

The search for gravitational wave bursts
in data from the second LIGO science run

by
Shourov Keith Chatterji

S.B. Physics
Massachusetts Institute of Technology, 1995


S.B. Electrical Engineering
Massachusetts Institute of Technology, 1995


Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy


at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 2005


© Shourov Keith Chatterji, MMV. All rights reserved.

The author hereby grants to MIT permission to reproduce and
distribute publicly paper and electronic copies of this thesis document
in whole or in part.

Author 
Department of Physics
June 10, 2005

Certified by 
Erotokritos Katsavounidis
Assistant Professor of Physics
Thesis Supervisor

Accepted by 
Thomas J. Greytak
Professor of Physics
Associate Department Head for Education



ARCHIVES

The search for gravitational wave bursts in data from the second LIGO science run

by

Shourov Keith Chatterji

Submitted to the Department of Physics
on June 10, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The network of detectors comprising the Laser Interferometer Gravitational-wave Observatory (LIGO) are among a new generation of detectors that seek to make the first direct observation of gravitational waves. While providing strong support for the General Theory of Relativity, such observations will also permit new tests of physical theory in regions of strong space-time curvature and high matter-energy density. However, the observed signals are expected to occur near the limit of detector sensitivity. The problem of identifying such small signals is the primary focus of this work.

This work presents a novel method for the identification of astrophysically unmodeled bursts of gravitational radiation in data from networks of interferometric detectors. The method is based on the Q transform, a multiresolution time-frequency transform that efficiently targets waveforms within a finite region of time, frequency, and Q space. The method is also based on a modification of linear prediction that greatly simplifies the resulting statistical analysis by whitening interferometric detector data prior to Q transform analysis. Together, these techniques form the basis of a complete analysis pipeline that is equivalent to a template-based matched filter search for minimum uncertainty waveforms in the whitened data stream.

This method is then applied to search for gravitational-wave bursts with duration less than 1 second and frequency content between 64 and 1024 Hz in coincident data from two detectors during second LIGO science run. Although no gravitational-wave bursts are identified, upper bounds are reported for the rate of gravitational-wave bursts as a function of signal strength for isotropic and galactic populations of sources with both abstract and astrophysically motivated waveform. The results

indicate a maximum of 0.09 events per day at the 90% confidence level for bursts with characteristic strain amplitude in excess of 10^{-20} to 10^{-19} strain $\text{Hz}^{-1/2}$ depending on waveform. A comparison with previous searches demonstrates that this search is one of the most sensitive to date for gravitational-wave bursts of unknown waveform, and is inconsistent with recent indications for an statistical excess of events by the ROG collaboration at above the 99% confidence level.

Thesis Supervisor: Erotokritos Katsavounidis
Title: Assistant Professor of Physics

Acknowledgements

The completion of this work would not have been possible without the immeasurable support of a large number of people.

The sensitivity of this analysis is a testament to the abilities of the many people involved in the design, commissioning, and operation of the LIGO detectors. It is primarily through their effort, in constructing what is arguably the most sensitive network of detectors to date, that this work is able to achieve the strictest upper bound on the rate of gravitational-wave bursts to date.

I am also indebted to the members of the LIGO Scientific Collaboration, in particular its Burst Analysis Group, for providing an excellent forum for the exchange of ideas, many of which have had a profound impact on this work.

In particular, in the MIT LIGO Data Analysis Group, I have had the privilege to work with many extraordinarily talented people who have helped make this work possible. While I have appreciated the countless hours of discussion and assistance they have provided while this research was carried out, above all I have appreciated working in such an enjoyable group. In particular, I would like to thank Laura Cadonati, Katherine Rawlins, Lindy Blackburn, Michele Zanolin, Stefan Ballmer, Keith Bayer, and Marie Woods for their assistance. I look forward to continuing to work with them within the LIGO Scientific Collaboration.

On my thesis committee, I was fortunate to have Deepto Chakrabarty and Scott Hughes. Through their careful reading and suggestions, the content, quality, focus and accessibility of this work was greatly improved. In addition, parts of this work have also benefited from careful reading by Patrick Sutton, Rejean Dupuis, and Katherine Rawlins.

I would also like to thank Rai Weiss, who introduced me to the challenge of high precision experimental physics, and who was also instrumental in bringing me into the MIT graduate physics program.

I am particularly indebted to Stan Whitcomb and David Shoemaker, who kindly provided much needed support and encouragement during difficult times. The LIGO Laboratory is fortunate to have them.

I have had the great privilege of pursuing this research under the guidance of my thesis adviser, Erik Katsavounidis. I cannot thank him enough for the opportunity, support, and guidance he has provided. It was with his arrival at MIT, and his development of a strong gravitational-wave data analysis group here, that I was able to find a research topic which lead to an interesting and challenging thesis project. Since then, he has gone out of his way many times to help ensure the success of this project. And, for that, I am extremely thankful.

Finally, I would not have been able to accomplish this monumental task without the support of my family. All of them, particularly my wife Suzanne, have been a constant source of support, understanding, and encouragement throughout this very long process. And, for the last year and 9 months, my daughter Anna, with little effort, has provided more encouragement than she can possibly know. Now that this work is complete, I look forward to spending much more time with them.

This work was supported by the US National Science Foundation under Cooperative Agreement No. PHY-0107417.

Contents

- 1 Introduction** **13**

- 2 Gravitational radiation** **19**
 - 2.1 Gravitational waves 20
 - 2.2 Detectors 30
 - 2.3 Sources 38

- 3 Burst detection** **65**
 - 3.1 Parameterization of bursts 66
 - 3.2 Measurement of bursts 73
 - 3.3 Detection of bursts 84

- 4 Linear Prediction** **91**
 - 4.1 Definition 92
 - 4.2 Training 93

4.3	Properties	95
4.4	Application	98
4.5	Frequency response	99
4.6	Zero-phase filtering	102
4.7	Example	106
5	The Q Transform	109
5.1	The continuous Q transform	111
5.2	The discrete Q transform	113
5.3	Analysis window	117
5.4	Normalization	122
5.5	Statistics	138
5.6	Identification of events	155
5.7	The Q pipeline	157
5.8	The coherent Q pipeline	161
5.9	Example	167
6	Simulation	173
6.1	Well localized bursts	175

6.2	Stationary white noise	177
6.3	Simulated detector noise	186
7	All sky search	195
7.1	Data selection	197
7.2	Analysis pipeline	202
7.3	Event rates	205
7.4	Candidate events	213
7.5	Statistical analysis	234
7.6	Detection efficiencies	237
7.7	Systematic errors	245
7.8	Upper limits	246
7.9	Comparison of results	251
8	Conclusion	267
8.1	Summary	268
8.2	Future investigations	273
	Bibliography	279

Chapter 1

Introduction

Gravitational waves are linear perturbations of space and time, caused by the motion of matter, that travel outward through the universe at the speed of light. The existence of gravitational waves follows as a direct consequence of Einstein's General Theory of Relativity. As a result, the direct detection of their interaction with matter would constitute an important confirmation of theory.

There is already strong indirect evidence in support of the existence of gravitational waves. Pulsar timing observations of the relativistic binary system PSR 1913+16 by Hulse, Taylor, and colleagues indicate that the orbit of the system is decaying at a rate that is in very good agreement with the expected energy loss due to gravitational-wave emission[1, 2]. Since then, the discovery of the similar binary pulsar systems PSR B1534+12[3, 4], PSR 2127+11C[5, 6], and PSR J0737-3039[7] have provided further compelling evidence for the existence of gravitational radiation, and also permit an estimate of the population of such systems[8].

Nevertheless, direct observation of the interaction of gravitational waves and matter has not yet been achieved. Such an observation would not only lend strong support to the General Theory of Relativity, but would also usher in a new field of astrophysics that falls outside of the electromagnetic spectrum and is for the first time

directly sensitive to the dynamical motion of matter. Such observations will provide new tests of our understanding of physics in regions of strong space-time curvature and very high matter-energy density. In particular, in contrast to electromagnetic radiation, gravitational waves are not easily absorbed or scattered by matter. As a result, gravitational waves permit observation of the bulk motion of matter inside catastrophic events such as core collapse supernovae or the coalescence of binary compact objects such as neutron stars and black holes — regions that are currently inaccessible due to photon and neutrino scattering. In addition, the observation of a stochastic background of gravitational radiation from density fluctuations in the very early universe would provide important tests for theories of inflation or cosmic strings. Finally, there exists the possibility of discovering previously unknown phenomena, and the profound impact such a discovery would have on physical theory, as has typically been the case whenever the universe has been observed in a new region of the electromagnetic spectrum.

Starting with the pioneering work of Weber[9] on the first generation of resonant-mass detectors in the early 1960s, gravitational-wave detectors have now developed into a global network of detectors, with a number of new detectors currently nearing completion. The newest generation of these detectors, kilometer scale interferometers that were first studied in detail by Weiss[10] in the early 1970s, now promise greatly increased sensitivity over a much larger bandwidth than the existing generation of narrowband resonant-mass detectors. In particular, this work focuses on data from the Laser Interferometer Gravitational-wave Observatory (LIGO), which operates a network of three interferometric detectors at two separate locations in the United States[11, 12]. The LIGO detectors are now in the final stages of commissioning and have entered alternating periods of scientific observation and commissioning work as they approach their respective design sensitivities. Currently, the LIGO detectors have completed four coincident science runs with successively improving sensitivity, constituting the most sensitive searches for gravitational waves to date.

Despite many orders of magnitude improvement in detector sensitivity over the last

40 years, the detection of gravitational waves has so far remained an elusive goal. Due to the extremely weak interaction of gravitation waves with matter, the signals from potentially detectable sources are anticipated to be so small that they will occur near the sensitivity limit of existing detectors. Even as detectors improve, the set of potentially detectable sources will continue to be dominated by signals at the sensitivity limit of future detectors. Significant effort must therefore go into the task of robustly identifying such small signals in the presence of the numerous detector noise sources that constrain detector sensitivity. This last effort is the primary focus of the present work.

In the following chapters we propose a complete method for analyzing the data from networks of interferometric gravitational-wave observatories in order to identify and characterize astrophysically unmodeled bursts of gravitational radiation. We then apply this method to the search for gravitational-wave bursts in data from the second LIGO science run. Although we find no events of likely gravitational-wave origin, we also determine upper bounds on the possible rate of gravitational-wave bursts of various waveform. A comparison with the results of previous searches then indicates that this search is one of the most sensitive to date and excludes previous evidence for detection with a high degree of confidence.

The remainder of this work is organized as follows.

In chapter 2, we lay the foundations for this work by giving an overview of the current state of gravitational-wave physics. We begin by presenting gravitational waves as a consequence of Einstein's General Theory of Relativity and briefly describe the properties of such waves, including their generation and their interaction with matter. In the process, we also motivate our effort by highlighting the potential benefits that the detection of gravitational waves would bring to our understanding of nature. We then consider the detection of gravitational waves and give a brief overview of current efforts. Finally, we conclude with a survey of potentially detectable astrophysical sources of gravitational waves and review current predictions of the populations of

such sources.

In chapter 3, we first introduce a parameterization of gravitational-wave bursts that provides us with the language to describe both their measurement and detection. We then motivate the use of a multiresolution basis of waveforms with minimum time-frequency uncertainty as the ideal measurement basis for the detection of unmodeled bursts. Finally, we present a method for constructing such a basis in a way that optimally covers a targeted space of signals.

In chapter 4, we present a method, based on autoregressive modeling of stochastic processes, for whitening the data from interferometric gravitational-wave detectors prior to further analysis. We will find that by doing so we greatly simplify the resulting statistical analysis.

In chapter 5, we present the Q transform, a multiresolution time-frequency transform which forms the basis of the proposed search method. In the process, we consider the statistical properties of the Q transform applied to stationary white noise and demonstrate that it is the optimal detector for bursts with minimum time-frequency uncertainty in the presence of stationary white noise data. We then combine the Q transform with the techniques proposed in the previous chapters in order to form a comprehensive end-to-end analysis pipeline for the detection of gravitational-wave bursts. In addition, we present a method of coherently combining Q transform measurements from multiple detectors in a way that maximizes the combined signal to noise ratio of gravitational-wave bursts while taking into account potential differences in sensitivity between detectors and excluding inconsistent events.

In chapter 6, we present the performance of the proposed analysis pipeline for the special case of waveforms with minimum time-frequency uncertainty in the presence simulated detector noise. This controlled study permits a validation of the analysis pipeline by comparing the observed performance against the theoretically predicted performance. In doing so, we demonstrate that the proposed pipeline is equivalent to

the optimal search strategy for the special case of minimum uncertainty waveforms in the whitened data stream.

In chapter 7, we present the results of the proposed analysis pipeline applied to an all-sky search for gravitational-wave bursts in coincident data from the two colocated LIGO detectors during the second LIGO science run. We also perform a systematic study of the expected background event rate using non-physical time shifts and demonstrate that an observed excess of foreground events is not of likely gravitational-wave origin. The amplitude of the most significant foreground event is then used to determine upper bounds on the rate of gravitational-wave bursts arising from an isotropic distribution of sources of both abstract and astrophysically motivated waveform. Finally, we compare the resulting upper bounds with the those of previous searches and demonstrate that the resulting search is one of the most sensitive to date for unmodeled bursts of gravitational radiation. In addition, we find that the resulting bounds are highly inconsistent with the galactic population of sources necessary to explain recent claims of a statistical excess of events by the ROG collaboration[13, 14].

Finally, in chapter 8, we briefly summarize the significance of the work presented here and outline possible extensions of the proposed method for future investigation.

Chapter 2

Gravitational radiation

In this chapter, we present an overview of the current state of gravitational wave physics. In doing so, our goal is to provide the necessary background to motivate the search for gravitational waves and to understand the remainder of this work. As such, we do not aspire to a rigorous treatment of the subject. Instead, we follow the treatment of Saulson[15] and other recent review articles[16, 17, 18, 19] in providing an introduction to the field and in highlighting the potential contributions that the detection of gravitational waves could bring to physics. For a more detailed treatment of gravitational-wave theory, the reader is instead referred to the many comprehensive texts on the subject[20, 21, 22].

In what follows, we begin with an overview of the theory of gravitational waves. Specifically, we establish gravitational waves as a prediction of Einstein's General Theory of Relativity, briefly describing the properties of such waves, as well as their generation and their interaction with matter. From simple dimensional analysis and order of magnitude arguments we also identify the expected amplitude scale and frequency band of gravitational waves and are immediately led to conclude that the problem of gravitational wave detection is necessarily one of astrophysics.

Next, we present an overview of existing efforts to directly observe the interaction of gravitational waves with matter. In particular, we focus on the nascent field of interferometric gravitational-wave detection with an emphasis on the three detectors operated by the Laser Interferometric Gravitational-wave Observatory (LIGO) project, which are currently in the final stages of commissioning. In addition, we also briefly describe other efforts to observe gravitational waves, including the existing network of resonant mass detectors and future plans for space based interferometry.

Finally we consider a variety of potential astrophysical sources of gravitational radiation that may be detected by the first generation of interferometric gravitational wave observatories. We give emphasis, however, to potential sources of transient bursts of gravitational radiation, which is the primary focus of this work. In the process, we also motivate our effort by highlighting the potential benefits that the detection, or lack of detection, of gravitational waves will bring to our understanding to nature.

2.1 Gravitational waves

Motivated by the empirically observed equivalence between gravitational and inertial mass, Einstein's General Theory of Relativity postulates that, in sufficiently small regions of space-time and in the absence of any other force, it is impossible for an observer to detect the effect of gravitation. As a result, it is possible to construct a locally Lorentzian coordinate system at every point in space-time. The General Theory of Relativity then goes on to provide the mathematical framework necessary to describe the properties of space-time and the equations of classical physics in a way that is independent of the particular choice of coordinates used to identify events. In this approach, all of the necessary local properties of space-time are encoded by the metric $g_{\mu\nu}$, a symmetric matrix that provides a formalism for determining the invariant space-time interval between pairs of events. Specifically, given a particular

coordinate system x^μ , the proper distance between events is given by the expression

$$\delta s^2 = g_{\mu\nu} \delta x^\mu \delta x^\nu, \quad (2.1)$$

where we have assumed the usual Einstein convention of implicit summation over repeated indices; in this case μ and ν , which run over the four coordinates of space and time in units where the speed of light c is unity. In what follows, we make use of these conventions to simplify our discussion. However, we reintroduce the appropriate factors of c in the more important results.

From this approach, we can then determine the proper distance between arbitrary events by performing the integral

$$s = \int_{\lambda_1}^{\lambda_2} \left(g_{\mu\nu}(\lambda) \frac{dx^\mu}{d\lambda} \frac{dx^\nu}{d\lambda} \right)^{1/2} d\lambda, \quad (2.2)$$

where λ parameterizes the integration path between events λ_1 and λ_2 and the matrix elements $g_{\mu\nu}$ are in general a function of position. The motion of free particles is then that path λ that makes s an extremum. Such paths are referred to as geodesics in reference to their geometric interpretation as the shortest path between two points on a curved surface. By parameterizing the resulting path in terms of the proper distance s , we obtain the geodesic equation of motion for massive particles,

$$\frac{d^2 x^\alpha}{ds^2} + \Gamma_{\mu\nu}^\alpha \frac{dx^\mu}{ds} \frac{dx^\nu}{ds} = 0, \quad (2.3)$$

where $\Gamma_{\mu\nu}^\alpha$ is the so-called affine connection and may be defined in terms of first derivatives of the space-time metric by the relation

$$\Gamma_{\mu\nu}^\alpha = \frac{1}{2} g^{\alpha\beta} \left\{ \frac{\partial g_{\beta\nu}}{\partial x^\mu} + \frac{\partial g_{\beta\mu}}{\partial x^\nu} - \frac{\partial g_{\alpha\nu}}{\partial x^\beta} \right\}. \quad (2.4)$$

Although we are free to describe space-time using an arbitrary choice of coordinate system, there are still constraints imposed on the form of the space-time metric by the

General Theory of Relativity. In particular, there is a unique coordinate independent quantity R that describes the geometric curvature of space-time. This quantity can be computed from second derivatives of the metric by

$$R = g^{\mu\nu} R_{\mu\nu}, \quad (2.5)$$

where the $R_{\mu\nu}$ is the Ricci tensor and given by the contraction

$$R_{\mu\nu} = R^{\lambda}_{\mu\lambda\nu} \quad (2.6)$$

of the Riemann curvature tensor,

$$R^{\lambda}_{\mu\sigma\nu} = \frac{\partial\Gamma^{\lambda}_{\mu\nu}}{\partial x^{\sigma}} - \frac{\partial\Gamma^{\lambda}_{\mu\sigma}}{\partial x^{\nu}} + \Gamma^{\lambda}_{\sigma\alpha}\Gamma^{\alpha}_{\mu\nu} - \Gamma^{\lambda}_{\nu\alpha}\Gamma^{\alpha}_{\mu\sigma}. \quad (2.7)$$

In addition, it is also possible to express the effect of gravitation as a relationship between the space-time metric and the local density of energy and momentum in a way that is independent of our choice of coordinate system. This relationship is provided by the General Theory of Relativity in the form of the Einstein field equation,

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi GT_{\mu\nu}, \quad (2.8)$$

where $T_{\mu\nu}$ describes local density of energy and momentum and G is Newton's constant.

2.1.1 The wave equation

In general, the Einstein field equation is such that gravity by itself can effect the local curvature of space-time. As a result, we note that the resulting field equation comprises of a set of 10 coupled non-linear differential equations. Consequently, finding exact analytical solutions to the Einstein field equation is a formidable task. Fortunately, there are a few simple cases for which an analytical solution is possible. In

particular, in what follows we present the solution to the Einstein field equation in the absence of matter and in the limit of weak gravitational fields. That is we suppose that the space-time metric may be written as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (2.9)$$

where $\eta_{\mu\nu}$ is the familiar Minkowski flat space metric from the Special Theory of Relativity,

$$\eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (2.10)$$

and $h_{\mu\nu}$ represents a small perturbation such that its matrix elements obey the condition

$$|h_{\mu\nu}| \ll 1. \quad (2.11)$$

In this limit, we may solve the Einstein field equation by neglecting all terms that are higher than first order in $h_{\mu\nu}$. In addition, we are permitted a certain amount of freedom in our choice of coordinate system. While we don't develop the details here, we note that there exists a particular choice of coordinates, the so-called transverse-traceless gauge, for which $h_{\mu\nu}$ takes the simple form

$$h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (2.12)$$

and the Einstein field equation takes the simple form

$$\left(\nabla^2 - \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0, \quad (2.13)$$

This is easily recognized as the familiar wave equation with solutions of the form

$$h_{\mu\nu} = h_0 f(t - z). \quad (2.14)$$

Far from the source of gravitational radiation, these solutions are conveniently represented as a superposition of complex valued plane waves of the form

$$h(\mathbf{x}, t) = h_0 \exp [i(2\pi f t - \mathbf{k} \cdot \mathbf{x})], \quad (2.15)$$

where the wave vector \mathbf{k} defines the direction of wave propagation and, in more conventional units where c is not unity, has a magnitude given by

$$k = \frac{2\pi f}{c}. \quad (2.16)$$

We thus see that, in the far field limit, gravitational waves are plane waves that propagate at the speed of light. Of course, this is not an entirely unexpected result. Einstein's Special Theory of Relativity already prohibits action at a distance and postulates that the speed of light represents a fundamental limit on the rate at which information can propagate. Therefore, based solely on the Special Theory of Relativity, we could have anticipated, at least qualitatively, the need for gravitational waves as a mediator for changes in the gravitational field due to the motion of matter, much as electromagnetic waves mediate changes in the electromagnetic field due to the motion of charge. In fact, Schutz has shown that the most important effects of gravitational radiation can be predicted with reasonably good accuracy using only Newtonian gravity and special relativity[23]. The fact that gravitational waves propagate at the speed of light also indicates that the graviton, the fundamental particle which mediates the gravitational force, like the photon, has zero rest mass.

From Equation 2.12, we also note that there are only two independent components of the metric perturbation $h_{\mu\nu}$. This is a property that is independent of our choice of coordinate system and represents the fact that gravitational waves come in two or-

thogonal polarizations, which independently obey the wave equation of Equation 2.13. Despite the similarity with electromagnetism, these two polarizations are not entirely analogous to the polarization of electromagnetic waves. In particular, we will find that the two polarization represent a spin two field and are mapped into each other by a 45 degree rotation of our coordinate system in the plane normal to the direction of propagation.

We also note two other unique properties of gravitational waves that are not shared by their electromagnetic counterparts. As we will see, gravitational waves interact very weakly with matter. While this property is what makes their detection so challenging, it is also what makes it rewarding. Gravitational waves will be unaffected by intervening matter, allowing a direct probe of the motion of matter in regions, such as the core of supernovae or the very early universe, which are otherwise inaccessible by electromagnetic observation. In addition, the length scale of gravitational wave sources and detectors are typically much smaller than the wavelength of the wave itself. As a result, gravitational waves provide a measure of the coherent bulk motion of matter within the source. Moreover, gravitational wave detectors are directly sensitive to gravitational wave strain, which varies inversely with the distance to a source, rather than the corresponding energy flux, which varies inversely with the square of the distance to a source.

2.1.2 Interaction with matter

We now consider the effect of gravitational waves on the motion of free masses. There are two equivalent approaches to this problem, and as they are directly related to the detection of gravitational waves, we will find it instructive to consider both.

The first approach takes advantage of our use of the transverse-traceless gauge. Specifically, it turns out that with this choice of gauge, the spatial components of our coordinate system are fixed to the geodesics of freely falling masses. That is, the

coordinate position of free masses is a constant. However, this does not imply that a passing gravitational wave has no observable effects. This fact is evident if we instead consider the round trip travel T time of a photon traveling between two freely falling test masses as measured by an observer fixed to one of the masses. In this case, the proper interval s between reflections is necessarily zero for a photon. As a result, we find the relation

$$dt^2 = (1 + h_+)dx^2, \quad (2.17)$$

where for simplicity we are only considering masses separated in x . Since the end points of our path are fixed in the transverse-traceless gauge, we find that the round trip travel time of the photon is given by the integral

$$T = 2 \int_0^L (1 + h_+)^{1/2} dx, \quad (2.18)$$

where L is the x -directed separation of the two masses in the absence of any gravitational wave. Assuming that h_+ is sufficiently smaller than unity and approximately constant over the round trip travel time of the photon, we find that

$$T \simeq 2\left(1 + \frac{h_+}{2}\right)L. \quad (2.19)$$

In more conventional units where c is not unity, this corresponds to a additional round trip travel time of

$$\Delta T = \frac{h_+ L}{c} \quad (2.20)$$

due to the effect of the gravitational wave.

A similar result is obtained if we instead consider the proper interval between test masses evaluated at fixed coordinate time. In this case we find that the change in proper interval is given by the expression

$$\Delta L = \frac{h_+}{2} L. \quad (2.21)$$

Thus h_+ is a measure of the fractional change in proper distance between two free masses. For this reason, we refer to h_+ as a gravitational-wave strain.

In the above examples, we have considered only the effect of a gravitational wave on free masses separated in the x direction. Repeating the same analysis for free masses separated in y , we find the same result except for a sign change. As a result, the gravitational wave effects a differential motion of free masses along the two axis. Similarly, the h_\times matrix elements are also found to effect a differential motion of free masses, but in this case the motion is along an orthogonal pair of axis that are rotated by 45 degrees with respect to the x - y axis. In general, we see that there are two distinct orthogonal polarizations of gravitational plane waves, a first which acts in the x and y directions, and a second which acts at a 45 degree angle to the first. For this reason, the polarizations are termed *plus* and *cross* respectively. In Figure 2.1 we display the effect of these two polarizations when such waves are normally incident on an otherwise circular ring of free masses. It is precisely this differential motion of free masses which motivates our development of interferometric detectors in the next section.




Figure 2.1: Effect of the two polarizations of gravitational waves on an otherwise circular ring of free masses. Reading from left to right, the top row depicts the effect of the plus polarization, while the bottom row demonstrates the effects of the cross polarization.

An alternative approach to studying the interaction of gravitational waves with matter is to consider the relative acceleration of two free masses due to the presence of a passing gravitational wave. Here, however, we do not work in the transverse-

traceless gauge, but in a rigid coordinate system that is fixed to the center of mass between the two free masses. In this coordinate system, the two masses then follow geodesics described by Equation 2.3. If we then consider the relative acceleration of the two masses, we find that their motion is equivalent to what would be observed if gravitational waves produce an effective tidal force of

$$f = \frac{mL}{2} \frac{\partial^2 h_+}{\partial t^2}, \quad (2.22)$$

where again we have assumed a separation L in the x direction. If we also consider tidal forces in y direction, and also for both gravitational wave polarizations, we find the alternating differential tidal forces depicted in Figure 2.2. It is not difficult to show that in this case we recover the same differential changes in separation that we identified above by reflecting photons. However, our picture of gravitational wave producing an effective tidal force will be beneficial when we consider resonant mass detector in the next section.




Figure 2.2: Tidal force lines associated to a gravitational wave propagating in the z direction. This figure has been reproduced from reference [16].

2.1.3 Generation

We turn now to the generation of gravitational waves. In analogy with generation of electromagnetic waves from accelerating charges, we expect gravitational waves to be generated by the motion of matter. Continuing the analogy, we note that energy

conservation, like charge conservation, rules out the possibility of monopolar sources of gravitational radiation. However, in contrast to electromagnetic waves, we note that there is only one type of gravitational “charge”. As a result, it is also not possible to produce gravitational radiation from dipolar motions of matter, a fact which also follows from the conservation of linear and angular momentum. Instead, the first radiative term in the multipole expansion of a time varying mass distribution is due to the quadrupole moment

$$I_{\mu\nu} = \int_V \left(x_{\mu\nu} - \frac{1}{3} \delta_{\mu\nu} r^2 \right) \rho(\mathbf{r}) d^3r. \quad (2.23)$$

When such a source is included on the right hand side of the Einstein field equation, we find that the resulting gravitational wave strain amplitude at a distance d from the source is given in conventional units by the expression

$$h_{\mu\nu} = \frac{2G}{c^4} \frac{d^2 I_{\mu\nu}}{dt^2}. \quad (2.24)$$

It is instructive at this point to consider rough estimates of the possible strain amplitude for astrophysical sources of gravitational waves. To do so, we consider the optimistic case where the quadrupolar moment of the source is due to the relativistic motion of a mass M ,

$$\frac{d^2 I_{\mu\nu}}{dt^2} \sim M c^2. \quad (2.25)$$

Even in this best case scenario, we find that expected strain amplitude is given by

$$h \lesssim \frac{1}{d} \frac{2GM}{c^2} \lesssim 10^{-19} \left(\frac{M}{M_\odot} \right) \left(\frac{d}{\text{Mpc}} \right)^{-1}, \quad (2.26)$$

which we recognize as the ratio of the Schwarzschild radius of the source to its distance from an observer. However, this is a very optimistic estimate. Typical sources are expected to radiate a significantly smaller fraction of their rest mass energy. An immediate consequence of this result is that there is little hope in detecting terrestrial sources of gravitational radiation. As such, the detection of gravitational waves is

necessarily an astrophysical endeavor, where it still remains a formidable challenge.

It is also instructive to consider a rough estimate of the maximum frequency range of gravitational waves. In this case, we assume that the gravitational radiation is generated by the periodic motion of a relativistic compact object and note that the maximum oscillation frequency of the source is limited by the round trip light travel time across its extent, and that the extent of a massive object is fundamentally limited by its gravitational radius, $2GM/c^2$. As a result, we find an approximate upper bound of

$$f \lesssim \frac{c^3}{4\pi GM} \sim 16 \left(\frac{M}{M_\odot} \right)^{-1} \text{ kHz}. \quad (2.27)$$

Finally, we note that the energy flux carried by gravitational waves is given by

$$I = \frac{c^3}{16\pi G} \left[\left(\frac{dh_+}{dt} \right)^2 + \left(\frac{dh_\times}{dt} \right)^2 \right], \quad (2.28)$$

which, assuming isotropic emission from a source at distance d , corresponds to a gravitational wave luminosity of

$$\mathcal{L} = \frac{c^3 d^2}{4G} \left[\left(\frac{dh_+}{dt} \right)^2 + \left(\frac{dh_\times}{dt} \right)^2 \right]. \quad (2.29)$$

2.2 Detectors

In this section, we provide a brief overview of current efforts to detect gravitational waves, which we define as the direct measurement of their interaction with matter. Specifically, we describe two classes of detectors: resonant mass detectors, which attempt to sense the tidal forces exerted by gravitational waves on semi-rigid bodies; and interferometric detectors, which attempt to measure relative fluctuations in the round trip travel time of photons traveling between orthogonally oriented pairs of free masses.

2.2.1 Resonant mass detectors

Historically, the first efforts to detect gravitational waves were attempts to measure the tidal forces exerted by gravitational waves on a single large test mass. The premise is that a passing gravitational wave will excite the fundamental resonant frequency of the detector, producing a detectable vibration of the mass. At the most basic level, the fundamental mode of such detectors can be modeled as a coupled oscillator, which is then excited by the tidal force described by Equation 2.22. As a result, such resonant mass detectors achieve astrophysically interesting sensitivity, but only in a relatively narrow frequency band of ~ 10 Hz around their resonant frequency.

Starting with the pioneering work of Weber[9] on the first gravitational-wave detectors, resonant mass detectors have greatly matured and increased in sensitivity over time. In contrast to the initial room temperature detectors, the current generation of resonant mass detectors now consist of massive cryogenic aluminum cylinders that are suspended in vacuum by multi-stage pendula. Typically such detectors have an approximate mass of 2×10^3 kg, an approximate length of 3 meters, fundamental resonant frequencies on the order of 900 Hz, and operate at temperatures of a few kelvins. The oscillations of such detectors are then read out using an impedance matched transducer consisting of a small coupled oscillator mounted axially at one end of the cylindrical detector.

Current or recently operating resonant mass detectors include the ALLEGRO detector in Baton Rouge, Louisiana[24]; the AURIGA detector in Lengaro, Italy[25]; the EXPLORER detector in Geneva, Switzerland[26]; the NAUTILUS detector in Rome, Italy[27]; and the NIOBE detector in Perth, Australia[28]. For the purpose of performing coincident searches for gravitational-wave bursts, these detectors have also entered into occasional cooperative data sharing agreements under the auspices of the International Gravitational Event Collaboration (IGEC) [29]. For reference, in Table 2.1 we summarize the properties of these five detectors as they were configured

during the 1997-2000 run of the IGEC collaboration[30]. A similar coincident search has also been carried out during 2001[13, 14] by the Ricerca Onde Gravitazionali (ROG) collaboration[31], which consists of only the EXPLORER and NAUTILUS detectors. These collaborative searches represent two of the most sensitive resonant mass detector searches to date, and we will return to them in section 7.9, where we compare their results with the results of the search performed in this work. As an example of the sensitivity of resonant mass detectors, including recent improvements in sensitivity and bandwidth, in Figure 2.3 we display the detector noise spectrum of the AURIGA detector taken in December of 2003 during the same time as the second LIGO science run[32].

Table 2.1: Summary of resonant mass detector properties during the 1997-2000 run of the International Gravitational Event Collaboration (IGEC). This information has been taken from Table 1 of reference [30].

Detector	Material Alloy	Mass [kg]	Length [m]	Frequencies [Hz]	Temperature [K]
ALLEGRO	Al5056	2296	3.0	895, 920	4.2
AURIGA	Al5056	2230	2.9	912, 930	0.2
EXPLORER	Al5056	2270	3.0	905, 921	2.6
NAUTILUS	Al5056	2260	3.0	908, 924	0.1
NIOBE	Nb	1500	2.8	694, 713	5.0

The gravitational wave strain observed by a resonant mass detector is given by the linear combination

$$h(t) = F_+ h_+(t) + F_\times h_\times(t), \quad (2.30)$$

where the angular position and polarization dependent response of the detector are contained in the coefficients

$$F_+ = \sin^2 \theta \cos 2\psi \quad (2.31a)$$

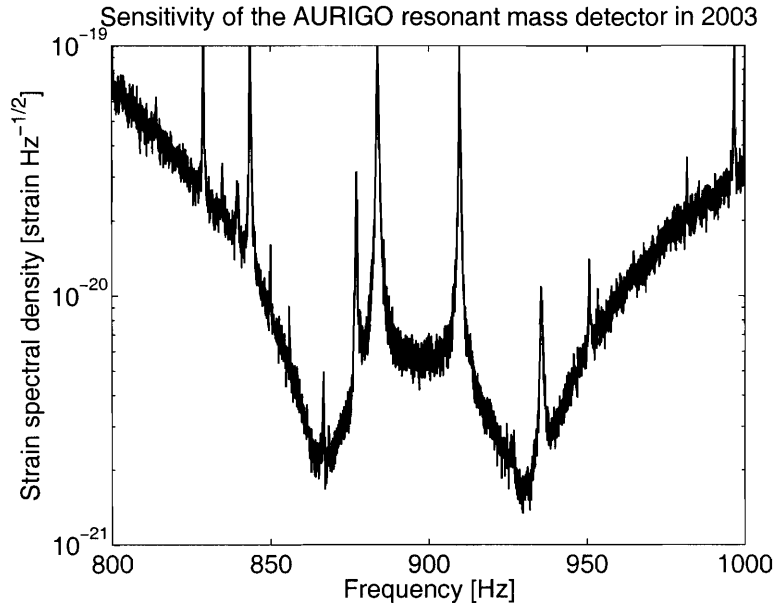


Figure 2.3: Strain spectral density for the AURIGA resonant mass detector at the same time as the second LIGO science run in December of 2003. The definition of strain spectral density used to characterize the detector sensitivity is described in more detail in section 3.3.1 and is the same quantity used to identify the sensitivity of interferometric detectors in Figure 2.7a and Figure 2.7b. The data for this figure is taken from Figure 1 of reference [32].

$$F_x = \sin^2 \theta \sin 2\psi. \quad (2.31b)$$

Here θ is the inclination angle of the source from the longitudinal axis of the detector, while ψ is angle that describes the choice of polarization convention. The resulting angular response of resonant mass detectors is shown Figure 2.4. The fact that resonant mass detectors exhibit poor directional sensitivity is evident from Figure 2.4. Although this makes it difficult to locate the position of a gravitational-wave source with a single detector, this wide angular acceptance is also beneficial in the sense that individual detectors are simultaneously sensitive to large portions of the sky.

Finally, we note that future plans for resonant mass detectors include incremental improvements in existing detectors to improve their sensitivity and slightly increase their bandwidth. In addition, spherical detectors have been proposed which would provide omni-directional sensitivity as well as being able to resolve the polarization

Angular response of resonant mass detectors




Figure 2.4: The angular response of resonant-mass gravitational wave detectors to gravitational waves with random linear polarization. Here the symmetry axis corresponds to the longitudinal axis of the bar. The response has a maximum value of $2/\pi$ for gravitational waves that are orthogonally incident on the bar and exhibits nulls for gravitational waves that are incident along the longitudinal axis of the bar.

of gravitational waves.

2.2.2 Interferometric detectors

The newest generation of gravitational wave detectors consist of interferometers.

The gravitational wave strain observed by an interferometric detector is given by

$$h(t) = F_+ h_+(t) + F_\times h_\times(t) \quad (2.32)$$

where the position and polarization dependent angular response of the detector is given by the coefficients F_+ and F_\times from the expressions

$$F_+ = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \cos 2\psi + \cos \theta \sin 2\phi \sin 2\psi \quad (2.33a)$$

$$F_\times = -\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi \sin 2\psi + \cos \theta \sin 2\phi \cos 2\psi \quad (2.33b)$$

Here, we have assumed standard spherical coordinates with an altitude coordinate θ ranging from 0 to π , with 0 indicating zenith, and an azimuthal coordinate ϕ ranging from 0 to 360. The polarization coordinate ψ is ranges from 0 to π , with 0 corresponding to alignment of the plus polarization with lines of constant ϕ and θ .

Angular response of interferometric detectors

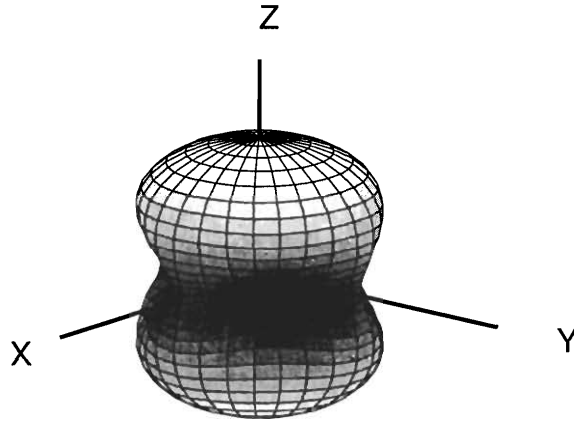


Figure 2.5: The angular response of interferometric gravitational wave detectors to gravitational waves with random linear polarization. Here the X and Y axes correspond to the orthogonal arms of an interferometric detector. The response has a maximum value of $2/\pi$ for sources at the zenith of the detector and exhibits nulls for gravitational waves that are incident in the plane of the detector at 45 degree angles to the interferometer arms.

A number of interferometric gravitational wave observatories have now been constructed around the world and are currently undergoing commissioning activities as they approach their design sensitivities. These include the two 4 km interferometers and one 2 km interferometer operated by the Laser Interferometer Gravitational-wave Observatory (LIGO) project in Hanford, Washington and Livingston, Louisiana[11]; the 600 meter GEO600 interferometer located near Hannover, Germany[33]; the 3 km Virgo interferometer located near Pisa, Italy[34], and the 300 meter TAMA detector located in Tokyo, Japan[35].

In combination with the existing resonant mass detectors, these detectors form a global network of gravitational-wave detectors that should allow strong verification tests for candidate events and permit the accurate recovery of the sky position, polar-

ization, and waveform of gravitational waves detected with sufficient signal to noise ratio. However, such multiple interferometer searches are only in their infancy as the global network of interferometric detectors comes on line. For now, we concentrate primarily on the network of the three LIGO interferometers.

The LIGO detectors are located at two sites in the United States. Two interferometers are located in Hanford, Washington. These interferometers share the same vacuum envelope and therefore co-aligned. However, one of the interferometers has 2 km long arms, while the other has 4 km long arms. As a result, these two detectors are expected to respond identically to the presence of a gravitational wave, except for an overall factor of two difference in detected strain. The third LIGO interferometer is located in Livingston, Louisiana and has 4 km long arms. The two sites are approximately aligned to respond to the same gravitational-wave polarization, except for slightly different tangent planes due to the curvature of the Earth's surface.

A simplified schematic of the LIGO interferometers is shown in Figure 2.6. At the most basic level, the LIGO detectors consist of kilometer scale Michelson interferometers illuminated by 10 Watt solid state Nd:YAG lasers with a wavelength of approximately 1 micron. However, in order to increase the accumulated phase delay resulting from incident gravitational waves, the interferometer arms are in practice comprised of Fabry-Perot resonant cavities such that the approximate light storage time in the arms is on the order of 1 millisecond. In addition, in order to decrease the fractional measurement error due to photon counting statistics, a higher circulating light power is maintained in the interferometer by introducing a recycling mirror at the interferometer input and operating the detector such that the output is maintained at a minimum of the interference pattern.

The gravitational wave strain incident on the LIGO interferometers is read out using a heterodyne method based on an extension of the Pound-Drever-Hall technique. In this method, a differential phase modulation is applied to the light in the interferometer at a radio frequency f_m on the order of 25 MHz. Due to operation at a minimum

in the interference pattern, in the absence of any gravitational-wave strain, this modulation results in fluctuations at a frequency $2f_m$ in the optical power observed at the interferometer output. However, in the presence of non-zero gravitational-wave strain, an output signal is detectable at the resonant frequency f_m . In practice, however, the LIGO interferometers are actively maintained by a feedback control system at an operating point corresponding to minimum optical power at the interferometer output and the gravitational-wave strain signal is determined from the control signals necessary to maintain this condition.

In order to isolate the LIGO detectors from terrestrial noise sources, the interferometers are maintained in a vacuum of between 10^{-8} – 10^{-9} mbar. In addition, all of the optical tables within the vacuum envelope are isolated from ground vibrations by 3 alternating stacks of stainless steel and dissipative bronze springs such that ground motion is attenuated by a factor of f^{-6} above the resonant frequencies of this passive seismic isolation, which fall in the range from 1 to 10 Hz. The test masses are further isolated by pendular suspensions such that they behave like free masses for excitations above their resonant frequency of approximately 1 Hz. Feedback control is achieved by applying a magnetic force to the test masses using a quadrupolar arrangement of electromagnetic coils and permanent magnets in order to cancel the effect of slowly varying ambient magnetic fields.

Finally, we note the LISA, the Laser Interferometer Space Antenna is currently in the development stage with a tentative launch in 2013. Due to its long baseline and the obvious absence of terrestrial noise sources, LISA will probe gravitational-waves in the frequency band from 10^{-4} to 10^{-1} Hz. In contrast to terrestrial detectors, however, the dominant low frequency noise source for LISA is expected to be a stochastic background due to the cumulative emission to many coalescing compact objects.




Figure 2.6: Simplified schematic of the LIGO detectors. The detectors consist of a modified Michelson interferometer with Fabry-Perot cavity arms in order to increase the differential phase delay due to incident gravitational waves. The detector is nominally operated such that the optical power at the anti-symmetric photodiode is minimized. At this operating point, the entire interferometer behaves like a single mirror and a recycling mirror allows for greater circulating light power in the detector.

2.3 Sources

We now turn to potential astrophysical sources of detectable gravitational waves. As noted in the previous two sections, laboratory or other terrestrial sources of gravitational radiation occur at amplitudes well below those that we may hope to detect. Instead, the only plausible sources of detectable gravitational radiation are those from the relativistic motion of extremely compact objects. As a result, the detection of gravitational waves is necessarily an endeavor in astrophysics. As such, our observations are limited to only that which nature provides. However, beyond the fact that astrophysical sources are the only sources we can hope to detect, there is also a strong astrophysical motivation to search for such sources: The observation of gravitational waves from astrophysical sources will provide us with the opportunity to study the behavior of matter at extremely high densities and strong gravity that we are not privy to in the laboratory, or by any other means currently available to astrophysicists.

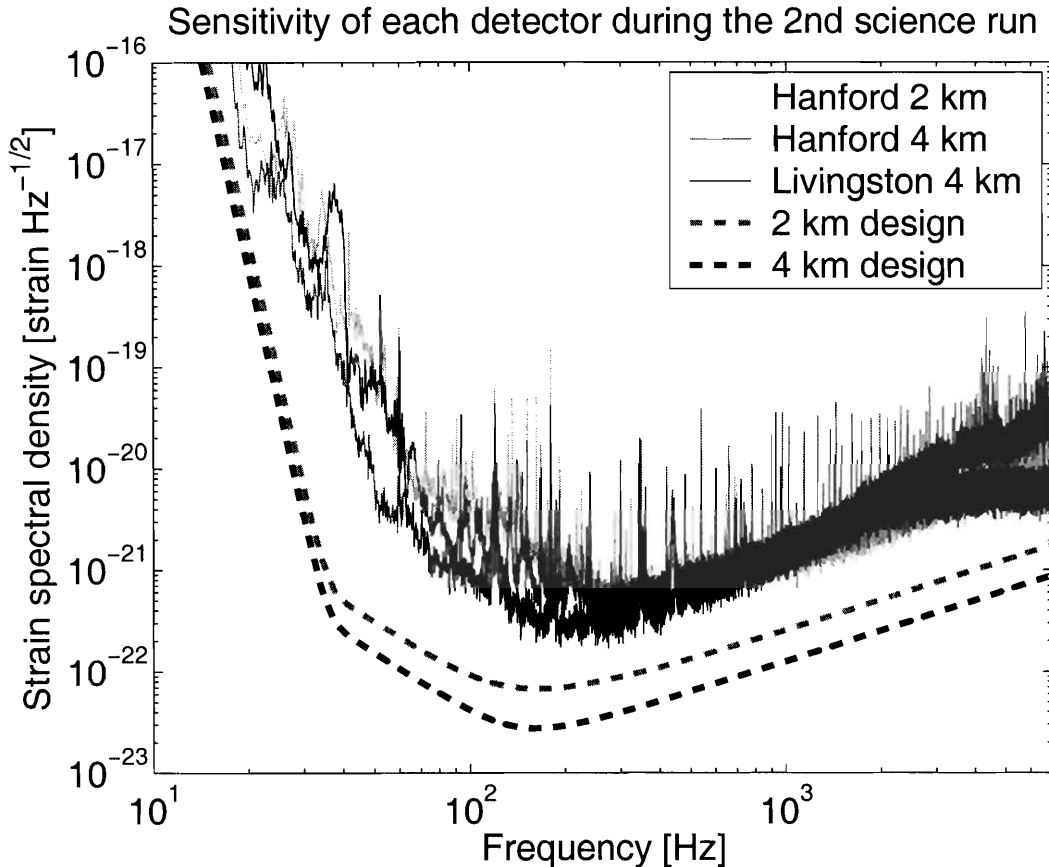


Figure 2.7a: The best sensitivity achieved by the Hanford 4km, Hanford 2km, and Livingston 4km detectors during the second LIGO science run are compared with the LIGO design sensitivity for 2 km and 4 km detectors. The definition of strain spectral density used to characterize the detector sensitivity is described in more detail in section 3.3.1 and is the same quantity used to characterize detector sensitivity in Figure 2.3.

In this section, we present an overview of a few of the more plausible astrophysical sources of gravitational waves that may potentially be detected by the first generation of interferometric ground based detectors. We begin by considering potential sources of transient gravitational radiation, which is the primary focus of this work. Such sources include the coalescence of compact objects such as binary neutron stars and binary black holes, core collapse supernovae of massive stars, ring down oscillations of perturbed black holes, instabilities of rotating neutron stars, gamma ray bursts, and possibly cosmic string cusps. For completeness, however, we also review potential sources of continuous gravitational radiation. In particular, we consider continuous

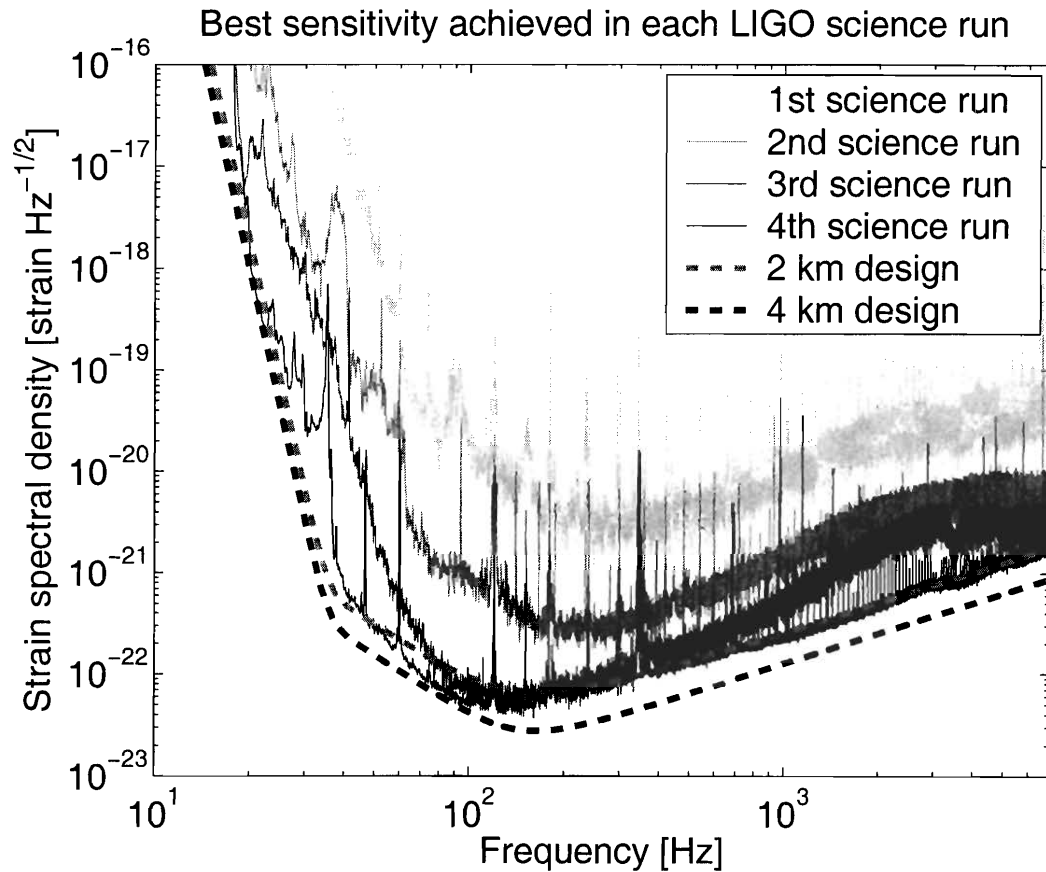


Figure 2.7b: The progression of detector sensitivity for the best performing detector during the first four LIGO science runs are compared with the LIGO design sensitivity for 2 km and 4 km detectors. The definition of strain spectral density used to characterize the detector sensitivity is described in more detail in section 3.3.1 and is the same quantity used to characterize detector sensitivity in Figure 2.3.

periodic emissions from pulsars, as well as a stochastic background of gravitational waves due to an ensemble of unresolved sources or relic gravitational waves from the very early universe. However, we should emphasize that our treatment is by no means exhaustive. Readers who are interested in a more exhaustive study of potential sources are referred to one of the many review articles on the subject[16, 17, 18, 19, 36] and the specific references listed below.

For each of these sources, we give a brief overview of what is known about the relevant physical processes that drive them and attempt to provide rough estimates of the expected amplitude and frequency scale of the resulting gravitational waves. When

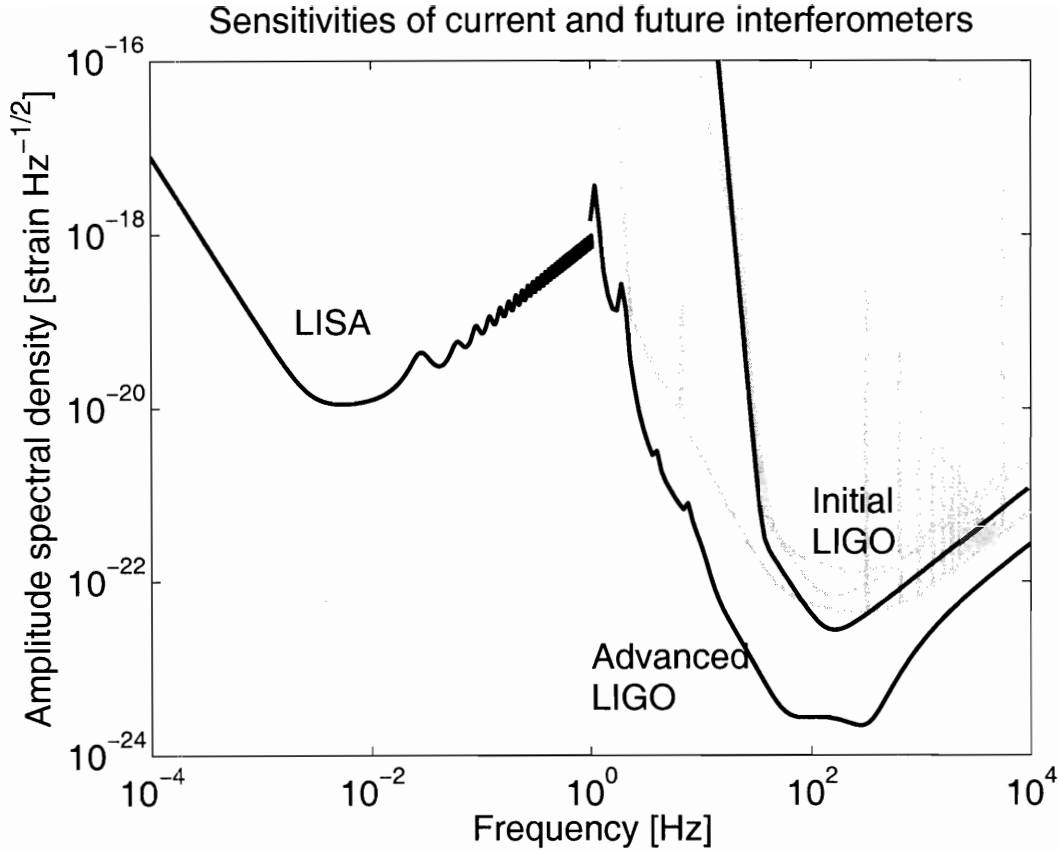


Figure 2.8: The design sensitivity of existing and future interferometric detectors. The initial design sensitivity of the 4 km LIGO detectors is compared with that of the planned second generation of LIGO detectors as well as the planned space-based interferometer LISA. For reference, we also show (in gray) the design sensitivities of the 2 km LIGO detector as well as the GEO600 and Virgo detectors (the lower of the three).

available, we also provide example waveforms, although often these will only be rough guesses of the expected behavior. In addition, we briefly describe the data analysis methods appropriate to different sources and provide estimates of the sensitivity such methods can hope to achieve. Finally, we review what is known about the populations of such sources, including any constraints from previous searches, as well as the prospects for detecting them with both current and future interferometric detectors.

2.3.1 Binary compact objects

Perhaps the most well known potential source of detectable gravitational waves is the coalescence of binary compact objects. Their importance results from the fact that detailed observations of the binary pulsar system PSR 1913+16 give strong evidence for the existence of gravitational waves as predicted by the General Theory of Relativity[2]. Their importance also results from the fact that the physics of the initial decay of the orbit is well understood and that we have accurate waveform predictions to guide our search. Finally, although numbering only a few, the observed number of binary pulsar systems and our understanding of the formation of such systems indicates that we have a good chance of discovering such sources with the first generation of interferometric detectors, and are virtually guaranteed a detection with more advanced detectors[17].

The coalescence of binary compact objects is typically considered in three different phases, based primarily on the different methods used to search for each phase.

The initial inspiral phase encompasses the gradual decay of the binary orbit as energy is lost by gravitational radiation. This phase of binary evolution is defined by the fact that their evolution can be described with very good accuracy by post-Newtonian methods and extends up to the inner most stable circular orbit, where General Relativity and the hydrodynamics of the individual objects become important. The predicted waveform for the inspiral phase of coalescing compact objects is well described by a chirp with an instantaneous frequency that varies as,

$$\frac{df}{dt} \propto M^{5/3} f^{11/3}, \quad (2.34)$$

and an amplitude that varies as

$$A \propto \frac{M^{5/3} f^{2/3}}{d}, \quad (2.35)$$

where M is the so-called chirp mass of the binary system,

$$\mathcal{M} = \frac{(M_1 M_2)^{3/5}}{(M_1 + M_2)^{1/5}}. \quad (2.36)$$

An example of a typical inspiral waveform is shown in Figure 2.9.

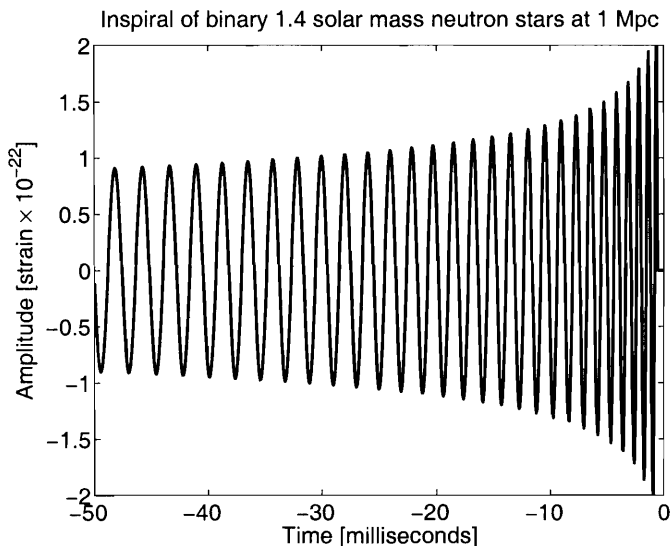


Figure 2.9: Simulated gravitational wave emission from the inspiral of 1.4 solar mass binary neutron stars at 1 Mpc as observed by an optimally oriented detector.

The second phase entails the actual dynamic merger of two-body system into a single object, and is the least understood of the three phases. For neutron star masses, the gravitational-wave signal produced by the merger phase is not well known, but is expected to occur above the sensitive frequency band of interferometric detectors. As a result, the search for the coalescence of binary neutron stars primarily targets the inspiral phase of the coalescence. However, for the merger of ~ 10 solar mass black holes, the signal from the merger and subsequent ringdown phase, although still poorly known, is expected to produce significant signal content in the frequency band of interferometric detectors.[37]. At the same time, the observation time for the inspiral phase is dramatically reduced due to the lower merger frequency and more rapid progression of the coalescence. As a result, a search for unmodeled bursts will form an important component in the detection of coalescing binary black holes than a search for the initial inspiral phase. Rough estimates of this expected gravitational

wave signal from the later stages of the merger and early ring-down of binary black holes has been simulated by Baker and collaborators[38, 39], an example of which is shown in Figure 2.10. However, the estimated waveforms are not predicted with nearly the same accuracy as for the inspiral phase. As a result, they do not comprise an appropriate set of template waveforms for a matched filter search. Nevertheless, they provide a rough idea of the signal space in which to search.

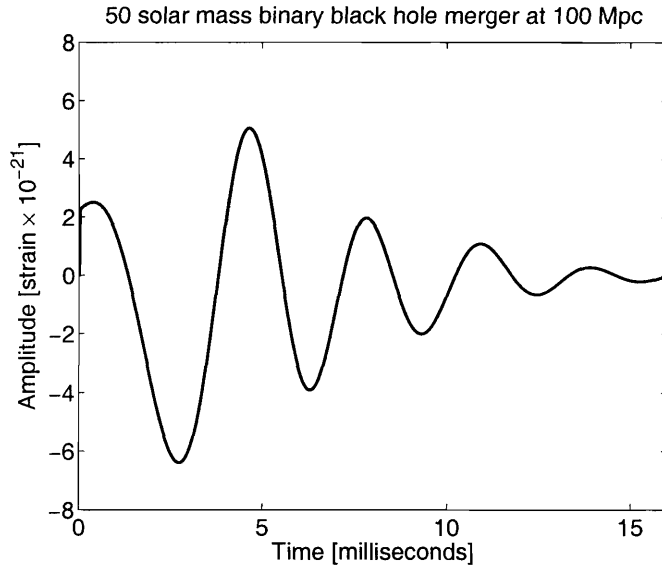


Figure 2.10: Example gravitational wave emission from the merger of binary black holes as simulated by Baker and collaborators[38, 39]. This waveform corresponds to an optimally oriented merger of two equal mass black holes with a combined mass of 50 solar masses at a distance of 100 Mpc.

The final phase occurs if the merger of two compact objects results in the formation of a new black hole. In this case, the initially perturbed black hole will undergo quasi-normal mode oscillations, relaxing back to a stationary rotating state via the emission of gravitational waves. However, since these black hole ring downs may also be produced by the formation of a new black hole in core collapse supernovae or other perturbative effects, we consider them separately in section 2.3.3.

Given the estimated galactic coalescence rates for both binary neutron stars and binary black holes, Cutler and Thorne[17] have estimated the potential detection rate for such events for both initial and advanced ground based interferometers. They es-

estimate that the coalescence of binary neutron stars should be detectable at a distance of 20 Mpc for initial LIGO and 300 Mpc for advanced LIGO, while the coalescence of binary black holes should be detectable at a distances of 100 Mpc for initial LIGO and cosmological redshifts z of 0.4 (~ 2 Gpc) for advanced LIGO. As a result, they report that a galactic binary neutron star coalescence rate of 1×10^{-6} to 5×10^{-4} per year corresponds to a detection rate of 3×10^{-4} to 0.3 per year for initial LIGO and 1 to 800 per year for advanced LIGO. In addition, they find that a galactic binary black hole coalescence rate of 1×10^{-7} to 1×10^{-5} per year corresponds to a detection rate of 4×10^{-3} to 0.6 per year for initial LIGO and 30 to 4000 per year for advanced LIGO.

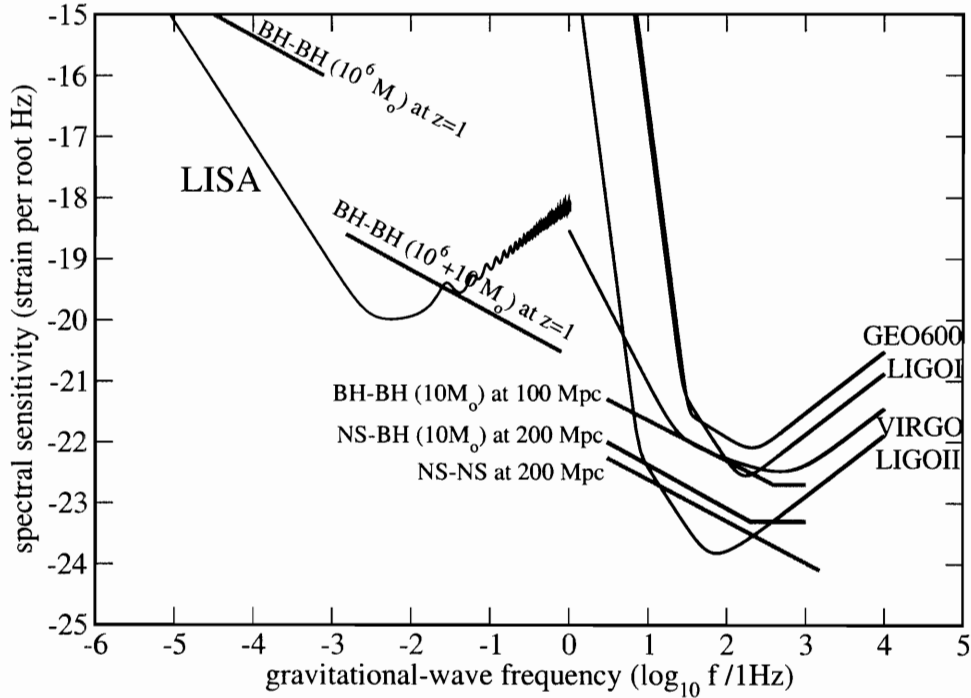


Figure 2.11: The expected signal amplitudes from the coalescence of binary compact objects of various masses are compared with the design sensitivities of current and future interferometric gravitational wave observatories. This figure has been reproduced from reference [19].

However, these rate estimates are based only on a known population of two binary systems occurring outside of globular clusters. As a result, there is a large uncertainty

due to the small statistical sample. In fact, with the recent discovery of the binary pulsar system PSR J0737-3039, Burgay and collaborators now predict a potential order of magnitude increase in the predicted merger rate. For the most favorable pulsar luminosity distribution model, the predicted detection rate of coalescing binary neutron stars by first generation interferometric detectors could be as high as 1 event per year at a 95 percent confidence level[7].

2.3.2 Core collapse supernovae

Aside from the coalescence of binary compact objects, one of the more plausible sources of detectable gravitational wave bursts is the rapid aspherical collapse and subsequent ejection of stellar matter associated with type II supernovae. This cataclysmic event is the typical fate of main sequence stars with initial masses in excess of roughly 9 solar masses[40]. When exhausted of their nuclear fuel, such massive stars contain a dense iron core that is initially supported against gravitational collapse by electron degeneracy pressure. However, as the core mass increases beyond the Chandrasekhar stability limit of 1.44 solar masses, photons in the core obtain sufficient energy to destroy heavy nuclei, a highly endothermic process resulting in loss of support and the production of free protons and neutrons. The free electrons supporting the core are then captured by these protons in a nuclear reaction releasing large amounts of energy in the form of neutrinos. As a result of this photodisintegration and electron capture, the electron degeneracy pressure supporting the core is lost and the core undergoes an extremely rapid collapse, reaching near relativistic velocities, and separating from the outer layers of the star.

There are two end scenarios for such a collapse. For stars with an initial mass in excess of approximately 40 solar masses, there is no known effect which can halt this process and the core completely collapses to form a black hole. For less massive stars, however, the collapse proceeds until the core reaches the density of nuclear matter, at which point neutron degeneracy pressure becomes sufficient to support the star. When

this occurs, the large kinetic energy of the collapsing core results in a bounce effect in which the infalling material is rapidly ejected. In the process, an immense amount of energy, approximately 10^{53} ergs, is released in the form of neutrinos. Optically, such supernovae are then observed when the ejected matter forms a shock front with infalling material from the outer layers of the star. Depending upon the mass of the initial star, the core remnant either stabilizes into a neutron star, or, for stars with initial mass in excess of 25 solar masses, collapses to form a black hole.

This model of core collapse supernovae has been largely confirmed by the fortuitous occurrence of supernova SN1987A in the nearby Large Magellanic Cloud and, in particular, the observation of neutrinos from SN1987A three hours prior to the onset of optical brightening[41].

The importance of core collapse supernovae to gravitational wave astrophysics results from the fact that such events provide a mechanism for a very large quantity of matter, on the order of one solar mass, to move at nearly relativistic velocities, up to a quarter of the speed of light. Such events, if even slightly aspherical, may then have a sufficient quadrupolar component to radiate a substantial amount of energy in the form of gravitational radiation. Furthermore, the observation of nearly coincident gravitational wave and neutrino bursts just prior to the telltale optical brightening of a supernova would provide an extremely strong confirmation of theory and an ideal laboratory to test our understanding of the hydrodynamics of rotational core collapse in the presence of very strong space-time curvature and nuclear densities.

Due to the importance of such a detection, a number of studies have been performed to try to predict the gravitational wave emission expected from core collapse supernovae[42, 43, 44]. Unfortunately, a full three dimensional treatment of the problem is not currently possible. Instead, such studies typically estimate the resulting waveform on the simplifying assumption of axisymmetric collapse of the rotating stellar core. In such models, the aspherical collapse required for gravitational wave emission results primarily from the effects of rotation at the time of

core bounce, but also due to convection and anisotropic neutrino emission. In order to study the variability of the predicted waveforms, the gravitational wave emission is typically simulated for families of progenitor stars, all of which differ slightly in their assumptions and initial conditions. In particular, Zwerger and collaborators[42] have performed axisymmetric Newtonian hydrodynamic simulations of 18 different rotating stellar models that differ by their amount of initial rotational kinetic energy, initial radial distribution of angular momentum, and adiabatic equation of state at sub-nuclear densities. In addition, Dimmelmeier and collaborators[43] have extended the work of Zwerger and collaborators to consider relativistic hydrodynamic simulations. Finally, Ott and collaborators[44] have considered the effect of a more realistic finite temperature nuclear equation of state, as well as more realistic massive star progenitor models. An example of two of the many waveforms predicted by such models is shown in Figure 2.12

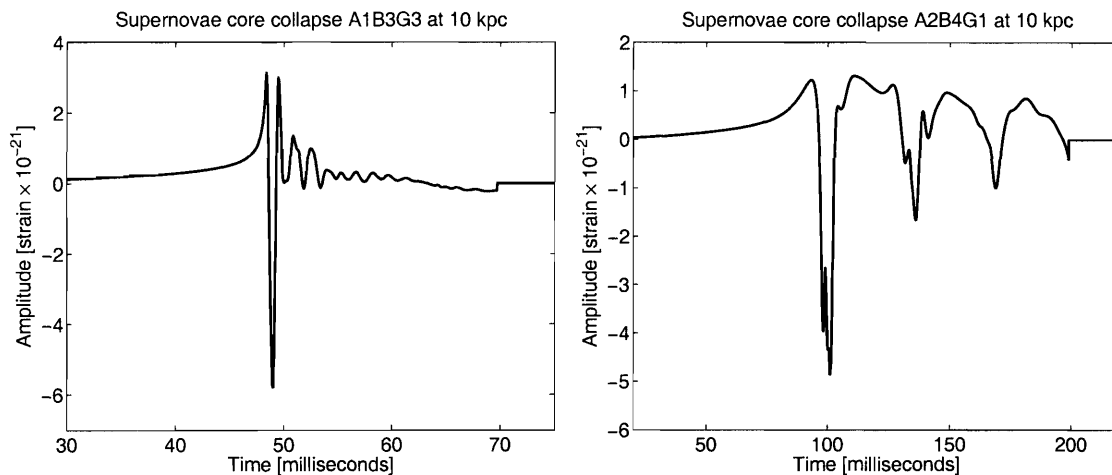


Figure 2.12: Example gravitational wave emission from core collapse supernovae simulated by Dimmelmeier and collaborators[43]. The two distinct waveforms are representative of two classes of waveforms observed in simulations and result from different assumptions regarding the length scale for differential rotation of the core, initial rotational kinetic energy, and the adiabatic index at sub-nuclear densities. Other simulations[42, 44] have produced a similar variety of waveforms with slightly different methods.

Collectively, these studies predict gravitational wave bursts with typical frequencies in the range from 50 to 3000 Hz, total emitted energy ranging from 1×10^{-11} to $4 \times 10^{-7} M_{\odot} c^2$, and peak strain amplitudes in the range from 2×10^{-23} to 4×10^{-20}

for bursts that are at a distance of 10 kpc. If such events occur within our own Galaxy, it is therefore likely that many of them would be detected by the initial generation of Earth based interferometric detectors. Even so, since the expected rate of core collapse supernovae within our galaxy is well known to be about 1 per 50 years, the likelihood that such an event will occur during the order 1 year observation times of the first generation of detectors is small.

It is also important to note that these studies are far from conclusive. While much progress has been made on individual fronts, an accurate treatment of the entire problem has not yet been feasible. Among the many outstanding issues are the need for a fully general relativistic treatment of the problem in three dimensions, the inclusion of magnetic fields, a more accurate treatment of the relevant microphysics and neutrino transport, and a more accurate accounting of progenitor stars. Indeed, the observation of gravitational waves from core collapse supernovae would shed much needed light on the physics of rotational core collapse and provide strong constraints on any future models.

A consequence of this uncertainty is that the gravitational waveforms of core collapse supernovae are not known to sufficient accuracy to permit the optimal linear search strategy of matched filtering, which we describe in more detail in section 3.2. Instead, any search for core collapse supernovae is necessarily a search for statistically significant events which exhibit consistency among multiple detectors. Although applicable to any transient source, such a search is the primary focus of this work. Nevertheless, the waveforms presented here are useful mainly in that they provide a very rough estimate of the behavior that may be expected from core collapse supernovae. In fact, in section 7.6, we characterize the performance of our search algorithm by evaluating its ability to detect such waveforms.

2.3.3 Black hole ring down

Besides the inspiral of binary compact objects, the only other potential source that is sufficiently well understood to permit a matched filter search is the relaxation of perturbed black holes. These are expected to produce oscillations at well defined quasi-normal mode frequencies that rapidly decay by the emission of gravitational radiation as the black hole returns to a stationary rotating state.

Such ring downs are expected as a result of the merger of binary compact objects when at least one member of the binary is a black hole, or when the merger results in a newly formed black hole. As a result, the detection of a ring down signal is an integral component of searches for binary compact objects of sufficiently large mass. In addition, detectable black hole ring downs may also be produced by core collapse supernovae of stars with an initial mass greater than 25 solar masses that result in the formation of a black hole.

In general, the expected gravitational waveform emitted by perturbed black holes can be obtained by solving the Einstein field equations for a perturbative expansion of the exact Kerr solution for spinning black holes[45, 46]. However, it has also been shown that the resulting waveforms are well approximated by a simple exponentially decaying sinusoidal signal[47],

$$h(t) = \begin{cases} h_0 \exp(-\pi ft/Q) \cos(2\pi ft) & t \geq 0, \\ 0 & t < 0. \end{cases} \quad (2.37)$$

Here, the frequency of the dominant quadrupolar quasi-normal mode is found to be well described by

$$f \simeq 32 \left(\frac{M}{M_\odot} \right)^{-1} [1 - 0.63(1 - a)^{3/10}] \quad \text{kHz}, \quad (2.38)$$

and the quality factor Q by

$$Q \simeq 2(1 - a)^{-9/20}. \quad (2.39)$$

Both of these parameters are expressed in terms of a dimensionless spin parameter a , which varies between 0 for non-spinning (Schwarzschild) black holes and 1 for maximally spinning (Kerr) black holes, and is related to the angular momentum J of the black hole by the expression

$$a = \frac{Jc}{GM^2}. \quad (2.40)$$

Finally, assuming that a fraction ϵ of the total mass energy of the black hole is radiation in the form of gravitational waves, the strain amplitude at a distance d is estimated to be

$$h_0 \simeq \frac{6 \times 10^{-21}}{[Q(1 - 0.63(1 - a)^{3/10})]^{1/2}} \left(\frac{r}{\text{Mpc}}\right)^{-1} \left(\frac{M}{M_\odot}\right) \left(\frac{\epsilon}{0.01}\right)^{1/2}. \quad (2.41)$$

From these expressions, we first note that the frequency range of Earth based gravitational wave observatories corresponds to the quasi-normal mode frequencies of black holes with masses between 10 and 600 solar masses. In addition, the expected peak strain amplitude is on the order of 2×10^{-21} for the ring down of 10 solar mass black holes at a distance of 20 Mpc, the approximate distance of the Virgo cluster.

It is intriguing to note that predictions by Flanagan and Hughes suggest that such events may actually be detectable at distances as far as 200 Mpc with the first generation of interferometric observatories and perhaps a z of 1 with advanced detectors[37]. Unfortunately, the expected rate of such events is poorly known. However, a rough estimate of the rate may be obtained by considering the population of binary black hole mergers, a primary candidate for detectable ring down events. While this population is estimated to be somewhat smaller than that of binary neutron stars, given the much greater distances at which black hole mergers and ring downs may be observed, the expected event rate may actually be as high as 0.5 per year for initial

LIGO and ranging between 10 and 2000 per year for advanced detectors[17].

It should be pointed out, however, that the search for ring downs is complicated by the expectation that exponentially decaying sinusoids are a common instrumental artifact. As a result, the detection of black hole ring downs requires careful consideration of instrument behavior, as well as corroborative evidence of a detection from multiple detectors. To date, efforts at identifying black hole ring downs have primarily consisted of the preliminary work by Creighton[48], Adhikari[49], and Tsunesada and collaborators[50] to develop the necessary methods for such a search. However, more comprehensive searches are anticipated as the first generation of interferometric detectors approach their design sensitivity.

Finally, it is interesting to note that the simple analytical form of the gravitational waveform expected from a ringing black holes permits us to readily determine both the black hole's mass and its angular momentum in the event of a detection. This is a particularly profound result, since such an observation would provide strong confirmation of some of the fundamental predictions of the General Theory of Relativity. In addition, the precursor waveform leading up to the observed ring down will provide insight into the causative perturbation and the interaction of the black hole with its surrounding environment, permitting the study of physics in a region of strongly curved space-time.

2.3.4 Gamma ray bursts

Gamma ray bursts are extremely energetic bursts of 1 keV to 100 MeV photons that last between 0.01 and 100 seconds[51]. Since their discovery in the late 1960s, these bursts have been observed by a series of satellites to be randomly distributed on the sky, consistent with an isotropic distribution. Such bursts are detected on average about once per day, and have been classified into two populations, those that are shorter than 2 seconds in duration, and those that last longer.

Based on the cumulative red shift measurements and absorption line studies of optical afterglow counterparts to long duration gamma ray bursts, such events are now thought to occur at cosmological distances with a z of approximately unity (~ 5 Gpc), although they have also been observed as close as 35 Mpc. Moreover, optical afterglow observations of the particularly bright long duration gamma ray burst GRB030329 revealed the characteristic spectrum of a type Ic core collapse supernova at a cosmological redshift z of 0.17 (~ 2 Gpc). As a result, the current consensus among astrophysicists is that such long duration gamma bursts are most likely due to the collapse of massive stars and the formation of a black hole. If this consensus is indeed true, then given their distance and our discussion in section 2.3.2 of gravitational waves from similar type II core collapse supernovae, the prospects for the detection of gravitational waves from such events seems remote.

However, evidence also exists that indicates that such events are likely strongly beamed, suggesting a much larger population of events that we do not observe in gamma rays, but may observe gravitationally. In addition, given their distance, such bursts must also emit an extraordinary amount of energy, on the order of 10^{51} ergs, in the form of gamma rays. Although it is speculated that this emission consists of highly relativistic jets along the rotational axis of the collapsing star, the exact mechanism for such emission remains to be understood. As a result, our search for a gravitational wave counterpart to gamma ray bursts is primarily motivated by the possibility that this unknown mechanism of gamma ray emission may also be capable of producing detectable gravitational radiation. In addition, accurate knowledge of the time and sky position of a gamma ray burst should reduce the signal to noise ratio required to detect a corresponding gravitational wave burst with a high level of confidence[17].

Much less observational evidence is available for short duration gamma ray bursts. Although it is postulated that these bursts may be due to the merger of binary neutron stars, their mechanism is far from understood. To date, accurate redshift or absorption line studies of optical afterglow have not been possible for these short duration bursts, although the recent launch of the SWIFT satellite may soon provide a wealth

of information in this respect. All that is currently known about these bursts is that they, like long duration bursts, are also distributed isotropically on the sky. Again, we are motivated by this uncertainty to search for a gravitational wave burst counterpart to short duration bursts, which may help shed much needed light on the mechanisms responsible for such events.

Since the expected gravitational waveform associated with a gamma ray burst is unknown, the search for such events is necessarily one of identifying coincident statistically significant deviations from the baseline detector noise in multiple gravitational wave detectors. In this case, however, we also have the advantage of a well defined time around which to search. In addition, for many gamma ray bursts, we are also provided a well defined position on the sky, permitting advanced knowledge of the expected time of flight delay and relative response of between gravitational wave detectors. Although this information does not permit an accounting of the effect of polarization, to first order this can be ignored due to the similar alignment of the LIGO detectors in Hanford and Livingston. The typical search method is then to search for a common signal in pairs of detectors by applying a cross-correlation based search[52, 53, 54]. Typically, such searches are performed over a time interval extending from two minutes prior to the onset of a gamma ray burst to one minute after, which encompasses all of the plausible models for gravitational wave emission from such objects. In addition, the relatively short duration of this window permits a much more thorough investigation of the search space than is computationally feasible in other searches for astrophysically unmodeled bursts.

Since it fortuitously occurred during the second LIGO science run, such a search has in fact been performed for GRB030329[55]. Although no gravitational wave burst was observed, efforts are currently underway to extend this search to a much larger set of gamma ray bursts. We also note that two previous searches have also been carried out by the ROG collaboration using data from the EXPLORER and NAUTILUS resonant mass detectors to search for possible gravitational wave counterparts to 434 gamma ray bursts observed by the BATSE and BeppoSAX missions between 1991

and 2001[56, 57]. The results of these two searches yield a 95 percent confidence level upper limit of $h_c \lesssim 2.5 \times 10^{-19}$ for the amplitude of such events within the sensitive frequency band of the two detectors.

2.3.5 Cosmic string cusps

Another conjectured source of detectable gravitational radiation results from the formation of cusps in cosmic strings or cosmic superstrings. These strings are linear topological defects that are hypothesized to have formed in the early universe as a result of symmetry breaking phase transitions[58]. Such strings are also conjectured to have extremely high densities and move at relativistic speeds. They are therefore an interesting candidate for gravitational wave emission. In particular, the cumulative gravitational wave emissions from networks of cosmic strings are expected to give rise to a stochastic background of gravitational radiation that may be detectable by the first generation of interferometric gravitational wave detectors, and which we will briefly consider in section 2.3.7.

More recently, however, studies by Damour and Vilenkin[59, 60] have shown that the expected distribution of gravitational wave strain amplitude from networks of cosmic strings will exhibit strongly non-Gaussian behavior as a function of time . In particular, the occasional formation of cosmic string cusps will produce sharp bursts of gravitational radiation that stand out from the stochastic background and may also be detectable by the first generation of interferometric gravitational wave detectors. In addition, Damour and Vilenkin have shown that similar predictions may apply to cosmic superstring theories as well[61].

A particularly attractive aspect of the search for gravitational wave bursts from cosmic string cusps is that their waveforms have been shown to have the surprisingly simple analytical form[59, 62],

$$h(t) \propto |t - t_0|^{1/3}, \tag{2.42}$$

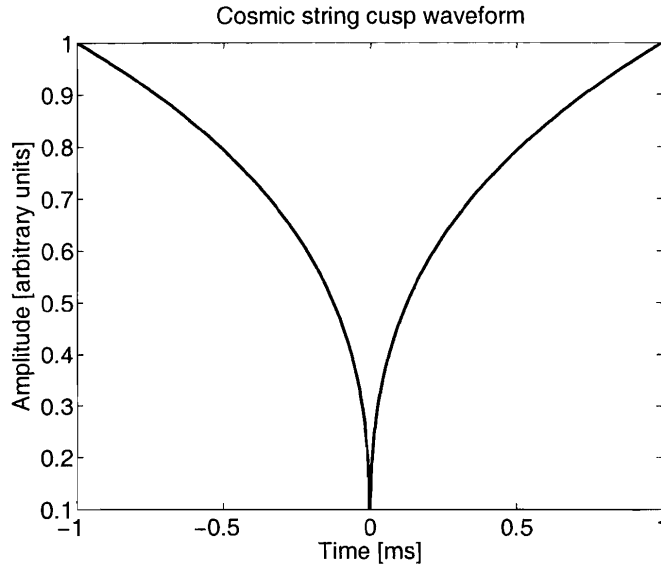


Figure 2.13: Example gravitational wave burst produced by a cosmic string cusp.

an example of which is shown in Figure 2.13. As a result, the detection of such bursts is most naturally suited to a matched filter search. Such a search is currently under development for the fourth LIGO science run and may soon make it possible to place constraints on more optimistic cosmic string and cosmic superstring theories.

2.3.6 Periodic sources

Rapidly rotating neutron stars, if they are not perfectly symmetric about their spin axis, may also be a potentially detectable source of gravitational waves.

There are a number of possible mechanisms by which these objects may develop sufficient asymmetry to radiate a detectable gravitational radiation. These mechanisms include small distortions of the shape of the neutron star away from axisymmetry, excitation of unstable neutron star oscillation modes, and neutron star spin precession.

The deviation of a neutron star from axisymmetry is typically measured by its ellipticity,

$$\epsilon = (I_{xx} - I_{yy})/I_{zz}, \quad (2.43)$$

where I is the quadrupole moment defined in Equation 2.23 and the component I_{zz} is the principal moment of inertia about the rotation axis. The detected gravitational wave strain amplitude due to a distorted neutron star at a distance d is directly proportional this ellipticity,

$$h = \frac{4\pi^2 G}{c^4} \frac{I_{zz} f_s^2}{d} \epsilon, \quad (2.44)$$

where f_s is twice the spin frequency of the neutron star.

The frequency scale of these sources is known by observations of pulsars to extend up to a maximum spin frequency of at least 641 Hz, corresponding to a gravitational wave frequency of 1282 Hz. As a result, a many rapidly rotating pulsars fall into the sensitive frequency band of interferometric gravitational wave detectors.

The best evidence that some spinning neutron stars are emitting gravitational radiation is due to Chakrabarty and collaborators[63, 64], who observe that the maximum spin frequency of neutron stars obtained by accretion in a low mass x-ray binary systems ~ 650 Hz is substantially smaller than the ~ 1.5 kHz limit imposed by neutron star breakup. This observation is in agreement with the suggestion of Bildsten[65, 66], who proposed that such systems reach an equilibrium with the spin-up due to accretion balanced by a spin down due to gravitational wave emission associated with accretion induced instabilities. Although not detectable by the initial generation of interferometric detectors, these systems are intriguing source for advanced detectors, which may be able to detect them if tuned for optimum sensitivity in the region around 600 Hz[17].

Similar to the ring down of black holes, neutron stars may also exhibit gravitational wave emission as a result of modal perturbations, which are typically associated with rotational instabilities of neutron stars. Dynamical instabilities are associated with non-axisymmetric perturbations with angular dependence $e^{im\phi}$ and grow on time scales associated with the sound crossing time of the neutron star, typically on the order of a rotation period, and may last many rotation periods depending upon the

differential rotation properties of the neutron star. The most well known of these instabilities is the $m = 2$ bar-mode instability, which may be excited newly formed neutron stars just after core bounce. Another class of instability, secular instabilities, are driven by fluid viscosity and gravitational wave emission. Of these, the f and r mode instabilities are considered the most likely to produce a detectable amount of gravitational radiation. In Figure 2.14, the potential gravitational wave emission from various rotational instabilities are compared with the design sensitivity of current and future interferometric detectors.




Figure 2.14: The gravitational wave signal amplitudes associated with neutron star instabilities are compared with the design sensitivities of current and future interferometric gravitational wave observatories. This figure has been reproduced from reference [19].

Contrary to searches for transient sources, searches for periodic sources may also benefit from the ability to coherently observe the signal over long periods of time. This approach results in an decrease in the signal to noise ratio necessary for detection that is inversely proportional to the square root of the integration time. In particular, searches for periodic signals are sensitive to gravitational wave strain on the order of

$$h_0 \simeq 11.4 \frac{S(f_s)}{T}, \quad (2.45)$$

where T is the observation time and $S(f_s)$ is the one-sided power spectral density of detector noise as defined by Equation 3.30 at twice the neutron star spin frequency. For sufficiently long observation times, searches for periodic signals can therefore detect signals with strain amplitudes well below that of transient signals.

While it is generally believed that the observed spin down of pulsars is due to magnetic dipole radiation, it is also possible to constrain the potential gravitational wave emission of pulsars by assuming that the observed spin down is entirely due to energy loss by gravitational radiation.

To date, searches for periodic emission of gravitational waves have fallen into two categories.

The first is a search for known pulsars. In this case, the expected gravitational wave frequency is well known from radio observations, as is the expected Doppler shift and detector response due to the time varying relative motion of the pulsar and its position and orientation on the sky. This knowledge therefore permits a heterodyne search for known pulsars in which the interferometric detector output is projected onto a compensated time varying sinusoid with a nominal frequency that is twice the known spin frequency of the pulsar.

Currently, using data from the second LIGO science run to search for 28 rapidly rotating isolated pulsars, the tightest bound obtained by this type of search is an upper limit of 4.5×10^{-6} for the ellipticity of PSR J2124-3358[67]. While this limit is

in excess of the maximum allowable ellipticity predicted by models of the neutron star crust, future LIGO science runs should produce substantially improved upper limits as the detectors approach their design sensitivity and observation times approach one year.

The second class of search is for unknown sources. However, such searches must still take into account the expected variability in the frequency and amplitude of periodic sources due to Doppler shifts and detector response. As a result, the parameter space for such a search is extremely large and presents a rather formidable computational challenge. An effort has recently been launched to perform this type of search using a global network of personal computers[68]. However, previous bandlimited searches for gravitational waves from unknown periodic sources have also been performed using resonant mass detectors[69, 70, 71]. These searches report upper bounds on the order of $h_c \lesssim 1 \times 10^{-23}$ for narrow frequency bands around 920 Hz.

2.3.7 Stochastic background

Another potential class of gravitational waves consist of an ensemble of random space-time fluctuations that collectively give rise to a stochastic background of gravitational radiation. Such a stochastic background can be produced by two possible mechanisms. One is due to the cumulative effect of an ensemble of relatively recent semi-transient sources, such as the coalescence of binary compact objects or core collapse supernovae. Indeed, a stochastic background from the early inspiral phase of binary compact objects is expected to be the limiting low frequency noise source for the LISA mission. The other possible source of a stochastic background is due to relic gravitational waves from the very early universe. Such a background would be analogous to the well known cosmic microwave background, which carries with it a wealth of information about the state of the universe at the time when electromagnetic radiation and matter decoupled. Similarly, since relic gravitational waves will experience negligible scattering by any intervening matter, we expect the observation

of a cosmological background of gravitational waves to provide important information about the universe at the time when gravitational radiation decoupled. Moreover, due to the extremely weak interaction of gravitation, the time scale for the decoupling of gravitational waves within the LIGO frequency band is estimated to be roughly 10^{-22} seconds after the formation of the universe[72, 73]. In contrast, the age of the universe at the time when electromagnetic radiation decoupled has been measured to be 3.8×10^5 years[74]. The observation of a cosmological background of gravitational radiation would therefore provide us with a direct view of the very early universe that is not currently accessible by any other means.

By convention, the spectrum of a stochastic background is specified in terms of the the dimensionless quantity

$$\Omega_{\text{gw}}(f) = \frac{f}{\rho_c} \frac{d\rho_{\text{gw}}}{df}, \quad (2.46)$$

where ρ_{gw} is the cumulative energy density spectrum of the stochastic background of gravitational waves and ρ_c is the current critical density required for a closed universe,

$$\rho_c = \frac{3c^2 H_0^2}{8\pi G}. \quad (2.47)$$

Here we take H_0 , the present day value of the Hubble constant, to be $71 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [74].

The quantity Ω_{gw} therefore describes the energy density of gravitational waves per unit logarithmic frequency interval relative to the critical density necessary for a closed universe. Given this definition, an energy density spectrum that falls as f^{-1} would have a equal energy per logarithmic frequency interval and therefore a constant value of $\Omega_{\text{gw}}(f) = \Omega_0$. Such a spectrum is predicted by both inflationary and cosmic string models in the frequency range of the LIGO detectors[72, 73].

Assuming a stochastic background that is isotropic, unpolarized, and stationary, it can be shown[75] that the corresponding one-sided power spectral density observed

by a gravitational-wave detector is given by

$$S_{\text{gw}}(f) = \frac{3H_0^2}{10\pi^2} f^{-3} \Omega_{\text{gw}}(f). \quad (2.48)$$

As a result, assuming a constant value of $\Omega_{\text{gw}}(f) = \Omega_0$, we expect that gravitational wave detectors will observe a strain spectral density that varies as $f^{-3/2}$, although we do not exclude the possibility of a frequency dependent $\Omega_{\text{gw}}(f)$.

A straightforward limit on the sensitivity of gravitational wave detectors to a stochastic background may then be obtained by comparing this predicted spectrum with the strain spectral density of detector noise. However, such constraints are naive in that they neglect the possible use of multiple detectors in order to differentiate between correlated and uncorrelated stochastic noise. Instead, a more sensitive measure of the stochastic background is obtained from the cross correlation of data from multiple detectors after taking into account their relative location and orientation, as well as any difference in their sensitivity[76, 75, 77]. Relative to a single detector measurement, such an approach offers an improvement in sensitivity by a factor of $(\Delta f T)^{1/2}$, where Δf is the characteristic bandwidth of the search and T is the integration time of the measurement. As a result, assuming one year of integration, initial LIGO is expected to be sensitive to a stochastic background of $\Omega_0 \sim 10^{-6}$, while advanced LIGO is anticipated to test Ω_0 at the level of 10^{-9} [78].

Although the first LIGO science run only achieved the value of $\Omega_0 \leq 41$ at the 90 percent confidence level, the second science run has achieved a preliminary constraint of $\Omega_0 \lesssim 0.03$, while the third science run is expected to achieve an approximate value of $\Omega_0 \lesssim 5 \times 10^{-4}$ [78]. For comparison, it is interesting to note that, within the LIGO frequency band, the standard model of big bang nucleosynthesis conservatively constrains cosmological backgrounds to $\Omega_0 \lesssim 10^{-5}$ [75, 79]. As a result, the sensitivity of initial LIGO should be sufficient to contribute to experimental bounds on the possible values of Ω_0 for cosmological backgrounds. It should also be noted that a stochastic background arising from more recent sources is not constrained by big bang




Figure 2.15: The stochastic background predicted by various cosmological models of the very early universe are compared with the constraints expected from interferometric gravitational wave observatories, as well as those from other methods. The vertical axis is presented in units that assume a present day Hubble constant of $100 \text{ km s}^{-1} \text{ Mpc}$. This figure has been reproduced from reference [73].

nucleosynthesis. Other constraints on stochastic gravitational wave backgrounds are derived by spacecraft ranging[80], pulsar timing measurements[81, 82], and large scale fluctuations in the cosmic microwave background known as the Sachs-Wolf effect[83, 84]. Even so, such constraints occur at very low frequency scales, and are well outside the frequency band of Earth based gravitational wave detectors. Finally, we note that neither initial nor advanced LIGO will likely test current models of inflation, which predict values of Ω_0 ranging from 10^{-15} to 10^{-13} . However, both detectors should begin to constrain some of the more optimistic predictions for stochastic gravitational wave backgrounds due to cosmic strings[72, 73]. In Figure 2.15, the approximate constraints expected from interferometric detectors are compared with those of other

measurements as well as with the expected stochastic background due to different cosmological models of the very early universe.

2.3.8 Expecting the unexpected

Finally, we note that there is of course the possibility of detecting the unexpected. This has been the case whenever the sky is observed with a novel or improved instrument. A few examples of such serendipitous discoveries include the discovery of the cosmic microwave background by Penzias and Wilson in 1964, the discovery of pulsars by Bell and Hewish in 1967, and the discovery of gamma ray bursts in the late 1960s by satellites tasked with monitoring the global ban on nuclear weapons testing. We expect that interferometric gravitational wave detectors will be no exception to this trend. They represent a significant advancement in a field that now appears poised to make the first direct detection of gravitational waves. Compared to the previous generation of resonant mass detectors, they offer nearly a factor of 10^3 increase in detector bandwidth, while at the same time providing almost a factor of 10^2 increase in sensitivity. Partly for this reason, the search methods presented in the remainder of this work attempt to make minimal assumptions about the waveform to be detected. While our approach has mainly been guided by our inability to accurately predict the waveform expected from many of the sources described above, we are also motivated by the desire to be sensitive to the unexpected.

Chapter 3

Burst detection

In this chapter, we present a parameterization for astrophysically unmodeled bursts of gravitational radiation that provides us with the language necessary to describe both their measurement and their detection.

In general, we define a gravitational-wave burst as time-varying strain in space that is sufficiently well localized in time that its time-domain amplitude is square integrable. Although this definition imposes no upper limit on the duration of gravitational-wave bursts, we also make the additional assumption that such bursts are shorter than one second. This choice conveniently constrains the space of possible signals, while still encompassing the majority of the potentially detectable transient astrophysical sources that fall within the sensitive frequency band of ground based interferometric detectors.

In what follows, we first define a characteristic gravitational-wave amplitude that plays a central role in predicting the detectability of gravitational-wave bursts. We also find it convenient to characterize bursts by their energy distribution in both time and frequency. In particular, our primary method of searching for bursts will be to create two-dimensional time-frequency maps that identify the time-varying fre-

quency content of potential bursts. We therefore define a number of parameters, such as central time, central frequency, duration, and bandwidth, which characterize the structure of bursts that are well localized in this time-frequency plane. These definitions then form a basis for characterizing bursts with more complex time-frequency structure, as well as allowing tests for consistency between multiple detectors.

Next, we describe the measurement of bursts by their projection onto a basis of functions that cover a specified region of signal space. In the process, we identify a multiresolution basis of complex-valued waveforms with minimum time-frequency uncertainty as most suited to the detection of astrophysically unmodeled bursts of gravitational radiation using one or more detectors. We then demonstrate a method of constructing such a basis that covers a targeted region of signal space with the minimum number of basis functions necessary to ensure a specified accuracy.

Finally, we consider the problem of measurement in the presence of detector noise and identify those parameters which are useful predictors of detectability and measures of detection confidence. In the process, we introduce a simple method for evaluating the approximate detectability of proposed astrophysical sources.

3.1 Parameterization of bursts

In what follows, we first present a representation independent measure of gravitational-wave burst amplitude. This then permits the definition of normalized time-domain and frequency-domain waveforms for arbitrary bursts, which in turn leads to a simple set of parameters that are ideally suited to describe bursts that are well localized in the time-frequency plane. We then briefly consider the applicability of this parameterization to bursts with more complicated time-frequency structure.

3.1.1 Characteristic strain amplitude

An arbitrary gravitational-wave burst has a time-domain and frequency-domain representation that form a Fourier transform pair¹,

$$h(t) = \int_{-\infty}^{+\infty} \tilde{h}(f) e^{+i2\pi ft} df \quad (3.1a)$$

$$\tilde{h}(f) = \int_{-\infty}^{+\infty} h(t) e^{-i2\pi ft} dt. \quad (3.1b)$$

Since we define gravitational-wave bursts to be square integrable in time, we may also define the characteristic strain amplitude, $\|h\|$, which, due to Parseval's theorem, may be computed in either the time-domain or the frequency-domain by

$$\|h\|^2 = \int_{-\infty}^{+\infty} |h(t)|^2 dt = \int_{-\infty}^{+\infty} |\tilde{h}(f)|^2 df. \quad (3.2)$$

The quantity $\|h\|$ is particularly convenient since it has units of dimensionless strain per square root Hz and is directly comparable to the amplitude spectral density of detector noise, which we will describe in more detail section 3.3.2. Consequently, the quantity $\|h\|^2$ has units of power spectral density and is conventionally referred to as the signal energy. However, this is not the same as the physical energy carried by the gravitational-wave, which we instead describe in section 3.3.3. As a result, care should be taken not to confuse the two.

3.1.2 Normalized wave-function

An arbitrary gravitational-wave burst may then be written in terms of its characteristic strain amplitude and a normalized time-domain or frequency-domain represen-

¹Note that in physics, one commonly encounters a different Fourier transform convention in which the exponents have the opposite sign. However, throughout this work we have chosen to use the convention that is more commonly found in the signal processing literature.

tation defined by the relations

$$h(t) = \|h\|\psi(t) \quad (3.3a)$$

$$\tilde{h}(f) = \|h\|\tilde{\psi}(f). \quad (3.3b)$$

Like the unnormalized representations, the normalized time-domain and frequency-domain representations also form a Fourier transform pair,

$$\psi(t) = \int_{-\infty}^{+\infty} \tilde{\psi}(f)e^{+i2\pi ft} df \quad (3.4a)$$

$$\tilde{\psi}(f) = \int_{-\infty}^{+\infty} \psi(t)e^{-i2\pi ft} dt. \quad (3.4b)$$

Moreover, the normalized representations are also defined to have unity characteristic strain amplitude,

$$\int_{-\infty}^{+\infty} |\psi(t)|^2 dt = \int_{-\infty}^{+\infty} |\tilde{\psi}(f)|^2 df = 1. \quad (3.5)$$

The resulting normalized waveforms therefore provide an amplitude independent means to describe the time-frequency structure of arbitrary gravitational-wave bursts and will prove convenient for the further parameterization of such bursts, as well as for evaluating the detectability of candidate waveforms in section 3.3.

3.1.3 Time-frequency parameterization

By interpreting the squared magnitude of the normalized time-domain and frequency-domain waveforms as probability density functions, we may also define the characteristic center time τ and characteristic center frequency ϕ of an arbitrary burst by the relations

$$\tau = \int_{-\infty}^{+\infty} t|\psi(t)|^2 dt \quad (3.6a)$$

$$\phi = 2 \int_0^{+\infty} f|\tilde{\psi}(f)|^2 df. \quad (3.6b)$$

Here, since we are primarily concerned with the distribution of signal energy, we integrate over only positive frequencies. As a result, we must be careful when applying this parameterization to bursts that have appreciable signal content at zero frequency, since such bursts will be subject to aliasing about zero frequency.

Similarly, we may define the characteristic duration and bandwidth of an arbitrary burst in terms of the second central moments in time and frequency.

$$\sigma_t^2 = \int_{-\infty}^{+\infty} (t - \tau)^2 |\psi(t)|^2 dt \quad (3.7a)$$

$$\sigma_f^2 = 2 \int_0^{+\infty} (f - \phi)^2 |\tilde{\psi}(f)|^2 df \quad (3.7b)$$

Finally, we may also define a dimensionless quality factor Q for bursts, which is just the ratio of center frequency to bandwidth,

$$Q = \frac{\phi}{\sigma_f}. \quad (3.8)$$

3.1.4 Localized bursts

It can be shown[85] that, for bursts with no zero frequency content, the characteristic duration and bandwidth as defined in Equation 3.7a and Equation 3.7b obey the uncertainty relation

$$\sigma_t \sigma_f \geq \frac{1}{4\pi}. \quad (3.9)$$

Therefore, a gravitational-wave burst cannot exhibit structure in the time-frequency plane with an effective time-frequency area less than this limiting value. We thus define bursts whose time-frequency area is on the order of this limiting value to be well localized.

Although computable for any square-integrable burst, the time-frequency parameters defined in the previous section are most appropriate to well localized bursts. In

particular, we note that for bursts that are well localized in time and frequency, the dimensionless quality factor Q is essentially a measure of the signal's aspect ratio in the time-frequency plane and is roughly equal to the number of oscillations of the time-domain waveform over the duration of the burst.

Furthermore, it can be shown[85] that the minimum time-frequency area given by Equation 3.9 is achieved by bursts whose time-domain and frequency-domain representations take the form of Gaussian windowed sinusoids,

$$h(t) = \|h\| \left(\frac{1}{2\pi\sigma_t^2} \right)^{1/4} \exp \left[-\frac{(t-\tau)^2}{4\sigma_t^2} \right] \exp \left[i2\pi\phi(t-\tau) \right] \quad (3.10a)$$

$$\tilde{h}(f) = \|h\| \left(\frac{1}{2\pi\sigma_f^2} \right)^{1/4} \exp \left[-\frac{(f-\phi)^2}{4\sigma_f^2} \right] \exp \left[-i2\pi\tau(f-\phi) \right]. \quad (3.10b)$$

Here, in order to allow for arbitrary phase, we have conveniently defined a complex valued waveform. This representation will also prove useful in section 3.2.2, in which a basis of such waveforms will be used to cover a finite region of the possible signal space. Real valued waveforms, which also achieve the minimum permissible time-frequency area, can be obtained by taking a linear combination of the real and imaginary parts of Equation 3.10a or a corresponding linear combination of the complex conjugate symmetric and anti-symmetric components of Equation 3.10b. Due to their minimum uncertainty nature and their relatively simple analytical form, these sinusoidal Gaussian bursts will prove to be a useful test case in chapter 6 and chapter 7, where they will be used to characterize the performance of our proposed search algorithm for gravitational-wave bursts. In the next section, we will also find that sinusoidal Gaussian bursts comprise a particularly useful measurement basis for the detection of unmodeled bursts gravitational radiation. As a result, such bursts play a central role in our proposed search algorithm. For reference, examples of such waveforms are shown in Figure 3.1 for a central frequency of 1 Hz and a quality factor of 10.

Note that our explicit use of the parameters ϕ and σ_f in our definition of sinusoidal Gaussian bursts implicitly assumes a two-sided frequency-domain interpreta-

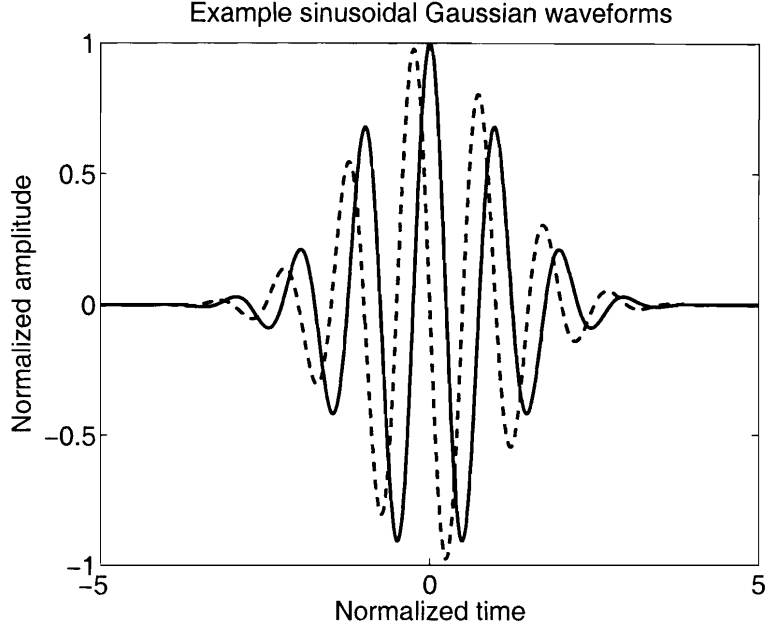


Figure 3.1: Example time-domain waveforms of a sinusoidal Gaussian burst with a central frequency of 1 Hz and quality factor Q of 10. Both even (solid) and odd (dashed) symmetric time-domain waveforms are shown, corresponding to the real and imaginary part of the complex valued waveforms given by Equation 3.11a and Equation 3.11b. Both waveforms are normalized to unity $\|h\|$.

tion, while the parameterization defined in Equation 3.6b and Equation 3.7b is one-sided in nature. These two parameterizations agree in the limit of large Q , where the resulting waveform has negligible signal content at zero frequency. However, for $Q \lesssim 3$, the two interpretations diverge due to aliasing effects at zero frequency. Although a Q dependent correction factor can be introduced to account for this discrepancy, in practice we will avoid this difficulty by only considering the measurement of bursts with $Q \gtrsim 3$. As a result, the detectability of bursts with $Q \lesssim 3$ will depend on their projection onto the space of signals with $Q \gtrsim 3$.

In order to evaluate the effect of this restriction, we therefore consider the limiting case of small Q , which results in simple Gaussian bursts of the form

$$h(t) = \|h\| \left(\frac{1}{2\pi\sigma_t^2} \right)^{1/4} \exp \left[-\frac{(t - \tau)^2}{4\sigma_t^2} \right] \quad (3.11a)$$

$$\tilde{h}(f) = \|h\| \left(2\pi\sigma_t^2 \right)^{1/4} \exp \left[-\pi^2\sigma_t^2 f^2 \right] \exp \left[-i2\pi\tau f \right]. \quad (3.11b)$$

As was the case for sinusoidal Gaussian bursts, simple Gaussian bursts also exhibit the minimum time-frequency uncertainty of Equation 3.9, but do so when viewed in the two-sided frequency-domain interpretation. Here we have avoided the difficulties of a one-sided vs. two-sided frequency-domain interpretation by expressing the frequency-domain waveform only in terms of the duration σ_t of the burst. For reference, application of the one-sided parameterization of section 3.1.3 yields the values

$$\phi = \left(\frac{1}{8\pi^3\sigma_t^2} \right)^{1/2} \quad (3.12a)$$

$$\sigma_f = \frac{1}{4\pi\sigma_t} \left(1 - \frac{2}{\pi} \right)^{1/2} \quad (3.12b)$$

$$Q = \left(\frac{\pi}{2} - 1 \right)^{-1}. \quad (3.12c)$$

Since such bursts fall outside of the targeted signal space of $Q \gtrsim 3$, they therefore provide an interesting test of detectability. As a result, we will also include such bursts in the set of waveforms used to characterize the performance of our search algorithm in chapter 7.

Finally, we note that it is occasionally convenient to refer to an alternative duration and bandwidth defined by

$$\Delta t = 2\pi^{1/2}\sigma_t \quad (3.13a)$$

$$\Delta f = 2\pi^{1/2}\sigma_f, \quad (3.13b)$$

such that the product, $\Delta t \Delta f$ is unity for minimum uncertainty bursts according to uncertainty relation of Equation 3.9. This definition will be particularly useful when quickly estimating the time-frequency overlap between minimum uncertainty bursts, where the number of possible non-overlapping minimum uncertainty bursts in a time T and bandwidth F cannot exceed the maximum number of independent measurements given by the product TF .

3.1.5 Non-localized bursts

While the simple time-frequency parameterization defined above adequately describes bursts that are well localized in time and frequency, it does not adequately parameterize the structure of more complicated bursts. However, even for more complicated bursts, their small scale time-frequency structure must obey the uncertainty relation of Equation 3.9. As a result, an appropriately selected basis of well localized waveforms is sufficient to resolve the structure of arbitrary bursts. In general, more complicated bursts can be decomposed into a linear combination of localized simple bursts, and a basis of such localized waveforms can be selected which adequately covers a specified region of parameter space. As a result, we do not attempt to parameterize the space of more general bursts and the detection of bursts with complicated time-frequency structure then depends upon the ability of our search algorithm to combine results from multiple time-frequency measurements.

3.2 Measurement of bursts

A number of search algorithms currently exist in order to identify transient sources of gravitational radiation. In general, all of these algorithms operate by linearly projecting the data under test onto a suitably chosen measurement basis. The resulting projections are then examined for statistically significant deviations from the baseline results expected in the absence of any gravitational-wave bursts. Based on our knowledge of potential astrophysical sources and the details of our data set, we may identify three types of searches that dictate the appropriate choice of measurement basis: those where the targeted sources are sufficiently well modeled that accurate waveforms exist, those where coincident data are available from multiple correlated detectors, and those where the waveform is poorly or completely unknown.

Matched filtering. In the first case, where the waveform is well modeled, the optimal measurement basis is well known[86]. In this case, the data under test are first whitened by a filter whose magnitude response is the inverse of the detector noise spectrum. The whitened data stream is then projected onto the set of targeted astrophysical waveforms, which have also been whitened by the inverse of the detector noise spectrum. This technique, known as matched filtering, is the optimal linear filter in the sense that it maximizes the observed signal to noise ratio, which is defined as the ratio of the maximum observed output when the targeted waveform is present in the data to the root mean square output when the targeted waveform is not present[87]. A number of astrophysical searches are currently being pursued using matched filtering. These include the search for inspiraling binary neutron stars[88, 89], inspiraling binary primordial black holes[90], inspiraling binary black holes[91], the ring down of perturbed black holes[48, 49, 50], and cusps from cosmic strings[62]. For well behaved noise, the expected performance of such searches are well known and therefore provide a benchmark for comparison with other search algorithms. We therefore return to matched filtering in section 3.3, where the signal to noise ratio achieved by such a search is used to characterize the detectability of arbitrary gravitational-wave bursts.

Cross-correlation. In the second case, where coincident data are available from multiple correlated detectors, the typical approach is to project an interval of data from one detector onto corresponding data from another detector. The resulting signal to noise ratio is maximized when the projected interval is similar in the duration to the duration of the burst. As a result, such searches typically consider a range of intervals in order to be optimally sensitive to gravitational-wave bursts at many time scales. Currently, due to the resulting computational cost, cross-correlation is only used to perform deep searches of short data segments when other information, such as the coincident observation of a gamma ray burst[55, 54, 53] or the previous identification of a candidate gravitational-wave burst[92, 93, 94], indicate an increased chance of detection. Such searches are limited, however, to pairs of detectors that are sufficiently aligned to be sensitive to the same gravitational-wave polarization. Extensions of

this approach to multiple non-aligned detectors are currently being developed[95, 96] based on initial work by Gursel and Tinto[97] and Flanagan and Hughes[98] on the inverse problem for gravitational-wave bursts.

Abstract bases. In the third case, where the waveform is poorly or completely unknown, the data under test are typically projected onto a convenient basis of abstract waveforms that are chosen to cover a targeted region of signal space. In this case, the detection of arbitrary gravitational-wave bursts depends upon their projection onto the particular basis under investigation. However, in searching for statistically significant events, such searches also typically employ methods for clustering the measurements from neighboring or overlapping basis functions in order to more optimally detect signals that are not well represented by the particular choice of basis. A number of different waveform basis have been proposed to search for unmodeled bursts of gravitational radiation, and such searches may be broadly classified into two categories: time-domain searches[99, 100, 101, 102], in which the primary basis consists of delta functions in time, and time-frequency searches[103, 104, 105, 106], in which the typical basis consists of windowed complex exponentials or wavelets. Such searches can easily be extended to multiple detectors by requiring coincident detection of statistically significant events, as well as consistency between the projections observed in each detector. In the special case of correlated detectors, such projections can additionally act as intermediaries to a direct cross-correlation comparison between the data from each detector. In this case, the use of an intermediate basis can also provide an efficient means of performing the comparison over a range of time and frequency scales.

3.2.1 Multiresolution analysis

In this work, we seek an algorithm for the detection of astrophysically unmodeled bursts of gravitational radiation that fall within the sensitive frequency band of ground-based interferometric detectors. We also desire a sufficiently general search

strategy that can be applied to one or more detectors regardless of their relative alignment. We therefore seek an abstract basis that efficiently covers a finite region of the time-frequency plane.

A number of well known time-frequency bases are already available from which to choose. These include the traditional short-time Fourier transform[85] as well as a large variety of more recently developed wavelet decompositions[107, 108]. Typically, in order to simplify the process of signal reconstruction, such bases are constructed to be both orthogonal and complete: the projection of any basis function onto any other basis function is zero, and the number of basis functions is exactly equal to the dimension of the signal space. Here, however, we are primarily interested in the detection of signals, not their reconstruction. The performance of our search is then determined by the maximum projection achieved by a single basis function. As a result, we will instead find it useful to implement an overcomplete basis in order to improve our prospects for identifying such a projection. In particular, we choose to implement an overcomplete multiresolution basis of complex-valued waveforms with minimum time-frequency uncertainty. That is, we take as our basis functions the complex-valued sinusoidal Gaussian bursts of Equation 3.10a and Equation 3.10b.

This choice is motivated by a number of considerations.

Recall that the time-frequency structure of gravitational-wave bursts is constrained by the minimum uncertainty relation of Equation 3.9. As a result, by selecting a basis of minimum uncertainty waveforms, we are able to maximally resolve the time-frequency structure of arbitrary bursts of gravitational radiation. At the same time, the requirement of minimum uncertainty does not separately constrain the time and frequency scale of such structure; only their product is constrained. As a result, our measurement basis must also incorporate multiple time scales in order to maximally resolve time-frequency structure over the time scales of potential gravitational-wave bursts. In the signal processing literature, such bases are commonly referred to as multiresolution bases. In our case, we simply allow the Q of our basis functions to be

specified independently their bandwidth.

Measurement with such a basis then provides the tightest possible time-frequency bounds on candidate events. This has two significant benefits. First, we may identify the best match set of non-overlapping minimum uncertainty waveforms that concentrate the total energy of candidate events into the fewest number of basis functions. As a result, we are able to maximize the observable signal to noise ratio for the minimum uncertainty structure of arbitrary gravitational-wave bursts. Second, we are able to apply the strictest possible tests of time-frequency coincidence between multiple detectors. As a result, we are able to minimize the rate of false detections resulting from the accidental coincidence of unrelated events.

A multiresolution time-frequency basis naturally includes the basis of time-domain delta functions in the limit of short duration and low Q . Although our targeted signal space does not extend below $Q \sim 3$ due to aliasing concerns, in practice such low Q bursts are nearly indistinguishable from $Q \sim 3$ bursts due to the masking effect of the low frequency seismic noise present in ground-based interferometric detector data.

Since the selected basis is overcomplete, the resulting basis functions are in general not orthogonal. As a result, we must be careful to account for their statistical interdependence when interpreting our results. However, we also note that the expected false detection rate is not affected by our adoption of an overcomplete basis, but is instead determined by the volume of the targeted signal space.

The choice of a complex-valued basis is a convenience that eliminates the need to separately include phase as an additional parameter of the search. Moreover, in chapter 5, we will take advantage of Fourier transform theory in order to implement such a complex-valued multiresolution basis in a computationally efficient way.

Finally, we note that the relatively simple analytical form of complex-valued sinusoidal Gaussians makes possible the optimal tiling of the targeted signal space considered in the next section.

3.2.2 Basis selection

Given the choice of a complex-valued multiresolution basis of minimum uncertainty waveforms, we now consider the required number and placement of basis functions in order to adequately cover a specified region of parameter space. There are two competing goals. On the one hand, for high detection efficiency, we wish to ensure that any well localized burst that falls within the targeted parameter space is closely matched to a basis function. On the other hand, for computational efficiency, we desire to minimize the number of required basis functions.

To accomplish these goals, we follow a procedure similar to that used in optimal matched filtering searches to select a set of template functions[109, 110, 111, 112, 48]. In particular, we tile the targeted signal space such that the mismatch between an arbitrary well localized burst and the closest measurement tile results in no more than a predefined fractional loss in the measured signal energy.

Fortunately, in contrast to the typical matched filtering case, we do not need to consider the effect of detector noise when selecting our basis. Instead, since our basis functions do not correspond to a specific astrophysical source, we are free to apply them to data that has first been whitened by the techniques described in chapter 4. As a result, we effectively perform a search of the whitened data stream rather than the gravitational-wave data stream. In practice, however, the difference is minor. For typical interferometric detector noise, the effective gravitational-wave basis functions are very similar to sinusoidal Gaussian bursts. In section 6.3, we demonstrate the negligible effect this difference by evaluating our proposed search algorithm on simulated detector noise, and find that the resulting performance is in very good agreement with the predicted performance based on the assumption of ideal white noise.

In order to construct our basis, we first rewrite the complex-valued sinusoidal Gaus-

sians bursts of Equation 3.10a and Equation 3.10b in their normalized form,

$$\psi(t; \tau, \phi, Q) = \left(\frac{8\pi\phi^2}{Q^2} \right)^{1/4} \exp \left[-\frac{4\pi^2\phi^2}{Q^2} (t - \tau)^2 \right] \exp [-i2\pi\phi(t - \tau)] \quad (3.14a)$$

$$\tilde{\psi}(f; \tau, \phi, Q) = \left(\frac{Q^2}{2\pi\phi^2} \right)^{1/4} \exp \left[-\frac{Q^2}{4\phi^2} (f - \phi)^2 \right] \exp [-i2\pi\tau(f - \phi)]. \quad (3.14b)$$

In doing so, we have also taken into account the minimum uncertainty property of sinusoidal Gaussian bursts, which requires that their duration and bandwidth cannot be independently specified. As a result, only three parameters are needed to describe our basis functions. Here, we take these parameters to be the center time τ , center frequency ϕ , and quality factor Q . Implicit in this choice is the assumption that the gravitational-wave bursts that are detectable by ground-based interferometric observatories primarily fall within a particular range of frequency and Q .

We begin by considering the inner product of two sinusoidal Gaussian bursts that differ in time, frequency, and Q by the amounts $\delta\tau$, $\delta\phi$, and δQ respectively. The resulting inner product, whose magnitude is a measure of the recovered signal amplitude, is a function of the mismatch between the two functions and is given by the expression,

$$\alpha(\delta\tau, \delta\phi, \delta Q) = \int_{-\infty}^{+\infty} \psi(t; \tau, \phi, Q) \psi^*(t; \tau + \delta\tau, \phi + \delta\phi, Q + \delta Q) dt. \quad (3.15)$$

Fortunately, for the case of complex-valued sinusoidal Gaussian bursts, an exact closed form analytical expression exists for the recovered signal amplitude. Even so, the derivation of this result is too extensive to include here. Instead, we simply note that the recovered signal energy is found to be

$$|\alpha(\delta\tau, \delta\phi, \delta Q)|^2 = \frac{2x_\phi x_Q}{x_\phi^2 + x_Q^2} \exp \left(-\frac{8\pi^2\phi^2 x_\phi^2 \delta\tau^2}{Q^2(x_\phi^2 + x_Q^2)} \right) \exp \left(-\frac{Q^2 x_Q^2 \delta\phi^2}{2\phi^2(x_\phi^2 + x_Q^2)} \right), \quad (3.16)$$

where we have also introduced the additional parameters

$$x_\phi = 1 + \delta\phi/\phi \quad (3.17a)$$

$$x_Q = 1 + \delta Q/Q \quad (3.17b)$$

in order to simplify our notation.

Next, we define the fractional energy loss due to mismatch by the expression

$$\mu(\delta\tau, \delta\phi, \delta Q) = 1 - |\alpha(\delta\tau, \delta\phi, \delta Q)|^2. \quad (3.18)$$

However, instead of substituting the exact expression of Equation 3.16 for the recovered signal energy, we will find it more convenient to use its second order expansion,

$$\mu(\delta\tau, \delta\phi, \delta Q) \simeq \frac{4\pi^2\phi^2}{Q^2} \delta\tau^2 + \frac{2+Q^2}{4\phi^2} \delta\phi^2 + \frac{1}{2Q^2} \delta Q^2 - \frac{1}{\phi Q} \delta\phi \delta Q. \quad (3.19)$$

Following the approach of Owen[109], this then suggests the definition of a metric,

$$\delta s^2 = g_{\tau\tau} \delta\tau^2 + g_{\phi\phi} \delta\phi^2 + g_{QQ} \delta Q^2 + 2g_{\phi Q} \delta\phi \delta Q, \quad (3.20)$$

to measure the space of sinusoidal Gaussian bursts in terms of the fractional energy loss due to mismatch. The components of this metric are then given by

$$g_{\tau\tau} = \frac{4\pi\phi^2}{Q^2} \quad (3.21a)$$

$$g_{\phi\phi} = \frac{2+Q^2}{4\phi^2} \quad (3.21b)$$

$$g_{QQ} = \frac{1}{2Q^2} \quad (3.21c)$$

$$g_{\phi Q} = -\frac{1}{2\phi Q}. \quad (3.21d)$$

Given this metric, we then seek a finite set of projections such that the worst case energy loss due to mismatch never exceeds a prescribed threshold μ_{\max} within the signal space of interest. At the same time, for computational efficiency, we wish to

minimize the necessary number of basis functions.

To do so, we choose to distribute our basis functions on a cubic lattice in the space defined by the metric of Equation 3.20. This situation is sketched in Figure 3.2. While it is known that this is not the optimal solution to the three dimensional close packing problem, it nevertheless provides a nearly optimal solution to the problem that is also particularly simple to implement. However, we have also neglected the effect of the off-diagonal term $g_{\phi Q}$, which would otherwise suggest a shear in the placement of our basis functions. Although it is possible to select an alternative parameterization of our basis which diagonalizes the resulting metric, we choose not to do so in order to preserve the physical significance and intuitive nature of our parameters. Instead, since it always acts to reduce the fractional energy loss incurred by an arbitrary burst, we ignore this off-diagonal term at the cost of a slightly less optimal tiling of our signal space.




Figure 3.2: The proposed cubic distribution of basis functions in the space of fractional energy loss described by the metric of Equation 3.20. Basis functions (black) are represented by the vertices in the cubic lattice, while the signal (white) corresponding to worst case mismatch occurs at the center of each cube. In order to ensure that the fractional energy loss encountered by this signal does not exceed μ_{\max} , the maximum mismatch allowed along any edge of the cube is given by Equation 3.22.

Given the choice of a cubic lattice sketched in Figure 3.2, we may then determine the maximum allowable distance between basis functions in order to ensure that the specified worst case energy loss of μ_{\max} is never exceeded. Here, the greatest energy loss occurs for a signal at the center of the cube. Ignoring the effect of the off-diagonal

metric component, this worst case signal is then equidistant from the eight vertices of the cube, representing the eight nearest basis functions. We thus seek the distance along each edge of the cube such that the distance from this worst case signal to any of the eight nearest basis functions is equal to the maximum allowable value of μ_{\max} . In this case, we find that the maximum allowable distance along each edge of the cube is given by

$$\delta s_{\max} = 2 \left(\frac{\mu_{\max}}{3} \right)^{1/2}. \quad (3.22)$$

To tile the targeted signal space, we then integrate separately over each metric component in order to obtain a measure of the cumulative mismatch distance over the signal space of interest:

$$s_{\tau} = \int_{\tau_{\min}}^{\tau_{\max}} g_{\tau\tau}^{1/2} d\tau = \frac{2\pi\phi}{Q} (\tau_{\max} - \tau_{\min}) \quad (3.23a)$$

$$s_{\phi} = \int_{\phi_{\min}}^{\phi_{\max}} g_{\phi\phi}^{1/2} d\phi = \frac{(2 + Q^2)^{1/2}}{2} \ln \left(\frac{\phi_{\max}}{\phi_{\min}} \right) \quad (3.23b)$$

$$s_Q = \int_{Q_{\min}}^{Q_{\max}} g_{QQ}^{1/2} dQ = \frac{1}{\sqrt{2}} \ln \left(\frac{Q_{\max}}{Q_{\min}} \right). \quad (3.23c)$$

Dividing this cumulative distance by the required distance between basis functions and rounding up then yields the necessary number of basis functions to cover the targeted signal space in each of the three parameters:

$$N_{\tau} = \left\lceil \frac{s_{\tau}}{\delta s_{\max}} \right\rceil \quad (3.24a)$$

$$N_{\phi} = \left\lceil \frac{s_{\phi}}{\delta s_{\max}} \right\rceil \quad (3.24b)$$

$$N_Q = \left\lceil \frac{s_Q}{\delta s_{\max}} \right\rceil. \quad (3.24c)$$

Finally, we tile the space of time, frequency, and Q by first selecting the discrete set of Q s given by

$$Q_n = Q_{\min} \exp \left[\sqrt{2} \left(n - \frac{1}{2} \right) \frac{s_Q}{N_Q} \right] \quad 1 \leq n \leq N_Q. \quad (3.25)$$

For each value of Q , we then identify the discrete set of frequencies given by

$$\phi_n = \phi_{\min} \exp \left[\frac{2}{(2 + Q^2)^{1/2}} \left(n - \frac{1}{2} \right) \frac{s_\phi}{N_\phi} \right] \quad 1 \leq n \leq N_\phi. \quad (3.26)$$

Finally, for each value of ϕ and Q , we select the discrete set of times given by

$$\tau_n = \tau_{\min} + \frac{Q}{2\pi\phi} \left(n - \frac{1}{2} \right) \frac{s_\tau}{N_\tau} \quad 1 \leq n \leq N_\tau. \quad (3.27)$$




Figure 3.3: An illustration of the optimal tiling of the space of time, frequency and Q . The tiling of the targeted signal space with the minimum number of basis functions required to ensure a specified worst case energy loss due to mismatch naturally leads to a multiresolution basis that is logarithmically distributed in Q , logarithmically distributed in frequency, and linearly distributed in time. For constant Q , the tiling of individual time-frequency planes is a generalization of the dyadic wavelet decomposition in which tiles are not restricted to frequencies that are related by powers of two.

We therefore find that the resulting basis functions are spaced logarithmically in Q , logarithmically in frequency, and linearly in time. That is, the optimal tiling of the targeted signal space naturally leads to a multiresolution basis! The resulting basis is sketched in Figure 3.3 and shows a strong resemblance to wavelet like tilings of the time-frequency plane. In particular, for planes of constant Q , the tiling of the time-frequency plane is similar to the more commonly encountered dyadic wavelet decomposition, except that tiles are not restricted to frequencies that are related by powers of two.

3.3 Detection of bursts

In this last section, we introduce the language necessary to describe the effects of detector noise on the measurement of gravitational-wave bursts. We first introduce the one-sided amplitude and power spectral densities of detector noise, which are the standard means for characterizing the sensitivity of gravitational-wave detectors. Next, we introduce the signal to noise ratio achieved by matched filtering on the assumption that the waveform to be detected is well known in advance. This ideal signal to noise ratio then forms our primary measure of burst detectability in a given detector. Finally, we consider the relationship between our detection based parameterization of gravitational-wave bursts and their actual astrophysical properties. Based on a few simplifying assumptions, we are then able to approximately relate the total energy emitted in the form of gravitational radiation with a distance to the source.

3.3.1 Detector noise

Throughout this work, we will find it useful to characterize stochastic processes, that is random time-series data such as detector noise $n(t)$, by their time-domain auto-correlation function

$$r_n(\tau) = \langle n(t)n(t - \tau) \rangle. \quad (3.28)$$

This is simply a measure of the statistical self-correlations within time-series data on a time scale τ [113]. In general, the auto-correlation of a signal may itself vary with time. Here, we have implicitly assumed that this is not the case, and that $n(t)$ is a stationary stochastic process such that $r_n(\tau)$ is independent of t . In what follows, we also assume that $n(t)$ has units of dimensionless strain, as is the case for strain data from gravitational-wave detectors. As a result, $r_n(\tau)$, which nominally has units of $n(t)^2$, has units of squared dimensionless strain. Finally, in order to gain a more intuitive feel for the auto-correlation, it is useful to note that $r_n(0)$ is simply the variance σ_n^2 of the process $n(t)$. Moreover, if every value of $n(t)$ is a statistically

independent random variable, then $r_n(\tau)$ is zero for non-zero τ . As a result, the auto-correlation of such noise is simply a delta function in τ with amplitude σ_n^2 .

An alternative but completely equivalent characterization of stationary stochastic processes is the two-sided power spectral density

$$G_n(f) = \int_{-\infty}^{+\infty} r_n(\tau) e^{-i2\pi f\tau} d\tau, \quad (3.29)$$

which is simply the Fourier transform of the time-domain auto-correlation function.

More commonly, however, we will describe the properties of stationary stochastic processes by their corresponding one-sided power spectral density,

$$S_n(f) = \begin{cases} 2G_n(f) & f \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (3.30)$$

The primary advantage of this representation is that it identifies the frequency dependence of detector noise. In particular, it can be shown that the variance σ_n^2 of the time-series $n(t)$ is given by the integral

$$\sigma_n^2 = \int_0^{\infty} S_n(f) df, \quad (3.31)$$

such that $S_n(f) df$ is the contribution to the overall variance of the resulting noise due to fluctuations in an infinitesimal frequency band df centered on frequency f [113].

Note that, in the special case where $r_n(\tau)$ is simply a delta function, then the resulting power spectral density is necessarily flat and has equal contributions to its variance from fluctuations at all frequencies. For this reason, noise that is comprised of a continuous stream of statistically independent random variables, is referred to as white noise.

Commonly, one also finds reference to the one-sided amplitude spectral density of

a stationary stochastic process. This is simply the square root, $S_n^{1/2}(f)$, of the corresponding one-sided power spectral density. As a result, instead of having units of squared dimensionless strain per Hz, as is the case for the power spectral density, the amplitude spectral density has units of dimensionless strain per $\text{Hz}^{-1/2}$. For gravitational-wave detectors, which are directly sensitive to the strain amplitude of gravitational-waves, this choice is generally more intuitive. In fact, the one-sided amplitude spectral density of detector noise is effectively the standard means for characterizing the sensitivity of gravitational-wave detectors. In Figure 2.3, Figure 2.7a, and Figure 2.7b, we have already described the sensitivity of gravitational-wave detectors in this way.

3.3.2 Signal to noise ratio

We now briefly consider the effect of detector noise on the measurement of gravitational-wave bursts.

In order to simplify our approach, we take advantage of the well known theory of matched filtering[87]. That is, we assume that the targeted astrophysical source is sufficiently well modeled that an accurate waveform is available in advance. In this case, we find that the maximum achievable signal to noise ratio[88] is given by

$$\rho_0 = \left[\int_0^\infty \frac{4|\tilde{h}(f)|^2}{S_n(f)} df \right]^{1/2}. \quad (3.32)$$

Although it is not in general the case that the waveform is accurately known in advance, Equation 3.32 nevertheless serves as a useful benchmark for comparison with other search methods. In particular, it represents the best case detectability of a gravitational-wave burst using what is effectively the optimal linear filter. As a result, we also take the matched filter signal to noise ratio as our primary measure of signal detectability.

If we also assume that the detector noise spectrum is approximately constant over the frequency band of a gravitational-wave burst, then by taking it outside of the integral we may further simplify our expression for the matched filter signal to noise ratio. In this case we find,

$$\rho_0 \simeq \left[\frac{2}{S_n} \int_{-\infty}^{+\infty} |\tilde{h}(f)|^2 df \right]^{1/2}, \quad (3.33)$$

where we have taken advantage of the symmetry properties of Fourier transforms to write the integral over all frequencies. We then recognize that this integral is simply our definition of characteristic strain energy $\|h\|^2$ from Equation 3.2. As a result, the matched filter signal to noise ratio is approximately given by

$$\rho_0 \simeq \sqrt{2}\|h\|/S_n^{1/2}. \quad (3.34)$$

We therefore find that the characteristic strain amplitude $\|h\|$ of a gravitational-wave burst can be directly compared to a detector's amplitude spectral density $S_n^{1/2}(f)$ in the bandwidth of the burst in order to quickly provide a rough estimate of a source's detectability.

3.3.3 Detectability

We now seek a relationship between the total gravitational-wave energy E emitted by an astrophysical source and the resulting characteristic amplitude $\|h\|$ observed by a detector.

We begin with the instantaneous energy flux carried by both polarizations of a gravitational-wave[22],

$$I(t) = \frac{c^3}{16\pi G} \left[|\dot{h}_+(t)|^2 + |\dot{h}_\times(t)|^2 \right]. \quad (3.35)$$

Next, we assume that this energy flux is radiated isotropically from a source at distance r . Alternatively, this is also equivalent to the mean emission from an ensemble of sources with random inclination relative to the line of sight. Integration over all time and a sphere of radius r then yields the total gravitational-wave energy emitted by the source,

$$E = \frac{c^3}{4G} r^2 \int_{-\infty}^{+\infty} \left[|\dot{h}_+(t)|^2 + |\dot{h}_\times(t)|^2 \right] dt. \quad (3.36)$$

Alternatively, Parseval's theorem and the differentiation property of Fourier transforms lead to the relation

$$\int_{-\infty}^{+\infty} \left[|\dot{h}_+(t)|^2 + |\dot{h}_\times(t)|^2 \right] dt = 4\pi^2 \int_{-\infty}^{+\infty} f^2 \left[|\tilde{h}_+(f)|^2 + |\tilde{h}_\times(f)|^2 \right] df, \quad (3.37)$$

which allows us to express the total gravitational-wave energy emitted by the source as the frequency domain integral

$$E = \frac{\pi^2 c^3}{G} r^2 \int_{-\infty}^{+\infty} f^2 \left[|\tilde{h}_+(f)|^2 + |\tilde{h}_\times(f)|^2 \right] df. \quad (3.38)$$

Of course, an interferometric gravitational-wave detector is not simultaneously sensitive to both gravitational-wave polarizations. In addition, the observed signal will also depend non-trivially on the relative orientation of the source and the detector. In what follows, we avoid these difficulties by assuming that the source emits only linearly polarization gravitational-waves and that our detector is optimally oriented to detect the maximum signal. As a result, our final result will reflect a best case measure of detectability.

Given these assumptions, the gravitational-wave energy emitted by the sources is then given by

$$E = \frac{\pi^2 c^3}{G} r^2 \int_{-\infty}^{+\infty} f^2 |\tilde{h}(f)|^2 df, \quad (3.39)$$

where $\tilde{h}(f)$ now represents the detected signal.

We then substitute the normalized frequency domain wavefunction of Equation 3.3b

and define the characteristic squared frequency

$$\langle f^2 \rangle = 2 \int_0^\infty f^2 |\tilde{\psi}(f)|^2 df. \quad (3.40)$$

As a result, we find the following expression for the total energy emitted by the source in the form of gravitational radiation.

$$E = \frac{\pi^2 c^3}{G} r^2 \|h\|^2 \langle f^2 \rangle \quad (3.41)$$

Not surprisingly, the total signal energy $\|h\|^2$ decreases as the square of the distance to the source. However, for a given total physical energy E , there is also a waveform dependent factor $\langle f^2 \rangle$ that favors the detection of low frequency signals.

We may also consider the matched filter signal to noise ratio of such a source in the presence of detector noise. In this case, substitution of the normalized frequency domain wavefunction of Equation 3.3b into the expression for matched filter signal to noise ratio of Equation 3.32 leads to

$$\rho_0^2 = \|h\|^2 \int_0^\infty \frac{4|\tilde{\psi}(f)|^2}{S_n(f)} df. \quad (3.42)$$

If we then define the characteristic detector noise over the band of the signal,

$$\langle S_n \rangle = \left[\int_0^\infty \frac{4|\tilde{\psi}(f)|^2}{S_n(f)} df \right]^{-1},$$

the matched filter signal to noise ratio becomes

$$\rho_0^2 = \frac{\|h\|^2}{\langle S_n(f) \rangle}. \quad (3.43)$$

By combining the above results, we find that the total gravitational-wave energy of an astrophysical source, its distance from the detector, and its matched filter signal

to noise ratio are all related by

$$E = \frac{\pi^2 c^3}{G} r^2 \rho_0^2 \langle S_n(f) \rangle \langle f^2 \rangle. \quad (3.44)$$

Finally, we assume that the total energy E emitted in the form of gravitational radiation is some fraction ϵ of the available rest mass energy Mc^2 of the source. As a result, we find that the approximate range to detectable sources is given by

$$r = \left[\frac{\epsilon GM}{\pi^2 c \rho_0^2 \langle S_n(f) \rangle \langle f^2 \rangle} \right]^{1/2}. \quad (3.45)$$

For the typical values relevant to ground-based interferometric detectors, this gives approximate detectable ranges on the order of

$$r \sim 14 \left(\frac{\epsilon}{0.01} \right)^{\frac{1}{2}} \left(\frac{M}{M_\odot} \right)^{\frac{1}{2}} \left(\frac{\rho_0}{5} \right)^{-1} \left(\frac{\langle S_n(f) \rangle}{10^{-44} \text{ Hz}^{-1}} \right)^{-\frac{1}{2}} \left(\frac{\langle f^2 \rangle}{10^4 \text{ Hz}^2} \right)^{-\frac{1}{2}} \text{ Mpc}, \quad (3.46)$$

where any variations in source frequency with mass and any corresponding variations in detector noise must still be taken into account.

Chapter 4

Linear Prediction

Prior to performing the multiresolution time-frequency analysis motivated in the previous chapter, it is first useful to whiten the input data stream such that individual samples of the resulting discrete time sequence are statistically independent random variables drawn from a common distribution. As we will see in section 5.5, such an approach greatly simplifies our subsequent statistical analysis by permitting the assumption of stationary white noise data. In addition, although we do not make use of it in this work, adequate whitening of gravitational-wave data is also beneficial prior to performing a cross-correlation based analysis[93]. Otherwise, coherent signal content, such as line sources and other parasitic resonances, will dominate the resulting correlation and obscure any potential transient correlated signals. For both types of searches, we therefore seek a method to identify and remove sample to sample correlations from the input data stream. This chapter presents linear predictor error filtering as a technique for accomplishing this task.

Linear prediction is a well known technique from stochastic signal processing that attempts to predict future values of time series data using an appropriate linear combination of previous measurements[114]. To do so, it assumes that the signal under investigation is due to white noise excitation of a filter whose output depends only

on its instantaneous input and previous values of its output. In the signal processing literature, such signals are commonly referred to as autoregressive processes. In order to recover the underlying white noise excitation, linear prediction seeks to identify and remove the effects of such autoregressive filters, which in our case includes coherent signal content due to both the detector and the environment. In doing so, we also take advantage of the fact that robust and computationally efficient solutions to the linear prediction problem are already well known.

In what follows, we develop the theory of linear prediction with the goal of removing the predictable content of stationary stochastic processes. In the process, we determine the necessary conditions to ensure that the resulting data stream is sample to sample uncorrelated on the time scale of any subsequent analysis. We then consider the possible side effects of whitening by linear prediction. In particular, we note the introduction of an arbitrary phase delay that could adversely affect coincident searches for bursts in data from networks of gravitational-wave detectors. To alleviate this problem, we present a novel zero-phase modification of linear predictive whitening that avoids this difficulty. Finally, in order to gain a more intuitive appreciation for linear prediction, we conclude with an example of the effect of zero-phase linear predictive whitening on a simulated gravitational wave burst.

4.1 Definition

Linear predictor error filtering assumes that the signal to be whitened is a stationary stochastic process that is well modeled by an autoregressive model of order M . That is, that the n^{th} sample of a discrete time sequence, $x[n]$, is well modeled by a linear combination of the previous M samples. Given this assumption, we define the predicted sequence $\tilde{x}[n]$ in terms of M undetermined coefficients $c[m]$ by the expression

$$\tilde{x}[n] = \sum_{m=1}^M c[m]x[n-m]. \quad (4.1)$$

Next, we define the prediction error sequence as the difference between the true sequence and the predicted sequence,

$$e[n] = x[n] - \tilde{x}[n]. \quad (4.2)$$

If the coefficients $c[m]$ are chosen correctly, the resulting prediction error sequence is generally composed of sample to sample uncorrelated white noise, but also contains any unpredictable non-stationarities that were present in the original data sequence. Thus, the prediction error sequence is the whitened data sequence which will be used in subsequent analysis. We therefore define the linear predictor error filter as the M^{th} order finite impulse response filter with coefficients $b[m]$, which when applied to a data sequence, $x[n]$, returns the corresponding prediction error sequence via

$$e[n] = \sum_{m=0}^M b[m]x[n - m]. \quad (4.3)$$

The coefficients of this filter follow directly from the prediction coefficients of the autoregressive model and are given by

$$b[m] = \begin{cases} 1 & m = 0 \\ -c[m] & 1 \leq m \leq M \\ 0 & \text{otherwise.} \end{cases} \quad (4.4)$$

4.2 Training

To determine the coefficients $c[m]$, and therefore $b[m]$, we first determine the mean squared prediction error, σ_e^2 , of the filter when applied to a representative data sequence of length N ,

$$\sigma_e^2 = \frac{1}{N} \sum_{n=1}^N e^2[n]. \quad (4.5)$$

The coefficients $c[m]$ are then chosen to minimize the mean squared prediction error in the least squares sense by requiring

$$\frac{\partial \sigma_e^2}{\partial c[k]} = 0 \quad 1 \leq k \leq M. \quad (4.6)$$

Assuming that $x[n]$ is a stationary stochastic process, this procedure, which is referred to as training, results in the well known Yule-Walker equations[114],

$$\sum_{m=1}^M r[m-k]c[m] = r[k] \quad 1 \leq k \leq M, \quad (4.7)$$

where $r[k]$ is the auto-correlation of the signal $x[n]$ evaluated at lag k ,

$$r[k] = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=-N/2}^{+N/2} x[n]x[n-k]. \quad (4.8)$$

Since the auto-correlation of real sequences is even-symmetric about zero lag, the Yule-Walker equations take on the particularly simple form of a symmetric Toeplitz matrix equation:

$$\begin{bmatrix} r[0] & r[1] & r[2] & \cdots & r[M-1] \\ r[1] & r[0] & r[1] & \cdots & r[M-2] \\ r[2] & r[1] & r[0] & \cdots & r[M-3] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ r[M-1] & r[M-2] & r[M-3] & \cdots & r[0] \end{bmatrix} \begin{bmatrix} c[1] \\ c[2] \\ c[3] \\ \vdots \\ c[M] \end{bmatrix} = \begin{bmatrix} r[1] \\ r[2] \\ r[3] \\ \vdots \\ r[M] \end{bmatrix}. \quad (4.9)$$

By taking advantage of the symmetry of this matrix equation, it is possible to solve for the M coefficients, $c[m]$, in $\mathcal{O}(M^2)$ operations, instead of the $\mathcal{O}(M^3)$ operations required for an arbitrary matrix. This leads to the recursion algorithm of Levinson and Durbin, which is widely available[115] and will not be reproduced in detail here.

In practice, due to the finite length N of the available training data, we actually estimate the auto-correlation by the expression

$$r[k] = \frac{1}{N} \sum_{n=|k|+1}^N x[n]x[n - |k|]. \quad (4.10)$$

Note that this estimate exhibits a triangular bias towards zero lag due to an effective weighting by $(N - |k|)/N$. Although this bias may be compensated for, it can be shown that this choice of estimate actually ensures that the resulting Yule-Walker equations are always non-singular[114]. Hence, except in the pathological case of an all zero input signal, a solution to the Yule-Walker equations is guaranteed to exist.

We also note that efficient estimation of the first M auto-correlation coefficients of the training data is possible via the fast Fourier transform. In this case, the necessary auto-correlation coefficients can be computed in $\mathcal{O}(N \log N)$ operations, instead of the $\mathcal{O}(NM)$ operations required by the more straightforward time-domain approach. In addition, we note that the length N of the training data must exceed the order M of the filter in order to allow estimation of the auto-correlation sequence up to a lag k equal to M . Finally, we note that for sufficiently small M , the computational cost of training is dominated by the $\mathcal{O}(N \log N)$ operations associated with estimating the auto-correlation sequence. However, for sufficiently large M , the $\mathcal{O}(M^2)$ operations required to solve the Yule-Walker equation may quickly become the dominant computational cost.

4.3 Properties

In addition to the guaranteed existence of a solution, the Yule-Walker equations have a number of useful properties.

Since the resulting linear predictor error filter has a finite impulse response, it consists

only of zeros (with poles at infinity) in the Z -plane and is by definition stable[116]. In addition, the resulting filter is also minimum phase[114]. That is, all of its zeros are inside the unit circle such that the inverse filter is also stable.

This suggests that whitening by linear predictor error filtering can be described by the two step process depicted in Figure 4.1. In the first step, training, the input signal is modeled as white noise shaped by an all-pole filter. In the second step, application, the inverse all-zero filter is applied, which restores the original white noise sequence.




Figure 4.1: Autoregressive model for whitening by linear predictor error filtering.

In general, however, it may not be possible to accurately model an arbitrary signal with an autoregressive model of order M . Instead, such cases require an autoregressive moving average model, where the signal is modeled as white noise shaped by an arbitrary infinite impulse response filter. If a stable inverse filter exists for such a model, in general it also has an infinite impulse response, and is therefore not exactly representable by the finite impulse response of linear predictor error filters.

Fortunately, we can approach the performance of the general case, while avoiding the difficulties of autoregressive moving average models, by allowing the linear predictor filter order, M , to approach infinity at the expense of increased computation time. In practice, we must truncate the filter at some finite, but large, value of M . This choice of M effects how well the filter can perform on arbitrary data, and has the following interpretation in terms of time-domain windowing.

Consider a stationary stochastic process produced by filtering white noise through an arbitrary infinite impulse response filter. Assume also that this filter has a stable

inverse, whose impulse response we desire in order to perfectly whiten the data. The solution of the Yule-Walker equations provides the best fit finite impulse response to this desired impulse response in the sense that it minimizes the mean squared prediction error. Qualitatively, this process produces a truncated version of the desired impulse response, which can also be interpreted as applying a rectangular window in the time-domain. In the frequency domain, this corresponds to convolving the desired transfer function with a sinc function, whose characteristic bandwidth, ϕ_p , is approximately

$$\phi_p \sim \frac{f_s}{M}. \quad (4.11)$$

Here f_s is the sample frequency of the input data. This has the effect of blurring the desired transfer function at a frequency resolution equal to the characteristic bandwidth, ϕ_p . As a result, the filter is not able to compensate for narrowband spectral features whose bandwidths are less than ϕ_p .

Fortunately, by increasing the filter order M , we can whiten to any arbitrary frequency resolution. In practice one selects the smallest order M such that the data are sufficiently white at the frequency resolution of any subsequent analysis. In the case of a time-frequency analysis, where the data is projected onto basis functions of maximum duration T , or in the case of a cross-correlation based analysis, where the cross-correlation is performed over segments of duration T , we simply choose M such that

$$M \gtrsim f_s T. \quad (4.12)$$

The training length, N , of a linear predictor error filter obeys a similar relationship. In particular, it is impossible for the filter to learn about narrowband spectral features whose bandwidth are less than a characteristic bandwidth, ϕ_t given by

$$\phi_t \sim \frac{f_s}{N}. \quad (4.13)$$

However, since N is typically much larger than M , the applicable frequency resolution of the filter is determined by ϕ_p rather than ϕ_t . Instead, the training time is usually

chosen based on computation cost and the minimum duration over which the input data stream is relatively stationary.

Finally, prior to training and applying a linear predictor error filter, it is also necessary to detrend data by high pass filtering it at a frequency equal to or greater than ϕ_p . Otherwise, signal content below this frequency introduces a bias during training since the resulting predictor cannot compensate for signal content whose period exceeds the predictor length. In this work, we apply the well known Butterworth high pass filter[116], whose frequency domain attenuation has the form

$$A(f) = \left[\frac{(f/f_c)^{2L}}{1 + (f/f_c)^{2L}} \right]^{1/2}, \quad (4.14)$$

where f_c is the characteristic cutoff frequency of the filter and L is the order of the filter.

In practice, due to the large low frequency signal content associated with seismic noise, data from Earth based gravitational-wave detectors are typically high pass filtered at a frequency around 50 Hz. For the typical linear predictor filter orders encountered in gravitational-wave data analysis, this turns out to be approximately equal to, or greater than, ϕ_p , such that the data is also sufficiently detrended for linear predictive whitening.

4.4 Application

It is evident from the preceding discussion that the typical linear predictor order may be much larger than that of other filters which are commonly used in the conditioning of gravitational-wave data. For example, consider a multiresolution time-frequency analysis with a maximum tile duration of 0.25 seconds applied to data sampled at 16384 Hz. Such an analysis would require a linear predictor error filter of order 4096 or greater. This would be prohibitively expensive to implement using the standard

time-domain convolution of Equation 4.3.

Fortunately, exact techniques exist to apply FIR filters in the frequency domain. These techniques, known as overlap-add or overlap-save, rely on piecing together the results of many cyclic convolutions performed in the frequency domain via fast Fourier transforms. This allows the filtering of arbitrarily long data streams, broken into blocks of length N , in $\mathcal{O}(N \log N)$ operations per block, instead of the $\mathcal{O}(NM)$ operations per block required by standard time-domain techniques. These techniques are well known and documented[116] and will not be reproduced here in detail.

4.5 Frequency response

In order to characterize the effects of linear predictive whitening, we now consider its performance on the simulated interferometric detector noise described in section 6.3. The ability of linear predictor error filtering to whiten this data is demonstrated in Figure 4.2. In this example, 64 seconds of data was first high pass filtered at 64 Hz using a 6th order Butterworth filter. Next, a 0.25 second linear predictor error filter was trained on the first 32 seconds of data and then applied to the subsequent 32 seconds of data. Figure 4.2 shows the amplitude spectra of the data before and after application of the linear predictor error filter. Both amplitude spectra were measured at a frequency resolution of 8 Hz, in agreement with the condition in Equation 4.12. The effectiveness of the linear predictor error filter to whiten the data at this frequency resolution is readily apparent from the resulting frequency spectrum. In particular, above the cutoff frequency of the high pass filter, the resulting spectrum is independent of frequency, as expected for stationary white noise.

In order to further characterize the linear predictor error filter, it is useful to define

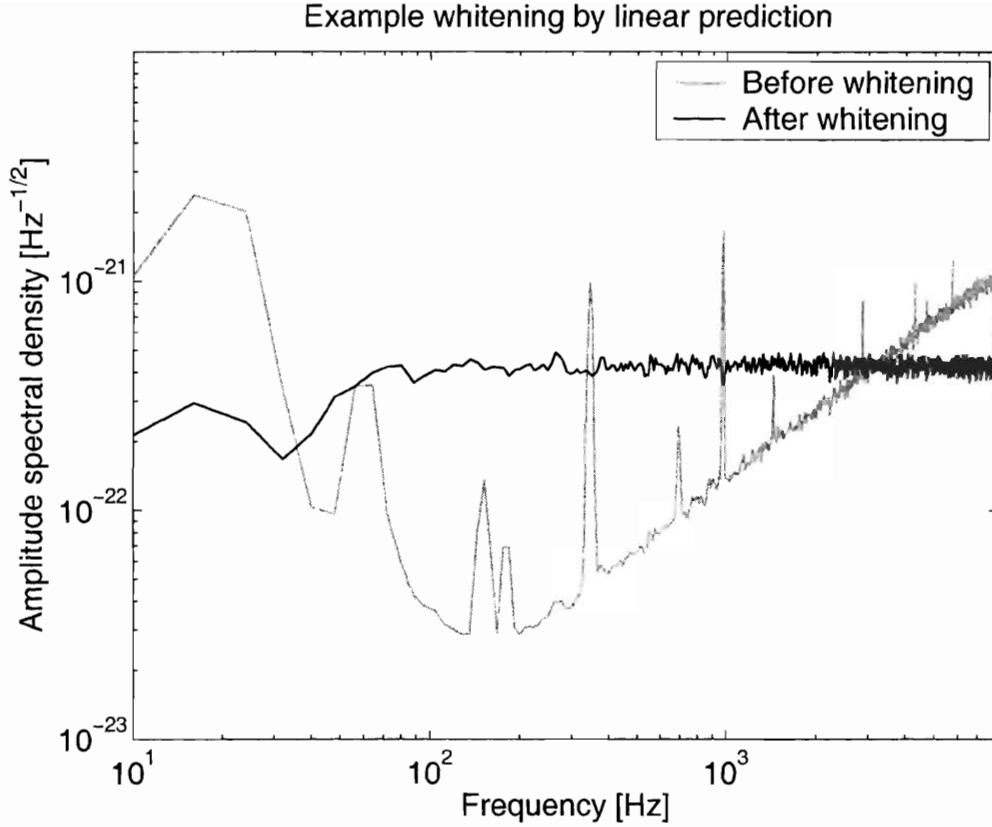


Figure 4.2: Amplitude spectra of simulated gravitational-wave data before and after whitening by a linear predictor error filter with a 4 Hz resolution. The amplitude spectra are shown at a frequency resolution of 8 Hz, consistent with the requirement of Equation 4.12. The simulated data is described in more detail in section 6.3. In both cases, the simulated data was first filtered by a 6th order Butterworth high pass filter with a cutoff frequency of 64 Hz. Aside from the lack of low frequency energy below the cutoff frequency of the high pass filter, the effectiveness of the linear predictive whitening is readily apparent.

its frequency response by its discrete time Fourier transform[116],

$$B(f) = \sum_{m=0}^{\infty} b[m]e^{-i2\pi fm/f_s}. \quad (4.15)$$

The frequency response of the example linear predictor error filter is then shown in Figure 4.3. As expected, the magnitude response of the filter is simply the inverse of the unwhitened amplitude spectrum seen in Figure 4.2, except for an arbitrary gain factor. However, the phase response of the filter is more problematic. Large phase changes are evident in the spectrum in conjunction with narrowband spectral features




Figure 4.3: Frequency response of linear predictor error filter trained on simulated gravitational-wave data with a 4 Hz resolution. The resulting magnitude response (top) is the inverse of the unwhitened spectrum observed in Figure 4.2. The associated phase response (middle) exhibits large fluctuations in proximity to narrow band spectral features in the unwhitened spectrum, which results in a group delay (bottom) for narrowband wavepackets that clearly exceed the light travel time between gravitational-wave detectors. Note that the group delay extends well off the bottom of the plot at the frequencies of narrow band spectral features.

in the original unwhitened spectrum. Unfortunately, as indicated below, such features will likely lead to unwanted dispersion.

A useful measure of the dispersion introduced by a filter is its group delay[116], which is also shown in Figure 4.3 for the example case under consideration. The concept of group delay is closely related to the more well known concept of group velocity and represents the time delay incurred by a narrowband wavepacket propagating through a communication channel. Since a uniform delay of τ_d seconds in the time-domain is equivalent to a linear phase shift in the frequency domain with a phase factor of $\exp(-i2\pi f\tau_d)$, the group delay, $\tau_d(f)$ at a particular frequency, f , is simply proportional to the local slope of the phase at that frequency,

$$\tau_d(f) = -\frac{1}{2\pi f_s} \frac{\partial}{\partial f} \arg[B(f)]. \quad (4.16)$$

Here the phase is determined by the complex argument of the filter's frequency domain response $B(f)$, and f_s is the sample frequency of the discrete time data.

As evidenced in Figure 4.3, the group delay near narrowband spectral features can be quite significant, leading to dispersion on the order of or greater than the speed of light travel time between gravitational-wave observatories on the Earth. This effect becomes more problematic at higher filter orders, where the linear predictor filter attempts to compensate for ever narrower spectral features, leading to ever more rapid fluctuations in the phase response. In the next section we consider the possible impact of this effect on multidetector analysis and propose a method for avoiding these difficulties.

4.6 Zero-phase filtering

The dispersion introduced by standard linear predictor error filtering poses a potential problem when conditioning data from multiple detectors in order to perform a coinci-

dent search for bursts. Differences in the predictable signal content of the individual detectors will invariably result in filters with varying amounts of dispersion. This difference in dispersion then results in a relative timing error between the whitened time series. If this timing error approaches or exceeds the speed of light travel time between gravitational-wave observatories, the time scale used to test for coincidence must be increased in order to avoid the false dismissal of true gravitational-wave events. This either results in an increase in the number of accidental coincidences or requires an increase in detection threshold and corresponding decrease in detection efficiency in order to maintain a fixed probability of false detection. In addition, the effects of dispersion make it impossible to perform a coherent search for bursts as described in section 5.8.

Fortunately, the problem of filter dispersion can be eliminated by the technique of zero-phase filtering. In this approach, a filter is first applied causally, then anti-causally in order to cancel the effects of any phase dispersion introduced on the forward pass. That is, we first convolve the input time series with the impulse response of the desired filter,

$$y[n] = \sum_{m=-\infty}^{+\infty} b[m]x[n - m], \quad (4.17)$$

then convolve the resulting time series with the time reversed impulse response of the same filter,

$$z[n] = \sum_{k=-\infty}^{+\infty} b[k]y[n + k]. \quad (4.18)$$

Alternatively, this is equivalent to applying a single filter,

$$z[n] = \sum_{l=-\infty}^{+\infty} r_b[l]x[n - l], \quad (4.19)$$

where the coefficients $r_b[l]$ are simply the unnormalized auto-correlation of the original filter coefficients,

$$r_b[l] = \sum_{m=-\infty}^{+\infty} b[m]b[m - l]. \quad (4.20)$$

The fact that the resulting filter has zero phase is directly related to the even symmetric property of the auto-correlation of a real sequence[116].

In the frequency domain, the proposed zero-phase filter has the much simpler form

$$Z(f) = B(f)B^*(f)X(f), \quad (4.21)$$

where $B(f)$ is the frequency domain response of the original filter as defined by Equation 4.15. Here, the zero-phase property of the resulting filter is immediately evident. However, we also note a side effect of zero-phase filtering: the magnitude response of the original filter is applied twice. While this is not a concern when high pass filtering the input time series, it is undesirable for whitening since it will result in the inverse, rather than white, spectrum.

Instead, we desire to apply the filter

$$E(f) = |B(f)|X(f). \quad (4.22)$$

Unfortunately, this filter cannot in general be implemented by a finite length impulse response of order M . However, we may approach the ideal frequency response by effectively increasing the filter order. This is accomplished by zero padding the impulse response of the original filter to a length L greater than M , computing the discrete Fourier transform of the zero padded time series, taking the absolute value of the resulting frequency series, and computing the inverse discrete Fourier transform to return to the time-domain. The resulting even symmetric filter coefficients then achieve the desired magnitude response in the limit as L goes to infinity. In Figure 4.4, we present an example of the accuracy with which the desired magnitude response is achieved as a function of the multiplicative increase in filter order. Again, we train the linear predictor error filter on the simulated detector noise from section 6.3. However, for clarity we have trained the filter for much coarser frequency resolution of 256 Hz.

In Figure 4.4, we first note the poor agreement between the magnitude response of the zero-phase filter and the desired filter with the same filter order. However, by increasing the order of the zero-phase filter by a factor 4, the magnitude response of the resulting zero-phase filter shows good agreement with the desired response. In practice, multiplicative factors on the order of 64 have been applied in this work, ensuring excellent agreement with the desired magnitude response.

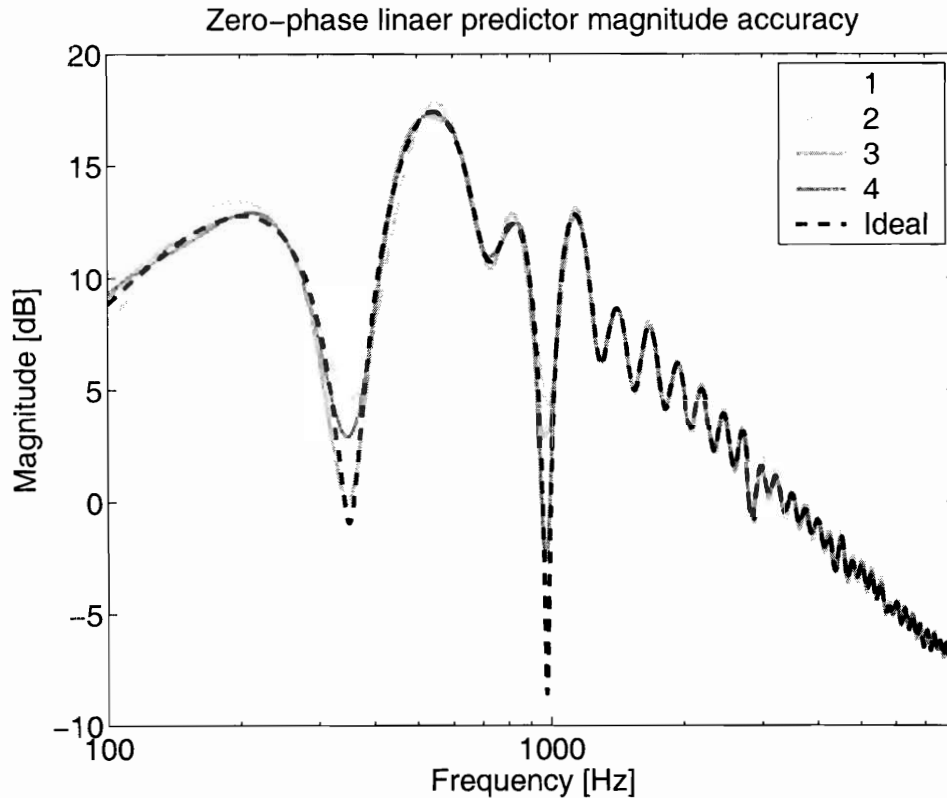


Figure 4.4: The accuracy with which the magnitude response of zero-phase linear predictor error filters achieve the desired magnitude response of the standard linear predictor as a function of the multiplicative increase in filter order. The desired magnitude response is similar to that of Figure 4.3, except that, for reasons of clarity, here the whitening filter has been designed for a much coarser frequency resolution of 256 Hz. Although noticeable disagreement is observed if the filter order is unchanged, an increase in filter order by a factor of 4 shows good agreement with the desired response.

4.7 Example

Finally, in order to gain a better understanding of the effects of zero-phase linear prediction, we conclude this chapter by considering a simple example of zero phase whitening applied to a simulated burst of gravitational radiation. In particular, we consider a sinusoidal Gaussian burst of the form

$$h(t; \tau, \phi, Q) = h_0 \exp \left[-4\pi^2 \phi^2 (t - \tau)^2 / Q^2 \right] \sin [2\pi \phi (t - \tau)], \quad (4.23)$$

with a central frequency of 256 Hz and a Q of 8.

This signal is then added to 32 seconds of simulated LIGO detector noise as described in section 6.3. Moreover, for clarity, this signal is injected with the reasonably large matched filter signal to noise ratio of 30. The resulting data are then zero-phase high pass filtered by a 6th order Butterworth filter with a cutoff frequency of 64 Hz and whitened by zero-phase linear prediction with a frequency resolution of 1 Hz. The resulting high pass filtered and whitened time series is shown in Figure 4.5 along with the high passed filtered time series prior to whitening and the true injected gravitational-wave burst. The benefit of zero-phase linear predictive whitening to remove predictable signal content is readily apparent. In addition, we also note the absence of any apparent time delay between the whitened sinusoidal Gaussian and the true signal. After developing the underlying theory of the Q transform, we will return to this example in section 5.9, where we will apply the Q transform to resulting whitened data stream in order to recover the time-frequency properties on the injected waveform.




Figure 4.5: Example zero-phase whitening of a sinusoidal Gaussian burst in the presence of simulated LIGO detector noise. In this case, the signal is a 256 Hz Gaussian wave packet with a Q of 8 injected into simulated LIGO detector noise with an optimal matched filter signal to noise ratio of 30. The time-domain waveform of the combined signal and noise are shown after zero-phase high pass filtering by an effective 12th order Butterworth filter with a cutoff frequency of 64 Hz (top), and whitening by a zero-phase linear prediction with a 1 Hz frequency resolution (middle). The original sinusoidal Gaussian burst is also shown for reference (bottom). Note the order of magnitude difference in the scale of the top plot relative to the other two. This is due to the large root mean square noise associated with coherent signal content in the input data stream.

Chapter 5

The Q Transform

In chapter 3 we motivated the use of a multiresolution basis of well localized functions to search for gravitational-wave bursts. In this chapter, we introduce the Q transform as a tool for performing such an analysis.

The Q transform is a modification of the standard short time Fourier transform[85] in which the analysis window duration varies inversely with frequency such that the time-frequency plane is covered by tiles of constant Q[117, 118]. In this sense, it is similar in construction to the continuous wavelet transform[108]. However, since reconstruction of the data sequence is not a concern, we permit violation of the zero mean admissibility requirement for wavelets. In addition, in contrast to the more commonly encountered dyadic wavelet decomposition, the discrete Q transform is not restricted to frequencies that are related by powers of two.

In what follows, to gain an understanding of the Q transform, we first introduce the more intuitive continuous form. Following that, we develop the corresponding version for the case of discrete data. In particular, we develop a frequency-domain form of the discrete Q transform which takes advantage of the computational efficiency of the fast Fourier transform. We also present a choice of analysis window that achieves

near minimum time-frequency uncertainty, and determine the necessary window normalization to ensure a complete and accurate accounting of signal energy.

We then focus on the interpretation of the resulting Q transform coefficients in order to identify statistically significant events. First, the statistical distribution of Q transform magnitude is determined assuming the data has been sufficiently whitened using the techniques presented in chapter 4. Next, an outlier rejection technique is presented which reduces the susceptibility of the estimated statistical parameters to non-stationarities in the input data. We then define the significance of Q transform coefficients as the probability of observing a coefficient of greater magnitude, given stationary white noise data. This permits the identification of statistically significant events using a detection threshold that is selected for a target event rate, again assuming stationary white noise. Finally, we present a simple exclusion algorithm that enables significant events from multiple overlapping Q transforms to be combined in a way that best determines the parameters of well localized bursts.

Next, we combine the methods proposed in this and the previous chapters in order to form a comprehensive end-to-end analysis pipeline for the detection of gravitational-wave bursts. In addition, we present a method of coherently combining Q transform measurements from multiple detectors in a way that maximizes the combined signal to noise ratio of gravitational-wave bursts while taking into account potential differences in sensitivity between the detectors.

Finally, in order to gain a more intuitive understanding of the Q transform, we conclude with a simple example of the Q transform applied to a simulated gravitational-wave burst in the presence of simulated detector noise that has been whitened using zero-phase linear prediction.

5.1 The continuous Q transform

In this section, we develop two forms of the continuous Q transform. The first is a straightforward time-domain projection of the input time series onto a basis of windowed sinusoids. The second is an alternative frequency-domain form, which will prove useful in developing the fast form of the discrete Q transform in section 5.2.2.

5.1.1 Direct form

In its most basic form, the continuous Q transform is simply the projection of the continuous time series $x(t)$ onto windowed complex exponentials of center frequency ϕ and quality factor Q . Mathematically, this is given by the expression

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} x(t) w(t - \tau, \phi, Q) e^{-i2\pi\phi t} dt, \quad (5.1)$$

where $w(t - \tau, \phi, Q)$ is a time-domain window centered on time τ with a duration that is proportional to Q and inversely proportional to the frequency ϕ under consideration.

Due to its simplicity, we will refer to Equation 5.1 as the direct form of the continuous Q transform. The resulting complex valued transform coefficients are a measure of the average signal amplitude and phase within a time-frequency region centered on time τ and frequency ϕ , whose shape and area are determined by the requested quality factor Q and the particular choice of analysis window. However, since our initial development of the Q transform does not depend upon the details of the analysis window, we defer discussion of the specific choice of window until section 5.3.

5.1.2 Alternative form

To develop an alternative frequency-domain form of continuous Q transform, we begin by defining the modulated input time series

$$v(t, \phi) = x(t) e^{-i2\pi\phi t}. \quad (5.2)$$

The transform of Equation 5.1 then becomes

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} v(t, \phi) w(t - \tau, \phi, Q) dt, \quad (5.3)$$

which we recognize as cross-correlation in time at a constant transform frequency ϕ and quality factor Q . This suggests an alternative Fourier space representation defined by

$$\tilde{v}(f, \phi) = \int_{-\infty}^{+\infty} v(t, \phi) e^{-i2\pi ft} dt \quad (5.4a)$$

$$\tilde{w}(f, \phi, Q) = \int_{-\infty}^{+\infty} w(t, \phi, Q) e^{-i2\pi ft} dt \quad (5.4b)$$

$$\tilde{X}(f, \phi, Q) = \int_{-\infty}^{+\infty} X(\tau, \phi, Q) e^{-i2\pi f\tau} d\tau, \quad (5.4c)$$

in which the Q transform is simply the product

$$\tilde{X}(f, \phi, Q) = \tilde{v}(f, \phi) \tilde{w}^*(f, \phi, Q), \quad (5.5)$$

with the superscript asterisk denoting complex conjugation. Here, in order to avoid confusion, we reiterate the fact that ϕ is the target Q transform frequency, while f is the frequency variable associated with the Fourier decomposition.

Based on Equation 5.2 and the frequency shift property of Fourier transforms, the Fourier transform in Equation 5.4a is simply

$$\tilde{v}(f, \phi) = \tilde{x}(f + \phi), \quad (5.6)$$

where

$$\tilde{x}(f) = \int_{-\infty}^{+\infty} x(t) e^{-i2\pi ft} dt, \quad (5.7)$$

is the standard Fourier transform of the original time series.

Transforming back to the time-domain then yields the desired alternative frequency-domain representation of the original time-frequency transform:

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} \tilde{x}(f + \phi) \tilde{w}^*(f, \phi, Q) e^{+i2\pi f\tau} df. \quad (5.8)$$

Thus, the Q transform at a specific frequency and quality factor is obtained by a simple Fourier transform of the original time series, a shift in frequency, multiplication by the appropriate frequency-domain window function, and an inverse Fourier transform. The benefit of Equation 5.8 is that the Fourier transform of the original time series need only be computed once. We then perform the inverse Fourier transform only for the logarithmically spaced frequencies and quality factors that we are interested in. From this point of view, the Q transform operates as a heterodyne detector for a band centered on the frequency ϕ with a bandwidth determined by the quality factor Q . That is, the input data stream is mixed down to near zero frequency by multiplication with a complex valued reference oscillator at frequency ϕ , then low pass filtered with cutoff frequency of approximately $\phi/2Q$. The resulting complex valued time series is therefore a measure of the signal amplitude and phase in a frequency band of width ϕ/Q centered on the frequency ϕ .

5.2 The discrete Q transform

Although the continuous Q transform can be evaluated for arbitrary time, frequency and Q , in practice this is not necessary. As demonstrated in chapter 3, basis functions that are closely spaced in time, frequency, and Q will produce highly correlated

projections. As a result, we need only evaluate the Q transform at a discrete set of parameters in order to accurately detect statistically significant events. This fact, coupled with the practical requirement of a finite length data sequence, motivates our development of the discrete Q transform.

Here we develop two forms of the discrete Q transform. The first is a discretization of the direct continuous Q transform and implements a straightforward projection of the input time series onto a basis of windowed complex exponentials. The second is a discretization of the frequency-domain form of the continuous Q transform that takes advantage of the computational efficiency of the fast Fourier transform.

Our development of the discrete case primarily proceeds in parallel to the continuous case. However, care must be taken to avoid the difficulties which are specific to discrete data. In particular, it is convenient to assume that all data sequences are infinite periodic sequences of period N . We may then apply the well developed theory of discrete Fourier transforms to the discrete Q transform. However, this assumption also introduces a number of complications, which will be pointed out in the course of the discussion.

5.2.1 Direct form

We first develop the direct form of the discrete Q transform, which is analogous to the transform of Equation 5.1. This is simply the projection of the discrete time series under test, $x[n]$, onto windowed complex exponentials of frequency index l and quality factor Q ,

$$X[m, l, Q] = \sum_{n=0}^{N-1} x[n] w[n - m, l, Q] e^{-i2\pi nl/N}. \quad (5.9)$$

Here $w[n, l, Q]$ is taken to be a periodic train of time-domain windows centered on time indices that are integer multiples of N with a uniform duration that is proportional

to Q and inversely proportional to the frequency index l under consideration.

The resulting complex valued transform coefficients are a measure of the average signal amplitude and phase within a time-frequency region centered on time index m and frequency index l , whose shape and area are determined by the requested Q and the particular choice of window. Again, since the initial development of the Q transform does not depend upon the details of the analysis window, we defer discussion of the specific choice of window until section 5.3.

Note that the assumption of periodic data sequences has also been applied to the window sequence. As a result, depending upon the particular choice of window, some number of discrete Q transform coefficients for time indices near the beginning and end of the data segment will be adversely effected by time-domain aliasing. In practice, when analyzing long data sequences using multiple Q transforms, such edge effects are avoided by evaluating the Q transform for overlapping blocks of data and discarding values near the edges of each block.

5.2.2 Fast form

We now develop an alternative frequency-domain form of the discrete Q transform that is analogous to the continuous transform of Equation 5.8. We begin by defining the modulated input time series

$$v[n, l] = x[n] e^{-i2\pi nl/N}. \quad (5.10)$$

The transform of Equation 5.9 then becomes

$$X[m, l, Q] = \sum_{n=0}^{N-1} v[n, l] w[n - m, l, Q]. \quad (5.11)$$

Assuming $v[n, l]$ and $w[n, l, Q]$ are periodic with period N , we recognize this as cyclic cross-correlation in time at a constant frequency index l and quality factor Q . This then suggests an alternative discrete Fourier space representation defined by

$$\tilde{v}[k, l] = \sum_{n=0}^{N-1} v[n, l] e^{-i2\pi nk/N} \quad (5.12a)$$

$$\tilde{w}[k, l, Q] = \sum_{n=0}^{N-1} w[n, l, Q] e^{-i2\pi nk/N} \quad (5.12b)$$

$$\tilde{X}[k, l, Q] = \sum_{m=0}^{N-1} X[m, l, Q] e^{-i2\pi mk/N}, \quad (5.12c)$$

in which the discrete Q transform is simply the product

$$\tilde{X}[k, l, Q] = \tilde{v}[k, l] \tilde{w}^*[k, l, Q]. \quad (5.13)$$

Note that we have chosen the convention where the factor of $1/N$, which necessarily appears in discrete Fourier transform pairs, is associated with the inverse transform rather than the forward transform.

Based on Equation 5.10 and the periodic frequency shift property of discrete Fourier transforms, the discrete Fourier transform in Equation 5.12a is simply

$$\tilde{v}[k, l] = \tilde{x}[k + l], \quad (5.14)$$

where

$$\tilde{x}[k] = \sum_{n=0}^{N-1} x[n] e^{-i2\pi nk/N} \quad (5.15)$$

is the standard discrete Fourier transform of the original time series.

Transforming back to the time-domain then yields the desired frequency-domain representation of the original time-frequency transform:

$$X[m, l, Q] = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{x}[k + l] \tilde{w}^*[k, l, Q] e^{+i2\pi mk/N} \quad (5.16)$$

As was the case with the alternative form of the continuous Q transform, this representation has the benefit that the Fourier transform of the original data need only be computed once. Moreover, for the discrete case, we may also take advantage of the computational efficiency of the fast Fourier transform to compute this initial transform, as well as the subsequent inverse Fourier transforms required for each desired value of l and Q . In addition, if we also choose a window that has finite non-zero extent in the frequency-domain, the length of these inverse transforms need only be equal to the number of non-zero frequency-domain window coefficients. As a result, depending upon the particular choice of window and the number of frequencies evaluated, the total computational cost of this form of the discrete Q transform is typically little more than order $N \log N$. For this reason, we refer to the transform of Equation 5.16 as the fast discrete Q transform.

Note that the frequency shift of Equation 5.14 is cyclic due to the assumed periodicity of the frequency-domain signal. As a result, depending upon the particular choice of window, some number of discrete Q transform coefficients for frequency indices near zero frequency or the Nyquist frequency will be adversely effected by frequency-domain aliasing. In practice, we avoid this effect by choosing a window with finite non-zero extent in the frequency-domain and then excluding Q transform coefficients for those frequency indices for which aliasing will occur.

5.3 Analysis window

Although our development of the continuous and discrete Q transforms did not consider a specific choice of analysis window, our discussion in chapter 3 motivated the use of a basis of functions that are well localized in the time-frequency plane. In theory, this suggests the use of Gaussian windows in time and frequency. However, in practice, this choice of window results in a number difficulties due to the Gaussian's infinite extent. In particular, the use of a Gaussian window is incompatible with the

assumption of a periodic window sequence as required by the discrete Q transform. As a result, there is no choice of analysis window that will achieve the theoretical minimum possible time-frequency uncertainty in the discrete case. Instead, we seek a realizable window with a time-frequency localization that comes close to the minimum possible uncertainty, while also basing our decision on a number of other desirable properties such as simplicity, computational efficiency, and minimal spectral leakage.

The topic of time-frequency windows is well studied and a large variety of windows exist that provide relatively good time-frequency localization[116, 119]. However, for our purposes, we will find it convenient to choose a frequency-domain implementation of the so-called bisquare or Connes window, which has the simple form

$$\tilde{w}(f) = \begin{cases} A(1 - (f/\Delta f)^2)^2 & |f| < \Delta f, \\ 0 & \text{otherwise.} \end{cases} \quad (5.17)$$

Here Δf is the half bandwidth of the window at zero magnitude, while A is an overall normalization factor to be determined in section 5.4. For simplicity, we have assumed that the window is centered on time zero and frequency zero.

This choice of window has a number of useful properties.

First, it has a particularly simple analytical form in the frequency-domain. Such a choice is desirable due to the frequency-domain representation of the fast discrete Q transform. In addition, this choice will prove convenient when deriving the appropriate normalization factors in section 5.4.

Second, it has finite non-zero extent in the frequency-domain. As pointed out at the end of section 5.2.2, such a choice allows for shorter inverse Fourier transforms when calculating the fast discrete Q transform, and is therefore advantageous from the point of view of computational efficiency. In addition, in section 5.4.4 we will find that a window with finite non-zero frequency-domain extent also makes possible a normalization that ensures an accurate accounting of total signal energy.

Third, it provides near minimum time-frequency uncertainty. In particular, if we apply the parameterization of section 3.1 to the bisquare window, we find that it achieves a time-frequency localization that is only 4.5 percent greater than the minimum possible time-frequency uncertainty associated with the ideal but unrealizable Gaussian window. We also find a simple relationship between Δf and our standard definition of bandwidth as given by Equation 3.7b. In particular, for a Q transform analysis at a frequency ϕ and a quality factor Q , we find the relationship

$$\Delta f = \frac{\phi\sqrt{11}}{Q}. \quad (5.18)$$

Finally, the bisquare window incurs relatively little energy leakage into time-domain side lobes. In particular, as seen in Figure 5.1, the squared magnitude of the first time-domain side lobe is less than the peak response of the bisquare window by a factor of ~ 600 . In addition, the squared magnitudes of subsequent side lobes drops off at a rate of $1/\delta t^3$, where δt is the relative time from the center of the window. In this respect, the bisquare window has a strong resemblance to the more commonly encountered Hann window. In fact, the bisquare window is simply a fourth order series expansion of the Hann window, adjusted to ensure a zero first derivative at the window edge. Here we have selected the bisquare window primarily because it simplifies calculation of the normalization in section 5.4.4. For comparison, in Figure 5.1 we present the time-domain and frequency-domain envelopes of the bisquare and Hann windows, as well as those of the ideal but unrealizable Gaussian window and the more straightforward rectangular window. In contrast to the bisquare window, the rectangular window exhibits an attenuation factor of only ~ 20 for the squared magnitude of the first time-domain side lobe and drops of at a rate of only $1/\delta t$. As a result, the bisquare window exhibits significantly reduced spectral leakage relative to the rectangular window. This dramatically improves the detection of relatively weak signals in the presence of nearby stronger signals. The similarity of the bi-square window to the ideal Gaussian window is also evident in Figure 5.2, where we compare example time-domain basis functions for each window assuming a typical Q value of 10.

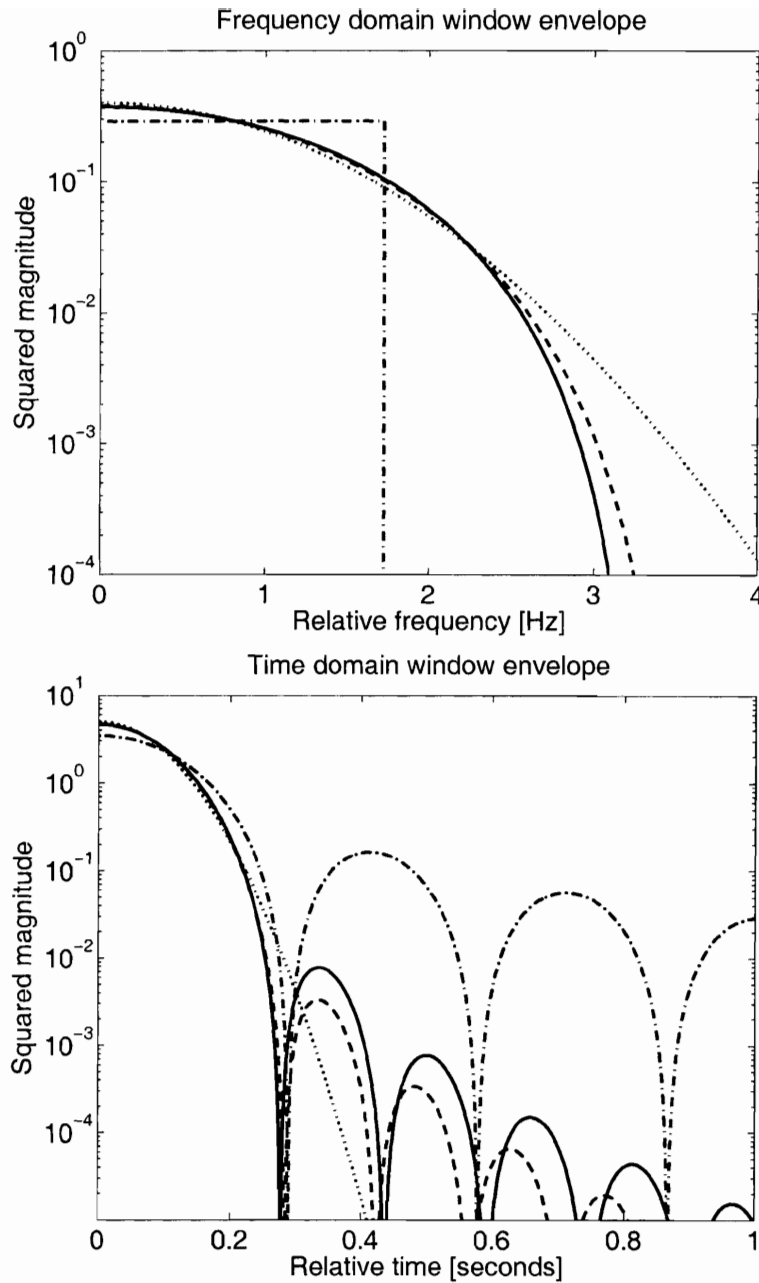


Figure 5.1: The one-sided frequency-domain (top) and one-sided time-domain (bottom) envelopes of the bisquare (solid) window are compared with those of the Gaussian (dotted), Hann (dashed) and rectangular (dashed-dotted) windows. For comparison, all of the windows have an identical bandwidth of 1 Hz as defined by Equation 3.7b and are normalized such that the integral of their squared magnitude is unity in both the time and frequency domains. Note that while the optimal Gaussian window is plotted for reference, it is not achievable in the discrete case due to its infinite extent. Also note the strong similarity between the bisquare and Hann windows and the fact that they have significantly less energy leakage into time-domain side lobes compared with the more straightforward choice of a rectangular window.

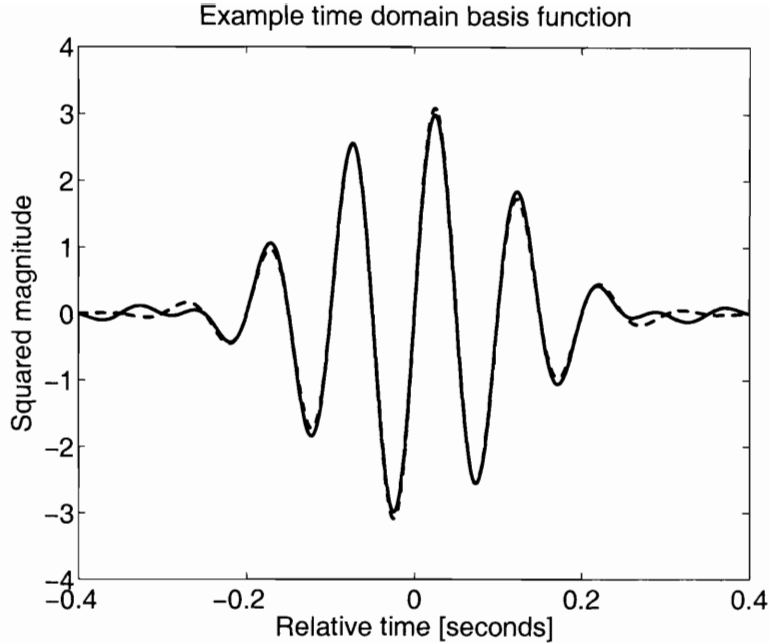


Figure 5.2: Example time-domain basis functions for both a Gaussian window (solid) and bisquare window (dashed). For comparison, both basis functions correspond to a time-frequency tile with a central frequency of 10 Hz and a bandwidth of 1 Hz, resulting in a Q of 10 as defined by Equation 3.6b, Equation 3.7b, and Equation 3.8. In addition, both are normalized such that the integral of their squared magnitude is unity.

For the case of the discrete Q transform, we simply choose a discretized form of the bisquare window consisting of periodic samples of Equation 5.17. At the end of the next section, our goal will be to find an appropriate normalization of this discretized form such that the resulting discrete Q transform coefficients effectively emulate samples of an ideal underlying continuous distribution. In doing so, we will find it simpler to first analytically normalize the continuous form of the bisquare window, and then apply the appropriate conversion factors for the discrete case.

We also note that, in the discrete case, the finite non-zero frequency-domain extent of the bisquare window enables us to strictly avoid frequency-domain aliasing. This is accomplished by identifying a minimum permissible Q and maximum permissible frequency index l such that the full zero to zero bandwidth of the frequency shifted window never extends beyond zero frequency or the Nyquist frequency. For the

bisquare window and a discrete time series of length N , we find the limits

$$Q \geq \sqrt{11} \tag{5.19a}$$

$$l \leq \frac{N}{2} \left(1 + \frac{\sqrt{11}}{Q} \right)^{-1}. \tag{5.19b}$$

At this point, the reader may be concerned that a different window has been selected for our analysis than was assumed in section 3.2.2 when determining the optimal tiling for a specified worst case fractional energy loss. However, since the time-frequency area of the bisquare window is only 4.5 percent greater than that of the Gaussian window, this discrepancy is of little practical import. To quantify this, in section 6.1 we will present a validation of our proposed tiling using the Q transform applied to Gaussian signals injected into stationary white noise. In doing so, we will find that the use of a bisquare window in conjunction with the proposed tiling of section 3.2.2 is still consistent with the specified worst case fractional energy loss.

5.4 Normalization

There are two different criteria we may use to determine the normalization of the analysis window. In what follows, we first determine the normalization necessary to exactly recover, in a single time-frequency tile, the energy of bursts which are well localized in the time-frequency plane. We also find that, when the Q transform is applied to a stationary stochastic process, this same normalization permits a useful interpretation of mean tile energy in terms of power spectral density. By combining these results, we then consider the recovery of well localized bursts in the presence of stationary detector noise.

Next, we take into account the energy overlap between time-frequency tiles. This allows us to develop an alternative normalization that ensures an accurate accounting

of total signal energy when integrated over all tiles. As a result, we exactly recover the energy of bursts that are poorly localized in the time-frequency plane.

Finally, we consider the appropriate normalization of the discretized analysis window such that the resulting discrete Q transform coefficients approximate samples of a hypothetical continuous Q transform.

5.4.1 Localized bursts

We first consider the appropriate normalization of the bisquare window such that when applied to the analysis of a well localized burst, the magnitude of the best match Q transform coefficient is an accurate measure of the burst's characteristic strain amplitude $\|h\|$ as defined in section 3.1.1. That is, we desire that

$$\int_{-\infty}^{+\infty} h(t; \tau, \phi, Q) w(t - \tau, \phi, Q) e^{-i2\pi\phi t} dt = \|h\| e^{i\theta}, \quad (5.20)$$

where we have assumed that the signal and the modulated analysis window have the same center time, center frequency, and Q, and that there is an arbitrary phase θ between them.

To accomplish this, we consider a real valued test signal which is also bisquare in form and has the same center time, center frequency, and Q as the modulated analysis window under test. In addition, we may simplify our notation without loss of generality by dropping the constant ϕ and Q from the list of window parameters and by taking τ to be zero. In this case, our test signal has a time-domain and a frequency-domain representation given by

$$h(t) = B \|h\| w(t) \cos(2\pi\phi t + \theta) \quad (5.21a)$$

$$\tilde{h}(f) = B \|h\| [\tilde{w}(f - \phi) e^{i\theta} + \tilde{w}(f + \phi) e^{-i\theta}] / 2. \quad (5.21b)$$

Here B is a constant normalization factor that is chosen in accordance with the definition of characteristic strain amplitude in Equation 3.2. If we also select Q in accordance with the condition of Equation 5.19a, then the two terms in the frequency-domain representation of our test signal are guaranteed not to overlap. As a result, we find that our test signal must obey the normalization condition

$$\frac{B^2}{2} \int_{-\infty}^{+\infty} |\tilde{w}(f)|^2 df = 1. \quad (5.22)$$

If we now apply the frequency-domain form of the continuous Q transform to our test signal, we find the expression

$$H(\tau, \phi, Q) = \frac{B}{2} \|h\| \int_{-\infty}^{+\infty} [|\tilde{w}(f)|^2 e^{i\theta} + \tilde{w}(f + 2\phi) \tilde{w}^*(f) e^{-i\theta}] df. \quad (5.23)$$

Again, assuming that Q is selected in accordance with Equation 5.19a, the second term in this integral vanishes due to the non-overlap of the multiplicands. As a result, in order to meet the requirement of Equation 5.20, we find the normalization condition

$$\frac{B}{2} \int_{-\infty}^{+\infty} |\tilde{w}(f)|^2 df = 1. \quad (5.24)$$

It is easy to see that the simultaneous normalization conditions of Equation 5.22 and Equation 5.24 are then satisfied, and the signal's characteristic strain amplitude is exactly recovered, if B is unity and the window obeys the normalization condition

$$\int_{-\infty}^{+\infty} |\tilde{w}(f)|^2 df = 2. \quad (5.25)$$

We thus seek the coefficient A in Equation 5.17 such that the one-sided frequency-domain integral of the window's squared magnitude is unity. Straightforward integration then yields a normalization factor of

$$A^2 = \frac{315}{128\sqrt{11}} \frac{Q}{\phi}, \quad (5.26)$$

where we have substituted the expression of Equation 5.18 for Δf .

Here, for convenience, we provide the final form of the normalized window for use in the alternative frequency-domain form of the continuous Q transform of Equation 5.8:

$$\tilde{w}(f, \phi, Q) = \begin{cases} \left(\frac{315}{128\sqrt{11}} \frac{Q}{\phi} \right)^{1/2} \left[1 - \left(\frac{fQ}{\phi\sqrt{11}} \right)^2 \right]^2 & |f| < \frac{\phi\sqrt{11}}{Q}, \\ 0 & \text{otherwise.} \end{cases} \quad (5.27)$$

5.4.2 Stationary stochastic processes

With this choice of normalization, the magnitudes of the resulting coefficients also have a useful interpretation when the Q transform is applied to the analysis of a stationary stochastic process. To develop this result, we begin with the direct form of the continuous Q transform applied to the stationary stochastic process $n(t)$,

$$N(\tau, \phi, Q) = \int_{-\infty}^{+\infty} n(t) w(t - \tau, \phi, Q) e^{-i2\pi\phi t} dt. \quad (5.28)$$

The squared magnitude of the resulting Q transform coefficients are then given by the double integral

$$|N(\tau, \phi, Q)|^2 = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} n(t) n(t') w(t - \tau) w(t' - \tau) e^{-i2\pi\phi(t-t')} dt dt'. \quad (5.29)$$

Again, we have taken ϕ and Q to be constant and dropped them from the list of window parameters in order to simplify our notation.

Due to the random nature of the data under test, the resulting squared magnitude of individual Q transform coefficients is necessarily a random variable. In section 5.5 we will determine the distribution of this random variable for the special case of stationary white noise. Here, however, we are only interested in its expectation value over an ensemble of many measurements. In the process, due to the stationarity of $n(t)$, we may also take τ to be zero without loss of generality. The resulting

expectation value is then given by the expression

$$\langle |N(\tau, \phi, Q)|^2 \rangle = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \langle n(t) n(t') \rangle w(t) w(t') e^{-i2\pi\phi(t-t')} dt dt'. \quad (5.30)$$

Next, we again make use of the stationarity of $n(t)$, which permits us to substitute the relative time lag T for the time difference $t - t'$, such that

$$\langle |N(\tau, \phi, Q)|^2 \rangle = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \langle n(t) n(t - T) \rangle w(t) w(t - T) e^{-i2\pi\phi T} dt dT. \quad (5.31)$$

We then recognize the expectation value on the right hand side as the stationary auto-correlation of $n(t)$ evaluated at time lag T ,

$$r_n(T) = \langle n(t) n(t - T) \rangle. \quad (5.32)$$

By substituting this relation and rearranging the resulting integral, we obtain the expression

$$\langle |N(\tau, \phi, Q)|^2 \rangle = \int_{-\infty}^{+\infty} r_n(T) \left(\int_{-\infty}^{+\infty} w(t) w(t - T) dt \right) e^{-i2\pi\phi T} dT. \quad (5.33)$$

Finally, we recognize this as the Fourier transform of a product of two time-domain functions, which is simply the frequency-domain convolution

$$\langle |N(\tau, \phi, Q)|^2 \rangle = \int_{-\infty}^{+\infty} G_n(f) |\tilde{w}(\phi - f)|^2 df. \quad (5.34)$$

Here $G_n(f)$ is the two-sided power spectral density of the stationary stochastic process $n(t)$ as defined by Equation 3.29 and we have used the Fourier transform relation

$$\int_{-\infty}^{+\infty} w(t) w(t - T) dt = \int_{-\infty}^{+\infty} |\tilde{w}(f)|^2 e^{-i2\pi\phi T} df, \quad (5.35)$$

for the auto-correlation of the time-domain window function.

To understand this result, we first choose Q in accordance with Equation 5.19a such

that the frequency shifted analysis window does not extend below zero frequency. This then allows us to rewrite Equation 5.34 in terms of the one-sided power spectral density $S_n(f)$ as defined by Equation 3.30, which gives the desired result

$$\langle |N(\tau, \phi, Q)|^2 \rangle = \frac{1}{2} \int_0^\infty S_n(f) |\tilde{w}(\phi - f)|^2 df. \quad (5.36)$$

We may then interpret the squared window function as a one-sided frequency-domain probability density function that is normalized according to Equation 5.25 such that the total energy of the signal is accounted for. Thus, when the Q transform is applied to the analysis of a stationary stochastic process, the expected value of the squared magnitude of the resulting transform coefficients is simply a measure of the signal's average one-sided power spectral density over the bandwidth of the frequency shifted analysis window. Note that in deriving this result, we have made no assumption other than that of a stationary stochastic process. If, however, the power spectral density of the stochastic process is approximately constant over the bandwidth of the analysis window, then the Q transform provides a direct measure of the local power spectral density,

$$\langle |N(\tau, \phi, Q)|^2 \rangle \simeq S_n(\phi). \quad (5.37)$$

5.4.3 Localized bursts in noise

It is instructive at this point to consider the case of well localized bursts in the presence of additive detector noise. That is, we consider the input time series

$$x(t) = n(t) + h(t - \tau, \phi, Q), \quad (5.38)$$

where $n(t)$ is a stationary stochastic process with one-sided power spectral density $S_n(f)$, while $h(t - \tau, \phi, Q)$ is a well localized burst with center time τ , center frequency ϕ , and quality factor Q .

Due to the linearity of the Q transform, the resulting complex valued Q transform coefficients are simply vector sums of the coefficients produced when the Q transform is applied separately to the burst and to the detector noise. As a result, the squared Q transform magnitude for the best match time-frequency tile is given by

$$|X(\tau, \phi, Q)|^2 = \|h\|^2 + |N(\tau, \phi, Q)|^2 + 2\|h\| |N(\tau, \phi, Q)| \cos \theta, \quad (5.39)$$

where $\|h\|$ is the characteristic strain amplitude of the well localized burst, $N(\tau, \phi, Q)$ is the Q transform applied only to the detector noise, and θ is a uniformly distributed random phase between the two complex valued transforms.

Here, however, we are not concerned with the value of a single measurement, but rather with the ensemble average of many such measurements. In this case the phase dependent term vanishes due to symmetry, leaving only the quadrature sum

$$\langle |X(\tau, \phi, Q)|^2 \rangle = \|h\|^2 + \frac{1}{2} \int_0^\infty S_n(f) |\tilde{w}(\phi - f)|^2 df. \quad (5.40)$$

If, in addition, the power spectral density of the detector noise is approximately constant over the bandwidth of the burst, then we obtain the simple result

$$\langle |X(\tau, \phi, Q)|^2 \rangle \simeq \|h\|^2 + S_n(f). \quad (5.41)$$

This motivates the definition of squared signal to noise ratio,

$$\rho^2 = \frac{\langle |X(\tau, \phi, Q)|^2 \rangle}{\langle |N(\tau, \phi, Q)|^2 \rangle} - 1, \quad (5.42)$$

which for a locally flat detector noise spectrum, yields the intuitive result

$$\rho^2 \simeq \frac{\|h\|^2}{S_n(f)}. \quad (5.43)$$

Thus we see that the signal to noise ratio of a well localized burst is directly proportional to its characteristic strain amplitude $\|h\|$, as long as the detector noise

spectrum is approximately constant over the bandwidth of the burst. As we will see in section 5.5, the squared signal to noise ratio of a well localized burst is also directly proportional to its significance, which in turn determines its detectability. In such cases, $\|h\|^2$ can be readily calculated for a postulated astrophysical source and then directly compared with noise power spectra $S_n(f)$ of a existing detectors in order to quickly evaluate a source's detectability.

We also note that the squared signal to noise ratio of Equation 5.43 is half of the optimal squared signal to noise ratio found from Equation 3.32 under similar conditions. However, this latter result is only achieved by matched filtering if the signal to be detected is completely known in advance. If instead, the waveform of the well localized burst is only known to within an arbitrary phase, then the signal to noise ratio achieved by matched filtering is in fact equal to the result of Equation 5.43[120]. As a result, for well localized bursts of unknown phase in the presence of a locally flat detector noise spectrum, the performance of the Q transform approaches that of the matched filter.

We should point out, however, that for poorly localized bursts, or for detector noise spectra that vary significantly over the bandwidth of a burst, the optimal signal to noise ratio will not be achieved by the Q transform. As a result, the detection of such bursts will be somewhat less than optimal, unless the Q transform is supplemented by other techniques. For poorly localized bursts, one option is to consider the joint statistical significance of multiple time-frequency tiles[103]. However, while such a clustering is briefly considered in section 5.4.4 to reconstruct the energy of poorly localized bursts, its application to their detection is beyond the intended scope of this work.

It is possible, however, to avoid the difficulties associated with the frequency dependence of detector noise by first whitening the data. However, in doing so we are actually performing a search of the whitened data stream rather than the gravitational-wave data stream. In fact, if the Q transform is applied to data that has been suffi-

ciently whitened using the techniques of chapter 4, then the resulting search is equivalent to an optimal matched filter search for well localized bursts within the whitened data stream. This is very similar to the well developed search for gravitational-waves from inspiraling binary neutron stars, for which the expected waveforms are well known in advance[88]. The only difference is the choice of signal space. In the first case, the whitened data stream is searched for well localized bursts of unknown phase that fall within a targeted region of time, frequency, and Q space. In the latter case, the gravitational-wave data stream is searched for inspiral waveforms of unknown phase corresponding to a targeted range of binary neutron star masses. Note that, since our measurement basis was not selected on astrophysical grounds, there is nothing problematic about searching for bursts in the whitened data stream rather than the gravitational-wave data stream.

Of course, in practice we are not provided with an ensemble of identically prepared bursts. Nor are we provided with separate signal and noise data. As a result, we cannot determine the theoretical signal to noise ratio defined by Equation 5.42. Instead, assuming stationary detector noise, we estimate the total energy of a single candidate burst by

$$\|\hat{h}\|^2 = |X(\tau, \phi, Q)|^2 - \langle |X(t, \phi, Q)|^2 \rangle, \quad (5.44)$$

and its corresponding signal to noise ratio by

$$\hat{\rho}^2 = \frac{|X(\tau, \phi, Q)|^2}{\langle |X(t, \phi, Q)|^2 \rangle} - 1. \quad (5.45)$$

Here $|X(\tau, \phi, Q)|^2$ is given by Equation 5.39 and $\langle |X(t, \phi, Q)|^2 \rangle$ is the average energy of a representative set of time-frequency tiles, all of which have the same frequency and Q as the candidate burst, but occur at times t other than τ . Unfortunately, with only one observation of the burst, the estimated energy and signal to noise ratio of the burst depend upon the instantaneous magnitude and phase of the detector noise, both of which are random variables. This results in an estimation error that depends upon the specific statistical properties of the detector noise and will be revisited in

section 5.5 for the special case of stationary white noise. For now, we note that the error in recovering the signal to noise ratio typically becomes important only for low signal to noise ratio bursts, where it is on the order of unity.

5.4.4 Non-localized bursts

Whereas the previous sections have all focused on the window normalization necessary to recover the total energy of a well localized burst within a single Q transform coefficient, this section presents an alternative window normalization which is also appropriate for bursts that are poorly localized in the time-frequency plane. In particular, we seek a generalization of Parseval's theorem to the Q transform such that we may exactly recover the total energy of an arbitrary burst by integrating its squared Q transform magnitude over the entire time-frequency plane. In doing so, we must take into account the Q dependent frequency-domain overlap between time-frequency tiles. This approach is particularly important for the case of non-localized bursts, which are by definition not well represented by a single time-frequency tile.

We first consider the integral over time of the squared Q transform magnitude. In doing so, we take advantage of the standard form of Parseval's theorem and the Fourier transform relationship of Equation 5.4c to express this integral in the Fourier domain as

$$\int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau = \int_{-\infty}^{+\infty} |\tilde{X}(f, \phi, Q)|^2 df. \quad (5.46)$$

Next, by substituting the Fourier domain form of the continuous Q transform from Equation 5.5 and Equation 5.6, we find the expression

$$\int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau = \int_{-\infty}^{+\infty} |\tilde{x}(f + \phi)|^2 |\tilde{w}(f, \phi, Q)|^2 df. \quad (5.47)$$

We then perform the change of variables, $\psi = f + \phi$, in order to isolate the ϕ

dependency of the right hand side to within the window function, such that

$$\int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau = \int_{-\infty}^{+\infty} |\tilde{x}(\psi)|^2 |\tilde{w}(\psi - \phi, \phi, Q)|^2 d\psi. \quad (5.48)$$

One-sided integration over positive analysis frequency ϕ then yields,

$$\int_0^{\infty} \int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau d\phi = \int_{-\infty}^{+\infty} |\tilde{x}(\psi)|^2 \int_0^{\infty} |\tilde{w}(\psi - \phi, \phi, Q)|^2 d\phi d\psi. \quad (5.49)$$

If we now assume that Q has been selected in accordance with Equation 5.19a such that the frequency shifted analysis window does not extend below zero frequency, then we find that the right hand side of the previous equation is non-zero only for positive ψ . If in addition we require the window normalization condition

$$\int_0^{\infty} |\tilde{w}(\psi - \phi, \phi, Q)|^2 d\phi = \begin{cases} 2 & \psi > 0, \\ 0 & \psi \leq 0, \end{cases} \quad (5.50)$$

then this previous equation becomes

$$\int_0^{\infty} \int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau d\phi = 2 \int_0^{\infty} |\tilde{x}(\psi)|^2 d\psi. \quad (5.51)$$

Finally, if the original time series $x(t)$ is real-valued, then its Fourier transform $\tilde{x}(\psi)$ is complex conjugate symmetric about zero ψ . As a result, the right hand side of the previous equation is simply the total signal energy as defined by Equation 3.2,

$$2 \int_0^{\infty} |\tilde{x}(\psi)|^2 d\psi = \int_{-\infty}^{+\infty} |\tilde{x}(\psi)|^2 d\psi = \|x\|^2. \quad (5.52)$$

Thus, if the normalization condition of Equation 5.50 is satisfied, then the integral of the squared Q transform magnitude over all time and positive frequencies is simply the squared characteristic amplitude or total signal energy

$$\int_0^{\infty} \int_{-\infty}^{+\infty} |X(\tau, \phi, Q)|^2 d\tau d\phi = \|x\|^2. \quad (5.53)$$

We therefore seek the coefficient A in Equation 5.17 such that bisquare window obeys the normalization condition of Equation 5.50. As a starting point, however, we take the previously determined window normalization of Equation 5.26 and seek a correction factor A_1 such that a normalization coefficient of the form

$$A^2 = \frac{315}{128\sqrt{11}} \frac{Q}{\phi} A_1^2, \quad (5.54)$$

satisfies the condition of Equation 5.50.

Note, however, that the integral of Equation 5.50 is over the analysis frequency ϕ and not the Fourier domain variable f , as was the case for the previous normalization condition of Equation 5.25. In addition, we also assume that A_1 is a function only of Q and not of ϕ . That is, we only take into account the Q dependent frequency-domain overlap between tiles, while preserving as much as possible the previously determined normalization. As a result, the same value of A_1 applies uniformly to all tiles within a time-frequency plane of constant Q .

The integral of Equation 5.50 for the specific case of the bisquare window is laborious and will not be considered here in detail. However, it should be noted that a solution is not guaranteed for all choices of analysis window. In particular, the necessary integral only converges for windows that have finite non-zero extent in the frequency domain. In addition, our primary reason for selecting the bisquare window, instead of the more commonly encountered Hann window, is due to its analytical simplicity, which in this case permits a closed form solution for the normalization correction factor A_1 :

$$A_1^2 = \frac{256}{315q} \left[(q^8 - 4q^6 + 6q^4 - 4q^2 + 1) \ln \left(\frac{q+1}{q-1} \right) + \left(-2q^7 + \frac{22}{3}q^5 - \frac{146}{15}q^3 + \frac{186}{35}q \right) \right]^{-1}. \quad (5.55)$$

Here, for notational convenience, we have introduced the alternative quality factor,

$$q = Q/\sqrt{11}, \quad (5.56)$$

which is simply the ratio of the canonical quality factor Q defined in Equation 3.8 to the minimum permissible Q from Equation 5.19a.

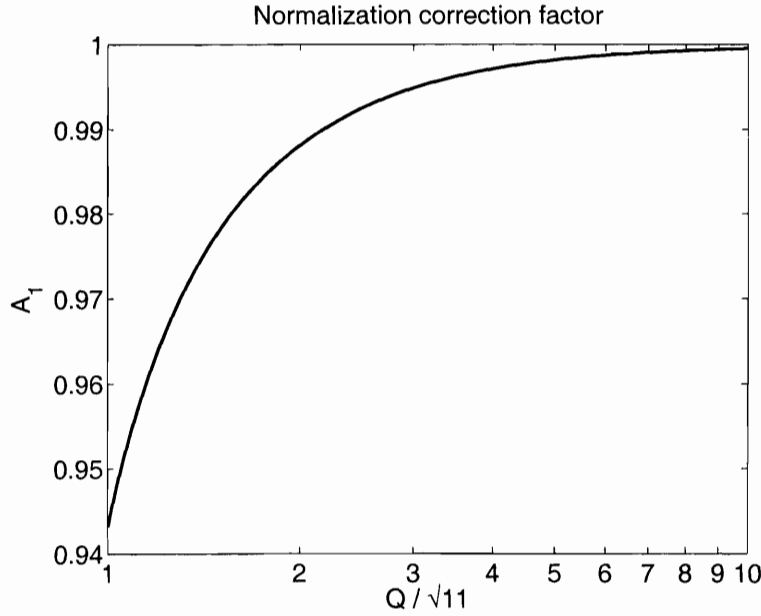


Figure 5.3: The necessary Q dependent normalization correction factor to allow for a complete accounting of the energy of non-localized bursts. Note that the correction is only marginally important at the minimum permissible Q , where the greatest frequency-domain overlap between tiles occurs. For larger values of Q , the correction factor quickly becomes negligible.

The behavior of this normalization correction factor is plotted in Figure 5.3 as a function of Q . Note that it has its greatest impact at the minimum permissible Q , where the frequency-domain overlap between tiles is at its maximum. In contrast, for larger values of Q this correction factor approaches unity and quickly becomes negligible as the frequency-domain overlap between tiles decreases. Even for the minimum permissible Q , however, the correction is only on the order of five percent. We also note that the calculation of A_1 as given by Equation 5.55 becomes numerically unstable for large values of Q . In this case, however, one can safely take A_1 to be unity without significant loss of accuracy.

Finally, we briefly consider the case of a non-localized burst in the presence of additive detector noise. Unfortunately, in this case the total signal energy of the noise has little meaning since it becomes infinite as the integral of Equation 5.53 is extended over all time and positive frequency. Instead, we only integrate over a finite region of the time-frequency plane, but one that is sufficiently large to encompass the entire burst. It is then possible to accurately recover the total signal energy of the burst. However, to do so, one must also know the integrated energy of the detector noise when no burst is present. In practice, this can be estimated by determining the mean integrated energy of similar time-frequency regions, assuming they themselves do not contain bursts. The total signal energy of the burst is then recovered by simply taking the difference between these foreground and background measurements. Note that this procedure is similar to the recovery of well localized bursts in the presence of additive detector noise as described at the end of section 5.4.3. As a result, we expect to encounter errors in the recovery of poorly localized bursts related to fluctuations in the instantaneous magnitude and phase of the detector noise. However, a detailed treatment of the issues involved in the detection of poorly localized bursts is beyond the scope of this work and is left to future investigation.

5.4.5 The discrete Q transform

Up to this point, we have only considered the window normalization appropriate to the continuous Q transform. In this section, however, we develop the normalization appropriate to the discrete case such that the resulting discrete Q transform coefficients approximate samples of an ideal underlying continuous distribution. To do so, we simply discretize the continuous frequency-domain window by periodic sampling, and then determine the necessary normalization correction factors such that the summations which occur in the fast discrete Q transform approximate continuous integrals. The resulting correction factors are equally applicable to the normalization developed in section 5.4.1 for the case of well localized bursts as they are to

the alternative normalization developed in section 5.4.4 for the case of non-localized bursts.

We begin by assuming that $x[n]$, the discrete time series under test, is produced by periodic sampling of an ideal continuous time series $x(t)$ such that

$$x[n] = x(n dt). \quad (5.57)$$

Here $dt = 1/f_s$ is the time-domain sample interval, while f_s is the corresponding sample frequency. In what follows, we assume that the sample frequency is sufficient to adequately resolve the original continuous time-domain signal.

Next, by considering both the continuous and the discrete Fourier transforms,

$$\tilde{x}(f) = \int_{-\infty}^{+\infty} x(t) e^{-i2\pi ft} dt \quad (5.58a)$$

$$\tilde{x}[k] = \sum_{n=0}^{N-1} x[n] e^{-i2\pi nk/N}, \quad (5.58b)$$

we find that the resulting continuous and discrete frequency series are approximately related by the expression

$$\tilde{x}[k] \simeq \frac{\tilde{x}(k df)}{dt}, \quad (5.59)$$

where $df = N/f_s$ is the frequency-domain sample interval.

If we then compare the frequency domain forms of the continuous and discrete Q transforms,

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} \tilde{x}(f + \phi) \tilde{w}^*(f, \phi, Q) e^{-i2\pi f\tau} df \quad (5.60a)$$

$$X[m, l, Q] = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{x}[k + l] \tilde{w}^*[k, l, Q] e^{-i2\pi mk/N}, \quad (5.60b)$$

we find that the resulting discrete transform coefficients will approximate samples of

the continuous transform if

$$\tilde{w}[k, l, Q] \simeq \tilde{w}(k df, l df, q) N dt df. \quad (5.61)$$

This is a surprisingly simple result since the factors $N dt df$ completely cancel each other out. Thus, we find that no correction factor is necessary for the discrete case, and we may discretize the continuous frequency-domain window simply by periodic sampling with a frequency-domain sample interval of $df = N/f_s$.

However, there is another correction factor which must be taken into account if we are to take full advantage of the computational efficiency of the fast discrete Q transform. As pointed out at the end of section 5.2.2, an analysis window with finite non-zero frequency-domain extent allows for increased computational efficiency by permitting shorter inverse Fourier transforms in the evaluation of the fast discrete Q transform. In particular, if the analysis window for a given value of l and Q contains only M non-zero values, then we may compute the discrete Q transform using an inverse Fourier transform of length M rather than the full length N by noting that

$$X[m \frac{M}{N}, l, Q] = \frac{M}{N} \frac{1}{M} \sum_k \tilde{x}[k + l] \tilde{w}^*[k, l, Q] e^{-i2\pi m(M/N)k/M}. \quad (5.62)$$

Here the summation is only over the M values of k in the set

$$0 \leq k \leq \frac{M}{2} \quad \text{and} \quad N - \frac{M}{2} + 1 \leq k \leq N - 1. \quad (5.63)$$

Note that, in this case, the resulting discrete Q transform will only have M time samples over the duration of the input time series instead of N . As a result, the discrete Q transform is effectively downsampled in time by a factor of N/M . In practice, both N and M are typically integer powers of two in order to take advantage of the computational efficiency of the fast Fourier transform. As a result, the downsampling ratio is also typically an integer power of two.

Note also that the definition of the downsampled transform of Equation 5.62 follows directly from the original definition of the fast discrete Q transform and the assumption of only M non-zero coefficients. In the process, we have been careful to ensure that, for those original time indices m that correspond to integer valued downsampled time indices mM/N , the values of the corresponding discrete Q transform coefficients remain the same. Thus, by inspection, we find that the appropriately normalized discrete frequency-domain window is given in terms of the continuous frequency-domain window of Equation 5.27 by the relation

$$\tilde{w}[k, l, Q] = \left(\frac{M}{N}\right) w\left(k\frac{f_s}{N}, l\frac{f_s}{N}, Q\right). \quad (5.64)$$

Finally, we note that M need not be exactly equal the number of non-zero window coefficients. Instead, we may also select larger values of M up to and including the length N of the input time series. This freedom to specify M then provides some measure of control over the temporal overlap of the resulting discrete Q transform tiles. As a result, it is possible to select combinations of l , Q , and M such that the targeted space of time, frequency, and Q is optimally tiled according to the procedure of section 3.2.2. That is, with the minimum number of tiles necessary to ensure that the mismatch between an arbitrarily well localized burst and the closest measurement tile results in no more than a pre-specified fractional loss in the measured signal energy.

5.5 Statistics

We now turn to the statistical interpretation of Q transform coefficients in order to quantify their significance in the presence of additive noise. Here we focus only on the discrete case, since in practice this is the transform that is applied. In doing so, our task is greatly simplified if the data has first been whitened using the techniques of

chapter 4. In this case, we will find that the distribution of the squared Q transform magnitudes takes on a particularly simple form. We then estimate the parameters of this distribution using a method that is insensitive to transient non-stationarities in the input data. The resulting distribution allows us to define the significance of Q transform coefficients as the probability of observing a coefficient with greater magnitude.

Next, we apply these concepts to the detection of well localized bursts. First, we quantify the error in recovering the signal to noise ratio of a burst due the presence of stationary white noise. Next, we consider the fraction of Q transform coefficients which exceed a specified white noise significance. Together, these results allow us to produce so-called receiver operating characteristics, which display the theoretical detection efficiency of localized bursts as a function of signal to noise ratio and white noise event rate. These predicted detection efficiencies then form a basis for comparison that will be useful in chapter 6 when validating our analysis pipeline.

5.5.1 Stationary white noise

In what follows, we assume that the input time series consists of ideal stationary white noise as described in chapter 4. That is, we assume that each sample is an independent random variable drawn from a common distribution. For simplicity, we also assume that the input data has been detrended and normalized such that this underlying distribution has zero mean and unity standard deviation. We do not, however, make any other assumption about the particular distribution from which the samples are drawn.

The Q transform then projects this input data onto a basis of complex windowed exponentials. To develop the statistics of the resulting coefficients, it is useful to consider their real and imaginary components separately. That is, we write the Q

transform as

$$X[m, l, Q] = X_R[m, l, Q] + iX_I[m, l, Q], \quad (5.65)$$

where the real and imaginary components are given by the projections

$$X_R[m, l, Q] = \sum_{n=0}^{N-1} x[n] w[n - m, l, Q] \cos(2\pi nl/N), \quad (5.66a)$$

$$X_I[m, l, Q] = \sum_{n=0}^{N-1} x[n] w[n - m, l, Q] \sin(2\pi nl/N). \quad (5.66b)$$

It is also useful to explicitly define the corresponding basis functions,

$$b_R[n; m, l, Q] = w[n - m, l, Q] \cos(2\pi nl/N), \quad (5.67a)$$

$$b_I[n; m, l, Q] = w[n - m, l, Q] \sin(2\pi nl/N). \quad (5.67b)$$

If these basis functions are of sufficient duration, then the well known central limit theorem provides us with the distributions of $X_R[m, l, Q]$ and $X_I[m, l, Q]$. In particular, both projections are Gaussian distributed with zero mean and with standard deviations of

$$\sigma_R^2[m, l, Q] = \sum_{n=0}^{N-1} |b_R[n; m, l, Q]|^2, \quad (5.68a)$$

$$\sigma_I^2[m, l, Q] = \sum_{n=0}^{N-1} |b_I[n; m, l, Q]|^2. \quad (5.68b)$$

This is a particularly powerful result: the resulting distributions do not depend upon the specific distribution of the input noise. However, it is important to note that this result is only approximate, and that this approximation is poor for basis functions of very short duration. Fortunately, for non-pathological noise, the minimum Q of Equation 5.19a and the maximum analysis frequency of Equation 5.19b are typically sufficient to ensure reasonable agreement with a Gaussian distribution.

We also note that, for the above results to hold, the input data need only be uncorrelated on time scales equal to or shorter than the duration of the longest basis function

under consideration. For this reason, according to the discussion of section 4.3, the order M of the linear predictor error filter used to whiten the data is typically set equal to the effective length of this longest basis function.

For notational clarity, we now omit the indices m , l , and Q whenever their presence is understood from context. According to the central limit theorem, the distributions of the real and imaginary Q transform coefficients are then

$$f(X_R) dX_R = \left(\frac{1}{2\pi\sigma_R^2} \right)^{1/2} \exp \left(-\frac{X_R^2}{2\sigma_R^2} \right) dX_R, \quad (5.69a)$$

$$f(X_I) dX_I = \left(\frac{1}{2\pi\sigma_I^2} \right)^{1/2} \exp \left(-\frac{X_I^2}{2\sigma_I^2} \right) dX_I. \quad (5.69b)$$

We now seek the corresponding distribution of Q transform magnitudes. This is greatly simplified if the basis functions $b_R[n]$ and $b_I[n]$ are orthogonal. Fortunately, for the bisquare window, this is yet another property that is guaranteed by constraints of Equation 5.19a and Equation 5.19b. As a result, the random variables X_R and X_I are guaranteed to be independent and their joint probability distribution is simply the product of their individual distributions. In addition, the constraints of Equation 5.19a and Equation 5.19b are also sufficient to ensure the equality of σ_R and σ_I , both of which we will denote by σ_X . Combining these results allows us to write the joint probability distribution as

$$f(X_R, X_I) dX_R dX_I = \frac{1}{2\pi\sigma_X^2} \exp \left(-\frac{X_R^2 + X_I^2}{2\sigma_X^2} \right) dX_R dX_I. \quad (5.70)$$

Finally, if we perform a change of variables to the polar coordinate system,

$$|X| = (X_R^2 + X_I^2)^{1/2}, \quad (5.71a)$$

$$\angle X = \tan^{-1} (X_I/X_R), \quad (5.71b)$$

we find that the Q transform magnitude $|X|$ is distributed as

$$f(|X|) d|X| = \frac{|X|}{\sigma_X^2} \exp\left(-\frac{|X|^2}{2\sigma_X^2}\right) d|X|, \quad (5.72)$$

which is the well known Rayleigh distribution.

Alternatively, the distribution of the squared Q transform magnitude,

$$f(|X|^2) d|X|^2 = \frac{1}{2\sigma^2} \exp\left(-\frac{|X|^2}{2\sigma^2}\right) d|X|^2, \quad (5.73)$$

is simply an exponential distribution.

Not surprisingly, these are the same distributions that are found for the magnitude and squared magnitude of discrete Fourier transform coefficients for the case of stationary white noise[121]. Although we make no use of it here, it is interesting to note that this similarity with the discrete Fourier transform also extends to coherent signal content. In this case, we find that the resulting Q transform coefficients obey the Rician distribution

$$f(|X|) d|X| = \frac{|X|}{\sigma_X^2} \exp\left(-\frac{|X|^2 + X_0}{2\sigma_X^2}\right) I_0\left(\frac{|X|X_0}{\sigma_X^2}\right) d|X|, \quad (5.74)$$

where I_0 is a modified Bessel function of the first kind and X_0 is the effective amplitude of any coherent signal content. For the case of no coherent signal content, this expression simplifies to the Rayleigh distribution of Equation 5.72.

We may also consider the phase of Q transform coefficients. Unfortunately, due to the independence of X_R and X_I for stationary white noise, the phase of Q transform coefficients is distributed uniformly over all angles. As a result, Q transform phase provides no useful information for the detection of bursts within a single detector. However, for multiple detectors, the relative Q transform phase between detectors may potentially be applied to both the detection and validation of coincident bursts. Such an approach will be developed in section 5.8, where we consider the coherent

detection of bursts in data from networks of detectors.

5.5.2 Significance

We now define a number of quantities that are useful measures of a burst's detectability in the presence of stationary white noise. Starting with the exponential distribution of Equation 5.73, we find that the mean squared Q transform magnitude is

$$\langle |X|^2 \rangle = 2\sigma_X^2. \quad (5.75)$$

This suggests the definition of the normalized energy for Q transform coefficients,

$$Z = |X|^2 / \langle |X|^2 \rangle, \quad (5.76)$$

for which the exponential distribution of Equation 5.73 takes the particularly simple form,

$$f(Z) dZ = \exp(-Z) dZ. \quad (5.77)$$

It is then straightforward to define the significance of a Q transform coefficient as the probability of observing a coefficient of greater normalized energy,

$$P(Z' > Z) = \exp(-Z). \quad (5.78)$$

We also note that the normalized energy of a Q transform coefficient is closely related to its squared signal to noise ratio as estimated by Equation 5.45,

$$\hat{\rho}^2 = Z - 1. \quad (5.79)$$

We therefore find that the normalized energy, significance, and signal to noise ratio of a single Q transform coefficient are all monotonically related. As a result, we may apply a threshold on any of these three quantities in order to identify candidate bursts.

In practice, however, the estimated signal to noise ratio is somewhat problematic since its squared value can be negative in the absence of a strong signal. In addition, since the significance of a candidate burst is typically a very small number, it is conventional to instead report the negative logarithm of a burst's significance. Conveniently, this is given by the normalized energy,

$$Z = -\ln P(Z' > Z). \quad (5.80)$$

In what follows, we will therefore take the normalized energy of a time-frequency tile as the primary measure of the tile's significance.

Up to this point, we have focused only on the detection of bursts that are well localized in the time-frequency plane. As a result, the detectability of a poorly localized burst is limited by the statistical significance of its well localized substructure. To improve the detectability of such bursts, we may also consider the combined white noise significance of clusters of multiple time-frequency tiles. This task is greatly simplified if we assume that the tiles comprising the cluster are statistically independent of one another. In section 5.6 we will describe an exclusion algorithm which justifies this assumption by identifying the set of non-overlapping time-frequency tiles that best matches the well localized substructure of arbitrary bursts. Assuming statistical independence, the joint white noise probability distribution for a cluster of M tiles is simply the product

$$f(\mathbf{Z}) d\mathbf{Z} = \prod_{i=1}^M \exp(-Z_i) dZ_i. \quad (5.81)$$

We then define the joint normalized energy of such a cluster as the sum of single tile normalized energies,

$$Z_C = \sum_{i=1}^M Z_i. \quad (5.82)$$

The resulting probability distribution is then given by the well known chi-squared

distribution with $2M$ degrees of freedom,

$$f(Z_C) dZ_C = \frac{Z_C^{M-1} \exp(-Z_C)}{\Gamma(M)} dZ_C. \quad (5.83)$$

We may then define the joint white noise significance of such a cluster as the probability of observing a cluster with greater joint normalized energy,

$$P(Z'_C > Z_C) = \frac{\Gamma(M, Z_C)}{\Gamma(M)}, \quad (5.84)$$

where $\Gamma(M, Z_C)$ is the upper incomplete Gamma function defined by

$$\Gamma(M, Z_C) = \int_{Z_C}^{\infty} Z_C'^{M-1} \exp(-Z_C') dZ_C'. \quad (5.85)$$

Note that for the case of a single tile, this result reduces to the significance of Equation 5.78. However, while the normalized energy Z of single cluster is monotonically related to significance, the same is not true for clusters of tiles. In general, we must also consider the number of tiles M in the cluster and evaluate the full significance of Equation 5.85.

While the application of a joint significance detection threshold is rather straightforward, the necessary clustering of time-frequency tiles is beyond the scope of the search as implemented in this work. We therefore include it in chapter 8 as a possibility for future analysis. However, it is worth noting that the above result is also applicable to the joint white noise significance of tiles in coincidence between multiple detectors.

5.5.3 Detection

We now consider the effect of stationary white noise on the recovery and detectability of well localized bursts. To do so, we begin with the estimated signal to noise ratio of Equation 5.45 and substitute the expansion of Equation 5.39. As a result, we find that the estimated signal to noise ratio of a single well localized burst in the presence

of additive detector noise is given by the expression

$$\hat{\rho}^2 = \frac{\|h\|^2 + |N(\tau, \phi, Q)|^2 + 2\|h\| |N(\tau, \phi, Q)| \cos \theta}{\langle |N(t, \phi, Q)|^2 \rangle} - 1. \quad (5.86)$$

Here $\|h\|^2$ and $|N(\tau, \phi, Q)|^2$ are the true but unknown energies of the signal and noise in the time-frequency tile which best matches the well localized burst, while $\langle |N(t, \phi, Q)|^2 \rangle$ is the average noise energy observed in time-frequency tiles with the same ϕ and Q , but with times t not equal to the time τ of the burst. In addition, an unknown relative phase θ between the signal and noise gives rise to an interference term in the resulting expression.

We may also define the true signal to noise ratio,

$$\rho^2 = \frac{\|h\|^2}{\langle |N(\tau, \phi, Q)|^2 \rangle}, \quad (5.87)$$

where $\langle |N(\tau, \phi, Q)|^2 \rangle$ is the true but unknown mean energy of the detector noise in the best match time-frequency tile.

From these two expressions, we can identify two unavoidable sources of error that impact our ability to accurately recover the true signal to noise ratio of a well localized burst. The first is due to random fluctuations of the instantaneous detector noise energy $|N(\tau, \phi, Q)|^2$ relative to the true mean detector noise energy $\langle |N(\tau, \phi, Q)|^2 \rangle$. The second is due to the random phase θ between the well localized burst and the instantaneous detector noise within the same time-frequency tile.

Fortunately, the statistical distributions associated with both of these sources of error are readily available. In section 5.5.1, assuming stationary white noise, we demonstrated that the distribution of $|N(\tau, \phi, Q)|^2$ is simply an exponential distribution with mean $\langle |N(\tau, \phi, Q)|^2 \rangle$. In addition, assuming the detector noise is independent of any signal, the distribution of θ is simply a uniform distribution over all angles.

We may also identify two secondary sources of error in estimating the true signal to noise ratio of a well localized burst. The first is due to random fluctuations of the empirical average $\langle |N(t, \phi, Q)|^2 \rangle$ about the true mean detector noise energy $\langle |N(\tau, \phi, Q)|^2 \rangle$. The second is due to the loss in measured signal energy resulting from a mismatch between an arbitrary well localized burst and the nearest measurement tile, which results in a bias towards underestimating the true signal to noise ratio.

Theoretically, however, both of these secondary errors can be made arbitrarily small. In the first case, the central limit theorem predicts that $\langle |N(t, \phi, Q)|^2 \rangle$ is Gaussian distributed about the true mean energy $\langle |N(\tau, \phi, Q)|^2 \rangle$ with a standard deviation that is inversely proportional to the square root of the number of independent time-frequency tiles contributing to the average. Thus, this error can be made negligible by increasing the duration of data used to estimate the mean detector noise energy. In the second case, the distribution depends upon the tiling of the search space as described in section 3.2.2 and can be made negligible by increasing the overlap between time-frequency tiles. Of course, in practice, the maximum data duration used to estimate the mean detector noise energy is limited by stationarity concerns, while the allowable overlap between time-frequency tiles is limited by computational cost. However, even with these constraints, both of these errors can be made sufficiently small that their effect is typically negligible at the signal levels of interest in comparison with the unavoidable errors identified above.

In aggregate, the statistical fluctuations due to these four sources of error then give rise to a distribution of observed signal to noise ratios. Unfortunately, a closed form analytical expression for this distribution is not readily determined. Instead, we must turn to Monte Carlo methods to empirically determine the expected distribution. In doing so, we make the following assumptions. We assume that 100 independent time-frequency tiles are used to compute the estimated average detector noise energy $\langle |N(t, \phi, Q)|^2 \rangle$. We also assume that the energy loss due to measurement mismatch is uniformly distributed between 0 and 10 percent, which roughly corresponds to the empirically observed distribution of Figure 6.1 for the case of a 20 percent worst case

energy loss.

For each value of the true signal to noise ratio, we therefore simulate an ensemble of measurements drawn from an exponential distribution of instantaneous detector noise energies $|N(\tau, \phi, Q)|^2$, a uniform distribution of relative detector noise phase θ , a Gaussian distribution of the estimated mean detector noise energy $\langle |N(t, \phi, Q)|^2 \rangle$, and a uniform distribution of worst case energy loss due to mismatch. Substitution of these simulated values into Equation 5.86 then leads to an empirical distribution for the observed signal to noise ratio $\hat{\rho}^2$ as a function of true signal to noise ratio.

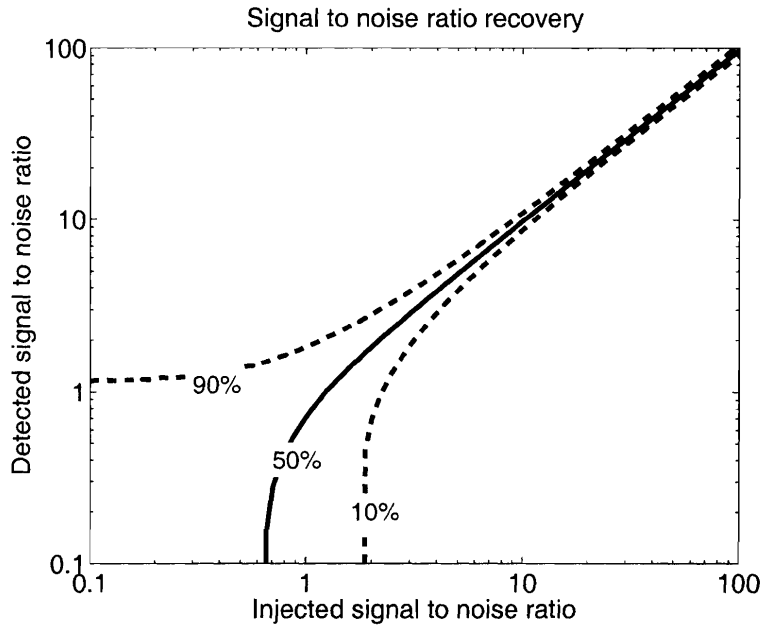


Figure 5.4: The accuracy with which the signal to noise ratio of a single well localized burst can be recovered in the presence of stationary white noise using the prescription of Equation 5.45. The range of observed signal to noise ratios is shown as a function of the theoretical signal to noise ratio given by Equation 5.42. The solid curve corresponds to the median observed value, while the dashed curves correspond to the 10th and 90th percentile values.

The resulting error in recovering the true signal to noise ratio of a well localized burst is then shown in Figure 5.4 in terms of the 10th, 50th and 90th percentile contours of the observed signal to noise ratio distribution. As expected, the contribution of detector noise dominates the resulting distribution at signal to noise ratios on the order of unity or smaller, making it virtually impossible to accurately recover the true signal

to noise ratio of such bursts. At high signal to noise ratios, however, the accuracy in recovering the true signal to noise ratio is limited by the two secondary sources of error identified above. That is, the use of only 100 independent time-frequency tiles to estimate the average detector noise energy and the allowed 20 percent worst case energy loss due to mismatch. In addition, this latter error also results in an approximate 2.5 percent bias towards under-measurement of the injected signal to noise ratio, although such a bias is too small to be observed in Figure 5.4. Finally, we note that the apparent departure of the median and 10th percentile contours towards an observed signal to noise ratio of zero occurs because the instantaneous total energy in a time-frequency tile can occasionally be smaller than the estimated average energy in similar time-frequency tiles. In such cases, the observed signal to noise ratios become imaginary, which we take to indicate an observed signal to noise ratio of zero.

Given the results of our Monte Carlo simulation, we may also determine the effect of stationary white noise on the detectability of well localized bursts. Here we define a detected event as one for which the observed signal to noise ratio of the event exceeds a predetermined threshold. We may then define the detection efficiency of a particular search algorithm as the fraction of events that are detected from a theoretical population of events, all of which have the same true signal to noise ratio.

As a result, given a predefined threshold on the observed signal to noise ratio, we may determine the corresponding white noise detection efficiencies directly from contours such as those in Figure 5.4. To do so, we simply draw a horizontal line at the selected threshold and, for each true signal to noise ratio of interest, we identify the fraction of events that exceed the identified threshold. By following this procedure, we may then construct families of detection efficiency curves for different choices of observed signal to noise ratio threshold. Here, we do so for observed signal to noise ratio thresholds of 3, 4, 5, and 6, and present the resulting white noise detection efficiency curves in Figure 5.5.

Note that, even in the absence of a signal, our definition of detection still results in a




Figure 5.5: Theoretically predicted detection efficiency curves for well localized bursts in the presence of stationary white noise. The detection efficiency at a particular injected signal to noise ratio is determined by identifying the fraction of events that exceed a specified threshold on the observed signal to noise ratio. Here, we have considered observed signal to noise ratio thresholds of 3, 4, 5, and 6.

non-zero detection efficiency due to the false acceptance of events arising from random fluctuations in detector noise. For the special case of stationary white noise, the probability of such false events follows directly from Equation 5.78 and Equation 5.79 and is given by

$$P(\hat{\rho}' > \hat{\rho}) = \exp(-\hat{\rho}^2 - 1). \quad (5.88)$$

As we will see in section 7.5, it is the presence of such false events which limits the sensitivity of a search. Of particular interest, therefore, is not just the probability of such events, but the number occurring within a given observation time. However, estimating the expected number of such events also requires knowledge of the total number of measurements within the given observation time. In general, this depends upon the specifics of a search, but for the Q transform is well approximated by the proper volume of the targeted signal space evaluated using the metric of Equation 3.20. Here, however, we make the additional assumption that our search covers of a sufficiently large signal space such that all of the available information content of the data is

exercised. Assuming stationary white noise data, the total number of measurements in an observation time T is then simply the product $f_s T$, where f_s is the sample frequency of the data under test. As a result, we find that the expected number of false events due to stationary white noise data is given by

$$\langle N(\hat{\rho}, f_s, T) \rangle = \exp(-\hat{\rho}^2 - 1) f_s T. \quad (5.89)$$

In Figure 5.6, we plot the resulting expected number of false events as a function of detection threshold assuming one year of stationary white noise data sampled at 16384 Hz.




Figure 5.6: The expected number of single detector false events in one year of observation time are shown as a function of the observed signal to noise ratio threshold for detection. The expected number of false events is estimated according to Equation 5.89 and assumes maximal use of the available data, which consists of stationary white noise sampled at 16384 Hz. As a result, this represents the maximum stationary white noise false rate for a given detection threshold. Any increase the targeted signal space would have no effect on the resulting number of false events.

Finally, we combine these results to determine the expected detection efficiency for well localized bursts as a function of false event rate. The resulting curves, known in signal detection theory as a receiver operating characteristic, provide a comprehensive description of the performance of a search algorithm for a particular waveform.




Figure 5.7: The theoretically predicted receiver operating characteristic curves for well localized bursts in the presence of stationary white noise. Each curve displays the relationship between detection efficiency and false alarm rate for a particular choice of injected signal to noise ratio. Here, for comparison with the results of chapter 6, we have tested true signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0. Note, however, that these are specified in terms of the matched filter signal to noise ratio defined by Equation 3.32 rather than that of Equation 5.45. In addition, the estimated false rates assume maximal use of the available data, which consists of stationary white noise sampled at 16384 Hz.

Here, for comparison with the results of chapter 6, we determine the expected receiver operating characteristic for well-localized bursts with matched filter signal to noise ratios of 3.0, 3.5, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32. In doing so, we also assume that our search exercises the entire information content of stationary white detector noise that has been sampled at 16384 Hz. The resulting receiver operating characteristic is then plotted in Figure 5.7.

Finally, in Figure 5.7 we have in effect determined the theoretical receiver operating characteristic of a templated matched filter search for minimum uncertainty waveforms of unknown phase in stationary white noise, with the additional assumption that the full information content of the signal is exercised. Consequently, these predictions then form a benchmark for comparison, which will be useful in chapter 6

when we validate our analysis pipeline by experimentally reproducing the theoretical receiver operating characteristic presented here.

5.5.4 Outlier rejection

The above statistical arguments have been developed assuming stationary detector noise. However, typical gravitational-wave detectors are characterized by numerous non-stationarities. In particular, large transients of gravitational-wave or non-gravitational-wave origin could dramatically effect the estimation of mean tile energy.

To avoid this problem, we follow the suggestion of Cadonati[122] and apply the box-plot prescription of Tukey[123] in order to identify and remove outlying values prior to estimation of the mean tile energy. This is accomplished by first sorting the tiles in each frequency band in order of increasing energy. Next, we identify the lower and upper quartile energies, $|X|_{\text{lower}}^2$ and $|X|_{\text{upper}}^2$, as the tile energies corresponding to the 25th and 75th percentile values. The difference between these energies is designated the interquartile range,

$$|X|_{\text{range}}^2 = |X|_{\text{upper}}^2 - |X|_{\text{lower}}^2. \quad (5.90)$$

We then define a threshold value, the so-called whisker of the box-plot, which is a constant multiple of the interquartile range above the upper quartile energy,

$$|X|_{\text{outlier}}^2 = |X|_{\text{upper}}^2 + \alpha |X|_{\text{range}}^2. \quad (5.91)$$

Tiles whose energies exceed this threshold are identified as outliers, and such tiles are then excluded from the subsequent calculation of the mean tile energy. Note that in Tukey's original formulation, α was taken to be 1.5. However, here we allow specification of α as a parameter of the search. In practice, this parameter determines the amount of non-stationary detector noise to incorporate into the estimate of the

background tile energy.

One difficulty with this approach, however, is that it produces a bias in the resulting estimate of the mean tile energy. This is due to the fact that some tile energies from the underlying exponential distribution will also be excluded based on the outlier rejection formalism. Since this only occurs for large tile energies, the resulting estimate of the mean tile energy will exhibit a bias toward lower energies. Fortunately, it is possible to correct for this bias for the case of stationary white noise data by considering its effect on an ideal exponential distribution of tile energies. To do so, we first express the lower and upper quartile energies as the 25th and 75th percentile energies of an exponential distribution with a true mean energy $\langle |X|^2 \rangle$,

$$\int_0^{|X|_{\text{lower}}^2} \frac{1}{\langle |X|^2 \rangle} \exp\left(-\frac{|X|^2}{\langle |X|^2 \rangle}\right) d|X|^2 = \frac{1}{4} \quad (5.92a)$$

$$\int_0^{|X|_{\text{upper}}^2} \frac{1}{\langle |X|^2 \rangle} \exp\left(-\frac{|X|^2}{\langle |X|^2 \rangle}\right) d|X|^2 = \frac{3}{4}. \quad (5.92b)$$

This leads to expressions for the lower and upper quartile energies, the interquartile energy range, and the threshold energy for identifying outliers of

$$|X|_{\text{lower}}^2 = \langle |X|^2 \rangle \ln(4/3) \quad (5.93a)$$

$$|X|_{\text{upper}}^2 = \langle |X|^2 \rangle \ln 4 \quad (5.93b)$$

$$|X|_{\text{range}}^2 = \langle |X|^2 \rangle \ln 3 \quad (5.93c)$$

$$|X|_{\text{outlier}}^2 = \langle |X|^2 \rangle \ln(4 \cdot 3^\alpha). \quad (5.93d)$$

We then determine the expected biased measurement of the mean tile energy by considering the mean of the truncated exponential distribution consisting of all tiles with energy less than the identified threshold, such that

$$\langle |X|^2 \rangle_{\text{biased}} = \frac{\int_0^{|X|_{\text{outlier}}^2} \frac{|X|^2}{\langle |X|^2 \rangle} \exp\left(-\frac{|X|^2}{\langle |X|^2 \rangle}\right) d|X|^2}{\int_0^{|X|_{\text{outlier}}^2} \frac{1}{\langle |X|^2 \rangle} \exp\left(-\frac{|X|^2}{\langle |X|^2 \rangle}\right) d|X|^2}. \quad (5.94)$$

As a result, we find that the required correction factor for estimating the mean tile energy, as a function of the outlier specification factor α , is given by

$$\langle |X|^2 \rangle = \left[1 - \frac{\ln(4 \cdot 3^\alpha)}{4 \cdot 3^\alpha - 1} \right]^{-1} \langle |X|^2 \rangle_{\text{biased}}. \quad (5.95)$$

By excluding outliers and compensating for the resulting bias, this technique then affords a robust unbiased estimate of the mean of the underlying exponential distribution of tile energies, dramatically reducing the sensitivity of our analysis to non-stationarities in the input data stream.

5.6 Identification of events

As we noted in chapter 3, it is possible to select a measurement basis that covers a finite region of time, frequency, and Q space such that the mismatch between an arbitrary well localized burst and the closest basis function results in no more than a specified worst case fractional loss in the measured signal energy. In doing so, we necessarily cover the targeted signal space using multiple Q planes consisting of overlapping highly correlated time-frequency tiles. As a result, a typical event is likely to produce many overlapping significant tiles using the formalism described in section 5.5.

This section introduces a simple exclusion algorithm with the goal of reducing the set of significant tiles by identifying the most significant non-overlapping tiles among all Q planes. This has two benefits. First, by selecting the most significant tiles, we identify the best match set of tiles for the time-frequency structure of candidate events. This then provides the tightest possible bounds on their time-frequency structure, maximizing the observable signal to noise ratio and minimizing the coincident false rate between multiple detectors. Secondly, by dramatically reducing the number of significant tiles, we substantially reduce the subsequent computational cost of testing

for coincidence between multiple detectors.

One obvious approach is to simply exclude the least significant tile from all overlapping pairs of tiles. However, in order to preserve the structure of less localized bursts and to enable the detection of weak bursts in the presence of stronger nearby transients, we must be careful when encountering the situation depicted in Figure 5.8. Here three overlapping tiles are shown with the most significant on top and the least significant on the bottom. In such cases we wish to exclude the intermediary tile since it is overlapped by the most significant tile. However, we do not wish to exclude the least significant tile since it is not overlapped by the most significant tile and the intermediate tile has already been excluded.




Figure 5.8: Example set of overlapping significant tiles used to motivate the exclusion algorithm for identifying the most significant non-overlapping tiles. Here the tiles are sorted by significance with the most significant tile (1) shown on top. In this case, we seek an algorithm which excludes the intermediary tile (2) while preserving the most significant (1) and least significant (3) tiles.

This then suggests the following algorithm. We first sort the tiles in order of decreasing significance. Then, starting with the most significant tile and proceeding to the least significant, we exclude any tile that is overlapped by a more significant tile that has not already been excluded. Thus, for each well localized event, only the single tile which best represents the event’s parameters is reported.

We conclude by noting that the approach we have described does not exclude the possibility of clustering tiles from poorly localized bursts. In particular, the tiles which pass both the initial significance threshold as well as the subsequent exclusion algorithm will best represent the strong localized features of such bursts. This then suggests the possibility of testing the joint significance of all tiles within a predefined time-frequency window in order to more optimally detect such poorly localized bursts. This approach, however, is beyond the planned scope of this work and is left for

subsequent investigation.

5.7 The Q pipeline

At this point, we have developed the necessary tools in order to construct a comprehensive end-to-end analysis pipeline for the detection of gravitational-wave bursts in data from a single interferometric detector. This analysis pipeline consists of whitening by zero-phase linear prediction, application of the discrete Q transform, thresholding on the white noise significance of Q transform coefficients, and identification of the most significant set of non-overlapping time-frequency tiles. In addition, we also incorporate a final stage that excludes all but the most significant time-frequency tile within a specified time window in order to prevent the redundant reporting of candidate events. The resulting pipeline is shown in Figure 5.9.

While this single detector pipeline is designed for the detection of gravitational-wave bursts, it is important to note that it is also equally applicable to the detection of statistically significant events in data from environmental monitors and auxiliary detector channels. This then permits an important test for environmental disturbances and pathological detector behavior that could otherwise be mistaken for a gravitational-wave burst. As an example of the utility of this approach, in section 7.1.2 we successfully apply it to microphone data during the second LIGO science run in order to exclude events that were due to airplane overflights of the Hanford observatory.

The computational cost of the single detector Q pipeline depends primarily on the density of tiles necessary to ensure that the requested worst case energy loss is not violated, but also depends on the size of the targeted signal space as well as the sample frequency of the input data stream. In Figure 5.11, we display the estimated number of floating point operations required to analyze a typical block of interferometric detector data as a function of the requested worst case energy loss and the Q range of




Figure 5.9: The proposed end-to-end analysis pipeline for the identification of candidate gravitational-wave bursts in data from a single interferometric detector. The same analysis pipeline is also applicable to the detection of statistically significant events in both environmental and auxiliary detector channels. This permits a test for environmental disturbances or pathological detector behavior in order to exclude events of non-gravitational-wave origin.

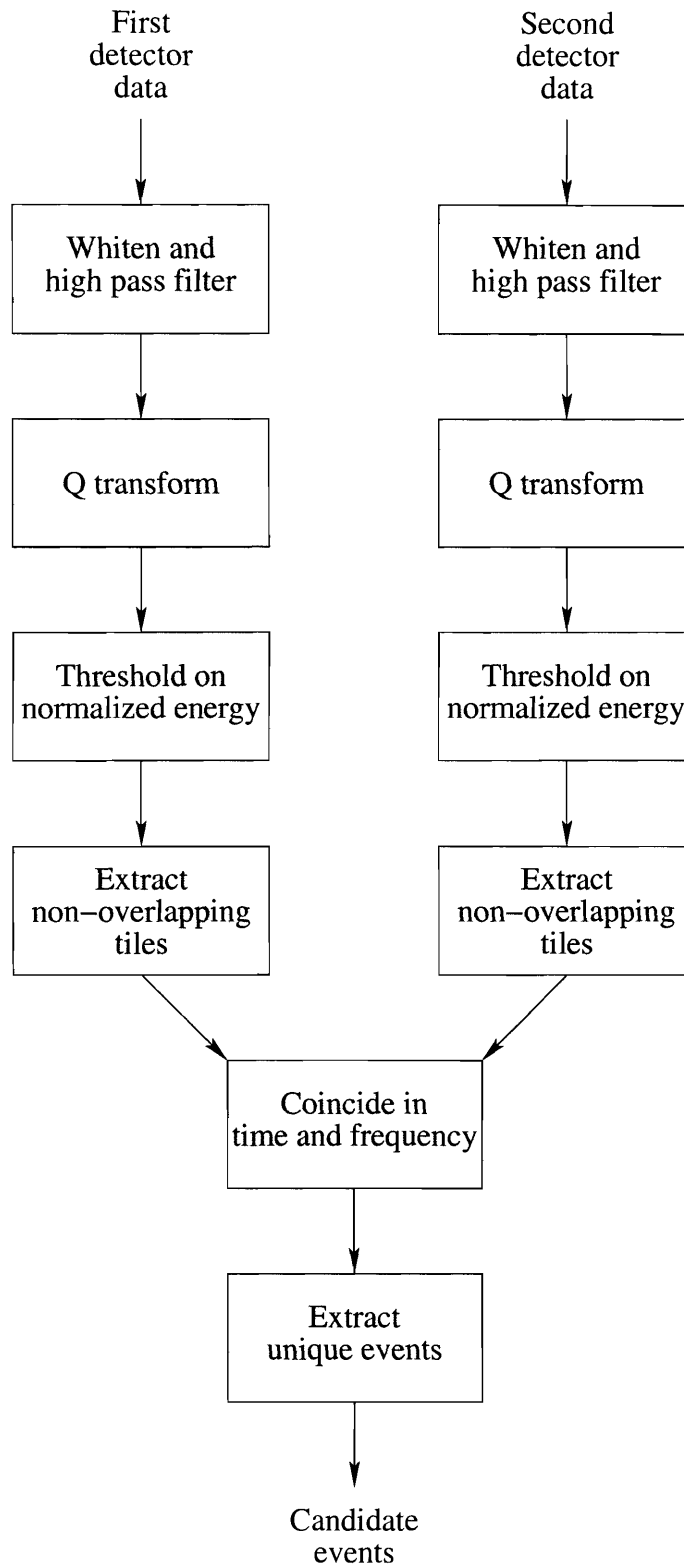


Figure 5.10: The proposed end-to-end analysis pipeline for the identification of coincident candidate gravitational-wave bursts in data from multiple interferometric detectors.

the search. In this case we have assumed that the data to be analyzed is 64 seconds in duration and sampled at 16384 Hz. The targeted signal space is assumed to extend from 64 to 1024 Hz, but the targeted Q range is allowed to vary from a minimum value of 4 to a maximum value of 4, 8, 16, 32, or 64.

As we noted in section 5.2.2, the computational cost of the fast form of the discrete Q transform is dominated by an initial fast Fourier transform of the entire data block, followed by multiple shorter length inverse fast Fourier transforms for each central frequency and bandwidth pair of interest. It is the number and size of these inverse transforms which in many cases determines the overall computational cost of the search. This is typically the case at high Q , where a larger number of inverse transforms are required to cover the time-frequency plane. This is also the case at high frequency, since the length of each inverse transform is proportional to its bandwidth, which is in turn proportional to frequency for fixed Q . As a result, the computational cost increases notably with both the maximum frequency and Q of the analysis. This increase is also compounded by the need to oversample the targeted signal space in order to achieve a particular worst case energy loss due to mismatch.

For relatively large values of the worst case energy loss, the resulting computational cost is instead dominated by the cost of zero-phase linear predictive whitening of the input data stream. As was the case for the fast discrete Q transform, the computational cost of linear predictive whitening is partly due to fast Fourier transforms of the entire data block. Such transforms are necessary to compute the autocorrelation of the input data stream in order to train the filter, as well as to apply the resulting filter in the frequency domain. In fact, some advantage is gained by reusing the results of intermediate transforms in the discrete Q transform. The computational cost of linear predictive whitening is also due to the cost of determining the required filter coefficients using Levinson-Durbin recursion. Since this scales as the square of the filter order, this cost of training becomes important for sufficiently long filters, which in turn depends upon the maximum Q and minimum frequency of the search. The cost of whitening, as well as its dependence on the maximum Q of the search is

evident in Figure 5.11 in the limit of larger worst case energy loss.

Finally, we note that the estimated computational cost of Figure 5.11 is only approximate has been computed assuming $N \ln N$ floating point operations for each fast Fourier transform of length N and M^2 operations for Levinson-Durbin recursion of order M . For reference, the operating point corresponding to the single detector version of the search described in section 7.2 is also noted in the figure.

Up to this point, we have only considered the detection of gravitational-wave bursts in data from a single detector. However the simultaneous observation using multiple detectors offers the possibility of a dramatically reduced false event rate due to environmental disturbances, pathological detector behavior, or statistical fluctuations in detector noise. A number of techniques have been proposed to search for unmodeled bursts using networks of interferometric detectors[93, 97, 103, 124, 125, 126]. However, the most straightforward extension of our single detector analysis to a network of multiple detectors consists of simply testing for time-frequency coincidence between the candidate events identified in each of the detectors. Such an analysis pipeline is shown in Figure 5.10 for the case of two detectors. In this case, however, care must be taken to account for the different detector response of geographically separated detectors. In particular, the speed of light travel time between detectors must be taken into considered when testing for coincidence. The computational cost of such a combination scales linearly with the number of detectors since the time-frequency coincidence of candidate single detector events is computationally negligible for typical event rates.

5.8 The coherent Q pipeline

In this section, we consider an alternative multiple detector analysis pipeline that is appropriate to the special case of collocated interferometric detectors, such as the two

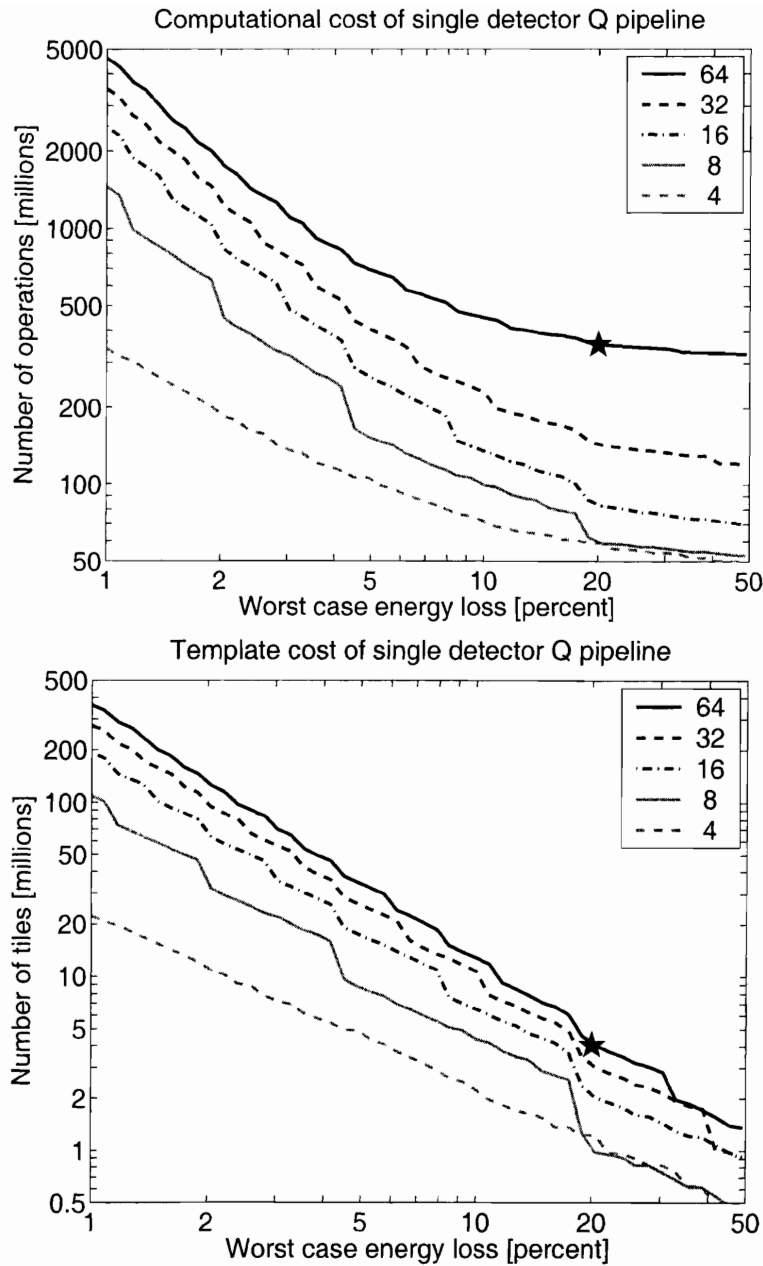


Figure 5.11: The estimated computational and template cost of a single detector Q pipeline as a function of the requested worst case energy loss due to mismatch and the Q range of the search. The estimate includes the cost of both whitening and Q transform analysis for a 64 second block of data that has been sampled at 16384 Hz. The frequency range of the search extends from 64 to 1024 Hz and Q range extends from 4 to maximum values of 4, 8, 16, 32, and 64. At large fraction energy loss, the computation cost is dominated by the cost of whitening. At small fractional energy loss, the increased computational cost is due to the high oversampling of the targeted signal space necessary to ensure that the requested worst case energy loss is not exceeded. The operating point for the search described in section 7.2 is also indicated.

LIGO Hanford detectors. Specifically, we develop a method to coherently combine the Q transform coefficients from collocated detectors in a way that maximizes the expected signal to noise ratio of well localized bursts. In addition, since collocated detectors necessarily have the same response to gravitational-wave bursts, we take advantage of the expected signal consistency to provide a more sensitive test for gravitational-wave bursts while excluding statistically significant events that exhibit inconsistent behavior.

We begin by considering the Q transform of a single gravitational-wave burst as observed by N detectors. The Q transform coefficient in a specific time-frequency tile of the n th detector is then

$$X_n(\tau, \phi, Q) = h(\tau, \phi, Q) + N_n(\tau, \phi, Q), \quad (5.96)$$

where $h(\tau, \phi, Q)$ is the true but unknown Q transform of the signal in this tile and $N_n(\tau, \phi, Q)$ is the true but unknown Q transform of the noise in the n th detector.

We then define the coherent Q transform $X^{(N)}(\tau, \phi, Q)$ as the weighted average of the Q transform coefficients from each of the N individual detectors,

$$X^{(N)}(\tau, \phi, Q) = \sum_{n=1}^N C_n(\tau, \phi, Q) X_n(\tau, \phi, Q), \quad (5.97)$$

where the real-valued coefficients $C_n(\tau, \phi, Q)$ obey the normalization condition

$$\sum_{n=1}^N C_n(\tau, \phi, Q) = 1. \quad (5.98)$$

If we also assume that the noise in each detector is due to an independent stationary stochastic process, then the mean squared magnitude of the coherent Q transform is simply

$$\langle |X^{(N)}(\tau, \phi, Q)|^2 \rangle = |h(\tau, \phi, Q)|^2 + \langle |N^{(N)}(\tau, \phi, Q)|^2 \rangle, \quad (5.99)$$

where true mean noise energy is given by

$$\langle |N^{(N)}(\tau, \phi, Q)|^2 \rangle = \sum_{n=1}^N C_n^2(\tau, \phi, Q) \langle |N_n(\tau, \phi, Q)|^2 \rangle. \quad (5.100)$$

We therefore define the true coherent signal to noise ratio

$$\rho^{(N)} = \left[\frac{|h(\tau, \phi, Q)|^2}{\langle |N^{(N)}(\tau, \phi, Q)|^2 \rangle} \right]^{1/2}, \quad (5.101)$$

and seek the coefficients $C_n(\tau, \phi, Q)$ that maximize this quantity subject to the constraint of Equation 5.98. Straightforward maximization using the method of Lagrange multipliers then leads to the coefficients

$$C_n(\tau, \phi, Q) = \left(\sum_{m=1}^N \frac{1}{\langle |N_m(\tau, \phi, Q)|^2 \rangle} \right)^{-1} \frac{1}{\langle |N_n(\tau, \phi, Q)|^2 \rangle}, \quad (5.102)$$

such that the coherent Q transform is given by,

$$X^{(N)}(\tau, \phi, Q) = \left(\sum_{n=1}^N \frac{1}{\langle |N_n(\tau, \phi, Q)|^2 \rangle} \right)^{-1} \sum_{n=1}^N \frac{X_n(\tau, \phi, Q)}{\langle |N_n(\tau, \phi, Q)|^2 \rangle}. \quad (5.103)$$

If, following Equation 5.36, we then interpret the quantity $\langle |N_n(\tau, \phi, Q)|^2 \rangle$ as a measure of the average power spectral density over the bandwidth of the tile, we see that the coherent Q transform is simply the average Q transform weighted by the power spectral density of the noise in each detector. As expected, this then implies that more sensitive detectors preferentially contribute to the resulting coherent transform. Furthermore, we also note that the result of Equation 5.102 allows for a unique set of weighting factors in each time-frequency tile. As a result, the coherent Q transform naturally accounts for any potential frequency dependent variation in the relative sensitivity between detectors. Finally, we note that for the ideal case of N identical detectors, the resulting coherent signal to noise ratio is a factor of \sqrt{N} greater than the signal to noise ratio in any single detector.

In practice, however, we have no knowledge of the true mean noise energy in each detector. Instead, assuming stationary detector noise, we estimate the mean noise energy in each detector by

$$\langle |N_n(\tau, \phi, Q)|^2 \rangle \sim \langle |X_n(t, \phi, Q)|^2 \rangle, \quad (5.104)$$

where the average on the right hand side is the mean of a representative set of time-frequency tiles, all of which have the same frequency and Q as the tile under consideration, but occur at times t sufficiently far removed from the time τ of the burst that they measure only detector noise. Similarly, we estimate the optimal detector weighting by

$$C_n(\tau, \phi, Q) \sim \left(\sum_{m=1}^N \frac{1}{\langle |X_m(t, \phi, Q)|^2 \rangle} \right)^{-1} \frac{1}{\langle |X_n(t, \phi, Q)|^2 \rangle}, \quad (5.105)$$

such that the coherent Q transform becomes,

$$X^{(N)}(\tau, \phi, Q) \sim \left(\sum_{n=1}^N \frac{1}{\langle |X_n(t, \phi, Q)|^2 \rangle} \right)^{-1} \sum_{n=1}^N \frac{X_n(\tau, \phi, Q)}{\langle |X_n(t, \phi, Q)|^2 \rangle}. \quad (5.106)$$

Lastly, we may define a corresponding coherent normalized energy by

$$Z^{(N)} = \frac{|X^{(N)}(\tau, \phi, Q)|^2}{\langle |X^{(N)}(t, \phi, Q)|^2 \rangle}, \quad (5.107)$$

where the expectation value in the denominator is again an empirical average over time-frequency tiles with the same frequency and Q , but at times t sufficiently far removed from the time τ of the burst.

In order to determine the statistical significance of coherent Q transform coefficients, we recall that in section 5.5 we demonstrated that the real and imaginary components of individual Q transform coefficients are Gaussian distributed for the special case of stationary white noise data. According to the central limit theorem, it then follows

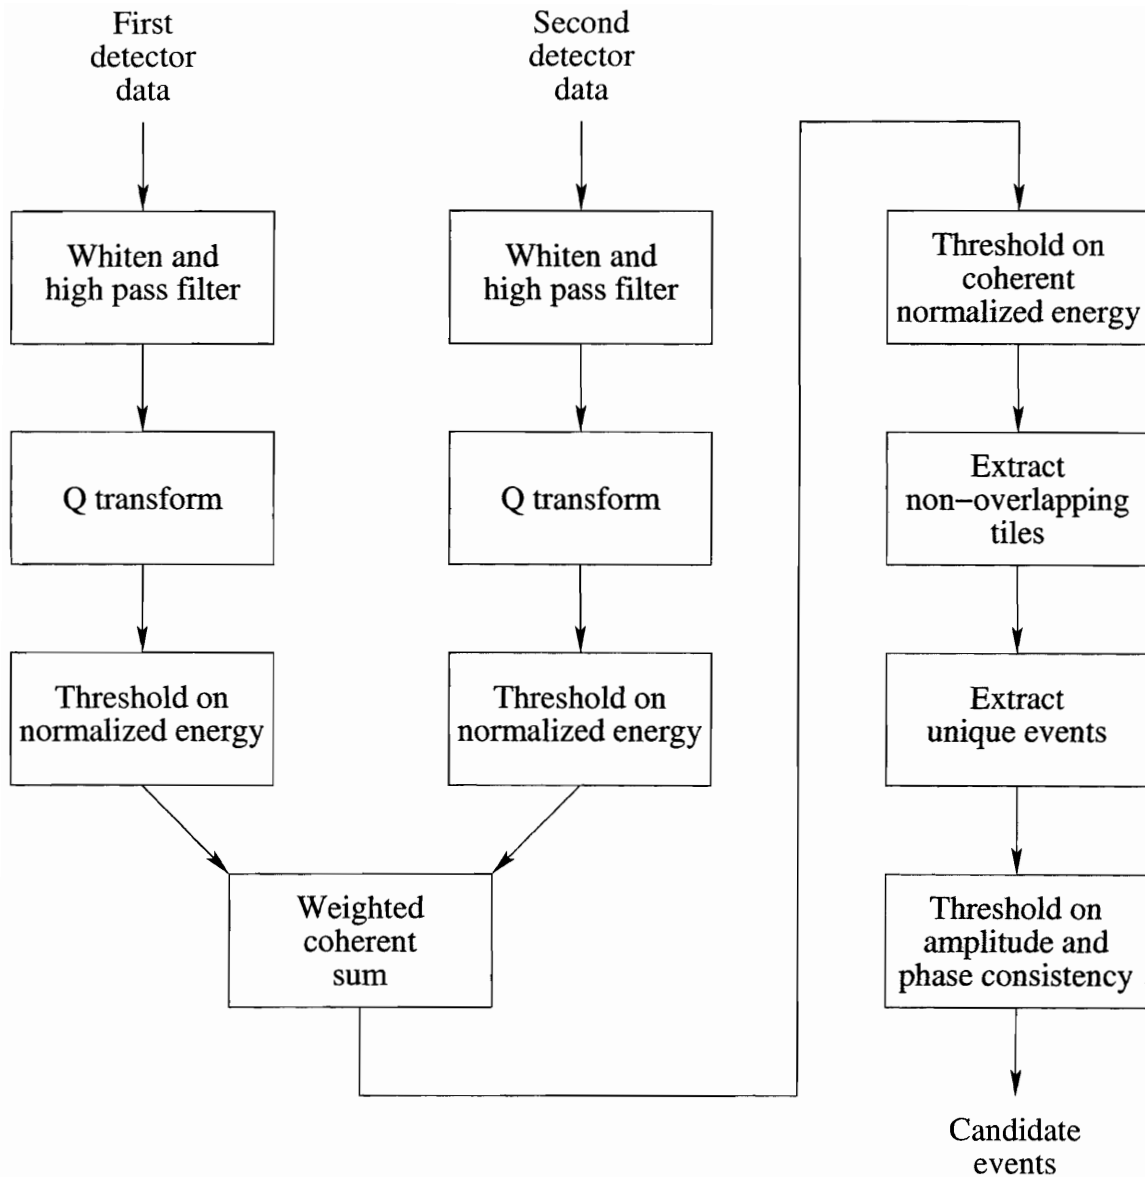


Figure 5.12: The proposed coherent end-to-end analysis pipeline for the identification of candidate gravitational-wave bursts in data from multiple collocated interferometric detectors. The same analysis pipeline is also applicable to a directional search for gravitational-wave bursts using networks of non-collocated detectors. In this case, the data are first scaled and shifted in time such that they effect a virtual network of collocated detectors for gravitational-wave bursts originating from a specific position on the sky.

that the real and imaginary components of any linear combination of such transforms are also Gaussian distributed. As a result, for independent stationary white noise, we find that the coherent Q transform obeys the same statistics that apply to individual Q transform coefficients. In particular, we note that the measured coherent normalized

energy of Equation 5.107 is also an exponentially distributed random variable.

Lastly, we note that, in addition to testing the statistical significance of coherent Q transform coefficients, the expected agreement between colocated detectors also permits simple tests of both amplitude and phase consistency between detectors. This is particularly useful since the presence of a strong signal in even a single detector will nevertheless appear as a statistically significant signal in the coherent Q transform. However, by testing for amplitude and phase consistency between the Q transform coefficients observed in multiple detectors, we may exclude such false events from our search. By combining the coherent Q transform with such amplitude and phase consistency tests, we can then produce a coherent end-to-end analysis pipeline for the identification of gravitational-wave bursts in data from multiple colocated detectors. The resulting coherent analysis pipeline is shown in Figure 5.12, and is also the pipeline that we apply in chapter 7 to the search for gravitational-wave bursts in data from the two colocated LIGO Hanford detectors. Finally, we point out that the coherent analysis pipeline is also applicable to a directional search for gravitational-wave bursts using networks of non-colocated detectors. In this case, the data must first be scaled and shifted in time such that they emulate data from a virtual network of colocated detectors for bursts originating from a specific position on the sky. However, such an approach is beyond the scope of this work and is left for future investigation.

5.9 Example

In order to gain a more intuitive understanding of the Q transform, we conclude this chapter by considering a simple example of the Q transform applied to a simulated burst of gravitational radiation. In particular, we consider a sinusoidal Gaussian burst

of the form

$$h(t; \tau, \phi, Q) = h_0 \exp \left[-4\pi^2 \phi^2 (t - \tau)^2 / Q^2 \right] \sin [2\pi \phi (t - \tau)], \quad (5.108)$$

with a central frequency of 256 Hz and a Q of 8.

This signal is then added to 32 seconds of simulated LIGO detector noise as described in section 6.3. Moreover, for clarity, this signal is injected with the reasonably large matched filter signal to noise ratio of 10, corresponding to a normalized energy of 51. Note that, except for a smaller signal to noise ratio, this is the same signal used to demonstrate the effectiveness of zero-phase linear predictive whitening in the example of section 4.7. Here, however, in addition to whitening by linear prediction, we also apply the Q transform and the exclusion algorithm of the previous section in order to identify the best match set of non-overlapping time-frequency tiles for our example burst.

We first apply the Q transform to the injected signal. Moreover, in order to understand the choices involved, we do so for three different logarithmically spaced values of Q and for two different values of the worst case energy loss due to mismatch. Specifically, we test Q s of 4, 8, and 16, and worst case energy losses of 1 percent and 20 percent. The resulting normalized energy spectrograms are then presented in Figure 5.13a.

We first note that a Q of 8 exactly corresponds to the Q of the injected signal. As a result, those spectrograms computed with this value of Q exhibit a greater normalized energy in their best match time-frequency tile than those spectrograms computed using other values of Q . In addition, the finer sampling of the time-frequency plane necessary to ensure no more than 1 percent energy loss is clearly evident when compared with the coarser sampling of the time frequency plane required for a 20 percent worst case energy loss. It is important to note, however, that this finer sampling does not correspond to an improved time-frequency resolution. Each tile still repre-

sents a region of minimum time-frequency uncertainty. Instead, the finer sampling corresponds to a greater overlap between adjacent time-frequency tiles. In fact, for the studies conducted in this work, we typically choose to tile the time-frequency plane for a less computationally intensive 20 percent worst case energy loss, and only consider a finer tiling when retrospectively examining interesting events.

Next, we demonstrate the ability of the exclusion algorithm of the previous section to identify the best match set of non-overlapping time-frequency tiles for our example burst. In this case, we have first applied the Q transform to analyze the frequencies between 64 and 1024 Hz and Q s between 4 and 64 with a tiling sufficient to ensure a 20 percent worst case energy loss due to mismatch. Significant tiles are then identified as those tiles that exhibit a normalized energy greater than 19, corresponding to a matched filter signal to noise ratio threshold of 6. We then apply the exclusion algorithm of the previous section. The resulting significant tiles, from both before and after application of the proposed exclusion algorithm, are shown in Figure 5.13b.

As expected, due to the minimum uncertainty nature of our example burst, application of the exclusion algorithm results in a single best match time-frequency tile. In general, however, a less localized burst will most likely result in multiple non-overlapping significant tiles. Finally, we also note that the resulting best match time-frequency tile shows very good agreement with the injected signal. In particular, the central time of the example burst is recovered to within 0.4 milliseconds, the frequency is recovered exactly, the Q is recovered to within a factor of $\sqrt{2}$, and the observed normalized energy represents a 9 percent lost — well within the specified worst case loss of 20 percent.

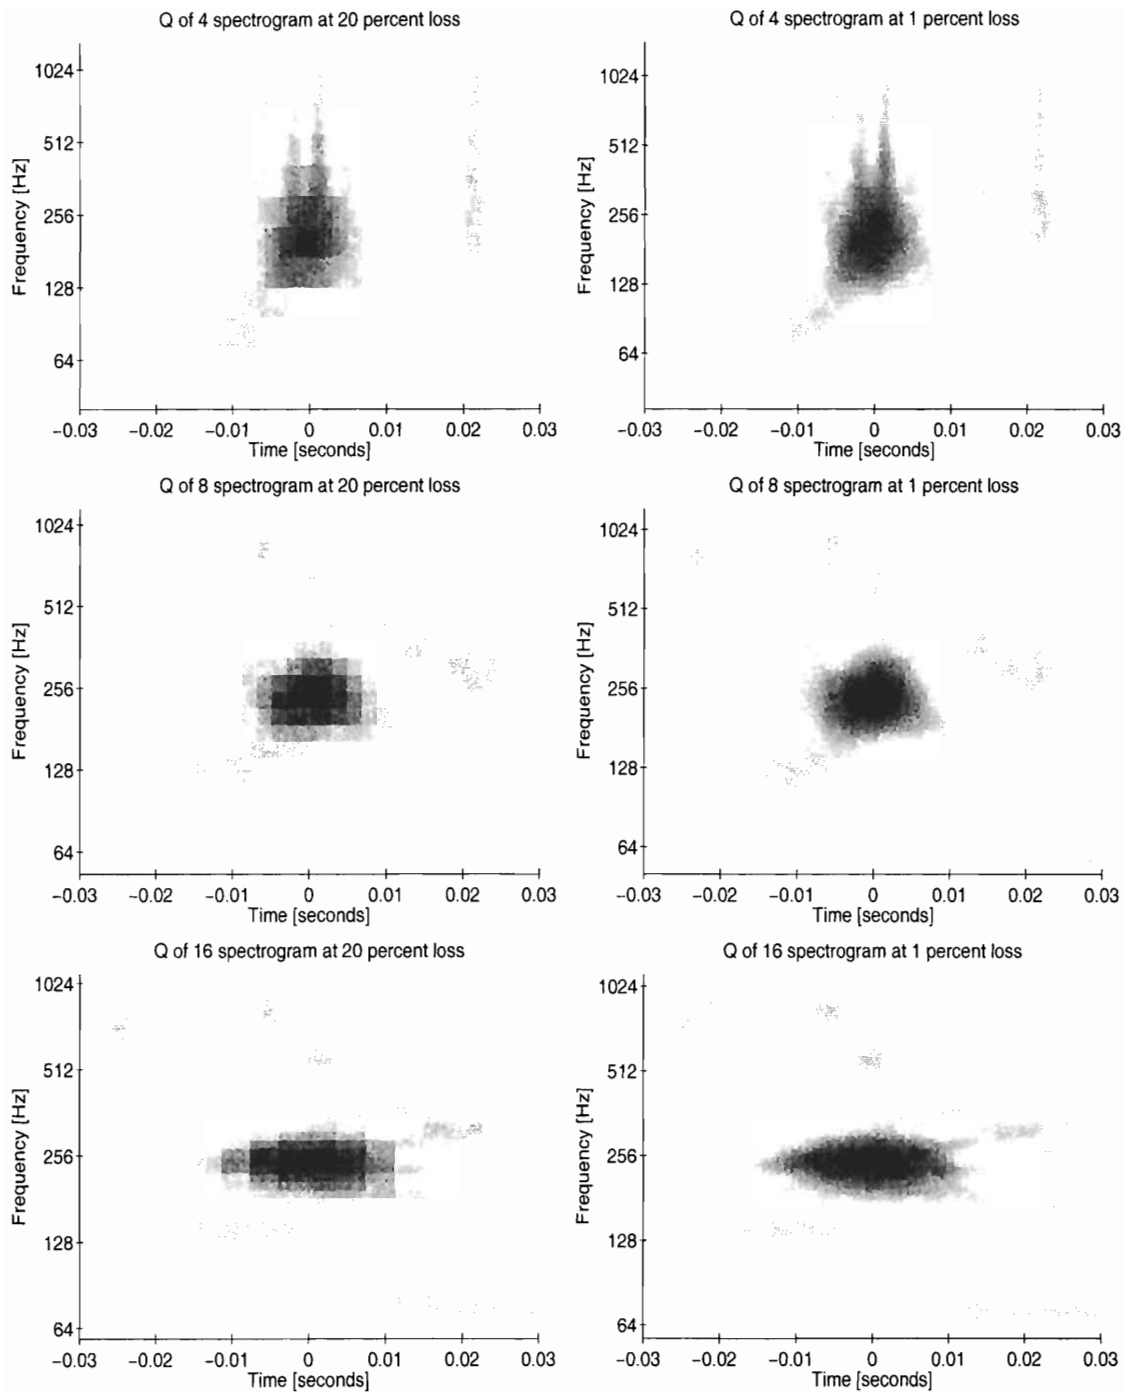


Figure 5.13a: Six different constant Q time-frequency spectrograms of the same simulated gravitational-wave burst. Each spectrogram displays the normalized energy as observed in tiles with Q s of 4 (top), 8 (middle), and 16 (bottom). For a typical search, the signal space is tiled with a worst case energy loss due to mismatch of 20 percent (left). However, for clarity, the same spectrograms are shown tiled at a computationally more intensive 1 percent worst case energy loss due to mismatch (right). In this case, the signal is a 256 Hz Gaussian wave packet with a Q of 8 injected into simulated LIGO detector noise with an optimal matched filter signal to noise ratio of 10.




Figure 5.13b: Significant tiles corresponding to the simulated gravitational-wave burst of Figure 5.13a before (left) and after (right) excluding the least significant of overlapping tiles. Significant tiles were identified at an initial normalized energy threshold of 19.

Chapter 6

Simulation

In this chapter, we present a validation of the proposed single detector pipeline by comparing its performance on simulated gravitational-wave data with the theoretically predicted performance of the Q transform from section 5.5. In doing so, our goal is not to evaluate the performance of the proposed analysis pipeline for arbitrary gravitational-wave bursts or non-ideal detector noise. Such a study is postponed until chapter 7, where the proposed pipeline is applied to the search for astrophysically unmodeled bursts of gravitational radiation in data from the second LIGO science run. Here, we instead seek to validate the single detector implementation of the pipeline used in chapter 7 by applying it to a number of simple test cases for which the expected performance can readily be predicted from first principles arguments.

We first consider the simple case of well localized bursts in the absence of any background noise. This permits us to validate the tiling of the space of time, frequency, and Q proposed in section 3.2.2 by demonstrating that the energy loss incurred by an arbitrary well localized burst within the targeted signal space never exceeds the specified worst case value.

Next, we consider the case of well localized bursts in the presence of stationary white

noise. In this case, the expected response of our pipeline is well known, permitting a simple test of its performance. In particular, we determine the accuracy with which the signal to noise ratio of well localized bursts are recovered and demonstrate that the results are in good agreement with the theoretical performance predicted in Figure 5.4. In the process, we also determine the distribution of the error incurred in recovering the time, frequency, and Q of such well localized bursts in order to characterize the accuracy with which such properties are measured.

We then consider the performance of the proposed analysis pipeline on stationary white noise data in the absence of any bursts and demonstrate that the observed false event rate is in reasonable agreement with prediction of Equation 5.89, confirming that the targeted signal space is sufficiently large that the full information content of the data is exercised. In addition, we also histogram the properties of such false events in order to verify uniform coverage of the targeted signal space.

Finally, we extend our validation study to also consider stationary colored noise. Specifically, we evaluate the performance of our analysis pipeline when applied to simulated detector noise with an amplitude spectrum that corresponds to the design sensitivity of the first generation of LIGO detectors. Such a study permits a more thorough validation of our analysis pipeline by also testing its ability to sufficiently whiten colored data prior to Q transform analysis, justifying our subsequent assumption of stationary white noise when interpreting results. By measuring the detection efficiency of this search for a number of different sinusoidal Gaussian bursts, we are then able to construct a family of receiver operating characteristics that describe the sensitivity of our pipeline over the targeted signal space and that demonstrate good agreement with the theoretical predictions of Figure 5.7.

6.1 Well localized bursts

In this section, we present a validation of the proposed tiling of section 3.2.2, in which the space of time, frequency, Q is tiled such that an arbitrary well localized burst incurs no more than a specified worst case energy loss as a result of mismatch with the nearest measurement tile. To do so, we consider the response of the Q transform to sinusoidal Gaussian bursts of the form

$$h(t; \tau, \phi, Q) = h_0 \exp \left[-\frac{4\pi^2 \phi^2 (t - \tau)^2}{Q^2} \right] \sin [2\pi \phi (t - \tau) + \theta], \quad (6.1)$$

where the peak amplitude h_0 of the Gaussian envelope is related to the characteristic signal energy $\|h\|^2$ by the expression

$$\|h\|^2 = h_0 \left(\frac{Q^2}{32\pi\phi^2} \right)^{1/2} \left[1 - \cos(2\theta) \exp \left(-\frac{Q^2}{2} \right) \right]. \quad (6.2)$$

Here, in order to measure the energy loss incurred by such bursts, we fix this characteristic signal energy to be unity. We then choose to test many such bursts, with each burst occurring at a random time τ once every 64 seconds, a logarithmically distributed random frequency ϕ between 64 and 1024 Hz, a logarithmically distributed random Q between 4 and 64, and a uniformly distributed random phase θ between 0 and 2π . We then apply the Q transform to identify bursts within this same signal space, which we tile for the specific worst case energy loss under test. Finally, for each 64 second data block, we identify the Q transform coefficient with the greatest squared magnitude as the best match measurement of the injected signal and determine the corresponding energy loss due to any mismatch.

In Figure 6.1 we display the cumulative distributions of the fractional energy loss observed for three different tilings of the targeted signal space that correspond to requested worst case energy losses of 10, 20, and 40 percent. Here, each distribution represents the measured energies of 40 thousand different bursts covering the targeted signal space.




Figure 6.1: The cumulative distribution of the fractional energy loss incurred by arbitrary well localized bursts in the targeted signal space as a result of mismatch with the nearest measurement tile. The resulting distributions are clearly consistent with the requested worst case fractional energy losses of 10 (dashed), 20 (solid), and 40 (dash-dot) percent. In all cases, this worst case loss is never exceeded in 10^4 measurements.

It is immediately clear from Figure 6.1 that the requested worst case energy loss is never violated, even given the large number of bursts considered here. In fact, the overwhelming majority of well localized bursts incur a fractional energy loss which is significantly smaller than the specified worst case value. In particular, for the 20 percent worst case energy loss that we apply in the remainder of this work, over 90 percent of well localized bursts incur an energy loss that is less than 10 percent. Finally, we note that while the fractional energy loss metric in section 3.2.2 was developed for the specific case of mismatched sinusoidal Gaussian bursts, it appears to be equally successful at predicting the energy loss due mismatch between sinusoidal Gaussian bursts and the frequency domain bisquare waveforms which comprise the Q transform.

6.2 Stationary white noise

We now consider the special case of stationary white noise data that has been sampled at the same 16384 Hz rate produced by the LIGO detectors. In the following tests, the Q transform is applied to 64 second segments of such data in order to identify bursts with central frequencies between 64 to 1024 Hz and Q s between 4 to 64. In addition, this targeted signal space is tiled for a 20 percent worst case energy loss due to mismatch between an arbitrary well localized burst and the nearest measurement tile. For reference, we note that these are the same search parameters that are applied in the next section to the search for bursts in simulated detector noise, as well as in chapter 7 to the search for bursts in data from the second LIGO science run.

Into this data, we then inject our standard test signal, the sinusoidal Gaussian bursts described by Equation 6.1. Like before, we inject these bursts with random time, frequency, phase, and Q within the targeted signal space. In this case, however, we also choose to inject such bursts with a logarithmically distributed random signal to noise ratio between 0.1 and 100 as defined by Equation 5.87. We then determine the observed signal to noise ratio for each injected burst by applying the formalism of Equation 5.76 and Equation 5.79 to the expected best match time-frequency tile. In this case, since we are not concerned with detection of such bursts, we simply determine this best match time-frequency tile from *a priori* knowledge of the parameters of the injected burst. Finally, we plot the resulting observed signal to noise ratios against the corresponding injected signal to noise ratios.

The scatter plot resulting from the application of this test to 1000 well localized bursts is shown in Figure 6.2 and represents the accuracy with which the signal to noise ratio of such well localized bursts can be recovered in the presence of stationary white noise. Finally, we note that this is the same relationship that we derived from first principle arguments in section 5.5 and it is evident that the resulting scatter plot shows good agreement with the theoretical predictions presented in Figure 5.4.




Figure 6.2: The relationship between injected and observed signal to noise ratio for 1000 well localized bursts injected into stationary white noise. This scatter plot represents the accuracy with which such bursts can be recovered and shows good agreement with the theoretically predicted distribution shown in Figure 5.4. Note that, for small values of injected signal to noise ratio, the energy in the best match time-frequency tile may occasionally be less than average energy of similar tiles. As a result, the observed signal to noise ratio given by Equation 5.79 can become imaginary. Here, such cases have been included by plotting them along the bottom axis of the figure.

Using a similar infrastructure, we are also able to study the accuracy with which other parameters of well localized bursts are recovered in the presence of stationary white noise. To do so, we again inject such bursts with random time, frequency, phase, and Q . In this case, however, we only inject these signals at a single fixed signal to noise ratio of 10 according to the definition of Equation 5.87. The injection of such reasonably strong bursts then permits us to easily identify the corresponding best match time-frequency tile as the one with the greatest squared Q transform magnitude in each 64 second block of data. We then compare the time, frequency, duration, bandwidth, Q , and signal to noise ratio of this best match time-frequency tile with the true properties of the injected burst in order to determine the absolute and relative errors incurred in the recovery of these parameters. Note that, here we define the duration and bandwidth of well localized bursts in accordance with

Equation 3.13a and Equation 3.13b such that their product is unity. This procedure is then repeated for many such bursts in order to determine the average accuracy with which such properties are recovered and to identify any potential measurement bias. In this case, we have done so for 3000 different bursts, and histograms of the resulting absolute and relative errors are shown in Figure 6.3a and Figure 6.3b, along with their mean and standard deviation.

Here we have defined the absolute errors as simply the differences between the observed and injected values of each parameter. Such errors are particularly relevant to the choice of thresholds used to test for time-frequency coincidence between detectors, as well as for thresholds used for other possible consistency tests. In particular, it is interesting to note that the standard deviation of the observed timing error is approximately 3.5 milliseconds, which is sufficient to allow reasonably tight temporal coincidence tests between detectors in order to exclude coincident bursts of non-gravitational-wave origin that are not consistent with the expected 10 millisecond speed of light travel time between LIGO sites. On the other hand, such a timing resolution is only sufficient to provide very coarse sky position information, and not sufficient to permit more than the simplest of consistency tests.

However, it is important to note that the absolute errors reported in Figure 6.3a reflect the aggregate error incurred over the entire targeted signal space, but that these errors may also vary considerably across this space. For example, we note that the timing resolution of our search is expected to be worst for those bursts with low frequency and high Q , which exhibit the longest durations, and best for those bursts with high frequency and low Q , which exhibit the shortest durations.

As a result, we also define dimensionless relative errors that take into account the expected spread in measurements due to the properties of the injected bursts. Specifically, we define the relative timing error as the ratio of the absolute timing error to the true duration of the injected burst. Similarly, we define the relative frequency error as the ratio of the absolute frequency error to the true frequency of the injected

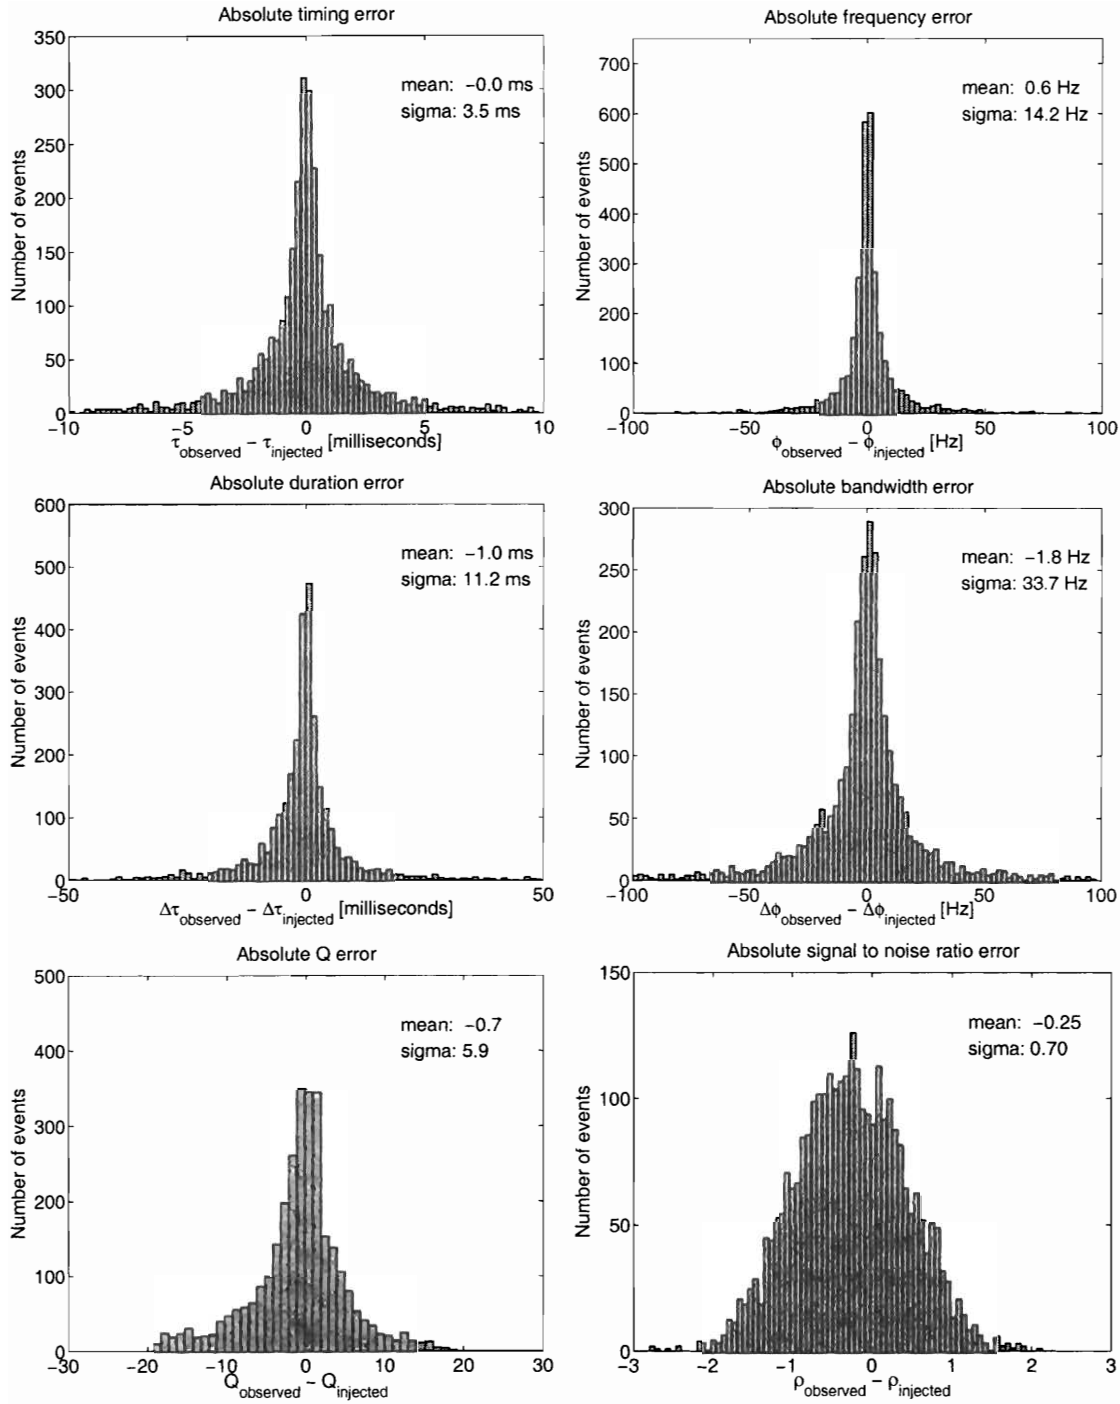


Figure 6.3a: Histograms of the absolute accuracy with which the time, frequency, duration, bandwidth, Q , and signal to noise ratio are recovered for well localized bursts in the presence of stationary white noise. In this case 3000 bursts were injected with a fixed signal to noise ratio of 10 according the definition of Equation 5.87, but with random time, frequency, phase, and Q within the targeted signal space. The absolute error is simply the difference between the observed and injected values for each property. Here we have reported the duration and bandwidth of well localized bursts as defined by Equation 3.13a and Equation 3.13a such that their product is unity.

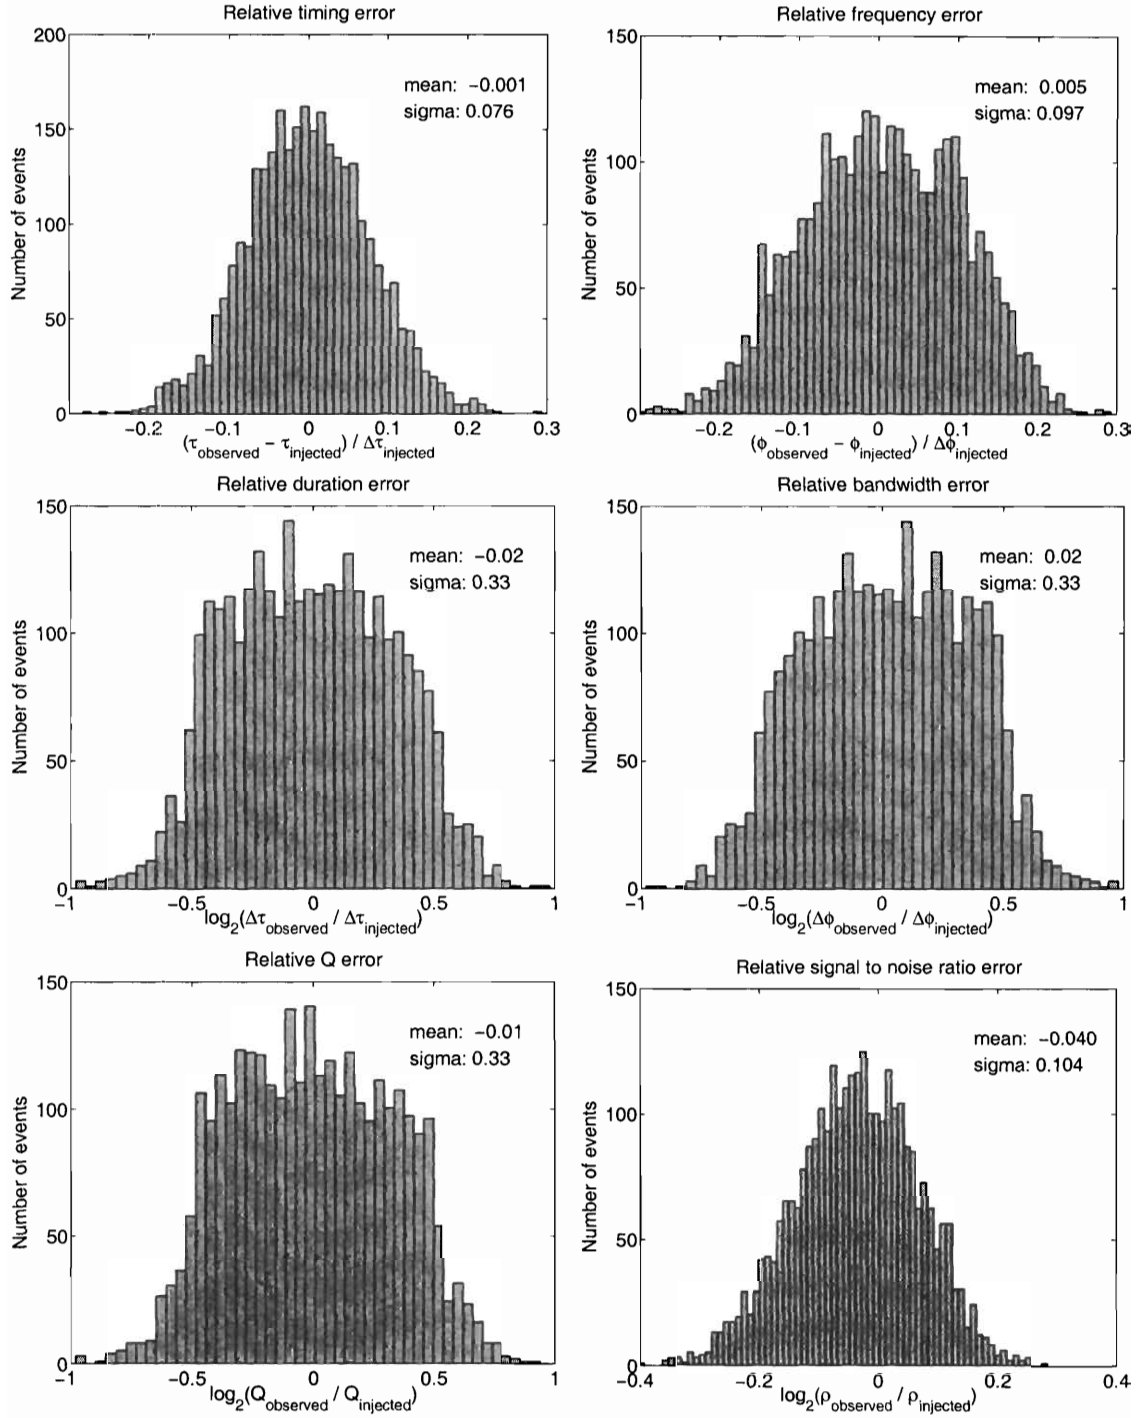


Figure 6.3b: Histograms of the relative accuracy with which the time, frequency, duration, bandwidth, Q, and signal to noise ratio are recovered for well localized bursts in the presence of stationary white noise. In this case 3000 bursts were injected with a fixed signal to noise ratio of 10 according the definition of Equation 5.87, but with random time, frequency, phase, and Q within the targeted signal space. The definition of relative error varies between parameters and is described in more detail in the text. Here we have reported the duration and bandwidth of well localized bursts as defined by Equation 3.13a and Equation 3.13a such that their product is unity.

burst. For the remaining parameters, whose measurement errors are expected to be proportional to their injected values, we find it convenient to define the relative error as the logarithm of the ratio of the observed value to the injected value. These errors then give a better way of characterizing the accuracy of our pipeline in a way that does not depend upon the details of the targeted signal space.

It is then evident from Figure 6.3b that the central time and central frequency of well localized bursts are recovered to within approximately 10 percent of their duration and frequency, respectively. In addition, the duration, bandwidth, and Q of the majority of well localized bursts are recovered with an approximate multiplicative factor of 1.5 of their true value. Finally, the injected signal to noise ratio of 10 is recovered to within an approximate multiplicative factor of 1.15.

It is also evident that, while none of the other parameters exhibit a significant measurement bias, the observed signal to noise ratio exhibits an apparent bias towards under-measurement of the true value. However, we also note that the observed negative bias, in this case approximately 2.5 percent of the true value, is consistent with the average amplitude loss expected when tiling the time frequency plane for a worst case energy loss of 20 percent for well localized bursts. As a result, given knowledge of the expected distribution of energy loss due to mismatch, such a bias is easily accounted for. However, since we are only concerned with the detection of such bursts, and since a similar bias is expected in all detectors, we make no attempt to correct for it here. In addition, the method of section 5.4.4 can also be applied to accurately recover the total signal energy of both well-localized and non-localized bursts by integrating the squared Q transform magnitude over a finite region of the time-frequency plane.

Finally, we consider the case of stationary white noise in the absence of any bursts. This permits us to verify the expected false event rate as a function of detection threshold, as well as to characterize the distribution of false event properties and verify uniform coverage of our search over the targeted signal space. To do so, we simply

apply the Q transform to stationary white noise data using the search parameters described above. The resulting Q transform coefficients are then thresholded at an observed signal to noise ratio of 2.95 in order to produce an approximate false event rate of 1 Hz according to the predictions of Equation 5.89. In addition, we also apply the exclusion algorithm of section 5.6 in order to eliminate any redundant events and avoid any possible overcounting.

An empirical estimate of the false event rate is then determined from the number of remaining Q transform coefficients with observed signal to noise ratio greater than a specified detection threshold, divided by the total observation time. This is easily determined as a function of detection threshold from the cumulative distribution of the observed signal to noise ratios for the remaining events. Application of this procedure to 62000 seconds of stationary white noise data then results in the measured false event distribution shown in Figure 6.4. In addition, we also plot the theoretically expected false event rate from Equation 5.89, assuming that the targeted signal space is sufficiently large that the entire information content of data is exercised. While the resulting distributions differ slightly, they nonetheless show remarkably good agreement over a wide range of false event rates considering the simple nature of our first principle arguments and the complexity of our analysis pipeline.

Finally, we consider the distributions of the other properties of false events in order to verify the uniform coverage of our search over the targeted signal space. In Figure 6.5, we present histograms of the observed central time, central frequency, duration, bandwidth, Q , and signal to noise ratio of 80084 false events observed in 62000 seconds of stationary white noise with a signal to noise ratio threshold greater than 2.95.

As expected, the observed central time of such false events is uniformly distributed over the duration of the each data block, consistent with uniform coverage of the targeted signal space. Unfortunately, the observed distributions of central frequency,




Figure 6.4: The observed false event rate due to stationary white noise (solid) is compared with the predicted false event rate of Equation 5.89 (dashed) assuming that the targeted signal space is sufficiently large that the entire information content of the data, 16384 independent measurements per second, is exercised. Considering the simplicity of our theoretical prediction and the complexity of our analysis pipeline, the two distributions show remarkably good agreement over five decades of false event rates.

duration, and bandwidth are more difficult to verify due to our discrete logarithmic tiling of frequency and Q . However, since the distributions of observed frequency and bandwidth are directly determined from the distributions of duration and Q , we do not attempt to interpret them here. Instead, we simply note that the observed distribution of durations is approximately consistent with a distribution that varies inversely with duration, and that such a distribution is expected due to the greater number of short duration bursts within a given observation time and finite range of Q . In addition, we also note that the distribution of observed Q clearly shows the four discrete values of Q tested by our analysis pipeline as a result of tiling the targeted signal space for a 20 percent worst case energy loss due to mismatch. Here, the slightly greater number of events observed at the extrema of the tested range of Q is presumably due to the detection of false events with best match time-frequency tiles outside of the targeted signal space. Finally, the resulting distribution of observed signal to noise ratio is consistent with the expected exponential distribution of normalized energies

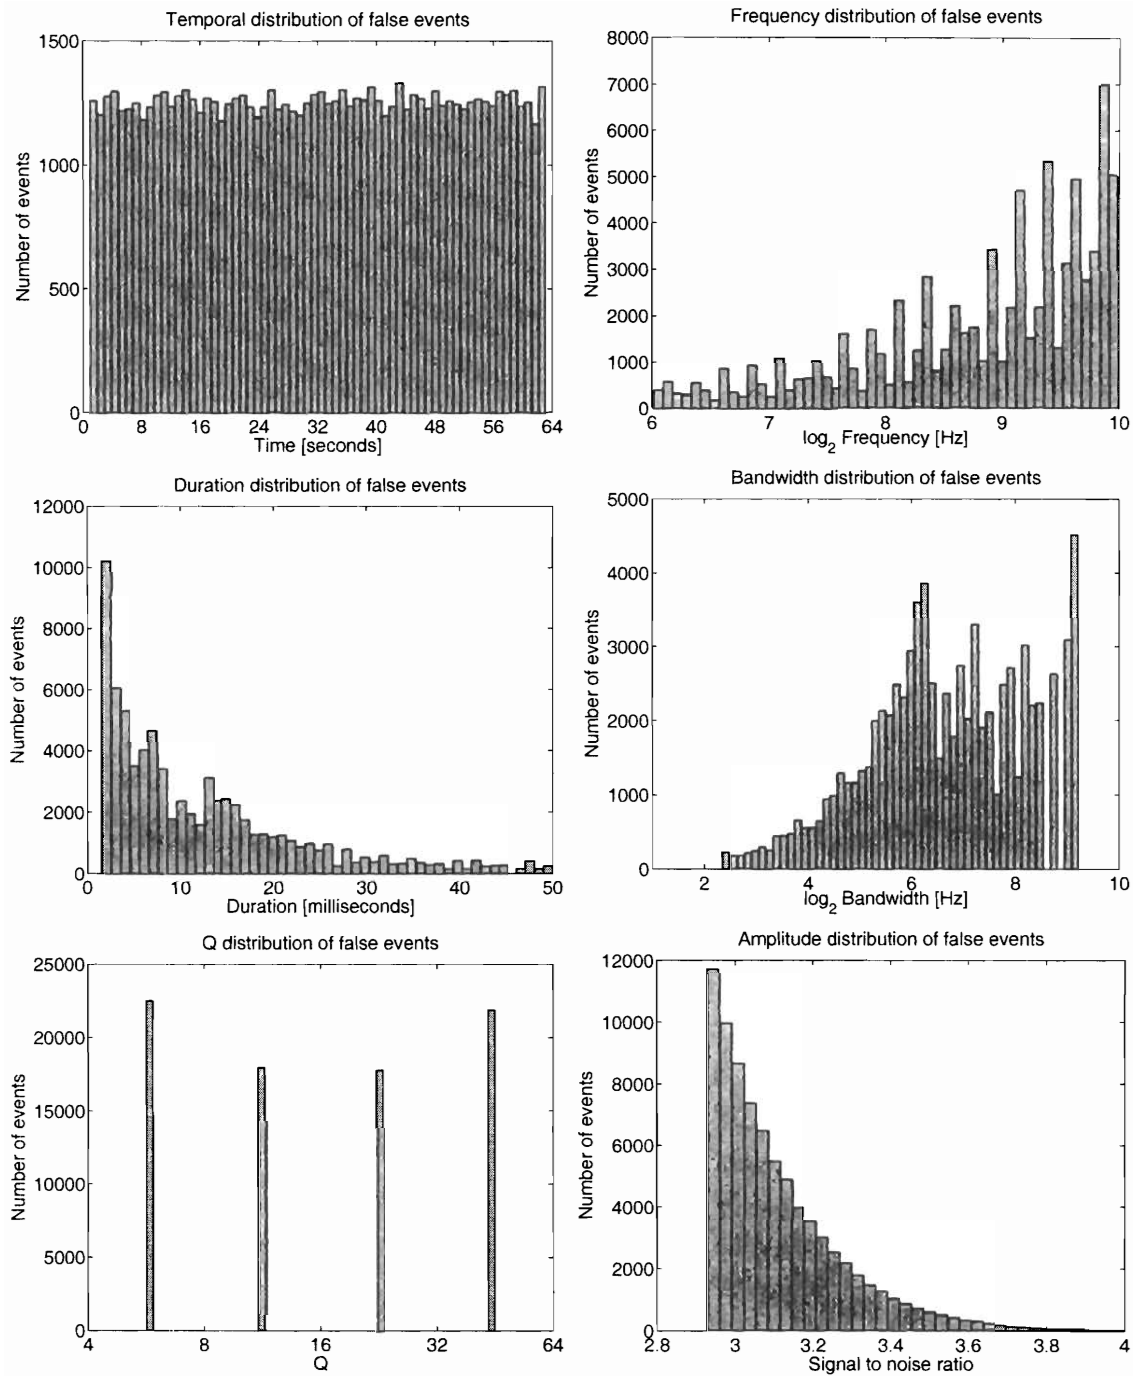


Figure 6.5: The distributions of the observed properties of 80084 false events identified in 62000 seconds of stationary white noise. After taking into account the discrete tiling of the targeted signal space and the expected larger number of independent short duration time-frequency tiles within a given observation time and finite range of Q , the observed distributions are consistent with uniform coverage of the targeted signal space.

predicted by Equation 5.77 and the definition of observed signal to noise ratio given by Equation 5.79.

6.3 Simulated detector noise

We now consider the performance of our single detector pipeline for the case simulated detector noise. Such a study extends our validation from the previous section by also testing the ability of linear prediction to sufficiently whiten data in order to justify our later assumption of stationary white noise when interpreting results. As such, it offers a comprehensive end-to-end validation of our single detector pipeline.

The simulated detector noise considered in this section has been produced by filtering stationary white Gaussian noise such that the resulting colored noise has an amplitude spectrum matching the design sensitivity of the initial LIGO detectors. In addition, to further test our ability to whiten colored noise, we also simulate a number of stationary narrowband signals at frequencies corresponding to some of the expected mechanical resonances of the LIGO detector as well as the 60 Hz and harmonic frequencies associated with electrical power distribution in the United States. However, such resonances are only approximately simulated by applying random phase modulation to sinusoidal signals in order to achieve a bandwidth of roughly 1 Hz. The resulting spectrum is shown in Figure 6.6, where it is compared with initial sensitivity goal of the 4 km LIGO detectors.

It is important to note, however, that this simulated detector noise does not attempt to model the non-stationary behavior expected from actual interferometric detectors. Again, our aim here is not to predict the astrophysical sensitivity of our analysis pipeline. Instead, we seek to validate the proposed single detector analysis pipeline by evaluating its performance under conditions which permit a comparison with the performance predicted from first principle arguments in section 5.5. We therefore




Figure 6.6: Amplitude spectral density of the simulated LIGO detector noise used to characterize the performance of the single detector Q pipeline. The noise consists of stationary white Gaussian noise that has been shaped by a series of filters in order to correspond to the design strain sensitivity of the LIGO 4 km detectors. A select number of resonant line sources are also simulated by random phase modulation of pure sinusoids. However, no attempt is made to simulate the non-stationary behavior that is observed in currently existing interferometric gravitational-wave detectors. Here the amplitude spectral density of the simulated detector noise (gray) is compared with the design sensitivity (black, dashed) of the LIGO 4 km interferometers. Note that, to avoid numerical accuracy problems, the simulated detector noise spectrum flattens out with some ringing below 20 Hz.

postpone the application of our analysis pipeline to actual interferometer data until chapter 7, where we apply it to the search for gravitational-wave bursts in data from the second LIGO science run. Nevertheless, the simulated data set described here provides a convenient benchmark for evaluating and comparing the performance of alternative search algorithms[127, 128].

As was the case for our stationary white noise tests, we again evaluate the ability of the proposed single detector pipeline to detect sinusoidal Gaussian bursts of the form described by Equation 6.1. In this case, however, we test a set of 25 such waveforms which span the targeted space of frequency from 64 to 1024 Hz and Q

from 4 to 64. Specifically, we inject sinusoidal Gaussian bursts with logarithmically spaced center frequencies of 64, 128, 256, 512, and 1024 Hz, and logarithmically spaced Q s of 4, 8, 16, 32, and 64. In addition, these waveforms are injected into the simulated detector noise with a matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 according to the definition of Equation 3.32. We then apply the single detector Q pipeline to search this signal space, which we tile for a worst case energy loss of 20 percent. The resulting list of candidate events is then compared with the list of injected events in order to determine the detection efficiency and the false event rate of our search as a function of significance threshold. Finally, these results are used to produce receiver operating characteristics for each waveform. The resulting receiver operating characteristics are shown in Figure 6.7a through Figure 6.7e and characterize the coverage of our search across the targeted signal space. From these results it is apparent that the proposed single detector pipeline provides similar coverage over the targeted signal space, but that some degradation in performance is seen for waveform on the edge of this signal space.

Finally, we also consider the aggregate performance of our analysis pipeline for all of the sinusoidal Gaussian bursts considered here. The resulting receiver operating characteristics are shown in Figure 6.8, where they are also compared with the theoretically predicted performance from Figure 5.7. We note that the results show very good agreement considering that the present study consisted of stationary colored noise rather than the stationary white noise assumed by the predictions. Finally, it is also interesting to note from Figure 6.5 that the four tested values of Q corresponding to a 20 percent worst case energy loss happen to fall geometrically exactly in between the Q s of the injected signals. Thus, although not by choice, our validation study represents the case of maximum mismatch in Q , making the observed agreement with prediction all the more remarkable.

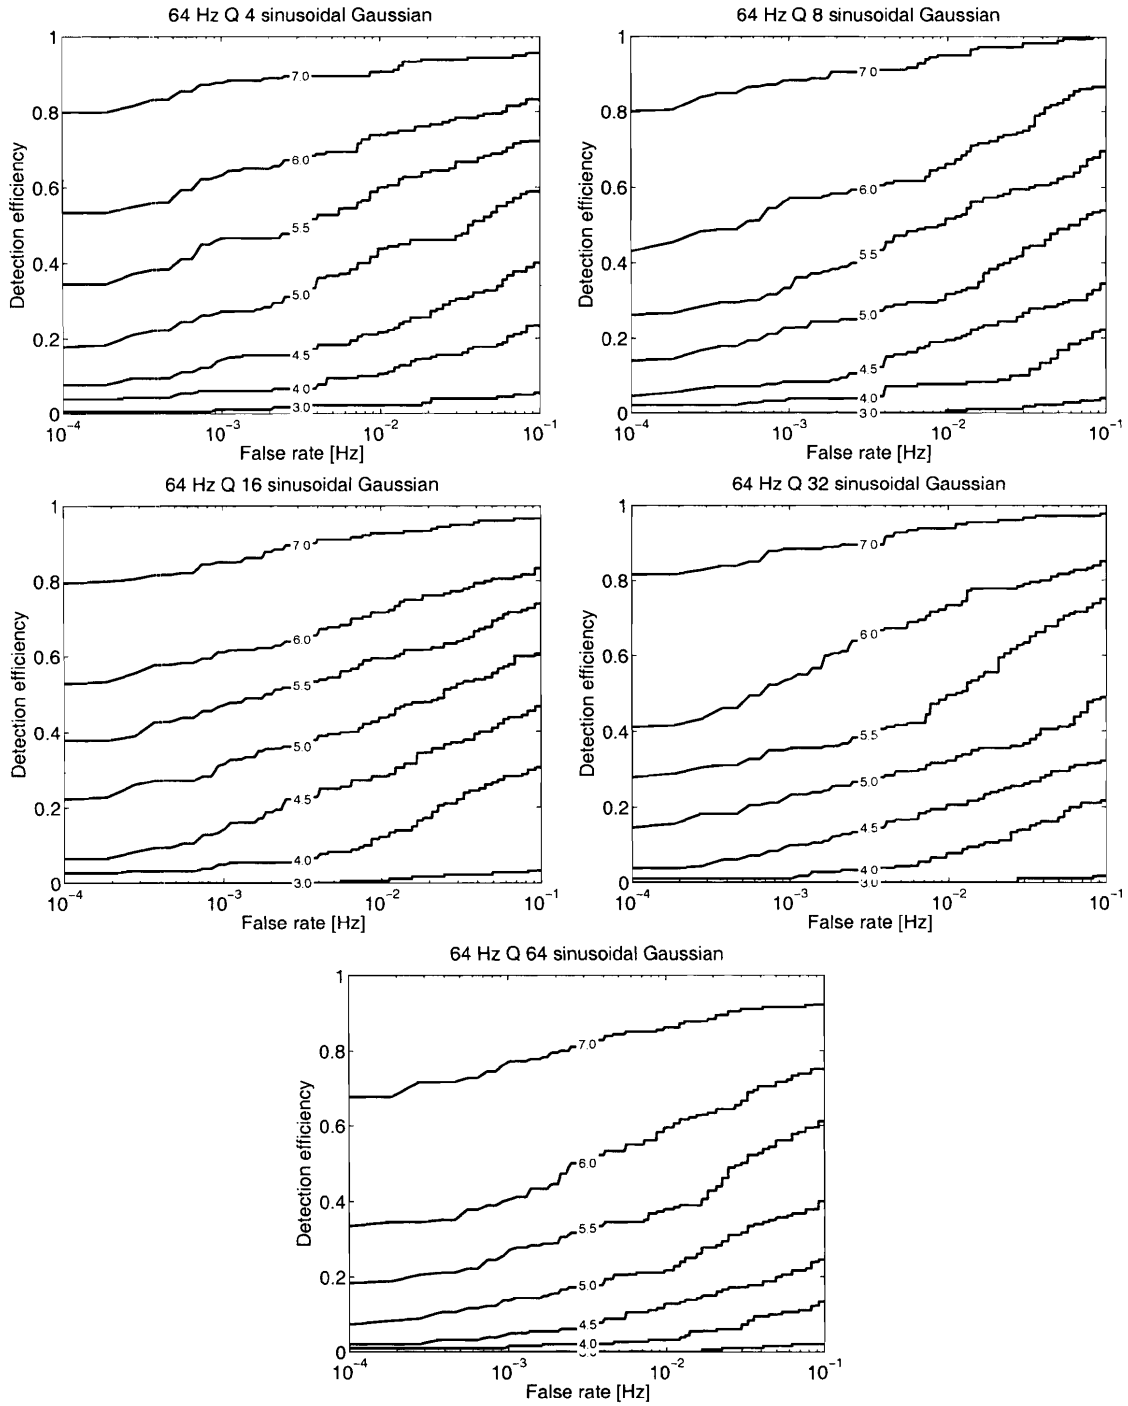


Figure 6.7a: Receiver operating characteristics of the single detector Q pipeline applied to 64 Hz sinusoidal Gaussian gravitational-wave bursts with Qs of 4, 8, 16, 32, and 64 injected into simulated detector noise at matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32.

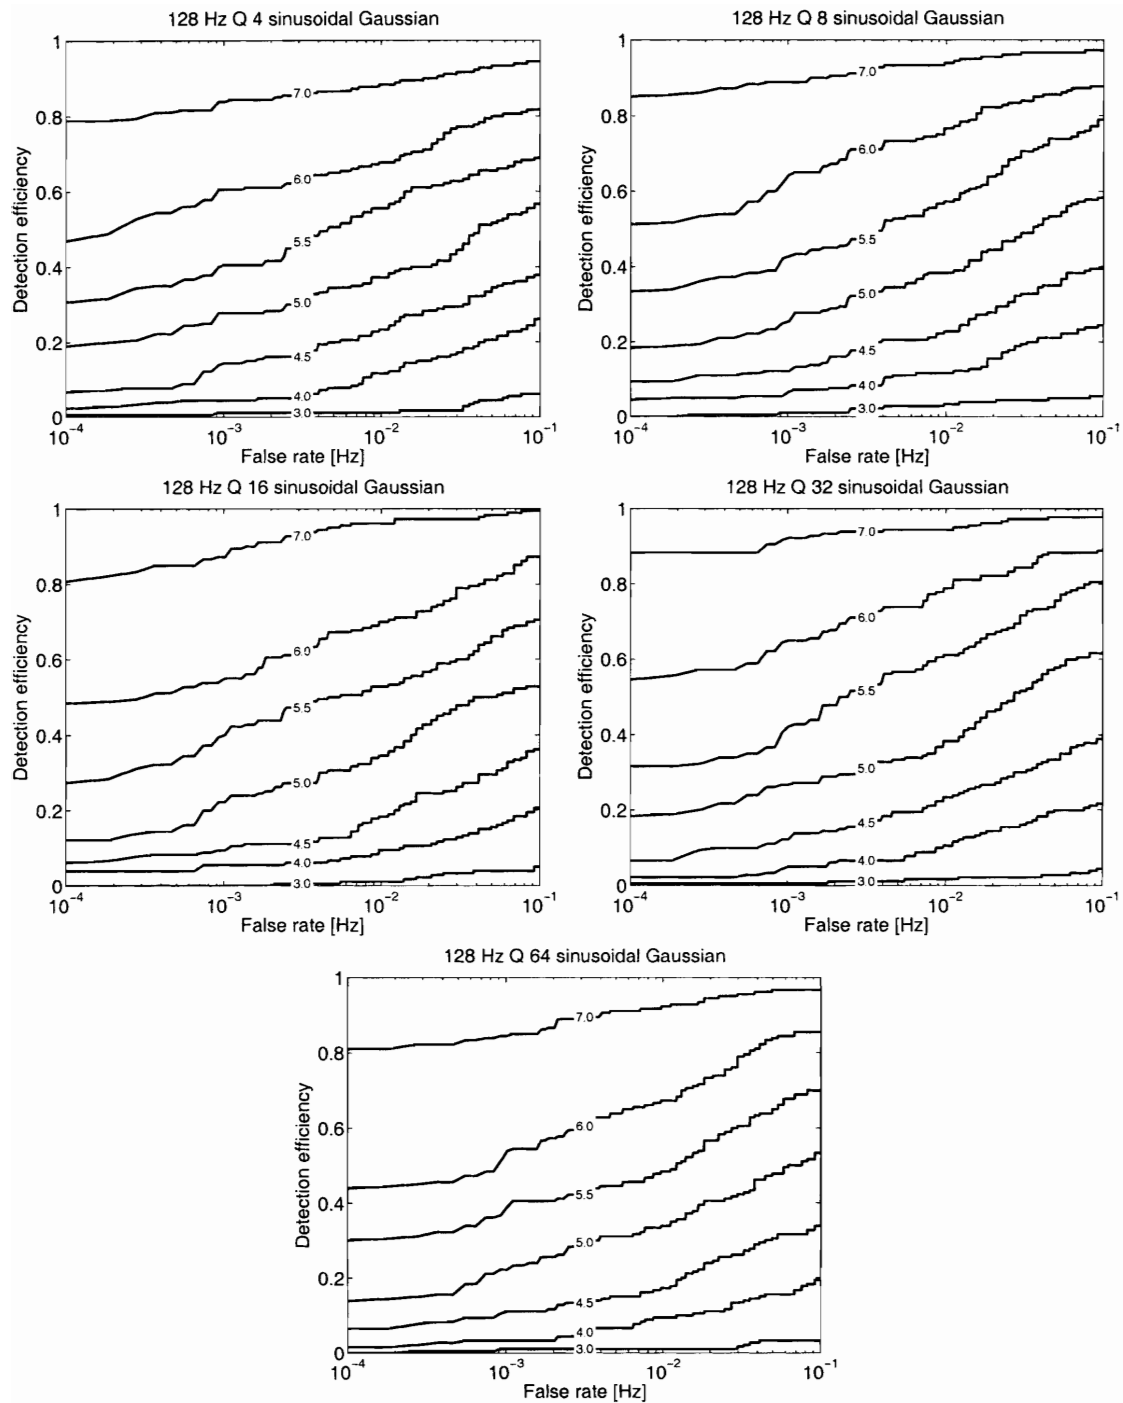


Figure 6.7b: Receiver operating characteristics of the single detector Q pipeline applied to 128 Hz sinusoidal Gaussian gravitational-wave bursts with Qs of 4, 8, 16, 32, and 64 injected into simulated detector noise at matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32.

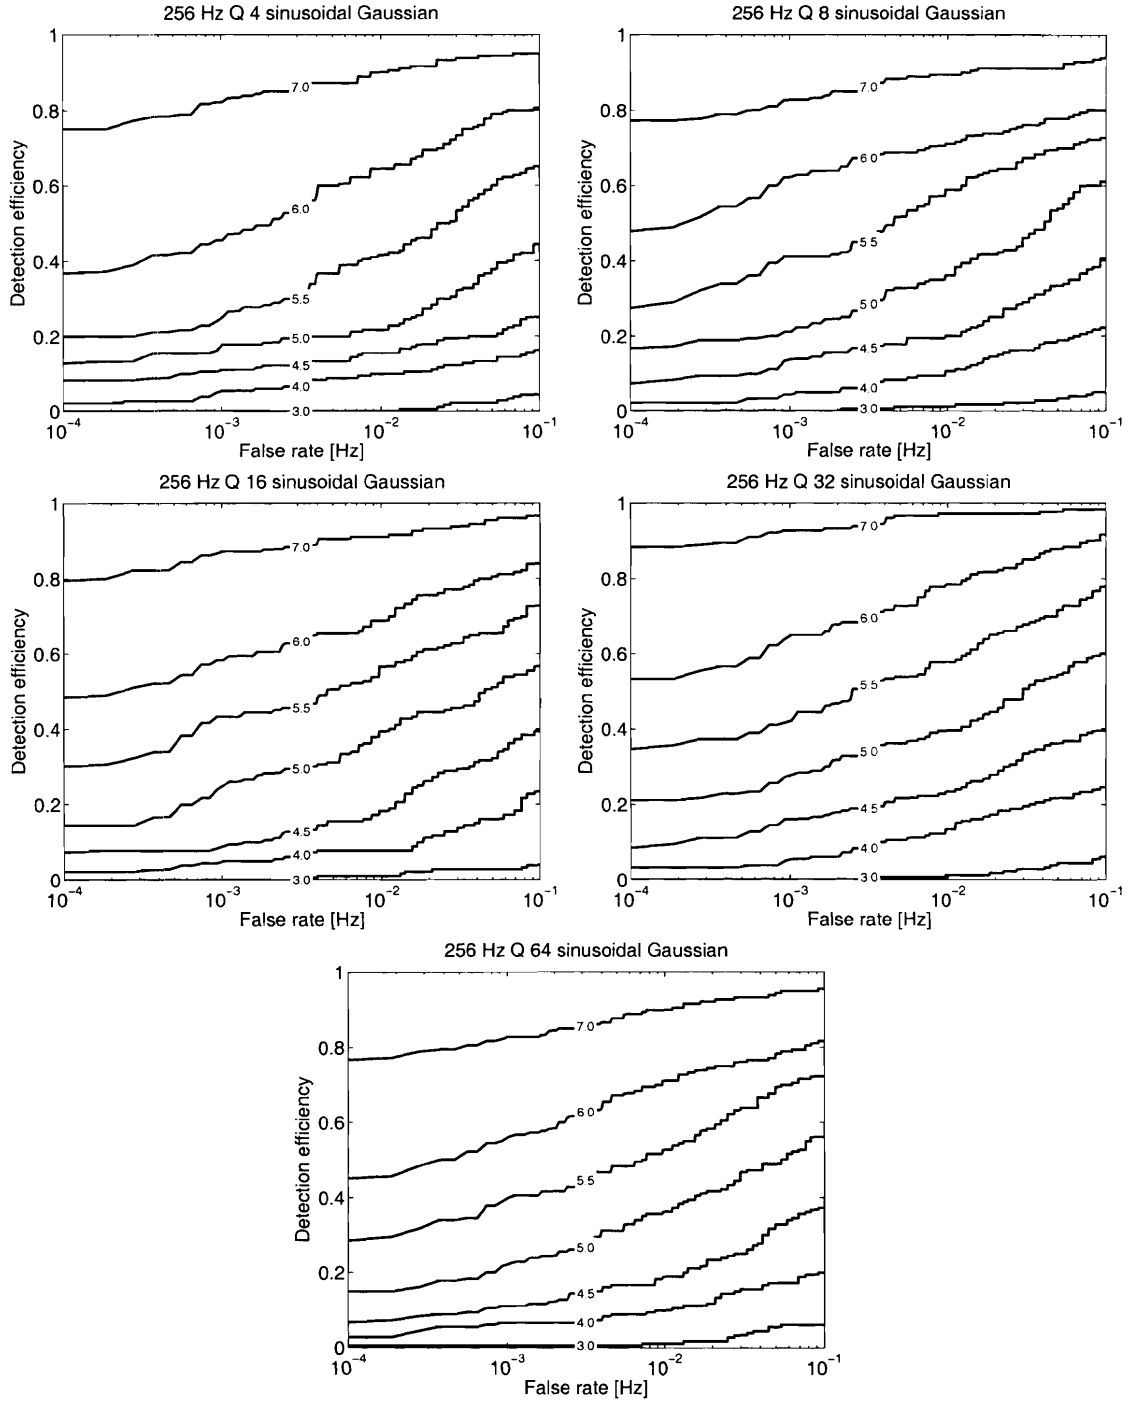


Figure 6.7c: Receiver operating characteristics of the single detector Q pipeline applied to 256 Hz sinusoidal Gaussian gravitational-wave bursts with Qs of 4, 8, 16, 32, and 64 injected into simulated detector noise at matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32.

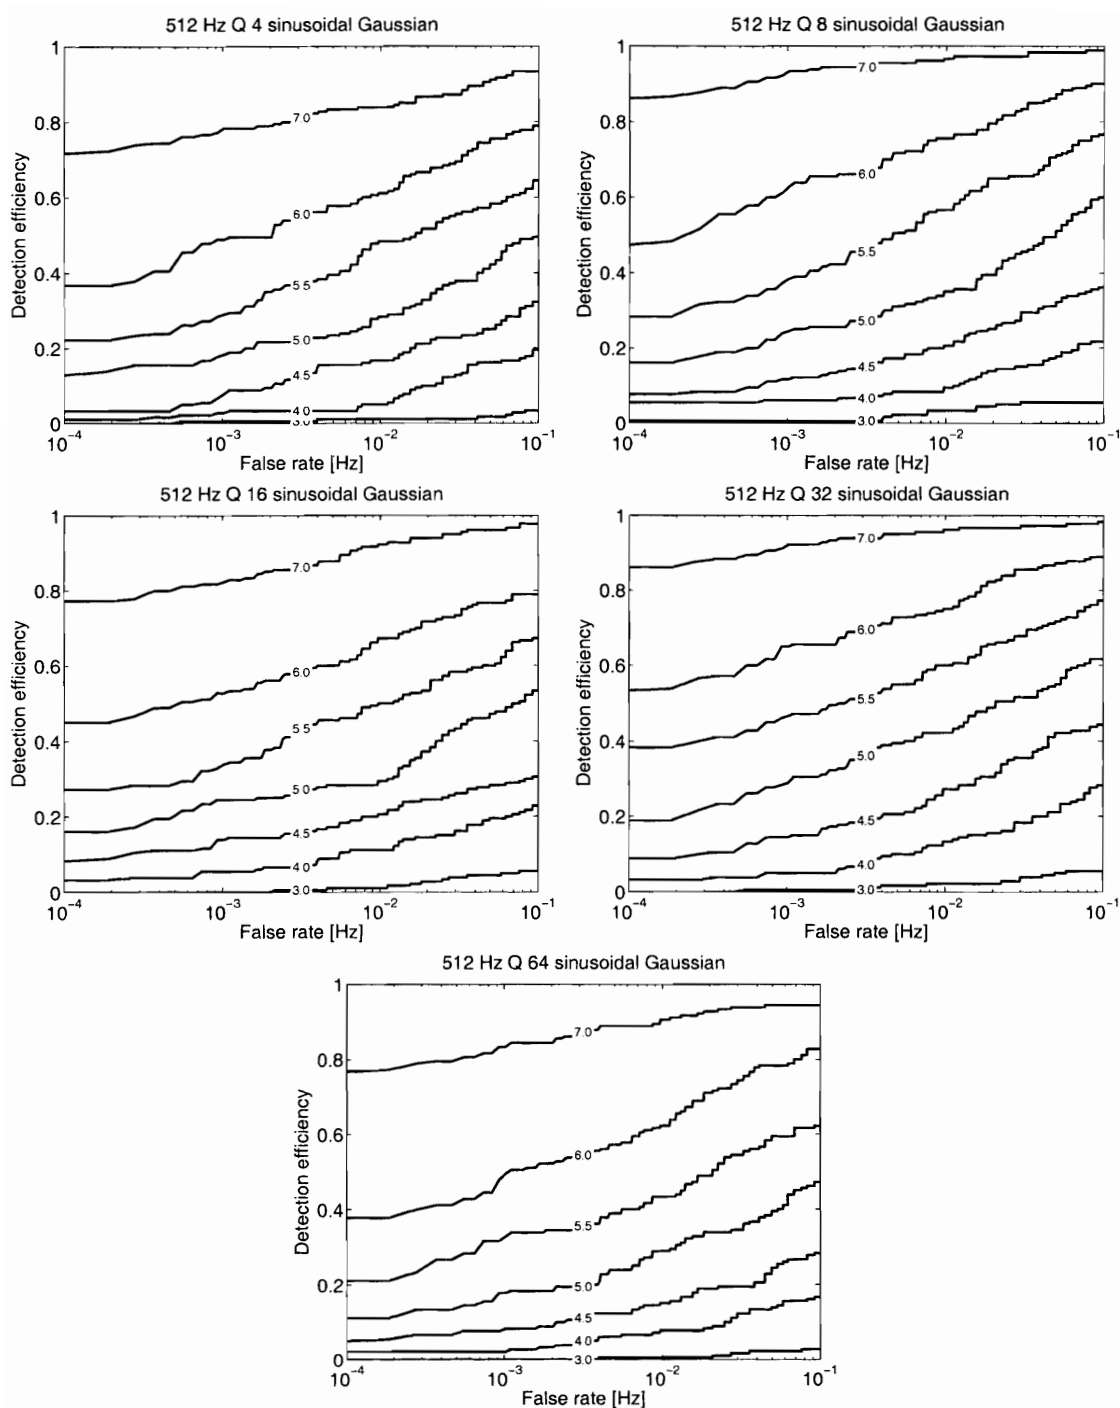


Figure 6.7d: Receiver operating characteristics of the single detector Q pipeline applied to 512 Hz sinusoidal Gaussian gravitational-wave bursts with Qs of 4, 8, 16, 32, and 64 injected into simulated detector noise at matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32.

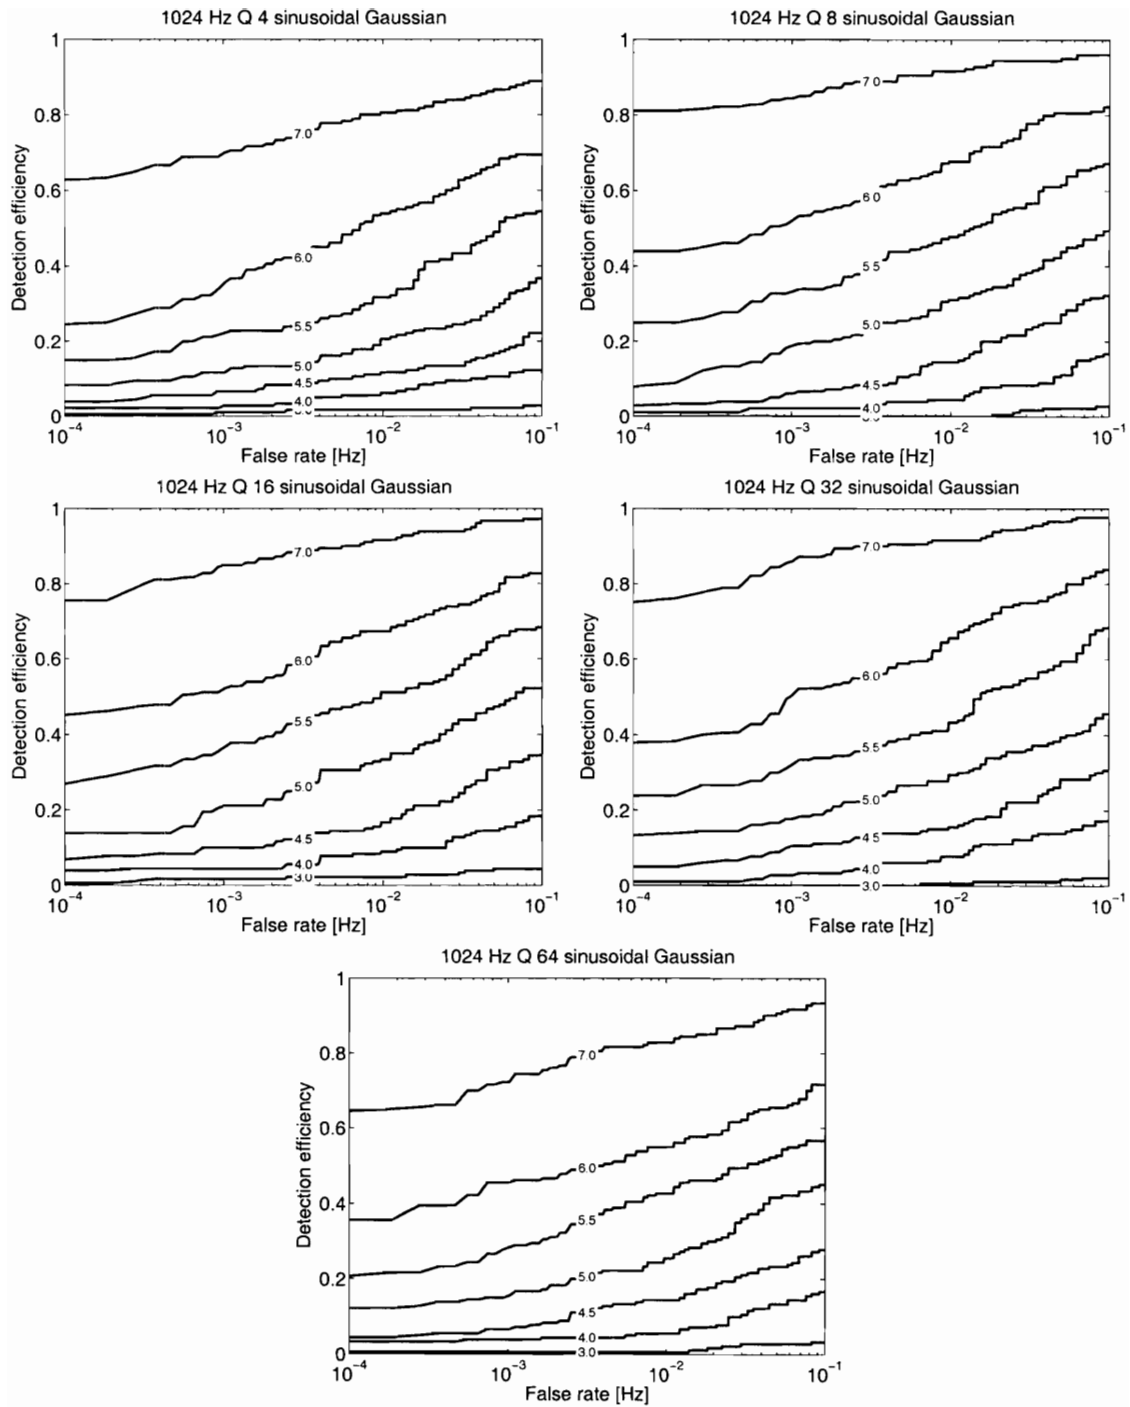


Figure 6.7e: Receiver operating characteristics of the single detector Q pipeline applied to 1024 Hz sinusoidal Gaussian gravitational-wave bursts with Qs of 4, 8, 16, 32, and 64 injected into simulated detector noise at matched filter signal to noise ratios of 3.0, 4.0, 4.5, 5.0, 5.5, 6.0, and 7.0 as defined by Equation 3.32.




Figure 6.8: Aggregate receiver operating characteristic for all sinusoidal Gaussian waveforms. Each curve represents one value of injected signal to noise ratio. The empirically observed receiver operating characteristic (solid) shows reasonably good agreement with the theoretically predicted performance from Figure 5.7, which is also reproduced here (dashed) for comparison.

Chapter 7

All sky search

We now apply the methods of the previous chapters to the all-sky search for unmodulated bursts of gravitational radiation in data from the second LIGO science run.

For this search, we choose to analyze only a subset of this data set that consists of coincident data from the two Hanford detectors. In what follows, we first describe the details of this data set, identifying the advantages and disadvantages of such a choice. In addition, we also identify a number of data quality issues that justify the exclusion of a portion of this data set from our analysis.

We then define the details of the analysis pipeline to be applied to this data set, including the signal space to search and the significance and consistency thresholds for detection. We also identify a procedure, based on non-physical time shifts, for estimating the expected number of background events due to the chance coincidence of non-stationary detector noise. Application of this pipeline to the data then yields a number of candidate events. Although the resulting distribution of these events indicates a statistically significant excess of foreground events, we show that all of the candidate events are inconsistent with the signal expected from a gravitational-wave burst.

Next, we apply the loudest event formalism of Brady, Creighton, and Wiseman[129] to the present search in order to determine an upper bound on the rate of gravitational-wave bursts at a detection threshold corresponding to the amplitude of the most significant foreground event. In order to interpret this result, we also determine the sensitivity of the present search to an isotropic distribution of simulated bursts of both abstract as well as astrophysically motivated waveform. For each waveform, the resulting detection efficiencies are then used to determine an upper bound on the rate of gravitational-wave bursts as a function of signal strength. By doing so, we find that the cumulative observation time of the present search permits a limiting upper bound of 0.086 events per day at a 90 percent confidence level for large amplitude bursts. In addition, depending upon the particular waveform under consideration, we demonstrate that the search achieves 50 percent detection efficiency for isotropic populations of bursts with characteristic strain amplitude $\|h\|$ in the range from 10^{-20} to 10^{-19} strain $\text{Hz}^{-1/2}$ as measured by an optimally oriented detector.

Finally, we compare the resulting upper bounds on the rate of gravitational-wave bursts with the results of four previous searches: the triple coincident all-sky search from the first and second LIGO science runs[130, 92], the cumulative observations from 1997 through 2000 by the network of resonant mass detectors comprising the International Gravitational Event Collaboration (IGEC)[30], and the 2001 run of the EXPLORER and NAUTILUS resonant mass detectors by the Ricerca Onde Gravitazionali (ROG) collaboration[13, 14]. By doing so, we demonstrate that the search presented here is one of the most sensitive to date for unmodeled bursts of gravitational radiation. In addition, we demonstrate at a confidence level in excess of 99 percent that an anomalous excess of events observed by the ROG collaboration is unlikely to be due to a population of sources in the galactic plane.

7.1 Data selection

The second LIGO science run consisted of 1415 hours of joint data collection by all three LIGO interferometers between February 14, 2003 and April 14, 2003. During this period, the Hanford 4 km detector was operational 74 percent of the time, the Hanford 2 km detector was operational 58 percent of the time, and the Livingston 4 km detector was operational 38 percent of the time. In combination, at least one detector was operational 85 percent of the time, two or more detectors were operational 61 percent of the time, and all three detectors were operational 22 percent of the time.

For the present search, we focus only on those times during which both of the Hanford detectors were operational. Due to the smaller duty cycle of the Livingston interferometer, this double coincident Hanford data set accounts for a large majority of the total double coincident observation time. In particular, the two Hanford detectors were in simultaneous operation approximately 50 percent of the time, providing a total of 701 hours of double coincident observation.

This choice of data set has both advantages and disadvantages. Most importantly, it allows for a significantly longer observation time, by a factor of 2.3, than does the corresponding triple coincident analysis. As a result, there is a corresponding increase in the opportunity for detection. In addition, in the absence of non-gravitational-wave background events, such a search is expected to yield an upper bound on the rate of gravitational-wave bursts that is lower, by a similar factor of 2.3, than that which would be produced by a triple coincident search, assuming . Such a search also benefits from the coincident analysis of two detectors, which dramatically reduces the rate of accidental events compared with a single detector search. On the other hand, neglecting the possibility of an additional detector comes at the potential cost of reduced sensitivity in order to maintain an event rate comparable to that of the triple coincidence search.

It is also important to note that the two Hanford detectors share a common environment, including the same geometric orientation for gravitational waves. This also has its advantages and disadvantages. On the one hand, there is an increased probability of observing coincident events due to environmental causes other than gravitational radiation. Such events may result in a statistically significant excess number of foreground events that are not of gravitational-wave origin. On the other hand, the fact that collocated detectors should observe identical signals due to gravitational-waves allows for much stricter consistency tests than those that are possible for non-collocated detectors. In particular, coincident signals of non-gravitational-wave origin are unlikely to produce signals with consistent amplitude and phase between the two detectors. Finally, we note that the choice of collocated detectors will also be beneficial to our analysis in section 7.9.4, where it will simplify the interpretation of our results for the case of non-isotropic distributions of sources on the sky.

We also note that a search of the double coincident Hanford data set does not preclude the ability to search for the presence of a candidate event in other detectors that were operational at the time. In addition, investigation of environmental or auxiliary detector data around the time of candidate events also offers a means of testing the validity of such events.

As we will show, the stricter coincidence testing permitted by the collocated double coincident search effectively compensates for the potential increase in coincident events due to a common environment and the use of only two detectors. As a result, the observed sensitivity to gravitational-wave bursts is comparable to that of the triple coincident search[92]. However, in the final analysis, the added benefit of increased observation time allows one to set a stricter bound on the rate of gravitational-wave bursts than is possible with the triple coincident search.

7.1.1 Data quality

In addition to focusing only on those periods during which both Hanford detectors were operational, it was also necessary to further exclude some portions of this data set based on a number of data quality issues. In particular, extensive investigations by a number of people involved in the commissioning and operation of the LIGO detectors have revealed occasional periods of scientifically unreliable data during the second science run[131]. These include periods of time during which the following problems were identified.

- Invalid or missing calibration information. The time varying response of the interferometer to gravitational-waves is normally monitored by the continuous injection of sinusoidal signals at a few select frequencies. There were times, however, when some of these calibration lines were either absent or significantly weaker than normal, leading to periods of invalid or missing calibration information.
- Invalid timing. The acquisition of data is normally synchronized via the global positioning system. Short periods of time were identified, however, during which the data acquisition system lost timing synchronization, leading to periods with unreliable timing information.
- Photodiode saturation. A mechanism was identified by which occasional large low frequency excitations could cause photodiode saturation at the interferometer's anti-symmetric port. The resulting non-linear response produced occasional bursts of excess noise in the detector output.
- Elevated acoustic noise. Acoustic noise was observed to couple into the interferometer via the input and output optics tables located outside the vacuum envelope. As a result, periods of elevated acoustic noise led to increased noise in the detector.

- Anomalous detector noise. A marginally stable servo loop in the 4 km Hanford detector occasionally produced periods of increased detector noise. However, in order to prevent the exclusion of potential gravitational-waves, such periods were excluded from the analysis only if they exceeded 5 minutes in duration.

After excluding those periods of time that were identified as unreliable, the remaining data set consisted of 650 hours of double coincident data. This corresponds to a 7 percent loss in the available observation time, but still represents an overall 46 percent duty cycle. It is also interesting to note that, after all periods of questionable detector performance are excluded, the double coincident Hanford data set is larger than the triple coincident data set by a factor of 2.6.

7.1.2 Acoustic veto

Although the data quality concerns identified in the previous section included a test for elevated acoustic noise, subsequent investigation found that such tests failed to exclude a number of transient acoustic events that were also observed to couple into the detector. In particular, aircraft flying in proximity to the Hanford site were observed to produce a characteristic time-frequency signature in the interferometer output. An example of such an event is shown in Figure 7.1, where the same characteristic Doppler curve is easily identified in data from both gravitational-wave detectors as well data from microphones located in the corner station and the x-arm end station.

In order to identify such events and remove them from the data stream, the single detector search algorithm proposed in Figure 5.9 was applied to data from a microphone that was located in the corner station. This microphone data was searched for transient events with central frequencies between 64 and 1024 Hz and Q s between 4 and 64. The search was performed at a resolution corresponding to a 20 percent worst case energy loss due to mismatch. The resulting transform coefficients were identified as significant if their normalized energy exceeded 19, corresponding to a

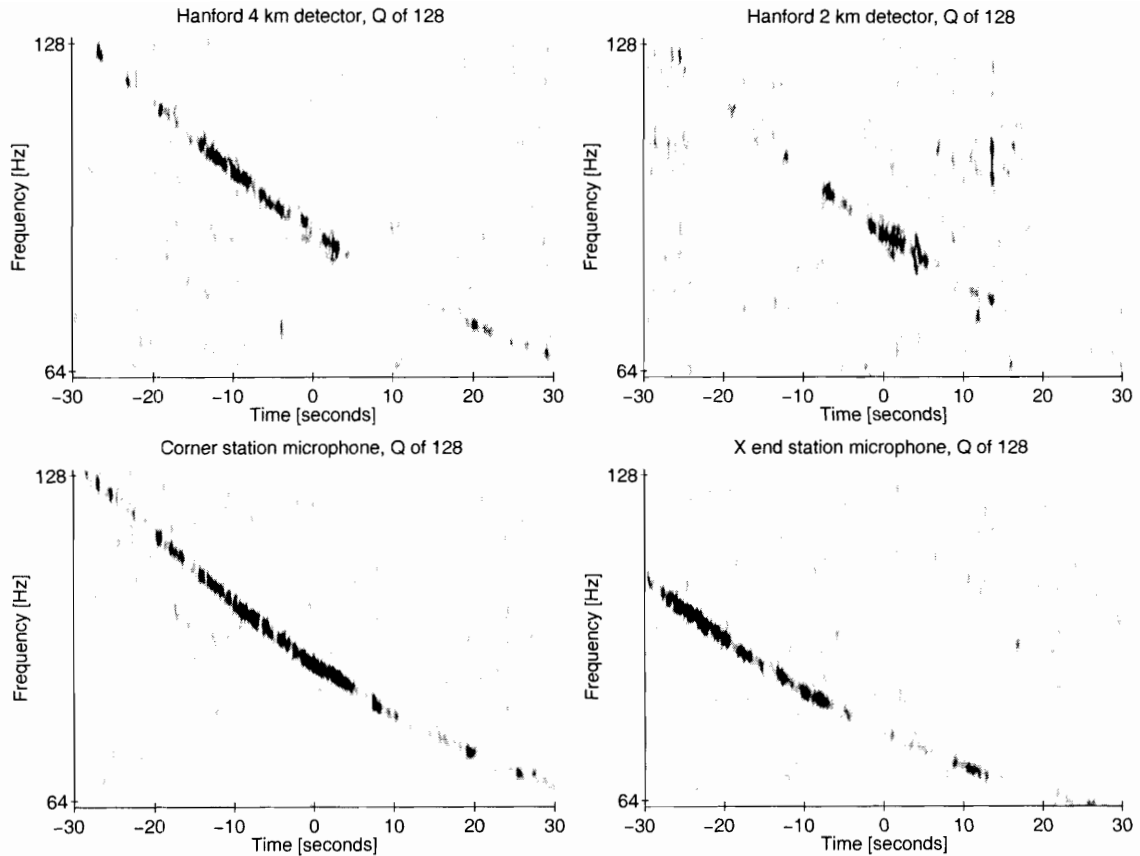


Figure 7.1: The constant Q time-frequency spectrogram of a typical airplane overflight as observed by the Hanford 4km (top left) and 2 km (top right) interferometers as well as microphones located in the corner station (bottom left) and x-arm end station (bottom right). The passing airplane produces a characteristic Doppler curve in the time-frequency plane, with the maximum signal amplitude and maximum frequency rate of change occurring at the point of closest approach. The similarity in timing between the signal observed in the corner station microphone and both interferometers confirms that the primary acoustic coupling occurs in the corner station, where the input and output optics are located outside of the vacuum envelope. All of the spectrograms were computed at a fixed Q of 128 and are shown with a colormap set to saturate at a fixed normalized energy of 25.5, corresponding to an matched filter signal to noise ratio of 7.0 for minimum uncertainty waveforms. The origin in time corresponds to the time of the maximum signal observed in the corner station microphone.

matched filter signal to noise ratio of 6 for minimum uncertainty waveforms. The resulting significant tiles were then sorted and filtered to identify the most significant set of non-overlapping tiles according to the algorithm of section 5.6. Finally, candidate airplane events were identified as those events for which the best match time-frequency tile had a central frequency less than 128 Hz and a Q in excess of 32. These last two conditions were applied in order to preferentially identify airplane

overflights rather than the occasional broadband transient that showed little evidence of coupling into the interferometers.

The resulting event list consisted of 290 candidate airplane overflights during the 650 hours of observation time that survived the data quality considerations of the previous section. We then excluded all data within a 30 second interval on either side of each transient microphone event. The remaining data set consisted of 645 hours of available observation time, corresponding to a decrease of only 0.75 percent due to the exclusion of the candidate airplane events.

It is important to note that the decision to identify and exclude airplane events from the current search is based partly on the results of the triple coincident search for bursts during the second LIGO science run[92]. In that search, the most significant observed event was due to an airplane overflight of the Hanford site in coincidence with a weak transient of unknown origin in the Livingston detector. While the coupling of airplane overflights into the detector has been well known since before the first LIGO science run, no effort was made to exclude such events from the triple coincident search due in part to the low probability of a coincident transient event in a geographically distant detector. Here, based partly on the experience of the triple coincident search and partly on the expectation of coincident events due to a common environment, we have chosen to exclude such events from the analysis.

Finally, we should also note that an aggressive acoustic mitigation effort was carried out after the second LIGO science run. As a result, acoustic coupling has been dramatically reduced for future science runs.

7.2 Analysis pipeline

One of the primary benefits of the present search is the stricter consistency testing afforded by collocated detectors. To take advantage of this, we apply the coherent

analysis pipeline of Figure 5.12 to the current search. In what follows, we briefly describe the specific values that were used for the various search parameters.

The data streams from each detector were first separated into overlapping 64 second blocks of data. This choice reflects the constraints imposed by the available computational resources as well as a compromise between the desire for sufficient statistics and concerns regarding the stationarity of the data. In addition, a minimum overlap of 4 seconds was required in order to discard edge effects associated with filter startup transients and time-domain aliasing in the Q transform.

In order to suppress large low frequency noise outside of the band of interest, the data from both detectors were filtered by a 6th order Butterworth high pass filter with a cutoff frequency of 64 Hz[116]. This filter was first applied causally and then acausally in order to cancel the phase delay introduced by a single pass. As a consequence, the low frequency noise was effectively attenuated by the squared magnitude of the original filter's frequency response. The resulting filter response increases as f^{12} for frequencies less than 64 Hz, achieves 50 percent attenuation at this cutoff frequency, and has an asymptotically flat response for higher frequencies.

The resulting high pass filtered data was then whitened by zero-phase linear prediction using the method described in chapter 4. In order to ensure sufficiently white data for the subsequent search, the whitening was performed using one second linear predictor error filters such that the resulting data were stationary and white on time-scales shorter than one second. In addition, to account for slowly varying changes in the detector noise spectra in a way that is insensitive to the presence of transient events, the linear predictor error filters were retrained on the entirety of each 64 second block of data, with the exception of the filter transients introduced by the zero-phase high pass filter. The resulting linear predictor error filters were then applied to the same 64 second blocks of data used for training.

The Q transform was then applied to the whitened data in order to search for transient

events with central frequencies between 64 and 1024 Hz and Q s between 4 and 64. This choice of frequency range is primarily motivated by the sensitive frequency band of the LIGO Hanford detectors during the second LIGO science run, while the choice of Q ranges from just above the minimum permissible value of Equation 5.19a to a value that yields a maximum signal duration on the order of one second. The resulting signal space encompasses the expected duration and dominant frequency band of gravitational wave bursts from core collapse supernovae as well as the merger of binary compact objects and the ring down of black holes with masses in the range from roughly 1 to 100 solar masses. This targeted signal space was then covered with a sufficient number of overlapping time-frequency tiles to guarantee no more than 20 percent energy loss due to mismatch between an arbitrary minimum uncertainty burst and the nearest measurement tile. The mean and normalized energies of the resulting transform coefficients were then determined after applying the formalism of section 5.5.4 to exclude outliers at an α value of 2.0.

Coincident time frequency tiles were then identified as significant if the observed single detector normalized energies were in excess of 20 in both detectors and the observed coherent normalized energy was in excess of 30. This corresponds to single detector matched filter signal to noise ratio thresholds of 6.2 and a coherent detector matched filter signal to noise ratio threshold of 7.6 for minimum uncertainty waveforms.

The observed set of coincident significant tiles were then sorted and filtered in order to identify the most significant set of non-overlapping tiles according to the algorithm of section 5.6. The resulting set of coincident tiles were then identified as consistent if the ratio of characteristic strain amplitudes observed in both detectors was less than 5 and the differences in phase observed in both detectors was less than 35 degrees. Finally, in order to avoid the redundant reporting of events, the resulting set of consistent tiles were filtered to exclude all but the most significant event on one second time-scales. Here we take the coherent normalized energy to be the primary measure of a tiles significance.

It is worth noting that the search described here was performed in approximately 1.5 hours using a cluster of 290 machines located at the Caltech LIGO Laboratory[132]. Each machine consisted of dual 2.66 GHz Intel Xeon processors with a total shared memory of 2 GB running the RedHat Fedora Core 3 Linux-based operating system. The distribution of the search across this cluster of machines was managed by the Condor[133] batch management system. We also note that, on a single 2.66 GHz Intel Xeon processor, the coherent double coincident search can be performed at a rate approximately 1.75 times faster than real time.

7.3 Event rates

Application of the proposed analysis pipeline and acoustic veto to the selected data set yields a total of 27 coincident events, 10 of which exhibit amplitude and phase consistency as defined in the previous section. These events were identified during a total observation time of 641 hours, which corresponds to an overall duty cycle of 45 percent and takes into account the time lost due to a minimum analyzable segment length of 64 seconds, as well as technical difficulties encountered when carrying out the search.

This number of events is extremely unlikely given the assumption of independent stationary white noise data. It is less surprising, however, when the transient non-stationarities observed in both detectors are considered along with the fact that the two detectors share a common environment. In order to evaluate the significance of our result, we therefore require a method of estimating the expected event rate from such effects. Here we consider only the expected event rate due to the random coincidence of transient non-stationarities. This is accomplished by repeatedly searching the same data set while introducing an artificial time shift between the two detectors. In particular, we choose to test time shifts that are much larger than the expected duration of the gravitational-wave bursts we aim to detect. For the present search,

we examine time shifts of ± 1 , ± 5 , ± 10 , and ± 20 seconds.

The results of this search are summarized in Table 7.1 and Table 7.2, which list the number of unshifted foreground and time-shifted background events observed before and after application of the acoustic veto. Note that the observation time is different for each time shift due to the loss of some data at the boundaries of segments as well as occasional technical difficulties encountered when carrying out the search. As a result, the numbers of events observed at each time shift should not be compared directly. Instead, both tables list the corresponding estimated event rates, as well as an estimated range of event rates corresponding to a one standard deviation error on the true but unknown event rate.

The effectiveness of the acoustic veto is immediately apparent upon comparison of the two tables. In particular, we note that 23 out of 50 coincident foreground events and 6 out of 16 consistent foreground events were excluded as a result of the acoustic veto. For reference, we recall that the acoustic veto represented only a 0.75 percent loss of data. In addition, due to the temporal extent of airplane events, which is typically on the order of 30 seconds, such events are also identified and excluded in the time shifted background searches. In total, 33 coincident events and 14 consistent events were excluded from the cumulative background as a result of the acoustic veto. Many of these events, however, are due to redundant detections of the same acoustic events at multiple time shifts. We also note that, although the acoustic veto appears to be extremely effective, it does not exclude the most significant event observed in the unshifted foreground. As a result, application of the acoustic veto increases our confidence that the surviving events may be gravitational-wave bursts, but does not affect the upper bound on the rate of gravitational-wave events set by the present search.

We now focus only on those events which survive the amplitude and phase consistency tests as well as the acoustic veto. In Figure 7.2, we display the rate of such events as a function of time shift, including the unshifted foreground. Again, we also include

Table 7.1: The observed foreground and background event rates before application of the acoustic veto. The table lists the unshifted foreground event rates as well as the background event rates observed at eight different nonphysical time shifts. For each time shift, the number of events and the corresponding event rate are listed for each of three different thresholds. Coincident events are those events which exceed a single detector normalized energy threshold of 20 and an initial coherent normalized energy threshold of 30. Consistent events are those coincident events which also exhibit amplitude consistency within a factor of 5 and phase consistency within 35 degrees. Finally, significant events are those consistent events which also have a coherent normalized energy that exceeds that of the most significant event observed in the unshifted foreground. For each entry, the table also lists the estimated range of event rates corresponding to one standard deviation errors on the true but unknown event rate.

time shift [seconds]	coincident events	consistent events	significant events	coincident event rate [μHz]	consistent event rate [μHz]	significant event rate [μHz]	observation time [seconds]
0	50 (43.4, 57.6)	16 (12.5, 20.5)	0 (0.0, 1.0)	21.5 (18.7, 24.8)	6.9 (5.4, 8.8)	0.0 (0.0, 0.4)	2325389
-1	28 (23.2, 33.8)	8 (5.6, 11.4)	1 (0.4, 2.6)	12.0 (10.0, 14.5)	3.4 (2.4, 4.9)	0.4 (0.2, 1.1)	2324888
+1	29 (24.1, 34.9)	10 (7.3, 13.7)	0 (0.0, 1.0)	12.5 (10.4, 15.1)	4.3 (3.2, 5.9)	0.0 (0.0, 0.4)	2316182
-5	21 (16.9, 26.1)	10 (7.3, 13.7)	1 (0.4, 2.6)	9.1 (7.3, 11.3)	4.3 (3.2, 5.9)	0.4 (0.2, 1.1)	2316625
+5	20 (16.0, 25.0)	13 (9.9, 17.1)	1 (0.4, 2.6)	8.7 (6.9, 10.8)	5.6 (4.3, 7.4)	0.4 (0.2, 1.1)	2305567
-10	11 (8.1, 14.9)	4 (2.4, 6.6)	0 (0.0, 1.0)	4.8 (3.5, 6.4)	1.7 (1.1, 2.8)	0.0 (0.0, 0.4)	2306568
+10	14 (10.7, 18.3)	5 (3.2, 7.8)	0 (0.0, 1.0)	6.1 (4.6, 7.9)	2.2 (1.4, 3.4)	0.0 (0.0, 0.4)	2310972
-20	11 (8.1, 14.9)	2 (1.0, 4.0)	0 (0.0, 1.0)	4.8 (3.6, 6.5)	0.9 (0.4, 1.7)	0.0 (0.0, 0.4)	2293168
+20	9 (6.5, 12.5)	4 (2.4, 6.6)	0 (0.0, 1.0)	3.9 (2.8, 5.5)	1.7 (1.1, 2.9)	0.0 (0.0, 0.4)	2286997
$\neq 0$	143 (132, 156)	56 (49.0, 64.0)	3 (1.7, 5.3)	7.75 (7.1, 8.4)	3.0 (2.7, 3.5)	0.2 (0.1, 0.3)	18460967

Table 7.2: The observed foreground and background event rates after application of the acoustic veto. The table lists the unshifted foreground event rates as well as the background event rates observed at eight different nonphysical time shifts. For each time shift, the number of events and the corresponding event rate are listed for each of three different thresholds. Coincident events are those events which exceed a single detector normalized energy threshold of 20 and an initial coherent normalized energy threshold of 30. Consistent events are those coincident events which also exhibit amplitude consistency within a factor of 2 and phase consistency within 35 degrees. Finally, significant events are those consistent events which also have a coherent normalized energy that exceeds that of the most significant event observed in the unshifted foreground. For each entry, the table also lists the estimated range of event rates corresponding to one standard deviation errors on the true but unknown event rate.

time shift [seconds]	coincident events	consistent events	significant events	coincident event rate [μHz]	consistent event rate [μHz]	significant event rate [μHz]	observation time [seconds]
0	27 (22.3, 32.7)	10 (7.3, 13.7)	0 (0.0, 1.0)	11.7 (9.7, 14.2)	4.3 (3.2, 5.9)	0.0 (0.0, 0.4)	230820
-1	17 (13.3, 21.7)	5 (3.2, 7.8)	1 (0.4, 2.6)	7.4 (5.8, 9.4)	2.2 (1.4, 3.4)	0.4 (0.2, 1.1)	230780
+1	15 (11.6, 19.4)	3 (1.7, 5.3)	0 (0.0, 1.0)	6.5 (5.0, 8.4)	1.3 (0.7, 2.3)	0.0 (0.0, 0.4)	229900
-5	18 (14.2, 22.8)	9 (6.5, 12.5)	1 (0.4, 2.6)	7.8 (6.2, 9.9)	3.9 (2.8, 5.5)	0.4 (0.2, 1.1)	229960
+5	18 (14.2, 22.8)	12 (9.0, 16.0)	1 (0.4, 2.6)	7.9 (6.2, 9.9)	5.2 (3.9, 7.0)	0.4 (0.2, 1.1)	228870
-10	10 (7.3, 13.7)	3 (1.7, 5.3)	0 (0.0, 1.0)	4.4 (3.2, 6.0)	1.3 (0.7, 2.3)	0.0 (0.0, 0.4)	228960
+10	12 (9.0, 16.0)	4 (2.4, 6.6)	0 (0.0, 1.0)	5.2 (3.9, 7.0)	1.7 (1.1, 2.9)	0.0 (0.0, 0.4)	229400
-20	11 (8.1, 14.9)	2 (1.0, 4.0)	0 (0.0, 1.0)	4.8 (3.6, 6.5)	0.9 (0.4, 1.8)	0.0 (0.0, 0.4)	227660
+20	9 (6.5, 12.5)	4 (2.4, 6.6)	0 (0.0, 1.0)	4.0 (2.8, 5.5)	1.8 (1.1, 2.9)	0.0 (0.0, 0.4)	227040
$\neq 0$	110 (100, 121)	42 (36.0, 49.0)	3 (1.7, 5.3)	6.0 (5.5, 6.6)	2.3 (2.0, 2.7)	0.2 (0.1, 0.3)	1832610




Figure 7.2: The number of candidate events observed in the foreground and background as a function of the time-shift of the search. An estimated range of event rates is also shown for each time shift and corresponds to one standard deviation errors on the true but unknown event rate. A statistically significant excess number of events is evident in the unshifted foreground relative to the majority of the time-shifted background searches. However, a similar excess of events is also observed in the background at time shifts of +5 and -5 seconds. Such an excess may be evidence of environmental correlations on a similar time scale.

an estimated range of event rates that corresponds to a one standard deviation error on the true but unknown event rate.

In order to obtain a more accurate estimate of the background event rate, we also consider the aggregate of all non-zero time shift experiments. Such a measurement has an observation time that is larger, by a factor of approximately 8, than the unshifted foreground observation time. Normalizing the aggregate background event rate by this ratio then yields an estimate of the background event rate with an error that is smaller, by a factor of approximately $\sqrt{8}$, than the error from a single time shift experiment. In Figure 7.3, we compare the cumulative coherent normalized energy distribution of background events obtained in this way with the corresponding distribution of foreground events.

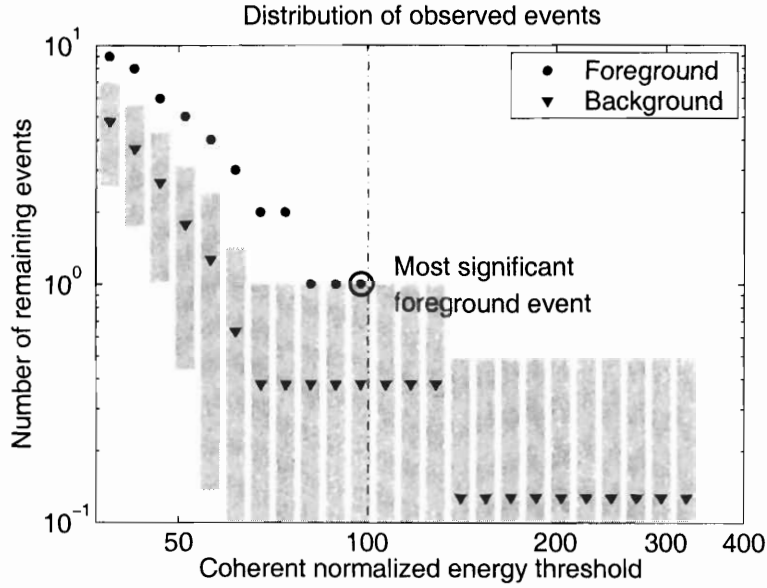


Figure 7.3: The coherent normalized energy distribution of the ten candidate foreground events listed in Table 7.3 superposed on the cumulative distribution of background events observed at all time shifts. Here the background distribution has been normalized by the ratio of the cumulative background observation time to the foreground observation time. For comparison, the background distribution is also shown with error bars corresponding to a one standard deviation variation in the expected foreground rate. Although an apparent statistically significant excess number of foreground events is observed relative to the cumulative background, such an excess may be due to the common environment of the two detectors and is not necessarily indicative of a gravitational-wave foreground.

From both Figure 7.2 and Figure 7.3, it is apparent that the unshifted foreground search exhibits a statistically significant excess event rate relative to the time-shifted background. However, such an excess event rate is not necessarily unexpected. Our method of estimating the background event rate with time shifts only accounts for the accidental coincidence of transient non-stationary events. It does not account for the possibility of coincident events due to a common environmental disturbance.

We note, however, that a statistically significant excess event rate is also observed in the background at time shifts of +5 and -5 seconds. This latter result, coupled with the success of the acoustic veto at non-zero time shifts, suggests that our background estimate does indeed incorporate some knowledge of environmental events. Nevertheless, there remains the possibility of observing an excess foreground due to transient

environmental disturbances with durations shorter than one second.

Unfortunately, distinguishing between environmental events of gravitational-wave and non-gravitational-wave origin is not a simple task and requires an extensive investigation of both environmental and auxiliary detector data around the time of an event. Here we postpone these considerations until the next section, where we consider the likelihood that the identified events are due to gravitational-wave bursts. For now, due to the apparent excess of foreground events, we will refer to the ten remaining events as candidate events. For reference, we list the detected properties of these events in Table 7.3.

Finally, from both Table 7.2 and Figure 7.3, we note that the most significant foreground event occurred at a coherent normalized energy of 100. We will make use of this result again in section 7.6 and section 7.8 where the detection efficiency of our search at this coherent normalized energy threshold essentially determines our upper bound on the rate of gravitational-wave bursts from the present search.

Table 7.3: The measured properties of the ten most significant consistent events observed in the unshifted foreground search. The events are listed in decreasing order of coherent normalized energy, with the most significant event listed first. For all but three of the events, the Livingston 4 km detector was not taking scientifically reliable data. For the remaining three events, no consistent significant event was evident upon manual inspection of the Livingston data. In addition, manual inspection of environmental auxiliary data from the Hanford detectors reveals that none of the events are of likely gravitational-wave origin. Note that the properties are those of the single minimum uncertainty time-frequency tile that best matches each burst. In general, however, they are not representative of the properties of the entire burst.

event rank	central frequency [Hz]	bandwidth [Hz]	duration [ms]	coherent normalized energy	coherent amplitude [strain Hz ^{-1/2}]	4 km normalized energy	2 km normalized energy	amplitude ratio	phase difference [deg]
1.	104.3	8.2	122.4	100.3	3.43×10^{-20}	74.2	26.1	1.03	1.3
2.	104.3	8.2	122.4	76.1	2.43×10^{-20}	29.7	63.8	2.03	33.3
3.	72.5	5.7	176.0	62.2	4.71×10^{-20}	37.6	27.5	1.58	3.4
4.	111.7	8.7	114.3	60.8	1.17×10^{-20}	47.2	21.2	2.40	17.9
5.	104.3	8.2	122.4	51.4	2.08×10^{-20}	23.1	29.1	1.28	2.3
6.	79.4	6.2	160.7	47.9	2.04×10^{-20}	23.7	27.7	1.66	11.1
7.	82.2	12.9	77.7	44.0	1.98×10^{-20}	21.7	22.8	1.10	10.9
8.	82.2	12.9	77.7	43.8	1.95×10^{-20}	24.8	20.4	1.12	18.9
9.	107.9	16.9	59.1	41.2	2.21×10^{-20}	24.0	21.0	1.82	11.6
10.	106.7	8.4	119.6	34.2	2.61×10^{-20}	21.4	23.4	2.68	34.4

7.4 Candidate events

In the previous section we noted a statistically significant excess number of events in the unshifted foreground relative to the distribution of events in the time-shifted background. While such an excess is not unexpected due to the common environment of the two detectors, it is instructive to individually investigate all of the candidate events listed in Table 7.3 in an attempt to identify an underlying cause. Here we first seek to identify an obvious environmental cause for each event. However, if such a cause cannot be identified, we also seek evidence of instrumental anomalies at the time of the event. To do so, we perform an extensive manual search of environmental as well as auxiliary interferometer data in proximity to each of the candidate events. In the process, we demonstrate that none of the candidate events is of likely gravitational-wave origin. In what follows, we consider each event individually, starting with the most significant.

Event 1 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows strong evidence of instrumental anomalies in coincidence with the observed event. These included strong signals in both the Michelson and power recycling control signals. Such signals are typically not observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. In addition, while the feature near 100 Hz shows good agreement between both detectors, the 2 km detector contains additional features that are not consistent with the signal observed in the 4 km detector. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.4.

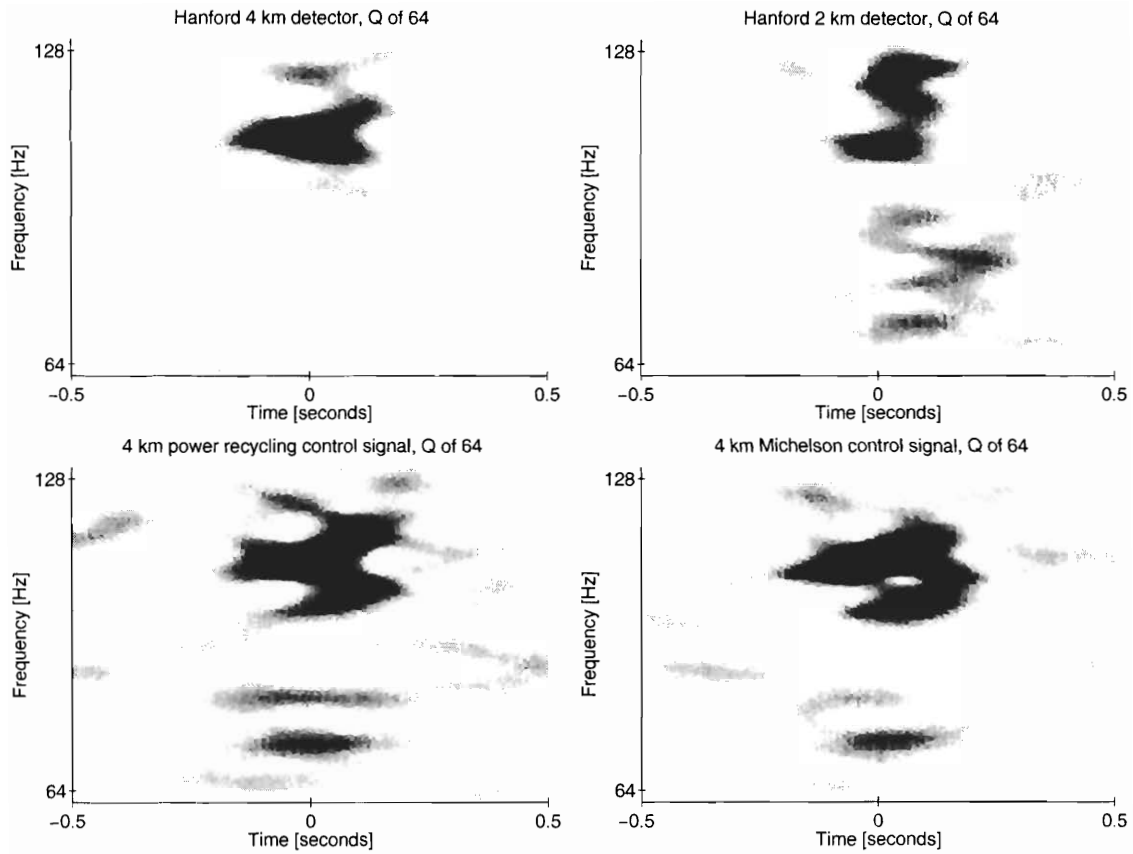


Figure 7.4: Event 1. The constant Q time-frequency spectrogram of the most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower left) and Michelson control signal (lower right) of the Hanford 4 km detector. All four spectrograms were computed at a fixed Q of 64 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 2 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows strong evidence of instrumental anomalies in coincidence with the observed event. In particular, a strong coincident signal was observed in the power recycling control signal. Such signals are not typically observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. In addition, while the feature near 100 Hz shows good agreement between both detectors, the 2 km detector contains additional features that are not consistent with the signal observed in the 4 km detector. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.5.




Figure 7.5: Event 2. The constant Q time-frequency spectrogram of the second most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower) of the Hanford 4 km detector. All three spectrograms were computed at a fixed Q of 64 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 3 This event was clearly coincident with a strong transient signal that appears in seismometer data local to the corner station. However, no corresponding signal was observed in data from seismometers located at the mid-stations or end-stations, nor was any signal present in data from microphones located in the corner station. The event appears to be due to an impulsive seismic event of unknown origin in proximity to the corner station. In addition, while the feature near 70 Hz shows somewhat good agreement between both detectors, the 2 km detector contains additional strong features near 90 Hz that are not consistent with the signal observed in the 4 km detector. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as seismometer data from the corner station are presented in Figure 7.6.




Figure 7.6: Event 3. The constant Q time-frequency spectrogram of the third most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is evident in two different seismometers located in the corner station (lower left and right). The gravitational-wave spectrograms were computed at a fixed Q of 64, while the seismometer spectrograms were computed at a fixed Q of 8. Note that the frequency scale differs between the gravitational-wave and seismometer spectrograms. All four spectrograms are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 4 This event was clearly coincident with an airplane overflight that failed to be detected by both the elevated acoustic noise test of section 7.1.1 and the transient acoustic event test of section 7.1.2. The observed event occurred near the peak rate of change of the Doppler curve, corresponding to the time of closest approach of the airplane to the corner station. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as microphone data from the corner station are presented in Figure 7.7. The extended duration of the signal is clearly evident in both detectors.

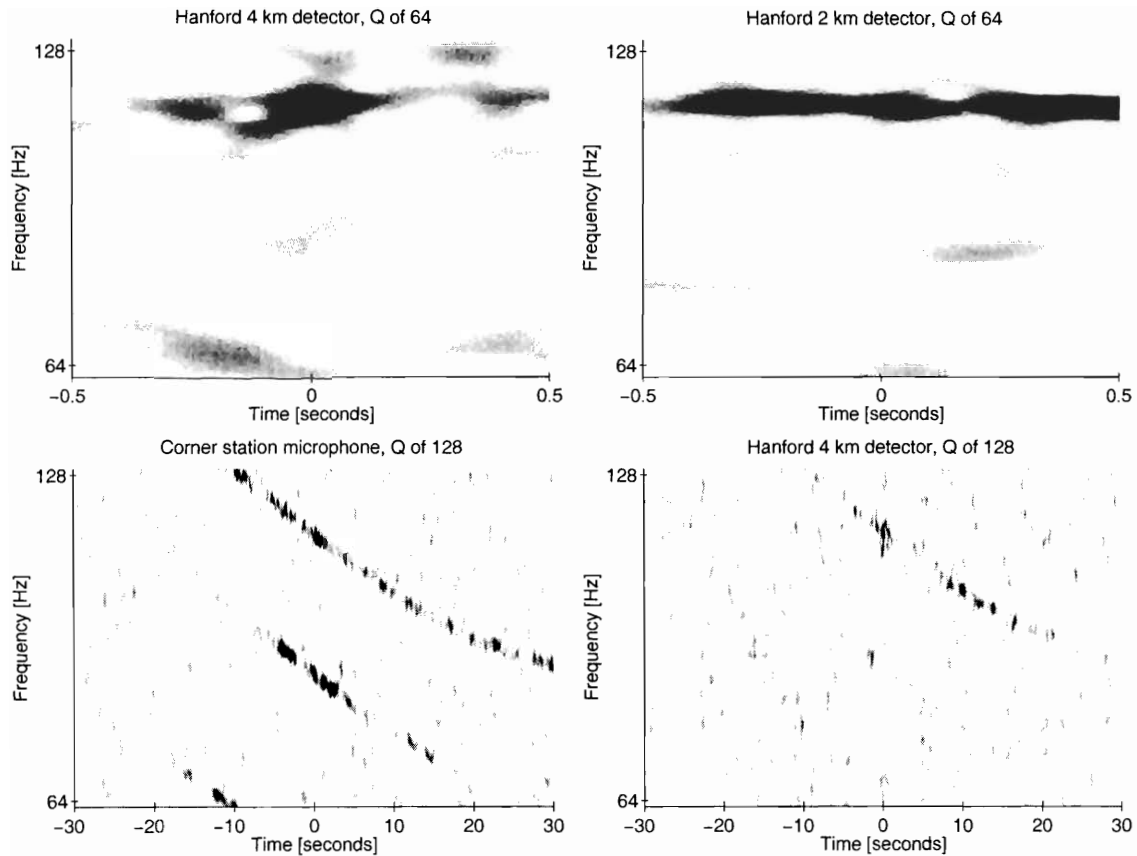


Figure 7.7: Event 4. The constant Q time-frequency spectrogram of the fourth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. The signal is coincident with an airplane overflight, whose characteristic Doppler curve is evident in the spectrogram of a microphone located in the corner station (lower left) as well as an expanded view of data from the 4 km detector (lower right). The gravitational-wave spectrograms were computed at a fixed Q of 64, while the microphone spectrogram and expanded gravitational-wave spectrogram were computed at a fixed Q of 128. All four spectrograms are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 5 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows strong evidence of instrumental anomalies in coincidence with the observed event. These included strong signals in both the Michelson and power recycling control signals. Such signals are typically not observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. In addition, although the Livingston detector was in operation during this event, no event was observed in the Livingston detector. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.8.




Figure 7.8: Event 5. The constant Q time-frequency spectrogram of the fifth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower left) and Michelson control signal (lower right) of the Hanford 4 km detector. All four spectrograms were computed at a fixed Q of 64 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 6 No obvious environmental cause was found for this event, nor was there any obvious sign of instrumental anomalies in auxiliary interferometer data. However, while the cause of this coincident event remains unknown in both detectors, the 2 km detector exhibits extended time-frequency structure that is clearly inconsistent with the signal observed in the 4 km detector. In addition, although the Livingston detector was in operation during this event, no event was observed in the Livingston detector. Constant Q spectrograms of both gravitational-wave signals are presented in Figure 7.9.




Figure 7.9: Event 6. The constant Q time-frequency spectrogram of the sixth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. Both spectrograms were computed at a fixed Q of 64 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 7 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows evidence of instrumental anomalies in proximity to the observed event. In particular, a strong signal was observed in the power recycling control signal that was in coincidence with a portion of the observed burst. Such signals are not typically observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.10.

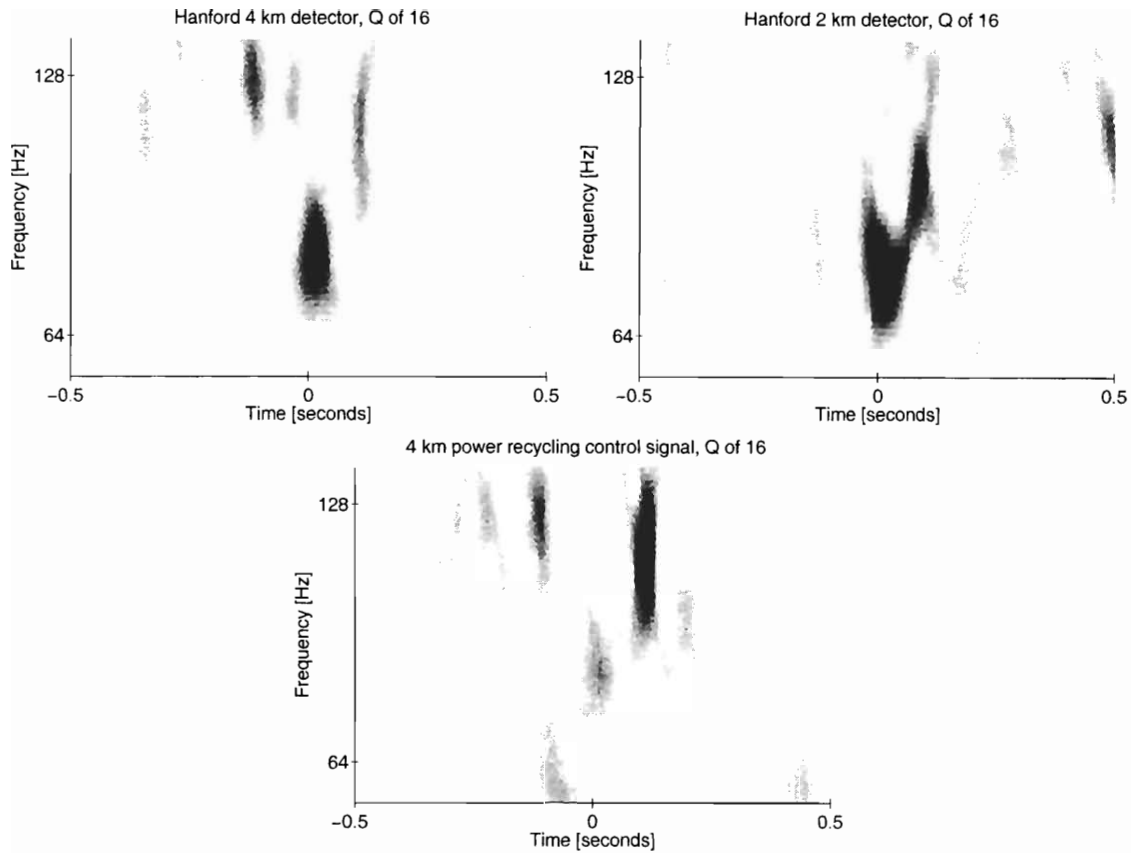


Figure 7.10: Event 7. The constant Q time-frequency spectrogram of the seventh most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower) of the Hanford 4 km detector. All three spectrograms were computed at a fixed Q of 16 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 8 This event was clearly coincident with a strong transient signal that appears in seismometer data local to the corner station. However, no corresponding signal was observed in data from seismometers located at the mid-stations or end-stations, nor was any signal present in data from microphones located in the corner station. The event appears to be due to an impulsive seismic event of unknown origin in proximity to the corner station. In addition, the 2 km detector contains additional strong features that are not consistent with the signal observed in the 4 km detector. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as seismometer data from the corner station are presented in Figure 7.11.




Figure 7.11: Event 8. The constant Q time-frequency spectrogram of the eighth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is evident in two different seismometers located in the corner station (lower left and right). The gravitational-wave spectrograms were computed at a fixed Q of 32, while the seismometer spectrograms were computed at a fixed Q of 16. Note that the frequency scale differs between the gravitational-wave and seismometer spectrograms. All four spectrograms are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 9 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows strong evidence of instrumental anomalies in coincidence with the observed event. In particular, a strong coincident signal was observed in the power recycling control signal. Such signals are not typically observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. In addition, although the Livingston detector was in operation during this event, no event was observed in the Livingston detector. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.12.

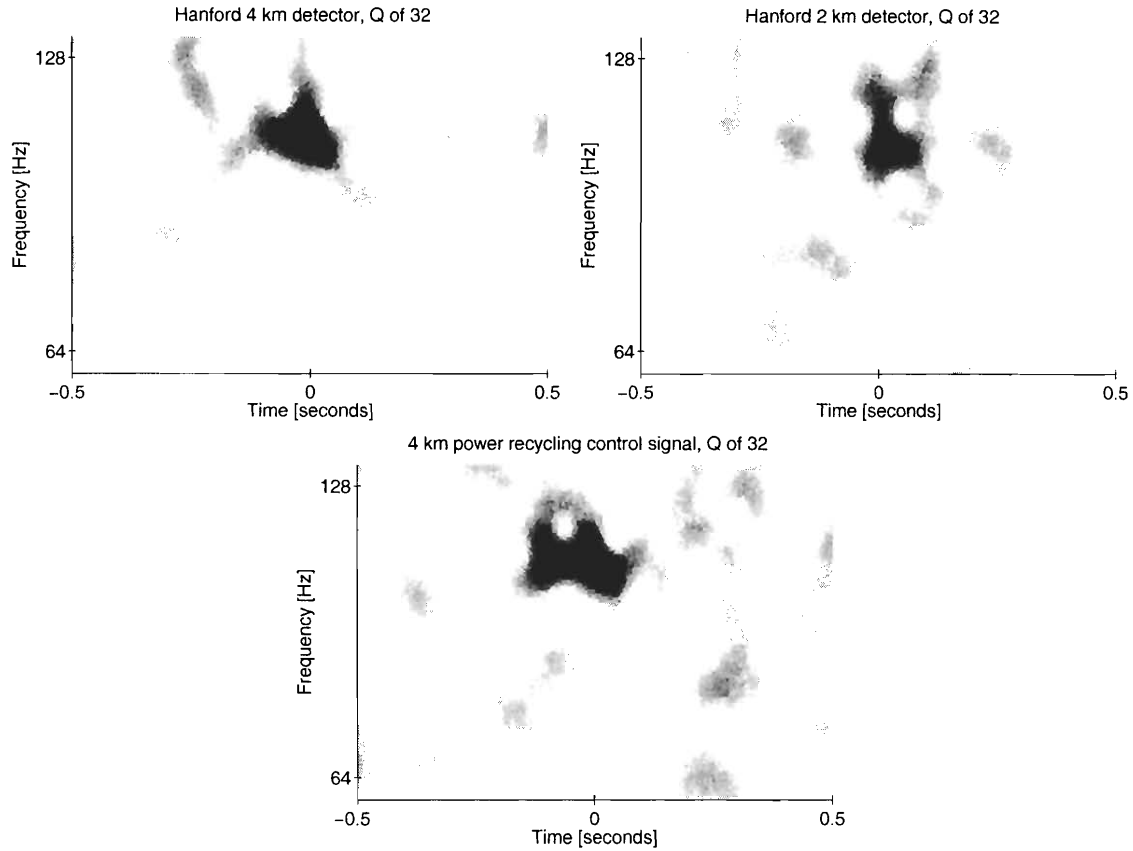


Figure 7.12: Event 9. The constant Q time-frequency spectrogram of the ninth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower) of the Hanford 4 km detector. All three spectrograms were computed at a fixed Q of 32 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

Event 10 No obvious environmental cause was found for this event. However, auxiliary interferometer data from the 4 km detector shows strong evidence of instrumental anomalies in coincidence with the observed event. These included strong signals in both the Michelson and power recycling control signals. Such signals are typically not observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Similar anomalies, however, were not observed in the 2 km detector, and the cause of the coincident event in the 2 km detector remains unknown. The Livingston 4 km detector was not operational during this event. Constant Q spectrograms of both gravitational-wave signals as well as auxiliary data from the 4 km detector are presented in Figure 7.13.




Figure 7.13: Event 10. The constant Q time-frequency spectrogram of the tenth most significant foreground event as observed by the Hanford 4km (upper left) and 2 km (upper right) detectors. A strong coincident signal is also evident in the power recycling control signal (lower left) and Michelson control signal (lower right) of the Hanford 4 km detector. All four spectrograms were computed at a fixed Q of 64 and are shown with a colormap that is set to saturate at a fixed normalized energy of 20, corresponding to a matched filter signal to noise ratio of 6.2 for minimum uncertainty waveforms.

We therefore find that all ten of the candidate events listed in Table 7.3 are not of likely gravitational-wave origin. For three of the events, an obvious environmental cause has been identified. Two of them are of seismic origin, and one is due to an airplane overflight. The remaining seven events are excluded for a combination of reasons. Three of the events occurred while the Livingston 4 km detector was in operation, but a search of the Livingston data failed to yield evidence of the burst. Six of the events are excluded on the basis of strong coincident signals in auxiliary interferometer data that are not typically observed when simulated gravitational-wave bursts are injected into the interferometer by magnetically pushing on the end test masses. Finally, three of the events are also excluded based on the observation of inconsistent time-frequency structure between the gravitational-wave signal in the two detectors.

7.5 Statistical analysis

We now apply the loudest event formalism of Brady, Creighton and Wiseman[129] in order to determine an upper bound on the rate of gravitational-wave bursts. In what follows, we briefly review their approach, and then apply the resulting method to the present search.

We first postulate a source population with the following properties. We assume that gravitational-wave bursts occur according to a Poisson process with a true rate r . The probability of n events occurring in an observation time T is then

$$P(n|r, T) = \frac{(rT)^n e^{-rT}}{n!}. \quad (7.1)$$

We further assume that such events are detected according to a binomial process with a probability of success $\varepsilon(Z)$. The probability of detecting m events given an

occurrence of n events is then

$$P(m|n) = \binom{n}{m} \varepsilon(Z)^m [1 - \varepsilon(Z)]^{n-m}. \quad (7.2)$$

Here $\varepsilon(Z)$ is simply the detection efficiency of the search evaluated for the specified source population at a coherent normalized energy threshold of Z .

Given these assumptions, the joint probability of n events occurring in an observation time T and the subsequent detection of m of these events is then

$$P(m, n|r, T, \varepsilon(Z)) = \binom{n}{m} \varepsilon(Z)^m [1 - \varepsilon(Z)]^{n-m} \frac{(rT)^n e^{-rT}}{n!}. \quad (7.3)$$

The probability of observing zero events during an observation time T is then determined by marginalizing this distribution over all possible n and setting m to zero:

$$P(m = 0|r, T, \varepsilon(Z)) = \sum_{n=0}^{\infty} [1 - \varepsilon(Z)]^n \frac{(rT)^n e^{-rT}}{n!} \quad (7.4)$$

Finally, this result may be simplified by recognizing it as a Taylor series expansion of the exponential function,

$$P(m = 0|r, T, \varepsilon(Z)) = e^{-rT\varepsilon(Z)}. \quad (7.5)$$

We may then construct a classical frequentist upper bound on the rate of gravitational-wave bursts by noting that this result is also the probability $P(Z' < Z|r, T, \varepsilon(Z))$ that the most significant event occurs at a coherent normalized energy Z' that is less than the threshold value of Z at which the search is performed. Setting this probability equal to $1 - p$, where p is the desired confidence level, then yields a frequentist upper bound on the rate of gravitational-wave bursts,

$$r_p = \frac{-\ln(1 - p)}{T\varepsilon(Z)}, \quad (7.6)$$

where the detection efficiency $\varepsilon(Z)$ is to be evaluated at the coherent normalized

energy Z of the most significant event. By construction, the resulting upper bound is chosen such that if the same procedure is repeated for a hypothetical ensemble of identical experiments, then the resulting upper bound is guaranteed not to exclude the true rate of gravitational-wave bursts in a fraction p of those experiments. We thus state with confidence p that the true rate of gravitational-wave bursts is less than r_p .

Note that in developing this result, we have assumed that the most significant event is due only to the assumed source population and has zero probability of arising from a non-gravitational-wave background. Although an expected background can be incorporated into the above formalism, we have decided not to do so for the present search. This choice is based primarily on the uncertainty in our estimate of the background event rate. However, it is also interesting to note that this decision actually leads to a more conservative statement in the sense that the resulting upper bound is not violated by failure to include the background.

Application of Equation 7.6 to the current search then yields the upper bounds listed in Table 7.4. To simplify the comparison of these results with those of other searches, we list the resulting upper bounds in units of events per day for confidence levels of 90, 95, and 99 percent. Note, however, that these are uninterpreted bounds. That is, we have not yet specified a particular source population. Instead, the event rate bounds reported in Table 7.4 correspond to the special case of a source population for which the detection efficiency is unity at the specified threshold. For now, we note that this is the lowest achievable upper bound on the rate of gravitational-wave bursts given the present search. We will revisit this result in section 7.8 after determining the detection efficiencies for a variety of waveforms at a coherent normalized energy threshold equal to that of the most significant foreground event.

Table 7.4: The uninterpreted upper bound on the rate of gravitational-wave bursts from the present search. For comparison with previous searches, the upper bound is given for confidence levels of 90, 95, and 99 percent and listed in units of events per day. The listed event rates correspond to a source population for which the detection efficiency is unity at a coherent normalized energy threshold equal to that of the most significant foreground event. In the next two sections, we derive interpreted upper bounds that also incorporate the detection efficiency of the present search for a variety of simulated gravitational-wave bursts as a function of characteristic strain amplitude and distance.

confidence level	upper bound number of events	upper bound events per day
90%	2.303	8.62×10^{-2}
95%	2.996	1.12×10^{-1}
99%	4.605	1.72×10^{-1}

7.6 Detection efficiencies

In the previous section, we determined an upper bound on the rate of gravitational-wave bursts in terms of the detection efficiency for a particular source population. In order to apply this result, we now determine the detection efficiency of the present search for a variety of simulated gravitational-wave bursts of both abstract and astrophysically motivated waveform. Following the prescription of the previous section, these detection efficiencies are evaluated at a coherent normalized energy threshold equal to that of the most significant foreground event.

This is accomplished by coherently injecting simulated populations of each waveform into the existing double coincident data set and repeating the search described in section 7.2. In doing so, we must make assumptions regarding the distribution of gravitational-wave bursts on the sky. However, since we have little *a priori* knowledge of the population of gravitational-wave bursts, we simply choose the least informative assumption. That is, we assume that the bursts originate from an isotropic distri-

bution of sources with random linear polarization. In order to determine an upper bound on the rate of such bursts as a function of amplitude, we further assume that all of the bursts emitted from such a population occur with the same characteristic strain amplitude $\|h\|$ as observed by an optimally oriented detector. We then determine the detection efficiency for many such populations, each with a different value of $\|h\|$.

To simulate each population, the proposed waveform is then repeatedly added to the data at randomly selected times with an amplitude determined by the characteristic strain amplitude of the population and an attenuation factor consistent with the angular response of the detector and a randomly selected sky position and polarization. We then determine the detection efficiency for each such population as the fraction of the injected events that are detected at the specified threshold.

For each waveform, the resulting detection efficiencies are then presented as a function of characteristic strain amplitude.

In general, the resulting measurements comprise a monotonically increasing function of signal amplitude and are conveniently described by an asymmetric sigmoid function of the form[92]

$$\varepsilon(\|h\|) = \frac{1}{1 + (\|h\|/h_c)^{-x(\|h\|)}}, \quad (7.7)$$

where the function $x(\|h\|)$ is given by

$$x(\|h\|) = \frac{1 + a \tanh(\|h\|/h_c)}{w}. \quad (7.8)$$

Here h_c is the approximate characteristic strain amplitude at the center of the transition region, w is a dimensionless parameter that describes the approximate width of the transition region in decades, and a is a dimensionless asymmetry parameter. In what follows, we fit the observed detection efficiencies for each waveform to a function of this form in order to obtain an accurate estimate of the detection efficiency as a continuous function of characteristic strain amplitude. Although there is no rigor-

ous mathematical basis for this parameterization, it has empirically been observed to provide a good fit to the measured detection efficiencies for a wide range of search methods and simulated waveforms.

We now consider the detection efficiencies for four families of waveforms. Two of the families are abstract and consist of simple Gaussian bursts and sinusoidal Gaussian bursts. The primary motivation for injecting such waveforms is to provide a widely distributed test of the performance of the search over the targeted signal space, as well as to provide a basis for quickly estimating the detectability of other proposed waveforms. Unfortunately, due to limited computational resources, the number of such waveforms that were simulated is insufficient to completely cover the space of targeted signals. As a result, we rely both on the sparse test afforded by the these simulated waveforms as well as the more comprehensive validation study presented in chapter 6 in order to provide confidence in the ability of our search to cover the targeted signal space.

We first consider source populations consisting of simple Gaussian bursts of the form

$$h(t) = \|h\| \left(\frac{1}{2\pi\sigma_t^2} \right)^{1/4} \exp \left[\frac{-(t - \tau)^2}{4\sigma_t^2} \right]. \quad (7.9)$$

Populations of such Gaussian bursts were injected with durations σ_t of 0.05, 0.25, 0.5, 2.0, and 3.0 milliseconds and with twenty different values of characteristic strain amplitudes $\|h\|$ as measured by an optimally oriented detector. The resulting detection efficiencies for isotropic distributions of such sources are shown in Figure 7.14.

We also consider source populations consisting of sinusoidal Gaussian bursts of the form

$$h(t) = \|h\| \left(\frac{32\pi\phi^2}{Q^2} \right)^{1/4} \exp \left[-\frac{4\pi^2\phi^2(t - \tau)^2}{Q^2} \right] \sin [2\pi\phi(t - \tau)]. \quad (7.10)$$

Populations of such bursts were injected with central frequencies ϕ of 100, 153, 235, 361, 554, and 849 Hz and a single Q of 12.7. As was the case for the Gaussian bursts,




Figure 7.14: The detection efficiency as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of simple Gaussian bursts. Detection efficiencies are shown for bursts with duration σ_t ranging from 0.05 to 3.0 milliseconds at a detection threshold corresponding to the coherent normalized energy of the most significant foreground event.

source populations were simulated for twenty different values of characteristic strain amplitude $\|h\|$ as measured by an optimally oriented detector. The resulting detection efficiencies for isotropic distributions of such sources are shown in Figure 7.15.

In addition to these two abstract families of waveforms, we also inject two families of astrophysically motivated waveforms. These are binary black hole mergers and axisymmetric core collapse supernovae. Here we make use of the waveforms produced by the various simulation efforts that are described in chapter 2. We note, however, that such simulations remain an active topic of research and that the corresponding waveforms, although plausible, are not predicted with nearly the same degree of certainty as that of inspiraling binary neutron stars. As a result, their astrophysical importance is limited. Nevertheless, such waveforms are included because they provide a richer set of waveforms on which to test the robustness of the search, and also because they provide a rough estimate of the sensitivity of the search to two of the most plausible sources that LIGO may detect. In addition, such waveforms also




Figure 7.15: The detection efficiency as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of sinusoidal Gaussian bursts. Detection efficiencies are shown for bursts with a Q of 12.7 and central frequency ϕ ranging from 100 to 849 Hz at a detection threshold corresponding to the coherent normalized energy of the most significant foreground event.

provide a useful basis of comparison with other algorithms and detectors.

We first consider source populations consisting of the simulated binary black hole merger waveforms from Baker, Campanelli, Lousto, and Takahashi[38, 39]. In this case, we have injected populations of binary black hole mergers with total binary masses of 10, 30, 50, 70, and 90 solar masses. Again, source populations were simulated for twenty different values of characteristic strain amplitude $\|h\|$. However, we also note an important difference between the injected populations of binary black hole mergers and the abstract waveforms considered previously. Here, we are provided with two waveforms, one for each gravitational-wave polarization, rather than just a single waveform. In general, the injected waveform will then depend non-trivially on both the position of the source on the sky as well as the orientation of the source relative to the line of sight. For simplicity, however, we inject the binary black hole merger waveforms with equal signal energy in both gravitational-wave polarizations. Although we still allow for a random polarization angle, this choice corresponds to a

fixed 60 degree inclination of the rotation axis of the binary black hole system relative to the line of sight. In addition, we characterize the amplitude of each population by the quadrature sum of the signal amplitudes in both gravitational-wave polarizations incident on the Earth. Roughly speaking, we expect these choices to yield a signal amplitude similar to the average signal amplitude expected from a randomly oriented population of sources that are all at the same distance. However, here we do not attempt to accurately account for a truly random distribution. Instead, we simply present our results for populations of fixed inclination sources that are isotropically distributed on the sky, all at the same characteristic strain amplitude. The resulting detection efficiencies for such source populations are shown in Figure 7.16.




Figure 7.16: The detection efficiency as a function of characteristic strain amplitude $\|h\|$ for isotropic distributions of binary black hole mergers with fixed inclination. Detection efficiencies are shown for simulated binary black hole merger waveforms[38] with a total binary masses of 10, 30, 50, 70, and 90 solar masses at a detection threshold corresponding to the coherent normalized energy of the most significant foreground event. Here, the rotation axis of the binary black hole systems are all fixed at a 60 degree inclination relative to the line of sight, and the reported signal amplitudes correspond to quadrature sums of the signal amplitudes in both gravitational-wave polarization incident on the Earth.

We also consider source populations consisting of the simulated axisymmetric supernovae core collapse waveforms from Zwerger and Mueller[42], Dimmelmeier, Font,

and Mueller[43], and Ott, Burrows, Livne, and Walder[44]. However, due to the large variety of waveforms in each of the three studies, we do not consider each waveform separately. Instead we choose to consider each of the three sets of simulations as a single population. As a result, since the amplitude of the waveforms vary throughout the population, it is also more natural to parameterize these source populations by a distance rather than a characteristic strain amplitude. Here we inject them at twenty different distances ranging from 10 to 1000 parsecs. In addition, as was the case for the binary black hole merger waveforms, the supernovae are injected with a fixed inclination relative to the line of sight. In this case, however, we choose an inclination angle corresponding to maximum gravitational-wave emission in the direction of the Earth. This also corresponds to the case of a single gravitational-wave polarization, which we take to be randomly oriented on the sky. Roughly speaking, we expect these assumption to yield a signal amplitude on the order of 2 times larger than the average signal amplitude expected from a randomly oriented population. However, we do not attempt to accurately account for this difference here. Instead, we simply present our results for populations of optimally oriented sources that are isotropically distributed on the sky at a fixed distance. The resulting detection efficiencies for such source populations are shown in Figure 7.17.

In this case, we find that the resulting asymmetric sigmoid fits are of poorer quality than those of the other waveforms that we have considered. This is primarily a result of the large variety of waveforms included in each family of simulated supernovae waveforms, and the fact that a significant number of these waveforms have appreciable spectral content outside of the sensitive frequency band of the LIGO Hanford detectors during the second LIGO science run. As a side effect, the necessary range of signal amplitudes required to fully study the detectability of such bursts was also somewhat underestimated, leading to greater uncertainty in the resulting fit, particularly at high signal amplitudes.

We also recall that the simulated binary black hole merger waveforms and simulated core collapse supernovae waveforms considered here are not known to nearly the same




Figure 7.17: The detection efficiency as a function of distance for isotropic distributions of optimally oriented supernovae. Detection efficiencies are shown for three different families of simulated axisymmetric core collapse supernovae waveforms[42, 43, 44] at a detection threshold corresponding to the most significant observed event. Here, the symmetry axis of the supernovae are all assumed to be perpendicular to the line of sight, such that the maximum signal energy is emitted in the direction of the Earth.

level of accuracy as that of the inspiral phase of binary neutron stars. Nevertheless, the resulting asymmetric sigmoid curves permit a rough order of magnitude measure of the detectability of such sources, as well as providing a benchmark for comparison between different searches and search algorithms.

Finally, we note that the set of simulated waveforms that were injected in the present search were initially developed for use in the triple coincident search of data from the second LIGO science run[92]. Although this choice was primarily one of convenience, it also allows for a straightforward comparison between the two searches. In addition, a portion of the simulated waveforms also overlap with the set of waveforms considered in the triple coincident search of data from the first LIGO science run. However, although the waveforms were the same, the parameters used to describe them vary slightly from those used in present search. In particular, we note that the value of Q as defined by Equation 3.8 is a factor of $\sqrt{2}$ larger than the value of Q as defined in

the triple coincident searches. In addition, the value of σ_t , defined by Equation 3.7a and used here to describe the duration of Gaussian bursts, is a factor of 2 smaller than the duration τ used to describe such bursts in the triple coincident searches. In section 7.9, however, all comparisons to the results of other searches have been performed by first translating such results into the notation of the present work.

7.7 Systematic errors

There are two primary sources of systematic error that effect our estimate of detection efficiency. One of these sources of error is due to our use of simulated waveforms from the triple coincident search. A consequence of this choice is that detection efficiencies have only been estimated for those periods of time during which the Livingston detector was operational, which accounts for only 40 percent of the available observation time. This choice may introduce a bias into our measurement of detection efficiency if, due to nearby human activity, diurnal variations in the sensitivity of the Hanford detectors are correlated with variations in the operational status of the Livingston detector. In order to estimate the bias due to this effect, we have separated the injected signals into two sets: those injected on weekdays between 6am and 6pm, and those injected at other times. By doing so, we find at worst a 20 percent increase in the signal amplitude corresponding to 50 percent detection efficiency as a result of daytime activity. This discrepancy is somewhat ameliorated, however, by taking into account the approximate day to night duty cycle ratios of 1:2 for the present search and 1:5 for the injected waveforms. As a result, we anticipate a worst case bias of 4 percent due to this effect. The other primary source of systematic error in our analysis is due to our uncertainty in calibrating the response of the interferometer to gravitational-waves[134]. Here, the total systematic error due to this uncertainty has been judged to be less than 8 percent[92]. Combining both uncertainties, we conservatively estimate our overall systematic error to be no more than 10 percent.

7.8 Upper limits

In order to determine interpreted upper bounds on the rate of gravitational-wave bursts, we now apply the statistical analysis of section 7.5 to the detection efficiencies from section 7.6. In the process, we also account for the systematic uncertainties identified in the previous section by introducing a 10 percent increase in the reported signal amplitude. The resulting rate limits are displayed in Figure 7.18 through Figure 7.21 at the 90 percent confidence level. For large amplitude waveforms, where the measured detection efficiencies approach unity, the resulting interpreted upper bounds are limited by the observation time of the search and approach a lowest possible bound of 0.086 events per day at the 90 percent confidence level as listed in Table 7.4. At lower signal amplitudes, the resulting upper bounds increase with the decreasing detection efficiency of the search. In Table 7.5 through Table 7.8 we characterize the onset of this transition by the signal amplitude required for 50 percent detection efficiency and find typical transition amplitudes of 10^{-20} to 10^{-19} strain $\text{Hz}^{-1/2}$ for the sinusoidal Gaussian and binary black hole merger waveforms, 10^{-20} to 10^{-18} strain $\text{Hz}^{-1/2}$ for the simple Gaussian waveforms, and 20 to 70 parsecs for the simulated core collapse supernovae waveforms. For the simple Gaussian and sinusoidal Gaussian bursts, we also report the corresponding total energy emitted in the form of gravitational radiation on the assumption of isotropic emission from sources at a galactic distance of 10 kpc according to the prescription of Equation 3.41. In this case, we find that the typical energies required for 50 percent detection efficiency fall in the range from 10^{-5} to 10^{-3} solar masses, depending on waveform.

Table 7.5: *Simple Gaussian bursts.* The characteristic strain amplitude $\|h\|$ and corresponding total energy E required for 50 percent detection efficiency of an isotropic population of simple Gaussian bursts emitting isotropically at a galactic distance of 10 kpc with duration σ_t ranging from 0.05 to 3.0 milliseconds.

Duration [ms]	Amplitude [strain Hz ^{-1/2}]	Energy at 10 kpc [$M_\odot c^2$]
0.05	4.9×10^{-20}	4.1×10^{-3}
0.25	3.2×10^{-20}	7.0×10^{-5}
0.5	4.1×10^{-20}	2.9×10^{-5}
1.25	1.5×10^{-19}	6.1×10^{-5}
2.0	3.6×10^{-19}	1.4×10^{-4}
3.0	1.1×10^{-18}	5.7×10^{-4}

Table 7.6: *Sinusoidal Gaussian bursts.* The characteristic strain amplitude $\|h\|$ and corresponding total energy E required for 50 percent detection efficiency of an isotropic population of sinusoidal Gaussian bursts emitting isotropically at a galactic distance of 10 kpc with central frequencies ϕ ranging from 100 to 849 Hz and a Q of 12.7.

Frequency [Hz]	Amplitude [strain Hz ^{-1/2}]	Energy at 10 kpc [$M_\odot c^2$]
100	9.7×10^{-20}	6.4×10^{-5}
153	4.2×10^{-20}	2.8×10^{-5}
235	2.2×10^{-20}	1.8×10^{-5}
361	2.0×10^{-20}	3.5×10^{-5}
554	2.9×10^{-20}	1.7×10^{-4}
849	4.0×10^{-20}	7.8×10^{-4}

Table 7.7: *Binary black hole mergers.* The characteristic strain amplitude $\|h\|$ required for 50 percent detection efficiency of an isotropic population of simulated equal mass binary black hole mergers[38, 39] at a 60 degree inclination relative to the line of site and with total binary mass ranging from 10 to 90 solar masses.

Total mass [M_{\odot}]	Amplitude [strain $\text{Hz}^{-1/2}$]
10	5.4×10^{-20}
30	2.2×10^{-20}
50	2.0×10^{-20}
70	3.0×10^{-20}
90	4.4×10^{-20}

Table 7.8: *Core collapse supernovae.* The source distance required for 50 percent detection efficiency of an isotropic population of axisymmetric core collapse supernovae with optimal inclination angle relative to the line of site. Characteristic distances are listed for three different families of simulated waveforms[42, 43, 44].

Waveform family	Distance [parsecs]
Zwenger, <i>et al.</i>	62
Dimmelmeier, <i>et al.</i>	56
Ott, <i>et al.</i>	22




Figure 7.18: Upper limits on the rate of gravitational-wave bursts as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of Gaussian bursts. The upper limits are shown for bursts with duration σ_t ranging from 0.05 to 3.0 milliseconds and are reported at a 90 percent confidence level.




Figure 7.19: Upper limits on the rate of gravitational-wave bursts as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of sinusoidal Gaussian bursts. The upper limits are shown for bursts with a Q of 12.7 and central frequency ϕ ranging from 100 to 849 Hz and are reported at a 90 percent confidence level.




Figure 7.20: Upper limits on the rate of gravitational-wave bursts as a function of characteristic strain amplitude $\|h\|$ for isotropic distributions of binary black hole mergers at a 60 degree inclination relative to the line of sight. Upper limits are shown for the simulated binary black hole merger waveforms[38, 39] with total binary masses of 10, 30, 50, 70, and 90 solar masses and are reported at a 90 percent confidence level.




Figure 7.21: Upper limits on the rate of gravitational-wave bursts as a function of distance for isotropic distributions of optimally oriented supernovae. Upper limits are shown for three different families of simulated axisymmetric core collapse supernovae waveforms[42, 43, 44] and are reported at a 90 percent confidence level.

7.9 Comparison of results

We conclude this chapter by comparing the interpreted upper bounds set by the present search with the results of previous searches for gravitational-wave bursts. In particular, we consider the all-sky triple coincident search for bursts from both the first and second LIGO science runs[130, 92], the cumulative observations from 1997 through 2000 by the network of resonant mass detectors comprising the International Gravitational Event Collaboration (IGEC)[30], and the 2001 run of the EXPLORER and NAUTILUS resonant mass detectors[13, 14].

7.9.1 First LIGO science run

The search for bursts in data from the first LIGO science run consisted of 35.5 hours of triple coincident observation time between August 23 and September 9, 2002[130]. For this search, a straightforward comparison with the present search is made possible by a convenient overlap in the choice of waveforms and source populations used to characterize both searches. Here we choose to compare two waveforms, one from each of two families, for which the previous search was most sensitive. These are a simple Gaussian burst with a duration of 0.5 milliseconds and a sinusoidal Gaussian burst with a central frequency of 361 Hz and a Q of 12.7. However, we recall that, although the waveforms are the same, a slightly different notation was used in the triple coincident search. Specifically, a duration of 0.5 milliseconds and a Q of 12.7 in the present search correspond to a duration of 1 millisecond and a Q of 9 in the triple coincident search. In Figure 7.22, we compare the upper bounds achieved by the two searches for both of the waveforms under consideration.

A substantial improvement in both sensitivity and rate is clearly evident between the two searches. This improvement is due to a combination of three factors. As evidenced by Figure 2.7b, the increased sensitivity provided by the present search is in large part




Figure 7.22: A comparison of the upper bounds achieved by the present search with the upper bounds set by triple coincident search of data from the first LIGO science run. The resulting upper bounds are given at the 90 percent confidence level as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of 0.5 millisecond simple Gaussian bursts and 361 Hz sinusoidal Gaussian bursts with a Q of 12.7. The stricter bounds provided by the present search are primarily due to improvements in detector sensitivity between the two science runs as well as the significantly longer observation time of the second science run. However, the difference in sensitivity also reflects improvements in the search algorithms used to identify gravitational-wave bursts.

due to a substantial improvement, by approximately an order of magnitude, in the sensitivity of the instrument between the two science runs. However, a comparison of the characteristic strain amplitude corresponding to 50 percent detection efficiency for the two waveforms indicates an improvement by a factor of 28 for the 0.5 millisecond Gaussian bursts and a factor of 21 for the 361 Hz sinusoidal Gaussian bursts. As a result, some of the increased sensitivity, approximately a factor of 2, appears to be due to improvements in the search algorithm as well. Finally, the significantly lower rate limit set by the present search is a direct consequence of the substantial increase in observation time, by a factor of 18, relative to the first LIGO science run.

7.9.2 Triple coincident search

We also compare the upper bounds achieved by the present double coincident search with that of the corresponding triple coincident search of data from the second LIGO science run.[92]. Again, a straightforward comparison is possible due to the common set of simulated waveforms used to characterize the performance of both searches, and again we choose to compare the performance of the two searches for simple Gaussian bursts of duration 0.5 milliseconds and sinusoidal Gaussian bursts with a central frequency of 361 Hz and a Q of 12.7. The upper bounds achieved by the double and triple coincident searches for these two waveforms are compared in Figure 7.23.




Figure 7.23: A comparison of the upper bounds achieved by the present search with the upper bounds set by corresponding triple coincident search of data from the second LIGO science run. The resulting upper bounds are given at the 90 percent confidence level as a function of characteristic strain amplitude $\|h\|$ for isotropic populations of 0.5 millisecond simple Gaussian bursts and 361 Hz sinusoidal Gaussian bursts with a Q of 12.7. Despite the use of only two detectors, the present search is able to achieve comparable sensitivity to the triple coincident search by coherently searching for bursts with consistent amplitude and phase. Furthermore, the greater observation time afforded by the double coincident search results in an upper bound which is stricter by approximately a factor of three.

We first note that the two searches have very similar sensitivity for the two waveforms

under consideration. This is true despite the higher false event rate expected from the use of only two detectors and the increased potential for coincident false events due to a shared environment. In this case, however, the double coincident search was able to achieve comparable sensitivity by taking advantage of the collocated nature of the Hanford detectors by performing a coherent search for bursts with consistent amplitude and phase. Moreover, as we noted in section 7.1, the double coincident search also affords a significantly greater observation time than the corresponding triple coincident search. As a result, the double coincident search achieves an upper bound on the rate of gravitational-wave bursts that is stricter, by approximately a factor of 3, than the upper bound set by the corresponding triple coincident search.

7.9.3 1997-2000 IGEC observations

We also compare the results of the present search with the cumulative results from the IGEC network of resonant mass detectors for the years between 1997 and 2000[30]. However, while the IGEC result is also presented as an upper bound on the rate of gravitational-wave bursts as a function of signal strength, there are a number of differences that must be taken into account in order to make a meaningful comparison.

We first note that the IGEC network consists of 5 resonant mass detectors with resonant frequencies ranging from 694 to 930 Hz and bandwidths on the order of a few Hz. Due to the narrower bandwidth of resonant mass detectors relative to interferometric detectors, any comparison will then necessarily depend upon the specific choice of waveform used for the comparison. In what follows, we choose to compare the response of the two searches to simple Gaussian bursts with a duration σ_t of 0.05 milliseconds. Although this choice of waveform has a relatively large bandwidth, it is also in approximate agreement with the signal model assumed by the IGEC analysis and corresponds to approximately uniform signal energy over the full range of resonant mass detector frequencies.

We also note that, by convention, the signal amplitude reported by the resonant mass detector community is not the characteristic strain amplitude $\|h\|$. Instead, results are reported in terms of the frequency-domain magnitude of the signal, assuming its spectrum is approximately uniform across the sensitive frequency band of the detector. Specifically, the IGEC reports a characteristic amplitude,

$$\tilde{h}_c = |\tilde{h}(f_0)|, \quad (7.11)$$

where $\tilde{h}(f_0)$ is the standard Fourier transform of the signal defined by Equation 3.1b evaluated at the resonant frequency f_0 of the detector. Given a particular choice of waveform, it is then straightforward to translate between the two definitions. For a Gaussian burst, the relationship follows directly from its Fourier transform,

$$|\tilde{h}(f)| = \|h\| (8\pi\sigma_t^2)^{1/4} \exp(-4\pi^2\sigma_t^2 f^2). \quad (7.12)$$

In particular, evaluating this expression in the frequency range of the resonant mass detectors and for a duration σ_t of 0.05 milliseconds yields,

$$\tilde{h}_c \sim 1.5 \times 10^{-2} \|h\| \quad 694 \leq f \leq 930 \text{ Hz}. \quad (7.13)$$

In addition, the signal amplitudes reported by the IGEC collaboration correspond only to the portion of a gravitational-wave burst's energy that couples into the IGEC detectors, all of which were operated in a similar orientation. As a result, the upper limits presented in Figure 13 of the IGEC paper[30] can be interpreted as an upper bound on the rate of gravitational-wave bursts from populations of sources that are optimally oriented on the sky. In order to compare the IGEC result with the results of the present search, we must therefore determine the response of the present search to such optimally oriented source populations rather than the isotropic source distributions considered in the previous section. Fortunately, for the case of random linear polarization, the necessary information is provided by the amplitudes of the injected

gravitational-wave bursts after taking into account their position on the sky and the corresponding response of the Hanford detectors. In this case, however, the injected amplitudes take on a continuum of values rather than discrete populations. In order to determine the detection efficiency of the present search as a function of injected signal amplitude, we therefore consider populations of bursts that encompass finite ranges of injected amplitudes. The resulting detection efficiencies for populations of optimally oriented Gaussian bursts with duration σ_t of 0.05 milliseconds are shown in Figure 7.24.




Figure 7.24: The detection efficiency as a function of characteristic strain amplitude $\|h\|$ for optimally oriented populations of simple Gaussian bursts with duration σ_t of 0.05 milliseconds. The detection efficiency is shown at a detection threshold corresponding to the coherent normalized energy of the most significant foreground event.

Finally, we note that the upper limits reported by the IGEC are presented at a 95 percent confidence level. We therefore apply the statistical analysis of section 7.5 to the detection efficiencies presented in Figure 7.24 in order to determine an upper bound on the rate of optimally oriented 0.05 millisecond Gaussian bursts at this same confidence level. After accounting for the systematic uncertainties identified in section 7.7, the resulting upper bounds are presented in Figure 7.25 along with the corresponding results from the IGEC analysis.




Figure 7.25: The upper limits on the rate of gravitational-wave bursts as a function of characteristic strain amplitude $\|h\|$ for optimally oriented 0.05 millisecond Gaussian bursts at the 95 percent confidence level. The results of the present search are compared with the corresponding upper limits achieved by the 1997-2000 IGEC analysis. The stricter upper limit achieved by the IGEC analysis at large signal amplitudes is due to the much longer observation time of the IGEC analysis. However, the stricter constraints achieved by the present search at small signal amplitudes is due to the greater sensitivity of the present search. Due to the narrowband nature of resonant mass detectors, such a comparison is necessarily waveform specific. However, the choice of a 0.05 millisecond Gaussian burst is consistent with the signal model assumed by the IGEC analysis. In this case, the benefit of the wider bandwidth of interferometric detectors is readily apparent.

From Figure 7.25, we first note that the significantly longer observation time of the IGEC analysis results in a substantially lower upper bound on the rate of large amplitude gravitational-wave bursts. However, the significantly greater sensitivity of the present search to 0.05 millisecond Gaussian bursts becomes apparent at characteristic strain amplitudes less than approximately 2×10^{-19} strain $\text{Hz}^{-1/2}$. In this case, the difference in sensitivity is primarily due to the wider bandwidth of interferometric detectors. While such a comparison is admittedly waveform specific, we also note that similar results are expected for any non-pathological waveform with sufficient bandwidth to encompass the resonant frequencies of all of the IGEC detectors. In addition, interferometric detectors are also sensitive to the wider class of potential

gravitational-wave signals that fall outside the frequency band of resonant mass detectors. Finally, we note that future LIGO science runs are expected to provide coincident observation times of at least one year in duration and approximately an order of magnitude improvement in sensitivity relative to the present search.

Nevertheless, for narrow bands around their resonant frequencies, the sensitivities of the existing generation of resonant mass detectors are comparable to the sensitivity of the present search. While this is expected to change as the first generation of interferometric detectors approach their design sensitivities, the sensitivities of improved resonant mass detectors are still expected to approach that of interferometric detectors in these narrow frequency bands. As a result, resonant mass detectors will continue to be a valuable complement to interferometric detectors, providing increased detection confidence via the possibility of an independent consistency test for large amplitude candidate gravitational-wave events.

7.9.4 2001 EXPLORER/NAUTILUS observations

Finally, we compare the results of the present search with the results reported by the ROG collaboration for data from the EXPLORER and NAUTILUS resonant mass detectors during 2001[13, 14]. Specifically, we seek to compare the upper limits from the present search with a potential excess of events from the galactic plane reported by the ROG collaboration. However, to do so, we must first consider the sensitivity of the present search to a distribution of bursts in the galactic plane.

In general, such a comparison would best be performed by directly testing the sensitivity of the present search to a simulated distribution of events in the galactic plane. However, due to time and computational constraints, this approach is not considered here and is left for future investigation. Instead, we apply an approximate method proposed by Katsavounidis and developed by Rawlins to modify the results of the present search[135]. In doing so, we are aided by our exclusive use of collocated de-

tector data. This greatly simplifies the necessary modifications, since both of the Hanford detectors necessarily have the same response to arbitrary distributions of sources on the sky.

In what follows, we consider two idealized source distributions. One is a uniformly distributed ring of sources, centered on the Earth, and lying in the galactic plane. The other is a concentration of sources in the direction of the galactic center. In both cases, the sources are also assumed to be distributed with random linear polarization.

In order to estimate the sensitivity of the present search to such distributions, we first determine the average attenuation incurred by a such populations due to the position dependent response of the detectors. For non-isotropic source distributions, this is necessarily a function of sidereal time as the source populations sweep across the sky. Here, we have used a Monte Carlo model in order to estimate the average attenuation incurred by these source distributions as seen by the LIGO Hanford detectors. The resulting average attenuation of both distributions are shown in Figure 7.26 as a function of sidereal hour and are also compared with the average attenuation incurred by an isotropic distribution of sources.

In order to determine the sensitivity of the present search for non-isotropic source distributions, it is also necessary to account for the sidereal exposure of the analyzed data set. Here we simply histogram the analyzed observation time as a function of sidereal hour. The resulting distribution is displayed in Figure 7.27. The noticeable sidereal variation of the resulting histogram is due to diurnal variations in the duty cycle of the two Hanford detectors and the limited extent of the second science run.

Using this information, we then estimate the detection efficiency of our search for the proposed galactic distributions by applying the following procedure.

1. We first consider an isotropic population of sources with random linear polarization and a characteristic strain amplitude $\|h_0\|$. The detection efficiency




Figure 7.26: The mean attenuation due to the antenna response of the Hanford detectors incurred by three different source distributions as a function of sidereal hour: a uniformly distributed ring of sources in the galactic plane (black), a distribution of sources along the line of sight to the galactic center (gray), and an isotropic distribution of sources (dotted). In all three cases, the sources are also assumed to be distributed with random linear polarization.




Figure 7.27: The sidereal hour distribution of analyzed observation time for the double coincident search using the LIGO Hanford detectors. Since the second LIGO science run only lasted for two months, the diurnal variation of the double coincident duty cycle is apparent in the resulting sidereal hour distribution.

$\varepsilon(\|h_0\|)$ has already been determined in section 7.6 for such populations as a function of $\|h_0\|$.

2. From Figure 7.26, we then determine a sidereal dependent correction factor $\eta(\tau)$, which is equal to the ratio of the average attenuation incurred by the desired source distribution to the average attenuation incurred by an isotropic source distribution.
3. Following the table lookup procedure outlined in Figure 7.28, we then determine the sidereal dependent detection efficiency $\varepsilon'(\|h_0\|, \tau)$ for sources with amplitude $\|h_0\|$ that are distributed according to the desired distribution:

$$\varepsilon'(\|h_0\|, \tau) = \varepsilon(\eta(\tau)\|h_0\|). \quad (7.14)$$

4. Finally, we determine the sidereal averaged detection efficiency $\varepsilon'(\|h_0\|)$ for a population of sources with the desired distribution and amplitude $\|h_0\|$, weighted by the sidereal distribution $T(\tau)$ of the analyzed observation time shown in Figure 7.27:

$$\varepsilon'(\|h_0\|) = \frac{\int T(\tau) \varepsilon'(\|h_0\|, \tau) d\tau}{\int T(\tau) d\tau}. \quad (7.15)$$

This process is then repeated for multiple choices of $\|h_0\|$ in order to produce an estimate of the detection efficiency $\varepsilon'(\|h\|)$ for the desired source distribution as a function of $\|h\|$.

In order to perform a meaningful comparison of results, we must also convert the signal amplitude reported by the ROG collaboration into the units used in the present search. As was the case for the comparison with the IGEC results, such a comparison is necessarily waveform dependent. Here we choose to compare the two searches primarily on the assumption of a simple Gaussian burst with duration σ_t of 0.05 milliseconds. Again, this choice of signal has a uniform energy spectrum extending




Figure 7.28: The table lookup procedure used to estimate the approximate sidereal dependent efficiency of the search to a source population with characteristic strain amplitude $\|h_0\|$ that is distributed according to a non-isotropic source distribution. Here $\varepsilon(\|h_0\|)$ is the detection efficiency for an isotropic source distribution, while $\varepsilon'(\|h_0\|, \tau)$ is the approximate detection efficiency of the non-isotropic source distribution at sidereal time τ . The correction factor $\eta(\tau)$ is determined from the ratio of the attenuations of the non-isotropic distribution relative to the isotropic distribution as shown in Figure 7.26.

to approximately 1 kHz, which is in approximate agreement with the signal model assumed by the ROG collaboration[14]. However, in order to evaluate the effect of waveform choice on the resulting comparison, we also consider a sinusoidal Gaussian burst with a central frequency of 849 Hz and a Q of 12.7. Since this latter waveform exhibits a greater concentration of signal energy in the sensitive band of the EXPLORER and NAUTILUS detectors, it therefore provides a more conservative comparison between the two searches. However, we also note that this latter choice of waveform is in poorer agreement with the signal model assumed by the ROG analysis and that the existence of such relatively narrowband bursts near the resonant frequency of the ROG detectors would be somewhat fortuitous.

In converting the ROG observations into the units of the present search, we choose to evaluate the characteristic strain amplitude corresponding to the average of the signal energies reported in Table 3 of reference [13]. In this case we find an average

signal energy of 120 mK with a standard deviation of 53 mK. According to Equation 6 in reference [14], this then corresponds to an average one-sided Fourier transform magnitude of

$$\tilde{h}_c \sim 2.7 \times 10^{-21} \text{ strain Hz}^{-1}. \quad (7.16)$$

Application of Equation 7.12 for a Gaussian burst of duration 0.05 milliseconds and an average resonant frequency of 914 Hz for the EXPLORER and NAUTILUS detectors then yields a characteristic strain amplitude of

$$\|h\| \sim 1.9 \times 10^{-19} \text{ strain Hz}^{-1/2}. \quad (7.17)$$

We also note that a duration σ_t of 0.05 milliseconds is close to the optimal duration of 0.09 milliseconds, which yields the minimum $\|h\|$ for Gaussian waveforms of 1.2×10^{-19} strain Hz^{-1/2}.

In order to determine the corresponding characteristic strain amplitude of 849 Hz sinusoidal Gaussian bursts, we apply a similar procedure. In this case, however, we make use of the approximate relationship,

$$|\tilde{h}(f)| \sim \|h\| \left(\frac{Q^2}{4\pi\phi^2} \right)^{1/4} \exp \left(-\frac{(f - \phi)^2 Q^2}{2\phi^2} \right), \quad (7.18)$$

which follows from the Fourier transform of sinusoidal Gaussian bursts under the assumption of positive frequency and $Q \gtrsim 3$. Evaluating this expression at the average resonant frequency of the ROG detectors for sinusoidal Gaussian bursts with a central frequency ϕ of 849 Hz and a Q of 12.7 then yields a characteristic strain amplitude of approximately

$$\|h\| \sim 6.7 \times 10^{-20} \text{ strain Hz}^{-1/2}. \quad (7.19)$$

Finally, based on the observed sidereal distribution and number of excess events, the ROG collaboration has estimated an associated event rate of 0.55 events per day for such events[14].

We now apply the approximate method of Katsavounidis and Rawlins to the all-sky detection efficiencies from section 7.6 for the two waveforms under consideration. The resulting modified efficiency curves are shown in Figure 7.29 for distributions of sources in the galactic plane and Figure 7.30 for sources in the direction of the galactic center. It is interesting to note that the resulting efficiency curves are very similar to the corresponding all-sky efficiency curves. This similarity with the all-sky result is not unexpected, however, and is primarily due to the wide angular acceptance of interferometric detectors. This fact is also evident in Figure 7.26, where it is apparent that the sidereal averaged sensitivity to source distributions in the galaxy is very similar to that of an isotropic distribution.

Finally, we apply the statistical analysis of section 7.5 to the detection efficiencies presented in Figure 7.29 and Figure 7.30. In this case, however, we choose to evaluate the resulting upper bounds at the 99 percent confidence level in order to provide a more conservative comparison with the potential excess of events postulated by the ROG collaboration. After accounting for the systematic uncertainties identified in section 7.7, the resulting upper bounds are presented in Figure 7.31 and Figure 7.32 along with the approximate location of the 2001 EXPLORER and NAUTILUS results in the units of the present search. Barring the fortuitous existence of narrowband bursts near the resonant frequency of the EXPLORER and NAUTILUS detectors, it is relatively clear that the present search did not detect a similar population of events, even given the approximate nature of this comparison. However, we intend to repeat this analysis with data from the third LIGO science run, and expect to make a more definitive statement on the ROG result by implementing a more accurate treatment of a galactic source population and by taking advantage of an expected improvement in detector sensitivity and a slightly longer observation time.




Figure 7.29: The estimated detection efficiency as a function of characteristic strain amplitude $\|h\|$ for a population of bursts in the galactic plane. Detection efficiencies are shown for simple Gaussian bursts with a duration σ_t of 0.05 milliseconds (black) and sinusoidal Gaussian bursts with a central frequency of 849 Hz and Q of 12.7 (gray) at a detection threshold corresponding to the most significant foreground event.




Figure 7.30: The estimated detection efficiency as a function of characteristic strain amplitude $\|h\|$ for a population of bursts in the direction of the galactic center. Detection efficiencies are shown for simple Gaussian bursts with a duration σ_t of 0.05 milliseconds (black) and sinusoidal Gaussian bursts with a central frequency of 849 Hz and a Q of 12.7 (gray) at a detection threshold corresponding to the most significant foreground event.




Figure 7.31: Upper bounds on the rate gravitational-wave bursts in the galactic plane. Upper limits are shown at a 99 percent confidence level for simple Gaussian bursts with a duration σ_t of 0.05 milliseconds (black) and sinusoidal Gaussian bursts with a central frequency of 849 Hz and a Q of 12.7 (gray), and compared with the postulated excess of events reported by the ROG collaboration[13, 14].




Figure 7.32: Upper bounds on the rate gravitational-wave bursts in the direction of the galactic center. Upper limits are shown at a 99 percent confidence level for simple Gaussian bursts with a duration σ_t of 0.05 milliseconds (black) and sinusoidal Gaussian bursts with a central frequency of 849 Hz and a Q of 12.7 (gray), and compared with the postulated excess of events reported by the ROG collaboration[13, 14].

Chapter 8

Conclusion

In this work, we have considered the problem of identifying gravitational-wave bursts of *a priori* unknown waveform that occur near the sensitivity limit of interferometric gravitational-wave detectors. To confront this problem, we have developed a complete analysis pipeline for the coherent identification of gravitational-wave bursts using coincident data from multiple interferometric detectors. Through extensive testing on simulated interferometer data, we have also performed a comprehensive validation of the proposed method and demonstrated that it is equivalent to the optimal search strategy for the special case of minimum uncertainty waveforms in stationary white noise. Finally, we have applied the proposed method to the search for gravitational-wave bursts with duration less than 1 second and frequency content in the range from 64 to 1024 Hz in data from the second LIGO science run. Although no bursts of likely gravitational-wave origin were identified, we have determined upper bounds on the possible rate of gravitational-wave bursts from isotropic and galactic populations of sources with both abstract and astrophysically motivated waveform. The resulting search achieves a limiting upper bound of 0.086 events per day at the 90 percent confidence level for bursts with characteristic strain amplitude in excess of 10^{-20} to 10^{-19} strain $\text{Hz}^{-1/2}$, depending upon the particular waveform under consideration. A comparison with previous searches then demonstrates that the present search is one

of the most sensitive to date for astrophysically unmodeled bursts of gravitational radiation. In addition, the resulting upper bounds are shown to be inconsistent at above the 99 percent confidence level with a population of sources in the galactic plane proposed to explain an excess of events observed by the ROG collaboration in data from the 2001 run of the EXPLORER and NAUTILUS resonant mass detectors[13, 14].

In what follows, we first summarize the importance of this work, highlighting its novel aspects as well as its astrophysical significance. We then identify a number of possible extensions to the proposed method, including alternative astrophysical searches, and conclude by considering potential prospects from future LIGO science runs.

8.1 Summary

One of the two primary results of this work has been the development of a complete end-to-end data analysis pipeline for the identification of gravitational-wave bursts. While this pipeline has many of its origins in existing methods, it also brings a number of significant improvements to the search for gravitational-wave bursts.

- We have motivated the use of a multiresolution basis of minimum uncertainty waveforms in order to best resolve the time-frequency structure of gravitational-wave bursts. While other multiresolution techniques have been applied to the search for unmodeled bursts[103, 105], the present search is the first to apply the template placement methods of Owen and others[109, 48]. In doing so, we present a method of optimally tiling a well defined signal space with the minimum number of basis functions to guarantee a specified minimum overlap between an arbitrary minimum uncertainty burst and the nearest basis function.
- We have demonstrated the effectiveness of linear prediction to whiten interferometric detector data, which greatly simplifies our resulting statistical analy-

sis. While others have proposed the use of such autoregressive techniques for whitening gravitational-wave data[136, 137], we have also introduced a novel zero-phase implementation that preserves the relative timing of bursts between detectors. In addition, we have also determined the minimum filter order required to ensure that the resulting data are sufficiently white on the time-scales of the subsequent analysis.

- We have introduced the Q transform, a time-frequency projection of the input data stream onto a multiresolution basis of minimum uncertainty waveforms. In addition, we have presented a discrete form of this transform that takes advantage of the computational efficiency of the fast Fourier transform, as well as a normalization that allows for the accurate recovery of total signal energy.
- We have implemented an outlier rejection technique, based on the box-plot formalism of Tukey and first suggested for gravitational-wave data analysis by Cadonati[122], that dramatically reduces the susceptibility of our statistical analysis to transient non-stationarities in interferometric detector data.
- We have performed an extensive validation of the proposed single detector pipeline by demonstrating very good agreement between its performance on minimum uncertainty waveforms in simulated gravitational-wave data and the corresponding performance predicted by a Monte Carlo model based on first principles arguments. In the process, we have demonstrated that the proposed single detector pipeline is equivalent to the optimal search strategy of a template-based matched filter search for the special case of minimum uncertainty bursts in the whitened data stream.
- We have demonstrated a method of coherently combining data from collocated interferometric detectors that maximizes the combined signal to noise ratio of minimum uncertainty bursts while taking into account the frequency dependent difference in sensitivity between detectors. At the same time, we have introduced amplitude and phase consistency tests to reduce the false event rate

associated with transient non-stationarities in a single detector. In combination, these techniques provide a time-frequency compliment to existing time-domain coherence tests[93].

- Finally, we note that the computational efficiency of our analysis pipeline is sufficient to search coincident data from two collocated detectors at a rate approximately 1.75 times faster than real time using a single 2.66 GHz Intel Xeon processor.

The second primary result of the present work has been the application of the proposed analysis pipeline to the search for gravitational-wave bursts in data from the second LIGO science run. In addition to the proposed analysis pipeline, this search differs from previous ones in two unique ways.

- Whereas previous searches have typically relied on networks of geographically distant detectors in order to reduce the expected false alarm rate[130, 92, 30, 13], the present work represents the first search for gravitational-wave bursts in data exclusively from two collocated detectors. Despite the use of only two detectors and the potential for coincident false events due to a shared environment, the present search is able to achieve comparable sensitivity to the corresponding triple coincident search[92] by coherently searching for bursts with both amplitude and phase consistency. In addition, due to the increased observation time afforded by the use of only two detectors, we are able to set a stricter upper bound on the possible rate of gravitational-wave bursts than that of the corresponding triple coincident search.
- We have also demonstrated the efficacy of applying the proposed single detector pipeline to environmental and auxiliary detector data in order to identify significant events of non-gravitational-wave origin. While similar efforts have been attempted in previous searches[130, 92], the present search represents the first time such a veto effort has proven particularly effective.

Although the resulting search identified no bursts of likely gravitational-wave origin, a comparison with previous searches indicates that it nonetheless represents one of the most sensitive searches to date.

- Both the sensitivity and observation time of the present search exhibit a substantial improvement over the that of the triple coincident search for bursts in the first LIGO science run[130]. In particular, we note a factor of 21 to 28 improvement in sensitivity and a factor of 18 improvement in observation time, resulting in a substantially tighter upper bound on the rate of gravitational-wave bursts at all amplitudes.
- Surprisingly, the sensitivity of the present search is comparable to the that of the corresponding triple coincident search[92]. Moreover, the greater observation time of the present search results in a upper bound on the rate of gravitational-wave bursts that is stricter, by approximately a factor of three, than the upper bound set by the triple coincident search.
- We have also performed a comparison of the present search with the cumulative results of the IGEC network of resonant mass detectors between 1997 and 2001[30]. Due to the narrowband sensitivity of resonant mass detectors, such a comparison is necessarily waveform specific. To perform a relatively conservative comparison, we have select a waveform that is consistent with the assumptions of the IGEC analysis. While the substantially longer observation time of the IGEC search results in a stricter upper bound on the rate of large amplitude gravitational-wave bursts, we find that the greater sensitivity of the present search results in a stricter upper bound on the rate of bursts with amplitude smaller than approximately 2×10^{-19} strain $\text{Hz}^{-1/2}$. This improved sensitivity of interferometric detectors becomes all the more apparent for less conservative choices of waveform.
- Lastly, we have compared the sensitivity of the present search with the observed statistical excess of events reported by the ROG collaboration using coincident

data collected by the EXPLORER and NAUTILUS resonant mass detectors during 2001[13, 14]. While our comparison is both approximate and waveform specific, it nonetheless appears to contradict their assertion of gravitational-wave bursts from the galactic disk with a high degree of confidence. However, we intend to repeat this comparison with data from the third LIGO science run, and expect to make a more definitive statement on the ROG collaboration result by performing a more accurate and comprehensive comparison and by taking advantage of an expected improvement in detector sensitivity and a slightly longer observation time.

Finally, in order to evaluate the astrophysical significance of the present work, we recall the sensitivity of our search to the simulated supernovae core collapse and binary black hole merger waveforms reported in section 7.8. We reiterate, however, that such waveforms are not predicted to nearly the same degree of certainty as that of the inspiral of binary neutron stars, nor have we simulated a realistic population of such sources. As a result, the astrophysical import of such simulations are limited. Nevertheless, we may still obtain a rough order of magnitude estimate for the detectable range to such sources. For the present search, we find a detectable range of roughly 50 parsecs for gravitational-waves from axisymmetric core collapse supernovae and roughly 5 Mpc for gravitational-waves from the merger of binary black holes with a total binary mass on the order to 50 solar masses. As a result, we expect that the present search is only sensitive to supernovae within a small fraction of our galaxy, while the detectable range for the merger of binary black holes likely encompasses the local group of galaxies. Unfortunately, given the anticipated population and rate of such events, these distances correspond to a very low probability of detection over the course of the second LIGO science run. As a result, the absence of likely gravitational-wave bursts at the sensitivity of the present search does not provide a test of our present understanding of potential gravitational-wave sources. However, as we note below, future LIGO science runs are expected to provide a significant improvement in both sensitivity and observation time.

8.2 Future investigations

Throughout this document, in addition to the advances summarized above, we have also identified a number of potential extensions to the present work. For reference, we collect them here.

- We first note that our treatment has been somewhat limited for bursts that are poorly localized in the time-frequency plane. In particular, we have only considered the statistical significance of tiles with minimum time-frequency uncertainty. As a result, an improvement in the detectability of poorly localized bursts should be possible if we also consider the combined statistical significance of clusters of time-frequency tiles. However, it is interesting to note that the sensitivity of the present search to simulated supernovae and binary black hole merger waveforms is comparable to that of the corresponding triple coincident search[92]. This is somewhat surprising given the fact that neither waveform family exhibits minimum time-frequency uncertainty, and that the triple coincident analysis pipeline already considers the statistical significance of clusters of time-frequency tiles.
- Similarly, when testing the consistency of coincident events, we have only required amplitude and phase consistency within a single time-frequency tile. By doing so, we have ignored the possibility of inconsistent amplitude or phase in any nearby time-frequency tiles that also exhibit strong normalized energy. A natural extension of our current approach is then to consider the overall amplitude and phase consistency of clusters of jointly significant time-frequency tiles. In addition, such tests should also take into account the expected increase in measurement accuracy with signal strength and apply stricter consistency tests at higher signal to noise ratios. From our discussion of the ten most significant events identified by the present search, it is evident that such tests would substantially reduce the number of false detections by excluding non-

gravitational-wave events that exhibit significant but inconsistent structure in the time-frequency plane.

- We also note that, while our treatment of different interferometers took into account differences in their stationary noise, it did not take into account differences in their non-stationary behavior. In particular, detectors with similar noise spectral density may nonetheless exhibit very different rates of non-gravitational-wave events due to the presence of transient non-stationarities. As a result, it may be possible to decrease the coincident false alarm rate at a given detection efficiency by allowing independent specification of single detector normalized energy thresholds. A systematic method of tuning our analysis pipeline to account for such differences between detectors then offers the possibility of improved performance.
- It has also been pointed out that, when evaluating the statistical significance of clusters of time-frequency tiles or testing for time-frequency coincidence between detectors, a more accurate treatment of the overlap between time-frequency tiles can be obtained by applying the mismatch formalism from section 3.2.2[138].
- Once a candidate gravitational-wave event has been identified, the methods presented in this work can also be applied in order to retrospectively examine such events in greater detail. In particular, the linear prediction methods presented in chapter 4 can be extended in a way that aids in the reconstruction of candidate signals, while the Q transform presented in chapter 5 can be applied at high resolution in order to more accurately characterize the parameters of such events. Such a hierarchical approach then allows for more accurate tests of candidate events without the computational cost of performing such tests on the entire data set.
- In section 5.8, we noted the possibility of searches that target a specific object of interest on the sky. Such a search may be accomplished by an extension of our coherent analysis pipeline to handle data from networks of non-collocated

detectors using a formalism similar to that proposed by Gursel and Tinto for identifying the sky position of candidate gravitational-wave bursts[97]. Potential targets of interest include the center of the Milky Way and Andromeda galaxies, large globular clusters such as 47 Tucanae, and the Virgo cluster of galaxies.

- We also note that, as the LIGO detectors improve in sensitivity, it will be necessary to implement more realistic populations of simulated gravitational-wave bursts in order to more accurately evaluate the astrophysical significance of our search. Such populations should not only include truly random distributions of source inclination, but should also conform to the expected mass distribution of nearby galaxies.
- Lastly, although this work has focused on the detection of gravitational-wave bursts, we note that many of the techniques developed here are also applicable to the more general problem of detecting transient signals in the presence of additive colored noise. As a result, many of these methods may also be relevant to problems other than that of gravitational-wave detection.

Finally, we consider the prospects of applying the proposed algorithm to future LIGO science runs.

As of the writing of this document, the LIGO project has just completed data taking for its fourth science run. In the interim, the LIGO detectors have also undergone two periods of commissioning activities resulting in significant improvements in both their sensitivity and operational duty cycle, as well as significantly reduced environmental coupling to seismic noise, acoustic noise, and radio frequency interference. In addition, improvements in calibrating the response of the LIGO interferometers are expected to allow for stricter tests on both amplitude and phase consistency. As a result, both the third and fourth LIGO science runs offer the prospect of significantly improved sensitivity over the results of the present search. In particular, preliminary studies

indicate that the third LIGO science run may provide a factor of 2 to 5 improvement in sensitivity and a 20 percent increase in observation time compared to the present search, with the greatest improvement in sensitivity for signals around 100 Hz. For the fourth LIGO science run, early indications are that it may already yield sensitivities within a factor of a 2 to 3 from the LIGO design sensitivity, although the available observation time will be slightly shorter than previous science runs due its planned duration of only one month. Finally, we note that goal of the fifth LIGO science run, tentatively scheduled to commence in the Fall of 2005, is to provide at least one year of coincident data at the LIGO design sensitivity[139].

Given the significant improvements in sensitivity since the second LIGO science run, as well as the much longer observation times that we can expect in the near future, we can expect future searches to produce significantly tighter bounds on the rate of gravitational-wave bursts. In particular, it is interesting to note that the initial LIGO design sensitivity represents an improvement by a factor of roughly 20 compared to the sensitivity of the present search. Since this directly corresponds to a similar increase in the detectable range to sources, we estimate that LIGO may soon be sensitive to binary black hole mergers beyond the Virgo cluster of galaxies. In addition, slightly more optimistic assumptions of the gravitational-wave energy emitted by core collapse supernovae indicate that such events may soon be detectable over a significant fraction of our galaxy. Although the prospects for the detection of supernovae core collapse and binary black hole mergers are significantly improved, they are still not guaranteed at these sensitivity levels and observation times. Nevertheless, the potential benefit of a detection is compelling, and even in the absence of a detection, the expected sensitivities and observation times are such that it should be possible to begin to constrain more optimistic models of binary compact object formation and evolution[140].

Finally, we note that a number of other interferometric gravitational-wave observatories are currently undergoing commissioning[33, 35, 34], and offer the possibility of applying the algorithms presented in this work to data from even larger networks

of detectors. In addition, beyond the initial generation of interferometric detectors, advanced ground based detectors offer a further order of magnitude improvement in sensitivity, virtually guaranteeing the detection of gravitational waves within the next decade, while planned space based detectors offer similar sensitivity in frequency bands that are not accessible to ground based detectors.

Bibliography

- [1] R. A. Hulse and J. H. Taylor, “Discovery of a pulsar in a binary system,” *Astrophys. J.*, vol. 195, pp. L51–L53, 1975, 1975ApJ.195L.51H.
- [2] J. H. Taylor, “Binary pulsars and relativistic gravity,” *Rev. Mod. Phys.*, vol. 66, p. 711, 1994, 1994RvMP.66.711T.
- [3] A. Wolszczan, “A nearby 37.9-ms radio pulsar in a relativistic binary system,” *Nature*, vol. 350, pp. 688–690, April 1991, 1991Natur.350.688W.
- [4] I. H. Stairs, S. E. Thorsett, J. H. Taylor, and A. Wolszczan, “Studies of the relativistic binary pulsar psr b1534+12: I. timing analysis,” *Astrophys. J.*, vol. 581, pp. 501–508, 2002, astro-ph/0208357.
- [5] S. B. Anderson, P. W. Gorham, S. R. Kulkarni, T. A. Prince, and A. Wolszczan, “Discovery of two radio pulsars in the globular cluster m15,” *Nature*, vol. 346, pp. 42–44, July 1990, 1990Natur.346.42A.
- [6] T. A. Prince, S. B. Anderson, S. R. Kulkarni, and A. Wolszczan, “Timing observations of the 8 hour binary pulsar 2127 + 11c in the globular cluster m15,” *Astrophys. J. Letters*, vol. 374, pp. L41–L44, June 1991, 1991ApJ.374L.41P.
- [7] M. Burgay *et al.*, “An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system,” *Nature.*, vol. 426, pp. 531–533, 2003, astro-ph/0312071.

- [8] V. Kalogera *et al.*, “The cosmic coalescence rates for double neutron star binaries,” *Astrophys. J. Letters*, vol. 601, pp. L179–L182, 2004, astro-ph/0312101. Erratum published in Vol. 614, Iss. 2, pp. L137–L138.
- [9] J. Weber, “Detection and generation of gravitational waves,” *Physical Review*, vol. 117, p. 306, 1960, 1960PhysRev.117.306.
- [10] R. Weiss, “Electromagnetically coupled broadband gravitational antenna,” *Quarterly Progress Report of the Research Laboratory of Electronics*, vol. 105, p. 54, April 1972.
- [11] <http://www.ligo.caltech.edu/>.
- [12] A. Abramovici *et al.*, “Ligo: The laser interferometer gravitational wave observatory,” *Science*, vol. 256, pp. 325–333, 1992.
- [13] P. Astone *et al.*, “Study of the coincidences between the gravitational wave detectors explorer and nautilus in 2001,” *Class. Quantum Grav.*, vol. 19, pp. 5449–5463, 2002, gr-qc/0210053.
- [14] E. Coccia, F. Dubath, and M. Maggiore, “On the possible sources of gravitational wave bursts detectable today,” *Phys. Rev. D*, vol. 70, p. 084010, 2004, gr-qc/0405047.
- [15] P. R. Saulson, *Fundamentals of interferometric gravitational wave detectors*. World Scientific Publishing, 1994.
- [16] K. Thorne, “Gravitational radiation,” in *300 Years of Gravitation* (S. Hawking and W. Israel, eds.), Cambridge University Press, 1987.
- [17] C. Cutler and K. S. Thorne, “An overview of gravitational-wave sources.” review article, 2002, gr-qc/0204090.
- [18] E. E. Flanagan and S. A. Hughes, “The basics of gravitational wave theory,” *submitted to New Journal of Physics*, vol. special issue: Spacetime 100 Years Later, 2005, gr-qc/0501041.

- [19] N. Andersson and K. D. Kokkotas, “Gravitational-wave astronomy: The high-frequency window,” *Lectures presented at the 2nd Aegean Summer School on the Early Universe*, 2004, gr-qc/0403087.
- [20] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. W. H. Freeman, 1970.
- [21] S. Weinberg, *Gravitation and cosmology: principles and applications of the general theory of relativity*. John Wiley and Sons, 1972.
- [22] B. F. Schutz, *A first course in general relativity*. Cambridge University Press, 1990.
- [23] B. Schutz, “Gravitational waves on the back of an envelope,” *American Journal of Physics*, vol. 52, pp. 412–419, May 1984, 1984AmJPh.52.412S.
- [24] <http://gravity.phys.lsu.edu/>.
- [25] <http://www.auriga.lnl.infn.it/>.
- [26] <http://www.roma1.infn.it/rog/explorer/explorer.html>.
- [27] <http://www.lnf.infn.it/esperimenti/rog/NAUTILUS/>.
- [28] <http://www.gravity.pd.uwa.edu.au/>.
- [29] <http://igec.lnl.infn.it/>.
- [30] P. Astone *et al.*, “Methods and results of the igec search for burst gravitational waves in the years 1997-2000,” *Phys. Rev. D*, vol. 68, p. 022001, 2003.
- [31] <http://www.roma1.infn.it/rog/>.
- [32] L. Cadonati *et al.*, “The auriga-ligo joint burst search,” *submitted to Class. Quantum Grav.*, 2005.
- [33] <http://www.geo600.uni-hannover.de/>.

- [34] <http://www.virgo.infn.it/>.
- [35] <http://tamago.mtk.nao.ac.jp/>.
- [36] .
- [37] E. E. Flanagan and S. A. Hughes, “Measuring gravitational waves from binary black hole coalescences. i: Signal to noise for inspiral, merger, and ringdown,” *Phys. Rev.*, vol. D57, pp. 4535–4565, 1998, gr-qc/9701039.
- [38] J. G. Baker, M. Campanelli, C. O. Lousto, and R. Takahashi, “Modeling gravitational radiation from coalescing binary black holes,” *Phys. Rev.*, vol. D65, p. 124012, 2002, astro-ph/0202469.
- [39] J. Baker, M. Campanelli, C. O. Lousto, and R. Takahashi, “Coalescence remnant of spinning binary black holes,” *Phys. Rev.*, vol. D69, p. 027505, 2004, astro-ph/0305287.
- [40] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, “How massive single stars end their life,” *Astrophys. J.*, vol. 591, pp. 288–300, 2003, astro-ph/0212469.
- [41] W. Arnett, J. Bahcall, R. Kirschner, and S. Woosley, “Supernova sn1987a,” *Annual Reviews of Astronomy and Astrophysics*, vol. 27, p. 629, 1989.
- [42] T. Zwerger and E. Mueller, “Dynamics and gravitational wave signature of axisymmetric rotational core collapse,” *Astron. Astrophys.*, vol. 320, pp. 209–227, 1997.
- [43] H. Dimmelmeier, J. A. Font, and E. Muller, “Relativistic simulations of rotational core collapse. ii. collapse dynamics and gravitational radiation,” *Astron. Astrophys.*, vol. 393, pp. 523–542, 2002, astro-ph/0204289.
- [44] C. D. Ott, A. Burrows, E. Livne, and R. Walder, “Gravitational waves from axisymmetric, rotational stellar core collapse,” *Astrophys. J.*, vol. 600, pp. 834–864, 2004, astro-ph/0307472.

- [45] S. A. Teukolsky, “Rotating black holes - separable wave equations for gravitational and electromagnetic perturbations,” *Phys. Rev. Lett.*, vol. 29, pp. 1114–1118, 1972.
- [46] E. W. Leaver, “An analytic representation for the quasi-normal modes of kerr black holes,” *Proceedings of the Royal Society of London, Series A (Mathematical and Physical Sciences)*, vol. 402, pp. 285–298, 1985.
- [47] F. Echeverria, “Gravitational wave measurements of the mass and angular momentum of a black hole,” *Phys. Rev.*, vol. D40, pp. 3194–3203, 1989.
- [48] J. D. E. Creighton, “Search techniques for gravitational waves from black-hole ringdowns,” *Phys. Rev.*, vol. D60, p. 022001, 1999, gr-qc/9901084.
- [49] R. Adhikari, *Sensitivity and noise analysis of 4 km laser interferometric gravitational wave antennae*. PhD dissertation, Massachusetts Institute of Technology, Department of Physics, 2004.
- [50] Y. Tsunesada *et al.*, “On detection of black hole quasi-normal ringdowns: Detection efficiency and waveform parameter determination in matched filtering,” 2004, gr-qc/0410037.
- [51] B. Zhang and P. Meszaros, “Gamma-ray bursts: Progress, problems & prospects,” *Int. J. Mod. Phys.*, vol. A19, pp. 2385–2472, 2004, astro-ph/0311321.
- [52] L. S. Finn, S. D. Mohanty, and J. D. Romano, “Detecting an association between gamma ray and gravitational wave bursts,” *Phys. Rev.*, vol. D60, p. 121101, 1999, gr-qc/9903101.
- [53] S. D. Mohanty, S. Marka, R. Rahkola, S. Mukherjee, I. Leonor, R. Frey, J. Cannizzo, and J. Camp, “Gamma ray bursts and gravitational waves: triggered search strategy in the ligo science runs,” *Class. Quantum Grav.*, vol. 21, pp. S765–S774, 2004.

- [54] S. D. Mohanty, S. Marka, R. Rahkola, S. Mukherjee, I. Leonor, R. Frey, J. Cannizzo, and J. Camp, “Search algorithm for a gravitational wave signal in association with gamma ray burst grb030329 using the ligo detectors,” *Class. Quant. Grav.*, vol. 21, pp. S1831–S1838, 2004, gr-qc/0407057.
- [55] B. Abbott *et al.*, “A search for gravitational waves associated with the gamma ray burst grb030329 using the ligo detectors,” *To be submitted to Phys. Rev. D*, 2005, gr-qc/0501068.
- [56] P. Astone *et al.*, “Search for correlation between grb’s detected by beposax and gravitational wave detectors explorer and nautilus,” *Phys. Rev.*, vol. D66, p. 102002, 2002, astro-ph/0206431.
- [57] P. Astone *et al.*, “Cumulative analysis of the association between the gravitational wave detectors nautilus and explorer and the gamma ray bursts detected by batse and beposax,” 2004, astro-ph/0206431.
- [58] A. Vilenkin and E. P. S. Shellard, *Cosmic strings and other topological defects*. Cambridge University Press, 2000.
- [59] T. Damour and A. Vilenkin, “Gravitational wave bursts from cosmic strings,” *Phys. Rev. Lett.*, vol. 85, pp. 3761–3764, 2000, gr-qc/0004075.
- [60] T. Damour and A. Vilenkin, “Gravitational wave bursts from cusps and kinks on cosmic strings,” *Phys. Rev.*, vol. D64, p. 064008, 2001, gr-qc/0104026.
- [61] T. Damour and A. Vilenkin, “Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows,” *Phys. Rev.*, vol. D71, p. 063510, 2005, hep-th/0410222.
- [62] X. Siemens and K. D. Olum, “Cosmic string cusps with small-scale structure: Their forms and gravitational waveforms,” *Phys. Rev.*, vol. D68, p. 085017, 2003, gr-qc/0307113.

- [63] D. Chakrabarty *et al.*, “Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars,” *Nature*, vol. 424, pp. 42–44, 2003, astro-ph/0307029.
- [64] R. Wijnands *et al.*, “Quasi-periodic x-ray brightness fluctuations in an accreting millisecond pulsar,” *Nature*, vol. 424, pp. 44–47, 2003, astro-ph/0307123.
- [65] L. Bildsten, “Gravitational radiation and rotation of accreting neutron stars,” *Astrophys. J.*, vol. 501, p. L89, 1998, astro-ph/9804325.
- [66] L. Bildsten, “Arresting accretion torques with gravitational radiation,” 2002, astro-ph/0212004.
- [67] B. Abbott *et al.*, “Limits on gravitational wave emission from selected pulsars using ligo data,” *Phys. Rev. Lett.*, vol. 94, p. 181103, 2005, gr-qc/0410007.
- [68] <http://einstin.phys.uwm.edu/>.
- [69] P. Astone *et al.*, “Search for periodic gravitational wave sources with the explorer detector,” *Phys. Rev.*, vol. D65, p. 022001, 2002, gr-qc/0011072.
- [70] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Krolak, “Data analysis of gravitational-wave signals from spinning neutron stars. iv: An all-sky search,” *Phys. Rev.*, vol. D65, p. 042003, 2002, gr-qc/0012108.
- [71] E. Mauceli, M. P. McHugh, W. O. Hamilton, W. W. Johnson, and A. Morse, “Search for periodic gravitational radiation with the allegro gravitational wave detector,” 2000, gr-qc/0007023.
- [72] B. Allen, “The stochastic gravity-wave background: Sources and detection,” *Proceedings of the Les Houches School on Astrophysical Sources of Gravitational Waves*, 1996, gr-qc/9604033.
- [73] R. A. Battye and E. P. S. Shellard, “Primordial gravitational waves: A probe of the very early universe,” 1996, astro-ph/9604059.

- [74] C. L. Bennett *et al.*, “First year wilkinson microwave anisotropy probe (wmap) observations: Preliminary maps and basic results,” *Astrophys. J. Suppl.*, vol. 148, p. 1, 2003, astro-ph/0302207.
- [75] B. Allen and J. Romano, “Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities,” *Phys. Rev.*, vol. D59, p. 102001, 1999, gr-qc/9710117.
- [76] N. Christensen, “Optimal detection strategies for measuring the stochastic gravitational radiation background with laser interferometric antennas,” *Phys. Rev.*, vol. D55, pp. 448–454, 1997.
- [77] B. Abbott *et al.*, “Analysis of first ligo science data for stochastic gravitational waves,” *Phys. Rev.*, vol. D69, p. 122004, 2004, gr-qc/0312088.
- [78] P. f. t. L. S. C. Fritschel, “The search for a stochastic gravitational wave background.” Presented at the April 2004 meeting of the American Physical Society, May 2004, <http://www.ligo.caltech.edu/docs/G/G040236-02/G040236-02.pdf>.
- [79] E. W. Kolb and M. S. Turner, “The early universe,” Redwood City, USA: Addison-Wesley (1990) 547 p. (Frontiers in physics, 69).
- [80] J. W. Armstrong, L. Iess, P. Tortora, and B. Bertotti, “Stochastic gravitational wave background: Upper limits in the 10^{-6} -hz 10^{-3} -hz band,” *Astrophys. J.*, vol. 599, pp. 806–813, 2003.
- [81] V. M. Kaspi, J. H. Taylor, and M. F. Ryba, “High - precision timing of millisecond pulsars. 3: Long - term monitoring of psrs b1855+09 and b1937+21,” *Astrophys. J.*, vol. 428, p. 713, 1994.
- [82] M. P. McHugh, G. Zalamansky, F. Vernotte, and E. Lantz, “Pulsar timing and the upper limits on a gravitational wave background: A bayesian approach,” *Phys. Rev.*, vol. D54, pp. 5993–6000, 1996.

- [83] C. L. Bennett *et al.*, “4-year cobe dmr cosmic microwave background observations: Maps and basic results,” *Astrophys. J.*, vol. 464, pp. L1–L4, 1996, astro-ph/9601067.
- [84] R. K. Sachs and A. M. Wolfe, “Perturbations of a cosmological model and angular variations of the microwave background,” *Astrophys. J.*, vol. 147, pp. 73–90, 1967.
- [85] D. Gabor, “Theory of communication,” *Journal of the Institution of Electrical Engineers*, vol. 93, no. 26, pp. 429–457, 1946.
- [86] K. Thorne, “Gravitational waves,” in *Proceedings of the Snowmass 95 Summer Study on Particle and Nuclear Astrophysics*, World Scientific, 1995.
- [87] C. W. Helstrom, *Statistical theory of signal detection*. Pergamon Press, second ed., 1968.
- [88] B. Abbott *et al.*, “Analysis of ligo data for gravitational waves from binary neutron stars,” *Phys. Rev.*, vol. D69, p. 122001, 2004, gr-qc/0308069.
- [89] B. Abbott *et al.*, “Search for gravitational waves from galactic and extra-galactic binary neutron stars,” *to be submitted to Phys. Rev. D*, 2005, gr-qc/0505041.
- [90] B. Abbott *et al.*, “Search for gravitational waves from primordial black hole binary coalescences in the galactic halo,” *to be submitted to Phys. Rev. D*, 2005, gr-qc/0505042.
- [91] B. Abbott *et al.*, “Search for gravitational waves from binary black hole inspirals in ligo data,” *to be submitted to Phys. Rev. D*, 2005.
- [92] B. Abbott *et al.*, “Upper limits on gravitational wave bursts in ligo’s second science run,” *to be submitted to Phys. Rev. D*, 2005, gr-qc/0505029.
- [93] L. Cadonati, “Coherent waveform consistency test for ligo burst candidates,” *Class. Quantum Grav.*, vol. 21, pp. S1695–S1703, 2004, gr-qc/0407031.

- [94] L. Cadonati and S. Marka, “Corrpower: a cross-correlation based algorithm for triggered and untriggered gravitational-wave burst searches,” *submitted to Class. Quantum Grav.*, 2005.
- [95] L. Wen and B. F. Schutz, “Coherent network detection of gravitational waves: The redundancy veto,” *submitted to Class. Quantum Grav.*, 2005.
- [96] M. Rakhmanov and S. Klimenko, “Cross-correlation analysis of signals for burst searches with misaligned gravitational-wave detectors,” *submitted to Class. Quantum Grav.*, 2005.
- [97] Y. Gursel and M. Tinto, “Solution of the inverse problem for gravitational wave bursts,” *Phys. Rev. D*, vol. 40, p. 3884, 1989.
- [98] E. E. Flanagan and S. A. Hughes, “Measuring gravitational waves from binary black hole coalescences. ii: The waves’ information and its extraction, with and without templates,” *Phys. Rev.*, vol. D57, pp. 4566–4587, 1998, gr-qc/9710129.
- [99] N. Arnaud *et al.*, “Comparison of filters for detecting gravitational wave bursts in interferometric detectors,” *Phys. Rev.*, vol. D67, p. 062004, 2003, gr-qc/0210098.
- [100] T. Pradier *et al.*, “An efficient filter for detecting gravitational wave bursts in interferometric detectors,” *Phys. Rev.*, vol. D63, p. 042002, 2001, gr-qc/0010037.
- [101] N. Arnaud, F. Cavalier, M. Davier, and P. Hello, “Detection of gravitational wave bursts by interferometric detectors,” *Phys. Rev.*, vol. D59, p. 082002, 1999, gr-qc/9812015.
- [102] J. McNabb, M. Ashley, L. S. Finn, E. Rotthoff, A. Stuver, T. Summerscales, P. Sutton, M. Tibbits, K. Thorne, and K. Zaleski, “Overview of the blocknormal event trigger generator,” *Class. Quantum Grav.*, vol. 21, pp. S1705–S1710, 2004.
- [103] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and E. E. Flanagan, “An excess power statistic for detection of burst sources of gravitational radiation,” *Phys. Rev.*, vol. D63, p. 042003, 2001, gr-qc/0008066.

- [104] J. Sylvestre, “Time-frequency detection algorithm for gravitational wave bursts,” *Phys. Rev.*, vol. D66, p. 102004, 2002, gr-qc/0210043.
- [105] S. Klimenko and G. Mitselmakher, “A wavelet method for detection of gravitational wave bursts,” *Class. Quantum Grav.*, vol. 21, pp. S1819–S1830, 2004.
- [106] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, “Multiresolution techniques for the detection of gravitational-wave bursts,” *Class. Quantum Grav.*, vol. 21, pp. S1809–S1818, 2004, gr-qc/0412119.
- [107] I. Daubechies, *Ten Lectures on Wavelets*. SIAM, 1992.
- [108] S. Mallat, *Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [109] B. J. Owen and B. S. Sathyaprakash, “Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement,” *Phys. Rev.*, vol. D60, p. 022002, 1999, gr-qc/9808076.
- [110] S. V. Dhurandhar and B. S. Sathyaprakash, “Choice of filters for the detection of gravitational waves from coalescing binaries. 2. detection in colored noise,” *Phys. Rev.*, vol. D49, pp. 1707–1722, 1994.
- [111] B. S. Sathyaprakash and S. V. Dhurandhar, “Choice of filters for the detection of gravitational waves from coalescing binaries,” *Phys. Rev.*, vol. D44, pp. 3819–3834, 1991.
- [112] N. Arnaud *et al.*, “An elliptical tiling method to generate a 2-dimensional set of templates for gravitational wave search,” *Phys. Rev.*, vol. D67, p. 102003, 2003, gr-qc/0211064.
- [113] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, fourth ed., 2002.
- [114] J. Makhoul, “Linear prediction: A tutorial review,” *Proc. IEEE*, vol. 63, 1975.
- [115] W. Press *et al.*, *Numerical Recipes in C*. Cambridge University Press, 1992.

- [116] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Prentice-Hall, first ed., 1989.
- [117] R. G. Stockwell, L. Mansinha, and R. P. Lowe, “Localization of the complex spectrum: the s-transform,” *IEEE Transactions on Signal Processing*, vol. 44, no. 4, pp. 998–1001, 1996.
- [118] J. C. Brown, “Calculation of a constant q spectral transform,” *The Journal of the Acoustical Society of America*, vol. 89, no. 1, pp. 425–434, 1991.
- [119] F. J. Harris, “On the use of windows for harmonic analysis with the discrete fourier transform,” *Proceedings of the IEEE*, vol. 66, January 1978.
- [120] B. Allen, “A chi-squared time-frequency discriminator for gravitational wave detection.” review article, 2004, gr-qc/0405045.
- [121] A. Papoulis, *The Fourier Integral and its Applications*. McGraw-Hill, 1962.
- [122] L. Cadonati personal communication.
- [123] D. C. Hoaglin, F. Mosteller, and J. W. Tukey, . *Understanding Robust and Exploratory Data Analysis*. Wiley, 1983.
- [124] J. Sylvestre, “Optimal generalization of power filters for gravitational wave bursts, from single to multiple detectors,” *Phys. Rev.*, vol. D68, p. 102005, 2003, gr-qc/0308062.
- [125] N. Arnaud *et al.*, “Detection in coincidence of gravitational wave bursts with a network of interferometric detectors. i: Geometric acceptance and timing,” *Phys. Rev.*, vol. D65, p. 042004, 2002, gr-qc/0107081.
- [126] N. Arnaud *et al.*, “Coincidence and coherent data analysis methods for gravitational wave bursts in a network of interferometric detectors,” *Phys. Rev.*, vol. D68, p. 102001, 2003, gr-qc/0307100.

- [127] L. Blackburn *et al.*, “A first comparison of search methods for gravitational wave bursts using ligo and virgo simulated data,” *submitted to Class. Quantum Grav.*, 2005, gr-qc/0504060.
- [128] L. Blackburn *et al.*, “A first comparison between ligo and virgo inspiral search pipelines,” *submitted to Class. Quantum Grav.*, 2005, gr-qc/0504050.
- [129] P. R. Brady, J. D. E. Creighton, and A. G. Wiseman, “Upper limits on gravitational-wave signals based on loudest events,” *Class. Quantum Grav.*, vol. 21, pp. S1775–S1781, 2004, gr-qc/0405044.
- [130] B. Abbott *et al.*, “First upper limits from ligo on gravitational wave bursts,” *Phys. Rev.*, vol. D69, p. 102001, 2004, gr-qc/0308050.
- [131] <http://gallatin.physics.lsa.umich.edu/~keithr/S2DQ/>.
- [132] <http://www.ldas-cit.ligo.caltech.edu/>.
- [133] <http://www.cs.wisc.edu/condor/>.
- [134] G. Gonzalez, M. Landry, B. O’Reilly, and H. Radkins, “Calibration of the ligo detectors for s2,” 2004, <http://www.ligo.caltech.edu/docs/T/T040060-00.pdf>.
- [135] K. Rawlins, “Antenna patterns and the galactic plane,” *unpublished*.
- [136] E. Cuoco *et al.*, “Noise parametric identification and whitening for ligo 40- meter interferometer data,” *Phys. Rev.*, vol. D64, p. 122002, 2001, gr-qc/0104071.
- [137] A. C. Searle, S. M. Scott, and D. E. McClelland, “Spectral line removal in the ligo data analysis system (ldas),” *Classical and Quantum Gravity*, vol. 20, no. 17, pp. S721–S730, 2003.
- [138] L. Blackburn personal communication.
- [139] S. Whitcomb, “State of the ligo laboratory,” 2005, <http://www.ligo.caltech.edu/docs/G/G050124-00/G050124-00.pdf>.

- [140] P. Nutzman, V. Kalogera, L. S. Finn, C. Hendrickson, and K. Belczynski, “Gravitational waves from extragalactic inspiraling binaries: Selection effects and expected detection rates,” *Astrophys. J.*, vol. 612, pp. 364–374, 2004, astro-ph/0402091.
- [141] S. Haykin, *Adaptive Filter Theory*. Prentice-Hall, fourth ed., 2002.
- [142] B. Abbott *et al.*, “Detector description and performance for the first coincidence observations between ligo and geo,” *Nucl. Instrum. Meth.*, vol. A517, pp. 154–179, 2004, gr-qc/0308043.
- [143] <http://www.ligo.org/>.
- [144] <http://www.anu.edu.au/Physics/ACIGA/>.
- [145] D. Sigg, “Gravitational waves,” in *Neutrinos in physics and astrophysics, from 10^{-33} to 10^{28} cm (TASI 98)* (P. Langacker, ed.), World Scientific, 1998, <http://www.ligo.caltech.edu/docs/P/P980007-00.pdf>.
- [146] J. H. Taylor and J. M. Weisberg, “Further experimental tests of relativistic gravity using the binary pulsar psr 1913+16,” *Astrophys. J.*, vol. 345, pp. 434–450, 1989.
- [147] A. Buonanno, M. Maggiore, and C. Ungarelli, “Spectrum of relic gravitational waves in string cosmology,” *Phys. Rev.*, vol. D55, pp. 3330–3336, 1997, gr-qc/9605072.
- [148] A. Drake, *Fundamentals of Applied Probability Theory*. McGraw-Hill, 1900.
- [149] B. Abbott *et al.*, “Setting upper limits on the strength of periodic gravitational waves using the first science data from the geo 600 and ligo detectors,” *Phys. Rev.*, vol. D69, p. 082004, 2004, gr-qc/0308050.
- [150] B. J. Owen, “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing,” *Phys. Rev.*, vol. D53, pp. 6749–6761, 1996, gr-qc/9511032.

[151] S. Klimenko, I. Yakushin, M. Rakhmanov, and G. Mitselmakher, “Performance of the waveburst algorithm on ligo data,” *Class. Quantum Grav.*, vol. 21, pp. S1685–S1694, 2004.

[152] http://www.ligo.caltech.edu/~lazz/distribution/LSC_Data/.