0.9688	0.9902	0.9994	0.9858	0.9914	0.9885	1.0000	0.9690	0.9800
4PN	0.9937	0.9973	0.9643	0.9663	0.9975	0.9996	0.9903	0.9914	0.9968	0.9992	0.9769	0.9795
4.5PN	0.9980	0.9999	0.9671	0.9690	0.9967	1.0000	0.9902	0.9913	0.9996	1.0000	0.9787	0.9801
5PN	0.9968	0.9979	0.9661	0.9667	0.9994	0.9994	0.9913	0.9914	0.9992	0.9991	0.9790	0.9797

Table 2.2: Faithfulness of *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ templates in the test mass limit. Overlaps are calculated assuming a flat spectrum for the detector noise (white-noise).

Order (n)	$(1M_\odot, 10M_\odot)$				$(1M_\odot, 50M_\odot)$				$(1M_\odot, 100M_\odot)$			
	<i>TaylorTl</i>		<i>TaylorFl</i>		<i>TaylorTl</i>		<i>TaylorFl</i>		<i>TaylorTl</i>		<i>TaylorFl</i>	
	S	C	S	C	S	C	S	C	S	C	S	C
0PN	0.6124	0.1475	0.6088	0.1446	0.2045	0.4683	0.2104	0.4750	0.2098	0.4534	0.2208	0.4641
1PN	0.1322	0.1433	0.1350	0.1461	0.1182	0.1446	0.1236	0.1508	0.1395	0.1901	0.1432	0.1994
1.5PN	0.5227	0.4005	0.5241	0.3967	0.3444	0.3947	0.3505	0.3866	0.3260	0.7869	0.3399	0.7700
2PN	0.7687	0.5707	0.7680	0.5689	0.5518	0.6871	0.5535	0.6827	0.4377	0.8528	0.4506	0.8486
2.5PN	0.4735	0.5268	0.4748	0.5278	0.2874	0.3561	0.2933	0.3625	0.2787	0.4001	0.2918	0.4133
3PN	0.8629	0.8165	0.8932	0.8277	0.9420	0.6317	0.9334	0.6222	0.7579	0.8407	0.7570	0.8194
3.5PN	0.9309	0.9979	0.9194	0.9609	0.6689	0.9681	0.6695	0.9632	0.5740	0.9425	0.5805	0.9383
4PN	0.9174	0.9303	0.9087	0.9176	0.6693	0.7227	0.6701	0.7230	0.6129	0.7112	0.6236	0.7159
4.5PN	0.9525	0.9744	0.9330	0.9415	0.7829	0.9242	0.7827	0.9229	0.7286	0.9689	0.7308	0.9632
5PN	0.9370	0.9392	0.9225	0.9241	0.7275	0.7417	0.7276	0.7420	0.6972	0.7409	0.7027	0.7500

White-noise case

First, we explore the general question of the closeness of the standard adiabatic and complete adiabatic templates to the exact waveform assuming flat power spectral density for the detector noise. Fig. 2.2–2.4 show the effectualness and faithfulness of various PN templates for three archetypal binaries with component masses $(1M_{\odot}, 10M_{\odot})$, $(1M_{\odot}, 50M_{\odot})$ and $(1M_{\odot}, 100M_{\odot})$, respectively⁸.

The central result of this study is that *complete adiabatic approximants bring about a remarkable improvement in the effectualness for all systems at low PN orders ($< 3PN$)*. The complete adiabatic approximants converge to the adiabatic exact waveform at lower PN orders than the standard adiabatic approximants. This indicates that at these orders general relativistic corrections to the conservative dynamics of the binary are quite important contrary to the assumption employed in the standard post-Newtonian treatment of the phasing formula. On the other hand, the difference in effectualness between the standard and complete adiabatic approximants at orders greater than 3PN is very small. Thus, *if we have a sufficiently accurate (order $\geq 3PN$) T -approximant of the flux function, the standard adiabatic approximation is nearly as good as the complete adiabatic approximation for construction of effectual templates*. Thus at all orders the standard adiabatic approximation provides a good lower bound to the complete adiabatic approximation for the construction of *effectual* templates.

The faithfulness of both the approximants fluctuates as we go from one PN order to the next and is generally much smaller than our target value of 0.965. The fluctuation continues all the way up to 5PN order reflecting the oscillatory approach of the flux function to the exact flux function at different PN orders. It is again interesting to note that complete adiabatic approximants are generally more faithful than the standard adiabatic approximants. It is certainly worth exploring, in a future study, the anomalous cases where it performs worse than the standard.

Noise spectra of specific detectors

Having addressed the general question concerning the closeness of standard adiabatic and complete adiabatic templates to the exact waveforms, we now compare the overlaps in the case of specific detectors. The effectualness and faithfulness of various test mass PN templates for the

⁸For the sake of convenience we also tabulate the results shown in Fig. 2.2–2.4 in Tables 2.1 and 2.2.

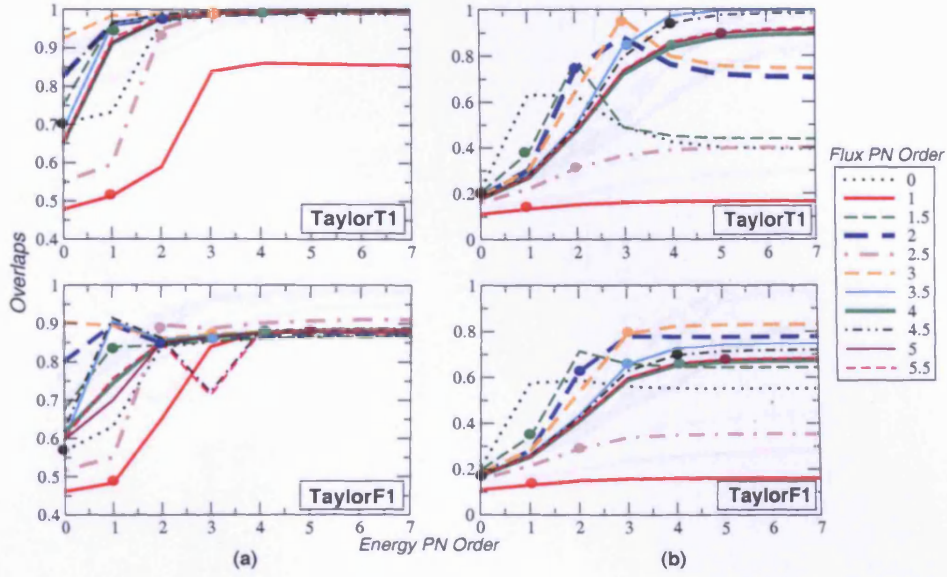


Figure 2.5: Effectualness (left panels) and faithfulness (right panels) of various test mass *TaylorT1* and *TaylorF1* templates in detecting a signal from a $(1M_{\odot}, 10M_{\odot})$, calculated for the Initial LIGO noise PSD. Different lines in the panels correspond to different orders of the flux function. Each line shows how the overlaps are evolving as a function of the accuracy of the energy function. Standard adiabatic approximants $T(E_{[n]}, \mathcal{F}_n)$ are marked with thick dots. All values are max-max overlaps.

$(1M_{\odot}, 10M_{\odot})$ binary and $(1M_{\odot}, 50M_{\odot})$ binary are plotted in Fig. 2.5 and Fig. 2.6, respectively, and are shown in Tables 2.3, 2.4 and 2.5. As in the case of white noise, here too we see that standard adiabatic approximants of order less than 3PN have considerably lower overlaps than the corresponding complete adiabatic approximants and the difference in overlaps between standard adiabatic and complete adiabatic approximants of order greater than 3PN is very small. Thus, if we have a sufficiently accurate (order ≥ 3 PN) T-approximant of the flux function, we conclude that the standard adiabatic approximation provides a good lower bound to the complete adiabatic approximation for the construction of effectual templates. The plots and Table 2.5 indicate that the *faithfulness of PN templates greatly improves in a complete adiabatic treatment* for all orders studied in the case of Initial LIGO, and for the vast majority of orders in the case of Advanced LIGO and Virgo. However, for the latter two detectors, as in the white noise case, there were certain orders for which this was not the case.

We also calculate the bias in the estimation of parameters while maximizing the overlaps over the intrinsic parameters of the binary. The (percentage) bias in the estimation of the param-

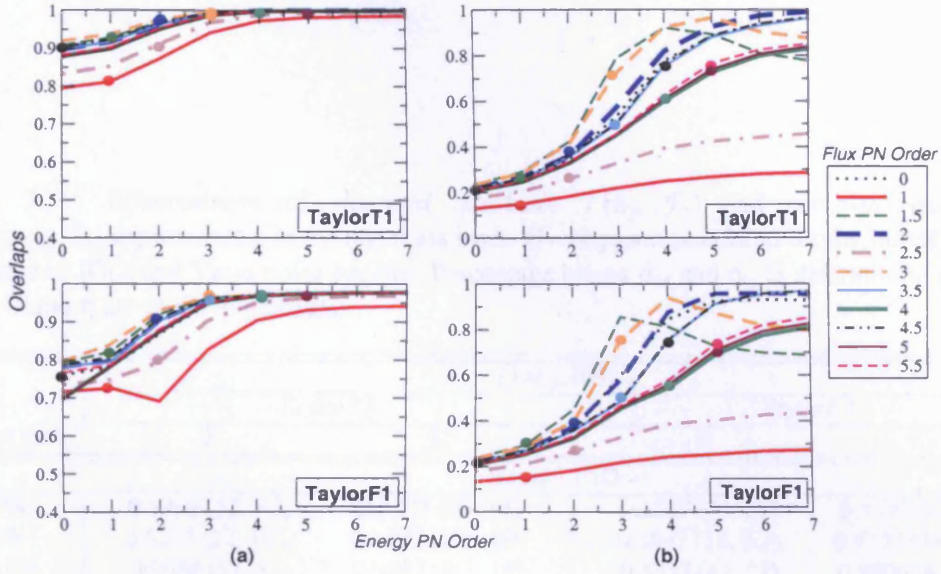


Figure 2.6: As in Fig. 2.5 except that the signal corresponds to a $(1M_{\odot}, 50M_{\odot})$ binary.

eter p is defined as

$$\sigma_p \equiv \frac{|p_{\max} - p|}{p} \times 100, \quad (2.13)$$

where p_{\max} is the value of the parameter p which gives the maximum overlap. Along with the maximized overlaps (effectualness), the bias in the estimation of the parameters M and η are also quoted in Tables 2.3, 2.4. It can be seen that at lower PN orders (order < 3 PN) the complete adiabatic approximants generally show significantly lower biases, although there are some anomalous cases where the standard adiabatic approximant shows lower biases than the complete adiabatic approximant. Even at higher PN orders complete adiabatic approximants are generally less-biased than the corresponding standard adiabatic approximants.

Accuracy of energy function Vs. flux function

In most of the cases, TaylorT1 and TaylorF1 templates show trends of smoothly increasing overlaps as the accuracy of the energy function is increased keeping the accuracy of the flux function constant. This is because the T-approximants of the energy function smoothly converge to the exact energy as we go to higher orders (see Fig. 2.1a). On the other hand, if we improve the accuracy of the flux function for a fixed order of energy, the overlaps do not show such a smoothly converging behavior. This can be understood in terms of the oscillatory nature of the T-approximants of the flux function. For example, templates constructed from 1PN and 2.5PN

Table 2.3: Effectualness of *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants in the test mass limit. Overlaps are calculated for the Initial LIGO, Advanced LIGO, and Virgo noise spectra. Percentage biases σ_M and σ_η in determining parameters M and η are given in brackets.

Order (n)	$(1M_\odot, 10M_\odot)$			
	<i>TaylorTl</i>		<i>TaylorFl</i>	
	S	C	S	C
Initial LIGO				
0PN	0.5910 (12, 5.7)	0.9707 (36, 45)	0.5527 (31, 28)	0.8395 (48, 53)
1PN	0.5232 (22, 105)	0.8397 (125, 69)	0.4847 (18, 9.7)	0.8393 (147, 74)
1.5PN	0.9688 (52, 51)	0.9887 (8.3, 15)	0.8398 (61, 57)	0.8606 (4.7, 10)
2PN	0.9781 (18, 25)	0.9942 (0.4, 0.6)	0.8485 (32, 40)	0.8693 (15, 22)
2.5PN	0.9490 (96, 68)	0.9923 (26, 32)	0.8963 (123, 75)	0.9071 (49, 50)
3PN	0.9942 (0.3, 1.1)	0.9989 (3.7, 6.2)	0.8713 (16, 23)	0.8822 (12, 18)
3.5PN	0.9940 (6.9, 11)	0.9998 (0.6, 1.4)	0.8685 (23, 31)	0.8834 (17, 25)
4PN	0.9974 (6.2, 11)	0.9996 (3.9, 6.9)	0.8746 (23, 30)	0.8817 (21, 28)
4.5PN	0.9988 (3.3, 5.5)	1.0000 (0.8, 1.6)	0.8795 (19, 27)	0.8868 (18, 26)
5PN	0.9992 (4.0, 6.9)	0.9997 (3.5, 5.7)	0.8792 (21, 29)	0.8825 (20, 28)
Advanced LIGO				
0PN	0.4281(9.18, 2.63)	0.8960(31.7, 42.1)	0.4061(8.55, 4.62)	0.7264(32.5, 42.4)
1PN	0.3498(28.2, 8.90)	0.7258(156, 75.2)	0.3378(8.82, 17.1)	0.7336(166, 76.9)
1.5PN	0.9010(47.9, 48.8)	0.9653(11.1, 20.8)	0.7496(50.9, 50.5)	0.7704(2.91, 3.61)
2PN	0.9266(13.5, 19.9)	0.9814(2.55, 4.18)	0.7529(16.2, 23.2)	0.7767(9.09, 15.2)
2.5PN	0.8917(88.5, 66.0)	0.9913(25.7, 31.7)	0.7209(76.4, 62.0)	0.8180(43.2, 45.8)
3PN	0.9913(0.73, 1.56)	0.9989(3.91, 7.34)	0.7842(10.1, 15.9)	0.8023(4.45, 7.35)
3.5PN	0.9816(4.45, 7.35)	0.9994(0.36, 0.33)	0.7788(14.1, 20.3)	0.7996(9.00, 14.3)
4PN	0.9895(4.18, 7.13)	0.9970(3.00, 5.28)	0.7842(14.3, 21.2)	0.7944(11.4, 16.8)
4.5PN	0.9965(2.09, 3.62)	0.9999(0.82, 1.63)	0.7943(10.7, 16.4)	0.8031(10.1, 15.1)
5PN	0.9954(2.91, 5.20)	0.9977(2.55, 4.01)	0.7924(11.4, 16.8)	0.7958(11.5, 16.9)
Virgo				
0PN	0.3894(42.0, 41.2)	0.7256(0.82, 3.81)	0.3699(10.7, 11.3)	0.7275(16.6, 27.6)
1PN	0.2956(10.9, 6.51)	0.6876(187, 80.1)	0.2926(1.18, 21.5)	0.6951(191, 80.4)
1.5PN	0.8474(31.2, 36.6)	0.9487(11.5, 22.2)	0.7688(31.5, 36.8)	0.8079(10.2, 18.8)
2PN	0.8933(9.91, 14.9)	0.9711(3.00, 4.61)	0.7855(10.0, 15.0)	0.8197(0.73, 0.66)
2.5PN	0.8179(68.7, 58.5)	0.9864(25.6, 31.7)	0.7581(67.6, 58.2)	0.8375(30.9, 36.5)
3PN	0.9845(0.64, 1.48)	0.9970(3.82, 7.25)	0.8312(2.18, 3.70)	0.8441(0.09, 0.08)
3.5PN	0.9722(4.27, 7.20)	0.9991(0.36, 0.33)	0.8222(6.73, 10.9)	0.8397(4.00, 6.98)
4PN	0.9829(4.09, 7.06)	0.9955(2.91, 5.20)	0.8266(6.73, 10.9)	0.8339(6.27, 9.66)
4.5PN	0.9937(2.00, 3.54)	0.9999(0.82, 1.63)	0.8343(5.45, 9.02)	0.8418(5.00, 7.79)
5PN	0.9920(3.00, 5.28)	0.9967(2.64, 4.09)	0.8323(6.27, 9.66)	0.8353(6.00, 9.45)

Table 2.4: Same as Table 2.3 except that the values corresponds to the $(1M_{\odot}, 50M_{\odot})$ binary.

Order (n)	$(1M_{\odot}, 50M_{\odot})$			
	<i>TaylorTl</i>		<i>TaylorFl</i>	
	<i>S</i>	<i>C</i>	<i>S</i>	<i>C</i>
Initial LIGO				
0PN	0.8748 (24, 29)	0.9471 (19, 14)	0.8294 (21, 34)	0.8974 (17, 13)
1PN	0.8101 (28, 104)	0.9392 (19, 40)	0.7662 (23, 116)	0.8898 (18, 43)
1.5PN	0.9254 (21, 4.1)	0.9996 (6.7, 20)	0.8772 (18, 0.2)	0.9590 (6.4, 20)
2PN	0.9610 (18, 16)	0.9993 (7.5, 16)	0.9113 (16, 14)	0.9583 (7.7, 17)
2.5PN	0.9104 (21, 6.9)	0.9940 (8.3, 0.7)	0.8630 (19, 8.7)	0.9574 (9.1, 1.9)
3PN	0.9968 (11, 21)	0.9992 (2.6, 10)	0.9500 (11, 21)	0.9648 (2.7, 11)
3.5PN	0.9923 (13, 19)	0.9997 (2.4, 5.2)	0.9445 (12, 18)	0.9679 (2.8, 6.5)
4PN	0.9979 (8.8, 13)	0.9995 (3.5, 4.3)	0.9560 (8.9, 14)	0.9672 (3.9, 5.6)
4.5PN	0.9995 (7.1, 14)	1.0000 (0.9, 1.9)	0.9590 (7.0, 14)	0.9698 (1.5, 3.4)
5PN	0.9994 (5.2, 7.7)	0.9990 (2.6, 2.4)	0.9634 (5.9, 10)	0.9690 (3.4, 5.1)
Advanced LIGO				
0PN	0.6461(26.6, 21.9)	0.8099(47.8, 53.5)	0.6314(20.1, 34.7)	0.7187(31.7, 38.8)
1PN	0.6200(25.0, 123)	0.7093(27.2, 12.9)	0.6039(20.6, 129)	0.6820(41.4, 14.9)
1.5PN	0.6919(27.3, 19.6)	0.9532(2.02, 8.68)	0.6759(21.7, 11.0)	0.9567(1.16, 6.94)
2PN	0.8835(30.7, 39.1)	0.9833(6.29, 13.2)	0.7915(26.4, 33.9)	0.9751(5.02, 10.3)
2.5PN	0.6720(25.9, 6.17)	0.9194(17.2, 21.2)	0.6573(20.2, 3.10)	0.9432(18.8, 23.1)
3PN	0.9645(8.43, 15.8)	0.9740(1.39, 1.38)	0.9630(7.20, 13.9)	0.9738(1.47, 1.46)
3.5PN	0.9875(14.0, 22.5)	0.9987(2.00, 3.85)	0.9693(13.6, 22.2)	0.9854(1.71, 3.57)
4PN	0.9967(9.49, 15.7)	0.9973(4.37, 6.93)	0.9799(9.41, 15.7)	0.9859(4.67, 8.13)
4.5PN	0.9932(6.14, 11.2)	1.0000(0.92, 1.87)	0.9799(5.37, 9.67)	0.9868(0.80, 1.75)
5PN	0.9986(5.69, 9.93)	0.9960(3.55, 6.21)	0.9843(5.82, 10.0)	0.9866(3.86, 6.48)
Virgo				
0PN	0.6004(50.4, 25.2)	0.8689(50.5, 55.7)	0.5878(44.2, 17.2)	0.8577(50.8, 55.7)
1PN	0.5498(51.4, 30.0)	0.7217(51.6, 28.4)	0.5328(52.6, 23.9)	0.7190(59.3, 36.1)
1.5PN	0.7308(56.0, 52.7)	0.9619(1.06, 6.85)	0.7652(68.1, 60.3)	0.9664(0.67, 5.52)
2PN	0.9291(34.4, 42.9)	0.9854(5.39, 11.6)	0.9141(34.4, 42.9)	0.9829(4.82, 10.1)
2.5PN	0.6579(48.5, 40.9)	0.9446(19.1, 23.2)	0.6551(57.1, 48.6)	0.9591(19.5, 24.3)
3PN	0.9697(7.41, 14.1)	0.9818(1.47, 1.46)	0.9725(6.71, 12.6)	0.9824(1.78, 1.78)
3.5PN	0.9885(13.7, 22.3)	0.9980(1.90, 3.76)	0.9794(13.3, 21.1)	0.9912(1.53, 3.41)
4PN	0.9971(9.45, 15.7)	0.9973(4.35, 6.92)	0.9878(9.41, 15.7)	0.9915(4.71, 8.16)
4.5PN	0.9926(6.04, 11.2)	1.0000(0.90, 1.85)	0.9874(5.25, 9.57)	0.9920(0.82, 1.77)
5PN	0.9987(5.73, 9.96)	0.9960(3.53, 6.19)	0.9906(5.84, 10.1)	0.9918(3.98, 6.59)

Table 2.5: Faithfulness of *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ templates in the test mass limit. Overlaps are calculated for the Initial LIGO, Advanced LIGO and Virgo noise spectra.

Order (n)	$(1M_{\odot}, 10M_{\odot})$				$(1M_{\odot}, 50M_{\odot})$			
	<i>TaylorTl</i>		<i>TaylorFl</i>		<i>TaylorTl</i>		<i>TaylorFl</i>	
	S	C	S	C	S	C	S	C
Initial LIGO								
0PN	0.2186	0.6272	0.2108	0.5879	0.2134	0.3498	0.2145	0.3593
1PN	0.1342	0.1615	0.1308	0.1563	0.1511	0.2196	0.1527	0.2210
1.5PN	0.3788	0.4492	0.3449	0.6471	0.2915	0.9223	0.2956	0.9195
2PN	0.7449	0.7633	0.6279	0.7782	0.3613	0.8157	0.3674	0.8318
2.5PN	0.3115	0.3970	0.2905	0.3532	0.2608	0.4233	0.2606	0.4161
3PN	0.9633	0.7566	0.7913	0.8297	0.7194	0.9686	0.7057	0.9323
3.5PN	0.8385	0.9984	0.6582	0.7464	0.4941	0.9273	0.5046	0.9442
4PN	0.8356	0.8909	0.6527	0.6725	0.5960	0.7934	0.5864	0.8131
4.5PN	0.9395	0.9851	0.6967	0.7195	0.7594	0.9644	0.7605	0.9614
5PN	0.8960	0.9129	0.6770	0.6821	0.7344	0.8350	0.7432	0.8579
Advanced LIGO								
0PN	0.1456	0.4915	0.1454	0.4139	0.1608	0.2955	0.1630	0.2929
1PN	0.0853	0.1041	0.0821	0.0985	0.1159	0.1609	0.1169	0.1623
1.5PN	0.2711	0.3063	0.2626	0.4711	0.2187	0.6735	0.2212	0.6677
2PN	0.6998	0.6140	0.6280	0.6407	0.2765	0.8403	0.2768	0.8406
2.5PN	0.2143	0.2710	0.2108	0.2623	0.1961	0.3094	0.1985	0.3120
3PN	0.8889	0.5791	0.7593	0.7473	0.7252	0.6971	0.7257	0.6820
3.5PN	0.7476	0.9985	0.6210	0.7196	0.3852	0.9087	0.3868	0.9215
4PN	0.7314	0.8144	0.6099	0.6290	0.4404	0.5761	0.4414	0.5754
4.5PN	0.9001	0.9718	0.6564	0.6798	0.5714	0.9078	0.5657	0.9190
5PN	0.8273	0.8518	0.6357	0.6400	0.5303	0.6166	0.5260	0.6232
Virgo								
0PN	0.1384	0.3644	0.1376	0.3188	0.1265	0.3881	0.1270	0.3900
1PN	0.0682	0.0818	0.0679	0.0813	0.0887	0.1205	0.0891	0.1211
1.5PN	0.2524	0.2348	0.2517	0.3575	0.1859	0.5783	0.1858	0.5730
2PN	0.7451	0.4617	0.6941	0.5536	0.2514	0.8597	0.2525	0.8588
2.5PN	0.2003	0.2496	0.1984	0.2469	0.1612	0.2420	0.1602	0.2412
3PN	0.8339	0.5745	0.8012	0.7451	0.7978	0.6210	0.7990	0.6131
3.5PN	0.7684	0.9968	0.6901	0.7842	0.3821	0.9259	0.3836	0.9348
4PN	0.7501	0.7892	0.6778	0.6981	0.4024	0.5306	0.4030	0.5295
4.5PN	0.8753	0.9595	0.7239	0.7458	0.5298	0.9132	0.5309	0.9225
5PN	0.8033	0.8232	0.7048	0.7093	0.4968	0.5617	0.4962	0.5626

flux functions can be seen to have considerably lower overlaps than the other ones. This is because of the poor ability of the 1PN and 2.5PN T-approximants to mimic the behavior of the exact flux function (see Fig. 2.1b). This inadequacy of the 1PN and 2.5PN T-approximants is prevalent in both test mass and comparable mass cases. Hence it is not a good strategy to use the T-approximants at these orders for the construction of templates. On the other hand, 3.5PN and 4.5PN T-approximants are greatly successful in following the exact flux function in the test mass case, and consequently lead to larger overlaps.

We have found that in the test mass case if we improve the accuracy of energy function from 3PN to 4PN, keeping the flux function at order 3PN, the increase in effectualness (respectively, faithfulness) is $\simeq 0.32\%$ (-16%). The same improvement in the energy function for the 3.5PN flux will produce an increase of $\simeq 5.9\%$ (-2.0%). On the other hand, if we improve the accuracy of flux function from 3.5PN to 4PN, keeping the energy function at order 3PN, the increase in effectualness (respectively, faithfulness) is $\simeq -0.070\%$ (-12%). The values quoted are calculated using the TaylorT1 approximant for the $(1M_{\odot}, 10M_{\odot})$ binary for the Initial LIGO noise PSD. *If the comparable mass case is qualitatively similar to the test mass case, this should imply that neither the improvement in the accuracy of the energy function from 3PN to 4PN nor the improvement in the accuracy of the flux function from 3.5PN to 4PN will produce significant improvement in the overlaps in the comparable mass case.* The trends in the faithfulness are very different for different binaries, so it is difficult to make any statement about the improvement in faithfulness.

Number of gravitational wave cycles

The number of GW cycles accumulated by a template is defined as [71]

$$\mathcal{N}_{\text{tot}} \equiv \frac{1}{2\pi} (\varphi_{\text{iso}} - \varphi_{\text{low}}) = \int_{F_{\text{low}}}^{F_{\text{iso}}} dF \frac{N(F)}{F}, \quad (2.14)$$

where φ_{iso} and φ_{low} are the GW phases corresponding to the last stable orbit and the low frequency cutoff, respectively, and $N(F) \equiv F^2/\dot{F}$ is the *instantaneous number of cycles* spent near some instantaneous frequency F (as usual, \dot{F} is the time derivative of F). However, it has been noticed that [71], the large number \mathcal{N}_{tot} is not significant because the only really *useful* cycles are those that contribute most to the signal-to-noise ratio (SNR). The *useful cycles* are defined

as [71]

$$\mathcal{N}_{\text{useful}} \equiv \left(\int_{F_{\text{low}}}^{F_{\text{Iso}}} \frac{df}{f} w(f) N(f) \right) \left(\int_{F_{\text{low}}}^{F_{\text{Iso}}} \frac{df}{f} w(f) \right)^{-1}, \quad (2.15)$$

where $w(f) \equiv a^2(f)/h_n^2(f)$. If $S_n(f)$ is the two-sided PSD of the detector noise, $h_n(f)$ is defined by $h_n^2(f) \equiv f S_n(f)$, while $a(f)$ is defined by $|H(f)| \simeq a(t_f)/[\dot{F}(t_f)]^{1/2}$ where $H(f)$ is the Fourier transform of the time-domain waveform $h(t)$ (See Eqs.(2.2) and (2.3)) and t_f is the time when the instantaneous frequency $F(t)$ reaches the value f of the Fourier variable.

The total numbers of GW cycles accumulated by various standard adiabatic $T(E_{[n]}, \mathcal{F}_n)$ and complete adiabatic $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants in the test mass limit are tabulated in Table 2.6 along with the number of useful cycles calculated for the Initial LIGO noise PSD. We use Eq.(2.4) to calculate \dot{F} and numerically evaluate the integrals in Eq.(2.15) to compute the number of useful cycles. In order to compute the total number of cycles, we numerically integrate the integral in Eq.(2.14) ⁹.

It can be seen that all complete adiabatic approximants accumulate fewer number of (total and useful) cycles than the corresponding standard adiabatic approximants. This is because the additional conservative terms in the complete adiabatic approximants add extra acceleration to the test mass which, in the presence of radiation reaction, would mean that the test body has to coalesce faster, and therefore such templates accumulate fewer number of cycles. Noticeably enough, approximants (like 3PN and 4.5PN) producing the highest overlaps with the exact waveform, accumulate the closest number of cycles as accumulated by the exact waveform. This is indicative that the phase evolution of these approximants is closer to that of the exact waveform. On the other hand, the fractional absolute difference in the number of cycles of the approximants producing the lowest overlaps (like 0PN, 1PN and 2.5PN) as compared to the exact waveform is the greatest, which indicates that these templates follow a completely different phase evolution.

In order to illustrate the correlation between the number of (total/useful) cycles accumulated by an approximant and its overlap with the exact waveform, we introduce a quantity $\delta\mathcal{N}_{\mathcal{G}} = \frac{|\mathcal{N}_{\mathcal{G}} - \mathcal{N}_{\text{exact}}|}{\mathcal{N}_{\text{exact}}}$ which is the fractional absolute difference between the number of (total/useful) cycles

⁹Note that the method we are using to calculate the number of cycles is different from the method used by some authors, for example, Blanchet. et. al [54] who use a PN expansion of the GW phasing to calculate the ϕ_{Iso} and ϕ_{low} appearing in Eq.(2.14). This method is not appropriate for our study as it involves further re-expansion of $E'(v)/\mathcal{F}(v)$. Since different approaches to the phasing formula can produce slightly different phasings, one should not be surprised to see any differences in the results obtained by these different methods.

accumulated by a template and the exact waveform. Here \mathcal{N}_b and $\mathcal{N}_{\text{exact}}$ are the number of (total/useful) cycles accumulated by the n PN approximant and exact waveform, respectively. In Fig. 2.7, we plot $\delta\mathcal{N}_b$ of various standard adiabatic and complete adiabatic approximants against the corresponding overlaps in the case of a $(1M_\odot, 10M_\odot)$ binary.

The following points may be noted when comparing the results quoted here for the number of cycles with those of other works eg. Refs [54, 55, 56]. As emphasized in Ref. [23] one can get very different results for the phasing depending on whether one consistently re-expands the constituent energy and flux functions or evaluates them without re-expansion. In the computation of the number of useful cycles different authors treat the function \dot{F} differently; some re-expand, whereas others do not, leading to differences in the results. The other important feature we would like to comment on is a result which appears, at first, counter-intuitive. This is the fact that, in some cases, the number of useful GW cycles is greater than the total number of GW cycles. A closer examination reveals that, while for most cases of interest this does not happen, in principle its occurrence is determined by the ratio $f_r \equiv F_{\text{low}}/F_{\text{iso}}$. To understand this in more detail let us consider the ratio \mathcal{N}_c of the number of useful cycles to the total number of cycles in the case of white noise (in a frequency band F_{low} to F_{iso} , for which

$$\mathcal{N}_c \equiv \frac{\mathcal{N}_{\text{useful}}}{\mathcal{N}_{\text{tot}}} = \frac{5}{12} f_r^{1/3} \frac{1 - f_r^{4/3}}{1 - f_r^{1/3}} \quad (2.16)$$

For $f_r \ll 1$, $\mathcal{N}_c < 1$. However, as f_r increases to about $f_r = 0.52$, \mathcal{N}_c transits from being less than one to becoming greater than one. Essentially this arises due to the details of the scalings of the various quantities involved and the point of transition depends on the PN order and the precise form of the noise PSD. For $f_r \simeq 1$, the calculation of useful cycles does not make much physical sense. This explains the absence of $\mathcal{N}_{\text{useful}}$ results for the $(1M_\odot, 10M_\odot)$ binary in table 2.6.

2.3 Non-adiabatic models

Before introducing new non-adiabatic models in this section, let us recapitulate our point of view in summary. Contrary to the standard adiabatic approximant which is constructed from consistent PN expansions of the energy and flux functions to the same relative PN order, we

Table 2.6: Number of GW cycles accumulated by various *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants in the test mass limit. The number of useful cycles calculated for the Initial LIGO noise PSD is also shown in brackets. We choose a low frequency cutoff of 20Hz.

Order (n)	$(1M_{\odot}, 10M_{\odot})$		$(1M_{\odot}, 50M_{\odot})$		$(1M_{\odot}, 100M_{\odot})$	
	S	C	S	C	S	C
0PN	481 (92.3)	424 (74.6)	118 (110)	77.8 (64.4)	13.6	6.7
1PN	560 (117)	526 (102)	180 (186)	124 (104)	25.7	10.6
1.5PN	457 (81.7)	433 (71.8)	88.8 (76.3)	58.5 (38.2)	8.4	2.3
2PN	447 (77.7)	440 (74.0)	77.0 (61.8)	62.5 (41.5)	6.1	2.6
2.5PN	464 (84.5)	454 (79.6)	96.8 (85.5)	74.5 (50.5)	9.7	2.9
3PN	442 (74.7)	440 (73.3)	64.5 (45.2)	58.1 (35.5)	3.4	1.6
3.5PN	445 (76.1)	442 (74.5)	68.7 (49.7)	60.6 (36.8)	4.0	1.4
4PN	445 (75.8)	443 (75.2)	66.4 (45.1)	62.9 (39.0)	2.9	1.6
4.5PN	443 (75.1)	442 (74.5)	63.7 (42.0)	60.0 (35.6)	2.5	1.2
5PN	444 (75.3)	443 (75.0)	63.8 (40.9)	62.2 (37.8)	2.1	1.4
Exact	442 (74.1)		59.1 (34.3)		0.9	

considered a new complete adiabatic approximant (still based on PN expansions of the energy and flux functions but of different PN orders) but consistent with a complete PN acceleration. Viewed in terms of the acceleration terms they include, the standard adiabatic approximation is inconsistent by neglect of some intermediate PN order terms in the acceleration. The complete adiabatic approximation on the other hand is constructed to consistently include all the relevant PN acceleration terms neglected in the associated standard approximant. These models were a prelude to phasing models constructed from the dynamical equations of motion considered in this section. However, we have worked solely within the adiabatic approximation. It is then pertinent to ask whether one can construct natural non-adiabatic extensions of both the standard and complete adiabatic approximants. And if so, how do their performances compare? Indeed, the work of Buonanno and Damour [24] within the effective one-body approach to the dynamics did find differences between the adiabatic and non-adiabatic solutions. In this Section we investigate whether it is possible to introduce non-adiabatic formulations of the standard and complete approximants considered in the previous Section.

The Lagrangian models studied by Buonanno, Chen and Vallisneri [65] seem to be the natural candidates for the purpose since they are specified by the acceleration experienced by the binary system. The Lagrangian models considered in Ref. [65] can be thought of as the *standard non-adiabatic* approximants, since, following standard choices, they lead to gaps in




Figure 2.7: The fractional absolute difference $\delta \mathcal{N}_b$ between the number of cycles accumulated by various approximants and the exact waveform, plotted against the corresponding overlaps. *Standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ approximants are marked with lighter markers and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants are marked with darker markers. Top panels show $\delta \mathcal{N}_b$ for the total number of cycles and bottom panels show $\delta \mathcal{N}_b$ for the number of useful cycles. The number of useful cycles are calculated for the Initial LIGO noise PSD and the low frequency cutoff is chosen to be 20Hz. The plotted results are for a $(1M_\odot, 10M_\odot)$ binary.

the post-Newtonian expansion of the acceleration. Generalizing these Lagrangian models so that there are no missing PN terms, or gaps, in the acceleration we can construct the *complete non-adiabatic* approximants. With a non-adiabatic variant of the standard and complete approximants we can then compare their relative performances. However, we will be limited in this investigation because of two reasons: Firstly, the Lagrangian models are available only up to 3.5PN order, and higher order PN accelerations are as yet unavailable. Secondly, the only *exact* waveform we have has, however, been constructed only in the *adiabatic* approximation. Even in the test mass limit, the exact waveform is not known beyond the adiabatic approximation. Due to lack of anything better, we continue to use the exact waveform in the adiabatic approximation to measure the effectualness and faithfulness of the non-adiabatic approximants.

Thus to obtain non-adiabatic approximants, the signal is constructed by integrating the equations of motion directly using a Lagrangian formalism. The equations are schematically written as:

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}, \quad \frac{d\mathbf{v}}{dt} = \mathbf{a} \quad (2.17)$$

For the complete non-adiabatic model of order n , all terms in the PN expansion for acceleration are retained up to order $n + 2.5$ without any gaps. For the standard non-adiabatic models, on the other hand, only terms in the acceleration consistent with the treatment of standard phasing are retained in the acceleration, resulting in gaps in the acceleration corresponding to intermediate PN terms neglected in the treatment. E.g. for $n = 0$ the standard non-adiabatic approximation includes only the \mathbf{a}_N and $\mathbf{a}_{2.5PN}$ while the complete non-adiabatic approximation would include in addition the \mathbf{a}_{1PN} and \mathbf{a}_{2PN} ¹⁰. Given the current status of knowledge of the two-body equations of motion, we have only two complete non-adiabatic approximants, at 0PN and 1PN retaining all acceleration terms up to 2.5PN and 3.5PN, respectively. The associated 0PN (1PN) standard non-adiabatic approximation retains acceleration terms corresponding to 0PN and 2.5PN (0PN, 1PN, 2.5PN and 3.5PN).

The explicit terms for accelerations for each order are given as follows [65, 73, 74]

$$\mathbf{a}_N = -\frac{M}{r^2} \hat{\mathbf{n}}, \quad (2.18)$$

$$\mathbf{a}_{1PN} = -\frac{M}{r^2} \left\{ \hat{\mathbf{n}} \left[(1 + 3\eta)v^2 - 2(2 + \eta)\frac{M}{r} - \frac{3}{2}\eta\dot{r}^2 \right] - 2(2 - \eta)\dot{r}\mathbf{v} \right\}, \quad (2.19)$$

$$\begin{aligned} \mathbf{a}_{2PN} = & -\frac{M}{r^2} \left\{ \hat{\mathbf{n}} \left[\frac{3}{4}(12 + 29\eta) \left(\frac{M}{r} \right)^2 + \eta(3 - 4\eta)v^4 + \frac{15}{8}\eta(1 - 3\eta)\dot{r}^4 \right. \right. \\ & - \frac{3}{2}\eta(3 - 4\eta)v^2\dot{r}^2 - \frac{1}{2}\eta(13 - 4\eta)\frac{M}{r}v^2 - (2 + 25\eta + 2\eta^2)\frac{M}{r}\dot{r}^2 \Big] \\ & \left. - \frac{1}{2}\dot{r}\mathbf{v} \left[\eta(15 + 4\eta)v^2 - (4 + 41\eta + 8\eta^2)\frac{M}{r} - 3\eta(3 + 2\eta)\dot{r}^2 \right] \right\}, \quad (2.20) \end{aligned}$$

$$\begin{aligned} \mathbf{a}_{3PN} = & -\frac{M}{r^2} \left\{ \hat{\mathbf{n}} \left[-\frac{35\dot{r}^6\eta}{16} + \frac{175\dot{r}^6\eta^2}{16} - \frac{175\dot{r}^6\eta^3}{16} + \frac{15\dot{r}^4\eta v^2}{2} - \frac{135\dot{r}^4\eta^2 v^2}{4} + \frac{255\dot{r}^4\eta^3 v^2}{8} \right. \right. \\ & - \frac{15\dot{r}^2\eta v^4}{2} + \frac{237\dot{r}^2\eta^2 v^4}{8} - \frac{45\dot{r}^2\eta^3 v^4}{2} + \frac{11\eta v^6}{4} - \frac{49\eta^2 v^6}{4} + 13\eta^3 v^6 \\ & + \frac{M}{r} \left(79\dot{r}^4\eta - \frac{69\dot{r}^4\eta^2}{2} - 30\dot{r}^4\eta^3 - 121\dot{r}^2\eta v^2 + 16\dot{r}^2\eta^2 v^2 + 20\dot{r}^2\eta^3 v^2 \right. \\ & + \frac{75\eta v^4}{4} + 8\eta^2 v^4 - 10\eta^3 v^4 \Big) + \frac{M^2}{r^2} \left(\dot{r}^2 + \frac{22717\dot{r}^2\eta}{168} + \frac{11\dot{r}^2\eta^2}{8} - 7\dot{r}^2\eta^3 \right. \\ & + \frac{615\dot{r}^2\eta\pi^2}{64} - \frac{20827\eta v^2}{840} + \eta^3 v^2 - \frac{123\eta\pi^2 v^2}{64} \Big) \\ & + \frac{M^3}{r^3} \left(-16 - \frac{1399\eta}{12} - \frac{71\eta^2}{2} + \frac{41\eta\pi^2}{16} \right) \Big] + \mathbf{v} \left[-\frac{45\dot{r}^5\eta}{8} + 15\dot{r}^5\eta^2 \right. \\ & + \frac{15\dot{r}^5\eta^3}{4} + 12\dot{r}^3\eta v^2 - \frac{111\dot{r}^3\eta^2 v^2}{4} - 12\dot{r}^3\eta^3 v^2 - \frac{65\dot{r}\eta v^4}{8} + 19\dot{r}\eta^2 v^4 + 6\dot{r}\eta^3 v^4 \Big] \end{aligned}$$

¹⁰Note that there is no 1.5PN acceleration term, as dissipative terms due to radiation reaction first enter at 2.5PN order.

$$\begin{aligned}
& + \frac{M}{r} \left(\frac{329\dot{r}^3\eta}{6} + \frac{59\dot{r}^3v^2}{2} + 18\dot{r}^3\eta^3 - 15\dot{r}\eta v^2 - 27\dot{r}\eta^2v^2 - 10\dot{r}\eta^3v^2 \right) \\
& + \frac{M^2}{r^2} \left(-4\dot{r} - \frac{5849\dot{r}\eta}{840} + 25\dot{r}\eta^2 + 8\dot{r}\eta^3 - \frac{123\dot{r}\eta\pi^2}{32} \right) \Big] \Big\}, \quad (2.21)
\end{aligned}$$

$$\mathbf{a}_{2.5\text{RR}} = \frac{8}{5}\eta \frac{M^2}{r^3} \left\{ \dot{r}\hat{\mathbf{n}} \left[18v^2 + \frac{2M}{3r} - 25\dot{r}^2 \right] - \mathbf{v} \left[6v^2 - 2\frac{M}{r} - 15\dot{r}^2 \right] \right\}, \quad (2.22)$$

$$\begin{aligned}
\mathbf{a}_{3.5\text{RR}} = & \frac{8}{5}\eta \frac{M^2}{r^3} \left\{ \dot{r}\hat{\mathbf{n}} \left[\left(\frac{87}{14} - 48\eta \right) v^4 - \left(\frac{5379}{28} + \frac{136}{3}\eta \right) v^2 \frac{M}{r} + \frac{25}{2}(1+5\eta)v^2\dot{r}^2 \right. \right. \\
& + \left. \left(\frac{1353}{4} + 133\eta \right) \dot{r}^2 \frac{M}{r} - \frac{35}{2}(1-\eta)\dot{r}^4 + \left(\frac{160}{7} + \frac{55}{3}\eta \right) \left(\frac{M}{r} \right)^2 \right] \\
& - \mathbf{v} \left[-\frac{27}{14}v^4 - \left(\frac{4861}{84} + \frac{58}{3}\eta \right) v^2 \frac{M}{r} + \frac{3}{2}(13-37\eta)v^2\dot{r}^2 \right. \\
& + \left. \left. \left(\frac{2591}{12} + 97\eta \right) \dot{r}^2 \frac{M}{r} - \frac{25}{2}(1-7\eta)\dot{r}^4 + \frac{1}{3} \left(\frac{776}{7} + 55\eta \right) \left(\frac{M}{r} \right)^2 \right] \right\}, \quad (2.23)
\end{aligned}$$

where $\hat{\mathbf{n}} = \mathbf{r}/r$. In the above Eq.(2.21), the logarithmic terms present in $\mathbf{a}_{3\text{PN}}$ in [73] have been transformed away by an infinitesimal gauge transformation following [74]¹¹.

The above equations are solved numerically to attain the dynamics of the system. Then the orbital phase $\phi(t)$ is calculated by numerically solving the equations

$$\frac{d\phi}{dt} = \omega, \quad v^2 = \omega^2 r^2. \quad (2.24)$$

where the calculation of the orbital angular frequency assumes that the orbit is circular. Once we have the orbital phase, the waveform is generated using Eq. 2.2 since the orbital phase ϕ is related to the GW phase φ by $\varphi = 2\phi$.

2.3.1 Standard and complete non-adiabatic approximants in the test mass case

In the following we discuss the results of our study for the non-adiabatic waveforms in the test mass limit. To determine the appropriate expressions for the acceleration in the test mass case, we start with the general expression for the acceleration and set $\eta = 0$ in the conservative 1PN, 2PN and 3PN terms. Since doing the same in the dissipative terms at 2.5PN and 3.5PN orders prevents the orbit from decaying, we retain terms linear in η at these orders and set to zero terms of higher order in η . In the first part we discuss results found on the problem of the closeness of the standard and complete non-adiabatic template families with the fiducial exact waveform. In

¹¹We thank Luc Blanchet for pointing this to us and providing us this form

Table 2.7: Effectualness of the Lagrangian templates in the test mass case for the white noise.

Order (n)	$(1M_{\odot}, 10M_{\odot})$		$(1M_{\odot}, 50M_{\odot})$		$(1M_{\odot}, 100M_{\odot})$	
	S	C	S	C	S	C
OPN	0.5521	0.4985	0.5553	0.8399	0.6775	0.7789
1PN	0.4415	0.4702	0.5760	0.8327	0.6557	0.7591

Table 2.8: Faithfulness of the Lagrangian templates in the test mass case for the white noise.

Order (n)	$(1M_{\odot}, 10M_{\odot})$		$(1M_{\odot}, 50M_{\odot})$		$(1M_{\odot}, 100M_{\odot})$	
	S	C	S	C	S	C
OPN	0.0450	0.0441	0.1778	0.0991	0.5959	0.1851
1PN	0.0471	0.0474	0.3195	0.1235	0.3646	0.2404

the second part we extend our results to the noise spectrum expected in Initial LIGO, Advanced LIGO and Virgo.

White noise

First, we explore the general question as to the closeness of the standard non-adiabatic and complete non-adiabatic templates assuming a flat power spectral density for the detector noise. Tables 2.7 and 2.8 show the effectualness and faithfulness of Lagrangian models for the same three archetypal binaries as before: $(1M_{\odot}, 10M_{\odot})$, $(1M_{\odot}, 50M_{\odot})$, and $(1M_{\odot}, 100M_{\odot})$ binaries. At present, the results are available at too few PN orders to make statements of general trends in effectualness and faithfulness. However, the main result obtained for the adiabatic approximants seems to hold good again for the non-adiabatic approximants: the effectualness is higher for the complete non-adiabatic model as opposed to the standard non-adiabatic model. This is indicative of the fact that corrections coming from the conservative part of the dynamics (i.e. the well-known general relativistic effects at 1PN and 2PN) make an improvement of the effectualness. However, as in the adiabatic case, the faithfulness of both standard and complete non-adiabatic models is very poor. But, in sharp contrast to the adiabatic case, here it appears that the complete non-adiabatic approximation results in a decrease in the faithfulness of the templates.

Table 2.9: Effectualness of the Lagrangian templates in the test mass case for the Initial LIGO, Advanced LIGO, and Virgo noise PSDs. Percentage biases σ_M and σ_η in determining parameters M and η are given in brackets.

Order (n)	$(1M_\odot, 10M_\odot)$		$(1M_\odot, 50M_\odot)$	
	S	C	S	C
Initial LIGO				
OPN	0.5848(30, 26)	0.9496(55, 107)	0.8741(3.3, 9.2)	0.9835(35, 4.2)
1PN	0.6762(37, 49)	0.9273(3.1, 27)	0.8530(34, 191)	0.9784(24, 28)
Advanced LIGO				
OPN	0.4259 (31.09, 29.04)	0.8682 (55.0, 105.01)	0.6384 (4.71, 4.02)	0.9360 (40.20, 23.15)
1PN	0.5256 (52.55, 58.95)	0.8280 (0.36, 30.55)	0.6080 (28.96, 124.8)	0.9211 (17.61, 0.47)
Virgo				
OPN	0.3720 (22.18, 0.95)	0.7631 (55.18, 107.0)	0.5985 (20.49, 30.68)	0.9618 (40.35, 23.46)
1PN	0.3599 (3.81, 21.17)	0.7386 (3.45, 34.8)	0.5777 (21.74, 76.08)	0.9525 (17.59, 0.49)

Table 2.10: Faithfulness of the Lagrangian templates in the test mass case for the Initial LIGO, Advanced LIGO, and Virgo PSDs.

Order (n)	$(1M_\odot, 10M_\odot)$		$(1M_\odot, 50M_\odot)$	
	S	C	S	C
Initial LIGO				
OPN	0.2463	0.1216	0.5048	0.1747
1PN	0.4393	0.1823	0.3650	0.3119
Advanced LIGO				
OPN	0.2636	0.0754	0.4857	0.0999
1PN	0.3419	0.1368	0.4493	0.2038
Virgo				
OPN	0.1991	0.0570	0.4947	0.1097
1PN	0.2499	0.0911	0.5259	0.1954

Noise spectra of particular detectors

Having addressed the question concerning the closeness of standard non-adiabatic and complete non-adiabatic templates to exact waveforms, we now compare the overlaps in the case of the Initial LIGO, Advanced LIGO and Virgo detectors.

Tables 2.9 and 2.10 show the effectualness and faithfulness, respectively, of Lagrangian templates for the $(1M_\odot, 10M_\odot)$ and $(1M_\odot, 50M_\odot)$ binaries. In this case, we see that the effectualness gets significantly improved in the complete non-adiabatic approximation and is greater than 0.9 for all the systems studied in this Chapter. Faithfulness appears to be decreased by the use of the complete non-adiabatic approximation (this result is in sharp contrast with the corresponding adiabatic case where we find that the complete adiabatic approximation gener-

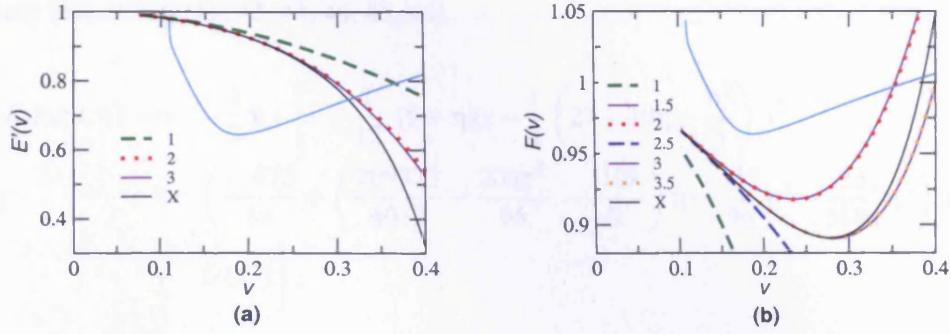


Figure 2.8: Various T-approximants of Newton-normalized (v -derivative of) energy function $E'_T(v)/E'_N(v)$ (left) and flux function $F_T(v)/F_N(v)$ (right) in the comparable mass case, along with the corresponding fiducial ‘exact’ functions (denoted by X). Also plotted is the amplitude spectral density (per $\sqrt{\text{Hz}}$) of Initial LIGO noise in arbitrary units.

ally brings about a significant improvement in faithfulness), but again there is no indication that either standard or complete templates are reliable in extracting the parameters of the system.

2.4 Comparable Mass Waveforms

In the case of comparable mass binaries there is no *exact* template available and the best we can do is to compare the performance of the standard adiabatic and complete adiabatic templates by studying their overlaps with some plausible fiducial ‘exact’ waveform. As in the case of the test masses, here too we will consider all possible combinations of T-approximants of energy and flux functions, construct PN templates and calculate the overlaps of these templates with the fiducial ‘exact’ waveform. In all cases, the fiducial ‘exact’ waveform is constructed by numerically integrating the phasing formula in the time-domain (TaylorT1 approximant), and terminating the waveforms (‘exact’ and approximate) at $v_{lso} = \frac{1}{\sqrt{6}}$ which corresponds to $F \simeq 1570$ Hz for a $(1.4M_\odot, 1.4M_\odot)$ binary and $F \simeq 220$ Hz for a $(10M_\odot, 10M_\odot)$ binary. Also, we take the noise PSD of the detector to be infinite below a lower frequency cut-off $F_{\text{low}} = 40$ Hz.

2.4.1 The energy function

Unlike in the test mass limit, the energy function $E(x; \eta)$ is not known exactly in the comparable mass case but only a post-Newtonian expansion, which has been computed at present up to 3PN

accuracy [38, 39, 40, 41, 42, 43, 44, 45, 46].

$$\begin{aligned}
E_{3PN}(x; \eta) = & -\frac{1}{2} \eta x \left[1 - \frac{1}{12} (9 + \eta) x - \frac{1}{8} \left(27 - 19\eta + \frac{\eta^2}{3} \right) x^2 \right. \\
& + \left(\frac{-675}{64} + \left(\frac{209323}{4032} - \frac{205\pi^2}{96} - \frac{110\lambda}{9} \right) \eta - \frac{155}{96} \eta^2 - \frac{35}{5184} \eta^3 \right) x^3 \\
& \left. + O(x^4) \right], \tag{2.25}
\end{aligned}$$

where $\lambda = -1987/3080$ [44, 45, 46, 47]. The corresponding $E'(v; \eta)$ appearing in the phasing formula reads,

$$\begin{aligned}
E'_{3PN}(v; \eta) = & -\eta v \left[1 - \frac{1}{6} (9 + \eta) v^2 - \frac{3}{8} \left(27 - 19\eta + \frac{\eta^2}{3} \right) v^4 \right. \\
& + 4 \left(\frac{-675}{64} + \left(\frac{209323}{4032} - \frac{205\pi^2}{96} - \frac{110\lambda}{9} \right) \eta - \frac{155}{96} \eta^2 - \frac{35}{5184} \eta^3 \right) v^6 \\
& \left. + O(v^8) \right]. \tag{2.26}
\end{aligned}$$

We use this expression truncated at the necessary orders to construct the various approximate templates. To compute a fiducial ‘exact’ waveform, we use the exact energy function in the test mass limit supplemented by the finite mass corrections up to 3PN in the spirit of the hybrid approximation [75, 76]. In other words, the fiducial ‘exact’ energy $E'(v; \eta)$ will then look like

$$\begin{aligned}
E'_{\text{exact}}(v; \eta) = & -\eta v \left[\frac{-E'_{\text{exact}}(v)}{\eta v} - \frac{\eta}{6} v^2 - \frac{3}{8} \left(-19\eta + \frac{\eta^2}{3} \right) v^4 \right. \\
& \left. + 4 \left(\left(\frac{209323}{4032} - \frac{205\pi^2}{96} - \frac{110\lambda}{9} \right) \eta - \frac{155}{96} \eta^2 - \frac{35}{5184} \eta^3 \right) v^6 \right] \tag{2.27}
\end{aligned}$$

where $E'_{\text{exact}}(v)$ is the v -derivative of the exact energy function in the test mass limit given by Eq. (2.9). The T-approximants of the energy function $E'_T(v; \eta)$ as well as the fiducial exact energy $E'_{\text{exact}}(v; \eta)$ are plotted in Fig. 2.8a. The v_{lso} corresponding to the fiducial ‘exact’ energy function can be determined by solving $E'_{\text{exact}}(v; \eta) = 0$. This will yield a value $v_{\text{lso}}^{3PN\text{-hybrid}} \simeq 0.4294$ against the $v_{\text{lso}} \simeq 0.4082$ in the test-mass case (more precisely it is the v_{MECO} [65]). If the η -corrections are included only up to 2PN instead of 3PN, $v_{\text{lso}}^{2PN\text{-hybrid}} \simeq 0.4113$. It is worth pointing out that $v_{\text{lso}}^{2PN\text{-Pade}} \simeq 0.4456$ [70] and it is not unreasonable to suspect that, with 3PN corrections the differences between the different ways of determining the lso converge. (For the purposes of our analysis, we have checked that there is no drastic change in our conclusions due

to these differences and hence we use uniformly the value $v_{iso} = 0.4082$).

2.4.2 The flux function

The flux function in the case of comparable masses has been calculated up to 3.5PN accuracy [48, 49, 50, 51, 52, 53, 54], and is given by:

$$\mathcal{F}(v; \eta) = \frac{32}{5} \eta^2 v^{10} \left[\sum_{k=0}^7 A_k(\eta) v^k + B_6(\eta) v^6 \ln v + O(v^8) \right], \quad (2.28)$$

where

$$\begin{aligned} A_0(\eta) &= 1, \quad A_1(\eta) = 0, \quad A_2(\eta) = -\frac{1247}{336} - \frac{35}{12} \eta, \quad A_3(\eta) = 4\pi, \\ A_4(\eta) &= -\frac{44711}{9072} + \frac{9271}{504} \eta + \frac{65}{18} \eta^2, \quad A_5(\eta) = -\left(\frac{8191}{672} + \frac{583}{24} \eta \right) \pi, \\ A_6(\eta) &= \frac{6643739519}{69854400} + \frac{16\pi^2}{3} - \frac{1712}{105} \gamma + \left(-\frac{11497453}{272160} + \frac{41\pi^2}{48} + \frac{176\lambda}{9} - \frac{88\Theta}{3} \right) \eta \\ &\quad - \frac{94403}{3024} \eta^2 - \frac{775}{324} \eta^3 - \frac{1712}{105} \ln 4, \\ A_7(\eta) &= \left(-\frac{16285}{504} + \frac{214745}{1728} \eta + \frac{193385}{3024} \eta^2 \right) \pi, \quad B_6(\eta) = -\frac{1712}{105}, \end{aligned} \quad (2.29)$$

and the value of Θ has been recently calculated to be $\simeq -1.28$ [56] by dimensional regularization.

To construct our fiducial ‘exact’ waveform, we will use the energy function given by Eq. (2.27) and the flux function

$$\begin{aligned} \mathcal{F}_{\text{exact}}(v; \eta) &= \frac{32}{5} \eta^2 v^{10} \left[\mathcal{F}_{\text{exact}}(v) - \sum_{k=0}^7 \left(A_k v^k + B_6 v^6 \ln v \right) \right. \\ &\quad \left. + \sum_{k=0}^7 \left(A_k(\eta) v^k + B_6(\eta) v^6 \ln v \right) \right], \end{aligned} \quad (2.30)$$

where $\mathcal{F}_{\text{exact}}(v)$ is the Newton-normalized (numerical) exact flux in the test-mass limit. The expansion coefficients A_k ’s and B_6 refer to the test-mass case and $A_k(\eta)$ ’s and $B_6(\eta)$ refer to the comparable mass case. The exact flux function is thus constructed by superposing all that we know in the test mass case from perturbation methods and the two body case by post-Newtonian methods. It supplements the exact flux function in the test body limit by all the η -dependent corrections known up to 3.5PN order in the comparable mass case. The T-approximants of the




Figure 2.9: Effectualness (Fig. *a*) and Faithfulness (Fig. *b*) of various TaylorT1 and TaylorF1 templates in the comparable mass case for the Initial LIGO noise spectrum. Different lines in the panels correspond to different orders of the flux function. Each line shows how the overlaps are evolving as a function of the accuracy of the energy function. Standard adiabatic approximants $T(E_{[n]}, \mathcal{F}_n)$ are marked with thick dots. Label **A** represents the $(10M_{\odot}, 10M_{\odot})$ binary and label **C** represents the $(1.4M_{\odot}, 1.4M_{\odot})$ binary. All values are max-max overlaps.

flux function $\mathcal{F}_T(v; \eta)$ and the fiducial exact flux $\mathcal{F}_{\text{exact}}(v; \eta)$ are plotted in Fig. 2.8*b*.

2.4.3 Comparable mass results in the adiabatic approximation

The effectualness and faithfulness of various PN templates in the case of comparable mass binaries are plotted in Fig. 2.9*a* and Fig. 2.9*b*, respectively, and are tabulated in Tables 2.11 and 2.12. The overlaps of the fiducial ‘exact’ waveform in the figures are calculated with the TaylorT1 and TaylorF1 approximants using the Initial LIGO noise spectrum. Different lines in the panels of Fig. 2.9 correspond to different PN orders of the flux function. Let us note that in the case of comparable mass binaries the complete adiabatic approximants can be calculated, at present, at most up to 1PN order. From the Tables 2.11 and 2.12 one can see that the complete adiabatic approximation generally improves the *effectualness* of the templates at 0PN and 1PN orders. But, as far as *faithfulness* is concerned, it is hard to conclude that one approximation is better than the other at these PN orders.

Even though complete adiabatic approximants are not calculated for higher PN orders, the general conclusion one can make from Fig. 2.9 is that the complete adiabatic approximation of the phasing will not result in a significant improvement in overlaps if we have a flux function of order ≥ 1.5 PN. We, thus, conclude that, provided we have a sufficiently accurate (order ≥ 1.5 PN) T-approximant of the flux function, the standard standard adiabatic approximation

Table 2.11: Effectualness of *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ approximants in the comparable-mass case. Overlaps are calculated for the Initial LIGO, Advanced LIGO and Virgo noise spectra. Percentage biases σ_M and σ_η in determining parameters M and η are given in brackets.

Order (n)	$(10M_\odot, 10M_\odot)$		$(1.4M_\odot, 1.4M_\odot)$	
	S	C	S	C
Initial LIGO				
0PN	0.8815 (14, 0.2)	0.9515 (3.7, 0.1)	0.8636 (1.4, 0.2)	0.6993 (4.3, 0.2)
1PN	0.8457 (59, 0.1)	0.8957 (45, 12)	0.5398 (5.0, 0.1)	0.5639 (4.3, 0.2)
1.5PN	0.9536 (3.9, 0.3)		0.9516 (0.4, 0.2)	
2PN	0.9833 (0.4, 0.2)		0.8751 (0.0, 0.1)	
2.5PN	0.8728 (14, 0.1)		0.8517 (0.4, 0.1)	
3PN	0.9822 (1.5, 0.0)		0.9955 (0.0, 0.3)	
3.5PN	0.9843 (1.4, 0.0)		0.9968 (0.0, 0.3)	
Advanced LIGO				
0PN	0.7606(8.45, 0.06)	0.9132(5.00, 0.28)	0.8347(1.43, 0.05)	0.5809(4.29, 0.09)
1PN	0.7110(56.8, 0.57)	0.7360(38.8, 0.61)	0.3959(6.07, 0.01)	0.4194(5.00, 0.12)
1.5PN	0.8741(2.40, 0.12)		0.9034(0.00, 0.18)	
2PN	0.9803(0.80, 0.16)		0.8179(0.36, 0.03)	
2.5PN	0.7705(7.30, 0.05)		0.7826(0.36, 0.03)	
3PN	0.9626(0.50, 0.01)		0.9981(0.36, 0.28)	
3.5PN	0.9683(1.25, 1.51)		0.9977(0.36, 0.28)	
Virgo				
0PN	0.7009(5.25, 0.00)	0.9280(4.70, 0.93)	0.7119(1.43, 0.08)	0.4405(5.00, 0.02)
1PN	0.5834(55.5, 0.28)	0.6148(30.3, 0.20)	0.2808(3.93, 0.10)	0.2968(2.86, 0.04)
1.5PN	0.8698(1.25, 0.00)		0.7724(0.36, 0.67)	
2PN	0.9815(0.75, 0.24)		0.6420(0.00, 0.01)	
2.5PN	0.7299(4.75, 0.00)		0.6266(0.00, 0.08)	
3PN	0.9624(0.50, 0.11)		0.9822(0.00, 0.25)	
3.5PN	0.9627(0.45, 0.14)		0.9823(0.00, 0.25)	

provides a good lower bound to the complete adiabatic approximation for the construction of both effectual and faithful templates in the case of comparable mass binaries. It should be kept in mind that unlike the test mass case where the exact energy and flux functions are known leading to an exact waveform in the adiabatic approximation in the comparable mass case we are only talking about *fiducial* energy and flux functions constructed from what is known. Probably, the exact waveform in this case has nothing much to do with the exact waveform predicted by general relativity.

Table 2.12: Faithfulness of the *standard adiabatic* $T(E_{[n]}, \mathcal{F}_n)$ and *complete adiabatic* $T(E_{[n+2.5]}, \mathcal{F}_n)$ templates in the comparable-mass case. The overlaps are calculated for the Initial LIGO, Advanced LIGO and Virgo noise spectra.

Order (n)	$(10M_\odot, 10M_\odot)$		$(1.4M_\odot, 1.4M_\odot)$	
	S	C	S	C
Initial LIGO				
0PN	0.5603	0.8560	0.3783	0.1624
1PN	0.3026	0.3491	0.1520	0.1615
1.5PN	0.7949		0.7259	
2PN	0.9777		0.5565	
2.5PN	0.5687		0.5934	
3PN	0.9440		0.9888	
3.5PN	0.9522		0.9916	
Advanced LIGO				
0PN	0.3902	0.7030	0.3731	0.1300
1PN	0.1944	0.2248	0.1054	0.1128
1.5PN	0.6362		0.5735	
2PN	0.8895		0.3964	
2.5PN	0.4125		0.4407	
3PN	0.9117		0.9947	
3.5PN	0.9106		0.9952	
Virgo				
0PN	0.4262	0.5490	0.3138	0.0794
1PN	0.1574	0.1798	0.0686	0.0732
1.5PN	0.5950		0.3986	
2PN	0.8120		0.3027	
2.5PN	0.3842		0.3726	
3PN	0.9169		0.9668	
3.5PN	0.9177		0.9686	

Table 2.13: Effectualness of the Lagrangian templates in the comparable mass case for Initial LIGO, Advanced LIGO and Virgo noise spectra. Percentage biases σ_M and σ_η in determining parameters M and η are given in brackets.

Order (n)	$(1.4M_\odot, 1.4M_\odot)$		$(10M_\odot, 10M_\odot)$	
	S	C	S	C
Initial LIGO				
0PN	0.9282(27.43, 1.56)	0.5848(31.60, 3.61)	0.8533(14.35, 2.51)	0.9433(34.50, 8.87)
1PN	0.5472(22.22, 1.25)	0.6439(24.31, 3.09)	0.8137(3.60, 0.17)	0.9329(11.15, 7.38)
Advanced LIGO				
0PN	0.9147 (25.35, 1.93)	0.4338 (28.93, 3.84)	0.7417 (18.60, 2.21)	0.8322 (35.1, 3.36)
1PN	0.3937 (20.00, 3.51)	0.5132 (21.43, 2.98)	0.7433 (4.40, 0.001)	0.8158 (13.4, 1.36)
Virgo				
0PN	0.8142 (25.71, 1.81)	0.3113 (29.64, 1.13)	0.6895 (20.65, 2.31)	0.7341 (32.9, 5.97)
1PN	0.2880 (21.42, 2.68)	0.3944 (19.28, 1.22)	0.6807 (4.85, 0.04)	0.7420 (14.5, 4.00)

2.4.4 Comparable mass results beyond the adiabatic approximation

Finally, for the comparable mass case, non-adiabatic waveforms were generated in the Lagrangian formalism, using the complete equations Eq. (2.17) – Eq. (2.23).

Tables 2.13 and 2.14 show the effectualness and faithfulness of the standard and complete non-adiabatic Lagrangian waveforms for the Initial LIGO, Advanced LIGO and Virgo detectors. The results are more mixed in this case than for the test mass case. For the 0PN order, for the NS-NS binary, the standard non-adiabatic approach seems to be more effectual and faithful than its complete non-adiabatic counterpart. However, in the BH-BH case, the complete non-adiabatic seems to be more effectual; but less faithful. At the 1PN order, the effectualness is always higher for the complete non-adiabatic case; but faithfulness is always lower. It is interesting to note that the effectualness trends shown by the adiabatic and non-adiabatic approximants are the same at orders 0PN and 1PN. However, further work will be necessary to make very strong statements in this case.

2.5 Summary and Conclusion

The *standard adiabatic* approximation to phasing of gravitational waves from inspiralling compact binaries is based on the post-Newtonian expansions of the binding energy and gravitational wave flux both truncated at the *same relative* post-Newtonian order. To go beyond the adiabatic approximation one must view the problem as the dynamics of a binary under conservative rela-

Table 2.14: Faithfulness of the Lagrangian templates in the comparable mass case for Initial LIGO, Advanced LIGO and Virgo noise spectra.

Order (n)	$(1.4M_{\odot}, 1.4M_{\odot})$		$(10M_{\odot}, 10M_{\odot})$	
	S	C	S	C
Initial LIGO				
OPN	0.0717	0.0658	0.6689	0.3146
1PN	0.0810	0.0771	0.7380	0.6568
Advanced LIGO				
OPN	0.0637	0.0546	0.4662	0.2192
1PN	0.0794	0.0703	0.6594	0.4788
Virgo				
OPN	0.0414	0.0334	0.3439	0.1661
1PN	0.0803	0.0463	0.5709	0.3704

tivistic forces and gravitation radiation damping. In this viewpoint the standard approximation at leading order is equivalent to considering the 0PN and 2.5PN terms in the acceleration and neglecting the intermediate 1PN and 2PN terms. A complete treatment of the acceleration at leading order should include *all* PN terms upto 2.5PN. These define the *standard* and *complete non-adiabatic* approximants respectively. A new post-Newtonian *complete adiabatic* approximant based on energy and flux functions is proposed. At the leading order it uses the 2PN energy function rather than the 0PN one in the standard approximation so that loosely it does not miss any intermediate post-Newtonian terms in the acceleration. We have evaluated the performance of the standard adiabatic vis-a-vis complete adiabatic approximants, in terms of their *effectualness* (i.e. larger overlaps with the exact signal) and *faithfulness* (i.e. smaller bias in estimation of parameters). We restricted our study only to the inspiral part of the signal neglecting the plunge and quasi-normal mode ringing phases of the binary [66, 67, 24, 23, 65, 68, 69]. We have studied the problem for the white-noise case, as well as the Initial LIGO, Advanced LIGO and Virgo noise curves.

The main result of this study is that the conservative corrections to the dynamics of a binary that are usually neglected in the standard treatment of the phasing formula are rather important at low PN orders. At the low PN orders, they lead to significant improvement in the overlaps between the approximate template and the exact waveform. In both the white-noise case, and the case of specific detectors, we found that at low (< 3 PN) PN orders the effectualness of the approximants significantly improves in the complete adiabatic approximation. However, standard adiabatic approximants of order ≥ 3 PN are nearly as good as the complete adiabatic

approximants for the construction of effectual templates.

In the white-noise case, the *faithfulness* of both the approximants fluctuates as we go from one PN order to the next and is generally much smaller than our target value of 0.965. The fluctuation continues all the way up to 5PN order probably reflecting the oscillatory approach of the flux function to the exact flux function with increasing PN order. Poor faithfulness also means that the parameters extracted using these approximants will be biased. It is again interesting to note that complete adiabatic approximants are generally more faithful than the standard adiabatic approximants. For the Initial LIGO noise case on the other hand, the faithfulness of the *complete adiabatic* approximants is vastly better at all orders.

To the extent possible, we also tried to investigate this problem in the case of comparable mass binaries by studying the overlaps of all the approximants with a fiducial ‘exact’ waveform. It is shown that, provided we have a T-approximant of the flux function of order $\geq 1.5\text{PN}$, the *standard adiabatic* approximation provides a good lower bound to the *complete adiabatic* approximation for the construction of both *effectual* and *faithful* templates. This result is in contrast with the test mass case where we found that the *complete adiabatic* approximation brings about a significant improvement in *effectualness* up to 2.5PN order, and a significant improvement in *faithfulness* at all orders. To achieve the target sensitivity of 0.965 in effectualness, standard adiabatic approximants of order 2PN and 3PN are required for $(10M_{\odot}, 10M_{\odot})$ and $(1.4M_{\odot}, 1.4M_{\odot})$ binaries. Whether the complete adiabatic approximant achieves this at an earlier PN order is an interesting question. It is worth stressing that this result is only relevant for the family of inspiral waveforms. In the real physical case of BH-BH binaries the inspiral family would not be adequate and must be supplemented by the plunge part of the waveform as first discussed in [23, 24] and later in [65, 68]. A discussion of plunge requires a 3PN description of dynamics so that the 2PN templates are no longer adequate. This is an example of the second variety of questions one can study in this area referred to in our introduction related to whether a template family indeed represents GWs from a specific astrophysical system.

We have also constructed both standard and complete non-adiabatic approximants using the Lagrangian models in Ref. [65]. However, we were limited in this investigation because of two reasons: Firstly, the Lagrangian models are available only up to 3.5PN order, and higher order PN accelerations are as yet unavailable which makes it impossible to calculate the *complete*

non-adiabatic approximants of order ≥ 1 PN. Secondly, the only *exact* waveform we have has, however, been constructed only in the *adiabatic* approximation. So we are unable to make strong statements of general trends and view this effort only as a first step towards a more thorough investigation. From the non-adiabatic models studied, the conclusion one can draw is that while complete non-adiabatic approximation improves the effectualness, it results in a decrease in faithfulness.

There is a limitation to our approach which we should point out: complete adiabatic models can be very well tested in the test mass case where both approximate and exact expressions are available for the various quantities. However, complete models cannot be worked out to high orders in the comparable mass case since they need the energy function to be computed to 2.5PN order greater than the flux and currently the energy function is only known to 3PN accuracy. Also, due to the lack of an exact waveform, one is constrained to depend upon some fiducial exact waveform constructed from the approximants themselves. Though, in the present Chapter, we have used the new approximants to construct waveform templates, one can envisage applications to discuss the dynamics of the binary using numerical integration of the equations of motion.

During the course of this study, we also attempted to assess the relative importance of improving the accuracy of the energy function and the flux function by systematically studying the approach of the adiabatic PN templates constructed with different orders of the energy and the flux functions to the exact waveforms. From the study of test-mass templates we also conclude that, provided the comparable mass case is qualitatively similar to the test mass case, neither the improvement of the accuracy of the energy function from 3PN to 4PN, nor the improvement of the flux function from 3.5PN to 4PN will result in a significant improvement in effectualness in the comparable mass case.

Chapter 3

Developing a low-latency search for inspiralling compact binaries

Searching for inspiral compact binary sources in gravitational wave data is a challenging task. For the current generation of ground-based detectors, the signals from even strong sources will be buried within the detector noise. To locate the signals, it is therefore necessary to filter the data, and utilise various techniques for removing spurious events caused by noise. This process is very computationally intensive. Accordingly, for a search to be carried out in a feasible amount of time, it is necessary to design search pipelines which run across several nodes on large computing clusters.

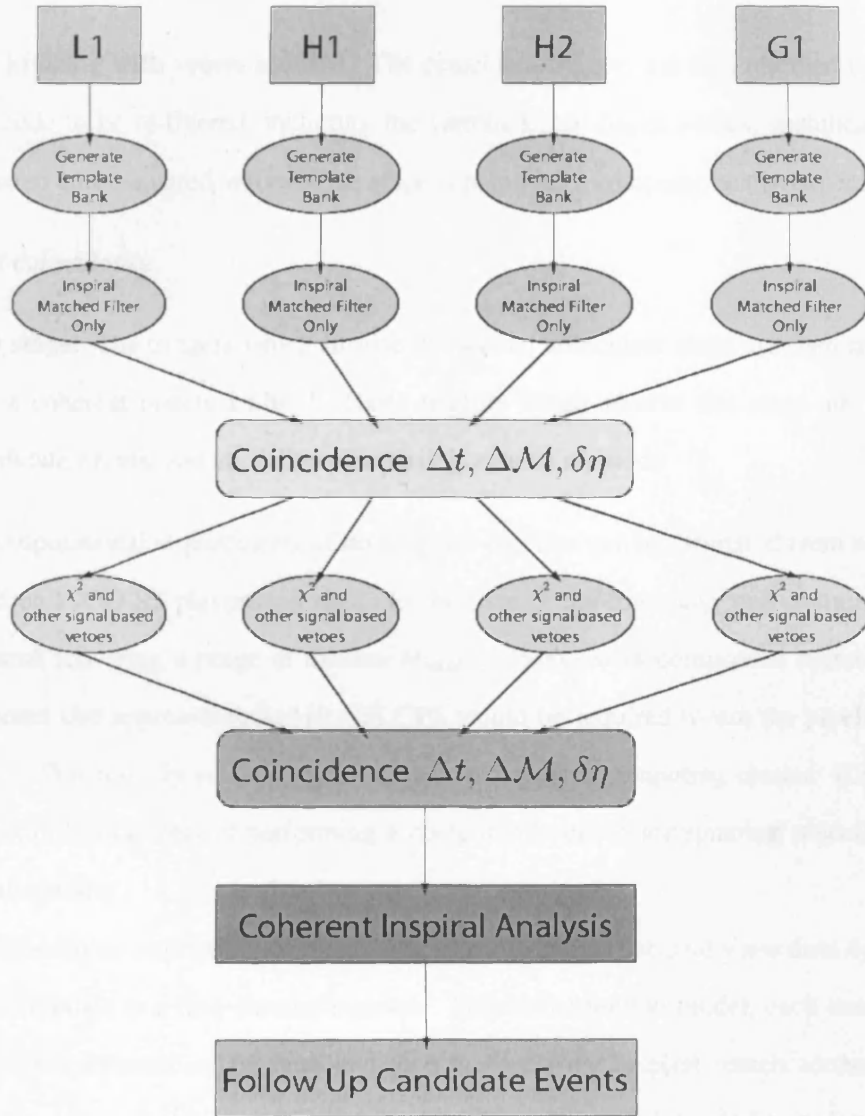
The design of the inspiral pipeline currently used in the flagship LIGO searches consists of a number of stages, as illustrated in Fig. 3.1. Each stage of the pipeline is run in an embarrassingly parallel manner. The various stages of the pipeline perform the following tasks [77]:

Generate a template bank for each chunk of data. The bank generated depends upon the choice of mass range, the requested minimal match, and the noise power spectrum for that chunk.

Match filter the data using the inspiral code. This takes the templates produced in the previous stage and matched filters them against the data stream. Any time that the signal to noise ratio exceeds a specified threshold, we record a trigger. This stage is the most computationally intensive part of the pipeline.

Search for coincident triggers between different detectors. The triggers generated by the pre-

Figure 3.1: Topology of the compact binary inspiral search pipeline [77].



vious stage are read in, and triggers from different detectors are compared by their parameters (for example, τ_0, τ_3, t_C). If the parameters agree to within a certain tolerance, the event is deemed coincident.

Trigbank. Coincident events are processed to create template banks for the following stage.

This ensures that only relevant areas of the parameter space are re-filtered using the computationally costly vetoes.

Matched filtering with vetoes applied. The coincident triggers are then returned to the inspiral code to be re-filtered, including the various signal-based vetoes, specifically the chi squared and r-squared vetoes. This stage is potentially computationally expensive.

Repeated coincidence.

Coherent stage. The triggers which survive the second coincident stage are then analysed using a coherent matched filter.¹ Those triggers which survive this stage are considered candidate events, and are followed up using various methods.

The computational requirements of running this pipeline can be estimated from runs already performed on LIGO S5 playground data. In the case of the low-mass non-spinning compact binary search (covering a range of masses $M_{total} < 35M_\odot$, with component masses $> 1M_\odot$), it is estimated that approximately 410 GFLOPS would be required to run the pipeline in real-time[78].² This roughly equates to 200 CPUs on a typical computing cluster. It is expected that the computational cost of performing a comparable search for spinning objects would be significantly greater.

In performing an inspiralling compact binary search in gravitational wave data, typically the search is distributed in a data-parallel manner.³ In this distribution model, each computational node receives a different set of data, and then performs the inspiral search across the entire search space. Although this leads to efficient use of computational resources, it also introduces a large latency into the search. This latency is not necessarily a problem for offline searches

¹It should be noted that, for the searches being performed on the first calendar year of the LIGO S5 data set, the coherent stage is not run as part of the pipeline. However, it will be used as one of the follow-ups on candidate events.

²This number neglects additional computing time required to perform Monte-Carlo injections of simulated signals, which are used for pipeline tuning, and measurement of the efficiency

³An exception is the online search (D. Brown et al).

across relatively short data sets; however, it can have a number of drawbacks. Firstly, when the detectors are engaged in very long, or continuous science runs, it may be desirable to obtain search results as quickly as possible, so as to provide an on-line summary of how the run is progressing. The latency involved in the data-parallel search may preclude the generation of such a summary as promptly as may be desired. Additionally, should the need arise to follow up a source seen in electromagnetic observations, it would be useful to obtain a result as quickly as possible.⁴ Also, if a strong candidate source is detected in gravitational waves, it may be beneficial to follow it up with electromagnetic observations. This is only feasible if the results from gravitational wave searches are obtained with a low latency.

This leads to the idea of a low latency search, wherein the search is distributed in a data-serial, parameter space-parallel manner. In this case, for the matched-filtering parts of the search pipeline, each computational node receives the same data, but processes a different area of the parameter space. Thus the latency is greatly reduced.

A question arises about how best to split the parameter space between computational nodes. For certain searches (the spinning BBH case in particular), parts of the parameter space may be more computationally intensive than others. Also, if running on a heterogeneous set of resources, it may be the case that the intrinsic performance of a particular resource may differ from the others. Thus, it was necessary to develop a way of balancing the load so that each node took the approximately same amount of computational time for each data set.

In this Chapter, two approaches at attacking this problem are detailed. The first is a simple *step-wise load balancing algorithm*. In this case, the parameter space is initially split naively (i.e., each node receives the same number of templates). Utilising the time taken for each job to run in this case, the splitting of the template bank is adjusted for the next run. Thus the computational load is balanced via an iterative process, where the splitting of the parameter space of the next run depends on the timings for the previous run. In the event of a job failure or unexpected delay, the algorithm is designed not to wait indefinitely, but ‘march on ahead’ when a user specified fraction of the jobs have finished.

The second approach consists of a more dynamical load-balancing mechanism. In this case, slave nodes will request templates from a controller node whenever they become idle. This

⁴An example of such a source is GRB070201, the position of which on the sky indicated that it could be located in Andromeda. If this were the case, and assuming the source was a compact-binary coalescence, this would be a strong candidate for gravitational wave detection.

process continues until the entire list of templates is exhausted. Since, in this approach, the slave nodes will always be kept occupied, the CPU load will automatically be balanced.

3.1 Step-wise load balancing algorithm

While developing a low-latency pipeline, there are a number of considerations to take into account in deciding which approach to take. The first is the computational efficiency of the pipeline: we would ideally like the pipeline to be as efficient as possible. However, there is another, potentially competing consideration: the necessity to significantly modify existing analysis code, and develop new code. If an approach can be found so that the computational cost can be kept to a ‘good enough’ level, without significant modifications to code, it could be argued that such an approach should be pursued, rather than scientists spending large amounts of time re-writing code to make further, potentially unnecessary improvements.

The step-wise load balancing algorithm was designed with the latter consideration in mind. It required very little (if any) modification to the main data-analysis codes used in the inspiral pipeline to obtain a good level of balance for the load on each node.

3.1.1 How the algorithm works

The problem of balancing the load on each node for running the inspiral analysis more or less boils down to assessing the fraction of templates which should be given to each compute node⁵. In the case of the step-wise load balancing algorithm, the run initially starts by naively splitting the template bank, so that each node receives roughly the same number of templates. For subsequent steps of the run, the details of the splitting of the bank is determined by the timing information of the previous steps. Note that this means that subsequent steps cannot be run until the previous step has been completed. However, since the main motivation for such an algorithm is for generating online results with low latency, this is not necessarily a problem.

The load is balanced for subsequent steps in the following way. If there are n compute nodes, and node i receives P_i out of P_{total} templates for analysis, and takes time T_i to analyse

⁵This assumes that approximately the same point in different template banks will correspond to the same area of the parameter space. Since all template banks within a run are generated using the same placement algorithm, this is a fair assumption.

them, one can compute the average analysis time per template for each node

$$t_i = \frac{T_i}{P_i}. \quad (3.1)$$

To perfectly balance the nodes, one wishes for each node to take the same amount of time to complete its analysis, i.e.,

$$t_1 P_1 = t_2 P_2 = \dots = t_n P_n. \quad (3.2)$$

To achieve this aim, for the next step, each node is allocated P_i^{next} out of P_{total}^{next} templates according to the following:

$$P_i^{next} = \frac{k}{t_i} P_{total}^{next}, \quad (3.3)$$

where

$$k = \left(\sum \frac{1}{t_i} \right)^{-1}. \quad (3.4)$$

By the definition given above, the numbers P_i^{next} are not whole numbers; therefore, for the actual number of templates, they are rounded to the nearest whole number. Any discrepancy between $\sum P_i^{next}$ and P_{total}^{next} is then taken care of by adding to, or subtracting from, random P_i^{next} 's as appropriate.

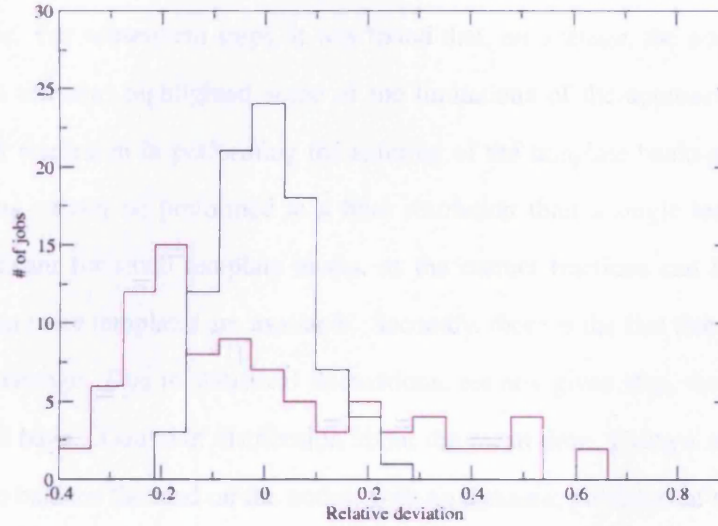
Once more than one step of the pipeline has been run, the algorithm can instead use the average time per template for each node, \bar{t}_i . This is more robust against being thrown off-course by the occasional 'rogue' job.

3.1.2 Marching ahead: how we deal with failures

In performing a long analysis run, it is inevitable that occasionally, certain jobs may fail. This could happen due to a node crashing, data becoming inaccessible, or another, unknown reason. In addition, due to compute nodes becoming busy with other tasks, or a rogue stretch of data having different characteristics to those expected, it can be the case that, for a particular step, one node may take much longer than the others to complete its analysis. To prevent the algorithm failing, or subsequent steps getting held up by a rogue job, it is necessary to have a means of dealing with these eventualities.

To this end, it was decided that, in the event of a failure, or a rogue long-running job, the

Figure 3.2: Histogram showing the deviation from the average time taken per compute node for each compute node. The red histogram shows the distribution for an initially unbalanced cluster. The black histogram shows the distribution obtained using the step-wise load-balancing algorithm.



algorithm would march on ahead regardless. In the event of a job failure, the failure would be logged, so that it could be rescued at a later time. In the event of a long running job, the job would be passed to a monitor running in a separate thread, to keep track of its eventual completion or failure.

In practice, this was implemented as follows: the load-balancing script would go ahead and commence running the next step after a certain fraction of the nodes had completed (typically 95%, but this fraction could be specified by the user). The time per template t_i for a long-running job, or one which failed on the previous step, was taken to be the average, \bar{t}_i calculated from the prior steps; or in case of the first step of the run, it would be taken to be the average over the other nodes. For long-running jobs, the time taken for the job would be incorporated into the average upon completion.

3.1.3 Benefits of using the algorithm


To test the efficacy of using such an algorithm for balancing the cluster, an implementation was written in Python, to be used in conjunction with the Condor job management system. To perform the tests, the distribution of templates amongst the nodes was initially in an unbalanced state. It was found that, even after one step, the algorithm had brought the nodes into a far more balanced state. For subsequent steps, it was found that, *on average*, the nodes were balanced. However, the test also highlighted some of the limitations of the approach. Firstly, there is only a certain resolution in performing the splitting of the template banks amongst the nodes, as the splitting cannot be preformed at a finer resolution than a single template. This issue is only significant for small template banks, as the correct fractions can be more accurately achieved when more templates are available. Secondly, there is the fact that the nodes are only balanced on average. Due to statistical fluctuations, for any given step, the actual timings for the nodes will have a Gaussian distribution about the mean time. Using a step-wise algorithm such as this to balance the load on the nodes, with no dynamic provision of templates to nodes, it is impossible to obviate this limitation. However, given that using this approach requires no modification of the analysis code, and does balance the load on average, this may be good enough for our purposes.

There is another issue with this approach which should be of some concern. A fundamental assumption in this approach is that the time taken to run a particular analysis job is dominated by the time taken to perform the analysis itself. Time spent performing I/O operations is assumed to be negligible. This may be a safe assumption in the case of a cluster in a relatively idle state, or for jobs which require huge numbers of templates to be filtered. However, for the case of a heavily loaded cluster, or jobs which filter only a relatively small number of templates, this assumption may not be valid. In such situations where the time spent performing I/O is significant, the performance of the load-balancing algorithm will be significantly reduced.

3.2 Dynamic load-balancing algorithm

In an optimal situation, one would expect that all compute nodes involved in the analysis of a stretch of data would be completely occupied while the data is being analyzed. Although proven to be reasonably successful, the step-wise balancing algorithm cannot guarantee that this will

Figure 3.3: A schematic of the mechanism for centralised load balancing.




be the case for any given job. To do this, a dynamic way of providing templates to nodes would be necessary. To provide such functionality, it may be necessary to make significant re-writes to existing analysis codes; a potentially costly exercise. However, the long-term benefit may justify the investments.

The approach taken to this problem was a centralised dynamic approach. In this approach, a master process holds the collection of templates which need to be analyzed, and provides the slave nodes with templates to analyse. Whenever a slave node becomes idle, it requests further work, and the master node provides it with further templates for analysis. This process continues until the template list is exhausted. When this occurs, upon completion of their analysis, the slave nodes are sent a termination message, and the master node collects the results.

One potential drawback of using such an approach is that the master node can only process one request at a time. This may result in a bottleneck if many slave nodes make requests simultaneously. However, this should not be a problem provided the number of slave nodes are not too many, and the tasks are computationally intensive [79]. The latter condition can be achieved by a judicious choice of the number of templates sent by the master at any one time. For current searches, the former should not be a problem; however, for runs across larger parameter spaces,

Figure 3.4: A plot of the time taken to perform an analysis against the number of compute nodes used for dynamic load-balancing.




where searches may need to be run across huge numbers of nodes, the centralised approach may fail, and a move to a decentralised approach may be necessary.

3.2.1 Benefits of using the algorithm

The results of using the dynamic load-balancing algorithm for performing an analysis of a stretch of data using ~ 3500 templates can be seen in Fig. 3.4. It can be seen that there is a marked reduction in the time taken to perform the analysis as the number of nodes is increased. As can be seen in Fig. 3.5, the time taken approximately follows a power law $(n-1)^{-s}$, where n is the number of nodes being used, and $s \lesssim 1$. This is to be expected given the nature of distributing the job by the allocation of templates. The deviation of s from 1 as the value of n increases is due to a combination of the increase in the likelihood of communication bottlenecks due to simultaneous requests from compute nodes as n increases, and the increasing significance of the unchanging start-up time as the analysis time is reduced. The cost of communication can be at least partially mitigated by reducing the number of requests made by compute nodes. This can be achieved by tuning the number of templates sent by the master node in response to a request.

Figure 3.5: A log-log plot of the time taken to perform an analysis against the number of compute nodes used for dynamic load-balancing.



3.3 Conclusions

The search for gravitational wave signals from compact binary sources is a computationally intensive exercise, requiring analysis to be carried out across several processors in order for the search to be completed in a reasonable amount of time. The standard approach to the distribution of work amongst the processors is data-parallel, in which each processor performs the same work on a different stretch of data. Although computationally efficient, this approach introduces a large latency in the acquisition of results. It is expected that, in the future, acquiring results from a search as quickly as possible will become more important. To this end, we have investigated two approaches to distributing the work in a parameter-space parallel manner: step-wise load-balancing, and dynamic load balancing.

The step-wise load-balancing algorithm has the advantage of requiring little or no modifications to existing analysis codes, and undoubtedly speeds up the time taken to process the data. However, due to the changing characteristics of the data as a run progresses, as well as the changing loads on computational resources, it is only possible in this approach to balance the time taken per node on average. Therefore, using this approach, there will always be some com-

pute nodes which finish faster than others, making it inevitable that compute nodes will spend some time sitting idle.

The dynamical load-balancing algorithm is balanced by default, with slave nodes requesting more work off a master node when they become idle. The performance of the algorithm has been shown to be close to $t \propto (n - 1)^{-s}$, where $s \lesssim 1$ for up to 10 processors, which is very close to the naively expected optimal behaviour for distribution of the work in a parameter-space parallel manner. However, such an approach requires significant modification of analysis codes, warranting much investment of time and effort. In addition, due to the need for slave nodes to request and receive work from a master node, the use of such an approach introduces inter-process communication, which is not required in the step-wise approach. Although such communication may not introduce a large penalty if the analysis is run on a single computing cluster, the penalty will become more significant if the analysis is distributed across several resources located in different locations, which may be the case in a Grid environment.

Chapter 4

A geometric algorithm for efficient coincident detection of gravitational waves

4.1 Motivation

Long baseline interferometric gravitational wave detectors, such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) [80], Virgo [81], and GEO 600 [82], are currently acquiring the best data ever. The data sets from the different detectors can be either brought together and analyzed phase coherently [83, 84, 85, 86], or analyzed separately followed up by a coincidence analysis [87, 88, 85, 89, 90, 91, 92, 93, 94] of the triggers from different detectors. Coherent analysis maximizes signal visibility (i.e., gives the best possible signal-to-noise ratio in the likelihood sense) while the goal of coincidence analysis is to reduce and mitigate the non-stationary and non-Gaussian background noise. A recent comparison of coherent analysis vis-a-vis coincidence analysis under the assumption that the background noise is Gaussian and stationary has concluded that coherent analysis, as one might expect, is far better than coincidence analysis. However, there are two reasons why current data analysis pipelines prefer the latter over former. Firstly, since the detector noise is neither Gaussian nor stationary, coincidence analysis can potentially reduce the background rate far greater than one might think otherwise. Secondly, coherent analysis is computationally far more expensive than coincidence

analysis and it is presently not practicable to employ coherent analysis.

4.1.1 The problem of coincident detection

In coincidence analysis (see for example, Refs. [91, 95, 92, 93, 96, 94]), data sets from each detector will be analyzed separately and the triggers from the end of the pipeline from different detectors compared with one another to identify triggers that might be in coincidence with one another. More precisely, the goal is to find if the parameters (e.g., in the case of a coalescing binary the time of merger, the component masses and spins) of a trigger from one detector are identical to those from another. Since the background noise corrupts any inherent signal it is highly improbable that the same gravitational wave in different detectors can be associated with exactly the same set of parameters. However, it should be possible to detect signals in coincidence by demanding that the measured parameters lie in a sufficiently small range [91, 95, 92, 93, 96, 94]. Thus, we can revise the coincidence criteria as follows: *triggers from different detectors are said to be in coincidence if their parameters all lie within a certain range*. Events that pass the coincidence test are subject to further scrutiny but we shall focus in this Chapter on the coincidence test itself.

From the above discussion it is clear that an important aspect of coincidence analysis is the determination of the range of parameter values to be associated with each trigger. To this end, until recently, the LIGO Scientific Collaboration (LSC) has deployed a phenomenological method for assigning the ranges [91, 95, 92, 93, 96, 94]. More precisely, one performs a large number of simulations in which a signal with a known set of parameters is added in software to the data which is then passed through the analysis pipeline. The pipeline identifies the most probable parameters with each injected signal and the ensemble of injected and measured parameters gives the distribution of the errors incurred in the measurement process. Given the distribution of the errors, one can choose a range for each parameter such that more than, say, 95% of the injected signals are detected in coincidence. Choosing wider windows will enable greater detection probability but also increase the rate of accidental triggers. On the contrary, smaller windows decrease the false alarm rate but also reduce the detection probability.

4.1.2 A geometric approach to choosing coincident windows

In this Chapter we propose a new algorithm based on the metric (equivalently, the information matrix) defined on the signal manifold. The idea is very simple, even obvious, but leads to a great reduction in the background trigger rate. The advantages of the new algorithm are better appreciated by listing certain drawbacks of the phenomenological method. The drawbacks are quite naturally remedied in the new approach.

First, because the current method uses rectangular windows it ignores the correlations between different parameters. For instance, in the case of a chirping signal from a black hole binary the shape of the signal depends, among others, on the component masses. However, not all combinations of the two masses lead to signals that are easily distinguishable from one another. Indeed, at the lowest post-Newtonian order the waveform depends only on a certain combination of the masses called the *chirp mass*; binaries of different values for the two masses but the same chirp mass produce essentially the same signal. This degeneracy is broken when post-Newtonian corrections are included. Nevertheless, the two mass parameters continue to be highly correlated.

The second drawback is that the method employs windows of the same size throughout the parameter space while we know that errors in the measurement of the parameters depends, in some cases quite sensitively, on the parameters. Drawing again from our example of a binary, the error in the estimation of the chirp mass can vary by more than two orders of magnitude across the parameter space of interest in the case of systems that LIGO is expected to observe (see, e.g., [31, 26, 97, 98, 99, 100, 101, 102, 103, 66, 104]). Clearly, it is not optimal to deploy windows of the same size all over the parameter space.

Thirdly, by not taking into account parameter covariances, the method entails independent tuning of several parameters at the same time. This could be a horrendous problem when dealing with signals characterized by many parameters. For instance, continuous radiation from a pulsar is characterized by the location of the pulsar, its spin frequency, the derivative of the frequency and so on. These physical parameters are all not independent; the existence of covariances among them means that not all variations of the parameters leads to distinct signals. This further implies that it may not be necessary to tune each parameter separately, rather it should be enough to tune only a subset of the parameters or, more precisely, only the principal

components. Furthermore, the method does not provide a unique set of windows, rather several possibilities could be worked out.

Finally, by using windows of the same size irrespective of the signal-to-noise ratio of the trigger, the method suffers from an undesirably high false alarm rate, particularly in the tail of the SNR distribution. Needless to say, a successful detection of gravitational waves necessitates as clean a distribution of the SNRs as possible, with little contamination of the tails. One way of reducing the false alarm rate is by using tighter windows at higher SNRs. This is well-motivated since true high-SNR events will be associated with smaller errors.

The geometric algorithm proposed in this Chapter quite naturally overcomes the drawbacks of the phenomenological method. The algorithm takes into account the correlations amongst the various parameters and deploys parameter- and SNR-dependent ellipsoidal windows defined by the Fisher information matrix using a single parameter. The most important consequence of the new algorithm is a great reduction in the background rate.

4.1.3 Organization of the Chapter

In Sec. 4.2 we present and discuss the new algorithm to identify events in coincidence. The algorithm comprises two steps. The first step consists in associating each trigger with a p -dimensional ellipsoid. In the second step one tests if the ellipsoid associated with a trigger from one detector overlaps, or at least touches, an ellipsoid associated with a trigger from another detector. In Sec. 4.5 we apply the algorithm developed in Sec. 4.2 to the case of a transient chirp signal from a binary black hole. This will help us assess the extent to which the algorithm is helpful in reducing the background. Sec. 4.6 concludes by summarizing the application of the new algorithm in real data analysis pipelines and future prospects.

4.2 A geometric coincidence algorithm

This Section begins with a brief introduction to the geometric formulation of signal manifold and metric introducing the terminology needed in later Sections. The metric so defined helps us in identifying ellipsoidal regions with a given point on the manifold whose size is chosen depending on the signal-to-noise ratio (SNR) and the parameter space region where the point lies. We then compare the volume of the ellipsoid with that of a proper rectangular box enclosing

the ellipsoid and aligned along the coordinate lines.

4.2.1 Scalar Product, Signal Manifold and Metric

The problem of gravitational wave data analysis was addressed in a geometric framework with the intention of understanding parameter estimation [102, 103] and computational requirements for matched filtering [105, 106, 107]. In this framework, one thinks of the outputs of an ensemble of detectors as either finite- or infinite-dimensional vectors depending on whether one considers data streams as a discrete sampled set or the continuum limit of the same, respectively. For the sake of convenience, in this Chapter we shall deal with the continuum limit. However, all our results are applicable to the more realistic case in which detector outputs are treated as finite dimensional vectors. It is easy to see that the set of all detector outputs form a vector space satisfying the usual axioms of a vector space. The starting point of our discussion is the definition of the scalar product. Given any two functions $x(t)$ and $y(t)$, their scalar product $\langle x, y \rangle$ is defined as [26, 97, 98, 31]

$$\langle x, y \rangle = 2 \int_0^\infty \frac{df}{S_h(f)} [X(f)Y^*(f) + X^*(f)Y(f)], \quad (4.1)$$

where $X(f) \equiv \int_{-\infty}^\infty dt x(t) \exp(2\pi i f t)$ is the Fourier transform of the function $x(t)$ (and similarly, $Y(f)$) and $S_h(f)$ is the one-sided power-spectral density of the detector. Scalar product Eq. 4.1 is motivated by the optimal detection statistic of a known signal buried in Gaussian, stationary background [108].

Amongst all vectors, of particular interest are those corresponding to gravitational waves from a given astronomical source. While every signal can be thought of as a vector in the infinite-dimensional vector space of the detector outputs, the set of all such signal vectors don't, by themselves, form a vector space. One can immediately see that the norm of a signal h (i.e., the square-root of the scalar product of a signal with itself) gives the SNR ρ for a signal that is filtered using an optimal template [28, 109]:

$$\rho \equiv \langle h, h \rangle^{1/2} = 2 \left[\int_0^\infty \frac{df}{S_h(f)} |H(f)|^2 \right]^{1/2}, \quad (4.2)$$

where $H(f)$ is the Fourier transform of the signal $h(t)$. In particular, we can define signals \hat{h} of

unit norm:

$$\hat{h} \equiv \frac{h}{\sqrt{\langle h, h \rangle}} = \frac{h}{\rho}, \quad \langle \hat{h}, \hat{h} \rangle = 1. \quad (4.3)$$

Although the set of all signal vectors do not form a vector space, the set of all normed signal vectors (i.e., signal vectors of unit norm) do form a manifold, the parameters of the signal serving as a coordinate system [102, 103, 106, 107]. Thus, each class of astronomical source forms a n -dimensional manifold \mathcal{S}_n , where n is the number of independent parameters characterizing the source. For instance, the set of all signals from a binary on a quasi-circular orbit inclined to the line-of-sight at an angle ι , consisting of non-spinning black holes of masses m_1 , and m_2 , located a distance D from the Earth, initially in the direction (θ, φ) and expected to merge at a time t_C with the phase of the signal at merger φ_C , forms a nine-dimensional manifold with coordinates $\{D, \theta, \varphi, m_1, m_2, t_C, \varphi_C, \iota, \psi\}$, where ψ is the polarization angle of the signal. In the general case of a signal characterized by n parameters we shall denote the parameters by p^α , where $\alpha = 1, \dots, n$.

The manifold \mathcal{S}_n can be endowed with a metric $g_{\alpha\beta}$ that is induced by the scalar product defined in Eq. 4.1. The components of the metric in a coordinate system p^α are given by:

$$g_{\alpha\beta} \equiv \langle \partial_\alpha \hat{h}, \partial_\beta \hat{h} \rangle, \quad \partial_\alpha \hat{h} \equiv \frac{\partial \hat{h}}{\partial p^\alpha}. \quad (4.4)$$

The metric can then be used on the signal manifold as a measure of the proper distance $d\ell$ between near by signals with coordinates p^α and $p^\alpha + dp^\alpha$, that is signals $\hat{h}(p^\alpha)$ and $\hat{h}(p^\alpha + dp^\alpha)$,

$$d\ell^2 = g_{\alpha\beta} dp^\alpha dp^\beta. \quad (4.5)$$

Now, by Taylor expanding $\hat{h}(p^\alpha + dp^\alpha)$ around p^α , and keeping only terms to second order in dp^α , it is straightforward to see that the overlap O of two infinitesimally close by signals can be computed using the metric:

$$\begin{aligned} O(dp^\alpha; p^\alpha) &\equiv \langle \hat{h}(p^\alpha), \hat{h}(p^\alpha + dp^\alpha) \rangle \\ &= 1 - g_{\alpha\beta} dp^\alpha dp^\beta, \end{aligned} \quad (4.6)$$

The metric on the signal manifold is nothing but the well-known Fisher information ma-

trix usually denoted $\Gamma_{\alpha\beta}$, (see, e.g., [108]) but scaled down by the square of the SNR, i.e., $g_{\alpha\beta} = \rho^{-2}\Gamma_{\alpha\beta}$. The information matrix is itself the inverse of the covariance matrix $C_{\alpha\beta}$ and is a very useful quantity in signal analysis. The ambiguity function $\mathcal{A}(\mathbf{d}p^\alpha; p^\alpha)$, familiar to signal analysts, is the overlap function defined above: $\mathcal{A}(\mathbf{d}p^\alpha; p^\alpha) = O(\mathbf{d}p^\alpha; p^\alpha)$. Thus, the equation

$$\mathcal{A}(\mathbf{d}p^\alpha; p^\alpha) = \varepsilon, \text{ or } O(\mathbf{d}p^\alpha; p^\alpha) = \varepsilon, \quad (4.7)$$

where ε ($0 < \varepsilon < 1$) is a constant, defines the ambiguity surface, or level surface. In gravitational wave literature ε , which measures the overlap between two mis-matched signals, is also called the *match*. Using the expression for the overlap O [cf. Eq. 4.6] in Eq. 4.7, we can see that the coordinate distance $\mathbf{d}p^\alpha$ to the ambiguity surface from the coordinate point p^α is related to the proper distance¹ by:

$$g_{\alpha\beta} \mathbf{d}p^\alpha \mathbf{d}p^\beta = 1 - \varepsilon, \quad (4.8)$$

Equivalently, $d\ell = \sqrt{1 - \varepsilon}$. For a given value of the match ε the above equation defines a $(n - 1)$ -dimensional ellipsoid in the n -dimensional signal manifold. Every signal with parameters $p^\alpha + \mathbf{d}p^\alpha$ on the ellipsoid has an overlap ε with the reference signal at p^α .

4.2.2 Coincidence windows

Having defined the metric (equivalently, the information matrix) and the ambiguity function, we next consider the application of the geometric formalism in the estimation of statistical errors involved in the measurement of the parameters and then discuss how that information may be used in coincidence analysis. We closely follow the notation of Finn and Chernoff [26, 97, 98] to introduce the basic ideas and apply their results in the choice of coincidence windows.

Let us suppose a signal of known shape with parameters p^α is buried in background noise that is Gaussian and stationary. Since the signal shape is known one can use matched filtering to dig the signal out of noise. The measured parameters \bar{p}^α will, in general, differ from the true parameters of the signal². Geometrically speaking, the noise vector displaces the signal vector and the process of matched filtering projects the (noise + signal) vector back on to the signal

¹Here the proper distance refers to the distance between the signal $\hat{h}(p^\alpha)$ at the coordinate point p^α and a signal $\hat{h}(p^\alpha + \mathbf{d}p^\alpha)$ with coordinates $p^\alpha + \mathbf{d}p^\alpha$ on the ambiguity surface.

²In what follows we shall use an over-line to distinguish the measured parameters from true parameters p^α .

manifold. Thus, any non-zero noise will make it impossible to measure the true parameters of the signal. The best one can hope for is a proper estimation of the influence of noise.

In general, the posterior probability density function \mathcal{P} of the parameters \bar{p}^α is given by a multi-variate Gaussian distribution:

$$\mathcal{P}(\Delta p^\alpha) d^n \Delta p = \frac{d^n \Delta p}{(2\pi)^{n/2} \sqrt{C}} \exp \left[-\frac{1}{2} C_{\alpha\beta}^{-1} \Delta p^\alpha \Delta p^\beta \right], \quad (4.9)$$

where n is the number of parameters, $\Delta p^\alpha = p^\alpha - \bar{p}^\alpha$, and $C_{\alpha\beta}$ is the covariance matrix, C being its determinant. Noting that $C_{\alpha\beta}^{-1} = \rho^2 g_{\alpha\beta}$, we can re-write the above distribution as:

$$\mathcal{P}(\Delta p^\alpha) d^n \Delta p = \frac{\rho^n \sqrt{g} d^n \Delta p}{(2\pi)^{n/2}} \exp \left[-\frac{\rho^2}{2} g_{\alpha\beta} \Delta p^\alpha \Delta p^\beta \right]. \quad (4.10)$$

where we have used the fact that $C = 1/(\rho^{2n} g)$, g being the determinant of the metric $g_{\alpha\beta}$. Note that if we define new parameters $p'^\alpha = \rho p^\alpha$, then we have exactly the *same* distribution function for all SNRs, except the deviations Δp^α are scaled by ρ .

Let us first specialize to one-dimension to illustrate what region of the parameter space one should associate with a given trigger. In one-dimension the distribution of the deviation from the mean of the measured value of the parameter p is given by:

$$\mathcal{P}(\Delta p) d\Delta p = \frac{d\Delta p}{\sqrt{2\pi}\sigma} \exp \left(-\frac{\Delta p^2}{2\sigma^2} \right) = \frac{\rho \sqrt{g_{pp}} d\Delta p}{\sqrt{2\pi}} \exp \left(-\frac{\rho^2}{2} g_{pp} \Delta p^2 \right), \quad (4.11)$$

where, analogous to the n -dimensional case, we have used $\sigma^2 = 1/(\rho^2 g_{pp})$. Now, at a given SNR, what is the volume V_P in the parameter space such that the probability of finding the measured parameters \bar{p} inside this volume is P ? This volume is defined by:

$$P = \int_{\Delta p \in V_P} \mathcal{P}(\Delta p) d\Delta p. \quad (4.12)$$

Although V_P is not unique it is customary to choose it to be centered around $\Delta p = 0$:

$$P = \int_{(\Delta p/\sigma)^2 \leq r^2(P)} \frac{d\Delta p}{\sqrt{2\pi}\sigma} \exp \left(-\frac{\Delta p^2}{2\sigma^2} \right) = \int_{\rho^2 g_{pp} \Delta p^2 \leq r^2(P)} \frac{\rho \sqrt{g_{pp}} d\Delta p}{\sqrt{2\pi}} \exp \left(-\frac{\rho^2 g_{pp} \Delta p^2}{2} \right), \quad (4.13)$$

where given P the above equation can be used to solve for $r(P)$ and it determines the range of

integration. For instance, the volumes V_P corresponding to $P \simeq 0.683, 0.954, 0.997, \dots$, are the familiar intervals $[-\sigma, \sigma]$, $[-2\sigma, 2\sigma]$, $[-3\sigma, 3\sigma]$, \dots , and the corresponding values of r are 1, 2, 3. Since $\sigma = 1/\sqrt{\rho^2 g_{pp}}$ we see that in terms of g_{pp} the above intervals translate to

$$\frac{1}{\rho} \left[-\frac{1}{\sqrt{g_{pp}}}, \frac{1}{\sqrt{g_{pp}}} \right], \frac{1}{\rho} \left[-\frac{2}{\sqrt{g_{pp}}}, \frac{2}{\sqrt{g_{pp}}} \right], \frac{1}{\rho} \left[-\frac{3}{\sqrt{g_{pp}}}, \frac{3}{\sqrt{g_{pp}}} \right], \dots \quad (4.14)$$

Thus, for a given probability P , the volume V_P shrinks as $1/\rho$. The maximum distance d_{\max} within which we can expect to find “triggers” at a given P depends inversely on the SNR ρ : $d\ell = \sqrt{g_{pp}\Delta p^2} = r/\rho$. Therefore, for $P \simeq 0.954$, $r = 2$ and at an SNR of 5 the maximum distance is 0.4, which corresponds to a match of $\varepsilon = 1 - d\ell^2 = 0.84$. In other words, in one-dimension 95% of the time we expect our triggers to come from templates that have an overlap greater than or equal to 0.84 with the buried signal when the SNR is 5. This interpretation in terms of the match is a good approximation as long as $d\ell \ll 1$, which will be true for large SNR events. However, for weaker signals and/or greater values of P we can’t interpret the results in terms of the match although, of course, the foregoing equation can be used to determine $r(P)$. As an example, at $P \simeq 0.997$, $r = 3$ and at an SNR of $\rho = 4$ the maximum distance is $d\ell = 0.75$ and the match is $\varepsilon = 7/16$, which is significantly smaller than 1 and the quadratic approximation is not good enough to compute the match.

These results generalize to n dimensions too. In n -dimensions the volume V_P is defined by

$$P = \int_{\Delta p^\alpha \in V_P} \mathcal{P}(\Delta p^\alpha) d^n \Delta p. \quad (4.15)$$

Again, V_P is not unique but it is customary to centre the volume around the point $\Delta p^\alpha = 0$:

$$P = \int_{\rho^2 g_{\alpha\beta} \Delta p^\alpha \Delta p^\beta \leq r^2(P,n)} \frac{\rho^n \sqrt{g} d^n \Delta p}{(2\pi)^{n/2}} \exp \left[-\frac{\rho^2}{2} g_{\alpha\beta} \Delta p^\alpha \Delta p^\beta \right]. \quad (4.16)$$

Given P and the parameter space dimension n , one can iteratively solve the above equation for $r(P, n)$. The volume V_P is the surface defined by the equation

$$g_{\alpha\beta} \Delta p^\alpha \Delta p^\beta = \left(\frac{r}{\rho} \right)^2. \quad (4.17)$$

This is the same as the ellipsoid in Eq. 4.8 except that its size is defined by r/ρ . Let us note

Table 4.1: The value of the (squared) distance $d\ell^2 = r^2/\rho^2$ for several values of P and the corresponding smallest match that can be expected between templates and the signal at different values of the SNR.

	$P = 0.683$		$P = 0.954$		$P = 0.997$	
ρ	$d\ell^2$	ϵ_{MM}	$d\ell^2$	ϵ_{MM}	$d\ell^2$	ϵ_{MM}
$n = 1$						
5	0.04	0.9798	0.16	0.9165	0.36	0.8000
10	0.01	0.9950	0.04	0.9798	0.09	0.9539
20	0.0025	0.9987	0.01	0.9950	0.0225	0.9887
$n = 2$						
5	0.092	0.9529	0.2470	0.8677	0.4800	0.7211
10	0.023	0.9884	0.0618	0.9686	0.1200	0.9381
20	0.00575	0.9971	0.0154	0.9922	0.0300	0.9849
$n = 3$						
5	0.1412	0.9267	0.32	0.8246	0.568	0.6572
10	0.0353	0.9822	0.08	0.9592	0.142	0.9263
20	0.00883	0.9956	0.02	0.9899	0.0355	0.9821

the generalization of a result discussed previously, namely that the size of the ellipsoid is not small enough for all combinations of P and ρ and, therefore, it is not always possible to interpret the distance from the centre of the ellipsoid to its surface in terms of the overlap or match of the signals at the two locations except when the distance is close to zero. This is because the expression for the match in terms of the metric is based on the quadratic approximation which breaks down when the matches are small. However, the region defined by Eq. 4.17 always corresponds to the probability P and there is no approximation here (except that the detector noise is Gaussian).

When the SNR ρ is large and $1 - P$ is not close to zero, the triggers are found from the signal with matches greater than or equal to $1 - (r(P, n)/\rho)^2$. Table 4.1 lists the value of r for several values of P in one-, two- and three-dimensions and the minimum match ϵ_{MM} for SNRs 5, 10 and 20. Table 4.1 should be interpreted in the light of the fact that triggers come from an analysis pipeline in which the templates are laid out with a certain minimal match and one cannot, therefore, expect the triggers from different detectors to be matched better than the minimal match. From the Table, we see that when the SNR is large (say greater than about 10) the dependence of the match ϵ_{MM} on n is very weak; in other words, irrespective of the number of dimensions we expect the match between the trigger and the true signal (and for our

purposes the match between triggers from different instruments) to be pretty close to 1, and mostly larger than a minimal match of about 0.95 that is typically used in a search. Even when the SNR is in the region of 5, for low P again there is a weak dependence of ϵ_{MM} on the number of parameters. For large P and low SNR, however, the dependence of ϵ_{MM} on the number of dimensions becomes important. At an SNR of 5 and $P \simeq 0.954$, $\epsilon_{\text{MM}} = 0.80, 0.72, 0.66$ for $n = 1, 2, 3$ dimensions, respectively.

In general, for a given probability P the size of the ellipsoid at an SNR ρ is smaller by a factor ρ compared to that at $\rho = 1$. Thus, the volume in the parameter space in which the measured parameters will lie at a given probability P will scale with the SNR as ρ^{-n} . Therefore, if the goal of an experiment is to have false dismissal probability that is no greater than $1 - P$ then the ellipsoidal windows given by Eq. 4.17 could be employed when testing triggers from different detectors for coincidences. We now have our first result which states that:

When performing coincidence analysis of triggers one should test to see if the associated ellipsoids overlap with each other. These ellipsoids describe the smallest possible volume within which the false dismissal probability is no more than a pre-specified value.

Notice also that the false alarm rate would then go down by ρ^{-n} . Thus, given the false dismissal probability $1 - P$ the size of the ellipsoid further depends on the SNR of the events that are being subject to coincidence analysis, the size shrinking sharply as a function of the event's SNR. Thus we have the second of our results:

The size of the ellipsoids should be chosen in inverse proportion to the signal-to-noise ratio.

However, this latter feature has not yet been implemented in current gravitational wave searches and will be a priority for implementation in future versions of the search pipeline. The final, and practically speaking probably the most important, result is the following:

Our coincidence algorithm reduces the number of tunable parameters from n (where n is the number of parameters) to 1, irrespective of the dimensionality of the signal parameter space.

This parameter μ introduced in Eq. 4.19 essentially scales the volume of the ambiguity ellipsoid – the shape and orientation of which is entirely determined from the metric components. The appropriate value of this parameter can be determined by extensive Monte-Carlo tests where one injects GW signals in interferometer noise and by optimising the detection efficiency vis-a-vis false alarm rate, an acceptable value of μ is arrived at. Having just one parameter greatly simplifies this tuning procedure. Note that, as argued before, μ is SNR-dependent: loud signals with high SNRs are expected to be more consistent in their parameters in different detectors. Thus the ellipsoids associated with these high SNR triggers are expected to overlap (and hence pass coincidence) even if they each have a smaller volume. On the other hand, for weaker signals we need to associate larger ellipsoids in order for them to overlap.

4.3 Overlap of ellipsoids

A key tool in determining coincidences of triggers from two or more detectors is a mathematical algorithm to determine if the ellipsoids associated with triggers either touch or overlap with each other. This algorithm forms the workhorse for identifying coincidence of triggers from two or more detectors.

As stated in Section 4.2.1, triggers resulting from the analysis pipeline are projections of the data by normed signal vectors onto an n -dimensional space \mathcal{S}_n , where n is the number of independent parameters characterizing the source. In the foregoing Section we introduced ellipsoidal regions in the n -dimensional parameter space with their centers at the location of the signal. When we analyze the data, however, we will not know before hand if a signal is present in the data and even when there is one we would not know where its location in the parameter space is. We will have, nevertheless, the knowledge of the location of the triggers in the parameter space. Let us denote the coordinates of a trigger from a detector A as q_A^α , where α is the index on the parameter space. The coincidence analysis proceeds in the following manner. Define an ellipsoidal region $\mathcal{E}(\mathbf{p}_A, \bar{g})$ around each trigger q_A^α by

$$\mathcal{E}(\mathbf{p}_A, \bar{g}) = \{ \mathbf{p}_A \in \mathcal{S}_n \mid (\mathbf{p}_A - \mathbf{q}_A)^T \bar{g} (\mathbf{p}_A - \mathbf{q}_A) \leq 1 \}, \quad (4.18)$$

where $\mathbf{q}_A \in \mathcal{S}_n$ is the position vector of the center of the ellipsoid (i.e., the location of the trigger

from detector A) and \bar{g} is the rescaled metric which we shall refer to as the *shape matrix*. It is related to the metric by

$$\bar{g}_{\alpha\beta} = \mu^2 g_{\alpha\beta} \quad (4.19)$$

where μ^2 is a numerical scaling factor used to expand the linear distances of the ellipsoid while holding the position of the center and the spatial orientation constant. Eq. 4.17 allows us to interpret the parameter μ in terms of the probability P with which the trigger can be expected to be found within the ellipsoid $\mathcal{E}(\mathbf{p}_A, \bar{g})$:

$$\mu^2 = \frac{\rho^2}{r^2(P)}. \quad (4.20)$$

Further, the probabilities P associated with a given μ can be found using Eq. 4.16 when the background noise is Gaussian. However, most detector noise is non-Gaussian and non-stationary and in those cases μ serves as a parameter that must be tuned to achieve a certain detection efficiency or, alternatively, a certain false alarm rate.

Thus, the shape matrix is the scaled metric and encodes the local correlations between the parameters in the neighborhood of the trigger center. It is trivial to check that when $\mu = 1$, Eq. 4.18 defines the interior of the ambiguity ellipsoid previously defined in Eq. 4.7.

Once an ellipsoidal model for the trigger is established, following [110] one can construct a contact function $\mathcal{F}_{AB}(\lambda)$ of two ellipsoids $\mathcal{E}(\mathbf{q}_A, \bar{g}_A)$ and $\mathcal{E}(\mathbf{q}_B, \bar{g}_B)$ (basically, triggers from detectors A and B) as

$$\mathcal{F}_{AB}(\lambda) = \lambda(1-\lambda) \mathbf{r}_{AB}^T [\lambda \bar{g}_B^{-1} + (1-\lambda) \bar{g}_A^{-1}]^{-1} \mathbf{r}_{AB}, \quad (4.21)$$

where $\mathbf{r}_{AB} = \mathbf{q}_B - \mathbf{q}_A$ and $\lambda \in [0, 1]$ is a scalar parameter. The maximum of the contact function over λ in the interval $[0, 1]$ can be shown [110] to be unique. It can also be shown that for two overlapping ellipsoids, the maximum of the contact function is less than 1, i.e.,

$$F = \max_{0 \leq \lambda \leq 1} [\mathcal{F}_{AB}(\lambda)] < 1. \quad (4.22)$$

When $F = 1$, the two ellipsoids 'touch' each other externally.

In the 'coincidence' data-analysis paradigm, given triggers from N detectors ($N \geq 2$), one

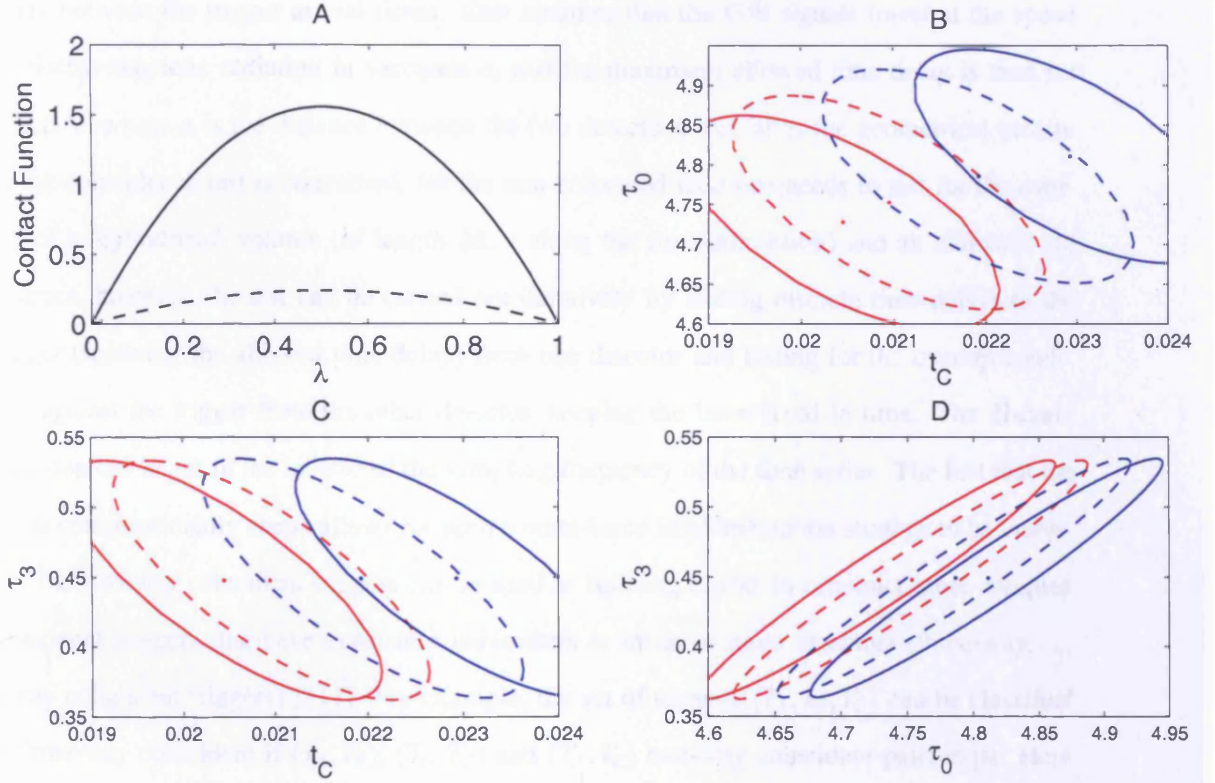


Figure 4.1: Panel A plots the contact function Eq. 4.21 for two pairs of three-dimensional ellipsoids taken from a search for binaries consisting of non-spinning compact objects characterized by parameters (t_c, τ_0, τ_3) [see Sec. 4.5, in particular Eq. 4.34]. Panels B, C and D are the projections of the ellipsoids in (t_c, τ_0) , (t_c, τ_3) and (τ_0, τ_3) orthogonal planes, respectively. Solid lines refer to the case of non-overlapping ellipsoids and dashed lines are for over-lapping (i.e., coincident) triggers. Note that in the latter case the maximum of the contact function is ≤ 1 , which is the test that is carried out to determine if a pair of triggers are in coincidence.

draws up a list of ‘coincident triggers’ for further analysis to test their significance. The simplest coincident triggers consist of those which have ‘consistent’ parameters in two detectors (two-way coincidence). Testing for two-way coincidences for triggers from co-located detectors (e.g., the two LIGO detectors at Hanford) can be accomplished by a single test of Eq. 4.22 on a pair of triggers.

When the detectors are non-co-located, one needs to allow for a non-zero ‘time-of-flight’ delay between the trigger arrival times. One assumes that the GW signals travel at the speed of electro-magnetic radiation in vacuum c , and the maximum allowed time delay is then set to $\pm\Delta/c$, where Δ is the distance between the two detectors. As far as the geometrical picture of the coincidence test is concerned, for the non-co-located case one needs to test for the overlap of a ‘cylindrical’ volume (of length $2\Delta/c$ along the time-dimension) and an ellipsoid. In practice, however, the test can be carried out iteratively by adding discrete time delays to the trigger (spanning the allowed time-delay) from one detector and testing for the overlap condition against the trigger from the other detector, keeping the latter fixed in time. The discrete time step can be set to the inverse of the sampling frequency of the time series. The fact that the test is computationally cheap allows for such a brute-force implementation strategy to be viable.

The two-way coincident triggers can be used as building blocks to construct more complex coincident triggers that have consistent parameters in three or more detectors (three-way, ..., n -way coincident triggers) [111]. For example, the set of triggers (T_A, T_B, T_C) can be classified as three-way coincident if (T_A, T_B) , (T_B, T_C) and (T_A, T_C) two-way coincident pairs exist. Here again, the subscripts $A \neq B \neq C$ are labels on the detectors. This idea can be generalised to determine the list of n -IFO coincident triggers given the list of $(n - 1)$ -way coincidences. It is useful to note that Eq. 4.22 is the only test we need in order to build the entire hierarchy of coincident triggers.

We conclude this Section by drawing attention to two practical issues in implementing this geometrical coincidence test. The first has to do with the algorithm one uses to draw up two-way coincidences. Given the set of triggers from two detectors, one can (a) work with time-ordered triggers and (b) find the maximum length of the bounding box of the ellipsoid along the time dimension over all the triggers such that for any trigger from one detector, the test for overlap is carried out only if a trigger from the other detector occurs at a time that is within twice this

interval. This approach greatly reduces the overall number of overlap tests required to find two-way coincidences. The expression for the length of the sides of the bounding box can be algebraically determined given the shape matrix of the triggers and is explicitly given for 2- and 3-dimensions in the next Section.

The second point is on the numerical implementation of the test of the overlap of ellipsoids where we maximize the contact function over a single parameter λ . Evaluation of the contact function involves matrix inversion which can be quite expensive computationally. Under these circumstances, prior knowledge of the inverse of trigger shape matrices can prove to be more efficient than on-the-fly computation. Brent's minimization method [112, 113] is particularly suitable for fast convergence to the maxima given the well behaved nature of the contact function and is available as part of the GNU Scientific Library [114].

4.4 Expected reduction in false alarm rate

Next, let us consider the reduction in the false alarm rate as a result of using ellipsoidal windows as opposed to rectangular windows. In order to achieve false dismissal probability less than or equal to $1 - P$, a rectangular window has to be at least as large as the box that encloses the ellipsoid. Now the volume of an n -dimensional ellipsoid ($n \geq 2$) whose semi-axes are a_k , $k = 1, \dots, n$, is given by a recursive formula:

$$V_n = \frac{2\pi V_{n-2}}{n} \prod_{k=1}^n a_k, \text{ where } V_0 = 1, V_1 = 2. \quad (4.23)$$

On the other hand, the smallest volume an n -dimensional box that encloses the ellipsoid would be

$$U_n = \prod_{k=1}^n (2a_k) = 2^n \prod_{k=1}^n a_k, \quad (4.24)$$

where a factor of 2 arises since a_k are semi-major axes and the side-lengths of the enclosing box will be twice that value. Thus, the rectangular box's volume is larger than that of the ellipsoid by the factor

$$w \equiv \frac{U_n}{V_n} = \frac{n 2^{n-1}}{\pi V_{n-2}}. \quad (4.25)$$



Thus, in 2-, 3- and 4-dimensions the saving is $4/\pi$, $6/\pi$ and $32/\pi^2$, respectively. However, the real factor could be far greater as the error ellipsoids are generally not oriented along the coordinate axes.

When the ellipsoid is not aligned with the coordinate axes the side-lengths of the bounding box are given by maximizing each coordinate axis over the entire ellipsoidal surface as follows. Starting from Eq. 4.17 one can express the first of the coordinates $p \equiv p^1$ in terms of the other coordinates:

$$g_{11}p^2 + 2g_{1i}p p^i + g_{ij}p^i p^j - \left(\frac{r}{\rho}\right)^2 = 0, \quad i, j = 2, \dots, n, \quad (4.26)$$

which can be solved to obtain

$$p_{\pm} = \frac{1}{g_{11}} \left[-g_{1i} p^i \pm \sqrt{(g_{1i}g_{1j} - g_{11}g_{ij}) p^i p^j + (g_{11}r^2/\rho^2)} \right]. \quad (4.27)$$

For our purposes we only need the ‘plus’ solution. One can then set-up $n - 1$ equations in as many variables by demanding that $\partial p_+ / \partial p^k = 0$, which gives

$$\left(\frac{(g_{1i}g_{1k} - g_{11}g_{ik})(g_{1j}g_{1k} - g_{11}g_{jk})}{g_{1k}^2} - (g_{1i}g_{1j} - g_{11}g_{ij}) \right) p^i p^j = g_{11} \frac{r^2}{\rho^2}. \quad (4.28)$$

These are again quadratic equations that must be solved (simultaneously) for the coordinates p^j , $j = 2, \dots, n$. The resulting (positive) roots, denoted p_1^j can be substituted in Eq. 4.26 to obtain the half-side-length of the ellipse. We shall next give explicit expressions for the side-lengths of the enclosing box in two and three dimensions. In higher dimensions the expressions are rather cumbersome but the general procedure outlined above can be used to compute the volume of the bounding box in all cases.

The side-lengths of the bounding box are given in two dimensions by

$$x = 2\sqrt{\frac{\bar{g}_{22}}{|\bar{g}|}}, \quad y = 2\sqrt{\frac{\bar{g}_{11}}{|\bar{g}|}}, \quad (4.29)$$

and in three dimensions by:

$$x = 2\sqrt{\frac{(\bar{g}_{23}^2 - \bar{g}_{22}\bar{g}_{33})\bar{g}_{22}}{(\bar{g}_{12}\bar{g}_{23} - \bar{g}_{22}\bar{g}_{13})^2 - (\bar{g}_{23}^2 - \bar{g}_{22}\bar{g}_{33})(\bar{g}_{12}^2 - \bar{g}_{11}\bar{g}_{22})}},$$

$$\begin{aligned}
y &= 2\sqrt{\frac{(\bar{g}_{13}^2 - \bar{g}_{11}\bar{g}_{33})\bar{g}_{11}}{(\bar{g}_{12}\bar{g}_{13} - \bar{g}_{11}\bar{g}_{23})^2 - (\bar{g}_{13}^2 - \bar{g}_{11}\bar{g}_{33})(\bar{g}_{12}^2 - \bar{g}_{11}\bar{g}_{22})}}, \\
z &= 2\sqrt{\frac{(\bar{g}_{12}^2 - \bar{g}_{11}\bar{g}_{22})\bar{g}_{11}}{(\bar{g}_{12}\bar{g}_{13} - \bar{g}_{11}\bar{g}_{23})^2 - (\bar{g}_{12}^2 - \bar{g}_{11}\bar{g}_{22})(\bar{g}_{13}^2 - \bar{g}_{11}\bar{g}_{33})}}.
\end{aligned} \tag{4.30}$$

4.5 Application to coalescing binaries

Inspiralling compact binaries are one of the most promising candidates for detection by the laser interferometric detectors. It will, therefore, be interesting to investigate the gains of using the new coincidence method in such searches. For the purpose of our discussion, it will suffice to use a simple model of the signal. We shall use the Fourier representation of the waveform from a binary consisting of non-spinning compact objects on a quasi-circular orbit in which post-Newtonian (PN) corrections to the amplitude are neglected, but corrections to the phase are included to the desired order. This waveform is calculated using the stationary phase approximation, and is of the form:

$$\tilde{h}(f) = \frac{AM^{5/6}}{D\pi^{2/3}} \sqrt{\frac{5\eta}{24}} f^{-7/6} \exp \left[i\Psi(f; t_C, \phi_C, k) + i\frac{\pi}{4} \right], \tag{4.31}$$

$$\Psi(f) = 2\pi f t_C + \phi_C + \sum_k \lambda_k f^{(k-5)/3}. \tag{4.32}$$

Waveforms of this type at 2PN order [49, 48] have been used in previous searches for binary neutron star inspirals [92], and are currently being used in searches for compact binary inspirals with a total mass of $< 35M_\odot$. Moreover, the metric computed for such a waveform has been shown to be approximately valid for a range of physical approximants [115, 29]. At the 2PN order, the coefficients λ_k are given by the following expressions:

$$\begin{aligned}
\lambda_0 &= \frac{3}{128\eta(\pi M)^{5/3}}, \quad \lambda_1 = 0, \quad \lambda_2 = \frac{5}{96\pi\eta M} \left(\frac{743}{336} + \frac{11}{4}\eta \right), \\
\lambda_3 &= \frac{-3\pi^{1/3}}{8\eta M^{2/3}}, \quad \lambda_4 = \frac{15}{64\eta(\pi M)^{1/3}} \left(\frac{3058673}{1016064} + \frac{5429}{1008}\eta + \frac{617}{144}\eta^2 \right),
\end{aligned} \tag{4.33}$$

where M is the total mass of the system, and η is the symmetric mass ratio, which is defined as $\eta \equiv m_1 m_2 / M^2$.

The metric required for determining coincidence in the case of non-spinning binaries is

that in the 3-dimensional space of (t_C, τ_0, τ_3) , where τ_0 and τ_3 are the chirp times, which are a convenient way of parameterizing the masses of the binary system. They are given by

$$\tau_0 = \frac{5}{256\pi f_L \eta} (\pi M f_L)^{-5/3}, \quad \tau_3 = \frac{1}{8f_L \eta} (\pi M f_L)^{-2/3}, \quad (4.34)$$

where f_L is the low frequency cutoff.

In obtaining the metric, it proves to be more convenient to use parameters $(t_C, \theta_1, \theta_2)$, where $\theta_1 \equiv 2\pi f_L \tau_0$, and $\theta_2 \equiv 2\pi f_L \tau_3$. This metric was obtained by Owen in [106]. Here, Eq. 4.6 was used, and the phase ϕ_C maximized over to give the expression for the metric:

$$g_{\alpha\beta} = \frac{1}{2} \left(\mathcal{J} [\Psi_\alpha \Psi_\beta] - \mathcal{J} [\Psi_\alpha] \mathcal{J} [\Psi_\beta] \right), \quad (4.35)$$

where Ψ_α is the derivative of the Fourier phase of the inspiral waveform with respect to parameter θ_α . \mathcal{J} is the moment functional of the noise PSD, which is defined for any function $a(x)$ as:

$$\mathcal{J}(a) \equiv \frac{1}{I(7)} \int_{x_L}^{x_U} \frac{a(x) x^{-7/3}}{S_h(x)} dx. \quad (4.36)$$

$I(q)$ is the q th moment of the noise PSD, which is defined by:

$$I(q) \equiv S_h(f_0) \int_{x_L}^{x_U} \frac{x^{-q/3}}{S_h(x)} dx, \quad (4.37)$$

where $x \equiv f/f_0$, f_0 being a fiducial frequency used to set the range of the numerical values of the functions contained in the integrals. The value of x_L is chosen so that the contribution to the integral for values below x_L would be negligible. $x_U \equiv f_U/f_0$, where f_U is the ending frequency of the inspiral waveform in question. In deriving the explicit expression for the metric, the starting point is the Fourier phase of the waveform in the form [115]:

$$\begin{aligned} \Psi(f; t_C, \theta_1, \theta_2) &= 2\pi f t_C + a_{01} \theta_1 x^{-5/3} + \left[a_{21} (\theta_1/\theta_2) + a_{22} (\theta_1 \theta_2^2)^{1/3} \right] x^{-1} + a_{31} \theta_2 x^{-2/3} \\ &+ \left[a_{41} (\theta_1/\theta_2^2) + a_{42} (\theta_1/\theta_2)^{1/3} + a_{43} (\theta_2^4/\theta_1)^{1/3} \right] x^{-1/3}, \end{aligned} \quad (4.38)$$

where the coefficients a_{km} are given by

$$\begin{aligned} a_{01} = \frac{3}{5}, \quad a_{21} = \frac{11\pi}{12}, \quad a_{22} = \frac{743}{2016} \left(\frac{25}{2\pi^2} \right)^{1/3}, \quad a_{31} = \frac{-3}{2}, \quad a_{41} = \frac{617}{384} \pi^2, \\ a_{42} = \frac{5429}{5376} \left(\frac{25\pi}{2} \right)^{1/3}, \quad a_{43} = \frac{15293365}{10838016} \left(\frac{5}{4\pi^4} \right)^{1/3}. \end{aligned} \quad (4.39)$$

Using the above in Eq. 4.35, one can find an explicit expression for the metric. This expression is too unwieldy to write here, but it can be obtained by utilising the fact that, since the Fourier phase is a polynomial function, \mathcal{J} can be expanded in terms of normalised moments J , where

$$J(p) \equiv \frac{I(p)}{I(7)}. \quad (4.40)$$

To assess the potential gains of using this coincidence method for inspiral analysis, it is useful to consider the difference in volume between the ellipsoidal region defined by \bar{g} , and its bounding box aligned with the co-ordinate axes (t_C, τ_0, τ_3) . This ratio can be calculated with the help of Eqs. (4.30). Fig. 4.2 shows how this ratio varies across the (τ_0, τ_3) space in the case of Initial and Advanced LIGO, Virgo and Einstein Telescope (a third generation European detector that is currently being designed). It can be seen that for most of the parameter space, the volume of the bounding box is an order of magnitude larger than the volume of the ellipsoid; however, in certain regions, corresponding to high masses, this ratio can be as large as two orders of magnitude. This suggests that significant reductions of the background can be achieved by using ellipsoidal windows. Runs on example data sets suggest that in practice, the reduction in background coincident triggers due to using such a coincidence method will be a factor of ~ 10 .

To assess the improvement in the confidence in any candidate detection, it is helpful to look at how reducing the background rate by a factor of k will improve the odds O of a detection

$$O(h|D) = \frac{P(h|D)}{P(0|D)}, \quad (4.41)$$

where $P(h|D)$ is the posterior probability of a signal h being present given the set of triggers D has been obtained, and $P(0|D)$ is the probability of there being no signal given D . We take the accidental trigger rate to be a Poisson process, with a trigger rate prior to reduction λ . Assuming that the detection efficiency is not affected by the reduction in the trigger rate, we see that the




Figure 4.2: The \log_{10} of the ratio of the volume of the bounding box to the volume of the ellipsoid as a function of location in (τ_0, τ_3) space. The plots shown are (clockwise from top-right) for the initial LIGO, advanced LIGO, Virgo and Einstein Telescope. The low frequency cutoff is chosen to be 20 Hz.

odds improves by the following factor:

$$\frac{O(h|D)_{\lambda/k}}{O(h|D)_\lambda} = \frac{1 - e^{-\lambda T}}{1 - e^{-\lambda T/k}}, \quad (4.42)$$

where T is the duration of the run.

The factor by which the odds of a signal being present improves by reducing the false alarm rate by a factor of k depends on how high the false alarm rate was to start with. If the initial false alarm rate is low ($\lambda T \ll 1$), the improvement in the odds approaches the factor k . However, for high false alarm rates, the improvement becomes less marked, tending to a factor of 1 as $\lambda T \rightarrow \infty$.

4.6 Summary and Conclusions

A new method of coincidence analysis is proposed in which, instead of the rectangular windows on parameters conventionally used, ellipsoidal windows are employed based on the metric defined on the signal manifold. This allows us to use windows of appropriate size depending on the location in the parameter space, instead of using a phenomenological ‘best fit’ choice of windows across the entire space. The algorithm has a massive practical advantage in that it requires the tuning of only one parameter irrespective of the number of dimensions of the parameters. This contrasts with the conventional method that required us to tune nearly as many parameters as the dimension of the parameter space. In addition, the method allows us to take into account covariances between parameters, thus significantly reducing the volume enclosed within the windows. In particular, for the case of non-spinning compact binary coalescences in Initial LIGO, it is expected that the use of such a method will reduce the background rate of coincident triggers by roughly an order of magnitude. By also incorporating SNR-dependence into the size of the windows, the background of high SNR events can be reduced even further.

Chapter 5

Searching for high-mass binary coalescences in LIGO S5 data

The current generation of interferometric gravitational wave detectors are taking data at unprecedented sensitivities. Throughout the fifth science run (S5), the LIGO detectors have been taking data at their design sensitivity. Therefore, there is a genuine chance of a detection occurring in the data currently being analysed. Among the most promising sources for detection are compact binary coalescences (CBCs) consisting of neutron stars and/or black holes. The signals emitted by such systems consist of 3 distinct phases: the adiabatic inspiral phase; the merger phase; and the ringdown of the final black hole. The inspiral phase is well modelled using the post-Newtonian approximation, and the ringdown phase can be modelled as the ringing of the quasinormal modes of the final black hole. However, obtaining the description of the merger phase requires solving the full non-linear equations of general relativity to a high degree of accuracy. Recent advances in numerical relativity have allowed us to solve the merger phase of a coalescing binary system, and provide detection templates for use in future searches. This can enhance the reach of the detectors by a factor of 2, leading to an increase in the event rate of a factor of 8. An analytic approach to obtaining the approximate form of the merger signal - the effective one-body (EOB) approach - was taken by Buonanno and Damour[61, 24]. The recently obtained full numerical solutions have been seen to agree remarkably well with the predictions of the EOB approach, as can be seen in Fig. 5.2.

The rate at which signals from CBCs are expected to be detected depends on a number




Figure 5.1: The horizon distance for a binary of equal component masses plotted as a function of the total mass of the binary, for the first 3 months of S5. The horizon distance is defined as the distance at which an optimally oriented, optimally located binary will have an expected SNR of 8. The horizon distance has been plotted for integrating up to the last stable orbit (LSO), and up to the light ring. It can be seen that integrating up to the light ring greatly enhances the sensitivity to more massive systems.




Figure 5.2: The left-hand panel shows the waveform produced by a coalescing binary with equal mass, non-spinning components, generated at 2.5PN order using the effective one-body approach[24]. A ringdown signal for the final black hole consisting of 1 quasinormal mode has also been included. The right-hand panel shows the numerical waveform for an equal-mass binary, where the components had small spins, obtained by Pretorius[116]. (Figures taken from [117].)




Figure 5.3: The search for inspiralling compact binary systems is split into two regions. The low mass region (enclosed by the solid black line) is searched using frequency-domain templates in the stationary phase approximation at the 2PN order. The high mass region (enclosed by the dotted red line) is searched using 2PN effective one-body templates generated in the time-domain. Using the effective one-body approach allows us to include the merger phase in the template.

of factors, including the sensitivity of the detectors, and the astrophysical model used for the distribution of such sources in the universe. The typical assumption is that the rate of compact binary coalescences in a given volume is proportional to the star formation rate, and will therefore follow the blue light luminosity. For binary neutron star (BNS) systems, it is possible to infer merger rates per L_{10} , where L_{10} is 10^{10} times the blue Solar luminosity, based on the BNS systems which have been observed. However, the absence of observational evidence for systems consisting of a neutron star and a black hole (BHNS), and two black holes (BBH), introduces a large degree of uncertainty in any expected rate, although population synthesis models suggest that such systems should exist. Moreover, there is also uncertainty about the assumption that merger rates are proportional to blue light luminosity; in particular, it has been suggested that this assumption will underestimate the contribution to merger rates from elliptical galaxies[118]. Nonetheless, despite the uncertainties, the current best estimates of merger rates are $10 - 170 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ for BNS, $0.1 - 15 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ for BHNS, and $0.15 - 10 \times 10^{-6} \text{yr}^{-1} L_{10}^{-1}$ for BBH.

There are currently a number of search efforts underway looking for signals from CBCs.

Firstly, there are the inspiral searches, which are carried out by matched filtering the data with templates based on the inspiral part of the waveform. These searches are performed over two overlapping mass regions, as shown in Fig. 5.3. The low mass region contains BNS, BHNS, and BBH systems up to a total mass of $35M_{\odot}$. It is searched using templates based on second order restricted post-Newtonian inspiral waveforms in the stationary phase approximation. These templates are of the same type as in the case of the search for BNS in the S3 and S4 data sets[27]. The high-mass region contains BHNS and BBH systems with total mass in the range $25M_{\odot} \leq M \leq 100M_{\odot}$. Such signals typically have an ending frequency near the the most sensitive region of the detector noise curve. Additionally such signals may only appear in the detectors sensitive bandwidth for a very short time. Therefore, it is important to integrate the templates as far as possible, and include the plunge waveform. To this end, the search uses effective one-body (EOB) templates integrated in the time-domain.

In addition to these blind searches, there are also targetted searches being pursued. It has been suggested that short hard gamma-ray bursts (GRBs) could be associated with coalescences of BNS or BHNS systems. Thus, there is also a search taking place which looks for compact binary coalescences in coincidence with GRB events. The known sky location and time of such signals allows one to dig deeper into the noise, and obtain a higher detection efficiency for a given false alarm rate than would otherwise be possible. A description of a recently performed ‘quick look’ search performed on LIGO data coincident with GRB 070201, whose electromagnetically determined sky position was coincident with the spiral arms of the Andromeda galaxy, can be found in [119].

The searches detailed above are all taking place using templates that are derived assuming the objects within the binary are non-spinning. However, it is known that, for BHNS and BBH systems, the effect of spin on the amplitude and phase of the observed gravitational radiation can be significant. Very little is currently known about the distribution of spins of BHs in binaries; they could potentially be very large. Therefore, it could prove to be of great importance to include spin effects in templates to search for such systems. In practice, due to the computational cost of implementation, it is currently impossible to perform a search using a template bank across all physical parameters required to fully take into account the spin of the objects. To reduce this cost for current searches incorporating spin effects, two approaches have been

taken to reduce the effective dimensionality: introduction of a detection template family (DTF), which captures the essential characteristics of the inspiral signal, while depending on fewer (phenomenological) parameters[120], as was done for the search for systems containing spinning objects in S3 LIGO data[121, 122]; and to use a quasi-physical template family based on the PN equations of motion for a binary system with a single spinning object[123, 124]. Both approaches are under assessment and development for current and future searches; in particular, for the case of the DTF, it is important to assess whether the increased false alarm rate which results from using such templates is a price worth paying for the associated gain in detection efficiency.

As well as the searches above, there is also a search for black-hole ringdowns. For this, templates are used which model the ringing of quasinormal modes of the final black hole following merger. The ringdown search becomes particularly useful for systems $\gtrsim 100M_{\odot}$.

The focus of this Chapter will be on the search for non-spinning high-mass CBCs in data from the first calendar year of the LIGO S5 run. This search is currently an ongoing endeavour, so here we give the current status of various aspects of this search. In particular, results given in this Chapter will only involve a subset of the data, known as the *playground data*. The organisation of the Chapter will be as follows: first we will discuss the details of the data set being analysed, and the pipeline used to perform the search. Next, we will discuss the tuning of the various parameters involved in performing the search. Following this, we will discuss the accuracy of parameter recovery for simulated signals injected into the data. Finally, we will discuss the current status of the search, and give a projection of the expected final result.

5.1 The S5 first calendar year

5.1.1 Details of the data analysed

The S5 run commenced on 4th November 2005 with only the H1 and H2 detectors running. They were joined in the run by the L1 detector on 14th November 2005. The first calendar year searches analyse data from the start of the run up to the anniversary of L1 joining the run (i.e. up to 14th November 2006). For data to be analysed, it is required that at least two instruments have data suitable for analysis for that time. For initial investigations and tuning purposes, data from a fraction of the total time, the *playground data* is used. The playground data consists of

approximately 10% of the total data, and is chosen by selecting 600 seconds of data every 6370 seconds. More formally, the playground data set \mathcal{T} is defined as [125]

$$\mathcal{T} \equiv \{[t_n, t_n + 600) : t_n = 729273613 + 6370n, n \in \mathbb{Z}\}. \quad (5.1)$$

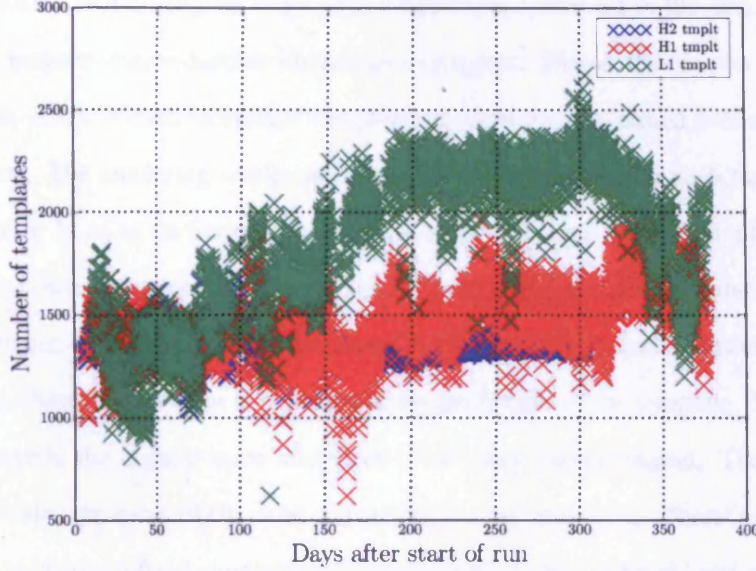
To avoid any bias which may be introduced via the tuning process, any final upper limits will be calculated using only non-playground data. In searching for potential detection candidates, all data will be analysed.

A list of times for which the instruments have poor data quality has been produced. These times are grouped into different categories, depending on the cause of the poor data quality. Depending on the category the times belong to, the data for that particular instrument will either not be analysed, or any triggers produced will be vetoed post-generation. In certain circumstances, it is possible that a signal arriving during the vetoed times could nonetheless be detected and validated; however, these vetoed times will be ignored in computing final upper limits. In addition to these, there is a further set of data quality flags which are not significant enough to veto any triggers produced, but may prove useful in following up candidate events.

5.2 Structure of the pipeline

A brief description of the pipeline used for inspiral searches was given in Chapter 3, and illustrated in Fig. 3.1. We therefore will focus on the specific details of the pipeline regarding this particular search, rather than the general topology of the pipeline itself. Firstly, it is necessary to break the data up into analysis segments. This is done by breaking up the data into stretches of 2048 s in duration. To perform the analysis, these stretches are then broken up into 15 overlapping segments of 256 s duration. To avoid issues with corruption of data, the first and last 64 s in a segment are not used for trigger generation. To ensure this doesn't lead to data ignored in the analysis, the 2048 s stretches are selected to overlap by 128 s in contiguous data. At the start and end of a science segment, to avoid any transients which might arise, we ignore the first and last 8 s of data respectively. Note that this means that science segments of less than 2064 s will not be analysed. For the analysis of playground data, as described previously, triggers are only kept for 600 s stretches of data; however, for the calculation of the power-spectral density,

Figure 5.4: A plot of the number of templates required for the high-mass coalescing binary search vs time for S5 playground data.



2048 s stretches are still used.

Once we have the data segments, we generate the template banks which will be used in the search. Template banks are generated for each detector, and for each 2048 s data stretch, to take into account fluctuations of the detector power-spectral densities. For the high-mass coalescing binary search, the banks are generated to cover the range of component masses $1M_{\odot} \leq m \leq 99M_{\odot}$, with a total mass range $25M_{\odot} \leq M \leq 100M_{\odot}$, with a minimal match of 0.97. The templates are placed on a hexagonal lattice using the metric defined on the parameter space, as described in [29]. Figure 5.4 shows how the number of templates required in each detector changed with time for the times being analysed. The fluctuation of this number highlights the non-stationarity of the noise.

5.2.1 Filtering

Following the generation of the templates, the data is then filtered as described in Section 1.5.1. For this particular search, we use 2PN effective one-body templates, generated in the time-domain, to model the expected waveform. The process of filtering each template will produce

an SNR time-series, $\rho(t)$. If this exceeds the chosen threshold, ρ^* , we keep it as a *gravitational wave trigger*. In the case of the high-mass search this threshold is set to a value of 6. Performing such filtering across the entire template bank will produce a prohibitively large number of triggers, which will mostly be due to noise transients contained in the data. Therefore, it is necessary to perform data reduction on this set of triggers. This is done in two ways: firstly, for the SNR time-series of each individual template, clustering is performed over a sliding window of 3 s duration. The clustering works as follows: if we have a trigger such that $\rho > \rho^*$, if this trigger is within 3 s of an earlier trigger from the same template, which had a larger value of ρ , we discard the current trigger; however, if the earlier trigger has a smaller value of ρ , we discard the earlier trigger. Note that in other searches for compact binary coalescences on S5 data, the length of the clustering window is determined by the length of the template. However, for the high-mass search, the highest mass templates are of very short duration. These highest mass templates are also the most likely to be triggered by noise transients. Therefore, in this case, it is necessary to choose a fixed clustering window to achieve the required level of data reduction.

Once this clustering has been performed, the surviving triggers are then clustered across the parameters (t_C, τ_0, τ_3) , where t_C is the coalescence time, and τ_0 and τ_3 are the chirp times, which are a parametrization of the masses of the system. The clustering works by making use of the metric on the parameter space, and assigning triggers which are ‘near’ to each other in this space as being part of a cluster. The trigger with the highest SNR in each cluster is then kept, while the rest are thrown away. A more complete description of this clustering algorithm can be found in [30].

5.2.2 Coincidence analysis

To reduce the number of false triggers further, those which survive the filtering stage are then subjected to a coincidence test. For triggers in different detectors to be considered coincident, we require that the parameters of the triggers agree to within a certain tolerance. In previous searches, this was achieved by performing large-scale Monte-Carlo simulations, to tune windows on each of the individual parameters. However, in this case, we use the coincidence algorithm detailed in Chapter 4. This approach has two advantages: firstly, it reduces the background by a factor of ~ 10 for comparable detection efficiency, and secondly, it only requires the

tuning of a single parameter, the parameter μ^2 in Eq. 4.19¹. In fact, we actually choose to tune $v \equiv 1/\mu^2$, which, in the current implementation within the LIGO analysis pipeline, is known as the *e-thinca parameter*. For triggers to survive this stage, we require coincidence in at least two detectors.

5.2.3 Signal-based vetoes

If there are any triggers which survive the first coincidence stage, they are then re-filtered, this time with signal-based vetoes applied to further reduce the number of false triggers. So far, two such tests have been applied on the high-mass search: the χ^2 veto, and the r^2 veto. Here, we briefly describe each of these tests.

The χ^2 veto

The chi-squared test can be thought of as a ‘goodness-of-fit’ test between the data which generated the trigger, and the template itself [126, 127, 128, 27, 129]. The veto works as follows: firstly, we divide up the template into p frequency bands, such that each band contains the same quantity of power, i.e., for orthogonal templates $h_c(t)$ and $h_s(t)$, we break them up into bands u_j and v_j respectively, where $j = 1, \dots, p$, such that

$$\langle u_i, u_k \rangle = \frac{1}{p} \delta_{ik}, \quad (5.2)$$

$$\langle v_i, v_k \rangle = \frac{1}{p} \delta_{ik}, \quad (5.3)$$

$$\langle u_i, v_k \rangle = 0. \quad (5.4)$$

$$(5.5)$$

By filtering each of these bands against the data $x(t)$, we obtain the quantities

$$y_i = \langle x, u_i \rangle, \quad (5.6)$$

$$z_i = \langle x, v_i \rangle. \quad (5.7)$$

¹In the previous chapter, the parameter μ^2 depended on ρ^2 . However, the current implementation within the LIGO analysis pipeline ignores the dependence on ρ^2 .

By definition, it is obvious that

$$\frac{(y^2 + z^2)}{\sigma^2} = \rho^2, \quad y = \sum_{i=1}^p y_i, \quad z = \sum_{i=1}^p z_i. \quad (5.8)$$

We then define the quantities Δy_i and Δz_i as

$$\Delta y_i = y_i - \frac{y}{p}, \quad (5.9)$$

$$\Delta z_i = z_i - \frac{z}{p}. \quad (5.10)$$

We then define the χ^2 statistic as

$$\chi^2 = \frac{p}{\sigma^2} \sum_{i=1}^p [(\Delta y)^2 + (\Delta z)^2]. \quad (5.11)$$

It is clear that, for a perfectly matched signal and template, in the absence of noise, the value of χ^2 will be zero. In the presence of noise only $x(t) = n(t)$, the statistic will be χ^2 -distributed with $2p - 2$ degrees of freedom. In fact, it can be shown that, even in the presence of a signal and noise, $x(t) = n(t) + h(t)$, the statistic still obeys a χ^2 -distribution with $2p - 2$ degrees of freedom. However, if the data contains a spurious artefact which is picked up with a high SNR, but does not perfectly match the template, this will lead to a high value of χ^2 . Thus, this statistic is a useful discriminant for distinguishing between signals and noise[126].

So far, we have assumed that a signal and template are perfectly matched; however, this is very unlikely to be the case. Firstly, this is due to the discrete nature of the template banks used in the search. Secondly, there will be differences between any real signal detected, and the templates used to approximate them. To take into account these differences, in practice, we threshold on the value χ_*^2 , where

$$\chi^2 < \chi_*^2(p + \delta^2 \rho^2), \quad (5.12)$$

where δ^2 is a parameter used to take into account the mismatches between templates and potential signals. The values of δ^2 and χ_*^2 used for the search are tuned using Monte-Carlo simulations. The status of this tuning will be discussed in Section 5.3.

The r^2 veto

The r^2 -veto allows us to further distinguish between true signals and instrumental transients.

The veto works by examining the r^2 time-series, which is defined as

$$r^2(t) = \frac{\chi^2(t)}{p}. \quad (5.13)$$

It utilises the fact that the r^2 time-series behaves differently for signals and instrumental glitches as the coalescence time of the trigger is approached. The r^2 -veto works by looking at the r^2 time-series for a duration Δt_* leading up to the coalescence time of a trigger. If the r^2 time-series remains above a threshold value r_*^2 for longer than a duration Δt within the interval Δt_* , the trigger will be vetoed. Further details of how the r^2 veto works can be found in [130].

Typically, the tuning of the r^2 -veto consists of two regimes. Below a certain threshold SNR, the search requires a fixed Δt to be specified. Above this SNR, the threshold value $\Delta t \equiv A\rho^n$, where the r^2 coefficient A , and the power n , both need to be tuned using Monte-Carlo simulations.

Second coincidence

Following the filtering of the data with application of signal-based vetoes, the remaining triggers are then run through the coincidence stage again. Those triggers which survive in at least double coincidence will constitute the output of the search. These surviving coincident triggers are then clustered over a 10s window, keeping the coincident trigger with the maximum combined SNR ρ_c . For previous searches [27], the combined SNR was constructed from effective SNRs, which combined ρ and χ^2 values in a certain way. Such a statistic has not yet been properly assessed for the high-mass search; therefore, where ρ_c has been used in the playground tuning runs, it has been defined by

$$(\rho_c)^2 = \sum_i^N \rho_i^2, \quad (5.14)$$

where ρ_i is the SNR of the trigger in the i th detector (H1, H2 or L1). The clustered coincident triggers constitute the final output of the pipeline, known as the *in-time* coincident triggers.

5.3 Monte-Carlo simulations and background

We now discuss the process of tuning the parameters for each stage of the pipeline. This is done by performing the following actions:

Monte-Carlo injections of signals. Signals are injected into the data streams, which we then attempt to detect using the pipeline. Due to the uncertainties in the waveform during the last few cycles, for the high-mass search, it is particularly important to inject signals of many different families of waveforms, including waveforms from spinning objects.


Estimation of the background. To assess the probability of an in-time coincident trigger being the result of an accidental coincidence of noise triggers, it is necessary to obtain an estimate of the rate of background events. This is done by performing the coincidence analysis with the data streams of the detectors shifted in time relative to one another. In the playground data, we performed 30 time-shifts forward, and the same number backward, and took these to be 60 independent trials from which to estimate the background. H2 and L1 were shifted in time relative to H1 by multiples of 5 s and 10 s respectively.

The goal of the tuning procedure is to obtain the maximum detection efficiency for as low a background rate as possible. Due to the many parameters involved, many of which may not be independent, obtaining an optimal tuning is a very difficult task. Many ongoing projects are underway to attack this problem, using tools such as neural networks. However, for the time being, the approach taken is a phenomenological one, with each parameter being tuned individually by looking at the triggers from injections and time-shifts. The general philosophy is to tune the parameters so that cuts are relatively loose, with as many triggers associated with injections surviving as possible, provided this doesn't also let a prohibitively large number of time-shifted triggers through.

5.3.1 Coincidence step

In tuning the coincidence step, we only need to tune a single parameter, the e-thinca parameter v . In accordance with the philosophy specified above, we tune this parameter such that the cut will be loose, thereby letting as many signals through as is reasonable, while leaving any resulting accidental triggers to be dealt with by further vetoes. As a guide for tuning this parameter, we

Figure 5.5: A plot showing the effective distance of found (blue crosses) and missed (red crosses) injected effective one-body waveforms vs chirp mass.




note the fact that the parameter v corresponds to $1 - m$, where m is the match. The theoretical studies in stationary, Gaussian noise in Chapter 4 indicated that, at an SNR of 5, we would need to use an $v \sim 0.35$ for 99.7% of triggers associated with signals to be correctly identified as such. However, in practice, the noise is not stationary or Gaussian, and we therefore need to use an v which is significantly larger than this. We choose the value of v to be 0.5 for the high-mass search.

Fig. 5.5 shows the effective distance of found and missed injected effective one-body waveforms against chirp mass, M_{chirp} , defined as

$$M_{chirp} = M\eta^{3/5}, \quad \eta = \frac{m_1 m_2}{M^2}, \quad (5.15)$$

where M is the total mass of the binary system, and m_i are the component masses. It can be seen that, for $M_{chirp} \sim 6M_\odot$, there appears to be an excess of missed events at relatively small effective distances. This feature has not been seen in previous searches. The reason for this phenomenon is that, in the high-mass search, we encounter signals with high total mass, but also small symmetric mass ratio η . These factors conspire to give signals from such systems a

Figure 5.6: A plot showing the χ^2 value vs ρ for time-shifted triggers (black crosses), plus found injected signals for a range of signal approximants (red crosses). Triggers plotted here are for all detectors: H1, H2 and L1. The proposed veto is given by the cyan line.




relatively low amplitude, while also occupying only a small part of the sensitive bandwidth of the detector. Thus, we expect to be less sensitive to these systems.

5.3.2 The χ^2 veto

We will now discuss the tuning of the parameters for the χ^2 veto, namely χ_*^2 and δ^2 . To perform this task, Monte-Carlo simulations of many different families of waveform were performed, as well as performing estimations of the background using time-shifted triggers. These runs were performed using a very loose χ^2 cut, and their output was used to choose the parameters for the veto.

In previous searches, the χ^2 statistic has been calculated for $p = 16$. However, since the high-mass search will look for signals which will only be in the sensitive bandwidth of the detector for a short time, and will only cover a small bandwidth, here we used $p = 10$. Fig. 5.6 shows the χ^2 vs ρ for time-shifted triggers, plus found injections for a range of signal approximants. Since the general philosophy adopted is to tune the veto parameters to be the same for all detectors,

Figure 5.7: A plot showing the r^2 duration Δt vs ρ for time-shifted triggers (black crosses), plus found injected signals for a range of signal approximants (red crosses). Triggers plotted here are for all detectors: H1, H2 and L1.




we plot triggers from H1, H2 and L1 on this figure. It can be seen that there is a lot of mixing between time-shifted triggers and those which are associated with injections; therefore, the χ^2 veto will not be as efficient a discriminator as has been the case in other searches, including the low-mass compact binary search on the first calendar year of S5. However, it is possible to apply a veto which will remove a number of background events, including some with very high values of ρ , while barely affecting the number of found injected signals. The proposed values for use in the χ^2 cut are a χ^2_* of 10, and a δ^2 value of 0.15.

5.3.3 The r^2 veto

Once we have an approximate tuning of the χ^2 , we then set about tuning the r^2 veto. Firstly, we need to choose the r^2 threshold value, r_*^2 , and the time duration Δt_* for which we will examine the r^2 time-series. We require the value of r_*^2 to be higher than the threshold value for the χ^2 veto. Therefore, we choose r_*^2 to be 15, and Δt_* to be 6s.

Fig. 5.7 shows the duration Δt over which the r^2 time series remains above r_*^2 within the time

Figure 5.8: A plot showing $\nu \equiv \Delta t / (0.0075 \rho^{1.124764})$ vs ρ for time-shifted triggers (black crosses), plus found injected signals for a range of signal approximants (red crosses). Triggers plotted here are for all detectors: H1, H2 and L1. The proposed veto is given by the cyan line.



interval Δt_* plotted against ρ . The two regimes detailed above can clearly be seen. As was the case for the χ^2 veto, the separation of the triggers from injected signals and time-shifted triggers is not as clean as for previous searches. However, with appropriate choices of parameters, it is clear that the r^2 veto will be very useful in removing large numbers of background triggers. The proposal for the veto is to throw away triggers with $\Delta t > 0.0075\rho^{1.124764}$. This cut is better illustrated in Fig. 5.8.

5.4 Estimation of parameters

The primary purpose of the S5 high-mass search is the detection of signals. Good estimations of the parameters of the system can be obtained by following up any interesting candidate events with further analysis, such as Markov Chain Monte-Carlo (MCMC) methods. However, it is interesting to assess how well the search itself is performing the estimation of parameters, as an aid in diagnosing potential problems which may exist in search codes; and also as a measure of the reliability of the parameters passed as input to further investigations undertaken on candidate events. Histograms of the errors in the measurement of M_{chirp} , t_C and η for EOB injections can be found in Figs. 5.9, 5.10 and 5.11 respectively.

Fig. 5.12 shows the chirp mass accuracy for EOB injections as a function of the injected chirp mass, for the H1 detector. There are a number of curious features which are highlighted in this plot. The most obvious feature is the systematic underestimation of M_{chirp} when the injected $M_{chirp} \gtrsim 25M_\odot$. One reason for this is that the region of the parameter space in question is covered by only a few templates in the template bank. There is another feature present in this plot, namely a propensity to overestimate the chirp mass for injected chirp masses of $5M_\odot \lesssim M_{chirp} \lesssim 15M_\odot$. Investigations have shown that this is due to asymmetric high-mass systems being detected by templates of similar total mass, but more symmetric (hence greater M_{chirp}).

To see where these biases arise, it is informative to look at the errors in the estimation of τ_0 and τ_3 , defined in Eq. 4.34, as a function of the location in the (τ_0, τ_3) space. Figs. 5.13 and 5.14 show the errors in τ_0 and τ_3 respectively as a function of injected (τ_0, τ_3) . It can be seen in these plots that underestimations in τ_0 and τ_3 , corresponding to overestimations of M_{chirp} come from the left-hand side of the (τ_0, τ_3) space, getting worse moving from bottom to top. This corresponds to the highest mass end of the parameter space, getting worse as

Figure 5.9: A histogram of the errors in M_{chirp} for the H1 detector.

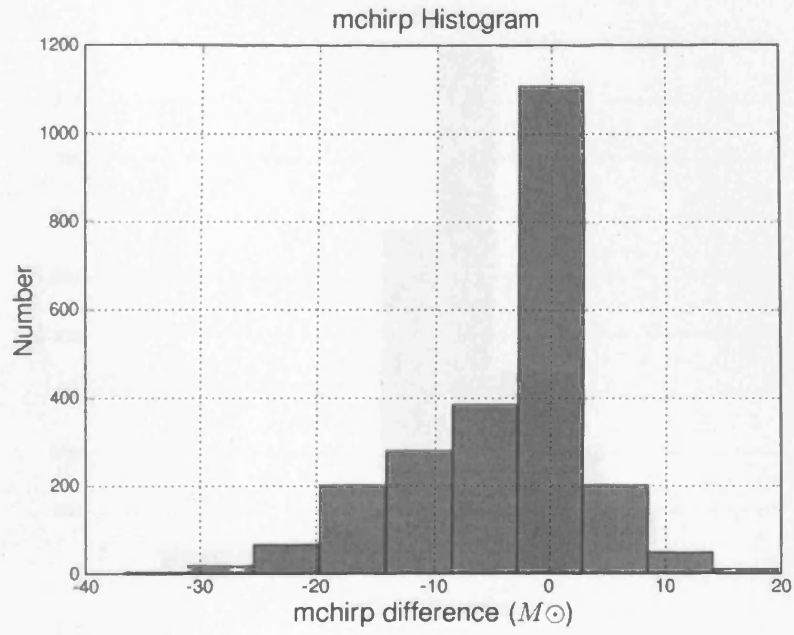


Figure 5.10: A histogram of the errors in t_c for the H1 detector.




Figure 5.11: A histogram of the errors in η for the H1 detector.

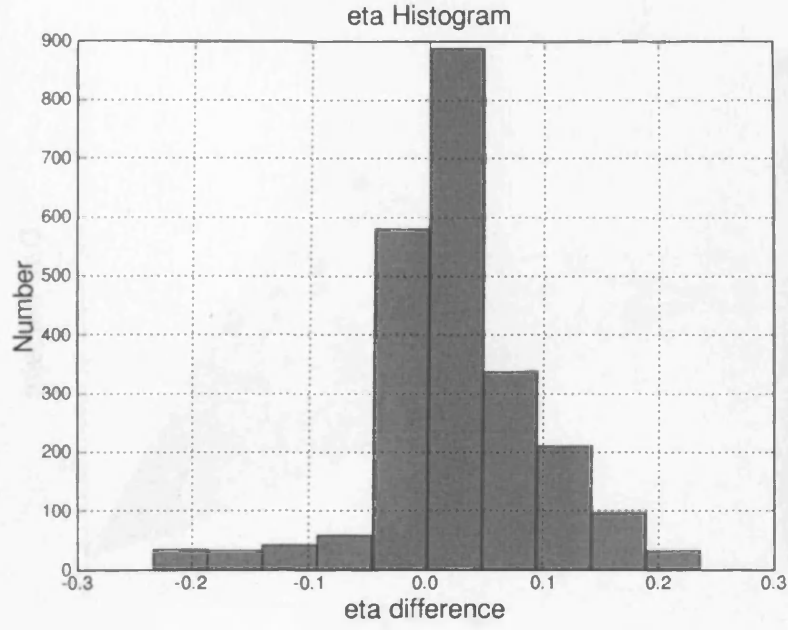


Figure 5.12: A plot of the error in M_{chirp} against the injected M_{chirp}




Figure 5.13: Errors in the measurement of τ_0 as a function of injected (τ_0, τ_3) .




Figure 5.14: Errors in the measurement of τ_3 as a function of injected (τ_0, τ_3) .




Figure 5.15: Errors in the measurement of τ_0 as a function of detected (τ_0, τ_3) .





Figure 5.16: Errors in the measurement of τ_3 as a function of detected (τ_0, τ_3) .



the systems become more asymmetric. Figs. 5.15 and 5.16, which show the errors in τ_0 and τ_3 respectively as a function of detected (τ_0, τ_3) , clearly indicate that these systems are being detected by templates in the bottom-left corner of the (τ_0, τ_3) space, which correspond to high-mass symmetric templates.

5.5 The S5 first-calendar year playground results

Here we discuss the current results for the S5 playground data-set. As the tuning of the parameters for the χ^2 and r^2 vetoes is only preliminary at present, the vetoes as described above are not applied in obtaining the following results; instead, a very loose χ^2 veto was used, with $\chi_*^2 = 100$, and $p = 16$. For the purposes of this discussion, we only use triple-coincident data.

Fig. 5.17 shows the cumulative histograms for the various combinations of double-coincident triggers. Fig. 5.18 shows SNR scatter plots of the double-coincident triggers. It can be seen that there are no double-coincident triggers which are not consistent with the background. There is a statistical excess of H1H2 in-time triggers. This is due to the correlation of noise sources in H1 and H2 due to their location on the same site, which the estimation of the background via time-shifts fails to take into account. It is notable that the scatter plots exhibit high-SNR tails. These have been seen in previous searches, where they have been suppressed by using an effective SNR, which takes into account the values of ρ and χ^2 . The presence of the tails here indicates that such a statistic is worth pursuing in the current search.

There were no triple-coincident triggers found in the time-shifted data used to estimate the background. However, there were two in-time triple-coincident events, which can be seen in Fig. 5.19. This immediately identified these events as candidates for various follow-up checks. These checks indicated that these triggers occurred at times of heightened seismic activity at Hanford in one case, and both Hanford and Livingston in the other case. This, coupled with other checks which have been carried out, reduce the plausibility of either event being identified with a genuine inspiral signal. In fact, the apparent significance of these events likely has the same cause as the excess of H1H2 in-time coincidences, namely the failure to take into account correlated noise sources in H1H2 when generating the background estimate. Any plausible events surviving the final vetoes in the full data set would be subject to rigorous follow-up procedures.




Figure 5.17: Cumulative histograms for H1H2, H1L1, and H2L1 double-coincident in-time triggers (circles), and estimated background from accidental coincidences (squares and 1-standard-deviation ranges) for the preliminary run-through of the high-mass compact binary search in S5 first calendar-year playground data.

Figure 5.18: Scatter plots of ρ_1 vs ρ_2 for H1H2, H1L1, and H2L1 double-coincident in-time triggers (red circles), and time-shifted triggers (black crosses), for the preliminary run-through of the high-mass compact binary search in S5 first calendar-year playground data.





Figure 5.19: Scatter plots of ρ_1 vs ρ_2 for the triple-coincident events, projected onto the H1H2, H1L1, and H2L1 planes. There were no time-shifted triple-coincident events.



5.5.1 Efficiency and projected upper limit

Assuming there are no signals detected in the search, we will then set an upper limit on the rate of compact binary coalescence in the universe for the mass range in question. In this section, we give a brief discussion of how such an upper limit is set, along with giving a projected estimate of the upper limit we may obtain. It is important to stress that the upper limit that will be obtained upon completion of the search will depend on the combined SNR of the loudest event seen, as well as its probability of being a background event. Therefore, the projected upper limit given here should in no way pre-empt that which will be obtained when the search is completed. However, it may serve as an indicator of the order of magnitude we hope to achieve.

The Bayesian upper limit, at a confidence level α , if we have a uniform prior on the rate R , is given by [131, 132]:

$$1 - \alpha = e^{-RT C_L(\rho_{C,max})} \left[1 + \left(\frac{\Lambda}{1 + \Lambda} \right) RT C_L(\rho_{C,max}) \right], \quad (5.16)$$

where $C_L(\rho_{C,max})$ is the cumulative blue luminosity which we are sensitive to at a combined SNR of $\rho_{C,max}$, T is the observation time, and Λ is a measure of the likelihood that the loudest event is due to the foreground. If the loudest event candidate is very likely to be a background event, we get

$$R_{90\%} = \frac{2.3}{T C_L(\rho_{C,max})}, \quad (5.17)$$

where $R_{90\%}$ is the upper limit at a 90% confidence level.

The cumulative luminosity C_L is dependent on the efficiency ϵ of detecting a particular event with combined SNR $> \rho_{C,max}$, and the predicted source luminosity as a function of the effective distances and the chirp mass, $L(\mathcal{D}_H, \mathcal{D}_L, M_{chirp})$. It is obtained by integrating ϵ times $L(\mathcal{D}_H, \mathcal{D}_L, M_{chirp})$ over effective distance and mass. In practice, it is usual to remove the dependence on M_{chirp} by using chirp distances $\mathcal{D}_{chirp} \equiv \mathcal{D}(M_{chirp,o}/M_{chirp})^{5/6}$, where $M_{chirp,o}$ is some fiducial chirp mass. We then have

$$C_L(\rho) = \int_0^\infty \int_0^\infty \epsilon(\mathcal{D}_{chirp,H} \mathcal{D}_{chirp,L}, \rho) L(\mathcal{D}_{chirp,H} \mathcal{D}_{chirp,L}) d\mathcal{D}_{chirp,H} d\mathcal{D}_{chirp,L}. \quad (5.18)$$




Figure 5.20: A plot showing the efficiency of the search evaluated at the threshold combined SNR, estimated by injecting simulated signals into the playground data, as a function of effective distance.

Table 5.1: The amount of time analysed in S5 when at least two detectors were switched on.

Detector combination	Time / hrs
H1-H2-L1	4107
H1-H2	1881
H1-L1	345
H2-L1	383

The efficiency is evaluated by injecting large numbers of simulated signals into the data. A plot of the efficiency as a function of $(\mathcal{D}_H, \mathcal{D}_L)$, evaluated for the threshold combined SNR, can be seen in Fig. 5.20. The distribution of blue luminosity as a function of distance is based on the catalog described in [133]. To obtain the luminosity as a function of \mathcal{D}_{chirp} , it is necessary to assume some form for the astrophysical distribution of the masses of the sources. For the purposes of generating the predicted upper limit, we assume that the total mass of a binary system is uniformly distributed. In this case, we estimate the cumulative luminosity we are sensitive to, $C_L \sim 13700 L_{10}$ for the threshold SNR. Taking the amount of time analysed in the first year of S5 to be as detailed in Table 5.1, we obtain a predicted upper limit $\mathcal{R}_{90\%} \sim 2.2 \times 10^{-4} \text{ yr}^{-1} L_{10}^{-1}$. Although this estimate should be taken with a grain of salt, it indicates that we can expect to do significantly better than the upper limit of $\mathcal{R}_{90\%} = 0.5 \text{ yr}^{-1} L_{10}^{-1}$ which was obtained for S4[27].

5.6 Status and future plans

The S5 high-mass search is still an on-going endeavour. So far, we have used the playground data to assess of the performance of the search, gain a measure of the accuracy of the estimation of parameters, and obtain estimates of the expected background. These have been used to tune the parameters for the various steps of the pipeline, such as the coincidence test, χ^2 veto, and r^2 veto, with the intention of reducing the background as much as possible, while still obtaining a good detection efficiency. We have also looked at the in-time coincident triggers in the playground, obtained using very loose parameters for vetoes, to assess the character of the data.

Going on from here, there are many things which still need to be done. Of prime importance is to investigate the possibility of using an effective SNR in the search. In addition to this, a number of vetoes which may prove useful need to be assessed; for example, requiring consistency in measurements of effective distance in H1 and H2.

Following this, larger scale runs will be required, performing simulations using injected signals, and time-shifted data, using the full data set. This is done for three reasons: to confirm the tuning performed on the playground; to assess the efficiency of the search; and to get a better estimate of the expected background. Once this has been performed, the in-time full data set will be analysed, and any interesting candidate events will be followed up. Assuming the full

data set produces no detections, the final result of the search will be an upper-limit on the rate of high-mass compact binary coalescences. Preliminary predictions of the expected upper limit suggest that we can expect to do significantly better than that which was obtained for S4.

Chapter 6

Conclusion

The search for gravitational waves poses many challenges, theoretically and computationally. This thesis has touched on a number of these issues. We now present a summary of the the main points discussed in this thesis.

In Chapter 2, we investigated a new class of approximants for binary inspiral waveforms. The standard adiabatic approximant uses the post-Newtonian expansions of the energy and gravitational wave flux truncated to the same relative post-Newtonian order. However, when looking at the dynamics of the system, we know that the leading order radiation reaction term enters at 2.5PN. From the viewpoint of the dynamics of the system, the standard approximation at leading order is thus equivalent to keeping the 0PN and 2.5PN terms in the acceleration, but neglecting the 1PN and 2PN terms. This motivated the construction of the new complete adiabatic approximants, which at the leading order, use the 2PN energy function, and therefore correspond in spirit with the dynamics of the system with no missing terms in the post-Newtonian expansion of the acceleration.

To compare the performances of the standard and complete approximants, we compared their overlaps in the test mass case with the exact waveform in the adiabatic approximation. We found that the complete approximants generally brought about a remarkable improvement in the effectualness (i.e. larger overlaps with the exact signal when maximizing over the template parameters) for orders $< 3\text{PN}$. However, for orders $\geq 3\text{PN}$, the standard adiabatic approximants were nearly as good as the complete adiabatic approximants for the creation of effectual templates. The faithfulness (i.e. overlaps with the exact signal, with the parameters of the tem-

plate kept the same as those of the signal) was also generally better for the complete adiabatic approximants than for the standard adiabatic approximants.

We then compared standard and complete non-adiabatic approximants using the Lagrangian models of Buonanno, Chen and Vallisneri in the test mass limit. In this case, we were limited in the scope of our investigation, due to the terms in the post-Newtonian expansion of the acceleration only being available up to 3.5PN order. In addition, in the absence of an exact non-adiabatic signal, we had to compare overlaps with the exact signal in the adiabatic approximation. Given these limitations, one should be wary about making general statements about the trends in effectualness and faithfulness. However, at the orders studied, we have seen that the effectualness of the complete non-adiabatic approximant is usually better than the standard non-adiabatic approximant; but the faithfulness is generally better for the standard non-adiabatic approximant.

A limited extension to the comparable mass case was also provided, where the performances of the standard and complete approximants were assessed by comparing their overlaps with a fiducial waveform. We found that the standard adiabatic approximants achieve the target effectualness of 0.965 at order 3PN. Moreover, if we assume that the comparable mass case is qualitatively similar to the test mass case, then improving the accuracy of the energy function from 3PN to 4PN and/or improving the accuracy of the flux function from 3.5PN to 4PN will not result in a significant improvement in effectualness in the comparable mass case for terrestrial laser interferometric detectors.

In Chapter 3 we looked at a study of two approaches to performing searches for inspiralling compact binaries in a time-critical manner. The first took a step-wise approach to splitting the work amongst the compute nodes. In this approach, the search templates are initially distributed naively amongst the compute nodes, with each node receiving the same number of templates. The splitting of the template bank amongst the nodes for the next step is then adjusted according to the time taken for each job to finish. Thus, the computational load on each node is balanced via an iterative process, where the splitting of the template bank for the next run depends on the timings for the previous run. This simplistic approach has the advantage of requiring little or no modification to the analysis codes, and it is found to do a reasonably good job; however, due to the changing characteristics of the data as a run progresses, as well as changing loads on

computational resources, it is only possible to balance the time taken for each node on average.

The second approach involved a dynamic load-balancing algorithm, in which slave nodes request work of a master node whenever they become idle. Such an approach will balance the load amongst the worker nodes by default. However, implementing such an approach requires significant modifications to analysis codes. In addition, the use of such an approach introduces inter-process communication, which may introduce a performance penalty if the analysis is distributed across several computing resources found in different locations.

In Chapter 4 we developed a new method for coincidence analysis. The algorithm makes optimal use of the variances and covariances that exist amongst the different parameters of the signal in a coincident detection of events. We associate with each trigger ellipsoidal regions which are defined by the metric on the parameter space, or, equivalently, the Fisher information matrix. Triggers from different detectors are said to be in coincidence if their ellipsoids are found to overlap. The use of such a method offers a significant reduction in the background for a comparable detection efficiency. An implementation of such an approach, with the ellipsoid scaling factor being independent of SNR, is currently being employed in the search for compact binary coalescences in LIGO S5 data.

Following this, in Chapter 5 we looked at the current status of the search for high-mass compact binary coalescences in the first calendar year of data from the fifth LIGO science run (S5). This search is using effective one-body (EOB) templates to search for systems with masses in the range $25M_{\odot} \leq M_{total} \leq 100M_{\odot}$. We looked at the preliminary tuning of the parameters of the search. We also looked at the playground data using very loose parameters for vetoes, to assess the character of the data. Even in the absence of a detection, when completed, the search promises to dramatically improve on the upper limit on compact binary coalescences obtained for the previous science run.

Following the completion of the S5 run, commissioning of a series of major improvements and upgrades, known as Enhanced LIGO is about to begin. These improvements will improve the sensitivity of the instruments by a factor of 2. Beyond this, improvements to the Virgo detector, and the commissioning of Advanced LIGO suggest that a concrete detection of gravitational waves is likely to be made in the not too distant future. The first detection will be a major scientific breakthrough, and will no doubt lead to further major discoveries, and allow us to test

our current theories. If the first detection is a binary black hole, it will be the first direct confirmation of the existence of such systems. Moreover, the information contained in the waveform will allow us to perform strong-field tests of general relativity. In conclusion, it promises to be a very exciting time, as a new window on the Universe is opened.

Bibliography

- [1] http://en.wikipedia.org/wiki/Arthur_St Stanley_Eddington. 1
- [2] F.W. Dyson, A.S. Eddington, and C.R. Davidson. A determination of the deflection of light by the sun's gravitational field, from observations made at the total eclipse of may 29, 1919. *Phil. Trans. of the R. Soc. A*, 220:291–333, 1920. 1
- [3] Albert Einstein. The foundation of the general theory of relativity. *Annalen der Physik*, 49:769–822, 1916. 1
- [4] P Ajith, B R Iyer, C A K Robinson, and B S Sathyaprakash. A new class of post-newtonian approximants to the waveform templates of inspiralling compact binaries: Test-mass in the schwarzschild spacetime. *Phys. Rev. D*, 71:044029, 2005. Erratum-ibid 72, 049902(E) (2005). 1
- [5] P. Ajith, Bala R. Iyer, C. A. K. Robinson, and B. S. Sathyaprakash. Complete adiabatic waveform templates for a test-mass in the schwarzschild spacetime: Virgo and advanced ligo studies. *Class. Quant. Grav.*, 22:S1179–S1188, 2005. 1
- [6] L.D. Landau and E.M. Lifshitz. *The classical theory of fields*. Pergamon, Oxford, 1971. 1.1.1, 1.1.2
- [7] Ray d'Inverno. *Introducing Einstein's Relativity*. Oxford University Press, 1992. 1.1.1
- [8] Bernard Schutz. 2005. Post-graduate Lecture Notes, Cardiff University. 1.2, 1.2.2, 1.3
- [9] Jim Hough, Sheila Rowan, and B.S. Sathyaprakash. The search for gravitational waves. *J. Phys. B: At. Mol. Opt. Phys.*, 38(9):S497–S519, 2005. 1.2
- [10] <http://www.geo600.uni-hannover.de>. 1.2, 2.1

- [11] <http://www.ligo.caltech.edu>. 1.2, 2.1
- [12] <http://www.virgo.infn.it>. 1.2, 2.1
- [13] tamago.mtk.nao.ac.jp. 1.2, 2.1
- [14] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophys. J.*, 195:L51–L53, January 1975. 1.2.1
- [15] J.M. Weisberg and J.H. Taylor. Relativistic binary pulsar b1913+16: Thirty years of observations and analysis. In F.H. Rasio and I.H. Stairs, editors, *Binary Radio Pulsars*. ASP Conference Series, 2005. 1.2.1
- [16] Duncan A. Brown. *Searching for Gravitational Radiation from Binary Black Hole MACHOs in the Galactic Halo*. Phd thesis, University of Wisconsin-Milwaukee, 2004. 1.2.2, 1.2.2
- [17] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, Singapore, 1994. 1.2.2
- [18] R. Adhikari. *Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae*. Phd thesis, MIT, Cambridge, MA, 2004. 1.2.2
- [19] C Cutler and K.S. Thorne. An overview of gravitational-wave sources. In N.T Bishop and S.D Maharaj, editors, *Proceedings of GR-16*. World Scientific, 2002. 1.3, 2.1.6
- [20] Sanjit Mitra, Sanjeev Dhurandhar, Tarun Souradeep, Albert Lazzarini, Vuk Mandic, Sukanta Bose, and Steffan Ballmer. Gravitational wave radiometry: Mapping a stochastic gravitational wave background. *To be submitted to Phys. Rev. D*, 2007. 1.3
- [21] V. Kalogera, C. Kim, D.R. Lorimer, M. Burgay, N. D’Amico, A. Possenti, R.N. Manchester, A.G. Lyne, B.C. Joshi, M.A. McLaughlin, M. Kramer, J.M. Sarkissian, and F. Camilo. The cosmic coalescence rates for double neutron star binaries. *Astrophys. J.*, 601:L179–L182, 2004. Erratum, —it ibid 614, L137. 1.4
- [22] R. O’Shaughnessy, C. Kim, T. Frakgos, V Kalogera, and K. Belczynski. Constraining population synthesis models via the binary neutron star population. *Astrophys. J.*, 633:1076–1084, 2005. 1.4

- [23] 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. 1.4, 1.4, 1.4, 1.4, 2.1.1, 2.1.1, 2.1.2, 2.1.4, 2.1.5, 2.1.6, 2.1.6, 6, 2.2.3, 2.5
- [24] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D*, 62:064015, 2000. 1.4, 2.1.3, 2.1.4, 2.3, 2.5, 5, 5.2
- [25] L A Wainstein and V D Zubakov. *Extraction of Signals from Noise*. Prentice-Hall, Englewood Cliffs, 1962. 1.5.1, 1.5.1
- [26] L.S. Finn. Detection, measurement, and gravitational radiation. *Phys. Rev. D*, 46:5236, 1992. 1.5.1, 4.1.2, 4.2.1, 4.2.2
- [27] LIGO Scientific Collaboration. Search for gravitational waves from binary inspirals in s3 and s4 ligo data. 2007. gr-qc/0704.3368v2. 1.5.1, 5, 5.2.3, 5.2.3, 5.5.1
- [28] 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. 1.5.1, 4.2.1
- [29] Thomas Cokelaer. Gravitational waves from inspiralling compact binaries: hexagonal template placement and its efficiency in detecting physical signals. 2007. 1.5.1, 4.5, 5.2
- [30] Anand S. Sengupta, J.A. Gupchup, and Craig A.K. Robinson. Trigscan: A new clustering algorithm for binary inspiral search pipelines. In preparation. 1.5.1, 5.2.1
- [31] C. Cutler and E.E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Phys. Rev. D*, 49:2658–2697, 1994. 2.1.1, 4.1.2, 4.2.1
- [32] E. Poisson. Gravitational-radiation from a particle in circular orbit around a black-hole. i. analytical results for the nonrotating case. *Phys. Rev. D*, 47:1497–1510, 1993. 2.1.1
- [33] E. Poisson. Gravitational radiation from a particle in circular orbit around a black-hole. vi. accuracy of the post-newtonian expansion. *Phys. Rev. D*, 52:5719–5723, 1995. Erratum *Phys. Rev. D* **55**, 7980, (1997). 2.1.1, 2.2.2

- [34] H. Tagoshi and T. Nakamura. Gravitational waves from a point particle in circular orbit around a black hole: Logarithmic terms in the post-newtonian expansion. *Phys. Rev. D*, 49:4016–4022, 1994. 2.1.1
- [35] Masaru Shibata. Gravitational waves induced by a particle orbiting around a rotating black hole: Spin-orbit interaction effect. *Phys. Rev. D*, 48:663–666, 1993. 2.1.1
- [36] T. Tanaka, H. Tagoshi, and M. Sasaki. Gravitational waves by a particle in circular orbit around a schwarzschild black hole. *Prog. Theor. Phys.*, 96:1087–1101, 1996. 2.1.1, 2.2.2, 2.2.2
- [37] M. Sasaki and H. Tagoshi. Analytic black hole perturbation approach to gravitational radiation. *Living Rev. Relativity*, 6:6, 2003. 2.1.1
- [38] T. Damour, P. Jaranowski, and G. Schäfer. Poincaré invariance in the adm hamiltonian approach to the general relativistic two-body problem. *Phys. Rev. D*, 62:021501(R), 2000. Erratum-ibid **63**, 029903(E) (2000). 2.1.1, 2.4.1
- [39] T. Damour, P. Jaranowski, and G. Schäfer. Equivalence between the adm-hamiltonian and the harmonic-coordinates approaches to the third post-newtonian dynamics of compact binaries. *Phys. Rev. D*, 63:044021, 2001. Erratum-ibid **66**, 029901(E) (2002). 2.1.1, 2.4.1
- [40] Thibault Damour, Piotr Jaranowski, and Gerhard Schäfer. On the determination of the last stable orbit for circular general relativistic binaries at the third post-newtonian approximation. *Phys. Rev. D*, 62:084011, 2000. 2.1.1, 2.4.1
- [41] Luc Blanchet and Guillaume Faye. Equations of motion of point-particle binaries at the third post-newtonian order. *Phys. Lett. A*, 271:58, 2000. 2.1.1, 2.4.1
- [42] Luc Blanchet and Guillaume Faye. General relativistic dynamics of compact binaries at the third post-newtonian order. *Phys. Rev. D*, 63:062005, 2001. 2.1.1, 2.4.1
- [43] V.C. de Andrade, L. Blanchet, and G. Faye. Third post-newtonian dynamics of compact binaries: Noetherian conserved quantities and equivalence between the harmonic-coordinate and adm-hamiltonian formalisms. *Class. Quantum Grav.*, 18:753–778, 2001. 2.1.1, 2.4.1

- [44] T. Damour, P. Jaranowski, and G. Schäfer. Dimensional regularization of the gravitational interaction of point masses. *Phys. Lett. B*, 513:147–155, 2001. 2.1.1, 2.4.1, 2.4.1
- [45] Luc Blanchet, Thibault Damour, and Gilles Esposito-Farèse. Dimensional regularization of the third post-newtonian dynamics of point particles in harmonic coordinates. *Phys. Rev. D*, 69:124007, 2004. 2.1.1, 2.4.1, 2.4.1
- [46] Yousuke Itoh and Toshifumi Futamase. New derivation of a third post-newtonian equation of motion for relativistic compact binaries without ambiguity. *Phys. Rev. D*, 68:121501(R), 2003. 2.1.1, 2.4.1, 2.4.1
- [47] Yousuke Itoh. Equation of motion for relativistic compact binaries with the strong field point particle limit: Third post-Newtonian order. *Phys. Rev. D*, 69:064018, 2004. 2.1.1, 2.4.1
- [48] Luc Blanchet, Thibault Damour, Bala R. Iyer, Clifford M. Will, and Alan. G. Wiseman. Gravitational radiation damping of compact binary systems to second post-newtonian order. *Phys. Rev. Lett.*, 74:3515–3518, 1995. 2.1.1, 2.4.2, 4.5
- [49] Luc Blanchet, Thibault Damour, and Bala R. Iyer. Gravitational waves from inspiralling compact binaries: Energy loss and wave form to second post-newtonian order. *Phys. Rev. D*, 51:5360–5386, 1995. 2.1.1, 2.4.2, 4.5
- [50] C.M. Will and A.G. Wiseman. Gravitational radiation from compact binary systems: Gravitational waveforms and energy loss to second post-newtonian order. *Phys. Rev. D*, 54:4813–4848, 1996. 2.1.1, 2.4.2
- [51] Luc Blanchet, Bala R. Iyer, Clifford M. Will, and Alan G. Wiseman. Gravitational wave forms from inspiralling compact binaries to second-post-newtonian order. *Class. Quantum Grav.*, 13:575–584, 1996. 2.1.1, 2.4.2
- [52] Luc Blanchet. Energy losses by gravitational radiation in inspiralling compact binaries to five halves post-newtonian order. *Phys. Rev. D*, 54:1417–1438, 1996. Erratum-ibid. **71**, 129904(E) (2005). 2.1.1, 2.4.2

- [53] Luc Blanchet, Bala R. Iyer, and Benoit Joguet. Gravitational waves from inspiralling compact binaries: Energy flux to third post-newtonian order. *Phys. Rev. D*, 65:064005, 2002. Erratum-ibid **71**, 129903(E) (2005). 2.1.1, 2.4.2
- [54] 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). 2.1.1, 9, 2.2.3, 2.4.2
- [55] 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). 2.1.1, 2.2.3
- [56] Luc Blanchet, Thibault Damour, Gilles Esposito-Farèse, and Bala R. Iyer. Gravitational radiation from inspiralling compact binaries completed at the third post-newtonian order. *Phys. Rev. Lett.*, 93:091101, 2004. 2.1.1, 2.2.3, 2.4.2
- [57] Luc Blanchet and Bala R. Iyer. Hadamard regularization of the third post-newtonian gravitational wave generation of two point masses. *Phys. Rev. D*, 71:024004, 2004. 2.1.1
- [58] Luc Blanchet, Thibault Damour, and Bala R. Iyer. Surface-integral expressions for the multipole moments of post-newtonian sources and the boosted schwarzschild solution. *Class. Quantum Grav.*, 22:155, 2005. 2.1.1
- [59] Luc Blanchet. Gravitational radiation from post-newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 5:3, 2002. 2.1.1
- [60] Luc Blanchet. On the accuracy of the post-newtonian approximation. 2002. In *A relativistic spacetime odyssey*, edited by I. Ciufolini *et al*, World Scientific, p. 411. 4
- [61] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, 1999. 2.1.3, 5
- [62] Thibault Damour, Eric Gourgoulhon, and Philippe Grandclement. Circular orbits of corotating binary black holes: Comparison between analytical and numerical results. *Phys. Rev. D*, 66:024007, 2002. 2.1.3

- [63] E. Gourgoulhon, P. Grandclément, and S. Bonazzola. *Phys. Rev. D*, 65:044020, 2002. 2.1.3
- [64] P. Grandclément, E. Gourgoulhon, and S. Bonazzola. *Phys. Rev. D*, 65:044021, 2002. 2.1.3
- [65] A. Buonanno, Y. Chen, and M. Vallisneri. Detection template families for gravitational waves from the final stages of binary black-holes binaries: Nonspinning case. *Phys. Rev. D*, 67:024016, 2003. Erratum-ibid. **D** 74, 029903(E) (2006). 2.1.3, 2.1.4, 2.3, 2.3, 2.4.1, 2.5
- [66] Eanna E. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences. ii: The waves' information and its extraction, with and without templates. *Phys. Rev.*, D57:4566–4587, 1998. 2.1.4, 2.5, 4.1.2
- [67] Éanna É. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences. i. signal to noise for inspiral, merger, and ringdown. *Phys. Rev. D*, 57(8):4535–4565, Apr 1998. 2.1.4, 2.5
- [68] Thibault Damour, Bala R. Iyer, Piotr Jaranowski, and B. S. Sathyaprakash. Gravitational waves from black hole binary inspiral and merger: The span of third post-newtonian effective-one-body templates. *Phys. Rev. D*, 67:064028, 2003. 2.1.4, 2.5
- [69] Y. Tsunesada et al. On detection of black hole quasi-normal ringdowns: Detection efficiency and waveform parameter determination in matched filtering. *Phys. Rev.*, D71:103005, 2005. 2.1.4, 2.5
- [70] T. Damour, B R Iyer, and B S Sathyaprakash. Improved filters for gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 57:885–907, 1998. 2.1.5, 2.4.1
- [71] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. Frequency-domain p-approximant filters for time-truncated inspiral gravitational wave signals from compact binaries. *Phys. Rev. D*, 62:084036, 2000. 2.1.5, 2.2.3, 2.2.3
- [72] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. A comparison of search templates for gravitational waves from binary inspiral: 3.5pn update. *Phys. Rev. D*, 66:027502, 2002. Erratum-ibid. **D** 72 (2005) 029901. 2.1.5

- [73] Luc Blanchet and Bala R. Iyer. Third post-newtonian dynamics of compact binaries: Equations of motion in the center-of-mass frame. *Class. Quantum Grav.*, 20:755, 2003. 2.3, 2.3
- [74] Thierry Mora and Clifford M. Will. A post-newtonian diagnostic of quasi-equilibrium binary configurations of compact objects. *Phys. Rev. D*, 69:104021, 2004. 2.3, 2.3
- [75] C.M. Will L.E. Kidder and A.G. Wiseman. Innermost stable orbits for coalescing binary systems of compact objects. *Class. Quantum Grav.*, 9:L125–L131, 1992. 2.4.1
- [76] L.E. Kidder, C.M. Will, and A.G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. *Phys. Rev. D*, 47:R4183–R4187, 1993. 2.4.1
- [77] Steve Fairhurst et al. The ligo scientific collaboration’s inspiral search pipeline. In preparation. 3, 3.1
- [78] Duncan Brown. Inspiral computing requirements for advanced ligo. 2006. Internal LSC document. 3
- [79] Barry Wilkinson and Michael Allen. *Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers*. Pearson Education Ltd., London, 2nd edition, 2005. 3.2
- [80] B. Abbott et al. Detector description and performance for the first coincidence observations between ligo and geo. *Nucl. Instrum. Meth.*, A517:154–179, 2004. 4.1
- [81] F. Acernese et al. The virgo status. *Class. Quant. Grav.*, 23:S635–S642, 2006. 4.1
- [82] H. Luck et al. Status of the geo600 detector. *Class. Quant. Grav.*, 23:S71–S78, 2006. 4.1
- [83] 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. 4.1
- [84] Sukanta Bose, Sanjeev V. Dhurandhar, and Archana Pai. Detection of gravitational waves using a network of detectors. *Pramana*, 53:1125–1136, 1999. 4.1

- [85] Lee Samuel Finn. Aperture synthesis for gravitational-wave data analysis: Deterministic sources. *Phys. Rev.*, D63:102001, 2001. 4.1
- [86] Nicolas Arnaud et al. Coincidence and coherent data analysis methods for gravitational wave bursts in a network of interferometric detectors. *Phys. Rev.*, D68:102001, 2003. 4.1
- [87] P. Jaranowski and A. Krolak. Optimal solution to the inverse problem for the gravitational wave signal of a coalescing compact binary. *Phys. Rev.*, D49:1723–1739, 1994. 4.1
- [88] P. Jaranowski and A. Krolak. *Class. Quantum Grav.*, 13:1279, 1996. 4.1
- [89] Nicolas Arnaud et al. Detection in coincidence of gravitational wave bursts with a network of interferometric detectors. i: Geometric acceptance and timing. *Phys. Rev.*, D65:042004, 2002. 4.1
- [90] Hideyuki Tagoshi et al. Detecting gravitational waves from inspiraling binaries with a network of detectors: Coherent strategies by correlated detectors. *Phys. Rev.*, D75:087306, 2007. 4.1
- [91] B. Abbott et al. Analysis of ligo data for gravitational waves from binary neutron stars. *Phys. Rev.*, D69:122001, 2004. 4.1, 4.1.1
- [92] B. Abbott et al. Search for gravitational waves from galactic and extra- galactic binary neutron stars. *Phys. Rev.*, D72:082001, 2005. 4.1, 4.1.1, 4.5
- [93] B. Abbott et al. Search for gravitational waves from primordial black hole binary coalescences in the galactic halo. *Phys. Rev.*, D72:082002, 2005. 4.1, 4.1.1
- [94] B. Abbott et al. Search for gravitational waves from binary black hole inspirals in ligo data. *Phys. Rev.*, D73:062001, 2006. 4.1, 4.1.1
- [95] B. Abbott et al. Upper limits on gravitational wave bursts in ligo’s second science run. *Phys. Rev.*, D72:062001, 2005. 4.1.1
- [96] B. Abbott et al. Search for gravitational wave bursts in ligo’s third science run. *Class. Quant. Grav.*, 23:S29–S39, 2006. 4.1.1

- [97] L.S. Finn and D.F. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. *Phys. Rev. D*, 47:2198–2219, 1993. 4.1.2, 4.2.1, 4.2.2
- [98] David F. Chernoff and Lee Samuel Finn. Gravitational radiation, inspiraling binaries, and cosmology. *Astrophys. J.*, 411:L5–L8, 1993. 4.1.2, 4.2.1, 4.2.2
- [99] K.D. Kokkotas, A. Królak, and G. Tsebas. Statistical analysis of the estimators of the parameters of the gravitational-wave signal from a coalescing binary. *Class. Quantum. Grav.*, 11:1901, 1994. 4.1.2
- [100] A. Królak, K.D. Kokkotas, and G. Schäfer. Estimation of the post-newtonian parameters in the gravitational-wave emission of a coalescing binary. *Phys. Rev. D*, 52:2089–2111, 1995. 4.1.2
- [101] E. Poisson and C.M. Will. Gravitational waves from inspiralling compact binaries - parameter-estimation using second-post-newtonian wave-forms. *Phys. Rev. D*, 52:848–855, 1995. 4.1.2
- [102] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Estimation of parameters of gravitational waves from coalescing binaries. *Pramana*, 45:L463, 1995. 4.1.2, 4.2.1, 4.2.1
- [103] 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). 4.1.2, 4.2.1, 4.2.1
- [104] K. G. Arun, B. R. Iyer, B. S. Sathyaprakash, and P. A. Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-newtonian gravitational wave phasing: The non-spinning case. *Phys. Rev. D*, 71:084008, 2005. Erratum-ibid. *D* 72, 069903 (2005). 2.1.6, 4.1.2
- [105] B. S. Sathyaprakash. Filtering post-newtonian gravitational waves from coalescing binaries. *Phys. Rev. D*, 50(12):R7111–R7115, Dec 1994. 4.2.1
- [106] B. Owen. *Phys. Rev.*, *D* 53:6749–6761, 1996. 4.2.1, 4.2.1, 4.5
- [107] B. Owen and B. S. Sathyaprakash. *Phys. Rev.*, *D* 60:022002, 1998. 4.2.1, 4.2.1

- [108] C.W. Helström. *Statistical Theory of Signal Detection*, volume 9 of *International Series of Monographs in Electronics and Instrumentation*. Pergamon Press, Oxford, U.K., New York, U.S.A., 2nd edition, 1968. 4.2.1, 4.2.1
- [109] B.F. Schutz. In D Blair, editor, *The detection of gravitational waves*, England, 1989. Cambridge University Press. 4.2.1
- [110] J.W. Perram and M.S. Wertheim. Statistical mechanics of hard ellipsoids. i. overlap algorithm and the contact function. *J. Comp. Phys.*, **58**:409–416, 1985. 4.3, 4.3
- [111] Stephen Fairhurst and Others. 4.3
- [112] Wikipedia article: Brent’s method. [http://en.wikipedia.org/wiki/Brent’s_method](http://en.wikipedia.org/wiki/Brent's_method). 4.3
- [113] R.P. Brent. *Algorithms for Minimization without Derivatives*, chapter 4. Prentice-Hall, Englewood Cliffs, NJ., 1973. 4.3
- [114] Gsl - the gnu scientific library is a free numerical library licensed under the gnu gpl. 4.3
- [115] 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. 4.5, 4.5
- [116] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius. Inspiral, merger and ring-down of equal-mass black-hole binaries. *Accepted by Phys. Rev. D*, 2006. arXiv:gr-qc/0610122v2. 5.2
- [117] Alessandra Buonanno. Gravitational waves. *To appear in the Proceedings of Les Houches Summer School, Particle Physics and Cosmology: The Fabric of Spacetime, Les Houches, France, 31 Jul - 25 Aug 2006*, 2007. 5.2
- [118] J.A. de Freitas Pacheco, T. Regimbau, S. Vincent, and A. Spallicci. Expected coalescence rates of ns-ns binaries for laser beam interferometers. *Int. J. Mod. Phys.*, **D15**:235–250, 2006. 5
- [119] LIGO Scientific Collaboration and K. Hurley. Implications for the origin of grb 070201 from ligo observations. *To be submitted to Ap.J.*, 2007. arXiv:0711.1163v2 [astro-ph]. 5

- [120] 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). 5
- [121] B. Abbott et al. Search of s3 ligo data for gravitational wave signals from spinning black hole and neutron star binary inspirals. In preparation. 5
- [122] G.W. Jones. Phd thesis, Cardiff University, 2007. 5
- [123] Y. Pan, A. Buonanno, Y. Chen, and M. Vallisneri. Physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single-spin binaries. *Phys. Rev. D*, 69(10):104017–+, 2004. Erratum-ibid. **D** 74, 029905(E) (2006). 5
- [124] A. Buonanno, Y. Chen, Y. Pan, and M. Vallisneri. Quasiphysical family of gravity-wave templates for precessing binaries of spinning compact objects: Application to double-spin precessing binaries. *Phys. Rev. D*, 70(10), November 2004. Erratum-ibid. **D** 74, 029902(E) (2006). 5
- [125] L. S. Finn. Ligo technical document. <http://www.ligo.caltech.edu/docs/T/T030256-00.pdf>, 2001. 5.1.1
- [126] Bruce Allen. A chi-squared time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005. 5.2.3, 5.2.3
- [127] B. Abbott et al. Analysis of ligo data for gravitational waves from binary neutron stars. *Phys. Rev. D*, 69:122001, 2004. 5.2.3
- [128] B. Abbott et al. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Phys. Rev. D*, 72:082001, 2005. 5.2.3
- [129] S. Babak, H. Grote, M. Hewitson, H. Lück, and K.A. Strain. Signal based vetoes for the detection of gravitational waves from inspiralling compact objects. *Phys. Rev. D*, 72:022002, 2005. 5.2.3
- [130] Andres Rodriguez. Reducing false alarms in searches for gravitational waves from coalescing binary systems. Masters thesis, Louisiana State University, 2007. 5.2.3

- [131] Stephen Fairhurst and Patrick Brady. Interpreting the results of searches for gravitational waves from coalescing binaries. 2007. arXiv:0707.2410v1 [gr-qc]. 5.5.1
- [132] Rahul Biswas, Patrick R. Brady, Jolien D.E. Creighton, and Stephen Fairhurst. The loudest event statistic: General formulation, properties and applications. 2007. arXiv:0710.0465v1 [gr-qc]. 5.5.1
- [133] Ravi Kumar Kopparapu, Chad R. Hanna, Vicky Kalogera, Richard O'Shaughnessy, Gabriela Gonzalez, Patrick R. Brady, and Stephen Fairhurst. Host galaxies catalog used in ligo searches for compact binary coalescence events. *Submitted to Astrophys. J.*, 2007. arXiv:0706.1283v1 [astro-ph]. 5.5.1

