School of Physics and Astronomy

Improving the sensitivity of searches for gravitational waves from compact binary coalescences

Duncan Michael Macleod

Declaration of Authorship

- DECLARATION:
  This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

  Signed: .... (candidate) Date: ....

- STATEMENT 1:
  This thesis is being submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (PhD).

  Signed: .... (candidate) Date: ....

- STATEMENT 2:
  This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references. The views expressed are my own.

  Signed: .... (candidate) Date: ....

- STATEMENT 3:
  I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

  Signed: .... (candidate) Date: ....

- STATEMENT 4: PREVIOUSLY APPROVED BAR ON ACCESS
  I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loans **after expiry of a bar on access previously approved by the Academic Standards & Quality Committee.**

  Signed: .... (candidate) Date: ....

Summary of thesis

The detection of gravitational waves from the coalescence of two compact objects has been brought to within touching distance by the construction and operation of a global network of laser-interferometer detectors. However, the amplitude of the radiation from these events is so low that direct detection will require the combined innovations of advanced interferometry and detector characterisation, along with powerful methods of extracting weak, but modelled, signals from the background detector noise. This work focuses on enhancing the probability of such detection through improved identification of noise artefacts in the instrumental data, and improved signal processing and extraction.

We begin with a recap of the theory of gravitational waves as derived from Einstein's theory of gravity, and the mechanisms that allow propagation of this radiation away from a source. We also catalogue a number of promising astrophysical progenitors, with a focus on compact binary coalescences.

We detail the interactions between gravitational waves and an observer, and describe the layout of the large-scale laser interferometers that have been built to enable direct detection. A description of the operation of these detectors during the last science run is given, focusing on their stability and sensitivity, isolating a number of key instrumental noise mechanisms and how they affected astrophysical searches over the data. Additionally, we illustrate a new method to improve the identification of seismic noise bursts, allowing their removal from search data, improving search sensitivity.

The LIGO and Virgo gravitational-wave detectors operated as a network during the last joint science run. A summary is given of the analysis pipeline used to search for gravitational waves signals from compact binary coalescences using a coincidence-based method, including details of the results of that analysis. Details are also given of the pipeline used to search for gravitational waves associated with short, hard gamma-ray bursts, in which a new coherent method was tuned to search over the reduced parameter space constrained by the electromagnetic counterpart. Finally, we present a new pipeline adapting the coherent method to the blind, all-sky, all-time search, allowing for a more sensitive analysis, as demonstrated by direct comparison.

Acknowledgements

I am very thankful to a huge number of people, without any of whom I would have had a much harder time trying to find my way towards a completed Ph.D. First, and greatest, thanks go to my parents, for setting me on my way, and allowing me to become their 'saddest child' by studying physics, instead of biology or music or French. They have always provided encourage and support (not infrequently financial) to help me on my way, and so this work is dedicated to them.

A huge amount of any thanks must go to Steve (aka. 'boss') for accepting me as a post-graduate student in the first place, then quickly shipping me off to the middle of nowhere to learn something that I didn't sign up for. Steve has been open to me wandering off and working on my own thing, whilst also providing the guidance I needed to turn that work into something useful for me, and then into something useful for other people. He has always been a great teacher, carefully explaining any situation, and asking little in return (beer excepted); I have been lucky to have Steve as a mentor, and a friend, and hope we can continue to work together in the future.

Thank you also to Mark Hannam, Sathyaprakash, and Patrick Sutton for guidance and direction whenever Steve was off helping somebody else - it has been great to have such a helpful team, always (looking) happy to answer any questions I had. I have been very lucky to have their company, and that of the other members of the Cardiff Gravitational Physics group: Thomas Adams, James Clark, Thomas Dent, Mark Edwards, Ian Harry, Paul Hopkins, Gareth Jones, Ioannis Kamaretsos, Sebastian Khan, Erin Macdonald, Dave McKechan, Chris Messenger, Laura Nuttall, Frank Ohme, Valeriu Predoi, Michal Pürrer, Craig Robinson, Patricia Schmidt, John Vietch. In particular, special commendations are awarded to Ian – for putting up with my annoying questions, and letting me steal a little of the credit for his work – and to Laura for putting with my annoying habits at both Hanford and Livingston, and not complaining when I didn't stop working at a reasonable hour.

Additionally, I would like to thank the various members of the LIGO-Virgo Compact Binary Coalescences group for helping me along the way, and pointing me in the right direction for explanations and assistance. Thanks also to everyone from the Detector Characterisation group for their help, most notably Nelson Christensen, Gaby Gonzalez, Andy Lundgren, Jess McIver, Chris Pankow, Jamie Rollins, and Josh Smith. Gaby deserves a special mention, not only for being a great help in getting me into detector characterisation, but

Contents

III   Searches for signals from compact binary coalescences

List of Figures

List of Tables

Co-authored papers

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

- Chapter 4 is in preparation for publication as 'Characterization of the LIGO detectors during their sixth science run' [26]. This is a collaborative work of the LIGO Scientific Collaboration and the Virgo collaboration, for which D. M. Macleod is corresponding author.

- Chapter 5 was published as 'Reducing the effect of seismic noise in LIGO searches by targeted veto generation' [27]. This is a collaborative work on which D. M. Macleod is lead author.

- Chapter 6 presents the methods and results of the search for low-mass binary coalescences in data from the LIGO and Virgo detectors during the joint science run S6/VSR2/3 [18]. The author was a member of the team performing this search, and coordinated this analysis during the final three months of the run ('S6D'). Additionally, the author coordinated data quality studies throughout the full run, and was heavily involved in the analysis of candidate event GW100916.

- Chapter 7 presents the methods and results of the search for GWs from compact binary coalescences (CBCs) associated with GRB detections [19]. The author was heavily involved in the development of the search pipeline, especially for methods used in the analysis of *Fermi*-detected bursts, and was lead analysts on a number of target analyses, including that of GRB 090802A.

To Mum, Dad, and Granny.
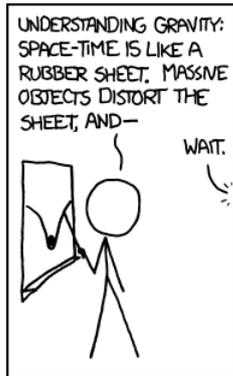
Chapter 0

Introduction

The general theory of relativity is the most complete and robust theory of gravity currently available. Formulated by Einstein during the beginning of the twentieth century [1], it combines the special theory of relativity [2] with Newton's laws of gravitation to produce a geometrical description of gravity as a property of spacetime in the presence of mass.

General relativity (GR) was revolutionary not only in its description and conclusion of known problems, for example the advance of the orbital perihelion of Mercury, but also in its prediction of unknown phenomena. For example, it predicts that electromagnetic radiation interacts with a gravitational field, and that light is attracted towards massive bodies. This effect is the foundation of gravitational time dilation and gravitational redshift, giving us an understanding of gravitational lensing, for example, and is is the reason the Global Positioning System (GPS) works today.

Another important prediction of GR is the existence of gravitational waves (GWs). This form of radiation is emitted from accelerating massive objects, and travels through spacetime at the speed of light, analogous to electromagnetic radiation from accelerating charged objects. As a result, any accelerating or asymmetric rotating system will emit radiation that could be detectable from the Earth, giving observers a new window onto the universe. The Earth-Sun binary system is predicted to emit GWs with an strain amplitude of 1 part in $\sim 10^{-26}$ at a distance of 1 light-year, meaning that detecting such waves would require measuring displacements of less than the width of a proton. Such small amplitudes lead many early theorists to believe GWs were undetectable.

However, in the 1960s, Joe Weber built a series of aluminium bar detectors, capable of detecting strains of one part in 10^{16} [3, 4], and in 1968 claimed the

first tentative detection of GWs [5]. These claims were eventually rejected by the larger scientific community, and it was not until the 1980s that a new generation of GW detectors took shape.

The GEO project was born out of collaboration between the Max Planck institutes for Astrophysics (Munich) and Quantum Optics (Garching) and the University of Glasgow with plans for the first kilometre-scale laser interferometer [6]. While the scale of the original proposal was reduced due to funding cuts, an instrument with 600 metre arms [7] was completed in 2001 [8]. Alongside this, a Japanese collaboration designed and built the TAMA300 detector in Japan, with 300 metre arms [9], achieving their first science data in 1999 [10].

Since then, the first true kilometre-scale instruments were constructed as part of the Laser Interferometer Gravitational-Wave Observatory (LIGO) project [11], with three detectors completed in 2002, and the Virgo detector [12], completed in 2003.

The global network of detectors, including LIGO [13], GEO600 [14], and Virgo [15], recently completed a second joint data-taking run[1], with sensitivity to GWs from events occurring more than 20 megaparsecs from the Earth [17]. These data have been analysed for both short- and long-duration GW signals. Short-duration signals include those from compact binary coalescences (CBCs) [18, 19, 20], and unmodelled bursts [19, 21]. The long-duration sources include pulsars [22, 23, 24] and a stochastic gravitational-wave background (SGWB) [25].

This work focuses on the sensitivity of searches for gravitational waves from CBCs, with particular attention paid to data quality and the impact of noise transients on the search, and improving the ability to distinguish true signals from these noises through a fully-coherent matched filter and associated signal-consistency tests.

Part I concentrates on the theory of gravitational waves. Chapter 1 details the prediction of gravitational waves as part of GR, and the propagation of these waves from a source. Chapter 2 describes source of GWs, including the Hulse-Taylor pulsar, with a focus developing towards compact binary coalescences.

In part II we shift focus onto detection of GWs with ground-based laser-interferometers. Chapter 3 introduces these instruments as a means of detecting GWs through their interaction with a laser beam, while chapters 4 and 5

---

[1]The TAMA300 detector has been decommissioned to make way for the Kamioka Gravitational Wave Detector (KAGRA) (formerly the Large-scale Cryogenic Gravitational Wave Telescope (LCGT)) [16].

describe the efforts during the most recent LIGO science run to characterise the noise output of these instruments and mitigate the impact of this noise on searches of the data [26, 27].

These astrophysical searches are the core of part III. The methods and results of the coincidence-based search for GWs from CBCs are described in chapter 6, while the methods and results of the coherent search for CBCs signals associated with gamma-ray bursts (GRBs) are described in chapter 7. Finally, Chapter 8 presents the first application of a fully-coherent all-sky, all-time search for CBC signals.

Part I

Theory and sources

Chapter 1

Gravitational waves: theory and sources

This chapter gives a brief introduction to the theory of gravitational waves as predicted by general relativity. A number of texts give a more complete study of these ideas [28, 29, 30], while in this work we focus on those concepts crucial to the later discussions in parts II and III, of GW detectors and astrophysical searches respectively.

Section 1.1 begins with the linearised approximation to GR, before section 1.2 introduces a particular gauge in which GWs take on their simplest form. In sections 1.3 and 1.4 we describe the propagation of this radiation away from a source, towards an observer.

1.1 The linearised general theory of relativity

We begin with the action for motion in a gravitational field [28],

$$S = S_E + S_M, \tag{1.1}$$

where,

$$S_M = \int d^4x \sqrt{-g} \mathcal{L}_M, \tag{1.2}$$

is the matter action, and

$$S_E = \frac{c^3}{16\pi G} \int d^4x \sqrt{-g} R, \tag{1.3}$$

is the gravitational, or Einstein, action. $\mathcal{L}_M$ is the Lagrangian associated with the matter of the system, R is the Ricci scalar, and g is the metric – containing

all information about the curvature of the spacetime. Variation of this action with respect to the metric gives,

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu}, \tag{1.4}$$

where $T_{\mu\nu}$ is the energy-momentum tensor of matter, defined specifically from variation of the matter action (equation (1.2)) under a change of the metric.

Equation (1.4) defines the Einstein Field Equations, the mathematical statement of general relativity. The left-hand side is a geometric object, and describes how the curvature of space will affect matter, while the right hand side is the energy-momentum tensor, describing how mass alters the curvature of space.

The full theory of gravity is non-linear, which is mathematically too complex to allow for simple discussion of its implications. However, in many applications of GR, including the beginnings of GW physics, it is enough to approximate the full metric, $g_{\mu\nu}$, as the Minkowski metric, $\eta_{\mu\nu}$, plus small perturbations. The effect of gravity is then described as a small perturbation $h_{\mu\nu}$ on top of flat space. We write

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \tag{1.5}$$

where $|h_{\mu\nu}| \ll 1$. This is the weak-field approximation, which we can use to expand the Einstein Field Equations to linear order in $h_{\mu\nu}$.

If we start with the Christoffel symbols [29],

$$\Gamma^\rho{}_{\mu\nu} = \frac{1}{2}g^{\rho\sigma}\left(\frac{\partial}{\partial x^\mu}g_{\sigma\nu} + \frac{\partial}{\partial x^\nu}g_{\sigma\mu} - \frac{\partial}{\partial x^\sigma}g_{\mu\nu}\right), \tag{1.6}$$

we can define the Reimann curvature tensor,

$$R^\mu{}_{\nu\rho\sigma} = \frac{\partial}{\partial x^\rho}\Gamma^\mu{}_{\nu\sigma} - \frac{\partial}{\partial x^\sigma}\Gamma^\mu{}_{\nu\rho} + \Gamma^\mu{}_{\alpha\rho}\Gamma^\alpha{}_{\nu\sigma} - \Gamma^\mu{}_{\alpha\sigma}\Gamma^\alpha{}_{\nu\rho}. \tag{1.7}$$

If we then expand the Christoffel symbols to first-order in h we have

$$\Gamma^\rho{}_{\mu\nu} = \frac{1}{2}\left(\frac{\partial}{\partial x^\mu}h^\rho{}_\nu + \frac{\partial}{\partial x^\nu}h^\rho{}_\mu - \frac{\partial}{\partial x_\rho}h_{\mu\nu}\right), \tag{1.8}$$

and from this we can calculate the first order Riemann tensor, and contract it

to form the Ricci tensor

$$R_{\mu\nu} = R^{\rho}{}_{\mu\rho\nu} = -\frac{1}{2}\left(\Box h_{\mu\nu} - \frac{\partial^2}{\partial x^\mu \partial x^\rho}h^\rho{}_\nu - \frac{\partial^2}{\partial x^\nu \partial x^\rho}h^\rho{}_\mu + \frac{\partial^2}{\partial x^\mu \partial x^\nu}h^\rho{}_\rho\right). \quad (1.9)$$

If we insert these simplifications into equation (1.4), the linearised Einstein Field Equations read

$$\Box h_{\mu\nu} + \frac{\partial^2}{\partial x^\mu \partial x^\nu}h^\rho{}_\rho - \frac{\partial^2}{\partial x^\mu \partial x^\rho}h^\rho{}_\nu - \frac{\partial^2}{\partial x^\nu \partial x^\rho}h^\rho{}_\mu -$$
$$\eta_{\mu\nu}\left(\Box h^\rho{}_\rho - \frac{\partial^2}{\partial x^\rho \partial x^\sigma}h^{\rho\sigma}\right) = -\frac{16\pi G}{c^4}T_{\mu\nu}. \quad (1.10)$$

The increased complexity of equation (1.10) compared to equation (1.4) may seem like a step backwards, but only temporarily. Following the convention of [29], we can use

$$h = \eta^{\mu\nu}h_{\mu\nu}, \quad (1.11)$$

to define the trace reverse of $h_{\mu\nu}$

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h. \quad (1.12)$$

This new form holds the same information as before, but allows the equation (1.10) to be written in a simpler form,

$$\Box\bar{h}_{\mu\nu} + \eta_{\mu\nu}\frac{\partial^2}{\partial x_\rho \partial x_\sigma}\bar{h}_{\rho\sigma} - \frac{\partial^2}{\partial x_\rho \partial x^\nu}\bar{h}_{\mu\rho} - \frac{\partial^2}{\partial x_\rho \partial x^\mu}\bar{h}_{\nu\rho} = -\frac{16\pi G}{c^4}T_{\mu\nu}. \quad (1.13)$$

The first term is the flat-space Laplacian operator, with the other terms only maintaining the gauge invariance of the system. Because of this invariance the linearised equations cannot give unique solutions, since given any solution we can always perform a suitable coordinate transformation to generate others. But, we can make use this gauge freedom to simplify the form of equation (1.13), allowing us to more easily expose the important physical content of the theory.

The most general coordinate transformation that preserves the weak-field approximation (equation (1.5)), is

$$x^\mu \rightarrow x'^\mu = x^\mu + \varepsilon^\mu(x), \quad (1.14)$$

where the $\partial\varepsilon^\mu/\partial x^\nu$ are at most the same order of magnitude as $h_{\mu\nu}$. Under a

general coordinate transform the metric transforms as

$$g_{\mu\nu}(x) \rightarrow g'_{\mu\nu}(x') = \frac{\partial x^\rho}{\partial x'^\mu} \frac{\partial x^\sigma}{\partial x'^\nu} g_{\rho\sigma}(x), \tag{1.15}$$

and so, applying equation (1.14), the transformation of $h_{\mu\nu}$ to linear order is

$$h_{\mu\nu}(x) \rightarrow h'_{\mu\nu}(x') = h_{\mu\nu}(x) - \left(\frac{\partial \varepsilon_\nu}{\partial x^\mu} + \frac{\partial \varepsilon_\mu}{\partial x^\nu} \right). \tag{1.16}$$

This highlights the gauge freedom of the system, where we can choose a particular gauge in which to work in order to simplify our calculations. We choose the *harmonic gauge* [28] in which,

$$g^{\mu\nu} \Gamma^\rho{}_{\mu\nu} = 0, \tag{1.17}$$

or, using equation (1.8), we have

$$\frac{\partial}{\partial x^\mu} h^\mu{}_\nu = \frac{1}{2} \frac{\partial}{\partial x^\nu} h^\mu{}_\mu, \text{ or} \tag{1.18a}$$

$$\frac{\partial}{\partial x_\nu} \bar{h}_{\mu\nu} = 0. \tag{1.18b}$$

Applying the harmonic gauge condition (equation (1.18)) to the linearised field equations (1.13) gives a wave equation:

$$\Box \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}. \tag{1.19}$$

This simple (looking) expression predicts the existence of gravitational waves.

1.2 The transverse-traceless gauge

The wave equation (1.19) is the result from which all analytical study of gravitational radiation stems. However, we can impose another approximation to further simplify the form of this radiation.

Away from the source, we can approximate $T_{\mu\nu} = 0$, resulting in a homogeneous equation for gravitational waves,

$$\Box \bar{h}_{\mu\nu} = 0, \tag{1.20}$$

whose solution, satisfying the harmonic condition (equation (1.18)), is a linear

superposition of plane waves of the form

$$\bar{h}_{\mu\nu}(x) = \mathcal{E}_{\mu\nu} \exp\left(ik_\rho x^\rho\right). \tag{1.21}$$

where k_ρ is the null wave-vector, and the $\mathcal{E}_{\mu\nu}$ are constants. In the harmonic gauge we have

$$k_\mu \mathcal{E}^\mu{}_\nu = 0, \tag{1.22}$$

imposing four conditions on the system, reducing the number of degrees of freedom from ten to six. We can reduce this to just two by manipulating further coordinate freedoms still present.

Recall the general co-ordinate transformation,

$$x^\mu \to x'^\mu = x^\mu + \varepsilon^\mu(x). \tag{1.14}$$

The gravitational perturbation is transformed as

$$\bar{h}'_{\mu\nu}(x) = \mathcal{E}'_{\mu\nu} \exp(ik_\rho x^\rho), \tag{1.23}$$

where

$$\mathcal{E}'_{\mu\nu} = \mathcal{E}_{\mu\nu} + k_\mu \varepsilon_\nu + k_\nu \varepsilon_\mu. \tag{1.24}$$

We are free to fix the gauge, and construct ε such that the only non-zero elements of $\mathcal{E}$ are

$$\mathcal{E}_{11} = -\mathcal{E}_{22}, \tag{1.25a}$$

$$\mathcal{E}_{12} = \mathcal{E}_{21}. \tag{1.25b}$$

We now choose to relabel $\mathcal{E}_{11} = h_+$ and $\mathcal{E}_{12} = h_\times$, and so the solution of the homogeneous gravitational wave equation (1.20) is

$$h_{\mu\nu}^{\mathrm{TT}}(x) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} \exp\left(ik_\rho x^\rho\right). \tag{1.26}$$

We have chosen a gauge in which the waves are transverse – with no wave amplitude in the direction of propagation – and traceless, and the restriction to two degrees of freedom is clear. h_+ and $h_\times$ are known as the 'plus' and 'cross' polarisations of the gravitational wave, as depicted in figure 1.1.
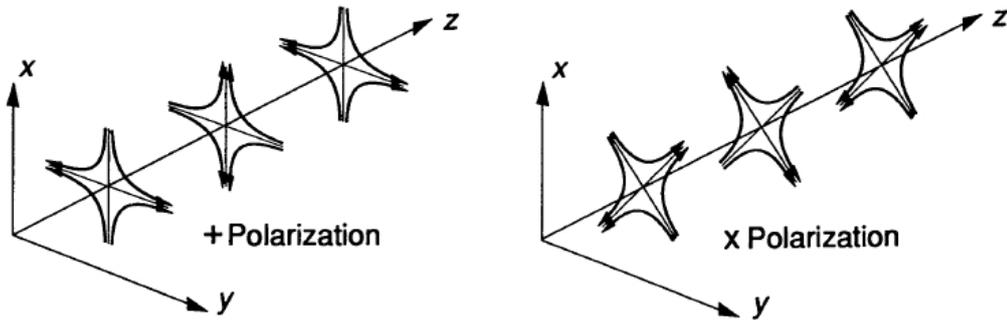
Figure 1.1: Polarisation amplitudes of a gravitational wave [11].

1.3   The multipole expansion

The transverse, traceless (TT) gauge solution (equation (1.26)) is only valid for the homogeneous field equations (1.20), that is, in the vacuum. Unsurprisingly, gravitational waves are generated by a source with non-zero energy and momentum, so we must consider the solution of the full wave equation (1.19) [30],

$$\bar{h}_{\mu\nu}(t, \mathbf{r}) = \frac{4G}{c^4} \int d^3x' \frac{T_{\mu\nu}\left(t - \frac{|\mathbf{r}-\mathbf{x}'|}{c}, \mathbf{x}'\right)}{|\mathbf{r} - \mathbf{x}'|}, \tag{1.27}$$

where $\mathbf{r} = D\hat{\mathbf{n}}$ is the wave position from the centre of the source, and $\mathbf{x}'$ a point within the source.

We still wish to consider detection far from the source, where $D \gg x'$ (the radius of the source), and so we can expand,

$$|\mathbf{r} - \mathbf{x}'| = D - \mathbf{x}' \cdot \hat{\mathbf{n}} + O\left(\frac{x'^2}{D}\right), \tag{1.28}$$

and consider the internal motions of the source small enough to be neglected[1]. The retarded potential (1.27) can then be expanded in powers of the small parameter $(\mathbf{x}' \cdot \hat{\mathbf{n}})/c$.

We begin this expansion by defining the moments of the energy-momentum

---

[1]We assume $v \ll c$, where v is the typical velocity of internal motion.

tensor [30],

$$S^{ij}(t) = \int d^3x \, T^{ij}(t, \mathbf{x}), \tag{1.29a}$$

$$S^{ij,k}(t) = \int d^3x \, T^{ij}(t, \mathbf{x})x^k, \tag{1.29b}$$

$$S^{ij,kl}(t) = \int d^3x \, T^{ij}(t, \mathbf{x})x^k x^l, \tag{1.29c}$$

$$\vdots$$

where $i, j, \ldots$ are the spatial indices. We can then define the *multipole expansion*[2],

$$h_{ij}(t, \mathbf{r}) = \frac{4G}{Dc^4} \left[S^{kl} + \frac{1}{c} n_m \dot{S}^{kl,m} + \cdots \right]. \tag{1.30}$$

This expansion is also known as the *post-Newtonian expansion*.

In order to better show the physical meaning of all of the terms in the multipole expansion, it is beneficial to eliminate the moments of T^{ij} in favour of the moments of the mass-energy density, T^{00}, and the linear momentum, T^{0i}.

We define the mass moments as

$$M = \frac{1}{c^2} \int d^3x \, T^{00}(t, \mathbf{x}), \tag{1.31a}$$

$$M^i = \frac{1}{c^2} \int d^3x \, T^{00}(t, \mathbf{x})x^i, \tag{1.31b}$$

$$M^{ij} = \frac{1}{c^2} \int d^3x \, T^{00}(t, \mathbf{x})x^i x^j, \tag{1.31c}$$

$$\vdots$$

and the momentum moments as

$$P^i = \frac{1}{c} \int d^3x \, T^{0i}(t, \mathbf{x}), \tag{1.32a}$$

$$P^{ij} = \frac{1}{c} \int d^3x \, T^{0i}(t, \mathbf{x})x^j, \tag{1.32b}$$

$$\vdots$$

In this form, the mass monopole, M, represents the total mass in the system. If we combine equations (1.18) and (1.19) we see that the stress-

---

[2]Each factor x^i is $O(d)$ and each time derivative introduces a term $O(\omega_s)$. Thus each additional term in equation (1.30) is a correction to S^{kl} of higher order in (v/c).

energy tensor is divergence free

$$\frac{d}{dx_\mu}T^{\mu\nu} = 0 \tag{1.33}$$

Using equation (1.33) we can relate the mass and momentum moments,

$$\dot{M}^i = P^i, \tag{1.34a}$$

$$\dot{M}^{ij} = P^{ij} + P^{ji}, \tag{1.34b}$$

$$\dot{P}^{ij} = S^{ij}. \tag{1.34c}$$

and so, given that S_{ij} is symmetric,

$$S^{ij} = \frac{1}{2}\ddot{M}^{ij}. \tag{1.35}$$

Using equation (1.35) we can define a relatively simple expression for the leading order term[3],

$$h_{ij}^{\text{TT}}(t, \mathbf{r}) = \frac{2G}{Dc^4}\ddot{M}_{kl}\left(t - \frac{r}{c}\right), \tag{1.36}$$

From equation (1.36) we see the lowest order term in the multipole expansion of linearised GR is the quadrupole, and that neither monopole or dipole radiation is present. You may be forgiven for thinking that this is a consequence of using linearised gravity, which lead us originally to the wave equation (1.19) and on to the expansion in equation (1.30), but this is fact a general property of the full non-linear theory of gravitational radiation.

1.4 Propagation of gravitational waves

The above expression for the quadrupole term of gravitational radiation is expressed in a coordinate system configured for the source. When considering observation of GWs far from the source, it is necessary to transform into a new coordinate frame better suited to radiation in the direction of an observer.

We can generalise to any coordinate system whilst preserving the relative simplicity of the TT gauge formalism, by introducing the projection, $\Lambda_{ij,kl}$, into the TT gauge[4] [30]. This projection, allows us to directly express the two

---

[3]Recall in the TT frame, $h_{0\mu} = 0$

[4]From this point forward, all expressions for the GW polarisations (h_+ and $h_\times$) will be implicitly evaluated in the TT gauge.
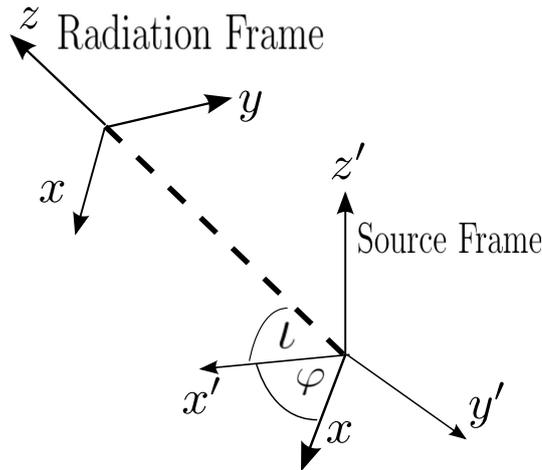
Figure 1.2: The relationship between the source frame, and the observer's line-of-sight (radiation frame) [31].

polarisation amplitudes in any coordinate frame as

$$h_+ = \frac{G}{Dc^4} \left(\ddot{M}_{11} - \ddot{M}_{22} \right), \tag{1.37a}$$

$$h_\times = \frac{2G}{Dc^4} \ddot{M}_{12}. \tag{1.37b}$$

We can then transform into an arbitrary 'radiation' frame, in which the radiation propagates towards an observer. If the radiation coordinates are (x, y, z), the polarisation amplitudes in the new frame will be

$$h_+ = \frac{G}{Dc^4} \left(\ddot{M}'_{11} - \ddot{M}'_{22} \right), \tag{1.38a}$$

$$h_\times = \frac{2G}{Dc^4} \ddot{M}'_{12}. \tag{1.38b}$$

where (x', y', z') are the source frame coordinates. We can then express the mass quadrupole moment in this radiation frame, M_{ij}, in terms of the source moment, M'_{ij}, through a rotation,

$$M_{ij} = R_i^k M'_{kl} R_j^l. \tag{1.39}$$

As shown in figure 1.2 this is done by introducing two extrinsic parameters: the twist angle, φ, relating the x-direction in each frame and the x-direction in the radiation frame; and the inclination angle, ι, defining the $z' \to z$ tilt between these two frames. Then, in the radiation frame we have

$$h_+ = \frac{G}{Dc^4} \Big[\ddot{M}'_{11}(\cos^2\varphi - \sin^2\varphi \sin^2\iota)$$

$$+ \ddot{M}'_{22}(\sin^2\varphi - \cos^2\varphi \sin^2\iota)$$

$$- \ddot{M}'_{33}\cos^2\iota$$

$$- \ddot{M}'_{12}\sin 2\varphi(1 + \sin^2\iota)$$

$$+ \ddot{M}'_{13}\sin\varphi \sin 2\iota$$

$$+ \ddot{M}'_{23}\cos\varphi \sin 2\iota \Big], \tag{1.40a}$$

$$h_\times = \frac{G}{Dc^4} \Big[\left(\ddot{M}'_{11} - \ddot{M}'_{22} \right) \sin 2\varphi \sin\iota$$

$$+ 2\ddot{M}'_{12}\cos 2\varphi \sin\iota$$

$$- 2\ddot{M}'_{13}\cos\varphi \cos\iota$$

$$+ 2\ddot{M}'_{23}\sin\varphi \cos\iota \Big]. \tag{1.40b}$$

These equations give us the full form of the gravitational radiation waveform in an arbitrary frame in terms of the mass quadrupole moment components, M'_{ij} in the source frame. In the next chapter we will examine a number of astrophysical sources, including the coalescence of two compact stars – either neutron stars or black holes, using these equations to derive the emission waveform for this event.

1.5   Energy of gravitational radiation

The energy carried by gravitational waves can be seen by evaluating the energy-momentum tensor of the GW itself, allowing evaluation of the flux, and so the emission power. Maggiore [30] evaluates the flux, and shows it can be integrated over the emission sphere to give the power in terms of the polarisation amplitudes,

$$\frac{dE}{dt} = \frac{c^3 r^2}{32\pi G} \int d\Omega \left\langle \dot{h}_+^2 + \dot{h}_\times^2 \right\rangle \tag{1.41}$$

where the angle brackets indicate an average over time. It is then shown that, using the quadrupole expansion to first order, the power carried by gravitational waves can be given in terms of the mass quadrupole moment as

$$\frac{dE}{dt} = \frac{G}{c^5} \left\langle \dddot{M}_{ij}\dddot{M}_{ij} - \frac{1}{3}(\dddot{M}_{kk})^2 \right\rangle. \tag{1.42}$$

Chapter 2

Sources of gravitational waves

The opening chapter outlined the theory of gravitational wave emission, and the form these waves taken when propagating to an observer. We now turn to the astrophysical and cosmological sources of this emission, first recalling one particular source giving the strongest indirect evidence that the theory is correct, before turning to four categories of sources.

For each of a stochastic gravitational-wave background (SGWB), isolated spinning neutron stars, and GW burst progenitors, a short description of the waveform is given along with a summary of the latest search results. For compact binary coalescences (CBCs), descriptions of modelling each section of the event are described, with an emphasis on the analytical inspiral models, with the methods and results for searching for inspiral signals in GW detector data left for later chapters.

2.1 Evidence for astrophysical gravitational wave emission: the Hulse-Taylor binary

In 1974, Hulse and Taylor discovered a new pulsar during a survey with the Arecibo radio telescope [32], for which anomalies in the pulse timing identified it to be part of a binary system [33]. This was the first observation of a pulsar as part of a binary system.

Further analysis of the pulsar's emissions detected an advanced in the time of orbital periastron, pointing to a decay in the binary orbit due to energy emission. This decay was found to match exactly the prediction of GR assuming emission of gravitational waves from a binary neutron star (BNS) [34].

The pulses from B1913+16, as it is known, have been observed for over
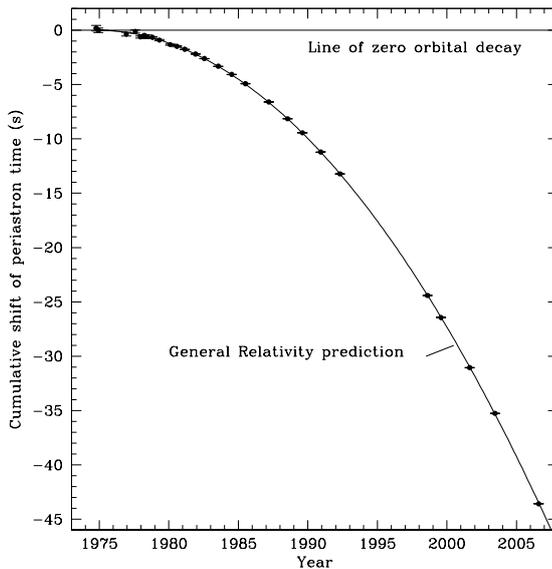
Figure 2.1: Cumulative shift of the periastron of B1913+16 over thirty years of observations [35]. The points are observational data, and the solid line gives the prediction from GR assuming emission of GWs.

thirty years, with the advance in the time of periastron compared against theoretical predictions from GR [35]. This comparison, shown in figure 2.1, is the strongest indication that binary stars do indeed emit gravitational waves as predicted.

While this kind of indirect detection allows a number of tests of GR, and study of the progenitor object, in the remainder of this chapter we discuss a number of sources of directly detectable gravitational waves. This kind of detection will allow observations in a new regime of astrophysics, complementing and expanding on knowledge gained from electromagnetic (EM) instruments.

2.2   A stochastic gravitational wave background

The Hulse-Taylor pulsar is continuously emitting GWs, however the amplitude of this emission is such that the signal will be detectable from ground-based instruments for only a very short time - as the binary components inspiral, the amplitude will increase continually, and only reach a detectable level immediately before the binary merger. We consider now a potentially detectable long-duration source, from which a GW signal should be continuously observable.

A background of GWs is expected to exist from a mix of low-amplitude

astrophysical signals [36, 37, 38, 39] and a cosmological echo of the early universe [40]. The SGWB is usually described in terms of a dimensionless energy density [25],

$$\Omega_{\mathrm{GW}}(f) = \frac{f}{\rho_c} \frac{d\rho_{\mathrm{GW}}}{df}. \tag{2.1}$$

where $d\rho_{\mathrm{GW}}$ is the energy density of GWs in the frequency range f to $f + df$, and ρ_c is critical energy density of the universe. This signal is expected to be a Gaussian random variable whose amplitude is well below the detector noise, and so a network of at least two detectors must be used to identify the signal as a correlated 'noise', distinguishing it from the uncorrelated detector noise background.

The most recent analysis of GW detector data taken during LIGO Science Run 5 (S5) used a signal model of the form

$$\Omega_{\mathrm{GW}}(f) = \Omega_\alpha \left(\frac{f}{f_R}\right)^\alpha, \tag{2.2}$$

where α is the spectral index, and f_R a reference frequency such that $\Omega_\alpha = \Omega_{\mathrm{GW}}(f_R)$. The results presented in [25] used a reference frequency of $100\,\mathrm{Hz}$ with a flat spectrum ($\alpha = 0$), and calculated an upper limit of $\Omega_{\mathrm{GW}} < 6.9 \times 10^{-6}$ with 95% confidence. This upper limit improves on the restricted limit on a cosmological GW background from cosmic microwave background data [41], and improves the previous limit on an astrophysical background from data taken during LIGO Science Run 4 (S4).

2.3   Isolated spinning neutron stars

Where a SGWB manifests as a correlated, random signal in multiple detectors, an ideal source of stable, continuous GW emission is a non-axisymmetric, rapidly-spinning neutron star. Many spinning neutron stars have been observed as pulsars, and, assuming some deviation from axisymmetry, can be expected to emit GWs to some extent.

The dominant mode of gravitational wave emission is expected at twice the rotation frequency, with small drifts in frequency due to either spin-down (due to energy emission) or spin-up (due to matter accretion). Current searches for

GWs from known pulsars assume a model of the form [42]

$$h_+(t) = h_0 \left(\frac{1 + \cos^2 \iota}{2} \right) \cos \Phi(t), \qquad (2.3a)$$

$$h_\times(t) = h_0 \cos \iota \sin \Phi(t), \qquad (2.3b)$$

where the amplitude h_0 is given by

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz} \epsilon f^2}{D}. \qquad (2.4)$$

I_{zz} is the rotational moment of inertia and D the distance to the source, while f is the GW emission frequency and ϵ the ellipticity of the neutron star.

Using data from S5, a blind search for these signals was performed, making no detections, and placing worst-case [1]upper limits on GW strain amplitude of 10^{-24} around $150\,\mathrm{Hz}$ and 3.8×10^{-24} at higher frequencies [43].

Additionally, searches have been targeted at known pulsars, using knowledge of the sky location, pulse frequency, and spin-down (or spin-up) rate to allow a more sensitive analysis. The data from S5 were searched for 116 known pulsars, and with no detections, upper limits were set for each source on the rate of spin-down due to GW emission [23] (Table 1).

Other searches have been specifically targeted at the Crab [22] and Vela pulsars [24]. In the results of each of these searches, astrophysically interesting upper-limits have been placed on the amount of pulsar spin-down that can be attributed to GW emission. For the Crab pulsar, the upper limit on strain amplitude restricts the GW emission restricts the GW emission to $\lesssim 6\%$ of the total available power, while for the Vela pulsar that restriction was $\lesssim 35\%$.

2.4   Unmodelled burst sources

A large number of astrophysical events are predicted to emit transient bursts gravitational waves, but whose emission is, as yet, unmodelled. Core-collapse supernovae [44], pulsar glitches [45], magnetars [46] including soft gamma repeaters (SGRs) [47], long GRBs [19] and a host of other highly-energetic events are expected to result in detectable emission, and have been included in the latest searches for GW bursts (see [21] for a full list of possible sources).

These search algorithms make no assumption of the waveform, and detect

---

[1]Here 'worst-case' refers to the minimum of the detector response over all signal polarisations and sky locations
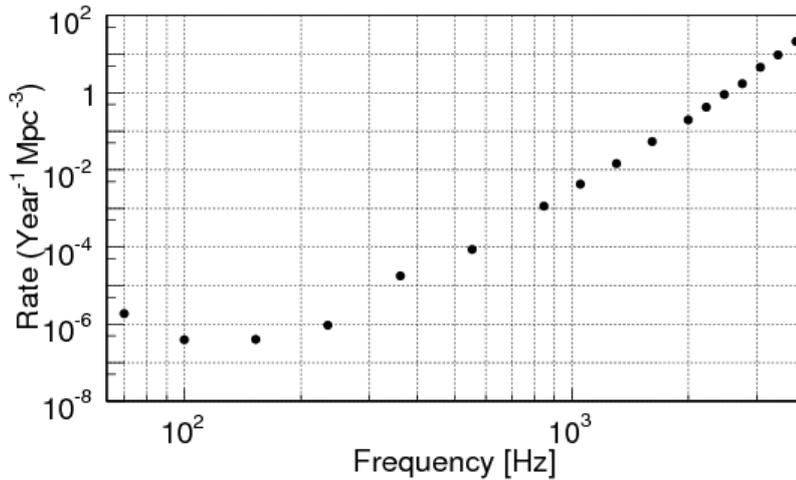
Figure 2.2: Rate density limit for GW bursts as determined by a search of the LIGO-Virgo data during the second joint science run assuming an energy of $1M_\odot c^2$ [21].

events by identifying time-frequency pixels with excess energy correlated between detectors [48]. Separation of signals from noise is highly dependent on detector network consistency checks, and aggressive data quality cuts. The latest analyses produced no detection of GWs, and so placed an upper limit on the rate of such bursts [21]. Figure 2.2 shows the 90% confidence upper limit on rate density as a function of frequency, assuming a linearly-polarised standard-candle source with $E_{GW} = M_\odot c^2$.

These energies, and so rates, are optimistic when compared against typical models for long GRBs, with energies of $\sim 10^{-2} M_\odot c^2$ [21], and core-collapse supernovae with $E_{GW} \sim 10^{-8} M_\odot c^2$ [44].

2.5 Compact binary stars

Perhaps the most well-known source of GWs is the coalescence of two compact objects, either neutron stars or black holes. Two compact stars in gravitationally-bound rotation form a simple accelerating quadrupole moment, and so will emit GWs. This emission causes the orbit to decay to the point where the bodies will coalesce into a new star that settles into a stable state or further decays into a black hole.

The Hulse-Taylor binary pulsar from section 2.1 is one such example, with its orbital decay rate in accurate agreement with the predictions of the theory. Additional to this system, a number of other pulsars have been discovered in

binary systems including white dwarf stars, neutron stars, and higher mass companions [49, 50], informing estimates of rates of binary mergers [51, 52].

2.5.1 Modelling compact binary inspiral

The initial phase of the event, known as the inspiral, carries the majority of the signal energy, as seen by ground-based gravitational wave detectors, and so we now focus on it relative to the later, shorter phases.

The inspiral can be well modelled using the post-Newtonian (PN) expansion outlined in chapter 1, which relies on some basic assumptions that are maintained through the remainder of this work. First, we model all component stars as point masses; this ignores any tidal effects in mergers involving a neutron star. Second, we assume all orbits are circular; for nearly equal mass binaries this assumption is motivated by the prediction that any orbital ellipticity would have been radiated away before the signal entered the detection band of ground-based detectors [53, 54]. Additionally, we assume the binary has no spin, removing 6 dimensions from the signal parameter space (one three-dimensional spin vector for each binary component), allowing for greatly simplified models.

With these assumptions in place, we can model the binary as two point-like companions, with masses m_1 and m_2, and positions $\mathbf{x}_1$ and $\mathbf{x}_2$ respectively, moving on a trajectory described by Newtonian mechanics. We can then choose our (x, y, z) coordinate system to be the centre-of-mass frame where the z-direction is aligned with the rotational axis, and the orbital motion is given by

$$x(t) = r(t) \cos \left(\int_0^t \omega(t') \, dt' \right), \tag{2.5a}$$

$$y(t) = r(t) \sin \left(\int_0^t \omega(t') \, dt' \right). \tag{2.5b}$$

For this system we know that the quadrupole moment is [28]

$$M^{ij}(t) = \mu x^i(t) x^j(t), \tag{2.6}$$

where μ is the reduced mass of the binary,

$$\mu = \frac{m_1 m_2}{m_1 + m_2}, \tag{2.7}$$

and so the acceleration of this moment is given by

$$\ddot{M}_{11}(t) = -\ddot{M}_{22}(t) = 2\mu r^2(t) f^2(t) \cos\left(2 \int_0^t \omega(t')\, dt'\right), \qquad (2.8a)$$

$$\ddot{M}_{12}(t) = \ddot{M}_{21}(t) = -2\mu r^2(t) f^2(t) \sin\left(2 \int_0^t \omega(t')\, dt'\right), \qquad (2.8b)$$

where all other terms are zero.

We can simplify these expressions by recalling Kepler's Law for circular orbits,

$$\omega^2 = \frac{GM}{r^3}, \qquad (2.9)$$

where $M = m_1 + m_2$ is the total mass of the system, allowing the amplitudes of the above expression to be written as

$$\mu r^2 \omega^2 = \mathcal{M}_c^{5/3} G^{2/3} \omega^{2/3}, \qquad (2.10)$$

where we have now introduced the *chirp mass*

$$\mathcal{M}_c = \mu^{3/5} M^{2/5}. \qquad (2.11)$$

We can then insert the simplified accelerations into the waveform polarisation amplitudes from equations (1.40) to see

$$h_+(t) = \frac{4}{D}\left(\frac{G\mathcal{M}_c}{c^2}\right)^{5/3}\left(\frac{\omega(t)}{c}\right)^{2/3}\left(\frac{1 + \cos^2(\iota)}{2}\right)\cos(\Phi + 2\psi), \qquad (2.12a)$$

$$h_\times(t) = \frac{4}{D}\left(\frac{G\mathcal{M}_c}{c^2}\right)^{5/3}\left(\frac{\omega(t)}{c}\right)^{2/3}\cos\iota\,\sin(\Phi + 2\psi). \qquad (2.12b)$$

These expressions describe the emitted radiation in terms of the chirp mass of the source binary, $\mathcal{M}_c$; the orbital frequency of the radiation, ω; the inclination angle of the orbital axis relative to the observer's line of sight, ι; the phase of the radiation, $\Phi = 2\int_0^t \omega(t')\,dt'$; the polarisation angle, ψ; and the distance to the source, D.

Waveforms, like those above, were used as analytical models in the latest search for binary inspiral signals in LIGO and Virgo data. A search for coalescences of low-mass binaries (neutron star-neutron-star and neutron star-black hole binaries up to total mass of $25\,\mathrm{M_\odot}$) has has been completed [18], with no detections made, placing a rate density upper limit of $1.3 \times 10^{-4}\,\mathrm{yr}^{-1}\,\mathrm{Mpc}^{-3}$ for BNS coalescences – more than one thousand times higher than optimistic

estimates [52].

Further details on the methods and results of this analysis will be given in chapters 6 and 7.

2.5.2 Compact binary merger and ringdown

The above derivation has described an intuitive, analytical method for modelling the inspiral of compact stars. This model is dependent on the validity of the multipole expansion, which breaks down when the orbital motions reach relativistic velocities close to merger. As a result numerical methods must be used to solve the Einstein equations in full, non-linear GR during the merger phase. For binary black holes, numerical methods have successfully modelled the merger for near-equal-mass systems (see [55] for a recent review), however when a neutron star is involved, the matter in the system must be modelled also, further complicating any numerical modelling; for more details see [56, 57] and references therein.

The ringdown phase begins once the binary components have coalesced into a new, deformed black hole, which settles into a steady state by emitting radiation in a number of quasi-normal modes. Black hole perturbation theory can be used to model these modes [58], using a damped sinusoid of the form [59, 60]

$$h(t) = \mathcal{A}\frac{GM}{c^2 D}e^{-\pi f_0 t/Q}\cos(2\pi f_0 t). \tag{2.13}$$

where f_0 is the frequency of the fundamental resonant mode, and Q the damping factor.

Figure 2.3 shows a model of the full inspiral-merger-ringdown waveform emitted during the coalescence of two stellar mass black holes, at a distance of $100\,\mathrm{Mpc}$. This model [61] matches a PN inspiral approximation to the output of a number of numerical simulations to produce an analytical model for the full coalescence event waveform.

The data from GW detectors operating during S4 have been searched for signals from black hole ring-downs by matched-filtering a bank of waveforms against the data. This search modelled black holes with mass between $10\,\mathrm{M}_\odot$ and $500\,\mathrm{M}_\odot$, with no detections made, and a 90% confidence upper limit on the rate density of $3.2 \times 10^{-5}\,\mathrm{yr}^{-1}\,\mathrm{Mpc}^{-3}$ was set.

Similarly, a search for full GW inspiral-merger-ringdown (IMR) signals from binary black hole (BBH) coalescences with total mass in the range $25-100\,\mathrm{M}_\odot$ has been completed using more recent data from the LIGO-Virgo network
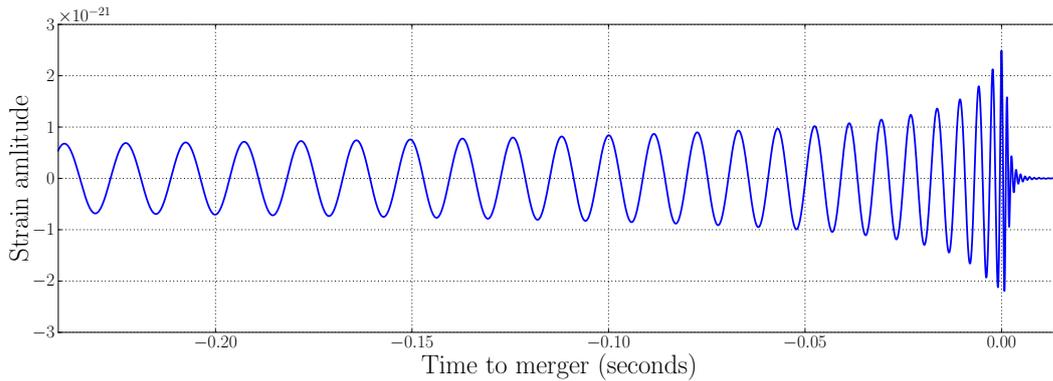
Figure 2.3: The gravitational waveform emitted during the coalescence of two equal-mass, non-spinning black holes [61]. The inspiral, merger, and ringdown components, as described in the text, can all be seen.

during the joint run S6/VSR2/3. This search also made no detections and placed rate density upper limits as a function of component mass [20] (figure 5).

2.5.3   Binary event rate

While CBCs are theoretically a plausible source of GWs, with a firm model for the waveform, building a detector to search for them would be ill-advised if these events were prohibitively rare. A number of methods have been used to estimate rates of binary mergers detectable by both the first- and second-generation ground-based GW detectors [52] placing estimates on BNS merger events between $0.01 - 50\,\mathrm{Myr}^{-1}\,\mathrm{Mpc}^{-3}$ [62, 63]. The range of these estimates highlights the inherent uncertainty in predicting event rates, with theoretical models of binary evolution compared with extrapolations based on the observed BNS populations providing a range of several orders of magnitude between pessimistic and maximally optimistic rates. Readers are referred to [52] for a review of these methods and their implications for event rates.

Given the sensitivities of the first-generation instruments during the joint S5/VSR1 data-taking run, these astrophysical rates can be converted to detection rates of between $10^{-4} - 0.6$ per year – one can see that detection with the first-generation would have been fortuitous. However, when the sensitivity of the forthcoming second-generation GW detectors is folded into these estimates, a detection rate between $0.4 - 400$ (realistically ~ 40) BNS mergers per

year is expected[2]. Even with pessimistic rate estimates, regular detection of GW signals from BNS coalescences is expected in the Advanced Detector Era (ADE).

---

[2]These estimates assume a network of detectors operating at their designed sensitivity. This is expected ~ 2018, rather than during their first data-taking runs [64].

Part II

Detectors and performance characterisation

Chapter 3

Gravitational wave detectors

In part I we saw the prediction of gravitational waves, and the form they take when propagating from a source. We also saw a number of likely astrophysical source objects, briefly touched on their signal models, and skimmed over the latest search results.

We now move on to the more practical issue of detecting gravitational waves, with our focus on ground-based laser interferometers. A number of these instruments were constructed and formed the so-called first-generation detector network, including TAMA300 [65], GEO600 [14], LIGO [13], and Virgo [15].

This chapter outlines the theoretical approach to detecting gravitational waves using laser interferometers, whereby astrophysical signals introduce a differential phase offset between a laser beam in each of the instrumental arms. We include descriptions of the optical layout and operation of such instruments, and detail their sensitivity as limited by three fundamental noise sources. Finally, a brief description of the eagerly-awaited second-generation instruments is given.

3.1 The effect of gravitational waves on an observer

Throughout the derivations of gravitational-wave propagation from chapter 1, the transverse, traceless gauge was used to simplify the mathematics. This choice allowed identification of the 'plus' and 'cross' polarisations of the waveforms, and allowed us to de-couple the signal into these two parts.

When considering the detection of these signals, we use the fact that the

TT gauge was constructed such that velocities in this gauge, $dx^i/d\tau$, are zero at all times [30], meaning the coordinate separation of two points will not change. However, we can show that a passing gravitational wave will change the time taken for light to travel between these points.

Consider, for simplicity, a linearly polarised gravitational wave ($h_\times = 0$) propagating in the z-direction. In the TT frame we have

$$
\begin{aligned}
ds^2 = g_{\mu\nu}dx^\mu dx^\nu &= \left(\eta_{\mu\nu} + h_{\mu\nu}\right)dx^\mu dx^\nu \\
&= -c^2 dt^2 + dx^2 \left(1 + h_+^{TT}\right) + dy^2 \left(1 - h_+^{TT}\right). \quad (3.1)
\end{aligned}
$$

Now, if we consider an observer sending a beam of photons travelling in the x-direction, we have to first order in h,

$$
dt = \frac{dx}{c}\left(1 + \frac{1}{2}h_+(t)\right). \quad (3.2)
$$

and so the time taken to travel between two points, from $\mathbf{x}_0$ at time t_0 to $\mathbf{x}_1$ at time t_1, can be obtained by integrating equation (3.2):

$$
\Delta t = t_1 - t_0 = \frac{L}{c} + \frac{1}{2}\int_{t_0}^{t_1} dt' h(t') \quad (3.3)
$$

where L is the fixed coordinate distance between the two points. So, we see the GW will introduce a change in the travel time.

3.2 Laser interferometers as gravitational-wave detectors

Given the above interaction between light and gravitational waves, laser interferometers are considered an ideal instrument with which to measure a light travel time change in two orthogonal directions, and directly detect a strain induced by GWs.

We can adapt the simple example of the previous section to that of a Michelson-Morley laser interferometer [66], in which a single laser beam is separated at a 50%-transmissive 'beam-splitter' mirror, with one beam propagating in one direction (known as an 'arm') and another propagating in a perpendicular direction. Each beam is fully reflected at the end of the arm and returns to recombine at the beam-splitter, when some part of each beam will be reflected or transmitted back towards the input, and some towards
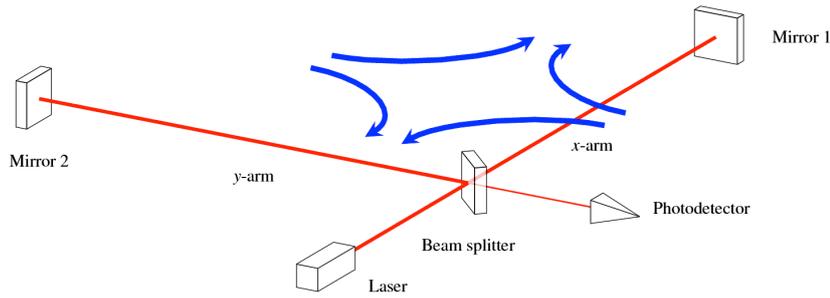
Figure 3.1: Layout of an idealised Michelson interferometer [67]. Overlaid in blue is the interaction of a linearly polarised gravitational wave travelling in the vertical direction, as described in the text.

a photo-detector. Predictably, we choose to orient our coordinates so that one arm points in the x-direction, and the other in the y; figure 3.1 shows a schematic of the Michelson interferometer, including a representation of a linearly-polarised gravitational wave travelling in the z-direction.

If each of the mirrors can be shielded from external forces – by inserting in a vacuum, and suspending from a pendulum, for example – they are good approximations to the ideal test masses considered in the toy model of section 3.1[1]. The change in light travel time induced by GWs manifests as a phase offset in the beams travelling in each arm, changing the interference pattern between them as they recombine at the beam-splitter and propagate towards the output.

For the beam travelling in the x-direction, we can adapt equation (3.3) to measure the time taken for a round-trip down the arm (time t_0 to t_1) and back (time t_1 to t_2),

$$\begin{aligned}
\Delta t &= [t_2 - t_1] + [t_1 - t_0] \\
&= \left[L_x/c + \frac{1}{2} \int_{t_1}^{t_2} dt' \, h_+(t') \right] + \left[L_x/c + \frac{1}{2} \int_{t_0}^{t_1} dt' \, h_+(t') \right] \\
&= \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_2} dt' \, h_+(t')
\end{aligned} \tag{3.4}$$

where L_x is now the length of the interferometer arm.

We can consider a GW model of the form

$$h_+(t) = h_0 \cos\left(\omega_{\mathrm{gw}} t\right) \tag{3.5}$$

---

[1]It must be noted, however, that some external forces cannot be shielded, gravity gradient noise for example [68], and that their impact is truly limiting.

and so equation (3.4) gives (as always, to first order in h_0)

$$\Delta t = \frac{2L_x}{c} + \frac{\sin(w_{\mathrm{gw}}L_x/c)}{w_{\mathrm{gw}}} \, h_0 \cos\left[w_{\mathrm{gw}}(t - L_x/c)\right]. \tag{3.6}$$

The resultant phase shift in the x-arm due to gravitational waves is then

$$\Delta\phi_x(t) = \frac{\omega_L}{\omega_{\mathrm{gw}}} \sin(w_{\mathrm{gw}}L_x/c) \, h_0 \cos\left[w_{\mathrm{gw}}(t - L_x/c)\right], \tag{3.7}$$

where ω_L is the frequency of the laser beam.

If the two interferometer arms are of equal length then the phase shift in the y-arm is the negative of that in the x[2], and so the differential phase shift induced by GWs in an interferometer is

$$\Delta\phi_{\mathrm{Mich}} = \Delta\phi_x - \Delta\phi_y = 2\Delta\phi_x. \tag{3.8}$$

The electric field at the output can be expressed as [30]

$$E_{\mathrm{out}}(t) = E_x(t) + E_y(t) = -iE_0 e^{-i\omega_L(t - 2L/c)} \sin[\phi_0 + \Delta\phi_x(t)] \tag{3.9}$$

where ϕ_0 is the phase of the light at the output in the absence of any gravitational waves. This phase can be chosen as part of the experimental configuration, as described in section 3.3.2.

Detecting modulations of the electric field, its amplitude, and its phase, can allow the detection of gravitational waves. However, as we will see in the remainder of this chapter, all interferometer components – analogue and digital, environmental and instrumental – will constrain the sensitivity of the instrument by adding their own amplitude and phase modulations as noise on the output photodetector. Understanding and attenuation of these noises is key to maximising the scientific output of the search for gravitational waves, as discussed in more detail in chapters 4 and 5.

3.3 First-generation laser interferometers

Based on the above methods, a number of kilometre-scale laser interferometer detectors have been constructed and operated. The first large-scale laser interferometer detector, TAMA300 [65], took its earliest science data in 1999,

---

[2]The phase change upon reflection from the front of the beam-splitter, for the y-arm beam, relative to the reflection from the back of the beam-splitter after the round trip for the x-arm beam, introduces an overall sign change [69].

with 300 metre arms, while the GEO600 detector [14] was completed in 2001, with a 600 metre arm length. The Laser Interferometer Gravitational-Wave Observatory project (LIGO) [13] finished construction in 2002 of three interferometers at two sites, including two with 4-kilometre arm lengths and a third (co-located with one of the others) with 2 km arms. Each of the TAMA300 and LIGO instruments used additional mirrors at the beginning of the arms to create a Fabry-Perot optical cavity, causing the laser light to circulate in the arms for a much longer time before recombination at the beam-splitter [11]. The Virgo detector [15] joined the network in 2007, with 3 km Fabry-Perot arm cavities.

All of the instruments above had a similar overall design, but each included a number of unique components that allowed increased sensitivity to various sources. The smaller scale of GEO600 in particular has driven several innovative techniques to be first tested on that instrument, in order to improve sensitivity relative to the larger detectors, including triple pendulum suspensions [70]; signal recycling [71]; DC readout with an output mode cleaner [72, 73]; and quantum squeezing [74]. The following description is for the LIGO detectors as operated during LIGO Science Run 6 (S6).

3.3.1 Optical layout

The first-generation LIGO instruments were augmented versions of the Michelson-Morley instrument, as shown in figure 3.2. In this design, the primary light source was a 35 Watt laser system, employing feedback loops to stabilise both the intensity and frequency outputs [76], injecting at most 20 W of light at 1064 nm into the main interferometer[3]. To shield components from acoustic noise and to stop the beam scattering off gas particles, most components of each of the large scale interferometers are housed inside an ultra-high vacuum system [77]. An electro-optic modulator added radio frequency (RF) sidebands to the input beam, used for control of the instrument at various points, before it was incident on a triangular mode cleaner. This optical cavity filtered higher-order optical modes from the light to to ensure maximum transmission of a fundamental Gaussian beam profile before the beam entered the full interferometer.

The simple Michelson layout was enhanced with a power recycling mirror, forming a secondary optical cavity with the input test masses; the laser

---

[3]With a higher input laser power, maintaining resonance in the arm cavities is more difficult, and so the full 35 W input was not used [26].
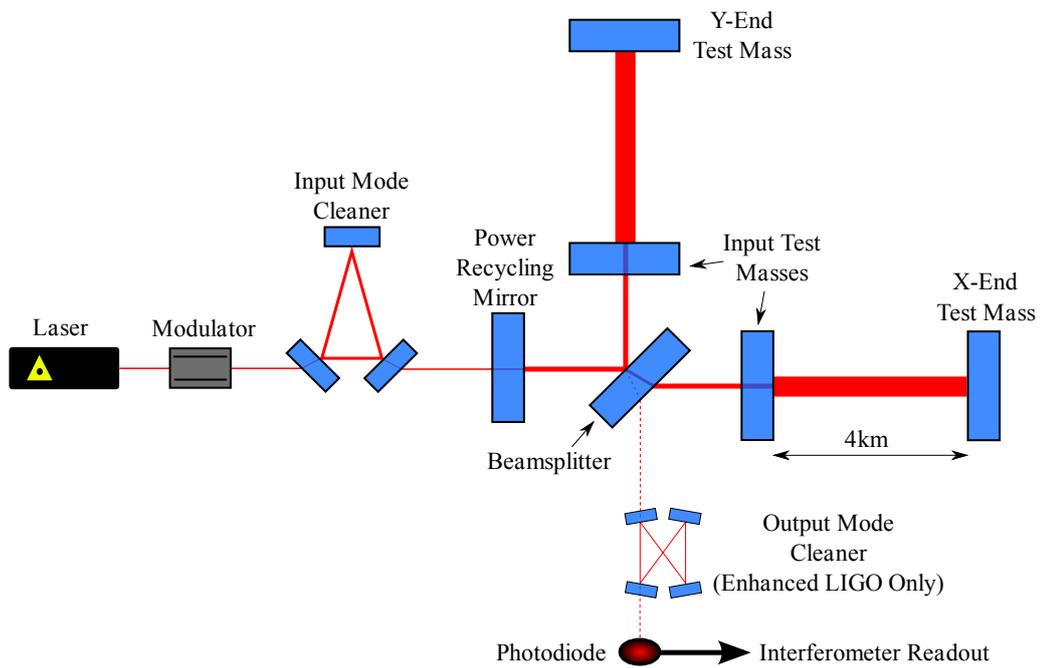
Figure 3.2: Optical layout of the LIGO interferometers during S6 [75]. The original Michelson-Morley design has been expanded with optical mode cleaners (both input and output), a power recycling mirror, and input test masses.

power reflected from the beam-splitter towards the input was recycled into the arms, effectively increasing the input laser power by a factor of ~ 40 [78]. As mentioned above, the input test masses (see figure 3.2) were used to form Fabry-Perot cavities in the arms, allowing power to build up in resonance and acting to store light for a longer period. If we recall equation (3.8), we see the maximal phase shift due to GWs occurs when $\omega_{\mathrm{gw}} L/c = \pi/2$ or

$$L = \frac{\pi}{2} \frac{c}{w_{\mathrm{gw}}} = \frac{\lambda_{\mathrm{gw}}}{4}. \tag{3.10}$$

So, for a gravitational wave at $500\,\mathrm{Hz}$ the optimal arm length is $150\,\mathrm{kilometres}$. The Fabry-Perot cavity allows this to be achieved in a practical manner by effectively increasing the arm length of the instrument.

Light eventually reflected towards the output system was further cleaned by an output mode cleaner, which also acted to reduce the power incident on the detection photodiodes which recorded the signal [79].

3.3.2 Detector operation

As described above, the instruments detect GWs by measuring the amplitude and phase of the light incident on the output photodiode. In order to maintain resonance in the various optical cavities, and allow signal power to build up in the instrument, the detector is operated such that each cavity length is an integer multiple of the laser wavelength. The two arms are stabilised through sensing of their common and differential motion, whilst the power recycling cavity length is stabilised separately.

Additionally, during the first five LIGO Science Runs (S1–S5), the relative distances of each input test mass from the beam-splitter were controlled to ensure the light on the output photodiode was on a dark fringe. In this mode, the main laser light from each arm interferes deconstructively at the beam-splitter, whereby GW-induced phase-modulation sidebands will interfere constructively (recall the 180° phase shift described in footnote 2), and register as power on the output photodiode.

For S6, LIGO adopted a new detection system, in which the Michelson phase is stabilised away from a dark fringe, and the main laser beam is used to carry GW power to the output [72, 80]. A more complete description of this *DC readout* system, and other upgrades installed for S6, dubbed Enhanced LIGO, will be given in chapter 4.

3.4 Detector sensitivity

The sensitivity of all laser interferometer detectors is measured in terms of the GW strain amplitude spectral density (ASD) as a function of frequency; figure 3.3 shows the sensitivities of the LIGO detectors during the last science run (S6). The sensitivity is built-up through combinations of a vast array of
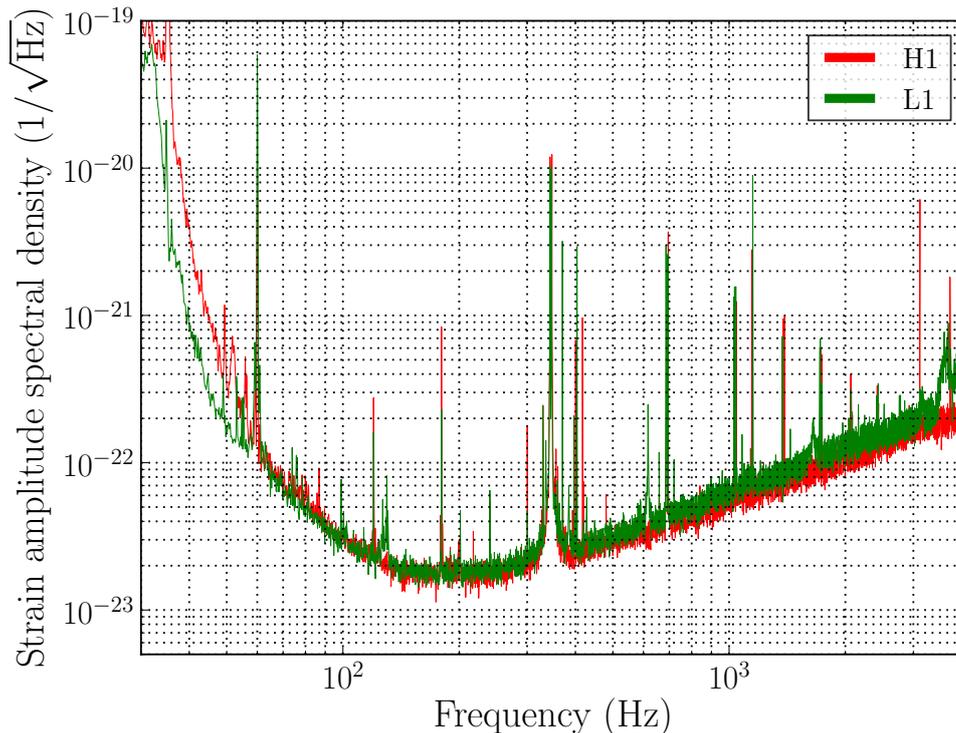
Figure 3.3: Representative strain noise amplitude spectra of the first-generation LIGO interferometers [81]. Here the prefixes 'H1' and 'L1' refer the 4-km instruments at each of the two LIGO sites. The overall sensitivity is limited in the low- ($\lesssim 40\,\text{Hz}$), mid- ($\sim 40 - 200\,\text{Hz}$), and high-frequency bands ($\gtrsim 200\,\text{Hz}$) by seismic, thermal, and shot noise respectively.

system components, with the three dominant contributions described below.

3.4.1 Seismic noise

At the lowest frequencies, up to $\sim 40\,\text{Hz}$, the dominant noise source is seismic motion. The instrument is built to measure microscopic displacements of macroscopic mirrors, meaning any extra motion of components caused by an auxiliary source will overwhelm a GW signal. Figure 3.4 shows the displacement noise caused by seismic motion at each LIGO site, outlining the level of
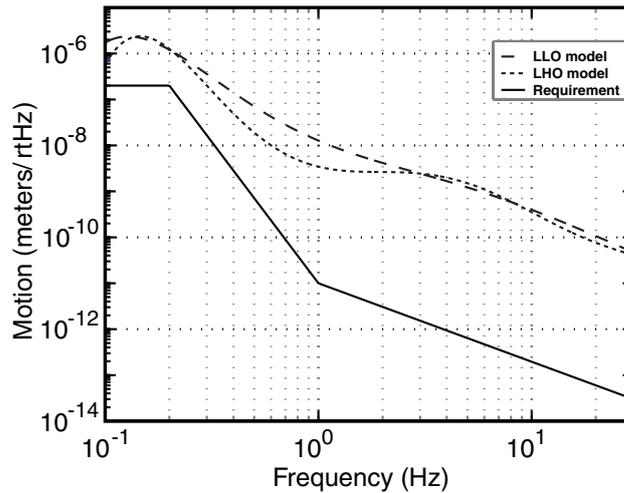
Figure 3.4: Models of the seismic displacement noise spectra at each of the LIGO sites [82], and the requirement for aLIGO design sensitivity. Whilst the seismic noise level drops with frequency, the sensitivity goal requires an increasing level of damping, already at a factor of ~ 500 above a few Hertz. Further damping via passive suspensions is required to get from this system requirement to the full interferometer sensitivity.

attenuation required for the Advanced LIGO (aLIGO) seismic isolation system, especially at higher frequencies.

This seismic noise is mitigated by a combination of active and passive isolation systems. In the early science runs, a simple pre-isolation system was used whereby the signal from a seismometer attached to the vacuum tank containing each arm cavity mirror was fed forward to actuators on the sides of the tank to mitigate motion. At LIGO Livingston Observatory (LLO) for S6, an advanced hydraulic system was used, where a triple seismometer signal was used as input for eight actuators controlling the six degrees of freedom of the in-vacuum optical tables [83]. In addition, a multiple-pendulum stack is used to passively damp motion above the resonance frequency of the pendula; figure 3.5 shows the response of a pendulum to an external stimulus: the response is flat up to the resonance ($\sim 1\,\mathrm{Hz}$) at which any stimulus is greatly magnified, and above which the amplitude falls away with frequency. The materials of the suspensions, and the methods by which they are connected to the mirrors were chosen to have the resonances at low frequency (below the detection band of the full interferometer), and with high quality factor, providing a very narrow peak. In first-generation LIGO core optics were hung from triple suspensions, while quadruple stacks will be in use for Advanced LIGO (see section 3.5.1).

3.4.2 Thermal noise

There is no reason that seismic noise cannot be damped further, and further, until the resulting motion is below the required levels. However, due to mechanical thermal excitations, the test mass suspensions – and the masses themselves – will move, creating another noise source [84]. These noises are fundamentally limited by the equipartition theorem, which puts a lower limit on the random motion of the mirror surfaces at the atomic level [85]. The thermal noise of the mirror suspensions is concentrated around the vibrational resonances as shown in figure 3.6 (c.f. figure 3.5).

Dissipative losses from the suspensions can be reduced by making the wires as thin as possible, or using a less dissipative material; first-generation iLIGO used steel wires, while aLIGO will use fused silica [87]. In addition the resonances are tuned to be outside of the detection band of the instrument as much as is practicable. Thermal noise in the test masses is dominated by the coatings used to provide a smooth surface.

3.4.3 Quantum noises

The third main noise category is roughly split into two parts, with radiation pressure noise another dominant component at mid frequencies, while quantum shot noise dominates fundamentally at high frequencies [88].

Radiation pressure noise is caused by quanta of light in the laser beam transferring momentum onto the mirrors upon reflection. This noise is most prevalent around 50 Hz – and will be limiting for aLIGO – but can be mitigated with larger, heavier mirrors; however, higher laser power (as used in S6 and planned for aLIGO) will exacerbate the problem.

Quantum shot noise is the dominant noise above a few hundred Hertz, and comes as a result of statistical fluctuations in the rate of photons incident on the output photodiode. The magnitude of this noise is ruled by the Planck constant, h, and laser power in the interferometer. Thus it is fundamental for a given laser, but can be reduced by moving to a higher power – at the expense of higher radiation pressure noise.

3.5 Second-generation detectors

The first-generation LIGO detectors described above stopped operation in October 2010, ending the second joint science run with the Virgo detector. These
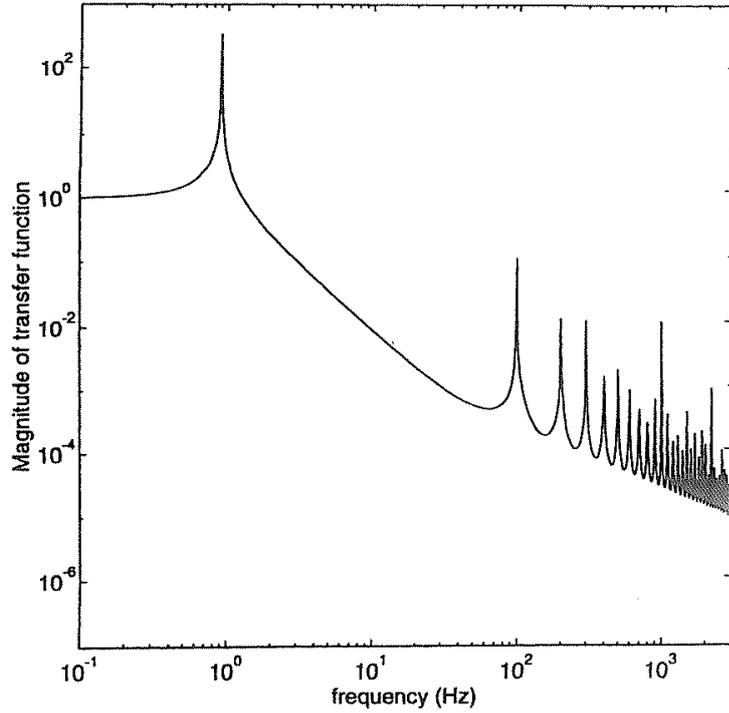
Figure 3.5: Example of the transfer function for a pendulum [86]. The response to a stimulus falls away with frequency above the fundamental resonant peak. Higher order resonances are hidden by other noise sources in all GW detectors.
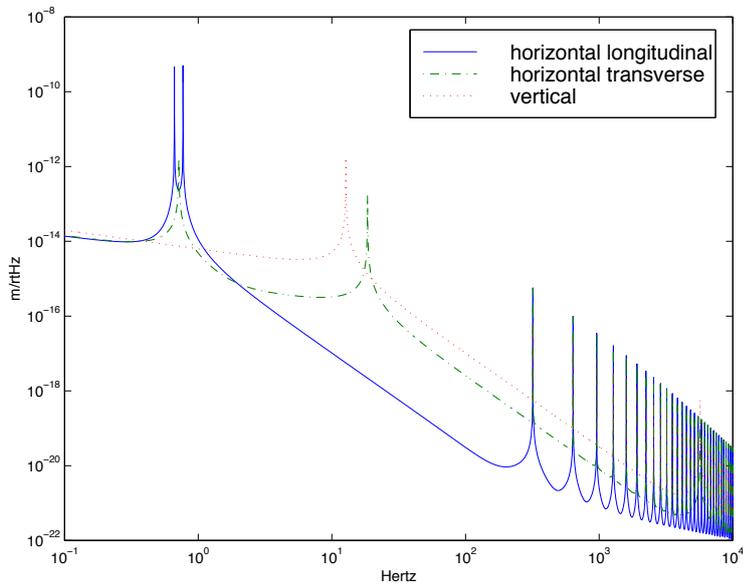
Figure 3.6: Thermal displacement noise for the Initial LIGO (iLIGO) mirrors [84]. The thermal noise from the test mass suspensions is heavily coupled to the resonant modes of oscillation (cf. figure 3.5).
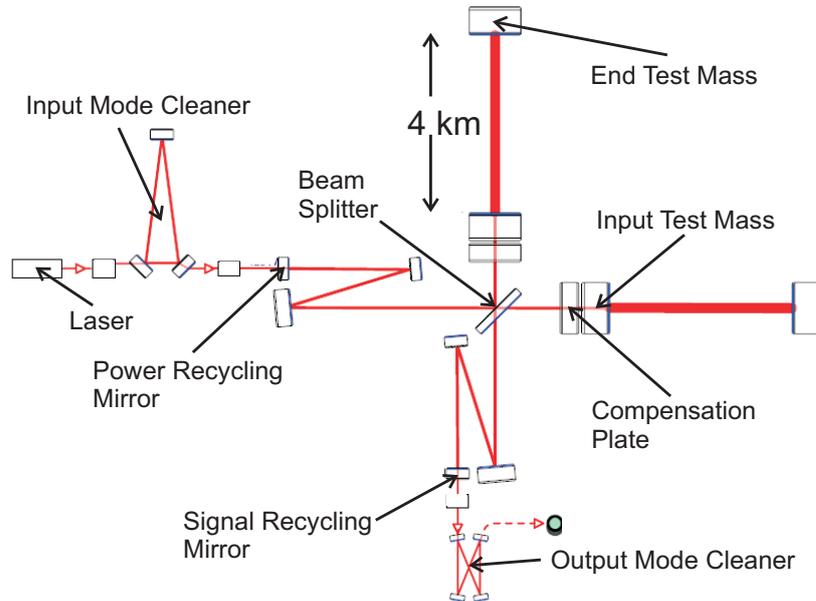
Figure 3.7: The simplified optical layout of the aLIGO interferometers [89]. The improved design includes a signal-recycling cavity relative to initial LIGO (cf. figure 4.1).

instruments were decommissioned in preparation for the installation of the second-generation Advanced LIGO interferometers [89]. Virgo remained in operation until late in 2011, operating a joint run with GEO600, before it too was decommissioned to make way for Advanced Virgo [90, 91].

The Advanced LIGO detectors are designed to be an order of magnitude more sensitive to GW strain amplitude than the first-generation, representing an ×1000 increase in the rate of detectable events.

3.5.1 Key hardware improvements

The aLIGO instruments include a huge number of upgrades and improvements relative to the first-generation LIGO design [89, 92]. We detail a number of the key upgrades that improve performance relative to each of the three fundamental noise sources described above. The designed layout of the new instruments is shown in figure 3.7.

Seismic isolation

For aLIGO, the seismic isolation system is split into two designs, separately for the large chambers containing suspended core optics – the beam-splitter and

test mass mirrors – and the small chambers with auxiliary optics – steering mirrors and optical periscopes, for example – mounted on tables [82, 83].

Both styles of chamber will be fitted with the hydraulic actuation system installed at LLO during S6 (see chapter 4 for more details). A combination of feed-forward controls from seismometers on the ground and feed-back controls from various sensors attached to vacuum chambers will be used to provide a factor of 10 decrease in seismic noise up to 5 Hz. Additionally, an in-vacuum feed-back system will control two stages of the isolation platform from which the optics are suspended, to mitigate motion above 0.2 Hz.

Suspensions and optics

The core optics (test masses and the beam-splitter), will be suspended from a quadruple pendulum system, adapted from that used first with GEO600 [93]. In this design, the final stage suspension wires, the mirror suspension bonds ('ears'), and the test mass mirrors themselves will all be made from the same materials, and bonded together to create a monolithic final stage with much lower loss than that used in initial LIGO [87, 94]. It is predicted that the thermal noise from this system will sit below that of radiation pressure noise (contrary to the initial LIGO design), however this depends on the laser power used in operations.

The test mass optics themselves will be 40 kg (10 kg for initial LIGO) silica cylinders [89]; greater mass and larger radius relative to the first-generation will reduce radiation pressure noise whilst still being practical. Improved coatings have also been developed to reduce thermal noise and scattering [95].

To compensate for the increased thermal lensing effects of a higher laser power, ring heaters are being fitted around the input test masses, designed to radiate onto the outer barrel of the optics and reduce the lensing curvature [96]. A CO_2 laser system similar to that from initial LIGO will also heat reaction masses immediately behind the test masses ('compensation plates' in figure 3.7) to dynamically shape the optic through radiative heating.

Input laser

A more powerful laser system will be used to provide the main beam into the interferometer. Similar to initial LIGO, a 1064 nm laser will be used, but now with a peak input power of 180 W [97]. This system builds on the 35 W input used in S6, with an extra ring oscillator stage to build power to the maximum. Improved amplitude- and frequency-stabilisation systems have

also been developed. At the time of writing, this system has already been installed, and is in performance testing, at both of the LIGO sites.

Signal recycling

While the first-generation LIGO detectors were power-recycled interferometers, the aLIGO detectors will be dual-recycled, with the inclusion of a signal-recycling mirror [98]. This mirror constructs another auxiliary optical cavity used to recycle signal power from the output port back into the main interferometer. Dual-recycling has successfully been implemented for the GEO600 detector [71].

3.5.2 Sensitivity improvement

As noted above, these improved components will give the second-generation instruments a factor of ten increase in sensitive distance, if they run in their designed configuration. However, the design sensitivity will not immediately be achieved in the first advanced science run; rather a series of shorter data-taking periods, with ever-increasing sensitivity, will lead to this level by the year 2019 [64]. Figure 3.8 shows the projected evolution of the spectral sensitivity over a number of observing periods.

As the sensitivity increases, the detectable rate of events will rise. While the first-generation instruments, at the peak of their sensitivity, had a plausible detection rate of 1 event per 50 years (for a binary neutron star coalescence), the second-generation detectors running at their design sensitivity should see around 40 such events every year [52] (recall the substantial error bars from section 2.5.3).
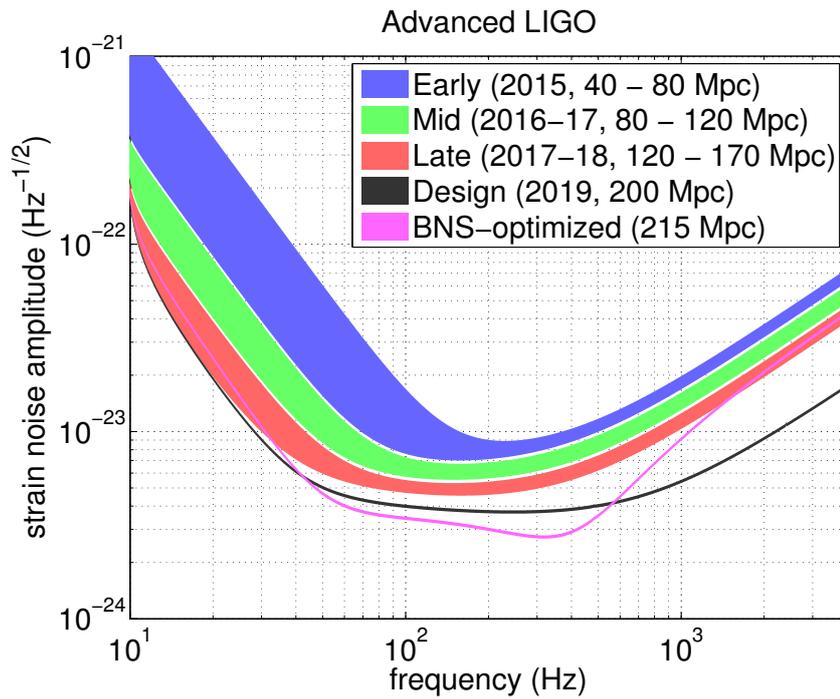
Figure 3.8: Evolution of the strain sensitivity of the aLIGO detectors through 'early, 'middle', and 'late' commissionging phases, including the designed sensitivity [52]. The pink curve shows the sensitivity in the case of the full instrumental configuration tuned for maximal sensitivity to a binary neutron star merger. For each stage, the average sensitive distance to such a BNS merger event is given in Mpc.

Chapter 4

Characterisation of the LIGO detectors during their sixth science run

4.1 Introduction

Between July 2009 and August 2010, the LIGO [13] operated two kilometre-scale laser interferometers as part of a global network aiming to detect and study GWs of astrophysical origin. These detectors, at LIGO Hanford Observatory, WA (LHO), and LIGO Livingston Observatory, LA (LLO) – operating beyond their initial design with greater sensitivity – took data in collaboration with GEO600 [14] and Virgo [15].

The data from each of these detectors have been searched for GW signals from a number of sources, including CBCs [18, 19, 20], unmodelled GW bursts [19, 21], non-axisymmetric spinning neutron stars [42], and a SGWB [99]. The ideal performance of each of these analyses is determined by the searched volume of the universe times the searched time duration. However, long and short duration artefacts in real data, such as glitches and spectral lines, further restrict the sensitivity of GW searches.

Searches for transient GW signals including CBCs and GW bursts are sensitive to many short-duration noise events (known as glitches), the most common artefacts seen in the LIGO detectors, coming from a number of environmental, mechanical, and electronic mechanisms. Each search pipeline employs powerful signal-based methods to distinguish a GW event from noise [48, 100, 101, 102], but also relies on careful studies of the detector behaviour to provide information that leads to improved data quality through 'vetoes' that remove data

likely to contain noise artefacts from the analyses. Searches for long-duration continuous waves (CWs) and a SGWB are sensitive to disturbances from spectral lines and other sustained noise artefacts. These effects cause elevated noise at a given frequency and so impair any search over these data.

This paper describes the work done to characterize the LIGO detectors and their data during the S6, and estimates the improvement in sensitivity for analyses resulting from detector improvements and data quality vetoes. This work follows from previous studies of LIGO data during S5 [103, 104] and S6 [105, 106]. Similar studies have also been performed for the Virgo detector relating to data taking during Virgo Science Runs (VSRs) 2, 3 and 4 [107, 108].

Section 4.2 details the configuration of the LIGO detectors during S6, and section 4.3 details their performance over this period, outlining some of the problems observed and improvements seen. Section 4.4 describes examples of important noise sources that were identified at each site and steps taken to mitigate them. In section 4.5, we present the performance of data-quality category vetoes when applied to each of two astrophysical data searches: the ihope CBC pipeline and Coherent WaveBurst (cWB) burst pipeline. A short conclusion is given in section 4.6, along with plans for characterization of the next-generation Advanced LIGO detectors, currently under construction.

4.2 Configuration of the LIGO detectors during the sixth science run

The first-generation LIGO instruments were augmented versions of the Michelson interferometer [66], with which GW amplitude is measured as a strain of the 4-kilometre arm length, as shown in figure 4.1 [75]. In this layout, a diode-pumped, power-amplified Nd:YAG laser generated a carrier beam in a single longitudinal mode at 1064 nm [109]. This beam passed through an electro-optic modulator which added a pair of RF sidebands used for sensing and control of the test mass positions, before the modulated beam entered a triangular optical cavity. This cavity (the 'input mode cleaner') was configured to filter out higher-order spatial modes from the main beam before it entered the main interferometer.

The initial Michelson design was enhanced with the addition of a power-recycling mirror, which reflected any light returned towards the input system back into the interferometer. The resonant build-up in the power recycling cavity (PRC) acted to effectively increase the power of the input laser beam [78].
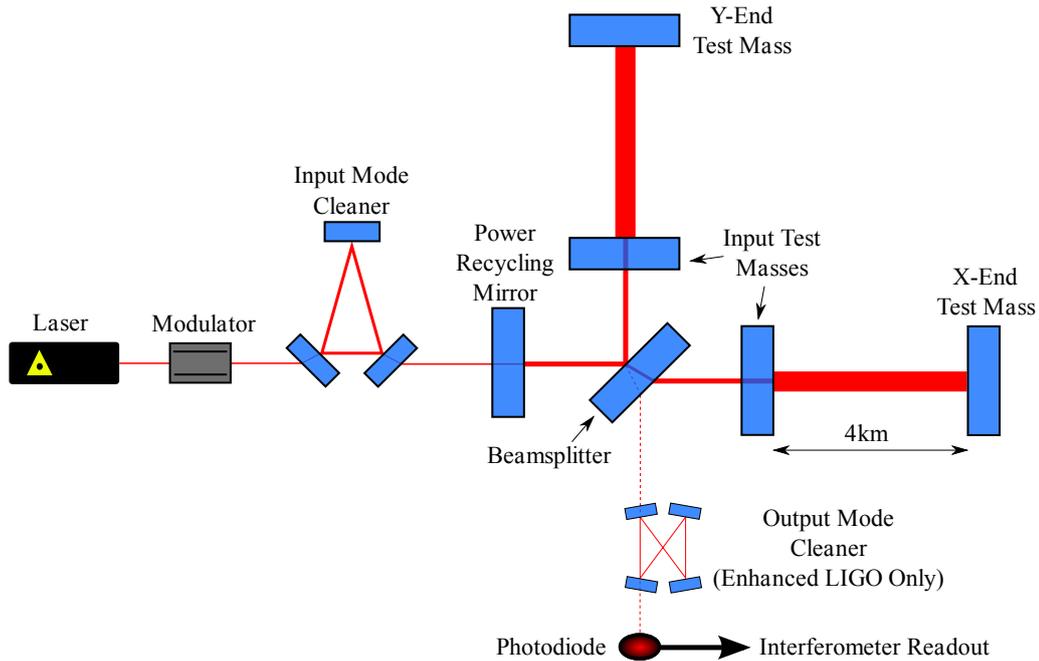
Figure 4.1: Optical layout of the LIGO interferometers during S6 [75]. The layout differs from that used in S5 with the addition of the output mode cleaner.

Additionally, the effective lengths of the interferometer arms were increased with the addition of input test masses; these formed Fabry-Perot cavities in the arms, storing the light for a much longer time. During S5, the relative lengths of each arm were controlled to ensure that the light exiting each arm cavity interfered destructively at the output photodiode, and all power was returned towards the input. In such 'dark fringe' operation, the phase modulation sidebands induced by interaction with GWs would interfere constructively at the output, recording a GW strain in the demodulated signal. In this configuration, the LIGO instruments achieved their design sensitivity goal over a 2-year run. A thorough description of the initial design is given in [13].

For the sixth science run (S6), a number of new systems were implemented to improve sensitivity and to prototype upgrades for the second-generation aLIGO detectors [75, 89]. The initial input laser system was upgraded from a 10 W output to a maximum of 35 W, with the installation of new master ring oscillator and power amplifier systems [110]. The higher input laser power from this system improved the sensitivity of the detectors at high frequencies ($> 150\,\text{Hz}$) and allowed prototyping of several key components for the aLIGO laser system [111]. Additionally, an improved CO_2-laser thermal-compensation

system was installed [112, 113] to counteract thermal lensing caused by expansion of the test mass coating substrate due to heat from absorption of the main beam.

An alternative GW detection system was installed, replacing the initial heterodyne readout scheme [114]. A special form of homodyne detection, known as *DC readout*, was implemented, whereby a local oscillator field was introduced at the same frequency as the main laser beam [80]. In this system, GW-induced phase modulations would interfere with this field to produce power variations on the output photodiode, without the need for demodulating the output signal. In order to improve the quality of the light incident on the output photodiode in this new readout system, an output mode cleaner (OMC) cavity was installed to filter out the higher-order mode content of the output beam [115]. The OMC was required to be in-vacuum, but also highly stable, and so a new single-stage seismic isolation system was designed and installed for the output optical platform [116], from which the OMC was suspended.

Futhermore, controls for seismic feed-forward to a hydraulic actuation system were implemented at LLO to combat the higher level of seismic noise at that site [117]. This system used signals from seismometers at the Michelson vertex, and at ends of each of the arms, to suppress the effect of low-frequency (below ~ 10 Hz) seismic motion on the instrument.

4.3   Detector sensitivity during S6

The maximum sensitivity of any GW search, such as those cited in section 4.1, is determined by the amount of coincident multi-detector operation time and astrophysical reach of each detector. In searches for transient signals these factors determine the number of sources that could be detected during a science run, while in those for continuous signals they determine the accumulated signal power over a data-taking run.

The S6 run took place between July 2009 and October 2010, with each detector recording over seven months of data in that period. The data-taking was split into four epochs, A–D, identifying distinct analysis periods set by changes in detector performance or the detector network itself. Epochs A and B ran alongside the second Virgo Science Run (VSR2), before that detector was taken off-line for a major upgrade [108], with S6A running for ~ 2 months before a month-long instrumental commissioning break; S6B ran to the end of 2009, before another commissioning break. The final 2 epochs, C and D,
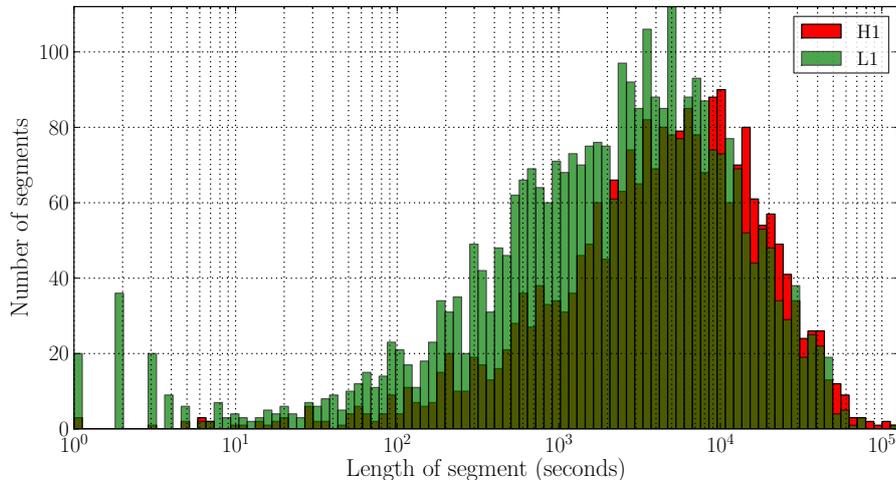
Figure 4.2: A histogram of the duration of each science segment for the LIGO detectors during S6. The distribution is centred around ~ 1 hour.

spanned a continuous period of detector operation, over nine months in all, with the distinction marking the start of VSR3 and the return of a three-detector network.

Instrumental stability over these epochs is measured by the detector duty cycle – the fraction of the total run time during which science-quality data was recorded. Each individual data-taking stretch is known as a *science segment*, when the interferometer arm cavities are locked in resonance and the spectral sensitivity is deemed acceptable by an operator on duty. A science segment is typically ended by a critically large noise level in the instrument meaning arm-cavity resonance cannot be maintained by the electronic control system (known as *lock-loss*), however a small number of segments are ended manually during clean data in order to perform a measurement for calibration purposes, for example. Figure 4.2 shows a histogram of science segment duration over the run; as seen, the majority of segments span several hours, but there are a significant number of shorter segments, symptomatic of interferometer instability. In particular, for L1 the rate of shorter segments is higher than that for H1, a result of poor detector stability during the early part of the run, especially during S6B.

Table 4.1 summarises the science segments for each site over the four run epochs. Both sites saw an increase in duty cycle, that of H1 increasing by 30%, and L1 by over 40% between epochs A and D. Additionally, the median duration of a single science-quality data segment more than doubled at both sites between the opening epochs (S6A and S6B) and the end of the run. These

Epoch	Median duration (mins)	Longest duration (hours)	Total live time (days)	Duty cycle (%)
S6A	54.0	13.4	27.5	49.1
S6B	75.2	19.0	59.2	54.3
S6C	82.0	17.0	82.8	51.4
S6D	123.4	35.2	74.7	63.9

(a) H1

Epoch	Median duration (mins)	Longest duration (hours)	Total live time (days)	Duty cycle (%)
S6A	39.3	11.8	25.6	45.7
S6B	17.3	21.3	40.0	38.0
S6C	67.5	21.4	82.3	51.1
S6D	58.2	32.6	75.2	64.3

(b) L1

Table 4.1: Science segment statistics for the LIGO detectors over the four epochs of S6.

increases in stability highlight the developments in understanding of the critical noise couplings and how they affect operation of the instruments, as well as improvements in the control system used to maintain cavity resonance.

The sensitivity to GWs of a single detector is typically measured as as a strain amplitude spectral density of the calibrated detector output. This is determined by a combination of noise components, some fundamental to the design of the instruments, and some from additional noise coupling from instrumental and environmental sources. Figure 4.3 shows the representative amplitude sensitivity of the LIGO detectors during S6. The dominant contribution below 40 Hz is noise from seismically-driven motion of the core interferometer optics, and from the servos used to control their alignment. The reduced level of the seismic wall at L1 relative to H1 can be, in part, attributed to the prototype hydraulic isolation installed at that observatory [117]. Intermediate frequencies, 50-150 Hz, have significant contributions from thermal effects: mechanical excitations of the test masses and their suspensions due to thermal motion [118, 95]. Above 150 Hz, quantum shot noise from discrete photon detection at the output port is the dominant fundamental noise source [119]. The spectral sensitivity gives a time-averaged view of detector performance, and so is sensitive to the long-duration noise sources and signals, but is insensitive to transient noise and signals.

The standard measure of a detector's astrophysical reach is the distance to which that instrument could detect GW emission from the inspiral of a binary neutron star system with a signal-to-noise ratio (SNR) of 8 [120, 17], averaged over source sky locations and orientations. Figure 4.4 shows the evolution of
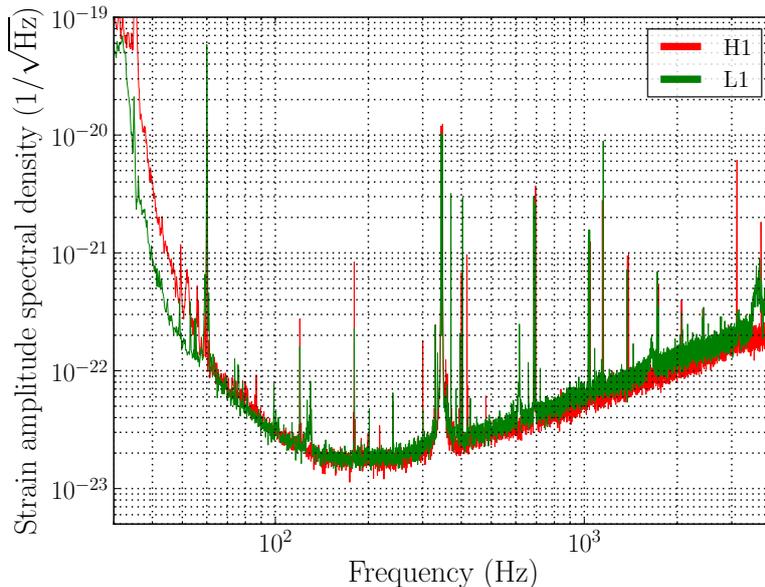
Figure 4.3: Representative strain amplitude sensitivity of the LIGO detectors during S6.

this value over the science run, with each data point representing an average over 2048 s of data. Over the course of the run, the detection range of H1 increased from $\sim$ 16–20 Mpc, and of L1 from $\sim$ 14–20 Mpc. The instability of S6B at L1 can be seen between days 80–190, with the lower duty cycle (also seen in table 4.1) and low detection range; this period included extensive commissioning of the seismic feed-forward system at LLO [117].

For sources uniform in volume (such as binary neutron stars), the expected detectable number for an instrumental network is determined as the product of the searched volume of the universe and the searched duration. From the increased amplitude sensitivity, and improve duty cycle seen between S6 epochs A and D, the rate of detectable sources for the LIGO detectors was greatly increased.

4.4 Data quality issues

Detector data frequently contains components from non-fundamental noise sources, that can have a significant impact on search sensitivity. The challenge, then, is to identify the sources of the important non-Gaussian noise components and reduce them or to veto the times or spectral frequencies at which they occur. One of the methods for doing this is to correlate noise in the detector outputs with one or more of the $> 20,000$ auxiliary signals recorded
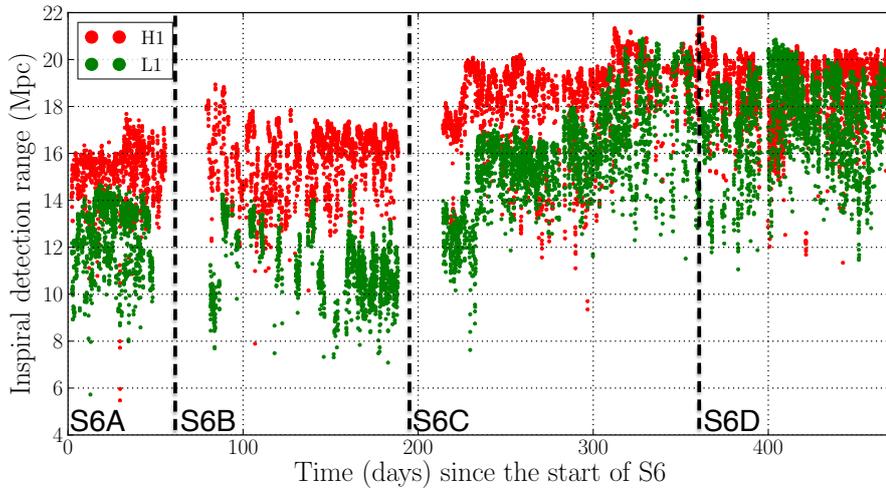
Figure 4.4: The inspiral detection range of the LIGO detectors throughout
S6 to a binary neutron star merger, averaged over sky location
and orientation. The rapid improvements between epochs can be
attributed to hardware and control changes implemented during
commissioning periods.

by instrumental and environmental sensors at each observatory.

Each noise investigation probes different detector components and requires different methods, but the basic form is as follows. A new noise source is identified in the GW data, either by its effect on the sensitivity of an astro-physical search, or a change in the detector power spectral density (PSD), and logged for a follow-up investigation. This noise is then correlated to an aux-iliary instrumental or environmental signal, one that would not see the effect of a GW signal, and the coupling mechanism hopefully identified. During S6 tens of thousands of instrumental control signals were recorded, along with a number of acoustic, magnetic, and seismic sensors monitoring the physical environment around the instrument [121].

This coupling is identified by scanning auxiliary data around the noisy time, or by studying statistical correlations between glitches in these data and the GW data. Ideally this auxiliary system can be modified in hardware, or in control software, to remove this noise source from current and future data. However, when such a fix cannot be implemented, or when contaminated data will corrupt an analysis, noisy periods in this auxiliary channel are noted and recorded as likely to adversely affect the GW data. These data quality (DQ) flags and their associated time segments are recorded; each analysis can construct lists of times with elevated noise to use in deciding which time to analyse, or which detection candidates to reject as likely noise artefacts. In

section 4.5, we show the positive impact of identifying noisy times, even if an instrumental fix cannot be implemented.

Special care must be taken to ensure that a flag will not be activated under the influence of an astrophysical event, and so all flags are tested against simulated GW signals known as *hardware injections*. These are performed by mechanically exciting one of the interferometer test masses to mimic a passing GW, testing the instrument calibration, analysis efficiency, and DQ flag safety. If a statistically significant number of hardware injections are associated with a given flag, it is considered unsafe and is not used.

The remainder of this section outlines a number of specific issues that were present for some time during S6 at LHO or LLO, some of which were fixed at source, some which were tagged to be vetoed, and one which was not resolved.

4.4.1 Seismic noise

In the initial design for LIGO, seismic noise was predicted to limit the sensitivity to GWs below 40Hz. As described above, seismically driven motion did, in fact, limit the spectral sensitivity during S6. However, seismic noise was also observed to be strongly correlated with glitches in the detector output, often at higher frequencies ($\sim 100 - 200\,\text{Hz}$) than those limited by linear coupling of seismic noise.

During S6, the low-latency KleineWelle and the Ω-pipeline [122] glitch monitors identified seismic noise as a regular source of short-duration glitches during working hours at both sites. The top panels of figure 4.5 shows traces and a time-frequency map of seismic motion on the floor of the LHO instrumental area, with the lower left and right panels highlighting the correlation with events recorded by low-latency burst and inspiral searches respectively. Crucially, during high seismic noise, the low-latency inspiral pipeline 'daily ihope' [102] produced candidate events (triggers) across the full range of signal templates. This behaviour of these DQ event generators was echoed by the full analyses for burst and CBC, increasing the significance of the noise backgrounds and limiting the sensitivity of the searches.

A new method of targeting the the Ω-pipeline at noise events in seismometer data was developed in order to identify times of high coupling into the GW readout [27]. This method allowed identification of time-segments during which large number of glitches in the GW readout were correlated with seismic noise events, in a relatively short amount of search time.
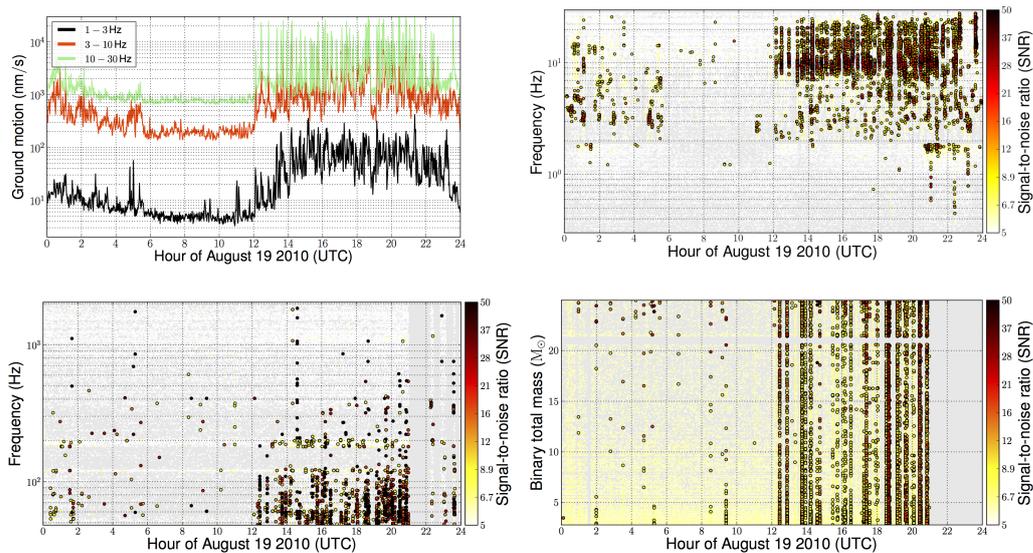
Figure 4.5: Seismic motion of the laboratory floor in various frequency bands (top-left) and as seen by the the Ω-pipeline (top-right), and its correlation into low-latency GW burst (bottom-left) and inspiral (bottom-right) analyses.

4.4.2   Seismically-driven length sensing glitches

During S6B, the majority of glitches at L1 were correlated with noise in the length control signals of two short length degrees of freedom: the PRC, and the short Michelson formed by the beam-splitter and the input test masses (MICH). Both of these length controls were glitching simultaneously, and these glitches were correlated with upwards of 70% of the glitches in the GW data.

The length signals for both MICH and PRC are derived from a photodiode signal reading light picked-off at the beam splitter, however, the glitches were absent in that photodiode signal, indicating that they were not due to a physical length change. However, during the winter, the microseismic noise is particularly high, driving large instabilities of the power recycling cavity that caused significant drops in the circulating power and resulted in large glitches in both the MICH and PRC length controls. These actuation signals, applied to the main interferometer optics, then coupled into the detector output.

In an attempt to decrease the rate of these glitches, additional low-pass filtering was added to the power recycling gain signal to smooth the fluctuations. This was not successful at reducing the glitch rate, possibly because the low-passing was not aggressive enough to eliminate the largest dropouts in the gain signal.

These glitches were eliminated in mid December 2009 by commissioning

work on a seismic feed-forward system [117]. This decreased the PRC motion by a factor of three by recording low-frequency seismic motion on the ground and feeding this signal to a system of hydraulic actuators. With the cavity stabilized, the power recycling gain did not fluctuate as much and this type of glitch disappeared. The glitchy data before the fix were identified by both the HierarchichalVeto (HVeto) and Used Percentage Veto (UPV) algorithms [123, 124] – used to rank auxiliary signals according to the statistical significance of glitch coincidence with the GW data – with the relevant time segments used by the searches to dismiss noise artefacts from their results (more in section 4.5).

4.4.3 Upconversion of low-frequency noise due to the Barkhausen effect

In the early LIGO science runs, reductions in interferometer sensitivity were noticed in the 40–200 Hz band associated with increased levels of ground motion below 10 Hz produced by passing trucks, distant construction activities, seasonal increases in water flow over dams, high wind, and earthquakes [125, 75, 104, 27]. This noise, termed seismic upconversion noise, often limited the sensitive distance of the instruments. Figure 4.6 shows that a reduction in the range to binary neutron star inspirals was contemporaneous with the workday increase in anthropogenic seismic noise.

Experiments subsequently showed that seismic upconversion noise levels correlated better with the amplitudes of the currents to the electromagnets that held the test masses in place as the ground moved than with the actual motion of the test masses or of the ground. An empirical, frequency-dependent function was developed to estimate upconversion noise from the low-frequency test mass actuation currents. This function was used to produce flags that indicated time periods that were expected to have high levels of seismic upconversion noise.

In addition to average reductions in sensitivity, seismic upconversion noise transients further reduced sensitivity to unmodelled GW bursts. Figure 4.7 shows that the rate of low-SNR glitches – in a frequency band above that for seismic noise – was correlated with the test mass actuation current, suggesting that seismic upconversion was the source of a low-SNR noise background that limited GW burst detection.

Investigations found that seismic upconversion noise bursts were clustered in periods of high slope in the amplitude of the current to the magnetic actuators. This was evidence that the seismic upconversion noise was Barkhausen
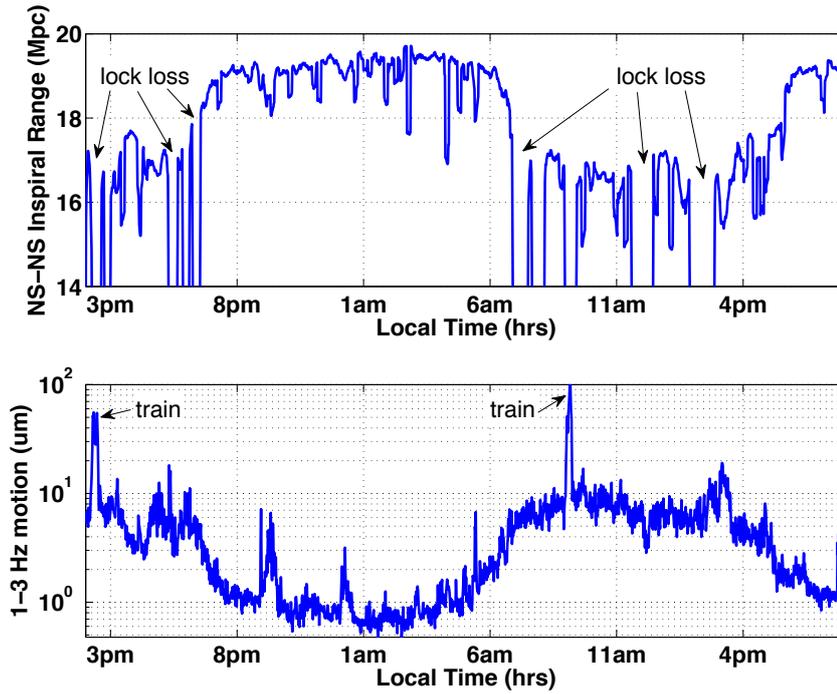
Figure 4.6: Range to a binary neutron star (top) and ground motion in the 1–3 Hz band (bottom) for a day at LLO. The inverse relationship is believed to be due to non-linear up-conversion of low frequency seismic ground motions to higher frequency ($\sim 50-250\,\mathrm{Hz}$) noise in the GW output.
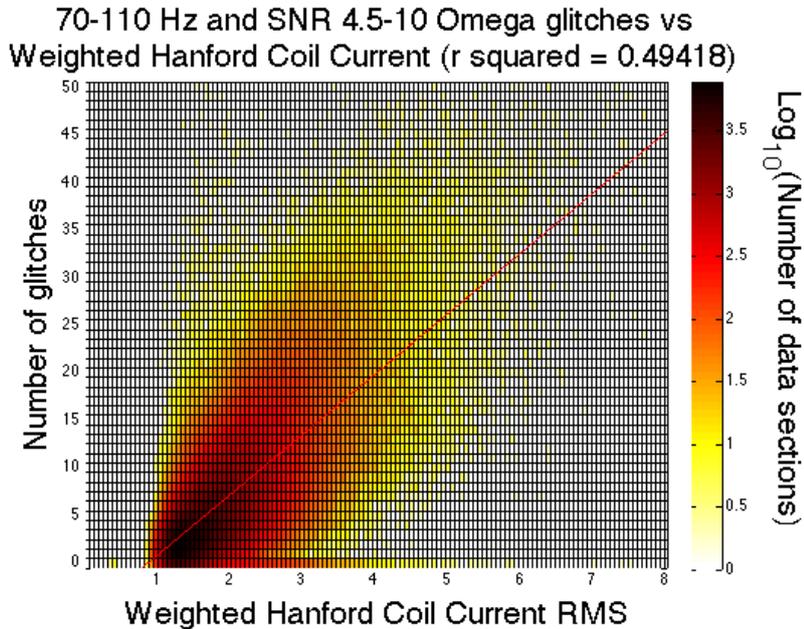
Figure 4.7: Correlation between low SNR glitches in the GW output, and current in the test mass coil at H1. The r^2 value measures the normalized cross-correlation through a least-squares fit. This correlation is indicative of the Barkhausen effect.

noise [126]: magnetic field fluctuations produced by avalanches of magnetic domains in ferromagnetic materials that occur when the domains align with changing magnetic fields. The Barkhausen noise hypothesis was supported by investigations in which the noise spectrum was reproduced by magnetic fields that were generated by a separate system. These investigations also suggested that the putative source of the Barkhausen noise was near or inside the test mass actuators. It was found that fasteners inside the magnetic actuator, made of grade 303 steel, were ferromagnetic, probably because they were shaped or cut when cold. For aLIGO, grade 316 steel, which is much less ferromagnetic after cold working, is being used at the most sensitive locations.

4.4.4 Beam jitter

Throughout the science run, misalignment of the primary laser beam relative to any optical cavities in the instrument resulted in excess noise on the output signal. The arm cavities are aligned to allow maximum transmission of the fundamental TEM_{00} mode of the input laser light, and to damp the amplitude of other, higher-order modes, helped by both input and output mode cleaners respectively before and after the main interferometer.

However, motion of the main interferometer test masses, auxiliary mirrors steering the output beam onto the OMC, and the OMC itself, caused transmission of non-TEM_{00} modes, and introduced noise sidebands attached to existing alignment noise [80]. Such noise was seen at both sites during S6, with seismic noise and vibration of optical tables introducing low-frequency sidebands around existing higher frequency spectral lines. This instability can be seen causing glitches at $180 - 200\,Hz$ in figure 4.5.

The mitigation of jitter glitches involved ensuring that the alignments of all optical cavities were controlled in the most stable manner possible. Detector operators modified these controls as required to minimise the transmission of higher order optical modes into the output. Further details can be found in [80].

4.4.5 Mechanical glitching at the reflected port

A set of glitches was produced by faults in the servo actuators used to stabilize the pointing of the beam at the reflected port of the H1 interferometer. This table holds a number of photodiodes used to detect light reflected from the beam-splitter mirror towards the input, used to sense and control angular
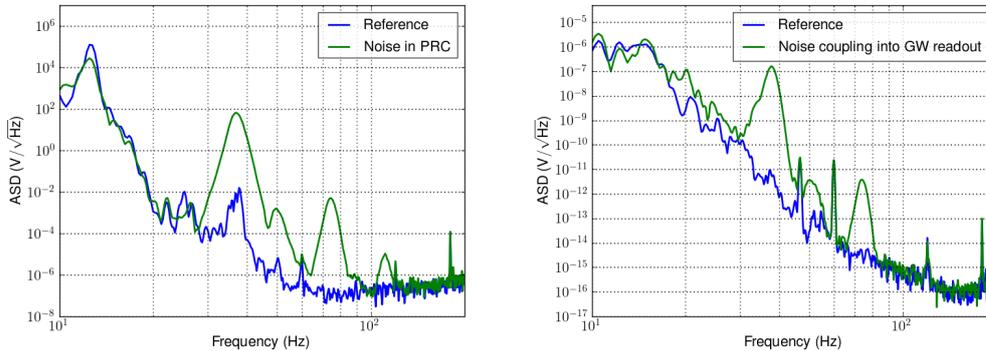
Figure 4.8: Broad noise peaks centred at $37\,\text{Hz}$ and its harmonics in the power recycling cavity length signal (left) and the GW output error signal (right). The green curves show the noise in comparison with a reference taken from clean data.

motion of the arm-cavity mirrors.

The source of the glitches was identified with the help of HVeto, which discovered that a number of angular and length sensing channels derived from photodiodes at the reflected port were strongly coupled with events in the GW data. Figure 4.8 shows the broad peaks in the spectra of one length sensing channel and the un-calibrated GW readout compared to a quiet reference time. In addition, accelerometer signals from the optical table at the reflected port were found to be coupling strongly, having weak but coincident glitches. The accelerometer coincidences suggested the possibility that the glitches were produced by mechanical motions of steering mirrors resulting from a faulty piezo-electric actuation system. The piezoelectric power supplies were shut off and, when this stopped the glitches, the steering servo was deemed unnecessary and decommissioned.

4.4.6   Broadband noise bursts from poor electrical connections

Another example of glitches caused by an electronics problem were dubbed the 'grid glitches', appearing in H1 data toward the end of S6. This noise was characterised by periods of repeated glitching lasting from minutes to hours, generating short, broadband triggers with SNR ~ 20 in the the Ω-pipeline [127], as shown in figure 4.9.

The main diagnostic clues were coincident but louder glitches in all of the quadrant photo-diodes (QPDs) monitoring the OMC; two of these QPDs each
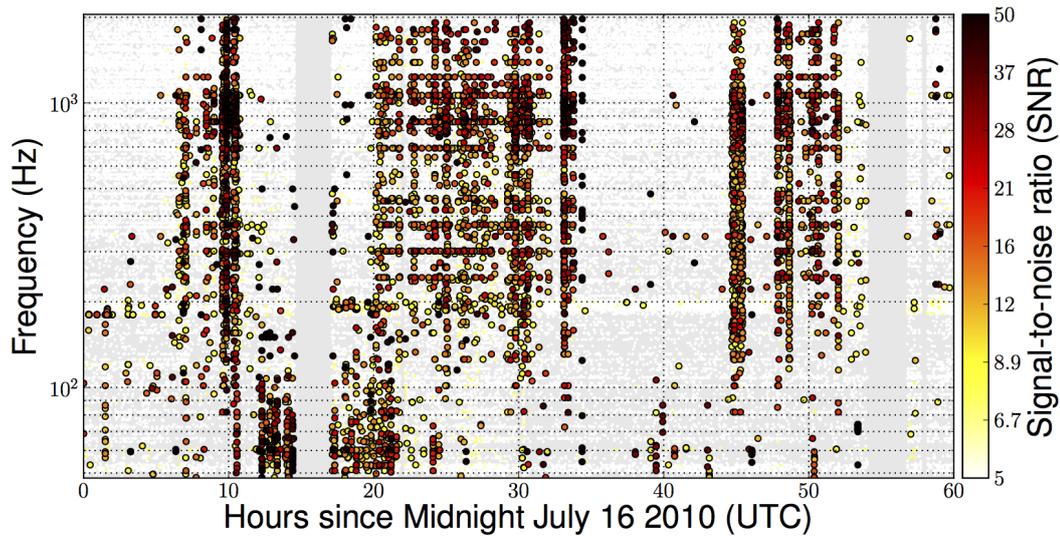
Figure 4.9: Noise events recorded by the the Ω-pipeline over a 60-hour period at LHO. The high SNR events above 100 Hz were caused by broadband noise from a faulty electrical connection. The grid-like nature of these events is due to the discrete tiling in frequency by the trigger generator.

monitor the beam incident on the OMC and that reflected from it. The shape of the signal on the QPD was very sharp and unnatural, making it unlikely that the noisy QPDs could detect a glitch in the beam more sensitively than the GW photodiode. The prime suspect then became the electronics involved with reading these photodiodes, or electronics in proximity to the readout, causing glitches that were looped back into the OMC controls.

In the process of isolating the cause, several other electronics boards in the output mode cleaner were inspected, re-soldered, and/or swapped for spares. The problem was finally solved by re-soldering the connections on the electronics board that provided the high-voltage power supply to drive a piezo-electric transducer.

4.4.7 The 'spike' glitch

The spike glitch was the name given to a class of very loud transients seen in the instrument at LLO. They were characterized by a distinctive shape in the time series of the light on the GW output photodiode, beginning with a narrow but smooth dip (lasting ~ 1 ms) before a period of damped oscillation; these oscillations are noticeably longer than the beginning spike, with a half-width of about 3 milliseconds. The amplitude of the glitch is extremely large, meaning it is visible in the raw time series (which is dominated by low-frequency seismic
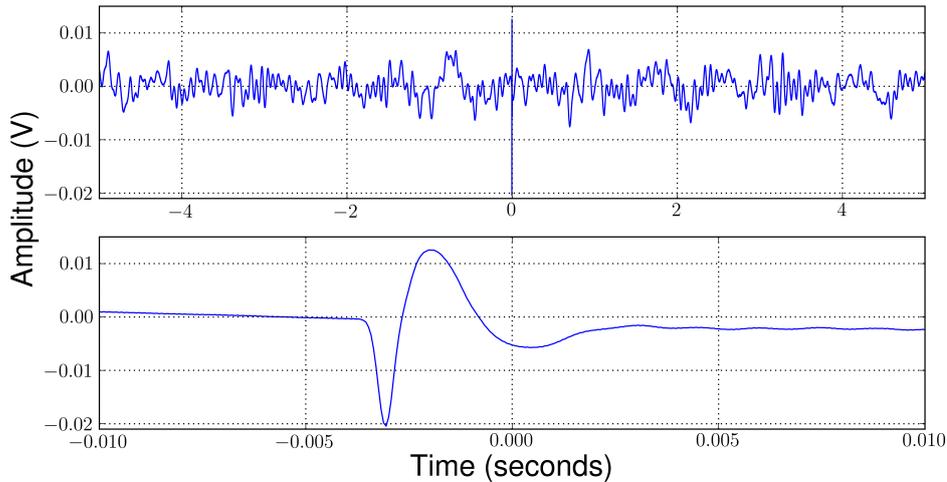
Figure 4.10: A spike glitch in the raw GW photodiode signal for L1. The top panel shows the glitch in context with 10 s of data, while the bottom shows the glitch profile as described in the text.

motion), as shown in Figure 4.10; other types of glitches, except for severe seismic disturbances, are not visible in the time series. The SNRs as seen by a burst search range from 200 to well over 20,000. There may be spike glitches at lower SNRs that are not identifiable as such because they do not stand out clearly in the noise. These glitches were surprisingly common given their amplitude, with about 10 to 30 per day at apparently random times.

The cause of the spike glitch was never determined despite many investigations. The Fabry-Perot arms of the interferometer act as a low-pass filter, so the size and rapidness of the initial dip suggests that the glitch must have a source after the light has left the beam splitter on its way to the readout. The damped oscillations after the initial dip, however, are likely due to the response of the length control loop of the interferometer. So an actual or apparent sudden dip in the light on the output photodiode could explain the entire shape of the spike glitch. To investigate this possibility, the interferometer was run in a configuration where the light did not enter the arm cavities, but went almost directly into the output mode cleaner. This removed the length and angular servos from consideration. Sharp downward dips in the light were seen during this test, although they were 0.2 milliseconds wide, much narrower than the initial dips of the spike glitches.

4.5 The impact of data quality on gravitational wave searches

The impact of non-Gaussian noise in the LIGO detectors on searches for GWs is significant. In Gaussian noise, an event with SNR = 8 is rare, occurring only a few times per day, while glitches with SNR over $1,000$, such as the spike glitch described previously, would effectively never occur. Such loud glitches can mask or greatly disrupt transient GW signals present in the data at the same time, while high rates of lower SNR glitches can significantly increase the background in searches for these sources. Additionally, spectral lines and continued glitching in a given frequency range reduces the sensitivity of searches for long-duration signals at those frequencies. Both noise sources have a notable effect on search sensitivity if not mitigated.

Non-Gaussian noise in the detector outputs that can be correlated with auxiliary signals that have negligible sensitivity to GWs can be used to create flags for noisy data; these flags can then be used in astrophysical searches to remove artefacts and improve sensitivity. With transient noise, the flags are used to identify time segments in which the data may contain glitches; spectral lines are recorded as frequencies, or narrow frequency bands, at which the detector sensitivity is reduced.

In this section, the impact of DQ is measured by its effect on the primary analyses of the LIGO-Virgo CBC and burst all-sky searches [18, 21]: the low-mass 'ihope' [102] and cWB [48] pipelines respectively. The ihope search is a coincidence-based analysis, where data from each detector is matched-filtered against a bank of binary inspiral template signals, producing a SNR time-series for each. Peaks in the SNR between multiple detectors are considered coincident if they are within a certain time (set by the maximum time it would take for a GW to travel between sites) and stellar component mass (of the matched filter template) separation. The cWB algorithm calculates a multi-detector statistic by clustering time-frequency pixels with significant energy that are coherent across the detector network. In both cases, the multi-detector events identified are then subject to a number of consistency tests before being considered detection candidates.

The background of each search is determined by shifting detector data in time before performing the search. These time shifts are much greater than the time taken for a GW to travel between sites, ensuring that any multi-detector events in these data cannot have been produced by a single astrophysical signal.

Although both searches require signal power in at least two detectors, strong glitches in a single detector coupled with Gaussian noise in others still contributed significantly to the search background during S6. DQ flags were highly effective in removing these noise artefacts from the analyses. The effect of a time-domain DQ flag can be described by its *deadtime*, the fractional amount of analysis time that has been vetoed; and *efficiency*, the fractional number of GW candidate events removed by a veto in this deadtime.

Additionally, the *use percentage* – the fraction of auxiliary channel glitches which veto a GW candidate event – allows a measure the strength of the correlation between the auxiliary and GW channel data. Flag performances are determined by their efficiency-to-deadtime ratio (EDR): random flagging and vetoing of data gives EDR $\simeq 1$, whereas effective removal of glitches gives a much higher value.

Each analysis chose to apply a unique set of DQ flags in order to minimise deadtime whilst maximising search sensitivity; for example the CBC search teams did not use a number of flags correlated with very short, high-frequency disturbances, as these do not trigger their search algorithm, while these were used in searches for unmodelled GW bursts.

We present the effect of three categories of veto on each of the above searches in terms of reduction in analysable time and removal of noise artefacts from the search backgrounds. Only brief category definitions are given, for full descriptions see [104].

Category 1 vetoes

The most egregious interferometer stability problems are flagged as category 1. These flags denote times during data taking when the instrument was not running under the designed configuration, and so should not be included in any analysis.

The Data Monitoring Tool (DMT) automatically identified certain problems in real time, including losses of cavity resonance, and errors in the $h(t)$ calibration. Additionally, scientists monitoring detector operation in the control room at each observatory manually flagged individual time segments from observed instrumental issues and errors.

All LIGO-Virgo search groups use category 1 vetoes to down-select analysis segments; as a result their primary effect is in the reduction in analysable time over which searches are performed. This impact is magnified by search requirements on the duration for analysed segments, with the cWB and ihope

	Absolute deadtime % (seconds)		Search deadtime % (seconds)	
Observatory	cWB	ihope	cWB	ihope
H1	0.3% (53318)	0.4% (176079)	0.4% (77617)	3.8% (786284)
L1	0.4% (75016)	0.1% (20915)	0.7% (137115)	6.2% (1180976)

Table 4.2: Summary of the reduction in all time and analysable time by category 1 veto segments during S6

searches requiring a minimum of 316 and 2064 seconds of contiguous data respectively to accurately measure the detector PSD. Table 4.2 outlines the absolute deadtime (fraction of science-quality data removed) and the search deadtime (fractional reduction in analysable time after category 1 vetoes and segment selection). At both sites the amount of science-quality time flagged as category 1 is less that half of one percent, highlighting the stability of the instrument and its calibration. However, the deadtime introduced by segment selection is significantly higher, especially for the CBC analysis. The long segment duration requirement imposed by the ihope pipeline results in an order of magnitude increase in search deadtime relative to absolute deadtime.

Categories 2 and 3

The higher category flags are used to identify events produced by the search pipelines as likely noise artefacts. Both pipelines produce a first set of candidate event triggers after application of category 2 vetoes, and a reduced set after application of category 3.

The majority of category 2 veto segments are also generated in real time by the DMT and include photodiode saturations; digital overflows; and high seismic, or other environmental, noise. These are all generated from data signals whose correlation with the GW readout has been firmly demonstrated by instrumental commissioning and investigations.

Category 3 includes a set of DMT flags identifying lower-level couplings, as well as a large set of flags generated by statistical correlation algorithms. The HVeto [123] and UPV [124] algorithms are used, by the Bursts and CBC groups respectively, to identify coupling between auxiliary data and the GW readout.

Table 4.3 gives the absolute, relative, and cumulative deadtimes of these categories after category 1 and segment selection, outlining the amount of analysed time during which event triggers were removed. Similarly to category 1, category 2 vetoes have deadtime $\mathcal{O}(1)\%$, but with significantly higher appli-

cation at L1 compared to H1. This is largely due to the relative abundance of short science-quality segments for L1 and the application of a DQ flag removing the last 30 seconds before lock-loss. Additionally photodiode saturations and computational timing errors were more prevalent at the LLO site than at LHO and so contribute to higher relative deadtime.

Category 3 flags contributed $\mathcal{O}(10)\%$ deadtime for each instrument. Both statistical algorithms employed are designed to only output those veto segments that are highly efficient; HVeto was restricted to those correlations producing vetoes with EDR > 10 on its training set, for example. As a result, the relatively high deadtime for this category is acceptable to the search groups.

Deadtime type	Cat.	H1		L1	
		cWB	ihope	cWB	ihope
Absolute % (s)	2	0.26%	0.77%	1.59%	1.53%
	3	7.90%	9.26%	8.54%	7.03%
Relative % (s)	3	7.73%	9.00%	7.06%	6.10%
Cumulative % (s)	3	7.97%	9.71%	8.54%	7.54%

Table 4.3: Summary of the absolute, relative, and cumulative deadtimes introduced by category 2 and 3 veto segments during S6. The relative deadtime is the additional time removed by category 3 not vetoed by category 2, and cumulative deadtime gives the total time removed from the analysis.

Figure 4.11 shows the effect of category 3 vetoes on the background events from the cWB pipeline; these events are those single detector events that were identified in the background from time-slides. This search applies category 2 vetoes in memory, and does not record any events before this step, so efficiency statements are only available for category 3. The results are shown after the application of a number of network- and signal-consistency checks that reject a large number of the loud events, meaning the background is dominated by low SNR events, with a small number of loud outliers. At both sites, DQ vetoes have cumulative EDR ≥ 5 at SNR 3, with those at L1 removing the tail above SNR 20. However, despite the reduction, this search was still severely limited by those events in the multi-detector background distribution, identified as from the middle of the single-detector distributions shown above [21].

Figure 4.12 shows the effect of category 2 and 3 vetoes on the background from the ihope pipeline; this search sees a background extending to higher SNR. As shown, the background is highly suppressed by DQ vetoes, with an efficiency of 50% above SNR 8, and 80% above ~ 100 at both sites. The re-weighted SNR statistic, as defined in [102], is highly effective in down-ranking
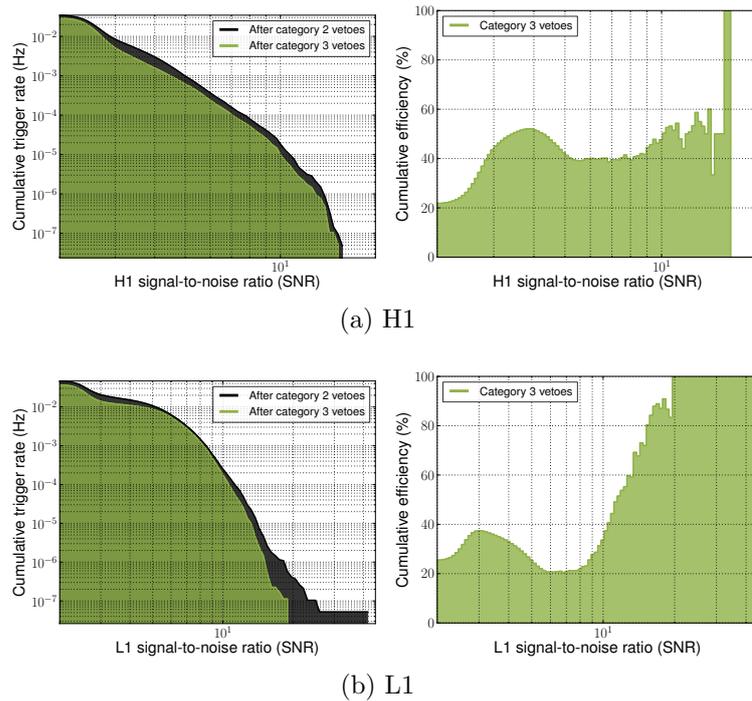
Figure 4.11: The effect of category 3 vetoes on the cWB pipeline for (a) H1 and (b) L1. The left panels show the reduction in event rate, while the right panels show the cumulative veto efficiency, both as a function of single-detector SNR.

the majority of outliers with high matched-filter SNR, but a non-Gaussian tail is still present at both sites. Category 3 vetoes successfully remove this tail, reducing the loudest event at H1 (L1) from 16.0 (15.3) to 11.1 (11.2). Search sensitive distance is roughly inversely proportional to the χ^2-weighted SNR of the loudest event, and so reducing the loudest event by $\sim 30\%$ with $\sim 10\%$ deadtime can be estimated as a factor of ~ 2.5 increase in detectable event rate.

The above figures, and their descriptions, show the considerable effect LIGO data quality flags have on the searches. Coupled with a number of signal-based consistency tests developed by the analysis groups, the noise background is greatly reduced.

4.6 Conclusions and outlook for Advanced LIGO

The LIGO instruments, at both Hanford and Livingston, are regularly effected by both non-Gaussian noise transients and long-duration spectral features. Throughout S6 a number of problems were identified as detrimental to stable and sensitive data-taking at the observatories, as well as the astrophysical

(a) H1

(b) L1
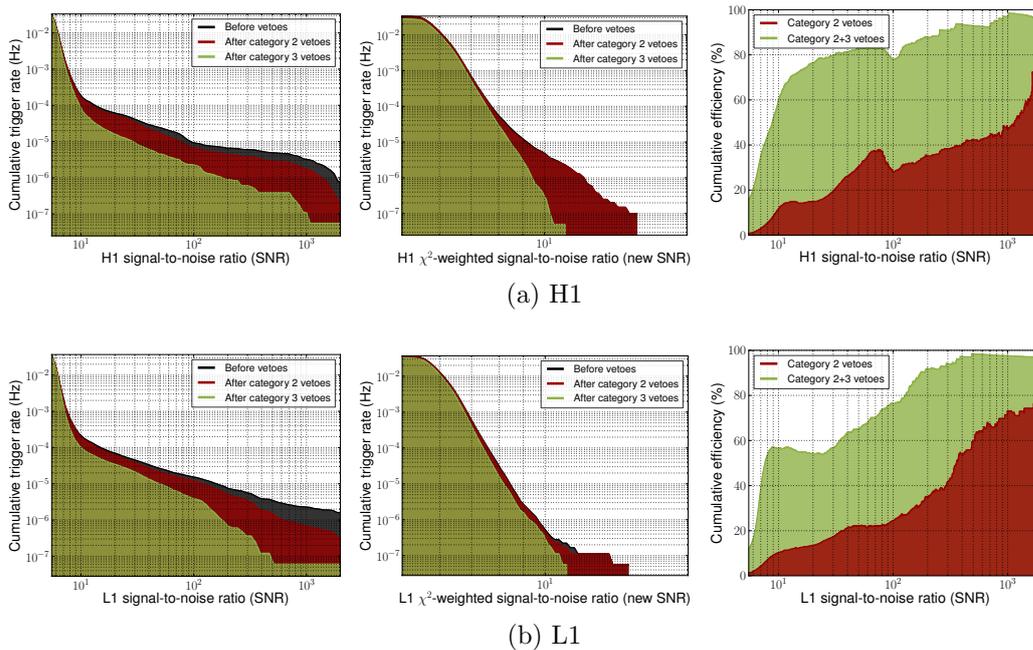
Figure 4.12: The effect of category 2 and 3 vetoes on the CBC ihope pipeline for (a) H1 and (b) L1. The left panels show the reduction in event rate as a function of SNR, the centre panels show the reduction in event rate as a function of the χ^2-weighted SNR, and the right panels show the cumulative efficiency as a function of SNR.

searches performed on the data. These issues have been presented, including their instrumental causes and the measures taken to mitigate them, with their impact on data quality the key factor in each study.

Instrumental fixes and general improvements, employed throughout the science run, resulted in an increasingly stable and sensitive instrument. Median segment duration and overall duty cycle improved from epoch to epoch (table 4.1) and the detection range to the canonical binary neutron star increased by a significant factor (figure 4.4). Data quality flags, used to identify known correlations between noise in auxiliary systems and the GW data, allowed for a significant reduction in the event background of both core transient searches, ihope and cWB (figures 4.11 and 4.12). An EDR above 5 for both searches, at both sites, allowed for a significant increase in the sensitivity of the search, improving the upper limits on event rate for both astrophysical source types.

However, a tail of high SNR events was still present in the cWB search for bursts of GWs, requiring deeper study of their morphology and improved identification methods. A large number of these remaining transient noise sources are still undiagnosed, meaning a large effort must be undertaken to mitigate similar effects in the second-generation instruments.

The first-generation LIGO instruments were decommissioned shortly following the end of S6, with installation and early testing of aLIGO systems under way [89]. With the next data-taking run scheduled for 2015 [64], many methods and tools developed during the last run are set to be upgraded to further improve instrument and data quality. Improvements are in place for each of the event trigger generators, allowing for more accurate detection of transient noise in all channels, and work is ongoing for the HVeto and UPV statistical veto generators, enabling more efficient identification of noise in the GW data. In addition, multi-variate statistical classifiers are being developed for use in glitch identification, using more information produced by the event trigger generators (ETGs) to improve veto efficiency and identification of false-alarms with minimal deadtime.

One of the major goals of the aLIGO project is multi-messenger astronomy – the collaboration between GW observatories and EM telescopes [128, 129]. With this in mind, a large part of the development in detector characterisation in the LIGO Scientific Collaboration (LSC) is toward on-line analysis of instrumental data, including the GW output and all auxiliary channels. Both the burst and CBC search working groups are developing low-latency analyses from which to trigger followup with partner EM telescopes, requiring a much greater effort in low-latency characterisation of the data. An Online Detector Characterisation (ODC) system is being developed for aLIGO in which the status of all instrumental and environmental systems are recorded in real-time to allow fast identification of false alarms in these on-line analyses, and reduce the latency of follow-up requests.

A great effort will be required in commissioning these new instruments, where a strong understanding of the noise sources will be required before the first advanced science run. Estimates predict ~ 40 detections of GWs from binary neutron stars per year at design sensitivity [52], and so fast, accurate, and efficient data quality information and study will be essential to achieve this goal.

Chapter 5

Reducing the effect of seismic noise in LIGO searches by targeted veto generation

The primary output of each LIGO detector is a single data stream that in general contains some combination of a GW signal and detector noise. Glitches can mask or mimic astrophysical signals, thus limiting the sensitivity of any search that can be performed over these data [103, 105, 26]. In the searches for short-duration GW signals the noise background is dominated by glitches, requiring intense effort from analysis groups and detector scientists to understand the physical origins and eliminate them. Throughout the lifetime of LIGO up to and including S6, search sensitivity has been improved by careful use of *vetoes*. Vetoes allow analysts to tune and operate search pipelines using a subset of cleaner data, increasing the chance of extracting a signal from the noise [104].

The detrimental effect of seismic noise has been known to be a key limiting factor to the sensitivity of GW detectors at low frequencies (below a hundred Hz [125]). However it is also a common cause of glitches at higher frequencies due to non-linear coupling of low-frequency seismic noise into the gravitational wave readout. Previous methods to generate veto segments for times of high seismic noise have proven ineffective. In this chapter, we introduce a new method of constructing vetoes for the specific case of seismic noise that has proven highly effective when used in the latest searches for transient GW signals. We find a large statistical correlation between *triggers* from the low-latency CBC search and seismic noise, vetoing 60% of all triggers in 6% of time for H1 and 6% of triggers in 0.6% of time for L1.

Frequency (Hz)	Distance (km)	Source
$0.01 - 1$	10^3	Distant earthquakes Microseism
$1 - 3$	10^1	Far anthropogenic noise Close earthquakes Wind
$3 - 10$	10^0	Anthropogenic noise Wind
$10 - 100$	10^{-1}	Close anthropogenic noise

Table 5.1: Description of the main seismic frequency bands and their sources

The chapter is set out as follows. Section 5.1 describes the seismic environment at each of the LIGO sites, and the effect it has on detector sensitivity. In section 5.2 we outline existing veto methods used in S5 and S6. In section 5.3 we describe our new method for identifying and vetoing noise in seismometer data. In section 5.4 we present the results in terms of veto efficiency and deadtime. Finally, section 5.5 presents a brief discussion of implications and further applications of the method.

5.1 Seismic noise in LIGO

5.1.1 LIGO seismic environment

The two LIGO sites were chosen to be far from urbanised areas, thus reducing the incident seismic noise, whilst their separation provides a long baseline helpful in sky-localisation of astrophysical signals [130]. The various types of seismic noise to which they are subject, as characterised by their source, can be separated into the four frequency bands given in Table 5.1, and the variability of noise during evenings and weekends relative to standard working hours is shown in Figure 5.1. The strain sensitivity of the two LIGO detectors is shown in figure 4.3, with seismic noise the limiting factor below 50 Hz.

LHO is located 15 km from the United States Department of Energy (US-DOE) Hanford Site, in which several working areas include use of heavy earth-moving machinery, and the Tri-Cities area begins roughly 20 km away, both contributing heavily to the amount of anthropogenic seismic noise incident on the detector. The site is also subject to high winds up to 40 m/s, causing motion of the buildings and the concrete slabs supporting the instruments. These relatively local sources generate noise in the higher bands in Table 5.1, above
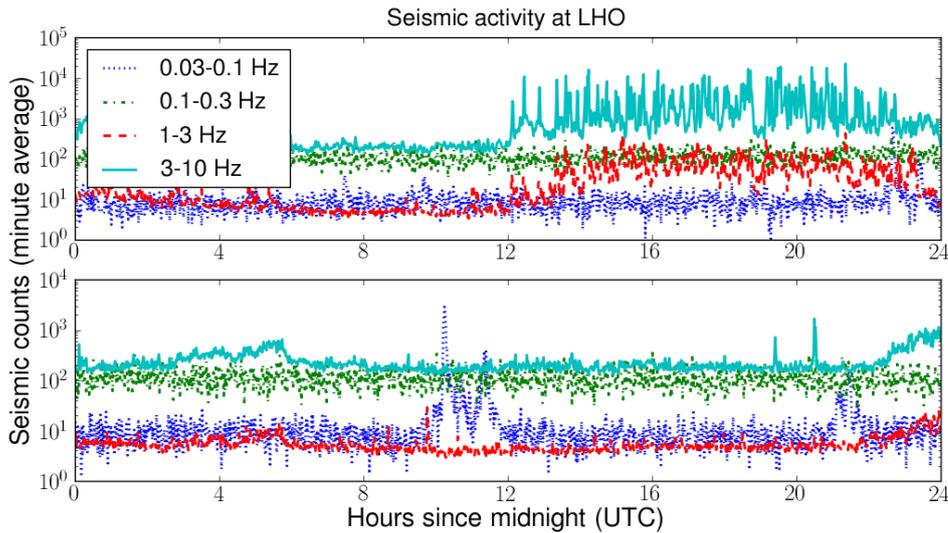
Figure 5.1: Examples of the difference between seismic-induced acceleration incident at the LHO site during a full 24-hour span on a weekday (top) and a weekend (bottom), as measured by a seismometer. The higher frequency bands show elevated noise on a weekday between 12:00–24:00 UTC (05:00-17:00 local time) from working-day traffic.

$\sim 1\,\text{Hz}$.

LIGO Livingston Observatory (LLO) is located 7 km from the town of Livingston, and only 3 km from a railway line used daily by cargo trains [125], plus, the land surrounding LLO is used for timber harvesting. The site is only 130 km from the Gulf of Mexico, and is subject to violent rain and windstorms.

Both sites are also subject to noise from earthquakes occurring almost anywhere on Earth, and to microseismic noise from oceanic activity due to their relatively short separation from the nearest coastline. These distant events are sources of noise below $\sim 1\,\text{Hz}$.

Due to the softer composition of the surrounding geological landscape, seismic noise was worse at LLO during early science runs, so the decision was taken to install an active seismic isolation system on L1 before S4. The hydraulic external pre-isolation (HEPI) feed-forward system [117] damps low-frequency noise by using signals from the onsite seismometers to control movement of the vacuum chambers for the end test masses. This particular system was not installed at LHO before S6 – although other isolation systems were used – but will be installed as part of the aLIGO project [82].

The LIGO instruments are designed to be sensitive in the range $40 - 7000\,\text{Hz}$ [13], so one may be forgiven for assuming that seismic noise below 30 Hz should not affect sensitivity in the detection band. However, *upconversion* has been a

problem throughout each science run, caused by a number of factors related to ground motion, for example scattered light [131] and the Barkhausen effect (see chapter 4). This effect contaminates the sensitive band of the LIGO detectors, meaning seismic noise is an even greater problem than it would be otherwise.

5.2 Existing veto methods

As described in chapter 4, the GW data stream is not the only information drawn from the LIGO detectors. Thousands of auxiliary data channels are recorded, containing control and error signals from instrumental systems, and measurements from the physical environment. These data are analysed in order to study and improve detector performance, but also to identify and remove glitches that can mimic GWs. Veto segments can be constructed around excess noise if it is known to couple into the detection channel.

In S5 and S6, vetoes were constructed by two methods [103, 26, 104, 124, 123]. The first method relied on known physical couplings between an auxiliary subsystem and the GW data, whereby when a correlation is understood, the time stream of a particular auxiliary channel is analysed, and times for which a certain threshold was exceeded are recorded. Simple, but highly effective examples include overflows in analogue-to-digital converters, and light dips.

The second method replaces known couplings with statistics, applying the Kleine-Welle (KW) wavelet-based algorithm [122] to data in auxiliary channels with negligible sensitivity to GWs, producing lists of triggers. These events are then tested for time-coincidence with triggers in the GW data, indicating whether that candidate GW event was likely to be of astrophysical origin. Vetoes are constructed around a subset of triggers in the auxiliary data chosen to maximise *efficiency* whilst minimising *deadtime*.

In both methods, highly effective vetoes are those with a high ratio of efficiency to deadtime. The HVeto and UPV algorithms [124, 123], as mentioned in chapter 4, were used in searches for unmodelled bursts [132] and CBCs [133] during S5, producing comparable results [103].

5.3 Targeted veto methods

The methods described in the previous section are subject to shortcomings when applied to seismometer data. Simple thresholds have to be placed high enough to catch only the worst noise spikes, and so have low efficiency over

weekends and evenings (around lower seismic noise as shown in Figure 5.1). Similarly, the KW algorithm is tuned for high-frequency gravitational-wave bursts (GWBs), with limited sensitivity to the low-frequency signature of seismic noise, resulting in low trigger numbers and poor statistical significance of coincidences.

In order to produce effective vetoes, we have devised a novel method to explicitly identify low-frequency seismic events, and construct veto segments to remove this noise from GW searches. This method uses the *Ω-pipeline* tuned specifically for low-frequency performance to generate trigger lists highlighting seismic events, and the low-latency inspiral pipeline *Daily iHope* to generate triggers from CBC template matched-filtering. The two are combined by the HVeto algorithm into lists of time segments during which seismic noise has polluted the GW analysis.

5.3.1   The Ω-pipeline

The *Ω-Pipeline* is a burst detection algorithm developed within LIGO as a combination of the *Q Pipeline* [127] and *X-Pipeline* [100] and used, during S6, for low-latency detection of GW events to trigger electromagnetic followup [128]. The single-detector triggers were also used for DQ investigations.

The algorithm is based on the Q Transform [134] and projects detector data, $s(t)$, onto a bank of windowed complex exponentials of the following form:

$$S(\tau, f, Q) = \int_{-\infty}^{\infty} s(t)w(t - \tau, f, Q) \exp(-i2\pi f t) \, dt, \qquad (5.1)$$

where w is a time-domain window centred on τ, f is the central frequency, and Q is the quality factor. An example of the output of the Ω-Pipeline applied to low-latency gravitational wave readout data is show in Figure 5.2. A high density of low SNR (white) triggers is expected from Gaussian noise, but the higher SNR events (black) indicate increased noise at low frequencies, known to be correlated with the high seismic activity shown in Figure 5.1.

As a result of using both KW and Ω-Pipeline for DQ studies, direct comparisons were drawn on the performance of each, especially in frequency reconstruction at low frequency. It was found that even before tuning, the Ω-Pipeline gave much greater low frequency sensitivity, and better frequency resolution.
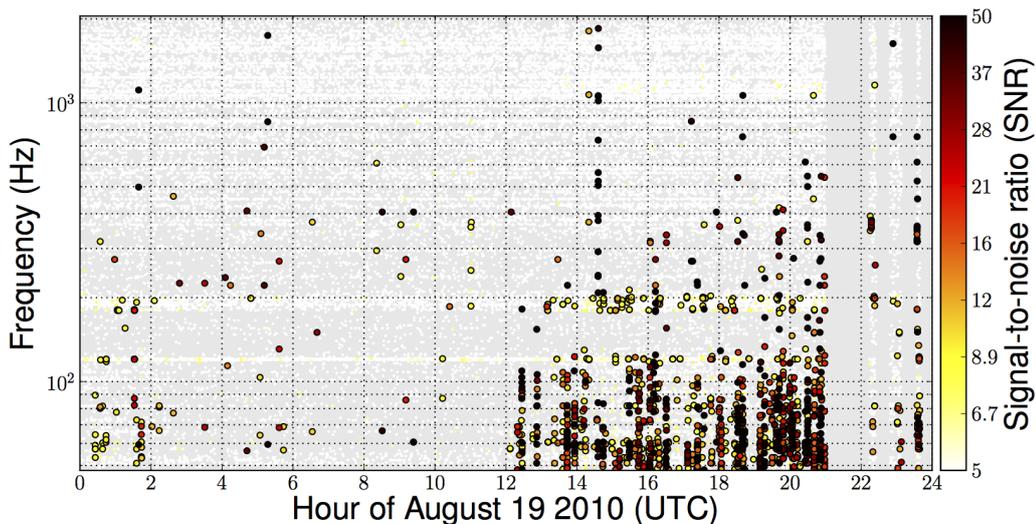
Figure 5.2: A 24-hour time-frequency map from the Ω-pipeline of data from the GW readout of the H1 detector for the same period as figure 5.1 (top). 12:00 UTC (05:00 local time) onwards are subject to increased noise at the lowest frequencies due to excess seismic noise from the working day. The stripes with no events at the end of the period indicate that the detector was not operational.

5.3.2 Parameterisation of the Ω-Pipeline for seismic noise

The low-latency Ω-Pipeline analysis used a parameter set tuned for performance in the detection band, with a frequency range of 48–2048 Hz, and analyses performed using 64-second-long data segments to estimate the background noise spectrum. As can be seen in Figure 5.2, the frequency range is such that the seismic band is almost completely ignored. However, significant SNR is recorded up to around 200 Hz that can be attributed, by time-coincidence, to seismic noise upconversion.

In order to improve performance when applied to seismometer data, the parameter space was split into two sets: the anthropogenic band, above 2 Hz, and very low frequency seismic activity, below 2 Hz. The following paragraphs detail the changes made to tune the Ω-Pipeline algorithm for each frequency band, describing three key parameters. The *sampling rate* defines the highest frequency (half the sample rate) of the data to be filtered, and the *frequency range* gives the complete span of frequencies searched. The *block duration* defines the length of period used to estimate the PSD of the detector, set as a power of 2 (for ease of computation). In order that a small number of loud events do not affect the measurement of the PSD, we use a duration significantly longer than the longest resolvable events, and use a median-mean

Parameter	Untuned Value	Tuned value	
		$< 2\,\mathrm{Hz}$	$> 2\,\mathrm{Hz}$
Sample frequency	$4096\,\mathrm{Hz}$	$4\,\mathrm{Hz}$	$64\,\mathrm{Hz}$
Frequency range	$48 - 2048\,\mathrm{Hz}$	$0.01 - 2\,\mathrm{Hz}$	$2 - 32\,\mathrm{Hz}$
Block duration	$64\,\mathrm{s}$	$65536\,\mathrm{s}$	$4096\,\mathrm{s}$

Table 5.2: The parameter sets applied to the Ω-Pipeline search algorithm before and after tuning for low-frequency performance.

average method reduce sensitivity to transient noise.

The anthropogenic band, $> 2\,\mathrm{Hz}$

As described in Section 5.1 the seismic band extends upwards in frequency to a few tens of Hz. Short, high-frequency events may corrupt the calculation of the background around them for a longer time, shadowing lower-frequency, lower-amplitude events. Lowering the sampling frequency to $64\,\mathrm{Hz}$[1] filtered out any high frequency seismic noise, allowing longer time-scale events to be triggered by the search. The power spectrum was drawn from blocks of $4096\,\mathrm{s}$, meaning a number of discrete seismic events above $2\,\mathrm{Hz}$ could be individually resolved above the background.

The earthquake band, $< 2\,\mathrm{Hz}$

This band was chosen specifically to target long-distance earthquakes, that, as described in Table 5.1, add noise down to $0.01\,\mathrm{Hz}$ for up to several hours. Here the sampling frequency could be reduced to $4\,\mathrm{Hz}$, eliminating higher frequency disturbances, with a minimum frequency of $0.01\,\mathrm{Hz}$, while blocks of 65536 seconds were used to allow accurate PSD estimation in the presence of hour-long earthquake events.

This method was applied to the four main seismometers at LHO: EX, EY, LVEA and VAULT[2]; and the three at LLO: EX, EY and LVEA (LLO has no VAULT seismometer). As can be seen in Figure 5.3b, the new parameter sets allows a huge increase in the number of triggers produced by the Ω-Pipeline.

---

[1]The Ω-Pipeline search for GWs downsamples the readout data to $4096\,\mathrm{Hz}$, while the seismometers are only sampled at $256\,\mathrm{Hz}$

[2]The EX and EY seismometers sit outside of the vacuum chambers containing the end test masses for the X- and Y-arms respectively, the Large Vacuum Equipment Area (LVEA) seismometer sits beside the chamber housing the GW readout photodetector, and the VAULT seismometer is in an underground chamber a small distance from the LVEA.
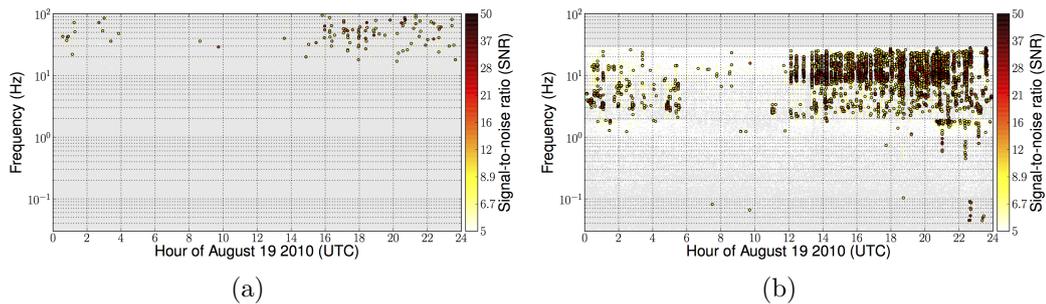
Figure 5.3: Examples of the untuned KW and tuned the Ω-pipeline algorithms applied to seismometer data. The left figure, (a), shows the lack of sensitivity in the untuned KW analysis, while that on the right, (b), has many orders of magnitude more events. Comparing to figure 5.2 we can see loud triggers around 10 Hz after 12:00 UTC, but also triggers with SNR above 10 below 0.1 Hz.

The density of triggers has be greatly increased, especially around noisier times, with events recorded with frequencies as low as 0.03 Hz.

5.3.3 Low-latency inspiral triggers – *Daily iHope*

The joint LIGO-Virgo CBC group uses the *iHope* pipeline to search for GWs produced by binary coalescences. It is described more fully in [133, 135]. Seismic noise has been known to contribute significantly to the noise background estimates in these searches, and so creating good vetoes specifically for CBC searches was a major goal of this work. Here we summarize the key points and discuss the changes made for daily running of iHope in order to provide triggers to analyse alongside the Ω-Pipeline triggers from seismic data.

Daily iHope is a templated, matched-filter search using restricted, stationary phase, frequency-domain waveforms of the form

$$\tilde{h}(f; M, \eta) = \frac{2GM_\odot}{c^2 r} \left(\frac{5M\eta}{96M_\odot} \right)^{\frac{1}{2}} \left(\frac{M}{\pi^2 M_\odot} \right)^{\frac{1}{3}} f^{-\frac{7}{6}} \left(\frac{GM_\odot}{c^3} \right)^{-\frac{1}{6}} e^{i\Psi(f; M, \eta)}, \quad (5.2)$$

where $M = m_1 + m_2$ is the total mass of the binary, and $\eta = m_1 m_2 / M^2$ the symmetric mass ratio. A static bank of such templates spanning the total mass range from $2\,M_\odot - 25\,M_\odot$ was used for each interferometer, based on the layout at a quiet time in each instrument, with a minimal match of 0.95 for the region above a chirp mass $(M\eta^{3/5})$ of $3.46\,M_\odot$ and 0.5 below that. This distribution would not be good enough for an astrophysical search, but was shown to be adequate for identifying short-duration glitches that match
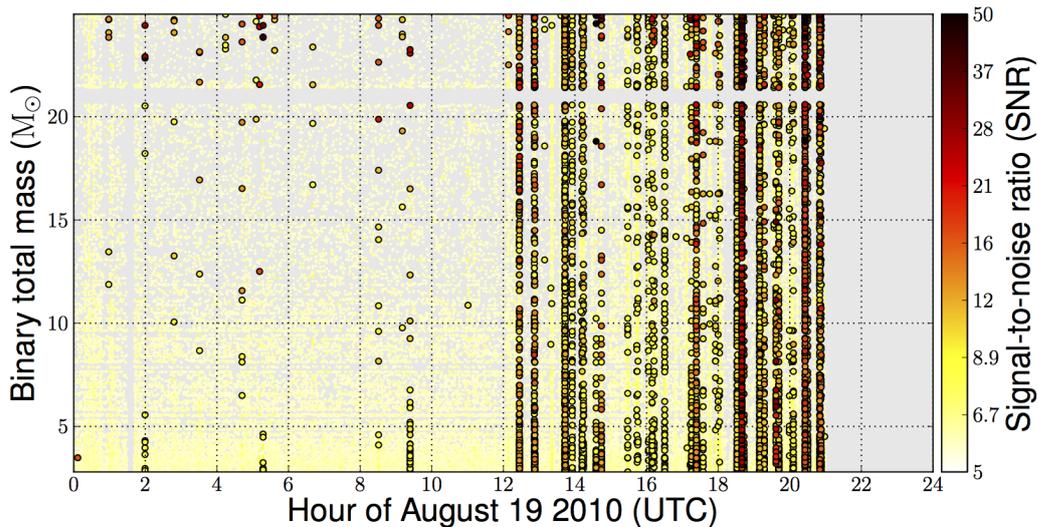
Figure 5.4: A 24-hour template mass versus time map of triggers from the Daily iHope pipeline. Comparing to figures 5.1 and 5.2, we see the same excess of triggers after 12:00 UTC, but here the excess noise at low frequency results in triggers across the entire template bank. No data is analysed after 21:00 UTC because ihope requires at least 2048 contiguous seconds to estimate the PSD and all data after this time was in smaller segments.

the higher-mass (shorter) templates better. This allowed for short-duration glitches, corresponding to higher-mass inspiral templates, to be flagged with large SNR. An example of the output of Daily iHope is shown in Figure 5.4.

5.3.4   Veto generation – HierarchichalVeto

The seismic triggers from the Ω-Pipeline, and the CBC triggers from Daily iHope were used to idenfity times of seismic noise using the statistical algorithm HVeto [123].

The HVeto algorithm tests the statistical significance of time-coincidence between triggers from one channel, nominally the GW data channel, and those from auxiliary channels. The significance statistic is defined as

$$S = -\log_{10}\left(\sum_{k=n}^{\infty} P(\mu, k)\right),\tag{5.3}$$

where n is the number of coincident events, μ is the expected number of random coincidences given the trigger rates in the two channels, and P is the Poisson probability distribution function. The significance is calculated for all

channels in a two-dimensional space of time-coincidence window, T_{win}[3], and SNR threshold, ρ_c.

The most significant point on the $(T_{\mathrm{win}}, \rho_c)$ plane is chosen for each auxiliary channel, with the loudest channel by significance selected. Veto times are constructed by generating segments of width T_{win} around all triggers with SNR above ρ_c in that auxiliary channel. These segments are then removed from the analysis – allowing the next round to be 'won' by a (generally) different auxiliary channel containing less significant coincidences – and the procedure repeated until the significance of the loudest channel does not exceed a given stopping point. In this way, vetoes are generated hierarchichally, allowing for little redundancy between different channels.

Several modifications were made to this algorithm in order to test and run on the new seismic triggers. Testing was completed in order to construct a new $(T_{\mathrm{win}}, \rho_c)$ plane relevant for the long-duration events from the seismic data, spanning $0.1 <= T_{\mathrm{win}} < 10\,\mathrm{s}$ and SNRs 10–300. Alongside this, as described in the caption to Figure 5.4, modifications were made to first read and understand the new Daily iHope triggers, and use the relevant new parameters.

5.4   Results

The method described above was used in the construction of the LIGO seismic veto dubbed SeisVeto for S6. Here we present the results for a test sample of data, spanning June 26 – August 6 2010. The results for LHO are shown in Table 5.3a and those for LLO in Table 5.3b with each row giving the statistics for the most significant channel in each frequency band, in addition to the cumulative results for the entire period[4].

For H1, the rounds contribute to give a cumulative efficiency of 62.5 %, with a cumulative deadtime of 5.94 %. This means that almost two thirds of all triggers produced by the low-latency inspiral pipeline are occurring in a small amount of time, which is coincident with high seismic noise. This statistic alone outlines the problem caused by seismic noise.

It should not be surprising that the most significant channel for the two higher frequency bands should be the LVEA seismometer. This building is

---

[3]Low-frequency events have a long duration whose maximum coupling time is not known. Also, many short-duration noise events in an auxiliary system can take a certain time to couple into the GW output.

[4]The cumulative results include all rounds passing selection criteria in each band, not just the most significant channel.

Freq. Band (Hz)	Loudest Channel	Significance	Efficiency (%)	Deadtime (%)
0-1	EX	1455.21	3.26	0.15
1-3	EY	355.37	3.19	0.71
3-10	LVEA	12024.98	22.11	1.24
10-32	LVEA	41042.78	35.99	1.04
Cumulative, all rounds			62.44	5.94

(a) Results for H1.

Freq. Band (Hz)	Loudest Channel	Significance	Efficiency (%)	Deadtime (%)
0-1	No channels passing selection criteria		0	0
1-3	LVEA	960.13	1.51	0.06
3-10	LVEA	420.55	0.88	0.06
10-32	EX	1601.22	2.29	0.07
Cumulative, all rounds			6.95	0.60

(b) Results for L1.

Table 5.3: HVeto results from coincidence between new Ω-Pipeline triggers from seismometer data and low-latency inspiral triggers for the LIGO detectors operating in S6. Shown are the most significant (loudest) channels for each frequency band, and the cumulative statistics for the entire analysis.

closer than any other to a major road, so experiences the highest magnitude of seismic noise from traffic and close anthropogenic noise, especially trucks serving the USDOE Hanford site, and also houses the majority of interferometer control optics and subsystems, notably the GW readout photodetector.

For L1 we can see much lower statistical significance of the correlation between seismic noise and the readout signal. This can be attributed in part to the improvements from the HEPI feedforward system for the Livingston instrument, but also to the different nature of the seismic environment relative to LHO. However, despite a lower efficiency, the ratio of efficiency to deadtime is still above 10, highlighting the statistical correlation between seismic noise and low-latency CBC triggers.

5.5   Summary

In this chapter, we have highlighted the problem of transient seismic noise in GW detection, and presented a new method to not only identify, but remove, times of high noise from short-duration GW searches. We have demonstrated a highly effective veto, with large efficiency-to-deadtime ratio, that has been crucial in removing the worst of the transient detector noise whilst leaving as

much searchable time as possible. This method was applied to the searches for GWB and CBC signals in the final part of S6, and was seen to have a dramatic effect on the background, [106].

Advanced LIGO (aLIGO) is a major upgrade program that will see the sensitive distance of the LIGO detectors increase by a factor of 10, giving a factor of 1000 in sensitive volume. This should mean regular detections of GW transients from CBC events [52]. However, it is likely that there will still be non-stationarities in the data from seismic events, and other sources. The method introduced here will allow us to remove them, increasing search sensitivity, and also gives a highly tuned means of directing site scientists to coupling noise sources in a newly commissioned machine.

Seismic noise was chosen as an obvious starting point, given the prevalence of glitches of seismic origin, and a prior lack of an effective veto method. This method can be generalised to any and all susbsystems of the next generation of interferometers by tuning a GW burst detection algorithm on the appropriate data channels and has the potential to lead to a great increase in search sensitivity as a result of the above benefits.

Part III

Searches for signals from
compact binary coalescences

Chapter 6

The multi-detector search for binary inspiral signals

This chapter describes the methodology of a search for gravitational waves from CBCs using the data from a network of detectors. As the discussion of part II showed, the noise background for a single detector includes many transient bursts with high amplitude, and so any analysis must correlate the data between multiple detectors to detect a common signal.

We start with discussions of the waveforms used in searches, and the optimal filter used to extract these signals from noisy detector data, and continue to outline the data processing pipeline 'ihope' [102] from applying this filter to the data with a bank of possible signal waveforms, through matching the output of different detectors, to a possible detection of GWs. We also present a summary of the analysis as performed on S6/VSR2/3 data, and the results obtained.

6.1   The binary inspiral waveform

In chapter 2 we introduced the binary inspiral waveform, using the model of two point masses orbiting around a common centre-of-mass; at this point, we can further simplify the functional form of the waveform.

6.1.1 Waveform evolution

To this end, we can determine how the radiation frequency will evolve over time; recall the mass quadrupole moment for the binary system (2.6),

$$M^{ij} = \mu x^i x^j. \tag{6.1}$$

We can calculate the third differential components of this moment as

$$\dddot{M}^{11} = -\dddot{M}^{22} = 4\mu r^2 \omega^3 \sin\left(\int_0^t \omega(t')\,dt'\right) \tag{6.2a}$$

$$\dddot{M}^{12} = \dddot{M}^{21} = 4\mu r^2 \omega^3 \cos\left(\int_0^t \omega(t')\,dt'\right) \tag{6.2b}$$

Now, inserting these into the expression for GW power (1.42), we calculate the power carried with the binary inspiral as

$$P = \frac{32}{5c^5}\left(G^4 \mathcal{M}_c^{10}\omega^{10}\right)^{1/3}. \tag{6.3}$$

The total energy of the system is simply calculated as

$$E_{\text{orbit}} = -\frac{m_1 m_2}{2r} = -\left(\frac{\mathcal{M}_c^5 \omega}{8G}\right)^{1/3}, \tag{6.4}$$

and so we can equate its derivative with the above power to see the evolution of the orbital frequency

$$\dot{\omega} = \frac{96}{5}\left(\frac{G^5 \mathcal{M}_c^5 \omega^{11}}{c^3}\right)^{1/3}. \tag{6.5}$$

The emission frequency then envolves as

$$\dot{f} = \frac{\dot{\omega}}{\pi} = \frac{96}{5}\left(\frac{G^5 \mathcal{M}_c^5 f^{11}}{c^3}\right)^{1/3}. \tag{6.6}$$

and so

$$f = \left(\frac{5}{256}\right)^{3/8}\frac{1}{\pi}\left(\frac{G\mathcal{M}_c}{c^3}\right)^{-5/8}\tau^{-3/8} \tag{6.7}$$

where we define the *coalescence time*, t_c, from the integration constant as that time when $f \to \infty$, and introduce

$$\tau = t_c - t \tag{6.8}$$

as the time to coalescence. The phase can also be re-expressed in terms of the coalescence time as

$$\Phi(\tau) = 2\pi \left(\frac{5G\mathcal{M}_c}{c^3}\right)^{-5/8} [\tau - t_c]^{5/8} \tag{6.9}$$

We can now write the polarisation waveforms (2.12) using their evolutions

$$h_+(\tau) = A(\tau)f(\tau)^{2/3} \left(\frac{1 + \cos^2 \iota}{2}\right) \cos(\Phi(\tau) + 2\psi) \tag{6.10a}$$

$$h_+(\tau) = A(\tau)f(\tau)^{2/3} \cos \iota \, \cos(\Phi(\tau) + 2\psi). \tag{6.10b}$$

6.1.2   Directional response of a laser interferometer

The waveforms above give the amplitude in a direct line of sight to the observer. However, the amplitude sensitivity of the GW detectors discussed in part II is sensitive to the source sky position and polarisation of the incoming signal.

If the angles (θ, ϕ) give the sky location of the source in the detector frame, and ψ is the polarization angle, the detected signal $h(\tau)$ recorded by an interferometer is now a weighted sum of the two wave polarisations [136],

$$h(\tau) = F_+ h_+(\tau) + F_\times h_\times(\tau), \tag{6.11}$$

where

$$F_+(\theta, \phi, \psi) = -\frac{1}{2}[1 + \cos^2(\theta)] \cos(2\phi) \cos(2\psi) - \cos(\theta) \sin(2\phi) \sin(2\psi) \tag{6.12a}$$

$$F_\times(\theta, \phi, \psi) = \frac{1}{2}[1 + \cos^2(\theta)] \cos(2\phi) \sin(2\psi) + \cos(\theta) \sin(2\phi) \cos(2\psi). \tag{6.12b}$$

6.1.3   Detection waveforms

Combining equations (6.12) with equations (6.10), it is possible to simplify the detection waveform (6.11) to

$$h(\tau) = A \, \mathcal{M}_c^{5/3} f^{5/3}(\tau) \cos(\Phi(\mathcal{M}_c, \tau) + \Phi_0), \tag{6.13}$$

where A and Φ_0 are functions of the extrinsic parameters of the waveform. From this, we can then separate the waveform into two components,

$$h(\tau) = h_0(\tau) \cos \Phi_0 + h_{\pi/2}(\tau) \sin \Phi_0, \tag{6.14}$$

where the two phases of the waveform are,

$$h_0(\tau) = A \, \mathcal{M}_c^{5/3} f^{5/3}(\tau) \cos\left(\Phi(\tau)\right) \tag{6.15a}$$

$$h_{\pi/2}(\tau) = A \, \mathcal{M}_c^{5/3} f^{5/3}(\tau) \sin\left(\Phi(\tau)\right). \tag{6.15b}$$

6.2 The matched filter

One of the most common problems in observational analysis is extracting a weak signal from noisy data – without a strong model it is difficult to distinguish signal energy from background noise, be that instrumental or environmental. However, knowledge of the signal can be used to construct a filter that will determine the amplitude of that model in the data.

In analysis of GW data, we can consider the generic detector output signal, $s(t)$, to be some combination of noise and signal,

$$s(t) = n(t) + h(t). \tag{6.16}$$

Here $n(t)$ is the detector noise, which we assume to be stationary and Gaussian with zero mean; and $h(t)$ is the signal for which we are searching, and for which we have a model.

We define a filter, $F(t)$, which we can apply to the data

$$\hat{s} = \int_{-\infty}^{\infty} F(t)\,s(t)\,dt = \int_{-\infty}^{\infty} \tilde{F}^*(f)\,\tilde{s}(f)\,df, \tag{6.17}$$

where we introduce the Fourier transform of a function $a(t)$ as

$$\tilde{a}(f) = \int_{-\infty}^{\infty} a(t)e^{-2\pi i f t}\,dt. \tag{6.18}$$

The challenge is to find the filter, F, that optimally extracts a signal from the noise. This is done by maximising the ratio H/N, where H is the expected value of $\hat{s}$ in the presence of only signal, and N the standard deviation in the

absence of signal.

For our stationary, Gaussian noise, the power spectral density (PSD) is defined by

$$\langle \tilde{n}(f)\,\tilde{n}(f')\rangle = \frac{1}{2}\delta\left(f - f'\right)S_n(f),$$ (6.19)

and so, in the presence of only noise we have,

$$
\begin{aligned}
\langle N^2 \rangle &= \left[\langle \hat{s}^2(t)\rangle - \langle \hat{s}(t)\rangle^2\right]_{h=0} \\
&= \langle \hat{s}^2(t)\rangle_{h=0} \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\tilde{F}^*(f)\,\tilde{F}(f')\,\langle \tilde{n}^*(f)\,\tilde{n}(f')\rangle\;df\,df' \\
&= \int_{-\infty}^{\infty} df\,|\tilde{F}(f)|^2\frac{1}{2}S_n(f).
\end{aligned}
$$ (6.20)

where in the first line we continue our assumption of the noise having zero mean. Then we have

$$\frac{H^2}{\langle N^2\rangle} = \frac{\left(\int_{-\infty}^{\infty}\tilde{F}^*(f)\,\tilde{h}(f)df\right)^2}{\int_{-\infty}^{\infty}df\,|\tilde{F}(f)|^2\frac{1}{2}S_n(f)}$$ (6.21)

In the numerator of this expression we can use the Cauchy-Schwartz inequality to see [137]

$$\left(\int_{-\infty}^{\infty}\tilde{F}^*(f)\,\tilde{h}(f)\,df\right)^2 = C\left(\int_{-\infty}^{\infty}|\tilde{F}^*(f)|^2\,S_n(f)\,df\right)\left(\int_{-\infty}^{\infty}|\tilde{h}(f)|^2/S_n(f)\,df\right)$$ (6.22)

where C is a constant, and so we can simplify equation (6.21) to

$$\frac{H^2}{\langle N^2\rangle} = C\int_{-\infty}^{\infty}\frac{\tilde{h}^*(f)\tilde{h}(f)}{S_n(f)}\,df$$ (6.23)

From this expression we find the optimal filter for extracting $h(t)$ from stationary, Gaussian noise as

$$\tilde{F}^*(f) \propto \frac{\tilde{h}(f)}{S_n(f)}.$$ (6.24)

We can use the above filter to define the inner product of the data, s, with the signal, h,

$$(h|s) = 4\,\mathrm{Re}\int_0^{\infty}df\,\frac{\tilde{h}^*(f)\tilde{s}(f)}{S_n(f)},$$ (6.25)

This result states that the best way to search for a signal in noisy data is to

filter with the signal itself, weighted by the noise spectrum. From this we can formally define the SNR of a signal h in detector data s as the normalised inner product of that signal with the data,

$$\text{SNR} = \frac{(s|h)}{(h|h)^{1/2}}. \tag{6.26}$$

This is then proportional to the amplitude of the signal in the data.

6.3 The multi-detector coincident pipeline

Armed with an analytical model of the signal, and a method of extracting that signal from detector data, we can start searching for gravitational waves. This section summarises the analysis pipeline used in the last joint LIGO-Virgo science run to search for GWs from low-mass binary inspiral signals [18, 102].

This pipeline, known as 'ihope', uses the matched-filter method to test for the presence of GW signals in the data from each detector, in the hope of identifying coincident peaks across the network signifying the presence of a gravitational wave. The basic structure is as follows:

1. construct a bank of templates with which to filter, spread out across the parameter space,

2. filter the data for each detector individually, giving a time-series of SNR,

3. match up those peaks in SNR from each detector that could correspond to a single event present in the network,

4. discard likely noise artefacts using data quality information and signal-consistency tests

5. work out the significance of the remaining coincidences compared to the noise background.

At the end of all this, we end up with a list of those events that have passed a number of cuts and have been ranked most significant relative to the background; if the significance of a single event is high enough, it could be considered a detection candidate. Full details of this method can be found in [102].

6.3.1 Event trigger generation

The pipeline can roughly be divided into two parts, the first of which deals only with extracting signals from the data of each detector. This is done by matching the data against a range of template waveforms, and recording peaks in SNR. In this coincidence-based method, events are recorded for each detector individually, with multi-detector combinations generated by combining events across the network.

Constructing a template bank

The matched-filter of section 6.2 will match the data against a single waveform model, with a specific choice of signal parameters. In order to perform a search that will accurately identify a signal with unknown parameters in the data we must filter using a number of template waveforms whose parameters span the parameter space.

We construct a bank of templates covering the required parameter space, distributing the templates to have as few as possible, whilst demanding that the loss in SNR for a signal that lies between templates does not exceed a given threshold. For a non-spinning binary system, the only intrinsic parameters of the waveform are the component masses – all other parameters are a result of the relative direction and orientation of the source and detector, and the timing of the coalescence; the extrinsic parameters act as amplitude and phase modifiers, and can be analytically marginalised during filtering [138].

In this case, the optimal method [139, 140, 138, 141] is to place templates in a hexagonal lattice in the (τ_0, τ_3) metric coordinates [102] such that the loss in SNR between templates will be no greater than 3%[1], as shown in figure 6.1.

Spectrum estimation and filtering

The pipeline analyses all times containing science-quality data – those times during which each instrument was fully operational, and no serious (category 1) data quality issues were recorded. For each of these time segments, the GW strain time-series for each detector are separated into blocks of 2048 seconds[2] over which the detector noise spectrum is measured. The full PSD (equation (6.19)) is estimated by taking the median average of 15 256 s PSDs from

---

[1]A 3% reduction in SNR gives $\lesssim$ 10% loss in signals.
[2]Detector noise can be approximated as stationary on this time-scale.
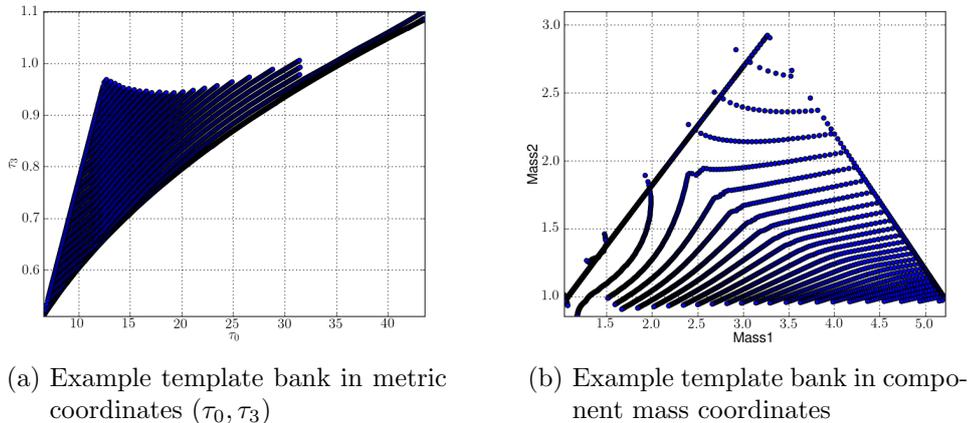
(a) Example template bank in metric coordinates (τ_0, τ_3)

(b) Example template bank in component mass coordinates

Figure 6.1: Example of the template bank used in the coincident search.

overlapping time segments.

The single-detector signal-to-noise ratio is calculated as the phase-maximised, matched-filter SNR, where we separate the template waveform into the 0 and $\pi/2$ phase components introduced in equations (6.15), and apply equation (6.26),

$$\rho^2 \big|_{\text{Max}(A,\Phi_0)} = \frac{(s|h)^2}{(h|h)} = \frac{(s|h_0)^2 + (s|h_{\pi/2})^2}{(h_0|h_0)}. \tag{6.27}$$

An inverse Fourier transform is applied to recover the SNR at all times,

$$(s|h)(t_c) = \int_{-\infty}^{\infty} \frac{\tilde{s}(f)\tilde{h}_0(f)^*}{S(f)} e^{2\pi i f t_c} \, df, \tag{6.28}$$

where $S(f)$ is the detector noise PSD, and t_c the coalescence time of the signal. Any peaks in the SNR time series are recorded as triggers, recording the time and amplitude of the peak, along with the parameters of the corresponding template waveform.

6.3.2 Multi-detector event selection

Once the data from each detector have been filtered over the template bank, the second stage of the pipeline down-selects the multi-detector triggers to identify those events that could represent a GW signal present in the detector network.

Coincidence

For the first- and second-generation GW detectors, an astrophysical signal is not likely to be resolvable using a single detector. The predicted amplitude of signals relative to the noise spectrum of each instrument means signal power in multiple detectors is required for a detection to be believed.

This constraint is applied to all triggers identified by the matched-filter; two triggers from different detectors are said to be 'coincident' if they are separated by a sufficiently small time window, and the templates with which they were best recovered are sufficiently similar in binary component mass [102].

Background estimation by time slides

Confident recovery of signals from the data – either real or simulated – is based on estimating how likely it is that each candidate could have been generated by a noise artefact in the detector network. As such, generation of a large set of background events that represent noise coincidences between detectors will allow a more accurate search.

Large amounts of background data can be constructed by shifting the data from one detector relative to another. Small time offsets between detectors could be explained by the signal arriving at one site before another as a result of its sky location; however, very large offsets can by used to ensure that any coincident triggers must have been produced by a noise artefact.

The coincidence-based analysis generates background triggers by shifting the triggers from one detector by a multiple of 5 seconds relative to another before the coincidence check. Sliding only the triggers, and not the raw data themselves, allows a large number of time slide backgrounds to be created with minimal computational cost.

Simulations

In order to test the pipeline's ability to distinguish GW signals from noise, simulated signals are added to the data before filtering. The data including these *software injections* are analysed identically to the regular data, with the results used to improve the analysis' ability to distinguish noise artefacts from real signals. This is done through tuning signal-consistency tests – identifying noise artefacts with a high matched-filter SNR – and selection of an optimal detection statistic used to rank foreground triggers against the background.

Additionally, in the case of a non-detection of GW signals, the results of the

simulation runs are used to determine search efficiency. The fractional number of simulations recovered as a function of simulated source distance are used to estimate the sensitive distance of the search, and construct upper limits on astrophysical event-rate density. Given that these signals are not present in the actual recorded data, a huge number of simulations, with signals spanning the full parameter space, can be performed using the same data, allowing large statistical samples to be generated.

6.3.3 Data quality and signal-based vetoes

To this point, the coincident pipeline has produced a list of multi-detector coincident triggers that pass a simple threshold on SNR. However, as we have seen in chapters 4 and 5, the presence of non-Gaussian noise in the data results in spikes of SNR in each detector not caused by GWs.

These events typically do not have the same spectral shape as an inspiral and so are identified using a spectral χ^2 statistic [142] that measures the mismatch between the template and the data over a number of independent bins in the frequency domain. A true GW signal will match the template well over each frequency bin, but a glitch will typically only match well in a narrow band. Figure 6.2 shows the distribution of the spectral χ^2 statistic versus SNR for background and simulations from the H1 detector during a two-week analysis of S6 data. Typically, the simulations maintain a low χ^2-statistic value for all values of SNR, but loud glitches match the template increasingly poorly with higher amplitude. Further signal-based cuts are not described here, for details see [102].

Additionally, as described in section 4.5, all triggers falling within time flagged by a category 2 DQ flag are identified and immediately discarded from the analysis. The remaining events are used to identify potential detection candidates worthy of individual follow-up. The category 3 DQ flags are then used to identify those triggers unlikely to be true GW events, with non-detection upper limits generated based on the distribution of events remaining after this stage. The effect of this process on the background events from S6 are shown in figure 4.12.

6.3.4 Detection statements and rate limits

The stages outlined in sections 6.3.1 to 6.3.3 are repeated for the foreground data (no simulations, no time shifts), each simulation run, and the time slide
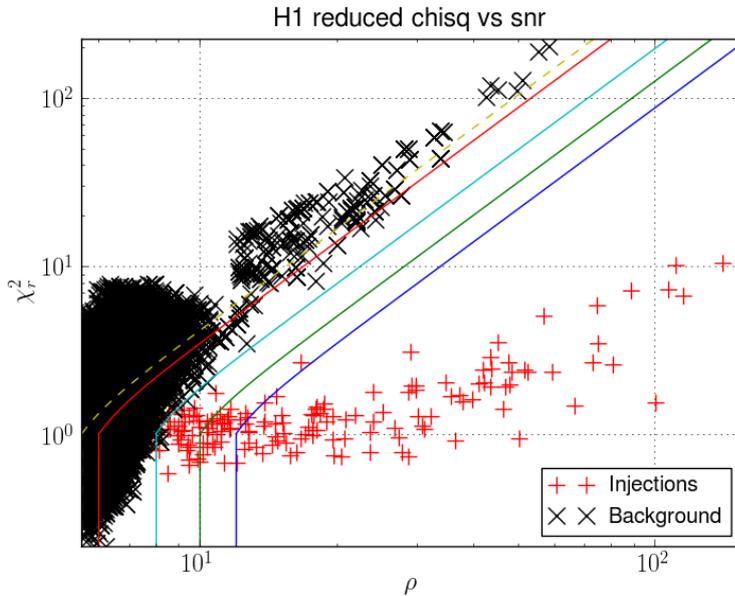
Figure 6.2: Distribution of the spectral χ^2-statistic versus ρ (SNR) for triggers from a two-week analysis of H1 data. Here the background (noise) events from time-slid analysis are in black, with simulated signals in red, and the contours represent constant values of the re-weighted SNR detection statistic.

background analyses. Figure 6.2 shows that ranking events based simply on their matched-filter SNR does not accurately distinguish signals from noise, and so a re-weighted SNR, incorporating the χ^2 signal-consistency test has been calculated as [102]

$$\rho_{\text{new}} = \begin{cases} \rho, & \chi^2 \leq n_{\text{dof}} \\ \dfrac{\rho}{\left[\left(1 + \left(\frac{\chi^2}{n_{\text{dof}}} \right)^3 \right) / 2 \right]^{1/6}}, & \chi^2 > n_{\text{dof}}. \end{cases} \tag{6.29}$$

where n_{dof} is the number of degrees of freedom of the χ^2 statistic. The contours shown in the figure represent constant values of this new statistic, clearly distinguishing the simulations from the background.

Events in the foreground are then assigned a false-alarm rate (FAR), a measure of how often an event at least this significant is seen in the background. Figure 6.3 shows the distribution of inverse FAR (so that more significant events have a large value) for an example two-week analysis, with the foreground events ranked against the background from 100 time-slide trials. If the FAR assigned to any event is sufficiently low, it may be claimed as a detec-
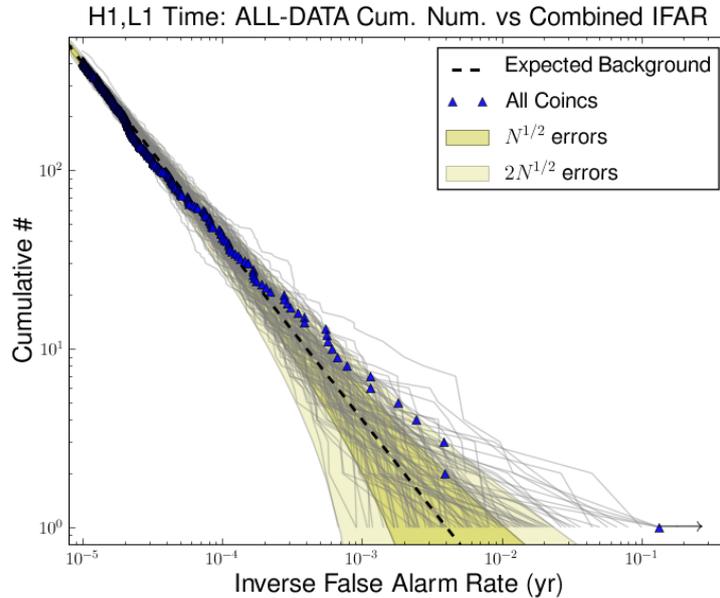
Figure 6.3: The inverse FAR (IFAR) distribution for an example two-week analysis using the coincidence-based pipeline. Each of the grey lines represents a single time-slid background trial, on top the shaded 1- and 2-σ error regions. The blue triangles show the search foreground; a simulated CBC signal appears on the far right with very high IFAR.

tion event. An example of the follow-up for such a candidate will be given in section 6.4.

The final result is a statistical statement of rate limits. Search efficiency is calculated as the fractional number of simulated signals that are successfully recovered as a function of their injected distance. Integrating this efficiency over distance gives an estimate of the sensitive volume of the analysis, which when combined with the duration of the search, yields an upper limit on the observable event rate density. In the event of a detection, the observed rate density can be used to raise the lower bound away from zero.

6.4 The all-sky coincident search during S6

During S6/VSR2/3, data from the LIGO-Virgo instrument network were searched for signals from low-mass compact binary coalescences. This section summarises the all-sky analysis that was performed, including the impact of data quality, the blind injection challenge, and final search results. The full analysis results were presented in [18].

6.4.1   Analysis summary

During the run, the data were separated into calendar fortnights, with each analysis performed as soon as all data, analysis segments, and DQ information were available and verified. This latency allowed quick identification of significant events in the foreground and any problems with the analysis.

The analyses were performed using a bank of non-spinning, inspiral-only waveforms generated at 3.5 PN order [143, 144] and placed across the mass parameter space using the method described in section 6.3.1. The templates were restricted to have a component mass of at least $1\,M_\odot$ with a maximum combined mass of $25\,M_\odot$, allowing for any combination of neutron stars and stellar-mass black holes. Given the first-generation LIGO sensitivity, the merger and ringdown phases of a CBC event with a plausible amplitude would have been outside of the sensitive band, allowing the fully analytical inspiral-only waveforms from section 2.5 to fully describe the detectable signals [18].

Several simulation campaigns were used to measure search efficiency, including a number of BNS, neutron star-black hole (NSBH), and BBH-specific configurations, and a set of spinning signals. No statement of the efficiency towards spinning signals was made in the results, rather these were used to inform decisions on analysis pipeline development.

6.4.2   Data selection

This search used data from the two LIGO detectors and the Virgo detector between July 7 2009 – October 20 2010, totalling 242 / 221 days of analysable data for the H1 and L1 detectors respectively and 50 days of data for the Virgo detector spanning VSRs 2 and 3.

These data were down-selected using category 1 vetoes, as outlined in chapter 4; these vetoes represented times during which the instruments were not operating in a stable manner, and as a result the data should not have be analysed. The resulting segments were further down-selected based on a 2064-second duration requirement, applied to ensure enough contiguous data were available to accurately measure the detector strain sensitivity at a given time. In the end, 233/208/49 days of data were analysed for the H1/L1/V1 detectors during the run.

All available data were analysed, whether or not data from another detector was available at the time, in order to construct larger background trials from the time-slid coincidences. A total of 138/138/41 days of data for each inter-
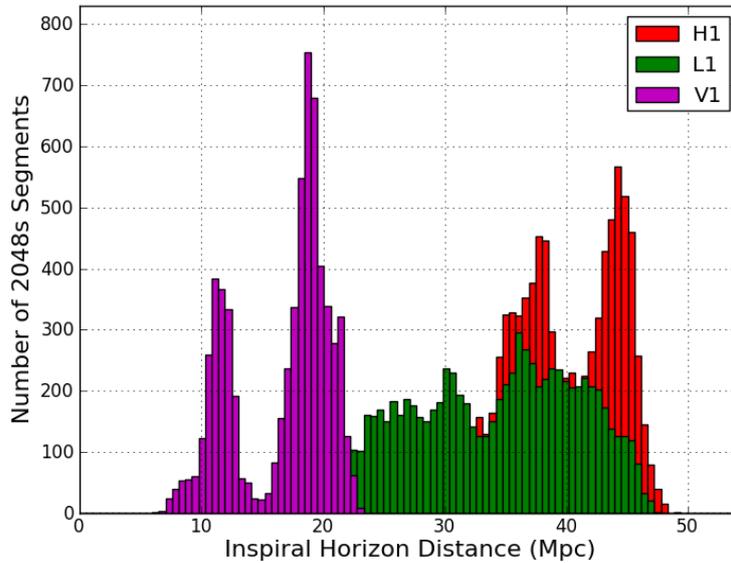
Figure 6.4: A histogram of sensitive distance for each analysable data segment for the three detectors during S6[17]. The horizon distance used here is a measure of the range to which each instrument could detect an optimally oriented and located BNS with an SNR of 8. The bi-modal distribution is due to poorer environmental conditions during working hours.

ferometer were found in at least two-detector coincident operation; it is from this time that detection candidates were identified. To construct the background trials, 100 time slide analyses were performed, allowing these events to be generated from effectively nearly 40 years of data for the LIGO detectors, and over ten years of data for Virgo.

6.4.3 Detector sensitivity and data quality

The general performance of the LIGO detectors during S6 was studied in chapter 4, however a number of issues central to the CBC all-sky search were not discussed or only briefly mentioned. Figure 6.4 highlights the stability of the three detectors in terms of the number of analysable segments achieving a given horizon distance sensitivity. All three detectors show a bi-modal distribution, understood to be the difference in the seismic environment around each site during the working day (lower sensitivity) and during evenings and weekends (higher), while L1 shows a more even spread due to its more unstable running during the early parts of the run.

Figure 6.5 gives the median range of each detector as a function of binary
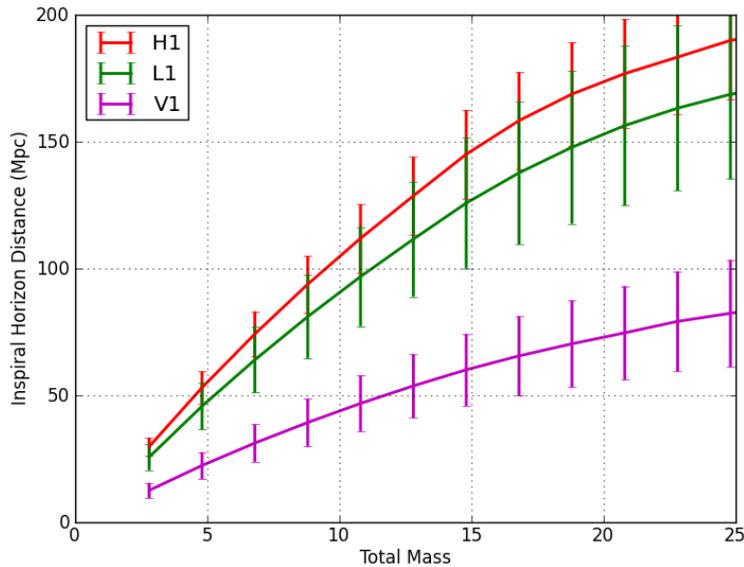
Figure 6.5: Sensitive distance to equal-mass binaries as a function of total mass for the LIGO-Virgo detectors during S6/VSR2/3 [17]. In this figure, the horizon distance is the sensitive distance assuming an optimally oriented and located source, for which the detector responses (equations (6.12)) are maximal.

total mass ($M_\odot$). The sensitive distance scales with the binary chirp mass as $\mathcal{M}_c^{5/6}$ and so the range increases monotonically from 2–25 $M_\odot$[3].

Apart from spectral sensitivity, search performance is determined by the background event distribution, that collection of events against which any detection candidates are ranked. A tail of loud events will result in any high SNR detection candidates being ranked as likely noise artefacts, and so not considered as a true detection.

Throughout the science run, any significant background events were identified with $2-4$ week latency. Weekly studies of DQ – using a low-latency, single-detector implementation of the ihope pipeline – and fortnightly studies of the data from the main analyses were performed, including the performance of those vetoes available at the time, with the loudest events studied in detail.

These studies included surveys of auxiliary data signals at the time of the loud events, including (for example): seismically-induced motion on the laboratory floor and on optical tables; angular beam motion at various points in the interferometer; and the power incident on photodiodes in the OMC. If any auxiliary signals were observed to correlate strongly to significant events

---

[3]For reference, the canonical BNS source used in figure 6.4 has a total mass of $2.8\,M_\odot$.

in the search, glitches in these channels could be used to remove times from the analysis that were likely to include noise artefacts.

Using this procedure, a number of new DQ flags were produced, including the targeted seismic-noise flag presented in chapter 5, and one flag used to remove 8 seconds on either side of any peak in the matched-filter SNR above 250 [18]. This flag was used in particular to identify the 'spike' glitches in the data from L1 described in **??**.

Figure 6.6 (reprinted in part from figure 4.12) shows the background event distribution for the LIGO-Virgo detectors before and after application of category 2 and 3 DQ vetoes. As described in section 4.5 for LIGO, category 2 vetoes have high efficiency in vetoing those events with high matched-filter SNR, but those are largely re-weighted to low significance by the spectral χ^2 statistic. The same is true of the Virgo detector background, except for this detector the background is very quiet to start with; a lower sensitivity for this detector throughout the analysis meant most of the background coincidences included loud transients from LIGO data with something innocuous in the Virgo data.

The category 3 vetoes had a huge impact on the analysis. In $\mathcal{O}(10)\%$ deadtime for each detector[4], the loudest event for each detector is reduced by tens of percent. Each of the LIGO detectors recorded a loudest event with χ^2-weighted SNR ~ 12, while that in Virgo data had a value of 8.8.

What is apparent is the use of the spectral χ^2 statistic to re-weight the matched-filter SNR has rendered category 2 vetoes almost obsolete. This category is defined to contain those analysable data containing noise in an auxiliary data signal whose correlation to the GW output is known. It seems those events, typically producing very high matched-filter SNR, are recognised by the χ^2 statistic as noise transients, and down-weighted appropriately. However, given the relative deadtime of category 2 vetoes, $\mathcal{O}(1)\%$ at each site, the amount of data lost is minimal.

6.4.4 Analysis of S6D

The final epoch of the run, S6D, ran from June 26 – October 20 2010. This period began with the scheduled return of the Virgo detector for VSR3, and finished when the first-generation LIGO interferometers were decommissioned. Technical issues led to the first analysis-quality data from the Virgo instrument

---

[4]Numbers can be found in section 4.5 for LIGO; for Virgo, category 2 vetoes contributed 1.8% deadtime while category 3 added another 5.4%.

(a) H1

(b) L1

(c) V1
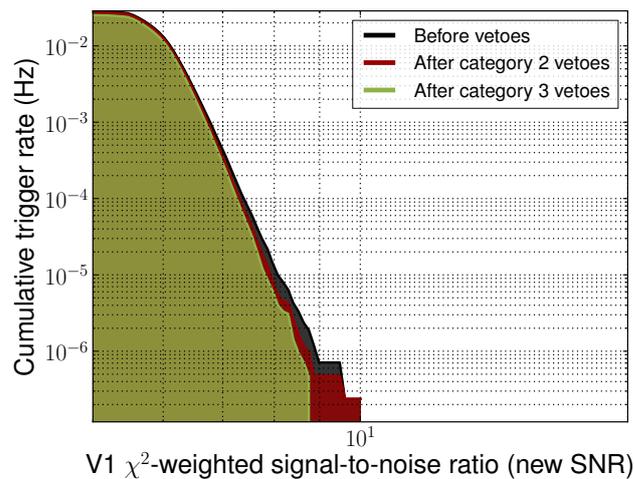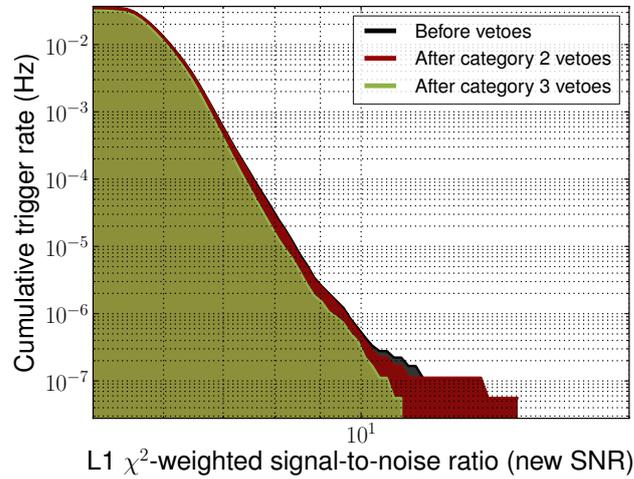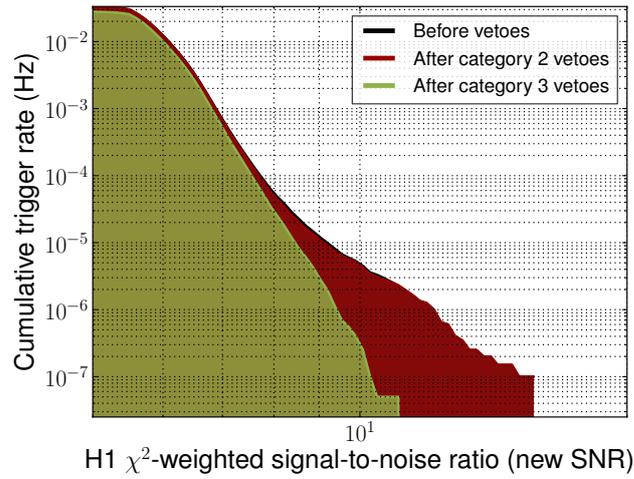
Figure 6.6: The background event distribution from the CBC all-sky search in S6/VSR2/3 before and after data quality vetoes.
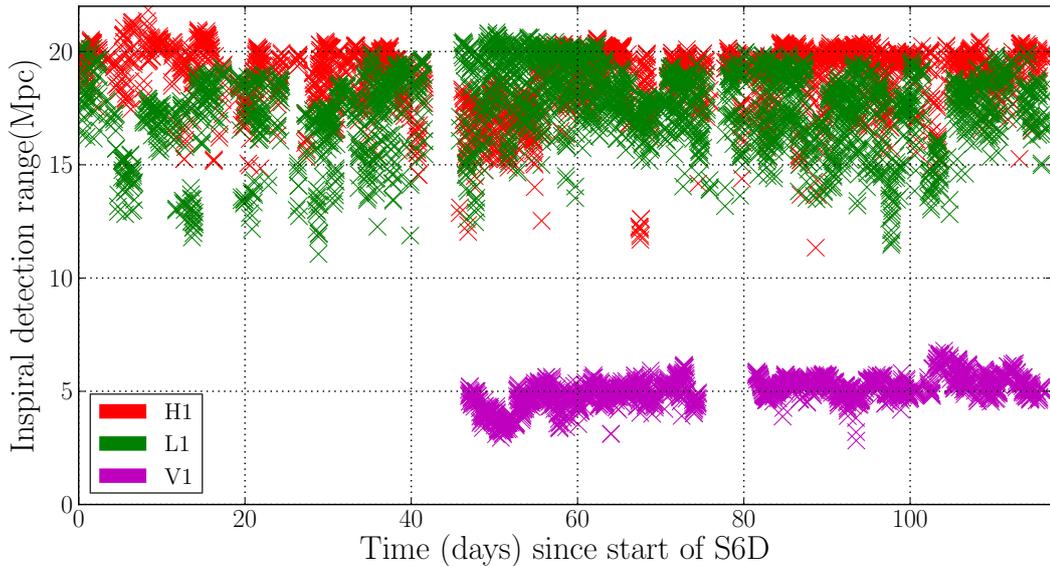
Figure 6.7: Inspiral sensitive distance for the H1L1V1 network during S6D averaged over all sky locations and source orientations.

in mid-August 2010.

The data were divided into 8 fortnightly analyses[5]. Figure 6.7 shows the maximal ('horizon') distance to which each detector could detect an optimally oriented and located BNS inspiral signal throughout this epoch. As shown, the LIGO detectors had a sensitive horizon distance around 40 (35) Mpc for H1 (L1), while the Virgo detector, only returning from its commissioning break around day 45, had that of around 10 Mpc. This was the most sensitive epoch for each of the LIGO instruments (cf. figure 4.4), however the Virgo detector performance during VSR3 was lower than that during VSR2 [108].

During analysis of one fortnight, an outlier was seen in the foreground with higher significance that anything in the background of 100 time-slides. Figure 6.8 shows time-frequency spectrograms of the H1 and L1 data at the time of the event, dubbed GW100916 (in the manner of GRB naming), clearly showing a chirp-like signal above the background. This signal, observed in both LIGO detectors but not the Virgo detector, was was recovered as a high-mass-ratio NSBH binary with a multi-detector, χ^2 re-weighted SNR of 12.5, with a false-alarm rate of 1 per 7000 years after an extended background analysis and inclusion of trials factors. That no event was seen in the Virgo data was consistent with it's sensitivity roughly 1/4 that of the H1 and L1 detectors which recorded individual re-weighted SNRs of 10.3 and 6.2 respectively.

This event also sparked an extensive review of all analysis and DQ informa-

---

[5]The final 'fortnight' spanned 19 days to include the last days of the science run.
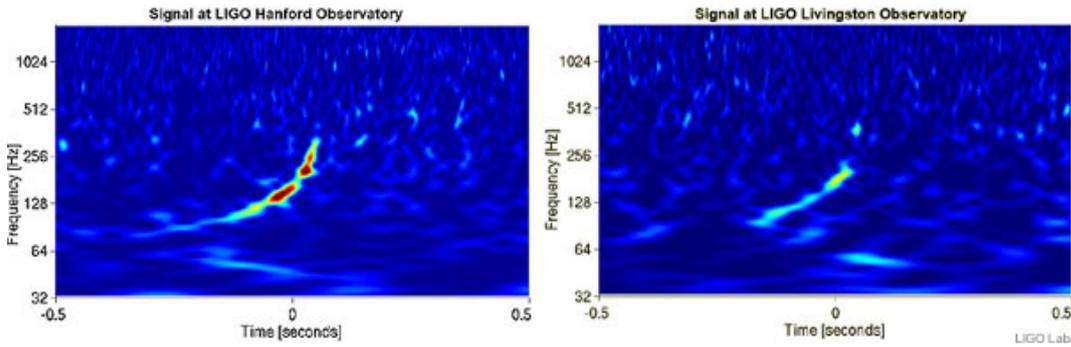
Figure 6.8: Time-frequency spectrograms of candidate GW100916 in the LIGO data [145].

tion. Auxiliary data signals were studied at the time of the event to determine any noise couplings that could have been falsely identified. Figure 6.9 shows the seismic noise level in the $1 - 3\,\mathrm{Hz}$ and $3 - 10\,\mathrm{Hz}$ bands at LHO; a vehicle was recorded on site around the time, but at can be seen, did not interfere with the candidate. Additionally, 9 seconds prior to the event in L1 data, a massive transient 'spike' glitch was observed. As noted above, these spikes can cause instability of the pipeline for a number of seconds around the glitch, however this event was not found in coincidence with data from the other detectors, and so did not affect this search.

After the analysis of this event was complete, it was revealed that this event was not a genuine GW, rather a signal added to the interferometer control feedback signal to simulate a GW. A full description of the Blind Injection Challenge, and the conclusion taken from the exercise, can be found in [145]. No other significant outliers were seen during this period, as is discussed in the coming section.

6.4.5 Search results

The fortnightly analyses, including those during S6D, were considered a 'first cut' at the data, and used to direct decision-making about search inputs – template parameters, DQ flags, etc. – and tune search outputs – including development of the χ^2 re-weighted SNR statistic. In the final analysis, the data were separated into larger chunks, each allowing for a larger background sample whilst remaining computationally practical.

These chunks were each analysed separately, with results studied from each, before the data were combined into a single results statement, and a calculation of search upper limits. The results are presented in full in [18].
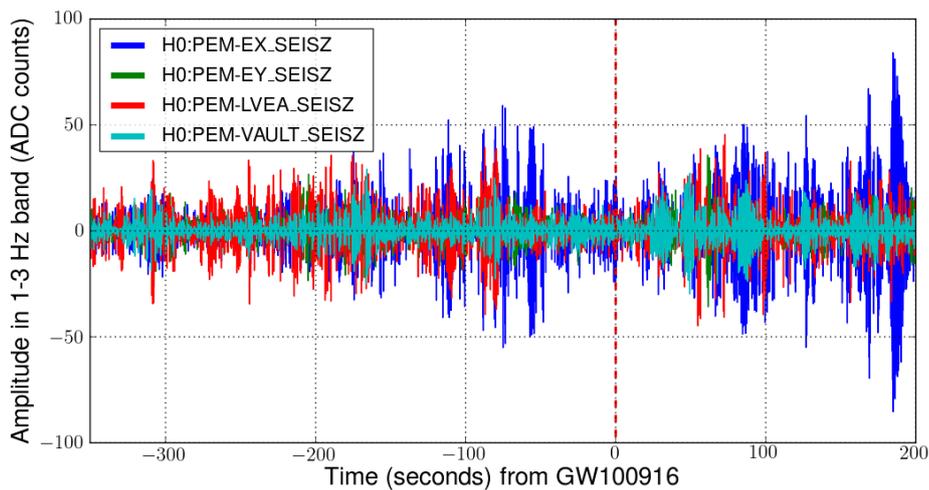
(a) $1 - 3\,\mathrm{Hz}$
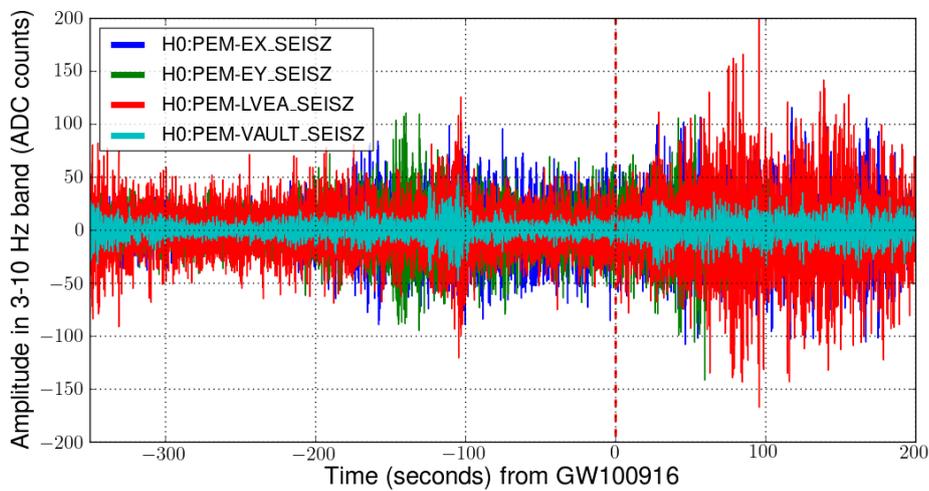
(b) $3 - 10\,\mathrm{Hz}$

Figure 6.9: Seismic noise around the time of the GW100916 candidate, in the (a) $1 - 3\,\mathrm{Hz}$ and (b) $3 - 10\,\mathrm{Hz}$ bands.
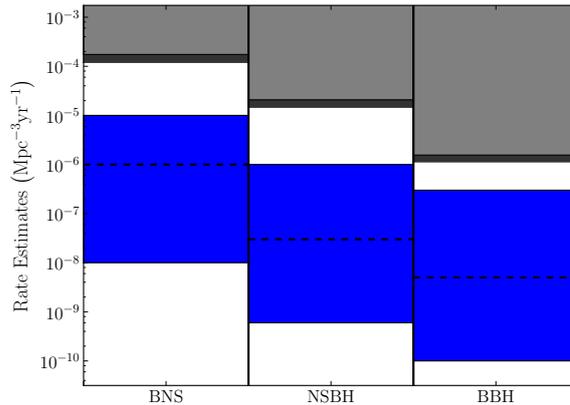
Figure 6.10: Upper limits on event rate density set by the S6/VSR2/3 search for low-mass CBC signals. The grey region represents the upper limit set during S5/VSR1; the black bar, the improvement set with S6/VSR2/3; while the blue region spans the pessimistic to optimistic rate estimates for each source [18].

No GW signals were observed in the data, with all foreground events consistent with the background from time-slid analysis. Upper limits were then set on the rate density of each of the three source types analysed – BNS ($m_1 = m_2 = 1.35 \pm 0.04 \, M_\odot$), NSBH ($m_1 = 1.35 \pm 0.04 \, M_\odot$, $m_2 = 5 \pm 1 \, M_\odot$) and BBH ($m_1 = m_2 = 5 \pm 1 \, M_\odot$) – as shown in figure 6.10. The upper limits for each source are still around one order of magnitude away from optimistic astrophysical rate estimates [52], however, with the sensitivity improvement predicted for the second-generation GW detectors, search sensitivity should improve by a factor of 1,000, allowing analyses to probe the range of realistic event rates shown.

Chapter 7

The coherent, targeted search for binary coalescence signals associated with gamma-ray bursts

In this chapter, we describe the concept of and results from a *targeted* search for gravitational waves associated with known GRB detections.

Any one of a number of partner telescopes (either electromagnetic or high-energy neutrino) or an automated alert service (such as *ATel* [146]) may indicate the detection of an astrophysical signal. If this event could have produced detectable GWs, a search around the same time, frequency, or sky location may be carried out. A number of such searches have been completed in the past, targeting GRBs [147], SGRs [47], high-energy neutrinos (HENs) [148], and pulsar glitches [149], for example.

For the specific case of GRBs, a single satellite, such as *Swift* [150] or Fermi [151], or a network of such satellites, such as the InterPlanetary Network (IPN) [152], will record tens or hundreds of observations per year, with information about each detection rapidly disseminated to interested parties by means of the Gamma-ray Coordinates Network (GCN) [153]. This information includes the detection time, the burst duration, the most likely sky location, and the radius of the 95% confidence area on the sky[1]. These data are used by the GW analysis groups to target a search over a subset of parameter space. The constraints placed by the target information allow greater distinction between background noise and signals, and so present an opportunity for a more

---

[1]Any other information included is not currently used to direct a GW search.

sensitive search.

It was for this type of search that a coherent matched-filtering method was first demonstrated for a CBC search [101], whereby the higher computational cost of a multi-detector matched-filter was reduced by the restricted sky location and timing provided by GRB detections.

This chapter presents an overview of the search for GWs associated with short, hard GRBs, thought to be produced by neutron-star binary mergers. We include a brief discussion of the suspected progenitor of these bursts, before describing the methodology of the coherent matched-filter. We describe the pipeline developed to use that method in a targeted search, including an extension of the analysis to search over the full 95% confidence sky area. Finally, the results of that search over LIGO and Virgo detector data during the joint S6/VSR2/3 run are summarised.

7.1 Short-hard gamma-ray bursts

Gamma-ray bursts have been seen since the late 1960s as some of the most highly-energetic events in the universe. Following a survey of detected bursts [154], a consensus has formed dividing GRBs into two classes, short and hard, and long and soft, where 'hard' and 'soft' refer to the hardness ratio, a measure of the relative photon energy in different bands. The bursts are characterised by their T_{90} – the time in which the central 90% of photon counts for a burst are observed – with those bursts with $T_{90} < 2$ seconds largely classed as short, and the others as long. Figure 7.1 shows the distribution of the T_{90} amongst a large sample of GRBs from the Burst And Transient Source Experiment (BATSE) catalogue [155].

It is believed that, specifically, short GRBs are produced during the coalescence of either BNSs [157, 158] or NSBH binaries [158, 159, 160]. In these models, a binary merger will form a new black hole surrounded by a disk of matter that very rapidly accretes onto the black hole, powering a violent burst of energy in the form of gamma-rays. Given the strong possibility of detecting GWs directly from the coalescence of two compact stars, the ability to target a search from a known sky location and time (and possibly distance) would greatly increase the confidence in that GW detection. Furthermore, combined detection, or non-detection, of the same event in multiple radiation types (rather then just multiple EM bands) would allow for a increased scientific yield [161, 162, 163].
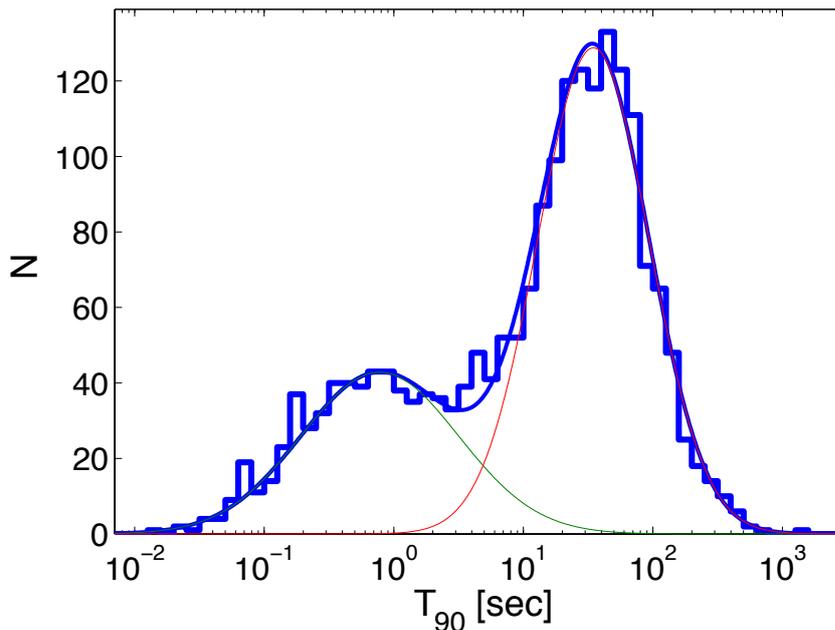
Figure 7.1: The distribution of the T_{90} duration of GRBs [156]. The blue histogram represents the data from GRBs from the BATSE catalogue [155], while the green and red traces give Gaussian models of short and long bursts respectively.

7.2 The multi-detector coherent matched-filter

Harry and Fairhurst [101] developed a coherent algorithm for use in a triggered search for GWs, whereby the data from multiple detectors is combined at the filtering stage to calculate a coherent network signal-to-noise ratio.

Assuming a binary waveform as described in equations (6.15), we can define the multi-detector inner product as a sum over detectors,

$$(\mathbf{a}|\mathbf{b}) = \sum_{X} \left(a^X | b^X \right).$$ (7.1)

Also, note that in transforming from the radiation frame to the detector frame, we introduced a polarisation angle, ψ. We choose to introduce an equivalent angle, χ, in transforming from the source frame to the radiation frame (see section 1.4), without placing any restrictions on the signal. This radiation polarisation can be chosen such that the two detector response functions, F_+ and $F_\times$, are orthogonal, effectively orienting our detector frame such that the interferometer network is much more sensitive to the + polarisation than the ×. In this 'dominant polarisation' frame [164], the multi-detector coherent

SNR is defined as [101]

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s}|\mathbf{F}_+\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_+\mathbf{h}_0|\mathbf{F}_+\mathbf{h}_0)} + \frac{(\mathbf{s}|\mathbf{F}_\times\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_\times\mathbf{h}_0|\mathbf{F}_\times\mathbf{h}_0)}. \quad (7.2)$$

Here s^X is the detector data stream for detector X, $F_{+,\times}^X$ are the detector response functions defined in section 6.1.2, and $h_{0,\pi/2}^X$ are the two phase amplitudes of the template gravitational waveform defined from equations (6.15).

We can see that this statistic, compared to the single-detector SNR of equation (6.27), uses information from both polarisations of the waveform, and accounts for the relative sensitivity of detectors for a given source sky location, allowing for an improved distinction between signals and noise.

Additionally, filtering a two-polarisation waveform against a network of three or more detectors allows for the construction of a *null stream*: a relative combination of detector data containing no signal power. This gives an additional signal-consistency check if the null stream is similarly filtered against the templates in the bank. Further details on signal-consistency tests used with the coherent statistic are given in the next sections.

7.3 The coherent search for gravitational waves associated with gamma-ray bursts

The remainder of this chapter describes the pipeline built around the coherent method. We begin with discussions of the search over *Swift*-detected GRB targets [150] – whose sky location is reported with greater accuracy than the GW detector network resolution – and the improvements that were implemented for those GRB targets detected by Gamma-ray Burst Monitor instrument aboard the Fermi satellite [151], where the uncertainty on the sky location is much greater than the GW network resolution.

7.3.1 Analysis of *Swift*-detected GRBs

The *Swift* Burst Alert Telescope (BAT) [150] was used for targeted GW searches during S5 and S6 [147, 19]. The typical directional accuracy for this instrument is 1-4′, which is greatly less than the sky location resolution of a GW detector [130]; this means the targeted search can be performed over a

single sky point at the centre of the localisation confidence interval.

Data selection

The information collected from the GRB detection notice is used to construct a search pipeline with the following search intervals:

- an *on-source* window of $[-5, +1)$ seconds around the GRB peak time is searched for a foreground GW signal,

- an *off-source* window of 2048 seconds around the peak time, not including the on-source and a 72 s buffer, is used to measure the background noise at the time. This is done by constructing 300 6-second background trials against which to rank the on-source.

The choice of on-source window reflects the fact that the GRB emission is expected as a result of post-merger accretion onto the newly formed black hole, but with no matter shell containing the burst (as in the case of long GRBs) [165], so GW emission should precede the GRB. Only those GRBs for which there is a multi-GW-detector network operating over the full off-source window are analysed.

Template generation

In a coincident analysis, where detectors are filtered individually, a template bank is generated for each detector, spanning the component mass space of the binary (see section 6.3.1); however, in a coherent search, a multi-detector template bank is required. At this time, a fully-coherent template bank placement algorithm (such as [166]) has not been implemented, but rather we simply use a bank generated for one detector.

Coherent filtering

The data from each detector are then filtered against each template in the bank, to generate single-detector complex SNR streams which are then combined to calculate the coherent SNR defined in section 7.2. The difference in signal arrival time (assuming GWs travel at the speed of light) is used to shift the single-detector SNR time series to represent a coherent sum for the recorded sky location.

In order to reduce computational cost, the coherent SNR is only calculated for those time points where the single-detector SNR in at least one detector

is above a threshold of 4. Futhermore, only those time points for which the coherent SNR is above 6 are recorded and passed to the next stage.

Signal consistency

In order to identify noise artefacts with high matched-filter SNR, four signal-consistency parameters are then calculated for each time point [101].

As well as a coherent version of the standard spectral χ^2 statistic mentioned previously, a second χ^2-distributed statistic is calculated to measure the overlap of the data with a fixed set of templates spread across the bank. Glitches in the data will likely have high SNR when filtered against a variety of templates – matching a small number of cycles with high amplitude – whereas a true signal should have markedly higher significance when filtered against the closest template relative to others in the bank.

Additionally, an auto-correlation χ^2 statistic was implemented to measure how well the distribution of SNR over time for a given template matches the prediction from filtering that template against itself. A glitch will result in high SNR for the duration of the template, as a result of its high match with a small number of cycles, whereas a true signal will sharply peak in SNR when the signal and template overlap best in time.

Figure 7.2 shows the distributions of the three χ^2 tests versus coherent SNR for the off-source, and a simulation run, for an example GRB as analysed during S6. The re-weighted SNR statistic defined in equation (6.29) is used to construct *signal-based vetoes* from the three χ^2 statistics, whereby any events whose spectral χ^2-, bank χ^2-, or auto-correlation χ^2-weighted SNR falls below 6 are vetoed from the analysis as likely glitches.

While the χ^2 statistic are similar to those used in the coincidence-based analysis, the coherent algorithm also allows the use of a new consistency check, the null SNR. This is measured for each event as the SNR for the relevant template filtered against the null stream, introduced in section 7.2, which by construction contains no true GW signals.

In addition to the signal-consistency checks, *category 2* data quality vetoes are used to clean the data of instrumental artefacts. These time segments are used to remove event triggers from both on-source and off-source data before any detection statements are made. Higher category vetoes are not required for the targeted search thanks to the improved distinction of foreground from background with the coherent matched-filter and the above signal-based vetoes.
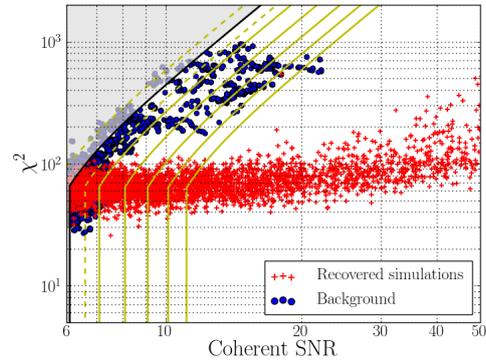
(a)
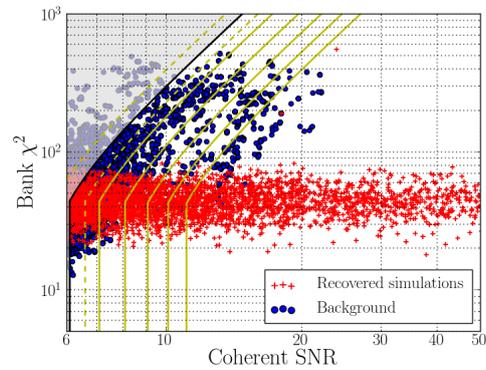
(b)
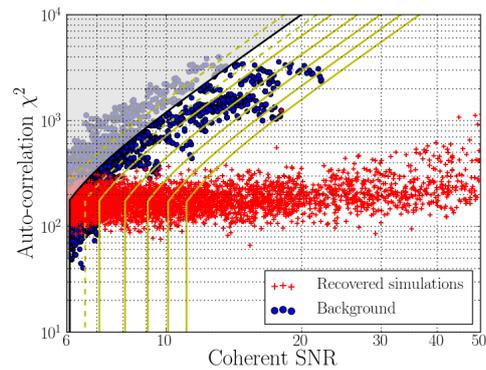
(c)

Figure 7.2: The distribution of the various χ^2 tests versus coherent SNR for an example GRB during S6. For each of (a) the standard spectral χ^2, (b) the bank χ^2, and (c) the auto-correlation χ^2 tests, the shaded region shows those events vetoed, while the contours show constant values of the reduced SNR, weighted by the relevant χ^2 statistic.

Simulation runs

For each GRB analysis, simulation runs were used to inject neutron star binary signals (BNS and NSBH) into the data, with neutron star masses normally distributed between $1 - 3\,\mathrm{M}_\odot$ and black hole masses between $2 - 25\,\mathrm{M}_\odot$. While the mass distributions were the same for all GRBs, the injection distances were tuned relative to the known sensitivity of the detector network at the target time, to best test analysis performance. Similarly, the source sky locations of the injections were distributed over a densely seeded region, centred on the reported most likely point, out to the reported 90% location uncertainty radius.

The recovery of these simulations were used to instruct tuning of cuts and the detection statistic, and to determine search sensitivity, as discussed below.

Background estimation

As noted under *Data selection*, background estimation is performed using detector data away from the on-source region to construct 300 equivalent trials. The loudest event in each of these trials is used to determine the significance of the loudest event in the on-source window, as a false-alarm probability.

In the coherent analysis, the time-shifts are applied to the raw data before filtering (cf. section 6.3.2) by re-ordering the 256-second analysis chunks between detectors. In this manner fully-coherent background data can be constructed, but with the penalty that each time-slide is computationally equivalent to the non-time-shifted analysis, presenting a massive increase in cost; for this reason, time slide analyses are only used to follow-up those GRBs whose loudest on-source event is more significant than any in the original 300 background trials.

Results: detection statements and distance exclusions

The detection statistic, ρ_{det}, uses a combination of χ^2 weighting, and a weighting by null SNR [101]. The χ^2 statistic is used to weight the coherent SNR as,

$$\rho_{\chi^2} = \begin{cases} \rho, & \chi^2 \leq n_{\mathrm{dof}} \\ \dfrac{\rho}{\left[\left(1 + \left(\frac{\chi^2}{n_{\mathrm{dof}}}\right)^{4/3}\right)/2 \right]^{1/4}}, & \chi^2 > n_{\mathrm{dof}} \end{cases} \tag{7.3}$$

This weighting differs slightly from that in equation (6.29), reflecting the different distributions of background between the targeted and all-sky searches. The null SNR is then used to further weight the statistic as,

$$\rho_{\text{det}} = \begin{cases} \rho_{\chi^2}, & \rho_{\text{Null}} \leq 4.25 \\ \dfrac{\rho_{\chi^2}}{\rho_{\text{Null}} - 3.25}, & 4.25 < \rho_{\text{Null}} \end{cases} \tag{7.4}$$

This additional step, relative to the coincidence-based analysis, down-ranks a small set of noise events whose χ^2-statistics do not identify it as likely noise, further separating foreground and background.

Each event is ranked by it's false-alarm probability (FAP), i.e. the likelihood that an event this loud could be a noise artefact, calculated as the fractional number of background events with a higher value of the detection statistic. If the loudest event in the foreground achieves a low enough FAP ($\lesssim 1\%$), or is louder than all of the background events, it would be followed-up as an interesting candidate.

Similar to the coincidence-based analysis, time-shifts can be used to provide a greater background sample from which to calculate a non-zero FAP for foreground events. With each time-shifted analysis providing a further 300 trials, enough data can be search to find a background event as (or more) significant than the candidate, given an estimate of how likely it would be that the candidate was produced by a noise artefact.

In the event of no interesting candidates, the simulations can be used to determine a distance exclusion statement for each GRB; a true GW signal would be recovered by the pipeline assuming it was louder than all other events in the on-source time.

Each simulation is marked as found if it is recovered with higher significance than the loudest event in the foreground, and the distance at which injection recovery falls below 90% is used to construct a exclusion limit for this GRB. This limit gives the distance inside which a GW sources would have a $\geq 90\%$ chance of producing a candidate event louder than the loudest observed event. Figure 7.3 shows the distance efficiency distribution for an example GRB.

7.3.2 Sky patches for larger localisation areas

Instruments on satellites other than *Swift*, such as the Fermi Gamma-ray Burst Monitor (GBM), give larger uncertainties on the source sky location, typically $\sim 10 \, \text{deg}^2$. If the search was performed using a single search point at the centre
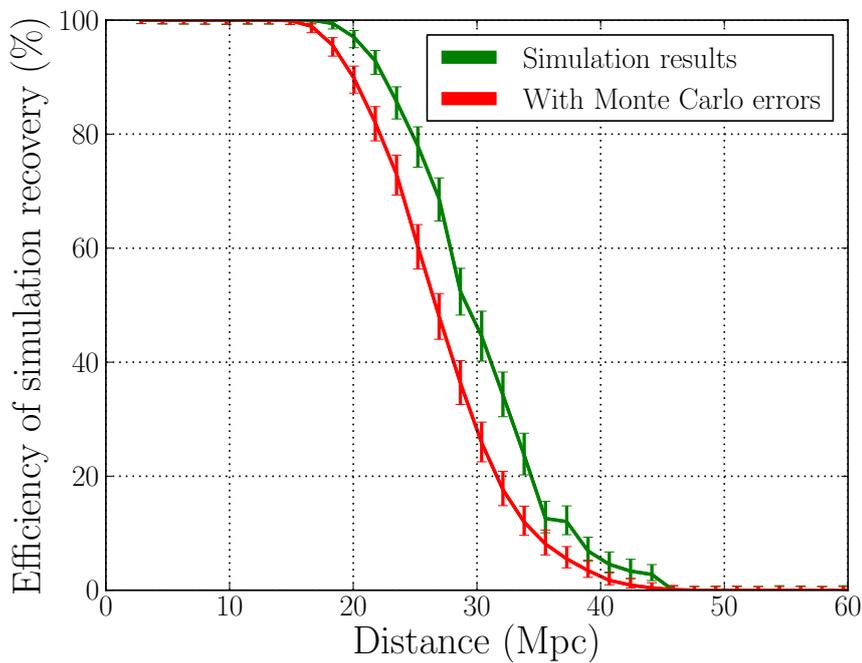
Figure 7.3: Injection recovery efficiency versus distance for GRB100404. The green curve represents the results of the simulation run, with the red curve presenting a more conservative limit after a Monte Carlo simulation of instrumental calibration errors. The exclusion distance lower limit is interpolated from these data at the 90% efficiency point (19.3 Mpc).
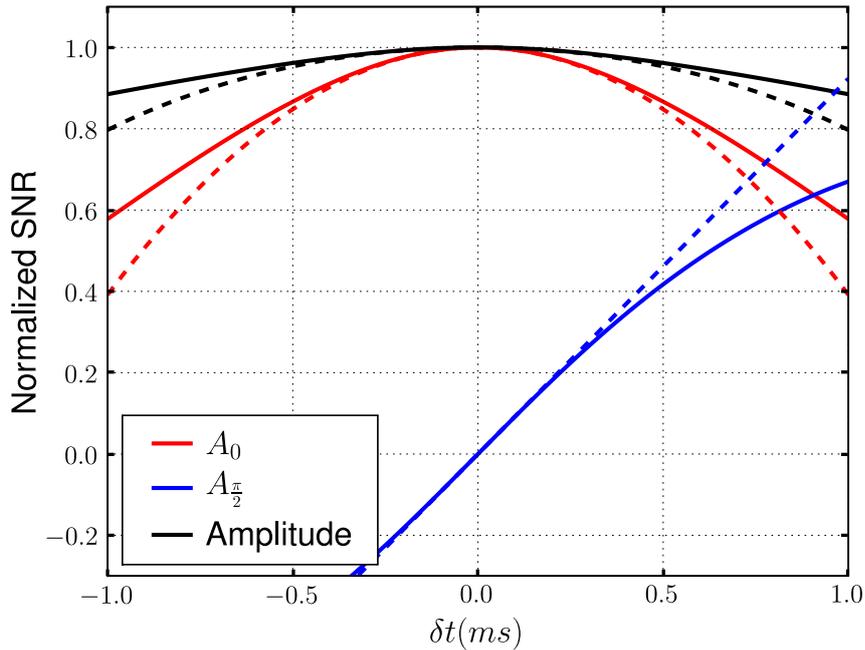
Figure 7.4: The normalised SNR as a function of timing error for initial
LIGO [130]. The solid lines show exact results for each the two
phases of the waveform, while the dashed lines use an approxima-
tion to second order in h.

of the region, a GW signal from the edge would be localised with a time-delay
error of up to several milliseconds. This error in time delay could reduce the
recovered coherent SNR by more than 10%, as shown in figure 7.4 [130]. As
a result, for the search for GWs associated with Fermi-detected GRBs during
S6, a method to tile the full 90% confidence region, and search coherently over
each sky point, was developed.

Considering a network of N detectors, $D_{\{1,\ldots,N\}}$, let t_i be the arrival time
of a GW signal from a GRB at detector i. Now, we can calculate the difference
in the arrival time of the signal at detectors i and j, τ_{ij} as [167],

$$\tau_{ij}(\phi,\theta) == t_j - t_i = \frac{1}{c} \left(\mathbf{r_j} - \mathbf{r_i}\right)^T \mathbf{w}(\phi,\theta), \qquad (7.5)$$

where $\mathbf{r}_i$ is the position of detector i, $\mathbf{w}$ is the unit wave vector, and (ϕ,θ)
are the standard coordinates in our spherical polar system on the unit sphere.
Additionally, if we define d_{ij} as the distance (through the Earth) between two
detectors $(d_{ij} = ||\mathbf{r_j} - \mathbf{r_i}||)$, we can define the *light travel time* between these
detectors as

$$T_{ij} == d_{ij}/c \qquad (7.6)$$
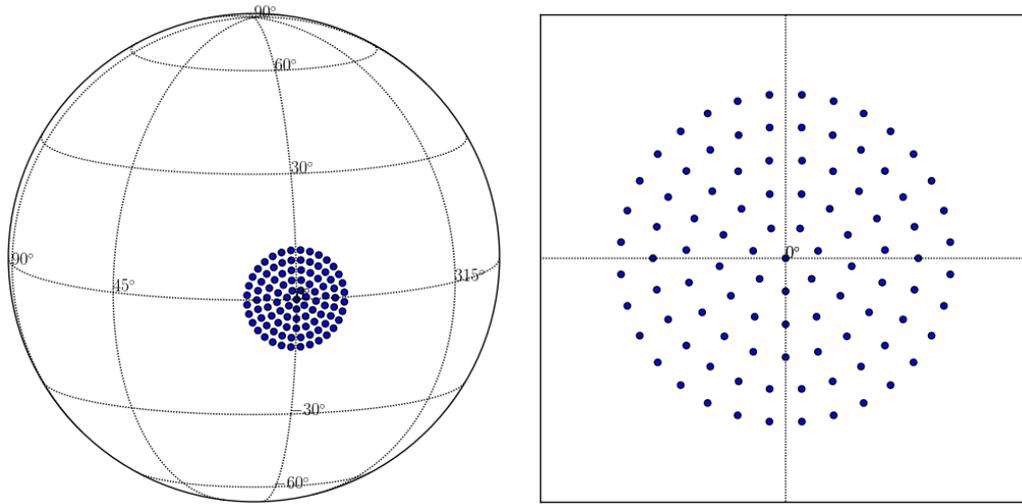
Figure 7.5: Example of a circular grid of sky points for a target with a confidence radius of 10° using the H1L1V1 detector network. In this example the angular resolution is 2.18° (0.5 ms).

In order to tile the necessary sky area, we construct a grid of concentric circles emanating from the target centre, encompassing the 90%-confidence radius from the GRB detection. The angular resolution for the grid is calculated as

$$\delta\alpha = \min_{i,j}\left[2 * \frac{\delta t}{\sqrt{T_{ij}^2 - \tau_{ij}^2}}\right]. \tag{7.7}$$

Here δt is the chosen timing resolution of the search, calculated as the timing error that would reduce the SNR by 3% due to phase-mismatching [130]. In practice, the resolution is determined by that detector pair (D_i, D_j) for which the GRB target location has smallest arrival time difference.

Using the calculated angular resolution, the circular grid can be generated by placing rings of points spaced by $\delta\alpha$, starting at the centre, with the final ring passing the 90% confidence radius. Each ring will have $\lceil 2\pi/\delta\alpha\rceil$ points. The result for a fictional GRB target is shown in figure 7.5.

In the analysis, each point in the grid is treated independently, with the coherent SNR calculated separately for each, applying the detector responses, F_+ and $F_\times$, for that sky point. The differences in signal arrival time between detectors for a given sky location are also used to shift the single-detector SNR streams into a coherent combination specific to that point. The re-calculation of the coherent SNR, and in turn the various signal-consistency checks (χ^2-like statistics, and null stream) for every sky position, introduces a significant increase in computational cost when using a large sky grid. This issue has still
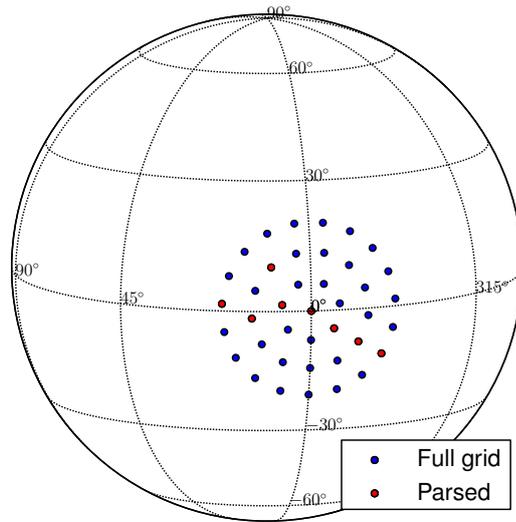
Figure 7.6: The effect of removing time-delay degenerate sky points from a circular map. The blue points show the full grid, while the red points are those few mapping to unique differences in signal arrival time between LHO and LLO. That the parsed points do not form a straight line is an artefact of the parsing routine, that has no effect on the grid reduction.

to be formally addressed, but was circumvented during the actual analysis by exploiting the reduced sky localisation accuracy of a two-site network.

Two-site time-delay degeneracy

In the specific case of a two-site detector network, as is the case for any LIGO-only configuration, the ability to resolve independent sky locations is vastly reduced. With a single baseline between sites, those sky locations separated by lines orthogonal to the baseline will have the same time delay [130]. Two such sky locations still map to unique values for the detector response functions, but with the difference over $10 \, \mathrm{deg}^2$ typically a few percent, these sky locations can be treated as degenerate.

In these searches, we consider any two points with the same time delay ($\pm \delta t / 2$) as degenerate, with only one point recorded, allowing a vast reduction in the size of the sky grids. Figure 7.6 shows the result of simply parsing the circular sky maps to remove degeneracies in time-delay. For the map shown, only 20% of the points are required to uniquely span the allowed time-delays between the LIGO sites, allowing a great speed up in the analysis for two-site GRB analyses.

7.4 Targeted, coherent searches during S6/VSR2/3

The coherent pipeline outlined in section 7.3 was run on GRB targets from the *Swift* and *Fermi* satellites during S6/VSR2/3. These satellites publish each observation through the GCN, informing analysis parameters including timing and sky location, but also including light-curves and any distance (redshift) information available. Both the X-Pipeline burst detection algorithm [100] and the coherent CBC pipeline were used to search for GWs associated with GRBs, with the combined results presented in [19].

This section presents a summary of the analysis of short GRBs using the coherent, targeted CBC pipeline, including the selection of GRBs for analysis, the special case of GRB 090802A, and the search results.

7.4.1 Target selection

In all, 407 GRB detections were announced between July 7 2009 – Oct 21 2010 (S6/VSR2/3). Only those GRBs that occurred when at least two GW detectors were operating in a stable configuration, and whose data were of a sufficient standard, were analysed for associated GW signals. As described in section 7.3.1, 2172 seconds of data around the GRB detection were used to measure the search background; if any category 1 DQ flags were active during that time, the GRB was not analysed. Similarly, if there were any active category 2 flags during the 6-second on-source search window, the GRB was not analysed.

After these criteria were applied, 150 GRBs were analysed by X-Pipeline, while the coherent CBC search analysed those 26 targets identified as *short GRBs*[2]. Of the selection searched for a CBC progenitor, 25 GRBs occurred during simultaneous operation of two detectors, with only GRB 090802A coincident with operation of the full three-detector LIGO-Virgo network.

7.4.2 Analysis summary

The S6 analysis was performed using the pipeline outlined in section 7.3. Analyses were coordinated in bulk once all GRBs targets were known, allowing time for pipeline development, including of the improved error regions for *Fermi* GRBs, and use of the most up-to-date DQ information.

---

[2]Of those targets analysed for CBC signals, 22 were also analysed by X-Pipeline.
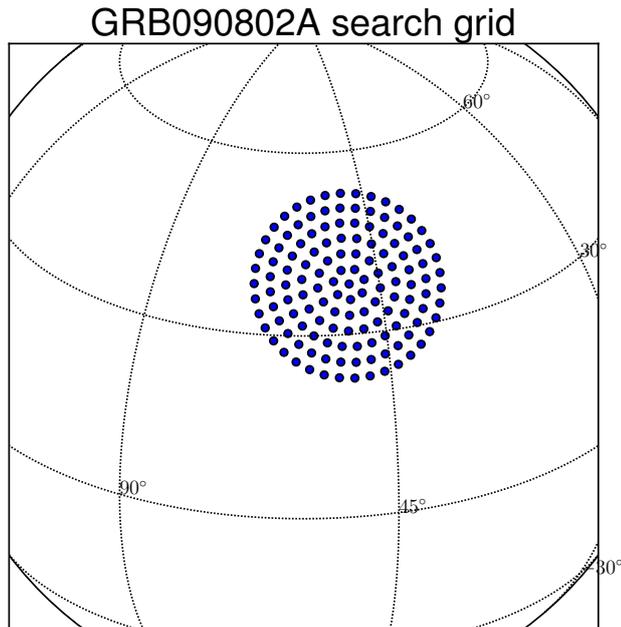
Figure 7.7: Initial search sky grid for GRB 090802A. Analysis of the full patch of 136 sky points would have been too computationally costly.

This search used a similar set of template waveforms to the all-sky coincidence-based analysis (section 6.4), namely non-spinning, 3.5 PN order, inspiral templates with component masses above $1\,M_\odot$ and total mass below $25\,M_\odot$.

A number of simulation campaigns were undertaken, testing the network sensitivity to both BNS and NSBH binary pairs. Given that short GRBs are likely to be beamed [168, 169, 170], each simulation type was divided into distributions with half-opening inclination angles of 10°, 30°, 45°, and 90°.

The 25 two-detector analyses used the sky tiling method set out in section 7.3.2, covering the sky region enclosed by the 90% confidence radius published by each satellite team. The same radius was used to generate the simulated signals, using a Gaussian distribution about the central point.

7.4.3   GRB 090802A

As noted above, one short GRB was detected during three-detector operations. GRB 090802A [171] was localised by the GBM instrument aboard the *Fermi* satellite with a 90% confidence radius of 14.8°. The sky tiling algorithm produced a grid as shown in figure 7.7, containing 136 sky points, presenting a prohibitively high computational cost.

As the analysis was being prepared, the *Fermi* GBM team published a revised estimate of the sky location, with a new triggered location almost
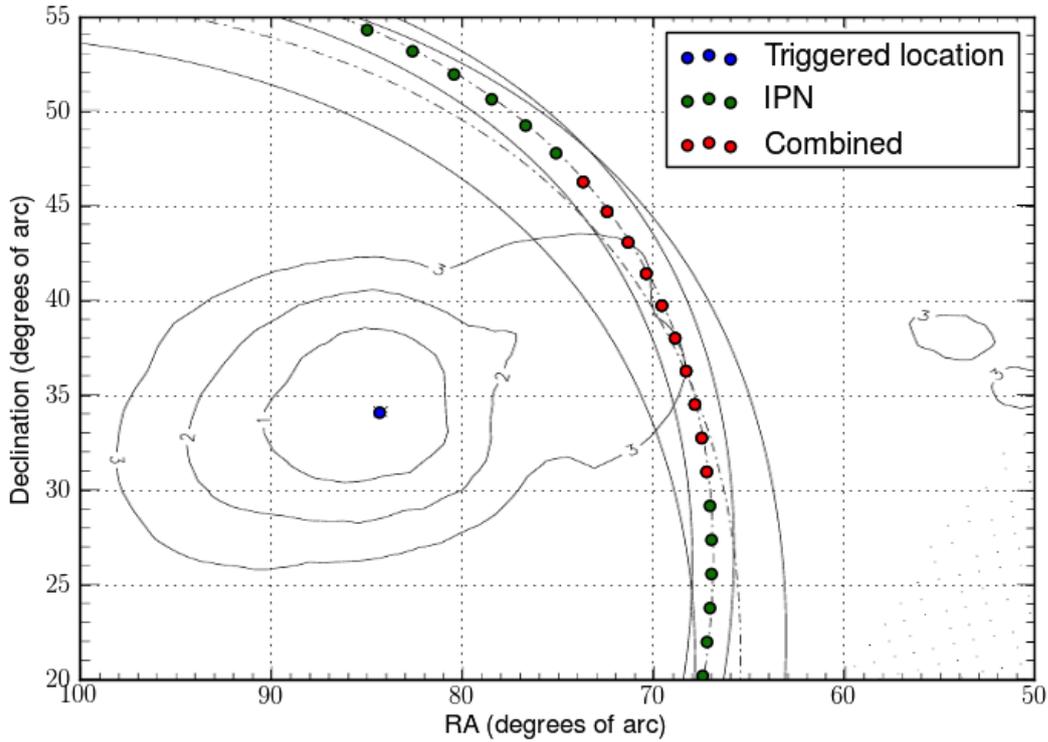
Figure 7.8: Combined IPN and *Fermi* localisation for GRB 090802A. The intersection of the large IPN contours and the updated 90% confidence interval from *Fermi* were used to place both search points and simulations. Here the blue dot represents the centre of the revised *Fermi* error region, while the green dots mark the IPN search grid in the case of an IPN-only analysis. The red dots then mark the search grid for the combined *Fermi*-IPN region. The original *Fermi* localisation peak can be seen on the far right of the figure as a 3-sigma confidence contour.

30° away from the original. The revised positioning did not come with a smaller error radius, however, so the decision was taken to combine this new information with that from the IPN's detection of the same event.

The IPN [152] uses triangulation from a network of detectors to locate a GRB to one or more ellipses on the celestial sphere. As shown in figure 7.8, the IPN contours overlapped with the updated 3-σ confidence contours supplied by the *Fermi* GBM team. The intersection of these two independent localisations of the same event allowed for a much narrower localisation band, used to seed both search grids and simulation grids. Using the same resolution as for the other GRBs, a search grid of ten points was generated manually, with this analysis configured independently of the others.

7.4.4 Data quality

In a similar manner to that of the all-sky search discussed in the previous section, DQ had a great impact on the coherent GRB analysis. As mentioned above, of the 407 GRB targets published between the start and end of S6, only 154 of those were analysed by either the CBC or GW Burst pipelines.

Of the total list, 242 triggers (59%) did not occur during stable, coincident operation of multiple GW detectors. From the summary of detector performance during S6 in chapter 4, this number is consistent with the overall duty cycle of the detectors, including large periods of no operation due to extended commissioning breaks, and the $\mathcal{O}(1)\%$ deadtime introduced by category 1 DQ vetoes.

From the remaining GRBs, only 3 were rejected due to poor DQ in the on-source segment. This is, again, consistent with the $\mathcal{O}(1)\%$ deadtime from category 2 flags throughout the analysis.

7.4.5 Search results

Each GRB was analysed independently, with the full results published, alongside those of the X-Pipeline search for GW bursts, in [19]. No associated GW detections were made by either search, with distance exclusion limits set for each GRB.

Figure 7.9 shows the 90%-confidence exclusion distances set for each GRB by the coherent CBC signal search. The distribution of exclusion distances is determined by the detector sensitivity and the detector response for each sky location. For BNS, the exclusions sit around $10-20\,\mathrm{Mpc}$, in agreement with the angle-averaged detector range throughout the runFor NSBH, the exclusions are higher (recall figure 6.5), with the histogram peaked sharply in the $30-40\,\mathrm{Mpc}$ distance bin.

In addition to per-GRB distance exclusions, limits were placed on the distribution of BNS and NSBH merger events (under the assumption that all analysed GRB targets had CBC progenitors). Figure 7.10 shows this result for each type, and an extrapolated distribution for aLIGO (at design sensitivity). With aLIGO sensitivity ($\sim$ 2018 [64]), detection of a GW counterpart to a GRB detection is possible, but even a non-detection will place constraints on GRB population models.
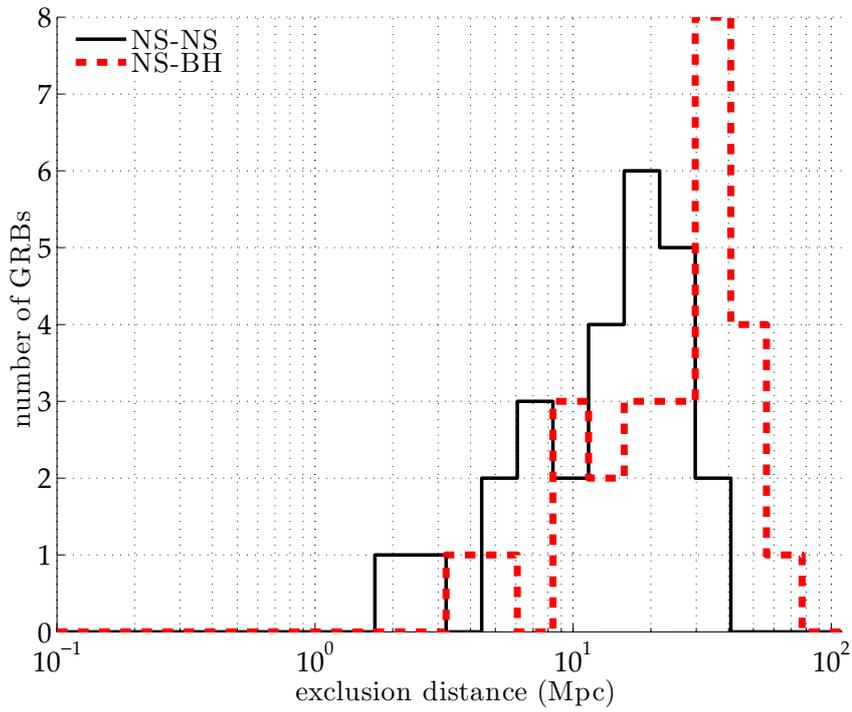
Figure 7.9: Histogram of the 90%-confidence distance exclusions set by the search for GWs from CBC associated with GRBs [19]. These data used simulations of neutron stars with masses Gaussian distributed around $1.4\,\mathrm{M}_\odot$ and black holes around $10\,\mathrm{M}_\odot$, and a GRB jet opening angle of 30°.
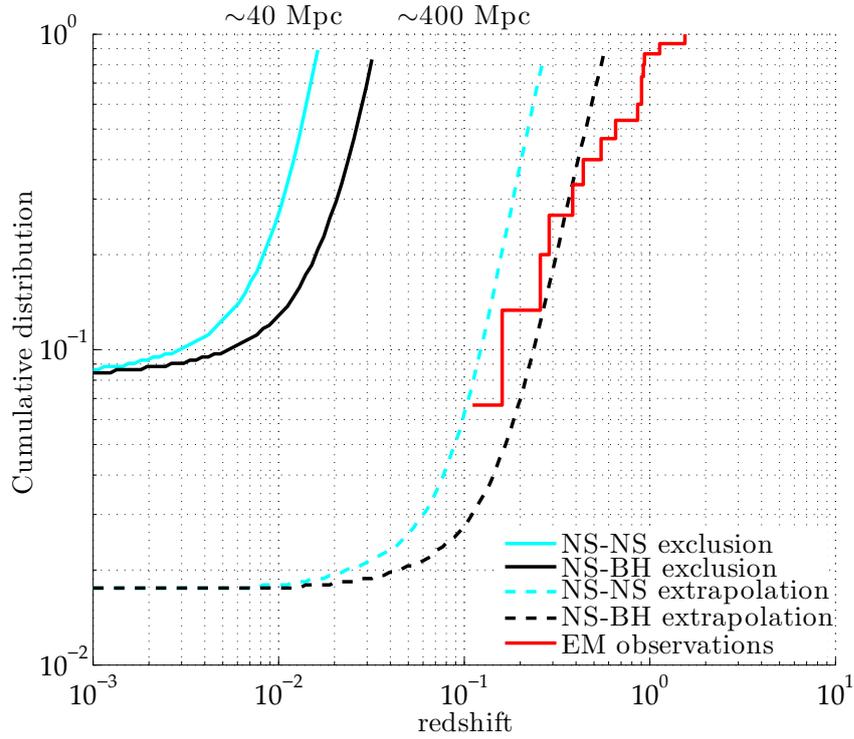
Figure 7.10: Cumulative redshift exclusion distribution from the search for GWs from CBCs associated with GRBs [19]. The extrapolated curves assume aLIGO operating at design sensitivity, while the red stepped curve shows the distribution from *Swift* GRB observations with known redshift.

Chapter 8

An all-sky, all-time, coherent search

In this chapter, we describe the first implementation of an all-sky, all-time coherent search for gravitational waves from compact binary coalescences, tying together many of the key ideas from each of the previous two chapters. In sections 8.1 and 8.2 we describe methods for tiling the sphere of source sky locations for two- and three-site networks, each based on a uniform tiling in cross-detector time-delays. Section 8.3 outlines the methods for a fully coherent analysis pipeline, using an example two-site analysis of data from S5. We compare new results with those from the coincidence-based pipeline.

8.1   Sky tiling for a two-site network

During the first year of LIGO Science Run 5 (S5), November 2005 – November 2008, the three LIGO detectors formed a two-site, three-detector network, with both LHO instruments taking part in the run. For this network, there exists only a single time-delay baseline between the observatories, restricting the accuracy of sky localisation [130]; any localisation by triangulation will reconstruct only the difference in time of arrival between the sites, mapping to a ring of source direction probability on the celestial sphere.

As discussed in the previous chapter, for a two-site network, lines of source sky positions running perpendicular to the baseline between sites all result in the same arrival time difference for a signal incident on the network. As a result, the most efficient tiling of the sky sphere for a two-site 'all-sky' search is given by the one-dimensional space of physically allowed time-delays.

While it is true that each of those sky locations degenerate in time delay will
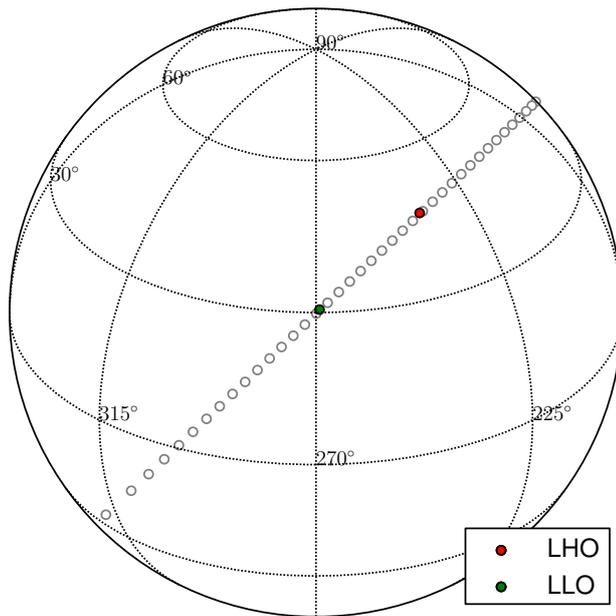
Figure 8.1: A two-site all-sky grid for the LIGO detector network. The point span the allowed time-delays between sites, forming a great circle passing overhead both.

have unique values of the detector response, the amplitude error introduced by this uncertainty averages over the sky and does not introduce an overall systematic error.

The allowed time delays are bounded by the light travel time, T, between sites; if we choose a temporal resolution δt, then the size of the grid is

$$N = \left\lfloor \frac{T}{\delta t} \right\rfloor. \tag{8.1}$$

The light travel time between LHO and LLO is $T = 10\,\text{ms}$, so, for a resolution of $500\,\mu\text{s}$, the two-site LIGO grid requires 40 sky points.

To seed this grid, we lay points with time delay τ in the range $[-T, T)$, evenly spaced by $2\delta t$, and project these onto the prime meridian (zero longitude) via,

$$\phi = 0 \tag{8.2a}$$

$$\theta = \cos^{-1}\left(\frac{\tau}{T}\right) \tag{8.2b}$$

This ring is then rotated onto the great circle containing both sites, as shown in figure 8.1. The grid then has greatest density orthogonal to the inter-detector baseline, where the time-delay is smallest, with density dropping symmetrically

in either direction, as the time-delay grows.

8.2   Sky tiling for a three-site network

During the second year of S5 and throughout S6, the LIGO-Virgo detectors operated a three-site network, allowing much more accurate time-delay triangulation, and better sky localisation [172]. This is likely to be the same network running during the middle years of the advanced detector era, after both Advanced LIGO and Advanced Virgo is operational, but before any other detectors are ready.

   With three or more detectors, the network can triangulate any signal to a single point in the hemisphere above the plane of the network, but a symmetry still exists in that plane, producing a second point in the other hemisphere.The unfortunate downside of better sky localisation is the requirement for much larger sky maps for a full coherent analysis, vastly increasing the computational cost of the search.

   In order to map the sky for three sites, we follow the analytical models of [167]; consider a network of D detectors, and define the time-delay vector

$$\boldsymbol{\tau} = [\tau_2, \ldots, \tau_D]^T,$$ (8.3)

where τ_n is the arrival time difference between detector 1 and detector n. If T_m is the light travel time between detector 1 and detector m, and we define α_{mn} as the angle separating the baselines between detector 1 and detectors m and n respectively, we can construct a bounding ellipse for the physically-admissible time delays,

$$\boldsymbol{\tau}^T \mathbf{A}_D \boldsymbol{\tau} \leq B_D,$$ (8.4)

where

$$\mathbf{A}_3 = \begin{bmatrix} \frac{T_3^2}{T_2^2} & -\frac{T_3}{T_2}\cos\left(\alpha_{23}\right) \\ -\frac{T_3}{T_2}\cos\left(\alpha_{23}\right) & 1 \end{bmatrix}$$ (8.5a)

$$B_3 = T_3^2 \sin^2\left(\alpha_{23}\right).$$ (8.5b)

Here, $\cos\alpha_{mn}$ effectively measures the correlation between the two time delays, where in the limit $\cos\alpha_{mn} = \pm 1$ (the three sites lie in a line) the time-delays are degenerate and the triangulation is no better than a two-site network. In the limit $\alpha_{mn} = \pi/2$, the time-delays are independent and the time-delay
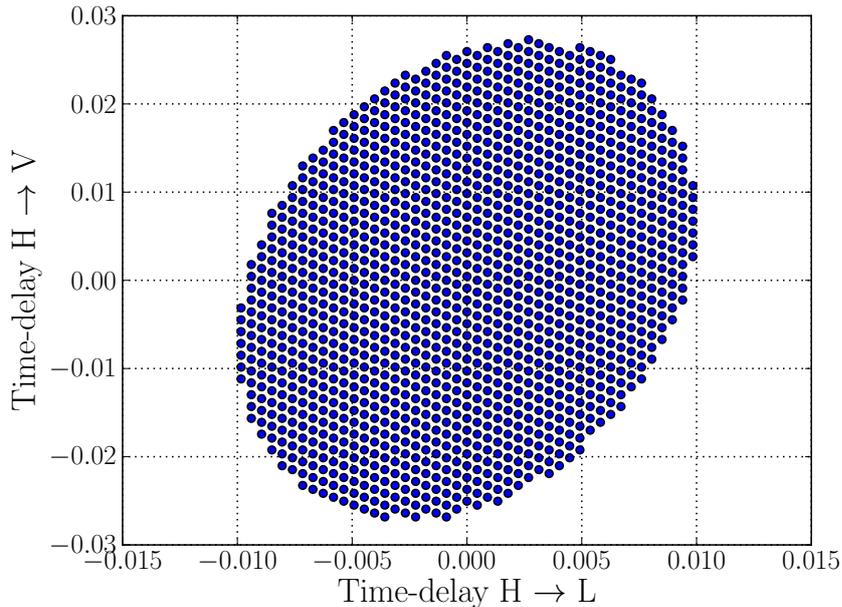
Figure 8.2: Time-delay tiles for the LIGO-Virgo three-site network. All phys-
ically admissible points in this space are laid in a hexagonal grid,
with a minimal 5μs spacing between neighbouring points.

baselines form the major and minor axes of the bounding ellipse.

Using the above conditions, a grid of hexagonal time-delay tiles is laid in
in (τ_2, τ_3) coordinates, such that the distance between any two points matches
the desired resolution. The resulting time-delay grid for the three-site LIGO-
Virgo network is shown in figure 8.2. This grid uses the same 0.5 ms resolution
as considered in chapter 7.

The time-delay grid is then projected onto the celestial sphere by construct-
ing a detector network coordinate system, as shown in figure 8.3, where the
time-delay coordinates are related to network longitude, ϕ, and latitude θ, via

$$\phi = \pm \cos^{-1}\left(-\frac{T_2\tau_3 - T_3\tau_2\cos(\alpha_{23})}{T_3\sqrt{T_2^2 - \tau_2^2}\sin(\alpha_{23})}\right), \tag{8.6a}$$

$$\theta = \cos^{-1}\left(-\frac{\tau_2}{T_2}\right). \tag{8.6b}$$

We can then relate the network coordinates (ϕ, θ) to earth-fixed longitude and
latitude with a simple rotation. This projection is done twice, once for each
hemisphere above and below the plane of the detector network.

Figure 8.4 shows an all-sky grid for the LIGO-Virgo network; 8.4a views
the grid from nearly overhead the plane of the network, where the grid is dens-
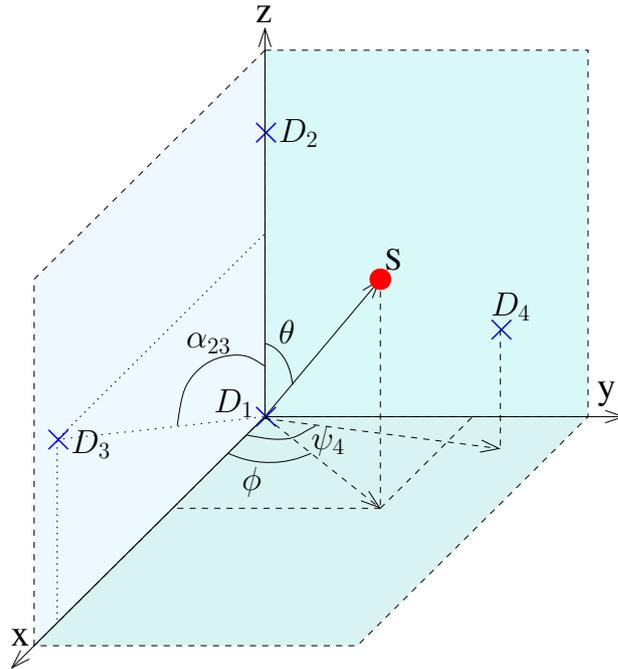
Figure 8.3: The network coordinate system used in projecting points in time-delay space onto the sky [167]. A three-site network defines a right-handed coordinate system, with a potential fourth-detector breaking the symmetry $x - z$ plane.
for a source on the sy

est, while 8.4b shows the relative sparsity of the grid parallel to the network plane. This grid, using a time-delay resolution of $\delta t = 0.5\,\mathrm{ms}$, contains over 2,700 points, representing a huge computational cost if applied naïvely to the coherent analysis.

8.3 Search implementation and testing

With the above methods for mapping all possible time-delays onto the sky for a given network, we can use the coherent matched-filter described in chapter 7 to construct an all-sky, all-time search. In this section we demonstrate the first implementation of a fully-coherent, all-sky search for signals from binary neutron star inspirals.

We analysed one month of data from S5, during which the three LIGO instruments formed a two-site, three-detector network, allowing a complete test of an all-sky coherent analysis with relatively small computational cost. This analysis, searching a reduced portion of the mass parameter space (compared to the previously published coincidence-based analysis described in chapter 6),
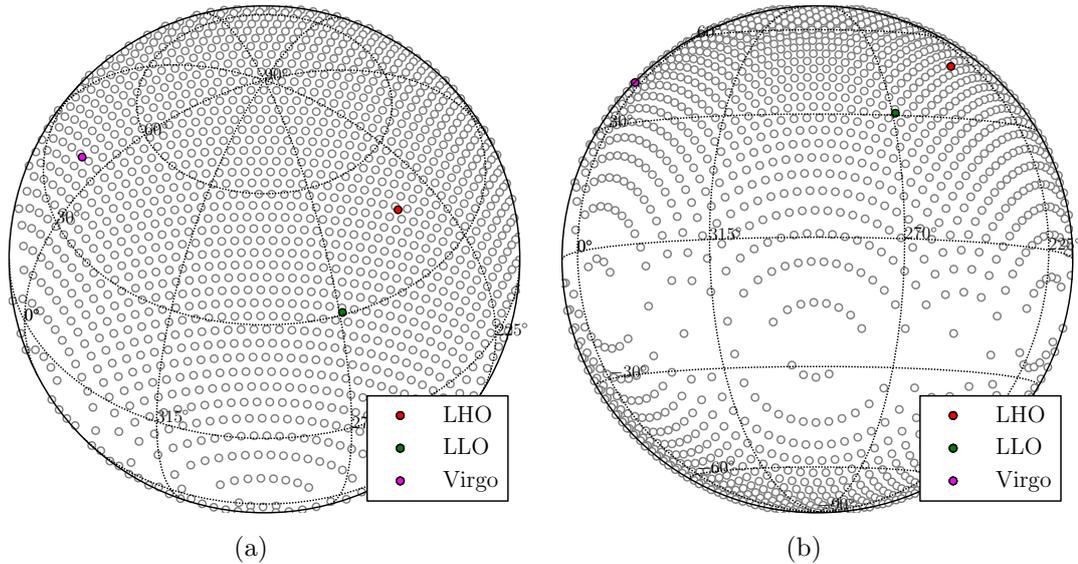
Figure 8.4: The three-site all-sky coherent search sky grid for the LIGO-Virgo
detector network. (a) and (b) view the same grid from different
angles.

and using only a two-site network, serves as a proof-of-principle demonstration
that a coherent analysis can be constructed and used to search current and
advanced detector data.

8.3.1 Data selection

The multi-detector coherent search is only run over those times when all de-
tectors in the network are operating, using category 1 vetoes to remove times
of sub-standard DQ for each (recall section 4.5 for more details). The detector
PSD is measured using 16 half-overlapping 256-second blocks, so all remaining
segments are down-selected to those longer than 2176 seconds; figure 8.5 shows
the segments selected for a month of data during S5. Of all available time,
48.9% falls in three-detector coincident operation (after applying category 1
vetoes), with 45.3% acceptable for analysis based on segment duration.

Figure 8.6 shows the detection range (averaged over source location and
orientation) for each detector in the network during this period. As seen, the
smaller H2 detector maintained a range between 6–7 Mpc throughout, while
the larger instruments improved as the run progressed, with H1 peaking at
16 Mpc.
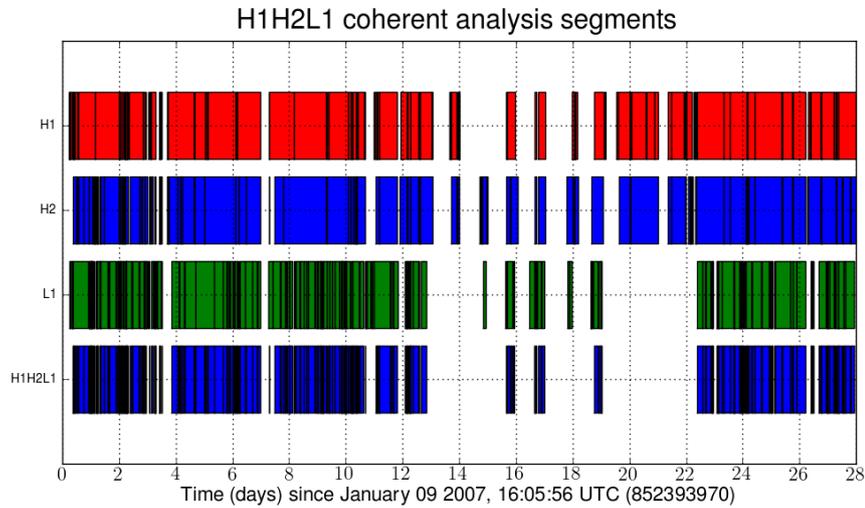
Figure 8.5: Analysis segments for coherent all-sky search of one month of S5 data.
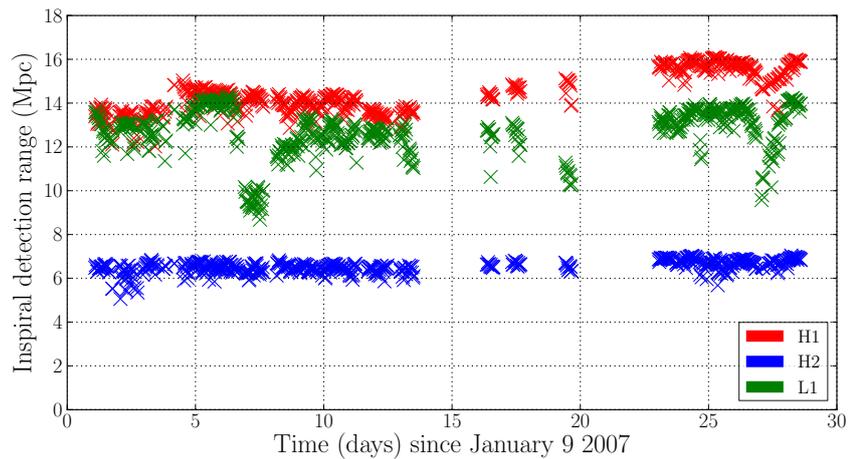
Figure 8.6: Inspiral averaged sensitive range for the LIGO network during one month of S5.
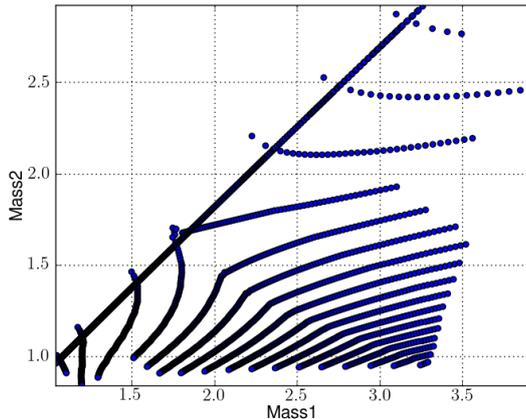
Figure 8.7: The template bank used for the coherent all-sky search of one month of S5 data.

8.3.2 Generating a template bank

A template bank is laid in the same method as used for the targeted search from section 7.3, with a single-detector bank for the L1 detector used as a simple approximation to a fully-coherent template bank. For the test analysis, non-spinning, 3.5 PN-order, BNS inspiral waveforms were placed with minimum component mass of $1\,\mathrm{M_\odot}$ and maximum combined mass of $6.2\,\mathrm{M_\odot}$, resulting in $\sim 2,200$ templates, as shown in figure 8.7.

8.3.3 Event trigger generation

For computational ease, each analysis segment is divided into chunks of 2176 seconds, overlapping by 64 seconds on each end; a template bank is generated for each of these chunks (due to changing detector sensitivity), with each chunk processed separately. Analogous to the targeted search associated with GRBs, the data from each detector are used to measure its PSD, using the average of sixteen 50%-overlapping 256-second sensitivity estimates.

For each of the 256-second blocks, the data from each detector are filtered against each template in turn, producing time-series of single-detector SNR. Each sky location in the map is used then to construct coherent combinations of these time-series to calculate the coherent SNR, with each of those peaks in coherent SNR recorded as a trigger.

Background estimation with time-slides

The background of noise is measured using time-slid data. As mentioned in chapter 7, since a fully-coherent search combines detector data at the filtering stage, each time-slide analysis requires re-computing the coherent matched-filter SNR. This represents a huge computational cost, in comparison to the coincidence-based analysis where time-shifts are performed on the single-detector triggers produced after filtering.

In this analysis, we construct 10 time-shifts, each sliding data from the L1 detector forward by a multiple of 128-seconds. Such large slides are computationally simple and performed by combining data from the LHO instruments with those from L1 in a different 256-second block. The slides are performed on a ring (formed by the sixteen analysis blocks for a single chunk), whereby any L1 data slid off the end of the analysis chunk is re-inserted at the start and filtered against the data from the LHO instruments in the first block.

Simulations

A set of non-spinning BNS inspiral simulations were used to inform tuning of the signal-consistency cuts and the detection statistic, and measure search performance through simulation recovery efficiency. These were injected into the data with simulated distances uniformly distributed between $1 - 60$ Mpc.

Throughout the following descriptions of signal-based and data quality cuts, we use the simulation results as a measure of the separation of foreground from background.

8.3.4   Event down-selection

The rate of events produced during the data processing is too high to be either useful or practical – as seen in chapter 4, large glitches will trigger across the template bank, producing multiple events from a single noise input.

Events are down-selected by identifying those most significant relative to the surrounding data. The full list of events are divided into 100 ms bins, with an event selected only it is the loudest in its own bin, and louder than all events in a 100 ms window around itself. This selection method typically reduces the event rate by a factor of 100 or more, by identifying those triggers that represent the peak of an excess power transient in the data (either noise or GW signal).

Signal-consistency cuts

Each of the signal-consistency tests outlined in section 7.3 are applied equally to the triggers from the un-slid foreground, each of the time-slide backgrounds, and the simulations. The following criteria are applied during the search:

- single-detector SNR ≥ 5 in one detector, and ≥ 4.5 in another

- coherent SNR ≥ 7,

- spectral χ^2 re-weighted SNR > 7,

- bank χ^2 re-weighted SNR ≥ 7,

- auto-correlation χ^2 re-weighted SNR ≥ 7,

- null SNR < 6

The thresholds used for these cuts are more aggressive than those used for the targeted search, in which single-detector SNR thresholds of 4 and re-weighted SNR thresholds of 6 were applied. The larger parameter space, including longer duration, and all-sky coverage, place less restrictions on the background, potentially reducing search sensitivity, and so stricter cuts are required to reduce the volume of triggers to a manageable level.

Additionally, the background distribution for the all-sky search more closely resembles that of the coincidence-based analysis from chapter 6, compared to the targeted search of chapter 7, and so the original tuning for the the χ^2 re-weighting was adopted:

$$
\rho_{\mathrm{new}} = \begin{cases} \rho, & \chi^2 \leq n_{\mathrm{dof}} \\ \dfrac{\rho}{\left[\left(1 + \left(\frac{\chi^2}{n_{\mathrm{dof}}} \right)^3 \right) / 2 \right]^{1/6}}, & \chi^2 > n_{\mathrm{dof}}. \end{cases}
$$

The first cut, on single-detector SNR, identifies those events with power in a single detector only, typically removing more than 90% of all events. Figure 8.8 shows the impact of this cut on the events from the S5 analysis, applied to both H1 and L1 data, including both the background from time slides and from the simulation set. Those background triggers (blue dots) on the diagonal in each figure represent events with power only in that detector, and fail the cut (black line) in the other detector. The simulated signals are
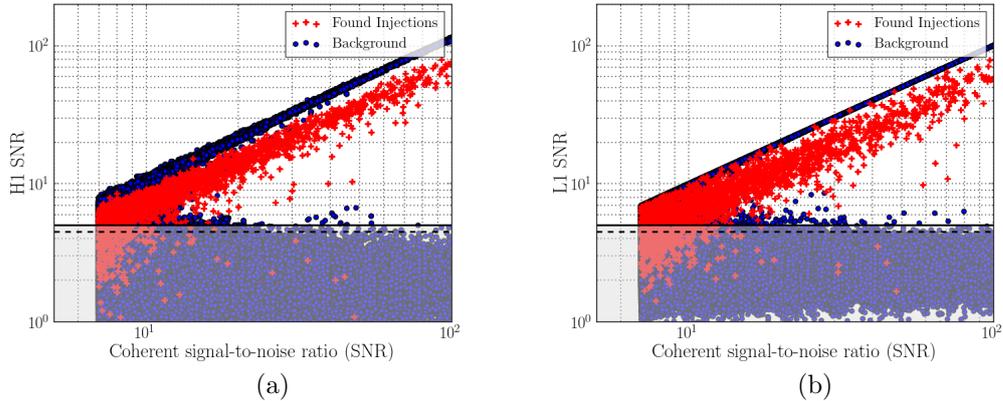
Figure 8.8: The impact of the single-detector SNR cut on events from an all-sky coherent search. The blue dots are those from the noise background, while the red pluses are those from simulated BNS signals. The shaded region represents the single-detector cut as applied. All events with power only in a single-detector are vetoed as likely noise artefacts.

below the diagonal because their coherent SNR is accumulated from a strong component in each detector.

Figures 8.9a to 8.9c show the impact of the three χ^2 consistency cuts, each evaluated after the single-detector SNR cut has been applied. These statistics clearly differentiate between the recovered simulations and the noise background, removing those events inconsistent with a true signal. Figure 8.9d shows the impact of the null SNR cut, similarly evaluated after the single-detector SNR cut has been applied. The null SNR cuts a relatively small number of noise events that are incoherent between detectors, but down-ranks a large number of others to a lower value of the detection statistic.

Data quality cuts

DQ vetoes are also used to identify noise artefacts in the data, using instrumental and environmental correlations, rather than signal-consistency. Category 2 and 3 vetoes are applied to data from each of the foreground, background and simulations such that if an event vetoed by any instrument is removed from the search. This should be compared to the coincidence-based analysis, whereby any event vetoed in a single-detector is still recorded as a two-detector event. Such events in the coherent search should be recorded and processed separately, but we leave such an extension to future work.

Figure 8.10 shows the impact of category 2 and 3 DQ vetoes on the back-
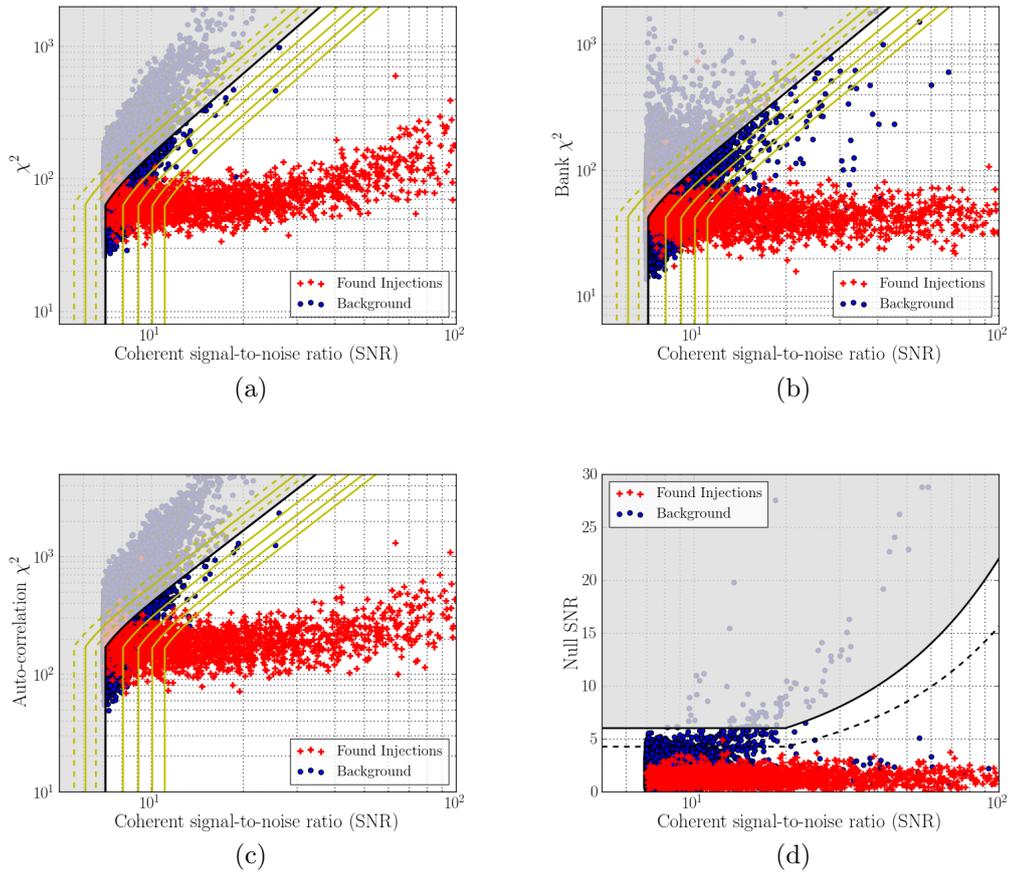
Figure 8.9: The impact of signal-consistency cuts on the background of an all-sky coherent search. The blue dots are background events from time-slid analysis, and the red pluses are events from BNS simulations, and the shaded region covers those events failing the signal-consistency test. In figures (a) to (c), the contours represent constant values of the re-weighted SNR (dashed lines for half-integers, solid for integers). The detection statistic is constructed from the χ^2 re-weighted SNR, and the null SNR, figure (d), for which events above the dashed line are down-ranked using equation (7.4).
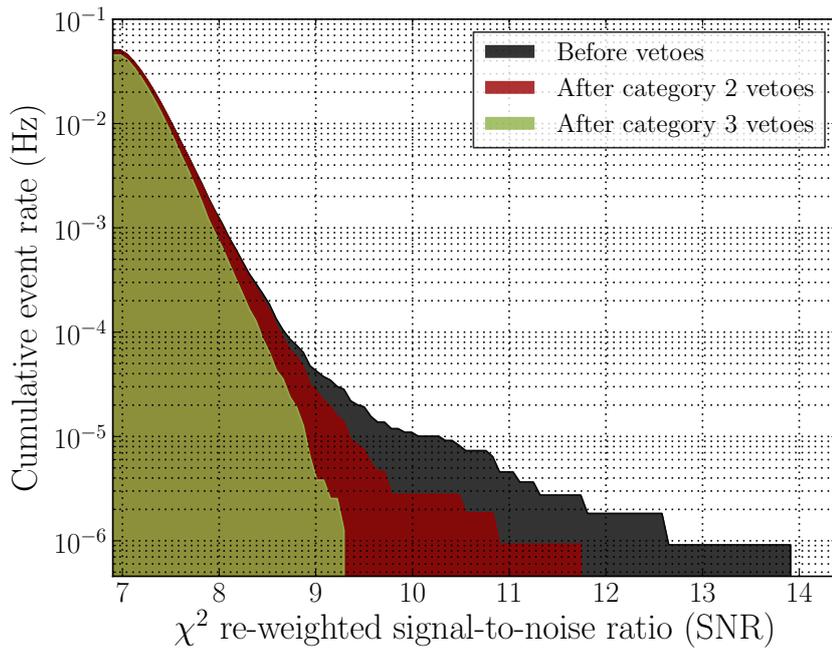
Figure 8.10: The impact of data quality vetoes on the background of an all-sky coherent search. As seen, the loudest event is reduced from a re-weighted SNR of 13.9 to 9.3.

ground events (from time slides), after the application of the signal-based vetoes. We use the same detection statistic as the previous coherent search, defined in equation (7.4). The category 2 vetoes are successful in removing the very loudest events, with the loudest event reduced from 13.9 to 11.8, with only 1% deadtime. Category 3 introduces a large amount of deadtime, above 26%, but eliminates the remaining tail of significant background events, with the loudest event reduced to a detection statistic value of 9.3.

We can qualify the overall impact of the combined signal-based and data quality vetoes, and the effectiveness of the chosen detection statistic, by comparing the high coherent SNRs seen in figure 8.8 to the final distribution in figure 8.10.

8.3.5   Search performance

The performance of the search is measured using the results of the simulation run, after all signal-based and data quality cuts have been applied. All simulations for which no event was recorded are classed as missed. Those simulations with an associated trigger that is ranked higher than all of the background events are classed as recovered, while those events with an associated triggers that is not louder than all background events are marginally
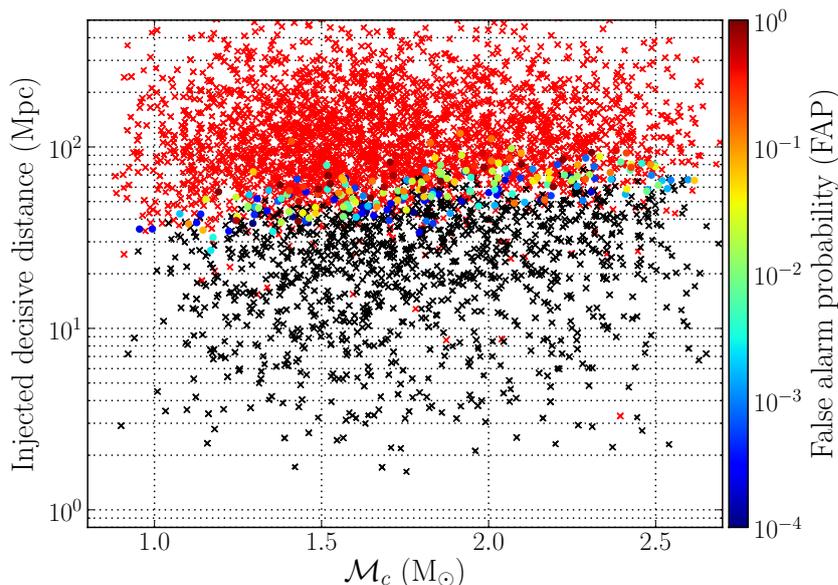
Figure 8.11: Recovery of simulated BNS signals during the S5 coherent all-sky analysis. Successfully recovered signals as marked with black crosses, missed signals with red crosses, and marginally-recovered signals with dots coloured by FAP. The 'decisive' distance is the second-largest effective distance for the network, as detailed in the text.

recovered.

The recovery of injections as a function of distance allows calculated of the volume sensitivity of the search, in order to calculate event rate upper limits. All results are constructed after application of the same signal-consistency and DQ cuts to each simulation run, as well as the background data.

Figure 8.11 shows the distribution of injections and their recovery as a function of the injected decisive distance and simulation chirp mass. The decisive distance is defined as the second largest effective distance for the detectors in the network, where the effective distance for a detector is the distance at which an optimally-oriented and located signal would have the produced the same SNR as the given simulation. This value gives a measure of how far in distance the network could have been expected to recover this event, given a requirement of power in at least two detectors.

At lower masses, the majority of simulations injected below $30\,\mathrm{Mpc}$ ($\sim 13\,\mathrm{Mpc}$ angle-averaged range) are successfully recovered, consistent with the network sensitive distance during S5 (figure 8.6)[1], with recovery improving as mass increases (recall figure 6.5). With the background highly cleaned by

---

[1]The angle-averaged range shown in figure 8.6 is calculated for a BNS with mass $m_1 = m_2 = 1.4\,\mathrm{M_\odot}$, for which $\mathcal{M}_c \simeq 1.2$.

the myriad cuts and vetoes, resulting in a low-significance loudest event, very few simulations are marginally recovered, with the transition rapidly made to completed missed signals at higher distances.

8.3.6   Comparison with the coincidence-based pipeline

In order to compare search performance against a benchmark, we have analysed the same data segments in S5 with the coincidence-based 'ihope' pipeline [102]. The same data quality veto method was used, whereby those events vetoed in a single detector are removed from the search. Finally, the search was performed using identical template parameters and simulation parameters, allowing a direct comparison of search efficiency with the new coherent pipeline.

Computational cost

Each of the coherent and coincident algorithm are limited by the speed of the fast Fourier transform (FFT) – the computational core of the matched-filter – and so with the inclusion of multiple sky points, and CPU-expensive calculations of the χ^2 statistics, the coherent algorithm is around a factor of 2 more expensive than its coincidence-based predecessor. Extension to the full three-site, all-sky grid will increase computational cost marginally, requiring more calculations but crucially no further FFT-based filtering.

However, with the current implementation of the background trials, each time-slide is computationally equivalent to the zero-lag foreground, resulting in a significant relative increase in computational cost for this new search. Recent developments in this method will reduce the cost of each time-slide, allowing ~ 30 unique background trials for the same cost as the zero-lag foreground, meaning a full detection search including $1,000$ trials would result in around a $\times 30$ computational cost in moving from coincidence-based to coherent.

Signal recovery

Figure 8.12 compares the efficiency of simulation recovery between the two analyses. The top and centre panels compare the injection recovery and cumulative volume respectively as functions of distance between the two searches; the bottom panel shows the relative volume improvement of the coherent search. The small reduction in sensitivity at low distances is within statistical errors, however, the near 25% improvement in sensitive volume over the
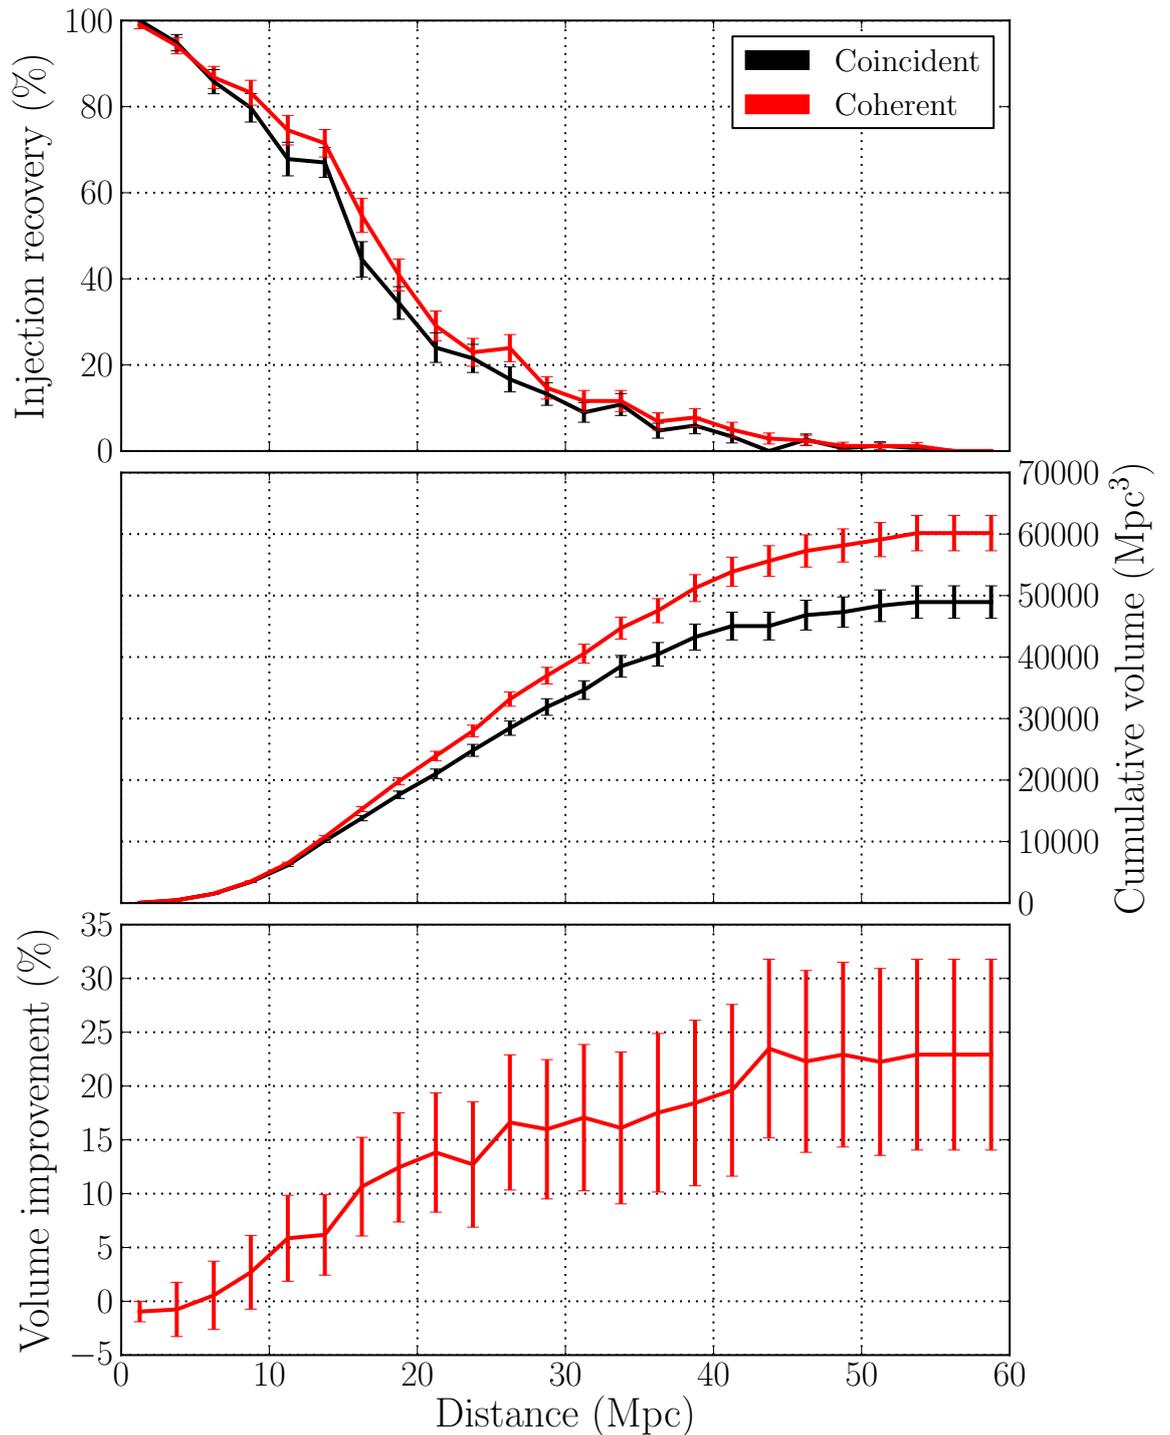
Figure 8.12: Comparison of search performance between the coincident and coherent all-sky searches for one month of S5 data. The top panel shows the injection recovery (efficiency) as a function of distance, while the second panel shows the the cumulative volume, comparing between the coincident (black) and coherent (red) searches. The bottom panel shows the relative volume improvement of the coherent search. The coherent search outperforms the coincident with nearly 25% greater volume sensitivity.

full simulation campaign highlights the advantage of this coherent pipeline compared to the previously published algorithm.

8.4 Discussion and future direction

In this chapter we have demonstrated the first implementation of a fully-coherent all-sky search for gravitational waves from the inspiral of two compact objects. Sections 8.1 and 8.2 detailed methods for tiling the physical space of time-delay for two-site and three-site networks respectively, while section 8.3 showed a proof-of-principle test of the search implementation.

The fully-coherent, two-site search was seen to improve sensitive volume by nearly 25% over a coincidence-based search of the same data. The detection rate estimates presented in [52] are associated with the sensitivity of the full coincident search of S5 data [135], predicting a rate of ~ 40 BNS detections per year with a design-sensitivity second-generation detector network; the above results show that using an all-sky coherent search over the same data could result in almost 50 detections per year.

Given the improvement seen in this example, it can be expected that a fully-coherent analysis using a three-site (or more) detector network will provide even greater improvements in sensitivity, as the power of the null stream increases, and a truly all-sky tiling method is used. With a two-site network likely to operate in the early years of the second-generation detector network, this implementation is applicable to those data, however the computational cost of a three-site search with a large background sample is currently very expensive; development is on-going to reduce the relative cost of coherent time-slides, as noted earlier. Furthermore, in order to construct a complete analysis pipeline, further testing will be required, including tuning of the detection statistic and signal-based vetoes over a larger mass parameter space and a larger background distribution.

While the above results are a preliminary statement only, the potential of a fully-coherent search for GWs from CBCs has been demonstrated, and presents an exciting path for future development.

Chapter 9

Conclusion

The first generation of gravitational wave detectors have been decommissioned after almost ten years of operation. While no direct detections were recorded, the knowledge gained from constructing and commissioning these instruments, and analysing the data they have produced, will be invaluable for the LIGO and Virgo collaborations as they work towards a second-generation network. Throughout this work we have concentrated on using the detector characterisation and data analysis skills developed during the last joint science run (S6/VSR2/3) to improve the sensitivity of searches over the data recorded by laser-interferometer detectors.

In part II we reviewed the operation and performance of the LIGO detectors during their sixth science run. This run saw the instruments at their most sensitive level, with a number of upgrades relative to S5 prototyping new and improved hardware looking ahead to Advanced LIGO. We showed the improvement in search sensitivity purely as a result of instrumental commissioning throughout the run (section 4.3), and as a result of the application of data quality vetoes in transient searches (section 4.5). While instrumental commissioning will play the crucial role over the next five years of developing the second-generation interferometers, DQ studies and veto methods will definitely be key to maximising the search output from the short science runs expected to begin in 2015.

The targeted veto method presented in chapter 5 showed the remarkable performance available from DQ vetoes if the auxiliary data are properly studied and understood, with $\sim 60\%$ background trigger elimination in $\sim 6\%$ deadtime at LHO. This improved identification of seismic noise bursts as a result of detailed tuning for low-frequency analysis demonstrated what could be possible if such efforts are extended to all interferometer systems over the coming years.

Characterisation of the aLIGO instruments is already underway at each observatory, with the input laser and optical systems installed and in commissioning. The lessons learned from the studies of S6 data will be used to ensure maximum performance of each of the instrumental sub-systems as they are installed and tested, to allow for high sensitivity as early in the advanced science runs as possible. We are confident that extension of the targeted noise event identification to all interferometer components will improve search sensitivity through lower deadtimes and better cleaning of the search data.

In part III we have discussed the core analyses of detector data searching for GW signals from compact binary coalescences. A coincidence-based analysis was used throughout the first-generation science runs as the primary method for all-sky, all-time searches, and for targeted searches up to and including S5. Over the course of the S6, this all-sky pipeline was greatly improved through better attention to DQ information, and development of a χ^2-weighted detection statistic. With these in place, the upper limits set from this analysis were improvements over S5, with the expectation that, as the instrumental sensitivity grows over the next few years, upper limits will begin to probe rates set by theoretical and observational estimates.

We have also reviewed the coherent matched-filter method, and described its adoption for the search associated with short GRBs observations. The improvements to this search over the S5 method, including mapping of the probability sky patches published by the Fermi satellite detections, offer a firm basis from which to build in directing second-generation search algorithms.

Finally, the first implementation of a fully-coherent, all-sky, all-time search was presented in chapter 8. This development has followed naturally from those of the coincidence-based all-sky pipeline, and the coherent targeted search, with a first analysis performed on S5 data using a two-site detector network. We have shown, on this one-month test using BNS inspiral signals, an improvement of more than 20% in search sensitive volume over the coincidence-based analysis run on the same data. It should be expected that, setting computational cost aside, a fully-coherent analysis with a three-site network will show even greater performance, with improved sky localisation, and a more powerful null stream consistency test.

While a coherent analysis with a large background sample is currently very costly, work is ongoing to improvement the efficiency of the computational infrastructure to minimise cost. Implementation and testing of such a search must be the goal ahead of completion of the Advanced Virgo interferometer, in

order to maximise the science goal in the early years of the second-generation science runs, when detections are likely to be made, but not in great numbers. The potential for increasing detection rate by 30% or more should be a driving force in development of this analysis.

Progress made over the next few years will define our field for the next decade. The first-generation experience has shown that direct detection gravitational waves is imminent, and all efforts must be made to commission and characterise not only the new instruments, but also new search methods. If these efforts are made, direct detections will be made by the second-generation network, allowing gravitational wave physics to become part of a multi-messenger era for astronomy.

Bibliography

[1] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Ann. Phys.*, 354(7):769–822, 1916.

[2] A. Einstein. Zur Elektrodynamik bewegter Körper. *Ann. Phys.*, 322(10):891–921, 1905.

[3] J. Weber. Detection and Generation of Gravitational Waves. *Phys. Rev.*, 117(1):306–313, January 1960.

[4] J. Weber. Observation of the Thermal Fluctuations of a Gravitational-Wave Detector. *Phys. Rev. Lett.*, 17(24):1228–1230, December 1966.

[5] J. Weber. Gravitational-Wave-Detector Events. *Phys. Rev. Lett.*, 20(23):1307–1308, June 1968.

[6] J. Hough, et al. Proposal for a Joint German-British Interferometric Gravitational Wave Detector. Technical Report MPQ 147, April 1989.

[7] K. Danzmann, et al. Proposal for a 600 m Laser Interferometric Gravitational Wave Antenna - GEO 600. Technical Report MPQ 190, March 1994.

[8] B. Willke, et al. The GEO 600 gravitational wave detector. *Class. Quantum Grav.*, 19(7):1377–1387, March 2002.

[9] K. Tsubono. 300-m laser interferometer gravitational wave detector (TAMA300) in Japan. In *Proceedings of the first Eduardo Amaldi Conference on Gravitational Wave Experiments*, pages 112–114, Roma, June 1994.

[10] M. Ando, et al. Stable Operation of a 300-m Laser Interferometer with Sufficient Sensitivity to Detect Gravitational-Wave Events within Our Galaxy. *Phys. Rev. Lett.*, 86(18):3950–3954, April 2001.

[11] A. Abramovici, et al. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256(5055):325–333, April 1992.

[12] C. Bradaschia, et al. The VIRGO Project: A wide band antenna for gravitational wave detection. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 289(3):518–525, April 1990.

[13] B. P. Abbott, et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Rept. Prog. Phys.*, 72(7):6901, July 2009.

[14] H. Grote and LIGO Scientific Collaboration. The status of GEO 600. *Class. Quantum Grav.*, 25(11):114043, May 2008.

[15] F. Acernese, et al. Status of Virgo. *Class. Quantum Grav.*, 25(11):114045, May 2008.

[16] K. Kuroda and the LCGT Collaboration. Status of LCGT. *Class. Quantum Grav.*, 27(8):084004, April 2010.

[17] The LIGO Scientific Collaboration and The Virgo Collaboration. Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO's Sixth and Virgo's Second and Third Science Runs. *ArXiv e-prints*, March 2012.

[18] J. Abadie, et al. Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3. *Phys. Rev. D*, 85(8):82002, April 2012.

[19] J. Abadie, et al. Search for Gravitational Waves Associated with Gamma-Ray Bursts during LIGO Science Run 6 and Virgo Science Runs 2 and 3. *ApJ*, 760(1):12, November 2012.

[20] J. Aasi, et al. Search for gravitational waves from binary black hole inspiral, merger, and ringdown in LIGO-Virgo data from 2009–2010. *Phys. Rev. D*, 87(2):022002, January 2013.

[21] J. Abadie, et al. All-sky search for gravitational-wave bursts in the second joint LIGO-Virgo run. *Phys. Rev. D*, 85(12):122007, June 2012.

[22] B. P. Abbott, et al. Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar. *ApJ*, 683(1):L45–L49, July 2008.

[23] B. P. Abbott, et al. Searches for Gravitational Waves From Known Pulsars with Science Run 5 Ligo Data. *Class. Quantum Grav.*, 713(1):671–685, March 2010.

[24] J. Abadie, et al. Beating the Spin-Down Limit on Gravitational Wave Emission From the Vela Pulsar. *ApJ*, 737(2):93, August 2011.

[25] B. P. Abbott, et al. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, 460(7258):990–994, August 2009.

[26] J. Aasi, et al. Charaterization of the LIGO detectors during their sixth science run. *In preparation*, August 2013.

[27] D. M. Macleod, et al. Reducing the effect of seismic noise in LIGO searches by targeted veto generation. *Class. Quantum Grav.*, 29(5):5006, March 2012.

[28] S. Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity* . John Wiley & Sons, January 1972.

[29] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation.* San Francisco: W.H. Freeman and Co., 1973.

[30] M. Maggiore. *Gravitational Waves, Volume 1.* Theory and Experiments. Oxford University Press, USA, 2008.

[31] I. W. Harry. *Can we hear black holes collide?* PhD thesis, Cardiff Univiersity, Cardiff University, June 2011.

[32] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *ApJ*, 195:L51, January 1975.

[33] J. H. Taylor, et al. Further observations of the binary pulsar PSR 1913+16. *ApJ*, 206:L53, June 1976.

[34] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar PSR1913 + 16. *Nature*, 277(5696):437–440, February 1979.

[35] J. M. Weisberg, D. J. Nice, and J. H. Taylor. Timing Measurements of the Relativistic Binary Pulsar Psr B1913+16. *ApJ*, 722(2):1030–1034, September 2010.

[36] T. Regimbau and J. A. de Freitas Pacheco. Cosmic background of gravitational waves from rotating neutron stars. *A&A*, 376(2):381–385, September 2001.

[37] T. Regimbau and J. A. de Freitas Pacheco. Gravitational wave background from magnetars. *A&A*, 447(1):1–7, February 2006.

[38] T. Regimbau and J. A. de Freitas Pacheco. Stochastic Background from Coalescences of Neutron Star–Neutron Star Binaries. *ApJ*, 642(1):455–461, May 2006.

[39] V. Ferrari, S. Matarrese, and R. Schneider. Gravitational wave background from a cosmological population of core-collapse supernovae. *Monthly Notices RAS*, 303(2):247–257, February 1999.

[40] Grishchuk, L.P. Rest Gravitons and Possibilities of their Detection. *Pisma Zh.Eksp.Teor.Fiz.*, 23:326–329, 1976.

[41] T. Smith, E. Pierpaoli, and M. Kamionkowski. New Cosmic Microwave Background Constraint to Primordial Gravitational Waves. *Phys. Rev. Lett.*, 97(2):021301, July 2006.

[42] J. Aasi, et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys. Rev. D*, 87(4):042001, February 2013.

[43] J. Abadie, et al. All-sky search for periodic gravitational waves in the full S5 LIGO data. *Phys. Rev. D*, 85(2):022001, January 2012.

[44] C. D. Ott. The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26(6):063001, February 2009.

[45] N. Andersson and G. Comer. Probing Neutron-Star Superfluidity with Gravitational-Wave Data. *Phys. Rev. Lett.*, 87(24):241101, November 2001.

[46] S. Mereghetti. The strongest cosmic magnets: soft gamma-ray repeaters and anomalous X-ray pulsars. *Astron Astrophys Rev*, 15(4):225–287, July 2008.

[47] J. Abadie, et al. Search for Gravitational Wave Bursts From Six Magnetars. *ApJ*, 734(2):L35, June 2011.

[48] S. Klimenko, et al. A coherent method for detection of gravitational wave bursts. *Class. Quantum Grav.*, 25(11):114029, May 2008.

[49] R. N. Manchester, et al. The Parkes multi-beam pulsar survey - I. Observing and data analysis systems, discovery and timing of 100 pulsars. *Monthly Notices RAS*, 328(1):17–35, November 2001.

[50] D. J. Morris, et al. The Parkes Multibeam Pulsar Survey - II. Discovery and timing of 120 pulsars. *Monthly Notices RAS*, 335(2):275–290, September 2002.

[51] K. Postnov and L. Yungelson. The Evolution of Compact Binary Star Systems. *Living Reviews in Relativity*, 9(6), December 2006.

[52] J. Abadie, et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Class. Quantum Grav.*, 27(17):173001, July 2010.

[53] P. Peters and J. Mathews. Gravitational Radiation from Point Masses in a Keplerian Orbit. *Phys. Rev. D*, 131(1):435–440, July 1963.

[54] P. Peters. Gravitational Radiation and the Motion of Two Point Masses. *Phys. Rev. D*, 136(4B):B1224–B1232, November 1964.

[55] M. Hannam. Status of black-hole-binary simulations for gravitational-wave detection. *Class. Quantum Grav.*, 26(11):114001, May 2009.

[56] M. Shibata and K. Taniguchi. Coalescence of Black Hole-Neutron Star Binaries. *Living Reviews in Relativity*, 14(6), August 2011.

[57] J. A. Faber and F. A. Rasio. Binary Neutron Star Mergers. *Living Reviews in Relativity*, 15(8), July 2012.

[58] S. Chandrasekhar and S. Detweiler. The Quasi-Normal Modes of the Schwarzschild Black Hole. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 344(1639):441–452, August 1975.

[59] B. P. Abbott et al. Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data. *Phys. Rev. D*, 80(6):062001, September 2009.

[60] S. E. Caudill. *Searches for Gravitational Waves from Perturbed Black Holes in Data from LIGO Detectors*. PhD thesis, Louisiana State University, July 2012.

[61] P. Ajith, et al. Inspiral-Merger-Ringdown Waveforms for Black-Hole Binaries with Nonprecessing Spins. *Phys. Rev. Lett.*, 106(24):241101, June 2011.

[62] V. Kalogera, et al. The Cosmic Coalescence Rates for Double Neutron Star Binaries. *ApJ*, 601(2):L179–L182, January 2004.

[63] C. Kim, V. Kalogera, and D. Lorimer. The effect of PSR J0737-3039 on the DNS merger rate and implications for gravity-wave detection. *New Astronomy Reviews*, 54(3-6):148–151, March 2010.

[64] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *ArXiv e-prints*, April 2013.

[65] R. Takahashi and t. T. Collaboration. Status of TAMA300. *Class. Quantum Grav.*, 21(5):S403–S408, February 2004.

[66] A. A. Michelson and E. W. Morley. On the Relative Motion of the Earth and of the Luminiferous Ether. *Sidereal Messenger, vol. 6, pp.306-310*, 6:306–310, November 1887.

[67] Michelson interferometer schematic. http://sci.esa.int/jump.cfm?oid=50150.

[68] S. Hughes and K. S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Phys. Rev. D*, 58(12):122002, November 1998.

[69] A. Freise and K. Strain. Interferometer Techniques for Gravitational-Wave Detection. *Living Reviews in Relativity*, 13:1, February 2010.

[70] M. V. Plissi, et al. GEO 600 triple pendulum suspension system: Seismic isolation and control. *Rev. Sci. Instrum.*, 71(6):2539, 2000.

[71] H. Grote, et al. Dual recycling for GEO 600. *Class. Quantum Grav.*, 21(5):S473–S480, February 2004.

[72] S. Hild, et al. DC-readout of a signal-recycled gravitational wave detector. *Class. Quantum Grav.*, 26(5):055012, February 2009.

[73] M. Prijatelj, et al. Control and automatic alignment of the output mode cleaner of GEO 600. *J. Phys.: Conf. Ser.*, 228:012014, July 2010.

[74] J. Abadie, et al. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nat Phys*, 7(12):962–965, September 2011.

[75] J. R. Smith. The path to the enhanced and advanced LIGO gravitational-wave detectors. *Class. Quantum Grav.*, 26(1):4013, June 2009.

[76] R. Abbott. Control System Design for the LIGO Pre-Stabilized Laser. In *The 8th International Conference on Accelerators and Large Experimental Physics Control Systems*, page 361, 2001.

[77] M. E. Zucker and S. E. Whitcomb. Measurement of Optical Path Fluctuations due to Residual Gas in the LIGO 40 Meter Interferometer. *Proceedings of the Seventh Marcel Grossman Meeting on recent developments in theoretical and experimental general relativity, gravitation, and relativistic field theories*, -1:1434, 1996.

[78] B. Meers. Recycling in laser-interferometric gravitational-wave detectors. *Phys. Rev. D*, 38(8):2317–2326, October 1988.

[79] R. L. Ward, et al. dc readout experiment at the Caltech 40m prototype interferometer. *Class. Quantum Grav.*, 25(11):114030, May 2008.

[80] T. T. Fricke, et al. DC readout experiment in Enhanced LIGO. *Class. Quantum Grav.*, 29(6):065005, February 2012.

[81] LIGO Laboratory Home Page for Interferometer Sensitivities. http://www.ligo.caltech.edu/~jzweizig/distribution/LSC_Data/.

[82] R. Abbott, et al. Seismic isolation for Advanced LIGO. *Class. Quantum Grav.*, 19(7):1591–1597, April 2002.

[83] R. Abbott, et al. Seismic isolation enhancements for initial and Advanced LIGO. *Class. Quantum Grav.*, 21(5):915, March 2004.

[84] G. González. Suspensions thermal noise in the LIGO gravitational wave detector. *Class. Quantum Grav.*, 17(2):4409–4435, November 2000.

[85] P. R. Saulson. Thermal noise in mechanical experiments. *Phys. Rev. D*, 42(8):2437–2445, October 1990.

[86] P. R. Saulson. *Fundamentals of interferometric gravitational wave detectors.* World Scientific, 1994.

[87] A. V. Cumming, et al. Design and development of the advanced LIGO monolithic fused silica suspension. *Class. Quantum Grav.*, 29(3):035003, January 2012.

[88] C. Caves. Quantum-mechanical noise in an interferometer. *Phys. Rev. D*, 23(8):1693–1708, April 1981.

[89] G. M. Harry and The LIGO Scientific Collaboration. Advanced LIGO: the next generation of gravitational wave detectors. *Class. Quantum Grav.*, 27(8):084006, April 2010.

[90] The Virgo Collaboration. Advanced Virgo Baseline Design. Technical Report VIR-0027A-09, August 2009.

[91] The Virgo Collaboration. Advanced Virgo Technical Design Report. Technical Report VIR-0128A-12, April 2012.

[92] The LIGO Scientific Collaboration. Advanced LIGO Reference Design. Technical Report M060056, March 2011.

[93] M. V. Plissi, et al. Aspects of the suspension system for GEO 600. *Rev. Sci. Instrum.*, 69(8):3055, 1998.

[94] S. M. Aston, et al. Update on quadruple suspension design for Advanced LIGO. *Class. Quantum Grav.*, 29(23):235004, October 2012.

[95] G. M. Harry, et al. Titania-doped tantala/silica coatings for gravitational-wave detection. *Class. Quantum Grav.*, 24(2):405–415, December 2006.

[96] R. Lawrence, et al. Adaptive thermal compensation of test masses in advanced LIGO. In *Proceedings of the 4th Edoardo Amaldi Conference on Gravitational Waves*, pages 1803–1812, March 2002.

[97] B. Willke, et al. Stabilized lasers for advanced gravitational wave detectors. *Class. Quantum Grav.*, 25(11):114040, May 2008.

[98] M. A. Arain and G. Mueller. Design of the Advanced LIGO recycling cavities. *Opt. Express*, 16(14):10018–10032, 2008.

[99] J. Abadie, et al. Upper limits on a stochastic gravitational-wave background using LIGO and Virgo interferometers at 600–1000 Hz. *Phys. Rev. D*, 85(12):122001, June 2012.

[100] P. J. Sutton, et al. X-Pipeline: an analysis package for autonomous gravitational-wave burst searches. *New J. Phys.*, 12(5):053034, May 2010.

[101] I. W. Harry and S. Fairhurst. Targeted coherent search for gravitational waves from compact binary coalescences. *Phys. Rev. D*, 83(8):084002, April 2011.

[102] S. Babak, et al. Searching for gravitational waves from binary coalescence. *Phys. Rev. D*, 87(2):24033, January 2013.

[103] L. Blackburn, et al. The LSC glitch group: monitoring noise transients during the fifth LIGO science run. *Class. Quantum Grav.*, 25(18):184004, September 2008.

[104] J. Slutsky, et al. Methods for reducing false alarms in searches for compact binary coalescences in LIGO data. *Class. Quantum Grav.*, 27(16):165023, July 2010.

[105] N. Christensen, The LIGO Scientific Collaboration, and The Virgo Collaboration. LIGO S6 detector characterization studies. *Class. Quantum Grav.*, 27(1):4010, October 2010.

[106] J. McIver. Data quality studies of enhanced interferometric gravitational wave detectors. In *Class. Quantum Grav.*, page 124010, June 2012.

[107] F. Robinet, The LIGO Scientific Collaboration, and The Virgo Collaboration. Data quality in gravitational wave bursts and inspiral searches in the second Virgo Science Run. *Class. Quantum Grav.*, 27(19):194012, September 2010.

[108] J. Aasi, et al. The characterization of Virgo data and its impact on gravitational-wave searches. *Class. Quantum Grav.*, 29(1):5002, August 2012.

[109] R. L. Savage, P. J. King, and S. U. Seel. A Highly Stabilized 10-Watt Nd: YAG Laser for the Laser Interferometer Gravitational-Wave Observatory (LIGO). *Laser Physics*, 8(3):679–685, 1998.

[110] R. Adhikari, P. Fritschel, and S. J. Waldman. Enhanced LIGO. Technical Report LIGO-T060156, LIGO Laboratory, July 2006.

[111] P. Kwee, et al. Stabilized high-power laser system for the gravitational wave detector advanced LIGO. *Opt. Express*, 20(10):10617–10634, 2012.

[112] S. Ballmer, et al. Thermal Compensation System Description. Technical Report LIGO-T050064, LIGO Laboratory, April 2005.

[113] R. S. Amin and J. A. Giaime. Gravitational-wave detector-derived error signals for the LIGO thermal compensation system. *Class. Quantum Grav.*, 27(21):215002, September 2010.

[114] P. Fritschel, et al. Readout and Control of a Power-Recycled Interferometric Gravitational-Wave Antenna. *Appl. Opt.*, 40(28):4988–4998, 2001.

[115] N. Smith-Lefebvre, et al. Optimal alignment sensing of a readout mode cleaner cavity. *Optics Letters*, 36:4365, November 2011.

[116] A. Bertolini, et al. Design and prototype tests of a seismic attenuation system for the advanced-LIGO output mode cleaner. *Class. Quantum Grav.*, 23(8):S111–S118, March 2006.

[117] R. DeRosa, et al. Global feed-forward vibration isolation in a km scale interferometer. *Class. Quantum Grav.*, 29(21):215008, October 2012.

[118] Y. Levin. Internal thermal noise in the LIGO test masses: A direct approach. *Phys. Rev. D*, 57(2):659–663, January 1998.

[119] A. Buonanno and Y. Chen. Optical noise correlations and beating the standard quantum limit in advanced gravitational-wave detectors. *Class. Quantum Grav.*, 18(15):L95–L101, July 2001.

[120] L. Finn and D. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. *Phys. Rev. D*, 47(6):2198–2219, March 1993.

[121] J. Aasi, et al. Environmental Influences on the LIGO Gravitational Wave Detectors During the 6th Science Run. *In preparation*, 2013.

[122] S. Chatterji, et al. Multiresolution techniques for the detection of gravitational-wave bursts. *Class. Quantum Grav.*, 21(20):S1809–S1818, September 2004.

[123] J. R. Smith, et al. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Class. Quantum Grav.*, 28(23):235005, November 2011.

[124] T. Isogai. Used percentage veto for LIGO and virgo binary inspiral searches. *J. Phys.: Conf. Ser.*, 243:012005, 2010.

[125] E. J. Daw, et al. Long-term study of the seismic environment at LIGO. *Class. Quantum Grav.*, 21(9):2255–2273, April 2004.

[126] G. Manson and G. H. d. Visme. The frequency spectrum of Barkhausen noise. *J. Phys. D: Appl. Phys.*, 5(8):1389–1395, May 2002.

[127] S. K. Chatterji. *The search for gravitational wave bursts in data from the second LIGO science run*. PhD thesis, Massachusetts Institute of Technology, 2005.

[128] J. Abadie, et al. Implementation and testing of the first prompt search for gravitational wave transients with electromagnetic counterparts. *A&A*, 539:A124, March 2012.

[129] P. A. Evans, et al. *Swift* follow-Up Observations of Candidate Gravitational-Wave Transient Events. *ApJS*, 203(2):28, November 2012.

[130] S. Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New J. Phys.*, 11(12):123006, December 2009.

[131] T. Accadia, et al. Noise from scattered light in Virgo's second science run data. *Class. Quantum Grav.*, 27(19):194011, September 2010.

[132] B. P. Abbott, et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev. D*, 80(10):102001, November 2009.

[133] B. P. Abbott, et al. Search for gravitational waves from low mass binary coalescences in the first year of LIGO's S5 data. *Phys. Rev. D*, 79(12):122001, June 2009.

[134] J. C. Brown. Calculation of a constant Q spectral transform. *The Journal of the Acoustical Society of America*, 89(1):425–434, January 1991.

[135] B. P. Abbott, et al. Search for gravitational waves from low mass compact binary coalescence in 186 days of LIGO's fifth science run. *Phys. Rev. D*, 80(4):047101, August 2009.

[136] K. S. Thorne. Gravitational Radiation. In S. W. Hawking and W. Israel, editors, *Three hundred years of gravitation*, pages 330–458. Cambridge University Press, Cambridge, 1987.

[137] D. G. Keppel. *Signatures and dynamics of compact binary coalescences and a search in LIGO's S5 data*. PhD thesis, California Institute of Technology, 2009.

[138] S. Babak, et al. A template bank to search for gravitational waves from inspiralling compact binaries: I. Physical models. *Class. Quantum Grav.*, 23(18):5477–5504, August 2006.

[139] B. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53(12):6749–6761, June 1996.

[140] B. Owen and B. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60(2):022002, June 1999.

[141] T. Cokelaer. A template bank to search for gravitational waves from inspiralling compact binaries: II. Phenomenological model. *Class. Quantum Grav.*, 24(24):6227–6242, November 2007.

[142] B. Allen. 2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71(6):062001, March 2005.

[143] L. Blanchet, et al. Gravitational-Radiation Damping of Compact Binary Systems to Second Post-Newtonian Order. *Phys. Rev. Lett.*, 74(18):3515–3518, May 1995.

[144] L. Blanchet, et al. Gravitational Radiation from Inspiralling Compact Binaries Completed at the Third Post-Newtonian Order. *Phys. Rev. Lett.*, 93(9):091101, August 2004.

[145] "Blind injection" stress-tests LIGO and virgo's search for gravitational waves. http://ligo.org/news/blind-injection.php. Accessed: 21/3/2013.

[146] The Astronomer's Telegram. http://www.astronomerstelegram.org/.

[147] B. P. Abbott, et al. Search For Gravitational-wave Bursts Associated with Gamma-ray Bursts using Data from LIGO Science Run 5 and Virgo Science Run 1. *ApJ*, 715(2):1438–1452, June 2010.

[148] S. Adrián-Martínez, et al. A first search for coincident gravitational waves and high energy neutrinos using LIGO, Virgo and ANTARES data from 2007. *J. Cosmol. Astropart. Phys.*, 2013(06):008–008, June 2013.

[149] J. Abadie, et al. Search for gravitational waves associated with the August 2006 timing glitch of the Vela pulsar. *Phys. Rev. D*, 83(4):042001, February 2011.

[150] S. D. Barthelmy, et al. The Burst Alert Telescope (BAT) on the SWIFT Midex Mission. *Space Sci Rev*, 120(3-4):143–164, October 2005.

[151] C. Meegan, et al. The Fermi Gamma-Ray Burst Monitor. *ApJ*, 702(1):791–804, August 2009.

[152] K. Hurley, et al. The Current Performance of the Third Interplanetary Network. In *AIP Conf.Proc.*, pages 473–476. UC Berkeley Space Sciences Laboratory, Berkeley, CA 94720-7450, January 2003.

[153] S. Barthelmy. GCN and VOEvent: A status report. *Astron. Nachr.*, 329(3):340–342, March 2008.

[154] C. Kouveliotou, et al. Identification of two classes of gamma-ray bursts. *ApJ*, 413:L101, August 1993.

[155] Batse: Burst and transient source experiment. http://www.batse.msfc.nasa.gov/batse/.

[156] E. Nakar. Short-hard gamma-ray bursts. *Phys. Rept.*, 442(1-6):166–236, April 2007.

[157] D. Eichler, et al. Nucleosynthesis, neutrino bursts and -rays from coalescing neutron stars. *Nature*, 340(6229):126–128, July 1989.

[158] R. Narayan, B. Paczynski, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *ApJ*, 395:L83, August 1992.

[159] B. Paczynski. Cosmological gamma-ray bursts. *Acta Astronomica*, 41:257–267, 1991.

[160] R. Mochkovitch, et al. Gamma-ray bursts as collimated jets from neutron star/black hole mergers. *Nature*, 361(6409):236–238, January 1993.

[161] J. Abadie, et al. Implications for the Origin of GRB 051103 from LIGO Observations. *ApJ*, 755(1):2, July 2012.

[162] B. P. Abbott, et al. Implications for the Origin of GRB 070201 from LIGO Observations. *ApJ*, 681(2):1419–1430, July 2008.

[163] S. Márka. Transient multimessenger astronomy with gravitational waves. *Class. Quantum Grav.*, 28(11):114013, May 2011.

[164] S. Klimenko, et al. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev. D*, 72(12):122002, December 2005.

[165] G. Vedrenne and J.-L. Atteia. *Gamma-Ray Bursts: The Brightest Explosions in the Universe.* Springer Praxis Books. Springer, February 2009.

[166] D. Keppel. Multidetector F-statistic metric for short-duration nonprecessing inspiral gravitational-wave signals. *Phys. Rev. D*, 86(12):123010, December 2012.

[167] O. Rabaste, E. Chassande-Mottin, and A. Pai. Sparse sky grid for the coherent detection of gravitational wave bursts. *ArXiv e-prints*, page 4832, May 2009.

[168] D. N. Burrows, et al. Jet Breaks in Short GammaRay Bursts. II. The Collimated Afterglow of GRB 051221A. *ApJ*, 653(1):468–473, December 2006.

[169] W. Fong, et al. A Jet Break in the X-Ray Light Curve of Short Grb 111020a: Implications for Energetics and Rates. *ApJ*, 756(2):189, August 2012.

[170] A. Nicuesa Guelbenzu, et al. Multi-color observations of short GRB afterglows: 20 events observed between 2007 and 2010. *A&A*, 548:A101, November 2012.

[171] GCN circular 9742. http://gcn.gsfc.nasa.gov/gcn3/9742.gcn3.

[172] F. Cavalier, et al. Reconstruction of source location in a network of gravitational wave interferometric detectors. *Phys. Rev. D*, 74(8):082004, October 2006.

[173] "Teaching Physics". http://xkcd.com/895/.