Abstract

Gravitational waves are a consequence of the general theory of relativity. Direct detection of such waves will provide a wealth of information about physics, astronomy, and cosmology. A worldwide effort is currently underway to make the first direct detection of gravitational waves. The global network of detectors includes the Laser Interferometer Gravitational-wave Observatory (LIGO), which recently completed its sixth science run.

A particularly promising source of gravitational waves is a binary system consisting of two neutron stars and/or black holes. As the objects orbit each other they emit gravitational radiation, lose energy, and spiral inwards. This produces a characteristic "chirp" signal for which we can search in the LIGO data. Currently this is done using matched-filter techniques, which correlate the detector data against analytic models of the emitted gravitational waves. Several choices must be made in constructing a search for signals from such binary coalescences.

Any discrepancy between the signals and the models used will reduce the effectiveness of the matched filter. However, the analytic models are based on approximations which are not valid through the entire evolution of the binary. In recent years numerical relativity has had impressive success in simulating the final phases of the coalescence of binary black holes. While numerical relativity is too computationally expensive to use directly in the search, this progress has made it possible to perform realistic tests of the LIGO searches. The results of such tests can be used to improve the efficiency of searches.

Conversely, noise in the LIGO and Virgo detectors can reduce the efficiency. This must be addressed by characterizing the quality of the data from the detectors, and removing from the analysis times that will be detrimental to the search.

In this thesis we utilize recent results from numerical relativity to study both the

degree to which analytic models match realistic waveforms and the ability of LIGO searches to make detections. We also apply the matched-filter search to the problem of removing times of excess noise from the search.

CHARACTERIZATION OF ENHANCED INTERFEROMETRIC
GRAVITATIONAL-WAVE DETECTORS AND STUDIES OF NUMERIC
SIMULATIONS FOR COMPACT-BINARY COALESCENCES

By

Larne Pekowsky
B.A. New York University, 1988
M.S. Yale University, 1991
B.S. Stony Brook University, 2006

DISSERTATION

Submitted in partial fulfillment of the requirements for the degree of Doctor of
Philosophy in Physics in the Graduate School of Syracuse University

December 2011

Approved _____
Professor Duncan Brown
Date _____

Contents

List of Tables

List of Figures

Preface

The work presented in this thesis stems from my participation in the LIGO Scientific Collaboration (LSC) and the NINJA Collaboration. This work does not reflect the scientific opinion of the LSC and it was not reviewed by the collaboration.

Acknowledgments

I have been fortunate to work with many wonderful people through my membership in the LIGO and NINJA collaborations. Sadly space makes it impossible to thank them all individually, and I hope the vast majority I must omit will forgive me.

I would first and foremost like to thank my advisor, Duncan Brown. I have benefited immensely from Duncan's vast knowledge of and passion for gravitational-wave astronomy. I could not have asked for a better mentor, nor did I ever expect to find one with whom I could share so many laughs.

I would also like to thank Andrew Lundgren, with whom I worked on detchar. Andy was great fun to work with, and his brilliance is matched only by his passion for squashing glitches in the detectors. Much of this work was also done with Josh Smith, to whom I am also grateful.

I am indebted to the other members of the Syracuse University gravitational-wave faculty, Peter Saulson and Stefan Ballmer, for their many excellent questions and suggestions regarding my work, and clear and detailed explanations of their own.

The NINJA project has introduced me to a great many numerical relativists from whom I have learned a great deal. I would like especially to thank Mike Boyle, with whom I worked on the studies in chapter 5, Harald Pfeiffer, and Deirdre Shoemaker.

On the data-analysis side of NINJA, I would like to thank Satya Mohapatra for his help in testing the data sets and numerous other contributions as well as many fascinating conversations at conferences.

It has been a joy to share an office with Collin Capano, Justin Garofoli, and Matt West, and I would like to thank them for their innumerable assistance and physics bull sessions.

I would also like to thank Ping Wei, with whom I worked on the segment database, and Peter Couvares for all his help with Condor and friendly chat about computer

science and other things.

None of this would have been possible without the love and support of my parents, to whom I am eternally grateful. Finally, my thanks to my long-time friends, known collectively as "the canetoads" for their words of encouragement and all the good times.

Conventions

We adopt the Einstein summation convention where repeated indices are summed over. Parenthesis are shorthand for the fully symmetric sum

$$A_{(\alpha\beta)} = \frac{1}{2}\left(A_{\alpha\beta} + A_{\beta\alpha}\right)$$

and square brackets are fully antisymmetric

$$A_{[\alpha\beta]} = \frac{1}{2}\left(A_{\alpha\beta} - A_{\beta\alpha}\right)$$

We take the signature of spacetime to be $(-1, +1, +1, +1)$.

Except where otherwise noted we work in geometric units where $G = c = 1$. We will often measure masses in multiples of the mass of the sun $1M_\odot \approx 1.99 \times 10^{30}$ kg

We define the Fourier transform of a function of time $g(t)$ to be $\tilde{g}(f)$, where

$$\tilde{g}(f) = \int_{-\infty}^{\infty} g(t)\, e^{-2\pi i f t}\, dt$$

and the inverse Fourier transform is

$$g(t) = \int_{-\infty}^{\infty} \tilde{g}(f)\, e^{2\pi i f t}\, dt$$

This convention differs from that used in some gravitational-wave literature, but is the adopted convention in the LIGO Scientific Collaboration.

The time-stamps of interferometer data are measured in Global Positioning System (GPS) seconds: seconds since 00:00.00 UTC January 6, 1980 as measured by an atomic clock.

Chapter 1

Introduction

Gravitational waves are one of the most remarkable predictions of Einstein's general theory of relativity [1]. In a sense they are an immediate consequence of reconciling gravity with special relativity. As nothing can travel faster than light, information about the change in position of an object can not propagate instantaneously. The gravitational field therefore becomes a dynamical entity, much like the electromagnetic field in Maxwell's equations.

There is currently compelling indirect evidence for the existence of gravitational waves. These waves carry energy away from their source, and the energy in a gravitationally-bound system of two stars or black holes correlates directly to the separation between them. It is therefore possible to track the change in separation of such a binary system, infer the power being radiated away, and compare this to the predicted power radiated by gravitational waves. Hulse and Taylor were awarded the Nobel prize in 1994 for the discovery of a system exhibiting such behavior [2, 3]. Continued observation of this system over the intervening years has shown that the loss of energy matches the prediction from general relativity to within 0.2% [4].

Direct detection of gravitational waves would do far more than further confirm the prediction of their existence. Such detections would open a new window on the universe. Radio, X-ray and microwave astronomy have each provided views of the Universe not available though the others, and each has taught us a great deal. The gravitational spectrum is completely new, and so far unexplored. Gravitational wave astronomy can probe regions that the electromagnetic spectrum can not, including black holes, the interiors of neutron stars and supernovae, and the very early universe.

Gravitational wave astronomy also offers a unique opportunity to test strong-field general relativity.

Attempts to directly detect gravitational waves are hampered by the fact that they interact very weakly with matter. This is a consequence of the fact that gravity is comparatively a very weak force; a small magnet can, over a short distance, balance the gravitational pull of the entire Earth. Nevertheless, direct detection is possible. There is currently a worldwide effort underway using interferometric techniques consisting of:

- The Laser Interferometer Gravitational wave Observatory (LIGO) in the US, with detectors in Hanford, WA and Livingston, LA [5].

- Virgo, in Cascina, Italy [6].

- GEO 600, near Sarstedt, Germany [7].

- The Large-scale Cryogenic Gravitational-wave Telescope (LCGT) is presently about to being construction in Japan [8].

The LIGO, Virgo and GEO detectors have all reached their initial design sensitivities enabling detection of gravitational waves from the final stages of systems like the Hulse-Taylor binary out to (in the case of LIGO) as far as 45 Megaparsecs (Mpc). As a point of reference, the Virgo supercluster extends about 33 Mpc and contains thousands of galaxies, including our own.

Although there are many potential sources of gravitational waves, a particularly promising one is the inspiral and coalescence of binary systems containing compact objects, neutron stars or black holes [9] (collectively called *compact binary coalescence*, henceforth CBC). Over a large range of masses the frequencies at which such systems emit the most power are the frequencies at which LIGO and Virgo are the most sensitive, from about 40 Hz to 1000 Hz. In chapter 2 we will show that the effect of gravitational waves is to change the distance between freely-falling objects. In this frequency band LIGO and Virgo are sensitive to changes in length to about 1 part in 10^{22}.

Despite the detectors' remarkable sensitivity, signals in the LIGO data will likely be quiet relative to the noise levels. We will have more to say about this noise in

Sec. 3.2.2 and chapter 10, but note here that it contains both background "static" which may be modeled as a *Gaussian random process* and isolated *non-Gaussian transients*. Sophisticated data analysis techniques will be required to extract the signals from noise and infer properties of their sources.

In addition to the achievements at LIGO and other sites, recent years have also seen major progress in the field of *numerical relativity*, the use of computers to accurately simulate systems governed by general relativity. Following the breakthroughs in in 2005 by Pretorius [10], researchers at Goddard [11] and researchers at the University of Texas at Brownville and Florida Atlantic University [11] there are now routine simulations of binary black hole systems with a wide variety of parameters, extending through several orbits. There are also an increasing number of simulations of systems containing matter, including both black holes in gaseous environments and neutron stars and supernovae, although we will not consider simulations containing matter further in this dissertation.

The goal of the research presented in this dissertation is to improve the efficiency of CBC searches. In order to claim a detection a candidate signal must stand out sufficiently above the noise. Efficiency can therefore be improved either by increasing the significance of the signal or by reducing the noise. We take both approaches here.

The remainder of this dissertation is organized as follows:

- In chapter 2 we present the theory of gravitational waves, starting from general relativity and deriving the effect of waves on matter. We also discuss both analytic and computational approaches to predicting the form of gravitational waves.

- In chapter 3 we start from the effect of gravitational waves on matter to discuss the principles of interferometric detectors.

- In chapter 4 we describe the search for CBC signals in data from interferometric detectors.

- Chapter 5 begins the process of utilizing predictions from numerical relativity to improve the efficiency of searches. We proceed by comparing the analytic models used in searches to a high-accuracy simulation from the Caltech-Cornell group.

Based on these comparisons we derive a number of possible improvements to the search.

- Chapter 6 introduces the first Numerical INJection Analysis (NINJA) project, a collaboration between numerical relativists and gravitational-wave astronomers. The goal of NINJA is to study the effectiveness of numerous search methods at detecting gravitational waves and extracting information about their sources. We apply the recommendations from the previous chapter to the NINJA data, which simulates the output from LIGO and contains signals from numerical relativity.

- Chapter 7 discusses some of the limitations of the first NINJA project and introduces NINJA-2. A notable feature of NINJA-2 is the requirements placed on the NR submissions. We discuss the comprehensive analyses that have been performed to ensure that the submissions meet these requirements.

- Chapter 8 discusses the construction of the simulated data sets for NINJA-2 and presents preliminary results from the CBC analysis of the data. We note some unexpected results that need further study and raise questions that we plan to address in NINJA-2.

- We then transition to the topic of detector noise. In chapter 9 we describe the infrastructure used to record information about the state of the detectors.

- Then in chapter 10 we present a reduced, simplified version of the CBC search. This was used in the latest runs of the LIGO and Virgo detectors to identify times of excess noise and remove these times from analysis. We include some open questions about the implementation details of the search revealed during these analyses.

Throughout this dissertation we focus on LIGO's 6th science run (denoted "S6", July 7, 2009 - October 20, 2010) which overlapped Virgo's second ("VSR2", July 7, 2009 - January 11, 2010), and third ("VSR3", August 11, 2010 - October 20, 2010) science runs. This corresponds to "Enhanced LIGO", so called because this run included a number of modifications intended for Advanced LIGO (to start in

2015) but which were available and could be added to the Initial LIGO configuration without major construction.

Chapter 2

Gravitational wave Theory and Source Modeling

In this chapter we briefly survey the theoretical issues behind gravitational-wave astronomy. We start in Sec. 2.1 with a review of general relativity, beginning with the relevant mathematics. In Sec. 2.2 we show how general relativity predicts the existence of gravitational radiation and discuss some of the properties of this radiation. Then in Sec. 2.3 we show how gravitational radiation affects freely-falling particles. This will motivate the design of the LIGO experiment to search for gravitational waves, an overview of which is presented in the next chapter. We then move to the generation of gravitational waves and discuss two approaches to modeling the waves produced by the inspiral and eventual merger of pairs of compact objects. These are analytic models, discussed in Sec. 2.5 and numerical models, discussed in Sec. 2.6.

2.1 General Relativity

We start with an overview of differential geometry and build to Einstein's equations. This is of necessity very brief, readers are referred to the textbooks by Misner, Thorne and Wheeler [12] and Carroll [13] for more complete treatments.

2.1.1 Elements of Differential Geometry

An n-dimensional (C^∞) manifold $\mathcal{M}$ is a set of points plus an *atlas*, a set of *charts* $\{\phi_i\}$ which are invertible maps from open subsets of $\mathcal{M}$ to open subsets of $\mathcal{R}^n$ such

that

- For all points $p \in \mathcal{M}$ there exists a ϕ_i such that p is in the domain of ϕ_i.

- The composition $\phi_i \circ \phi_j^{-1}$ on the intersections of the domains of ϕ_i and ϕ_j is a (C^∞) function from $\mathcal{R}^n \to \mathcal{R}^n$.

Two natural structures on a manifold are curves, maps from $\mathcal{R} \to \mathcal{M}$, and scalar functions, maps from $\mathcal{M} \to \mathcal{R}$. Compositing a function f with a curve $\gamma(\lambda)$ gives a map from $\mathcal{R} \to \mathcal{R}$ which may be differentiated in the usual way at a point p.

Geometrically, taking the derivative gives the tangent vector to the curve at a point p. It is possible to associate the set of such vectors with the set of directional derivatives, taking the partial derivatives along the coordinates as the basis. Henceforth this basis will be denoted both ∂_μ and $\vec{e}_\mu$. Note that the tangent to a curve is defined at the point p. Each point in the manifold possesses its own space of tangent vectors. These spaces are distinct, which will be important in what follows.

We next define *one-forms* as linear maps from vectors to $\mathcal{R}$. The set of one-forms at a point can be shown to form a vector space, a natural basis for which can be obtained by requiring

$$\vec{e}_\mu \tilde{\omega}^\nu = \delta_\mu^\nu$$

The components of an arbitrary form ω in this basis may be found by applying the form to the basis vectors.

We can then build up arbitrary $\binom{m}{n}$ tensors as linear maps from tensor products of m vectors and n one-forms to $\mathcal{R}$. The components of a tensor T in a choice of coordinates may found by applying it to combinations of the basis vectors and basis 1-forms. Finally, a $\binom{m}{n}$ tensor field is a map that associates to each point p in $\mathcal{M}$ an element in the space of $\binom{m}{n}$ tensors at p.

2.1.2 The Metric Tensor

A particularly important tensor in general relativity is the *metric*, a $\binom{2}{0}$ tensor that is symmetric ($g_{\mu\nu} = g_{\nu\mu}$) and non-degenerate (the determinant of g taken as a matrix $|g_{\mu\nu}| \neq 0$. The latter feature makes it possible to define the inverse metric $g^{\mu\nu}$ as

$$g^{\mu\rho} g_{\rho\nu} = \delta_\nu^\mu$$

Given a vector x^μ the object $g_{\mu\nu}x^\mu$ maps another vector to a real number, and is therefore a one-form. The metric therefore maps between the space of one-forms and the space of vectors at each point. Most importantly, the metric defines a notion of distance on the manifold. Infinitesimally

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu \tag{2.1}$$

In special relativity, in Cartesian coordinates, the metric has components $(-1, 1, 1, 1)$ along the diagonal, all other components are zero. The metric will be denoted $\eta_{\mu\nu}$ and called the *flat space metric.*

2.1.3 Covariant Derivatives

Since vectors are only defined at a point we need additional structure to define derivatives of vector fields, as there is no natural way to compare vectors that live in different spaces. We seek an operator ∇ with the following properties

- Maps $\binom{m}{n}$ tensors to $\binom{m}{n+1}$ tensors. This is so we may consider the directional derivative of a tensor T along a vector x as $x^\mu \nabla_\mu T$.

- Reduces to partial differentiation when applies to a scalar field.

- Linear.

- Satisfies the Leibniz rule, $\nabla(ab) = a\nabla b + b\nabla a$.

Such an operator applied to a vector field gives

$$\nabla_\mu(x^\nu \vec{e}_\nu) = (\partial_\mu x^\nu)\vec{e}_\nu + x^\nu(\nabla_\mu \vec{e}_\nu)$$

In flat space in Cartesian coordinates the basis vectors do not change and so the last term is zero. But in a curved space, or even flat space in non-Cartesian coordinates, they may. However, the new vector must be expressible as a linear combination of the original basis vectors. The components are called *connection coefficients* and are denoted as $\Gamma^\rho_{\nu\mu}$ so

$$\nabla_\mu(x^\nu \vec{e}_\nu) = (\partial_\mu x^\nu)\vec{e}_\nu + x^\nu \Gamma^\rho_{\mu\nu}\vec{e}_\rho \tag{2.2}$$

$$= (\partial_\mu x^\nu + x^\rho \Gamma^\nu_{\mu\rho})\vec{e}_\nu \tag{2.3}$$

$$\nabla_\mu x^\nu = \partial_\mu x^\nu + x^\rho \Gamma^\nu_{\mu\rho} \tag{2.4}$$

In general relativity the connection is usually assumed to be *torsion-free*, that is

$$\Gamma^{\nu}_{\mu\rho} = \Gamma^{\nu}_{\rho\mu} \tag{2.5}$$

which thus far has been borne out by experiment. However, it is possible to construct theories where this condition does not hold.

We will henceforth occasionally use commas and semicolons to denote partial and covariant differentiation, respectively:

$$x^{\mu}{}_{,\nu} \equiv \partial_{\nu} x^{\mu}$$
$$x^{\mu}{}_{;\nu} \equiv \nabla_{\nu} x^{\mu}$$

By considering the covariant derivative of a scalar constructed from a one-form acting on a vector, $\nabla(x^{\nu}\omega_{\nu})$, it can be shown that

$$\nabla_{\mu}\omega_{\nu} = \partial_{\mu}\omega_{\nu} - \Gamma^{\rho}_{\mu\nu}\omega_{\rho}$$

The covariant derivative of a $\binom{m}{n}$ tensor generalizes this and has a partial derivative term, m positive therms in Γ and n negative terms in Γ.

2.1.4 Parallel Transport

Covariant differentiation provides a way to "move a vector without changing it." We can *parallel transport* a vector v^{μ} infinitesimally along a curve whose tangent vector is u^{ν} by requiring

$$u^{\nu}\nabla_{\nu}v^{\mu} = 0$$

As an example of such transport, consider an arrow on the equator of the Earth pointing towards the north pole. This arrow can be carried halfway around the equator without rotating it, so it ends up on the other side of the globe, still pointing north. If the vector is then parallel transported northward to the pole and then continued until it returns to its starting point it will return pointing south. Although the vector was never rotated locally it has returned rotated. This is an indication that the underlying space is curved.

Of particular interest is the case where a vector is parallel-transported along itself

$$0 = v^{\mu}\nabla_{\mu}v^{\nu} = v^{\mu}(\partial_{\mu}v^{\nu} + \Gamma^{\nu}_{\mu\rho}v^{\rho})$$

Now consider a curve $x(\lambda)$ such that v is the tangent to this curve, $v^\mu = dx^\mu/d\lambda$, then

$$\frac{dx^\mu}{d\lambda}\frac{\partial}{\partial x^\mu}\frac{dx^\nu}{d\lambda} + \Gamma^\nu_{\mu\rho}\frac{dx^\mu}{d\lambda}\frac{dx^\rho}{d\lambda} = 0$$
$$\frac{d^2x^\nu}{d\lambda^2} + \Gamma^\nu_{\mu\rho}\frac{dx^\mu}{d\lambda}\frac{dx^\rho}{d\lambda} = 0$$

(2.6)

This is the *geodesic equation*, solutions to which are *geodesics*. The same equation can be derived by extremizing the path length, $\sqrt{g_{\mu\nu}dx^\mu dx^\nu}$. In general relativity test masses acting under the influence of gravity and no other forces follow geodesics.

2.1.5 The Christoffel Symbols

If we now require that scalars do not change under parallel transport we have, for arbitrary vectors fields u^α, v^β and x^μ

$$0 = x^\mu\nabla_\mu(g_{\alpha\beta}u^\alpha v^\beta)$$
$$= x^\mu(\nabla_\mu g_{\alpha\beta})u^\alpha v^\beta + g^{\alpha\beta}(x^\mu\nabla_\mu u^\alpha)v^\beta + g^{\alpha\beta}u^\alpha(x^\mu\nabla_\mu v^\beta)$$

We can now specialize such that u^α, v^β are constant and so the last two terms vanish, which implies the *metric compatibility* condition:

$$\nabla_\mu g_{\alpha\beta} = 0 \tag{2.7}$$

Equations 2.7 and 2.5 together fix the connection coefficients in terms of the metric:

$$\Gamma^\rho_{\mu\nu} = \frac{1}{2}g^{\rho\sigma}\left[\partial_\nu g_{\mu\sigma} + \partial_\mu g_{\nu\sigma} - \partial_\sigma g_{\mu\nu}\right] \tag{2.8}$$

Combining this with the previous section we see that the motion of particles are completely specified once we know the metric and their initial positions and velocities.

2.1.6 The Riemann Tensor

We now generalize the example given in Sec. 2.1.4, and ask how a vector A^μ changes as it is parallel-transported around an infinitesimal parallelogram with sides defined by the vectors B^μ and C^ν. Recalling that vectors and directional derivatives are the

same thing, it can be shown that this is equivalent to asking how covariant derivatives fail to commute. The result must be linear in A^μ and so we may write

$$[\nabla_\mu \nabla_\nu - \nabla_\nu \nabla_\mu] A^\rho = R^\rho_{\sigma\mu\nu} A^\sigma \qquad (2.9)$$

which defines the *Riemann tensor* $R^\rho_{\sigma\mu\nu}$. A number of properties follow from this definition (which are either obvious or may be proven by substituting the definition of the covariant derivative, eqn. 2.2). First, the symmetry properties

$$R_{\rho\sigma\mu\nu} = -R_{\sigma\rho\mu\nu} = -R_{\rho\sigma\nu\mu} = R_{\mu\nu\rho\sigma} \qquad (2.10)$$

which in turn imply

$$R^\rho_{[\sigma\mu\nu]} = 0 \qquad (2.11)$$

Second, the Bianchi identity,

$$R_{\rho\sigma\mu\nu;\alpha} + R_{\rho\sigma\nu\alpha;\mu} + R_{\rho\sigma\alpha\mu;\nu} = 0 \qquad (2.12)$$

We may now generalize Eqn. 2.9 and ask how an arbitrary tensor changes after being parallel-transported around a loop. It can be shown that

$$[\nabla_\mu \nabla_\nu - \nabla_\nu \nabla_\mu] B^{\rho_1\rho_2\dots\rho_n} = -R^{\rho_1}_{\sigma\mu\nu} B^{\sigma\rho_2\dots\rho_n} - R^{\rho_2}_{\sigma\mu\nu} B^{\rho_1\sigma\dots\rho_n} - \dots - -R^{\rho_n}_{\sigma\mu\nu} B^{\rho_1\rho_2\dots\sigma}$$
$$(2.13)$$

which may be proved by expanding

$$[\nabla_\mu \nabla_\nu - \nabla_\nu \nabla_\mu] (\vec{e}_\rho \otimes \vec{e}_\sigma)$$

and then proceeding by induction.

The symmetries of the Riemann tensor imply that there is, up to sign, only one non-trivial contraction

$$R_{\mu\nu} = R^\sigma_{\mu\sigma\nu} \qquad (2.14)$$

which defines the *Ricci tensor*. This may be contracted again

$$R = R^\mu_\mu \qquad (2.15)$$

to obtain the *Ricci scalar*.

Now, contracting the Bianchi identity twice gives

$$g^{\sigma\nu}\left(\nabla_\alpha R_{\sigma\nu} + \nabla^\rho R_{\rho\sigma\nu\alpha} + \nabla_\nu R^\mu_{\sigma\alpha\mu}\right) = 0$$

Using the symmetries of the Riemann tensor (eqn. 2.10) this can be written

$$\nabla_\alpha R - \nabla^\rho R_{\rho\alpha} - \nabla^\sigma R_{\sigma\alpha} = 0$$

Relabeling the dummy indices and using metric compatibility gives

$$\nabla^\rho\left(g_{\rho\alpha}R - 2R_{\rho\alpha}\right) = 0$$

This motivates the definition of the *Einstein Tensor* as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R \tag{2.16}$$

The previous result implies this is divergentless

$$\nabla^\nu G_{\mu\nu} = 0$$

Note also that G is symmetric, $G_{\mu\nu} = G_{\nu\mu}$.

We now relate this to physics by noting that the matter and energy content of a region is described by the stress-energy tensor $T_{\mu\nu}$ where each component is "the flow of μ momentum in the ν direction." For example, the $0,0$ component is energy density and the $0,i$ components are the i^{th} components of momentum.

Conservation of energy requires that the difference in momentum (p^i) across each face of a cube be balanced by a change of energy (ρ), within the cube,

$$\partial_t\rho = \partial_i p^i$$

In terms of the stress-energy tensor this becomes

$$0 = -\nabla^0 T_{00} + \nabla^i T_{0i} = \nabla^\nu T_{0\nu}$$

However the time direction is not uniquely specified as a change of coordinates will mix space and time components, so this must generalize to

$$\nabla^\nu T_{\mu\nu} = 0$$

That is, T is also divergentless, like G, and like G it is also symmetric. It is therefore reasonable to suggest the ansatz

$$G_{\mu\nu} \propto T_{\mu\nu}$$

Requiring agreement with Newton's law of gravity in the appropriate low-energy limit ($T_{00} \gg$ all other components) fixes the constant of proportionality and gives us *Einstein's field equation*

$$G_{\mu\nu} = 8\pi T_{\mu\nu} \tag{2.17}$$

Note that $G_{\mu\nu}$ is entirely determined by the metric. Equation 2.17 may therefore be thought of as a set of coupled, non-linear differential equations for $g_{\mu\nu}$.

2.2 Gravitational Radiation

We now move to the prediction of gravitational waves. We begin with Einstein's equation in empty space,

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 0$$

By taking the trace and substituting into Eqn. 2.17 it can be shown that this implies that in empty space $R_{\mu\nu} = 0$. Similarly, using the Bianchi identity and symmetries of the Riemann tensor gives, in empty space,

$$R_{\beta\delta;\gamma} - R_{\beta\gamma;\delta} = 0 \tag{2.18}$$

We next consider the application of the wave operator to the Riemann tensor. From the Bianchi identity (eqn. 2.12) this becomes

$$g^{\mu\nu}R_{\alpha\beta\gamma\delta;\mu\nu} = -g^{\mu\nu}\left[R_{\alpha\beta\delta\mu;\gamma\nu} + R_{\alpha\beta\mu\gamma;\delta\nu}\right]$$

Consider the first term on the right-hand side:

$$g^{\mu\nu}R_{\alpha\beta\delta\mu;\gamma\nu} = g^{\mu\nu}R_{\alpha\beta\delta\mu;\nu\gamma} + g^{\mu\nu}R_{\alpha\beta\delta\mu;\gamma\nu} - g^{\mu\nu}R_{\alpha\beta\delta\mu;\nu\gamma}$$
$$= g^{\mu\nu}R_{\alpha\beta\delta\mu;\nu\gamma} + g^{\mu\nu}\left[\nabla_\nu, \nabla_\gamma\right]R_{\alpha\beta\delta\mu}$$

The first term vanishes by Eqn. 2.18. The second term involves products of the Riemman tensor by 2.13. The second term on the right in Eqn. 2.2 has the same form.

We now specialize to the case where the Riemann tensor is small, so that terms involving multiple factors can be neglected. This is equivalent to considering the Riemann tensor as a field on a flat background. This gives

$$g^{\mu\nu} R_{\alpha\beta\gamma\delta;\mu\nu} = \Box R_{\alpha\beta\gamma\delta} = 0 \tag{2.19}$$

That is, each component of the Riemann tensor independently satisfies the vacuum wave equation. We can immediately write the solution:

$$R^{\alpha}_{\beta\gamma\delta} = \operatorname{Re} A^{\alpha}_{\beta\gamma\delta} \exp(ik_\mu x^\mu) \tag{2.20}$$

where $A^{\alpha}_{\beta\gamma\delta}$ is a set of amplitudes and k^μ is the wave vector. In a chosen coordinate system it has components (ω, k_x, k_y, k_z) where ω is the angular frequency and the spacial k components are wavelengths in each direction. It can be shown that

$$\nabla_{\vec{k}}\vec{k} = 0$$
$$k_\mu k^\mu = 0$$

which together imply that gravitational waves travel along geodesics at the speed of light.

2.3 Effect of Gravitational Waves

We now derive the effect of gravitational waves on matter. Consider two particles moving along world-lines A^μ and B^μ. Choose coordinates so that A remains fixed at the origin, $A^\mu = (1, 0, 0, 0)$. We may further specialize our coordinates such that at the origin $g_{\mu\nu} = \eta_{\mu\nu}$. It can be shown that we may also require that the first derivatives of the metric vanish at this point. We may not, however, make the second derivatives vanish in general. This corresponds to the fact that the Riemann tensor is defined in terms of second derivatives. We call the coordinate system thus constructed a *Local Lorentz Frame*.

We now define the separation between the two particles as

$$\xi^\mu = B^\mu - A^\mu$$

We fix ξ to be perpendicular to A, so that $\xi^0 = 0$. If space is curved it can readily be seen that ξ will not remain constant. For example, if the particles are initially at rest some distance from the surface of the Earth they will both move towards the center of the Earth and ξ will decrease. It can be shown that ξ obeys the equation of *geodesic deviation*,

$$\frac{d^2}{dt^2}\xi^\rho = -R^\rho_{\mu\nu\sigma}A^\mu\xi^\nu A^\sigma = -R^\rho_{0\nu0}\xi^\nu \tag{2.21}$$

Using the condition that ξ has no time component reduces this to

$$\frac{d^2}{dt^2}\xi^i = -R^\rho_{\mu\nu\sigma}A^\mu\xi^\nu A^\sigma = -R^i_{0j0}\xi^j \tag{2.22}$$

That is, the change in separation between two infinitesimally-separated test masses at rest with respect to each other in an arbitrary gravitational field is entirely specified by R^i_{0j0}. From the symmetries of the Reimann tensor this is symmetric in i and j, and hence appears to have 6 independent components. However, it can be shown that these can entirely be specified by two values, which without loss of generality we take to be R^x_{0x0}. R^x_{0y0}. We now recall that in empty space $R_{\mu\nu} = 0$, which implies that $R^y_{0y0} = -R^x_{0x0}$. We summarize this by saying that R is *traceless*.

We now specialize to the case of gravitational waves, so that the Riemann tensor satisfies Eqn. 2.19 and we choose coordinates such that the wave is traveling in the z direction. Using the fact that the speed of light is 1 in dimensionless units the solution can then be written

$$R_{i0j0} = A_{ij}(t-z)$$

where we have lowered the first index to simplify notation.

Now, using the fact that $\partial_x R_{i0j0} = \partial_y R_{i0j0} = 0$ and integrating the Bianchi identity we can show that $R_{x0y0} = R_{y0x0}$ and that all other components vanish. It can also be shown that in addition to being traceless R is *transverse*, $k^j R_{i0j0} = 0$. We denote these two facts by adding the superscript TT, and define the gravitational-wave field as

$$-\frac{1}{2}\frac{\partial^2 h^{TT}_{ij}}{\partial t^2} \equiv R^{TT}_{i0j0} \tag{2.23}$$

We now decompose the separation vector ξ into the initial separation and a time-dependant perturbation, $\xi = \xi_0 + \delta\xi$. In terms of this Eqn. 2.21 becomes

$$\frac{d^2}{dt^2}\delta\xi^i = -R^{0i0j}\xi^j_0 \tag{2.24}$$

where we drop the initial portion from the left-hand side because it is constant, and we drop the perturbation from the right hand side because it is much smaller than the initial portion. Comparing eqn. 2.23 and eqn. 2.24 we obtain the equation for the effect of a gravitational wave on free-falling test masses:

$$\delta\xi^i = \frac{1}{2}h_{ij}^{TT}\xi^j \tag{2.25}$$

We note in passing that this is the same result obtained in other treatments by expanding the metric in terms of the flat-space metric plus a perturbation, $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, substituting into the Einstein equation and expanding to first order in h, and then choosing a gauge in which h is transverse and traceless.

Now, define

$$h_+ \equiv h_{xx} = -h_{yy}$$
$$h_\times \equiv h_{xy} = h_{yx}$$

which we refer to as the *plus* (+) and *cross* (×) polarizations, respectively. Consider the case where $h_\times = 0$. If particle B is initially on the x axis then

$$\delta\xi^x = \frac{1}{2}h_{xx}^{TT}\xi^x$$
$$\delta\xi^y = \frac{1}{2}h_{xy}^{TT}\xi^y$$
$$= 0$$

The particle remains on the x axis. For an oscillatory wave the distance between the two particles likewise oscillates. We can describe this as an induced *strain*, $\Delta L/L$ where L is the initial separation. If B is initially on the y axis

$$\delta\xi^x = \frac{1}{2}h_{xx}^{TT}\xi^x$$
$$= 0$$
$$\delta\xi^y = \frac{1}{2}h_{yy}^{TT}\xi^y$$
$$= -\frac{1}{2}h_{xx}^{TT}\xi^y$$

The particle remains on the y axis and oscillates out of phase with a corresponding particle on the x axis. The net effect is that, after a quarter cycle, a set of masses

initially in a circle are moved into an ellipse flattened along one axis and stretched along the other such that the area remains constant. After another quarter cycle they return to a circle, in the next quarter cycle they are in an ellipse with the axes flipped, and so on. It is similarly straightforward to show that for a wave cross-polarized wave the eigendirections are on the lines $x = \pm y$. The effects are the same as for the $+$ polarization, rotated 45 degrees.

Now consider a thought experiment, originally due to Feynman, where we place a bead on a stick in the path of a gravitational wave. The wave will cause the bead to slide back and forth, heating the stick trough friction and imparting energy to the system. This implies that gravitational waves carry energy. This argument can be made precise [14].

2.4   Modeling Gravitational Waves

Having demonstrated the predicted existence of gravitational waves and their effects on matter we now turn to the question of their generation. From Eqn. 2.23 we can see that, since the components of the Riemann tensor obey the wave equation, the components of h^{TT} do as well. We now consider the solution of the wave equation when the source, the stress-energy tensor, is not zero. By analogy with electromagnetism we can immediately write down the solution in terms of retarded fields

$$h_{ij}^{TT} = 4 \left[\int \frac{T_{ij}(x', t - r)}{r} \, d^3x' \right]^{TT} \tag{2.26}$$

where we integrate over the source distribution x'. When the energy and momentum densities are small, so that the curvature of spacetime is likewise small, the coordinates may be taken to have their conventional flat-space meaning. We may also replace the covariant derivative by regular partial differentiation.

Starting from the conservation of energy and momentum written in terms of the stress-energy tensor

$$T^{00}{}_{,0} = -T^{0j}{}_{,j}$$
$$T^{i0}{}_{,0} = -T^{ij}{}_{,j}$$

Differentiating with respect to time and going through some algebra yields expressions for the first and second moments of the stress energy tensor, $T^{lm}{}_{,ml}x^j x^k$ and $T^{jm}{}_{,m}x^k$. These, plus Stoke's theorem, can be used to simplify Eqn. 2.26 to give

$$h_{ij}^{TT} = \frac{2}{r}\left[\int T^{00}{}_{,00}\, x^i x^j d^3 x'\right]^{TT}$$

We recognize T^{00} as the mass/energy density, and therefore this can be written as the *quadrapole formula*

$$h_{jk}^{TT} = \frac{2}{r}\ddot{\mathcal{I}}(t - r) \tag{2.27}$$

where $\mathcal{I}$ is the quadrupole moment of the source

$$\mathcal{I}_{ij} = \int \rho(\mathbf{x}) x_i x_j\, d^3 x$$

Any system of mass with a time-dependant quadrupole moment will give rise to gravitational radiation. However, restoring physical units to Eqn. 2.27 scales the right-hand side by G/c^4. We therefore need very large masses and/or rapid changes in order to generate waves we have any chance of detecting. There are many such sources of astrophysical and cosmological interest:

- supernovae

- rotating neutron stars with axial asymmetry

- processes in the early universe, which would have produced *relic* gravitational waves in principle still detectable today

- topological defects

- compact bodies, such as neutron stars and black holes, in orbit around each other.

We will focus on the last of these for the remainder of the thesis.

It is straightforward to start from the mass distribution of two point masses in orbit around their center:

$$\rho(\mathbf{x}) = m_1\delta(x - r_1\cos(\Omega t))\delta(y - r_1\sin(\Omega t))\delta(z)$$
$$+ m_2\delta(x + r_2\cos(\Omega t))\delta(y + r_2\sin(\Omega t))\delta(z)$$

and calculate h^{TT}. By choosing coordinates centered on a terrestrial gravitational-wave detector and a basis we can write the gravitational-wave strains as [15]

$$h_+(t) = -\frac{2G}{c^4 r}\mu(\pi GMf)^{\frac{2}{3}}(1 + \cos^2(\iota))\cos(2\pi ft - 2\phi_0) \tag{2.28}$$

$$h_\times(t) = -\frac{4G}{c^4 r}\mu(\pi GMf)^{\frac{2}{3}}\cos(\iota)\sin(2\pi ft - 2\phi_0)$$

where $M = m_1 + m2$, $\mu = m_1 m_2/M$, f is the *gravitational-wave frequency* which is twice the orbital frequency $f = 2\Omega/2\pi$, ϕ_0 is the orbital phase at a specified time, and ι is the *inclination* of the binary with respect to the detector, the angle between the normal to the plane of the binary and the line joining the detector to the binary.

We noted above that gravitational waves carry energy. This energy must be balanced by a loss of energy by the system. This energy can not be localized to any one point of the wave, since any point can be placed in a Local Lorentz Frame where there is no wave. However, by averaging over a cycle it can be shown [13] that

$$t_{\mu\nu} = \frac{1}{32\pi}\left\langle h^{TT}_{\rho\sigma,\mu}(h^{TT})^{\rho\sigma}_{,\nu}\right\rangle \tag{2.29}$$

where $t_{\mu\nu}$ is the *stress-energy pseudo-tensor*. The stress energy tensor itself is zero in a region of spacetime containing only gravitational radiation. However, t may be used to describe the energy content of such radiation, and we find the flux of energy in the radial direction out of a sphere enclosing a gravitating system is

$$\frac{dE}{dt} = T_{0r} = \frac{1}{8\pi r^2}\left\langle \dddot{\mathcal{I}}^{ij}\dddot{\mathcal{I}}_{ij}\right\rangle \tag{2.30}$$

When the separation between the masses, a is large and the masses are moving slowly the gravitational energy of the system is approximately given by the Newtonian expression

$$E = -\frac{1}{2}\frac{G\mu M}{a}$$

Therefore, orbiting bodies will emit gravitational radiation and *inspiral* until they inevitably collide.

2.5   Post-Newtonian Approximations

Substituting Eqn. 2.28 into Eqn. 2.30 shows that the power emitted goes as a^{-5}. Most of the power from an inspiral, and therefore our best chance of detecting such systems,

occurs when the masses have become close and are about to merge. The approxima-
tions we made in the previous section are no longer valid in this regime. In order
to obtain analytic models of gravitational waves from inspirals to a precision that
will be useful in searches we must therefore consider higher-order corrections beyond
Newtonian mechanics of the binary. This leads to *post-Newtonian* (pN) waveforms.

The key to doing this is the requirement that the energy flux carried away by
gravitational radiation must be balanced by loss of energy of the system

$$\frac{dE}{dt} = -\mathcal{F} \tag{2.31}$$

It can be shown [12] that the energy of a body of mass μ moving on a geodesic in the
Schwarzschild metric with mass M is

$$E = \mu \frac{1 - 2M}{\sqrt{1 - 3M/r}} \tag{2.32}$$

We now introduce the parameter $v = (\pi M f)$, which is the velocity in Newtonian
mechanics. We now consider the *adiabatic approximation*, where we treat the system
as moving through a series of circular orbits. On a circular geodesic $v = \sqrt{M/r}$ and
the energy becomes

$$E = \mu \frac{1 - 2v^2}{\sqrt{1 - 3v^2}} \tag{2.33}$$

We have written down the flux to first pN order above in Eqn. 2.30, in terms of
the parameter v it is

$$\mathcal{F} = \frac{32}{5} \left(\frac{\mu}{M} \right) v^{10} \tag{2.34}$$

Going to higher order requires techniques that are beyond the scope of this thesis.

Both the energy and flux may now be expanded in powers of v. We could then
obtain t as a function of v by rewriting Eqn. 2.31 and integrating. However, we will
instead obtain an expression relating time and the orbital phase $\phi = \int 2\pi f \, dt$ by
writing

$$\mathcal{F} = -\frac{dE}{dv} \frac{dv}{d\phi} \frac{d\phi}{dt}$$

which leads to the expression

$$\phi = \phi_0 + \frac{2}{M} \int_v^{v_{\text{ref}}} \frac{v^3 dE/dv}{\mathcal{F}} \, dv \tag{2.35}$$

where v_{ref} is the velocity at a given reference velocity, or equivalently reference fre-
quency. If $\mathcal{F}$ is given as a Taylor series in v then $\mathcal{F}^{-1}$ may be expanded to the same

order. dE/dv may be similarly expanded, and the integrand becomes the product of polynomials with rational coefficients.

Motivated by Eqn. 2.28 we model the waveform as a time-dependant amplitude and evolving phase:

$$h(t) = A(t)\cos(\phi(t)) = \frac{1}{2}A(t)\left(e^{i\phi(t)} + e^{-i\phi(t)}\right) \tag{2.36}$$

For reasons that will become clear in the next chapter it is often more convenient to work in the frequency domain, so we take the Fourier transform of Eqn. 2.36

$$\tilde{h}(f) = \frac{1}{2}\int A(t)\exp(2\pi ift + i\phi(t))\,dt \tag{2.37}$$
$$+ \frac{1}{2}\int A(t)\exp(2\pi ift - i\phi(t))\,dt$$

We next apply the *stationary-phase approximation*. Oscillatory integrands cancel except where the phase is stationary, at extrema of the exponent. We therefore expand the exponent around the time t_0 of the extremum and discard the linear term

$$2\pi ift + i\phi(t) \approx (2\pi ift_0 + i\phi(t_0)) + \frac{1}{2}i\ddot{\phi}(t_0)(t - t_0)^2 \tag{2.38}$$

Substituting back, the first term leads to a constant factor and the second leads to a Gaussian integral. Going through the calculation yields the approximate waveform

$$\tilde{f}(f) = \frac{2GM_\odot}{c^2 r}\left(\frac{5\mu}{96M_\odot}\right)^{\frac{1}{2}}\left(\frac{M}{\pi^2 M_\odot}\right)^{\frac{1}{3}}\left(\frac{GM_\odot}{c^3}\right)^{-\frac{1}{6}}f^{-\frac{7}{6}}\Theta(f - f_c)e^{i\Psi(f;M,mu)} \tag{2.39}$$

where $M_\odot$ is the mass of the sun and is introduced to so that M is measured in solar masses, a useful unit in astrophysical work. Ψ is the result of the expansion 2.35 written as a function of f. The coefficients coming from the expansion of the flux will depend on the mass M and ration μ or equivalently mass and *symmetric mass ratio* $\eta = \mu/M$. In general waveforms may also depend on other parameters such as the spins of the objects. We introduce the step function to terminate the waveform at a *cutoff frequency* f_c. This reflects the fact that we do not trust this approximation all the way through the evolution of the binary to merger. The appropriate value for the cutoff frequency is one of the questions we will address in chapter 5.

There are many different ways to perform post-Newtonian expansions, in both the time and frequency domains. These lead to different waveforms which have been

assigned standard names, a summary of the variations we will encounter throughout this thesis is given in Appendix A. We note here in particular that the stationary-phase approach discussed above leads to waveforms denoted *TaylorF2*. The degree to which these different waveforms agree has been studied in Ref. [16].

2.5.1 Effective One-Body

A relatively recent development in post-Newtonian theory is the *effective one-body* (EOB) formulation [17]. The basic idea is familiar from non-relativistic classical mechanics, where it is common to write the Hamiltonian for the Kepler problem in terms of a particle with reduced mass $\mu = m_1 m_2/(m_1 + m_2)$ moving in a central potential due to a fixed body of mass $M = m_1 + m_2$.

In general relativity we have noted that the energy, or equivalently the Hamiltonian, takes the form of Eqn. 2.32 for circular geodesics in a spherically symmetric spacetime. The EOB approach proceeds by seeking a map from the general two-body problem to an equivalent spherically symmetric metric

$$ds^2 = -A(R)c^2 dT^2 + \frac{D(R)}{A(R)}dR^2 + R^2(d\theta^2 + sin^2\theta d\varphi^2) \tag{2.40}$$

where the functions are written as expansions in $(GM/c^2 R)$. The Hamiltonian of geodesic motion in this metric captures the conservative portion of the motion (as can be seen by the fact that the metric has no T dependence). Radiation-reaction terms can then be added to this Hamiltonian to capture energy carried away by gravitational waves. This process results in coupled differential equations which can be evolved numerically and from which the gravitational waveforms can be obtained.

A key feature of the EOB approach is the use of *Padé resummation*. Given a function $f(x)$ with Taylor series $\sum a_i x^i$ one "resums" the series into the ratio of polynomials $\sum b_i x^i / \sum c^i x_i$ by expanding this ratio in a Taylor series and matching the coefficients to the original series, then solving for the b_i and c_i. The resulting function can converge to f faster than the Taylor series.

The expansion of the unknown functions in the metric to useful order leads to unknown coefficients that can not be determined from the problem alone. These must be found by fitting the results to a waveform obtained through other means. This is one application of waveforms from Numerical Relativity, discussed in the next section. The resulting waveforms are called *EOBNR* [18].

2.6 Numerical Relativity

Analytic methods themselves can not provide complete gravitational waveforms valid through the merger of the two objects and the evolution of the resulting single object to a steady state. Even where there is hope of being able to capture the full physics, as in EOB or phenomenological models that have been developed, the results must be tuned against complete waveforms which must be calculated using other means.

There is a long history of solving partial differential equations numerically on computers, and while Einstein's equations are in a particularly difficult class of such equations, we can hope to obtain complete solutions numerically. This proved to be a remarkably difficult problem, and progress through the 1990s was slow due to both conceptual difficulties and limited computational power. However, following Pretorius' successful simulation of two merging black holes [10] and breakthroughs by researchers at Goddard [11] and the University of Texas at Brownville and Florida Atlantic University [11] in 2005 the field has expanded rapidly. Readers are referred to the textbook by Alcubierre [19] and review articles by Husa [20] and Hindler [21] for more details than can be provided here.

The basis of most approaches in numeric relativity (hence "NR") is the "3+1" decomposition of Einstein's equation. There is a similar situation in electromagnetism. Maxwell's equations may be expressed in a geometric, four-dimensional form

$$F^{\mu\nu}_{,\nu} = j^\mu$$
$$\mathcal{F}^{\mu\nu}_{,\nu} = 0$$

where F is the electromagnetic field tensor and $\mathcal{F}$ is its dual. However, this is not the most usual form in which to do calculations. Instead we choose a time direction and rewrite in terms of separate spatial and temporal differential operators,

$$\nabla \cdot \mathbf{E} = 4\pi\rho \tag{2.41}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{2.42}$$

$$\partial_t \mathbf{E} = \nabla \times \mathbf{B} - 4\pi j^\mu \tag{2.43}$$

$$\partial_t \mathbf{B} = -\nabla \times \mathbf{E} \tag{2.44}$$

Note that equations 2.41 and 2.42 are *constraint equations* that impose conditions on valid solutions at any given time, and in particular for initial conditions. Equations 2.43 and 2.44 are *evolution equations* which determine how to evolve the system.

Wile much work in numerical relativity includes matter, for the purpose of this thesis we restrict our attention to *vacuum spacetimes*, that is, situations where $T_{\mu\nu} = 0$ everywhere. It is a remarkable fact that, although black holes are characterized by a mass parameter, a spacetime containing only black holes and gravitational waves is a vacuum. We need therefore only be concerned with general relativity, without considering the equations governing any other fields.

In general relativity the split into space+time is accomplished by writing the metric as

$$ds^2 = -(\alpha^2 - \gamma_{ij}\beta^i\beta^j)dt^2 + 2\gamma_{ij}\beta^j dt\, dx^i + \gamma_{ij}dx^i\, dx^j,$$

where γ_{ij} is a metric on the slices of constant time t, and the scalar function α and vector field β^i are commonly used to encode the freedom of coordinate choice. This results in a division of the Einstein equation. 2.17 into constraint and evolution equations. However, this split is not uniquely determined and care must be taken to ensure that the resulting systems are well-behaved. In particular, there may be issues ensuring that the constraints remain satisfied as the system evolves. While this is guaranteed mathematically in the continuum limit it may not be true once the system is discretized.

Once a scheme for performing the 3+1 split has been decided upon it is then necessary to find initial data that captures the situation of interest and satisfies the constraint equations. There are many ways to do this. Most approaches assume the spatial metric on the initial slice is *conformally flat*, meaning each point has a neighborhood which can be mapped to flat space. We have seen that at any one point the metric can be made to be the flat-space metric with vanishing first derivatives, but this does not extend into a neighborhood in general. In particular, NR systems do not include the history of gravitational waves that would have been produced by the system prior to the start of the simulation. This deviation from the true physics manifests as a burst of spurious "junk radiation" at the start of the system.

The creation of initial conditions is complicated for black holes by the presence of singularities, which can not be captured in a simulation. Most codes adopt the "moving puncture" approach. The idea is to consider black holes as "wormholes" with an

internal asymptotically flat region, and then compactify this region. A conceptually simpler approach is to excise the points near the singularity from the computational grid, but care must be taken to ensure that the excised region is causally disconnected from the computational domain in order to avoid unphysical results.

There are a number of techniques for performing the evolution. A common approach is to evaluate the metric on a grid and replace derivatives with differences. An alternative is to use *spectral methods*, which involve expanding the solution in terms of a set of basis functions, and then evolving the coefficients. In this scheme differentiation can be performed analytically. In either scheme there is a tradeoff between accuracy and performance. A finer grid will give more accurate results, but at the cost of higher computational cost. As simulations typically take weeks to months even on large computer clusters, improving the performance is desirable. Most codes therefore use some form of *adaptive mesh refinement*, employing a finer grid only in regions close to the holes where the metric is changing rapidly.

In order to extract gravitational-wave information most approaches use the *Newman-Penrose scalar*, which in vacuum may be written

$$\Psi_4 = R_{\alpha\beta\gamma\delta} n^\alpha \bar{m}^\beta n^\gamma \bar{m}^\delta$$

where m and n are vectors constructed from the basis vectors, which in spherical coordinates are

$$n = \frac{1}{\sqrt{2}} \left(\hat{t} - \hat{r} \right)$$
$$m = \frac{1}{\sqrt{2}} \left(\hat{\theta} + i\hat{\phi} \right)$$

and $\bar{m}$ is the complex conjugate of m. The gravitational-wave strain is related to Ψ_4 by two time derivatives,

$$\Psi_4 = \ddot{h}_+ - i\ddot{h}_\times$$

2.6.1 Hybrid Waveforms

As noted, NR simulations require a great deal of time and resources; typically a single simulation starting 10 orbits before the merger will require a few weeks of runtime on approximately 50-100 processors [22]. These requirements scale with the length of the waveform extracted. Long waveforms are therefore prohibitively

expensive; the longest currently available span about 30 cycles before merger. Systems of astrophysical interest less massive than about $36M_\odot$ will spend many more cycles in the frequency range to which Initial LIGO is most sensitive. We therefore desire much longer waveforms. Post-Newtonian waveforms can not be extended upwards into the late inspiral and merger, and NR waveforms can not be extended downwards to the early inspiral. However, we can consider "stitching" a pN waveform to an NR in order to create a *hybrid waveforms.*

This turns out to be possible, although as we will see in chapter 7 there are subtleties in doing so that have not yet been fully resolved. One approach described in Ref. [23] is to work with Ψ_4 and match the numerical waveform to the pN waveform by adjusting the time and phase offsets of the pN waveform to minimize the quantity

$$\Xi(\Delta t, \Delta \phi) = \int_{t_1}^{t_2} \left[\phi_{\mathrm{NR}}(t) - \phi_{\mathrm{pN}}(t - \Delta t) - \Delta \phi \right]^2 \, dt \; . \tag{2.45}$$

This technique has been shown [24] to give good results when the pN waveform is taken to be the time-domain *TaylorT4* waveform (see Appendix A) with terms up to 3.5-pN in phase and 3.0-pN in amplitude.

2.7   Conclusions

In this chapter we reviewed the basic properties of gravitational waves and methods used to model such waves from the inspiral and merger of systems of compact binaries. In the next chapter we discuss the principles behind the LIGO detectors, which are looking for gravitational-wave signals. Then in chapter 4 we discuss how pN waveforms are used to detect signals in the LIGO data.

Chapter 3

The LIGO Gravitational Wave Detectors

The fact that matter responds to gravitational waves as described in Sec. 2.3 offers the possibility of making a direct detection of such waves. Although several methods of detection have been proposed, we focus here on that used by LIGO and, with some minor differences, Virgo and GEO. Again, this presentation will be of necessity brief. We refer readers to the textbook by Saulson [25] for a much more comprehensive treatment.

3.1   A Toy Model

We motivate our discussion of the LIGO interferometers with a toy model. We wish to detect gravitational waves, and one method is suggested by the analysis of the preceding chapter; we look for the strain, $\Delta L/L$, by measuring the change in length of some system. In principle we could measure the changing distance between two freely-falling masses with a rigid ruler. The atoms in the ruler would not be freely-falling and therefore the distance change along the ruler would be different than that between the masses. However, the accuracy needed makes this approach impractical. A better approach is to measure distance by sending a projectile, say a marble, with known velocity, through the length and measuring the travel time. To avoid complex issues of synchronizing clocks at different points in general relativity we add a (hypothetical perfectly elastic) rubber wall at the far end, and measure how long

it takes to return. If the velocity of the marble is much larger than the velocity of the wall induced by the wave (that is, $v \gg h\omega L$ where h is the strain and ω the gravitational-wave frequency) then there is a simple relationship between the round-trip travel time and the amplitude of the wave. However, this requires unrealistic precision in measurement, the uncertainty in the marbles' launch time will swamp the small changes in length since, as we will see in the next chapter, h is typically very small.

Therefore, we instead construct a *null experiment* where we try to determine if a given quantity is exactly zero. We arrange two perpendicular paths, fire marbles down each at the same time and measure the difference in return times. As gravitational waves are rare, we "lock" the system by shifting one or both of the walls such that the marbles always collide exactly (say by measuring their recoil angle). Once locked, deviations in the length of either or both arms will cause the difference in arrive times to become non-zero, which can be determined by a change in the marbles' recoil angles or lack of collision entirely.

To obtain robust results we want ΔL to be as large as possible. Since gravitational waves are week, this means increasing L. Practical concerns may limit the ability to do this. For example, clearly the entire path from source to walls must be in vacuum in order for the marbles not to lose energy, and building large vacuum systems is difficult and expensive. We therefore use a trick and add a second set of walls between the source and reflectors, and arrange the paths so that the marbles bounce back and forth several times before returning to the source. This effectively extends L by the distance between the two surfaces multiplied by the number of bounces.

Finally, in this toy model the readout will be the number of marbles that land in a "bin" subtending a recoil angle. Our ability to detect a gravitational wave with statistical confidence then reduces to our ability to count marbles. We can model this as a Poisson process, where probability of observing N marbles is

$$p(N) = \frac{\bar{N}^N \exp(-\bar{N})}{N!}$$

where $\bar{N}$ is the average expected number of marbles per observation period. The error in estimating N from counting goes as $1/\sqrt{N}$, and it is therefore advantageous to send out as many marbles as possible.

3.2 Interferometric Gravitational Wave Detectors

The toy model presented above captures the essential principles behind LIGO, with the significant difference that light is used instead of marbles. A cartoon of the LIGO detectors is shown in Fig. 1. Ignoring the two `ITM` mirrors for the moment, the laser, beam splitter and two end mirrors (labeled `ETMX` and `ETMY` for "end test masses") would form a *Michelson interferometer*, and parallel the original toy model with marbles and two reflecting surfaces. The use of light actually simplifies the analysis because light travels on null geodesics, so

$$ds^2 = 0 = g_{\mu\nu}dx^\mu dx^\nu$$

We now consider a +-polarized gravitational wave travelling in the z direction, and place the arms on the x and y axes. We again assume the frequency of the wave is large compared to the travel time between the arms, which implies that, over the round trip, h_+ may be taken to be constant and the metric becomes

$$g_{\mu\nu} = -dt^2 + (1 + h_+)dx^2 + (1 - h_+)dy^2 + dz^2$$

Then for the x axis, restoring physical units,

$$c^2 dt^2 = (1 + h_+)dx^2$$

and the round-trip travel time is

$$t_x = \frac{2L}{c}\sqrt{1 + h_+} \approx \frac{2L}{c}\left(1 + \frac{h_+}{2}\right)$$

where we approximate the square root by its Taylor series and ignore higher-order terms in h_+. Similarly the round-trip light travel time along the y arm is

$$t_y = \frac{2L}{c}\left(1 - \frac{h_+}{2}\right)$$

Considering a single wavefront leaving the beam splitter, the difference in return time is

$$\Delta t = t_x - t_y = \frac{2L}{c}h_+$$

In interferometry we measure the difference in phase between returning wavefronts. This difference will cause an interference pattern that will serve as the readout. If we
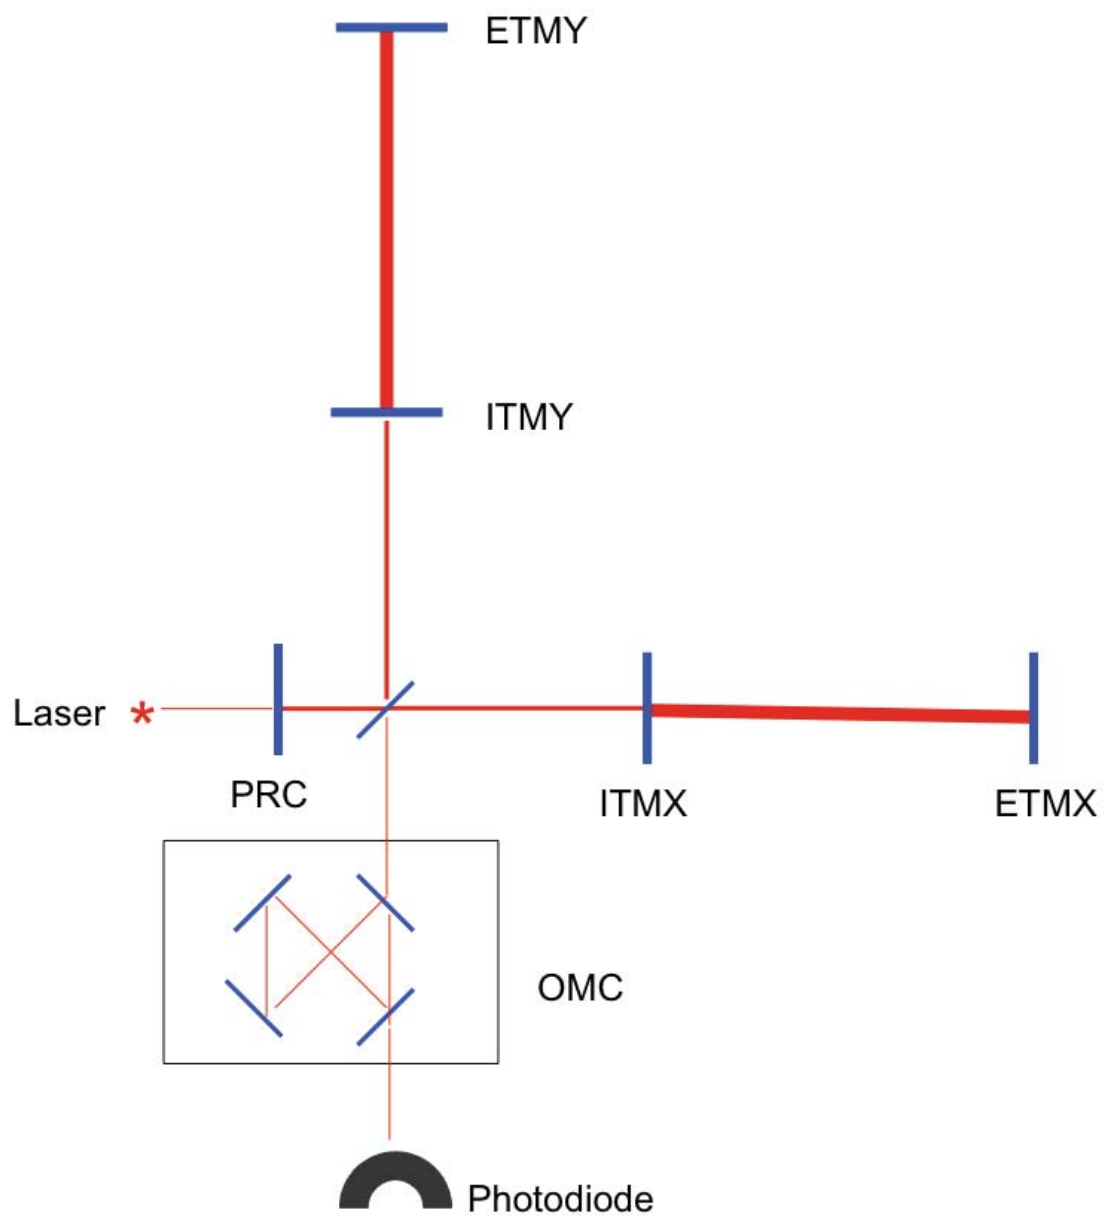
Figure 1 :   Block diagram of LIGO, see the text for description.

use laser light of frequency f and wavelength $\lambda = c/f$ then a time difference of Δt corresponds to a phase difference of

$$\Delta\Phi = \frac{2\pi f}{\Delta t} = \frac{4\pi fL}{c}h_+ = \frac{4\pi L}{\lambda}h_+$$

As in the marble example, we can increase the sensitivity of the detector by increasing L, but practical considerations prevent us from doing so. One of these considerations is in fact the same for marbles and light; the travel path must be in vacuum. We therefore employ the same trick and insert two additional mirrors, indicated as ITMX and ITMY (for "inner test masses"). The addition of these mirrors creates a *Fabry-Perot cavity* in each arm. By arranging the mirrors to be an integer number of wavelengths apart a resonance is built up that can trap the light for extended periods, approximately 200 bounces. It can be seen that if the mirrors are not appropriately spaced there will be destructive interference between the light moving in different directions. As the power in the beam must be conserved, this results in energy leaking out of the cavity, reducing the efficiency. The addition of the Fabry-Perot cavities provides an increase in phase difference of approximately two orders of magnitude.

In interferometry it is typically most useful to think of light as a wave. However, where in the toy example our ability to detect gravitational waves was limited by our ability to count marbles, in real LIGO we are limited by our ability to count photons. This manifests as a noise source called *shot noise*, to be discussed shortly. Photon number is related to laser energy by $E = h\nu$, so we therefore want to use as powerful a laser as possible. There are, however, technical obstacles to doing so. In the latest LIGO run the laser power was up to 14 W, although it was not always possible to run at this level.

In lieu of raising the laser power we can at least ensure that no power is wasted. LIGO is configured such that the beams interfere destructively when they recombine, we say the detector sits on a *dark fringe* [1]. By conservation of energy all the power emitted by the laser must go back towards the laser (neglecting the portion lost

---

[1]If we sat exactly on a dark fringe then any change in the arm lengths would cause an increase of light at the readout, and we would be unable to determine in which direction the mirrors were moving. We therefore sit a bit off the dark fringe. In addition, this condition is necessary for DC readout, to be discussed shortly.

to scattering, absorbed by the mirrors, etc). We can recover this power by adding another mirror, indicated on the diagram as `PRC`, or power-recycling cavity.

The final feature on Fig. 1 is the `OMC` (output mode cleaner), which was an important addition to the latest science run. A full description of this element is outside the scope of this thesis, but we note briefly that the cross-section of a laser beam can be decomposed into Hermite-Gaussian modes, as in Fig. 2. The higher-order modes do not contribute to the readout signal, however they do contribute to the shot noise. It is therefore advantageous to suppress such modes. In addition the OMC was necessary to enable *DC readout*, a scheme in which the power present at the output port of the interferometer is directly proportional to gravitational wave strain [26]. Such a scheme offers reduced shot noise and other advantages over *RF readout*, where the length changes are determined by the interaction of light at the frequency of the laser and light at sideband frequencies. We draw attention to the OMC and disregard other mode cleaners because the OMC tended to produce glitches, especially in the Livingston detector.

3.2.1   Readout

In order to operate correctly the two Fabry-Perot cavities and the power-recycling mirror must be positioned such that the light is resonant. The Michelson must likewise by arranged so that the output photodetector is on a dark fringe. When the detector is in this state we say it is *locked*, at that point a gravitational wave will perturb the system within limits and produce light out the output. However, the resonances must be very finely tuned and left untouched the system would quickly fall out of lock due to random motions of the mirrors. There is therefore a need for continuous, active corrections implemented by a system of sensors and servos throughout the detector.

Ignoring the OMC, in Fig. 1 the degrees of freedom are the two initial test masses, the two end masses, and the PRC. However, it is the lengths betweens elements that we are interested in, so the there are fewer degrees of freedom than there are optical elements. In addition it is convenient to work with linear combinations of these degrees of freedom:

- `DARM`: the differential arm length, (ETMX-ITMX) - (ETMY-ITMY)

- `CARM`: the common arm length, ((ETMX-ITMX) + (ETMY-ITMY) / 2
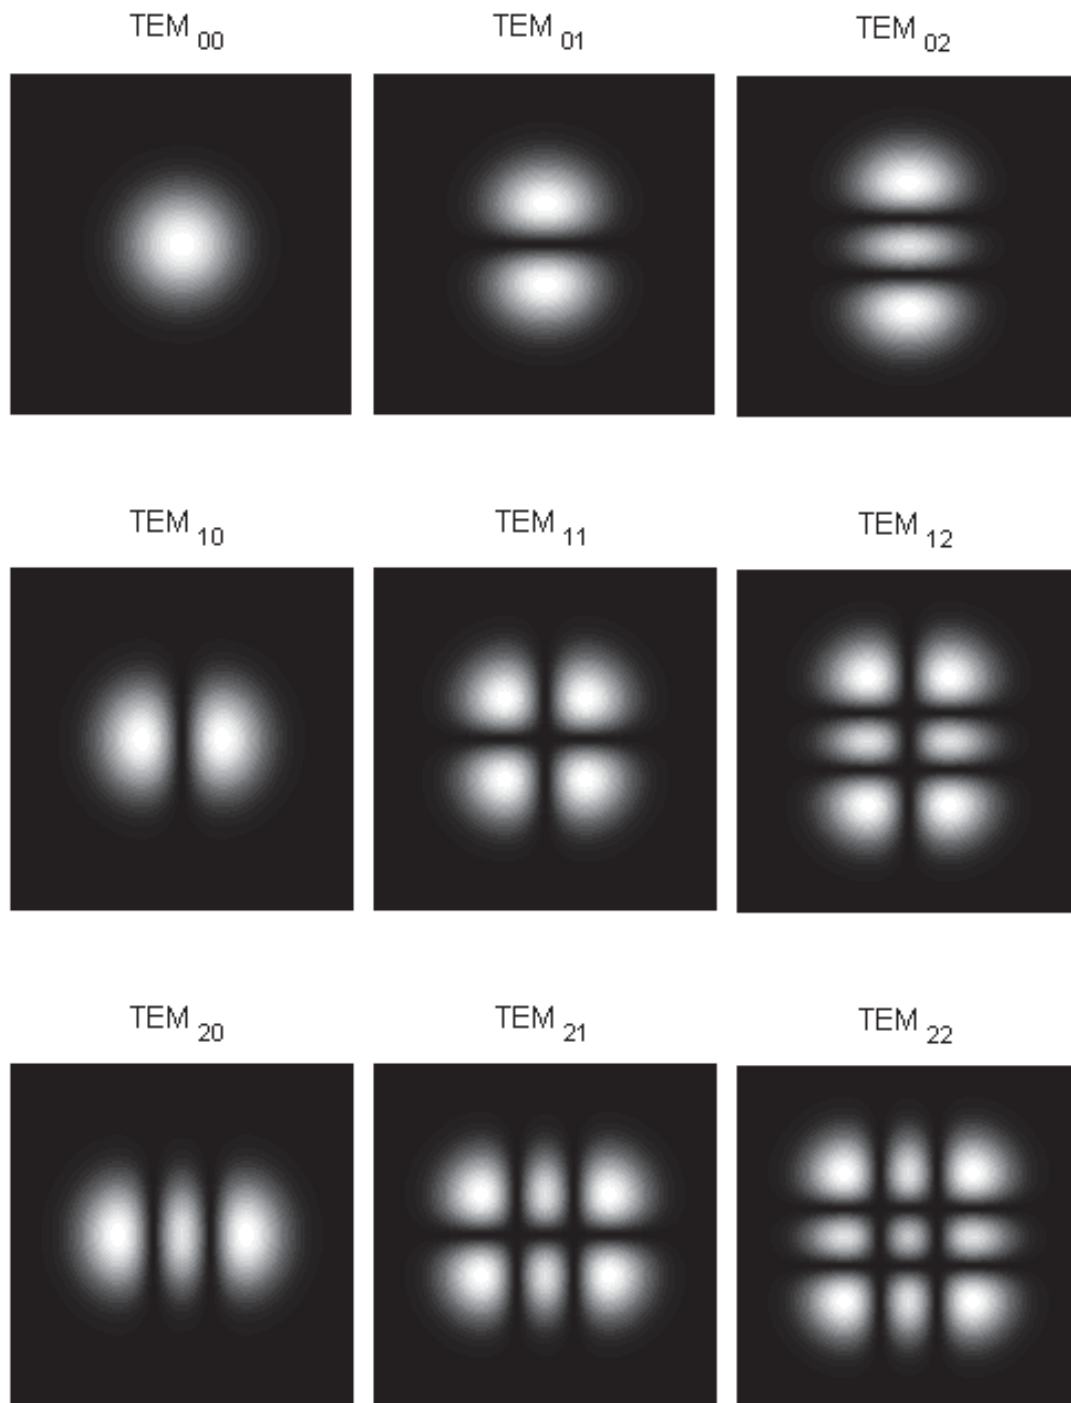
Figure 2 : Basis functions for the cross-sectional distribution of power in a laser beam. The lack of symmetry in the higher-order modes can induce instabilities in the LIGO optics. (*Public-domain image taken from Wikipedia [27]*)

- `PRC`: The common recycling cavity length

- `MICH`: the differential motion of the "small Michelson" comprising the ITMs and the beam splitter

One of the sensors used to keep the system on resonance is the output photodector. Along with this, note that `DARM` is the quantity of interest in the experiment, the change in length produced by a gravitational wave. It turns outs that rather than reporting the signal at the photodiode directly a better output is the extent to which this degree of freedom is off the resonance condition, which is recorded as `DARM_ERR`. Henceforth, and especially in chapter 10 we consider this "the output of the detector". It is not, however, the data stream in which we will search for gravitational waves. The detector output needs to be calibrated with respect to the detector's frequency response. This can be measured by injecting a sine wave of known amplitude into the system by actuating one of the mirrors, and measuring the amplitude and phase of `DARM_ERR`. The result is a complicated function of frequency. This can then be inverted to map `DARM_ERR` back to the true input, the result is stored as `LSC-STRAIN` (for *Length Sensing and Control*), and it is that channel on which gravitational-wave searches are performed.

3.2.2   Noise Sources

In addition to gravitational-wave sources the detector is subject to various other influences collectively known as noise. One class of noise can be modeled as a *Gaussian random process*, as we will discuss in the next chapter. These are best characterized by their frequency profiles. The dominant source are:

- *Seismic noise* is due to the coupling of the detector to the ground. Much work has been been done, and much research continues to be done, to isolate the mirrors from the environment. However, the isolation is not complete. This noise source dominates at low frequency, rising sharply below 40 Hz in Initial LIGO. In Advanced LIGO we hope to push this so-called *seismic wall* down to 10 Hz. This noise source includes the natural constant vibrations of the Earth, wind blowing over nearby structures, and anthropogenic sources such as vehicles

near the sites, logging activity, etc. Seismic noise is not Gaussian [2], however the deviation from Gaussianity will not impact any of the following analysis.

- *Thermal noise.* In any system each degree of freedom has an expected energy of $E = k_B T/2$ where T is the ambient temperature and k_B is Boltzmann's constant. In LIGO this energy manifests as random motion throughout the detector, although the motion from the test masses and the wires from which they are hung are the largest contributors to noise. The noise produced dominates from 40 Hz to approximately 200 Hz.

- *Shot noise* is the uncertainty inherent in counting photons. Consider the ideal case, where there are no other noise sources and the detector is locked off a dark fringe, so that we expect N photons in each sample interval. Due to counting uncertainty this value will actually vary according to the Poisson distribution. It can be shown that the output can be modeled as a time-series of delta functions with strengths proportional to $1/\sqrt{N}$. In the frequency domain this gives a white spectrum, however the arms of the interferometer act as a filter that amplifies this noise with increasing frequency. The result is that shot noise dominates above ~ 1 kHz.

In addition to these broad-band sources of noise there are also *lines*, particular frequencies at which the noise is much greater than the three sources above would produce. Two of the most significant are:

- *Electrical noise.* Despite shielding, at 60 Hz there is a sharp increase in the noise level due to the frequency of the US electrical grid.

- *Violin modes.* Although the wires suspending the mirrors vibrate over a range of frequencies due to thermal noise, the suspension system has a resonance at about 340 Hz, producing much more noise here.

There are also lines at higher harmonics of these frequencies.

---

[2]For Gaussian noise the power spectral density has an exponential distribution at each frequency. The mean equals the variance for an exponential distribution. It is therefore possible to measure the "degree of Gaussianity" in a noise source from the ratio of the mean to the variance of the PSD at each frequency. For Gaussian noise this ratio will be 1, however it is larger for seismic noise.

In addition to these continuous noise sources there are numerous *glitches*, short transient events. These glitches have a wide variety of different morphologies and originate from many different sources. Some glitches are reactions to external environmental conditions and others are internal to the detector. As one straightforward example, a heavy object dropped on a nearby road can, through seismic coupling, shake the mirrors and produce a sharp impulse in `DARM_ERR`.

At the most abstract level gravitational-wave searches entail looking for features in the data that deviate in some way from the continuous background. It is therefore not surprising that glitches can interfere with searches and hence must be removed to the extent possible. Chapter 9 will present the infrastructure used to store information about the state of the detector and environment, and chapter 10 will discuss part of the effort to identify glitches and remove them from the analysis.

3.3 Conclusions

In this chapter we reviewed the basic elements of interferometric gravitational-wave detectors. In the next chapter we discuss how we search the data obtained by these detectors for signals of the kind described in the previous chapter.

Chapter 4

Searching for Gravitational Waves from Compact-Binary Coalescences

In the previous chapter we showed that interferometric detectors, such as LIGO, are capable of encoding the presence of gravitational waves in the phase shifts of light. These shifts over time, or rather the motion of the servos to prevent them, are recorded as the data stream `DARM_ERR`. As noted, this data stream will also contain noise. This noise will contain both a continuous component which sets the overall sensitivity of the detector, and short glitches which can interfere with the search.

In this chapter we discuss a method to extract signals, if present, from the noise and hence allow the LIGO and Virgo collaborations to claim a direct detection of gravitational waves. For the most part in this chapter we will disregard implementation details. In particular, we will treat quantities as continuous in both time and magnitude. In reality the data is sampled at 16384 Hz, and all time integrals should be replaced by sums. Values are also stored in a discrete form, although the resolution is no coarser than the resolution of IEEE floating-point numbers. Analysis codes therefore do not need any special features to deal with this discretization, beyond the usual care taken not to accumulate numerical errors. In particular, Fourier transforms, special functions, and similar functionality can be provided by standard code libraries.

4.1 Detecting Signals in Noise

There are many ways to search for evidence of gravitational waves in noise and to a large extent the method used will depend on the nature of the signal being sought. In this chapter we focus on *matched filtering*, which is optimal given certain assumptions about the noise and if we have an accurate model of the signal for which we are searching. Such a method is ideal for searching for the coalescence of compact binaries, as discussed in chapter 2, as we can use the analytic approximations as the model. The question of how well these models match real signals then becomes critical, and we will return to this question in subsequent chapters.

4.1.1 Random Processes

We start by modelling the noise in the detectors, following the treatment in [28]. Seismic noise is not predictable in any detailed way, and thermal and shot noise are quantum-mechanical in nature and therefore can not be predicted even in principle. The noise is therefore a *random process*. Such a process is described by a collection of probability density functions of the form

$$p(y_n, t_n; y_{n-1}, t_{n-1}; \ldots y_0 t_0) dy_n dy_{n-1} \ldots dy_0$$

which gives the probability that the value at time t_0 will will lie between y_0 and $y_0 + dy_0$ and that the value at time t_1 will will lie between y_1 and $y_1 + dy_1$, and so on. By considering a hypothetical ensemble of such processes we can construct average values of functions of the process at one or more times. We denote the ensemble average of a quantity x by angle brackets, $\langle x \rangle$, for example the average value of the product $y(t_1)y(t_2)$ is

$$\langle y(t_1)y(t_2) \rangle = \int y_2 y_1 p(y_2, t_2; y_1, t_1) \, dy_2 \, dy_1$$

We now make the following assumptions about the LIGO noise:

- That it is *stationary*, the values of the functions 4.1.1 depend only on the differences between the times, and not any absolute external clock. This is not necessary a good model for seismic noise, but by breaking the analysis into small chunks of time it will be accurate enough over the span of each.

- That it is *Gaussian*, that is, the functions 4.1.1 are all of the form $\exp(-Y^T A Y)$ where Y is a vector of y values and A is a positive-definite matrix. This is not necessarily a good model for, eg, electrical noise, which is better modeled as a sine wave with time-dependant random amplitude. This is definitely not a good model for the numerous non-Gaussian glitch mechanisms. However, this assumption is again approximately correct. We will further assume that the mean is 0 (eg, that the mirrors are just as likely to swing one way as the other).

- That it satisfies the *ergodic hypothesis* so that we may replace ensemble averages with time averages.

Now consider the quantity $S_y(f)$ called the *power spectral density* (PSD), defined as

$$S_y(f) \equiv \lim_{T \to \infty} \frac{1}{T} \left| \int_{-T/2}^{T/2} (y(t) - \bar{y}(t)) e^{2\pi i f t} dt \right|^2 \tag{4.1}$$

Using Parceval's theorem it can be shown that

$$\int_0^\infty S_y(f) \, df = \sigma_y^2$$

That is, $S_y(f)$ measures the contribution of each frequency to the total variance of the process. It can also be shown that

$$\langle \tilde{y}(f) \tilde{y}(f') \rangle = \frac{1}{2} S_y(f) \delta(f - f') \tag{4.2}$$

4.1.2 The Matched Filter

We now specialize to the case of LIGO. We denote the noise in the detector as $n(t)$, and a gravitational-wave signal as $h(t)$. The output of the detector is then $s(t) = n(t) + h(t)$. We seek a filter on the data that we can use to infer the presence of the signal.

Consider the most general linear filter which in the discrete case would be

$$\hat{s} = \sum_i s_i K_i \tag{4.3}$$

Since we will want a real result, we require the K_i to be real. In the continuum limit this becomes

$$\hat{s} = \int_{-\infty}^\infty s(t) K(t) \, dt \tag{4.4}$$

We now define the signal strength S as the expected value of $\hat{s}$ when the signal is present:

$$S = \langle \hat{s} \rangle$$

$$= \left\langle \int_{-\infty}^{\infty} s(t)K(t)\,dt \right\rangle$$

$$= \int_{-\infty}^{\infty} \langle s(t)K(t) \rangle\,dt$$

$$= \int_{-\infty}^{\infty} \langle s(t) \rangle K(t)\,dt$$

$$= \int_{-\infty}^{\infty} \langle n(t) + h(t) \rangle K(t)\,dt$$

$$= \int_{-\infty}^{\infty} \left(\langle n(t) \rangle + \langle h(t) \rangle \right) K(t)\,dt$$

$$= \int_{-\infty}^{\infty} \left(0 + h(t) \right) K(t)\,dt$$

$$= \int_{-\infty}^{\infty} h(t)K(t)\,dt$$

Now, since $K(t)$ is real we can replace it by its complex conjugate and then apply Parseval's theorem to write this as

$$S = \int_{-\infty}^{\infty} h(t)K^{\star}(t)\,dt = \int_{-\infty}^{\infty} \tilde{h}(f)\tilde{K}^{\star}(f)\,df$$

We characterise the noise N as the rms value of $\hat{s}$ when the signal is absent:

$$N^2 = \langle \hat{s}^2(t) \rangle - \langle \hat{s}(t) \rangle^2 \tag{4.5}$$

$$= \langle \hat{s}^2(t) \rangle \tag{4.6}$$

$$= \int_{-\infty}^{\infty} K(t)K(t')\langle n(t)n(t') \rangle\,dt\,dt' \tag{4.7}$$

$$= \int_{-\infty}^{\infty} dt\,dt'\,K(t)K(t') \int_{-\infty}^{\infty} df\,df'\,e^{2\pi i f t}e^{-2\pi i f' t'}\,\langle \tilde{n}^{\star}(f)\tilde{n}(f') \rangle \tag{4.8}$$

$$\tag{4.9}$$

Using Eqn. 4.2 this becomes

$$N^2 = \int_{-\infty}^{\infty} df\,\frac{1}{2}S_n(f)|\tilde{K}(f)|^2 \tag{4.10}$$

We now introduce a new function

$$\tilde{u}(f) = \frac{1}{2} S_n(f) \tilde{K}(f) \tag{4.11}$$

In terms of which we can write the signal-to-noise ratio, which we henceforth denote as ρ as

$$\rho = \frac{\int_{-\infty}^{\infty} df \, \frac{\tilde{h}(f)\tilde{u}^\star(f)}{(1/2)S_n(f)}}{\left[\int_{-\infty}^{\infty} df \, \frac{\tilde{u}(f)\tilde{u}^\star(f)}{(1/2)S_n(f)}\right]^{1/2}} \tag{4.12}$$

This motivates the definition of an operator mapping pairs of functions to real numbers

$$(A \,|\, B) = 2 \int_{-\infty}^{\infty} df \, \frac{\tilde{A}(f)\tilde{B}^\star(f)}{S_n(f)} \tag{4.13}$$

This operator has the following properties

- Conjugate symmetry, $(x|y) = (y|x)^\star$

- Linearity in the first argument $(ax + by|z) = a(x|z) + b(y|z)$ for a, b numbers and x, y, z functions. This follows from the linearity of the Fourier transform.

- Positive-definiteness $(x|x) \geq 0$ and $(x|x) = 0$ iff $x = 0$. This follows from the positive-definiteness of the product $aa^\star$ for $a \in \mathcal{C}$.

The operator has all the properties of an inner product on the vector space of functions. We may therefore consider

$$\frac{u}{(u|u)}$$

to be a normalized vector.

An important feature of this inner product, which we will use repeatedly, is that the probability of obtaining any particular pattern $h(t)$ in the data is, up to normalization [29],

$$p(h) \propto \exp\left(-\frac{1}{2}\left(h \,|\, h\right)\right) \tag{4.14}$$

Now using Eqn. 4.13 the SNR becomes

$$\rho = \frac{(h \,|\, u)}{(u \,|\, u)^{1/2}} \tag{4.15}$$

We next seek a function u (and hence K) that will maximize ρ. In this form it is clear that u and h must be parallel, and therefore

$$\tilde{K}(f) \propto \frac{\tilde{h}}{S_n(f)} \tag{4.16}$$

The constant cancels, and may therefore be set to 1.

Combining these results, the value we will use to determine whether or not our data s contains the gravitational waveform h (henceforth called the *template* waveform) is

$$\rho = \frac{(s \mid h)}{\sqrt{(h \mid h)}} \tag{4.17}$$

The template has an unknown phase, denoted ϕ_0 in Eqn. 2.39. As this term has no frequency dependence it gives rise to a constant of the form $\exp(i\phi_0)$, which can be pulled out of the integral. We can then maximize over this value by taking the absolute value.

The result is valid for a segment of data of length t seconds and a waveform of the same length. However, there may be a signal in the data that does not end where the template does. We should therefore evaluate 4.17 repeatedly, sliding the template so that it ends at a different time for each iteration. However, this can be done in a single operation. If $\tilde{h}(f)$ is the Fourier transform of $h(t)$, then the Fourier transform of $h(t + \tau)$ is

$$\tilde{h}(f)' = \int h(t + \tau)e^{2\pi i f t}dt = \int h(t')e^{2\pi i f(t'-\tau)}dt' = e^{-2\pi i \tau}\tilde{h}(f)$$

Substituting this into the matched filter,

$$\rho(\tau) = \frac{2}{\sqrt{(h \mid h)}} \left| \int e^{-2\pi i f \tau} \frac{\tilde{s}(f)\tilde{h}^\star(f)}{S_n(f)} \, df \right| \tag{4.18}$$

where we have added the absolute value from the maximization over ϕ_0. This is just the inverse Fourier transform of the quantity $\tilde{s}(f)\tilde{h}^\star(f)/S_n(f)$.

4.1.3 The Overlap

Consider data consisting of noise and a signal h. Let ρ be the SNR obtained by filtering this data with template h. Then the loss in SNR incurred by filtering with

an incorrect template, h' is

$$\rho - \rho' = \frac{(s \mid h)}{\sqrt{(h \mid h)}} - \frac{(s \mid h')}{\sqrt{(h' \mid h')}} \tag{4.19}$$

$$= \frac{(n \mid h)}{\sqrt{(h \mid h)}} + \frac{(h \mid h)}{\sqrt{(h \mid h)}} - \frac{(n \mid h')}{\sqrt{(h' \mid h')}} - \frac{(h \mid h')}{\sqrt{(h' \mid h')}}$$

Taking the mean of both sides the first and third terms vanish because

$$\langle (s \mid h) \rangle = 2 \int_{-\infty}^{\infty} df \, \frac{\langle \tilde{n}(f) \rangle \tilde{h}^\star(f)}{S_n(f)} = 0$$

where the last step follows from our assumption that the noise is Gaussian with zero mean. Equation 4.19 then becomes

$$\langle \rho - \rho' \rangle = \frac{(h \mid h)}{\sqrt{(h \mid h)}} - \frac{(h \mid h')}{\sqrt{(h \mid h')}}$$

$$= \sqrt{(h \mid h)} \left(1 - \frac{(h \mid h')}{\sqrt{(h \mid h)(h' \mid h')}} \right)$$

This motivates the definition of the *overlap* between two waveforms h and h' with respect to a given PSD,

$$\langle h \mid h' \rangle \equiv \frac{(h \mid h')}{\sqrt{(h \mid h)(h' \mid h')}} \tag{4.20}$$

which is a measure of how similar the two waveforms are, or conversely how dissimilar they are and hence the fractional SNR we can expect to lose using one rather than the other. We will use this repeatedly in subsequent chapters.

4.1.4 Trigger Selection

From the SNR time series, $\rho(t)$ we wish to select *triggers*, discreet points whose time and ρ value correspond to the time and significance of a potential detection. This is done using an algorithm described in Ref. [30].

Equation 4.14 says that even in the absence of a signal there will be random fluctuations in $\rho(t)$. We therefore require a threshold on $\rho(t)$ that will eliminate most of the triggers arising from random fluctuations, but not so high as to prevent detection of realistic signals. LIGO has chosen 5.5 as this threshold. However, it

is unnecessary and would be wasteful to create a trigger from every point of $\rho(t)$ in excess of this threshold. $\rho(t)$ may be thought of as "sliding" the template against the signal. As the two line up $\rho(t)$ will increase, achieving a maximum when they are perfectly aligned, and then decreasing. We therefore need only create a trigger from the largest value of $\rho(t)$. Since glitches can also elevate $\rho(t)$, and indeed may have larger values than real signals we can not simply take the largest value across the entire analysis but must instead choose a *window*. A logical choice for the length of this window is the time length of the template. For pN waveforms described in Sec. 2.5 this is defined as the time from which the instantaneous frequency is 40 Hz to the time where it becomes infinite.

The full trigger selection algorithm is:

```
for each sample point j
  if ρ_j > threshold
    if there is no event yet
      event_start = j
      event_snr   = ρ_j
    else if ρ_j > event_snr
      event_start = j
      event_snr = ρ_j
    else if (j - event_start) == template length
      record event
      event_start = j
      event_snr   = ρ_j
```

This is illustrated in Fig. 3. We will discuss some implications of this algorithm in Sec. 10.6.

4.2 The χ^2 Test

Consider again the SNR time series, Eqn. 4.18, which we now write as

$$\rho(t) = \frac{2}{\sqrt{(h \mid h)}} \int e^{-2\pi i f t} \frac{\tilde{h}^{\star}(f)}{S_n(f)} \cdot \tilde{s}(f) \, df$$

By the convolution theorem this becomes

$$\rho(t) = \frac{2}{\sqrt{(h \mid h)}} h(-t) \star S(t) \star s(t)$$

where $S(t)$ is the inverse Fourier transform of $1/S_n(f)$. If the signal is an impulse, $s(t) = \delta(t)$, then the response is the time-reversed template ("fuzzed" by the noise curve). Given a sharp feature in the data the SNR will therefore be elevated, even though the data looks nothing like a gravitational wave. More generally this demonstrates that the SNR alone is not sufficient to distinguish signals from noise in the presence of non-Gaussian features of the data. Given the presence of glitches this is a very practical concern.

We therefore supplement the analysis with a χ^2 signal-consistency test [31]. The idea is to check not only that ρ is large, but also that the SNR was accumulated
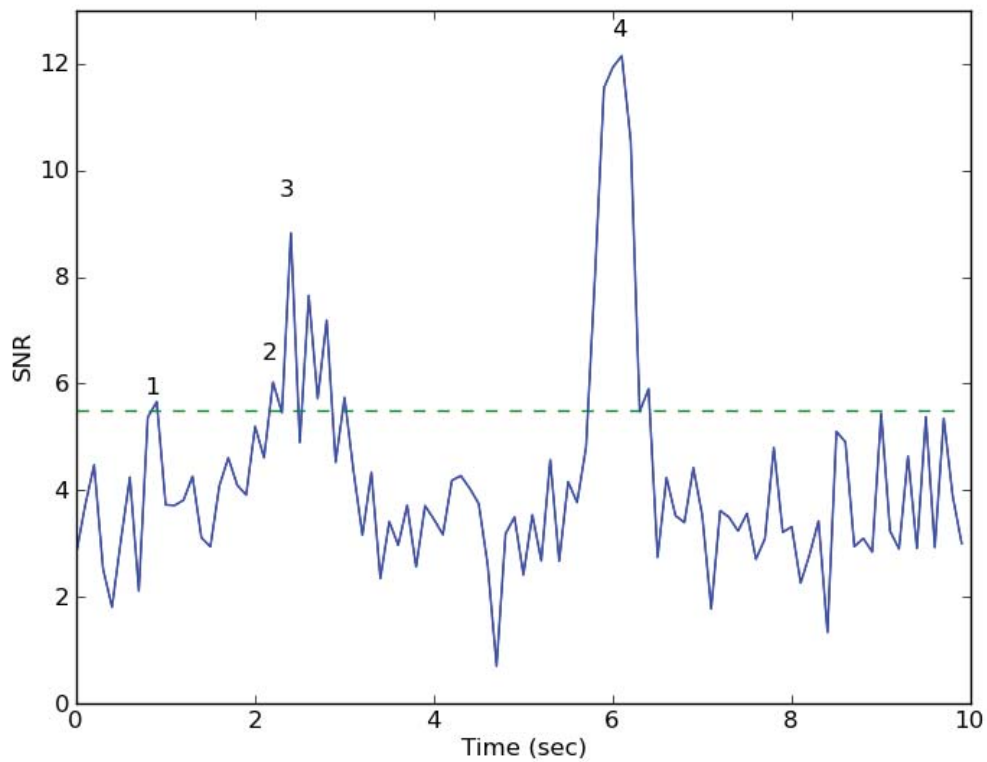
Figure 3 : A "cartoon" illustrating the trigger-selection algorithm. The graph shows an SNR time series of the kind produced by Eqn. 4.18. Assume the template is one second long. As the SNR rises above threshold just before point (1) the "window" opens. Point (1) is found as the nearest maximum. One second after point (1) the window closes and point (1) is recorded as a trigger. The SNR rises above threshold just before point point (2). This is more than one second away from point (1), so a new window opens. When point (3) is encountered it is recorded as the largest value, resetting the window. One second after point (3) the window closes and point (3) is recorded as a trigger. The two peaks following point (3) are within the one-second window, they are not recorded as triggers. Point (4) is likewise recorded as a trigger, as it is more than one second from (3).

over the frequency integration in a way consistent with a real signal. We do this by dividing the template into p sub-templates h_i, which match the original template between frequencies f_i and f_{i+1} and are zero elsewhere. The f_i are chosen such that

$$\int_{f_i}^{f_{i+1}} \frac{|h_i|^2}{S_n(f)}\, df = \frac{1}{p} \int_{f_0}^{f_c} \frac{|h|^2}{S_n(f)}\, df$$

where f_0 is the starting frequency (40 Hz in LIGO) and f_c is, as in Eqn. 2.39, the frequency at which we terminate the waveform. We then filter with these sub-templates independently to produce the p time series $\rho_i(t)$ and construct the quantity

$$\chi^2(t) = p \sum_{i=1}^{p} \left(\rho_i(t) - \frac{\rho(t)}{p} \right)^2 \tag{4.21}$$

If the signal exactly matches the template χ^2 will be zero, and it will increase as the signal and template deviate. Based on investigations in S5, we choose $p = 16$.

We incorporate this information in the analysis by constructing a quantity, ρ_{new} (new SNR), which downweights the SNR by χ^2,

$$\rho_{\text{new}}^2 = \begin{cases} \rho^2, & \chi_r^2 \leq 1, \\ \frac{\rho^2}{[(1+(\chi_r^2)^3)/2]^{1/6}}, & \chi_r^2 > 1, \end{cases} \tag{4.22}$$

where the first clause is added to avoid promoting triggers that happen to have anomalously low χ^2 values, and the form and factors of the second clause have been chosen based on studies to test the ability of the search to distinguish between simulated injected signals and glitches.

4.3 Choosing the Templates

We have thus far referred to the h as "the template", but recall from Eqn. 2.39 that the waveforms depend on the parameters of the binary, such as total mass M and symmetric mass ratio η. As discussed in Sec. 4.1.3, the difference between the signal and the template used in the matched filter corresponds directly to a reduction in ρ. We must therefore construct a *bank* of templates arranged to capture, to within some acceptable loss of SNR, all the signals of interest. To do so we consider the space of templates as a manifold, as in chapter 2, and we use the parameter values (or functions of these values) as coordinates. The distance between two points $\mathbf{p}$ and

q, with coordinates (=parameters) p_i and q_i, is defined in terms of the overlap of templates at those parameters

$$d = 1 - \langle h(\mathbf{p}) \,|\, h(\mathbf{q}) \rangle = \frac{(h(\mathbf{p}) \,|\, h(\mathbf{q}))}{\sqrt{(h(\mathbf{p}) \,|\, h(\mathbf{p})) \,(h(\mathbf{q}) \,|\, h(\mathbf{q}))}}$$

We then seek a set of templates, $\{h_i\}$ such that within a region of interest every point in the manifold is "sufficiently close" to at least one template:

$$\forall \mathbf{p}, \max_i \langle h(\mathbf{p}) \,|\, h_i \rangle \geq x$$

Such a bank will ensure that the loss in SNR from using the bank instead of testing the entire continuum of parameters is less than $x\%$. In LIGO we typically require no more than a 3% loss in SNR. As the distance to which we can detect signals scales with SNR and volume is the cube of distance, this corresponds to approximately a 10% loss in event rate.

It is possible to construct the bank *stochastically* [32]:

- Choose a point in the manifold at random

- If the overlap with any already-placed template is greater than the minimal match, discard it

- Otherwise, add it to the set and continue

- Repeat until there have been N consecutive rejected new candidates, where N is chosen to give a reasonable probability that the space has been covered.

In many cases we can do better then this. Since we are interested in distances on a manifold it makes sense to seek a metric on this space. Such a metric can be obtained by expanding the overlap function [33, 34]

$$\langle h(\lambda) \,|\, h(\lambda + \Delta\lambda) \rangle \approx 1 - \frac{1}{2} \frac{\partial^2 \langle h(\lambda) \,|\, h(\lambda + \Delta\lambda) \rangle}{\partial \Delta\lambda_i \partial \Delta\lambda_j}\bigg|_{\Delta\lambda = 0} (\Delta\lambda_i)(\Delta\lambda_j)$$

where we can drop the linear term because the overlap is a maximum at $\Delta\lambda = 0$. When using frequency-domain, stationary-phase templates of the form 2.39, the calculation of this metric can be done analytically and gives a result in terms of moments of the noise curve

$$I(q) = S_n(f_0) \int_{f_0}^{f_c} \frac{x^{-q/3}}{S_n(x f_0)} \, dx$$

where f_0 is a chosen reference frequency.

The metric can then be used to efficiently lay out a grid of templates [35]. In general it is impossible to lay out a grid on a curved surface such that every point is the same distance from its nearest neighbors, but the metric on the template manifold is approximately flat and, when written in appropriate coordinates, the coefficients are constant.

4.4 The Search Pipeline

We now discuss how the above elements are incorporated into a gravitational-wave search. The system is collectively known as a *pipeline*, as it may be thought of as a series of steps through which the data flows. In particular the LIGO-Virgo search for gravitational waves from the coalescence of compact binaries is called the *ihope* pipeline. There is a large gap between the basic principles outlined above and a functioning gravitation-wave search. More details can be found in Refs. [30, 31, 36, 35]. We now present a very high-level overview of how these elements were used in the search for gravitational waves from the coalescence of compact binaries. We focus in particular on LIGO's sixth science run which overlaps Virgo's second and third. For more details on this search see Ref. [37].

In the first step the data from each detector the analysis is broken into 2048-second segments. This is done for computational efficiency, as well as to restrict to timespans over which the PSD is nearly stationary. Next, we filter to remove power below 40 Hz. This is done because computers are able to store numbers with only finite range and precision. In particular we are unable to small differences above large values without losing information. As the seismic noise below 40 Hz is orders of magnitude greater than the noise at higher frequencies signals at higher frequencies would be completely lost due to rounding errors.

The PSD for each detector is then calculated by a variation of *Welch's method* [30]. The 2048 seconds are split into 15 256-second chunks, overlapping by 128 seconds. The PSD of each chunk is then calculated using the defining Eqn. 4.1. The final PSD is obtained by taking the median values of each frequency bin. This is done in order to make the PSD estimation robust against loud, short events. As an extreme example, if there were a loud gravitational wave in the data the PSD would be

elevated, paradoxically suppressing the SNR.

This PSD is then used to construct the template bank [33, 34, **?**]. In S6/VSR2,3 we used stationary-phase templates of the form of Eqn. 2.39. We take the phase evolution, Ψ to 3.5 pN order, corresponding to an expansion in v to seventh order. See chapter 5 for more on this choice.

We then break the 2048-second segment into 256-second chunks (to match the resolution of the PSD) and filter the data with each template in the bank. The Fourier transformation algorithm we use treats the data as periodic, it identifies the times $t = 0$ and $t = 256$. This correlates the data at the beginning and end of the segment, we therefore discard the first and last 64 seconds in each segment, and overlap the segments so as to cover the entire time. See Refs. [30, 15] for more on this issue. The χ^2 test is not enabled at this stage, as it is computationally expensive.

The PSD used at this stage is not exactly the PSD that was calculated, it has been modified to avoid wrap-around issues and prevent a loud impulse from corrupting the entire chunk [30]. We return to this issue in Sec. 10.6.1, but for the moment recall that the response of the matched filter to an impulse can elevate the SNR for an extended time.

The outcome of this stage is a set of triggers for each detector. We then apply a *coincidence test* [36] by testing whether each trigger is close (in the sense of the metric Eqn. 4.3) to triggers from the other detectors. We allow the trigger time to differ by the maximum time it could take a gravitational wave to travel between detectors.

There is now a second stage of the process. Templates in each bank that have not produced coincident triggers are discarded. The filtering is then repeated, now with χ^2 enabled. There is then another coincidence test and the resulting triggers, now identified as *foreground triggers* are assigned a *combined new SNR* value

$$\rho_{\text{new, combined}}^2 = \sum_i \rho_{\text{new},i}^2$$

where i ranges over the set of detectors.

In order to determine how significant a foreground trigger is it must be compared against the background distribution. In pure Gaussian noise the probability distribution for cumulative new SNR could be calculated from Eqn. 4.14. As the real data is not Gaussian, this distribution must instead be measured. We do so by performing *time slides* **??**. The data from each detector is slid by more than the light-travel time

between detectors. Assuming at most one gravitational wave in the data, this ensures that any coincidences are entirely due to triggers produced by noise. We repeat this 100 times in order to build up more statistics. At the end we have a probability (or equivalently a rate) of obtaining triggers at each combined new SNR value. The foreground triggers are then compared to this distribution, those with values that are sufficiently rare/improbable are potential detection candidates, which are then subjected to extensive follow up analysis.

The order to test and tune the pipeline we perform numerous *software injections*, where simulated signals are added to the data before the first filtering step. Changes to the pipeline can be tested by examining the efficiency of the search, the number of injections recovered as a fraction of the number injected in a given region of parameter space.

The LIGO and Virgo collaborations also perform end-to-end tests with *hardware injections*, in which signals are added to the data by moving the ETMX mirror. Generally the times and the parameters of the injections are known, allowing complete end-to-end tests of the pipeline. However, there have been "blind" injection challenges, signals injected at ETMY by a small group within the collaboration which are not announced. We will have much more to say about one such injection in Sec. 10.5.

4.5 Conclusions

We have derived the matched filter as the optimal search for a known signal in Gaussian noise. We then proceeded to construct an entire gravitational-wave search pipeline using matched filtering as the core concept.

The construction of this pipeline involved making many choices: the template waveforms, the frequency at which to terminate the templates, the spacing of the bank, the region of parameter space to cover, the SNR threshold, and so on. It is possible and useful to test these choices against pN signals, as is done in software and hardware injections. However, it is known that these waveforms to not capture the full physics of the systems for which we are looking. There is therefore strong impetus to test the pipeline against signals from numerical relativity. This topic is the basis for the next four chapters.

Chapter 5

Comparison of Numerical Simulations of Black-Hole Binaries with Post-Newtonian Waveforms

5.1 Introduction

In chapter 4 we outlined the CBC pipeline used to search LIGO/Virgo data for gravitational waves resulting from the inspiral of binary systems composed of neutron stars and/or black holes. In particular we noted the use of post-Newtonian templates of the form discussed in Sec. 2.5 to filter the data. The post-Newtonian approximation is valid so long as the velocities of the component masses are small compared to the speed of light. As the matched filter is performed in the frequency domain, it is convenient to work with a *cutoff frequency*, a frequency above which we do not trust the post-Newtonian approximations. The velocity depends on both the total mass of the system and the orbital frequency as

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

so we therefore require that the cutoff frequency vary inversely with mass. In current searches the cutoff frequency is taken to be the orbital frequency of a point particle around a Schwarzschild black hole (that is, a black hole without charge or spin) at the innermost stable circular orbit (ISCO). The radius of this orbit is $r_{\mathrm{ISCO}} = 6M$ where M is the mass of the black hole, and the corresponding frequency is $f_{\mathrm{ISCO}} =$

$1/(6\sqrt{6}\pi M)$.

For systems containing two neutron stars the Schwarzschild ISCO frequency is above 1000 Hz, even for a hypothetical system composed of two instances of the heaviest-known neutron star (PSR J1614-2230, at $1.97 \pm 0.04 M_\odot$) [38]). This is well outside LIGO's most sensitive band where most of the SNR will be accumulated, and we may therefore trust pN waveforms in a search for binary neutron stars (BNS). However, for systems including at least one black hole the ISCO frequency will be lower and the transition to plunge and merger may occur in the sensitive band. It is therefore important to test the ability of pN templates to detect such systems and, where possible, optimize their ability to do so. In order to evaluate the effectiveness of templates it is necessary to compare them against complete waveforms including the inspiral, merger and ringdown phases. This necessitates the use of hybrid pN-NR waveforms as described in Sec. 2.6.1. This chapter uses a high-accuracy waveform developed at Caltech and Cornell [1].

The most robust test would consider the ability of the full pipeline to detect BBH signals; this is the goal of the NINJA project introduced in the following chapter. Here we focus on a more fundamental question, the ability of pN waveforms to capture full physical signals. We test this by calculating the overlap (defined in eqn. 4.20) between the hybrid waveform and the pN approximation. If there are no parameters for which the chosen pN model provides a high overlap with the hybrid signal then the pipeline as a whole will be unable to detect such signals. A similar study has been performed by Pan *et al.* using numerical data from Pretorius and the Goddard groups [39]. Our main results are in agreement with their conclusion that a simple extension of the TaylorF2 waveforms currently in use by the CBC low-mass search yields high overlaps with numerical waveforms.

Throughout this chapter we use only the $(l, m) = (2, 2)$ component of the waveform $\Psi_4^{2,2}$ (as defined, e.g., in Ref. [23]). For convenience, we drop the superscript. Whenever possible, we use dimensionless quantities, like $r\,M\,|\Psi_4|$, where r is the areal radius of the observation sphere, and M is the total apparent-horizon mass of the holes in the initial data. However, for any calculation involving the LIGO noise

---

[1] The code used to produce these waveforms, SpEC, is now being developed and used by a collaboration including scientists at Caltech, Cornell and the Canadian Institute for Theoretical Astrophysics (CITA).

curve, we have a physical scale, and thus use standard mks units.

To weight the inner products in the overlap we use the following PSDs for Initial and Advanced LIGO: for Initial LIGO we use an analytic approximation to the LIGO design PSD given by

$$
\begin{aligned}
S_n(f) \;=\; & 3.136 \times 10^{-4} \left[\left(\frac{4.49f}{150.0} \right)^{-56.0} \right. \\
& + \; 0.16 \left(\frac{f}{150} \right)^{-4.52} + \left(\frac{f}{150.0} \right)^2 + 0.52 \Bigg]
\end{aligned}
\tag{5.1}
$$

All integrals start from 40 Hz. As shown in Fig. 5, at this frequency the noise is an order of magnitude higher than its lowest value, and below this frequency it rises rapidly as $\sim f^{-56}$. The region below 40 Hz therefore contributes very little signal power to the SNR [40].

For Advanced LIGO we use the GWINC program [41] to generate the PSD. GWINC reports the PSD in increments of 0.0124 Hz. When calculating discrete integrals against signals sampled at other frequencies we obtain values for the PSD by linearly interpolating between the values provided by GWINC. We start integrals at 10 Hz as that is the point where the noise has increased by two orders of magnitude above its minimum, as also shown in Fig. 5.

5.2   PN–NR Hybrid Waveform

In order to perform our comparison we need to construct a "true" black-hole binary waveform, which we might expect to observe with detectors. A numerical simulation provides the data for the crucial nonlinear merger phase. We carefully extract the data and extrapolate it to large radius, and investigate the effects of numerical error on the final result. Because this waveform is very computationally expensive to produce, it covers only about 32 cycles, which is not sufficient for a thorough investigation of the possibility of detecting it in searches of data from gravitational-wave detectors. Thus, we match the numerical waveform to a post-Newtonian waveform, producing a hybrid which extends for many thousands of cycles, covering the entire band of interest.

5.2.1 Numerical Simulation, Extraction, and Extrapolation

The numerical simulation is the same as that described in Refs. [24, 42]: an equal-mass, non-spinning, black-hole binary with reduced eccentricity [43], beginning roughly 16 orbits before merger, continuing through merger and ringdown [42]. It is performed with the Caltech–Cornell pseudospectral code, using boundary conditions designed to prevent constraint violations and gravitational radiation from entering the domain [44, 45].

Data is extracted from the simulation in the form of the Newman–Penrose scalar

$$\Psi_4 = -C_{\alpha\beta\gamma\delta} l^{\alpha} \bar{m}^{\beta} l^{\gamma} \bar{m}^{\delta} \, , \tag{5.2}$$

where l^{α} and the complex vector $\bar{m}^{\beta}$ are constructed with reference to the coordinate basis. Along the positive z axis, we have

$$l^{\alpha} = \frac{1}{\sqrt{2}} \left(t^{\alpha} - z^{\alpha} \right) \, , \tag{5.3}$$

$$\bar{m}^{\beta} = \frac{1}{\sqrt{2}} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right)^{\beta} \, . \tag{5.4}$$

Here, t^{α} is the timelike unit normal to the spatial hypersurface, and z^{α} is the unit vector in the positive z direction. The vectors $\partial/\partial x$ and $\partial/\partial y$ are the standard coordinate vectors, which are not normalized. Ψ_4 is extracted as a function of time, at various radii along the positive z axis. This is then extrapolated to large radii, as described in Ref. [24], and in greater detail in Ref. [46].

The measured (instantaneous) frequency at the beginning of the simulation is

$$f_{\text{initial}} = (1.08 \pm 0.01) \times 10^3 \, \text{Hz} \, \frac{M_{\odot}}{M} \, . \tag{5.5}$$

The measured ringdown frequency is

$$f_{\text{ringdown}} = (1.78 \pm 0.02) \times 10^4 \, \text{Hz} \, \frac{M_{\odot}}{M} \, . \tag{5.6}$$

The measured *Christodoulou* mass and spin of the final black hole are

$$M_{\chi,\text{final}} = (0.95162 \pm 0.00002) \, M_{\chi,\text{initial}} \, , \tag{5.7}$$

$$S_{\text{final}} = (0.68646 \pm 0.00004) \, M_{\chi,\text{final}}^2 \, . \tag{5.8}$$

Using this value for the spin, a quasi-analytic formula due to Echeverria [47] predicts a value of $1.77 \times 10^4 \, \text{Hz} \, \frac{M_{\odot}}{M}$, for the ringdown frequency, in close agreement with the measured frequency.
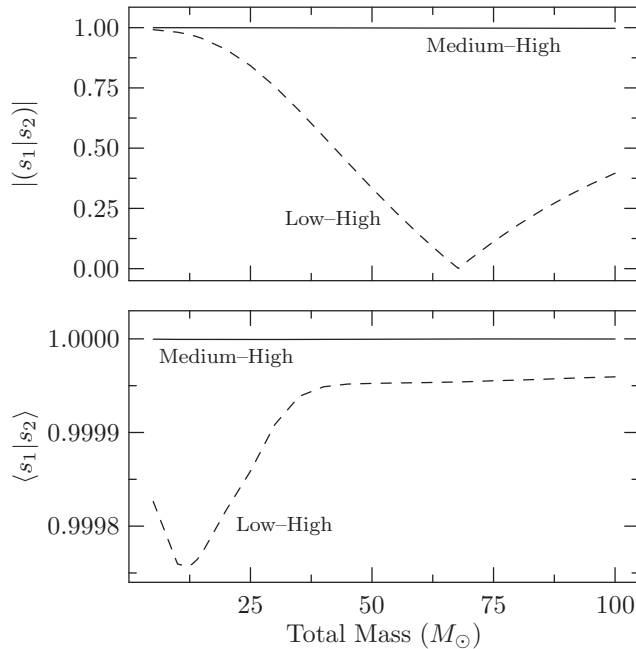
Figure 4 :   Convergence testing for numerical waveforms from a data-analysis perspective, using the match between waveforms computed at different numerical resolutions. The waveforms are scaled to various masses, and the Initial-LIGO noise curve is used in the calculation of the match. The upper panel shows the overlap without maximization over arrival time and phase; the lower panel shows the overlap after maximization. In each panel, the lower (dashed) line compares the lowest- and highest-resolution simulations, while the upper (solid) line compares the medium- and highest-resolution simulations. Note that this plot uses only numerical data, with no post-Newtonian contribution.

5.2.2   Accuracy of the Numerical Simulation

The numerical waveform will be the standard against which we will judge the *TaylorF2* waveforms used in LIGO data analysis (see Appendix A). To understand how precisely we should trust our final results, we need to understand the accuracy of the waveform itself. The most obvious measure of the error in this fiducial waveform is its convergence with increasing numerical resolution. Fig. 4 shows the overlap (Eq. (4.20)) between waveforms computed at different resolutions. The data used here are the extrapolated Ψ_4 waveforms, integrated in time twice.

Because of the short extent of the numerical waveforms, we need to be careful when using their Fourier transforms. The signal can be corrupted easily by the non-periodicity of the waveforms, and the discontinuous jumps that result. For Fig. 4

we mitigate this problem by increasing the sampling frequency of the input data, and restricting the Fourier transform to frequencies corresponding to instantaneous frequencies contained in the data. The input data can easily be upsampled in the time domain by interpolating the phase and amplitude of the complex data to a finer time grid. We then perform the transform, and explicitly set the data to zero at frequencies below f_{initial} and above f_{ringdown}, as given in Eqs. 5.5 and 5.6. While the results do depend on whether or not we impose these cutoffs, they do not depend sensitively on the actual cutoff frequencies.

The overlap between the lowest- and highest-resolution simulations (dashed lines) actually passes through zero, as shown in the upper panel. Presumably, this is because of loss of phase accuracy over the course of the simulation. All three simulations begin with the same initial data, so the waveforms are most similar at the beginning. Masses for which this is the most important segment (the lowest masses) will naturally have the highest overlap between resolutions. As the simulation progresses, numerical error accumulates—notably in the phase—so the overlap decreases with masses for which later segments dominate the overlap (higher masses). When the overlap is optimized over arrival time and phase (sec. 4.1) we can see that the overlap becomes much better, as shown in the lower panel, indicating sufficient accuracy within any frequency band for which phase coherence is required. In either case, the medium and highest resolutions are much more nearly the same. Without optimization, their overlap is within a few tenths of a percent of 1; after optimization, the overlap is within 10^{-6} of 1.

In the rest of our analysis we use the highest-resolution waveform. Because we always optimize over arrival time and phase, the lower panel of Fig. 4 is the most relevant, and shows that the waveform has converged to very high accuracy. The overlaps we quote below will only be given to three decimal places at most, because this is roughly the accuracy of the single-precision numerical methods used in the rest of the chapter. This accuracy is also sufficient for searches of gravitational-wave data. Thus, the truncation error of the simulated waveform is irrelevant for those purposes.

Other sources of error include residual eccentricity and spin, the influence of the outer boundary of the simulation, extrapolation errors, and coordinate effects, as discussed in Ref. [24]. The eccentricity had a disproportionately large effect on the error quoted in that paper because of the matching technique, which is not used

here. Restricting attention to the other effects of eccentricity, the uncertainty falls below that due to numerical error. Similarly, using the techniques of Ref. [48], the initial spins of the black holes have been measured more reliably, and found to be more than an order of magnitude smaller than previously determined, allowing us to reduce the estimate for that error to less than the numerical truncation error. The various coordinate effects were all estimated to be of roughly the same magnitude as the numerical error.

With the numerical error being many times more accurate than needed for this analysis, and the other sources of uncertainty being of roughly the same size, these considerations indicate that the overall error in our fiducial waveform is substantially less than the precision needed for this analysis.

5.3 Detection Efficiency of Gravitational-wave Templates

We now compare the signal described in the previous section to restricted, stationary phase *TaylorF2* post-Newtonian templates with terms up to order 2.0, order 3.5, and a "pseudo-4.0 pN-order" term recommended in Ref. [39]. Overlaps are calculated using the techniques of Sec. 4.1 with the signal s being the hybrid waveform described in the previous section, scaled to a range of masses. We consider both the Initial- and Advanced-LIGO noise curves.

Plots of the hybrid waveforms in comparison to the Initial-LIGO noise curve are shown in Fig. 5. The masses are chosen so that various frequencies of interest (the final stitching frequency, the ISCO, and the ringdown) occur at the "seismic wall" for Initial LIGO: 40 Hz. The waveforms $\tilde{s}$ are scaled to depict the detectability of the signal, typically quantified by the SNR introduced in (4.15), which may be written as

$$\rho^2 \equiv \int_0^\infty \frac{4\,\tilde{s}(f)\,\tilde{s}^*(f)}{S_n(f)}\,df = \int_0^\infty \frac{\left|2\,\tilde{s}(f)\,\sqrt{f}\right|^2}{S_n(f)}\,d\ln f \ . \tag{5.9}$$

In the final expression, the numerator and denominator have the same units, and are directly comparable. Because the square root of the denominator is familiar, we plot that along with the square root of the numerator. Plotting these two quantities together gives a graphical impression of the detectability of the waveform, and the relative importance of each part of the waveform, by its height above the noise curve. In Ref. [49], Brady and Creighton define a slightly different quantity, the characteristic
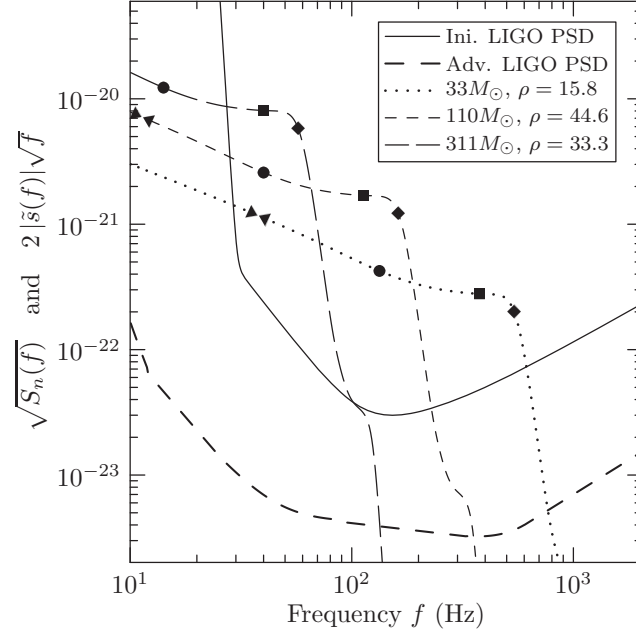
Figure 5 : Hybrid Caltech–Cornell waveform scaled to various total masses, with sources optimally oriented and placed at 100 Mpc, shown against the Initial- and Advanced-LIGO noise curves. Markers are placed along the lines at frequencies corresponding to various instantaneous frequencies of the waveforms. The triangles represent the beginning and end of the blending region; the circle represents the ISCO frequency; the square the light-ring; and the diamond the measured ringdown frequency. See the text for discussion of the normalization. The values given for ρ use the Initial-LIGO noise curve, with sources at a distance of 100 Mpc.

strain $h_{\text{char}} \equiv f \, |\tilde{s}(f)|$. The relative factor of $\sqrt{f}$ they use is present so that they can plot h_{char} against $\sqrt{f \, S_n(f)}$. Cutler and Thorne [50] define still another quantity, the signal strength $\tilde{h}_s(f)$, which is related to the Fourier transform by $\tilde{h}(f) = \sqrt{5} \frac{T}{N} \tilde{h}(s)$. The factor of $\sqrt{5}$ comes from averaging over the orientation of the binary, which we do not do. T/N is the ratio of the threshold to the rms noise at the endpoint of signal processing.

For each template family we initially optimize over signal mass M, symmetric mass ratio $\eta = m_1 m_2/(m_1 + m_2)^2$, and upper cutoff frequency f_c. The optimization is performed using a Nelder–Mead ("amoeba") algorithm [51]. The amoeba starts with a simplex in the parameter space and proceeds through a series of steps, each of which will improve the value of the function at at least one vertex. The algorithm terminates when all vertices have converged to the same point to within a specified tolerance. This process is deterministic, and amounts to an enhanced steepest-ascent algorithm. It is therefore only guaranteed to find a local maximum, and indeed we find that an amoeba instance started at a random point in the parameter space is most likely to converge to a point that does not give the highest possible overlap. We interpret this as being due to a large region in parameter space containing a local maximum and a relatively smaller region containing the global maximum. We therefore supplement the basic amoeba by running 300 instances with random starting values, and taking the best match obtained over all instances. In repeated runs the same optimal parameters were found by at least some of the amoebas, which supports the claim that this is the true maximum.

The results of optimizing over all of M, η and f_c for selected masses for Initial LIGO are given in Table 1 and summarized in Fig. 6. For Initial LIGO, in the range covered by the current CBC low-mass search ($M < 35M_\odot$) [52], the pseudo-4.0 pN *TaylorF2* waveforms achieve the highest overlaps, exceeding those obtained with 3.5 pN waveforms by $\sim 1\%$. Above $35M_\odot$ the 3.5 pN waveforms produce overlaps as much as 4% greater than those obtained with pseudo-4.0 pN waveforms over a range from 40–$80M_\odot$. With the Advanced-LIGO noise curve, in the CBC low-mass range, the 3.5 pN and pseudo-4.0 pN waveforms produce overlaps within 2% of each other, with 3.5 pN producing higher overlaps below 20 $M_\odot$ and pseudo-4.0 pN producing higher overlaps in the range 20–$35M_\odot$. Pseudo-4.0 pN continues to give the highest overlaps up to $60M_\odot$, producing overlaps as much as 4% greater than those obtained

	$(10+10)M_\odot$	$(20+20)M_\odot$	$(30+30)M_\odot$	$(50+50)M_\odot$
$\left\langle s^{\text{NR-CC}} \middle\| h^{\text{SPA}_c^{\text{ext}}(2.0)} \right\rangle$	0.99	0.98	0.97	0.96
$M/M_\odot$	$23.27^{+0.13}_{-0.12}$	$25.99^{+0.61}_{-0.56}$	$35.22^{+1.84}_{-1.89}$	$47.52^{+6.87}_{-4.73}$
η	$0.199^{+0.0030}_{-0.0030}$	$0.771^{+0.0490}_{-0.0420}$	$1.000_{-0.1390}$	$1.000_{-0.2490}$
f_{cut} (Hz)	$501.18^{+523.00}_{-153.00}$	$431.35^{+358.00}_{-77.00}$	$296.05^{+53.00}_{-31.00}$	$190.56^{+20.00}_{-14.00}$
$\left\langle s^{\text{NR-CC}} \middle\| h^{\text{SPA}_c^{\text{ext}}(3.5)} \right\rangle$	0.98	0.99	0.99	0.99
$M/M_\odot$	$18.75^{+0.10}_{-0.10}$	$31.88^{+0.77}_{-0.71}$	$47.15^{+4.37}_{-3.27}$	$259.89^{+0.00}_{-194.18}$
η	$0.290^{+0.0040}_{-0.0040}$	$0.493^{+0.0530}_{-0.0410}$	$0.756^{+0.2440}_{-0.2290}$	$0.954^{+0.0460}_{-0.2090}$
f_{cut} (Hz)	$506.50^{+518.00}_{-155.00}$	$448.80^{+576.00}_{-83.00}$	$324.74^{+145.00}_{-42.00}$	$197.17^{+24.00}_{-16.00}$
$\left\langle s^{\text{NR-CC}} \middle\| h^{\text{SPA}_c^{y}(4)} \right\rangle$	0.99	0.96	0.95	0.96
$M/M_\odot$	$23.64^{+0.13}_{-0.12}$	$47.90^{+1.28}_{-1.13}$	$61.81^{+8.68}_{-6.19}$	$89.93^{+20.44}_{-16.60}$
η	$0.182^{+0.0030}_{-0.0030}$	$0.181^{+0.0160}_{-0.0140}$	$0.523^{+0.4260}_{-0.1820}$	$0.529^{+0.4720}_{-0.3100}$
f_{cut} (Hz)	$509.47^{+654.00}_{-145.00}$	$352.44^{+73.00}_{-61.00}$	$309.53^{+72.00}_{-47.00}$	$195.63^{+21.00}_{-15.00}$

Table 1 : Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN templates using the Initial-LIGO noise curve. The first number in each block is the overlap; subsequent numbers are the template parameters that achieve this overlap. Parameter values within the specified ranges keep the overlap within 1% of the maximum by varying that parameter, while leaving others fixed. We restrict the search to $0 \leq \eta \leq 1.000$, so the upper error bounds when $\eta \sim 1.000$ may be artificially small.

with 3.5 pN waveforms. Above $60 M_\odot$ 3.5 pN waveforms again yield the best overlaps, by as much as 6% around 90 $M_\odot$.

A significant feature of Tables 1 and 2 is the size of the error bars on the cutoff frequencies. For $M = 20 M_\odot$ the cutoff frequency can vary as much as 128% above and 28% below the optimal value while losing no more than 1% of overlap. This leads us to consider the range of possible template parameters which may give high overlaps. In the next section we consider the reduction in overlap as the parameters f_c and η are independently varied from the optimal value.

5.3.1 Effect of Upper Frequency Cutoff

As shown in Fig. 5 the amplitude of the NR waveforms drops sharply at around the lightring frequency, which depends on the total mass of the binary. The *TaylorF2*

	$(10+10)M_\odot$	$(20+20)M_\odot$	$(30+30)M_\odot$	$(50+50)M_\odot$
$\left\langle s^{\text{NR-CC}}\middle|h^{\text{SPA}_c^{\text{ext}}(2.0)}\right\rangle$	0.98	0.92	0.91	0.94
$M/M_\odot$	$25.15^{+0.02}_{-0.02}$	$47.73^{+0.12}_{-0.11}$	$54.39^{+0.51}_{-0.43}$	$60.19^{+1.55}_{-1.29}$
η	$0.170^{+0.0010}_{-0.0010}$	$0.188^{+0.0010}_{-0.0010}$	$0.335^{+0.0080}_{-0.0070}$	$0.891^{+0.0660}_{-0.0490}$
f_{cut} (Hz)	$444.77^{+132.00}_{-115.00}$	$267.64^{+48.00}_{-50.00}$	$262.44^{+34.00}_{-36.00}$	$182.41^{+24.00}_{-18.00}$
$\left\langle s^{\text{NR-CC}}\middle|h^{\text{SPA}_c^{\text{ext}}(3.5)}\right\rangle$	0.97	0.92	0.92	0.96
$M/M_\odot$	$20.27^{+0.02}_{-0.02}$	$38.11^{+0.11}_{-0.09}$	$50.09^{+0.49}_{-0.42}$	$78.10^{+1.89}_{-1.50}$
η	$0.245^{+0.0010}_{-0.0010}$	$0.277^{+0.0020}_{-0.0020}$	$0.386^{+0.0130}_{-0.0100}$	$0.494^{+0.0760}_{-0.0330}$
f_{cut} (Hz)	$355.85^{+97.00}_{-88.00}$	$262.83^{+47.00}_{-48.00}$	$281.34^{+41.00}_{-37.00}$	$186.31^{+30.00}_{-19.00}$
$\left\langle s^{\text{NR-CC}}\middle|h^{\text{SPA}_c^{y}(4)}\right\rangle$	0.97	0.96	0.94	0.90
$M/M_\odot$	$22.24^{+0.02}_{-0.02}$	$46.57^{+0.11}_{-0.11}$	$72.06^{+0.35}_{-0.35}$	$118.50^{+1.99}_{-1.63}$
η	$0.208^{+0.0010}_{-0.0010}$	$0.190^{+0.0010}_{-0.0010}$	$0.177^{+0.0020}_{-0.0030}$	$0.186^{+0.0100}_{-0.0070}$
f_{cut} (Hz)	$473.49^{+551.00}_{-136.00}$	$353.18^{+73.00}_{-69.00}$	$242.43^{+37.00}_{-36.00}$	$152.16^{+19.00}_{-19.00}$

Table 2 : Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN templates using the Advanced-LIGO noise curve. The first number in each block is the overlap; subsequent numbers are the template parameters that achieve this overlap. Parameter values within the specified ranges keep the overlap within 1% of the maximum by varying that parameter, while leaving others fixed. We restrict the search to $0 \leq \eta \leq 1.000$, so the upper error bounds when $\eta \sim 1.000$ may be artificially small.
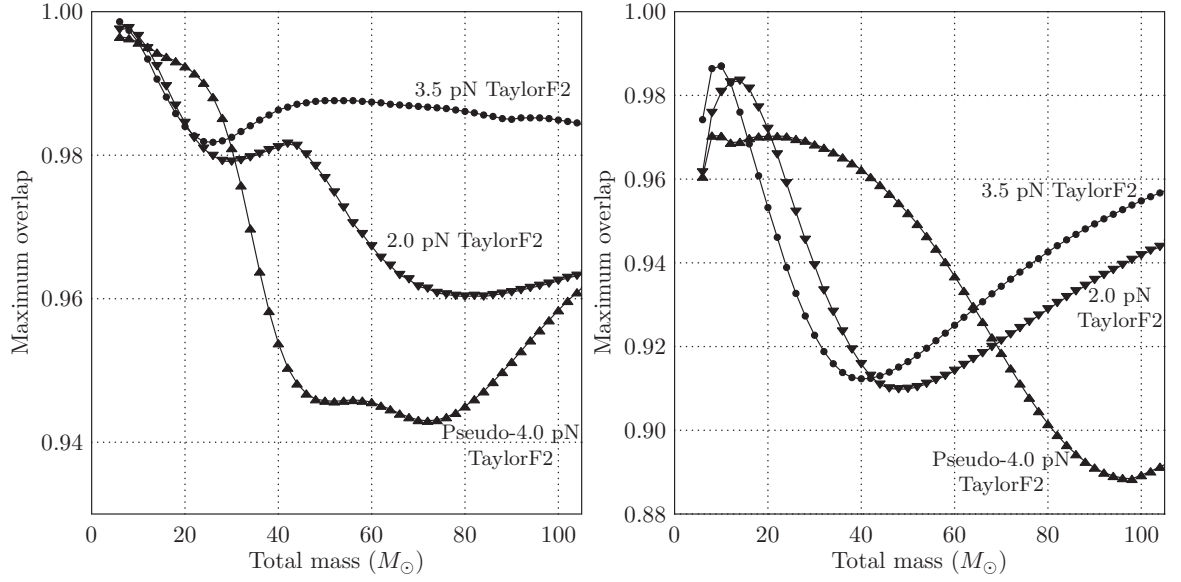
Figure 6 : Left: Overlaps between Caltech–Cornell hybrid waveforms, scaled to various masses, and restricted stationary-phase pN waveforms for Initial-LIGO PSD. Optimization is over M and η, which the cutoff frequency f_c is prescribed by the weighted average described below. The mass ratio η is allowed to range over unphysical values. The best-fit values found for the pseudo-4.0 pN templates are always physical in this case. See Sec. 5.3.2. Right: The same, for the Advanced-LIGO PSD
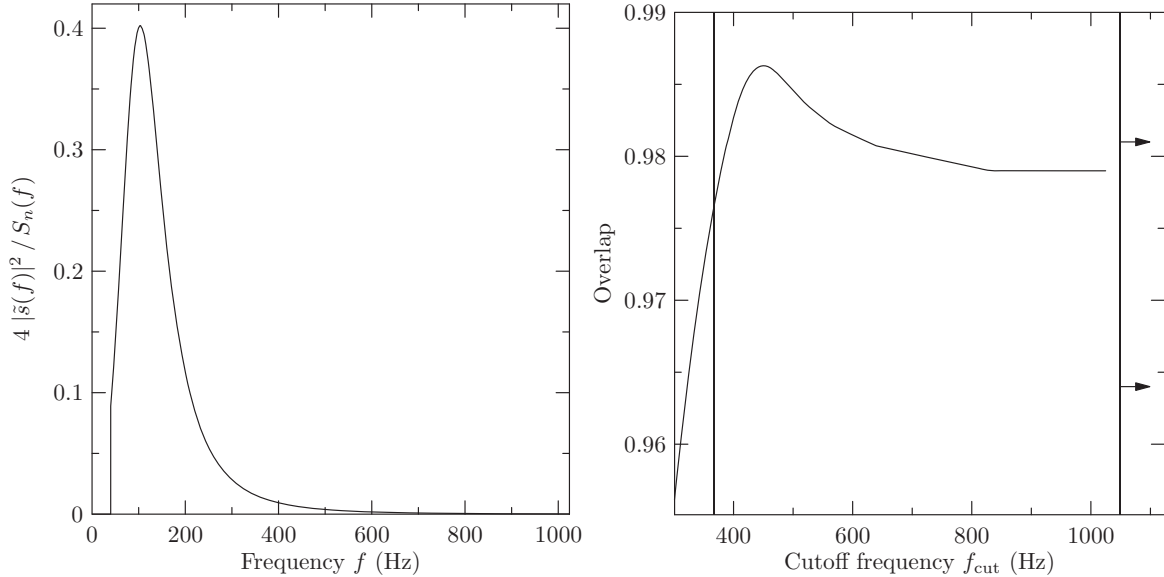
Figure 7 : Left: Integrand of Eq. (4.13) for a *TaylorF2*, 3.5 pN waveform with $M = 10$ and $\eta = 0.25$, at a distance of $100\,\mathrm{Mpc}$, using the Initial-LIGO noise curve. Note that the shape of this curve does not change as we change M and η; only the vertical scale changes. Right: Overlap between Caltech–Cornell waveform scaled to $M = 40\,M_\odot$ and restricted *TaylorF2*, 3.5 pN waveform using the best-match values for M and η, as a function of the cutoff frequency f_c, with the Initial-LIGO noise curve. The vertical bars are meant to delineate 1% loss. Note that the upper bound extends to higher frequencies indefinitely.

waveforms do not model the late inspiral, merger or ringdown and hence will continue to evolve as $f^{-7/6}$ at all frequencies, increasingly deviating from the NR waveform. This suggests that the upper frequency cutoff of the *TaylorF2* waveform should be chosen to be below the frequency at which the two diverge. However, the effect of the divergence is mitigated by the PSD. The denominator of the overlap, Eq. (4.20), depends on $(s|s)$ which is a constant, and $(h|h)$ which would increase without limit if not for the PSD. Fig. 7 shows $|\tilde{h}(f)|^2/S_n(f)$—the integrand of $(h|h)$—for the Initial-LIGO noise curve for an example *TaylorF2* waveform for an equal-mass $10\,M_\odot$ binary. We see that above about 450 Hz there is very little contribution to the integrand, and so extending the cutoff frequency above this will not impact the overlap.

The numerator of the overlap, $(s|h)$, can only increase as the cutoff frequency is raised, however frequencies above the lightring where the waveforms have diverged will contribute very little. The effect of including higher frequencies on the overlap is

therefore determined by the $(h|h)$ term in the denominator. For systems with ringdown frequencies well above the peak of the integrand in Fig. 7, this term will not significantly reduce the overlap. For example, binaries of total mass roughly $40\,M_\odot$ have ringdown frequencies at roughly $450\,\mathrm{Hz}$. Only a small fraction of the SNR comes from higher frequencies. Thus, we expect that systems with lower masses should not suffer great loss in overlap if the cutoff frequency is higher than ringdown. However for higher-mass systems the overlap can be significantly reduced if the upper frequency cutoff is too large. This is indeed what we find, as shown by a representative example on the right in Fig. 7. For this $40\,M_\odot$ system, using the Initial-LIGO noise curve, the optimal cutoff frequency is around $450\,\mathrm{Hz}$—roughly the ringdown frequency. Decreasing the cutoff quickly decreases the overlap. The cutoff may be increased almost indefinitely, however, with only 0.5% loss in overlap. This, of course, changes when using the Advanced-LIGO noise curve. We revisit this issue in Sec. 5.4.

5.3.2   Unrestricted η

The physical symmetric mass ratio is restricted to the range $0 < \eta \le 0.25$, values above this imply complex-valued masses. However the pN waveforms are well-behaved for $0 < \eta < 1.0$, and as seen from Tables 1 and 2, the highest overlaps are often obtained at unphysical values of η. In Fig. 8 we show the effect of limiting the optimization to physical η. At high masses, the limitation reduces the optimal overlap by up to 12%. *TaylorF2* waveforms with $\eta \le 1/4$ would not be expected to accurately model the late-inspiral and merger part of the waveform, as non-Newtonian effects are increasingly significant in this region. We find that allowing unphysical η broadens the space of waveforms covered by the *TaylorF2* approximation sufficiently to capture more of the late-inspiral and merger.

5.4   Recommendations for Improvements

Based on the analysis of the previous sections we propose a series of adjustments to searches using *TaylorF2* template waveforms to enhance the efficiency of those searches. First, as seen in Fig. 6 for Initial LIGO, adding terms up to 3.5 pN order produces overlaps as large or larger than the current 2.0 pN templates over most of the mass range, while the pseudo-4.0 pN templates recommended in Ref. [39] produce
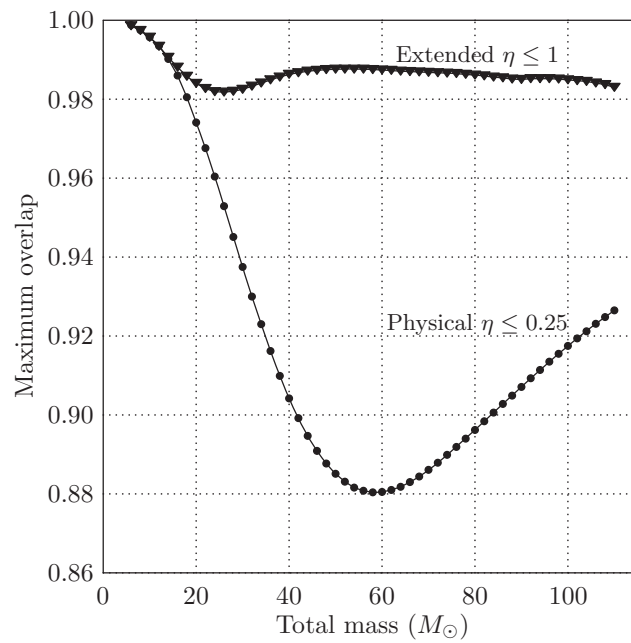
Figure 8 : Maximum overlaps obtained by allowing η to range over unphysical values, compared to those obtained by restricting the range of η. These overlaps are generated using 3.5 pN TaylorF2 templates, searching over values of the total mass and mass ratio. Extending to unphysical values of η improves the match by up to 11%.

slightly larger overlaps at masses near $20\,M_\odot$. Thus, we recommend pseudo-4.0 pN templates for the low mass range, $M < 35M_\odot$, and 3.5 pN templates for higher masses. The improvement due to 3.5 pN templates over 2.0 pN generally holds for Advanced LIGO as well. The 3.5 pN templates produce larger overlaps than 2.0 pN templates above $50\,M_\odot$ without a significant loss (within 1%) at lower masses. However, there is a large region for which the pseudo-4.0 pN term does significantly better. When using an Advanced-LIGO noise curve, we recommend 3.5 pN templates generally, 2.0 pN templates in the range 12–$21\,M_\odot$ and pseudo-4.0 pN templates for masses in the range 21–$65\,M_\odot$.

As a second improvement, we note from Fig. 8 that allowing η to range over unphysical values significantly improves matches with 3.5 pN templates above $30\,M_\odot$. In preliminary studies we have found that extending to $\eta \leq 1$ roughly doubles the size of the template bank, and the advantages must therefore be weighed against the increase in false alarm rate.

Our third recommendation involves the cutoff frequency used for the template waveform. Optimization over the cutoff frequency is too computationally intensive to be done in searches. Currently, the cutoff frequency is typically taken to be the Schwarzschild ISCO frequency. To examine the effect of this choice we vary f_c while keeping the mass and η at their optimal values, for each of the signal masses in our range. The result of one such variation is shown in Fig. 7 (right). Figs. 9 shows the variations for all masses, highlighting the regions within which the overlap drops by less than 1% (dark gray) and 3% (light gray) of the optimal value. This figure also shows the ISCO and ERD frequencies, neither of which stays within the 1% band for both Initial and Advanced LIGO. In particular, the ISCO is a poor choice for both Initial and Advanced LIGO except at very low masses, where the precise value of the cutoff is almost irrelevant.

The ISCO is often pointed to—somewhat arbitrarily—as a good estimate of the breakdown of post-Newtonian approximations [53]. So, for instance, if we were to match a pN template to a physical waveform, beginning at some point in the distant past, we might expect them to separate quite badly near the ISCO. Of course, for realistic black-hole binaries, the gravitational waves will only enter the LIGO band late in the inspiral—just before the ISCO for low-mass systems, or after the ISCO for high-mass systems. We can see from Fig. 5 that, for masses below about $30\,M_\odot$, the
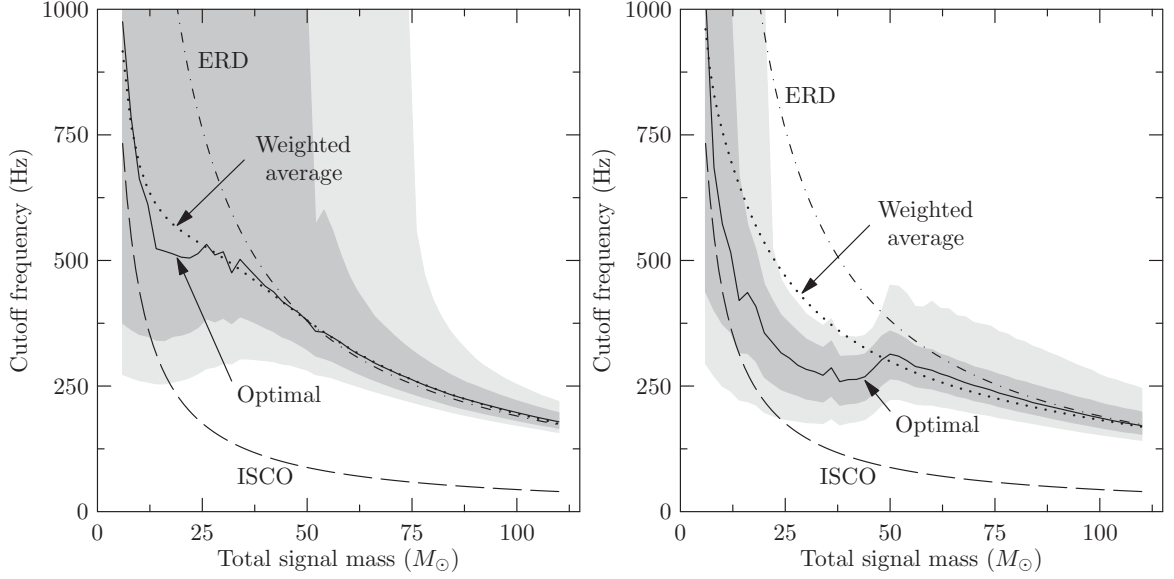
Figure 9 : Left: Candidate f_c values for 3.5 pN templates with Initial LIGO. The dark gray band contains cutoff frequencies with matches within 1% of the value at which the best overlap was obtained. The light gray band contains frequencies with matches within 3%. Right: Candidate f_c values for 3.5 pN templates with Advanced LIGO. The dark gray band contains cutoff frequencies with matches within 1% of the value at which the best overlap was obtained. The light gray band contains frequencies with matches within 3%. Note that the weighted-average cutoff extends past the 1% error bars for $12 < M/M_\odot < 40$. However, in that same region, the 3.5 pN templates do poorly overall, and we recommend pseudo-4.0 pN templates. The optimal cutoff frequency for pseudo-4.0 pN templates is much closer to the weighted-average cutoff in this mass range.

ISCO is high enough that lower-frequency parts of the waveform contribute the most to the SNR. For very high masses, however, this basically cuts the waveform down to nothing. In Initial LIGO, the ISCO is completely buried in seismic noise for masses above about $100\,M_\odot$. Thus, we must move the cutoff frequency up. We cannot push the cutoff far above ringdown, because the physical waveform simply ceases to exist (see Fig. 5). It has been suggested that an "effective ringdown" (ERD) frequency $f_{\mathrm{ERD}} \equiv 1.07\,f_{\mathrm{Ringdown}}$ is a useful upper limit [39]. For intermediate masses, we would like to interpolate somehow between these two extremes of ISCO and ERD. We suggest setting the cutoff frequency to a weighted average of the two, where the weights are the contributions to the SNR below the given frequency. If we assume coherent phasing between the template and the physical waveform, we can simply take the amplitudes of the two waveforms. Also, note that the restricted SPA approximation for the amplitude is reasonable. Thus, define

$$\rho_{\mathrm{ISCO}}^2 \;\equiv\; \int_0^{f_{\mathrm{ISCO}}} \frac{f^{-7/3}}{S_n(f)}\, df \;, \tag{5.10}$$

$$\rho_{\mathrm{ERD}}^2 \;\equiv\; \int_{f_{\mathrm{ISCO}}}^{f_{\mathrm{ERD}}} \frac{f^{-7/3}}{S_n(f)}\, df \;, \tag{5.11}$$

$$\rho_{\mathrm{tot}}^2 \;\equiv\; \int_0^{f_{\mathrm{ERD}}} \frac{f^{-7/3}}{S_n(f)}\, df \;, \tag{5.12}$$

$$f_{\mathrm{cut}} \;\equiv\; \frac{f_{\mathrm{ISCO}}\,\rho_{\mathrm{ISCO}} + f_{\mathrm{ERD}}\,\rho_{\mathrm{ERD}}}{\rho_{\mathrm{tot}}} \;. \tag{5.13}$$

We have already dropped constant factors in the expressions for ρ that will cancel out.

Note that these expressions only depend on the total mass by way of the limits of integrations—which are very simple, known functions of the mass—so these integrals could be done just once for a given noise curve, storing the intermediate values. When the cutoff needs to be calculated, the cumulative integral could be evaluated at the given ISCO and ringdown frequencies. Hence, this would be a fast way of calculating the cutoff, with no need to do the integrals each time the cutoff is needed.

We can test this recommended frequency by comparing it to the optimal cutoff frequency found by the amoeba search described in Sec. 5.3. For 3.5 pN templates in Initial LIGO, we find that it is an excellent match to the optimal frequency. Fig. 9 shows these two values, along with dark and light bands showing the regions in which changing f_c results in a loss of overlap of 1% and 3%, respectively. Of course, the

same figure shows that using the ERD recommendation would stay within the 1% error bounds. Nonetheless, the close match between this recommendation and the true optimum suggests that it is sound. Thus, our final recommendation is to use the weighted-average frequency cutoff throughout the entire mass range. While our analysis has been restricted to equal-mass systems, the cutoff frequency we have defined here could be applied to unequal-mass systems as well. It will be interesting to see how this cutoff fares in those situations.

Similar results hold for Advanced LIGO, when using our recommended template for each mass. That is, in regions where 3.5 pN templates do poorly (see Fig. 6), the weighted average is a poor predictor of the optimal cutoff frequency using those templates, as shown in Fig. 9. However, in those same regions—where pseudo-4.0 pN templates do well—the weighted average is a good predictor of the optimal cutoff frequency for 4.0 pN templates. Thus, again, we recommend using the weighted-average frequency cutoff throughout the entire mass range with Advanced LIGO.

By prescribing a cutoff frequency, the search does not need to extend over that parameter. Similarly, by prescribing a post-Newtonian order, we need use only one template for a given total mass. On the other hand, if these recommendations decrease the overlap found by too much when using them compared to the overlap found by an unconstrained search, it may be better to search the larger parameter space. We can evaluate the loss in overlap by comparing the results found using our recommendations to the results found when searching over the set of all three template families, and all masses, mass ratios, and cutoff frequencies. We have determined that this loss in overlap when using our recommendations is always less than 0.0025 for Initial LIGO, and less than 0.007 for Advanced LIGO.

5.5 Conclusions

We have compared high-accuracy NR waveforms for equal-mass binary black holes from the Caltech–Cornell group to stationary phase post-Newtonian waveforms. We examined a number of factors that influence the matches between the two, with the goal of optimizing the matches and hence improving the efficiency of templated searches in Initial and Advanced LIGO. We first considered the effect of the post-Newtonian order to which the phase evolution is taken, and found that adding terms

up to 3.5 pN or pseudo-4.0 pN to the currently-used 2.0 pN templates significantly improves the matches over a large range of masses, as shown in Fig. 6. We then studied the effect of varying the upper cutoff frequency of the templates. The frequency that achieves the optimal match is a function of mass, and we find this function is well-approximated by an average between ISCO and ERD, weighted by contribution to the SNR, as shown in Fig. 9. Finally, we allow the symmetric mass ratio η to range over unphysical values up to 1.0, and find that this dramatically improves matches, as shown in Fig. 8. Based on the results we recommend adjusting the searches using *TaylorF2* template waveforms by going up to 3.5 pN or 4.0 pN over most of the mass range, integrating up to our recommended cutoff, and allowing allowing η to extend up to 1. For Initial LIGO, the overlaps obtained using these parameters is always within 0.0025 of overlaps achievable by optimizing over all three parameters.

In future work we plan to extend this analysis to unequal-mass and spinning black-hole systems. We have found that allowing unphysical values of η roughly doubles the size of the template bank, and we also plan to study the impact of this change on the rate of spurious triggers from glitches.

Chapter 6

The First NINJA Project

Thus far, most searches for gravitational waves from BBH mergers have relied on post-Newtonian waveforms such as those discussed in Sec. 2.5, which are valid when the black holes are sufficiently far apart. Within its range of validity, post-Newtonian theory provides a convenient analytic description of the expected signals produced by binary systems. The numerical relativity results, on the other hand, have not yet been synthesised into an analytic model for the merger phase covering a broad range of parameters, i.e., a wide range of mass ratios, spins and if necessary, eccentricity; there has however been significant progress for the non-spinning case [54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65]. Similarly, despite significant progress, there is not yet a complete detailed description over the full parameter space of how post-Newtonian and numerical simulations are to be matched with each other. On the data analysis side, many pipelines, especially ones that rely on a detailed model for the signal waveform, have made a number of choices based on post-Newtonian results, such as the use of the Schwarzschild ISCO as a cutoff frequency, and it is important to verify that these choices are sufficiently robust. More generally, it is necessary to quantify the performance of these data analysis pipelines for both detection and parameter estimation. This is critical for setting astrophysical upper limits in case no detection has been made, for following up interesting detection candidates, and of course for interpreting direct detections. Numerical relativity now provides an important avenue for extending this beyond the early inspiral phase captured by post-Newtonian waveforms, to the late inspiral and merger phase.

There are significant challenges to be overcome before numerical relativity results

can be fully exploited in data-analysis pipelines. The Numerical INJection Analysis (NINJA) project was started in the spring of 2008 with the aim of addressing these challenges and fostering close collaboration between numerical relativists and data analysts. Participation in NINJA is open to all scientists interested in numerical simulations and gravitational-wave data analysis.

Several decisions were make to limit the scope of the project. NINJA chose to restrict attention to BBH simulations and have not used results from supernova simulations or simulations containing neutron stars [1]. For the first NINJA project (henceforth "NINJA-1") the waveform data came purely from numerical simulations and we did not attempt to extend numerical data using post-Newtonian waveforms. The NINJA-1 data set was constructed using Gaussian noise to model the response of the Initial LIGO and Virgo detectors — no attempt was made to include non-Gaussian noise transients found in real detector data. The comparisons and conclusions reported here are thus necessarily limited, and in many cases are only the first steps towards fully understanding the sensitivity of data-analysis pipelines to black hole signals. Further studies are needed regarding the accuracy and comparison of numerical waveforms, and of how systematic errors in these waveforms can affect parameter estimation. Some analyses of numerical waveforms with regard to gravitational-wave detection have already been performed [66, 67, 57, 65], accuracy standards have been developed for use of numerical waveforms in data analysis [68] and a detailed comparison of some of the waveforms used in NINJA-1 was performed in the related Samurai project [69]. The NINJA-2 project, discussed in the next chapter and currently ongoing at the time of writing, will build on these results to begin to address these issues.

Despite the limited scope of the NINJA-1, we are able to draw the following broad conclusion from this work. We conclude that the current data analysis pipelines used to search LIGO, Virgo and GEO600 data for black hole coalescence are able to detect numerical waveforms injected into the NINJA-1 data set at the expected sensitivities. Indeed, the standard pipeline is able to detect signals that lie outside the parameter space that they target. This is a non-trivial statement since most detectability estimates to date for these sources have relied on post-Newtonian waveforms, which are valid only when the black holes are sufficiently far apart. It should be noted,

---

[1]As of this writing a matter NINJA project is in an early planning stage

however, that the NINJA data set does not contain non-stationary noise transients so more work is needed to understand how detection performance is affected by the noise artifacts seen in real gravitational-wave detector data. NINJA has proven to be extremely valuable at framing the questions that need to be answered.

6.1 Numerical Waveforms

NINJA-1 studied BBH coalescence waveforms submitted by ten individuals and teams. Participation in NINJA-1 was open to anyone and the only restrictions were that each contribution: (i) was a numerical solution of the full Einstein equations, (ii) consisted of only two waveforms, or up to five waveforms if they were part of a one-parameter family.

No restrictions were placed on the accuracy of each waveform. All contributions followed the format specified in [70]. The waveforms are plotted in Figures 10 and 11. The contributed waveforms covered a variety of physical and numerical parameters. Most simulations modeled low-eccentricity inspiral, the mass ratio $q = m_1/m_2$ ranges from 1 to 4, and the simulations covered a range of spin configurations. The initial angular frequency of the $\ell = m = 2$ mode ranged from $0.033/M$ to $0.203/M$ (where M denotes the sum of the initial black-hole masses). This initial angular frequency marks where contributors considered the waveform sufficiently clean to represent the physical system (e.g. this will be chosen after initial unphysical radiation content, often referred to as "junk radiation" in numerical relativity, is radiated away). The length of the waveforms varied between a few 100M to over 4000M. The contributions naturally differed in accuracy, both regarding how well they captured the black-hole dynamics and in the extraction of the gravitational-wave signal.

Table 4 lists a few key parameters that distinguish the waveforms, and introduces the following tags for the different contributions and codes: BAM HHB [71, 72, 73, 74, 75] and BAM FAU [71, 72, 76, 75] are contributions using the `BAM` code, `CCATIE` is the AEI/LSU code [77, 78, 79, 80, 81], `Hahndol` is the Goddard Space Flight Center's code [82, 83], `LazEv` is the RIT code [84, 11, 85], `Lean` is Ulrich Sperhake's code [86, 87, 88], `MayaKranc` is the Georgia Tech/Penn State code [67, 89], PU stands for the Princeton University code [90, 10, 54, 91], `SpEC` for the Cornell/Caltech collaboration code [92, 93, 59, 22], and `UIUC` stands for the University of Illinois at

Urbana-Champaign team [94].

The codes listed above use different formulations of the Einstein equations, gauge conditions, mesh structures, initial data and wave extraction methods. Full details of each code are given in the references.

6.2   Construction of the NINJA Data Set

The data provided by the numerical relativity groups follows the format outlined in [70], which is based on the mode decomposition of the gravitational radiation field at large distances from the source. If we specify a gravitational waveform $h_{\mu\nu}$ in the Transverse-Traceless (TT) gauge, we only need the spatial components h_{ij}. We assume that we are sufficiently far away from the source so that the $1/r$ piece dominates:

$$h_{ij} = A_{ij}\frac{M}{r} + \mathcal{O}\left(r^{-2}\right), \tag{6.1}$$

where M is the total mass of the system, r is the distance from the source, and A_{ij} is a time-dependent TT tensor. In the TT gauge, h_{ij} has two independent polarisations denoted h_+ and $h_\times$ and the complex function $h_+ - ih_\times$ can be decomposed into modes using spin-weighted spherical harmonics $^{-2}Y_{lm}$ of weight -2:

$$h_+ - ih_\times = \frac{M}{r}\sum_{\ell=2}^{\infty}\sum_{m=-\ell}^{\ell} H_{\ell m}(t) \, ^{-2}Y_{\ell m}(\iota, \phi). \tag{6.2}$$

The expansion parameters H_{lm} are complex functions of the retarded time $t - r$, however if we fix r to be the radius of the sphere at which we extract waves then H_{lm} are functions of t only. The angles ι and ϕ are respectively the polar and azimuthal angles in a suitable coordinate system centred on the source. This decomposition is directly applicable to non-precessing binaries. Otherwise, a comparison of the waveforms requires a careful treatment of mode-mixing effects due to rotations of the frame; see for instance [103]. The numerical data contributed to NINJA is given in the form of an ASCII data file for each mode (ℓ, m), with accompanying metadata describing the simulation [70]. Only modes that contribute appreciably to the final waveform are included, at the discretion of the contributing group. Each data file consists of three columns: time in units of the total mass, and the real and imaginary parts of the mode coefficients $H_{\ell m}$ as a function of time. Note that the

Code Ref.	Run	q	$\vec{S}_i/m_i^2$	e	$\omega_{22}\,M$	D/M	eccentricity removal
BAM FAU [71, 72]	[76]	1	see caption	qc	0.06	$9.58\,\hat{y}$	T-PN [95, 96]
BAM HHB [71, 72]	S00 [73]	1	0	< 0.002	0.045	$12\,\hat{y}$	TR-PN [97]
	S25 [74]	1	$0.25\,\hat{z}$	≈ 0.006	0.045	$12\,\hat{y}$	T-PN [98]
	S50 [74]	1	$0.50\,\hat{z}$	≈ 0.006	0.052	$11\,\hat{y}$	$-\,"\,-$
	S75 [74]	1	$0.75\,\hat{z}$	≈ 0.006	0.06	$10\,\hat{y}$	$-\,"\,-$
	S85 [74]	1	$0.85\,\hat{z}$	≈ 0.006	0.06	$10\,\hat{y}$	$-\,"\,-$
CCATIE [77, 78, 79, 80]	r0 [80]	1	$0.6\,\hat{z},\ -0.6\,\hat{z}$	qc	0.079	$8\,\hat{x}$	TR-PN [97]
	r2 [80]	1	$0.6\,\hat{z},\ -0.3\,\hat{z}$	qc	0.078	$8\,\hat{x}$	$-\,"\,-$
	r4 [80]	1	$0.6\,\hat{z},\ 0$	qc	0.076	$8\,\hat{x}$	$-\,"\,-$
	r6 [80]	1	$0.6\,\hat{z},\ 0.3\,\hat{z}$	qc	0.075	$8\,\hat{x}$	$-\,"\,-$
	s6 [81]	1	$0.6\,\hat{z}$	qc	0.074	$8\,\hat{x}$	$-\,"\,-$

Table 3 : Initial conditions for numerical waveforms. The columns list: the name of the contribution, the name of the run if appropriate, the mass ratio $q = m_1/m_2$ where $m_1 \geq m_2$, the spins of the black holes (if only one spin is given, both spins are equal), an estimate of the initial eccentricity of the orbit ("qc" denotes cases where quasi-circular inspiral is modelled, but a value of the eccentricity has not been reported), the initial frequency of the $(2, 2)$ mode (rounded to three digits), the initial coordinate separation of either the black-hole punctures or the excision surfaces, and where appropriate the method of eccentricity removal. All binaries start in the xy-plane with momenta tangent to the xy-plane. See text for the identification of each contribution, and a description of the notation in the last column. The dimensionless spins of the BAM FAU run are $(-0.634, -0.223, 0.333)$ and $(-0.517, -0.542, 0.034)$.

Code Ref.	Run	q	$\vec{S}_i/m_i^2$	e	$\omega_{22} M$	D/M	eccentricity removal
`Hahndol`	kick	3	$0.2\,\hat{x}$, $0.022\,\hat{x}$	qc	0.078	$8.007\,\hat{y}$	T-PN [99]
[82, 83]	non	4	0	qc	0.070	$8.470\,\hat{y}$	$-$ ” $-$
`LazEv` [84, 11]	MH [85]	1	$0.92\,\hat{z}$	qc	0.07	$8.16\,\hat{x}$	T-PN [99, 100]
`Lean` [86]	c	4	0	qc	0.05	$10.93\,\hat{x}$	T-PN [71]
	2	1	$0.926\,\hat{z}$	qc	0.11	$6.02\,\hat{x}$	T-PN [99]
`MayaKranc`	e0 [89]	1	0	qc	0.05	$12\,\hat{x}$	TR-PN [97]
[67]	e02 [89]	1	0	0.2	0.05	$15.26\,\hat{x}$	n/a
PU [90, 10]	CP [54]	1	$0.063\,\hat{z}$	qc	0.07	$9.5\,\hat{x}$	T-ID [101]
	T52W [91]	1	0	≥ 0.5	0.07		n/a
`SpEC` [92]	q=1 [59, 22]	1	0	5×10^{-5}	0.033	$15\,\hat{x}$	TR-it [93]
UIUC [94]	cp [94]	1	0	qc	0.194	$4.790\,\hat{x}$	T-ID [101]
	punc [94]	1	0	qc	0.203	$4.369\,\hat{y}$	T-ID [102]

Table 4 : Initial conditions for numerical waveforms. The columns list: the name of the contribution, the name of the run if appropriate, the mass ratio $q = m_1/m_2$ where $m_1 \geq m_2$, the spins of the black holes (if only one spin is given, both spins are equal), an estimate of the initial eccentricity of the orbit ("qc" denotes cases where quasi-circular inspiral is modelled, but a value of the eccentricity has not been reported), the initial frequency of the $(2, 2)$ mode (rounded to three digits), the initial coordinate separation of either the black-hole punctures or the excision surfaces, and where appropriate the method of eccentricity removal. All binaries start in the xy-plane with momenta tangent to the xy-plane. See text for the identification of each contribution, and a description of the notation in the last column.

Code	Run	q	ΔT_{100} [s]	$f_{i,100}$ [Hz]	$M_{30Hz}[M_\odot]$
BAM HHB	S00	1	1.03	15	48
	S25	1	1.15	15	48
	S50	1	1.03	17	56
	S75	1	0.81	19	65
	S85	1	0.87	19	65
BAM FAU		1	0.54	19	65
CCATIE	r0	1	0.34	26	85
	r2	1	0.37	25	84
	r4	1	0.40	25	82
	r6	1	0.45	24	81
	s6	1	0.59	24	80
Hahndol	kick	3	0.25	25	84
	non	4	0.32	23	75
LazEv	MH	1	0.43	23	75
Lean	c	4	0.92	16	54
	2	1	0.20	36	118
MayaKranc	e0	1	1.23	16	54
	e02	1	0.74	16	54
PU	CP	1	0.29	23	75
	T52W	1	0.16	23	75
SpEC	q=1	1	1.96	11	36
UIUC	cp	1	0.10	63	209
	punc	1	0.10	66	219

Table 5 : Characteristic duration, mass and frequencies of the waveforms summarised in table 4. The columns ΔT_{100} and $f_{i,100}$ give the duration and initial frequency of the waveform when scaled to total mass $M = 100 M_\odot$. M_{30Hz} is the total mass of the waveform when it is scaled so that the initial frequency is 30Hz (this sets the lowest mass at which each waveform can be injected into the NINJA data).
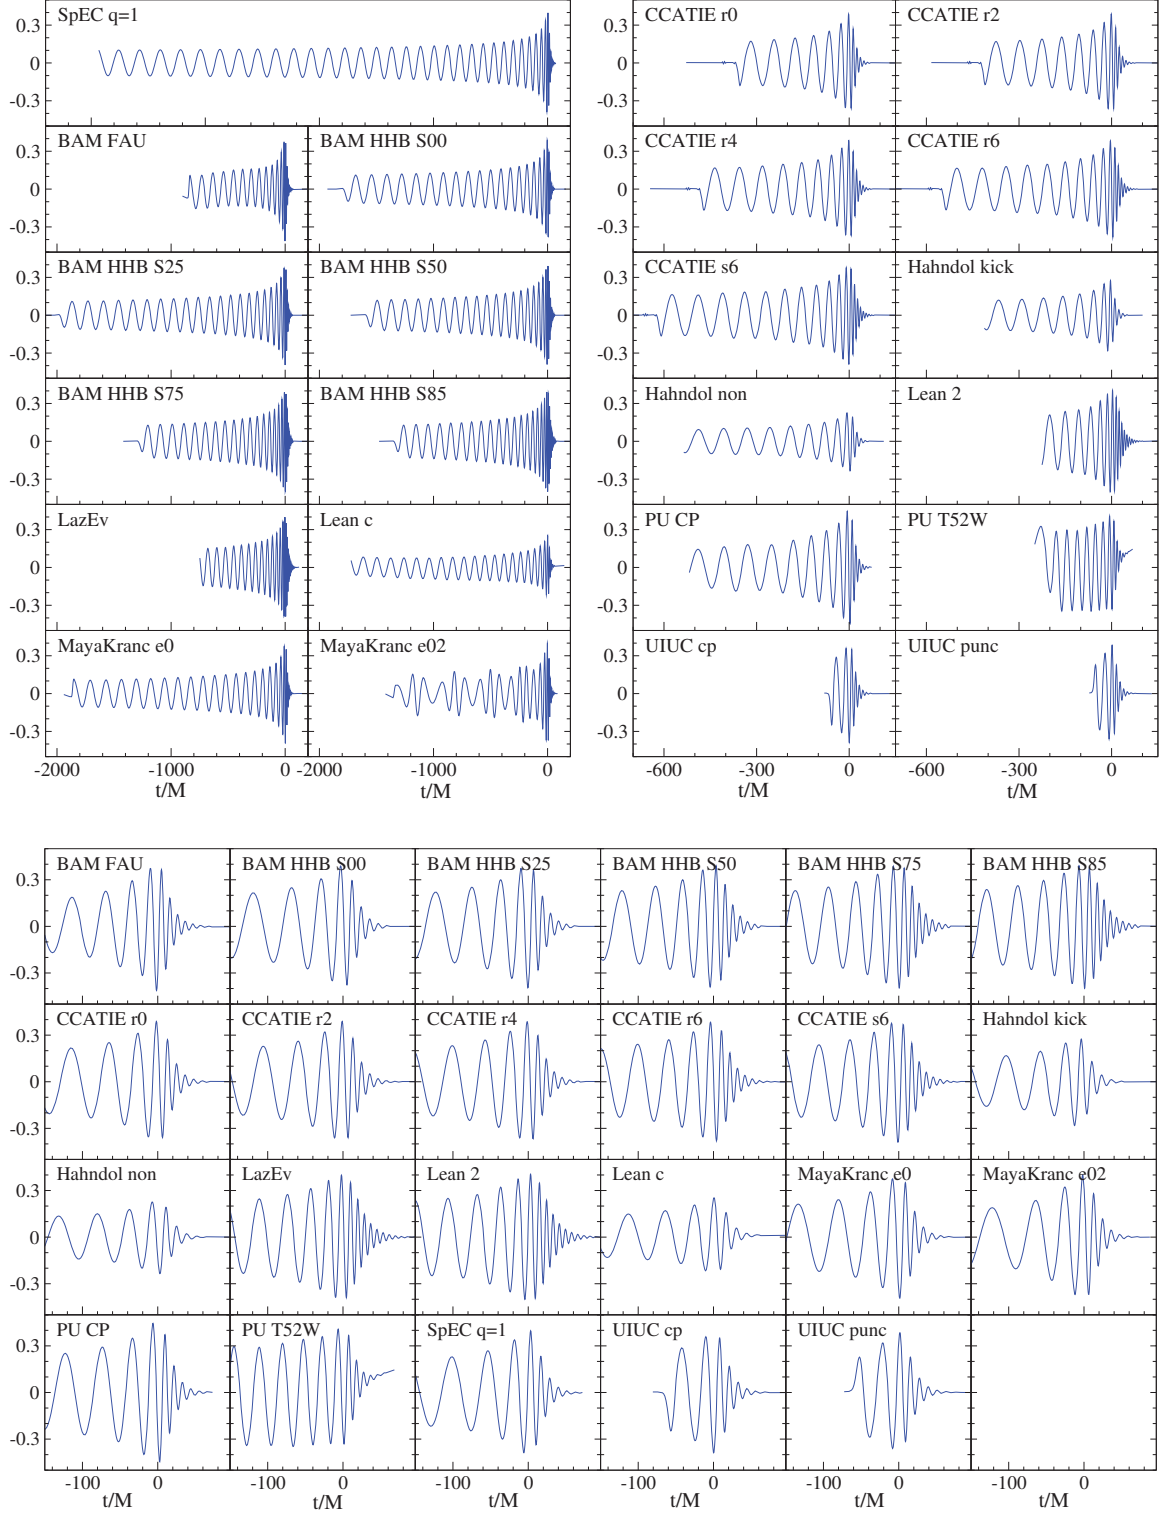
Figure 10 : Summary of all submitted numerical waveforms: $r/M\,\mathrm{Re}(h_{22})$ The x-axis shows time in units of M and the y-axis shows the real part of the $(\ell, m) = (2, 2)$ component of the dimensionless wave strain $rh = rh_+ - irh_\times$. The top panels show the complete waveforms: the top-left panel includes waveforms that last more than about $700M$, and the top-right panel includes waveforms shorter than about $700M$. The bottom panel shows an enlargement of the merger phase for all waveforms.
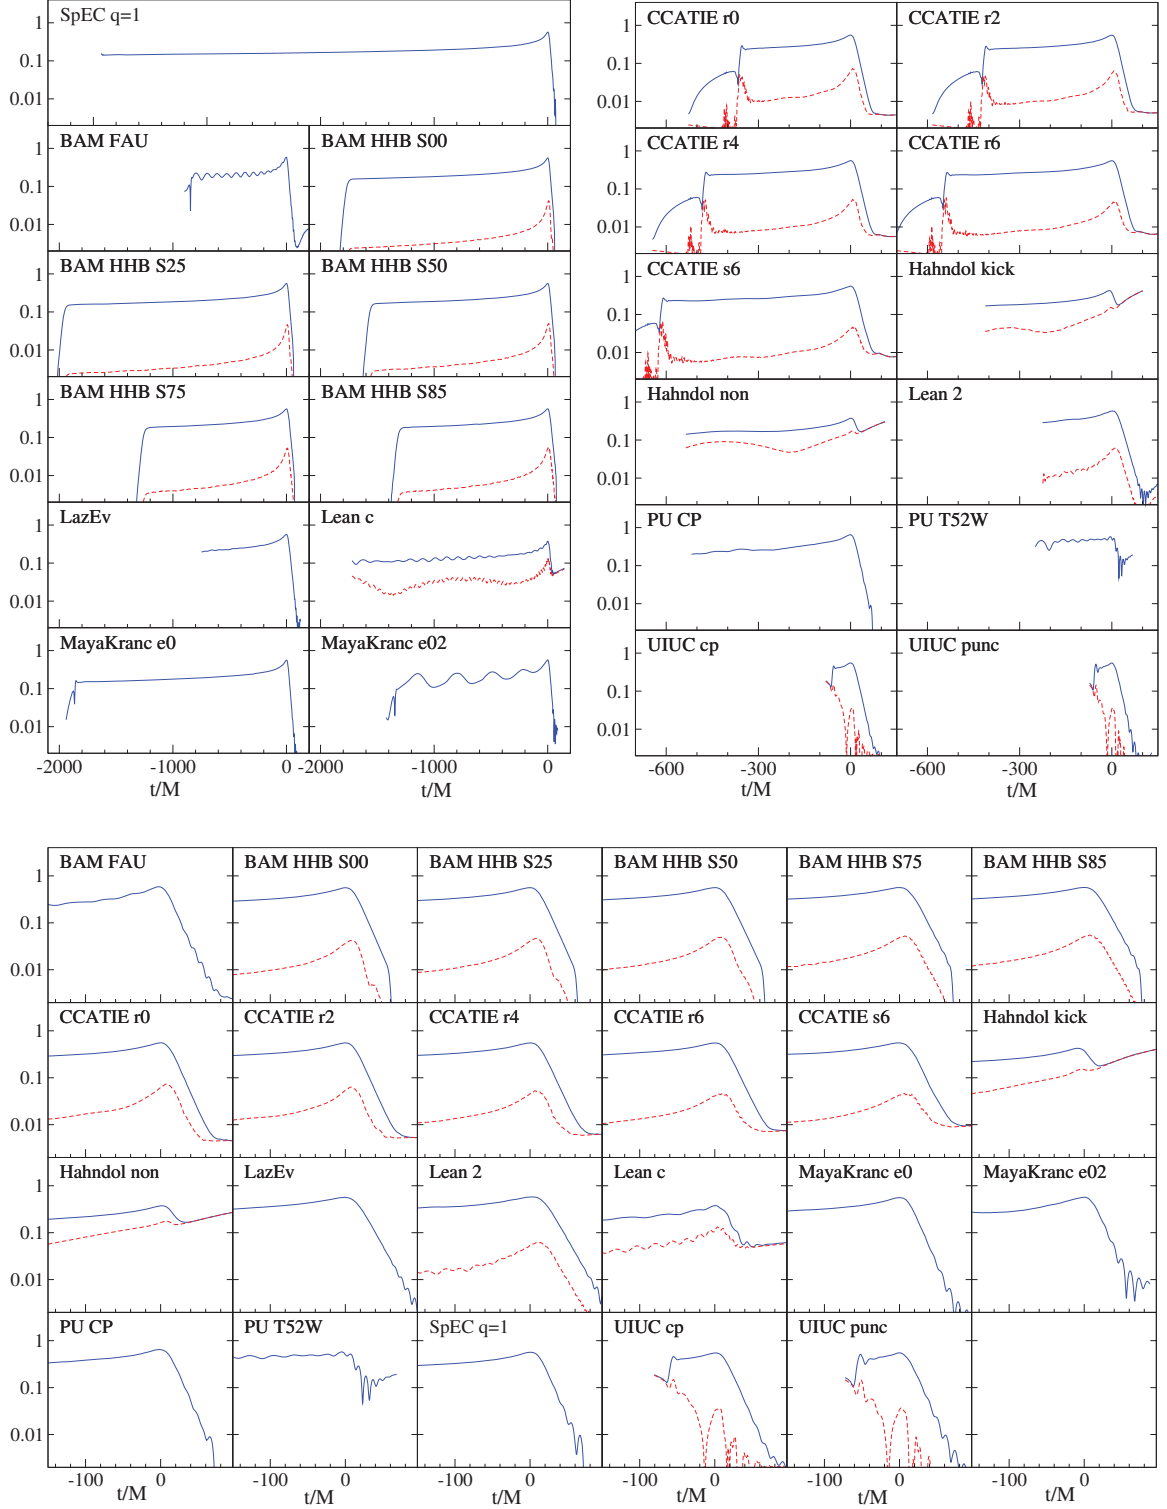
Figure 11 : Distribution of power into different spherical harmonics. The blue line shows $\left(\Sigma_{\ell,m}|h_{\ell m}\,r/M|^2\right)^{1/2}$. A dashed red line, if present, shows the same sum, but *excluding* the $(\ell,m) = (2,\pm2)$ modes. The separation between the two lines gives the relative importance of non $(2,\pm2)$ modes. If no red line is present for a certain run, then only the $(2,\pm2)$ modes were supplied. The layout is as in Fig. 10: The top panels show the complete waveforms, whereas the bottom panel shows an enlargement of the merger phase. The x-axis shows time in units of M.

total mass M scales both the time and the amplitude; thus the BBH waveforms for each simulation can be scaled to an arbitrary value of the mass. (This is not true in the case of simulations which include matter fields, but we do not consider such waveforms here.)

To model the signal seen by a gravitational-wave detector, we need to calculate the detector strain $h(t)$ from the above mode decomposition. To do this, we must choose particular values of the total mass, orientation and distance from the detector. Given the $H_{\ell m}$, the total mass, the distance to the source, and the angles (ι, ϕ), we calculate $h_{+,\times}$ using Equation. (6.2), and use the detector response functions $F_{+,\times}$ (see, for example, Ref. [9]) to calculate the observed strain

$$h(t) = h_+(t)F_+(\alpha, \delta, \psi) + h_\times(t)F_\times(\alpha, \delta, \psi). \qquad (6.3)$$

Here (α, δ) are sky-angles in the detector frame, ψ is the polarisation angle and the time t is measured in seconds. In this analysis, we wish to simulate signals that might be observed by the Initial LIGO and Virgo detectors. There are three LIGO detectors: a 4 km detector and a 2 km detector at the LIGO Hanford Observatory (called H1 and H2, respectively) and a 4 km detector at the LIGO Livingston Observatory (called L1). The Virgo detector is a 3 km detector in Cascina, Italy (called V1). We used the same two-letter codes for the simulated NINJA detectors. Since the location and alignment of the three observatories differ, we must use the appropriate detector response and arrival time to compute the strain waveform $h(t)$ seen at each observatory. This ensures that the waveforms are coherent between the detectors and simulate a true signal.

To model the detector noise, we generated independent Gaussian noise time series $n(t)$, sampled at 4096 Hz, for each detector. This sample rate was chosen to mimic that used in LIGO Scientific Collaboration (LSC)-Virgo searches and assures a tolerable loss in signal-to-noise ratio due to the discrete time steps [2] Stationary white noise time series are generated and coloured by a number of time-domain filters designed to mimic the design response of each of the LIGO and Virgo detectors. Figure 12 shows the one-sided amplitude spectral density $\sqrt{S_n(f)}$ of each time detector's time series.

---

[2]More careful study done in NINJA-2 revealed 4096 Hz to be insufficient, see chapter 7. However, we do not expect these issues to have significant effects in NINJA-1
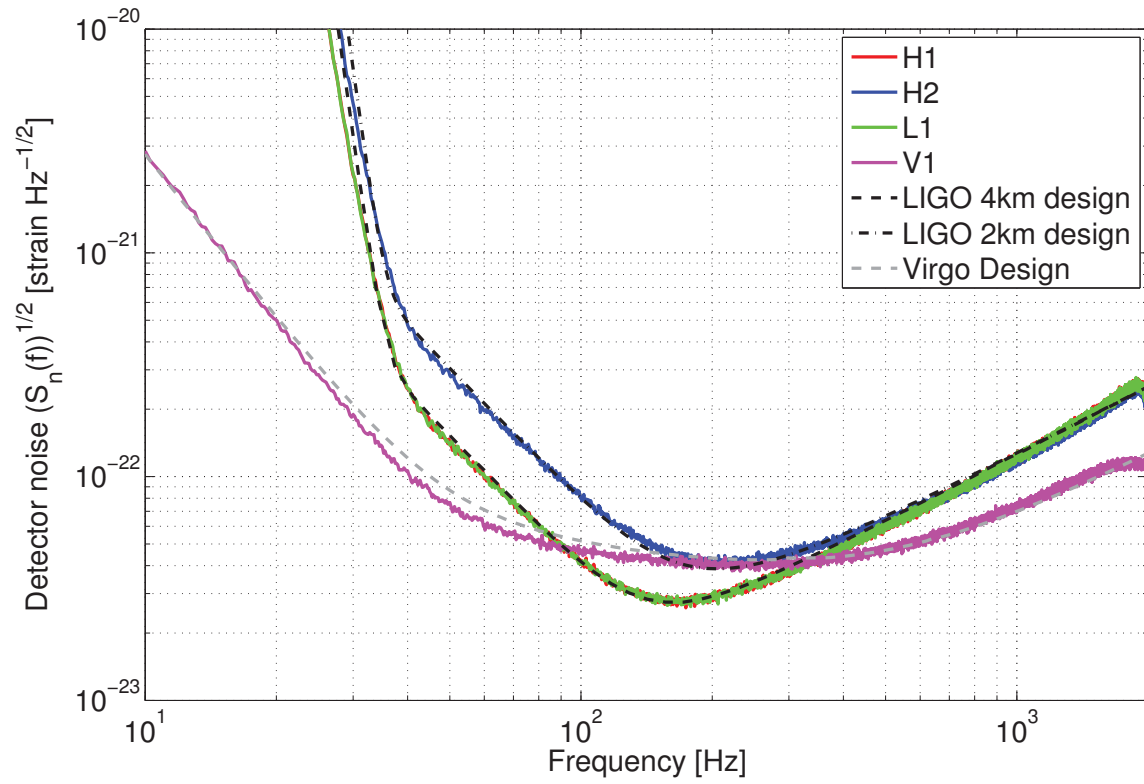
Figure 12 : The NINJA data noise curves and the design spectra of the first generation LIGO and Virgo detectors.

We see from Figure 12 that the noise power spectrum of the NINJA data set closely approximates the Initial LIGO design sensitivity in the frequency range of interest (30-10^3 Hz). There is a slight discrepancy with the Virgo design curve at low frequencies (between approximately 20 and 150 Hz), which is an artefact of the Virgo noise generation procedure. Narrow-band features such as the violin and mirror modes were removed from the detector response used to compute the NINJA data, but were included in the calculation of the Virgo design curve [104]. The $1/f$ tails of these narrow-band features are responsible for the small discrepancy.

Having generated the simulated detector data, we then generated a population of simulated signals using the numerical relativity data. This population was constructed to cover a broad range of masses and signal amplitudes. We required that the starting frequency of the dominant $\ell = m = 2$ mode of the signal was not more than 30 Hz, an appropriate threshold given the sensitivity curve of the Initial LIGO and Virgo detectors. This sets a minimum mass at which each waveform can be injected, which is given in Table 5. The minimum possible injection mass is therefore $36 M_\odot$. The maximum mass was chosen as $350 M_\odot$. To get a good sample of long injected waveforms, we systematically chose a lower range of masses for the longer waveforms. No restrictions were placed on the other simulation parameters, i.e., the spins, mass-ratios and eccentricities. We ensured that waveforms from all the participating groups were equitably represented by generating approximately 12 signals from the waveforms supplied by each group. The time interval between adjacent injected signals was chosen to be a random number in the range 700 ± 100 s.

Given these constraints, we generated the parameters of the signal population. The logarithm of the distance to the binary was drawn from a uniform distribution ranging from 50 Mpc to 500 Mpc, and the source locations and orientations were drawn from an isotropic distribution of angles. We then computed waveforms corresponding to this population and at the appropriate sampling rate. We required that the optimal matched filter signal-to-noise ratio of any injection be greater than five in at least one of the four simulated detectors. Any waveform that did not satisfy this constraint was discarded from the population. Subject to this condition, the distances of injected signals varied from 52 Mpc to 480 Mpc (median at 145 Mpc), the injected total mass range was $36 M_\odot \leq M \leq 346 M_\odot$ (median at $155 M_\odot$), with individual component masses in the range $11 M_\odot \leq m_i \leq 193 M_\odot$.
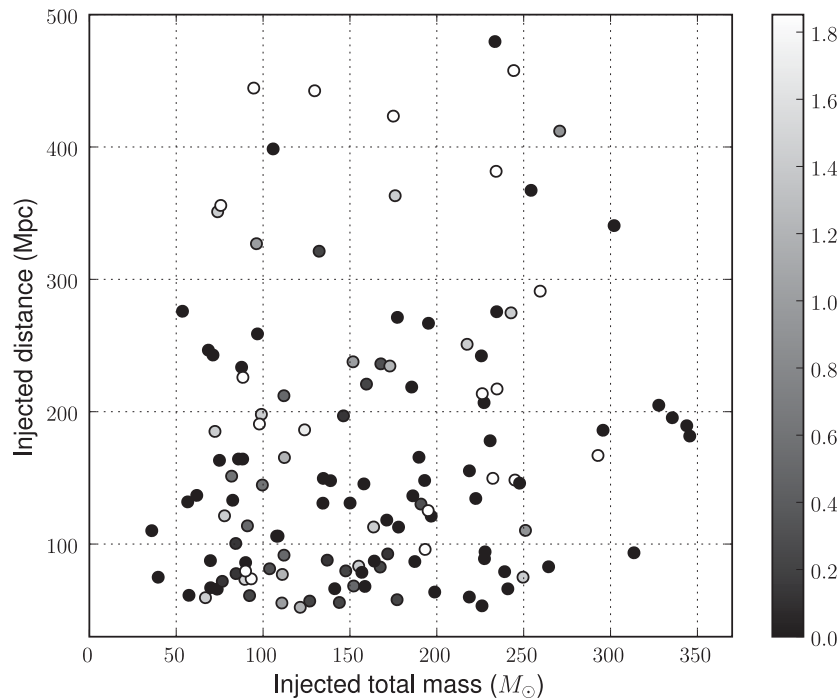
Figure 13 : The total mass and distance of the 126 NINJA injections. The grey scale encodes the sum of the dimensionless spins of the black holes, $|\vec{S_1}/m_1^2 + \vec{S_2}/m_2^2|$.

Finally, the waveforms $h(t)$ were added to the simulated detector noise $n(t)$ to generate the NINJA data set $s(t) = n(t) + h(t)$. As described above, care was taken to ensure that signals were coherently injected in the data streams from the four detectors. The software for carrying out this procedure is freely available as part of the LSC Algorithm Library (LAL) [105].

The data set used in this analysis consisted of a total of 126 signals injected in a total of 106 contiguous segments of noise each 1024 s long, thus spanning a duration of a little over 30 hours. Figure 13 shows the mass, spin and distance of the waveforms contained in the NINJA data set.

Group	Analysis
AEI	Phenomenological Waveforms in CBC pipeline
Birmingham	Bayesian Model Selection
Cardiff	Post-Newtonian (PN) Templates in CBC pipeline
Cardiff, Maryland	EOBNR waveforms in CBC pipeline
Goddard	Hilbert Huang Transform
Northwestern	Markov Chain Monte Carlo
Syracuse	Extended η PN Templates in CBC pipeline
UMass, Urbino	Q-pipeline analysis
UWM	PN templates in CBC pipeline, Neyman-Pearson criteria
UWM, UMass, Urbino	Ringdown analysis
UWM, UMass, Urbino	Inspiral, Merger, Ringdown combined search

Table 6 : The data-analysis contributions to the NINJA-1 project.

6.3   Data Analysis Results

Analysis of the NINJA-1 data was open to all and nine groups submitted contributions using a variety of analysis techniques. Participating groups were provided with the NINJA-1 data set containing signals embedded in noise and the parameters of the injected signals. Analysts were not given access to the raw numerical-relativity waveforms or noiseless injection data.

Methods used to analyse the NINJA-1 data include: matched-filter based searches, un-modelled waveform searches using excess-power techniques, and Bayesian model-selection and parameter-estimation techniques. Where possible, the performance of different searches was compared. The limited scope of the NINJA-1 data set makes detailed comparisons difficult, however. A list of the data-analysis contributions is shown in Table 6.

6.3.1   LIGO-only Searches

Henceforth we restrict attention to variations of the CBC pipeline, described in chapter 4. In particular, although many variations of this pipeline were tested in NINJA-1, we focus on testing the modifications to the template pN order, bank construction,

and termination frequency of the template waveforms suggested by the comparison studies in chapter 5. In this sense the NINJA-1 project can be seen as an evolution of those studies; having found a set of changes to the template waveforms that improve overlaps with numeric signals, the next step is to test these changes in searches. The results of these searches are summarised in Table 7, each column giving the results from a different search with a summary of the chosen parameters. We first describe the parameters varied between these analyses and then present a more detailed discussion of the results.

All NINJA analyses using *TaylorF2* waveforms (see Appendix A) used restricted templates (i.e. the amplitude is calculated to leading order), however the phase was calculated to various different post-Newtonian orders [106]. Phases were computed to either two [107, 108] or three point five post-Newtonian order [109, 110, 111] since these are, respectively, the order used in LSC-Virgo searches [112] as of the time of NINJA-1, and the highest order at which post-Newtonian corrections are known. The studies in chapter 5 show that 3.5 pN waveforms provide better overlaps with numeric waveforms than 2.0 pN.

After choosing a post-Newtonian order, one chooses a region of mass-parameter space to cover with the template bank. Figure 14 shows the boundaries of the template banks used in the analyses. One search used the range used by the LSC-Virgo "low-mass" search [112] ($m_1, m_2 \geq 1M_\odot, M \leq 35M_\odot$) and all other searches used templates with total masses in the range $20M_\odot \leq M \leq 90M_\odot$. These boundaries were chosen since there were no signals in the NINJA data with mass smaller than $36M_\odot$ and there is little, if any, inspiral power in the sensitive band of the NINJA data for signals with $M \gtrsim 100M_\odot$.

The standard LSC-Virgo template bank generation code [113] restricts template generation to signals with $\eta \leq 0.25$, since it is not possible to invert M and η to obtain real-valued component masses for $\eta > 0.25$. All but one of the searches enforced this constraint, with the $0.03 \leq \eta \leq 0.25$ for the low-mass CBC search and $0.1 \leq \eta \leq 0.25$ for the other "physical-η" searches. However, the comparison studies show that the overlaps obtain maximum values at unphysical values of η over much of the mass space. In particular, with 3.5 pN waveforms unphysical η values provide better overlaps above $40M_\odot$, which is the mass range covered by the NINJA-1 injections. However, these studies allowed allowed the values to vary continuously

and did not take into account the discretization imposed by a template bank, nor the parameter coincidence required at the second stage of the pipeline. It is therefore critical to test the extended η bank in searches.

Finally, it is necessary to specify a frequency at which to terminate the TaylorF2 waveform. In the LSC-Virgo analyses, this is chosen to be the innermost stable circular orbit (ISCO) frequency for a test mass in a Schwarzschild spacetime

$$f_{\text{ISCO}} = \frac{c^3}{6\sqrt{6}\pi GM}.$$ (6.4)

This cutoff was chosen as the point beyond which the TaylorF2 waveforms diverge significantly from the true evolution of the binary [106]. However, the studies reported in chapter 5 as well as those in [57] have shown that extending the waveforms up to higher frequencies improves the sensitivity of TaylorF2 templates to higher mass signals. The NINJA-1 TaylorF2 analyses use templates terminated at the ISCO frequency and two additional cut-off frequencies: the effective ringdown (ERD) frequency and a weighted ringdown ending (WRD) frequency. The ERD frequency was obtained by comparing post-Newtonian models to the Pretorius and Goddard waveforms [57]. The ERD almost coincides with the fundamental quasi-normal mode frequency of the black hole formed by the merger of an equal-mass non-spinning black-hole binary. The weighted ringdown ending (WRD) frequency is lies between ISCO and ERD, and found to close to optimal in chapter 5. It is calculated as in Eqn. 5.13

Analysis	(1)	(2)	(3)	(4)	(5)	(6)
Freq. Cutoff	ISCO	ISCO	ERD	ERD	WRD	WRD
PN Order	2 PN	2 PN	2 PN	3.5 PN	3.5 PN	3.5 PN
Total Mass $M_\odot$	2–35	20–90	20–90	20–90	20–90	20–90
η **range**	0.03–0.25	0.10–0.25	0.10–0.25	0.10–0.25	0.10–0.25	0.10–1
Found Single (H1, H2, L1)	69, 66, 75	72, 43, 66	83, 51, 81	91, 56, 87	90, 55, 88	90, 56, 88
Found Coincidence	49	59	79	82	82	84
Found Second Coincidence	48	59	77	81	81	81

Table 7 : Results of inspiral searches using TaylorF2 templates. There were 126 injections performed into the data. The table above shows the number of injections which were recovered from the three simulated LIGO detectors (H1, H2 and L1) using various different waveform families, termination frequencies f_{ISCO}, f_{ERD} and f_{WRD} (as described in the text), and post-Newtonian orders.

The results of these searches are reported in Table 7. The principal result is the number of injected signals detected by the search. For simplicity, we define a detected signal as one for which there is a candidate gravitational-wave signal observed within 50 ms of the coalescence time of the injection, determined by the maximum gravitational-wave strain of the injected signal. We do not impose any additional threshold on the measured SNR or effective SNR of the candidate. For a single detector, this will lead to a small number of falsely identified injections, but for coincidence results the false alarm rate is so low that we can be confident that the triggers are associated with the injection. We now describe these results in the order that they appear in Table 7.

Search (1) used second order post-Newtonian templates terminated at f_{ISCO} with a maximum mass of $M \leq 35M_{\odot}$. Despite the fact that no NINJA injections had a mass within the range of this search, a significant number of signals were still recovered in coincidence both before and after signal consistency tests. Although the templates are not a particularly good match to the injected signals, they are still similar enough to produce triggers at the time of the injections. Search (2) changed the boundary of the template bank to $20M_{\odot} \leq M \leq 90M_{\odot}$, but left all other parameters unchanged. The number of detected signals increases significantly as more signals now lie within the mass range searched.

Search (3) extended the upper cutoff frequency of the waveforms to f_{ERD}. The number of signals detected increased from 59 to 77, as expected since these waveforms can detect some of the power contained in the late inspiral or early merger part of the signal [57, 65]. Search (4) extends the post-Newtonian order to 3.5 PN, slightly increasing the number of detected signals to 81. With the limited number of simulations performed in this first NINJA analysis, it is difficult to draw a strong conclusion, although there does seem to be evidence that the higher post-Newtonian order waveforms perform better, consistent with previous comparisons of post-Newtonian and numerical relativity waveforms [57, 114, 73, 64, 65] and chapter 5.

Search (5) uses an upper-frequency cutoff of f_{WRD} for the templates. The number of injections found in coincidence for this search is the same as the search using 3.5 order templates with a cutoff of f_{ERD}, although there are slight differences in the number of found injections at the single detector level.

Search (6) extends the template bank of search (5) to unphysical values of the

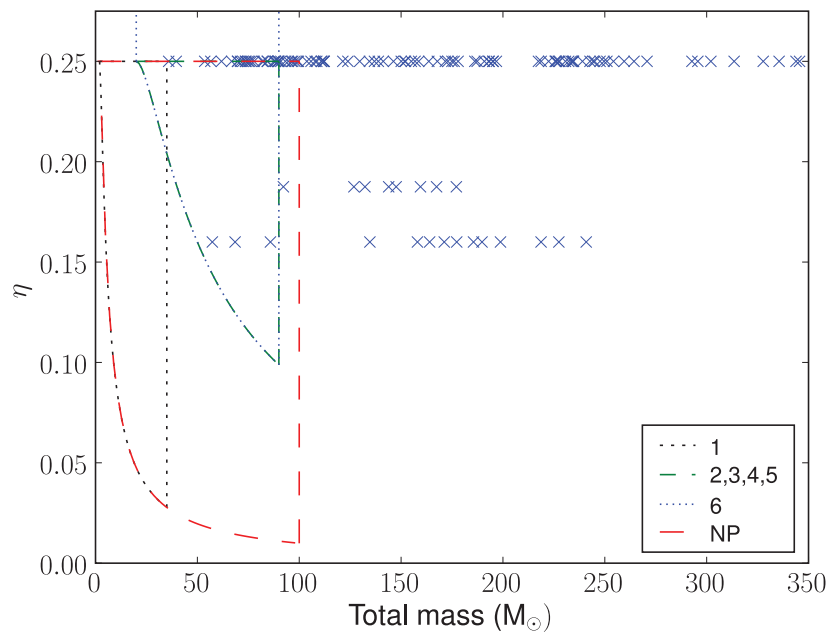Figure 14 : Boundaries of the template banks used in inspiral searches as a function of total mass M and symmetric mass ratio η. The crosses show the location of the injections in the NINJA data set. The numbers in the legend correspond to entries in table 7. Bank 6 extends in a rectangle up to $\eta = 1.00$, as indicated by the arrows. NP is the bank used in the Neyman-Pearson analysis not covered here.
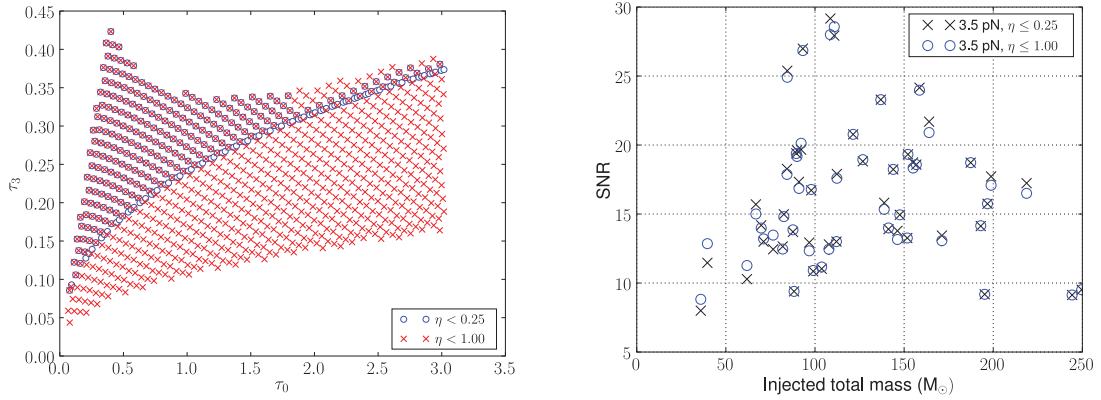
Figure 15 : **Left:** The template bank generated by the LSC-Virgo search pipeline (circles) and the bank obtained by extending to $\eta \leq 1.00$ (crosses). In this figure the bank is parametrised by τ_0 and τ_3 which are related to the binary masses by $\tau_0 = 5M/(256\eta v_0^8)$ and $\tau_3 = \pi M/(8\eta v_0^5)$, where $v_0 = (\pi M f_0)^{1/3}$ is a fiducial velocity parameter corresponding to a fiducial frequency $f_0 = 40.0 Hz$. **Right:** The signal-to-noise (SNR) ratio at which NINJA injections were recovered using the $\eta \leq 0.25$ bank (squares) and the $\eta \leq 1$ extended bank (circles) in the Hanford detectors, given by $\rho = (\rho_{H1}^2 + \rho_{H2}^2)^{1/2}$. The SNR of the signal recovered using the extended bank shows with significant ($> 10\%$) increases over the standard bank for certain injections.

symmetric mass ratio. Extending the bank to $\eta \leq 1$ increases the number of templates in the bank by a factor of ~ 2. The original and modified template banks are shown in Figure 15. With the extended template bank the number of injections found in coincidence remains the same as search (5) after signal-based vetoes are applied. However, many of the injections are recovered at a higher SNR, particular the low-mass signals, as shown in Figure 15. Some injections show a reduction in SNR; more work is needed to understand this effect.

Finally, we note that the majority of signals passed the χ^2 signal-based veto with the thresholds used in the LSC-Virgo pipeline. The last two lines of Table 7 show the number of recovered signals before and after these signal-based vetoes are performed. The post-Newtonian templates and numerical relativity signals are similar enough that virtually all of the injected signals survive the signal based vetoes.

To illustrate the results of these analyses in more detail, Figure 16 shows which signals were detected and which were missed by the 3.5 order post-Newtonian Tay-lorF2 templates terminated at f_{ERD}, as a function of injected total mass and effective distance of the binary (a measure of the amplitude of the signal in the detector),

defined by [30]

$$D_{\text{eff}} = d \left/ \sqrt{F_+^2 (1 + \cos^2 \iota)^2/4 + F_\times^2 \cos^2 \iota} \right. , \tag{6.5}$$

where d is the luminosity distance of the binary.

One signal, with total mass of $110 M_\odot$ and effective distance ~ 200 Mpc, was missed while others with similar parameters were found. This signal was one of the Princeton waveforms (labelled `PU-e0.5` in Figure 10) for which the maximum amplitude occurs at the start of the waveform rather than at coalescence[3], rendering our simple coincidence test invalid. The injection finding algorithm compares the peak time to the trigger time and, even though triggers are found at the time of the simulation, there are no triggers within the 50 ms window used to locate detected signals.

Figure 17 shows the accuracy with which the total mass and coalescence time of the binary are recovered when using the 3.5 post-Newtonian order Taylor F2 templates. The total mass fraction difference is computed as $(M_{\text{injected}} - M_{\text{detected}})/M_{\text{injected}}$. For lower mass signals, the end time is recovered reasonably accurately, with accuracy decreasing for the high mass systems. The total mass recovery is poor for the majority of signals, with good parameter estimation for only a few of the lowest mass simulations.

6.4 Conclusion

We have applied the alterations to the inspiral templates and bank suggested by the studies of the previous chapter to search for NR signals injected into Gaussian noise as part of the NINJA-1 project. The results indicate that these alterations are indeed advantageous. In particular, the second and last columns of table 7 compare the pipeline used in the S5 search (with a mass range extended to cover the range of injections) to the pipeline with all our recommendations implemented. By implementing these changes we recover approximately 25% more injections.

---

[3]That the maximum occurs at the start of the waveform is in part an "artifact" of the double-time integration from the Newman-Penrose scalar ψ_4 to the metric perturbation h, and in part a coordinate artifact. The two integration constants were chosen to remove a constant and linear-in-time piece for h, however, there is still a non-negligible quadratic component; we *suspect* this is purely gauge, though lacking a better understanding of this it was not removed from the waveform.
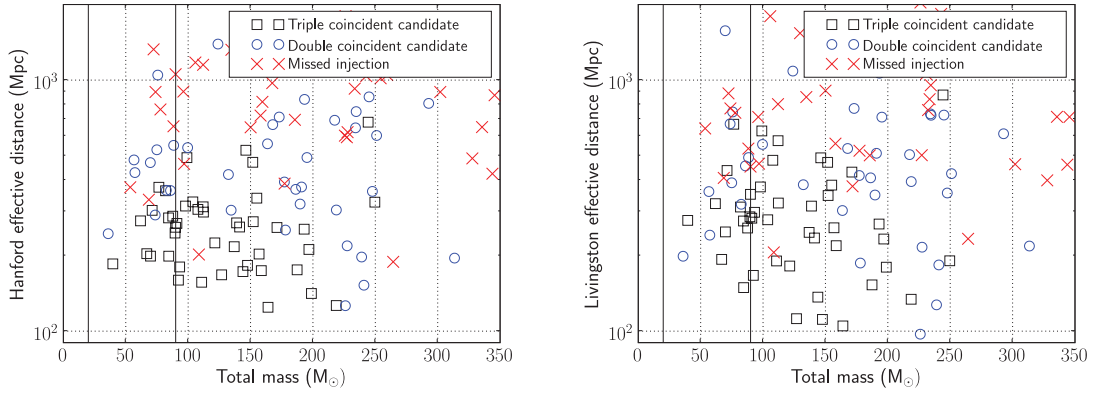
Figure 16 : Found and missed injections using TaylorF2 templates terminated at ERD, plotted as a function of the injected effective distance in Hanford (left) and Livingston (right) and the total mass of the injection. Since the LIGO Observatories are not exactly aligned, the effective distance of a signal can differ, depending on the sky location of the signal. The vertical bars mark the limits of the template bank used in the search. For the lower masses, we see that the majority of the closer injections are found in coincidence in all three of the detectors. There is then a band of injections which are found only in two detectors – H1 and L1 and not the less sensitive H2 detector. For higher masses, the results are less meaningful as the template bank was only taken to a total mass of $90M_{\odot}$.
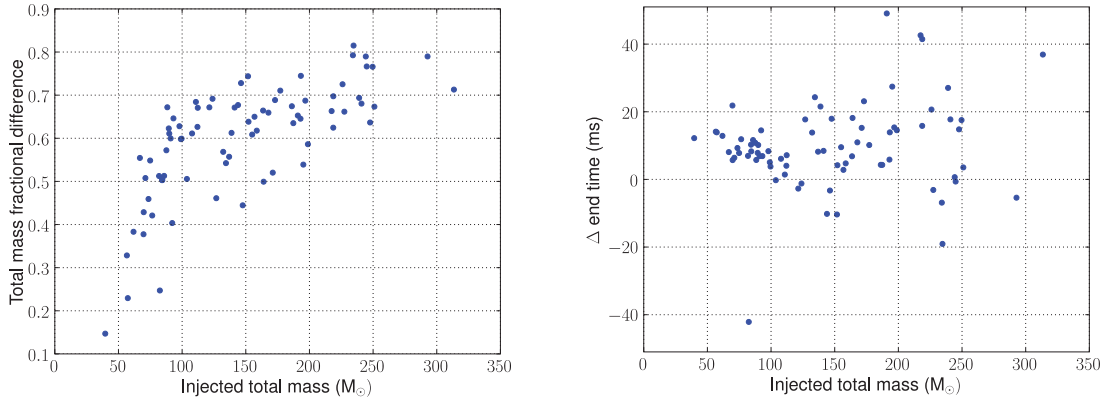
Figure 17 : Parameter accuracy using TaylorF2 templates terminated at ERD. **Left:** Accuracy with which the total mass is recovered. The template bank covers the region $20M_{\odot} \leq M \leq 90M_{\odot}$, hence the mass of injections with $M > 90M_{\odot}$ are always underestimated. Even within the region covered by the bank, the TaylorF2 templates systematically underestimate the mass of the injected signals and the total mass is recovered accurately only for a few injections. The vast majority of recovered signals have an error of 40% or greater. **Right:** Accuracy of determining the coalescence time of the injections. The end time is not recovered accurately, the timing error can become as large as 50ms, the limits of the injection window.

However, NINJA-1 took an open policy towards NR submissions and consequently there may be issues with the waveforms that limit our ability to draw definitive conclusions. The lower limit on mass at which waveforms could be injected certainly limits our ability to make statements on a low-mass search.

In order to refine these results it will be necessary to use longer waveforms that have been subjected to more rigorous validation. This is the goal of the second NINJA project, discussed in the next chapter.

Chapter 7

Waveform Analysis for the Second NINJA project

The NINJA-1 project was a huge success in bringing the numerical relativity and gravitational-wave astronomy communities together. The project also resulted in several intriguing qualitative results. Among these were promising indications that the alterations suggested by the studies of chapter 5 could indeed improve search efficiencies.

However, NINJA-1 only began the process of testing detection and parameter estimation pipelines against realistic signals. The follow-up project, NINJA-2, is ongoing as of the time of writing. NINJA-2 aims to remove some of the shortcomings of NINJA-1 and allow quantitative studies of the performance of gravitational-wave searches in various regions of signal parameter space. Specifically, NINJA-2 addresses issues with both the waveform submissions and the noise used to construct the data sets. This chapter describes the contributed waveforms and the studies that have been performed to validate them. The next chapter will discuss the construction of the data sets and will present preliminary results from the low- and high-mass CBC pipelines.

7.1 Contributed Waveforms

NINJA-1 had an open policy towards waveform submission in order to encourage wide participation. This meant there were no requirements on either waveform quality or

length. The lack of quality requirements allowed for the possibility of unphysical features in the waveforms. There were also no requirements to perform the kind of convergence testing reported in Sec. 5.2, although such validation is typically done by numerical relativists. The loose requirements limited the conclusions that could be drawn, for example it makes it difficult to say whether an injection was missed due to the parameters of the signal or an unintended feature of the waveform.

The lack of length requirement limited the available mass range to $M > 36 M_\odot$ for reasons that can be seen in Fig. 18. This plot shows the $q = 2$, non-spinning GA Tech submission to NINJA-2, along with the initial LIGO and advanced LIGO noise curves (see the discussion of Fig. 5 for an explanation of the choice of axes). Frequencies of interest are marked. Note the triangle in particular, which indicates the frequency at which the numeric waveform starts. If a waveform starts above 40 Hz then we will underestimate the SNR available from such a signal. If the waveform in Fig. 18 had no pN component the lowest mass at which it could be injected would be about 40 $M_\odot$.

To address these issues NINJA-2 specifies the following requirements on the submitted waveforms [115]. The raw numerical simulation should include at least five orbits of usable data before merger (i.e., not counting bursts of junk radiation or other significant noise). Given the computation cost of extending the NR waveforms, we instead require "stitching" to a post-Newtonian inspiral approximant, which should be performed at a GW frequency of $M\omega \leq 0.075$, where $M\omega$ is the frequency of the $(l = 2, m = \pm 2)$ harmonic. The full waveform should be long enough to be entirely within the sensitivity bands of LIGO and Virgo down to $10 M_\odot$ with a lower cutoff frequency of 10 Hz, which corresponds to a starting GW frequency of $M\omega = 0.003$. The numerical-waveform (before any hybridization) amplitude should be accurate to within 5%, and the phase (as a function of GW frequency) should have an accumulated uncertainty over the entire inspiral, merger and ringdown, of no more than 0.5 radian. The PN approximants used for hybridization should ideally use the highest PN orders available, both in phase and amplitude. These minimal accuracy requirements are motivated by the results of the Samurai project [69], and studies performed in preparation for the NR-AR collaboration project [116].

As noted in Eqn. 6.1 the complex function $h_+ - i h_\times$ can be decomposed into modes using spin-weighted spherical harmonics $^{-2}Y_{lm}$ of weight -2. Although most
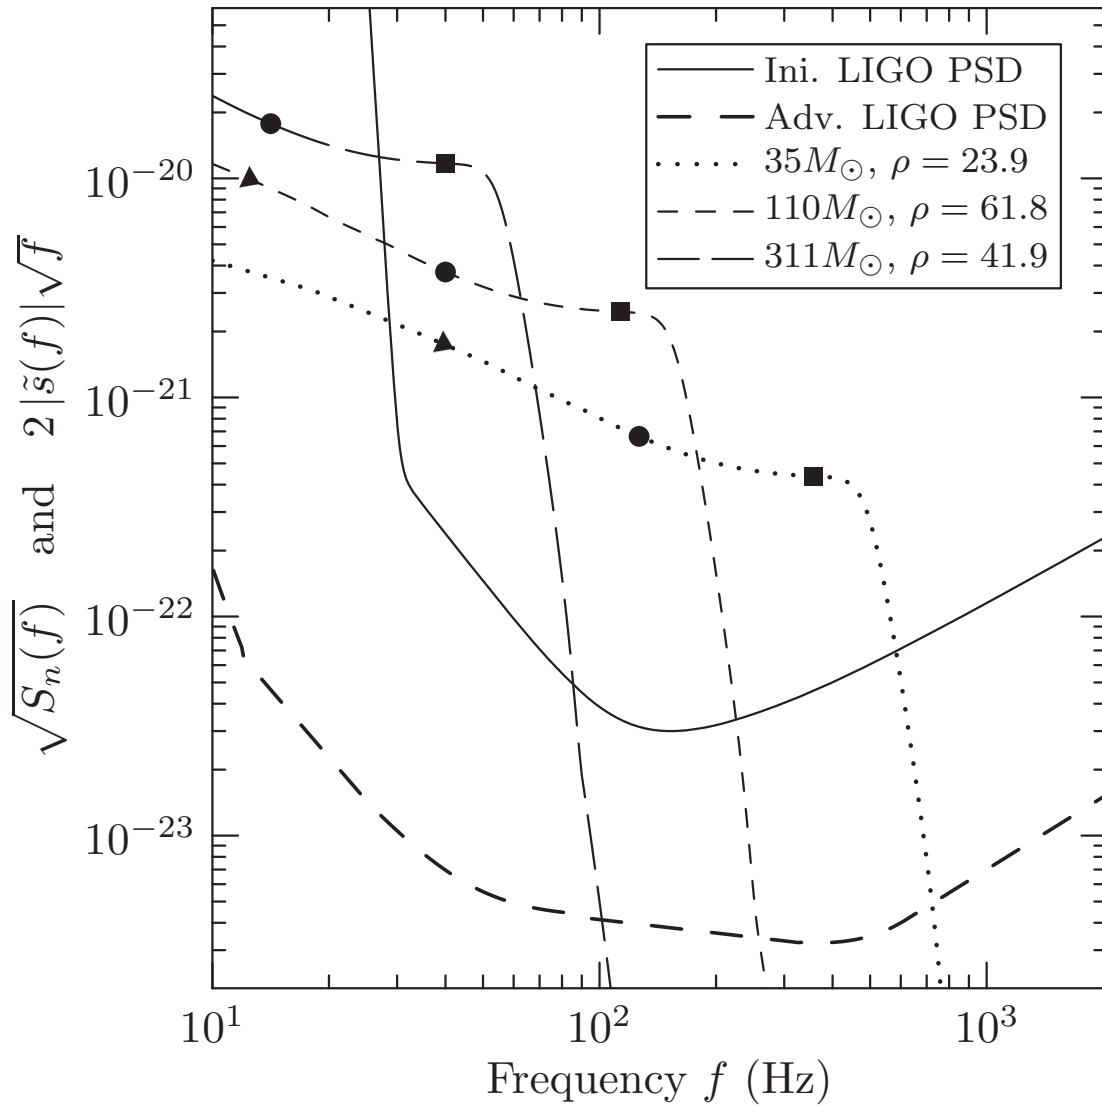
Figure 18 : The hybrid $q = 2$, non-spinning MayaKranc waveform scaled to various total masses, with sources optimally oriented and placed at $100\,\mathrm{Mpc}$, shown against the Initial- and Advanced-LIGO noise curves. Markers are placed along the lines at frequencies corresponding to various instantaneous frequencies of the waveforms. The triangles represent the frequency at which the numeric waveform begins; the circle represents the ISCO frequency; and the square the light-ring. The values given for ρ use the Initial-LIGO noise curve.

of the power is in the $(l, m) = (2, 2)$ mode, the NINJA-2 project encourages, although does not require, the inclusion of higher-order modes. We chose to restrict attention to non-spinning waveforms and waveforms with spins aligned or anti-aligned with the orbital angular momentum. Systems where the component spins are at other angles to the angular momentum exhibit precession, which complicates the waveform generation, the hybridization, and the data analysis significantly. There are sufficient open questions regarding these restricted cases to make this analysis interesting, without adding these additional complications. However, we plan to include precessing waveforms in a forthcoming NINJA-3.

A total of 60 waveforms from 8 groups were contributed to NINJA-2. These are summarized in tables 8, 9, 10, 11, 12, 13, 14, 15 and a map of the parameter values is shown in Fig. 19.

7.2 Verifying the Hybrid Waveforms

Each NR group verified that their waveforms met the minimum NINJA-2 requirements before submission. Once submitted, a series of checks were performed in order to validate the waveforms against each other. In the first stage the post-Newtonian expressions and codes were compared against each other and the literature. This required several iterations, but resulted in a set of codes in various languages that produce waveforms that all agree in both phase and amplitude. The results of this investigation have been included in the latest version of Ref. [70], and are included in this thesis as Appendix A.

7.2.1 Time-domain and Frequency-domain Checks

In the second stage of validation, the complete hybrid waveforms were examined. We first plotted the last 40 cycles of each waveform — enough to include the full NR portion, the hybridization region, and some of the pN portion — and looked for any anomalies such as those present in some of the NINJA-1 waveforms in Fig. 10. The amplitudes of the Fourier transform of the complete waveforms were also plotted. This analysis revealed unphysical features, primarily due to errors in the hybridization. An example is shown in Fig. 20, which shows a visible "kink" in the waveform at the hybridization frequency. This feature is no longer present in the waveform after the

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
BAM_D10spp85_80.T4.hyb.n2	1	0.85	0.85	TaylorT4	[74, 71]
BAM_D10spp85_80.T1.hyb.n2	1	0.85	0.85	TaylorT1	[74, 71]
BAM_D125smm50Nep_80.T1.hyb.n2	1	-0.50	-0.50	TaylorT1	[74, 71]
BAM_D125smm50Nep_80.T4.hyb.n2	1	-0.50	-0.50	TaylorT4	[74, 71]
BAM_D13smm75Nep_96.T4.hyb.n2	1	-0.75	-0.75	TaylorT4	[74, 71]
BAM_D13smm75Nep_96.T1.hyb.n2	1	-0.75	-0.75	TaylorT1	[74, 71]
BAM_D13smm85Nep_88.T4.hyb.n2	1	-0.85	-0.85	TaylorT4	[74, 71]
BAM_D13smm85Nep_88.T1.hyb.n2	1	-0.85	-0.85	TaylorT1	[74, 71]
BAM_D11spp50_96.T4.hyb.n2	1	0.50	0.50	TaylorT4	[74, 71]
BAM_D11spp50_96.T1.hyb.n2	1	0.50	0.50	TaylorT1	[74, 71]
BAM_D10spp75_80.T1.hyb.n2	1	0.75	0.75	TaylorT1	[74, 71]
BAM_D10spp75_80.T4.hyb.n2	1	0.75	0.75	TaylorT4	[74, 71]
BAM_D12smm25Nep_80.T4.hyb.n2	1	-0.25	-0.25	TaylorT4	[74, 71]
BAM_D12smm25Nep_80.T1.hyb.n2	1	-0.25	-0.25	TaylorT1	[74, 71]
BAM_EP_um4_D10-n96.T4.hyb.n2	4	0.00	0.00	TaylorT4	[74, 71]
BAM_EP_um4_D10-n96.T1.hyb.n2	4	0.00	0.00	TaylorT1	[74, 71]
BAM_um3_88.T4.hyb.n2	3	0.00	0.00	TaylorT4	[74, 71]
BAM_um3_88.T1.hyb.n2	3	0.00	0.00	TaylorT1	[74, 71]
BAM_um2_88.T1.hyb.n2	2	0.00	0.00	TaylorT1	[74, 71]
BAM_um2_88.T4.hyb.n2	2	0.00	0.00	TaylorT4	[74, 71]
BAM_R6_PN_80.T1.hyb.n2	1	0.00	0.00	TaylorT1	[74, 71]
BAM_R6_PN_80.T4.hyb.n2	1	0.00	0.00	TaylorT4	[74, 71]
BAM_D12spp25_96.T4.hyb.n2	1	0.25	0.25	TaylorT4	[74, 71]
BAM_D12spp25_96.T1.hyb.n2	1	0.25	0.25	TaylorT1	[74, 71]
BAM_q2a0a025_T_96_344.T1.hyb.n2.bbh	2	0.25	0.00		[71]
BAM_q2a0a025_T_96_344.T4.hyb.n2.bbh	2	0.25	0.00		[71]

Table 8 : BAM submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
BAM_hybrid_om0.025etmq3S0.4-0_0_S0.6_0_0_72	3	0.40	0.60	TaylorT4	[75, 96]

Table 9 : FAU submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
MayaKranc_D12_a0.00_m129_nj	1	0.00	0.00	TaylorT4	[117, 118]
MayaKranc_D10_a0.90_m129_nj	1	0.90	0.90	TaylorT4	[117, 118]
MayaKranc_D10_a0.20_m77_nj	1	0.20	0.20	TaylorT4	[117, 118]
MayaKranc_D10_a0.60_m77_nj	1	0.60	0.60	TaylorT4	[117, 118]
MayaKranc_D12_a0.60_m103_nj	1	0.60	0.60	TaylorT4	[117, 118]
MayaKranc_Sp02py0935th90_gr	1	0.80	0.00	TaylorT4	[117, 118]
MayaKranc_D12_a0.80_m103_nj	1	0.80	0.80	TaylorT4	[117, 118]
MayaKranc_D12_a0.00_q2_m90_nj	2	0.00	0.00	TaylorT4	[117, 118]
MayaKranc_D11_a0.20_q2_m90_nj	2	0.02	0.09	TaylorT4	[117, 118]
MayaKranc_D10_a0.40_m90_nj	1	0.40	0.40	TaylorT4	[117, 118]
MayaKranc_D10_a0.80_m90_nj	1	0.80	0.80	TaylorT4	[117, 118]
MayaKranc_D12_a0.40_m103_nj	1	0.40	0.40	TaylorT4	[117, 118]
MayaKranc_D12_a0.20_m103_nj	1	0.20	0.20	TaylorT4	[117, 118]

Table 10 : GATech submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
dq4	4	0.00	0.00	TaylorT1	[86]

Table 11 : LEAN submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
Llama_d550-h64-Hybrid	1	0.00	0.00	3.5pNTaylorF2	[119, 119]
Llama_d4d4-q1–D10-h64-r250.T4.hybrid	1	-0.40	-0.40	TaylorT4	[120, 121]
Llama_d4d4-q1–D10-h64-r250.T1.hybrid	1	-0.40	-0.40	TaylorT1	[120, 121]
Llama_u4u4-q1–D8-h64-r250.T1.hybrid	1	0.40	0.40	TaylorT1	[120, 121]
Llama_u4u4-q1–D8-h64-r250.T4.hybrid	1	0.40	0.40	TaylorT4	[120, 121]
Llama_d5q2-h016-Hybrid	2	0.00	0.00	3.5pNTaylorF2	[119]
Llama_u2u2-q1–D8-h64-r250.T1.hybrid	1	0.20	0.20	TaylorT1	[120, 121]
Llama_u2u2-q1–D8-h64-r250.T4.hybrid	1	0.20	0.20	TaylorT4	[120, 121]
Llama_d2d2-q1–D10-h64-r250.T1.hybrid	1	-0.20	-0.20	TaylorT1	[120, 121]
Llama_d2d2-q1–D10-h64-r250.T4.hybrid	1	-0.20	-0.20	TaylorT4	[120, 121]

Table 12 : Llama submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
LazEV_D8.4_10to1_nj_hybrid	10	0.00	0.00	TaylorT4	[11]

Table 13 : RIT submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
SpEC_q6s0	6	0.00	0.00	TaylorT1	[122]
SpEC_q4s0	4	0.00	0.00	TaylorT2	[122]
SpEC_EqualMassAntiAlignedSpins	1	-0.44	-0.44	NA	[123, 122]
SpEC_q1s-0.95	1	-0.95	-0.95	TaylorT1	[122]
SpEC_q2s0	2	0.00	0.00	TaylorT2	[122]
SpEC_EqualMassAlignedSpins	1	0.44	0.44	NA	[123, 122]
SpEC_q3s0	3	0.00	0.00	TaylorT2	[122]
SpEC_EqualMassNonspinning	1	0.00	0.00	TaylorT4	[22, 122]

Table 14 : SpEC submissions to NINJA-2

Run	q	Spin1$_z$	Spin2$_z$	pN Approx.	Refs
UIUC_spin_-0.25_om0.0528_22-HYBRID	1	-0.25	-0.25	NA	[124, 125, 126]
UIUC_spin_0.85_om0.0536_22-HYBRID	1	0.85	0.85	NA	[124, 125, 126]

Table 15 : UIUC submissions to NINJA-2
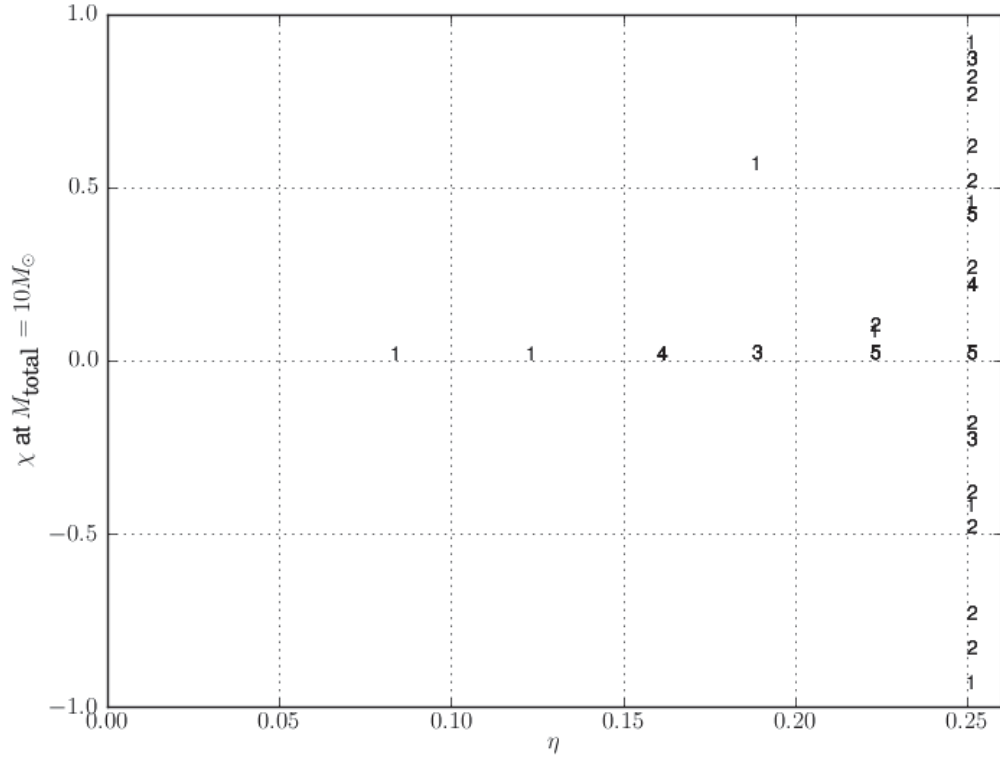

Figure 19 : Parameters of the NINJA-2 hybrid waveform submissions showing the symmetric mass ratio $\eta = m_1 m_2/(m_1 + m_2)^2$ and dimensionless spin parameter $\chi = (S_1/m_1 + S_2/m_2)/(m_1 + m_2)$ after scaling the waveforms to a reference total mass of 10 $M_\odot$. The numbers indicate how many distinct waveforms with the specified parameters were submitted.

hybrid was reconstructed correctly.

7.2.2 Overlap Comparisons

In this check the waveforms were compared against each other using standard data-analysis techniques, in particular the overlap defined in Eqn. 4.20 using the initial LIGO noise curve. The waveforms were grouped into sets with identical parameters. For each set one waveform was chosen as the reference and the overlap with all the other waveforms was calculated over a range of masses, optimizing over the unknown coalescence time and phase. This process was then repeated, taking each of the other waveforms as the reference in turn.

There is an important subtlety involved in calculating these overlaps. Recall from Sec. 4.1 that the output of the matched filter is a time series from which the point with the largest value is chosen in order to maximize over time. When comparing two very similar waveforms, such as two numeric simulations of the same system, the overlap function becomes very sharply peaked in time. It is therefore imperative that the sample rate used in calculating the overlap be large enough to find the true maximum.

This issue is demonstrated in Fig. 21. This figure shows the overlap function resulting from the comparison of two equal-mass, non-spinning waveforms sampled at four different rates. At 4096 Hz the peak of the function is missed and the overlap is underestimated. The consequence of this is illustrated in Fig. 22 which shows that the overlap as a function of mass exhibits oscillations. The mass determines the length of the waveform, as this length changes the waveforms slide with respect to the grid of sample points, and the overlap function traces out this beating. Consequently, all overlaps in this chapter were calculated at 32768 Hz. The LIGO/Virgo matched filter searches operate on data at 4096 Hz, however this issue is not a problem in these searches for reasons that can be seen in these two images. The loss of overlap by undersampling is no larger than 0.2%. In constructing the template bank (Sec. 4.3) we have already incurred a potential loss of 3%, this effect is therefore negligible.

For reference we first present comparisons between time-domain post-Newtonian waveforms of the kind used to construct hybrid waveforms. These are shown in Fig. 23. Overlaps between hybrid waveforms constructed using different pN approximants will not agree better than the values presented here at low mass, where the pN portion
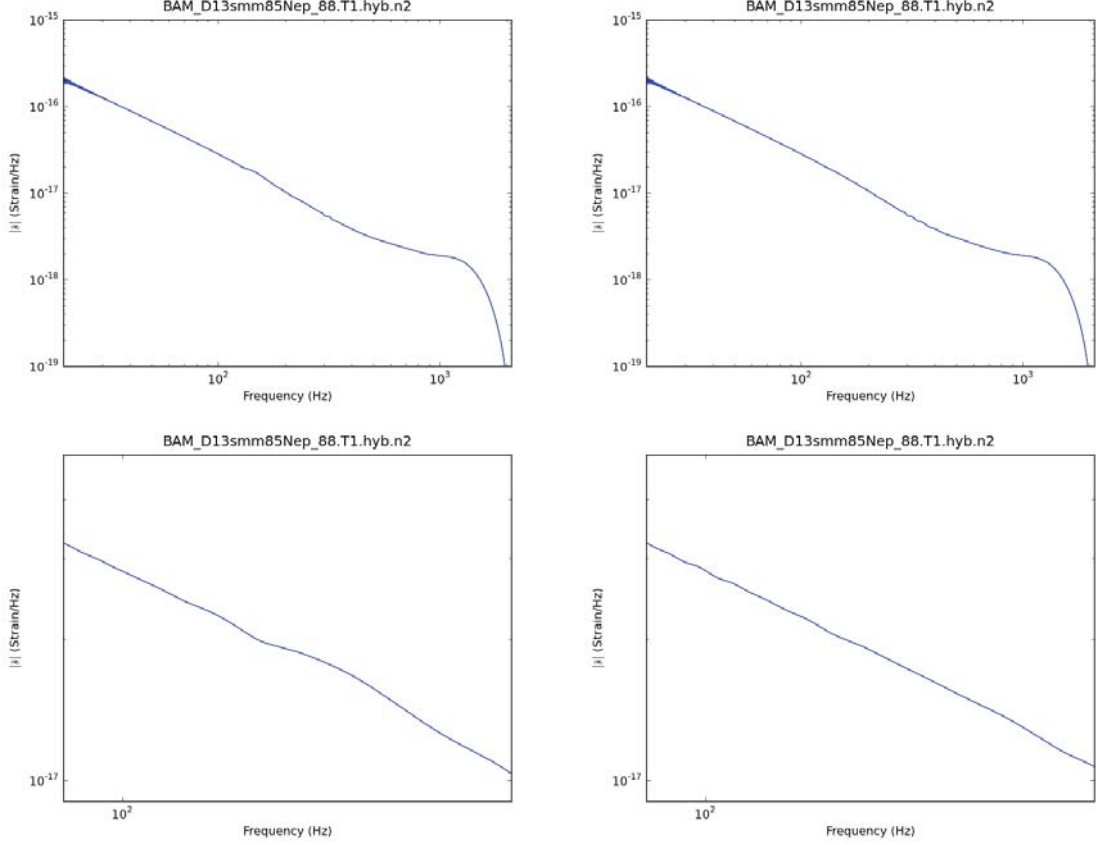
Figure 20 : Fourier amplitude of the (2,2) mode of a sample NINJA-2 hybrid waveform from the BAM/AEI group. The waveform has been scaled to 10 $M_\odot$ and placed 1 Mpc from the detector to give it physical units. The waveform on the top left is the version initially submitted, note there is a small visible "kink" in the waveform at around 100 Hz. The waveform on the top right has been re-hybridized and there is no longer a visible kink. On the bottom, the same waveform before and after rehybridization, zoomed into the region in question. This feature did not show up in the time domain view of the waveform. Recall that the in the post-Newtonian waveform Eq. 2.39 the amplitude evolved as $f^{-7/6}$. This can be seen here as a straight line (on the log scale) that extends up to ≈ 400 Hz. Above this point the slope changes as the system transitions to merger. The power then drops and approaches a constant frequency as the system rings down.
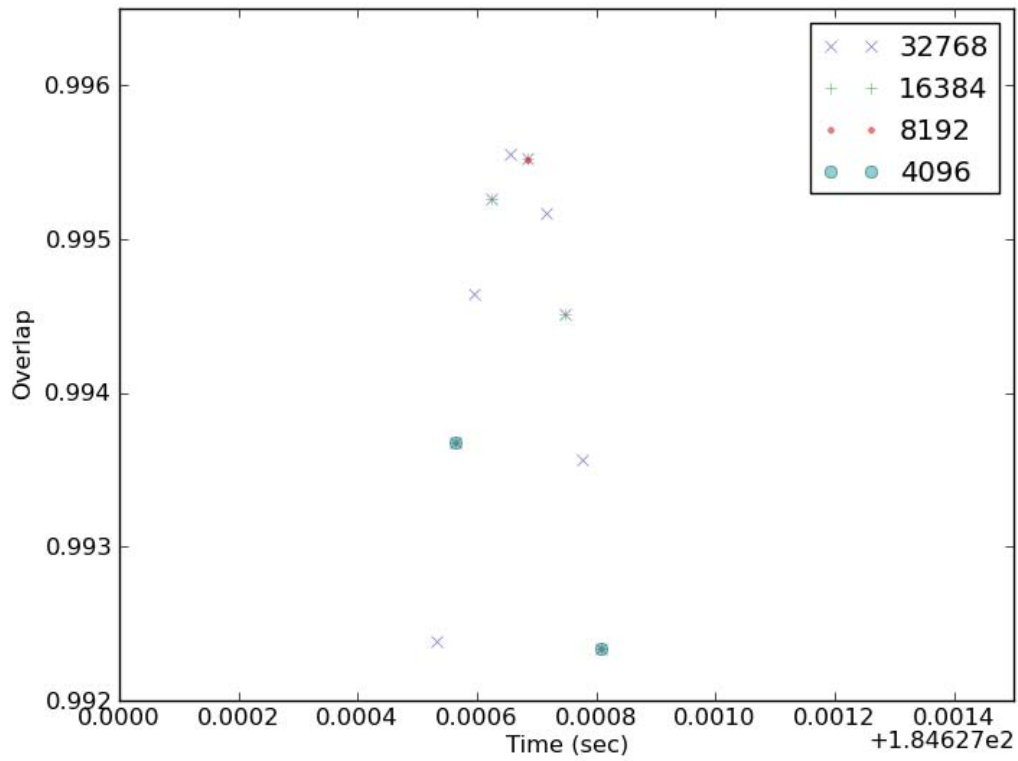
Figure 21 :   The overlap time series between the Georgia Tech and SpEC equal-mass, non-spinning waveforms at different sample rates. At 4096 Hz the sample point nearest the maximum is sufficiently far that the overlap is underestimated to an extent which is significant in doing such comparisons, where we are interested in differences to one part in 10^{-5}.
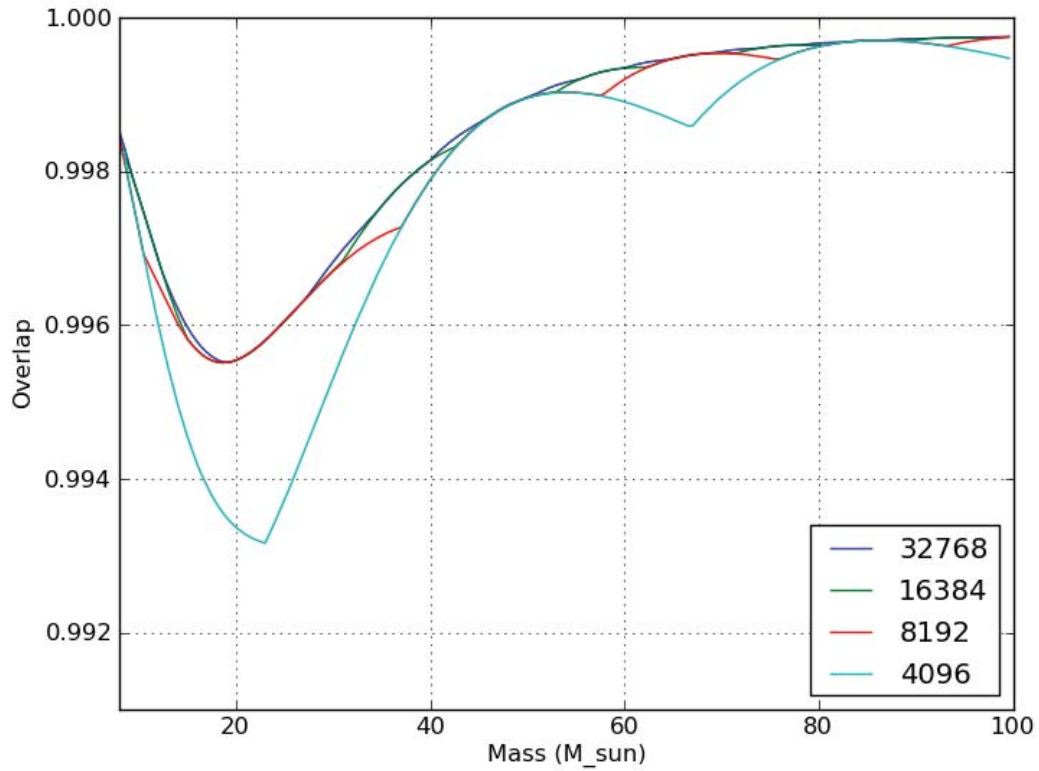
Figure 22 : Overlaps between the Georgia Tech and SpEC equal-mass, non-spinning waveforms, as a function of mass, at different sample rates. At 4096 Hz the Consequence of undersampling the overlap function. As the waveform lengths change with mass they beat against the griding caused by discrete sampling. This produces periodic beats, most evident in the overlaps sampled at 4096 Hz.

of the waveform extends through the most sensitive portion of the LIGO band. A sample set of overlaps, comparing all equal-mass, non-spinning waveforms submissions is shown in Fig. 24. At the high-mass end, where the numeric portion of the waveform dominates the overlap, the overlaps approach 1. Since all these waveforms model the same physical system, this is the expected behavior. This is nevertheless a significant result, as the waveforms were produced with different codes. At the low-mass end, where the overlap is dominated by the post-Newtonian portions of the waveform, the behavior is qualitatively as expected from Fig. 23. Although some of the overlaps, such as between the SpEC and BAM T1 waveforms, are notably lower than would be expected based on pN considerations alone. This is because different submissions hybridize at different frequencies, so in some cases there is a loss in overlap due to comparing pN waveforms against the hybridization.

This is more pronounced in the region of $\sim 20M_\odot$ where all the overlaps drop, some more significantly than others. This is the region where the hybridization is passing through the sensitive band. The goal of hybridization is to smoothly interpolate between the pN and NR waveforms, and in cases where the pN and NR separately agree we would expect this agreement to extend through the hybridization region. We see here, however, that this is not the case. Different choices of hybridization methods and parameters, as well as the frequencies over which the hybridization is performed, can lead to waveforms with significant mismatches.

These initial results prompted a number of the NR groups to revise their hybridization procedures, after which the overlaps were more in line with the expected values. The overlap plot produced with the updated waveforms is shown along with the initial results in Fig. 24. It is worth noting that even after rehybridizing there are still mismatches. In particular, the SpEC and MayaKranc submissions use the same hybridization method, the same pN approximant and are simulating the same physical system. The only differences are the length of the NR waveforms and the frequencies at which the hybridization is performed. The question of how many NR cycles are needed in order to produce a robust hybrid waveform is an area of active research [127].

The full set of comparisons between the final versions of the submissions appears at the end of the chapter, in figures 28 through 37.
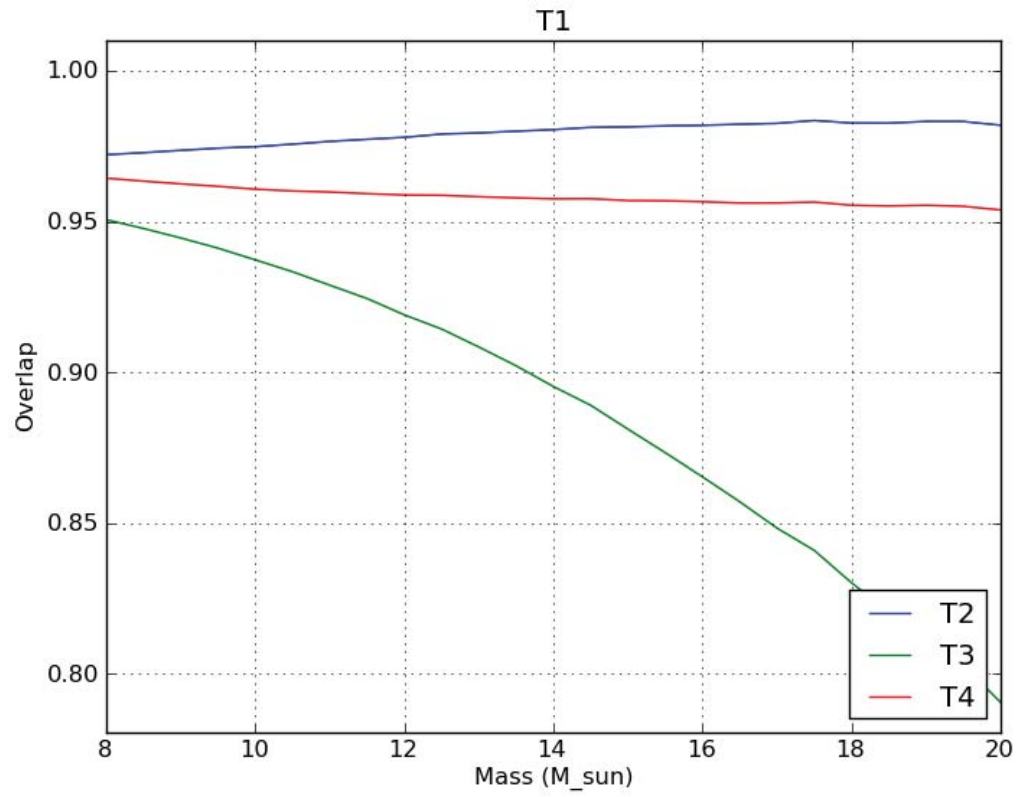
Figure 23 : Overlaps between equal-mass, non-spinning, post-Newtonian time-domain waveforms. Note in particular the discrepancy between T1 and T4, as these were used in the majority of the hybrid waveforms.
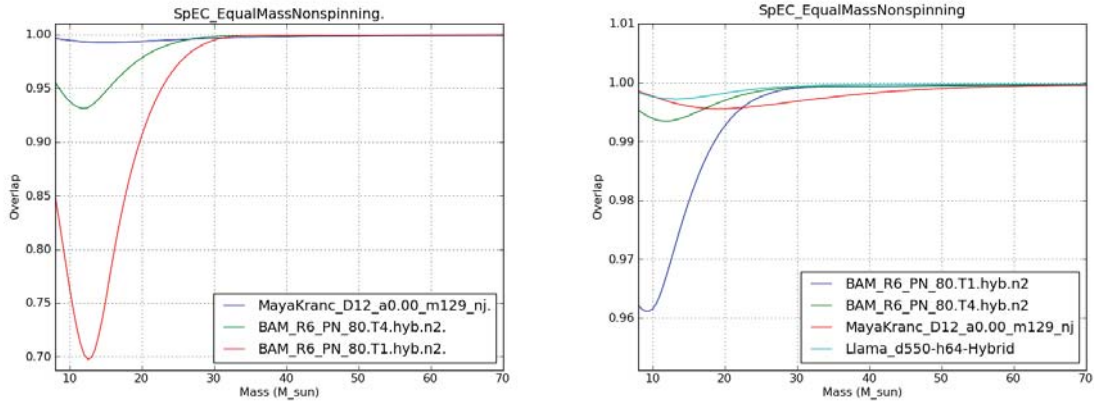
Figure 24 : Overlaps between the equal-mass, non-spinning NINJA-2 contributions, maximized over time and phase. For the original submissions (left) overlaps are as low as 0.70 between waveforms using different pN approximants and 0.94 for waveforms using the same approximant. After rehybridization (right) the waveforms achieve much higher overlaps, with minima above 0.94 for different approximants and above 0.98 for identical approximants. The residual differences between waveforms using TaylorT4 are due to hybridization details. The Llama waveform was accidentally omitted from the original runs.

The overlap plots discussed thus far address the "detection question." By validating the waveforms against each other we ensure that they are all capture the underlying physics well enough to provide high overlaps with existing templates, which in turn implies that these waveforms can be found in searches.

We also extended these overlap studies by maximizing over the mass of one of the waveforms, as well as the time and phase. This addresses the "parameter estimation question," the bias in the recovered mass gives an estimate of the minimum error that can be expected in parameter estimation pipelines. Example plots using the equal-mass, non-spinning MayaKranc waveform as the signal and BAM plus two different approximants as the template are shown in Fig. 25.

At the high-mass end the overlap is dominated by NR data, and as in Fig. 24 the overlaps are high without needing to move off the signal mass. At the low-mass end the same result would be expected in a pure pN/pN comparison although there is enough of the hybridization in-band to reduce the overlaps. However, changing the mass introduces a phase difference that accumulates over all the cycles in-band, and so higher overlaps can not be achieved. The result is optimal mass values close to the correct mass value, but with a low overlap. In the middle region these factors
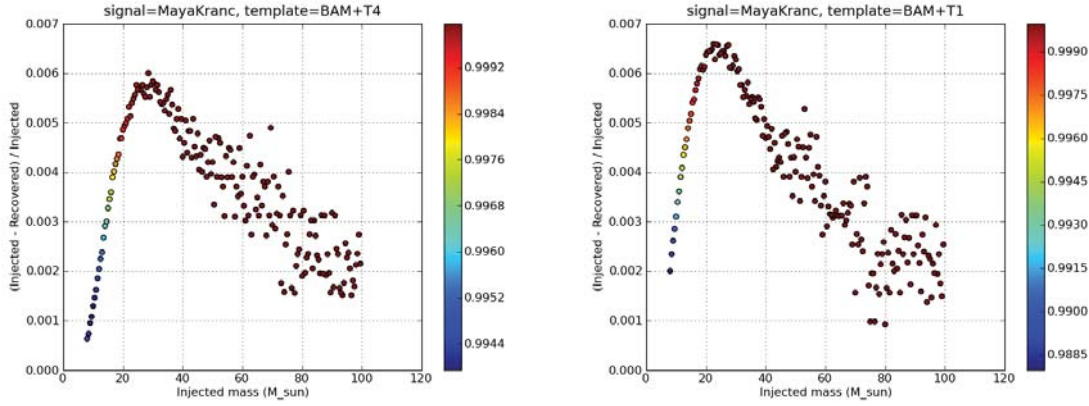
Figure 25 : Overlaps (indicated by color) between the equal-mass, non-spinning MayaKranc waveform taken as the signal, and the equal-mass, non-spinning BAM waveform hybridized with TaylorT4 (left) and TaylorT1 (right) taken as templates. Maximization is done over the mass of the template, as well as over time and phase. The x-axis gives the mass of the signal, the y-axis gives the fractional difference between the injected mass and the mass of the template that maximizes the overlap Note the lower overall overlaps and mass bias at the low-mass end of the figure on the right, where the two different pN waveforms dominate the overlap.

compete. At higher masses the overlap is reduced less by changing the mass and so the recovered value can stray further from the injected value. As the hybridization passes out of band this adjustment is no longer needed. The same general behavior can be seen in comparisons between non-spinning, unequal-mass ($q = 2$) waveforms, shown in Fig. 26. However, the overlaps drop slightly above $\sim 60 M_\odot$, suggesting a disagreement in the NR portion of the waveforms. The comparison between two equal-mass, spinning ($s_{1z} = s_{2z} = 0.4$) waveforms is shown in Fig. 27. Here there is a bias in recovered mass at the high-mass end, although the overlaps are high. This could suggest a problem with the overall scaling of one of the waveforms, but more study is needed.

7.3   Conclusion

We have reviewed the hybrid waveforms submitted to the second NINJA project. Although many groups have contributed many waveforms, these primarily cover only two lines in the (mass ratio, χ) plane, leaving large regions of parameter space unexplored. We also do not consider precessing signals, which adds several dimensions
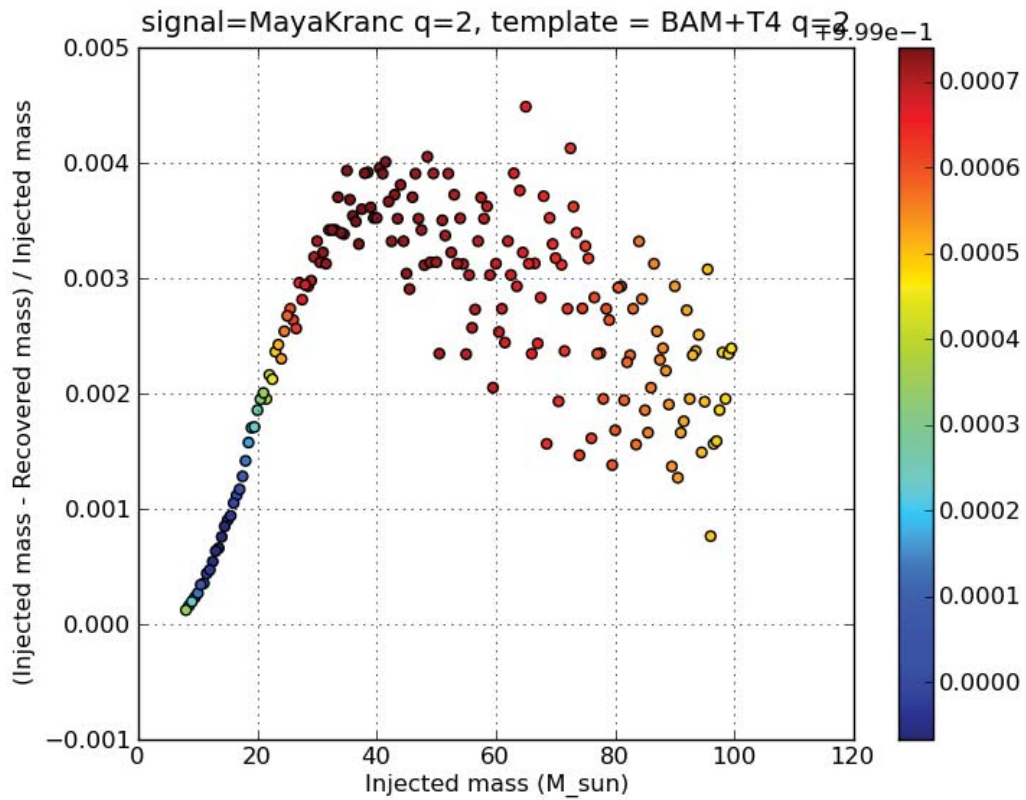
Figure 26 : Overlaps (indicated by color) between the $q = 2$ non-spinning MayaKranc waveform taken as the signal, and the $q = 2$, non-spinning BAM waveform hybridized with TaylorT4 taken as the templates. Maximization is done over the mass of the template, as well as over time and phase. The x-axis gives the mass of the signal, the y-axis gives the fractional difference between the injected mass and the mass of the template that maximizes the overlap. The low overlaps above $60M_{\odot}$ are due to differences in the NR waveforms.
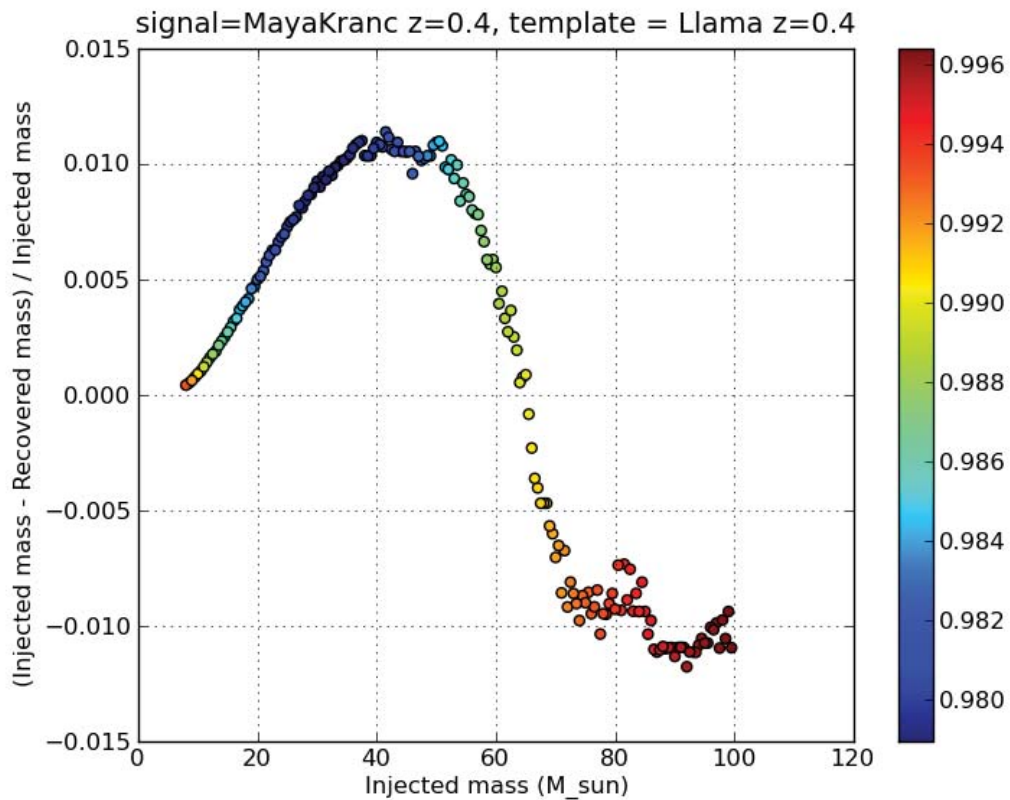
Figure 27 :   Overlaps (indicated by color) between the equal-mass $s_{1z} = s_{2z} = 0.4$ MayaKranc waveform taken as the signal, and the equal-mass $s_{1z} = s_{2z} = 0.4$ Llama waveform taken as the template. Maximization is done over the mass of the template, as well as over time and phase. The x-axis gives the mass of the signal, the y-axis gives the fractional difference between the injected mass and the mass of the template that maximizes the overlap. The low overlaps above $60 M_\odot$ are due to differences in the NR waveforms. The reason for the systematic bias at the high-mass end is under investogation.

to parameter space. These issues will be further explored in NINJA-3. A major feature of NINJA-2, not present in NINJA-1, is the validation and comparisons of the submissions discussed in this chapter. These efforts lead to dramatic improvements in the quality of the waveforms, which in turn will enable data analysts to draw more quantitative conclusions on the behavior of their pipelines with respect to the underlying physics. There is evidence that hybridization choices and methods effect both overlaps and parameter estimation. The degree to which these effects will bias searches is a question we hope NINJA-2 will be able to answer. In the next chapter we discuss the construction of data sets using these waveforms, and present preliminary results from the CBC low- and high-mass pipelines.

Figure 28 : Top left: overlap plot for $q = 1$ $S_{z1} = S_{z2} = -0.5$ Top right: overlap plot for $q = 1$ $S_{z1} = S_{z2} = -0.4$ Bottom left overlap plot for $q = 1$ $S_{z1} = S_{z2} = -0.2$ Bottom right: overlap plot for $q = 1$ $S_{z1} = S_{z2} = 0.8$

Figure 29 : Overlap plots for $q = 3$ $S_{z1} = S_{z2} = 0$

Figure 30 : Overlap plots for $q = 1$ $S_{z1} = S_{z2} = 0.2$
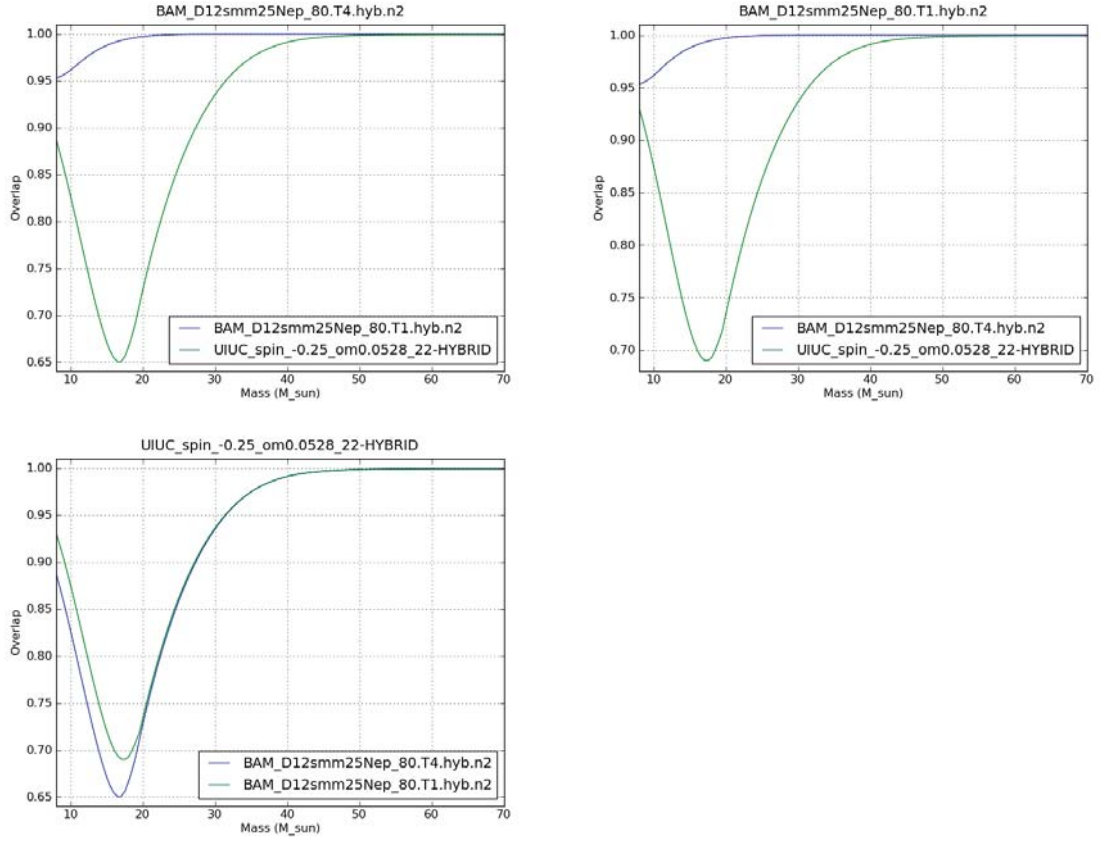
Figure 31 : Overlap plots for $q = 1$ $S_{z1} = S_{z2} = -0.25$
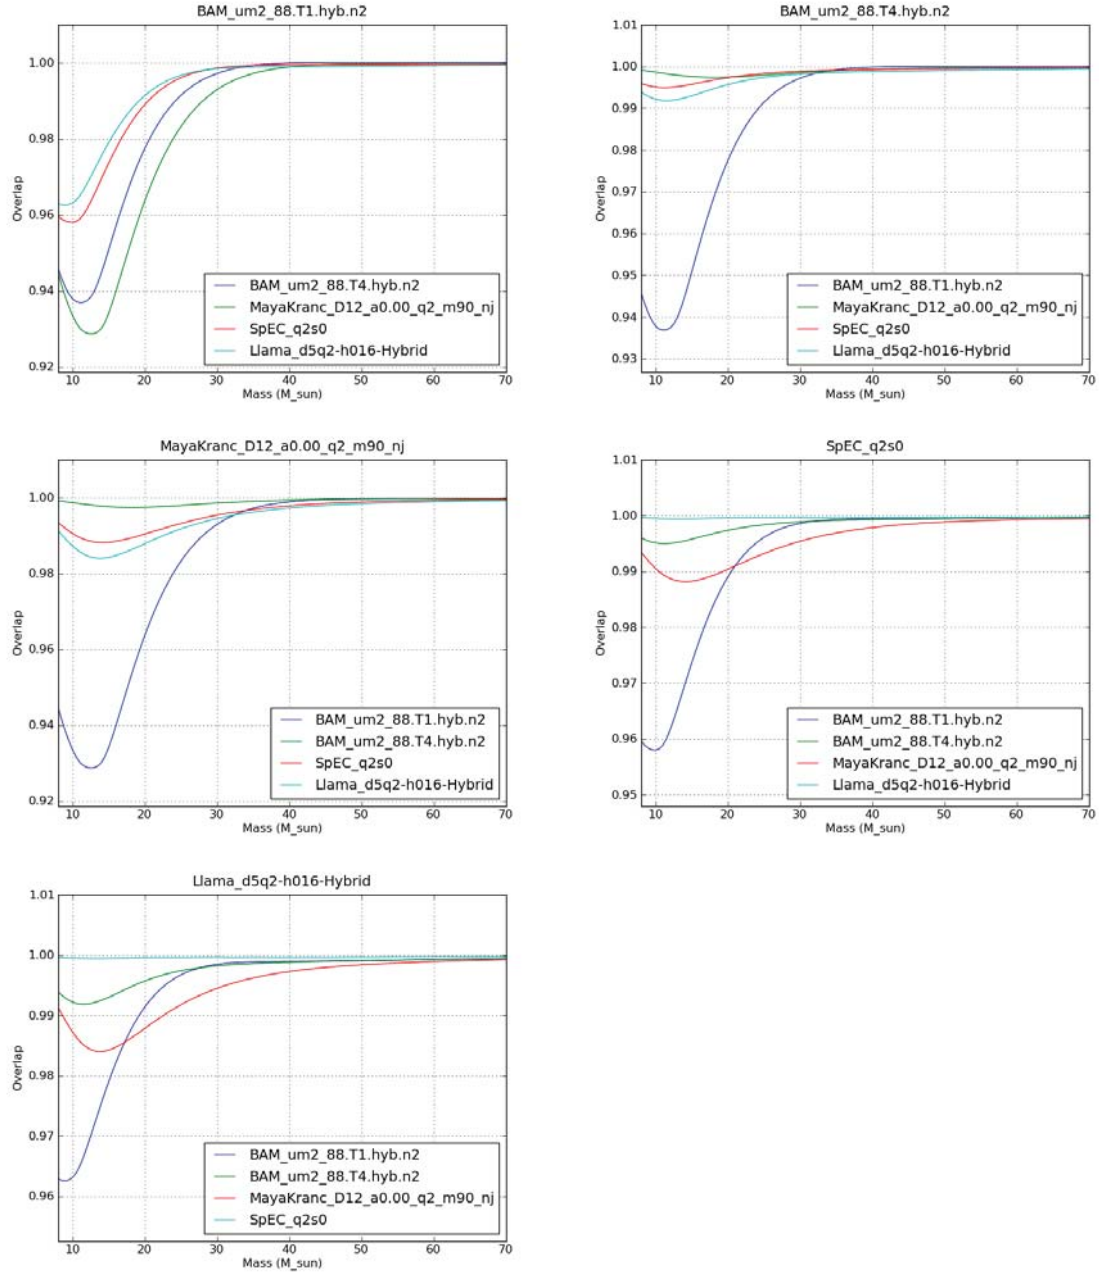
Figure 32 : Overlap plots for $q = 2$ $S_{z1} = S_{z2} = 0$
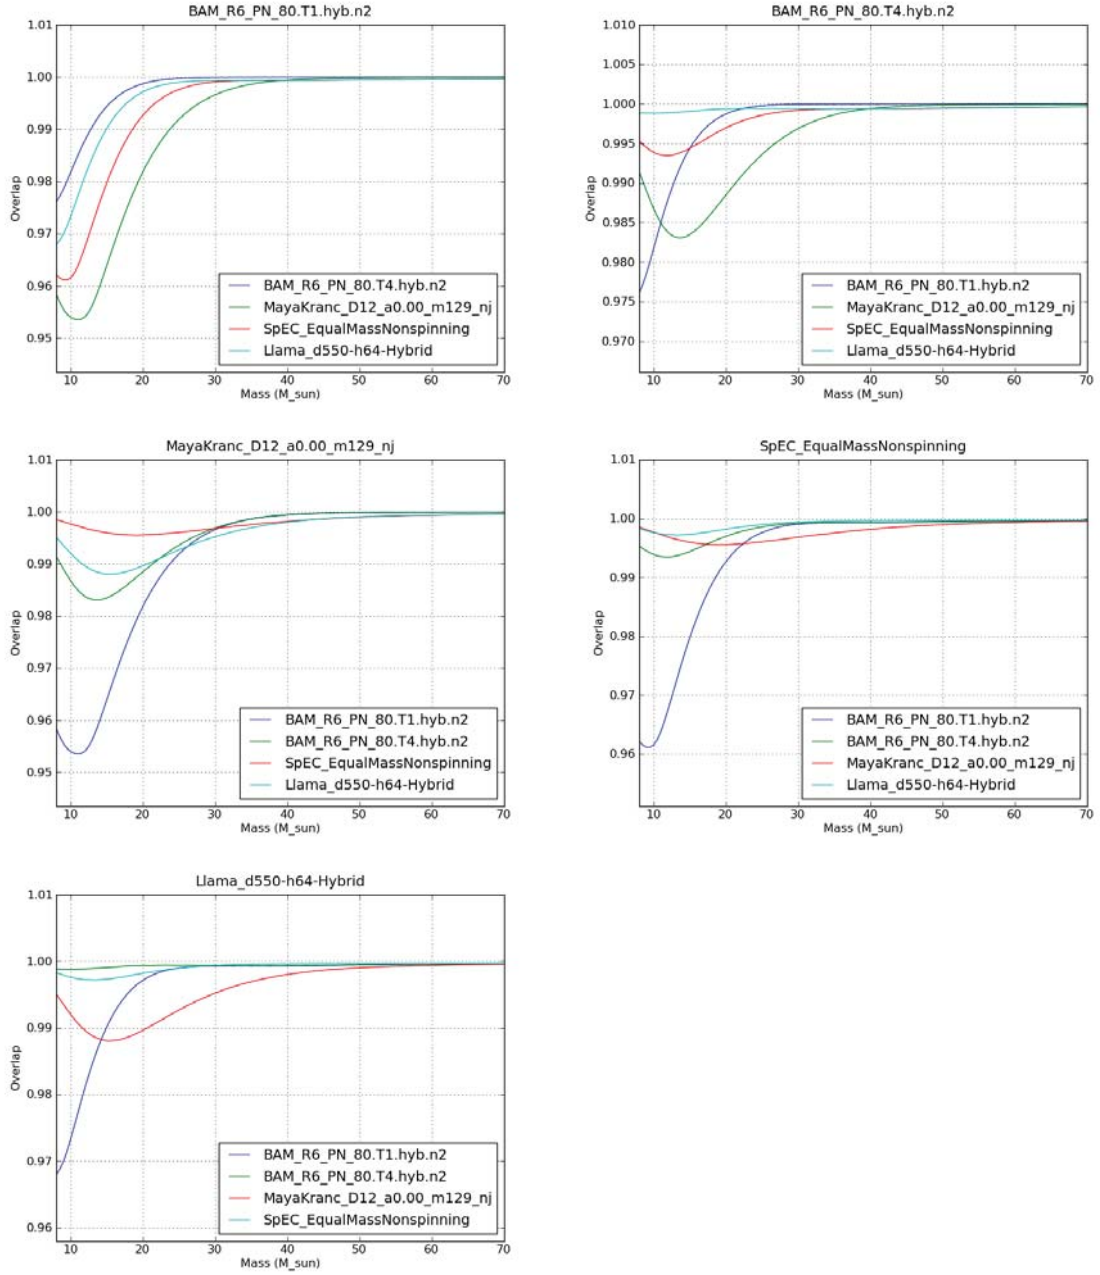
Figure 33 : Overlap plots for $q = 1$ $S_{z1} = S_{z2} = 0$
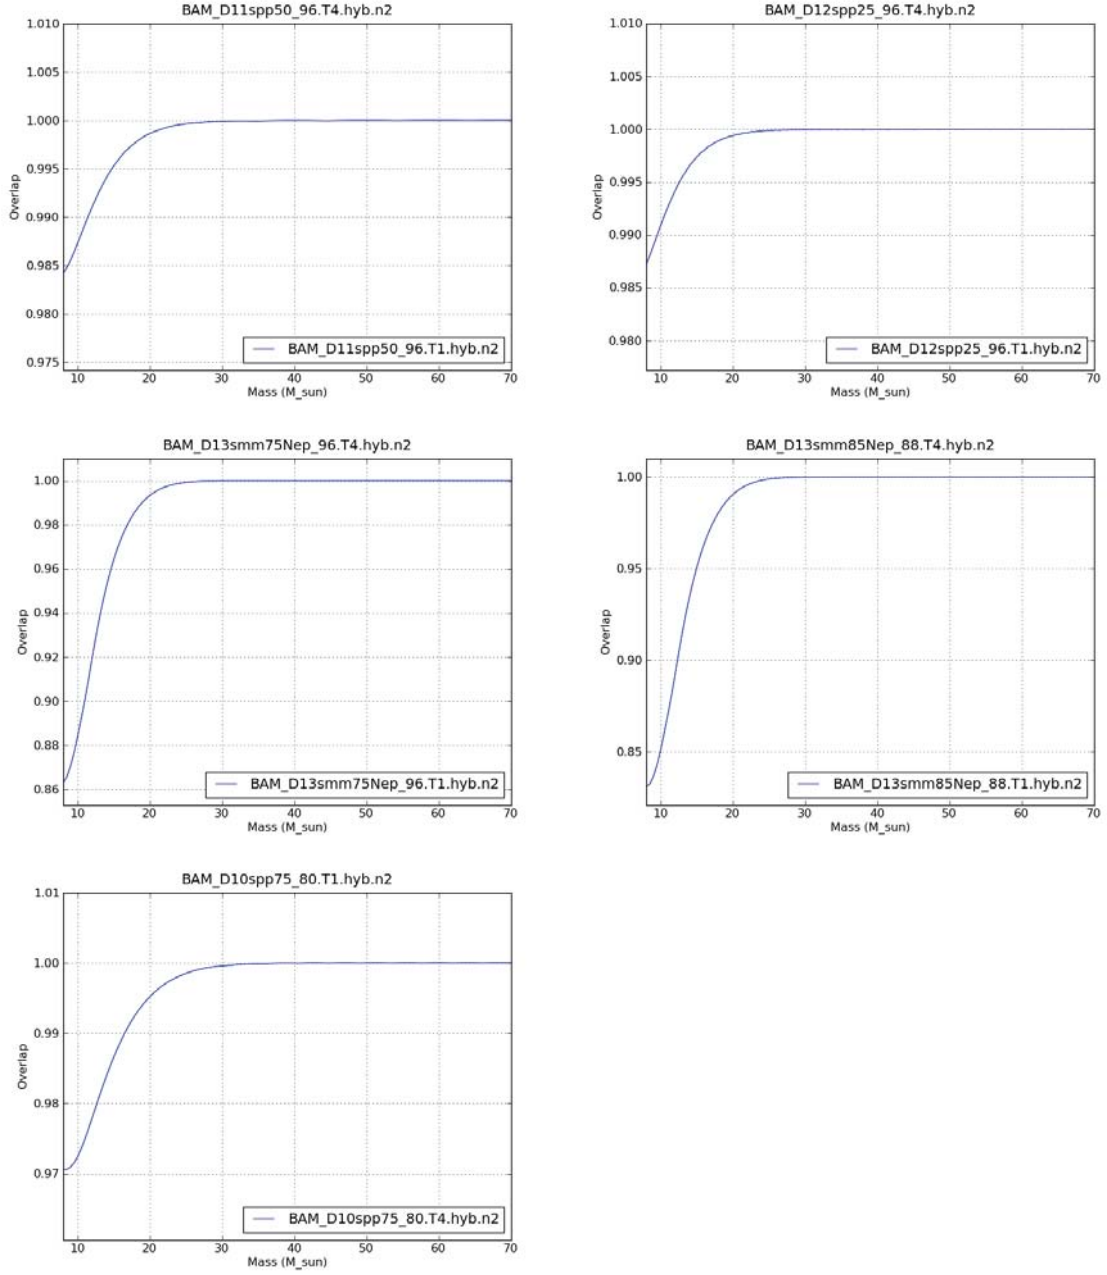
Figure 34 : Top left: overlap plot for $q = 1$ $S_{z1} = S_{z2} = 0.5$ Top right: overlap plot for $q = 1$ $S_{z1} = S_{z2} = 0.25$ Middl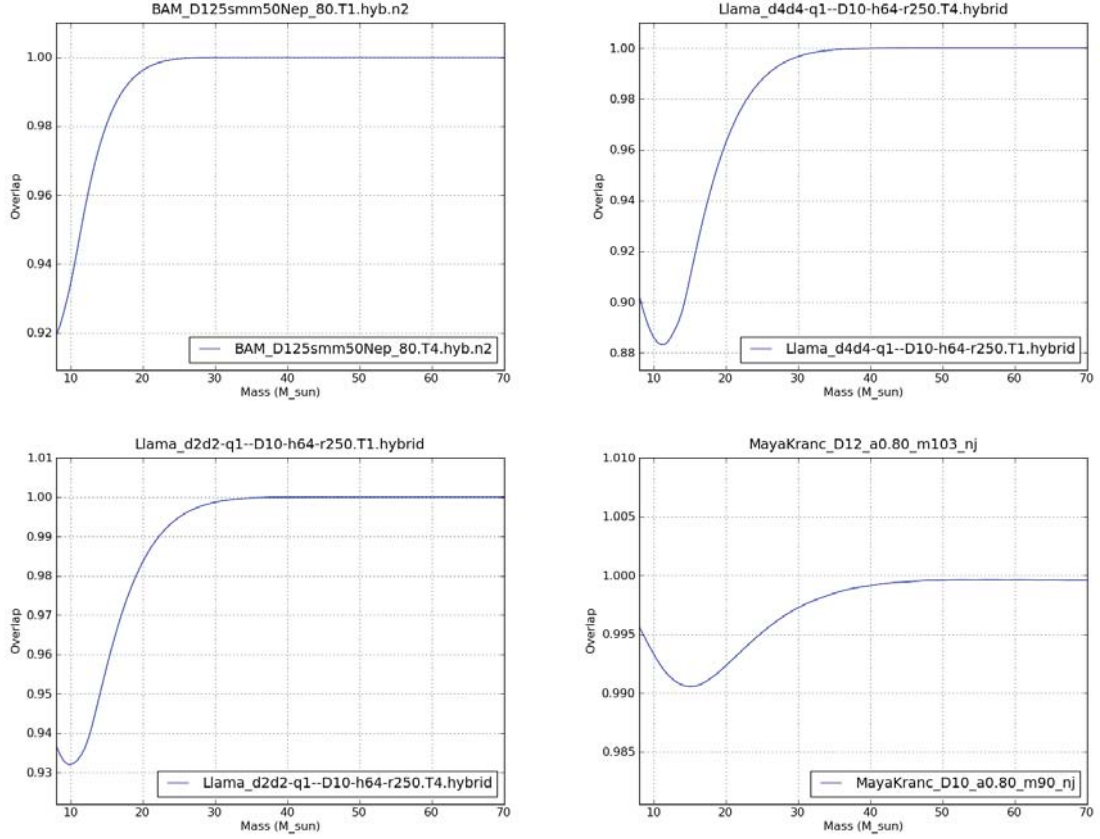e left: overlap plot for $q = 1$ $S_{z1} = S_{z2} = -0.75$ Middle right: overlap plot for $q = 1$ $S_{z1} = S_{z2} = -0.85$ Bottom left: overlap plot for $q = 1$ $S_{z1} = S_{z2} = 0.75$
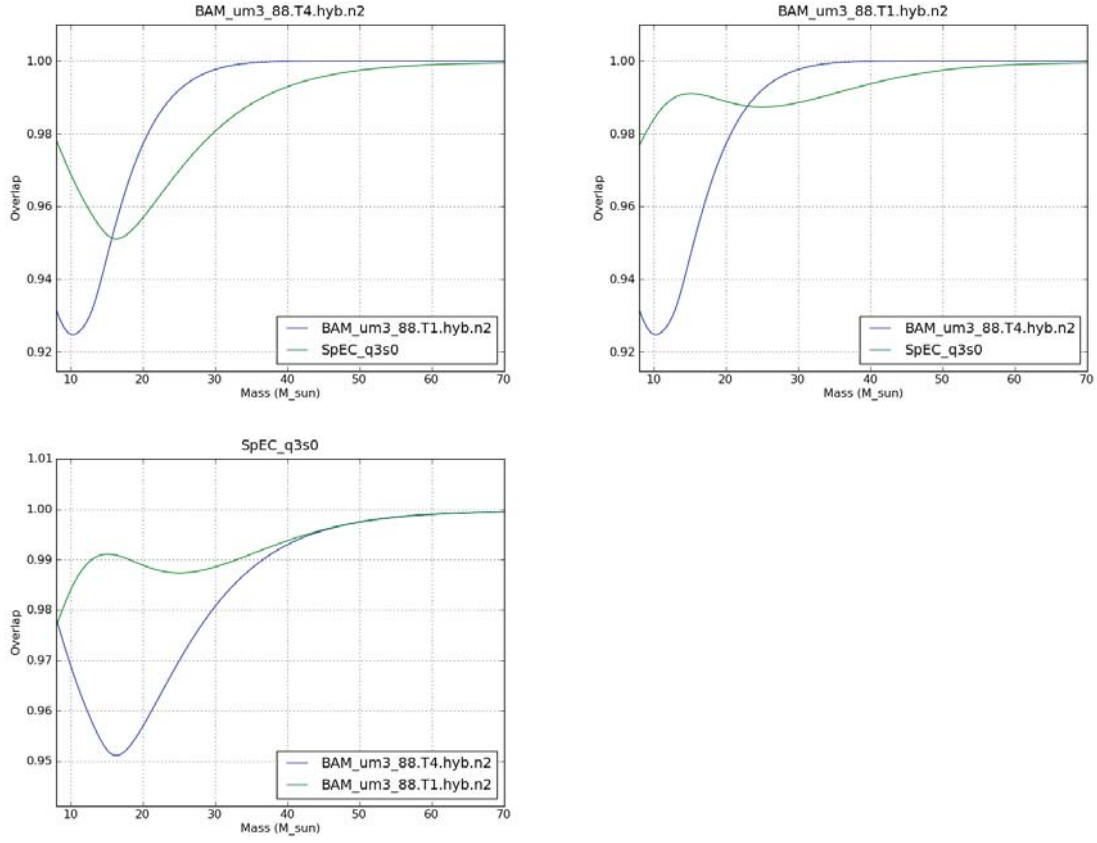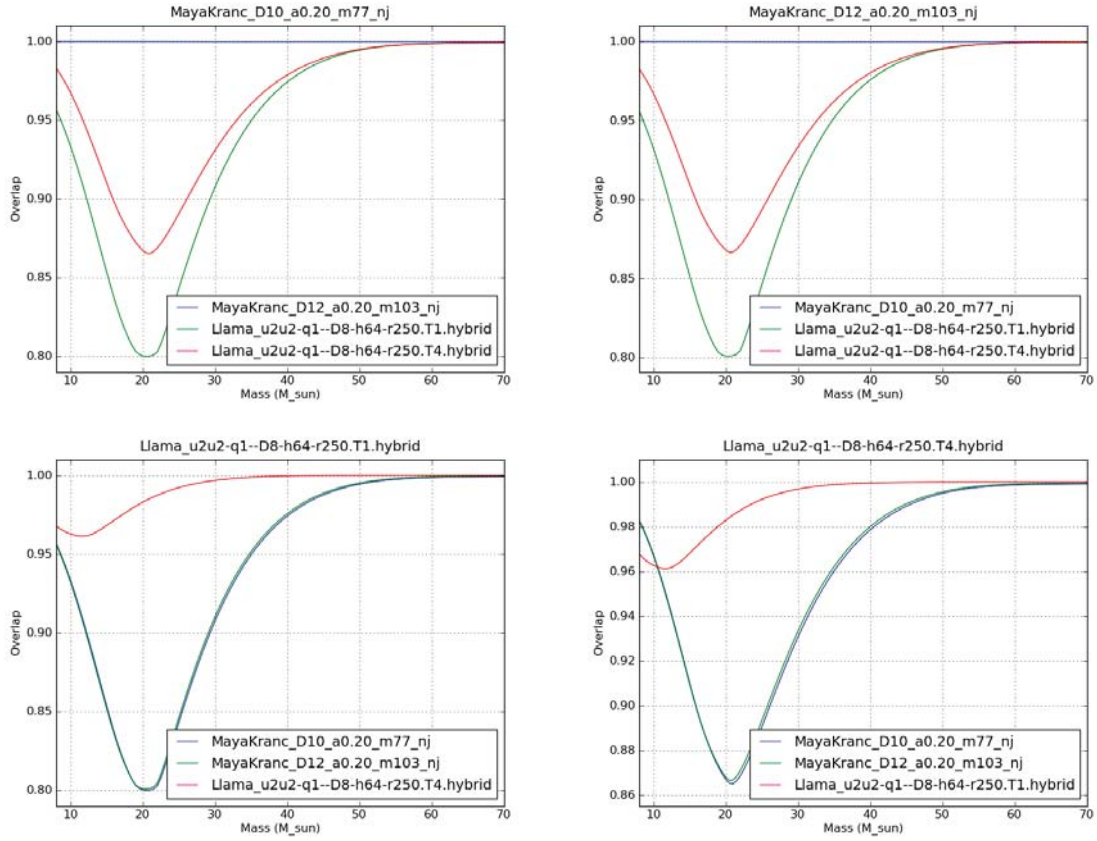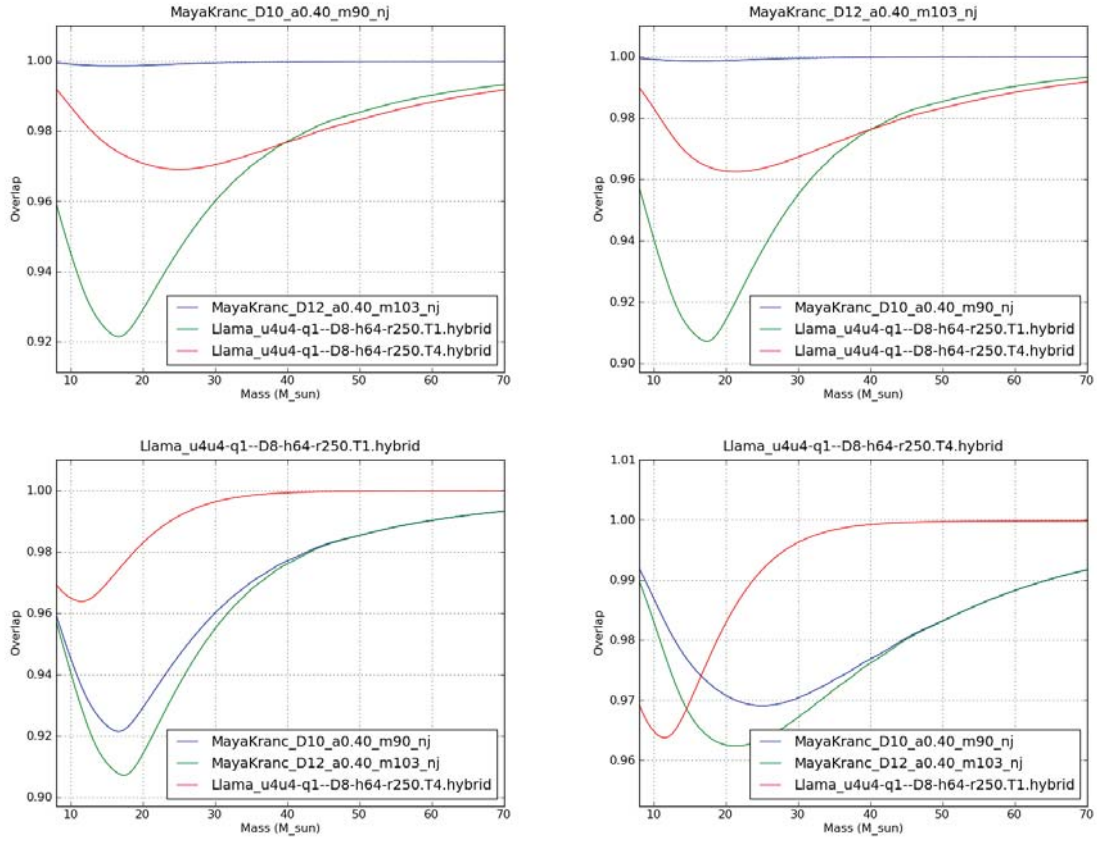
Figure 35 : Overlap plots for $q = 1$ $S_{z1} = S_{z2} = 0.4$
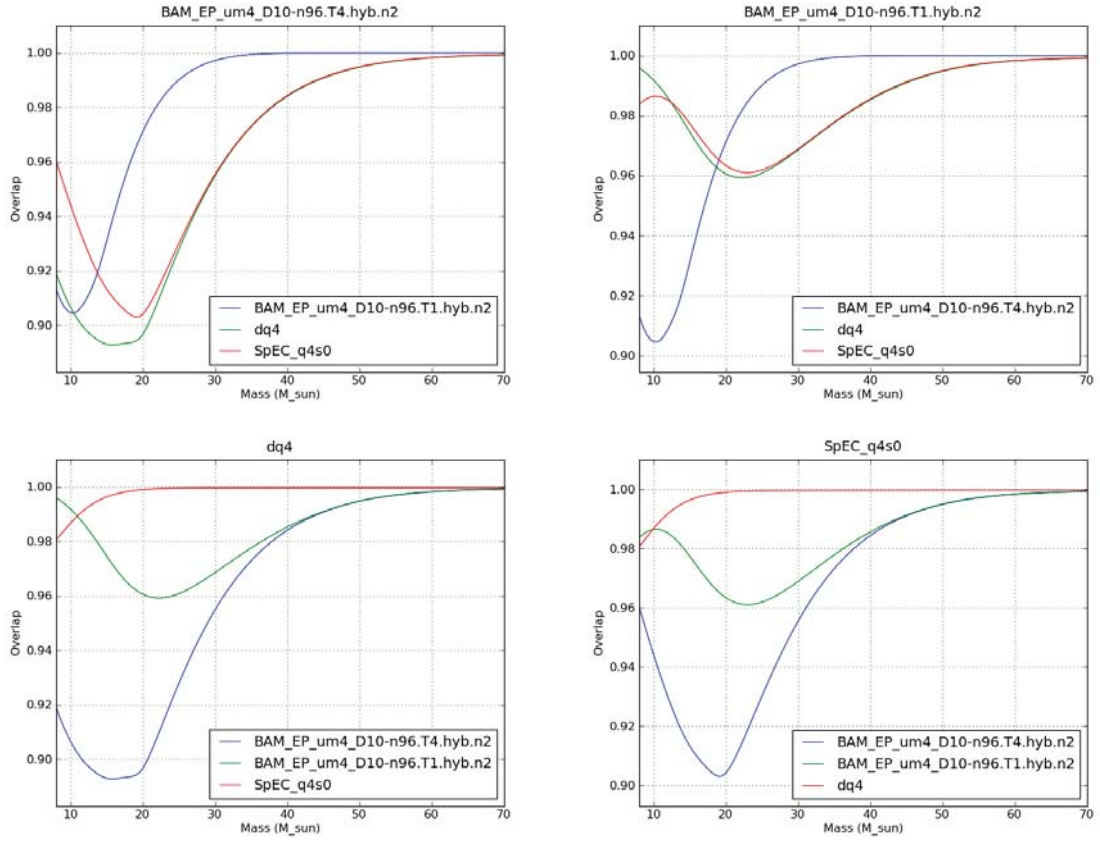
Figure 36 : Overlap plots for $q = 4$ $S_{z1} = S_{z2} = 0$
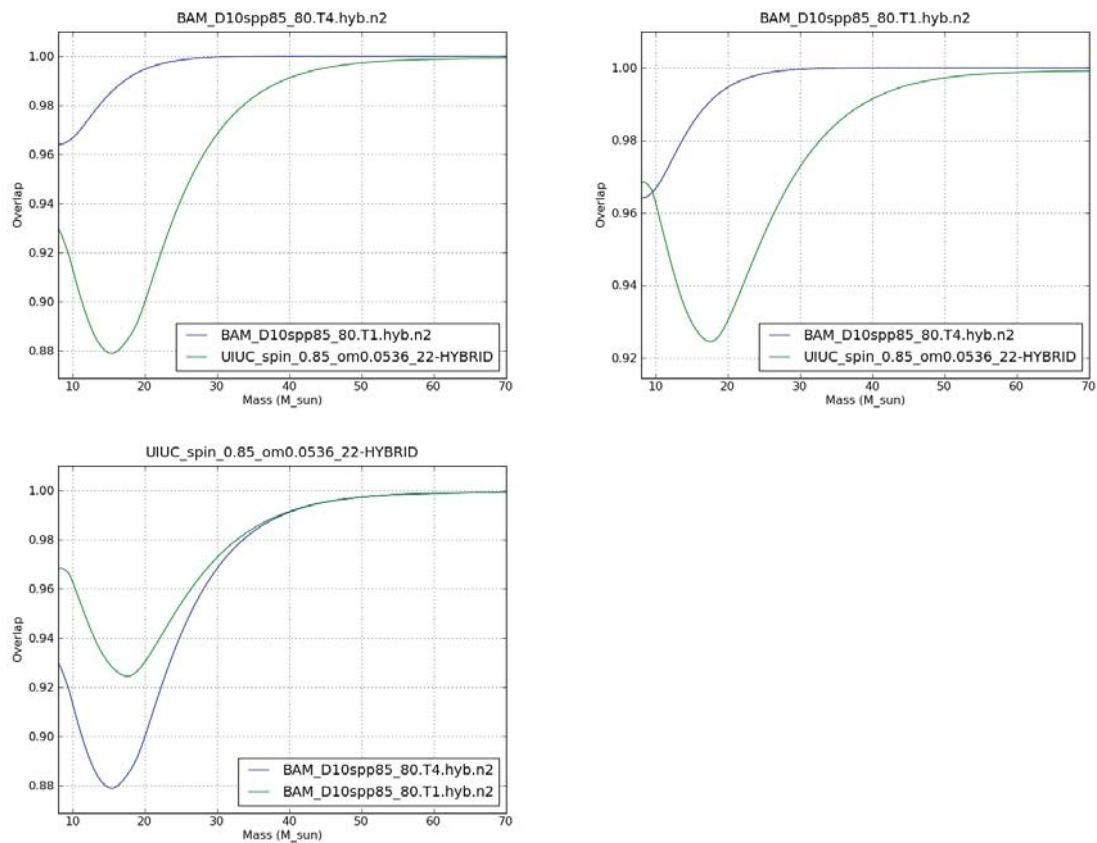
Figure 37 : Overlap plots for $q = 1$ $S_{z1} = S_{z2} = 0.85$

Chapter 8

Preliminary CBC Results from the Second NINJA project

In the previous chapter we discussed the hybrid pN/NR waveforms contributed to the NINJA-2 project, along with the studies performed to validate them. We now turn to the data analysis portion of NINJA-2, including the preliminary data sets and some results. We will be focusing on two CBC searches:

- The S6 "low mass" search, which uses TaylorF2 waveforms taken to 3.5 pN order and a bank spanning total masses from $2M_\odot$ to $25M_\odot$.

- The preliminary S6 "high mass" search, which uses EOBNR waveforms and a bank spanning $25M_\odot$ to $100M_\odot$. As of writing the tuning of this parameters in this search are ongoing. The final version of this search will be the one used for NINJA-2.

8.1 Construction of the Preliminary NINJA-2 Data Set

Thus far the NINJA-2 data sets have been used to test the codes used to construct these sets and to run preliminary sanity checks. The full run of NINJA-2 will use real detector data from S5 and VSR2. As we have noted in chapter 4, real data contains non-Gaussian features that can impact searches. These features were not captured in NINJA-1, but their inclusion is essential in order to derive realistic results.

In broad terms the plans for the NINJA-2 data sets follow those for NINJA-1

(ch. 6). Simulated Gaussian noise was generated to model the initial LIGO and Virgo noise curves, the spectra are identical to those in Fig. 12. Injection parameters, including choice of waveform, were then selected randomly. The injections were then added to the Gaussian noise and distributed to data analysis groups. However, several key changes were made in the details of this process in order to correct shortcomings in NINJA-1.

The NINJA-2 data was sampled at 16384 Hz rather than the 4096 Hz used by NINJA-1. This was done because investigations showed that there is power above 4096 Hz in the waveforms, which would get aliased down to lower frequencies if the sample rate is too low. This problem is illustrated in Fig. 38. Note that this is a different issue from the one that motivated performing the waveform overlaps at 32768 Hz, discussed in Sec. 7.2.2. The issue there was that undersampling could miss the maximum of the overlap function. The issue here is that undersampling can distort the end of the waveform due to aliasing. In the real search data is sampled at 16384 Hz, and although it is downsampled to 4096 before running the matched filters the downsampling utilizes an anti-aliasing filter that avoids this problem.

NINJA-1 consisted of only 127 injections in one day of data, which severely limited the ability to draw statistical conclusions on the behavior of the pipelines. To correct this in NINJA-2 we extended the duration to eight weeks. The density of injections was varied over this span: weeks 1-3 had one injection on average every 2000 seconds, week 4-6 had one injection on average every 14,400 seconds, and the final two weeks had one injection every on average every 216,000 seconds. The intent is that data analysts can tune and test their pipelines on the dense weeks, and then optionally perform a self-blinded test on the final two weeks.

In NINJA-1 the SNR was not chosen a priori but was determined by the other parameters. For NINJA-2 we draw the network SNR ($\sqrt{\sum_i \rho_i^2}$ where i ranges over the detectors) from a distribution and then scale the distance of the injection in order to achieve that SNR. For the first three weeks the distribution is linear from 6 to 130 in order to allow pipelines to test and tune out to large SNRs on the densest set of injections. For the remaining weeks the distribution falls as the reciprocal of the network SNR (uniform in log(SNR)) in order to better model the expected astrophysical distribution.
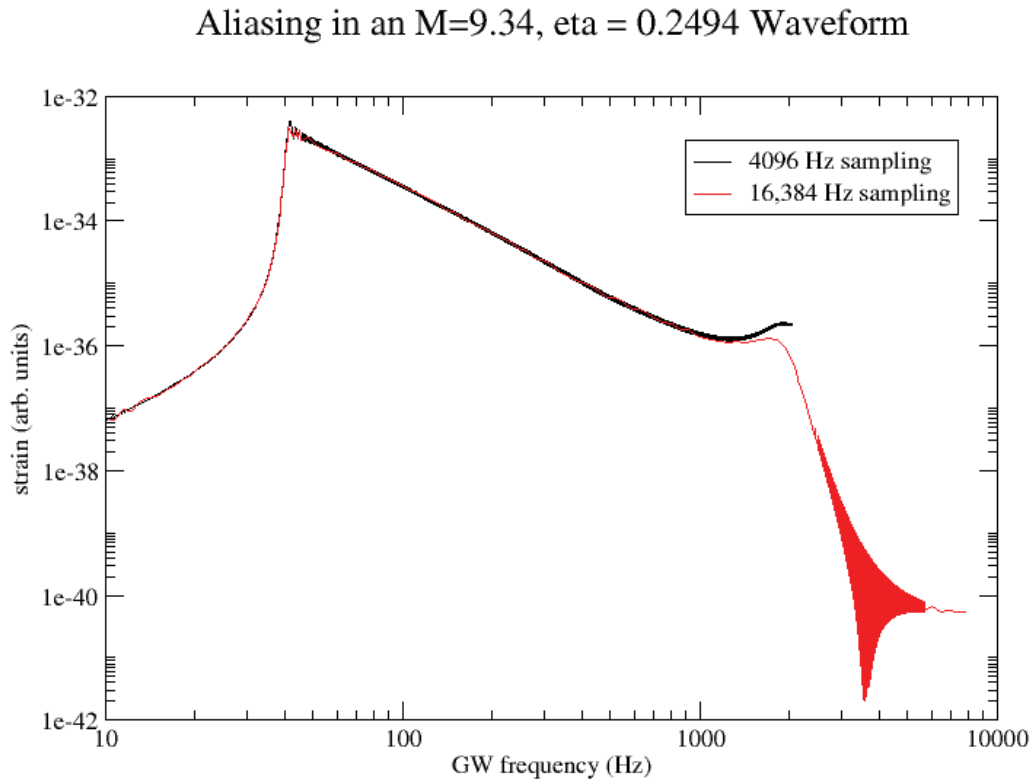
Figure 38 : Frequency-domain amplitudes of a NINJA-2 waveform at different sampling rates. At a sample rate of 4096 Hz the late portion of the waveform are distorted due to aliasing of power to lower frequencies.

The mass and waveform selection were also performed slightly differently in NINJA-2. For each injection a mass was first selected uniformly over the specified range; for the full 2-month run this range is from $10 - 350 M_\odot$. Then waveforms were selected at random until one was found that could be injected at the chosen mass such that the waveform turns on below 35 Hz. In practice this condition never caused any waveform to be rejected, as all submitted waveforms were long enough to be injected down to the lowest mass in the range. The mass ratio and spins are intrinsic to the waveforms, so choosing a submission amounts to a choice of these parameters as well. As in NINJA-1 the sky location and inclination were chosen uniformly at random.

Four NINJA-2 data sets were released:

- A test set consisting of one week of data was released on May 13, 2010. This included three separate sets spanning different mass regions; low-mass ($10 M_\odot - 40 M_\odot$), high-mass ($35 M_\odot - 100 M_\odot$) and a burst/ringdown set ($80 M_\odot - 350 M_\odot$). This purpose of this run was to shake out bugs in the injection code and waveforms, several of which were found.

- A second test set consisting of one week of data for each of the three mass bins was released on May 31, 2010. This was meant to test the fixes implemented after the first test set, and to do more careful sanity checks. Some results from this run are discussed in Sec. 8.2.

- Based on positive results from analyses of the second test week, full two-month data sets for all three mass bins were released on June 9, 2010. Unfortunately shortly after release a remaining major bug in the injection software was discovered. Data is stored in frame files spanning 4096 seconds. When an injection crossed frame boundaries this bug would cause the portion contained in the earlier frame to be omitted. This happened in enough cases to invalidate the entire data set.

- There was then a lengthy pause, in large part due to the LIGO-Virgo "blind injection challenge" (see Sec. 10.5). However, it was during this time that the validations discussed in the previous chapter were performed. It was also in this gap that studies discovered the need to move to 16 kHz sampling. This quadrupled the size of the data and made the release of 3 separate mass ranges

unfeasable. Therefore a single set spanning two months and containing injections from $10M_\odot - 350M_\odot$ was released on June 13, 2011. Analysis on this set has begun and some preliminary results will be discussed in Sec. 8.3

There are two motivations for constructing and distributing the data sets as NINJA-1 and NINJA-2 thus far have done. The first is to ensure that every group is looking at the same set of injections so that results can be compared. The second is due to the terms of the NINJA agreement, which restricted distribution of the raw NR waveforms. However, distribution of such static sets limits the ability of individual groups to tune their pipelines in optimal ways, and conceptually distributing a set of parameters would be sufficient to compare results across pipelines. In addition the size of the data sets makes distribution slow and complex. The NR groups within NINJA have therefore relaxed the conditions on their use of their waveforms. Consequently, subsequent NINJA-2 data sets will be distributed as sets of parameters, and data analysis groups will use the available code to either create data sets locally, or perform the injections "on the fly" as the analysis is performed. This will also allow groups to do special-purpose tuning runs or analyses by injecting specialized sets of injections into noise. The results of these studies may be published as short-author papers subject to the conditions of the NINJA agreement.

So far we have used simulated Gaussian noise, however as we have noted this is not a good model for the real detectors, which contain many "glitches" caused by both environmental factors and transient behavior of the detectors. These glitches produce a population of background triggers. In testing pipelines a critical issue is the ability to distinguish such background triggers from signals. If NINJA is to be able to to make definitive statements about the behavior of pipelines it is therefore imperative to use real detector noise. A memorandum of understanding has been signed between the NINJA collaboration and the LIGO and Virgo collaborations allowing the use of real noise from the 5th LIGO and first Virgo science runs. There are a number of technical issues to be resolved before this can be done, however the key results from NINJA-2 will come from on-the-fly injections into real noise.
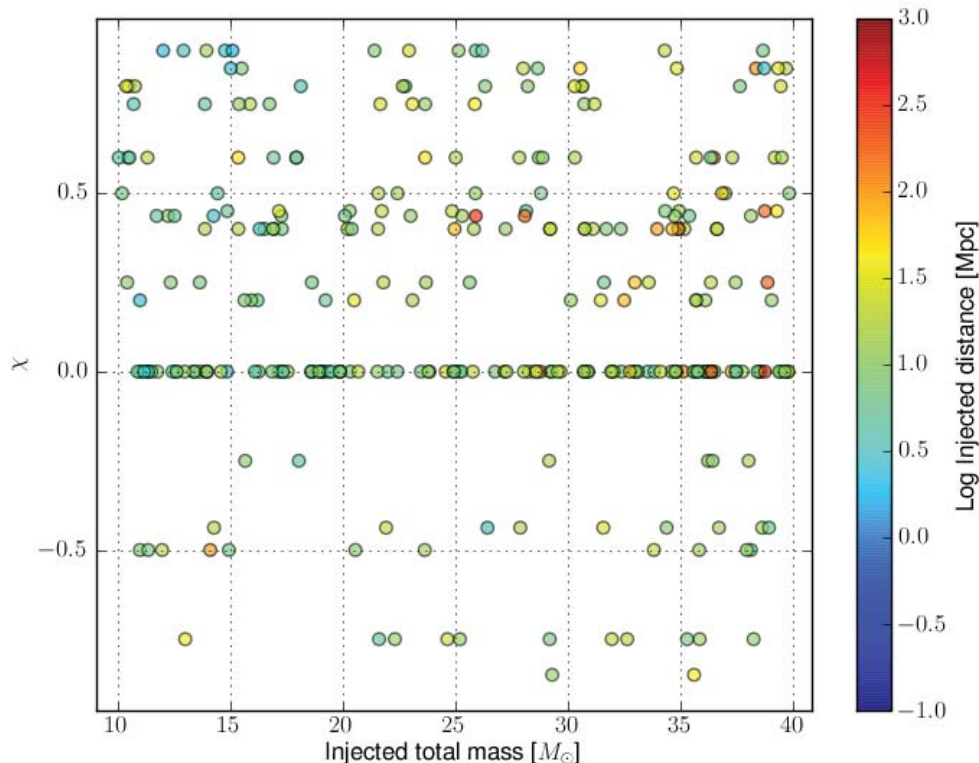
Figure 39 : Distribution of mass, spin and distance parameters in the one-week, Gaussian-noise test data set. Spin is reported as $\chi = (S_1/m_1 + S_2/m_2)/(m_1 + m_2)$

8.2 CBC Results from the Second One-week Data Set

As noted above, two one-week data sets were produced for testing the injection code and performing sanity checks on the results. Here we present the results of running the standard low-mass CBC pipeline on the low-mass data set from the second of these runs. As this was largely a testing and debugging exercise no effort was made to draw conclusions about the pipelines. In particular, no runs were performed with a bank extended to unphysical η or the other pipeline modifications suggested by the studies in chapter 5 and tested in NINJA-1. The injection parameters were chosen as described above, the selected masses, spins and distances are shown in Fig. 39.

The S6 version of the standard CBC low-mass pipeline was run over this data set. This version of the pipeline uses *TaylorF2* stationary-phase templates taken
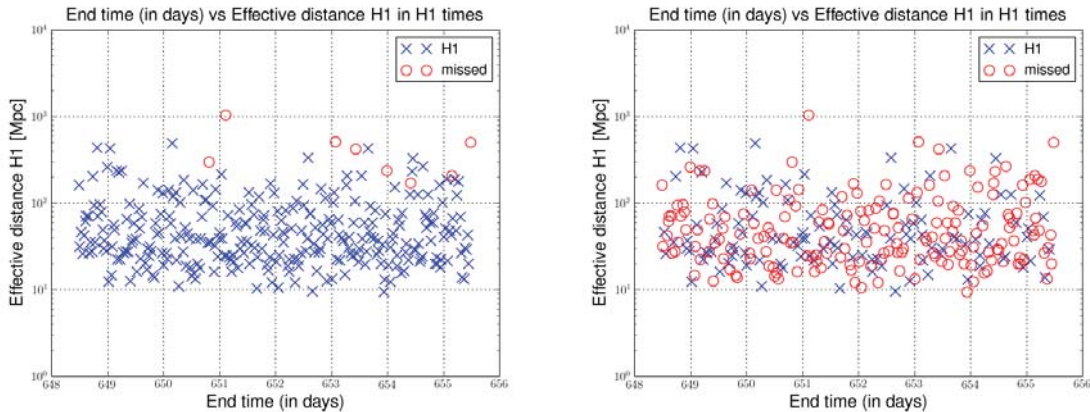
Figure 40 :   First-stage found/missed plots from the test data set. On the left, the results from the data set containing the injections, on the right the results from running on the noise-only data. Many signals appear in both, indicating that they are not really found, but only that there is a random background trigger within the time window.

to 3.5 pN order in phase evolution (see Appendix A). The first result of interest is the found/missed plots at the first stage, before the χ^2 test or coincidence between detectors has been applied. The result for the H1 detector is shown in Fig. 40. This can give some indication of how well the waveforms and the bank capture the signals. However, it can also be misleading. The pipeline considers an injection to be "found" if there is a trigger within 100 ms of the injection time. No parameter matching is required. This means a quiet trigger resulting from the noise can be mistaken for finding the injection. To see how often this can happen we also ran the analysis on frames containing only the noise, without any injected signals. This is also shown in Fig. 40, and indeed many of the reportedly-found injections can be seen to be coming from the noise. Both χ^2 and coincidence will cut down the number of false reports, in Fig. 41 we show the found/missed plots for all three detectors after the coincidence test, and many of the triggers coming from the background have been removed. At the second stage the results are sensible, the ability of the pipeline to recover the injections falls off as the effective distance increases. There are however a few close missed injections that would warrant follow up study in a full search.

In Fig. 42 we plot the SNR recovered by the pipeline versus the injected value. There is a distinct pattern exhibited, for low-mass signals the injected and recovered
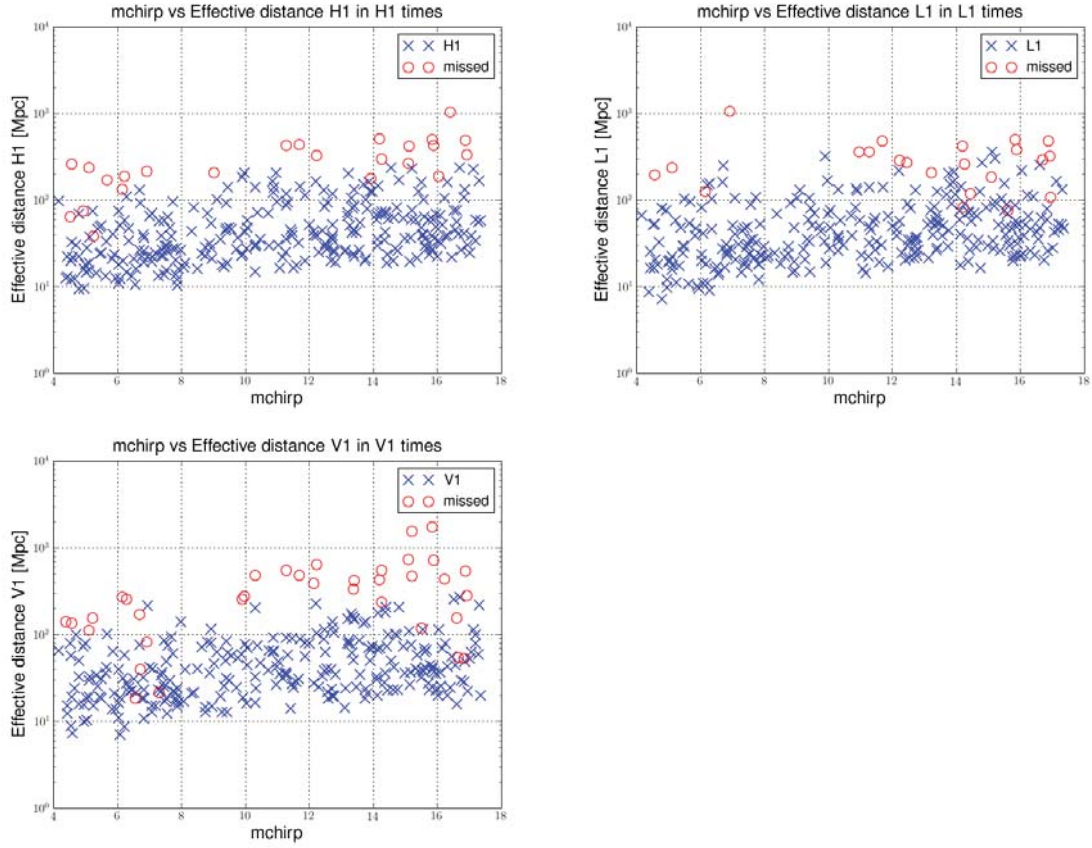
Figure 41 : Second-stage found/missed plots for the test data set, for all three detectors. Each plot shows signals that were reovered in that detector and at least one other. Requiring coincidence removes the background triggers seen in Fig. 40. The results are sensible: closer injections are more likely to be found than distant ones.
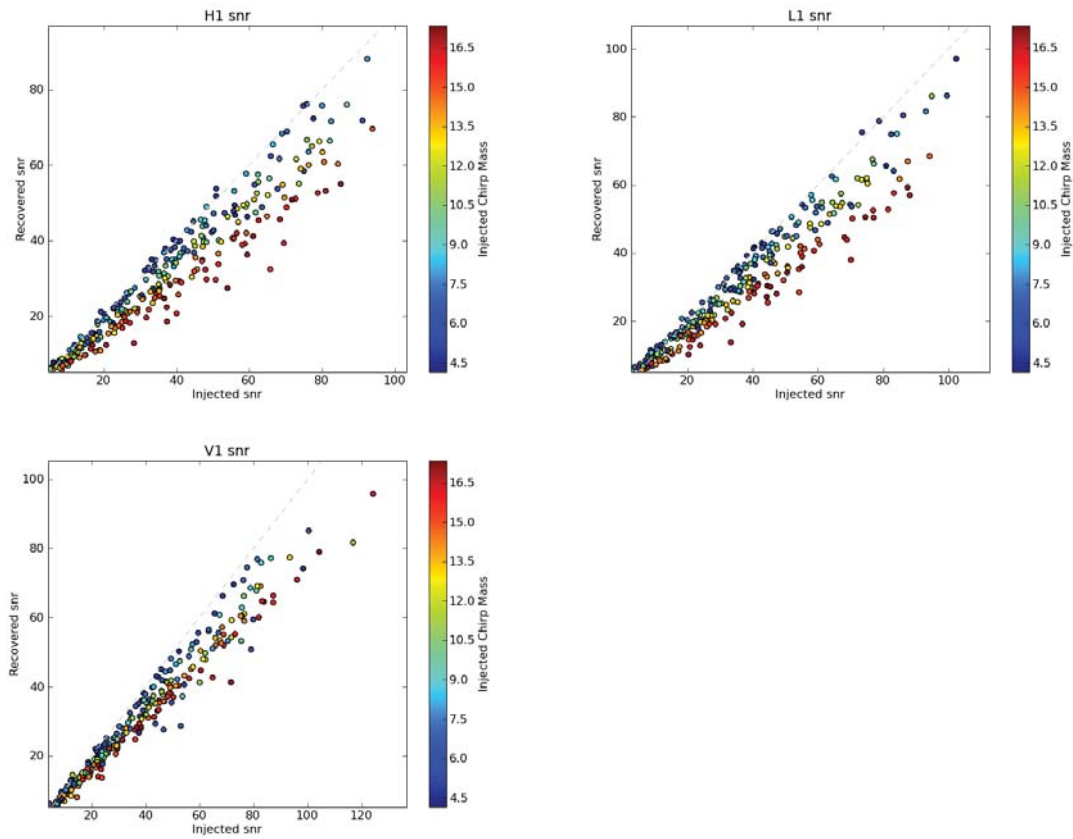
Figure 42 : SNR recovery for the test data set in all three detectors. The injected SNR is calculated from the entire waveform, the recovered SNR is calculated only from the inspiral up to the ISCO frequency. At higher masses this loses SNR as the late inspiral, merger, and ringdown pass through the detector sensitive bands. Note that V1, which has lower noise at low frequencies, recovers somewhat more SNR for the high-mass systems.

SNR match, but the recovered SNR drops off with increasing mass. This occurs because the injected SNR value is calculated using the entire waveform. By contrast the recovered value terminates the integration at the ISCO frequency. For higher-mass systems this means the integration cuts off in or before the sensitive band, while there is still power in the signal, and consequently the SNR is underestimated. While we can not hope to capture the full merger and ringdown with inspiral-only templates, this again confirms the results of chapter 5 and NINJA-1, which show that extending the templates to high frequencies can increase the SNR.

At the level of investigation performed the results from this data set appear reasonable. Although conversely this means that the waveform truncation issue discussed

above was not caught by this investigation. However, this and other analyses did indicate that there were no other more serious bugs in the injection software, which enabled us to move to the full two-month set.

8.3   CBC Results from the Two-month Data Set

We now turn to the second two-month data set in Gaussian noise. This is the latest set constructed, and includes many corrections and changes from earlier sets:

- The data is now sampled at 16,384 Hz.

- There is now only one set spanning the full mass range from $10 M_\odot - 350 M_\odot$.

- The waveforms have been updated, in particular many of the hybridizations were redone.

Apart from the change in mass range the injection parameters were chosen as described above. The masses, spins, and distances chosen are shown in Fig. 43.

This data set was analyzed with the standard CBC low-mass and high-mass pipelines. The parameters were exactly as in the S6/VSR2,3 runs, no changes were made to the configurations except for those relating to the names of the data files. The low-mass search uses Taylor F2 templates to 3.5 pN order in phase evolution (Sec. 2.5) in a mass region defined by minimum component masses of $1 M_\odot$ and maximum total mass of $25 M_\odot$. The high mass search uses EOBNR templates (Sec. 2.5.1) in a region defined by minimum component masses of $1 M_\odot$, minimum total mass of $25 M_\odot$, and maximum total mass of $100 M_\odot$.

Figure 44 shows the found/missed plots after the first stage in the high-mass search. As expected, distant signals are less likely to be found than close ones. As discussed above the loose coincidence test between injections and triggers means that many of these injections may not really be found. We therefore look at the second-stage results in Fig. 45, and as in Fig. 40 there are many fewer found injections.

The fraction of injections found decreases with increasing mass. The bank of the high-mass search extends only to $100 M_\odot$, but the injections extend up to $350 M_\odot$, so this result is not surprising. We can quantify the effect by plotting the efficiency, defined as this fraction, as a function of mass. These plots are shown in Fig. 46. While
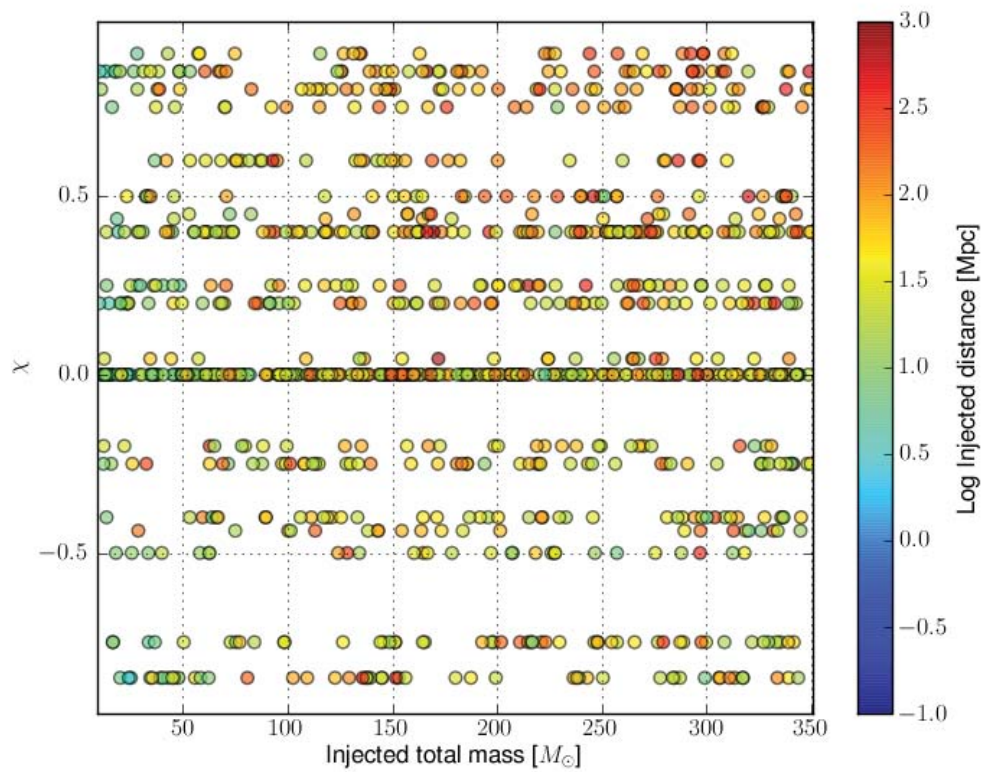
Figure 43 : Distribution of mass, spin and distance parameters in the two-month, Gaussian-noise data set.
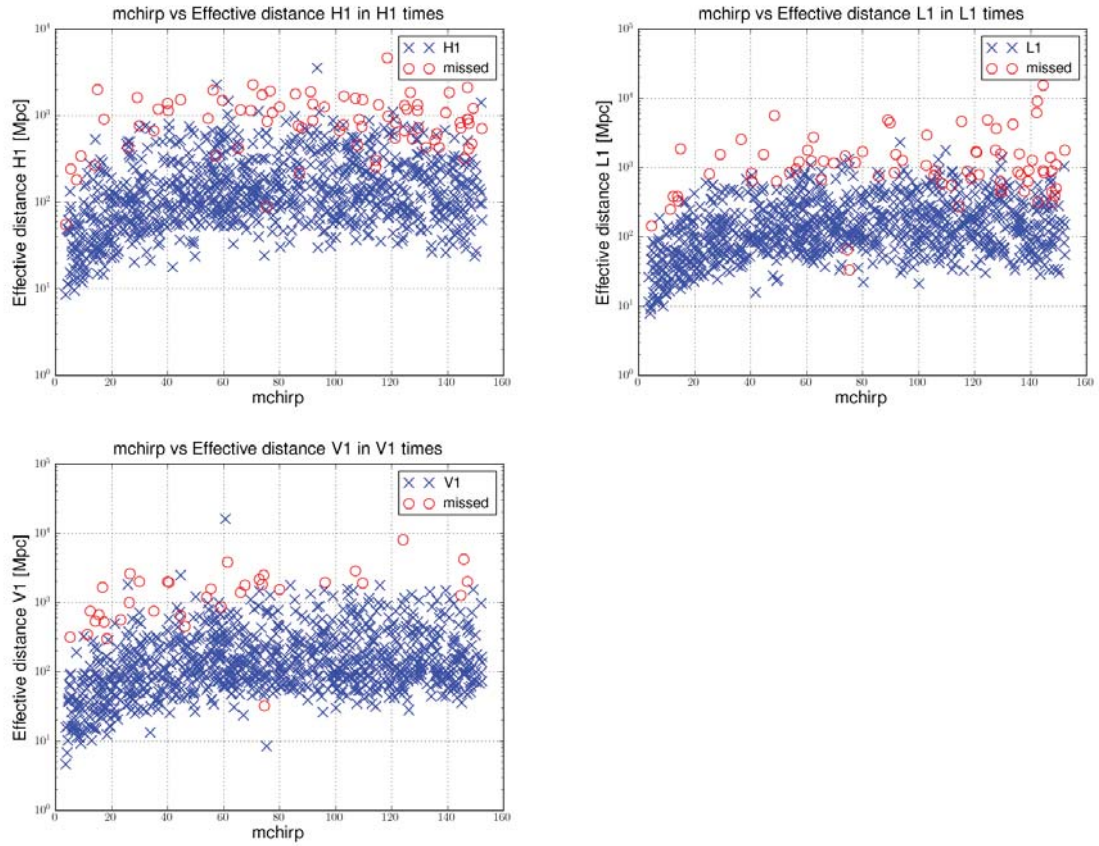
Figure 44 :   Preliminary first stage found/missed plots from the high-mass search. The behavior is as expected: distant signals are more likely to be missed than close ones.

Figure 45 : Preliminary second stage (after coincidence and χ^2) found/missed plots from the high-mass search. As expected, many of the signals reported as "found" after the first stage were due to background triggers within 100 ms of the injection. More signals are missed at higher masses because the high-mass bank only extends to $100M_\odot$ but there are injections up to $350M_\odot$. There is an anomaly in the Virgo results, see the text for discussion.

Figure 46 : Efficiencies of the high-mass pipeline as a function of mass. As can be seen in figure 45 the efficiencies decrease as mass increases. However, the scatter in these plots indicates that there are not enough injections to draw definitive conclusions.

these plots show the general trend, they are quite jagged, indicating that we do not have enough injections to draw statistical conclusions. Follow up work will remedy this by performing many thousands of "on the fly" injections. There is an unexpected feature in the V1 plot, which shows that at high mass more close injections are missed than distant ones. We will return to this issue below.

The standard CBC low-mass and high-mass searches both use non-spinning templates, and an important question is the ability of these pipelines to detect spinning signals. NINJA-2 is uniquely positioned to help answer this question, and we begin by plotting the recovery efficiency of the high-mass pipeline as a function of the spin parameter χ in Fig. 47. There is a general trend suggesting that the efficiency increases with χ, and in particular that the search performs worse on anti-aligned

Figure 47 : Efficiencies of the high-mass pipeline as a function of the spin parameter χ. There is a general trend suggesting the pipeline is more efficient at detecting aligned spins, but more injections are needed to verify and quantify this.

systems than aligned systems. Again, the scatter of the plots makes it impossible to draw definitive conclusions, this will be remedied in follow up studies with more injections.

We now consider the parallel analysis performed with the low-mass search. Figure 48 shows the found/missed plots after the first stage. Again, distant signals are more likely to be missed than close ones. The low-mass template bank extends to a total mass of $25M_{\odot}$, and correspondingly the efficiency decreases notably above a chirp mass of $\sim 80M_{\odot}$. The anomalous behavior in Virgo is more pronounced here than in the high-mass search, with most injections above a chirp mass of $100M_{\odot}$ and closer than 100 Mpc being missed. We return to this issue in Sec. 8.3.1. Figure 49 shows the found/missed plots after the second stage. Again, many of the injections
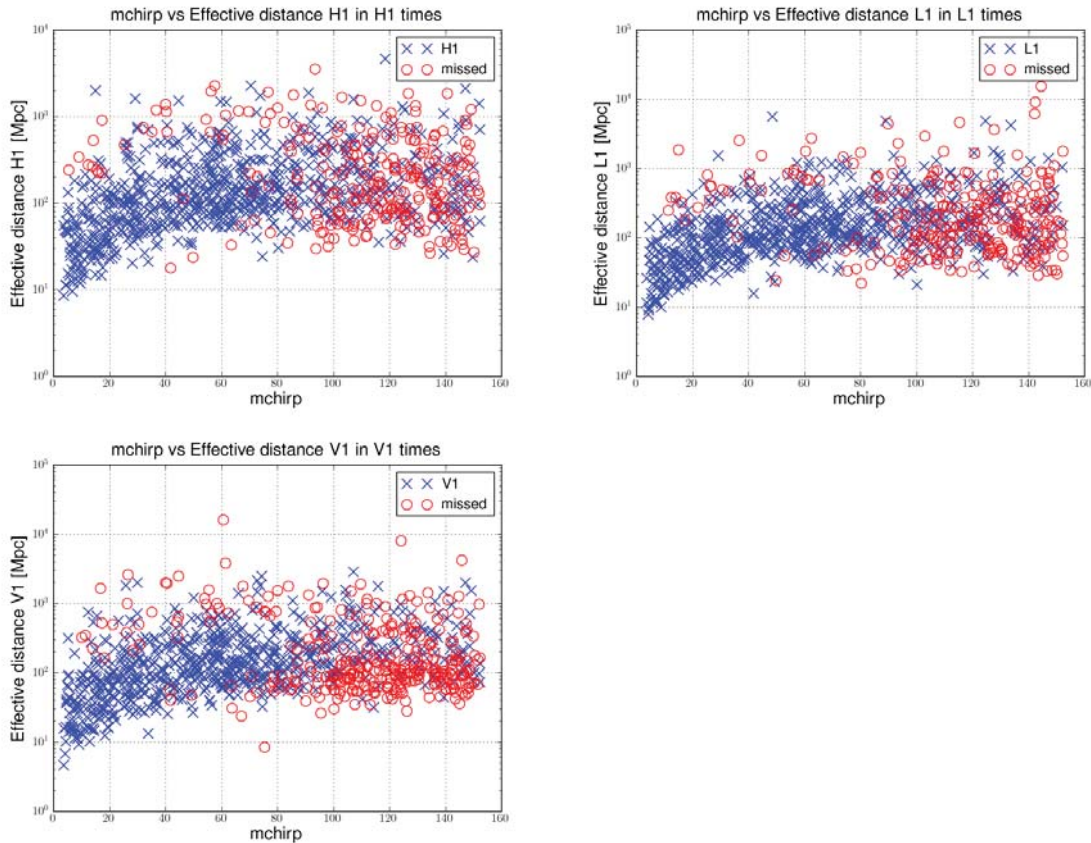
Figure 48 : Preliminary first stage found/missed plots from the low-mass search. The behavior is as expected: distant signals are more likely to be missed than close ones.
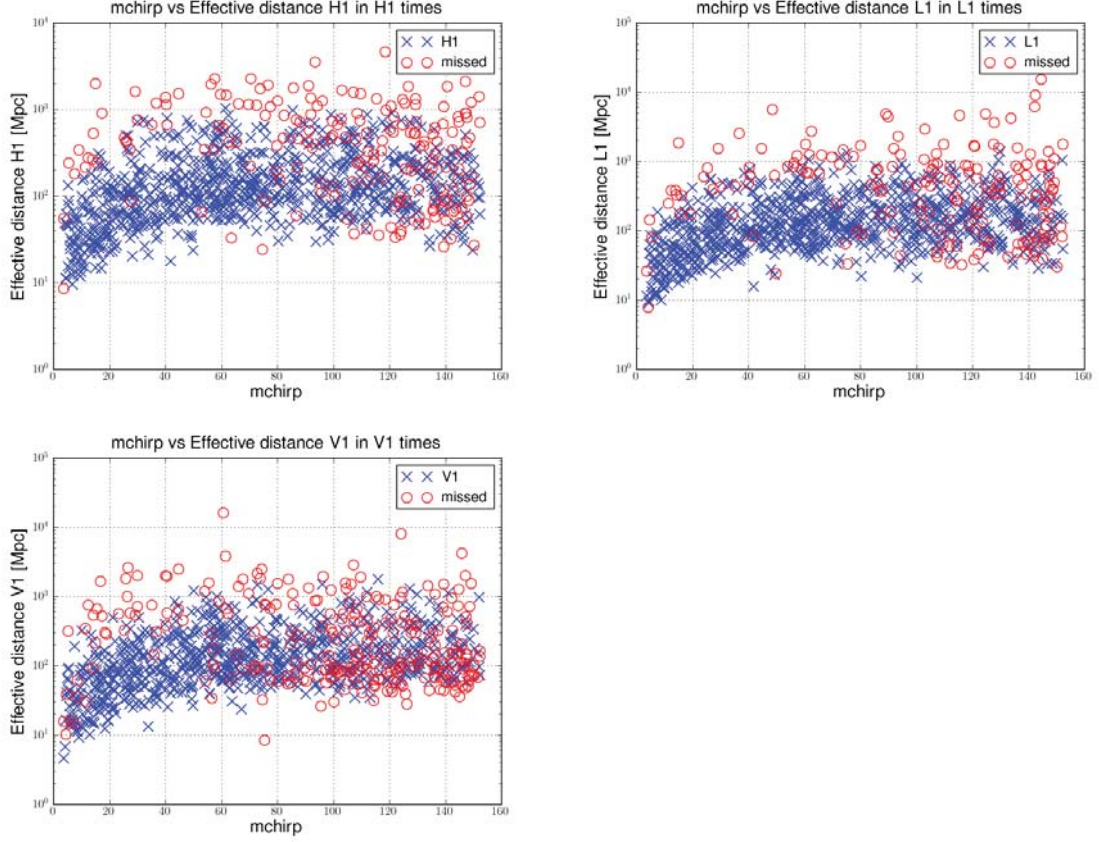
reported as found in the first stage are now reported as missed. The remaining found injections are clustered towards the low end of the mass range, as expected. However, in all detectors at this stage the remaining found injections at higher masses tend to be at farther distances, contrary to the expected behavior. In addition, the injections found in V1 are mostly confined to the region $M_{chirp} < 20M_{\odot}$, whereas those in H1 and L1 extend up to $\sim 60M_{\odot}$. This indicates that most of triggers above $20M_{\odot}$ are found in coincidence between H1 and L1. In turn this implies that signals above this point are either not really found in V1 at the first stage, or are seen with very different parameters that fail the coincident test with the H1 and L1 triggers. Again we defer further discussion to Sec. 8.3.1.
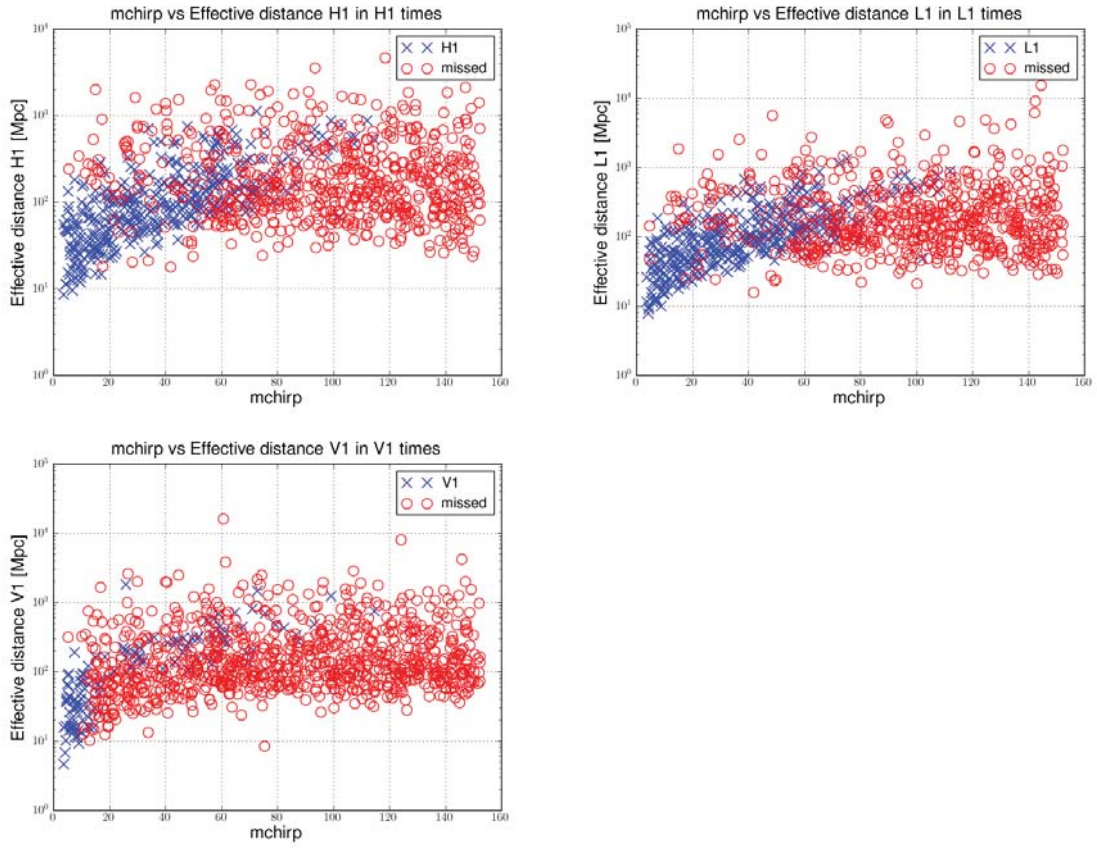
Figure 49 : Preliminary second stage found/missed plots from the low-mass search. Virgo appears to only find injections below a chirp mass of $20M_\odot$, all found injections above this point come from coincidences between Hanford and Livingston. In all detectors injections found at higher masses have larger effective distances. Further studies are needed to understand this behavior, see Sec. 8.3.1.
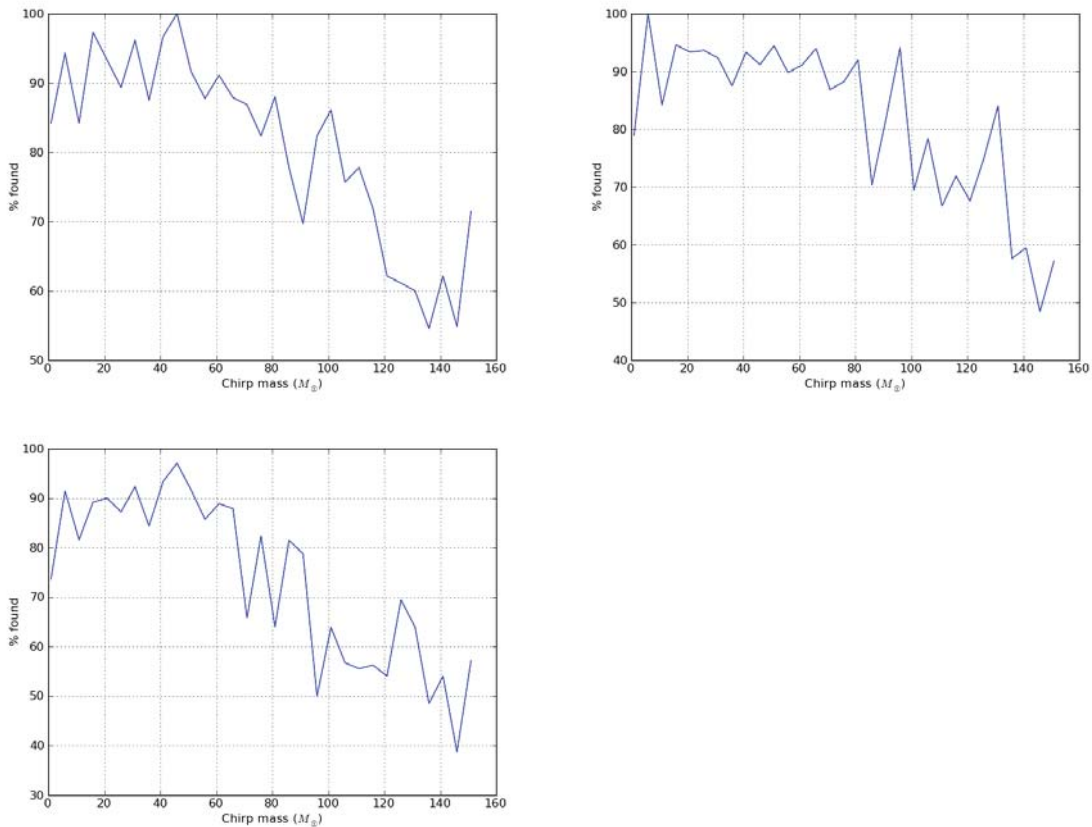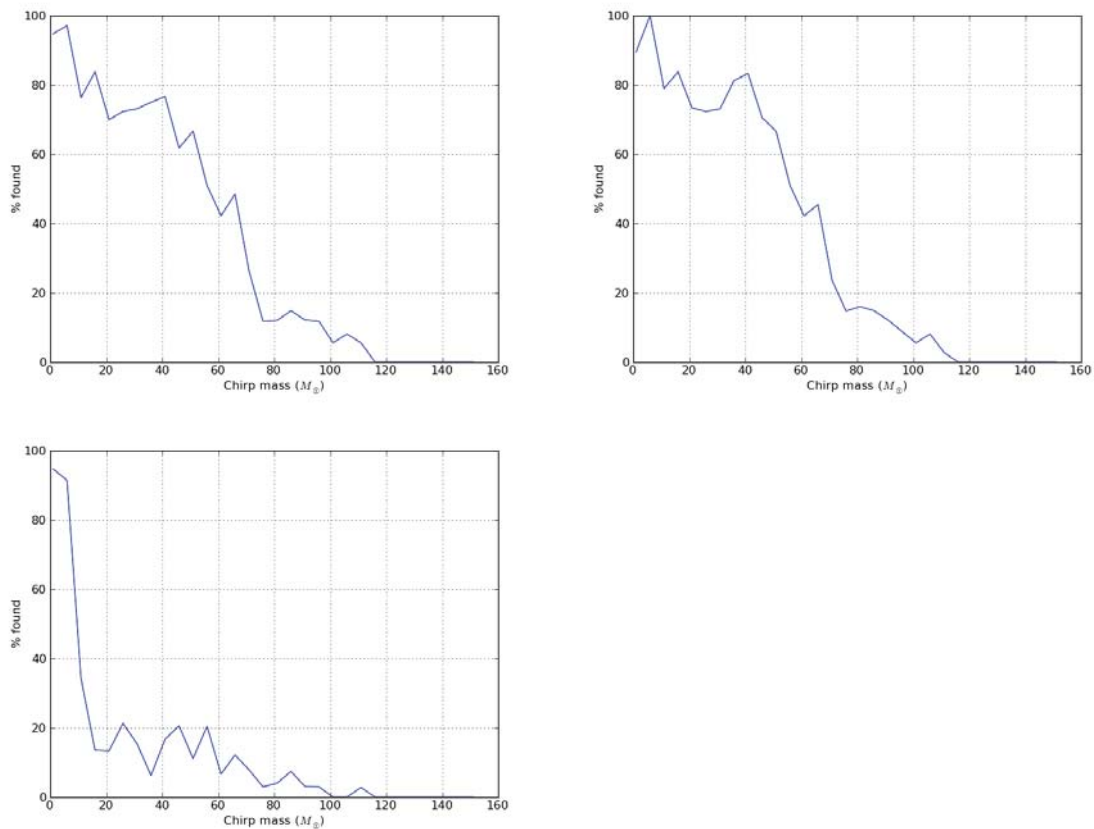
Figure 50 : Efficiencies of the low-mass pipeline as a function of mass. As can be seen in Fig. 49 the efficiencies decrease as mass increases. However, the scatter in these plots indicates that there are not enough injections to draw definitive conclusions.

The efficiencies of the pipeline as a function of mass is shown in Fig. 50, and as a function of spin in Fig. 51. The mass plots present the same information found in Fig. 49 in an alternate way, and as in the corresponding high-mass plots indicate that more injections are needed in order to refine the results. The spin plots show the same general trend in the high-mass plots, again indicating that the search is more efficient at detecting systems with aligned spins than anti-aligned.

8.3.1 Anomalous Virgo Results

In both the high-mass and low-mass searches we find that at higher chirp masses closer injections are missed while farther injections are found. This problem is evident in all three interferometers, but it is most extreme in Virgo and affects even the first-stage
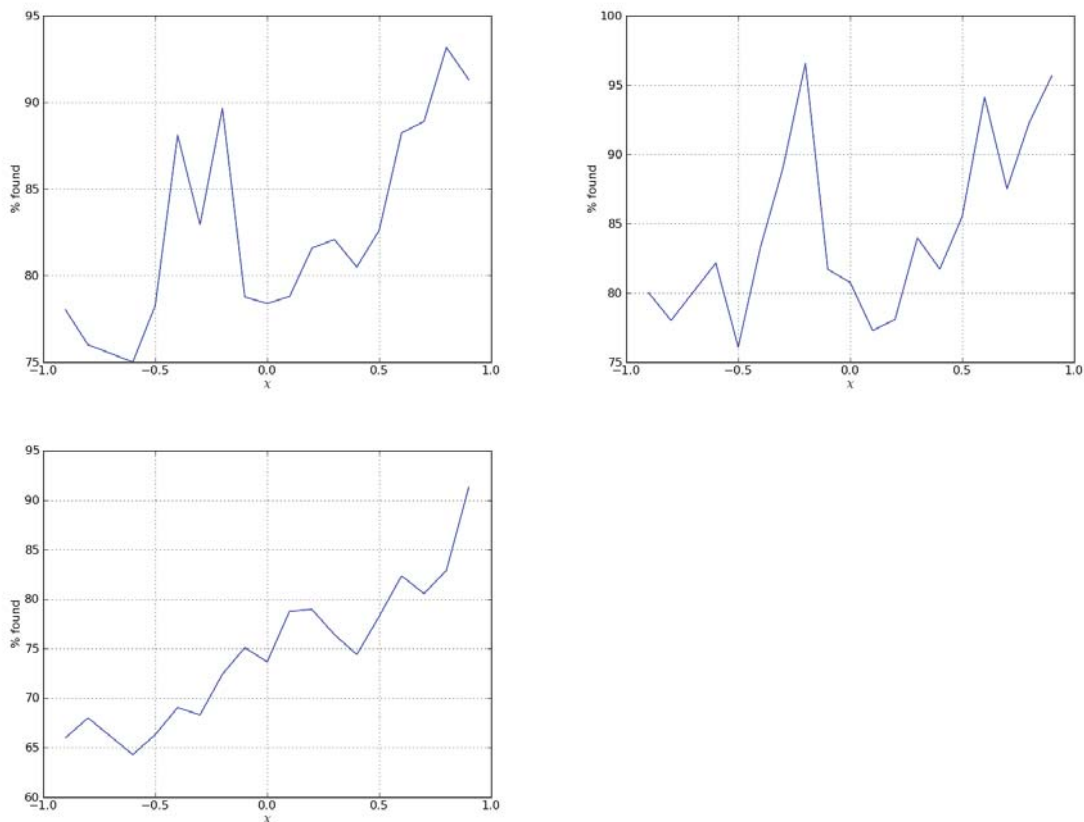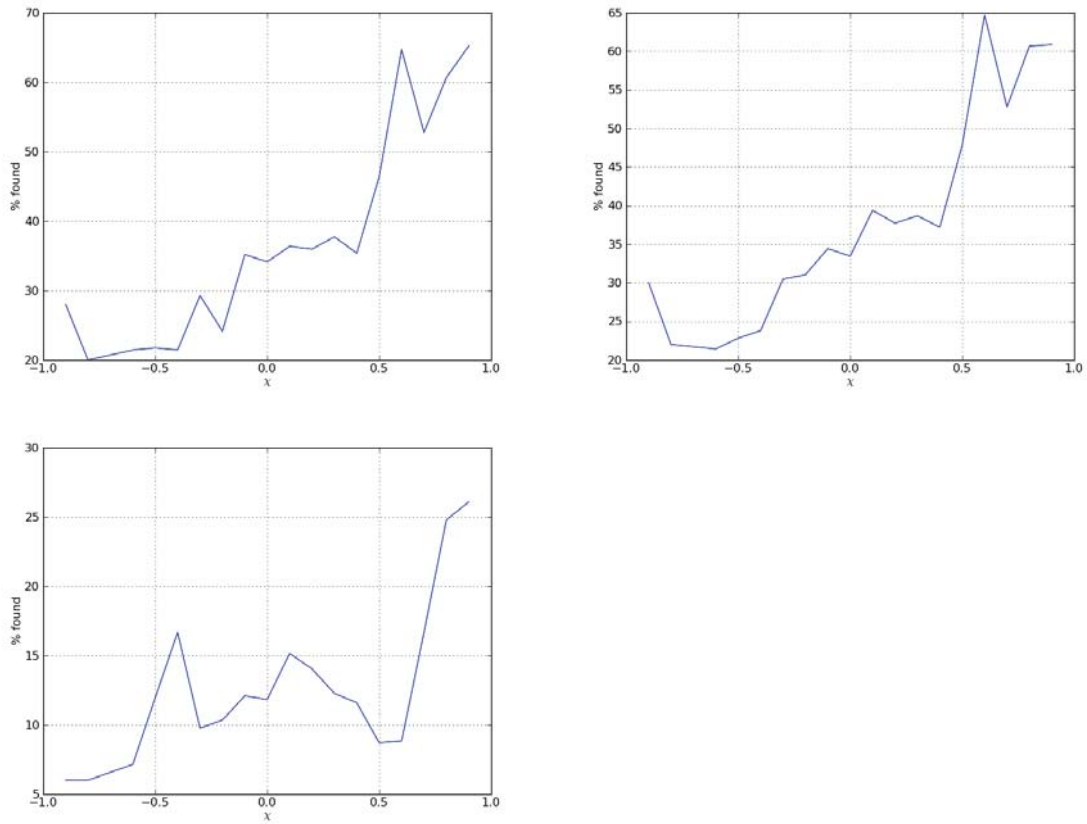
Figure 51 : Efficiencies of the low-mass pipeline as a function of the spin parameter χ. There is a general trend suggesting the pipeline is more efficient at detecting aligned spins, but more injections are needed to verify and quantify this.
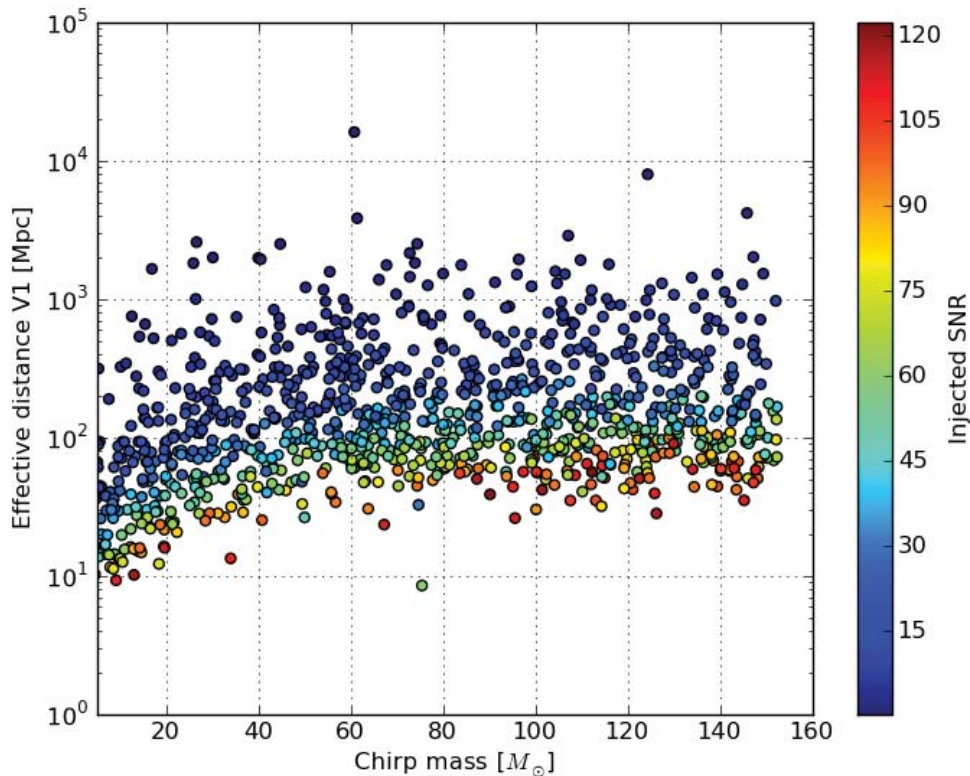
Figure 52 : Injections plotted as in the found/missed plots, color-coded by SNR. If high SNRs corresponded to points found in Virgo it would indicate that the effective distances were being miscalculated. However, this is not the case.

results. At the time of writing no explanation for this behavior has yet been found, however some possibilities have been ruled out. As the problem is most notable in Virgo we focus attention there.

First, we consider the possibility that the effective distance is simply being misreported. We check this by plotting the injections on the same axes as the found/missed plots, color-coding by injected SNR. Fig. 52 shows that the SNR and effective distance, which are calculated by different portions of the code, correlate. This rules out this possibility.

We next consider the possibility of a correlation between effective distance and total mass. These values are chosen randomly and independently, but if there were some correlation such that high effective distances corresponded to lower total masses
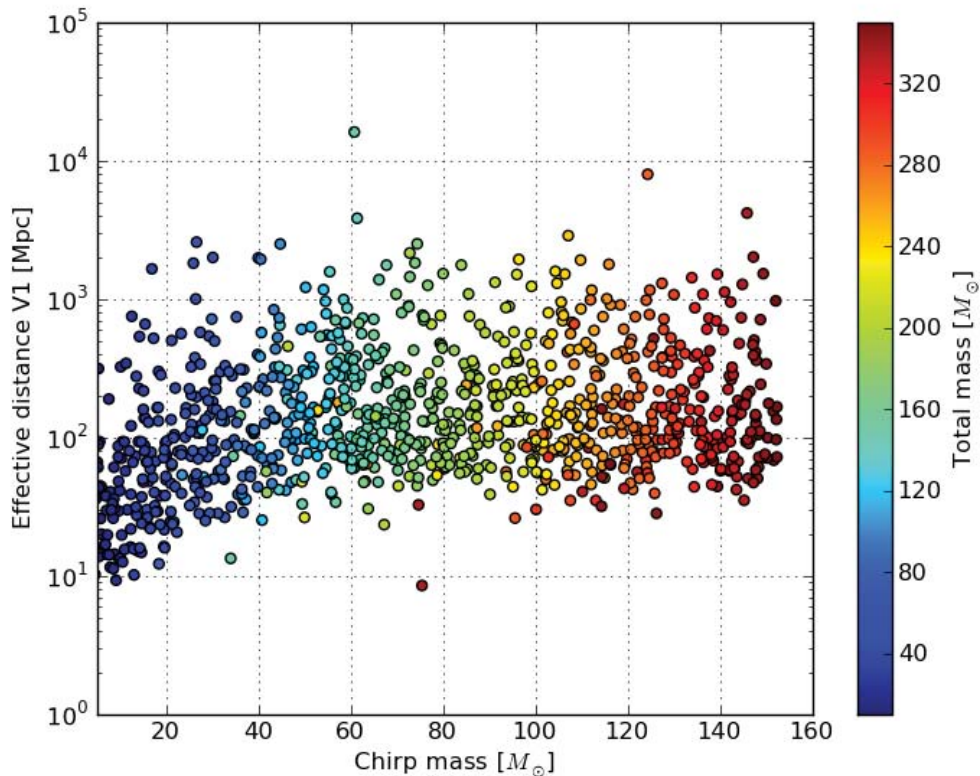
Figure 53 : Injections plotted as in the found/missed plots, color-coded by total mass. If lower masses corresponded to points found in Virgo it would indicate a correlation between effective distance and mass. However, this is not the case.

it would explain the results by indicating that more distant injections happened to be more likely to fall within the range of the bank. We again plot the injections as in the found/missed plots, now color-coding by total mass, in Fig. 53. Again, no correlations are seen.

Finally, we examine in detail one of the injections with chirp mass above $100 M_\odot$ that is found in Virgo at the first stage beyond 100 Mpc, looking for any indication that the time series are badly behaved. Figure 54 shows the injection, the filtered data, and the SNR time series arranged so that they all cover the same two-second interval. All three line up as expected, at the level of this test this injection seems to be found legitimately.
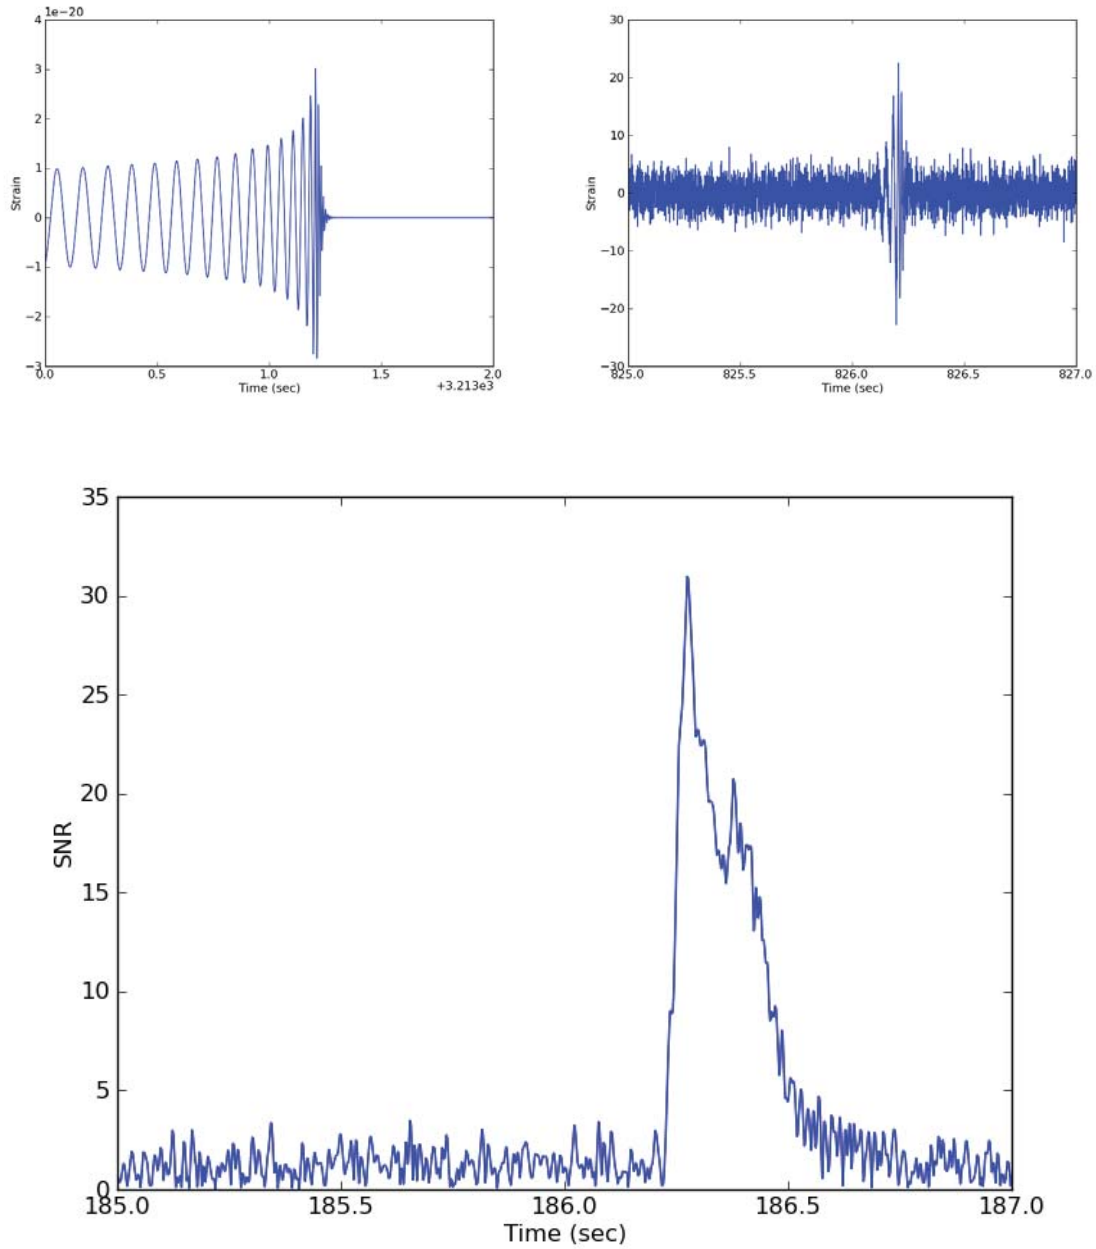
Figure 54 :   Time series plots of an injection found in Virgo, chirp mass $135 M_{\odot}$ at 157 Mpc. Top left shows the scaled waveform. Top right shows the data segment passed to the matched filter, after bandpassing to reduce low-frequency noise. Bottom shows the SNR time series. All plots line up as expected.

8.4 Open Questions for NINJA-2

Beyond resolving the unexplained behavior in the searches we intend to use NINJA-2 to address several important open questions. We briefly note these here.

Further tests of the pipeline modifications

We intend to rerun these analyses using a bank extended to unphysical η and terminating the waveforms at the WRD frequency, as in NINJA-1. The goal is to determine how this affects the efficiency of the search at lower masses than could be tested in NINJA-1, as well as on a wider range of parameter space. We will also repeat the analysis in real detector noise. A critical question is how the extended bank will affect the rate of background triggers. NINJA-1 indicates we can recover higher SNRs by making this change, but if this comes with an elevated background, making signals stand out less strongly, then this change will not be useful.

Transition between the low-mass and high-mass searches

The dividing line between the low- and high-mass searches is somewhat arbitrary at present. We hope to use efficiency plots such as those presented in this chapter, populated with many more injections, to determine the point at which the efficiencies cross. In particular, as will be discussed in the next chapter, much of the background in the low-mass search comes from the higher-mass end of the template bank. If we can reduce the upper limit of the low-mass search we can hope to clean up the background and make quiet signals more significant.

Effects of spin

Using efficiency plots such as those presented in this chapter we hope to determine the regions of parameter space that are not well-covered by the existing non-spinning searches. Such regions could then be supplemented by spinning searches currently in development. Conversely, by identifying regions where the current search performs adequately we can limit the range, and hence the background, over which spinning

searches need to run.

Effects of template placement

As discussed in Sec. 4.3, the current searches use a bank laid out according to a metric calculated from the *TaylorF2* stationary-phase waveform taken to 2.0 pN order in phase. Current searches use 3.5 pN waveforms, and the template placement is therefore incorrect. The extent to which it is incorrect, and the implications for the search, can be tested by running *bank simulations*. In such a simulation a bank is constructed and the maximum overlap between a signal and every template in the bank is found. By comparing this maximum against the maximum found by varying the parameters continuously (as in chapter 5) we can determine the loss in SNR due to discretizing the bank. We plan to run such analyses using the NINJA-2 signals with both 2.0 pN templates and 3.5 pN templates in order to determine the extent to which the incorrect metric decreases the efficiency of the search.

8.5   Conclusions

In this chapter we discussed the NINJA-2 data sets and presented results from a one-week test run on a data set containing only low mass signals, and a two-month set containing signals up to $350 M_\odot$. The test week shows no anomalies, however when the mass range is extended the low-mass pipeline shows unexpected behavior. We consider a few possibilities to explain this, but as of this writing no explanation has been found. It seems more likely to be a bug in the data set generation than in the analysis, as the latter has been much more extensively reviewed. In either case, more study is needed. Once these issues are resolved we will proceed with the full runs using real detector data.

Chapter 9

A Database for Instrumental Data Quality

We have made frequent references to *glitches* in the detectors, a type of noise consisting of short, transient features. In chapter 4 we have also noted that such features in the data can elevate $\rho(t)$ and produce spurious triggers. We will refer to triggers resulting from noise as *background triggers*.

In order to claim the detection of a gravitational wave we must have a candidate whose significance stands well above that of these background triggers. So far this thesis has focused on increasing the efficiency of the search to increase the significance of the signals. We now shift focus to reducing the number of triggers in the background. The key to doing this is ability to *veto* time, that is, remove from the analysis times which we believe will contribute excessively to the background and/or times in which we would be unable to confidently detect a real signal.

In this chapter we discuss the infrastructure that enables such vetoes. In the next chapter we look at a tool, *daily ihope*, which was used in S6 to characterize the behavior of the detectors and determine times that needed to be vetoed.

9.1 Data Quality Flags

The LIGO and Virgo detectors are extremely sophisticated, complex, and sensitive. Glitches can be produced in and by many parts of the detector, and may ultimately be caused by either environmental conditions or conditions internal to the detectors.

A great deal of information about the state of the detector is recorded in *channels*, which are time series sampled from 1 Hz to 16 kHz, depending on the channel. These include such information as

- the power at which the Laser is operating

- the output of every photodiode throughout the detector

- the power in seismic noise in various frequency bands

- the activity of the control servos

and a great deal more. When any of these exceed normal operating values it may indicate the presence of glitches, or even be their root cause. It is therefore extremely valuable to both data analysts and commissioners working on the detector to be able to *flag* any anomalous conditions.

The first line of defense against glitches is on-site as the detector is running. At all such times the control room is staffed by an operator, an employee of LIGO labs who is an expert in running the detector, and a science monitor ("SciMon") who is a member of the LIGO Scientific Collaboration (LSC). Either may flag times of unusual activity such as

- movement of heavy machinery on the site

- nearby storms

- modification of control parameters

The operator and SciMon also jointly decide when to flag time as suitable for analysis. This depends on a few conditions; there must be laser light in both arms and the detector must be *locked*, meaning the cavities are on resonance. When such conditions are met the operator and SciMon can choose to enable *science mode*. Any time such flagged will be included in the LIGO searches.

In addition to such manually-created flags, flags may be created automatically by scanning the channel data. This is done by a program called the *Data Monitoring Tool* (DMT) which records conditions such as

- an instrumental channel has exceeded a threshold

- the standard deviation of a channel has exceeded its expected value

- elevated seismic noise in a frequency band

- the detector is properly locked

Note that these are all yes-or-no conditions. Flags do not store values, such as the level of seismic noise. That information is available in the channel data, the flags are intended to provide the simplest possible summary of conditions. At any given time a flag can be in any of three states: on, off, or undefined. Flags are undefined during times when the responsible DMT process was not running.

In order to be useful flags must be represented in a standard format and accessible to members of the LSC. We start by identifying a flag as a triple of (detector id, flag name, version number). *detector id* identifies the detector with a two-letter abbreviation; H1, L1, V1 for Hanford, Livingston and Virgo respectively. *flag name* is a unique identifier, to prevent confusion with case sensitivity in processing tools, all DQ flag names were required to be in upper case case. Additionally, a 3-letter prefix was be added to each flag name to better identify the source where data came from. For example: the flag indicating high wind speed would be `DMT-WIND_OVER_30MPH`. *version number* is an integer starting from 1. Version numbers can change if, for example, a bug is found in one of the DMT processes. Once the bug is fixed the DMT can be rerun over the original channel data to produce the next subsequent version of the flag. However, by storing all prior versions we are able to reconstruct the results of earlier searches.

Conceptually, managing flag information consists of providing functions that

- map from a flag (ifo id, flag name, version number) to a set of times where the flag was defined

- map from a flag (ifo id, flag name, version number) to a set of times where the flag was active

The times during which the flag was inactive, if needed, can be obtained from the difference between these two. It is also useful to think of these functions as mapping (ifo id, flag name, version number, time) to an "active, inactive or undefined" indicator, or mapping a time to a set of flags that are active.

9.2   Implementation Details

The implementation of the maps between flags and sets of times is conceptually straightforward. We represent sets of times as sets of *segments*, which are half-open intervals aligned on GPS-second boundaries. We can assign to each flag triple a unique integer identifier, and then store segments as triples of (id, start time, end time). The functions above then consist of looking up the flag, determining the identifier, then returning all segments with matching identifiers. We will need two such segment stores, one for "defined" segments and one for "active" segments. We will call the map between flags and identifiers the `segment_definer`, the store of "defined" segments will be called the `segment_summary`, and we we reserve the name `segment` for the store of active segments.

We now have a structure in which data quality flags can be stored. In S6 this was implemented in two ways:

1. as files in eXtensible Markup Language (XML) format. XML is a scheme for adding structure to documents [128]. It is relatively simple in that XML files can be read and modified by standard text editors. The LSC utilizes a specialized form of XML called *LIGO light weight* or LIGO LW [129]. The names of most programs that produce or consume such files begin with the prefix `ligolw_`. See Sec. 9.3 for more on this format.

2. in a relational database called the *segment database*, hosted at CalTech. This allowed high-speed, distributed access to flag information. See Sec. 9.4 for more on the structure of the database.

The flow of flag information is illustrated in Fig. 9.2:

1. **DMT** The DMT handled the creation of all science segments and online data quality segments in S6. Every 60 seconds, the DMT wrote segment information in XML format to disk at each of the detectors. The DMT called the `dmtdq_seg_insert` program to insert the XML data into the segment database.

2. **Archival of DMT segment data:** The 60 second XML files generated by the DMT were replicated to Caltech and other LSC computing centers. Files were compressed using the gzip algorithm.
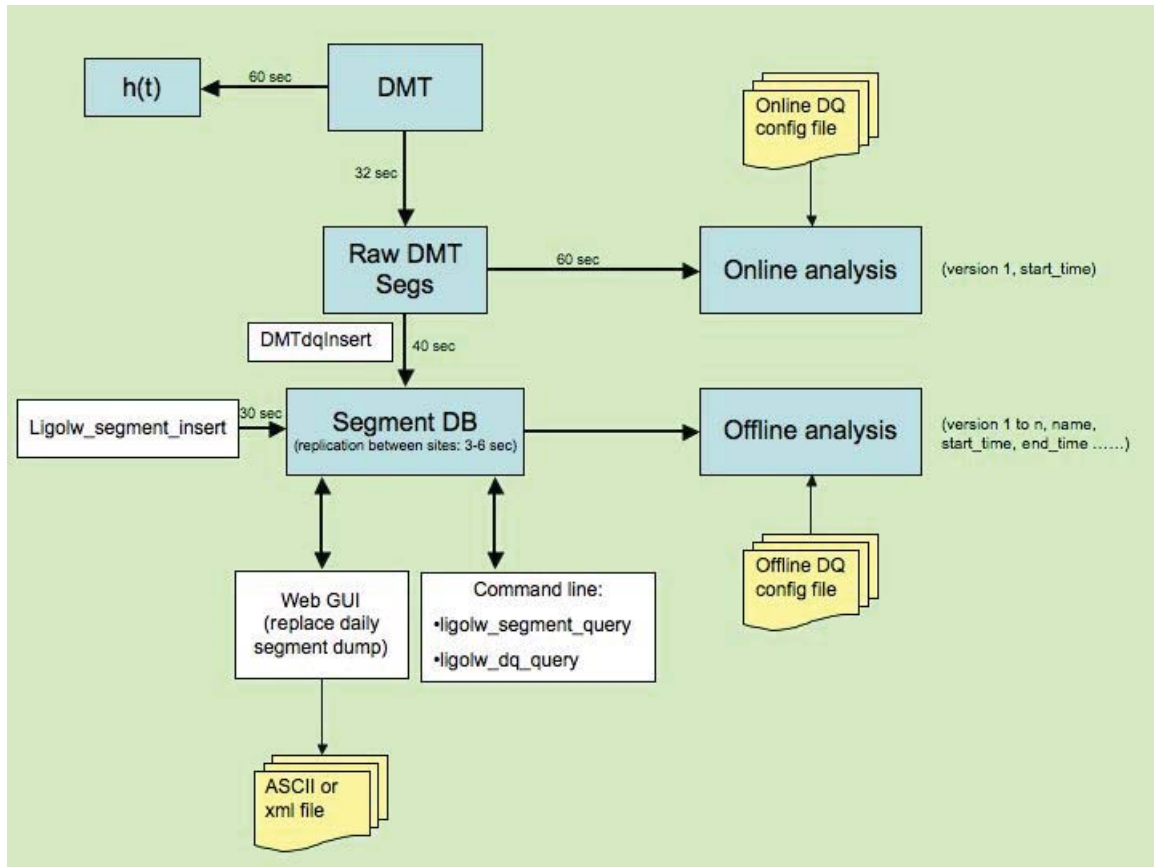
Figure 55 : Flow of S6 data quality segment information. Online data quality segments and science segment information is generated by the DMT. This can be directly queried for low-latency online analysis or inserted into a segment database for off-line or higher-latency analyses. Command line and web GUI tools were used to query and update the segment database.

3. Software was developed to combine the segment information from DMT XML files with a data quality categorization file managed by the search groups (see Sec. 9.5.2 below) to produce files containing veto information for a given search once per minute or so. The same software was be made available for interfacing with the segment database for offline searches and detector characterization (see Sec. 9.5.1 below).

4. **Expected Latencies:** as described in the implementation plan diagram, targeted latencies were:

   - DMT generates h(t) file: 60 seconds
   - From DMT to raw DMT segment disk: 60 seconds

- From raw DMT segment disk to segment database: 10 seconds

- From ligolw_segment_insert to segment database: 30 seconds

9.3 Segment Data File Format

To be ingested into the segment database, segment data must be in the format described in this section. Data should be in a valid LIGO LW XML file with `process`, `segment_definer`, `segment_summary` and `segment` tables. An optional `process_params` table can be used to store extra metadata about segment generation.

The `process` table should contain the columns given in the XML file below which describe the name, version, cvs repository and revision of the program used generate the data. The comment column can be used to add additional human-readable data. The `node`, `username`, `unix_procid`, `start_time` and `end_time` columns should store metadata describing who ran the process and where it ran. The `ifos` column should contain an alphabetical list of all ifo data using as input to the process. The `process_id` column is used to link the defined process to other rows in the file created by that process.

The `segment_definer` table should contain a definition of the segments included in the file. The `ifos`, `name` and `version` columns should contain the name and version of the segment. The name should be upper case for all segments. DMT-derived segments should be prefixed with the string `DMT-`, segments created by the detector characterization group should be prefixed with `DCH-` and segments from the Virgo database should be prefixed with `VDB-`. The comment column can be used to add a human-readable description of the segment. The `segment_definer` column is used to link this type of segment to the intervals in the `segment` and `segment_summary` tables.

The `segment` table should contain the GPS start and end time when the segment described in the `segment_definer` table was *active*. The `segment_summary` table should contain the GPS start time and end time of the interval for which the segment is *defined*. This will allow users to distinguish between the cases where a DQ segment is undefined for a particular time or simply inactive (i.e. it is *not* windy, as opposed to the wind monitor being down).

```
<?xml version="1.0"?>
<!DOCTYPE LIGO_LW SYSTEM "http://ldas-sw.ligo.caltech.edu/doc/ligolwAPI/html/ligolw_dtd.txt">
<LIGO_LW>
  <Table Name="processgroup:process:table">
    <Column Name="processgroup:process:program" Type="lstring"/>
    <Column Name="processgroup:process:version" Type="lstring"/>
    <Column Name="processgroup:process:cvs_repository" Type="lstring"/>
    <Column Name="processgroup:process:cvs_entry_time" Type="int_4s"/>
    <Column Name="processgroup:process:comment" Type="lstring"/>
    <Column Name="processgroup:process:node" Type="lstring"/>
    <Column Name="processgroup:process:username" Type="lstring"/>
    <Column Name="processgroup:process:unix_procid" Type="int_4s"/>
    <Column Name="processgroup:process:start_time" Type="int_4s"/>
    <Column Name="processgroup:process:end_time" Type="int_4s"/>
    <Column Name="processgroup:process:process_id" Type="ilwd:char"/>
    <Column Name="processgroup:process:ifos" Type="lstring"/>
    <Stream Name="processgroup:process:table" Type="Local" Delimiter=",">
      "SegGener","1.17",
      "/ldcg_server/common/repository_gds/gds/Monitors/SegGener/SegGener.cc\,v",
      865755895,"Segment generation from an OSC condition","granite","jzweizig",718,
      918756928,918836992,"process:process_id:0","H0H1H2"
    </Stream>
  </Table>
  <Table Name="segment_definergroup:segment_definer:table">
    <Column Name="segment_definergroup:segment_definer:process_id" Type="ilwd:char"/>
    <Column Name="segment_definergroup:segment_definer:segment_def_id" Type="ilwd:char"/>
    <Column Name="segment_definergroup:segment_definer:ifos" Type="lstring"/>
    <Column Name="segment_definergroup:segment_definer:name" Type="lstring"/>
    <Column Name="segment_definergroup:segment_definer:version" Type="int_4s"/>
    <Column Name="segment_definergroup:segment_definer:comment" Type="lstring"/>
    <Stream Name="segment_definergroup:segment_definer:table" Type="Local" Delimiter=",">
      "process:process_id:0","segment_definer:segment_def_id:35","H2","DMT-LIGHT",1,
      "H2 Light in arms from h(t) DQ flags",
      "process:process_id:0","segment_definer:segment_def_id:36","H2","DMT-SCIENCE",1,
      "H2 Science mode from h(t) DQ flags",
      "process:process_id:0","segment_definer:segment_def_id:37","H2","DMT-INJECTION",1,
      "H2 Injection mode from h(t) DQ flags",
      "process:process_id:0","segment_definer:segment_def_id:38","H2","DMT-UP",1,
      "H2 calibration OK in from h(t) DQ flags",
      "process:process_id:0","segment_definer:segment_def_id:39","H2","DMT-CALIBRATED",1,
      "H2 Calibration OK from h(t) DQ flags",
      "process:process_id:0","segment_definer:segment_def_id:40","H2","DMT-BADGAMMA",1,
      "H2 Bad gamma in h(t) DQ flags"
    </Stream>
  </Table>
  <Table Name="segmentgroup:segment:table">
    <Column Name="segmentgroup:segment:segment_id" Type="ilwd:char"/>
    <Column Name="segmentgroup:segment:start_time" Type="int_4s"/>
    <Column Name="segmentgroup:segment:end_time" Type="int_4s"/>
    <Column Name="segmentgroup:segment:segment_def_id" Type="ilwd:char"/>
    <Column Name="segmentgroup:segment:process_id" Type="ilwd:char"/>
    <Stream Name="segmentgroup:segment:table" Type="Local" Delimiter=",">
      "segment:segment_id:15",918836961,918836977,"segment_definer:segment_def_id:35",
      "process:process_id:0",
      "segment:segment_id:16",918836976,918836992,"segment_definer:segment_def_id:37",
      "process:process_id:0"
    </Stream>
  </Table>
  <Table Name="segment_summarygroup:segment_summary:table">
    <Column Name="segment_summarygroup:segment_summary:segment_sum_id" Type="ilwd:char"/>
    <Column Name="segment_summarygroup:segment_summary:start_time" Type="int_4s"/>
    <Column Name="segment_summarygroup:segment_summary:end_time" Type="int_4s"/>
    <Column Name="segment_summarygroup:segment_summary:comment" Type="lstring"/>
    <Column Name="segment_summarygroup:segment_summary:segment_def_id" Type="ilwd:char"/>
    <Column Name="segment_summarygroup:segment_summary:process_id" Type="ilwd:char"/>
    <Stream Name="segment_summarygroup:segment_summary:table" Type="Local" Delimiter=",">
      "segment_summary:segment_sum_id:5",918836976,918836992,"",
```

```
    "segment_definer:segment_def_id:40","process:process_id:0",
    "segment_summary:segment_sum_id:6",918836976,918836992,"",
    "segment_definer:segment_def_id:39","process:process_id:0",
    "segment_summary:segment_sum_id:11",918836976,918836992,"",
    "segment_definer:segment_def_id:37","process:process_id:0",
    "segment_summary:segment_sum_id:12",918836976,918836992,"",
    "segment_definer:segment_def_id:35","process:process_id:0",
    "segment_summary:segment_sum_id:42",918836976,918836992,"",
    "segment_definer:segment_def_id:36","process:process_id:0",
    "segment_summary:segment_sum_id:46",918836976,918836992,"",
    "segment_definer:segment_def_id:38","process:process_id:0"
  </Stream>
  </Table>
</LIGO_LW>
```

9.4   S6 Segment Database Design

The five tables used for the S6 segment database are shown in Fig. 9.4. Each table has a corresponding structure in the XML files. In addition to the `segment_definer`, `segment_summary` and `segment` tables discussed above, we add `process` and `process_params` tables that indicate the program that created each flag and how that program was run.

9.5   Command Line Tools

We now turn to the tools that implemented the various mappings between flags and sets of times discussed above. Except where noted these programs could all operate on either XML files or the segment database. Except for `ligolw_segment_insert` all programs could only read information from the database, they could not add or change existing data. Flags were specified in the form `ifo:name:version`. For many programs the version could be omitted and the program would report on the latest defined version.

9.5.1   ligolw_segment_query

`ligolw_segment_query` was designed to answer the following questions:

- What DQ flags exist in the database? `ligolw_segment_query --show-types`

- When was a given flag inserted? `ligolw_segment_query --query-types`

- When was a given DQ flag defined? `ligolw_segment_query --query-types`

S6 Segment Database Schema
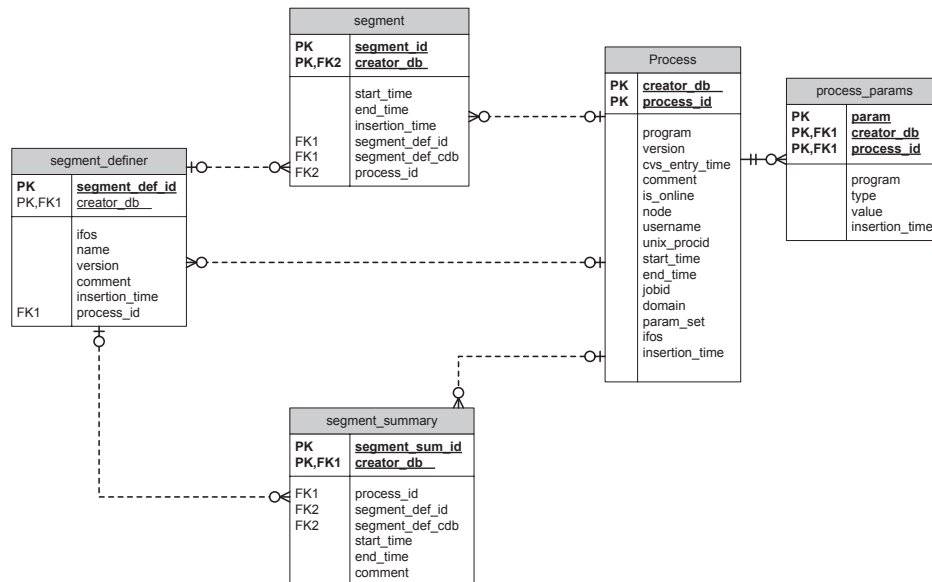
Figure 56 : Structure of the S6 segment database. Each box represents a *table* (a set of structured data), the name of which is in grey. The lines within the boxes are the names of *columns*, individual data elements within the table. For example, the `process_params` table contains a set of *rows*, each row contains data elements named `program`, `type`, `value` and `intertion_time`.

The lines between boxes indicate relations between tables. For example, the `segment_definer` is connected to the `segment_summary` and `segment` tables through the `segment_def_id`. This implements the conceptual relations between these entities discussed above.

- When was a given flag active? `ligolw_segment_query --query-segments`

The built-in help message, generated by `ligolw_segment_query --help` is shown below:

```
DESCRIPTION:
  --version           show program's version number and exit
  -h, --help          show this help message and exit
  -p, --ping          Ping the target server
  -y, --show-types    Returns a xml table containing segment type
                      information: ifos, name, version,
```

```
                        segment_definer.comment, segment_summary.start_time,
                        segment_summary.end_time, segment_summary.comment
-u, --query-types       Returns a ligolw document whose segment_definer table
                        includes all segment types defined in the given period
                        and included by include-segments and whose
                        segment_summary table indicates the times for which
                        those segments are defined.
-q, --query-segments    Returns a ligolw document whose segment table contains
                        the times included by the include-segments flag and
                        excluded by exclude-segments
-s gps_start_time, --gps-start-time=gps_start_time
                        Start of GPS time range
-e gps_end_time, --gps-end-time=gps_end_time
                        End of GPS time range
-t segment_url, --segment-url=segment_url
                        Segment URL. Users have to specify either 'https://'
                        for a secure connection or 'http://' for an insecure
                        connection in the segment database url. For example,
                        '--segment-url=https://segdb.ligo.caltech.edu'. No
                        need to specify port number.
-d, --database          use database specified by environment variable
                        S6_SEGMENT_SERVER. For example,
                        'S6_SEGMENT_SERVER=https://segdb.ligo.caltech.edu'
-f, --dmt-files         use files in directory specified by environment
                        variable ONLINEDQ, for example,
                        'ONLINEDQ=file:///path_to_dmt'. 'file://' is the
                        prefix, the acutal directory to DMT xml files starts
                        with '/'.
-a include_segments, --include-segments=include_segments
                        This option expects a comma separated list of a colon
                        separated sublist of interferometer, segment type, and
                        version. The union of segments from all types and
                        versions specified is returned. Use --show-types to
                        see what types are available.   For example:
                        --include-segment-types H1:DMT-SCIENCE:1,H1:DMT-
```

```
                         INJECTION:2 will return the segments for which H1 is
                         in either SCIENCE version 1 or INJECTION version 2
                         mode. If version information is not provided, the
                         union of the segments of the latest version of
                         requested segment type(s) will be returned.
  -b exclude_segments, --exclude-segments=exclude_segments
                         This option has to be used in conjunction with
                         --include-segment-types --exclude-segment-types
                         subtracts the union of unwanted segments from the
                         specified types from the results of --include-segment-
                         types. If version information is not provided,
                         --exclude-segment-types subtracts the union of
                         segments from the latest version of the specified
                         segment types. For example, --include-segment-types H1
                         :DMT-SCIENCE:1,H1:DMT-INJECTION:2 --exclude-segment-
                         types H1:DMT-WIND:1,H1:DMT-NOT_LOCKED:2,H2:DMT-
                         NOT_LOCKED:2 will subtract the union of segments which
                         H1 is in version 1 WIND and H1,H2 is version 2
                         NOT_LOCKED from the result of --include-segment-types
                         H1:DMT-SCIENCE:1,H1:DMT-INJECTION:2
  -S, --strict-off       The default behavior is to truncate segments so that
                         returned segments are entirely in the interval [gps-
                         start-time, gps-end-time).  However if this option is
                         given, the entire non-truncated segment is returned if
                         any part of it overlaps the interval.
  -o output_file, --output-file=output_file
                         File to which output should be written.  Defaults to
                         stdout.
```

ligolw_dq_query

Below are the questions that ligolw_dq_query was designed to answer:

- is a given flag active at a given time? ligolw_dq_query --active

- is a given flag defined at a given time? ligolw_dq_query --defined

- what is the status of all flags at a given time? `ligolw_dq_query --report`

The built-in help message, generated by `ligolw_dq_query --help` is shown below:

```
DESCRIPTION:
  --version             show program's version number and exit
  -h, --help            show this help message and exit
  -p, --ping            Ping the target server
  -y, --defined         Returns a segment summary table containing segments
                        defined at the given time(s).
  -u, --active          Returns a segment table containing segments active at
                        the given time(s).
  -q, --report          Prints which flags are defined/undefined at the given
                        time(s). For the flags which were defined, it
                        determines if the flag was active or inactive at that
                        time. For an active flag, it prints the start and end
                        time of the segment to which the active. For an
                        inactive flag, it prints the end time of the previous
                        adjacent active segment and the start time of the next
                        adjacent active segment
  -s start_pad, --start-pad=start_pad
                        Seconds before given time(s) to include in query
  -e end_pad, --end-pad=end_pad
                        Seconds after given time(s) to include in query
  -t segment_url, --segment-url=segment_url
                        Segment URL
  -d, --database        use database specified by environment variable
                        S6_SEGMENT_SERVER
  -f, --dmt-files       use files in directory specified by environment
                        variable ONLINEDQ
  -a include_segments, --include-segments=include_segments
                        This option expects a comma separated list of a colon
                        separated sublist of interferometer, segment type, and
                        version. The union of segments from all types and
                        versions specified is returned. Use --show-types to
                        see what types are available.   For example:
```

```
                        --include-segment-types H1:SCIENCE:1,H1:INJECTION:2
                        will return the segments for which H1 is in either
                        SCIENCE version 1 or INJECTION version 2 mode. If
                        version information is not provided, the union of the
                        segments of the latest version of requested segment
                        type(s) will be returned.
  -o output_file, --output-file=output_file
                        File to which output should be written.  Defaults to
                        stdout.
  -i, --in-segments-only
                        If set, report will only return segments that given
                        times were within
```

ligolw_segments_from_cats

`ligolw_segments_from_cats` read a veto definer file (see Sec. 9.5.2) and accessed either segment XML files or the segment database in order to produce a list of times to be vetoed. The built-in help message is displayed below:

```
DESCRIPTION:
  --version             show program's version number and exit
  -h, --help            show this help message and exit
  -v veto_file, --veto-file=veto_file
                        veto XML file (required).
  -o output_dir, --output-dir=output_dir
                        Directory to write output (default=cwd).
  -k, --keep-db         Keep sqlite database.
  -t segment_url, --segment-url=segment_url
                        Segment URL
  -d, --database        use database specified by environment variable
                        S6_SEGMENT_SERVER
  -f, --dmt-file        use files in directory specified by environment
                        variable ONLINEDQ
  -c, --cumulative-categories
                        If set the category N files will contain all segments
                        in categories <= N
```

```
-p, --separate-categories
                        If set the category N files will contain only category
                        N
-s gps_start_time, --gps-start-time=gps_start_time
                        Start of GPS time range
-e gps_end_time, --gps-end-time=gps_end_time
                        End of GPS time range
```

For example:

```
ligolw_segments_from_cats
  --gps-start-time 930960015
  --gps-end-time 931564887
  --segment-url https://segdb.ligo.caltech.edu:30015
  --cumulative-categories
  --veto-file H1L1V1-S6_CBC_LOWMASS_ONLINE-928271454-0.xml
```

ligolw_segment_insert

ligolw_segment_insert handled two tasks:

- Insert segments and/or segment types into the segment database.

- Append segments to the existing segment types.

This program uses the terminology IFO (*Interferometric Observatory*) instead of "detector." The built-in help message is displayed below:

```
DESCRIPTION:
  -h, --help            show this help message and exit
  -p, --ping            Ping the target server
  -t URL, --segment-url=URL
                        Users have to specify protocol 'https://' for a secure
                        connection in the segment database url. For example,
                        '--segment-url=https://segdb.ligo.caltech.edu'. No
                        need to specify port number'.
  -o FILE, --output=FILE
                        Write segments to FILE rather than the segment
```

```
                                database
  -j IDENTITY, --identity=IDENTITY
                                Set the subject line of the server's service
                                certificate to IDENTITY
  -I, --insert                  Insert segments to the segment database
  -A, --append                  Append segments to an existing segment type
  -i IFOS, --ifos=IFOS  Set the segment interferometer to IFOS (e.g. H1)
  -n NAME, --name=NAME  Set the name of the segment to NAME (e.g. DMT-
                                BADMONTH)
  -v VERSION, --version=VERSION
                                Set the numeric version of the segment to VERSION
                                (e.g. 1)
  -e EXPLAIN, --explain=EXPLAIN
                                Set the segment_definer:comment to COMMENT. This
                                should explaining WHAT this flag mean (e.g. "Light dip
                                10%"). Required when --Insert/-I is specified.
  -c COMMENT, --comment=COMMENT
                                Set the segment_summary:comment to COMMENT. This
                                should explaining WHY these segments were inserted
                                (e.g. "Created from hveto results")
  -S FILE, --summary-file=FILE
                                Read the segment_summary rows from FILE. This should
                                be a file containing the gps start and end times that
                                the flag was defined, deliminated by comma (i.e. the union of on
  -G FILE, --segment-file=FILE
                                Read the segment rows from FILE. This should containin
                                the gps start and end times when the flag was active deliminated
```

As an example, the command to insert segments of a new type would look like:

```
ligolw_segment_insert
  --segment-url https://segdb.ligo.caltech.edu
  --ifos 'H1'
  --name 'DCH-TEST'
  --version 1
  --comment 'testing if insert works'
```

```
  --explain 'test insert'
  --segment-file segment.txt
  --summary-file summary.txt
  --insert
```

To append segments to an existing segment, the command would look like:

```
ligolw_segment_insert
  --segment-url https://segdb.ligo.caltech.edu
  --ifos 'H1'
  --name 'DCH-TEST'
  --version 1
  --comment 'testing if append works'
  --segment-file append_segment.txt
  --summary-file append_summary.txt
  --append
```

9.5.2 The Veto Definer

Problems in the data may have differing levels of severity, and consequently we define several *veto categories* to characterize them. We defer discussion of the meaning of these categories until the next chapter, but note here that the categories are indicated by integers from 1 to 4, in decreasing levels of severity. These categories are commonly referred to as *CAT1* through *CAT4*. We also note that time marked with CAT1 is entirely excluded form the analysis. 4091 seconds of data with one second of CAT1 time in the middle will be split into two 2045-second chunks. Neither of these will be analyzable by the ihope pipeline, which requires 2048 seconds of data in order to compute the PSD as discussed in Sec. 4.4. Short CAT1 vetos can therefore effectively veto much larger times, and such vetoes are therefore to be avoided whenever possible.

In order to veto a chunk of time it is necessary to map a data quality flag to a veto category. This is accomplished with a *veto definer file*. Such files are specific to each search, as time that is problematic for one search may pose no issues for another. The `ligolw_segments_from_cats` program (see above) merges the information in this file with the set of active flag segments to produce *veto segments* at each veto level. Analysis is performed on times marked as science with no CAT1 vetoes. Triggers from times marked as CAT2, CAT3 and CAT4 are discarded from both foreground and background.

In some circumstances the search is more sensitive to anomalous conditions than the DMT and excess triggers may be produced outside the time covered by the flag. This is one motivation for using the search itself as a data quality tool, as discussed in the next chapter. The veto definer file therefore allows for *padding*, offsets which effectively extend the flags. The padding used will be specific to each search, as different searches will be sensitive to different classes of problems over different extents.

9.5.3 Example Veto Configuration File

The veto configuration file should contain a `process` table describing how it was created and a `veto_definer` table describing the flags to be applied at different levels. The comment column of the process table should contain the version of the file. The columns in the `veto_definer` table are as follows: `ifo`, `name` and `version` uniquely define a particular DQ/veto flag. `category` indicates at which category the veto should be applied, `start_time` and `end_time` denote the GPS time interval for which the DQ/veto flag should be applied (Note: if `end_time` is zero, then the current GPS time is assumed). `start_pad` and `end_pad` are the padding time (in seconds) applied to the start and end of the veto segments. Note that these are signed: if you want time vetoed to start time to be *earlier* than the start time listed in the database, then the emph `start_pad` should be *negative*. Similarly, if the time vetoed should extend *after* the end time stored in the database, then the value in `end_pad` should be *positive*. The `comment` column can be used to store an optional human-readable comment.

```
<?xml version='1.0' encoding='utf-8' ?>
<!DOCTYPE LIGO_LW SYSTEM "http://ldas-sw.ligo.caltech.edu/doc/ligolwAPI/html/ligolw_dtd.txt">
<LIGO_LW>
   <Table Name="process:table">
      <Column Name="process:process_id" Type="ilwd:char"/>
      <Column Name="process:program" Type="lstring"/>
      <Column Name="process:version" Type="lstring"/>
      <Column Name="process:cvs_repository" Type="lstring"/>
      <Column Name="process:cvs_entry_time" Type="int_4s"/>
      <Column Name="process:node" Type="lstring"/>
      <Column Name="process:username" Type="lstring"/>
      <Column Name="process:unix_procid" Type="int_4s"/>
      <Column Name="process:start_time" Type="int_4s"/>
      <Column Name="process:end_time" Type="int_4s"/>
      <Column Name="process:ifos" Type="lstring"/>
      <Column Name="process:comment" Type="lstring"/>
      <Stream Name="process:table" Type="Local" Delimiter=",">
      "process:process_id:0","ligolw_veto_file","1.1",
      "/usr/local/cvs/lscsoft/glue/bin/ligolw_veto_file,v",822908378,
      "ldas-grid.ligo.caltech.edu","jrsmith",16830,822879594,822879594,
      "H1","Example file by Josh"
      </Stream>
   </Table>
   <Table Name="veto_definer:table">
      <Column Name="veto_definer:process_id" Type="ilwd:char"/>
      <Column Name="veto_definer:ifo" Type="lstring"/>
```

```
        <Column Name="veto_definer:name" Type="lstring"/>
        <Column Name="veto_definer:version" Type="int_4s"/>
        <Column Name="veto_definer:category" Type="int_4s"/>
        <Column Name="veto_definer:start_time" Type="int_4s"/>
        <Column Name="veto_definer:end_time" Type="int_4s"/>
        <Column Name="veto_definer:start_pad" Type="int_4s"/>
        <Column Name="veto_definer:end_pad" Type="int_4s"/>
        <Column Name="veto_definer:comment" Type="lstring"/>
        <Stream Name="veto_definer:table" Type="Local" Delimiter=",">
        "process:process_id:0","H1","OUT_OF_LOCK",0,1,917985615,0,0,0,"",
        "process:process_id:0","H1","BADGAMMA",1,1,917985615,0,0,0,"",
        "process:process_id:0","H1","ASC_Overflow",0,2,917985615,0,-8,8,"ASC saturations are cat2",
        "process:process_id:0","H1","PD_Overflow",0,2,917985615,0,-8,8,"PD saturations are cat2",
        "process:process_id:0","H1","SEVERE_LSC_OVERFLOW",0,2,917985615,0,-8,8,"LSC saturations are cat2",
        "process:process_id:0","H1","INJECTION",1,2,917985615,0,-16,64,"Remove HW injections at cat1",
        "process:process_id:0","H1","ASC_Overflow",0,3,917985615,0,-8,25,"ASC saturations are cat3 with larger pad",
        "process:process_id:0","H1","SEVERE_LSC_OVERFLOW",0,3,917985615,0,-8,25,"LSC saturations are cat3 with larger pad",
        "process:process_id:0","H1","PD_Overflow",0,3,917985615,0,-8,25,"PD saturations are cat3 with larger pad",
        "process:process_id:0","H1","Wind_over_30MPH",0,3,917985615,0,-8,8,"Windy",
        "process:process_id:0","H1","LIGHTDIP_1_PERCENT",0,3,917985615,0,-2,2,"Exclude all lightdip segments",
        "process:process_id:0","H1","PRE_LOCKLOSS_10_SEC",0,3,917985615,0,0,0,"",
        "process:process_id:0","H1","PRE_LOCKLOSS_30_SEC",0,3,917985615,0,0,0,"",
        "process:process_id:0","H1","PRE_LOCKLOSS_60_SEC",0,3,917985615,0,0,0,"",
        "process:process_id:0","H1","PRE_LOCKLOSS_120_SEC",0,3,917985615,0,0,0,"",
        "process:process_id:0","H1","ASI_CORR_OVERFLOW",0,4,917985615,0,-8,25,"Does this make sense for DC readout",
        "process:process_id:0","H1","LSC_OVERFLOW",0,4,917985615,0,-8,25,"LSC saturations",
        "process:process_id:0","H1","PRE_LOCKLOSS_600_SEC",0,4,917985615,0,0,0,"",
        "process:process_id:0","H1","PRE_LOCKLOSS_1800_SEC",0,4,917985615,0,0,0,""
        </Stream>
    </Table>
</LIGO_LW>
```

Chapter 10

Characterization of a Gravitational Wave Detector

In chapter 3 we noted the presence of *glitches*, short transients, in the detector noise. In chapter 4 we noted that such glitches could produce triggers. We have also noted that, in order to claim a detection, a candidate gravitational-wave trigger must stand out above the background comprised of triggers from noise sources, including glitches. There is therefore a need to identify and remove glitches. Ideally glitches would be removed at the source, by addressing the underlying cause in the detector. However, by the time a glitch has been seen, and often long before a cause can be identified and addressed, it is already in the data. We must therefore devise ways to remove glitches from the data as part of the analysis. The activities of identifying glitches and glitch mechanisms and removing them from analyses fall into the category of *detector characterization.*

In the preceding chapter we discussed data quality flags, which characterize the state of the detectors at any time, the various repositories for storing these flags, and the tools for creating and accessing them. We also discussed the veto definer file, which maps flags to veto levels. By combining a veto definer with a set of data quality flags we produce a set of *veto segments*, intervals of time that should either not be analyzed at all or which should be removed from the analysis. We have thus far not discussed the construction of the veto definer file. This file may be thought of as the product of detector characterization studies. In broad terms the goal is to identify glitches, correlate them with existing flags or create new ones, and remove the glitches from analyses by associating the flag with a veto severity level through the veto definer file. One of the most significant features of the S6 run is that the data analysts working on detector characterization were in close contact

with the commissioners working on the detectors. This meant that, in addition to cleaning the search, information about problematic behavior in the detector could be used to fix the problems at the source, in turn allowing more time to be analyzed and the prospects for making a detection to improve.

Detector characterization in S6 was a major undertaking, much more information can be found in an upcoming paper by the LIGO Scientific Collaboration (LSC) and Virgo. Three primary tools were run either continuously or with low-latencies to provide views into the data, although these were supplemented with numerous short special-purpose programs. Two of these tools, the Klinewelle pipeline [130] and the Omega pipeline [131] project data onto a wavelet basis and identify as triggers times and frequencies with amplitudes above a given threshold. They thus look for "bursts" of excess power. Both these tools were run on many auxiliary channels in addition to the detector output. The third tool, daily ihope, is the focus of the current chapter. We proceed as follows:

In Sec. 10.1 we introduce the daily ihope pipeline, focusing on the differences between daily ihope and the full ihope pipeline discussed in 4.4.

We then move to general applications of the daily ihope results. Section 10.2 gives a walk-through of a web-based front end to daily ihope which provided numerous views into the behavior of the detectors over 24-hour periods. Section 10.3 discusses the changes to the low-mass pipeline that were implemented in part based on results from daily ihope. Section 10.4 focuses on the use of daily ihope in constructing the veto definer file used in the low-mass search.

We next discuss applications of daily ihope to the study of a specific event, a potential gravitational wave candidate that ultimately was revealed to be a hardware injection that was performed as a "blind injection challenge" meant to test the search pipelines. Until the nature of the candidate was revealed the collaboration followed it up with all due diligence, including many studies that needed to be performed rapidly. The studies that were enabled by daily ihope are discussed in Sec. 10.5.

Finally, in Sec. 10.6 we note some issues with the pipeline that were flagged by daily ihope, and discuss preliminary attempts to address them. As of this writing these problems remain unsolved and are areas of active research.

10.1   Daily Ihope

As discussed in chapter 4, the CBC group uses the *ihope* pipeline to search for gravitational waves produced by the inspiral of binary systems consisting of neutron stars and/or black

holes. During S6/VSR2,3 the goal of the CBC group was to run the full analysis every two weeks with latencies as low as possible. The daily ihope runs were conceived of as a way to characterize the detector in order to look for potential problems in advance of the full analysis, particularly problems specific to the CBC search.

It is important to stress that daily ihope was not itself a gravitational-wave search, rather it was a tool to characterise the detectors. This goal determined the specifics of the daily runs, but more significantly it leaves the full search unbiased.

10.1.1 The Daily Ihope Pipeline

Recall that ihope is a templated, matched-filter search. For the low-mass search, which was the focus of the daily runs, the templates are restricted, stationary-phase frequency-domain waveforms with phase evolution taken to 3.5 pN order (see Appendix A). The templates are laid out in a bank with 97% overlap between nearest neighbors, and the mass range is from 2 $M_\odot - 25M_\odot$. A χ^2 discriminator is constructed to better separate signals from glitches (eqn.4.21), and the information from this discriminator is used along with the SNR to construct the ρ_{new} detection statistic (eqn. 4.22).

By construction, a great deal of the information provided by the full bank search is redundant. This is necessary in the full search to avoid losing more than 3% of the SNR from a real signal. However, as daily ihope is not a search this requirement no longer holds. Further, the evaluation of each additional template and the initial layout of the bank incurs significant computational overhead. Therefore a number of simplifications were made to the daily ihope bank.

First, a static bank was used for each detector, based on the layout at a time in each detector when the inspiral range was high, no operational problems were noted by the scimon or operator, and the Omega, Klinewelle, and daily ihope pipelines showed no anomalous behavior. Second, the number of templates was reduced. Lower-mass systems have more cycles in the sensitive LIGO band, there is therefore ample information to distinguish between waveforms with close parameters. Conversely, this means that the low-mass end of the bank is very dense (in the sense of number of templates per unit square in parameter space, they are still equidistant in the sense of the bank metric). Such fine resolution is not needed when looking for glitches, therefore the minimal match between templates in the region below a chirp mass ($\mathcal{M} = M\eta^{3/5}$) of 3.46 $M_\odot$ was set to 0.5. At higher masses, up to a total mass of $25M_\odot$, the minimal match was set to 0.95. In addition to ensuring coverage in a mass region which is naturally sparser this ensures that short glitches, characteristic of many glitch mechanisms, were flagged with large SNRs. The resulting hybrid bank for the
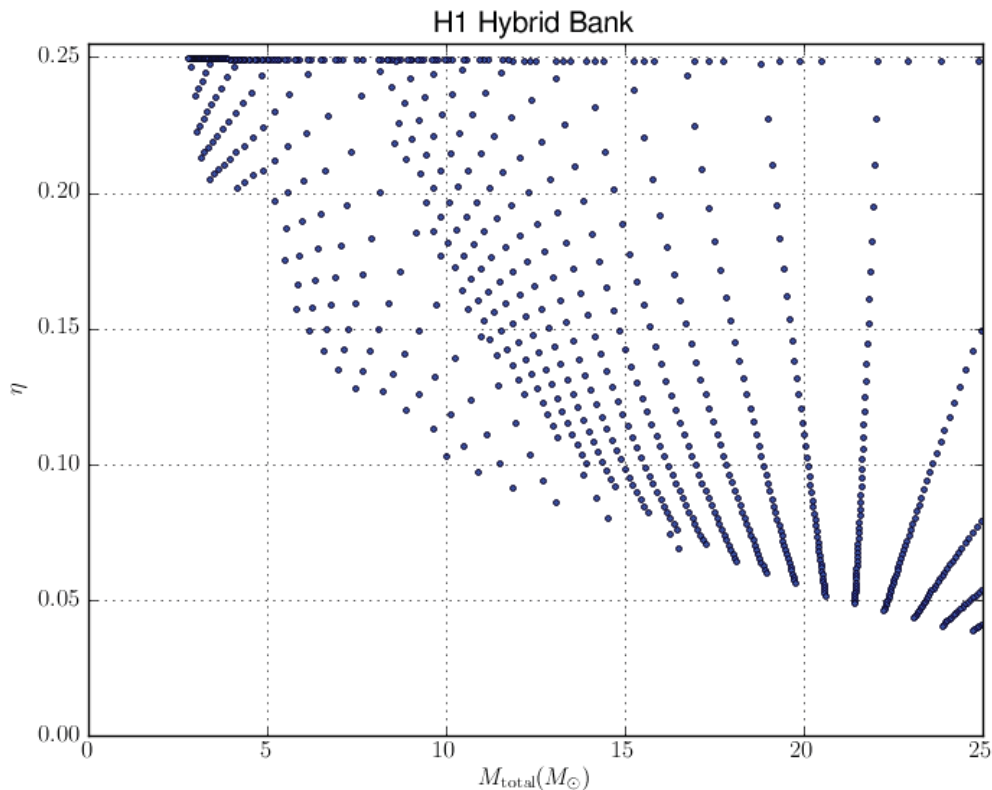
Figure 57 : The hybrid template bank used by daily ihope for the Hanford detector, the banks for L1 and V1 are similar. The cut is made at constant chirp mass, which is a curve in the total mass, η plane.

Hanford detector is shown in figure 57.

We are interested in not only the SNR, but also the extent to which glitches are able to mimic real signals. We therefore wish to calculate χ^2. However, this is computationally expensive, as it amounts to filtering with each template 16 times. Reducing the size of the bank helps, but we can obtain results even faster by splitting the bank into smaller sub-banks and filtering with each in parallel. As the templates are independent at this stage there is no complication in doing so.

Daily ihope processing ran at 03:00 GMT and examined data spanning the 24 hours ending at 00:00 GMT. Unlike full ihope, the analysis was done on all time marked as *science* by the operator and SciMon (see Sec. 9.1), including that which would be vetoed at CAT1 in the full analysis. This was done so that we could see the effect of the CAT1 vetoes. It is plausible, for example, that some such vetoes would be too aggressive and we could decide

based on the daily results to include such times back into the analysis. For consistency science time was denoted CAT0. In addition CAT3 vetoes to remove hardware injections were not applied, so that we could see how many injections would be lost by CAT4 vetoes. The daily pages therefore displayed results for categories 0, 1, 1+2, and 1+2+4.

For each 2048-second chunk of contiguous data, at each detector, the data was run through each template in the bank and triggers selected as described in Sec. 4.1.4. For each trigger χ^2 and new SNR were calculated and recorded. Clustering was then done in order to focus attention on glitches that were most likely to cause problems in the full search. Two sets of clustered triggers were recorded using windows of 30 milliseconds and 16 seconds. For each set the trigger with the largest value of new SNR across all templates within the window length was recorded.

A version of the veto definer file designated for low-latency use was placed on a web site from where it could be downloaded by daily ihope. This file was used to generate veto segments using the `ligolw_segments_from_cats` program discussed in Sec. 9.5.1. Any trigger that fell within the vetoed segments was removed.

The structure of daily ihope is summarized in Fig. 58. The result is, for each 2048-second block in each ifo, 3 cluster levels times 4 veto levels = 12 sets of triggers.

10.2 The Daily Ihope Report Pages

The output of daily ihope is a set of triggers. These triggers were used by many programs to produce or validate vetoes, some of these uses will be discussed in Sec. 10.4.

In addition to such programmatic usage, the daily pipeline also produced a large number of plots and reports that data analysts could use to spot potential problems before the full analysis was run in order to begin more detailed followup studies. These plots were organized into web pages that were available to the collaboration. We now present a walk-through of the daily pages.

Each report and plot is made for each combination of the following:

- Detector: H1, L1 and V1 for the Hanford, Livingston and Virgo detectors. Daily ihope refers to these as *Interferometric Observatories* or IFOs.

- Cluster level: Unclustered, 30 ms clustering, 16 second clustering

- Veto level: Show all triggers in science time (level "0"), only those not removed by category 1 vetos, only those not removed by categories 1 or 2, or only those note

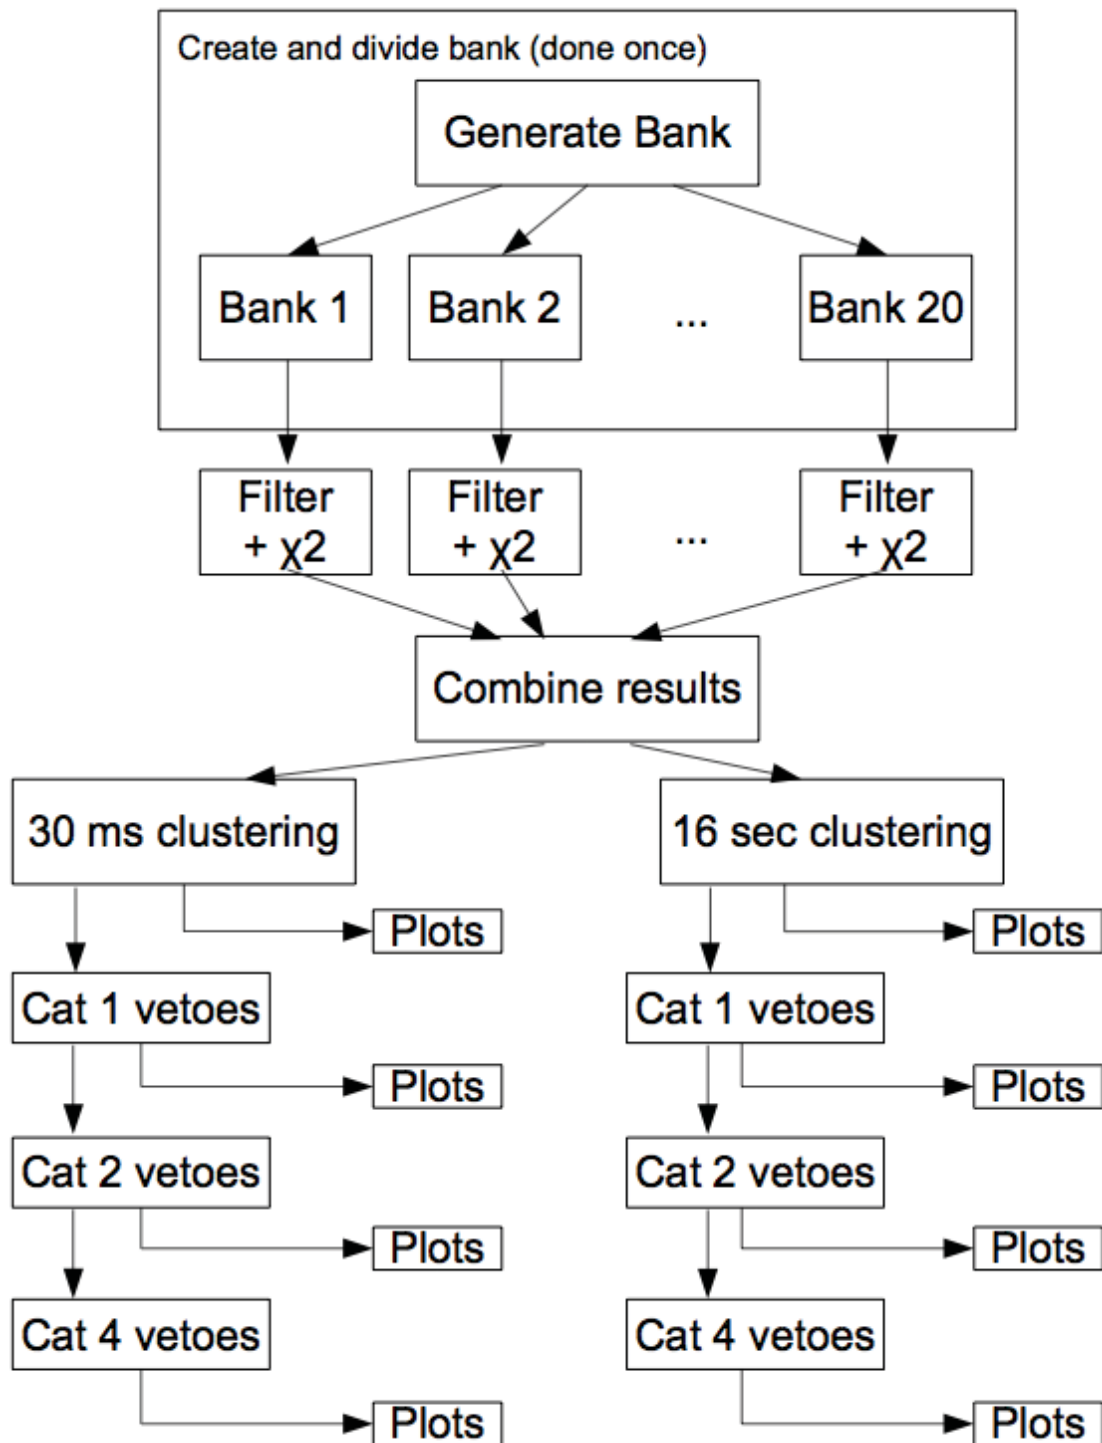Figure 58 : Structure of the daily iHope search. The bank has been generated and split into sub-banks to speed up the process, as described in the text. Each sub-bank is filtered with χ^2 turned on. Triggers are then grouped and clustered. The vetoes are applied sequentially, and at each stage plots are generated. The entire diagram is replicated for each 2048-second block and for each detector.
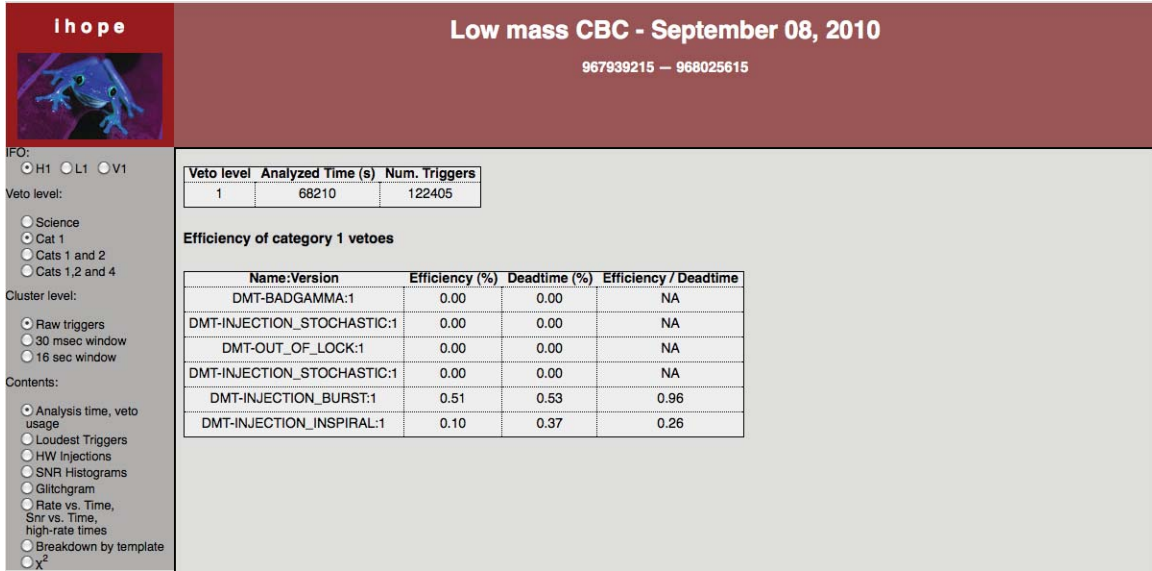
Figure 59 : Top-level web interface for daily ihope. Options can be selected with the controls on the left hand side, when a button is clicked the contents region is immediately replaced. The default view is the one shown here; the analysis time and veto efficiency at cat 1 for the Hanford detector.

removed by categories 1,2 and 4. Hardware injections vetos (category 3) were not applied so that we could determine whether category 4 vetoes were removing them.

Each of these features could be set independently. The top-level web interface for a sample day is shown in Fig. 59. The available reports can bee seen at the bottom of the list of controls on the left.

1. Analysis time and veto usage.

This report shows the total time analyzed, the vetoes applied beyond those applied at the previous level, the *efficiency* (percentage of triggers removed) and *deadtime* (percentage of time removed) by each veto. In addition the ratio of efficiency over deadtime is reported as a measure of quality of the veto. A random veto would result in an efficiency-over-deadtime of approximately 1, a finely-tuned veto that removes short, loud events that ring off the entire template bank would have a much higher ratio. An example is shown in Fig. 60.

2. Loudest triggers

This report was only run for 16-second-clustered triggers. Loud glitches tend to produce

Veto level	Analyzed Time (s)	Num. Triggers
4	63295	47633

Efficiency of category 4 vetoes on triggers that passed category 3

Name:Version	Efficiency (%)	Deadtime (%)	Efficiency / Deadtime
DMT-ASC_OVERFLOW:1	0.00	0.00	0.00
DMT-BRMS_SEISMIC_X_30_100_MHZ_HIGHTHRESH:1	0.00	0.00	NA
DMT-BRMS_SEISMIC_Y_30_100_MHZ_HIGHTHRESH:1	0.13	0.24	0.53
DMT-ISI_OVERFLOW:1	0.00	0.00	0.00
DMT-LIGHTDIP_6_PERCENT:1	0.08	0.01	11.39
DMT-LSC_OVERFLOW:1	0.00	0.00	0.00
DMT-OM1_OVERFLOW:1	0.00	0.00	NA
DMT-OM1_DCPD_OVERFLOW:1	0.00	0.00	0.00
DMT-OM2_OVERFLOW:1	0.00	0.00	NA
DMT-SEVERE_LSC_OVERFLOW:1	0.00	0.00	0.00
DMT-TCS_ITMX_OUT_LOW:1	0.00	0.00	NA
DMT-TCS_ITMX_OUT_MED:1	0.00	0.00	NA
DMT-TCS_ITMX_OUT_HIGH:1	0.00	0.00	NA
DMT-TCS_ITMY_OUT_LOW:1	0.00	0.00	NA
DMT-TCS_ITMY_OUT_MED:1	0.00	0.00	NA
DMT-TCS_ITMY_OUT_HIGH:1	0.00	0.00	NA
DMT-WIND_OVER_30MPH:1	0.00	0.00	NA
DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_2E3:1	62.66	14.51	4.32

Figure 60 :   A sample veto usage report, see text for explanation.   Note that 62.66% of all triggers were contained in 14.51% of time, and the DMT flagged this time with elevated seismic activity from 3-10 Hz in the Laser Vacuum Equipment Area (LVEA), as indicated by the DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_2E3 flag.

Glitches without associated DQ flag

ifo	end time UTC	end_time	end_time_ns	SNR	χ^2	New SNR	eff_distance	M_{chirp}	M_{total}	Ω scan	DQ flags
H1	Aug 19 2010 03:13:32 UTC	966222827	940673828	7.4	23	7.4	100	4.9	24	Ω scan	Light,Up,Calibrated,Science
H1	Aug 19 2010 05:17:22 UTC	966230257	88134765	41	1100	7.4	35	10	23	Ω scan	Light,Up,Calibrated,Science
H1	Aug 19 2010 12:15:00 UTC	966255315	423095703	7.6	35	7.3	74	2.5	10	Ω scan	Light,Up,Calibrated,Science
H1	Aug 19 2010 05:15:55 UTC	966230170	753417968	7.9	40	7.2	160	8.1	18	Ω scan	Light,Up,Calibrated,Science
H1	Aug 19 2010 11:58:34 UTC	966254329	31738281	7.1	27	7.1	100	3.6	8.3	Ω scan	Light,Up,Calibrated,Science
H1	Aug 19 2010 08:09:57 UTC	966240612	172607421	7.1	24	7.1	55	1.5	4.1	Ω scan	Light,Up,Calibrated,Science

Glitches with associated DQ flag

ifo	end time UTC	end_time	end_time_ns	SNR	χ^2	New SNR	eff_distance	M_{chirp}	M_{total}	Ω scan	DQ flags
H1	Aug 19 2010 00:06:49 UTC	966211624	802490234	16	59	12	22	1.5	3.4	Ω scan	**Light,Up,Calibrated,Science,Injection** DMT-INJECTION_INSPIRAL [966211616 966211632]
H1	Aug 19 2010 16:09:53 UTC	966269408	397460937	8.9	28	8.9	130	9.5	22	Ω scan	DMT-INSPIRAL_RANGE_STDEV_GT_0P75_MPC [966269408 966269424] DMT-INSPIRAL_RANGE_STDEV_GT_0P50_MPC [966269408 966269424] DMT-INSPIRAL_RANGE_STDEV_GT_1_MPC [966269408 966269424] DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_1E3 [966269408 966269414] **Light,Up,Calibrated,Science,Injection** DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_2E3 [966269408 966269414] DMT-AIRCRAFT_LIKELY [966269408 966269424] DMT-AIRCRAFT_VERY_LIKELY [966269408 966269424] DMT-INJECTION_INSPIRAL [966269408 966269422]
H1	Aug 19 2010 16:09:18 UTC	966269373	713134765	9.1	38	8.4	40	1.7	3.9	Ω scan	DMT-INSPIRAL_RANGE_STDEV_GT_1_MPC [966269360 966269376] DMT-INSPIRAL_RANGE_STDEV_GT_0P50_MPC [966269360 966269376] **Light,Up,Calibrated,Science,Injection** DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_1E3 [966269360 966269376] DMT-AIRCRAFT_LIKELY [966269364 966269376] DMT-AIRCRAFT_VERY_LIKELY [966269360 966269376] DMT-INSPIRAL_RANGE_STDEV_GT_0P75_MPC [966269360 966269376] DMT-INJECTION_INSPIRAL [966269360 966269376] DMT-BRMS_SEISMIC_LVEA_Z_3_10_HZ_THRESH_2E3 [966269360 966269376]

Figure 61 : A sample report on loudest triggers, see text for explanation. Note the rightmost column is populated by the `ligolw_dq_query` program in the `--report` mode. Note also that there are several triggers with no corresponding data quality flags. This would call for further investigation, the first step of which is to examine the associated omega scans, which we show in Fig. 62. Finally, note that of the loud triggers that do have data quality flags all are marked with the `DMT-INJECTION-INSPIRAL` flag, meaning these triggers are from hardware injections. This is a good sign, it means that the most significant triggers are simulated signals rather than glitches.

families of triggers, without clustering the loudest triggers from each day would likely result from one underlying event. This report considers two classes of triggers; those where no data quality flag was active and those where at least one flag was active. A sample report is shown in Fig. 61. For the five triggers with highest new SNR in each category a summary was presented along with a link to an *omega scan*. The omega pipeline is described in [131]. Omega scans are a time-frequency plot which are run on all auxiliary channels recorded by the detectors and therefore provide a visual aid to detect coupling between auxiliary channels and `DARM_ERR`. These can be used to suggest mechanisms behind glitches, especially those that were not already marked by a data quality flag. In addition, over time repeated shapes in the omega scan can pinpoint underlying problems in the detectors that need to be addressed. A few images from one of the glitches in Fig. 61 is shown in Fig. 62.
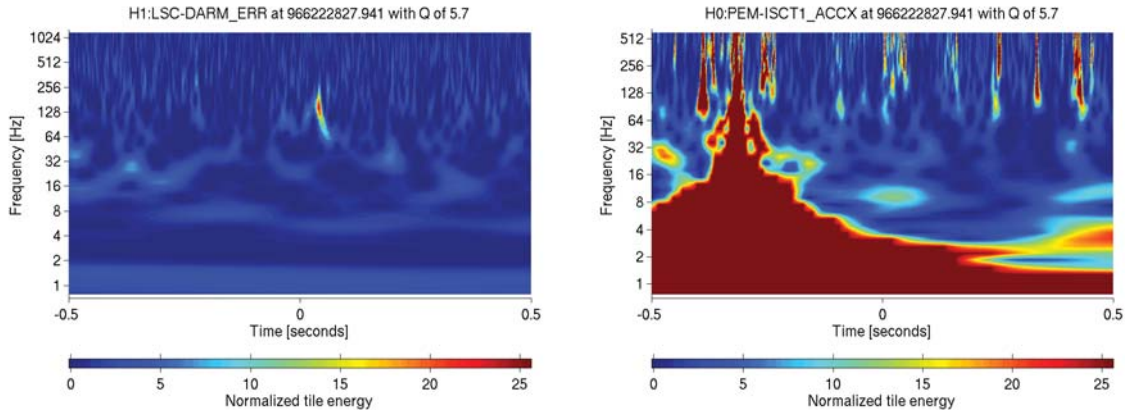
Figure 62 : Omega scans from the loudest trigger in the Hanford detector in figure 61. Note the trigger closely follows a loud event in one of the accelerometers on an instrument table.

3. Hardware injections

This report was generated by code written by John Veitch at Cardiff University. It compared the list of hardware injections, published as an XML file available from a web site, with the list of analysis times and triggers. The results were plotted to indicate whether each injection was found, missed, or not analyzed.

4. SNR histograms

These plots show the number of triggers as a function of ρ and ρ_{new}. As noted in Sec. 4.1, in Gaussian noise the number of triggers should be proportional to $\exp(-\rho^2/2)$. These plots therefore show the degree of "non-Gausianity" in the data. A common use for these plots was to compare different veto levels to get a sense of how well the cumulative vetoes were cleaning the data, as shown in figure 63.

Similarly, Fig. 64 shows histograms of the triggers in new SNR. The total number of triggers is greatly reduced because many high-SNR triggers have $\rho_{\text{new}} < 5$, where the plot cuts off. While triggers with such low ρ_{new} values are not excluded from later stages of analysis in the full pipeline, it is extremely unlikely that the resulting combined new SNR will be high enough to stand above background.

In addition to the lower number of triggers overall, the new SNR plot is closer to a straight line indicating behavior closer to that expected in Gaussian noise. However,
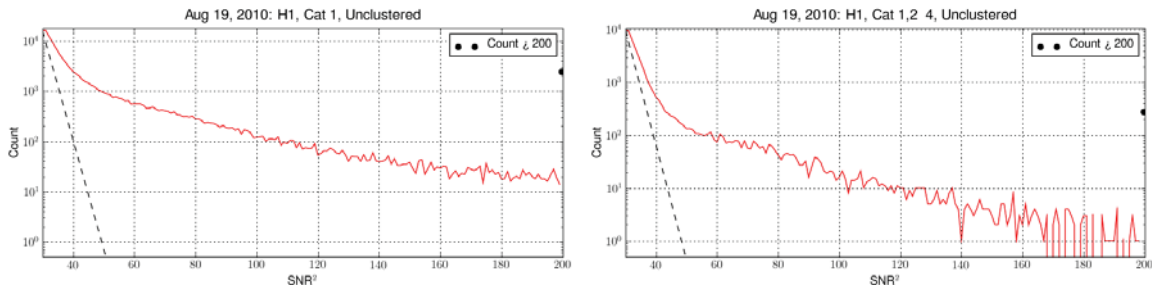
Figure 63 : Sample trigger histograms by SNR. The dashed line shows the expected values in Gaussian noise. The dot indicates the cumulative number of triggers with ρ_{new} greater than 200. Note that at veto category 4 (right) the histogram is closer to the expected line than it is at category 1 (left). This indicates the degree to which data quality has removed non-Guassian noise.
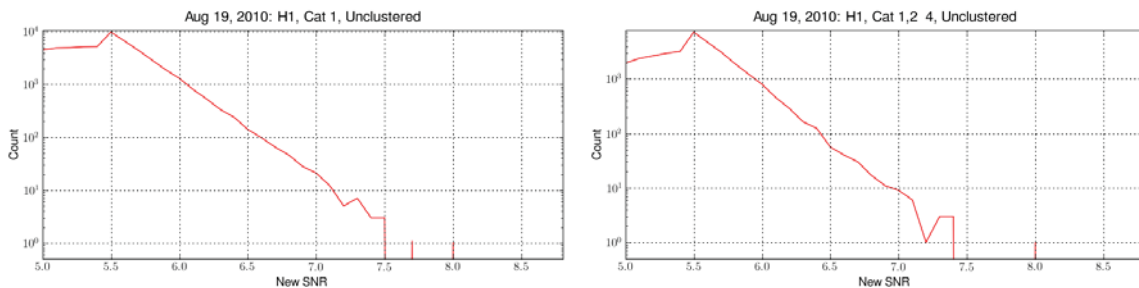

Figure 64 : Sample trigger histograms by new SNR. The data is much cleaner than the SNR histograms, but there is still an outlier at $\rho_{\mathrm{new}} = 8$.

detector characterization improves the situation still further, as the number of triggers is reduced at veto category 4.

However, there is one trigger with $\rho_{\mathrm{new}} = 8$ that is not removed by any veto. Such an outlier warrants further investigation, which here is provided by the loudest events page. An example of this page is shown in 61. The omega scan from the time of this event is shown in Fig. 65, and it clearly rules out a gravitational wave as the source of this trigger. This illustrates the important point that new SNR can be fooled, and hence continued human participation in detector characterization is necessary.

5. The "glitchgram"

This was an "at-a-glance" summary of the day, showing every trigger color-coded by SNR. This highlighted times of loud triggers. In addition dense regions of the plot indicated times of numerous quiet glitches, which might not show up on the rate-vs-time, snr-vs-time or loudest triggers pages. A sample is shown in figure 66.
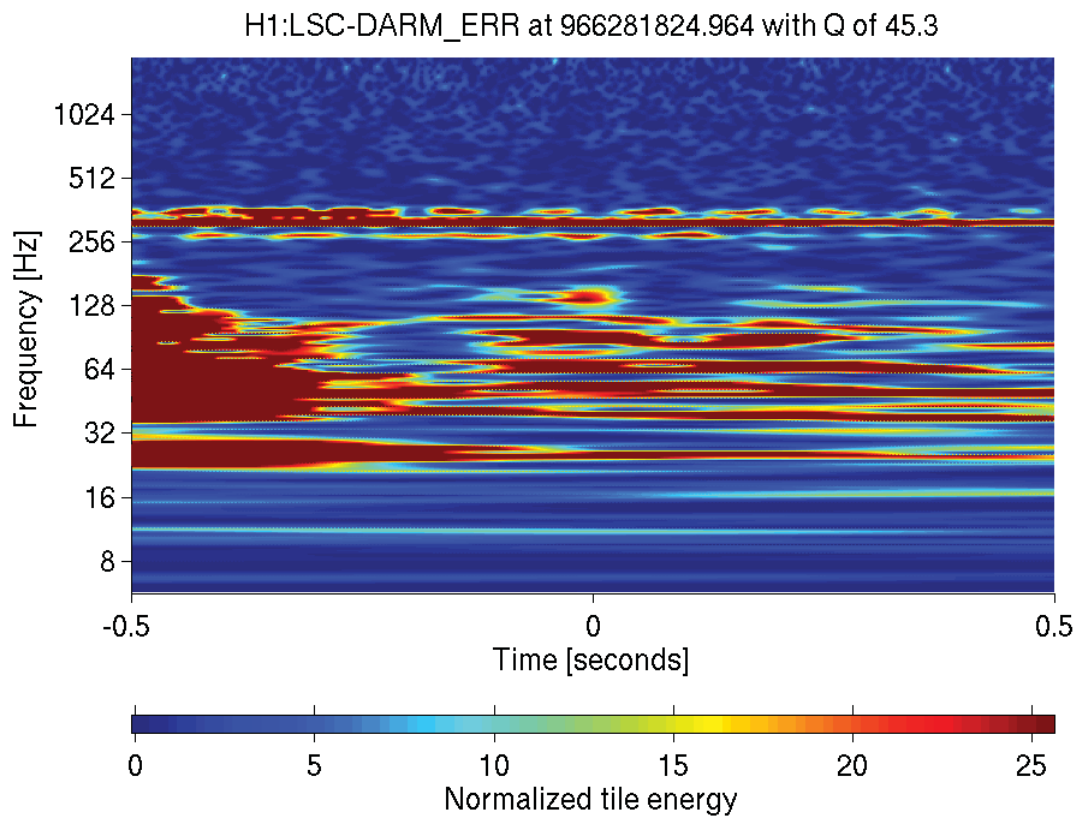
Figure 65 :   The omega scan from the outlier with $\rho_{\mathrm{new}} = 8$ that survives all automated data quality vetoes. Many auxiliary channels showed the same behavior at this time. This is clearly not a gravitational wave, but was not removed by either signal-based vetoes or data quality vetoes.
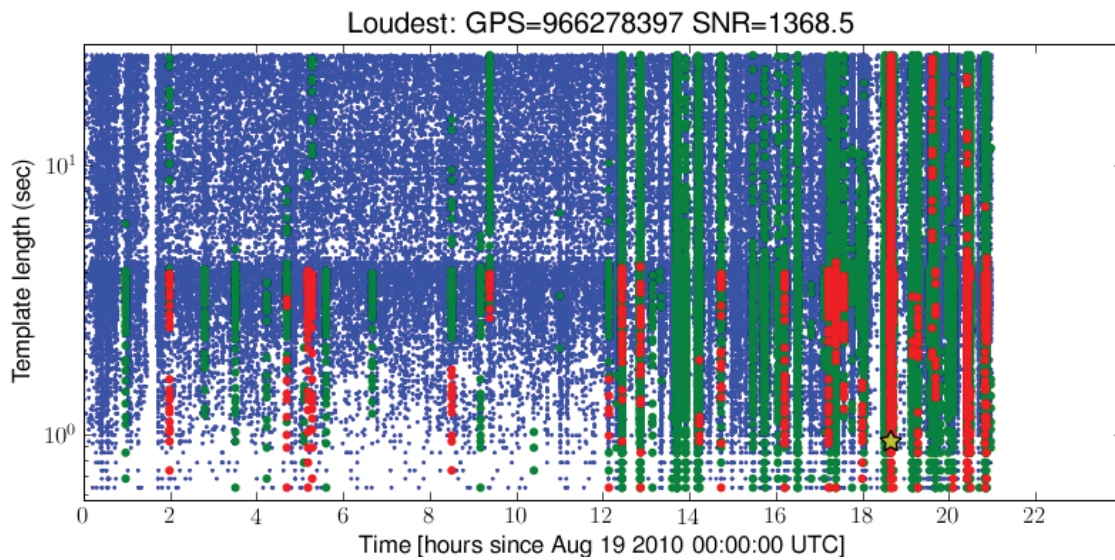
Figure 66 : A sample "glitchgram." Blue dots have ρ_{new} values below 8, green have values between 8 and 16, and red have values above 16. Template length was chosen as the Y axis in order to capture a feature of the templates that is not specific to gravitational-wave signals, such as chirp mass. Note the break at 4.3 seconds, corresponding to the chirp mass at which the bank switches from an overlap of 0.95 to 0.5. No data is analyzed after 21:00 because ihope requires at least 2048 contiguous seconds to estimate the PSD and all data after this time was in smaller segments.
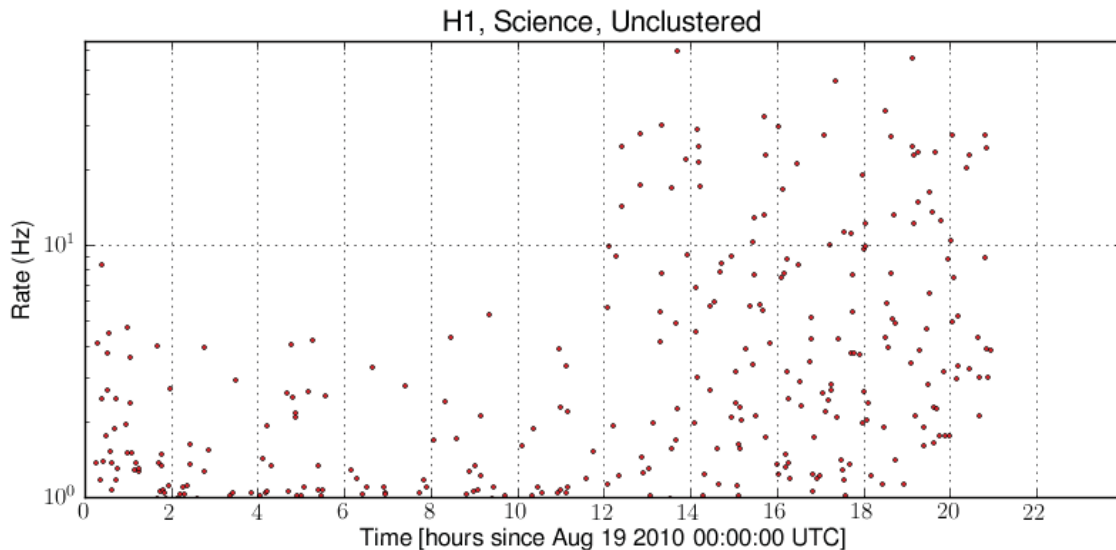
Figure 67 : Sample rate plot, showing rate in Hz (averaged over 1-minute intervals) as a function of time. Note that the rates increase after 12:00. This was due to increased seismic noise.

6. Rate vs. Time and SNR vs. Time

These plots complimented the glitchgram by breaking the triggers up differently. The rate plot (Fig. 67) showed average number of triggers over 1-minute intervals. The SNR plot (Fig. 68) showed the SNR of every trigger as a function of time. These tend to be correlated, as loud glitches ring off the entire bank and produce large numbers of triggers. The plots were accompanied by tables showing the times where the rate of triggers exceeded 500 Hz for more than one second, and exceeded 200 Hz for more than 10 seconds.

7. Breakdown by template

This page showed several histograms of number of triggers as a function of the length of templates in seconds. Examples of the standard histograms are shown in Fig. 69, they show that most of the triggers come from short templates, and templates shorter than 5 seconds produce up to six times as many triggers as shorter templates. Figure 70 shows the same information as a map of the template bank, color-coded to indicate how many triggers each template produced. Figure 71 breaks this same information up by time and template mass. Qualitatively these plots do not change much day-by-day, and so these plots were
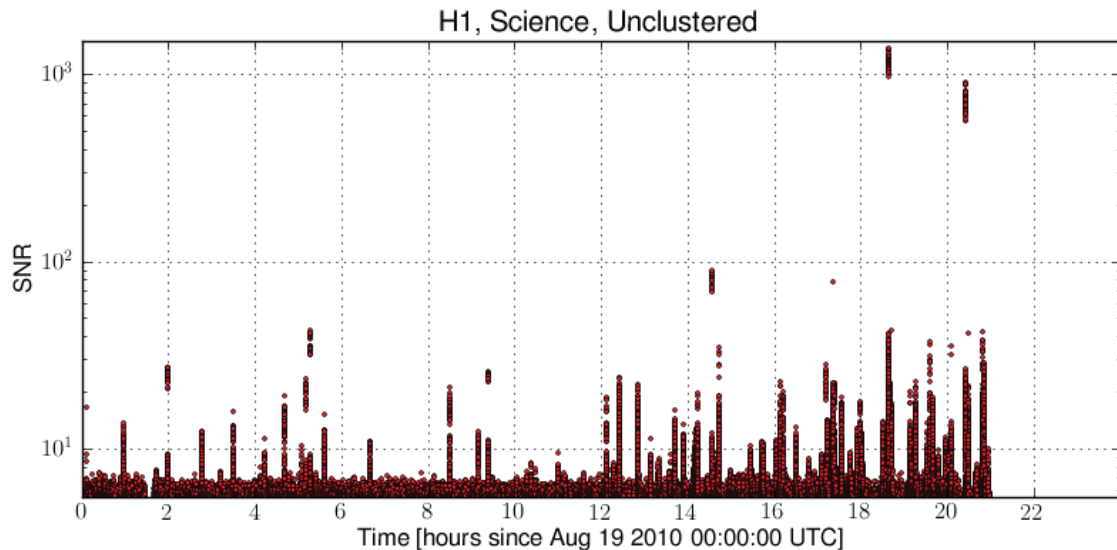
Figure 68 : Sample plot of SNR as a function of time. The density of triggers increases somewhat after 12:00 and there are more loud outliers. However the change in behavior is better seen in figure 67.

not often used. However, they did prompt a change to the analysis made early in the S6 run, see Sec. 10.3.

8. The χ^2 test

These plots showed all triggers with SNR values on the x-axis and $\chi^2/(2p - 2)$ (the reduced χ^2) on the y-axis. These plots are most useful when compared to a reference plot generated from a day of simulated Gaussian noise, shown in figure 72, and comparing the
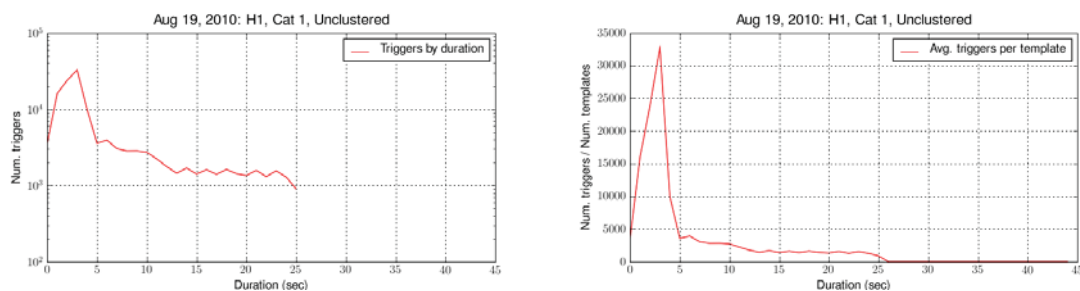

Figure 69 : Histograms of trigger rates by template length in daily ihope. The plot on the left combines all templates, the plot on the right normalizes by plotting (number of triggers resulting from templates of length x) divided by (number of templates of length x).
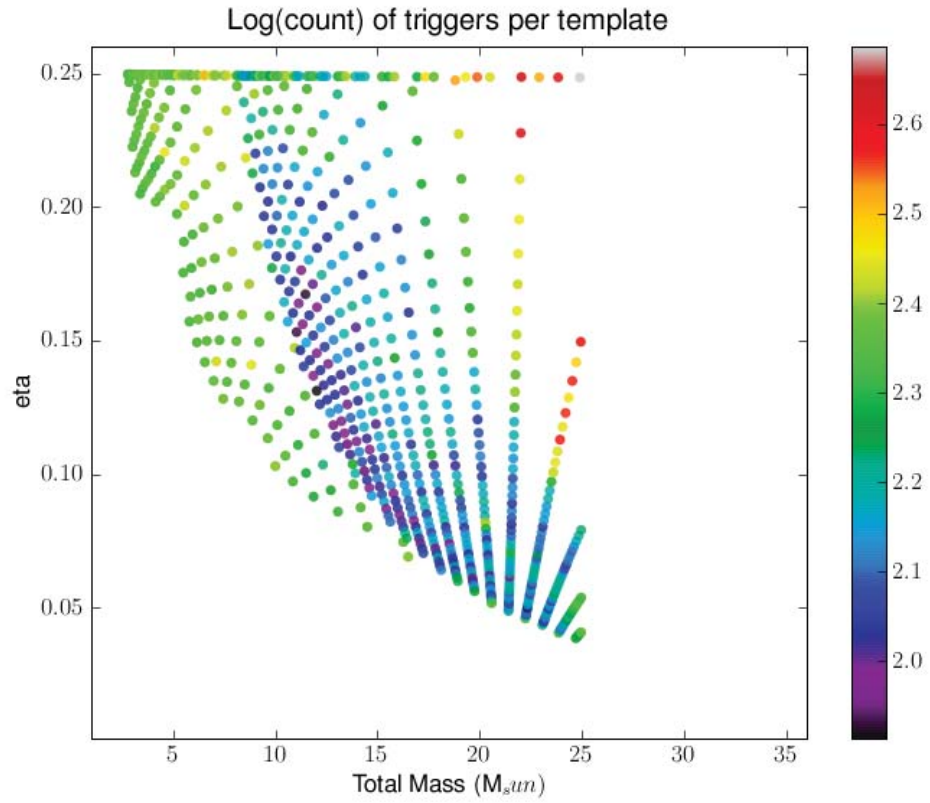
Figure 70 : The daily template bank, color-coded to show how many triggers each template produced. Note that the template in the upper-right corner, which is the template of shortest duration, produced approximately 1.3 times as many triggers as the next most active.
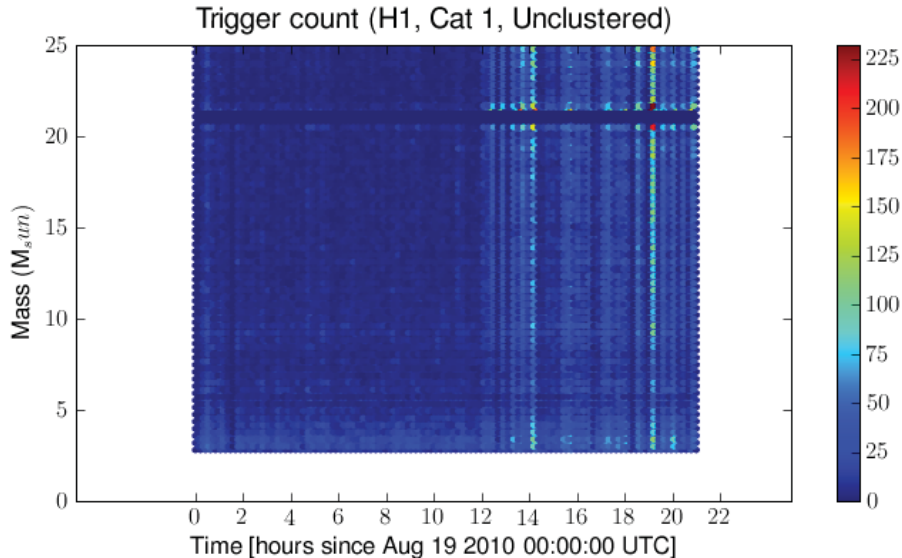
Figure 71 : Trigger rates as a function of time and template length. The elevated trigger rate after 12:00 is visible here as well. Note that particularly loud glitches, such as that around 19:30, ring off the entire bank.

plots at different veto levels, shown in figure 73.

10.3 Applications of Daily Ihope to Pipeline Tuning

The plots presented in the previous section provided a great deal of information about both the detectors and the response of the search software. These could be correlated with the behavior of the full search in order to identify the cause of problems and motivate solutions. We now give two examples.

Early in S6 there were frequent instances where we were unable to complete the full analysis. This was due to individual programs taking abnormally long to complete, consuming far more than the expected amount of memory, or failing outright. The problematic jobs tended to be individual runs of the `trigbank` program. This is the step at which triggers from the first stage are examined to determine which templates need to be used at the second stage (see Sec. 4.4. Comparing the times that caused problems to the daily pages immediately revealed a correlation – times over which `trigbank` were unable to run were those where the rates of triggers were abnormally high.

In S5 and the early weeks of S6 triggers were clustered using a method called *trigscan* [132] which attempts to collapse clusters of triggers that are close in time and parameters to a single most-significant trigger. Early versions of the daily page did this clustering as well,
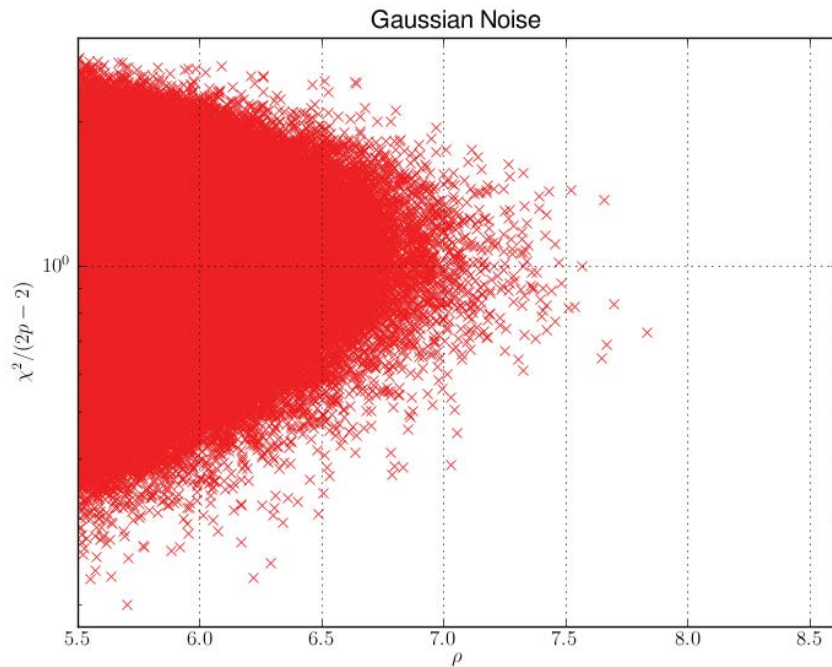
Figure 72 : The SNR/reduced χ^2 plane for a reference day of Gaussian noise. This is the product of a χ^2 distribution on the y axis and a Gaussian distribution on the x axis.
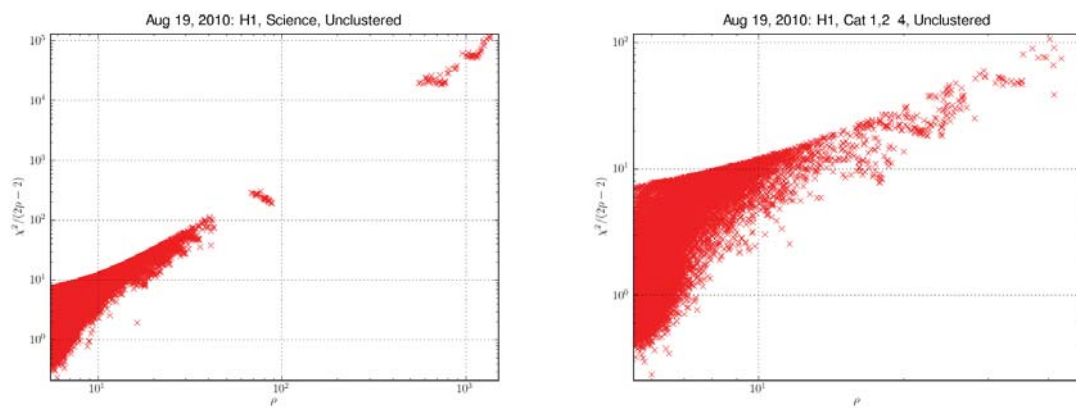
Figure 73 : SNR/reduced χ^2 plots of H1 data. The expected shape of figure 72 is discernible, but there are long tails of non-Gaussian glitches. The sharp cutoffs arise from thresholds within the inspiral code, see Sec. 4.1.4. There is a further population extending to the upper right at category 0 (left) that is removed by vetoes in category 4 (right).

Figure 74 : Problematic times identified by daily ihope. The plot on the left shows the rate of triggers over the day without any clustering. Note the sudden jump in rate at 5:40. On the right, the same plot after clustering triggers with the trigscan algorithm. The overall rate has been reduced by about a factor of two, indicated by the different scales on the y axes. However, the rate remains high enough to cause problems.
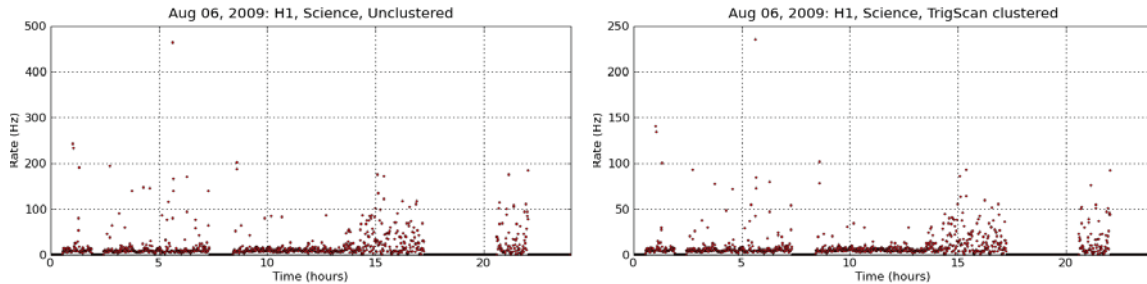
and comparison between the unclustered and trigscan-clustered triggers revealed that even after clustering periods of high trigger rates remained [1]. This is illustrated in Fig. 74. There is a short period of high trigger rate around 05:40 which remained high after clustering. A run of `trigbank` processing this time was unable to complete.

The possibility of vetoing such glitchy periods was raised, and this would have been easy to accomplish using the rate information from daily ihope. However, such vetoes would have needed to be CAT1 to avoid the problem, which would mean subdividing science segments and possibly losing short segments. Instead we replaced trigscan with fixed 30-millisecond clustering windows, after studies of software injections determined that this change did not harm the ability of the search to make detections.

At the start of S6 the range of the low-mass search extended up to $35M_\odot$, as it did in S5. However, along with times of high trigger rates, the daily ihope pages indicated that most of the triggers were coming from the high-mass end of the bank. This behavior is expected, as it is easier for short templates to match against glitches, but the trigger histograms highlighted the extent to which this was a problem. A sample plot of triggers per template from early S6 is shown on the left of Fig. 75. It shows that the shortest template in the bank, with $m_1 = m_2 = 17.5M_\odot$ was producing significantly more triggers than any other template. In part this was due to a bug in the template bank code that

---

[1] Trigscan worked well in S5, it is not known why it did not work as well in S6. It is possible that the detectors were simply more glitchy in the early days of S6. However, in order to group triggers together trigscan must use the bank metric, which was correct in S5 when 2.0 pN templates were used but incorrect in S6 when the analysis moved to 3.5 pN templates (see Sec. 4.3). Some preliminary investigations were performed, but results were inconclusive
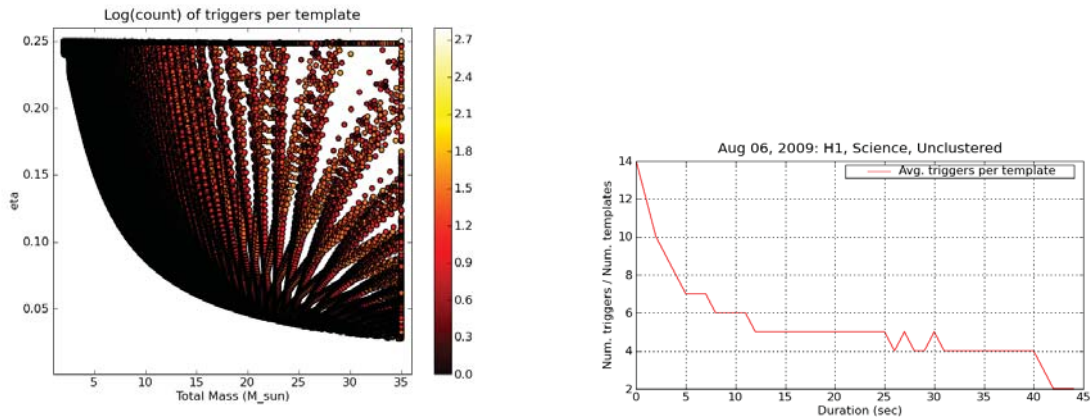
Figure 75 : Problematic templates seen in daily ihope. This shows the equivalent of figures 71 and 70 from an earlier version of daily ihope and an earlier version of the CBC search. Note the excess of triggers from the high-mass end of the bank, corresponding to shorter templates.

caused this template to appear twice in some banks. In part the abundance of triggers from this template is due to it appearing in every bank throughout the day, whereas other templates tend to get repositioned as the noise curve changes. Even taking these effects into account, most of the triggers come from templates shorter than 5 seconds, as seen on the right of Fig. 75.

The fixed clustering window means that only the loudest trigger in a 30-millisecond window will be passed to subsequent stages of the analysis. Given the numbers of triggers from short templates there was concern that a loud, short glitch could mask a quieter trigger from a binary neutron star coalescence. We therefore decided to limit the low-mass search to $M < 25M_\odot$ after comparing plots of injections in the same data made with the different mass ranges. These comparisons showed that the search with the smaller bank recovered injections in the mass range as well or better than the larger bank.

10.4 Applications of Daily Ihope to Individual Vetoes

Daily ihope was created primarily as a source of information to be used in construction of the veto definer file. We now define the veto categories used in S6 and present examples of the use of daily ihope in constructing each.

10.4.1   Category 1

Category 1 vetoes indicates time that should not have been marked as science mode. Typically attempting to analyze this time will adversely affect the entire 2048-second analysis chunk, for example by biasing the PSD (section 4.4). Note that it is possible to correct science time by creating a new version of the `DMT_SCIENCE` flag with an incremental version number. However, doing so is a more complicated process than creating a veto flag, and removing time by denoting it CAT1 carries additional information about the reason for the veto. As noted in section 9.5.2 category 1 vetoes are undesirable as they may break data into segments that are too short to be analyzed.

Glitches from the Thermal Compensation System

The design of the LIGO mirrors accounted for the fact that in operation they would be heated by the laser and the radius of curvature would change. However, the mirror's absorption of energy was larger than expected and the mirror overheated and deformed more than expected [133]. To correct for this a compensating laser was added that would heat the mirror in a ring in order to restore the radius of curvature to an optimal value. However, this laser would occasionally fluctuate, kicking the mirror through radiation pressure and producing a loud noise transient in the detector and loud triggers in the search.

The identification of these glitches was straightforward; they were often the loudest glitches of the day and were readily visible in the daily ihope "loudest trigger" report, as well as similar reports generated by Klinewelle and Omega. The cause was similarly easy to identify, as the omega scans generated by daily ihope showed egregious behavior in the TCS channels. An automated veto based on a threshold value in the TCS channel was added to the DMT at category 2. However, some instances were loud enough to interfere with the CBC pipeline in surrounding times for reasons we will discuss in Sec. 10.6. Hence it was decided to veto these times at category 1. These features are illustrated in Fig. 76.

Grid glitches

This was a loud, broad-band glitch that produced a grid of triggers in the Klinewelle pipeline's time-frequency plane, where the phenomena was first seen. They also showed up in the daily ihope rate and SNR plots as loud triggers accompanied by elevated trigger rates, as shown in Fig. 78. These also often appeared on the daily list of loudest glitches, the omega scans generated by daily ihope showed problems in magnetometers and the output
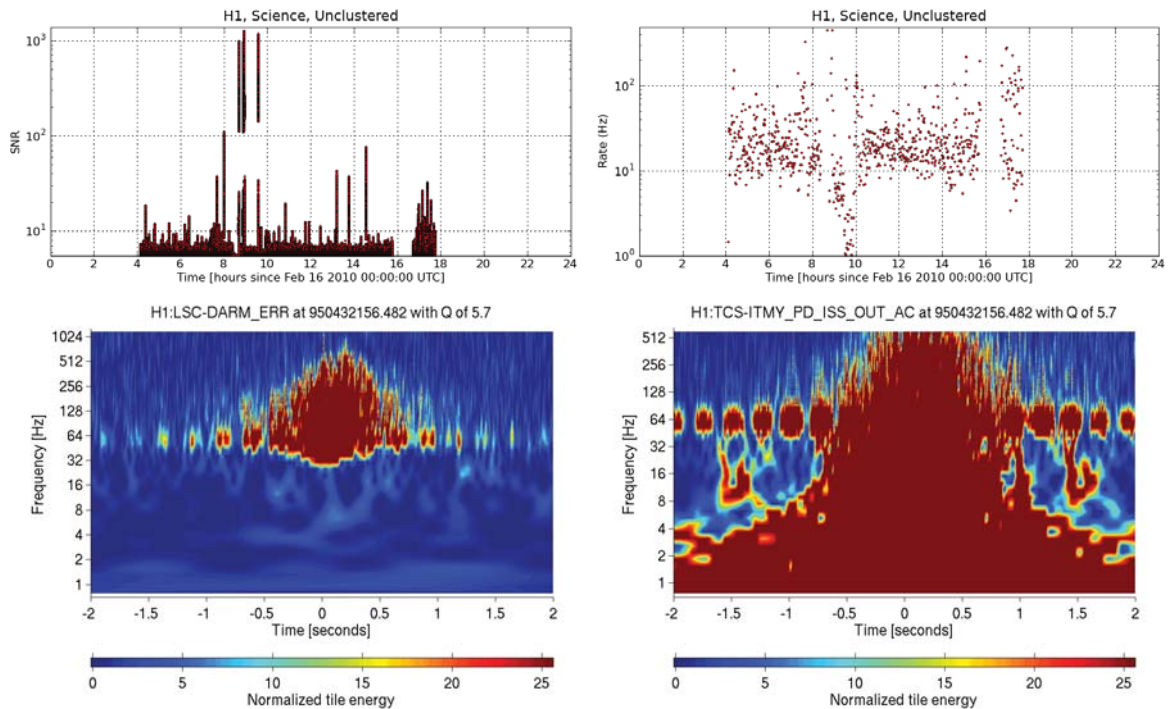
Figure 76 : TCS glitches in daily ihope and followup omega scans. The top left shows the trigger SNRs as a function of time, the loudest triggers are all TCS glitches. Note that each loud glitch comes in the form of a "tower" containing many triggers. This phenomena will be discussed further in Sec. 10.6.1. This is further illustrated by the trigger rate plot (top right), which shows elevated trigger rates at the times of the glitch. Note that around the time of the glitch the trigger rate actually *drops* significantly. This will be discussed in Sec. 10.6.2, but indicates that the glitches are loud enough to interfere with the analysis in surrounding times, justifying the use of a CAT 1 veto. The lower left plot shows the omega scan of the detector output channel at the time of the loudest glitch, clearly showing that this is not a gravitational wave. The lower right plot shows the omega scan of a channel monitoring the TCS, identifying it as the source of the glitch.

Figure 77 :   Grid glitches as seen by daily ihope. The plot on the left shows ρ_{new}, with the grid glitches indicated by arrows poin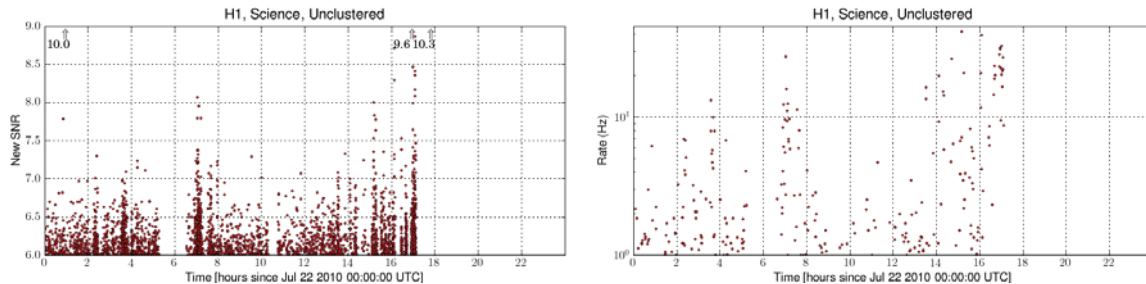ting above the range at which the plot cuts off (auto-scaling of this plot was removed from daily ihope to prevent a single loud glitch from squeezing all the other data to a thin region at the bottom of the plot). Note that the high value indicates that this glitch is fooling the χ^2 signal-based veto. The plot on the right shows the associated increase in trigger rate.

mode cleaner photodiodes, this is shown in Fig. 78. As with the TCS glitches, instances that showed up particularly loudly in daily ihope were hand-vetoed at category 1. The problem was eventually traced to a power supply, once this was resoldered the problem disappeared.

10.4.2   Category 2

Category 2 vetoes indicate times during which there was a problem, instrumental or environmental, with well-understood coupling into DARM_ERR. Such time can be analyzed without problem, but resulting triggers will be discarded from both the foreground and background.

The spike glitch

The spike glitch was a very loud, very short "bang" in the Livingston detector, characterized by a sudden drop in DARM_ERR immediately followed by a sharp jump. An example is shown in Fig. 79. In some cases there were a few cycles of ringing before the channel settled back down. These often produced SNR values of several thousand. The χ^2 values were typically large for the events themselves, resulting in negligible ρ_{new} values. However, the triggers in the tails (see Sec. 10.6.1) could have large new SNR values. A sample daily ihope plot with several spike glitches and representative omega scans are shown in Fig. 80.

Despite a great deal of effort the cause of these glitches was never identified. However, the distinctive shape allowed the creation of an automated veto. LSC_STRAIN was sampled every half-millisecond, giving a time series x_i. For each sample i the quantity
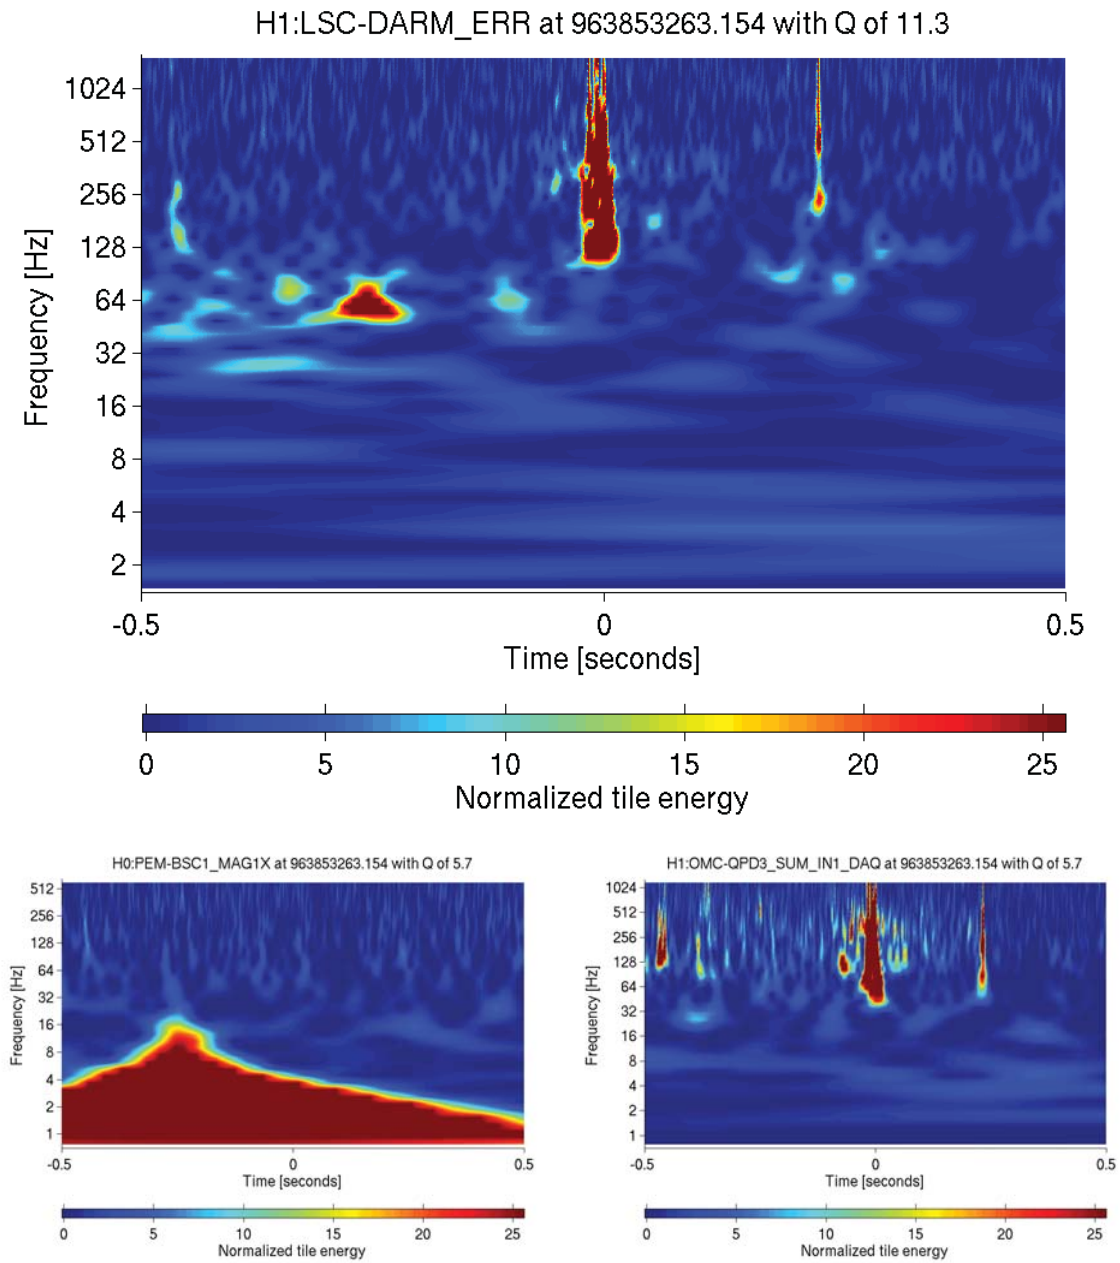
$$x_i - 2x_{i+1} + x_{i+2}$$

Figure 78 : Omega scans of the grid glitch flagged as the loudest trigger of the day and visible as an arrow in Fig. 77. At the top, the detector output channel, and at the bottom two channels that were associated with this class of glitch.
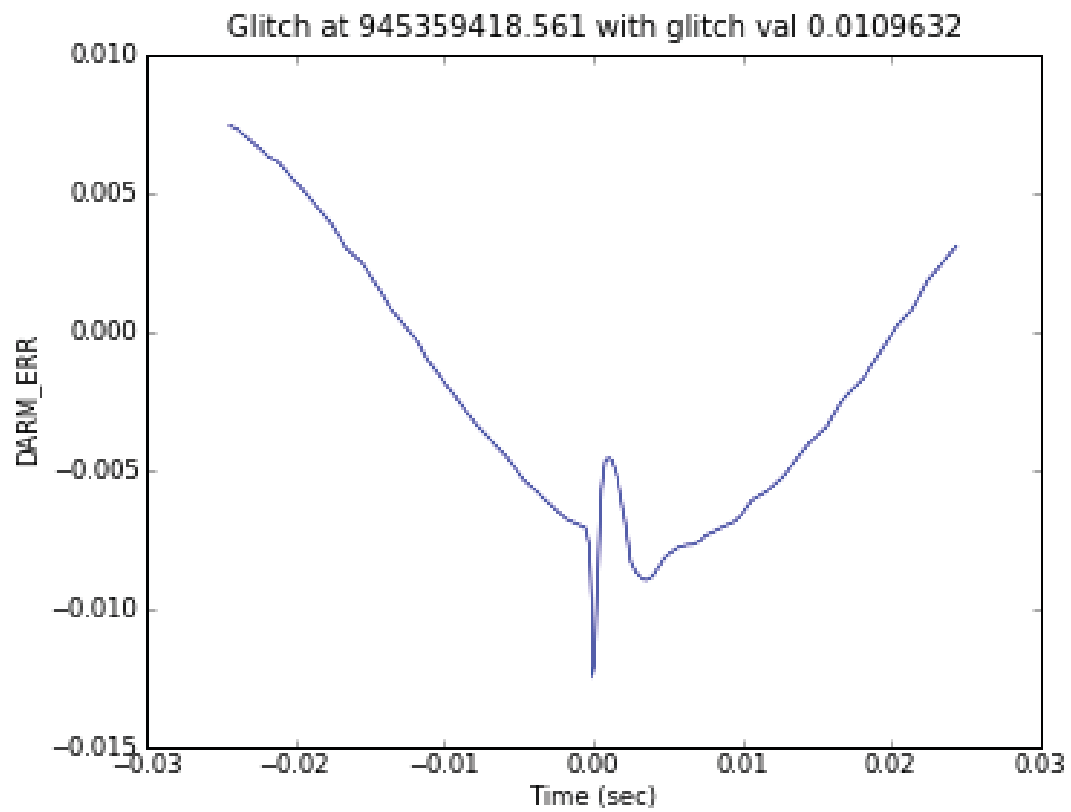
Figure 79 :   A sample spike glitches in DARM_ERR. Note the characteristic down-up-down pattern, and short timespan.

was calculated, values exceeding a threshold indicated a possible spike and the time was flagged in the segment database. Not every time so-flagged was a spike, but every instance indicated a potentially problematic rapid fluctuation in the data.

HEPI glitches

Livingston is subject to seismic activity, both natural and anthropogenic, that does not effect the Hanford detector. In order for the L1 detector to achieve the same low-frequency response it is therefore necessary to take additional steps. The Hydraulic External Pre-Isolation (*HEPI*) system sits between the ground and chambers containing optical components and provides a layer of active seismic isolation. This can reduce noise in the 1-3 Hz band by a factor of 10 [134].

In early 2010 it was noticed that the daily ihope triggers contained several instances with SNRs above 600 even after applying all known vetoes through category 4. This included removal of spike glitches and application of the use-percentage vetoes (discussed in the next section). Followup of these triggers in the daily loudest-glitch reports showed that they were all accompanied by loud glitches in the HEPI channels. Samples are shown in Fig. 81.

A trial veto was created by scanning the auxiliary channel time series for values exceeding 25,000. This veto was tested by applying it to the remaining daily ihope triggers, and it was found to be very effective. Histograms of triggers before and after application of this veto are shown in Fig. 82. After this confirmation the veto was applied at category 2 to the full search.

A great deal of work was done on HEPI during this period, and after Jan 10, 2010 the problem ceased.

10.4.3   Category 4

Category 3 vetoes remove hardware injections (Sec. 4.4). In the full analysis hardware injection vetoes do not have a category number and is denoted "hardware injections removed." As this veto category had a fixed meaning it was not used in daily ihope.

Category 4 is for time that appears to be ill-behaved according to data quality studies, but where there is no clearly-understood cause. In the full analysis this is denoted category
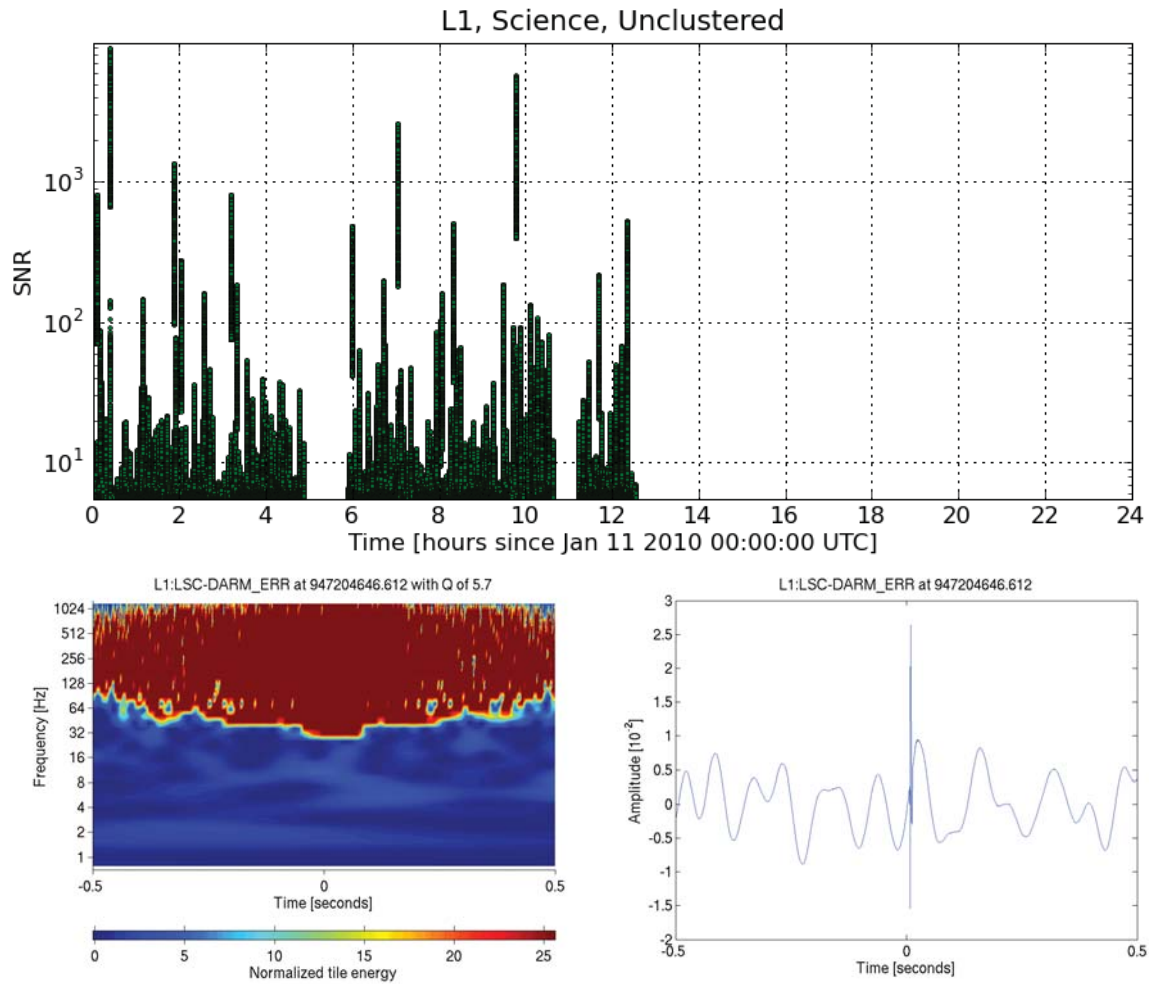
Figure 80 : Spike glitches in daily ihope and a typical omega scan. At the top a daily ihope SNR-vs-time plot, showing several spike glitches with SNRs of several thousand. On the bottom, the loudest of these spike glitches in an omega scan, showing the time-frequency plot (left) and unfiltered time series (right). The sharp drop-rise-drop of the time series behavior was reasonable well-captured by the filtered used to create data quality segments.
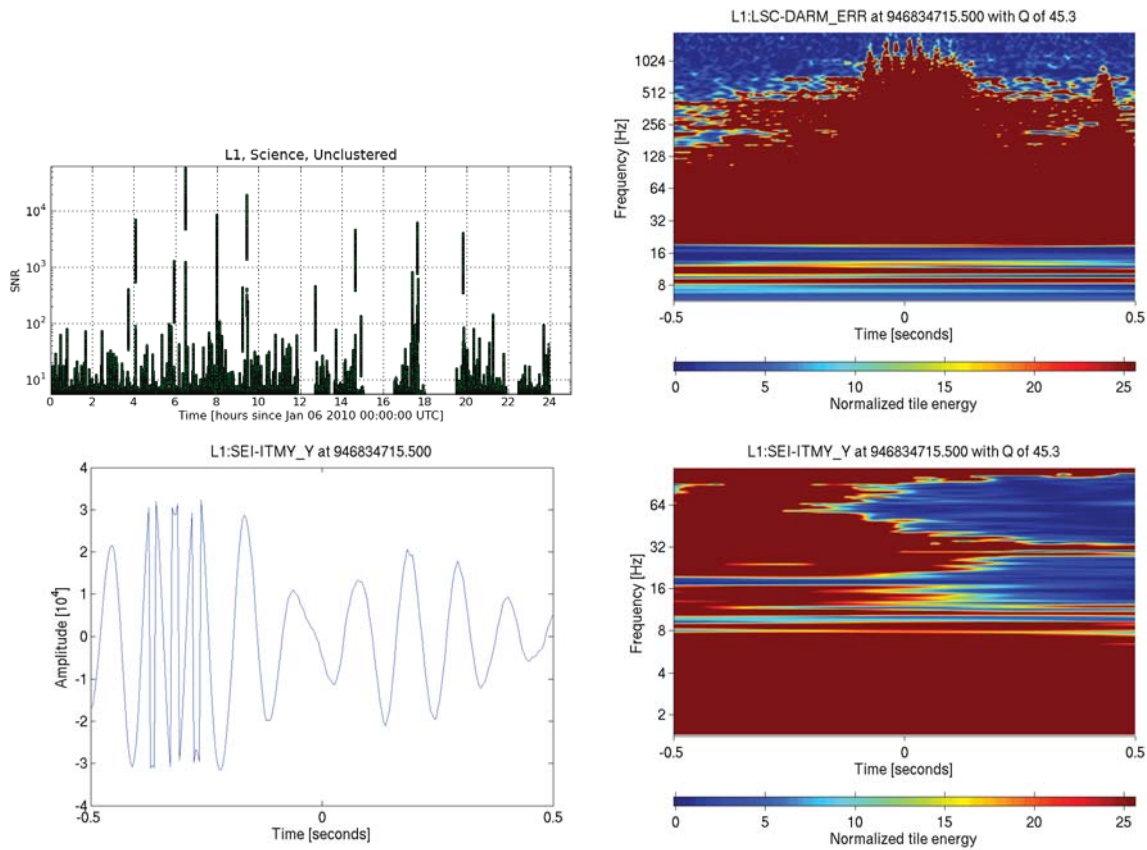
Figure 81 :   The HEPI glitch in daily ihope and Omega. On the top left, the SNR values for the day including a HEPI glitch at 17:38, as identified by the "loudest glitches" report. On the top right, this event in DARM_ERR. On the bottom left the time-frequency plot of a HEPI auxiliary channel, and on the bottom right the unfiltered time series of this same channel, showing that it has exceeded the threshold of 25,000
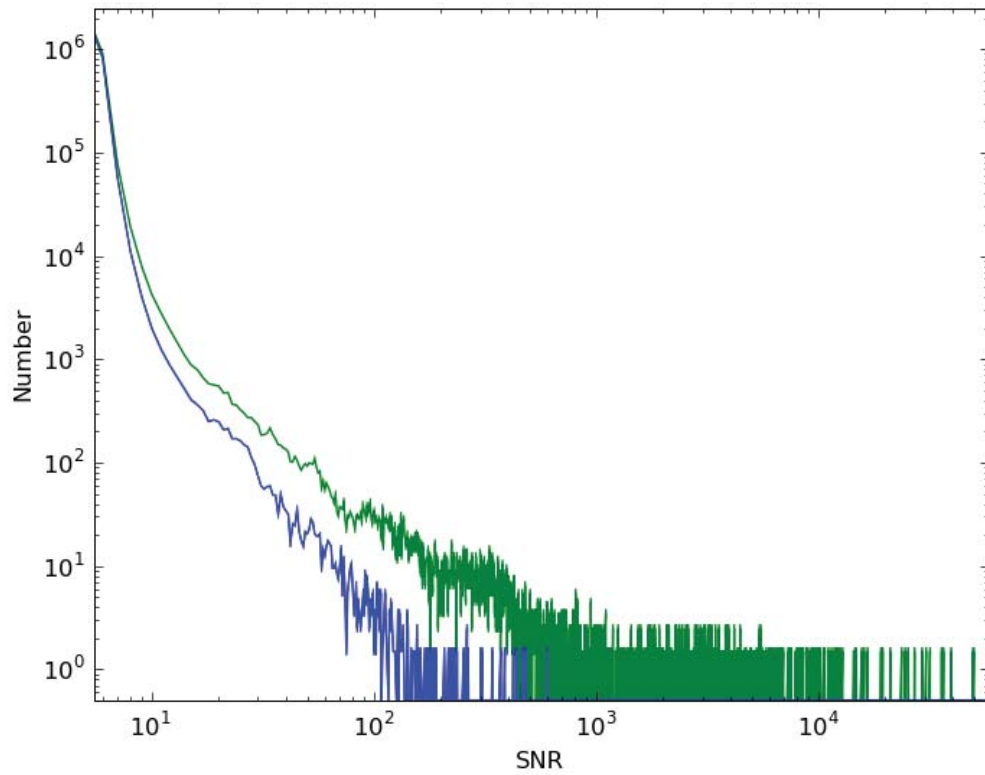
Figure 82 : Effectiveness of the HEPI veto. The green line shows the number of daily triggers after applying all known vetoes, including the removal of spike glitches. The blue line show counts after the additional HEPI veto is applied. Note that the highest SNR is reduced from 4,000 to 600, along with a reduction of triggers at SNRs down to ~10.

3.

Automated vetoes

There were three uses of the daily triggers through S6. The first two, the daily pages and ad-hoc veto studies, have already been discussed. The third was as input to systematic automated veto studies.

SeisVeto [135] compared low-frequency triggers from Omega running on seismic monitoring channels to daily ihope triggers. This was done using the *HVeto* [136] algorithm, which determines the significance of the correlation between two channels by computing

$$S = -\log 10 \left(\sum_{k=n}^{\infty} P(\mu, k) \right)$$

where n is the number of coincidences found between Omega triggers and daily ihope triggers, μ is the number expected to see by chance, and $P(\mu.k)$ is the Poisson probability of obtaining k instances when the expected value is μ. In other words, S is $-\log_{10}$ of the probability of observing as many or more coincidences between two series of random occurrences than were actually observed. Times of high correlation were flagged with category 4 vetoes, which could achieve up to 62% efficiency with only 6% deadtime.

The *Used Percentage Veto* (UPV) [137] program looked for correlations between triggers from `DARM_ERR` and auxiliary channels using a figure of merit defined as

$$\text{Used Percentage}(\rho) \equiv 100 \times N_{\text{coinc}}(\rho)/N_{\text{total}}(\rho)$$

where $N_{\text{total}}(\rho)$ is the total number of triggers from the auxiliary channel above significance threshold ρ in the analysis time, and $N_{\text{coinc}}(\rho)$ is the number of triggers from the auxiliary channel that lie within a second of a trigger from `DARM_ERR`. Intervals where this percentage exceeded 50% were entered into the segment database and applied as category 4 vetoes. Although UPV was originally developed using Klinewelle triggers for both `DARM_ERR` and auxiliary channels, it was modified during S6 to accept triggers from daily ihope. The resulting analysis was rerun weekly.

An example is shown in Fig. 83. During the week including this day UPV determined that the instrumental channel `ASC-WFS2_IY` had a high correlation with daily triggers. This is a quiet glitch, with an SNR of 6.5, and it would have taken considerable effort to track down manually. However, it is accompanied by 4 seconds of elevated trigger rates which could have impacted the search and should therefore be flagged at category 4.
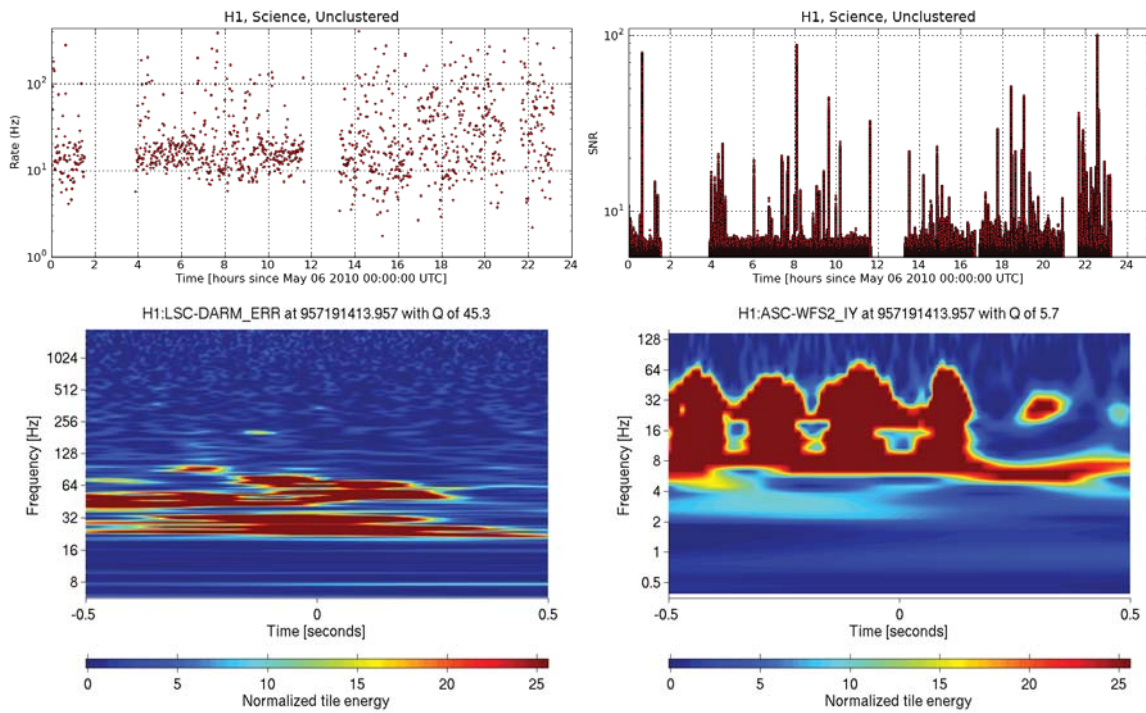
Figure 83 : A glitch found by UPV using the daily triggers. On the top the daily rate and SNR plots, including the glitch in question at 14:29. On the bottom the omega scan of DARM_ERR and the auxiliary channel whose Klinewelle triggers correlate with daily ihope DARM_ERR triggers.

ifo	Triggers (Count)	Vetoed (Count)	Efficiency	Time (sec)	Vetoed (sec)	Dead-time (sec)	Ratio
L1	2890507	12578	0.43	798720	880	0.11	3.95
H1	1692904	6452	0.38	647168	416	0.06	5.92

Table 16 : Effectiveness of the "SNR > 250" flag over the weeks 09/04/2010 to 09/17/2010. Note that the flag was used almost twice as often in L1 as H1, although there is only 23% more analysis time. This is another indication that L1 overall contained more non-Gaussian noise transients.

Loud SNR

Through the S6 run the *horizon distance*, the distance at which the coalescence of an optimally-oriented binary system consisting of two $1.4 M_\odot$ neutron stars would have an SNR of 8, was roughly 45 Mpc. The SNR scales inversely with distance, hence the distance at which we would expect to see such a system at, say, SNR 250 is 0.12 Mpc. Assuming uniform volume distribution, this makes an SNR 250 event 6.4×10^{-8} times less likely than an SNR 8 event. However, such loud glitches do occur in the data fairly often, in particular the spike glitch at L1.

Such loud glitches tend to have high χ^2 values which suppress them. However, around SNR 250 glitches tend to be accompanied by additional triggers spanning ± 8 seconds resulting from the interaction of the filter with the inverse spectrum truncation, see Sec. 10.6. Some of these auxiliary triggers can, by chance, have low χ^2 values and hence high new SNR values, and can potentially interfere with the search. This suggests a CAT 4 veto centered on times of triggers with SNRs exceeding 250 with 8 seconds of padding in both directions. Such a flag must be in place before the full run, making daily ihope the obvious choice for generating the flags.

This scheme was implemented starting on June 26, 2010, coinciding with the portion of the run designated S6D. An example of its effectiveness is shown in table 16 which shows the efficiency and deadtime off the flag applied to triggers from the full analysis after CAT 1, for the two weeks containing the blind injection. The efficiency to deadtime ratio is greater than 1, although still relatively small. Still, the flag was deemed useful as it removed triggers we could not have easily claimed were due to a gravitational wave.
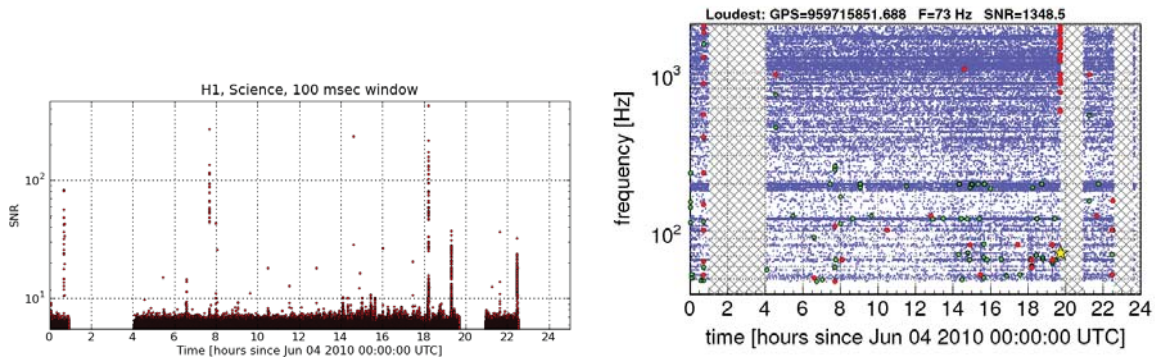
Figure 84 : Use of daily ihope to indicate that time did not need to be vetoed. The omega pipeline (left) shows a distinct feature at 200 Hz, within the sensitive band of LIGO. Daily ihope shows no excess of triggers or SNR, and hence the time did not need to be vetoed

10.4.4   Non-vetoed Time

It is at least as important to preserve time in which a detection could be made as it is to veto problematic time. On occasion there were potential problems seen in another pipeline or reported by SciMons or operators that daily ihope confirmed were not a problem for the CBC low-mass search (though they might have been problematic for other searches). One example is shown in Fig. 84. The Omega pipeline reports excess noise at 200Hz, which was traced to runoff from a nearby dam coupling into the detector. The corresponding daily ihope report shows no issues with excess or loud triggers, and the day was able to be analyzed.

10.5   Applications of Daily Ihope to the Blind Injection Challenge

We now move from general uses of daily ihope to its applications in following up a particular event. In addition to the hardware injections discussed in Sec. 4.4 it was known at the start of S6 that there would be any from zero to "a few" unannounced, blind hardware injections performed in order to provide an unbiased test of the search pipelines. One such injection was performed on Sep 16 2010 at 06:42 UTC, and showed up in multiple searches as a strong gravitational-wave candidate. This candidate was followed up by the CBC group to the point of writing a detection paper and submitting it to the collaboration for publication approval. Once approval had been granted the fact that it had been an injection was revealed. For more details on the event and how it was followed up, see the forthcoming S6

low-mass results paper, which is in preparation as of the time of writing.

Although daily ihope was not a search, the injection showed up on the page of loudest triggers for H1 and L1, with parameters shown in table 17 and omega scans shown in figure 85. The injection was not visible in V1 in daily ihope.

ifo	end_time	end_time_ns	SNR	χ^2	New SNR	Mchirp	DQ flags
H1	968654557	997314453	15	87	9.8	4.6	DMT-INSPIRAL_RANGE_STDEV_GT_0P50_MPC [968654544 968654560] DMT-INSPIRAL_RANGE_STDEV_GT_0P75_MPC [968654544 968654560] Light,Up,Calibrated,Science
L1	968654557	978027343	9.9	44	8.7	4.1	SCI-OTHER_ELOG [967120215 977875215] DMT-INSPIRAL_RANGE_STDEV_GT_1_MPC [968654544 968654560] DMT-INSPIRAL_RANGE_STDEV_GT_0P50_MPC [968654544 968654560] DMT-INSPIRAL_RANGE_STDEV_GT_0P75_MPC [968654544 968654560] SCI-FLAG_ERROR [967137346 977875215] Light,Up,Calibrated,Science

Table 17 : The blind injection as reported by daily ihope's "loudest triggers" page.
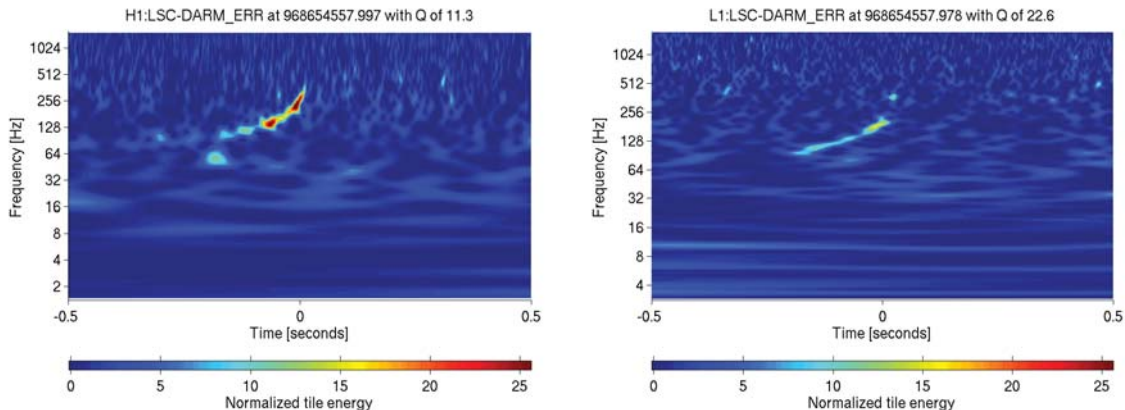
Figure 85 : H1 (left) and L1 (right) Omega scans of the injection as generated by daily ihope. Note the "chirp" shape which is the expected pattern from a compact binary inspiral.

10.5.1 False Alarm Rate Estimate

The potential detection occurred close to the end of S6 as the collaboration was preparing to take the LIGO detectors apart to install Advanced LIGO. The shutdown could potentially have been delayed if the existing configuration were necessary to vet the candidate. An ad-hoc committee was formed to determine whether this would be required, and while they found no immediate reason to delay the Advanced LIGO plans the report did include several recommendations, including:

> The committee urges a search in the S6 and S5 data for events with similar waveforms but lower SNR than the September 16 event in single detectors as well as in dual detector coincidence. Such an investigation has a dual purpose. First, it will establish some limits on the astrophysical source population producing the September 16 event and second, it will help in estimating the false alarm rate, although this will be better accomplished with more time slides [138].

The existing daily ihope triggers were ideal for this purpose, as they spanned all of S6 but with a reduced set of templates. This sampled the parameter space but produced a smaller set of triggers that made rapid significance studies computationally feasible. The significance was estimated by plotting histograms of all triggers throughout S6 and locating the candidate triggers in the resulting distribution. To parallel the full analysis the results were broken into mass bins. The low mass bin spans chirp masses up to $3.48 M_\odot$, the medium mass bin from $3.48 - 7.40 M_\odot$. Likewise, category 1,2 and 3 vetoes were applied
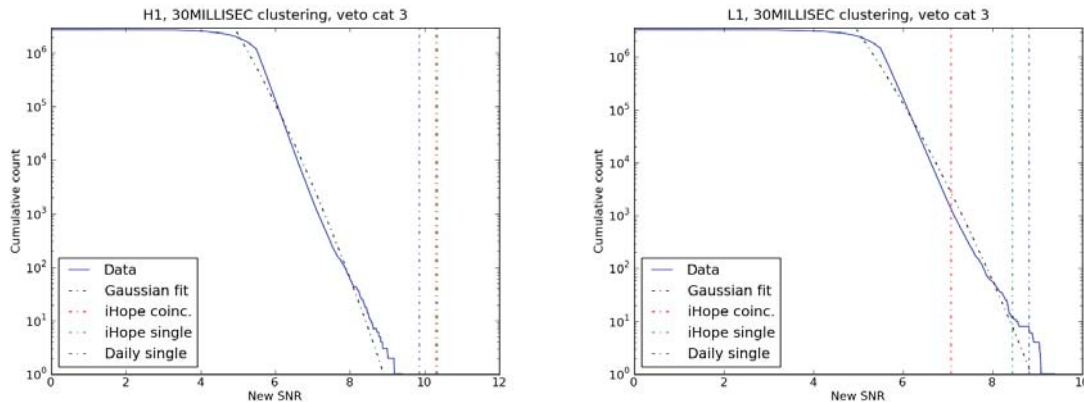
Figure 86 :   Significance of the blind injection in the low-mass bin in H1 (left) and L1 (right). Results are shown as cumulative histograms. The plots flatten out at low SNRs due the selection of triggers from the SNR time series, discussed in Sec. 4.1.4. The red lines are the new SNR reported by the full ihope run (after coincidence). The green lines are the loudest trigger in new SNR found at the first stage of the ihope analysis (before coincidence). The blue lines are the new SNR values of the injection found by daily iHope (no coincidence).

and triggers were clustered with a 30-millisecond window to to parallel the results of the full search. There are several ways of reporting the new SNR of the injection; the largest values reported by daily ihope, the largest single-detector values reported by the full search, and the component values of the largest combined new SNR reported by the full search. All of these options are included on the plot.

The results are shown in figure 86 for the low-mass bin and 87 for the medium-mass bin. The result in both bins is qualitatively the same. The injection is close to the loudest event in H1 for all measures of new SNR. The injection does not stand out as far in L1, which was known to contain more non-Gaussian noise transients over the course of S6.

10.5.2   Front-end Code Verification

Before the collaboration could claim a detection it was necessary to perform extensive checks to remove, or at least reduce, the possibility that the trigger was due to any source other than a gravitational wave. Consequently many components of the interferometer were subject to scrutiny. One such component was the *front-end control code*, which is responsible for providing the feedback loop between the photodiodes and servos that keeps the detector locked. This code is updated occasionally as new systems are added or bugs are found and fixed. To verify that the most recent change preceding the event did not significantly change the behavior of the detectors we compared histograms of triggers from
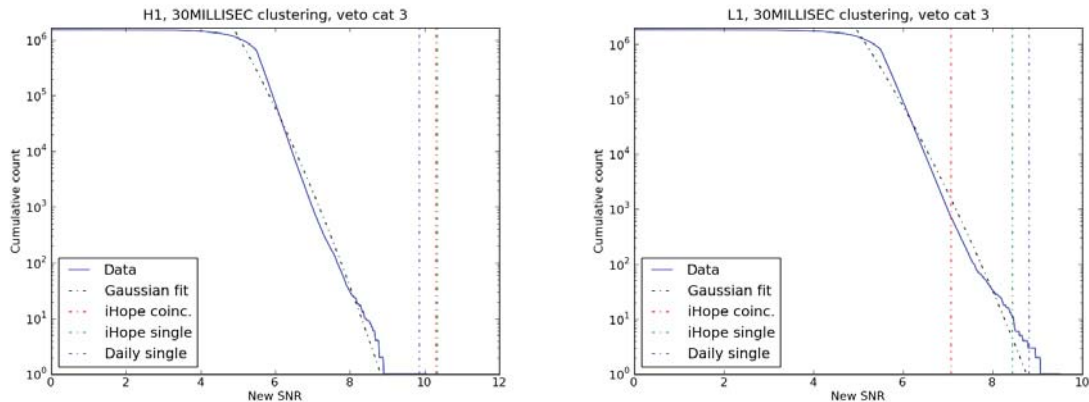
Figure 87 : Significance of the blind injection in the medium-mass bin in H1 (left) and L1 (right). Note the cumulative counts levels off around between 5 and 5.5, indicating that there are few triggers with smaller values.

daily ihope before and after these changes.

Two weeks prior to and following the most recent code changes at each site were selected. SNR histograms are shown in figure 88. There is a slight variation in H1, somewhat larger in L1. More rigorous testing could have been done, such as estimating the standard deviation in each SNR bin from several sample times before the code change and then checking whether the rates after the change fall within one sigma. However, the detection committee did not feel this level of analysis was necessary, and based on the plots in figure 88 concluded:

> Thus, to the extent allowed by the methods we adopted, there is no evidence
> for any malfunction in the front-end code of the interferometers [139].

10.6 Open Questions

We have seen in Sec. 10.3 that daily ihope was used to help identify and remedy problems in the low-mass search early in S6. We conclude this chapter by discussing two additional issues with the pipeline seen through daily ihope, both of which relate to the effect of loud glitches on the matched filter. Of course the ultimate goal is to remove such glitches at the source. However, it is likely that such glitches will continue to be present in the Advanced LIGO era and the search pipeline must be robust against them. We note here the problems and some initial attempts to remedy them, but more research will be needed to resolve them before Advanced LIGO comes on-line around 2015.
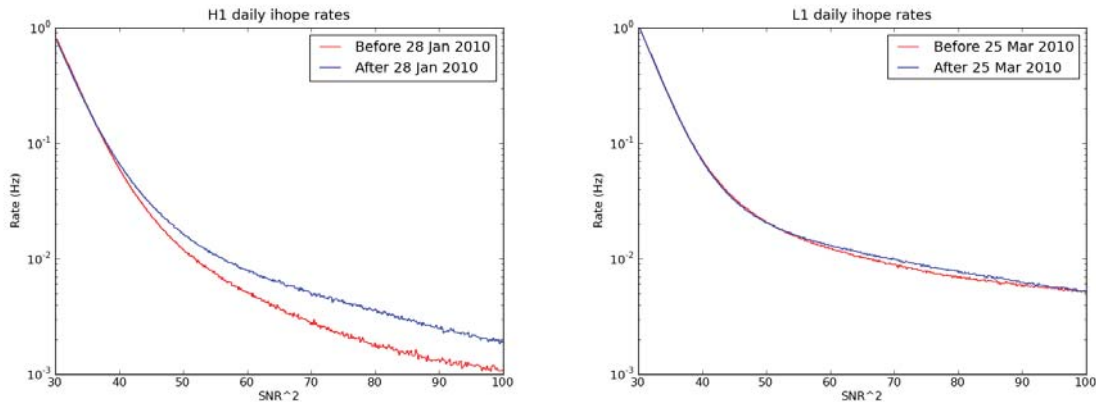
Figure 88 : SNR histograms comparing periods before and after front-end code changes at H1 (left) and L1 (right).

10.6.1 Excess Triggers Produced by Loud Glitches

From Fig. 82 we see that applying a veto to loud glitches does not only remove loud triggers, but also numerous triggers with lower SNRs. This same effect may be seen by comparing rate-vs-time and SNR-vs-time plots such as 77; loud glitches correlate with an increase in trigger rates. In part this behavior is expected. Many glitches, notably the spike glitch, are sharp enough that they may be roughly modeled as impulses. The impulse response of the matched filter (eqn. 4.13) is the time-reversed template convolved with a function of the noise curve. A loud glitch will therefore elevate the SNR time series for every template in the bank.

Recall from Sec. 4.1.4 that triggers are selected from the SNR time series by finding the largest value above threshold in a sliding window. The length of the window is taken to the length of the template, defined as the time required for the frequency of the pN waveform to go from 40 Hz to infinity. Based on this we would expect an impulse in the data to produce one trigger per template at approximately the time of the impulse. However this would not account for the number of templates seen or the length of time for which the trigger rate is elevated.

To study this in more detail simulated Gaussian noise was produced (as in the NINJA project) and the value of a single sample was increased to simulate a sharp glitch. The triggers produced for two values of the glitch amplitude are shown in Fig. 89. The expected behavior is seen in the rainbow "tower" at the top of both plots, successively longer templates have lower SNRs and trigger slightly later. Below this in both plots there is a
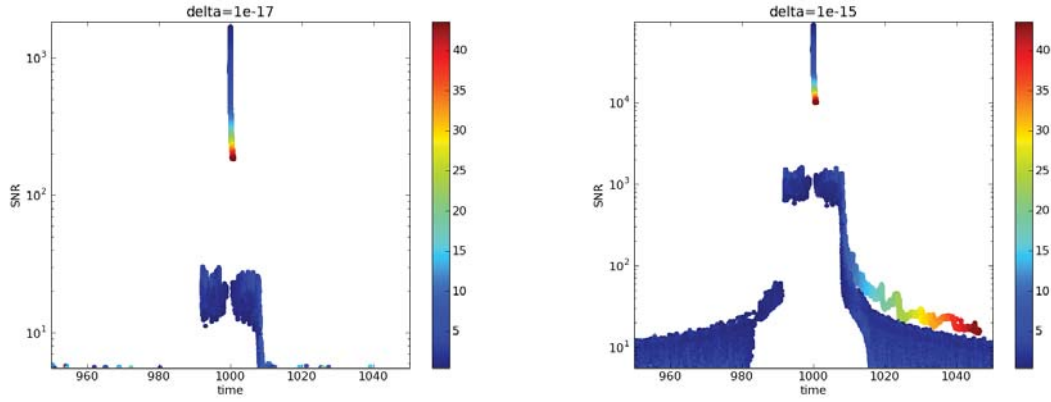
Figure 89 : Trigger SNRs as a function of time as the template bank responds to an impulse in the data. Color is the length of the template in seconds. On the left a single sample has been set to 10^{-17}, and on the right 10^{-15}. The expected response is visible, but there is a large number of additional triggers arranged in distinct features. See the text for discussion.

"plateau" of triggers from short templates. This results from the inverse spectrum truncation, described more fully in [140]. This behavior can be understood qualitatively as follows:

For simplicity denote the square root of the inverse PSD, $(S_n(|f|))^{-1/2}$ as $\tilde{S}(f)$, and its inverse Fourier transform in the time domain as $S(t)$. Likewise, denote the frequency-domain template as $\tilde{h}(f)$ as usual, and its inverse Fourier transform as $h(t)$. Finally, let $W(t)$ be a windowing function in the time domain with Fourier transform $\tilde{W}(f)$. In addition, denote multiplication of functions by $\cdot$ and convolution by $\star$. The application of the inverse spectrum truncation then proceeds as follows:

1. Calculate $\tilde{S}(f)$ and from it $S(t)$.

2. Apply the window in the time domain, giving $S(t) \cdot W(t)$.

3. Return to to the frequency domain, giving $\tilde{S}(f) \star \tilde{W}(f)$.

4. Square this (and correct the normalization, not shown here) giving

$$(\tilde{S}(f) \star \tilde{W}(f)) \cdot (\tilde{S}(f) \star \tilde{W}(f))$$

5. This replaces the $S_n(|f|)$ in the denominator of the match filter.

If the signal is then a delta function with strength M, $s(t) = M\delta(t)$ then $\tilde{s} = M$ and the
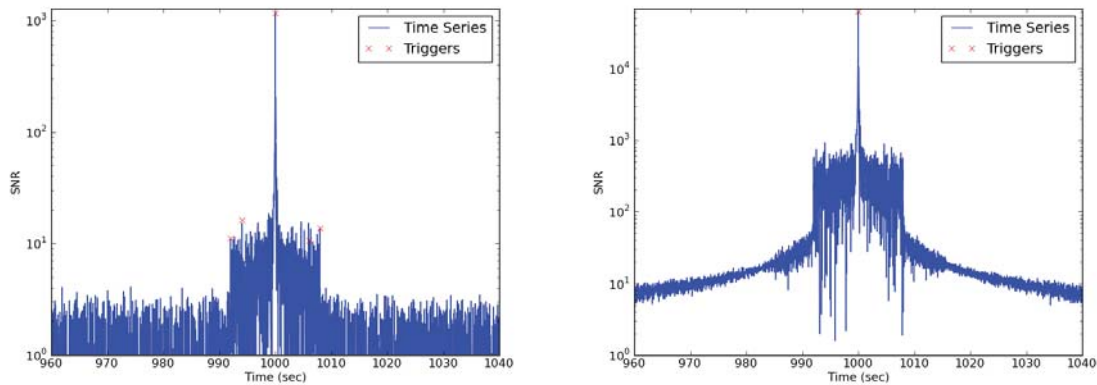
Figure 90 :   The SNR time series of a short (0.3 s) template responding to an impulse of strength 10^{-17} (left) and 10^{-15} (right). The elevation within the truncation window is clear, and is long enough to produce several triggers.

SNR time series is given by the matched filter,

$$\rho^2(t) = \int df \, e^{-2i\pi ift} M\tilde{h}^{\star}(f) \cdot (\tilde{S}(f) \star \tilde{W}(f)) \cdot (\tilde{S}(f) \star \tilde{W}(f))$$

$$= Mh(-t) \star (S(t) \cdot W(t)) \star (W(t) \cdot W(t))$$

Note that if the window function is zero outside a region then the elevated SNR from an impulse will likewise be bounded in time. This is the motivation for the truncation; without it a loud glitch would corrupt an entire analysis segment. However, when M becomes large the SNR value within the bounded region may exceed the trigger threshold. If the length used when scanning the SNR time series for triggers is less than the width of the truncation window then a loud glitch can produce several triggers. This explains both the 16-second width of the plateaus in Fig. 89 (as the window length is 16 seconds) and why they are composed of triggers from short templates. This may also be seen in the SNR time series shown in Fig. 90.

The plateau persists as the strength of the impulse is increased. Beyond a certain point we also get a second cluster of triggers from templates across the bank, seen to the right in Fig. 89. Note that the difference between triggers of the same color is precisely the length of template, identified by the same value in the colorbar. As the impulse-response of the filter is the time-reversed template, a sufficiently loud impulse will elevate the SNR for the length of the template, independently of how the spectrum is truncated. This can be seen in the SNR time series of a long template, shown in Fig. 91.

This second set of triggers from long templates suggests that the length used when
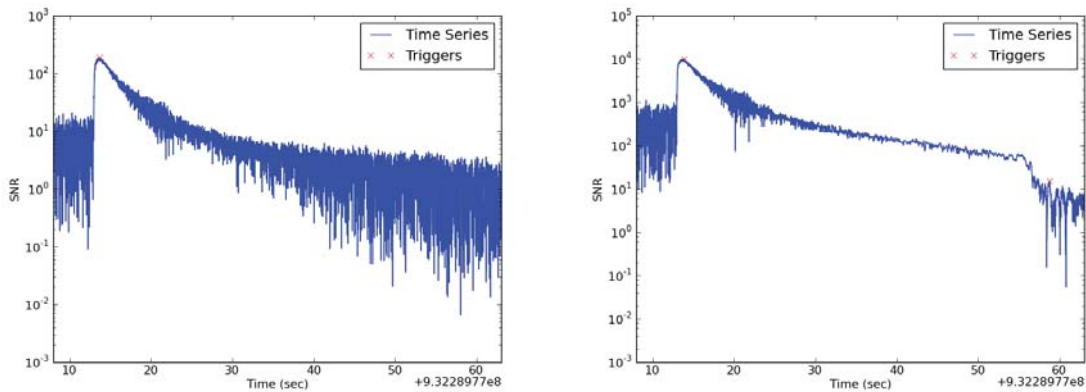
Figure 91 :   The SNR time series of a long (45 s) template responding to an impulse of strength 10^{-17} (left) and 10^{-15} (right). The SNR is elevated over the length of the template, and for the 10^{-15} impulse the SNR remains elevated long enough to generate a second trigger.

scanning the time series for is too short. This, combined with the excess of short triggers within the truncation window suggests replacing this length with the larger of the length of the truncation window or 1.1 times the currently used length (the exact value to be obtained by further investigation).

So far we have not used the χ^2 test. When χ^2 is enabled the trigger clustering algorithm (Sec. 4.1.4) is modified as follows

```
for each sample point j
  if ρⱼ > threshold
    if χ²ⱼ < chisq_thresh * (1 + ρ² * chisqfac ):
      if there is no event yet
        event_start = j
        event_snr   = ρⱼ
      else if ρⱼ > event_snr
        event_start = j
        event_snr = ρⱼ
      else if (j - event_start) == template length
        record event
        event_start = j
        event_snr   = ρⱼ
```

That is, an additional constraint is placed on triggers even before new SNR is calculated. We would expect that this would quash many, and hopefully all, triggers resulting from the
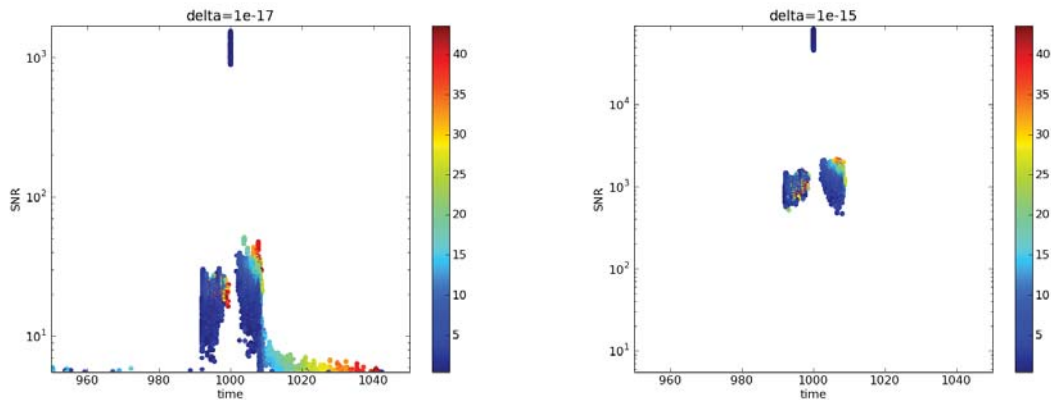
Figure 92 : Triggers produced by impulses of strength 10^{-17} (left) and 10^{-15} (right) when the χ^2 test is enabled. χ^2 removes many triggers, but many remain.

glitch. The results of rerunning with χ^2 enabled are shown in Fig. 92.

Turning on χ^2 does remove many triggers, in particular those in the original tower resulting from long templates. However, the triggers from short templates remain. This indicates that the χ^2 test is not as effective on short waveforms. The louder impulse on the right no longer produces the second rainbow of triggers. However, there are now long-template triggers in the plateau, and a set of low-SNR, loud-template triggers resulting from the quieter impulse. The reason for these is not clear, but they are have the potential to raise the background of the binary neutron-star search and are therefore problematic.

Despite passing the χ^2 clustering it is unlikely that these triggers, especially the ones from long templates, have good χ^2 values. We can attempt to remove them by altering the clustering algorithm as follows:
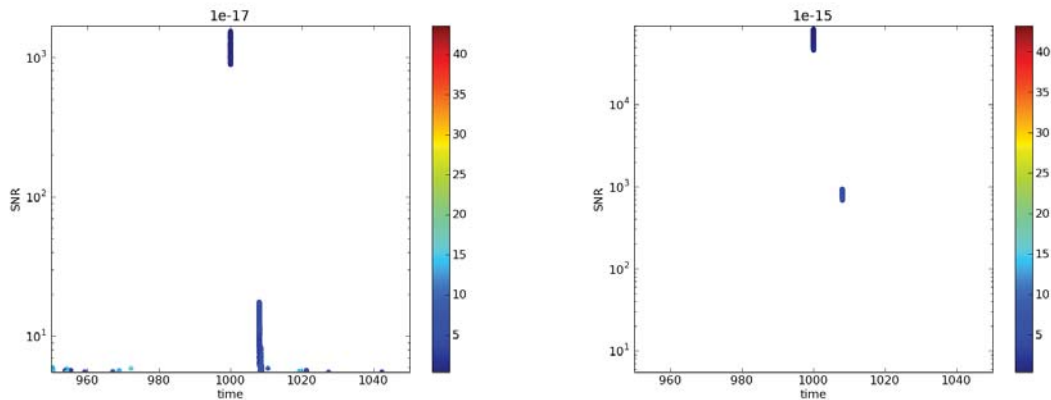
Figure 93 :   Triggers produced by the modified clustering algorithm which extends the clustering window and applies the χ^2 threshold after the candidate trigger has been found. Results are promising: only a relatively small number of triggers from short templates remain.

```
for each sample point j
  if ρⱼ > threshold
    if there is no event yet
      event_start = j
      event_snr   = ρⱼ
    else if ρⱼ > event_snr
      event_start = j
      event_snr = ρⱼ
    else if (j - event_start) == template length
      if χ²event_start < chisq thresh * ( 1 + ρ²event_start * chisqfac )
        record event
      event_start = j
      event_snr   = ρⱼ
```

Rather than applying the χ^2 test at each sample point, we cluster only on SNR and then use the χ^2 test to validate the candidate trigger. This was implemented, along with altering the clustering window as described above. The results are shown in Figure 93, and appear very promising. The plateaus have been removed almost entirely, leaving only the last trailing edge. Only a few triggers from short templates remain, and in particualr nothing that would interfere with the BNS search.

The revised code was then tested by performing injections of simulated signals into two weeks of real detector noise and examining the numbers of injections found and missed.
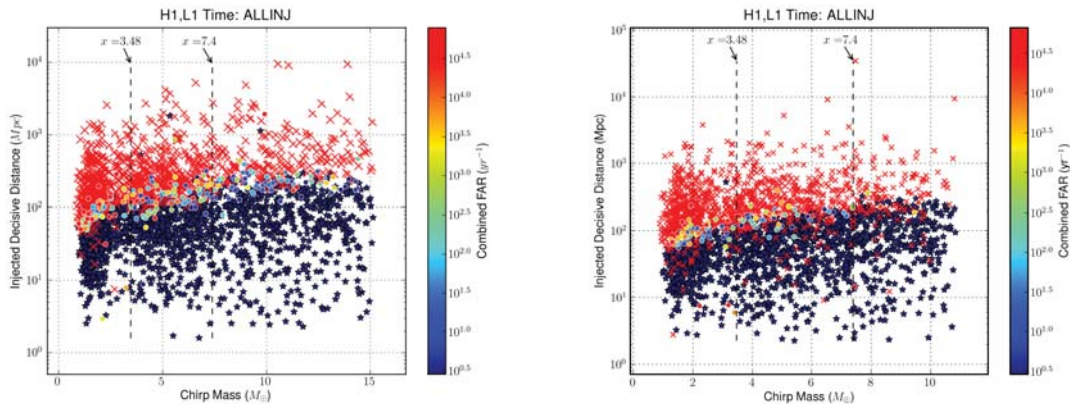
Figure 94 : Found and missed plots for simulated injection in two weeks of real noise. Results from the current code are on the left, the modified code (as described in the text) is on the right. The feature to focus on is found injections (blue dots) and missed injections (red crosses). There are many more missed injections using the modified code.

The new code misses many more injections, as shown in Fig. 94. To see why, consider a glitch with high SNR and high χ^2, within a clustering window of an injection with lower SNR and lower χ^2. In the original code the clustering window would not open at the glitch, because it would not pass the χ^2 test. The window would open at the time of the injection, and the trigger would be recorded. Under the new code, however, the higher SNR of the glitch will cause it to be the single event found in the window, but it will then be removed by the χ^2 test. This problem is exacerbated by extending the length of the window.

At present the issue of excess triggers from loud glitches remains unresolved. It is mitigated somewhat by the "loud SNR" veto at category 4, which extends ± 8 seconds in order to remove the excess triggers from the inverse spectrum truncation. However, as Fig. 89 shows, for very loud glitches excess triggers can be produced for as much as 45. An appealing possibility that has been discussed but not tested would be to run the current algorithm over a new SNR time series formed by combining the SNR and χ^2 series.

10.6.2   Bias of the PSD by Loud Glitches

Paradoxically, loud glitches are often accompanied by surrounding periods of decreased trigger rates. This can be seen, for example, in Fig. 76. This effect is confined to the 2048-second analysis segment containing the glitch, as can be seen from Fig. 95, which shows the trigger rate around a glitch as the analysis boundaries are changed.

This behavior is due the effect of a glitch on the PSD estimation for the analysis period.
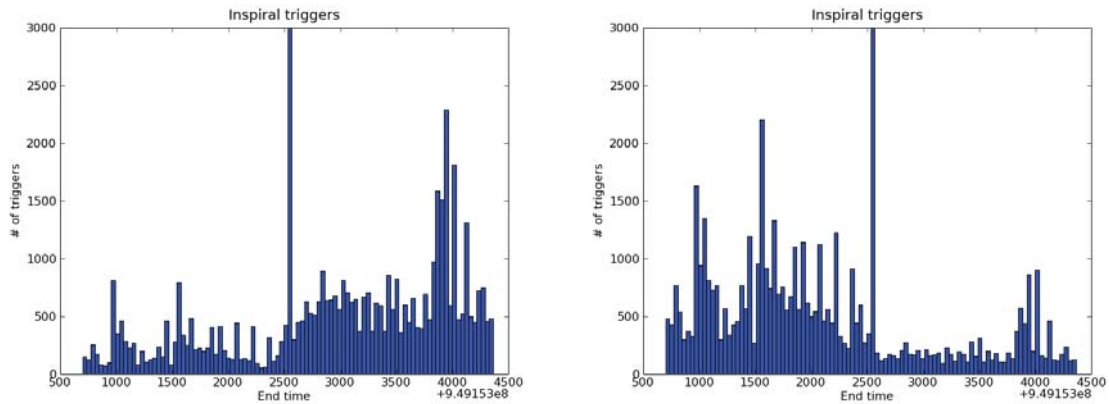
Figure 95 : Effect on trigger rates of a large glitch. The plots show the number of triggers in an H1 analysis in 1-second blocks around a large glitch centered at 2500. The analysis boundaries are arranged so that on the plot on the left the glitch falls at the end of the earlier segment, while on the right the glitch falls at the beginning of the later segment. In both cases the glitch has the same number of triggers, which is well in excess of the surrounding time for reasons discussed in the previous section. However, the segment containing the glitch shows a marked decrease in the number of triggers.

Recall from chapter 4 that the PSD is estimated by a variation of Welch's method, and the value of the PSD at each frequency f is the median over 15 256-second intervals, overlapped to span 2048 seconds. The median is used instead of the mean precisely because it is more robust against glitches. However, a value well outside the expected distribution can still bias the results to a significant extent. To demonstrate this we calculate the PSD in two ways, once by considering the median of all 15 PSDs and once by consider the median of 14, removing the one containing the glitch. We plot the fractional difference of each frequency bin in Fig. 96.

Some preliminary investigations using different-sized chunks and overlaps to compute the PSD were performed, but these have so far been inconclusive. At present the only way to ensure that a loud glitch won't suppress a quiet signal is to veto the time containing the glitch at category 1. This would be conceptually straightforward using something like the existing category 4 "loud SNR" veto, but it is far from an ideal solution. This is especially true in Advanced LIGO, where the low-mass templates will be significantly longer. A glitch in the middle of such a signal could cause the signal to be lost, as the category 1 veto would split the SNR accumulation into two disjoint segments. A better option would be to "gate" the data around a glitch; smoothly window out a second or so with a Tukey window or similar. This is an area of ongoing research.

Figure 96 : The bias in the PSD caused by a large glitch. This plot shows the difference between two PSDs, one calculated including t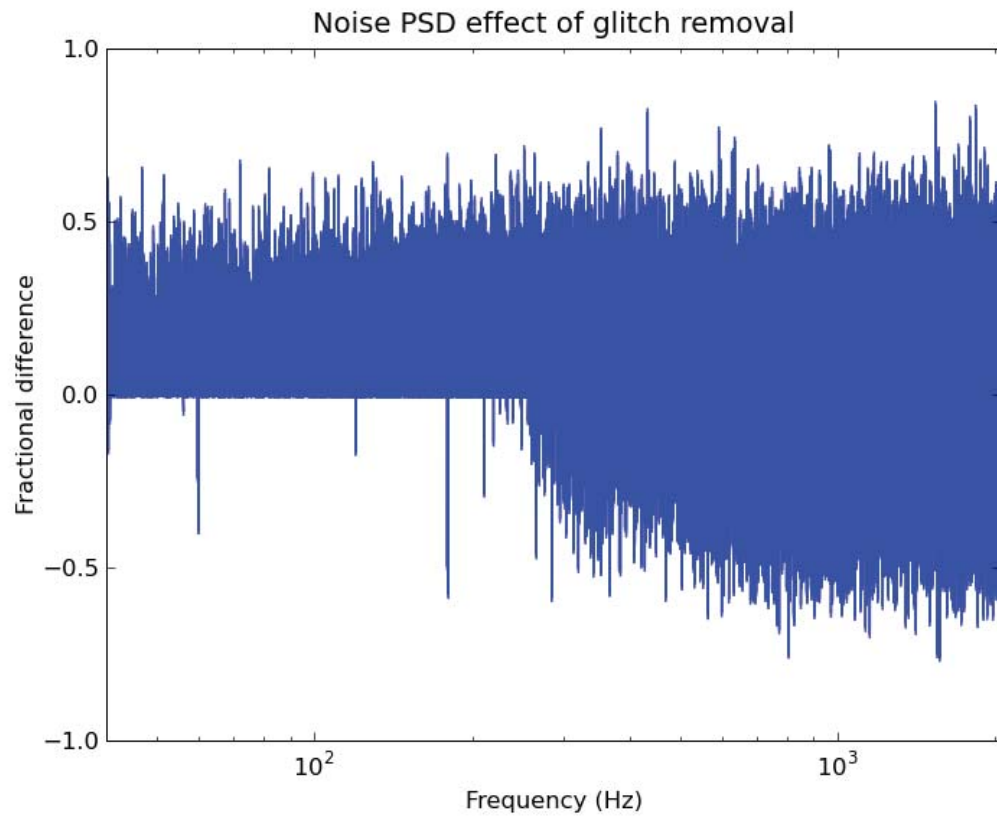he segment with the glitch and one without. The results are shown as the fraction difference in each mass bin, (with - without) / with. There is a notable bias upwards over much of the frequency range, and in particular over the most sensitive portion of the LIGO band.

Chapter 11

Conclusions

We have compared the post-Newtonian waveforms currently used to search for gravitational waves against a waveform produced by a full numeric simulation. Based on these studies we have identified several factors which can improve the overlaps between the two, and hence between analytic waveforms and the signals we expect to see.

We have tested these optimizations in the first NINJA project, which suggests that the improvement in overlaps carries through to improved SNR in a full search. This in turn implies that by implementing these changes we may be able to enhance the significance of signals in LIGO and Virgo.

We will further test these ideas in NINJA-2, which is currently on-going. In order to maximize the utility of NINJA-2 and correct for some of the limitations of NINJA-1 we have performed extensive validations of the hybrid pN/NR waveforms. A NINJA-2 data set is now available, and the unmodified CBC pipelines have been run. The results contain unexpected features which will need to be understood before we can use these results as a baseline against which to test our modifications. NINJA-2 will also enable many other studies. In particular we plan to use it to determine the optimum total mass at which the low mass and high mass searches should transition.

We have also shown how the tools and methods of the CBC search can be applied to detector characterization and used to identify noisy times in the detectors and remove them from analysis. The results from these detector characterization studies also proved to be tremendously useful in performing follow up studies of a candidate gravitational-wave event, although the event ultimately turned out to be a blind injection.

With the end of S6 the Initial LIGO era has ended. The interferometers are currently undergoing major upgrades in preparation for Advanced LIGO, to start around 2015. At

the design sensitivity Advanced LIGO will have ten times the range of Initial LIGO. This new era will present new challenges to data analysts, however the basic structure of the current templated, matched filter search will continue to be used. It will therefore be useful to continue optimizing the search against the Initial LIGO noise curve and Initial LIGO data, as NINJA-2 will do. In addition it is likely future NINJA projects will begin to use simulated Advanced LIGO noise. Similarly, daily ihope will continue to be run in Advanced LIGO, and perhaps even earlier, during engineering runs as the new detectors are being tested. We expect that continued studies with numeric waveforms and the use of CBC-specific detector characterization will contribute towards making the potential of LIGO as a new window to the Universe a reality.

Appendix A

Post-Newtonian Waveforms

A.1 Post-Newtonian waveforms in the adiabatic approximation

There are two basic elements to obtaining a post-Newtonian waveform: (1) finding the orbital phase of the binary, and (2) using that phase to find the so-called "amplitude" of the waveform. Previous references have been incomplete, or predate recent errata involving spin terms [141, 142]. Here, we gather together the most complete and current formulas for phase and amplitude when spins are *non-precessing* (i.e. spins aligned or anti-aligned with the orbital angular momentum), using consistent notation.

A.1.1 Phasing

The orbital phase evolution $\Phi(t)$ of the binary can be computed by considering energy conservation. The energy of the full system is accounted for in three parts, each computed as a function of the post-Newtonian expansion parameter

$$v := \left(M \frac{d\Phi}{dt} \right)^{1/3} . \tag{A.1}$$

The first part is the kinetic and gravitational binding energy of the binary—the orbital energy $E(v)$. As the system evolves, it gives off energy to infinity in the form of gravitational waves accounted for as the flux $\mathcal{F}(v)$ leaving the system. Finally, we must also account for the tide raised on each black hole by the other, and the flow of energy into the black holes due to the motion of these tides, given as a rate of change in the mass of the black holes $\dot{M}(v)$. Using this threefold accounting for the energy, we can express the conservation of

energy as

$$\frac{dE}{dt} + \mathcal{F} + \dot{M} = 0 \ . \tag{A.2}$$

Now, because the expression for the orbital energy is written in terms of v, we can straightforwardly differentiate to find $E'(v)$. With the chain rule, $dE/dt = E'(v)\,dv/dt$, we can rearrange this into a differential equation for v:

$$\frac{dv}{dt} = -\frac{\mathcal{F}(v) + \dot{M}(v)}{E'(v)} \ . \tag{A.3}$$

Given the expressions for $\mathcal{F}(v)$, $\dot{M}(v)$, and $E'(v)$, this equation can be integrated to find $v(t)$. Then, using the definition of v, we see that

$$\frac{d\Phi}{dt} = \frac{v^3}{M} \ , \tag{A.4}$$

which can be integrated in turn to find $\Phi(t)$. We now exhibit the formulas for $\mathcal{F}(v)$, $\dot{M}(v)$, and $E'(v)$, and discuss various methods for integrating the balance Eqn. (A.3).

Given the masses M_1 and M_2 and spin vectors $\mathbf{S}_1$ and $\mathbf{S}_2$, we define the following parameters:

$$M := M_1 + M_2 \ , \tag{A.5}$$

$$\eta := M_1 M_2 / M^2 \ , \tag{A.6}$$

$$\delta := (M_1 - M_2)/M \ , \tag{A.7}$$

$$\boldsymbol{\chi}_i := \mathbf{S}_i / M_i^2 \ , \tag{A.8}$$

$$\boldsymbol{\chi}_s := (\boldsymbol{\chi}_1 + \boldsymbol{\chi}_2)/2 \ , \tag{A.9}$$

$$\boldsymbol{\chi}_a := (\boldsymbol{\chi}_1 - \boldsymbol{\chi}_2)/2 \ . \tag{A.10}$$

We also define the quantities χ_s and χ_a to be the components of the spin vectors perpendicular to the orbital plane, namely $\chi_s := \boldsymbol{\chi}_s \cdot \boldsymbol{\ell}$ and $\chi_a := \boldsymbol{\chi}_a \cdot \boldsymbol{\ell}$, where $\boldsymbol{\ell}$ is the unit vector along the Newtonian angular momentum.

The orbital energy function can be written in terms of the PN expansion parameter v

defined above as [143, 144, 145, 142, 141, 146, 147][1]

$$E(v) = -\frac{M\eta v^2}{2}\left\{1 + v^2\left(-\frac{3}{4} - \frac{\eta}{12}\right) + v^3\left[\frac{8\,\delta\chi_a}{3} + \left(\frac{8}{3} - \frac{4\eta}{3}\right)\chi_s\right]\right.$$
$$+ v^4\left[-2\delta\chi_a\chi_s - \frac{\eta^2}{24} + (4\eta - 1)\chi_a^2 + \frac{19\eta}{8} - \chi_s^2 - \frac{27}{8}\right]$$
$$+ v^5\left[\chi_a\left(8\,\delta - \frac{31\delta\eta}{9}\right) + \left(\frac{2\eta^2}{9} - \frac{121\eta}{9} + 8\right)\chi_s\right]$$
$$\left.+ v^6\left[-\frac{35\eta^3}{5184} - \frac{155\eta^2}{96} + \left(\frac{34445}{576} - \frac{205\pi^2}{96}\right)\eta - \frac{675}{64}\right]\right\}.$$
(A.11)

We simply take the derivative of this formula with respect to v to find the energy function appearing in the phasing formula:

$$E'(v) = -M\eta v\left\{1 + v^2\left(-\frac{3}{2} - \frac{\eta}{6}\right) + v^3\left[\frac{20\delta\chi_a}{3} + \left(\frac{20}{3} - \frac{10\eta}{3}\right)\chi_s\right]\right.$$
$$+ v^4\left[-6\,\delta\chi_a\chi_s - \frac{\eta^2}{8} + (12\eta - 3)\chi_a^2 + \frac{57\eta}{8} - 3\chi_s^2 - \frac{81}{8}\right]$$
$$+ v^5\left[\chi_a\left(28\delta - \frac{217\delta\eta}{18}\right) + \left(\frac{7\eta^2}{9} - \frac{847\eta}{18} + 28\right)\chi_s\right]$$
$$\left.+ v^6\left[-\frac{35\eta^3}{1296} - \frac{155\eta^2}{24} + \left(\frac{34445}{144} - \frac{205\pi^2}{24}\right)\eta - \frac{675}{16}\right]\right\}.$$
(A.12)

Similarly, the flux function can be written as [143, 144, 145, 142, 141, 146, 147][2]

$$\mathcal{F}(v) = \frac{32}{5}v^{10}\eta^2\left\{1 + v^2\left(-\frac{1247}{336} - \frac{35}{12}\eta\right) + v^3\left[-\frac{11\delta\chi_a}{4} + \left(3\eta - \frac{11}{4}\right)\chi_s + 4\pi\right]\right.$$
$$+ v^4\left[\frac{33\delta\chi_a\chi_s}{8} + \frac{65\eta^2}{18} + \left(\frac{33}{16} - 8\eta\right)\chi_a^2 + \left(\frac{33}{16} - \frac{\eta}{4}\right)\chi_s^2 + \frac{9271\eta}{504} - \frac{44711}{9072}\right]$$
$$+ v^5\left[\left(\frac{701\delta\eta}{36} - \frac{59\delta}{16}\right)\chi_a + \left(-\frac{157\eta^2}{9} + \frac{227\eta}{9} - \frac{59}{16}\right)\chi_s - \frac{583\pi\eta}{24} - \frac{8191\pi}{672}\right]$$
$$+ v^6\left[-\frac{1712}{105}\ln(4v) - \frac{1712\gamma}{105} - \frac{775\eta^3}{324} - \frac{94403\eta^2}{3024} + \left(\frac{41\pi^2}{48} - \frac{134543}{7776}\right)\eta + \frac{16\pi^2}{3} + \frac{6643739519}{69854400}\right]$$
$$\left.+ v^7\left[\frac{193385\pi\eta^2}{3024} + \frac{214745\pi\eta}{1728} - \frac{16285\pi}{504}\right]\right\},$$
(A.13)

---

[1]The 1.5PN and 2.5PN spin terms were taken from Eq. (7.9) of [141]. The 2PN spin term and all nonspinning terms were taken from Eq. (C4) of [142]. Note that Eq. (C5) in the original published version of [142] is erroneous.

[2]The 1.5PN and 2.5PN spin terms were taken from Eq. (7.11) of [141]. The 2PN spin term and all nonspinning terms were taken from Eq. (C10) of [142], except that the term $\eta\left\{-\frac{103}{48}(\chi_s^2 - \chi_a^2) + \frac{289}{48}[(\chi_s \cdot \ell)^2 - (\chi_a \cdot \ell)^2]\right\}$ is omitted. (The authors of [142] have confirmed that this term should not be present.) Also note that Eq. (C11) in the original published version of [142] is erroneous.

where γ is the Euler Gamma.

Alvi [148] derived an expression for the transfer of energy from the orbit to each black hole by means of tidal heating. The calculation involves computing the deformation of each hole's horizon due to the Newtonian field of the other, then using that expression in formulas for energy absorption due to tidal deformation. In particular, his expression is applicable in the comparable-mass case. By combining the rates of mass change for both black holes, we obtain the total rate of change:

$$\dot{M}(v) = \frac{32}{5} v^{10} \eta^2 \left\{ -\frac{v^5}{4} \left[(1 - 3\eta)\chi_s(1 + 3\chi_s^2 + 9\chi_a^2) + (1 - \eta)\delta\chi_a(1 + 3\chi_a^2 + 9\chi_s^2) \right] \right\} .$$
(A.14)

The coefficient above is the leading-order term in the flux, meaning that this term is comparable to a relative 2.5PN spin effect in the flux. A similar calculation has been carried out in the extreme-mass-ratio limit [149], and agrees with this formula in that limit. Note that higher-order spin terms were calculated in [148], but are not included here, as they are at relative 3.5PN order, which is higher than the relative 2.5PN order to which other spin terms are known. Except for its explicit presence in the balance equation (A.3), we always treat the mass as a constant. This leads to additional errors at the 3.5PN spin level, which we ignore.

Below we will define some of the standard variants of computing the post-Newtonian phase from the energy and flux functions, using the naming convention of [150].

TaylorT1 phasing

The TaylorT1 approximant is computed by numerically integrating the ordinary differential equation for $v(t)$ in Eq. (A.3), using the expressions for orbital energy, flux, and mass change given in Eqs. (A.12), (A.13), and (A.14). The phase is then computed using this result for $v(t)$ in Eq. (A.4).

TaylorT4 phasing

The TaylorT4 approximant is similar to the TaylorT1 approximant, except that the ratio of the polynomials on the right-hand side of Eq. (A.3) is first expanded as a Taylor series,

and truncated at consistent PN order. Explicitly, the formula to be integrated is

$$
\begin{aligned}
\frac{dv}{dt} = \frac{32}{5M} v^9 \eta \bigg\{ & 1 + v^2 \left[-\frac{11\eta}{4} - \frac{743}{336} \right] + v^3 \left[-\frac{113\delta\chi_a}{12} + \left(\frac{19\eta}{3} - \frac{113}{12} \right) \chi_s + 4\pi \right] \\
& + v^4 \left[\frac{81\delta\chi_a\chi_s}{8} + \frac{59\eta^2}{18} + \left(\frac{81}{16} - 20\eta \right) \chi_a^2 + \left(\frac{81}{16} - \frac{\eta}{4} \right) \chi_s^2 + \frac{13661\eta}{2016} + \frac{34103}{18144} \right] \\
& + v^5 \left[-\frac{189\pi\eta}{8} - \frac{4159\pi}{672} + \left(\frac{3\eta}{4} - \frac{3}{4} \right) \delta\chi_a^3 + \left(\frac{9\eta}{4} - \frac{9}{4} \right) \delta\chi_a\chi_s^2 + \left(\frac{1165\eta}{24} - \frac{31571}{1008} \right) \delta\chi_a \right. \\
& \qquad \left. + \left(-\frac{79\eta^2}{3} + \frac{27\eta\chi_a^2}{4} - \frac{9\chi_a^2}{4} + \frac{5791\eta}{63} - \frac{31571}{1008} \right) \chi_s + \left(\frac{9\eta}{4} - \frac{3}{4} \right) \chi_s^3 \right] \\
& + v^6 \left[-\frac{1712\gamma}{105} - \frac{5605\eta^3}{2592} + \frac{541\eta^2}{896} + \left(\frac{451\pi^2}{48} - \frac{56198689}{217728} \right) \eta + \frac{16\pi^2}{3} + \frac{16447322263}{139708800} \right. \\
& \qquad - \frac{1712 \ln(4v)}{105} + \left(\frac{1517\eta^2}{72} - \frac{23441\eta}{288} + \frac{128495}{2016} \right) \chi_s^2 + \left(\frac{565\delta^2}{9} + \frac{89\eta^2}{3} - \frac{2435\eta}{224} + \frac{215}{224} \right) \chi_a^2 \\
& \qquad \left. + \left(\left(\frac{128495\delta}{1008} - \frac{12733\delta\eta}{144} \right) \chi_a + \frac{40\pi\eta}{3} - \frac{80\pi}{3} \right) \chi_s - \frac{80\pi\delta\chi_a}{3} \right] \\
& + v^7 \left[\frac{91495\pi\eta^2}{1512} + \frac{358675\pi\eta}{6048} - \frac{4415\pi}{4032} \right. \\
& \qquad + \left(-\frac{11\eta^2}{24} + \frac{979\eta}{24} - \frac{505}{8} \right) \chi_s^3 + \left(\frac{\delta\eta^2}{8} + \frac{742\delta\eta}{3} - \frac{505\delta}{8} \right) \chi_a^3 \\
& \qquad + \left(\left(\frac{3\eta^2}{8} + \frac{917\eta}{12} - \frac{1515}{8} \right) \delta\chi_a + 12\pi \right) \chi_s^2 \\
& \qquad + \left(\left(-124\delta^2 - \frac{3397\eta^2}{24} + \frac{7007\eta}{24} - \frac{523}{8} \right) \chi_s - 48\pi\eta + 12\pi \right) \chi_a^2 \\
& \qquad + \left(\frac{2045\eta^3}{216} - \frac{398017\eta^2}{2016} + \frac{10772921\eta}{54432} - \frac{2529407}{27216} \right) \chi_s + 24\pi\delta\chi_a\chi_s \\
& \qquad \left. \left. + \left(-\frac{41551\delta\eta^2}{864} + \frac{845827\delta\eta}{6048} - \frac{2529407\delta}{27216} \right) \chi_a \right] \right\} .
\end{aligned}
$$

$$(A.15)$$

Note that this expression does not include *all* of the spin-dependent terms at 3PN and 3.5PN, since the spin terms in the energy and flux functions are known only up to 2PN and 2.5PN, respectively. However, the 3PN and 3.5PN terms shown here will still be present in this formula when the higher-order terms are included in the energy and flux formulas.

TaylorT2 phasing

Expanding the inverse of Eq. (A.3) allows for the analytical integration of $t(v)$. The result reads

$$t(v) = t_0 - \frac{5M}{256\eta\, v^8} \left\{ 1 + v^2 \left[\frac{11\eta}{3} + \frac{743}{252} \right] + v^3 \left[-\frac{32\pi}{5} + \frac{226\delta\chi_a}{15} + \left(\frac{226}{15} - \frac{152\eta}{15} \right) \chi_s \right] \right.$$

$$+ v^4 \left[\frac{3058673}{508032} + \frac{5429\eta}{504} + \frac{617\eta^2}{72} - \frac{81}{4}\delta\chi_a\chi_s - \left(\frac{81}{8} - \frac{\eta}{2} \right)\chi_s^2 - \left(\frac{81}{8} - 40\eta \right)\chi_a^2 \right]$$

$$+ v^5 \left[-\frac{7729\pi}{252} - \frac{13\pi\eta}{3} + \left(\frac{147101}{756} - \frac{4906\eta}{27} - \frac{68\eta^2}{3} \right)\chi_s + \left(\frac{147101}{756} + \frac{26\eta}{3} \right)\delta\chi_a \right.$$

$$\left. + (6 - 6\eta)\delta\chi_s^2\chi_a + (6 - 18\eta)\chi_s\chi_a^2 + (2 - 6\eta)\chi_s^3 + (2 - 2\eta)\delta\chi_a^3 \right]$$

$$+ v^6 \left[\frac{6848\gamma}{105} - \frac{10052469856691}{23471078400} + \frac{128\pi^2}{3} + \left(\frac{3147553127}{3048192} - \frac{451\pi^2}{12} \right)\eta - \frac{15211\eta^2}{1728} + \frac{25565\eta^3}{1296} \right.$$

$$+ \frac{6848\ln(4v)}{105} - \left(\frac{584\pi}{3} - \frac{448\pi\eta}{3} \right)\chi_s - \frac{584\pi\delta\chi_a}{3} + \left(\frac{6845}{672} - \frac{43427\eta}{168} + \frac{245\eta^2}{3} \right)\chi_s^2$$

$$\left. + \left(\frac{6845}{672} - \frac{1541\eta}{12} + \frac{964\eta^2}{3} \right)\chi_a^2 + \left(\frac{6845}{336} - \frac{2077\eta}{6} \right)\delta\chi_s\chi_a \right]$$

$$+ v^7 \left[-\frac{15419335\pi}{127008} - \frac{75703\pi\eta}{756} + \frac{14809\pi\eta^2}{378} + \left(\frac{4074790483}{1524096} + \frac{30187\eta}{112} - \frac{115739\eta^2}{216} \right)\delta\chi_a \right.$$

$$+ \left(\frac{4074790483}{1524096} - \frac{869712071\eta}{381024} - \frac{2237903\eta^2}{1512} + \frac{14341\eta^3}{54} \right)\chi_s + (228\pi - 16\pi\eta)\chi_s^2$$

$$+ (228\pi - 896\pi\eta)\chi_a^2 + 456\pi\delta\chi_s\chi_a - \left(\frac{3237}{14} - \frac{14929\eta}{84} + \frac{362\eta^2}{3} \right)\chi_s^3$$

$$- \left(\frac{3237}{14} - \frac{87455\eta}{84} + 34\eta^2 \right)\delta\chi_a^3 - \left(\frac{9711}{14} - \frac{39625\eta}{84} + 102\eta^2 \right)\delta\chi_s^2\chi_a$$

$$\left. \left. - \left(\frac{9711}{14} - \frac{267527\eta}{84} + \frac{3574\eta^2}{3} \right)\chi_s\chi_a^2 \right] \right\} .$$

$$(A.16)$$

The comment made below Eq. (A.15) about spin contributions at 3PN and 3.5PN order is valid for Eq. (A.16) and the following expansions as well.

The orbital phase Φ can be integrated similarly to the time t. Eq. (A.4) and Eq. (A.3) yield

$$\frac{d\Phi}{dv} = \frac{v^3}{M}\frac{dt}{dv} = -\frac{v^3}{M}\frac{E'(v)}{\mathcal{F}(v) + \dot{M}(v)} \,, \tag{A.17}$$

which, after re-expanding in a Taylor series, can be integrated analytically. The final result reads

$$\Phi(v) = \Phi_0 - \frac{1}{32\eta\, v^5} \left\{ 1 + v^2 \left[\frac{3715}{1008} + \frac{55\eta}{12} \right] + v^3 \left[-10\pi + \frac{565\delta\chi_a}{24} + \left(\frac{565}{24} - \frac{95\eta}{6} \right)\chi_s \right] \right.$$

$$+ v^4 \left[\frac{15293365}{1016064} + \frac{27145\eta}{1008} + \frac{3085\eta^2}{144} - \frac{405}{8}\delta\chi_a\chi_s - \left(\frac{405}{16} - \frac{5\eta}{4} \right)\chi_s^2 - \left(\frac{405}{16} - 100\eta \right)\chi_a^2 \right]$$

$$
\begin{aligned}
+ v^5 \ln v &\left[\frac{38645\pi}{672} - \frac{65\pi\eta}{8} - \left(\frac{735505}{2016} - \frac{12265\eta}{36} - \frac{85\eta^2}{2} \right) \chi_s - \left(\frac{735505}{2016} + \frac{65\eta}{4} \right) \delta\chi_a \right. \\
&\left. - \left(\frac{45}{4} - \frac{45\eta}{4} \right) \delta\chi_s^2\chi_a - \left(\frac{45}{4} - \frac{135\eta}{4} \right) \chi_s\chi_a^2 - \left(\frac{15}{4} - \frac{45\eta}{4} \right) \chi_s^3 - \left(\frac{15}{4} - \frac{15\eta}{4} \right) \delta\chi_a^3 \right]
\end{aligned}
$$

$$\text{(A.18)}$$

$$
\begin{aligned}
+ v^6 &\left[\frac{12348611926451}{18776862720} - \frac{1712\gamma}{21} - \frac{160\pi^2}{3} - \left(\frac{15737765635}{12192768} - \frac{2255\pi^2}{48} \right)\eta + \frac{76055\eta^2}{6912} - \frac{127825\eta^3}{5184} \right. \\
&- \frac{1712\ln(4v)}{21} + \left(\frac{730\pi}{3} - \frac{560\pi\eta}{3} \right)\chi_s + \frac{730\pi\delta\chi_a}{3} - \left(\frac{34225}{2688} - \frac{217135\eta}{672} + \frac{1225\eta^2}{12} \right)\chi_s^2 \\
&\left. - \left(\frac{34225}{2688} - \frac{7705\eta}{48} + \frac{1205\eta^2}{3} \right)\chi_a^2 - \left(\frac{34225}{1344} - \frac{10385\eta}{24} \right)\delta\,\chi_s\chi_a \right] \\
+ v^7 &\left[\frac{77096675\pi}{2032128} + \frac{378515\pi\eta}{12096} - \frac{74045\pi\eta^2}{6048} - \left(\frac{20373952415}{24385536} + \frac{150935\eta}{1792} - \frac{578695\eta^2}{3456} \right)\delta\chi_a \right. \\
&- \left(\frac{20373952415}{24385536} - \frac{4348560355\eta}{6096384} - \frac{11189515\eta^2}{24192} + \frac{71705\eta^3}{864} \right)\chi_s - \left(\frac{285\pi}{4} - 5\pi\eta \right)\chi_s^2 \\
&- \left(\frac{285\pi}{4} - 280\pi\eta \right)\chi_a^2 - \frac{285\pi}{2}\delta\,\chi_s\chi_a + \left(\frac{16185}{224} - \frac{74645\eta}{1344} + \frac{905\eta^2}{24} \right)\chi_s^3 \\
&+ \left(\frac{16185}{224} - \frac{437275\eta}{1344} + \frac{85\eta^2}{8} \right)\delta\chi_a^3 + \left(\frac{48555}{224} - \frac{198125\eta}{1344} + \frac{255\eta^2}{8} \right)\delta\,\chi_s^2\chi_a \\
&\left.\left. + \left(\frac{48555}{224} - \frac{1337635\eta}{1344} + \frac{8935\eta^2}{24} \right)\chi_s\chi_a^2 \right] \right\}.
\end{aligned}
$$

Eq. (A.16) and Eq. (A.18) together define $\Phi(t)$ implicitly.

TaylorF2 phasing

Starting from the explicit expressions for time and orbital phase in the TaylorT2 approximant, it is possible to analytically construct the Fourier transform of the GW strain in the framework of the stationary phase approximation (SPA) [150, 151, 152]. Denoting the Fourier transform of the strain by $\tilde{A}_{\ell m}e^{i\psi_{\ell m}}$, the phase in the frequency domain can be approximated by

$$
\psi_{\ell m}(f) = 2\pi f\, t_f - m\Phi(t_f) - \frac{\pi}{4}. \tag{A.19}
$$

Here, f is the Fourier variable and t_f corresponds to the time when the instantaneous GW frequency coincides with f, i.e.,

$$
\frac{d(m\Phi)}{dt}(t_f) = 2\pi f \quad \Rightarrow \quad v(t_f) = \left(\frac{2\pi Mf}{m} \right)^{1/3}. \tag{A.20}
$$

The form of the Taylor series of $\psi_{\ell m}$ obviously depends on the spherical harmonic mode's m. For the sake of brevity, only the expansion for $m = 2$ is given below.

$$
\psi_{\ell 2}(v) = 2t_0 v^3 - 2\Phi_0 - \frac{\pi}{4} + \frac{3}{128\eta\, v^5}\left\{ 1 + v^2\left[\frac{3715}{756} + \frac{55\eta}{9} \right] + v^3\left[-16\pi + \left(\frac{113}{3} - \frac{76\eta}{3} \right)\chi_s + \frac{113\delta\chi_a}{3} \right] \right.
$$

$$
+ v^4 \left[\frac{15293365}{508032} + \frac{27145\eta}{504} + \frac{3085\eta^2}{72} - \left(\frac{405}{8} - \frac{5\eta}{2} \right) \chi_s^2 - \left(\frac{405}{8} - 200\eta \right) \chi_a^2 - \frac{405}{4} \delta \chi_s \chi_a \right]
$$

$$
+ v^5 \left(1 + 3 \ln v \right) \left[\frac{38645\pi}{756} - \frac{65\pi\eta}{9} - \left(\frac{735505}{2268} - \frac{24530\eta}{81} - \frac{340\eta^2}{9} \right) \chi_s - \left(\frac{735505}{2268} + \frac{130\eta}{9} \right) \delta \chi_a \right.
$$

$$
\left. - (10 - 10\eta)\delta \chi_s^2 \chi_a - (10 - 30\eta)\chi_s \chi_a^2 - \left(\frac{10}{3} - 10\eta \right) \chi_s^3 - \left(\frac{10}{3} - \frac{10\eta}{3} \right) \delta \chi_a^3 \right]
$$

$$
+ v^6 \left[\frac{11583231236531}{4694215680} - \frac{6848\gamma}{21} - \frac{640\pi^2}{3} - \left(\frac{15737765635}{3048192} - \frac{2255\pi^2}{12} \right) \eta + \frac{76055\eta^2}{1728} - \frac{127825\eta^3}{1296} \right.
$$

$$
- \frac{6848 \ln(4v)}{21} + \left(\frac{2920\pi}{3} - \frac{2240\pi\eta}{3} \right) \chi_s + \frac{2920\pi}{3}\delta \chi_a - \left(\frac{34225}{672} - \frac{217135\eta}{168} + \frac{1225\eta^2}{3} \right) \chi_s^2
$$

$$
\left. - \left(\frac{34225}{672} - \frac{7705\eta}{12} + \frac{4820\eta^2}{3} \right) \chi_a^2 - \left(\frac{34225}{336} - \frac{10385\eta}{6} \right) \delta \chi_s \chi_a \right]
$$

$$
+ v^7 \left[\frac{77096675\pi}{254016} + \frac{378515\pi\eta}{1512} - \frac{74045\pi\eta^2}{756} - \left(\frac{20373952415}{3048192} + \frac{150935\eta}{224} - \frac{578695\eta^2}{432} \right) \delta \chi_a \right.
$$

$$
- \left(\frac{20373952415}{3048192} - \frac{4348560355\eta}{762048} - \frac{11189515\eta^2}{3024} + \frac{71705\eta^3}{108} \right) \chi_s - (570\pi - 40\pi\eta)\chi_s^2
$$

$$
- (570\pi - 2240\pi\eta)\chi_a^2 - 1140\pi\delta \chi_s \chi_a + \left(\frac{16185}{28} - \frac{74645\eta}{168} + \frac{905\eta^2}{3} \right) \chi_s^3
$$

$$
+ \left(\frac{16185}{28} - \frac{437275\eta}{168} + 85\eta^2 \right) \delta \chi_a^3 + \left(\frac{48555}{28} - \frac{198125\eta}{168} + 255\eta^2 \right) \delta \chi_s^2 \chi_a
$$

$$
\left. \left. + \left(\frac{48555}{28} - \frac{1337635\eta}{168} + \frac{8935\eta^2}{3} \right) \chi_s \chi_a^2 \right] \right\}. \tag{A.21}
$$

According to Eq. (A.20), v should be understood as $v = (M \pi f)^{1/3}$ in the equation above.

A.1.2 Waveform amplitudes

Now, given the orbital phase Φ and the related post-Newtonian expansion parameter v defined in Eq. (A.1), we can obtain the waveform observed at infinity. Currently, the most complete expressions for the nonspinning parts of the waveform are found in [153]. In particular, Eqs. (9.3) and (9.4) of that reference give the decomposition of $h_+ - ih_\times$ into harmonics. Due to space considerations and the danger of transcription errors, we do not reproduce those equations here, but simply refer the reader to that paper. To these, we must add[3] the spin terms given most completely in [154]. There, the spin terms were not explicitly decomposed into harmonics, however, using Eq. (9.2) of [153], it is a simple matter to deduce them. Using Eqs. (F24) and (F25) of [154], and noting the overall sign error in

---

[3]Note that Refs. [153] and [154] share the notation set forth in [155], so that we can simply add the relevant terms. Also note that each of those references uses a different normalization for the variable H compared to the one used here.

Eq. (F25c), we obtain the only nonzero spin contributions to the harmonics:

$$H_{2,2} = -\frac{16}{3}\sqrt{\frac{\pi}{5}}v^5\eta\left[2\delta\chi_a + 2(1-\eta)\chi_s + 3v\eta\left(\chi_a^2 - \chi_s^2\right)\right]e^{-2i\Phi} \,, \tag{A.22}$$

$$H_{2,1} = 4i\sqrt{\frac{\pi}{5}}v^4\eta(\delta\chi_s + \chi_a)e^{-i\Phi} \,, \tag{A.23}$$

$$H_{3,2} = \frac{32}{3}\sqrt{\frac{\pi}{7}}v^5\eta^2\chi_s e^{-2i\Phi} \,. \tag{A.24}$$

In all cases, modes with negative values of m can be obtained from

$$H_{\ell,-m} = (-1)^\ell \bar{H}_{\ell,m} \,. \tag{A.25}$$

The appropriate SPA amplitude in Fourier space can easily be deduced from its time-domain description $A_{\ell m}$ by

$$\tilde{A}_{\ell m} = A_{\ell m}\sqrt{\frac{2\pi}{m\ddot{\Phi}}} = A_{\ell m}\sqrt{\frac{2\pi M}{3mv^2\,\dot{v}}} \,, \tag{A.26}$$

where $\dot{v}$ can be taken for instance from Eq. (A.15) and all arguments should be replaced according to Eq. (A.20).

Bibliography

[1] A. Einstein. Approximative Integration of the Field Equations of Gravitation. *Preuss. Akad. Weiss. Berlin*, page 688, 1916.

[2] R. A. Hulse. Nobel lecture: The discovery of the binary pulsar. 66:699, 1994.

[3] J. H. Taylor. Nobel lecture: Binary pulsars and relativistic gravity. 66:711, 1994.

[4] Joel M. Weisberg and Joseph H. Taylor. Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis. *ASP Conf. Ser.*, 328:25, 2005.

[5] D. Sigg and the LIGO Scientific Collaboration. Status of the LIGO detectors. *Classical and Quantum Gravity*, 25(11):114041–+, June 2008.

[6] F. Acernese et al. Status of Virgo. *Class. Quantum Grav.*, 25(11):114045–+, June 2008.

[7] H. Grote and the LIGO Scientific Collaboration. The status of GEO 600. *Classical and Quantum Gravity*, 25(11):114043–+, June 2008.

[8] K Kuroda and the LCGT Collaboration. Status of lcgt. *Classical and Quantum Gravity*, 27(8):084004, 2010.

[9] K S Thorne. Gravitational radiation. In S. W. Hawking and W. Israel, editors, *Three hundred years of gravitation*, chapter 9, pages 330–458. Cambridge University Press, Cambridge, 1987.

[10] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.

[11] Manuela Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96:111101, 2006.

[12] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation.* Freeman, New York, New York, 1973.

[13] Sean Carroll. *Spacetime and Geometry: An Introduction to General Relativity.* Addison Wesley, New York, 2003.

[14] Joseph Weber and John A. Wheeler. Reality of the cylindrical gravitational waves of einstein and rosen. *Rev. Mod. Phys.*, 29(3):509–515, Jul 1957.

[15] Duncan Brown. *Searching for Gravitational Radiation from Binary Black Hole MACHOs in the Galactic Halo.* PhD thesis, University of Wisconsin–Milwaukee, 2004.

[16] Alessandra Buonanno, Bala R. Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80(8):084043, October 2009.

[17] Alessandra Buonanno and Thibault Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, 1999.

[18] Alessandra Buonanno, Yi Pan, Harald P. Pfeiffer, Mark A. Scheel, Luisa T. Buchman, and Lawrence E. Kidder. Effective-one-body waveforms calibrated to numerical relativity simulations: coalescence of non-spinning, equal- mass black holes. 79:124028, 2009.

[19] Miguel Alcubierre. *Introduction to 3+1 Numerical Relativity.* Oxford Univerity Press, Oxford, 2008.

[20] Sascha Husa. Numerical modeling of black holes as sources of gravitational waves in a nutshell. *Eur. Phys. J. ST*, 152:183–207, 2007.

[21] Ian Hinder. The Current Status of Binary Black Hole Simulations in Numerical Relativity. *Class. Quant. Grav.*, 27:114004, 2010.

[22] Mark A. Scheel et al. High-accuracy waveforms for binary black hole inspiral, merger, and ringdown. 2008.

[23] Michael Boyle, Alessandra Buonanno, Lawrence E. Kidder, Abdul H. Mroué, Yi Pan, Harald P. Pfeiffer, and Mark A. Scheel. High-accuracy numerical simulation of black-hole binaries: Computation of the gravitational-wave energy flux and comparisons with post-Newtonian approximants. 78(12):104020, 2008.

[24] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. 76:124038, 2007.

[25] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific, Singapore, 1994.

[26] R L Ward, R Adhikari, B Abbott, R Abbott, D Barron, R Bork, T Fricke, V Frolov, J Heefner, A Ivanov, O Miyakawa, K McKenzie, B Slagmolen, M Smith, R Taylor, S Vass, S Waldman, and A Weinstein. dc readout experiment at the caltech 40m prototype interferometer. *Classical and Quantum Gravity*, 25(11):114030, 2008.

[27] http://en.wikipedia.org/wiki/File:TEMmn.png.

[28] Roger D. Blandford and Kip S. Thorne. *Applications of Classical Physics*. California Institute of Technology, 2004.

[29] L. S. Finn. Detection, measurement, and gravitational radiation. *Physical Review D*, 46(12):5236–5249, December 1992.

[30] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An Algorithm for Detection of Gravitational Waves from Inspiraling Compact Binaries. 2005.

[31] Bruce Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005.

[32] I. W. Harry, B. Allen, and B. S. Sathyaprakash. Stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D*, 80(10):104014, Nov 2009.

[33] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, 1996.

[34] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, 1999.

[35] Thomas Cokelaer. Gravitational wave from inspiralling compact binaries: hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D*, 76:102004, 2007.

[36] C. A. K. Robinson, B. S. Sathyaprakash, and Anand S. Sengupta. Geometric algorithm for efficient coincident detection of gravitational waves. *Phys. Rev. D*, 78(6):062002, 2008.

[37] C. Capano. *Searching for Gravitational Waves from Compact Binary Coalescence in S6*. PhD thesis, Syracuse University, Syracuse, NY, 2011.

[38] Paul Demorest, Tim Pennucci, Scott Ransom, Mallory Roberts, and Jason Hessels. Shapiro delay measurement of a two solar mass neutron star. *Nature*, 467:1081–1083, 2010.

[39] Yi Pan, Alessandra Buonanno, John G. Baker, Joan Centrella, Bernard J. Kelly, Sean T. McWilliams, Frans Pretorius, and James R. van Meter. Data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. 77(2):024014, 2008.

[40] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Phys. Rev. D*, 77:062002, 2008.

[41] GWINC: Gravitational Wave Interferometer Noise Calculator. v1 default parameters.

[42] Mark A. Scheel, Michael Boyle, Tony Chu, Lawrence E. Kidder, Keith D. Matthews, and Harald P. Pfeiffer. High-accuracy waveforms for binary black-hole inspiral, merger, and ringdown. 79:024003, 2009.

[43] Harald P. Pfeiffer, Duncan A. Brown, Lawrence E. Kidder, Lee Lindblom, Geoffrey Lovelace, and Mark A. Scheel. Reducing orbital eccentricity in binary black hole simulations. 24(12):S59–S81, 2007.

[44] Michael Holst, Lee Lindblom, Robert Owen, Harald P. Pfeiffer, Mark A. Scheel, and Lawrence E. Kidder. Optimal constraint projection for hyperbolic evolution systems. 70:084017, 2004.

[45] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne. A new generalized harmonic evolution system. 23:S447–S462, 2006.

[46] Michael Boyle and Abdul H. Mroué. Extrapolating gravitational-wave data from numerical simulations. arXiv:0905.3177 (gr-qc), 2009.

[47] F. Echeverria. Gravitational-wave measurements of the mass and angular momentum of a black hole. 40:3194, 1989.

[48] Geoffrey Lovelace, Robert Owen, Harald P. Pfeiffer, and Tony Chu. Binary-black-hole initial data with nearly-extremal spins. 78:084017, 2008.

[49] Patrick R. Brady and Jolien D. Creighton. Gravitational wave astronomy. In *Encyclopedia of Physical Science and Technology*, volume 7, pages 33–48. Academic Press, Inc, third edition, 2002.

[50] C. Cutler and K. S. Thorne. An overview of gravitational wave sources. In *Proceedings of the GR16 Conference on General Relativity and Gravitation*. World Scientific, 2002.

[51] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes: The Art of Scientific Computing.* Cambridge University Press, third edition, 2007.

[52] B. Abbott et al. Search for gravitational waves from low-mass binary coalescences in the first calendar year of S5 LIGO data, 2008. In preparation.

[53] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. 9, 2006.

[54] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius. Inspiral, merger and ring-down of equal-mass black-hole binaries. *Phys. Rev. D*, 75:124018, 2007.

[55] Emanuele Berti et al. Inspiral, merger and ringdown of unequal mass black hole binaries: A multipolar analysis. *Phys. Rev. D*, 76:064034, 2007.

[56] P. Ajith et al. A template bank for gravitational waveforms from coalescing binary black holes: I. non-spinning binaries. *Phys. Rev. D*, 77:104017, 2008.

[57] Yi Pan et al. A data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. *Phys. Rev. D*, 77:024014, 2008.

[58] Alessandra Buonanno et al. Toward faithful templates for non-spinning binary black holes using the effective-one-body approach. *Phys. Rev. D*, 76:104049, 2007.

[59] Michael Boyle et al. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev. D*, 76:124038, 2007.

[60] Parameswaran Ajith et al. Phenomenological template family for black-hole coalescence waveforms. *Class. Quantum Grav.*, 24:S689–S699, 2007.

[61] Thibault Damour and Alessandro Nagar. Comparing Effective-One-Body gravitational waveforms to accurate numerical data. *Phys. Rev. D*, 77:024043, 2008.

[62] Thibault Damour, Alessandro Nagar, Ernst Nils Dorband, Denis Pollney, and Luciano Rezzolla. Faithful Effective-One-Body waveforms of equal-mass coalescing black-hole binaries. *Phys. Rev. D*, 77:084017, 2008.

[63] Thibault Damour, Alessandro Nagar, Mark Hannam, Sascha Husa, and Bernd Bruegmann. Accurate Effective-One-Body waveforms of inspiralling and coalescing black-hole binaries. *Phys. Rev. D*, 78:044039, 2008.

[64] Michael Boyle et al. High-accuracy numerical simulation of black-hole binaries: Computation of the gravitational-wave energy flux and comparisons with post-Newtonian approximants. 78:104020, 2008.

[65] Michael Boyle, Duncan A. Brown, and Larne Pekowsky. Comparison of high-accuracy numerical simulations of black-hole binaries with stationary phase post-Newtonian template waveforms for Initial and Advanced LIGO. *Class. Quantum Grav.*, 26:114006, 2009.

[66] Thomas Baumgarte et al. Learning about compact binary merger: The interplay between numerical relativity and gravitational-wave astronomy. *Phys. Rev.*, D77:084009, 2008.

[67] Birjoo Vaishnav, Ian Hinder, Frank Herrmann, and Deirdre Shoemaker. Matched Filtering of Numerical Relativity Templates of Spinning Binary Black Holes. *Phys. Rev. D*, 76:084020, 2007.

[68] Lee Lindblom, Benjamin J. Owen, and Duncan A. Brown. Model Waveform Accuracy Standards for Gravitational Wave Data Analysis. *Phys. Rev. D*, 78:124020, 2008.

[69] Mark Hannam, Sascha Husa, John G. Baker, Michael Boyle, Bernd Bruegmann, Tony Chu, Nils Dorband, Frank Herrmann, Ian Hinder, Bernard J. Kelly, Lawrence E. Kidder, Pablo Laguna, Keith D. Matthews, James R. van Meter, Harald P. Pfeiffer, Denis Pollney, Christian Reisswig, Mark A. Scheel, and Deirdre Shoemaker. The Samurai Project: verifying the consistency of black- hole-binary waveforms for gravitational-wave detection. 79:084025, 2009.

[70] D. A. Brown et al. Data formats for numerical relativity waves. 2007.

[71] Bernd Bruegmann et al. Calibration of Moving Puncture Simulations. *Phys. Rev. D*, 77:024027, 2008.

[72] Sascha Husa, Jose A. Gonzalez, Mark Hannam, Bernd Bruegmann, and Ulrich Sperhake. Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing. *Class. Quant. Grav.*, 25:105006, 2008.

[73] Mark Hannam, Sascha Husa, Ulrich Sperhake, Bernd Bruegmann, and Jose A. Gonzalez. Where post-Newtonian and numerical-relativity waveforms meet. *Phys. Rev. D*, 77:044020, 2008.

[74] Mark Hannam, Sascha Husa, Bernd Brügmann, and Achamveedu Gopakumar. Comparison between numerical-relativity and post-Newtonian waveforms from spinning binaries: the orbital hang-up case. *Phys. Rev. D*, 78:104007, 2008.

[75] Bernd Bruegmann, Wolfgang Tichy, and Nina Jansen. Numerical simulation of orbiting black holes. *Phys. Rev. Lett.*, 92:211101, 2004.

[76] Wolfgang Tichy and Pedro Marronetti. The final mass and spin of black hole mergers. *Phys. Rev. D*, 78:081501, 2008.

[77] Miguel Alcubierre et al. Towards a stable numerical evolution of strongly gravitating systems in general relativity: The conformal treatments. *Phys. Rev. D*, 62:044034, 2000.

[78] Miguel Alcubierre et al. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003.

[79] Michael Koppitz et al. Getting a kick from equal-mass binary black hole mergers. *Phys. Rev. Lett.*, 99:041102, 2007.

[80] Denis Pollney et al. Recoil velocities from equal-mass binary black-hole mergers: a systematic investigation of spin-orbit aligned configurations. *Phys. Rev. D*, 76:124002, 2007.

[81] Luciano Rezzolla et al. Spin Diagrams for Equal-Mass Black-Hole Binaries with Aligned Spins. *Astrophys. J.*, 679:1422–1426, 2008.

[82] Breno Imbiriba et al. Evolving a puncture black hole with fixed mesh refinement. *Phys. Rev. D*, 70:124025, 2004.

[83] James R. van Meter, John G. Baker, Michael Koppitz, and Dae-Il Choi. How to move a black hole without excision: gauge conditions for the numerical evolution of a moving puncture. *Phys. Rev. D*, 73:124011, 2006.

[84] Y. Zlochower, J. G. Baker, Manuela Campanelli, and C. O. Lousto. Accurate black hole evolutions by fourth-order numerical relativity. *Phys. Rev. D*, 72:024021, 2005.

[85] Sergio Dain, Carlos O. Lousto, and Yosef Zlochower. Extra-Large Remnant Recoil Velocities and Spins from Near- Extremal-Bowen-York-Spin Black-Hole Binaries. *Phys. Rev. D*, 78:024039, 2008.

[86] Ulrich Sperhake. Binary black-hole evolutions of excision and puncture data. *Phys. Rev. D*, 76:104015, 2007.

[87] Ulrich Sperhake et al. Eccentric binary black-hole mergers: The transition from inspiral to plunge in general relativity. *Phys. Rev. D*, 78:064069, 2008.

[88] U. Sperhake, V. Cardoso, F. Pretorius, E. Berti, and J. A. Gonzalez. The high-energy collision of two black holes. 101:161101, 2008.

[89] Ian Hinder, Birjoo Vaishnav, Frank Herrmann, Deirdre Shoemaker, and Pablo Laguna. Universality and Final Spin in Eccentric Binary Black Hole Inspirals. *Phys. Rev. D*, 77:081502, 2008.

[90] Frans Pretorius. Numerical Relativity Using a Generalized Harmonic Decomposition. *Class. Quant. Grav.*, 22:425–452, 2005.

[91] Frans Pretorius and Deepak Khurana. Black hole mergers and unstable circular orbits. *Class. Quant. Grav.*, 24:S83–S108, 2007.

[92] Mark A. Scheel et al. Solving Einstein's equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.

[93] Harald P. Pfeiffer et al. Reducing orbital eccentricity in binary black hole simulations. *Class. Quant. Grav.*, 24:S59–S82, 2007.

[94] Zachariah B. Etienne, Joshua A. Faber, Yuk Tung Liu, Stuart L. Shapiro, and Thomas W. Baumgarte. Filling the holes: Evolving excised binary black hole initial data with puncture techniques. *Phys. Rev. D*, 76:101503, 2007.

[95] Pedro Marronetti et al. Binary black holes on a budget: Simulations using workstations. *Class. Quant. Grav.*, 24:S43–S58, 2007.

[96] Pedro Marronetti, Wolfgang Tichy, Bernd Brügmann, José Gonzalez, and Ulrich Sperhake. High-spin binary black hole mergers. *Phys. Rev. D*, 77:064010, 2008.

[97] Sascha Husa, Mark Hannam, Jose A. Gonzalez, Ulrich Sperhake, and Bernd Bruegmann. Reducing eccentricity in black-hole binary evolutions with initial parameters from post-Newtonian inspiral. *Phys. Rev. D*, 77:044037, 2008.

[98] Bernd Brügmann, José A. Gonzalez, Mark Hannam, Sascha Husa, and Ulrich Sperhake. Exploring black hole superkicks. *Phys. Rev. D*, 77:124047, 2008.

[99] Lawrence E. Kidder. Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects. *Phys. Rev.*, D52:821–847, 1995.

[100] Luc Blanchet. Equations of motion of compact binaries at the third post- Newtonian order. 1999.

[101] Gregory B. Cook and Harald P. Pfeiffer. Excision boundary conditions for black hole initial data. *Phys. Rev. D*, 70:104016, 2004.

[102] Wolfgang Tichy and Bernd Bruegmann. Quasi-equilibrium binary black hole sequences for puncture data derived from helical Killing vector conditions. *Phys. Rev. D*, 69:024006, 2004.

[103] Leonardo Gualtieri, Emanuele Berti, Vitor Cardoso, and Ulrich Sperhake. Transformation of the multipolar components of gravitational radiation under rotations and boosts. *Phys. Rev. D*, 78(4):044024, Aug 2008.

[104] Virgo Science Run 2 Sensitivity Curve. http://wwwcascina.virgo.infn.it/DataAnalysis/Calibration/Ser

[105] LIGO Scientific Collaboration. LSC Algorithm Library software packages LAL, LAL-WRAPPER, and LALAPPS.

[106] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 5:3, 2002.

[107] 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. Quant. Grav.*, 13:575–584, 1996.

[108] 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.

[109] 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.*, D65:061501, 2002.

[110] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet. Erratum: Gravitational-wave inspiral of compact binary systems to 7/2 post-newtonian order [phys. rev. d 65, 061501(r) (2002)]. *Phys. Rev. D*, 71(12):129902, Jun 2005.

[111] Luc Blanchet, Thibault Damour, Gilles Esposito-Farese, and Bala R. Iyer. Gravitational radiation from inspiralling compact binaries completed at the third post-Newtonian order. *Phys. Rev. Lett.*, 93:091101, 2004.

[112] B Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. 2009.

[113] 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.

[114] John G. Baker, James R. van Meter, Sean T. McWilliams, Joan Centrella, and Bernard J. Kelly. Consistency of post-Newtonian waveforms with numerical relativity. *Phys. Rev. Lett.*, 99:181101, 2007.

[115]

[116]

[117] F. Herrmann, I. Hinder, D. Shoemaker, and P. Laguna. Unequal mass binary black hole plunges and gravitational recoil. 24:S33–S42, 2007.

[118] D. J. A. McKechan, C. Robinson, and B. S. Sathyaprakash. A tapering window for time-domain templates and simulated signals in the detection of gravitational waves from coalescing compact binaries. *Class. Quantum Grav.*, 27:084020, 2010.

[119] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilagyi. Characteristic extraction in numerical relativity: binary black hole merger waveforms at null infinity. *Class. Quant. Grav.*, 27:075014, 2010.

[120] Denis Pollney and Christian Reisswig. Gravitational memory in binary black hole mergers. 2010.

[121] Denis Pollney, Christian Reisswig, Erik Schnetter, Nils Dorband, and Peter Diener. High accuracy binary black hole simulations with an extended wave zone. arXiv:0910.3803 (gr-qc), 2009.

[122] http://www.black-holes.org/SpEC.html.

[123] Tony Chu, Harald P. Pfeiffer, and Mark A. Scheel. High accuracy simulations of black hole binaries:spins anti-aligned with the orbital angular momentum. *arXiv.org:0909.1313*, 2009.

[124] Sascha Husa, Mark Hannam, José A. González, Ulrich Sperhake, and Bernd Brügmann. Reducing eccentricity in black-hole binary evolutions with initial parameters from post-Newtonian inspiral. 77:044037, 2008.

[125] P. Papadopoulos, E. Seidel, and L. Wild. Adaptive computation of gravitational waves from black hole interactions. 58:084002, 1998.

[126] Manuela Campanelli, Carlos O. Lousto, Hiroyuki Nakano, and Yosef Zlochower. Comparison of Numerical and Post-Newtonian Waveforms for Generic Precessing Black-Hole Binaries. *Phys. Rev.*, D79:084010, 2009.

[127] Ilana MacDonald, Samaya Nissanke, and Harald P Pfeiffer. Suitability of post-newtonian/numerical-relativity hybrid waveforms for gravitational wave detectors. *Classical and Quantum Gravity*, 28(13):134002, 2011.

[128] http://www.w3.org/XML/.

[129] R. Williams. Ldas lightweight data format specification, rev. 01. *Technical document LIGO-T990023-01-E*, 1999.

[130] L. Blackburn. Kleinewelle technical document. *Technical document LIGO-T060221-00-Z*, 2007.

[131] S.K. Chatterji. *The Search for gravitational-wave bursts in data from the second LIGO science run.* PhD thesis, MIT, Boston, MA, 2005.

[132] A. S. Sengupta, J. A. Gupchup, and C. A. K. Robinson. Trigscan: Using parameter correlations to cluster triggers in binary inspiral gravitational wave search pipeline. 2008. In Preparation.

[133] S. Ballmeri at. al. Thermal compensation system description. *Technical document LIGO-T050064-00-R*, 2005.

[134] Shyang Wen. *Improved Seismic Isolation For The Laser Interferometer Gravitational Wave Observatory With Hydraulic External Pre-isolator System.* PhD thesis, LSU, Baton Rouge, LA, 2009.

[135] D. Macleod et. al. Removing the effect of seismic noise from ligo data by targeted veto generation, 2011. LIGO Document number P1100031-v8; available on the web at https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=39453.

[136] Joshua R Smith, Thomas Abbott, Eiichi Hirose, Nicolas Leroy, Duncan Macleod, Jessica McIver, Peter Saulson, and Peter Shawhan. A hierarchical method for vetoing noise transients in gravitational-wave detectors. arXiv:1107.2948 (gr-qc), 2011.

[137] Tomoki Isogai, the Ligo Scientific Collaboration, and the Virgo Collaboration. Used percentage veto for ligo and virgo binary inspiral searches. *Journal of Physics: Conference Series*, 243(1):012005, 2010.

[138] R. Weiss et. al. Report of the adhoc committee on the september 16,2010 event, 2010. LIGO Document number L1000398-v1; available on the web at https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=22752.

[139] S. Whitcomb et. al. Detection committee report, 2011. LIGO Document number T1100114-v2; available on the web at https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=39122.

[140] Duncan A Brown. *Search for gravitational radiation from black hole MACHOs in the Galactic halo.* PhD thesis, University of Wisconsin–Milwaukee, 2004.

[141] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. Higher-order spin effects in the dynamics of compact binaries. II. Radiation field. 74(10):104034, November 2006. Note the two associated errata; the arXiv version has been corrected.

[142] K. G. Arun, Alessandra Buonanno, Guillaume Faye, and Evan Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. 79(10):104023, May 2009. Note that the 1.5 and 2.5pN spin terms in the flux and energy expressions do not account for an erratum from 2010.

[143] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet. Gravitational-wave inspiral of compact binary systems to 7/2 post-Newtonian order. 65(6):061501, Feb 2002.

[144] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet. Erratum: Gravitational-wave inspiral of compact binary systems to 7/2 post-Newtonian order [Phys. Rev. D 65, 061501(R) (2002)]. 71(12):129902, Jun 2005.

[145] Luc Blanchet, Bala R. Iyer, and Benoit Joguet. Erratum: Gravitational waves from inspiraling compact binaries: Energy flux to third post-Newtonian order [Phys. Rev. D 65, 064005 (2002)]. *Phys. Rev. D*, 71(12):129903, Jun 2005.

[146] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. Erratum: Higher-order spin effects in the dynamics of compact binaries. II. Radiation field [Phys. Rev. D 74, 104034 (2006)]. 75(4):049903, February 2007.

[147] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. Erratum: Higher-order spin effects in the dynamics of compact binaries. II. Radiation field [Phys. Rev. D 74, 104034 (2006)]. 81(8):089901, April 2010.

[148] Kashif Alvi. Energy and angular momentum flow into a black hole in a binary. 64(10):104020, Oct 2001.

[149] Hideyuki Tagoshi, Shuhei Mano, and Eiichi Takasugi. Post-Newtonian expansion of gravitational waves from a particle in circular orbits around a rotating black hole :Effects of black hole absorption. *gr-qc/9711072*, November 1997.

[150] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. Comparison of search templates for gravitational waves from binary inspiral. 63(4):044023, 2001.

[151] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. Comparison of search templates for gravitational waves from binary inspiral: 3.5PN update. 66(2):027502, 2002.

[152] K. G. Arun, Bala R. Iyer, B. S. Sathyaprakash, and Pranesh A. Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-Newtonian gravitational wave phasing: The non-spinning case. *Phys. Rev.*, D71:084008, 2005.

[153] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Siddhartha Sinha. The third post-Newtonian gravitational wave polarizations and associated spherical harmonic modes for inspiralling compact binaries in quasi-circular orbits. 25(16):165003, 2008.

[154] Clifford M. Will and Alan 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.

[155] Luc Blanchet, Bala R Iyer, Clifford M Will, and Alan G Wiseman. Gravitational wave-forms from inspiralling compact binaries to second-post-Newtonian order. 13(4):575, 1996.