ABSTRACT

Title of dissertation: LOOC UP:
SEEKING OPTICAL COUNTERPARTS
TO GRAVITATIONAL WAVE SIGNALS

Jonah Kanner, Doctor of Philosophy, 2011

Dissertation directed by: Professor Peter Shawhan
Department of Physics

Far-reaching investigations into astrophysics, precision measurements of cosmological parameters, and tests of fundamental physics are expected to be enabled through observations of gravitational wave (GW) signals. The gravitational waveform traces the bulk motion of matter in distant sources, and so observations of GWs would reveal the details of interesting astrophysical mechanisms that are inaccessible by any other means. Currently, the most sensitive GW instruments in the world are the Laser Interferometer Gravitational Wave Observatory (LIGO) and the Virgo observatory. These kilometer scale observatories make use of sensitive optics to record miniscule changes in the space-time metric induced by GW signals, and are located at three sites in the United States and Europe.

Some models for sources of observable, transient GW signals predict that an electromagnetic (EM) counterpart will accompany the GW signal. These counterparts would be identifiable as transients which fade over the course of hours or days.

Finding the EM counterpart to a GW signal would present significant benefits. The EM counterpart could confirm the astrophysical origin of the signal, and would also enhance the scientific investigations possible with the observation.

This work describes the first prompt search for EM counterparts to GW event candidates. For the search, carried out between December 2009 and October 2010, event candidates from the LIGO/Virgo network were identified with latencies of only a few minutes. The sky position of each potential source was estimated, and was delivered to a collection of radio, optical, and x-ray telescopes with roughly thirty minutes of latency. The telescopes then observed the estimated source position, in an attempt to discover an EM counterpart. The low-latency GW data analysis, the methods for estimating source positions, and the observing strategy that guides telescopes based on GW data are all novel features of the search. The ability of the LIGO/Virgo network to correctly localize sources on the sky was studied using a Monte Carlo simulation. These developments lay the groundwork for similar searches in the future with the next-generation GW detectors Advanced LIGO and Advanced Virgo.

LOOC UP: SEEKING OPTICAL COUNTERPARTS TO GRAVITATIONAL WAVE SIGNALS

by

Jonah Brian Kanner

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2011

Advisory Committee:
Professor Peter Shawhan, Chair/Advisor
Professor Jordan Goodman
Professor Ted Jacobson
Professor M. Coleman Miller
Professor Christopher Reynolds

Acknowledgments

Working on this search has been an extremely humbling experience. The progress described in this dissertation was only possible because of the incredible support provided by the gravitational wave community. The words "EM follow-up" and "low-latency analysis" became lightning rods for impressive levels of effort, focus, dedication, sacrifice, and bravery within the LIGO and Virgo collaborations. Sometimes, I don't feel like I constructed the low-latency search, so much as I set up these lightning rods and then got out of the way. I have dreaded writing this acknowledgement, because so many people have made sacrifices in the name of low-latency, that making a complete list would be nearly impossible. Though I have no way of making a complete list of the people who I owe a "thank you" and a good deal more, I will try to highlight some of the amazing people who have helped me personally in the last few years.

The LOOC UP project would have never gotten started without the enthusiasm and imagination of Szabi Marka and Erik Katsavounidis. I am thankful for their faith, both in the project, and in me. The project also benefited from courage and support from key members in the collaboration, including Dave Reitze and Stan Whitcomb - I am grateful for their warm words of encouragement.

Most of what I know about gravitational wave detection comes from conversations with members of the Burst working group. I have worked closely with a great many Burst group scientists, and have had only good experiences. I would

like to especially thank Patrick Sutton, Brennan Hughey, Shourov K. Chatterji, Jamie Rollins, Amber Stuver, and Antony Searle. In addition, this project posed some unique challenges for overworked review committees: Ben Owen, Stas Babak, and Nicolas Leroy spent a great deal of time ensuring that the low-latency software behaved as intended. I have also been very fortunate to work closely with a few members of the CBC group, especially Larry Price and Patrick Brady, who made the low-latency analysis a priority, and gave proper attention to coordination among the many moving parts.

I would also like to thank a number of people who have been instrumental in my personal growth, in some cases well before I ever wanted to attend graduate school. My wife, Erin, is an endless source of warmth, support, and wisdom. Our son, Jacob, has provided four months of much needed encouragement. My amazing parents have supported me in every intellectual pursuit I can remember, and continue to provide guidance. I am also happy to have the warm support of my wonderful in-laws, Chuck and Mary, as well as an incredible cast of brothers and, more recently, sisters. I suspect that, without the outstanding staff and faculty at University of Maryland, I would have abandoned physics long ago. The scientific community has Chris Lobb, Steve Anlage, Tom Cohen, and Tom Gleason to blame for my continued presence. I have also benefited from the wisdom of a great many other faculty members, and I would like to especially note those who served on my thesis committee: Cole Miller, Ted Jacobson, Chris Reynolds, and Jordan Goodman. I would also like to thank

some of my close friends in the Maryland physics department: Chris Bertrand, Jordan Horowitz, Ellie Twedt, and Jeff Temple.

Finally, this entire process would have been impossible without the sharp intellect, gentle manner, warm nature, broad knowledge, and endless, endless patience of my advisor, Peter Shawhan. I am at a loss to imagine any quality a student could wish for in an advisor that Peter lacks. The university is lucky to have him, and I am lucky to have had the chance to work with him so closely.

Table of Contents

List of Figures

List of Abbreviations

α, β	Indices indicating a gravitational wave detector
$\vec{r_\alpha}$	The location of an interferometer
$\vec{r_0}$	Location of the center of the earth
$\vec{\Theta}$	The direction of the source of a gravitational wave
θ, ϕ	The polar coordinates of the direction of propagation of a signal
$\tau_{\alpha\beta}$	The travel time between two locations
$h_{+,\times}$	The gravitational wave strain
ρ	Detection statistic for a generic ETG
η, Ω, ρ_c	The detection statistics of cWB, Omega Pipeline, and MBTA
s	The interval, or proper time, between two events in space-time
$\hat{x}, \hat{y}, \hat{z}$	Cartesian unit vectors
μ, ν	Indices for space-time coordinates, running from 0 to 3
$g_{\mu\nu}, \eta_{munu}, h_{munu}$	Space-time metrics
$I_{\mu\nu}$	Quadrupole mass moment
c	Speed of light
G	Gravitational constant
r_0	Separation distance
D	Distance between a source and an observer
$F_+, F_\times$	The antenna pattern response of an interferometer
t	A time coordinate
h_{rss}	A measure of the amplitude of a gravitational wave signal
h_{det}	The gravitational wave strain measured in a detector
σ	The uncertainty in a random variable
L	A likelihood based on GW data
P	A likelihood based on GW data and the galaxy catalog
M_{gal}	The blue light luminosity of a galaxy
FAR	False Alarm Rate
cWB	Coherent Wave Burst
MBTA	Multi-Band Template Analysis
IFO	Interferometer
SNR	Signal to noise ratio
ETG	Event trigger generator, software used to identify and reconstruct GW sources
GW	Gravitational wave
EM	Electromagnetic
NS	Neutron Star
FOV	Field of view

Chapter 1

Introduction

In recent years, the operation of large scale interferometers such as LIGO and Virgo has increased excitement surrounding the field of gravitational wave science. Perhaps most notably, installation of the much anticipated Advanced LIGO is now in progress, with Advanced Virgo to follow close behind. Both second-generation gravitational wave observatories are expected to reach sensitivities that will allow detection of inspiraling binary neutron stars, and possibly other gravitational wave sources as well, and have operations planned as early as 2014.

With the construction of this network, gravitational wave science is anticipated to transition from an experimental field, seeking a first discovery, to an observational field, providing detailed information on compact object source populations, galaxy formation, the dynamics of high energy astrophysical processes, cosmology, and even fundamental physics.

With the field of gravitational wave astronomy imminent, questions concerning the types of information obtainable with a network of interferometers are becoming significant. One theme is that, whatever information may be extracted from a gravitational wave signal, a great deal more information would be available if that signal is combined with a counterpart observation in the electromagnetic spectrum.

If catastrophic astrophysical events, such as mergers of compact objects, can be observed in both the gravitational and electromagnetic spectrums, then a wealth of insights into astronomy, cosmology, and fundamental physics may become available.

With this in mind, we have established the first campaign to actively seek optical counterparts to gravitational wave signals. Though gravitational and electromagnetic data have been combined in the past, this has always been done by comparing pre-recorded data. In our search, for the first time, astronomical instruments responded to triggers from gravitational wave observatories, collecting data from sky positions that may have hosted the gravitational wave sources. In order to discover optical counterparts that may be quickly fading, the gravitational wave data analysis was completed with low latency, providing astronomers and robotic telescopes with estimated source coordinates about thirty minutes after the arrival of the gravitational wave signal. The three-site low latency analysis, communication with robotic telescopes, and the position reconstruction algorithms were all novel developments for the search, which is described in this work. The low-latency, joint EM/GW search is also described in [1], which shares some content with this dissertation. The dissertation is organized as follows:

Chapter 1 presents an introduction to gravitational waves. Some fundamentals of gravitational wave physics are discussed, along with some astrophysical evidence for the existence of gravitational waves.

Chapter 2 provides an introduction to interferometer gravitational wave detec-

tors, with an emphasis on LIGO and Virgo, and provides an overview of the current global network of sensitive interferometers.

Chapter 3 presents an overview of potential sources of gravitational wave signals, followed by a discussion of various electromagnetic counterparts that may accompany a transient gravitational wave signal. The chapter includes some examples of the types of investigations that may be performed with joint GW/EM observations.

Chapter 4 provides a brief overview of the principles that guided our construction of the low-latency search for optical counterparts to gravitational wave triggers.

Chapter 5 presents estimates of the scale of angular errors that may be expected when using the LIGO/Virgo network to localize short duration signals. This is followed by an overview of the three algorithms that were used by the low-latency search to identify event candidates and estimate the source positions from GW data.

Chapter 6 describes some of the important considerations when choosing telescopes to search for electromagnetic counterparts to gravitational wave signals. The network of telescopes that searched for counterparts is briefly described.

Chapter 7 describes an algorithm implemented within the LUMIN software package to evaluate the uncertainty region for a GW trigger and select fields for telescopes to image.

Chapter 8 describes a Monte Carlo simulation which was used to test the

event reconstruction and LUMIN field selection. The results from the simulation demonstrate the ability of the pipeline to recover GW source positions.

Chapter 9 describes some of the details of the low-latency search, including the process for selecting triggers, and procedures for an on-call team that manually validated each trigger.

Chapter 10 presents the collection of gravitational wave event candidates that received optical follow-up observations, along with a preliminary analysis. An overview of the procedures that were used to analyze the optical data are presented, along with a discussion of issues related to identifying transients in optical data.

Finally, a summary of the work, with implications for advanced detectors, is presented in Chapter 11.

1.1 Gravitational Waves

General Relativity is the theoretical framework, first described by Albert Einstein, that is currently used to describe gravitational interactions. The centerpiece of the theory is the relationship between matter and space-time geometry contained in the Einstein field equations. Gravitational waves (GWs) are a consequence of this theory. Using the linearized Einstein equations, a small perturbation to space-time curvature can be shown to propagate through vacuum at the velocity of light. In this section, some of the basic features of gravitational waves are presented. For much of the discussion, we will follow the presentation in [82].

The study of relativity begins with the understanding that the only observer-independent measurement of the separation between events is a combination of space and time known as the interval, ds. In a region of space-time that is locally free of gravitational curvature (Minkowski space-time), the interval may be expressed in Cartesian coordinates as

$$ds^2 = -c^2dt^2 + dx^2 + dy^2 + dz^2 \tag{1.1}$$

where c represents the speed of light.

More generally, the interval may be calculated as

$$ds^2 = \sum_{\mu\nu} g_{\mu\nu} dx^\mu dx^\nu \tag{1.2}$$

Here, the indices μ and ν run from zero through three, representing the four space-time coordinates. For example, if Cartesian coordinates are used, then dx^μ represents each of $c\,dt$, dx, dy, and dz. The matrix $g_{\mu\nu}$ is the space-time metric, which contains all the information that relates our coordinates to the local measure of ds, and so defines the local space-time. Gravitational waves arise from small variations of the metric.

Comparing equations 1.1 and 1.2, one can see that the metric for flat space-

time, $\eta_{\mu\nu}$, may be expressed for Cartesian coordinates as the following matrix:

$$\eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{1.3}$$

To pursue our interest in gravitational waves, we will consider a metric that is *nearly* flat:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{1.4}$$

where $|h_{\mu\nu}| \ll 1$.

General Relativity allows a great deal of freedom in choosing coordinate definitions. A particular set of definitions, known as the "transverse traceless gauge", ties coordinates to the motion of imagined freely falling objects. If we choose this gauge, and take the weak-field limit of Einstein's field equations in vacuum, we find the remarkable result:

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0. \tag{1.5}$$

The above is the familiar wave equation. This is one of nature's most persistent forms; replacing $h_{\mu\nu}$ with an electric or magnetic field component would describe a radio signal, and with other adjustments the equation could be made to describe sound, seismic waves from an earthquake, or ripples in a shallow puddle. The wave equation allows solutions of the form $f(\vec{x} - \vec{v}t)$, where $\vec{v}$ is the phase velocity, and $|\vec{v}|$ is the speed of the wave, which is c for solutions to Equation 1.5. The fact that

this relationship may be derived from Einstein's field equations implies that the functional form that allows sound to reach our ears and light to reach our eyes also permits another communication across great distance, but this one a ripple of space itself.

For a gravitational plane wave traveling in the $\hat{z}$ direction, the choice of a transverse traceless gauge means that the perturbation must take the form

$$h_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{1.6}$$

where h_+ and $h_\times$ may be an arbitrary function of $(z - ct)$. h_+ and $h_\times$ define the two polarizations of a gravitational wave. A schematic of the effect of a gravitational wave on a ring of test particles may be seen in Figure 1.1. The h_+ polarization is seen to have opposite effects on the metric in the $\hat{x}$ and $\hat{y}$ directions. The h_+ polarization stretches spatial distances between freely falling objects in the x-direction, simultaneously shrinking them in the y-direction. The $h_\times$ polarization has a similar effect, but with a 45° rotation. A spherical wave front, if measured over a distance much smaller than the distance to the source, appears very similar to a plane wave, and so may be well described by Equation 1.6 as well. The measureable effects of h_+ and $h_\times$ will be discussed in Section 2.

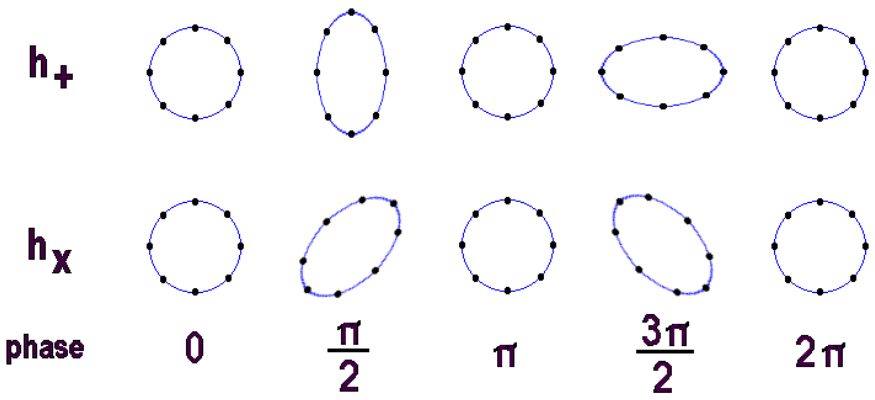

Figure 1.1: A representation of the effect of a passing gravitational wave on a ring of test particles. Each test particle is assumed to move freely in the plane of the page, with the wave incident from a direction normal to the page. The effect of the GW is greatly exaggerated so as to be visible. Image courtesy Wm. Robert Johnston.

1.2 Astrophysical Gravitational Waves

So far, our consideration of gravitational waves has been rather abstract - we have made no mention of potential sources, or even of the scale of perturbations we might expect.

In analogy with electromagnetic (EM) radiation, we can understand the expected gravitational radiation based on a multipole expansion of the source distribution. For EM waves, the sources are charges. Conservation of charge makes it impossible for any charge distribution to have a changing monopole moment. Then, for EM radiation, we look to the next highest moment, the dipole moment, as the dominant term when calculating EM radiation.

In the gravitational case, it will be the mass and energy distribution which will act as sources. As in the EM case, the monopole term cannot change, because energy is a conserved quantity. However, for a mass source distribution, unlike a charge distribution, conservation of momentum and angular momentum prevent radiation from dipole terms in an isolated system. For gravitational radiation, then, it is the next term, the quadrupolar moment $I_{\mu\nu}$, which is expected to give the dominant contribution:

$$I_{\mu\nu} \equiv \int dV(x_\mu x_\nu - \frac{1}{3}\delta_{\mu\nu}r^2)\rho(\mathbf{r}) \tag{1.7}$$

where $\mathbf{r}$ is the source position relative to some reference point, ρ is the mass density, and $\delta_{\mu\nu} = 1$ for $\mu = \nu$, but is zero otherwise. It is second derivatives of this moment

which ultimately lead to gravitational waves, following the relationship

$$h_{\mu\nu} = \frac{2G}{Dc^4}\ddot{I}_{\mu\nu} \tag{1.8}$$

where D is the distance from the source to the observer.

To set the scale for gravitational radiation from astrophysical sources, we imagine a simple model of a binary star system. Two spherical masses are in orbit about each other, under only the influence of their own gravity. The masses each have mass M, and are separated by a radius r_0. As an order-of-magnitude estimate, we can inspect equation 1.7 and see:

$$I_{\mu\nu} \sim Mr_0^2. \tag{1.9}$$

Looking at equation 1.8, we see that estimating the two time derivatives will require a time scale. The right time-scale in this system is the orbital period. Using Newtonian mechanics, and neglecting dimensionless coefficients, leads to $T^2 \sim \frac{r_0^3}{GM}$. So, the scale of the associated gravitational radiation is

$$h_{\mu\nu} \sim \frac{(GM)^2}{c^4 D r_0}. \tag{1.10}$$

It is possible to simplify this expression by noticing that the Schwarzschild radius is $2GM/c^2$. We denote this as r_S, and substitute this into equation 1.10 with continued disregard for dimensionless coefficients to find

$$h_{\mu\nu} \sim \frac{r_S}{D}\frac{r_S}{r_0}. \tag{1.11}$$

As an example of the sort of source LIGO hopes to detect, imagine a binary neutron star system in a nearby galaxy, say, 10 Mpc away (1 Mpc $= 3.1 \times 10^{19}$ km).

Neutron stars are the remnants of some classes of stars that have exhausted their nuclear fuel, and may be the densest observable objects in the universe. Neutron stars often have a mass of around 1.4 solar masses, which leads to a Schwarzschild radius of 4 km. If they are in a very tight orbit, their separation distance may be as small as 20 km. This leads to an estimate of the space-time perturbation reaching earth of

$$h_{\mu\nu} \sim 10^{-21} \tag{1.12}$$

Of course, to test for the existence of gravitational waves, and to calibrate gravitational wave detectors, a source of gravitational waves that could be controlled in the laboratory would be much more convenient than a binary neutron star system. The estimate used above may be applied to the sort of apparatus that one might hope to build to generate gravitational waves.

As a laboratory model of a neutron star binary system, imagine constructing something like the bar used for the bench press exercise at a gym. Such a bar could be about 1 m long, and include weights on each end of, say, 100 kg. The dumbbell could be mounted onto a motorized, rotating platform, and perhaps be made to spin as rapidly as 1000 Hz. Notice that this is rather extreme; the masses are moving with a velocity around Mach 9. To measure the effect as gravitational radiation, rather than just the local gravitational effects of moving the masses around, the detector should be placed at least one wavelength away, or about 300 km. This

leads to an estimate for the laboratory created gravitational wave amplitude of

$$h_{\mu\nu} \sim \frac{G}{c^4}\frac{Mr_0^2}{DT^2} \sim 10^{-41}. \tag{1.13}$$

The spinning dumbbell has created gravitational radiation which is weaker than the emission reaching earth from the binary neutron star system *by 20 orders of magnitude*. Other schemes have been imagined for manufacturing gravitational waves, but none have been proposed which could create a signal as strong as what is expected to arrive from binary neutron star mergers. Despite the inconvenience of depending on sources that cannot be controlled, sources of observable gravitational waves will have to be astrophysical, not terrestrial.

1.3 Evidence for Gravitational Waves

A handful of observed systems provide indirect evidence for the existence of gravitational waves [58]. The precision of some of these measurements is so unusual, and the data is so compelling, that few people doubt the signature of gravitational waves.

The best evidence for gravitational waves comes from observations of astrophysical systems known as binary pulsars. A pulsar is a neutron star which emits radio waves in a collimated beam. The beam sweeps through space as the neutron star rotates, in a fashion similar to the beam from a lighthouse. Pulsars are observable if the beam happens to pass across the earth. A binary pulsar system is simply a pulsar in orbit about another star. While a collection of binary pulsars has been

observed, this discussion will focus on the most dramatic evidence for gravitational waves, from binary pulsar B1913+16, known after its discoverers as the Hulse-Taylor pulsar.

The Hulse-Taylor pulsar rotates with remarkable speed: it delivers radio pulses to the earth at a rate of 17 times per second. Because the details of the pulses are tied to the rotational and orbital motion of the neutron star, observations of the pulsar allow measurements of the mass of both stars in the binary system, as well as a full description of their orbits. The two neutron stars have been found to orbit each other with a closest approach of around 750,000 km, and a period of 7.8 hours.

With all of these orbital parameters in hand, it becomes possible to predict the amount of orbital energy that exits the system through radiation in gravitational waves. For a Newtonian orbit, the total energy is $-GMm/2r$, and so as an orbit loses energy, the separation distance decreases, leading to a decrease in the orbital period. The period of the binary pulsar is observable in the radio emission, and so observing a change in the period over time could confirm or refute the prediction of general relativity.

The Hulse-Taylor pulsar was recognized as an important system when it was discovered in 1974, and so it has been carefully observed for over 30 years years [96, 95]. The stunning fact is that the observed shift in period matches the prediction based on energy loss due to gravitational radiation *to a precision of 0.2%*. A graphical representation of the orbital decay is shown in Figure 1.2, which is repro-

duced from [58]. It is difficult to imagine that this precise agreement between theory and observation is a coincidence. Moreover, a few other binary systems containing pulsars have been observed, and the observed orbital decay agrees with the prediction of general relativity in all cases [38, 53]. These systems argue very strongly for the existence of gravitational waves, their ability to carry energy, and for the existence of astrophysical systems which emit non-negligible amounts of gravitational radiation.

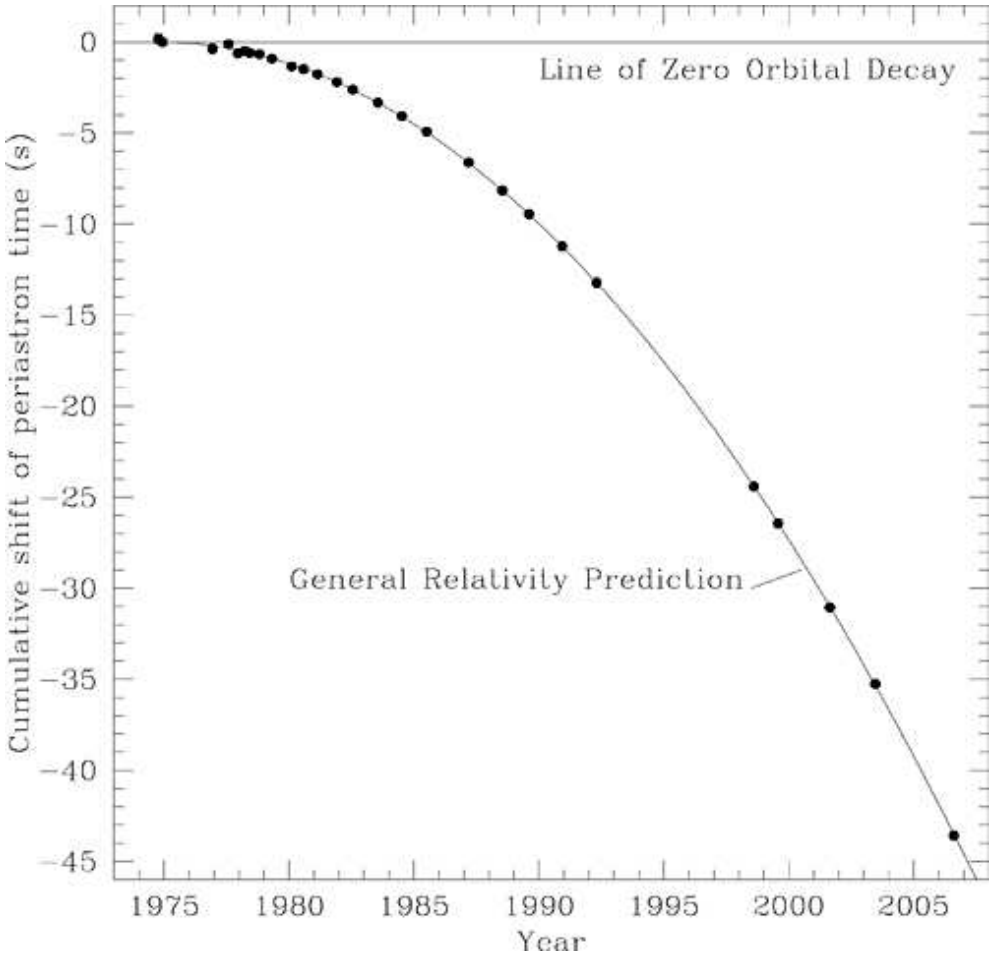

Figure 1.2: The orbital decay of the Hulse-Taylor pulsar. The points represent observations, and the solid curve is the prediction from general relativity, assuming that the only loss of energy is through gravitational waves. The prediction and thirty years of observation agree at the level of experimental accuracy, which is roughly 0.2%. The system presents strong evidence for the ability of gravitational waves to carry energy. Figure reproduced from [58].

Chapter 2

Interferometric Detectors

The search for gravitational waves spans six decades, and yet so far, no confirmed detections have been reported. The difficulty lies in the incredible subtlety of the sought signal, which leaves only the faintest impression on any kind of matter. However, the persistence of gravitational wave scientists has yielded continued progress, and detectors have become more and more sensitive with time. The current generation of detectors, led by the interferometers LIGO and Virgo, have demonstrated a level of sensitivity that implies a reach to merging neutron stars well beyond our own galaxy. The next generation of detectors is expected to observe gravitational waves from long-predicted sources. This chapter describes the mechanism that allows interferometers to respond to gravitational wave signals, and gives an overview of the current global interferometer network.

2.1 Response to Gravitational Wave

The first gravitational wave detector was designed and built by Joseph Weber, at the University of Maryland in College Park [94]. His detector was a solid aluminum cylinder, with piezo-electric crystals attached to detect energy deposited in the cylinder by passing gravitational waves. His work began a new experimental

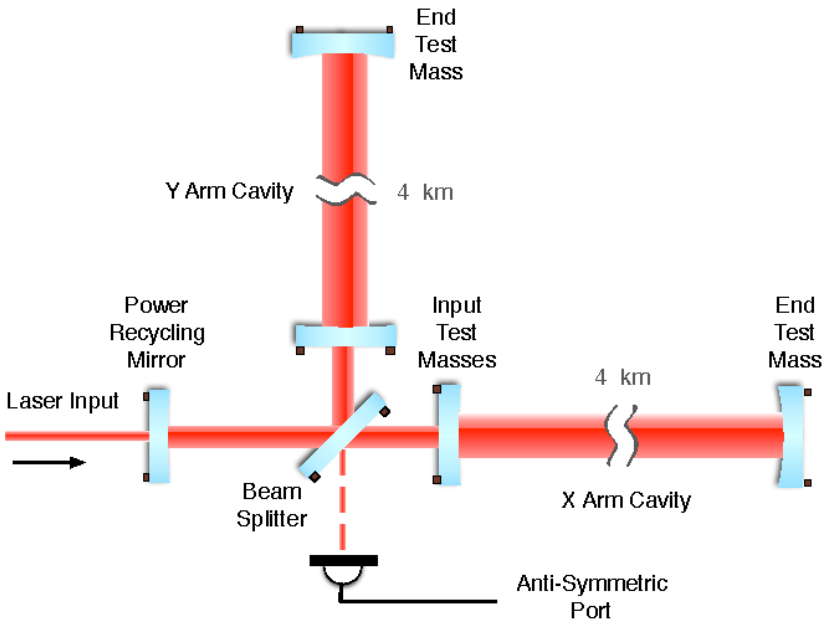

Figure 2.1: A simplified schematic of a LIGO interferometer. Laser light enters from the left, and is divided by the beam splitter into two arms, each containing a Fabry-Perot cavity. The light eventually recombines at the beam splitter. Differences in the path length along the x and y arms impact the phase of the recombining light, and so determine the output light reaching the photodiodes at the anti-symmetric port. The figure is reproduced from [2].

field, and a collection of similarly designed "bar detectors" were constructed at sites across the United States, Europe, and Australia.

Most of the recent effort to detect gravitational waves has been related to large scale interferometer detectors. A simplified schematic of one such observatory, the Laser Interferometer Gravitational Wave Observatory (LIGO), is shown in Figure 2.1 [2].

The schematic, along with some of the concepts developed in Section 1.1, allow an understanding of the basic principle that allows interferometers to respond to gravitational radiation.

At its most basic level, the LIGO detector is seen to be a Michelson interferometer. Infrared light, with a wavelength of 1064 nm, enters the detector, and is divided into two paths by a beamsplitter. The light then travels down two lengths of 4 km each, labeled the x and y "arms". At the end of each arm, the light encounters a highly reflective surface, which is a coating on the end test mass. The light then returns to the beamsplitter, where it recombines. Depending on the relative phase of the two beams, some of the recombined light exits the beam splitter in the direction of the anti-symmetric port, where photo-detectors record any incident power. Also shown in the diagram is the power recycling mirror which, as its name suggests, increases the laser power circulating in the observatory. The input test masses, along with the end test masses, form Fabry-Perot cavities in each arm which are held at resonance, further increasing the laser power circulating in the arms.

All four of the test masses are suspended by wires so that they may swing in the direction parallel to the beam path as pendulums. In this sense, they may be viewed as "freely falling" bodies, at least in one dimension. In the following discussion, we will imagine each test mass is free to move in the direction parallel to its arm length.

As with a simple Michelson interferometer, the flux of light exiting the anti-symmetric port depends on the difference of the phase of light striking the beam splitter from the y and x directions. The principle that makes an interferometer useful as an observatory is the fact that a passing GW will alter this phase difference.

To see how this works, we will imagine a further simplified scenario. Imagine removing the input test masses from the schematic, so what remains is a simple Michelson interferometer. Choose a coordinate system with the x-axis along one of the arms, and imagine a gravitational wave approaching from zenith, described by $h_+(t)$ and $h_\times(t)$, as in Equation 1.6. However, we will consider only a span of time much shorter than one period of the gravitational wave, so that both h_+ and $h_\times$ will be approximately constant. Since the mirrors are free to move parallel to the arm lengths, the use of a transverse, traceless gauge (as demanded by Equation 1.6) implies that the mirrors will remain at fixed coordinate locations.

With all of these simplifying assumptions in hand, note that the phase of light arriving at the beam splitter from the end test masses will be closely related to the travel time of a phase front down the arms and back. It is this travel time which we

wish to calculate. The interval between points along the world line of light traveling in vacuum must always be zero; this is a cornerstone of relativistic physics. So, for a phase front traveling the length of one arm, Equation 1.4 may be used to write

$$ds^2 = 0 = -c^2dt^2 + (1 + h_+)dx^2 + 2h_\times dxdy + (1 - h_+)dy^2 + dz^2 \tag{2.1}$$

For light traveling along the x-arm, $dy = dz = 0$, and so

$$c^2dt^2 = (1 + h_+)dx^2 \tag{2.2}$$

$$\frac{dt}{dx} = \pm\frac{1}{c}\sqrt{1 + h_+} = \pm\frac{1}{c}(1 + \frac{1}{2}h_+) \tag{2.3}$$

where we have approximated the square root, taking advantage of the fact that $h_+ \sim 10^{-21} \ll 1$. The sign on dt/dx should be chosen based on the direction of travel.

Imagine a phase front reaching the beam splitter from the left, and simultaneously embarking on journeys down the x and y arms. Let the arm length, as measured with no gravitational wave present, be written L. The time to reach the x-arm end test mass may be calculated as:

$$\tau_{x,out} = \int_0^{\tau_{x,out}} dt = \int_0^L \left(\frac{dt}{dx}\right) dx \tag{2.4}$$

$$\tau_{x,out} = \frac{L}{c}(1 + \frac{1}{2}h_+) \tag{2.5}$$

For the return trip, a similar calculation shows

$$\tau_{x,return} = \int_L^0 -\frac{1}{c}(1 + \frac{1}{2}h_+) = \frac{L}{c}(1 + \frac{1}{2}h_+) \tag{2.6}$$

Putting this together, the round trip time in the x-arm is:

$$\tau_{x,roundtrip} = \frac{L}{c}(2 + h_+). \tag{2.7}$$

Compared with the time expected in the absence of the gravitational wave, this shows an increase in round trip time of h_+L/c. Since the speed of light is independent of the metric, this difference in round trip time may be understood as a change in the length of the arm of the interferometer

$$\Delta L = \frac{1}{2}\left(\frac{h_+L}{c}\right)c = \frac{h_+L}{2} \tag{2.8}$$

$$\frac{\Delta L}{L} = \frac{h_+}{2} \tag{2.9}$$

The effect of the gravitational wave is seen to increase the distance between two freely falling objects, by an amount proportional to the initial distance. For this reason, the metric perturbations h_+ and $h_\times$ are referred to as strain.

Looking now at the y-arm, we find

$$\frac{dt}{dy} = \pm\frac{1}{c}(1 - \frac{1}{2}h_+) \tag{2.10}$$

which, after repeating the calculation used for the x-arm, leads to

$$\tau_{y,roundtrip} = \frac{L}{c}(2 - h_+) \tag{2.11}$$

To summarize the steps leading to Equations 2.7 and 2.11, the interferometer was initially placed in a configuration where the x and y arms were the same length (L), as measured with a flat metric. Under these conditions, the two equal phase

fronts that leave the beamsplitter at some moment should reunite upon their return to the beamsplitter.

However, the effect of the metric perturbation causes the equal phase fronts to arrive back at the beam splitter at different times, by an amount

$$\Delta\tau_{roundtrip} = \tau_{x,roundtrip} - \tau_{y,roundtrip} = \frac{2Lh}{c} \tag{2.12}$$

This time difference may be expressed as a difference of phase, $\Delta\gamma$, of the fronts that do arrive simultaneously:

$$\Delta\gamma = 2\pi\frac{\Delta\tau_{roundtrip}c}{\lambda} \tag{2.13}$$

The phase shift $\Delta\gamma$ imposed by the gravitational wave alters the so-called "fringe pattern" which is read out by the photodiode at the anti-symmetric port of the interferometer.

While this demonstrates the principle that allows an interferometer to respond to a gravitational wave, the actual response of the interferometer is dependent on the frequency of the gravitational signal. A number of features play into this frequency dependent response, and these are described in [2]. Moreover, in a real interferometer, each test mass's position is maintained through a set of magnets glued to the test masses, and actuated by currents in wire coils. These actuators, as well as the photodiodes at the anti-symmetric port, are part of a feedback loop that keeps the optical cavities on resonance. Above we derived a shift in phase assuming the test masses are free to swing, but in practice, the feedback loop keeps the optical path

length the same, and measuring the force applied by the coil currents allows the gravitational wave strain to be inferred. The modeled response of the interferometer is used to calibrate the output data to reflect the equivalent gravitational wave strain along the arms.

In addition to the frequency dependence, the response of the interferometer also depends on the direction that the gravitational wave is traveling. So far, we've only treated the case of a single GW polarization traveling in the $\hat{z}$ direction. Generally, coordinate rotations must be used to find the amount of strain along the interferometer arms, $h(t)$. For a GW traveling in direction θ, ϕ, the result may be expressed as

$$h(t) = F^{+}h_{+}(t) + F^{\times}h_{\times}(t) \tag{2.14}$$

$$F^{+} = \frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \cos 2\psi + \cos\theta \sin 2\phi \sin 2\psi \tag{2.15}$$

$$F^{\times} = -\frac{1}{2}(1 + \cos^2\theta)\cos 2\phi \sin 2\psi + \cos\theta \sin 2\phi \cos 2\psi. \tag{2.16}$$

$$\tag{2.17}$$

The coordinates θ, ϕ are the usual polar coordinates, with the positive $\hat{x}$ and $\hat{y}$ axes aligned along the arms of the interferometer. The polarization angle ψ ranges from 0 to π, with 0 indicating the + polarization is aligned with lines of constant θ and ϕ.

The functions F^{+} and $F^{\times}$ define the directional sensitivity of an interferometer, and are known as the antenna patterns. A single antenna pattern is a fairly broad pattern over the sky; as a result, interferometers are sometimes described as "all-

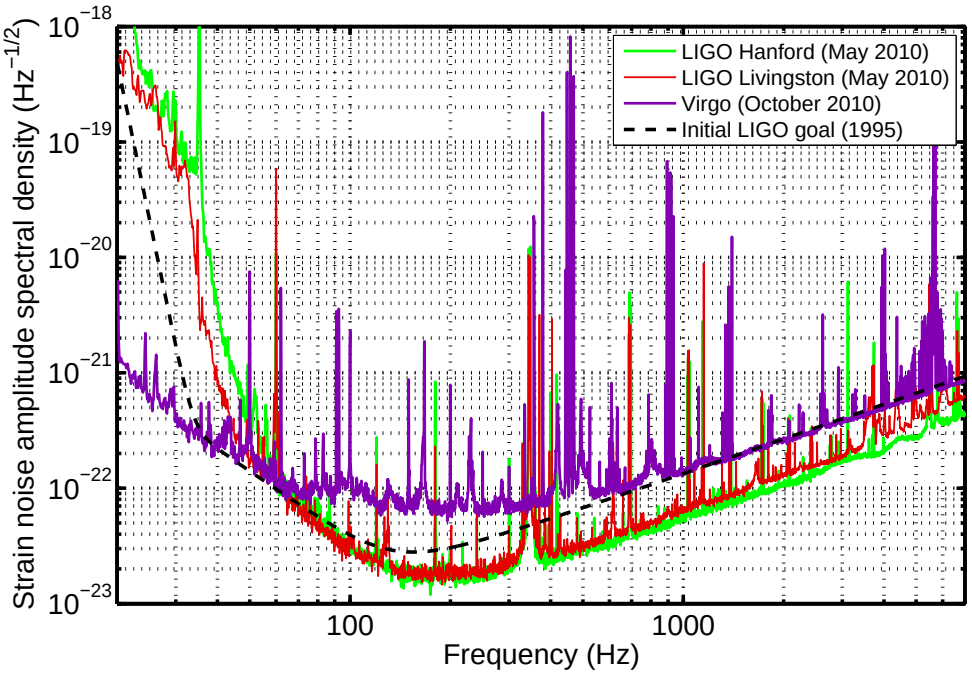

Figure 2.2: Sample noise curves of the two LIGO detectors and Virgo during the S6 and VSR3 runs. By the end of S6, LIGO had exceeded its initial sensitivity goal for most frequencies above 100 Hz.

sky" instruments. This broad antenna pattern has the advantage of allowing the observatories significant sensitivity to signals from most directions. On the other hand, for our present enterprise of seeking optical counterparts, the all-sky nature of interferometers will present a great challenge: a single interferometer is incapable of localizing a signal. For detectors similar to LIGO, localization will always depend on collecting data from a network of detectors.

2.2 Noise Sources

The interpretation that gravitational waves induce a differential strain on the interferometer arm lengths offers insight into the precision demanded of the experiment. In Chapter 1, the value of $h_{\mu\nu} \sim 10^{-21}$ was presented as the scale of gravitational waves we might hope to observe. Putting this scale into Equation 2.9 shows that the length change to be observed over the 4 km arm length is of order 10^{-18} m. This number is a bit startling at first glance – it shows that LIGO is expected to measure length changes smaller than 1% the diameter of an atomic nucleus. In fact, within the sensitive frequency band of LIGO and Virgo, this is exactly what happens.

Sample noise curves from 2010 are shown in Figure 2.2. The y-axis shows the square root of the power spectral density of the noise. A variety of noise sources must be suppressed to reach this level of sensitivity. The strain interpretation shows that motion of the test masses must be suppressed. A few mechanisms are used to isolate the test masses from seismic motion in the observing band. The test masses are suspended on wires, giving them a pendulum response that suppresses high frequencies. In addition, the wire suspension points are isolated from the ground using so-called seismic isolation stacks, which are systems of springs and masses, also designed to suppress high frequency motion. Finally, the 2009/2010 run included the use of a hydraulic system that uses a feedback loop to actively suppress seismic motion. In Figure 2.2, the seismic noise dominates at low frequencies, and is

primarily responsible for the noise below $\sim$ 60 Hz. At higher frequencies, above $\sim$ 400 Hz, the dominant noise source is "shot noise", which originates in the stochastic, quantum nature of photons. Excitations of vibration modes due to thermal noise are significant between these extremes. For example, the large peaks between 300 and 400 Hz in the LIGO noise curves are due to resonances in the wires that suspend the test masses.

2.3 Current and Future Detectors

LIGO, briefly described above, operates two interferometers with nearly identical designs at sites near Hanford, Washington and Livingston, LA, each with 4 km long arms. LIGO completed its fifth science run, S5, in October 2007, at design sensitivity, and ran its sixth science run in 2009 and 2010. During S5, the Hanford site included a second interferometer co-located with the first, though only half the length.

The 3 km Virgo observatory is located near Cascina, Italy. Virgo's first science run (VSR1) overlapped with LIGO's S5, allowing network analysis of the three site data (See, for example, [5]). Close collaboration between LIGO and Virgo scientists was arranged to jointly analyze data; a variety of common conventions, including sampling rates, data formats, and some data quality labeling conventions, were created and implemented. Over the course of S6, Virgo made two science runs, VSR2 and VSR3, with a commissioning break in between.

This work describes a low-latency analysis that ran during a subset of the time that LIGO and Virgo took coincident data during S6/VSR2/VSR3. The low-latency analysis ran during VSR2 between December 17, 2009 and January 8, 2010 (winter run) and during VSR3 between September 2 and October 20, 2010 (autumn run). At the time, these were the most sensitive broad-band GW detectors in the world.

In November of 2010, the LIGO interferometers were shut down for commissioning of the next generation interferometers, known as Advanced LIGO. Advanced LIGO is designed to detect signals with roughly 1/10 the amplitude of initial LIGO, giving it an order of magnitude increase in astrophysical reach. In addition, Advanced LIGO may include a return of the second interferometer at Hanford, though this time, extended to match the 4 km arms of the other detectors[1]. In autumn of 2011, Virgo will embark on installation of second-generation hardware, known as Advanced Virgo, designed to achieve sensitivity similar to Advanced LIGO.

In addition to LIGO and Virgo, the development of other planned and existing large interferometers is making important contributions to gravitational wave science. The GEO 600 detector is a 600 m detector in Germany. GEO is a full partner within the LIGO Scientific collaboration, and exchanges data openly with LIGO and Virgo. In addition, many of the technologies that are used in the LIGO observatories are developed and tested at GEO. An upgrade program is underway for the GEO 600 detector, known as GEO HF, which will utilize squeezed light to

[1]The 2nd Hanford detector may instead be installed at another location, as described later in this section

give GEO a sensitivity comparable to LIGO in the band between one and a few kHz.

The CLIO detector in Japan is the world's leading cryogenic gravitational wave interferometer [100]. Its test masses are separated by 100 m, and are cooled to 20 K during operation. CLIO is a prototype for the planned Large-scale Cryogenic Gravitational wave Telescope (LCGT)[2], which has a design sensitivity similar to Advanced LIGO, but with mirrors cooled to 20 K to reduce thermal noise. The LCGT is planned to be operational around 2018.

Other possible future projects include the Indian Initiative in Gravitational-wave Observations (IndIGO) [3], which seeks to build a large scale detector in the Asia-Pacific region. Discussions are in progress between LIGO, the National Science Foundation, and the government of Australia over the possibility of housing the optical equipment designed for one of the Hanford interferometers in an interferometer in western Australia. This proposal is known as LIGO Australia, and a decision regarding whether the hardware will be installed at Hanford, or be stored for use when an Australian site is built, will likely be reached in late 2011.

From the perspective of finding optical counterparts, having a variety of sites is extremely important. The localization precision improves with each site added to the network, and the difference between a three site and four site network can be dramatic [48, 84].

[2]http://gw.icrr.u-tokyo.ac.jp/lcgt/

[3]http://www.gw-indigo.org

Chapter 3

Astrophysical Sources

A variety of EM emission mechanisms, both observed and theoretical, may occur in association with observable GW sources. Characteristics of a few scenarios helped inform the design and execution of this search. This chapter presents the types of GW sources that are typically sought, as well as some likely models for characteristics of associated EM emission. The discussion also describes some of the scientific investigations that are possible with joint GW/EM observations.

3.1 Types of gravitational wave sources

To be effective, an analysis that searches for astrophysical signatures in the gravitational wave data must make some assumptions about the sought waveform. A given search must either look for a continuous signal, one that is persistently in the data over the length of the observation, or a transient signal which only exists for a short time. So far, most searches for transients in the LIGO data have assumed the signal is observable for a few minutes or less. Moreover, a search may attempt to use a model of a particular source to search as sensitively as possible, or may sacrifice some sensitivity to search broadly in a waveform agnostic, or unmodeled, manner. Here, we present a few examples of the types of sources that are currently

sought with the various approaches.

Continuous, unmodeled waveforms are known as "stochastic", or random, signals. These may be thought of as noise in the GW spectrum, though with astrophysical significance, analogous to the cosmic microwave background in the EM spectrum. In fact, there may also exist a cosmic background in the gravitational spectrum, with origins in the inflationary period in the very early universe; the measurement of a signal of this type has major implications for our understanding of events just after the Big Bang [61]. A less uncertain source of stochastic gravitational waves is the so-called "astrophysical background"; which is the sum of the signals of many distant supernovae, rotating neutron stars, and other gravitationally emitting sources [28, 29].

In the LIGO frequency band, the most likely sources of modeled, continuous signals are rotations of asymmetric neutron stars. In order to emit gravitationally, the NS must lack azimuthal symmetry about the axis of rotation [91]. This could be imagined as a slightly oblong NS, spinning as an egg on a countertop, or a NS with some distinct surface feature, analogous to a mountain on the earth. Past searches for these signals have placed limits on NS ellipticities of $\sim 10^{-6}$ out to distances of 1600 light-years, assuming rotation rates in the LIGO frequency range [7].

Turning now to GW transients, the unmodeled variety are known simply as "bursts." The distinguishing feature of a GW burst is an excess of signal power in a short span of time, typically defined as less than a second. Searches for bursts

attempt to be agnostic regarding waveform. Burst signals are thought to be associated with catastrophic astrophysical events, such as the core-collapse of massive stars. The last moments of a compact object merger, described below, may also be described as a burst signal. Additionally, because burst searches make few assumptions about the source mechanism, they provide an opportunity to discover unexpected sources.

The final class of signal, known as an inspiral or a compact binary coalescence (CBC), is the most widely anticipated, modeled, transient signal likely to be observed with LIGO. This source will play an important role in the development of our analysis, and so merits some description.

LIGO band inspiral signals originate with the final minutes of the deteriorating orbits of binary neutron stars, binary black holes, and binary systems comprised of a neutron star and a black hole. For example, the Hulse-Taylor pulsar discussed in section 1.3 is an example of a binary neutron star system, though currently with an orbital period far too long for observation with LIGO. As such a system evolves, however, the emitted gravitational waves carry away energy and angular momentum, so that over millions of years, the orbital radius becomes smaller, and the orbital frequency becomes faster [76]. The slowest orbital period that LIGO could observe is around 20 Hz, that is, 20 orbits per second. As orbits go, this is remarkably fast - think how strange our night sky would look if the moon made 20 complete revolutions about the earth every second! For a binary of NSs or stellar

mass black holes, such an orbit cannot be maintained for long. Gravitational waves would quickly carry away some of the remaining orbital energy, so that the orbit becomes faster and closer over a period of a few minutes or seconds. This phase of the binary's life, when the orbit is evolving under the influence of gravitational waves, is known as the inspiral. In the relativistic dynamics of this binary system, there is a minimum radius where stable orbits are allowed. Once the orbital radius reaches this minimum value, the stellar couple has reached the end of its very long engagement, and makes a final plunge towards cohabitation. This plunge is known as the merger phase, lasts well under a second for LIGO band objects, and results in the two objects coalescing into a single black hole. A final burst of gravitational wave energy is expected to exit the system during the merger.

Modeling the entire inspiral and merger process requires large-scale numeric simulations ([17] and references therein). However, for detection purposes, modeling the inspiral phase alone is often sufficient. This can be accomplished analytically using approximations of the full Einstein field equations [11]. The energy flux reaching Earth from an inspiral signal depends on the inclination angle of the binary orbit relative to the line of sight. The initial LIGO-Virgo network is sensitive to optimally oriented NS-NS mergers from as far away as 30 Mpc, and mergers between a NS and a $10\,M_{\odot}$ black hole out to 70 Mpc [4]. With advanced detectors, these range limits are expected to increase to 440 and 930 Mpc, respectively.

The low-latency analysis described in this work included searches for both

modeled, inspiral signals, and unmodeled GW bursts. Going forward, when we discuss GWs signals, we always mean *transient* GWs. This class of signal is often associated with some sort of catastrophic astrophysical event, large releases of gravitational binding energy, and, perhaps, some accompanying transient emission in the EM spectrum. It is towards the EM transients that may be associated with GW sources that we now turn our attention.

3.2 Short Gamma-ray Bursts

Gamma-ray bursts (GRBs) are short, extremely bright flashes detected a few times per month. As the name suggests, the bursts are observed initially as the very high frequency photons known as gamma-rays. Gamma-rays do not penetrate earth's atmosphere, so may only be directly detected [1] with observatories mounted on orbiting satellites, including Swift [2] and Fermi [3].

The initial, or "prompt", phase of a GRB may last only a few seconds. This prompt emission is thought to be directed in a narrow beam, so that any observed GRB must have its beaming axis directed toward earth. In addition to the prompt emission, coordinated observations of GRBs with a variety of telescopes has revealed that following the gamma-ray emission, lower energy photons may be observed for

[1] Gamma-rays are also observed indirectly through the particle showers that are created when they interact with the atmosphere by ground based observatories such as VERITAS

[2] http://heasarc.nasa.gov/docs/swift/swiftsc.html

[3] http://fermi.gsfc.nasa.gov/

a few minutes or even a few days (see, for example, [49]). This longer timescale emission is known as the afterglow, and has been observed in radio, infrared, optical, and x-ray bands, both with space and ground based telescopes [40, 33, 31].

The spectral and temporal properties of GRBs may be used to divide them into at least two classes, possibly with different source mechanisms. One feature commonly used to divide the classes is the length of time the prompt emission lasts. If more than 90% of the gamma-ray energy is contained in a burst of less than 2 seconds, it is classified as a short gamma-ray burst (SGRB), otherwise, it is called a long gamma-ray burst (LGRB).

A common, though unproven, theory is that the SGRBs are powered by compact object mergers [69, 65, 78]. In the popular model, the prompt, gamma-ray emission is nearly simultaneous with the merger event, with the beaming axis parallel to the binary's axis of rotation. If this is true, then the afterglows of SGRBs should also be observable as the optical counterparts to some inspiral GW signals.

The afterglows to SGRBs fade with a power law such that the flux is proportional to $t^{-\alpha}$, where α is between 1 and 1.5. We can use observed SGRBs with estimated source distance locations, which are typically further than the reach of initial LIGO, to estimate how bright such a source would appear if placed in the observable distance of LIGO. Taking observational data from [41], and placing a SGRB at 50 Mpc, we find the apparent optical magnitude would be between 12 and 20 for a source at 50 Mpc one day after the prompt emission.

Observing a population of GW inspiral signals linked to SGRBs would have a number of implications. The first would be to establish compact object mergers as the certain power source of SGRBs, a question that is still open. Moreover, GW observations would be able to determine the masses of the source objects, and so further distinguish between NSs and black holes as the progenitors of SGRBs. Another difficult to measure property of GRBs is the beaming angle, however, a simple comparison of the fraction of merger events with and without observed gamma-ray emission would provide a convincing estimate [51, 85].

Linking SGRBs with GWs does not necessarily require the program of optical follow-up observations described in this work. A more direct approach is to begin with GRB triggers from orbiting gamma-ray telescopes, and search in the GW data for coincident signals; in fact, such an approach has been used in the past (See, for example, [9]). However, there are some advantages that are gained by adding a program to search for afterglows based on GW data. The network of GRB satellites does not cover the whole sky at a given time, so a given GRB may be missed. Moreover, a GRB with an associated low-latency GW trigger would likely be of more interest to the astronomical community than other GRBs, and so would command more astronomical resources to observe potential afterglows.

Another interesting scenario is presented by the geometry of GRBs. While SGRBs are likely beamed, the optical afterglows are probably more isotropic, as is the GW emission. This means that a so-called "off-axis" GRB could be observable

in GW and optical, but not in gamma-rays. This presents the possibility that a GW trigger could lead to the first discovery of an orphan afterglow, that is, an afterglow without the observation of prompt emission.

3.3 Kilonovae

Even if a neutron star mergers are not the sources of SGRBs, they are expected to produce an observable optical counterpart by a completely different mechanism.

In a model first described by Li and Paczyński [57], during a NS-NS or NS-BH merger, some of the neutron star's mass is tidally ejected. As the ejected neutron-rich matter travels away from the center of the system, some of the neutrons become captured by nuclei. These neutron laden nuclei are unstable, and subsequently decay, releasing energy which heats the ejecta. The heating process brings the matter to a temperature where it emits thermally in the optical band. This mechanism, known as r-process nucleosynthesis, is known to produce heavy elements following supernovae. Because the heating mechanism is similar to that employed in a supernova, but in conditions leading to a dimmer optical transient, the proposed observable transient is known as a kilonova. Showing that this mechanism operates in nature, for example, through an observation of a kilonova in association with a GW inspiral trigger, would have major implications. The process could account for some fraction of the heavy element production in the universe.

The predicted optical emission is roughly isotropic, and so is observable re-

gardless of the orientation of the original binary system. This emission is expected to peak after about one day, around magnitude 18 for a source at 50 Mpc [67], and then fade over the course of a few days following the merger.

3.4 Pairing Inspiral Signals and EM Counterparts

Regardless if the optical counterparts to NS mergers fit the model of a kilonova, an SGRB afterglow, or some other description, there is little doubt that an observed population of compact object mergers and associated EM counterparts would be a powerful tool for studying astrophysics, cosmology, and fundamental physics.

In understanding the progenitor physics, the EM and GW signals are essentially complementary. The GW time series directly traces the bulk motion of mass in the source, whereas EM emissions arising from outflows or their interaction with the interstellar medium give only indirect information about source dynamics, requiring inference and modeling. This means that a GRB or kilonova, observed in association with an inspiral signal, would have the source mechanism immediately identified. On the other hand, observing an EM counterpart to a GW signal reduces the uncertainty in the source position from degrees to arcseconds. Knowledge of the source position is a great aid when estimating the parameters of an inspiral signal. For example, when attempting parameter estimation with a bank of templates and a single data stream, the source's distance, inclination angle, and angular position are largely degenerate. A precise source position from an EM counterpart would break

this degeneracy [22, 70]. The improved parameter estimation, in turn, leads to more precise knowledge of compact object populations; very good parameter estimation may even lead to constraints on the NS equation of state [21, 93, 30, 12].

The arcsecond localization of an inspiral source gives another advantage: it may lead to identification of the host galaxy. This would open up more studies in astrophysics, yielding insights into differences in stellar populations across galaxy type and over time. Moreover, an ensemble of inspirals with identified hosts could provide a novel measurement of cosmological parameters. Analysis of the well-modeled GW signal will provide a measurement of the luminosity distance to the source, while the redshift distance is measurable from the spectral properties of the galaxy. Taken together, they provide a direct measurement of the local Hubble constant, and one that is independent of the systematic errors that limit other methods [83, 64, 22, 70].

3.5 Stellar Core Collapse

During the bright phase of a star's life, internal pressure from the energy released in nuclear reactions maintains the star's volume against the pull of gravity. Eventually, the nuclear burning in a star's core slows, and the star quickly collapses, resulting in a large release of energy. In some cases, these stellar-core collapses are observed as supernovae. In other cases, they are thought to power the LGRBs discussed above [98, 60, 78, 99, 66]. In both scenarios, some amount of gravitational

radiation would be released, along with a very dramatic release of photons.

For core-collapse supernovae, large uncertainties still exist in the expected gravitational waveforms and energetics. Most models predict GW spectra that would be observable by initial LIGO and Virgo from distances within some fraction of the Milky Way, but not from the Mpc distances needed to observe GWs from another galaxy [73]. Neutrino detectors such as SuperKamiokande and IceCube should also detect a large number of neutrinos from a Galactic supernova [14, 36, 56]. Galactic supernovae normally would be very bright in the optical band, but could be less than obvious if obscured by dust or behind the Galactic center. Optical emission would first appear hours after the GW and neutrino signal and would peak days to weeks later, fading over the course of weeks or months.

Large uncertainties also exist for the GW emission associated with LGRBs. A large variety of possible GW emitting mechanisms within these systems have been proposed, with some models predicting GW spectra that would be observable from distances of a few Mpc with initial LIGO and Virgo [32, 50, 19, 79, 52, 46]. The afterglows of LGRBs, like the afterglows of SGRBs, typically show power law fading with $\alpha = 1 - 1.5$. However, the peak isotropic equivalent luminosity of LGRB afterglows is typically a factor of 10 brighter than SGRB afterglows [69, 40].

As with the SGRB case, an off-axis or LGRB may be observed as an orphan afterglow [35]. Another scenario allows for an observationally similar transient, known as a dirty fireball, that may be produced without accompanying gamma-ray

emission if the core collapse lacks sufficient energy [81]. These transients brighten over the course of several days or even weeks, depending on the observing band and viewing angle.

Core collapse scenarios, both those leading to supernovae and leading to LGRBs, contain complicated, dynamical physics, and have proven difficult to model. Competing models exist in the literature for many of the details. This is an area where GWs are uniquely poised to provide insight beyond what is possible with EM observations. Because the GW signal traces the mass distribution in the source, acoustic modes and other dynamical quantities could be directly probed with GW.

3.6 Cosmic String Sparks

A somewhat more exotic model of joint GW/EM emission arises from field theoretic or string theoretic considerations of the inflationary universe [86, 23, 92]. Some cosmological models predict that phase transitions in the early universe could lead to topological defects known as cosmic strings. When cosmic strings meet, there is some chance that a reconnection will occur, creating a cosmic string loop. The loops then radiate in GWs, removing energy from the string network.

The expected gravitational emission is a burst like signal, though with a distinct waveform. Some theories also expect a burst of EM emission to accompany the reconnection. In fact, there is at least one observational claim that seems to fit the proposed model. The radio band observation was a bright, 5 ms burst [59].

Dispersion in the signal suggested its origin was well outside our galaxy. If the burst originated with a cosmic string cusp, it suggests an exciting opportunity for a pairing of GW and EM information to discover a completely new phenomena, with large implications for cosmology.

3.7 Other Possible Sources

While GRB afterglows, kilonovae, and supernovae were the main models that guided our construction of the optical follow-up search, it is important to recognize that other types of optical signatures may accompany a GW transient. This suggests that a search for optical counterparts to GW transients should be somewhat broad in the characteristics of the sought optical signature.

In fact, the scaling laws derived from Equations 1.8 and 1.9 maybe be used to make a generic energy argument for the existence of optical counterparts to a source observed with LIGO.

We begin with the understanding that LIGO is sensitive to strains of roughly $\sim 10^{-21}$. Then, we see that our scaling laws suggest

$$h_{\mu\nu} \sim 10^{-21} \sim \frac{2G}{Dc^4}\frac{Mr_0^2}{T^2} \sim \frac{2G}{Dc^4}E_K \tag{3.1}$$

where E_K represents the kinetic energy of the source. Here, the units of energy are associated specifically with the *kinetic energy* of the source, because gravitational wave emission is linked with motion. More specifically, the emission of GWs depends on the time changing quadrupolar moment, so a source with a high degree of

symmetry would emit less efficiently in gravitational waves. Since we do not know the symmetry of the emitting source, we can take this as a lower limit on the kinetic energy of an observable source:

$$E_K > 10^{-21}\frac{Dc^4}{2G} = 10^{54}\text{ergs} \times \frac{D}{10\,\text{Mpc}} \tag{3.2}$$

So far, this argument is very generic - we've only assumed that the source is described by some mass, length, and time. To set the scale of an EM counterpart, we will further assume that some fraction of the kinetic energy, p, escapes as photons in the observing band, over a time T_{obs}. There is no good basis for choosing p and T_{obs}. We might argue, though, that an optimist could set p as high as a few percent, so we may be somewhat conservative in choosing $p = 10^{-6}$. The quantity T_{obs} represents the length of time we expect the optical transient to be visible. We might argue this should be around 10^5 seconds, since we hope to observe the transient within a day's time. This means that a generic optical counterpart may be described in the following way:

$$\text{Luminosity} \sim 10^{43}\frac{\text{ergs}}{\text{s}} \times \frac{10^5\text{s}}{T_{obs}}\frac{p}{10^{-6}}\frac{D}{10\,\text{Mpc}} \tag{3.3}$$

$$\text{Flux} \sim 10^{-9}\frac{\text{ergs}}{\text{s cm}^2} \times \frac{10^5\text{s}}{T_{obs}}\frac{p}{10^{-6}}\frac{10\,\text{Mpc}}{D} \tag{3.4}$$

This would appear as a transient with apparent magnitude of 12, easily observable for a modern telescope. Notice that the flux depends on one over the distance, rather than one over the distance squared, since a source that is further away must have more kinetic energy to be observable in GWs. This argument is nowhere near

a proof of the existence of EM transients accompanying GW transients. For example, the presence of a black hole in the source might imply that no photons escape, effectively setting p to zero. However, it does suggest EM counterparts are plausible even in the generic case. Since the burst searches used in low-latency search broadly for GW transients, it is reasonable to search broadly for the EM counterparts as well.

3.8 Pairing Unmodeled Sources and EM Counterparts

As in the case where the GW source is an inspiral, adding an observation of an EM counterpart to a GW burst greatly increases the astrophysical information that can be extracted.

For an unmodeled GW signal, it is not possible to extract the source distance from the GW data alone. So, in this case, the distance estimate could only be obtained if a counterpart was linked to a host galaxy. The distance estimate derived from a redshift measurement of the host galaxy could then be used to set the overall energy scale of the GW emission - one of the basic pieces of astrophysical information.

In addition, a generic GW signal and EM counterpart could allow the testing of some alternative theories of gravity - particularly those that predict gravitational waves to travel at a speed different from light [24].

Chapter 4

Search Principles and Motivation

4.1 Guiding Observations with Gravitational Waves

We have discussed in Chapter 3 some of the exciting questions that could be addressed if an astrophysical event were observed using both photons and GWs. Now we turn our attention to how joint observations may be coordinated.

In Chapter 2, we discussed how the broad antenna pattern of GW detectors makes them essentially all-sky monitors. For our search, we wished to exploit this all-sky capability to find exciting astrophysical events wherever, and whenever, they occur. Unlike optical telescopes, gravitational wave detectors observe during both night and day, are sensitive to most of the sky simultaneously, and respond to signals both from overhead and traveling "up" through the earth. Moreover, sources of observable GW transients necessarily have large amounts of kinetic energy, dense masses, and fast timescales - they are exciting and interesting events. In addition, any of the first observed GW sources will certainly be interesting, even if only because the detection of GWs will be a major discovery in itself.

Putting this all together, we conclude that GW detectors are potentially a sort of genie-in-a-bottle for the observing community: a GW network could select

only the most interesting astrophysical events in a given moment, and within the limited detection radius, find virtually *all of them!* To make exciting discoveries, the observer would need only to point his telescope in the indicated direction.

Attempting to unlock this genie for the first time, we began a project in 2007 to "Locate and Observe Optical Counterparts to Unmodeled GW Pulses" (LOOC UP) [42]. In the LOOC UP model, data from the LIGO/Virgo network is gathered and analyzed as it is collected. When triggers indicating a potential GW are found, automated programs assume the trigger represents a true GW signal, and estimate the direction of the source. This source position is then passed to astronomers, who may point their telescopes, and try to find the optical transient indicating the progenitor of the GW.

The search we designed and performed in 2009-2010 was the first to use GW data in this way. It is worth noting, however, that this is not the first time that a community of astronomers has been alerted to exciting astrophysical events as they happen. In particular, the Gamma-ray burst Coordinate Network (GCN), has been extremely successful in distributing sky positions of GRB events with low latency. This distribution of information has allowed the afterglows of GRBs to be observed in their earliest moments, and with a large variety of instruments, including radio, infrared, optical, ultra-violet, and x-ray observatories. This multi-messenger approach has informed our current understanding of GRBs, and delivered important evidence for developing the models of GRB progenitors discussed in Chapter 3.

In some ways, our search for common sources of gravitational and EM waves is patterned after the multi-messenger approach of the GCN.

4.2 Context for the search

The science runs of the LIGO and Virgo observatories in 2009 – 2010 (S6, VSR2, and VSR3) presented a fortuitous opportunity to implement this search. First, estimating the position of a GW source requires at least three detector sites. It is only in recent years that LIGO and Virgo have been sharing data, creating the first three site network of kilometer scale GW detectors. The S6 science run of LIGO, in addition to being only the second science run to operate in coincidence with Virgo, was also the last science run before the installation of Advanced LIGO. Implementing a low-latency search during S6 created an opportunity for the tools, protocols, and coordination needed to respond quickly to GW triggers to be developed before Advanced LIGO. Having overcome many of the obstacles, and jump-started work on others, it is now possible that some of the earliest detections with advanced interferometers will be accompanied by EM observations.

In addition, the timing of the S6 run seemed to match well onto current developments within the optical astronomy community. Source localization with the LIGO/Virgo network, discussed more in Chapter 5, is highly imprecise by astronomy standard: a given signal may have uncertainty regions of tens or even hundreds of square degrees. This suggests that finding associated optical transients may re-

quire imaging large patches of sky. Fortunately, there is currently a trend within the optical astronomy community towards building large area survey instruments, including the telescopes known as Pan-STARRS, SkyMapper, QUEST, and the Palomar Transient Factory (PTF). Some of this development in optical surveys is in anticipation of the construction of a next generation telescope, the Large Synoptic Survey Telescope (LSST), which is expected to survey the sky with a depth and cadence never before seen. All of these projects include in their science objectives the goal of finding transient objects in large areas of sky.

Finally, the current context of the search for gravitational wave signals suggested another reason that S6/VSR2/VSR3 was the right time to implement this search. The LIGO interferometers have been increasing in sensitivity over the past decade. During the S5 run, they met their design sensitivity, and in the S6 run, exceeded the sensitivity of the initial design. While astrophysical predictions suggest that second generation detectors, such as Advanced LIGO, are needed for GW detections to be assured, there was a sense that a first detection with initial LIGO and Virgo was possible. With this in mind, an exciting motivation for implementing the search for optical counterparts was the fact that an EM counterpart, if discovered, could build confidence that a signal in the GW data is truly of astrophysical origin. That is, discovering an EM counterpart could confirm an otherwise speculative GW detection. Burst searches, in particular, have been shown to be sensitive to detector artifacts that are difficult to distinguish from an unmodeled signal. Some people

have even speculated that virtually *any* GW signal discovered by a burst algorithm would be unconvincing as a first detection, due to the possibility that a similar signal could arise from detector noise. In this sense, the optical counterpart search may not only increase the knowledge we could extract from a detection, it may also effectively increase the sensitivity of the interferometer network.

To use EM transients for this purpose, the time-domain sky in the wavelength of interest must be well understood. Transients that are found in directional and time coincidence with GW triggers would increase confidence in the GW signal only if the chance of a similar, incidental coincidence is understood to be low [54].

4.3 Search Latency and Observing Cadence

To use GW triggers to guide observations, it is important to complete the GW analysis and deliver the estimated source coordinates before the sought optical transient has faded from view. Taking characteristics of GRB afterglows as a model for the expected transient implies that observing quickly presents a major advantage. SGRB afterglows have been observed as early as a few minutes after the prompt gamma-ray emission, and then fade quickly. The power law fading, observed for both SGRB and LGRB afterglows, suggests that one hour after the trigger time the transient would be ~ 20 times as bright as it would be a day after the trigger time, or a difference of ~ 3 in apparent magnitude. In order the facilitate observations soon after the trigger time, we set a goal of delivering the estimated source coordinates

within roughly thirty minutes of the trigger time.

On the other hand, observations spanning several nights after the trigger time also seem to be important. The theoretical kilonova is expected to take about a day to reach its peak luminosity. Supernovae are known to brighten over several days, and may not reach their peak luminosity until over a week after the initial collapse. Even for an optical transient whose peak brightness occurs soon after the GW trigger time, observations over subsequent nights are important in order to characterize the behavior of the light curve; that is, observations over multiple nights may be needed to identify interesting transients in the EM data. So, for our search, we planned to request a first observation of the estimated source location as soon as possible, meaning about thirty minutes after the trigger time, and then again over several nights. In some cases, observations were taken out to a month after the trigger time, in order to allow for the long brightening times that are seen in supernovae.

4.4 Galaxy Weighting

Knowing where to look for the counterpart to a GW trigger is challenging. Directional estimates of low signal to noise ratio (SNR) binary merger sources with the 2009/2010 GW interferometer network have uncertainties of several tens of square degrees [26] or more. This suggests using telescopes with a field of view (FOV) of at least a few square degrees if possible. Even with such a "wide field" instrument,

there is a striking mismatch between the large area needed to be searched, and the size of a single FOV.

The problem may be partially mitigated by making use of the known mass distribution in the nearby universe. Astrophysical sources of GWs should reside in galaxies, not in the vast stretches of space between galaxies. A search for GW counterparts may dramatically reduce the needed sky coverage by focusing observations on galaxies within the distance limits of the GW detectors [42, 71]. Limiting the search area to known galaxies may also improve the feasibility of identifying the true counterpart from among other objects with time-varying EM emissions [54]. Even within the Milky Way, a search may emphasize known targets by seeking counterparts within globular clusters, where binary systems of compact objects may form efficiently [72].

An emphasis on extragalactic sources has the potential drawback that any counterparts in our own Milky Way galaxy may be missed. Also, some neutron star mergers may occur at large distances from their host galaxies, and so would not be observed. However, this population should be small [15, 45].

Our selection of fields to observe was weighted towards areas containing known galaxies within 50 Mpc. The utilized catalog of nearby galaxies and globular clusters, and the process for selecting fields to observe, is described in Chapter 7.

Chapter 5

Gravitational Wave Event Reconstruction

The S6/VSR2/VSR3 low-latency search for GW candidates relied on three data analysis pipelines to identify reconstruct GW events. This chapter begins with an introduction to reconstructing source positions with a network of gravitational wave detectors, along with an overview of the types and scale of uncertainties that may be expected with the LIGO/Virgo network. Then, the three data analysis pipelines are introduced. Each pipeline uses a unique algorithm to assign every GW trigger a detection statistic, used to evaluate the significance of the trigger, and a skymap, used to estimate the source position.

5.1 Position Reconstruction using Triangulation

Determining the precision with which sources may be localized with the network of GW observatories is central to assessing the feasibility of follow-up observations. A variety of algorithms may be used to determine source position. However, every approach must ultimately rely on determining shifts in time-of-arrival of the waveform at each of the three detector sites. The measured travel times between sites, coupled with the known location of each observatory, leads to an estimate of the source direction.

This section presents a simple model of the triangulation technique to explore the relationship between timing errors in the interferometer and position errors on the sky. This, along with similar explorations in the literature, will provide a sense of the scale of the position errors to expect. The triangulation solution is described by Jaranowski and Krolak in [39]. In this model, a sharp feature of the gravitational waveform (say, the merger time of two neutron stars) is measured at a particular time, t_α, in each of the three interferometers, which are located at positions on the earth $\vec{r}_\alpha$.

The coordinates may be chosen to place the x-axis parallel to the line connecting interferometers 1 and 2, and the y-axis chosen so that the third interferometer also lies in the x-y plane. This may be expressed as

$$\vec{r_{12}} \equiv \vec{r_2} - \vec{r_1} \tag{5.1}$$

$$\vec{r_{13}} \equiv \vec{r_3} - \vec{r_1} \tag{5.2}$$

$$\vec{r_{12}}/c = (a, 0, 0) \tag{5.3}$$

$$\vec{r_{13}}/c = (b_1, b_2, 0) \tag{5.4}$$

where the coordinates are chosen so that a and b_2 are positive.

We imagine a GW signal is traveling in the direction (θ, ϕ). Then a vector $\vec{\Theta}$ pointing at the source may be written as

$$\vec{\Theta} = (-\sin\theta\sin\phi, -\sin\theta\cos\phi, -\cos\theta) \tag{5.5}$$

with θ and ϕ representing the polar coordinates about the $\hat{z}$ coordinates of the

propagation direction of the signal.

Generally, the travel time $\tau_{\alpha\beta}$ to a point $\vec{r_\alpha}$ from reference point $\vec{r_\beta}$ is

$$\tau_{\alpha\beta} = \vec{\Theta} \cdot \frac{1}{c}(\vec{r_\beta} - \vec{r_\alpha}) \tag{5.6}$$

The travel times between the interferometers may then be expressed as

$$\tau_{12} = t_1 - t_2 = \vec{\Theta} \cdot \vec{r}_{12}/c \tag{5.7}$$

$$\tau_{13} = t_1 - t_3 = \vec{\Theta} \cdot \vec{r}_{13}/c \tag{5.8}$$

The problem of source localization, then, is to determine the angles (θ, ϕ) from the measured values of τ_{12} and τ_{13}. Jaranowski and Krolak reported the following solution:

$$\sin\theta = \frac{1}{ab_2}\sqrt{\Delta} \tag{5.9}$$

$$\sin\phi = \frac{-b_2\tau_{12}}{\sqrt{\Delta}} \tag{5.10}$$

$$\Delta \equiv (b_1\tau_{12} - a\tau_{13})^2 + (b_2\tau_{12})^2 \tag{5.11}$$

Notice that these equations leave two possibilities for the value of θ, reflecting the fact that timing alone leads to two possible solutions, which are mirror images through the detector plane.

To see how a small error in measuring the time of arrival in detector α effects the reconstructed position, we can perturb these equations. The position error is:

$$\delta\theta = \sum_\alpha \frac{\partial\theta}{\partial t_\alpha}\delta t_\alpha \tag{5.12}$$

$$\delta\phi = \sum_{\alpha} \frac{\partial\phi}{\partial t_\alpha}(\sin\theta)\delta t_\alpha \tag{5.13}$$

where:

$$\frac{\partial\theta}{\partial t_1} = \frac{a\cos\phi - b_1\cos\phi - b_2\sin\phi}{ab_2\cos\theta} \tag{5.14}$$

$$\frac{\partial\theta}{\partial t_2} = \frac{b_1\cos\phi + b_2\sin\phi}{ab_2\cos\theta} \tag{5.15}$$

$$\frac{\partial\theta}{\partial t_3} = \frac{-a\cos\phi}{b_2\cos\theta} \tag{5.16}$$

$$\frac{\partial\phi}{\partial t_1} = \frac{-b_2 - \sin\phi[a\cos\phi - b_2\sin\phi - b_1\cos\phi]}{ab_2\sin\theta\cos\phi} \tag{5.17}$$

$$\frac{\partial\phi}{\partial t_2} = \frac{b_2 - \sin\phi[b_2\sin\phi + b_1\cos\phi]}{ab_2\sin\theta\cos\phi} \tag{5.18}$$

$$\frac{\partial\phi}{\partial t_3} = \frac{\sin\phi}{b_2\sin\theta} \tag{5.19}$$

At any point on the sky, θ and ϕ are orthogonal coordinates, so we may define the total angular error, $\delta\Theta$, as

$$\delta\Theta^2 = (\delta\theta)^2 + (\delta\phi\cos(\theta))^2 \tag{5.20}$$

Finally, we assume the timing errors in the interferometers are uncorrelated random variables with $< \delta t_1^2 > = < \delta t_2^2 > = < \delta t_3^2 > = < \delta t^2 >$. Then we may write the expectation value for the total angular error at each point on the sky as

$$\sqrt{< \delta\Theta^2 >} = \sqrt{\sum_{\alpha}\left(\frac{\partial\theta}{\partial t_\alpha}\right)^2 + \sum_{\alpha}\left(\frac{\partial\phi}{\partial t_\alpha}(\sin\theta)\right)^2} \times \sqrt{< \delta t^2 >}. \tag{5.21}$$

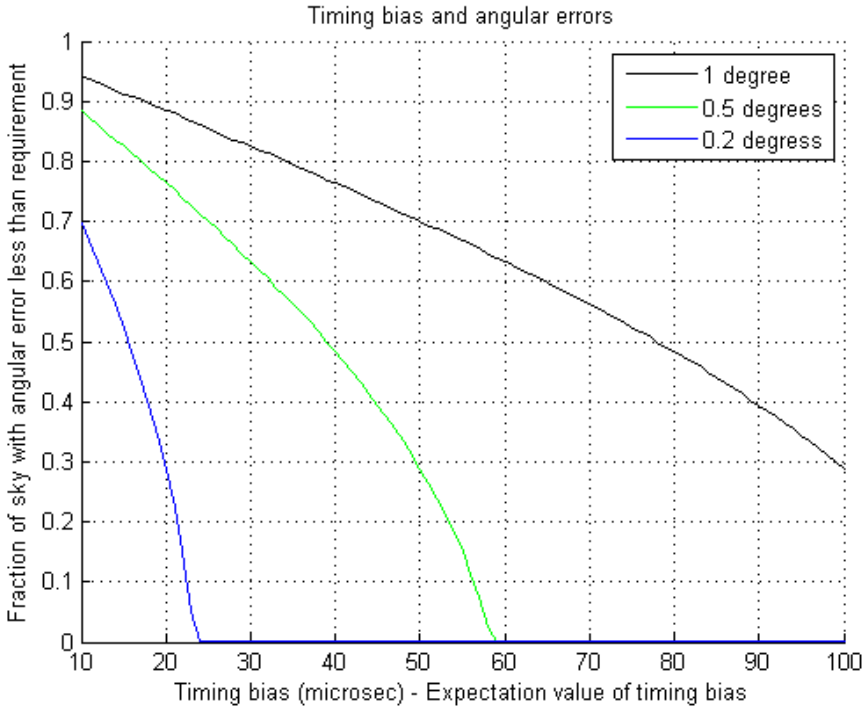

Figure 5.1: Figure showing the impact of timing errors on position reconstruction with the LIGO/Virgo network. The timing error shown on the x-axis is the expectation value of the timing error in all three observatories, assumed to be equal.

5.2 Expectations for the LIGO/Virgo Network

We can use the above considerations, as well as information from the literature, to develop expectations for position errors with the LIGO/Virgo network. The light travel time of 10 ms between the LIGO detectors may be used to construct an order-of-magnitude estimate of the relationship between timing errors and sky position errors. Plane waves traveling parallel to the baseline would take 10 ms to travel between the LIGO sites, while signals traveling perpindicular to the baseline would be detected simultaneously. A difference of 10 ms in measured time-of-arrival changes the estimated source position by 90 degrees, so in a naive picture, we may expect one degree of position error for each ~ 100 μs of timing error.

A more careful estimate is presented in Figure 5.1. This plot is a graphical representation of Equation 5.21, applied to the geometry of the LIGO/Virgo network. For the median sky position, 78 μs of timing error leads to 1 degree of position error. Or, as a convenient rule of thumb, we may look at the 70% sky position, and notice that $50\mu s$ of timing error leads to 1 degree of position error on the sky.

The errors in measuring the time of arrival of a GW signal may be divided into two categories: calibration errors and fitting errors. The calibration errors arise from uncertainties in how a signal propagates through the control loop in the interferometer. Fitting errors are errors in picking out the time of the feature of interest from a noisy signal. We would expect fitting errors to diminish for higher SNR signals, where calibration errors should depend only on our understanding of

the interferometers at a given time.

The calibration error budget for timing may, in turn, be divided into two categories: a time-domain offset known as the "timing error", and a frequency-dependent "phase error". The timing error budget is expected to be around 25 μs for the 2009/2010 network. The phase error varies across the spectrum, but is usually limited to between 5 and 10 degrees in the phase of the GW signal. For high frequencies, around 1000 Hz, this corresponds to small errors in timing. However, around the peak of the interferometer sensitivity, larger time offsets from phase errors are possible. For example, at 150 Hz, an 8 degree phase error implies 148 μs of timing offset. This suggests that calibration errors may lead to as much as 3 degrees of position error on the sky.

Estimates of fitting errors for various signals may be found in a paper by Steve Fairhurst [26]. For example, a signal from a binary neutron star system detected with SNR = 7 is estimated to have an arrival time that may be measured with a precision of 230 μs, and the arrival time of the same waveform at SNR = 10 may be measured to 180 μs. Typically, an SNR of 7 is considered near the threshold for detection. So, for these low SNR signals, we may expect fitting errors to be larger than calibration errors. Applying our rule of thumb, errors on the sky of around 4 degrees should be expected from these fitting errors. If this is interpreted as the radius of an error circle on the sky, then the implied area is around 50 square degrees.

As a final note, Fairhurst's paper also presents a useful scaling law for how the fitting errors relate to the signal properties. He finds

$$\sigma_t = \frac{1}{2\pi\rho\sigma_f} \tag{5.22}$$

where σ_t is the timing uncertainty, ρ is the SNR of the signal, and σ_f is the effective bandwidth of the signal. For low SNR signals, where the fitting errors dominate, this suggests angular error should scale with 1/SNR.

5.3 Event Reconstruction Algorithms

Searches for gravitational wave transients in large amounts of data require software packages to analyze the data, and identify times of unusual signal power, called triggers. These software packages are known collectively as event trigger generators (ETGs), or simply trigger generators.

The 2009/2010 low-latency search utilizes triggers from three ETGs. Each ETG software package includes algorithms for estimating the source position of each trigger. In all three cases, the source position information is reported as a list of trial positions, or pixels, from a grid that covers the sky. For each pixel, the ETGs assign a likelihood reflecting the probability that the direction represents the true source location. The resulting list of 0.4 x 0.4 degree pixels and associated likelihoods is known as the skymap. Only the likelihoods and positions of the 1000 most likely pixels are reported.

The ETGs are designed around different signal models, and so are sensitive to different types of waveforms, which partially justifies the use of multiple software packages. In this section, we briefly introduce the ETGs used in the search, the criteria used to generate triggers, and the three approaches to position reconstruction.

5.3.1 Coherent Wave Burst

Coherent Wave Burst (cWB) is a software package used to search for unmodeled, transient signals (bursts), and has been used in several previous analyses.

The signal energy in a GW burst signal, as it arrives at earth, may be characterized by the signal's "root sum squared" strain, h_{rss}:

$$h_{rss}^2 = \int_{-\infty}^{+\infty} (|h_+(t)|^2 + |h_\times(t)|^2) dt \tag{5.23}$$

where h_+ and $h_\times$ represent the strain in each of the two GW polarizations. The h_{rss} of a signal is independent of any detector, and $h_{rss}^2$is expressed in units of "strain-squared time". Strain is dimensionless, so h_{rss}^2 has units of seconds, or 1/Hz.

The signal observed in a given interferometer depends on the direction to the source and the orientation of the detector. This geometric dependency is expressed through the antenna pattern of each detector, F^+ and $F^\times$, so that the strain response of the interferometer is

$$h_{det}(t) = F^+(\vec{\Theta}) h_+(t) + F^\times(\vec{\Theta}) h_\times(t) \tag{5.24}$$

Notice that the antenna patterns depend on the source position direction, $\vec{\Theta}$.

Following the derivation in [18], the gravitational strain $h_{+,\times}$ of a plane wave from direction $\vec{\Theta}$ as measured at position $\vec{r_\alpha}$, is related to the strain at some reference position $\vec{r_0}$ by

$$h_{+,\times}(t + \tau_{\alpha 0}(\vec{\Theta}), \vec{r_\alpha}) = h_{+,\times}(t, \vec{r_0}) \tag{5.25}$$

where $\tau_{\alpha 0}(\vec{\Theta})$ is the propagation time from $\vec{r_0}$ to $\vec{r_\alpha}$, as in Equation 5.6. For a detector α at location $\vec{r_\alpha}$, we may then write the strain series of the output data as

$$d_\alpha(t + \tau_{\alpha 0}(\vec{\Theta})) = F_\alpha^+(\vec{\Theta}) h_+(t) + F_\alpha^\times(\vec{\Theta}) h_\times(t) + n_\alpha(t + \tau_{\alpha 0}(\vec{\Theta})) \tag{5.26}$$

where the term n_α represents detector noise, scaled appropriately to be represented as a strain. For the remainder of this section, we assume that the source position is known, and all data is time shifted appropriately as in Eqn 5.26, and so drop explicit references to $\tau_{\alpha 0}$ for brevity of notation.

Detector noise depends strongly on frequency, and so all searches for transient GWs depend on a transformation of the data from the time domain to the time-frequency domain. In cWB, this is accomplished using a wavelet transformation:

$$\tilde{d}_\alpha[m,n] = \int_{-\infty}^{\infty} \Psi_{mn} d_\alpha(t) dt \equiv \int_{-\infty}^{\infty} 2^{m/2} \Psi(2^m t - n) d_\alpha(t) dt \tag{5.27}$$

where the wavelet basis functions, Ψ_{mn}, are chosen so as to be an orthonormal basis for the Hilbert space of square-integrable functions. m and n are integers, and represent the frequency and time indices of the transformed data, $\tilde{d}_\alpha$. The wavelet transformation may be applied to both sides of Equation 5.26.

$$\tilde{d_\alpha}[m,n] = F_\alpha^+ \tilde{h}_+[m,n] + F_\alpha^\times \tilde{h}_\times[m,n] + n_\alpha[m,n]) \tag{5.28}$$

The symbol $\sigma_\alpha^2[m,n]$ is used to represent the variance of the transformed data when only noise is present, that is, $\sigma_\alpha[m,n]$ is the standard deviation of transformed noise. Using this definition, it is possible to "whiten" both sides of Equation 5.28, by dividing each term by $\sigma_\alpha[m,n]$.

$$\tilde{d}_{w\alpha} = F_{w\alpha}^+ \tilde{h}_+[m,n] + F_{w\alpha}^\times \tilde{h}_\times[m,n] + \tilde{n}_{w\alpha}[m,n] \tag{5.29}$$

$$\tilde{d}_{w\alpha} \equiv \frac{\tilde{d}_\alpha[m,n]}{\sigma[m,n]} \tag{5.30}$$

$$F_{w\alpha}^+(\vec{\Theta},[m,n]) \equiv \frac{F_\alpha^+(\vec{\Theta})}{\sigma[m,n]} \tag{5.31}$$

Notice that $F^{+}_{w\alpha}$ and $F^{\times}_{w\alpha}$ are functions of both the source position and frequency. For a network of three detectors, Equation 5.31 may be written in the equivalent form:

$$\begin{bmatrix} \tilde{d}_{w1} \\ \tilde{d}_{w2} \\ \tilde{d}_{w3} \end{bmatrix} = \begin{bmatrix} F^{+}_{w1} \\ F^{+}_{w2} \\ F^{+}_{w3} \end{bmatrix} \tilde{h}_{+} + \begin{bmatrix} F^{\times}_{w1} \\ F^{\times}_{w2} \\ F^{\times}_{w3} \end{bmatrix} \tilde{h}_{\times} + \begin{bmatrix} \tilde{n}_{w1} \\ \tilde{n}_{w2} \\ \tilde{n}_{w3} \end{bmatrix} \tag{5.32}$$

or

$$\tilde{\mathbf{d}}_{\mathbf{w}} = \mathbf{F}^{+}_{\mathbf{w}}\tilde{\mathbf{h}} + \mathbf{F}^{\times}_{\mathbf{w}}\tilde{\mathbf{h}} + \tilde{\mathbf{n}}_{\mathbf{w}} \tag{5.33}$$

where the boldface indicates vectors. Viewing Equation 5.32 as a vector equation, some important points become clear. On the left hand side, we see that the data vector exists in a three dimensional space, with each dimension corresponding to one detector. Looking at the right hand side, though, we see that in this three dimensional data space, the signal may only exist in the plane defined by the two vectors $\mathbf{F}^{+}_{\mathbf{w}}$ and $\mathbf{F}^{\times}_{\mathbf{w}}$. This means that we may remove some of the noise fluctuations in the data by simply projecting the data vector on to this two dimensional space.

To accomplish this, the + and × polarizations are defined in the dominant polarization wave frame [47], which has the property that $\mathbf{F}^{+}_{\mathbf{w}} \cdot \mathbf{F}^{\times}_{\mathbf{w}} = 0$. Then, a matrix may be formed as

$$\mathbf{P} = \begin{bmatrix} \mathbf{F}^{+T}_{\mathbf{w}} \\ \mathbf{F}^{\times T}_{\mathbf{w}} \end{bmatrix} \tag{5.34}$$

so that projection of the whitened data vector into the signal plane may be accom-

plished as $\mathbf{P}\tilde{\mathbf{d}}_{\mathbf{w}}$. The signal energy in this projected vector may be expressed as its squared magnitude, or

$$l_{\mathtt{max}}[m,n] = \tilde{\mathbf{d}}_{\mathbf{w}}[m,n]\mathbf{P}^{\mathbf{T}}[m,n]\mathbf{P}[m,n]\tilde{\mathbf{d}}_{\mathbf{w}}[m,n] \tag{5.35}$$

This may be summed over the samples, resulting in

$$L_{\mathtt{max}} = \sum_{m,n} \tilde{\mathbf{d}}_{\mathbf{w}}[m,n]\mathbf{P}^{\mathbf{T}}[m,n]\mathbf{P}[m,n]\tilde{\mathbf{d}}_{\mathbf{w}}[m,n] \tag{5.36}$$

$L_{\mathtt{max}}$ is the signal energy in the signal space projection of the data vector. It is also known as the maximum likelihood ratio [48], because it may be shown to be a near optimal statistic for identifying GW burst signals under some conditions. For this reason, the cWB detection statistic, η, is constructed from $L_{\mathtt{max}}$. Though not noted explicitly in Equation 5.36, recall that $L_{\mathtt{max}}$ depends on sky position both through the time shifts in the data vector, and in the direction dependent antenna patterns $F_\alpha^\times, F_\alpha^+$ that make up the matrix $\mathbf{P}$.

To see how cWB estimates the source coordinates of a signal, it is necessary to take a closer look at the separate terms within $L_{\mathtt{max}}$. We define the projection operator $\mathbf{Q} = \mathbf{P}^{\mathbf{T}}\mathbf{P}$, and write equation 5.36 as a summation over the detector indices α and β

$$L_{\mathtt{max}} = \sum_{m,n}\sum_{\alpha,\beta} \tilde{d}_{w\alpha}[m,n]Q_{\alpha\beta}[m,n]\tilde{d}_{w\beta}[m,n] \tag{5.37}$$

Notice from equation 5.37 that the projected signal energy, $L_{\mathtt{max}}$, contains "coherent" terms, that depend on cross correlations between the data streams in different detectors ($\alpha \neq \beta$), and "incoherent" terms that depend only on auto-

correlations of a single detector ($\alpha = \beta$). The coherent part is expected to have a strong dependence on sky position, and is given the label E_C.

$$E_c = \sum_{m,n} \sum_{\alpha,\beta} \tilde{d}_{w\alpha}[m,n] Q_{\alpha\beta}[m,n] \tilde{d}_{w\alpha}[m,n],\ \alpha \neq \beta \tag{5.38}$$

Also used in evaluation trial sky positions is a quantity E, representing the total energy in the whitened data stream before projecting into signal space.

$$E = \sum_{m,n} \sum_{\alpha} \tilde{d}_{w\alpha}[m,n] \tilde{d}_{w\alpha}[m,n] \tag{5.39}$$

Finally, we present the statistic, L_{sky}, that cWB uses to estimate source position [48]:

$$L_{sky} = \frac{L_{\texttt{max}} E_C}{E(E - L_{\texttt{max}} + |E_C|)} \tag{5.40}$$

In the denominator, the term $E - L_{\texttt{max}}$ is the energy in the data stream leftover after removing the signal energy. That is, this is the null energy, which can contain no GW signal energy, and is used to penalize the sky statistic. cWB calculates L_{sky} for each trial source direction in the skymap, and uses this value to rank source positions and calculate the probability distribution.

5.3.2 MBTA

The Multi-Band Template Analysis (MBTA) [63] searches for GW signals from compact binary systems. The analysis is an implementation of the matched template search used to search for inspiral signals [5]. Matched template searches exploit the fact that the sought waveform is known with good precision by cross-correlating the

whitened signal template with the whitened data. To describe the matched filter search, we follow the presentation in [11].

First, a "template bank" of the sought waveforms is constructed. For MBTA, the templates are constructed from a post-Newtonian approximation [16] to the solution of the Einstein equations. Each template, $h_m(t)$, represents the expected waveform from a class of compact binary systems, and depends only on the two masses of the component objects. In a real search, a finite collection of templates is used, so a single index m may be used to denote which template is represented.

Equation 1.8 shows that the amplitude of a GW signal scales as $1/D$, where D is the distance to the source. The amplitude of a signal detected in a particular interferometer will also depend on other factors, including the antenna pattern, $F_\alpha^+, F_\alpha^\times$, and the inclination angle of the binary system. These factors may all be combined to define the "effective distance", D_{eff}, which acts as a scaling factor for the amplitude of any detected inspiral signal.

$$D_{eff} = D\left[(F^+)^2\left(\frac{1+\cos^2\iota}{2}\right)^2 + (F^\times)^2\cos^2\iota\right]^{-1/2} \tag{5.41}$$

The inclination angle, ι, is the angle between the angular momentum axis of the binary system and the direction towards earth. D_{eff} is the distance at which an optimally located and oriented source would be placed to produce the observed amplitude.

In a given template bank, the amplitude of the templates are all scaled to have the same effective distance, typically set to $D_{eff} = 1$ Mpc. The templates, as

well as the data from a single gravitational wave detector, $d_\alpha(t)$, may be Fourier transformed with the definition

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt \tag{5.42}$$

so that we have $\tilde{h}_m(f)$ and $\tilde{d}_\alpha(f)$. To whiten both the detector data and template, so we define the one-sided power spectral density, S_α, of the noise in the detector data.

$$\langle \tilde{n}_\alpha \tilde{n}^*_\alpha \rangle = \frac{1}{2} \delta(f - f') S_\alpha(f) \tag{5.43}$$

To seek signals that match template h_m, the whitened data is cross-correlated with the whitened template

$$x(t) = 4\mathrm{Re} \int_0^\infty \frac{\tilde{d}(f)}{\sqrt{S(f)}} \frac{\tilde{h}^*_{\mathtt{m}}(f)}{\sqrt{S(f)}} e^{2\pi i f t} df = 4\mathrm{Re} \int_0^\infty \frac{\tilde{d}(f) \tilde{h}^*_{\mathtt{m}}(f)}{S(f)} e^{2\pi i f t} df \tag{5.44}$$

and the resulting time series, $x(t)$, is known as the filter output. MBTA is designed to run in low-latency, and calculates this integral quickly by dividing the analysis into two frequency bands, allowing more optimal frequency sampling.

Any given template may differ from an observed signal by a phase offset, related to the orbital phase of the binary. However, we can search for a template with any phase offset by taking the quadrature sum of the matched filters for two templates which differ only by a $\pi/2$ phase shift. This can be accomplished by simply keeping both the imaginary and real parts of the matched filter output

$$z(t) = x(t) + iy(t) = 4 \int_0^\infty \frac{\tilde{d}(f) \tilde{h}^*_{\mathtt{m}}(f)}{S(f)} e^{2\pi i f t} df \tag{5.45}$$

Then, $|z(t)|^2$ will be equivalent to the quadrature sum of the filter outputs associated with templates $\tilde{h}_m(f)$ and $i\tilde{h}_m(f)$.

To normalize the filter output, we define σ_m^2, which is the variance of $x_m(t)$ under the assumption that the data contains only stationary, Gaussian noise

$$\sigma_m^2 \equiv 4\int_0^\infty \frac{|\tilde{h}_m(f)|^2}{S(f)} df \tag{5.46}$$

σ_m is a measure of the expected magnitude of the filter output for detector noise, and so can be used to define the signal-to-noise ratio (SNR) of the signal, ρ_m

$$\rho_m(t) \equiv \frac{|z_m(t)|}{\sigma_m} \tag{5.47}$$

For stationary, Gaussian noise, ρ_m is unlikely to be much larger than one, and so searching for a GW signal is associated with searching for times when $\rho_m(t)$ is large. MBTA calculates $\rho_m(t)$ for each detector and each template, and identifies single detector triggers as times when $\rho_m(t) > 5.5$. It then clusters triggers across the template bank and across times, and looks for coincident triggers between detectors. Only times with coincident triggers in all three detectors are considered as possible event candidates. MBTA uses the combined SNR of the three detectors, ρ_c, as its detection statistic:

$$\rho_c = \sqrt{\rho_H^2 + \rho_L^2 + \rho_V^2}. \tag{5.48}$$

MBTA estimates source positions based on the time of arrival measured for each single detector trigger. The times are measured as the point when a reference frequency of 140 Hz is crossed in the signal. For each trial sky position, $\vec{\Theta}$, MBTA

calculates the expected delay times between detectors, $\hat{\tau}_{\alpha\beta}$. It then calculates the quadrature sum of the difference between the expected time delays and the measured time delays

$$\Delta\tau_{rss} = \sqrt{\sum_{\alpha\neq\beta} \left[\hat{\tau}_{\alpha\beta} - \tau_{\alpha\beta}\right]^2} \tag{5.49}$$

where $\tau_{\alpha\beta}$ is the measured time delay between two detectors. Pixels in the skymap are marked more likely for smaller values of $\Delta\tau_{rss}$. The actual relationship between timing errors and likelihood is determined by measuring the distribution of $\Delta\tau_{rss}$ in a collection of simulated signals.

In an attempt to break the mirror-image degeneracy of sky positions found by timing alone, MBTA uses an amplitude consistency check, defined as follows:

$$\Delta Q_{rss} = \sqrt{\sum_{\alpha\neq\beta} \left(\frac{D^2_{eff(\alpha)} - D^2_{eff(\beta)}}{D^2_{eff(\alpha)} + D^2_{eff(\beta)}} - \frac{Q^2_\alpha - Q^2_\beta}{Q^2_\alpha + Q^2_\beta} \right)^2} \tag{5.50}$$

The first term is the fractional difference in the measured values of D_{eff} between two interferometers. In fact, we would expect some non-zero difference in D_{eff} as measured in two different detectors because it depends on the orientation of the interferometer, expressed through the antenna pattern, $F^+, F^\times$ (see equations 5.41 and 2.17). The approximate expected difference in D_{eff} measured at two detectors is expressed in the second term in equation 5.50, where the quantity Q^2_α is defined as

$$Q^2_\alpha = \frac{1}{[F^+_\alpha(\vec{\Theta}, \psi = 0)]^2 + [F^\times_\alpha(\vec{\Theta}, \psi = 0)]^2} \tag{5.51}$$

As with the time delays, the distribution of ΔQ_{rss} is measured using simulated

signals. This distribution is used to create the skymap for each event candidate. In this way, both the consistency in timing, as measured through $\Delta\tau_{rss}$, and the consistency in amplitude, ΔQ_{rss}, found when assuming each trial sky location, is used to create the skymap.

5.3.3 Omega Pipeline

The Omega pipeline searches for burst signals in two stages: a coincidence stage and a coherent, "Bayesian" stage. The coincidence stage has been used in previous burst analyses [3]. The coherent stage uses cross-correlations between detectors to estimate the position of sources.

The first stage acts as a matched filter search for templates that are approximately sine-Gaussian waveforms in the whitened data, where a sine-Gaussian has the functional form:

$$h_{SG}(t) = h_0 \exp\left[\frac{-4\pi^2 f^2}{Q^2}(t-t_0)^2\right] \sin\left[2\pi f(t-t_0)\right] \tag{5.52}$$

where h_0 is a scaling constant, and t_0, f_0, and Q give the parameters of the template. A template bank of this form is chosen. Then, the whitened data from a single detector, $d_{w\alpha}$, is transformed into the space of these templates by the so-called "Q Transform"

$$X(t_0, f_0, Q) = \int_{-\infty}^{\infty} d_{w\alpha} h_{SG}(t, t_0, f_0, Q) dt \tag{5.53}$$

This filter output is normalized, so that the expectation value of X with Gaussian noise is unity. Then, as in the inspiral case, triggers are sought where a given

normalized filter output has a value much larger than one.

In the online version of Omega pipeline, coincident triggers across the three detectors are identified. The loudest set of triple coincident triggers in each 64 s block of data is then treated coherently by the Bayesian framework. This Bayesian framework calculates attempts to calculate the ratio of the likelihood functions for each triple coincidence trigger

$$\frac{p(\vec{d_w}|H_1)}{p(\vec{d_w}|H_0)} \tag{5.54}$$

where the numerator represents the probability of a signal leading to the whitened data vector, d_w, and the denominator represents the probability of noise alone giving rise to the observed data.

The analysis assumes that the waveform matches the sine-Gaussian template found by the first stage. It then calculates the likelihood functions in Equation 5.54 under the assumption that the waveform is coming from each direction in the skymap, while marginalizing over a range of possible amplitude values. These likelihood functions are then used to create the skymap. Summing the likelihood functions over the sky allows the pipeline to calculate a total likelihood for the data representing a true signal, known as "probSignal." This is used, along with other trigger information, to construct the detection statistic, Ω.

Chapter 6

Telescope Network

The 2009/2010 joint GW/EM low-latency search was designed and executed through a partnership between the LIGO Scientific Collaboration, the Virgo collaboration, and a number of groups within the observational community. This chapter describes some desirable features of a telescope network that aims observe EM counterparts to GW triggers. Also presented are issues related to observing cadences and coordination of search logistics. Finally, the network of radio, optical, UV, and X-ray telescopes that contributed to the search is described.

6.1 Providing Sky Coverage

One of the issues that arises when seeking optical data on rare, quickly fading events, is trying to ensure that at least one telescope will be available to take data when a trigger arrives. This is not an easy task - it means attempting to include enough telescopes in the network that essentially the whole sky will be covered at any one time. Of course, daylight prevents total sky coverage at any given moment, but we wanted as complete coverage as resources allowed.

A number of factors limit the range of sky, and the range of time, that a given telescope is able to observe. In addition to restrictions from daylight, a telescope's

viewing hours are impacted by weather conditions such as clouds or precipitation and technical problems that occasionally arise. The range of locations that a telescope may view is sometimes slightly impacted by the geometry of the mount and dome, however, the more fundamental limitation is atmosphere. Light is distorted and scattered as it passes through earth's atmosphere, degrading the quality of an observation. The effect is exacerbated as a telescope points at directions further from zenith, because at angles close to the horizon, the line of sight between the telescope and the target passes through a greater distance of atmosphere. For this reason, observatories typically restrict viewing to positions within about 70° of zenith.

Some examples of sky coverage with a network of telescopes are shown in Figures 6.1 and 6.2. The figures do not account for daylight, but merely show the sky positions that are within 70° of zenith for a particular location. Figure 6.1 shows the positions of the two TAROT telescopes in blue. While the two site TAROT network is able to observe both the northern and southern hemispheres, it is clear that at a given moment, much of the sky is not viewable, and there is only a small amount of redundancy. Figure 6.2 shows a network that includes TAROT, along with the four site ROTSE network, the Palomar Transient Factory, and SkyMapper. The magenta color indicates visibility from three or more telescopes, and is seen to cover much of the sky. In fact, with this network, nearly the entire sky is observable with at least two telescopes.

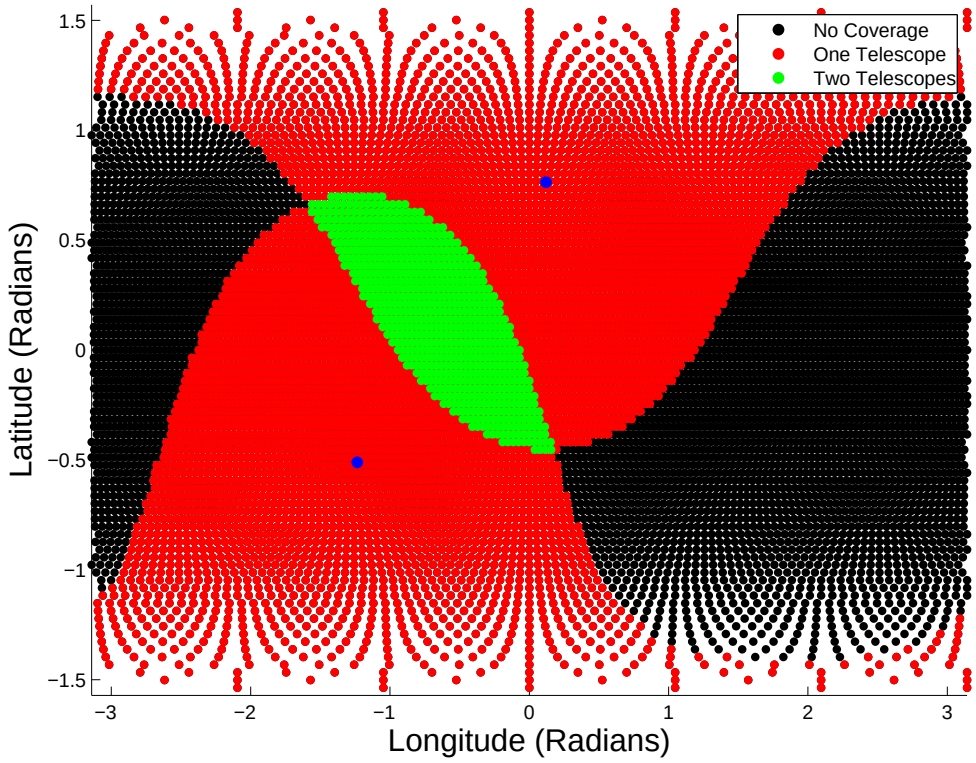

Figure 6.1: Figure showing the sky coverage by the two site TAROT network. The telescopes, marked as blue dots, are located near Calern, France, and La Silla, Chile. Red indicates coverage by one telescope, and green shows areas visible to both telescopes.

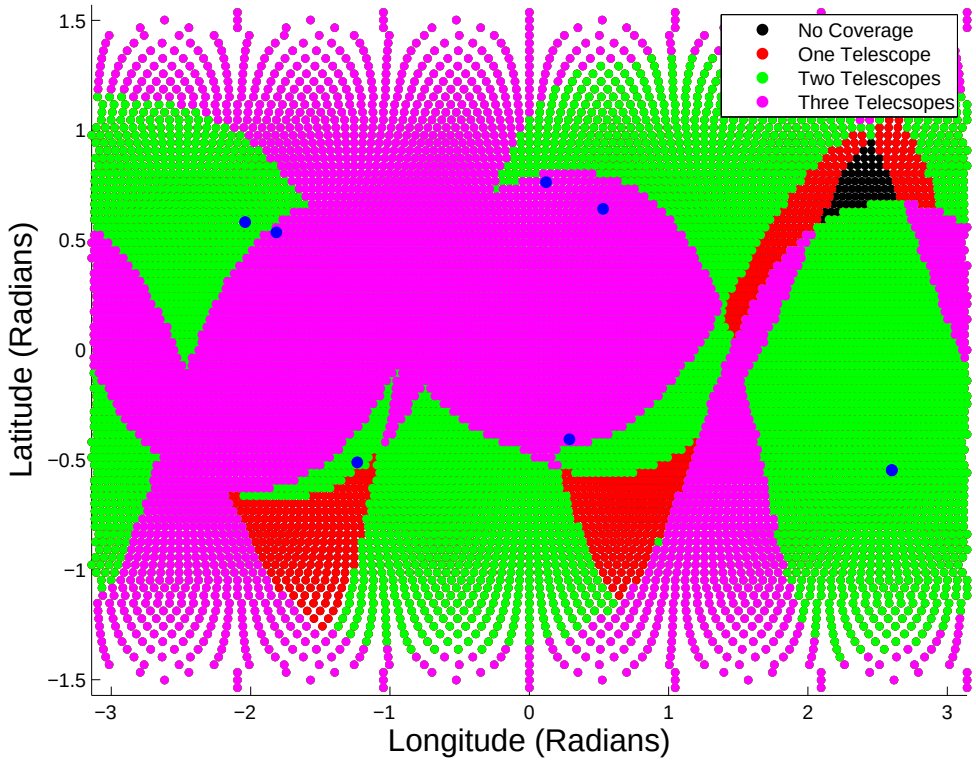

Figure 6.2: Figure showing the sky coverage of a telescope network comprised of the two TAROT telescopes, four ROTSE telescopes, Palomar Transient Factory, and SkyMapper. The magenta color indicates areas that are visible to three or more telescopes.

6.2 Telescope Selection

The large uncertainty in the gravitational wave position reconstruction presents a significant challenge to searches for EM counterparts. Where the GW uncertainty region may be tens of square degrees, a typical camera is designed to capture only a few arc minutes. Even a so-called "wide-field" camera might image only one or a few square degrees in each field, still not enough to cover the large, and irregularly shaped, GW uncertainty regions. For this reason, in seeking astronomer partners, we were led towards operators of wide-field cameras.

Another preferred feature for our search was automation of the telescope. So-called "robotic" telescopes are operated by software programs, and include self-maintenance features such as domes that close automatically in the event of bad weather. In seeking partners, we preferred robotic telescopes, both because they can respond to triggers almost immediately, and in order to avoid interrupting the observing plan of an astronomer who is on location.

These considerations led us to conclude that telescopes specializing in large area surveys seem to be a good fit. In addition to appropriate hardware, a partner with experience in large area surveys may also provide guidance and software tools for identifying transients. Modern surveys, such as Palomar Transient Factory and Pan-STARRS, collect large amounts of data each night. They employ sophisticated software for transient identification and classification in their hunts for the particular astrophysical targets of a given search.

In our search for partners, we were also attracted to projects that include GRB follow-up observations in their science objectives. GRB position information is sometimes imprecise, and so these telescopes also tend to be outfitted with wide field cameras. In addition, the "target of opportunity" mode of operation that our search demanded is very familiar to the GRB community. Not only are GRB afterglow telescopes almost always robotic, they already have software in place to receive real time alerts. In fact, in many cases, our software mimicked the GCN data format to interface smoothly with these telescopes.

In these considerations, there was no strong reason for preferring one observing band over another. For example, the afterglows to GRBs have been observed in every frequency range from radio to x-ray. However, the availability of observatories meeting our other preferences, namely wide field and robotic, was primarily in the optical. While radio band and x-ray band instruments were a minority in our list of partners, this is more a reflection of accessible resources than any scientifically motivated decision.

6.3 Arrangements with Partners

To orchestrate the close coordination between the LIGO/Virgo network and the observing partners, a number of details had to arranged. A basic feature of any observing arrangement is the amount of telescope time to be allotted. In fact, for this search the question was a bit complex, because it immediately led to asking

how the time would be used. For each GW event candidate, how many fields would be used in an attempt to cover the GW uncertainty region? Then, for each sky position, how many nights of follow-up images would be taken to trace the light curve? Finally, once the total number of fields for each trigger was known, we had to decide the rate of triggers the observatory would tolerate.

For the last question, we arrived at a suggestion that was accepted by most of our partners. We believed that a trigger with a false alarm rate of more than once a week could never be called a GW detection, no matter how spectacular an EM transient was found to be in association. We proposed to limit triggers to an average rate of between one per week and a few per week, and most of our partners were comfortable with the implied amount of telescope usage. The arrangements with each partner are described below.

Another important detail was a decision about what information would be delivered with each trigger. One could imagine sending the 1000 pixel skymap of likelihoods as the source position of each trigger. However, this would have put a burden on our astronomer partners to write software to make sense of this skymap. To alleviate this burden, we developed software that selects optimized pointings within a skymap, described in Chapter 7. For each trigger, the observers received the GPS time of the trigger, and the pointing information selected for the particular telescope. No other information about the trigger was included in the delivered message.

Other arrangements included the technical details of the trigger alerts, computer logins for accessing data and observing log files, and agreements about who would commit time to searching the collected EM images for transients. While we do not present these finer details here, it is worth noting that agreements on these issues were made on a case by case basis, so that both the technical and sociological interface with each partner was a custom fit. In many cases, a particular member or members of the LIGO/Virgo collaboration served as the liaison with a particular partner. As a final note on logistics, we highlight here that we had only completed arrangements with three of our astronomer partners, the QUEST, TAROT, and Swift collaborations, at the time of the first data collection period in December of 2009. See Figure 6.3 for a map of the complete network used in the second data collection period, which began in September of 2010. Each partner observatory is introduced briefly below.

6.4 Optical Instruments

The *Palomar Transient Factory* (PTF) [55, 80] operates a 7.3 square degree FOV camera mounted on the 1.2 m Oschin Telescope at Palomar Observatory. As the name suggests, PTF's primary science target are transients, including Supernovae, many of which are discovered in a survey mode. The PTF team preferred a lower rate of triggers, but devoted more fields per trigger, as compared with most of our other partners. For PTF, we targeted one trigger for every three weeks of cal-

endar time. PTF applied ten different sky positions to each trigger, and these were imaged over several nights with 60 s exposures, for an expected limiting magnitude of $R = 20.5$.

Pi of the Sky [62] observed using a camera with a 400 square degree FOV and exposures to limiting magnitude 11–12. It was located in Koczargi Stare, near Warsaw. The camera was a prototype for a planned system that will simultaneously image two steradians of sky. The target rate was approximately 1 per week in the autumn run, followed up with hundreds of 10 s exposures over several nights.

The *QUEST* camera [13], currently mounted on the 1 m ESO Schmidt Telescope at La Silla Observatory, views 9.4 square degrees of sky in each exposure. The camera was designed and built by the QUEST team, and is used for a variety of survey style science projects. Notably, a search for Kuiper Belt objects with the QUEST camera discovered one such object with a mass greater than Pluto's; in this way, the QUEST camera directly contributed to the demotion of our late ninth planet. The CCD's are arranged in four strips, with significant gaps in between, so that a roughly square field may be constructed by combining two dithers to form a nearly 20 square degree image. The telescope is capable of viewing to a limiting magnitude of $R \sim 20$. For each trigger, QUEST applied three such compiled fields, requiring a total of six pointings. Each pointing was observed twice the first night, with follow-up images possible by request on subsequent nights.

ROTSE III [10] is a collection of four robotic telescopes spread around the

world, each with a 0.45 m aperture and 3.4 square degree FOV. The ROTSE III network was built to respond to GRB triggers, and a combination of fast response times and nearly global sky coverage have made it extremely successful at finding GRB optical afterglows. No filters are used, so the spectral response is that of the CCDs, spanning roughly 400 to 900 nm. The equivalent R band limiting magnitude is about 17 in a 20 s exposure. The ROTSE team arranged for a series of thirty images for the first night, and several images on following nights, for each autumn run trigger, with a target rate of 1 trigger per week.

SkyMapper [68] is a survey telescope located at Siding Spring observatory in Australia. The mosaic camera covers 5.7 square degrees of sky in each field, and is mounted on a 1.35 m telescope with a collecting area equivalent to an unobscured 1.01 m aperture. SkyMapper's primary science goal is a complete digital survey of the Southern sky, and is exactly the type of survey telescope that seems to be a good fit for the large areas associated with GW triggers. It is designed to reach a limiting magnitude $g \sim 21$ (>7 sigma) in a 110 s exposure. SkyMapper accepted triggers in the autumn run with a target rate of 1 per week, with several fields collected for each trigger.

TAROT [49] operates two robotic 25 cm telescopes, one at La Silla in Chile and one in Calern, France. Like the ROTSE III system, each TAROT instrument has a 3.4 square degree FOV, and was designed to respond to GRB triggers and other exciting astrophysical events. A 180 second image with TAROT in ideal conditions

has a limiting R magnitude of 17.5. During the winter run, TAROT observed a single field during one night for each trigger. In the autumn run, the field selected for each trigger was observed over several nights. TAROT accepted triggers with a target rate of 1 per week.

Zadko Telescope [20] is a 1 m telescope located in Western Australia. The current CCD imager observes fields of 0.15 square degrees down to magnitude ~ 20 in the R band for a typical 180 s exposure. The control system for the Zadko telescope evolved over the course of developing the EM follow-up program. With help from the TAROT team, Zadko became a fully robotic telescope. For each accepted trigger in the autumn run, Zadko repeatedly observed the five galaxies considered most likely to host the source over several nights. The target trigger rate for Zadko was one trigger per week.

The *Liverpool telescope* [88] is a 2 m robotic telescope situated at the Observatorio del Roque de Los Muchachos on La Palma. For this project the RATCam instrument, with a 21 square arc minute field of view, was used. This instrumentation allows a five minute exposure to reach magnitude $r' = 21$. This was the narrowest FOV used in the search. This project was awarded 8 hours of target-of-opportunity time, which was split into 8 observations of 1 hour each, with a target rate of 1 trigger per week.

6.5 X-ray and Radio Band Instruments

LOFAR [27] is a dipole array radio telescope that observes the band between 10 and 250 MHz. It is part of a growing class of radio telescopes that rely on many, simple dipoles and fast processing power instead of large dishes. LOFAR is based in the Netherlands, and there are plans to build additional stations at various locations across Europe. During the autumn run, LOFAR accepted triggers at a target rate of 1 trigger per week, with a four hour observation for each accepted trigger in several frequency bands near 140 MHz, providing a $\sim$25 square degree field of view.

The *Swift* satellite [34] carries three instruments, each in different bands, and has been extremely successful observing both the prompt emission and afterglows of GRBs. Swift granted several target of opportunity observations with two of its instruments, the X-ray Telescope (XRT) and UV/Optical Telescope (UVOT), for the winter and autumn observing periods. The XRT is an imaging instrument with a 0.15 square degree FOV, sensitive to fluxes around 10^{-13} ergs/cm^2/s in the 0.5-10 keV band. A few fields were imaged for each trigger that Swift accepted.

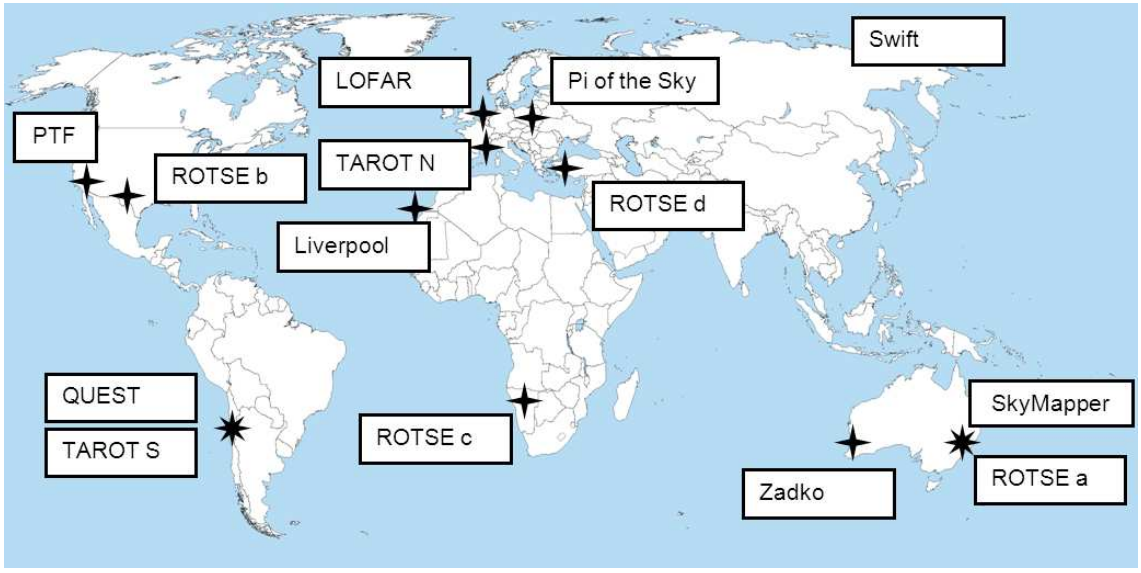

Figure 6.3: A map showing the approximate positions of telescopes that participated in the project. The Swift satellite observatory is noted at an arbitrary location. The image is adapted from a blank world map placed in the public domain by P. Dlouhý.

Chapter 7

Field Selection

The uncertainty in the position reconstruction of GW sources is expected to be larger than the FOV of a typical telescope. This creates a need to carefully choose pointings, or tiles, that optimize the probability of capturing the true gravitational wave source. For the EM counterpart search, we constructed a method of tile selection which is strongly weighted towards fields that include galaxies within 50 Mpc. This chapter describes the algorithm that was used to select pointings for each candidate event, and the catalog of nearby galaxies that was created for the search. The tile selection algorithm is included as part of the LOOC UP Management and INterface (LUMIN) software package, which was designed to support the EM counterpart search.

7.1 Galaxy Catalog

The Gravitational Wave Galaxy Catalog (GWGC) [97] was created to help this and future searches quickly identify nearby galaxies. The catalog contains up-to-date information compiled from the literature on sky position, distance, blue magnitude, major and minor diameters, position angle and galaxy type for 53,225 galaxies and 150 Milky Way globular clusters. The GWGC is complete out to

$\sim$40 Mpc at 15th magnitude and includes galaxies out to $\sim$100 Mpc. The catalog improves on the issue of multiple entries for the same galaxy suffered by previous catalogs by creating the GWGC from a subset of 4 large catalogs, each of which lists a unique Principal Galaxy Catalogue (PGC) number for every galaxy [74]. The catalogs used were: an updated version of the Tully Nearby Galaxies Catalog [89], the Catalog of Neighboring Galaxies [43], the V8k catalog [90], and HyperLEDA [75]. Also included is a list of 150 known Milky Way globular clusters [37]. These are all available freely online, but a local, homogeneous list is essential for rapid follow-up purposes.

7.2 Galaxy Weighting for Burst Triggers

The likelihood skymap, described in Chapter 5, is taken to be one factor in how probable each pixel on the sky is to be the true source of gravitational radiation. In addition, we expect the source to be associated with some galaxy within the 50 Mpc limit of the search. Some galaxies are more likely hosts of detectable GWs than others, and so a weighting scheme based on the galaxy's stellar mass and distance is employed to assign a likelihood to each galaxy [71].

Galaxies which contain more stars are assumed to contain more potential sources of GW transients. The weighting scheme uses a term proportional to the galaxy's blue light luminosity, M_{gal}, as a likelihood for being the correct source. Blue light luminosity is assumed to be a good tracer of the star formation rate in

a galaxy, and so follow the mass contained in neutron stars [6, 77]. The values of blue light luminosity contained in the catalog are expressed as an absolute magnitude. LUMIN converts this magnitude, m_b, to units of the Milky Way's blue light luminosity as

$$M_{gal} \equiv 10^{\frac{-20.5-m_b}{2.5}}. \tag{7.1}$$

Some objects in the catalog are listed without an estimate of the blue absolute magnitude. These objects are assigned a value of M_{gal} corresponding to 10% of the Milky Way luminosity.

Sources of weaker GW bursts are assumed to be more numerous in the universe than stronger burst sources. So, a nearby galaxy emitting weak GWs is expected to be more likely than a very distant galaxy emitting strong GWs. This is expressed in an ad hoc way by assuming a power law distribution for the number of GW sources as a function of intrinsic amplitude, A, scaled to be the h_{rss} of the signal as measured at one distance unit. The power law is then

$$\frac{dN_{sources}}{dA} = CA^{-\alpha}, \tag{7.2}$$

where C is an unknown constant.

In the simplest model of the detection process, assume there exists a minimum amplitude that can be detected at earth, h_{thresh}. Then, for any given distance D_{gal}, we may define the smallest intrinsic amplitude that we can detect, $A_{\texttt{MIN}}$, and relate it to the other quantities as

$$\frac{A_{\texttt{MIN}}}{D_{gal}} = h_{thresh} \tag{7.3}$$

$$A_{\texttt{MIN}} = h_{thresh} D_{gal}. \tag{7.4}$$

The distance prior to be assigned to each galaxy is proportional to the number of observable signals in that galaxy.

$$N_{obs}(D_{gal}) = \int_{A_{\texttt{MIN}}}^{\infty} \frac{dN}{dA} dA \tag{7.5}$$

N_{obs} is seen to be a function of D_{gal}, since $A_{\texttt{MIN}}$ is a function of D_{gal}.

Plugging Eqn 7.2 into the integral, assuming $\alpha > 1$, and carrying out the integration, we find:

$$N_{obs}(D_{gal}) = \frac{C}{-\alpha + 1} \left[\infty^{-\alpha+1} - A_{\texttt{MIN}}^{-\alpha+1} \right] \tag{7.6}$$

$$N_{obs} \propto A_{\texttt{MIN}}^{-\alpha+1} \tag{7.7}$$

Using Equation 7.4 leads to:

$$N_{obs} \propto D_{gal}^{-\alpha+1} \tag{7.8}$$

For the search, α is set to 2, and so a factor of D_{gal}^{-1} is applied to each galaxy's likelihood. Combining the mass and distance weightings, we define the final statistic used to evaluate each sky position

$$P \propto \frac{L M_{gal}}{D_{gal}} \tag{7.9}$$

where L is the likelihood calculated by the ETG based on the GW data.

This statistic, P, is evaluated for every pixel in the skymap. Pixels not associated with any object in the GWGC are given a probability of 0, and so are discarded. The process of associating pixels with galaxies is described below. In

practice, most galaxies are associated with several pixels, and the weighting factor M_{gal}/D_{gal} is divided among them.

7.3 Galaxy Weighting for MBTA Triggers

A different procedure is used to weight the MBTA triggers, since they are based on a specific source model. The distance weighting described in Section 7.2 does not apply to the distribution of NS merger signals. Instead, it seems more reasonable to assume that all NS mergers emit essentially the same strength GW signals, so there is no reason to prefer a closer galaxy to a farther one within the detector horizon distance. On the other hand, through the effective distance described in section 5.3.2, D_{eff}, each MBTA event candidate carries a measure of the maximum possible distance to the source. The MBTA position reconstruction algorithm uses the smallest effective distance measured with a single detector as an upper limit on the source distance. Galaxies within the GWGC at distances greater than this limit are discarded, and the abbreviated catalog is used to assign the statistic

$$P_i = L_i M_{gal} \tag{7.10}$$

to each pixel. This procedure occurs in the pipeline before the skymap is delivered to LUMIN, and so LUMIN applies no farther weighting procedure to MBTA skymaps.

7.4 LUMIN Handling for Large Galaxies

Galaxies are divided into two classes, large or small, determined by comparing the semi-major axis to the pixel resolution. The pixel resolution is the spacing between pixels in the skymap, and is approximately 0.4 degrees in all cases. Only a small fraction of objects in the GWGC meet the criteria for "large".

Large galaxies are imagined to be circles on the sky, with the semi-major axis as the radius. Every pixel with a center position inside this circle is associated with the galaxy. The prior $\frac{M_{gal}}{D_{gal}}$ is divided evenly over all of these pixels, so that each pixel receives a weight P_i

$$P_i = L_i \times \frac{M_{gal}}{D_{gal} N_{pixels}} \tag{7.11}$$

where N_{pixels} is the number of pixels falling within the circle of the galaxy.

7.5 Handling for Small Galaxies

Galaxies with a semi-major axis smaller than the pixel size may also have their prior divided over several pixels. The reason is that, in general, a galaxy will always lie between the center position of several pixels, and so cannot be unambiguously assigned to a single pixel. Moreover, the angular scale of the pixel resolution, 0.4 degrees, is less than the angular scale of the uncertainty in the skymap due to calibration uncertainty. This means, for example, that if a pixel with an especially large value of L is ~ 0.4 degrees away from a galaxy, we may reasonably associate

the pixel and the galaxy, even if some other pixel has a center position which is closer to the galaxy.

The procedure for small galaxies essentially "fuzzes out" the position of the galaxy, so that it is spread over an angular radius of approximately the pixel resolution, by convolving the location of the galaxy with a Gaussian function with a width, σ, of half the pixel resolution, or 0.2 degrees.

To implement the procedure, LUMIN first identifies all of the pixels within a box around the galaxy that is 0.8 degrees on a side. The angular distance between the center position of the galaxy and each pixel, a, is calculated. The weight applied to each pixel is then

$$W_i = C\mathrm{e}^{\frac{-a^2}{2\sigma^2}} \tag{7.12}$$

$$P_i = W_i L_i \tag{7.13}$$

where the normalization constant, C, is the same for each pixel associated with the galaxy, and is determined by the condition

$$\sum_i W_i = \frac{M_{gal}}{D_{gal}} \tag{7.14}$$

with the sum over all pixels associated with a single galaxy.

This spreads the prior associated with a single galaxy over the pixels near it, putting the most emphasis on pixels within the pixel resolution.

7.6 Choosing Tiles

The statistic, P, is calculated for each pixel as described above. For each telescope, then, the "best tile" is defined as the field that maximizes the total value of P captured within the field.

Given some trial tile, LUMIN counts the P value of a pixel as being included within the field if the center position of the pixel lies within the field. Note that this is a reasonable approach if the pixel size is much smaller than the FOV of the telescope. In the S6/VSR2+3 science run, most of the telescopes had fields between 1.8 and 4 degrees, compared with the 0.4 degree pixel resolution.

Since the pixels lie on the surface of a sphere (that is, the sky), the determination of which pixels lie within a field is slightly non-trivial. The rectangular field of the telescope must be projected onto the sphere of the sky.

First, the angular position of each pixel is expressed in Cartesian coordinates, with all of the pixels lying on a sphere of unit radius. Let θ_i^D and ϕ_i^R be the DEC and RA of pixel i. These positions may be expressed in Cartesian coordinates, x_i, y_i, z_i, as:

$$z_i = \sin(\theta_i^D) \tag{7.15}$$

$$x_i = \cos(\theta_i^D)\cos(\phi_i^R) \tag{7.16}$$

$$y_i = \cos(\theta_i^D)\sin(\phi_i^R) \tag{7.17}$$

(7.18)

A given pointing of a telescope with a rectangular FOV may be described by the DEC and RA of the center of the tile (θ_f^D, ϕ_f^R), along with the angular size of the field in both directions $(\Delta_{DEC}, \Delta_{RA})$. In general, a third angle is needed to account for the orientation of the field, but this is not currently implemented in LUMIN, because the telescopes in the network have equatorial mounts that maintain the orientation of the camera.

To easily find which pixels lie within the FOV, the coordinate system of the pixels, x_i, y_i, z_i, is rotated so that the z-axis is aligned with the center of the FOV. This is done using an Euler rotation.

First, the coordinate system is rotated around its own z-axis, so that the center of the field lies in the y-z plane.

$$\alpha = \phi_f^R - \frac{3\pi}{2} \tag{7.19}$$

$$\mathrm{D} = \begin{bmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{7.20}$$

Then, the coordinate system is rotated about the new x-axis, so that the positive z-axis pierces the center of the field

$$\beta = \frac{\pi}{2} - \theta_f^D \tag{7.21}$$

$$\mathrm{C} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\beta & \sin\beta \\ 0 & -\sin\beta & \cos\beta \end{bmatrix} \tag{7.22}$$

The location of the pixels in the new coordinate system (x_i', y_i', z_i') may be calculated as

$$\begin{bmatrix} x_i' & y_i' & z_i' \end{bmatrix} = C \times D \times \begin{bmatrix} x_i \\ y_i \\ z_i \end{bmatrix} \tag{7.23}$$

In the primed coordinate system, if we project the field of view of the telescope onto the unit sphere, then the center of the tile is located on the z-axis. From this point, moving in the north/south direction corresponds to moving in the y direction, and moving in the east/west direction corresponds to moving in the x direction. So, as long as the angular size of the field is a small fraction of the sky, $\Delta_{RA} \simeq \Delta_{DEC} \ll \pi$, a point can be seen to be inside the field of view only if it meets the following conditions:

$$x_i' < \sin(\Delta_{RA}/2) \approx \Delta_{RA}/2 \tag{7.24}$$

$$y_i' < \sin(\Delta_{DEC}/2) \approx \Delta_{DEC}/2 \tag{7.25}$$

$$z_i' > 0 \tag{7.26}$$

Notice that the last condition is important, as the other two conditions could be met by a point in the opposite direction of the tile.

The key advantage of rotating the coordinates in this manner is that the check

is equally valid for any choice of field location. Failing to rotate the coordinate system typically leads to algorithms that break down near the poles.

To find the "best tile", LUMIN begins with a list of possible center positions for the tile. The science goals associated with the QUEST and PTF telescopes include supernovae searches that repeatedly observe a fixed tiling of the viewable sky. Choosing fields which exactly overlap these supernovae search fields was desirable, as fields from the supernovae search could potentially be used as reference images for comparison when seeking transients associated with GW triggers. For these telescopes, LUMIN uses the entire list of supernovae fields as trial tiles. For the other telescopes, LUMIN uses the location of each pixel in the skymap as a trial tile center position. For each trial tile, LUMIN uses the above procedure to determine which pixels fall within the tile, and then sums the P values of all pixels within the tile to get a total P statistic. The tile with the highest P statistic is chosen as the "best tile". If multiple pointings are allowed with a telescope, LUMIN removes the pixels observed in the first tile from the skymap, and then repeats the procedure until the desired number of tiles are found.

Nowhere in this procedure does LUMIN account for the fact that some parts of the sky will always be invisible to a given telescope, due to such factors as the location of the telescope on the earth. This is an intentional choice, so that the selected best tile will always be the most likely sky position based on the GW data, regardless of the viewability of that location by a particular telescope. In cases where

all of the tiles selected were unviewable for a particular telescope, that telescope was not alerted to the GW trigger.

7.7 Special Handling for Small FOV Telescopes

The process for choosing the best tile described above is designed with the assumption that the FOV is larger than the pixel resolution. Trial fields, which are spaced to the pixel resolution, are intended to over-cover the sky. Moreover, most galaxies in the GWGC have an angular extent that is smaller than the pixel resolution, and so if the center position of the galaxy is within the field, then the entire galaxy is likely to be in the field.

On the other hand, if the FOV of a telescope is less than the pixel resolution, then we need some way to more precisely choose the position of the tile.

To do this, the field selection scheme treats the Zadko and Liverpool telescopes as special cases. Instead of the process described above, which is designed to account for fields which may contain several galaxies of interest, the procedure for small FOV telescopes selects fields to image individual galaxies of interest. The same weighting statistic is used as above,

$$P_{gal} = \frac{LM_{gal}}{D_{gal}}. \tag{7.27}$$

However, for the small FOV telescopes, instead of calculating P for each pixel in the skymap, P_{gal} is calculated for each galaxy. For each galaxy within 50 Mpc in the GWGC, and with an angular position roughly within the pixel resolution of a

non-zero pixel in the skymap, the procedure identifies L as the likelihood statistic of the pixel with the smallest angular distance from the galaxy. Using this and the galaxy's properties, P_{gal} is calculated for each galaxy. Then, the galaxies are placed in a sorted list based on P_{gal}. The Liverpool telescope imaged the highest ranked galaxy for each GW trigger, and the Zadko telescope imaged the top five galaxies for each trigger.

Chapter 8

Monte Carlo Simulation

A set of simulated gravitational wave signals was generated in order to test the ability of our software to localize sources. The simulation included ad hoc burst waveforms used to test cWB and Omega, and inspiral signals used to test MBTA. After event reconstruction by the pipelines, LUMIN selected tiles from each skymap using QUEST and TAROT sized fields. This chapter describes the simulation set, and presents the main results of the study.

8.1 Burst Simulation

8.1.1 Data Set

An ensemble of simulated GW signals was generated to measure the effectiveness of the reconstruction and follow-up procedures. For the Omega and cWB burst pipelines, these "software injections" were a mix of ad hoc sine-Gaussian, Gaussian, and white noise burst waveforms similar in type and distribution to those used in previous LIGO/Virgo all-sky analyses [8, 3]. While these waveforms are not based on specific astrophysical models, they do a good job of characterizing detector response for signals in specific frequency ranges (sine-Gaussians) and for broadband

signals (white-noise bursts).

Each sine-Gaussian waveform is defined by a central frequency, f, and a dimensionless quality factor, Q, and may be expressed in the time domain as

$$h(t) = h_0 \exp\left[\frac{-4\pi^2 f^2}{2Q^2} t^2\right] \sin\left[2\pi f t\right] \tag{8.1}$$

where h_0 is a amplitude factor. Notice this is similar, but not identical, to the definition of a sine-Gaussian used by Omega Pipeline, shown in Equation 5.52.

The testing set includes sine-Gaussian waveforms with $Q = 9$ and f set to 70, 100, 153, 235, 361, 554, 849, 1053, 1304, 1615, and 2000 Hz, and waveforms with $Q = 100$ and $Q = 3$ with central frequencies of 70, 235, 849, and 1615 Hz.

The Gaussian waveforms may be expressed as

$$h(t) = h_0 \exp\left(\frac{-t^2}{\tau^2}\right) \tag{8.2}$$

with τ chosen to be 0.1, 1.0, 2.5 or 4.0 ms.

The white noise bursts (WNB) were constructed by band passing white noise. Each WNB is generated with a filter band that starts at 100 Hz, 250 Hz, or 1000 Hz, and with bandwidth 10 Hz, 100 Hz, or 1000 Hz. In addition, the WNB are time windowed with a Gaussian envelope. Each envelope has a width of 100 ms, 10 ms, or 1 ms.

Each injection is normalized so that, from a distance of 15 Mpc, it would have an $h_{rss} = 2.5 \times 10^{-21}$. To study signals of a wide range of amplitudes, the entire injection set is multiplied by each of twelve scale factors: 0.0375, 0.075, 0.15, 0.3, 0.6, 1.2, 2.4, 4.8, 9.6, 19.2, 38.4, 76.8.

To emulate a realistic spatial distribution, each injection was calculated with a source distance and direction inside a randomly selected galaxy from the GWGC and the simulated GW amplitudes were weighted to be inversely proportional to distance. Only galaxies within 50 Mpc were included in the simulation, with weighting factors applied so that the probability of originating from each galaxy was proportional to its blue light luminosity. The simulation set and the analysis used the same catalog, so the results presented in this paper assume the GWGC is reasonably complete. Signals were superimposed on real LIGO-Virgo gravitational wave data taken between August and December 2009. This data set was processed by both cWB and Omega Pipeline, with the resulting triggers input to the LUMIN field selection algorithm. In addition, a second tiling selection software, known as the Gravitational Wave to Electromagnetic processor (GEM), ingested the skymaps. GEM was created as a cross-check for LUMIN, and to implement special case handling for the Swift observatory (see Section 9.7.2).

8.1.2 Results

Omega pipeline analyzed times in the simulated data of good data quality, and found 88,529 triggers associated with software injections. Likewise, cWB processed only times of good data quality, and also imposed a limit on the minimum detection statistic of triggers to retain ($\eta > 3.0$). This led to a set of 69,810 triggers at times of software injections.

To evaluate the uncertainty in the skymaps produced by cWB and Omega, we calculated the "search area" for each injection. This quantity is defined as the total angular area of the pixels with higher likelihood, L_i, than the pixel containing the true injection location. The sense of the statistic is that it represents the amount of sky area that would need to be searched before finding the true source, without aid from a galaxy catalog. The median search areas are shown in Figures 8.1 – 8.2 for a few waveforms.

The plots demonstrate the scale of the uncertainty in recovering Burst injections. For this data set, the threshold of detectability is around $h_{rss} \sim 10^{-21}$. We see that the uncertainty regions are typically between a few square degrees and a few tens of square degrees for signals near threshold. The Omega Pipeline position reconstruction algorithm assumes that the signal may be decomposed into a small number of sine-Gaussian basis functions. Accordingly, Omega Pipeline performs comparatively well on sine-Gaussian injections, but notably poorly on WNB. In fact, for some h_{rss} bins above $\sim 10^{-20}$, Omega pipeline failed to include the true source location anywhere in the 160 square degree skymap for more than half of the WNB injections, making the median search area an undefined quantity. For very loud signals, high frequency signals are localized more precisely than low frequency signals, consistent with the better time resolution of features in higher frequency signals.

Figures 8.3 - 8.6 show the success of the LUMIN tiling algorithm in recovering

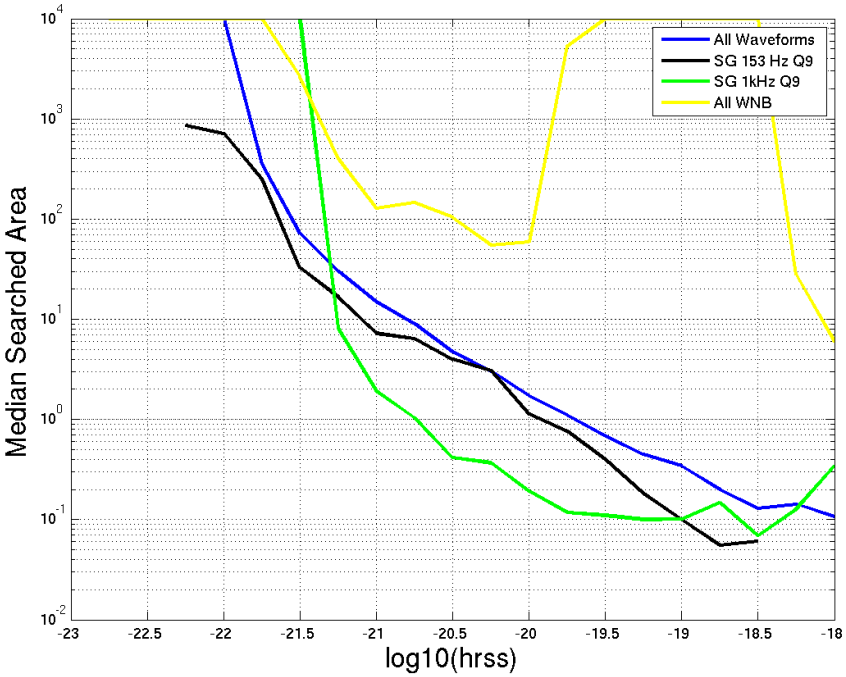

Figure 8.1: The median search area for a few waveforms in the simulation, using triggers from Omega Pipeline. The median search area for each class of injection is shown in units of square degrees.

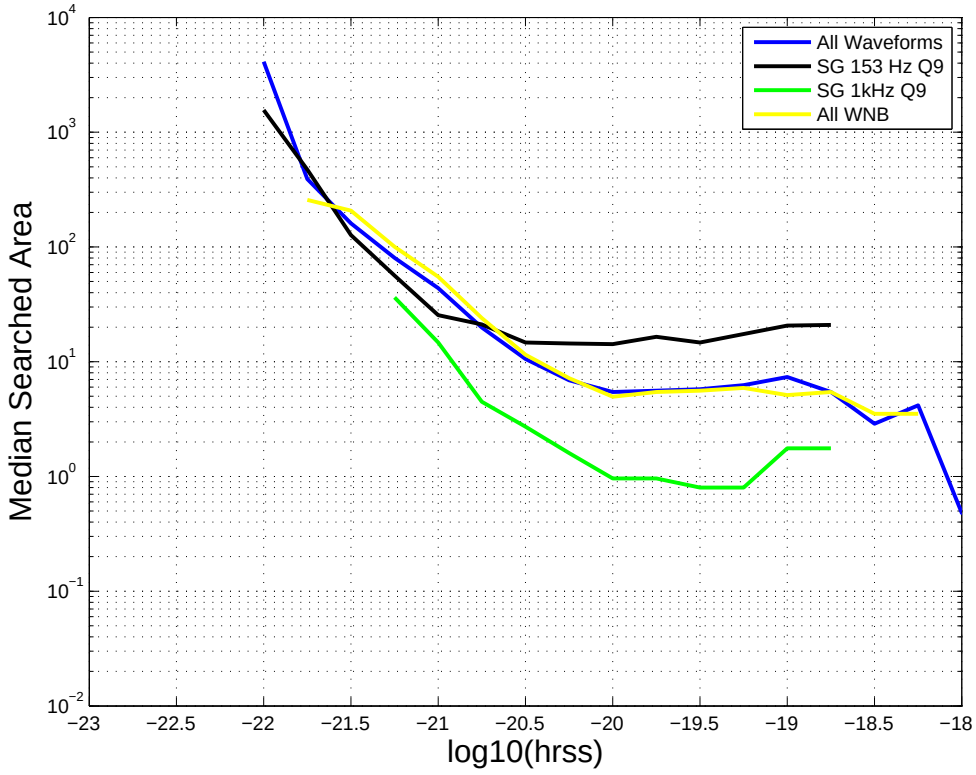

Figure 8.2: The median search area for various waveforms in the simulation, with triggers found by cWB. The median search area for each class of injection is shown in units of square degrees.

the correct sky position. Each injection's skymap was processed by LUMIN, assuming that a particular telescope was allowed to observe using one, two, or three fields. A fourth assumption, that the telescope was allowed to observe using fifty fields, was also tested, to allow LUMIN to cover the entire skymap. The figures show what fraction of injected source positions were imaged under the various assumptions. Results are shown both for cWB and Omega, and with a telescope with a 19 square degree FOV, representing the QUEST telescope, and a 3 square degree FOV, which is the size of both the ROTSE and TAROT FOV. The plots are shown as a function of the pipeline's native detection statistic, η or Ω. For reference, in the low-latency search, Omega events that triggered optical follow-up had $\Omega \sim 3$ or more, while for cWB, event candidates with $\eta > 3$ triggered optical follow-ups.

The curves show the fraction of event candidates in the simulation where the correct sky location was recovered by LUMIN with various values for the allowed number of tiles. Recall that, during the run, ROTSE and TAROT both allowed one field per event candidate, while QUEST accepted three. The extreme case, labeled as "Inf Shots", allows fifty tiles per event, typically enough to cover the entire skymap. The limitation that skymaps are no more than 160 square degrees means that this case does not cover the entire sky. Moreover, Omega and cWB only reported enough pixels so that the nominal likelihood, L, summed to 99%. This feature impacts performance, especially for loud signals, and probably accounts for the fact that both cWB and Omega approach only between an 80 and 90% success

rate with 50 tiles in the limit of high SNR signals.

Looking at the case of signals near the detection threshold, perhaps the case of most interest, we see that cWB recovers roughly 10% of signals with one tile, and around 20% if three tiles are used. The Omega numbers around $\Omega = 3$ are a bit stronger, recovering around 30% of signals with one tile, and 50% of signals with 3 tiles. This seemingly strong performance of Omega is likely influenced by the large fraction of sine-Gaussian waveforms in the simulation, which Omega reconstructs relatively well. Using the larger FOV on QUEST, compared with TAROT, is seen to make some improvement in the success rate. However, the impact is minor, a fact that highlights the success comes mainly from "guessing" the right galaxy, rather than from covering the entire skymap.

To further explore the effect of galaxy weighting, the cWB data was also processed by LUMIN without applying galaxy weighting. The result may be seen in Figure 8.7. With three TAROT tiles near the detection threshold, the success without galaxy targeting is seen to be around 10%, as compared to nearly 30% with galaxy targeting in Figure 8.4. Comparing these two plot shows that, with this data set, the galaxy weighting is reasonably effective in improving the success rate.

8.1.3 Calibration Errors

To explore the effects of calibration uncertainty, a second burst injection set was created and tested. These "miscalibrated" injections have intentional offsets

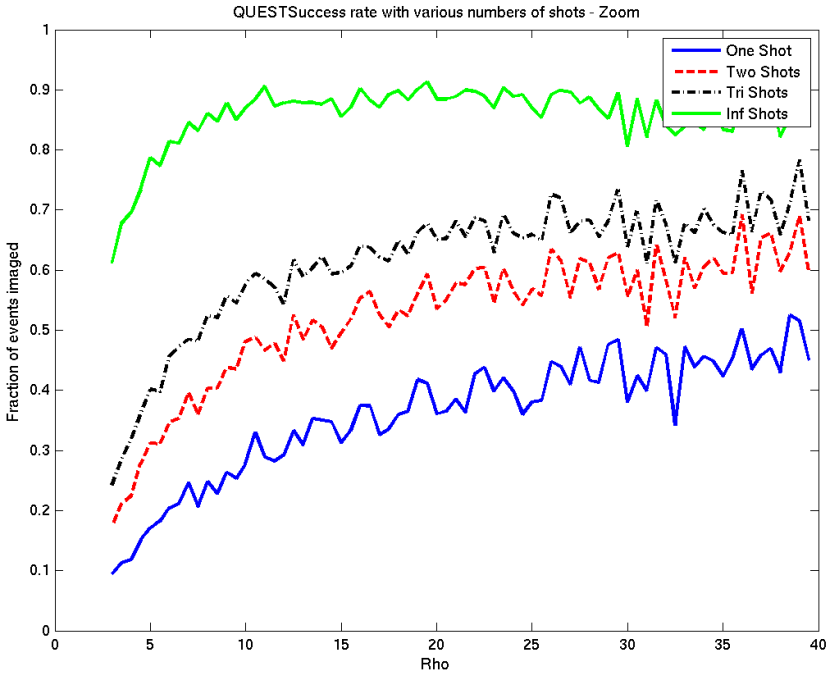

Figure 8.3: Success rate with cWB and the QUEST telescope.

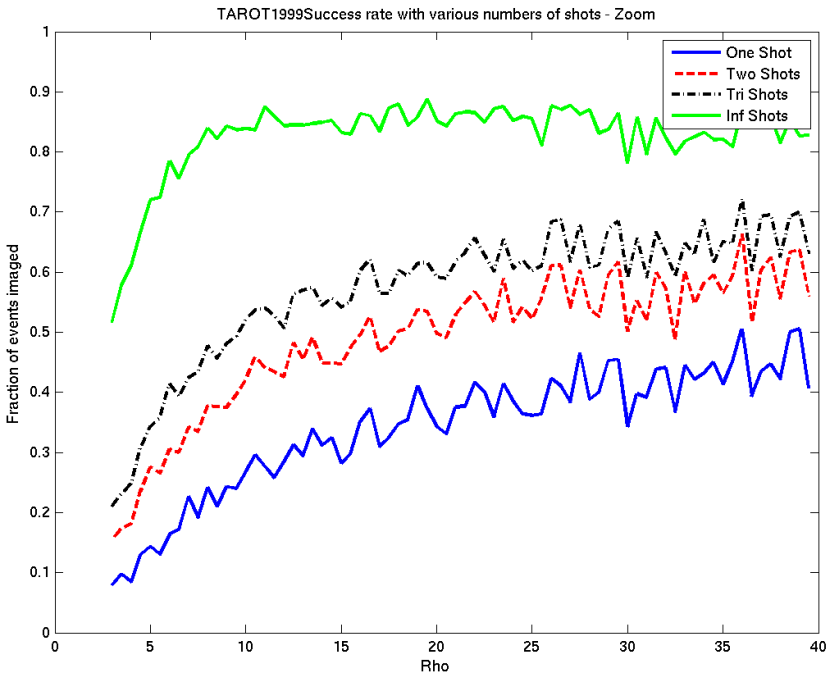

Figure 8.4: Success rate with cWB and the TAROT telescope.

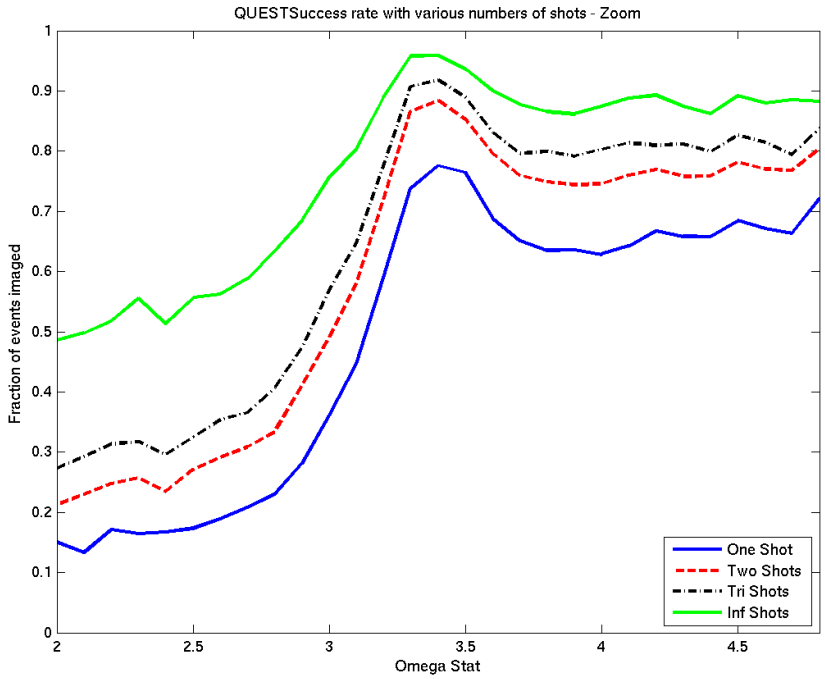

Figure 8.5: Success rate with Omega Pipeline and the QUEST telescope.

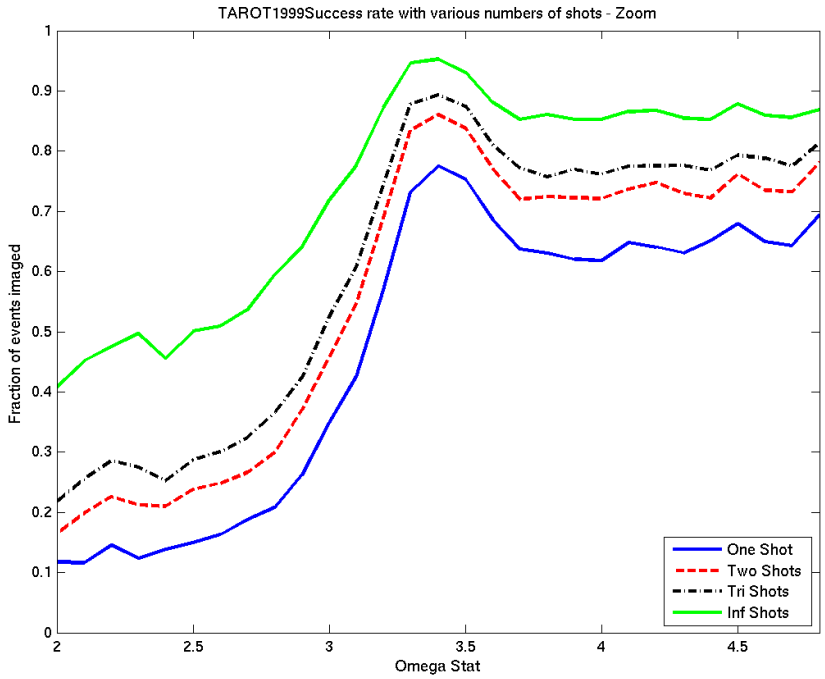

Figure 8.6: Success rate with Omega Pipeline and the TAROT telescope.

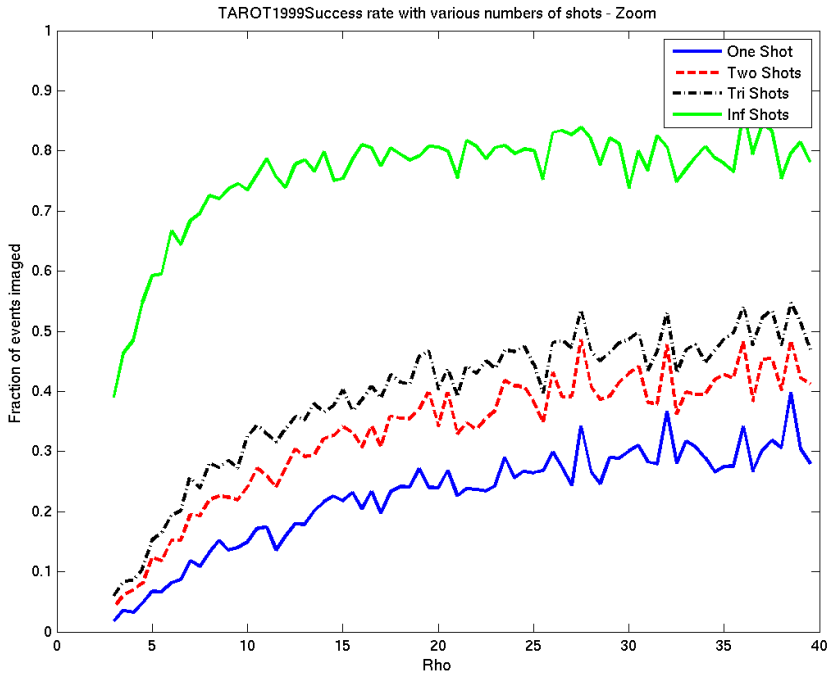

Figure 8.7: Success rate with cWB and a TAROT sized field. Here, no galaxy weighting was applied before tile selection.

added to the amplitude and central time of each injection in each detector. Each injection's amplitude is scaled in amplitude by a factor between 0.85 and 1.15, and has its central time shifted by between -150 μs and $+150$ μs. Both the amplitude factors and timing shifts are randomly selected from flat probability distributions. The limits of these errors were chosen to be consistent with preliminary expectations for the size of the calibration uncertainty in the S6/VSR2/VSR3 data set.

The simulated timing offsets are a simplification of the actual errors incurred due to calibration. They are intended to represent both errors in "phase", as well as "timing" offsets in the detector. True calibration "timing" errors are defined as uncertainties in the time domain difference between the time stamp on the data and the time a signal would have passed through the detector. "Phase" errors are defined as uncertainties in the frequency domain phase of the signal, arising from such places as uncertainty in calculating transfer functions at various points in the detector loop. The phase errors are, by their nature, frequency dependent, and are typically expected to be between 1 and 10 degrees. The limit of 150 μs is arrived at by taking the limit of phase uncertainty to be 8 degrees, and calculating its impact on a narrow-band signal around 150 Hz, near the most sensitive band for LIGO. The expected timing errors are smaller than this. So, the 150 μs uncertainty on timing is used to represent uncertainties from both timing and phase. The simulation data set that includes calibration errors was processed by cWB, LUMIN, and GEM.

The impact of calibration errors may be seen in Figure 8.8. The curves shown

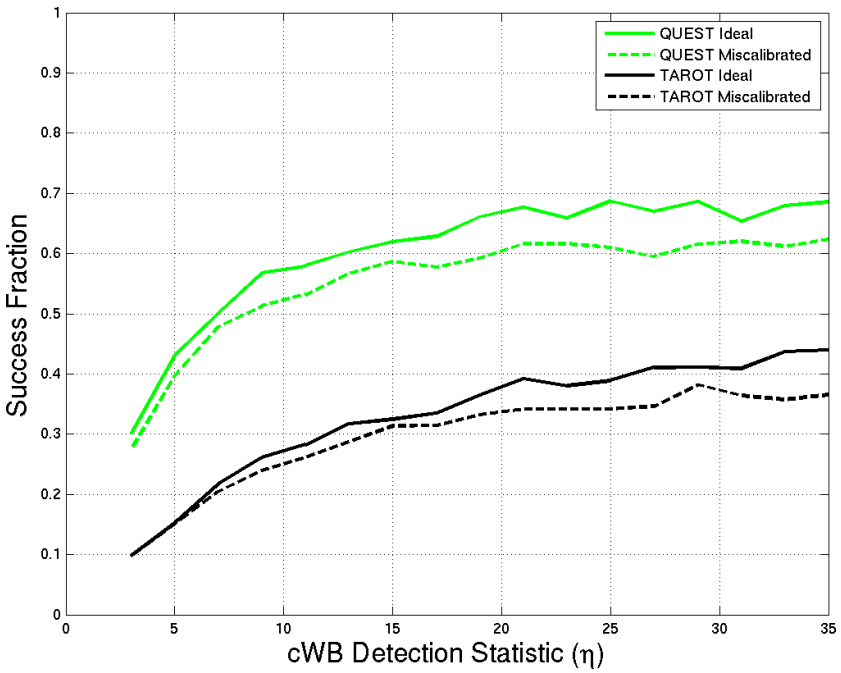

Figure 8.8: Success rate with cWB, using the QUEST and TAROT telescopes. The curves for QUEST assume three pointings for each event, and the curves for TAROT assume one pointing for each event, consistent with the actual observations. The dashed curves include the effects of simulated calibration errors, where the solid lines do not.

here may be taken as a summary of the results in this section. The plot assumes three fields for each QUEST event, and one field for each TAROT event, consistent with the observing schedule used in the 2009/2010 data collection periods. This difference in number of fields is the main reason the QUEST curves show much higher success rates than the TAROT curves.

For low SNR triggers, around $\eta = 5$, calibration errors are seen to have a relatively small impact. In this regime, the localization uncertainty that arises from difficulty identifying features of the signals in the noise is typically larger than the positional errors arising from calibration errors. At higher values of SNR, the effect of the calibration errors becomes more pronounced. This is expected, because errors arising from calibration error should not improve with higher SNR, whereas other errors should. In [26], a bound is derived that calibration error will not dominate if the maximum calibration error in radians, δp, is less than one over the single detector SNR. Since η is roughly the network SNR, we may approximate this limit as

$$\delta p \leq 1/(\eta/\sqrt{3}) \tag{8.3}$$

Then, since the maximum calibration uncertainty is 8 degrees, we expect that calibration errors could become significant for $\eta \geq 12$, which seems consistent with the shown results.

The sensitivity of Burst searches is typically measured using efficiency curves, shown for this simulation in Figure 8.9. Each curve includes all waveforms in the

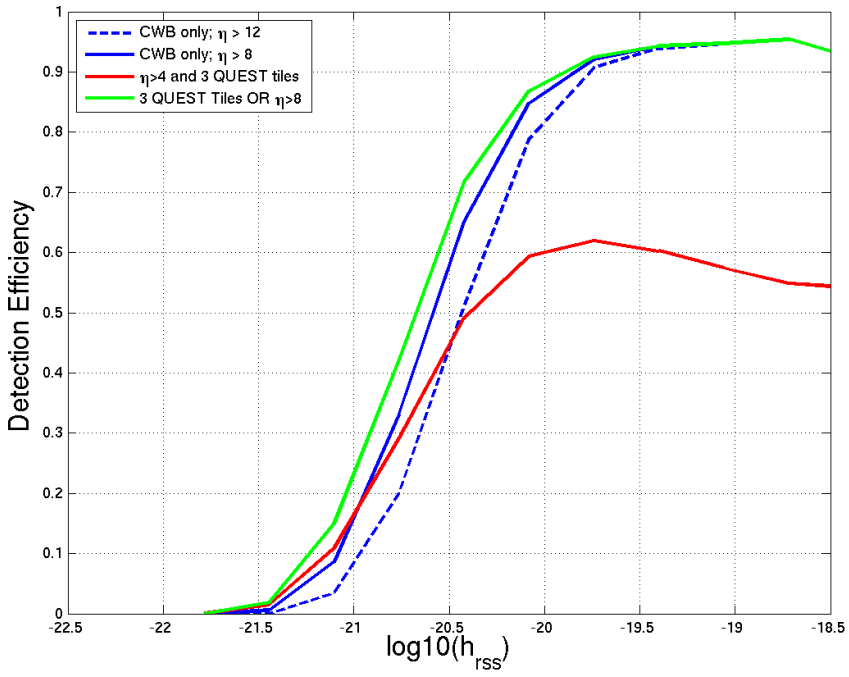

Figure 8.9: Efficiency curves for the burst simulation. The horizontal axis shows the strength of the signal when arriving at earth. The blue curves show efficiencies using nominal η thresholds that are similar to values used in previous Burst analyses. The red curve shows the efficiency with a low value of η, but also demands that the right sky location is recovered with three QUEST tiles. Finally, the green curve shows the efficiency if signals are recovered with $\eta > 8$ OR with $\eta > 4$ and a correctly recovered sky location.

simulation, and shows the fraction of injections recovered, with various assumptions about what it means to find a signal.

The blue curves are intended to represent an all-sky search without the aid of EM follow-up observations, similar to [3]. In past searches, various η thresholds have been used, depending on the details of the data set. In [3], thresholds of η used to set upper-limits range between 4 and 9. However, these values have been seen to be too low to declare a confident detection without validation from an EM signal. For example, a simulated candidate event with $\eta > 6$ was seen to be too weak to call a confident detection. Moreover, for louder signals, around $\eta > 6$, the background for cWB becomes relatively flat as a function of η in some cases, so it may take very large values of η to make a confident burst detection in the absence of corroborating evidence. For this reason, the curves representing the all-sky burst search are shown with η thresholds of 8 and 12, and are shown in blue in the figure.

These curves may compared with the red curve, which shows a nominal efficiency curve for a search that includes EM follow-up observations. This curve makes several assumptions which may not be realistic. In particular, it assumes that any burst signal will be accompanied by an easily recognizable EM transient. That is, we assume that if a GW signal is loud enough to trigger an optical follow-up, and that the correct sky location is recovered, then the EM signature will give enough confidence in the astrophysical nature of the signal to label it as a detection. In addition, effects that would prohibit an observation, such as bad weather or daylight,

are not included in the simulation.

With these caveats, the red curve shows the potential of EM follow-up observations to recover detections that would otherwise be difficult to distinguish from noise. Comparing with both all-sky curves, the efficiency curve that includes EM follow-ups is seen to be more sensitive to low amplitude signals, but less successful with high amplitude signals, due to the fact that some signals do not have their position correctly recovered, even at high amplitudes (See Figure 8.8).

Of course, searching for GW signals with an EM counterpart will never be a substitute for a GW only search. For this reason, the green curve is also presented, which shows the efficiency with a logical OR of the red and solid blue curves. For the green curve, an injection is considered recovered if it has an η value greater than eight, or it has $\eta > 4$ and the position recovered with three QUEST tiles.

By assuming a uniform source distribution, with each source having the same amount of GW output energy, we can compare the detection rates associated with each curve through the relation

$$\text{Detection Rate} \propto \int \epsilon(h_{rss}) h_{rss}^{-4} dh_{rss} \tag{8.4}$$

where $\epsilon(h_{rss})$ is the efficiency of the search. Using this metric, we find that the combination of a search with EM follow-ups and an all-sky search with $\eta = 8$ has 2.1 times the detection rate of the all-sky search alone. If we assume that a GW search alone requires an η threshold of 12, then we find that including the search for EM counterparts increases the detection rate by more than a factor of 5.

The modest improvement in amplitude sensitivity from optical follow-ups leads to significant increases in the volume of space that may be searched.

8.2 CBC Simulation

8.2.1 Data Set

Since MBTA is designed to detect different signals than the Burst pipelines, a separate set of software injections was needed for testing. The pipeline was tested using a week of real data added to 10,122 simulated inspiral signals, drawn from a population of binary neutron stars and neutron star-black hole binaries. Like the Burst simulation, each injection is placed in a galaxy from the GWGC, with a probability determined by the blue light luminosity of the galaxy. Galaxy selection was limited to a distance of 40 Mpc. In order to emphasize systems that may contain at least one neutron star, component masses of the binaries are drawn from a range between 1 and 15 $M_{\odot}$, with the total mass constrained to be between 2 and 20 $M_{\odot}$. MBTA recovered 736 of the injections in triple coincidence. For each of the recovered injections, the MBTA sky localization procedure determined a skymap, which was then used by LUMIN to select fields for a 1.85 degree FOV telescope. The galaxy weighting is applied by MBTA, not by LUMIN, and uses a slightly different method than the galaxy weighting for burst triggers, as described in Section 7.3.

8.2.2 Results

The results of the simulation are shown in Figures 8.10 and 8.11. The search area statistic is defined in the same way as for the Burst simulation, and shows the scale of the MBTA error box. Unlike the Burst searches, the MBTA localization procedure is based on differences in timing in single detectors, and does not use any coherence. For this reason, the MBTA error regions are always contiguous, rather than having multiple patches as is possible with the burst pipelines. So, the median search area for MBTA may be more correctly be interpreted as the size of a single region that would need to be searched to find the signal, without the use of galaxy targeting.

To evaluate the success rate corresponding to our search, the MBTA galaxy weighting was applied, and then TAROT sized tiles were selected by LUMIN. The result may be seen in in Figure 8.11. At the lowest shown bin centered at $\rho_c = 12$, the success rate is seen to be around 20% with a single shot, which is similar, but slightly better, than the success rate near the Burst detection threshold. For reference, during the 2009/2010 search, the single MBTA trigger to receive follow-up observations had $\rho_c = 10$, and so would be at the low end of this bin.

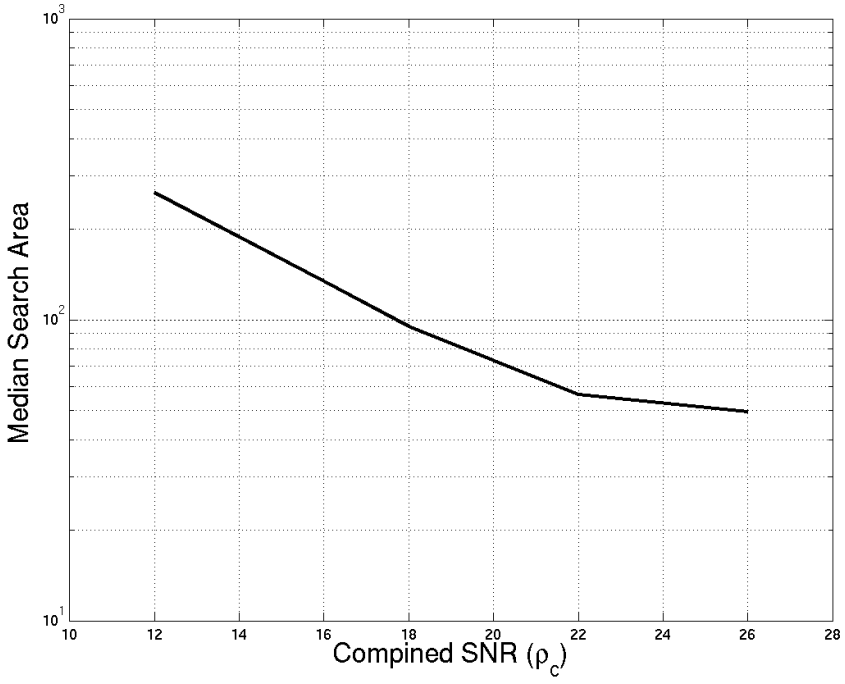

Figure 8.10: Search area with the MBTA Pipeline.

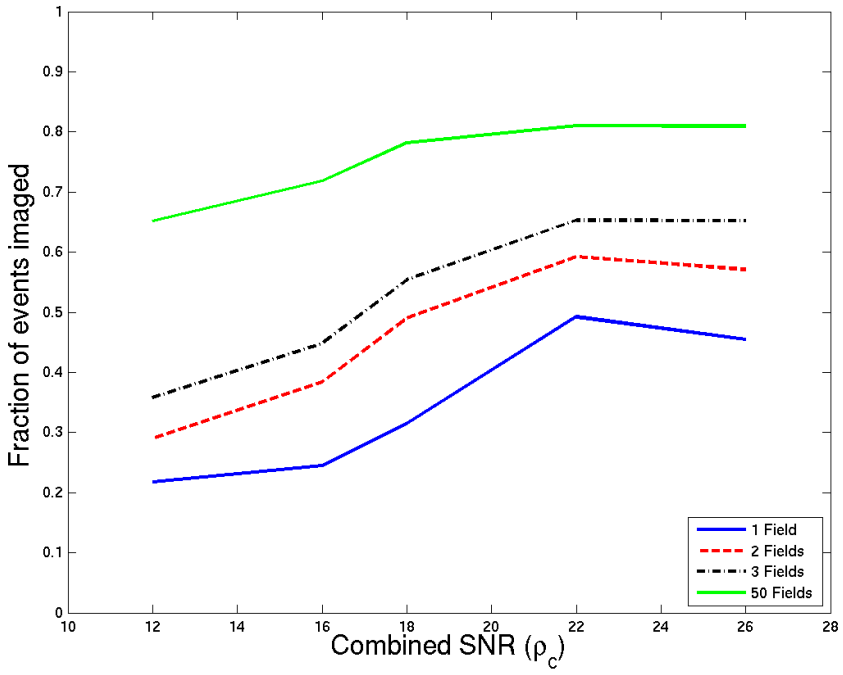

Figure 8.11: Success in localizing signals with the MBTA pipeline and a TAROT size FOV.

Chapter 9

Low Latency Search

Finding EM counterparts to GW transients requires performing a complete analysis before the EM counterpart has faded from view. This chapter describes the low-latency analysis that was implemented for the S6/VSR2/VSR3 search.

A graphical overview of the pipeline is presented in Figure 9.1. First, data and data quality information were copied from the three interferometer sites to central computing centers. This data was used by the Omega Pipeline, cWB, and MBTA to produce triggers, skymaps, and background estimates.

For each ETG, triggers with false alarm rates (FARs) more significant than roughly 1 trigger per hour were placed in the trigger archive, GraCEDb. GraCEDb then used the communication tool LVAlert to signal the presence of the new trigger to LUMIN and GEM.

LUMIN then took a number of steps to further process the trigger, including tile selection, checks on data quality, and application of a more stringent threshold on the FAR of the trigger. For triggers passing the the FAR cut, LUMIN then alerted an on-call team. The on-call team used the LUMIN web interface to manually validate the trigger. If validated, the team issued a command back into LUMIN, which then executed a variety of scripts to pass the selected fields to the robotic telescopes. This

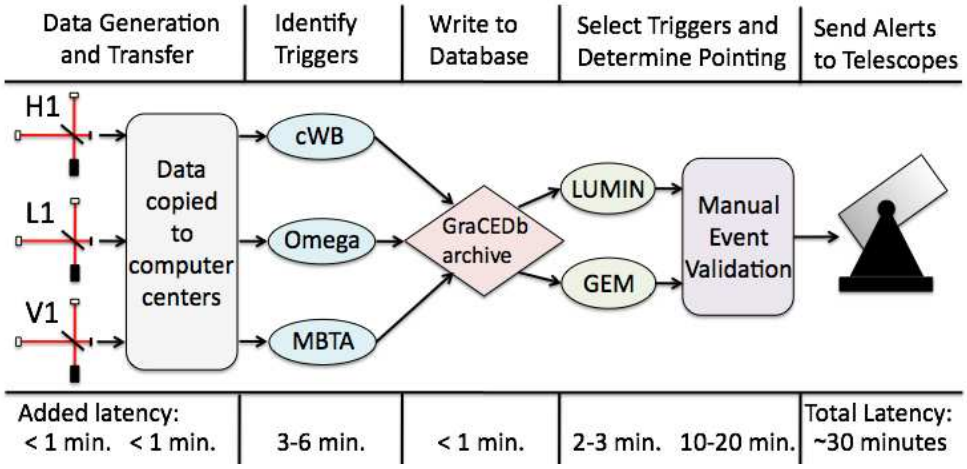

Figure 9.1: An overview of the low-latency pipeline. The data from three interferometer sites were searched for both burst and CBC signals. The GraCEDB archive stored significant triggers, and alerted the EM follow-up processors. LUMIN provided an interface for the manual validation, and tools to communicate with robotic telescopes.

entire chain of events typically executed in a time between thirty and fifty minutes, with the majority of the time spent in discussion among the on-call team. Details of the trigger selection process, manual validation, telescope alerts, and the LUMIN web interface are described in this chapter.

9.1 Data Quality

A number of common occurrences may make a stretch of interferometer data unsuitable for sensitive GW searches. Examples include times of large seismic disturbance, non-standard interferometer configurations, and temporary saturations of various photodiodes in the interferometer sensing and control system. To mark such

times, monitor programs analyze auxiliary data to produce lists of abnormal time segments with low latency.

Each monitored condition is known as a data quality flag. Typically, a GW search may be "tuned" by choosing to apply data quality flags that reject times that are found to have excesses of spurious triggers. A novel feature of the low-latency analysis was that data quality flags to be applied were chosen *before* the data collection period. Flags were chosen based on their studied performance on data collected before the low-latency search period.

Data quality flags were divided into categories based on their severity. The most severe, Category 1 flags, were times of unusable data. For LIGO, this means an interferometer was out of science mode, times within ten seconds of a loss of optical resonance in the Fabry-Perot cavities (lock loss), and times of poor calibration. Data marked with Category 1 flags were never analyzed. Category 2 flags were times when a well-understood problem was present. These included saturated photodiodes and saturated currents in the test mass actuators, and times of high seismic noise. Times marked with these flags were analyzed by the ETGs, but any triggers found in these times were automatically rejected. Possible data quality issues that were less well-understood were marked as Category 3. In the low-latency analysis, Category 3 flags were ignored by the automated selection of triggers, but were noted in the information delivered to the on-call team that validated triggers.

9.2 Trigger Selection

To search for EM counterparts to gravitational wave candidate events, one would like to select only the most promising triggers for follow-up observation. Telescope time is a limited resource, and so needs to be treated with proper care. However, in a low-latency analysis, the selection process must happen before all the triggers have been collected, so simply ranking the triggers is not sufficient; some criteria must be made to decide when a trigger is "good enough" to merit follow-up observation. Past gravitational wave searches, performed after the completion of the data collection periods, have not encountered this difficulty. As an added complication, the low-latency analysis was forced to choose among triggers from three different pipelines, which was also a novel feature of this search.

The main tool for addressing these challenges was low-latency background estimation. Each ETG, in addition to producing so-called "zero-lag" triggers based on the actual detector data, also produced "background" triggers based on data with artificial time-shifts between detector sites (see, for example, [3]). Since a true GW signal would not appear in coincidence in multiple detectors in time-shifted data, the vast majority of background triggers are thought to arise from detector noise. The information from the time-shifted analysis was used to estimate the rate of background triggers as a function of detection statistic, and so revealed the significance of a given zero-lag trigger.

Let ρ represent the detection statistic of a generic ETG. The false alarm rate,

$FAR(\rho_0)$ for a trigger with $\rho = \rho_0$ is defined as the average rate of background triggers with $\rho > \rho_0$. To calculate the FAR of a trigger with detection statistic ρ_0, an algorithm counts the total number of time-shifted triggers with $\rho > \rho_0$, $N_{\rho>\rho_0,B}$, and then divides by the total livetime represented by the time-shifted data, T_B:

$$FAR(\rho_0) = \frac{N_{\rho>\rho_0,B}}{T_B} \tag{9.1}$$

The FAR is a useful measure of a trigger's significance for two reasons. The goal of our trigger selection process was to choose triggers that were more likely to represent true gravitational wave signals. The FAR shows how likely a trigger is to have arisen from noise, in a pipeline independent manner. So, choosing triggers with low FARs allows us to compare triggers from different pipelines, and choose those that are more likely to represent true gravitational wave signals. The second reason is important from a strictly practical standpoint. Each partner observatory had requested a limited rate of triggers, so we had to pass along only event candidates which occurred more rarely than these target rates. With the sensitivity of initial LIGO and Virgo, we expected the large majority of triggers to be caused by noise events. This means that, in budgeting how often we made observations, a good estimate was to assume that all the triggers would be due to noise. With this assumption, to stay within the desired observing rate, a threshold should be applied that limits EM follow-ups to triggers with a FAR lower than the target rate for sending triggers to observatories.

Poisson statistics may be used to show the relationship between a triggers FAR

and its statistical significance. Once the FAR is known, the livetime of the actual analysis, T_A, may be used to estimate the number of triggers with $\rho > \rho_0$ that are expected in the zero-lag analysis due to noise alone:

$$\langle N_{\rho>\rho_0,A} \rangle = FAR(\rho_0) \times T_A \equiv \mu \tag{9.2}$$

where $\langle N_{\rho>\rho_0,A} \rangle$ is the expectation value of the number of triggers with $\rho > \rho_0$ that occur from noise over the livetime of the analysis. If we model the occurrence of noise trigger with $\rho > \rho_0$ as a Poisson process with rate $FAR(\rho_0)$, then the probability of finding exactly one event louder than ρ_0 in the zero-lag is

$$Pr[N_{\rho>\rho_0,A} = 1] = \mu e^{-\mu} \tag{9.3}$$

A trigger found in the zero-lag analysis should only be considered a possible GW candidate if we expect much less than one such trigger from background over the analysis time. In the limit $\mu \ll 1$, we find

$$Pr[N_{\rho>\rho_0,A} = 1] \approx \mu \tag{9.4}$$

This means that for interesting triggers, μ is the probability that an event with $\rho > \rho_0$ would occur due to noise alone.

In practice, T_A won't be known at the time of trigger generation, since this represents the total livetime of the analysis. However, we do expect T_A to be nearly the same value for all ETG's, since this is determined primarily by the state of the instruments. So, Equation 9.4 tells us that, to choose the trigger that is the most

likely to be a true GW signal, we need only choose the trigger with the lowest FAR. With this criteria, we can compare triggers from multiple ETG's, and choose the one most likely to be a true gravitational wave.

The conclusion, then, is that to both manage our telescope time, and choose triggers from among multiple pipelines, a reasonable strategy is to place a threshold on the FAR of triggers near the rate that the observing partners wish to receive them, and pass along triggers with a FAR lower than this threshold.

When working with data from real interferometers, the function $FAR(\rho_0)$ typically varies over time, due to evolving conditions within the instruments. A challenge for the low-latency analysis was measuring the FAR using data that is close enough to the trigger time to reflect the current state of the instruments. To address this challenge, the low-latency burst pipelines ran time-shift analyses in parallel with the unshifted, or zero-lag, analysis. Omega pipeline ran one hundred time-shifts, and cWB ran 4000 time shifts. The background analysis ran separately from the zero-lag analysis, and with higher latency. While zero-lag triggers were typically available in a few minutes, the time-shifted analysis typically ran with lag times between ten minutes and a few hours. This difference in lag time means that it was not possible to evaluate the FAR of a trigger using immediately adjacent data.

The low-latency analysis attempted a compromise between using background data as close to the zero-lag trigger time as possible, and using large statistical

Table 9.1: The FAR thresholds used by LUMIN to select event candidates for potential optical follow-up. Triggers which passed these cuts, and passed data quality conditions, caused alerts to the on-call team. The column "Zero-lag data" tells roughly the amount of data that was used in the time-shift analysis.

Run	ETG	FAR Thresholds	Zero-lag data
Winter	Omega	FAR $\leq$ 1 event per day	24 hours
		FAR $\leq$ 2 events per day	30 minutes
Winter	cWB	FAR $\leq$ 1 event per day	1 hour
		FAR $\leq$ 2 events per day	10 minutes
Autumn	Omega	FAR $\leq$ 0.25 events per day	24 hours
		FAR $\leq$ 2 events per day	30 minutes
Autumn	cWB	FAR $\leq$ 0.25 events per day	1 hour
		FAR $\leq$ 2 events per day	10 minutes

samples, by placing thresholds on the FAR of burst triggers based on both a shorter and longer span of data. If the noise statistical properties remained constant, then the measured FAR of a trigger should be nearly the same as evaluated in both spans of data. The cuts used in the analysis are shown in Table 9.1. In the table, the listed amount of time refers to the approximate amount of triple coincident, zero-lag livetime that was used (with time-shifts) to calculate the FAR. The same set of online data quality flags that were used to reject triggers in the zero-lag analysis were applied to the time-shifted analysis. This caused the actual amount of data used for background estimation to be somewhat smaller than the nominal values listed in the table, but calculated FAR were properly normalized with the actual amount of time-shifted livetime. Each time-slide is effectively a unique data set, so the FAR measurement typically drew on at least $\sim$ 1000 hours of artificially shifted data.

The target rate of triggers to receive follow-up observation was implemented as a threshold on the FAR as measured in the longer stretch of data. The FAR as measured in a shorter stretch of data used a looser threshold in all cases. This looser cut was used as a sanity check to reject triggers that occur in stretches of data with a rate of noise events that is much higher than in the surrounding time.

During the summer observation period, LUMIN accepted triggers from the MBTA pipeline, in addition to the burst pipelines. MBTA evaluated the FAR of each trigger without time shifts, by calculating the expected rate of coincident triggers from the rate of single trigger in each detector. The first step in the MBTA FAR estimation was collecting a list of the properties of the 200 ($N_{triggers}$) triggers in each detector that precedes the trigger of interest. These $N_{triggers}$ span times T_H, T_L, and T_V in the three detectors. Then, the single IFO trigger rate in each detector is

$$R_\alpha = \frac{N_{triggers}}{T_\alpha} \tag{9.5}$$

where α represents Hanford, Livingston, or Virgo. For a given trigger in IFO α, the probability that a single trigger from a second IFO β will be coincident within a window of time w is $R_j \times w$. This follows from equations 9.2 and 9.4. Then, the rate of coincident triggers is

$$R_{i+j} = R_i R_j w \tag{9.6}$$

For a real IFO network, the relevant coincidence window is twice the light travel time between the two detectors in question. Extending this to triple coinci-

dence, we find the rate for triple coincident triggers may be estimated as

$$R_{H+L+V} = \frac{N^3_{triggers}(2dt_{HL}2dt_{HV})}{T_H T_L T_V} \tag{9.7}$$

where dt_{HL} and dt_{HV} are the light travel times between Hanford and Livingston, and between Hanford and Virgo. This is the rate of total triple coincident triggers, however, it does not tell us the FAR above some threshold. To estimate what fraction of all triple coincident triggers will pass threshold, all possible combinations of one trigger from each IFO are formed. A cut on consistency of the estimated mass, that is also used in the zero-lag search, is applied to this set of $N^3_{triggers}$ events, and then events are cut with $\rho < \rho_0$. The fraction of events that pass the mass cut and are louder than the target event is then $\frac{N_{loud}}{N^3_{triggers}}$, where N_{loud} is the total number of triggers remaining after both the cuts on mass and ρ. Finally, we write the FAR of a given trigger as

$$FAR_{\rho_0} = \frac{N_{loud}(4dt_{HL}dt_{HV})}{T_H T_L T_V} \tag{9.8}$$

This expression is computationally cheaper to evaluate than the time-shift method used by the burst searches, and is calculated using the few minutes of data immediately preceding each zero-lag event.

9.3 Processing by LUMIN

When LUMIN received a new trigger from GraCEDb, via an LVAlert connection, it immediately performed a variety of tasks. A directory was created for the

trigger, and all LUMIN outputs associated with the trigger were stored in this event directory.

The information that LUMIN received from LVAlert was limited, but included a file or URL location where more detailed information was available. For each trigger, LUMIN read in the trigger properties, including the FAR and skymap information, which it found in the indicated directory. A different file reader was used for each ETG, but the properties were then stored in a common data structure. After this initial input of trigger information, LUMIN treated the three ETGs in essentially the same manner. The data structure that LUMIN used to store the trigger properties is known as the event array.

After reading in the basic trigger properties, LUMIN would query for any active data quality flags at the time of the trigger. The data quality information was stored in the event array, and also written to an output file in the event directory. If the data quality flags revealed the trigger to be associated with a hardware injection, then LUMIN checked the time against a locally stored list of scheduled burst hardware injections, and stored the injection location in the event array. This step allowed the generated skymaps to be easily compared with the intended injection location as an end-to-end check of the position reconstruction.

At this stage, LUMIN applied a threshold to determine if the event was significant enough to alert the on-call team, and potentially pass along for observation. To pass the threshold, the event must pass the data quality conditions, and pass

the FAR thresholds described in Section 9.2. If all of these conditions were satisfied, LUMIN began the process of sending out alerts to the on-call team, and marked the event as having an "alert" status in the event array.

LUMIN then launched Matlab jobs to calculate the optimal field selection for the trigger (See next section). A single Matlab job was launched for each trigger, with instructions to repeat the tiling calculation for each telescope in the network. The Matlab jobs took several minutes to complete, but ran in parallel to LUMIN's other operations, and so made no measurable contribution to the search latency. After the tiling code completed, LUMIN launched a second Matlab script to calculate the times when each selected tile was viewable by the associated telescope.

Finally, the entire event array was written to a file in the event directory. Information from the event directory would appear on a web page, the LUMIN "Events Page" shortly after the initial processing of the trigger (see Section 9.5).

9.4 Tile Selection

For each trigger that LUMIN received, it launched a Matlab job to choose the best tiles for each telescope, using algorithms described in Chapter 7. The choice of best tile was performed for each telescope in the network, with the idea that different sized FOVs may lead to different choices for the optimal tiling choice. There was some redundant calculation at this stage, since some of the telescopes, such as the four ROTSE III telescopes, have the same FOV. This did not present

any difficulties, though. While the Matlab jobs usually took several minutes to run, they ran in parallel to the other processes, and so add little or no latency to the pipeline.

The Matlab tiling procedure returned several outputs, which were all placed in the event directory, and were available to the human monitors through the LUMIN Events Page (see Section 9.5).

The "tile file" was labeled SCOPENAME_tiling.txt, and contained the RA, DEC, and P statistic of each selected tile. During the process of weighting the skymap with the galaxy catalog, the P statistic was normalized to unity for the entire skymap. Since the P statistic is the product of all the known likelihoods for picking the correct source position, the normalized P statistic is an attempt to calculate the probability that a given tile contains the right sky position, assuming that the signal is actually astrophysical in origin. So, if the tile file listed the P statistic as 0.50 for a particular choice of field, we might say there was a 50% chance that field contained the true source position. In fact, the P statistic involves too many dubious assumptions to be trusted as a true probability. In practice, though, it is a reasonable measure of how successful the tiling algorithm was in covering the skymap.

The "skymap plot" was labeled SCOPENAME_tiling.png, and was a PNG image that contained a plot of the skymap, selected tiles, and associated galaxy positions from the GWGC. This provided a visual sense of how spread out on the

sky the skymap was, where the key galaxies were located, and how the skymap compared with the FOV size.

A file labeled SCOPENAME_tiling_galaxies.txt listed the PGC number, name, RA, DEC, and M_{gal} for each galaxy that was associated with some non-zero pixel in the skymap. In addition, this file listed a *tile number* for each galaxy, indicating which, if any, of the selected tiles included the center position of this galaxy. The special handling for small FOV telescopes (See Chapter 7) used this file as an input, relying on the regular LUMIN tile selection code to create this relatively short list of galaxies that have positions consistent with the GW skymap.

After all of these expected output files were present in the event directory, LUMIN launched a Tcl module known as "WriteObserveFile", that calculated when each selected target location may be viewable. An object was assumed overhead during times when it was within 70 degrees of zenith. To determine the hours of night time, LUMIN assumed that the sky would be sufficiently dark when the sun was more than 12 degrees below the horizon, or 102 degrees from zenith of the telescope. The two time intervals, corresponding to hours of darkness and hours when the object was overhead for the telescope, were merged to calculate an interval representing the times when the object was overhead during night hours. It is possible that taking the union of these two segments results in two separate segments, in which case both time intervals were reported. The result was written to the "observation file" as a list of each tile selected for each telescope, and a list

of the time intervals, if any, when each tile's location on the sky was observable.

9.5 Web Pages and Alerts

A key feature of LUMIN was the web-based user interface. The requirement that each trigger be quickly, manually validated resulted in the creation of a rather large team to take shifts "on-call" to respond to triggers. These team members had to be trained to quickly get access to key trigger information, and to use available tools to document decisions and alert robotic telescopes. The size of this team, as large as twenty people, suggested that LUMIN needed to be relatively easy to access. Moreover, for each trigger, the on-call monitor needed to discuss the relevant information with the current science monitor (SciMon) at each site. Working as a SciMon is a condition of membership in the LIGO Scientific Collaboration. At a given moment, virtually any member of the 800+ LIGO Scientific Collaboration or Virgo Collaboration might be on shift. The people who happened to be in the role of SciMon for each trigger needed to access the trigger properties as well, demanding that LUMIN be accessible with essentially no training beyond brief, written instructions.

For most users, the point of entry for LUMIN was the LUMIN Events Page (see Fig. 9.2). During online running, the events page was updated roughly every minute. It contained a list of every trigger ingested by LUMIN within the past 24 hours, in a table format, with each line representing one trigger. The lines of the table

were color coded, so that significant triggers would be immediately recognizable. A line of the table with a green background indicated a trigger that had passed all of the automated cuts, including the FAR cut. These were the event candidates, and required validation by the on-call team. Triggers that were marked in purple occurred in times when a hardware injection was scheduled to enter the detectors. In Fig. 9.2, Omega Pipeline and both running versions of cWB all report triggers for the hardware injection that occurred around GPS time 968655006. Similarly, triggers which failed an automated data quality test were marked in red. For triggers which met none of the above criteria, the background color defaulted to black.

The table on the Events Page contained most of the basic properties of each trigger, with links to more detailed information.

Age The first column in the table showed how much time had elapsed since the trigger timestamp, in hours and minutes. The on-call team monitored this number as they discussed the event, and attempted to pass or veto events before the clock reached thirty minutes.

ID The ID number listed for each trigger was assigned by GraCEDb. This allowed a consistent numbering system across several pipelines. The GraCEDb ID was assigned sequentially in the order triggers are received by the archive. However, not all triggers in GraCEDb were reported to LUMIN, because GraCEDb also accepted triggers from at least one pipeline that was not included in the EM follow-up program.

GPS The GPS time of the event was listed.

DQ The DQ column displayed if the trigger failed any of the automated data quality checks. The information displayed on the page may reveal that the trigger occurred during a hardware injection, failed a category two or category three data quality test, or that the trigger was "clear" of any of these concerns.

Energy The energy column displayed the detection statistic for each trigger (see Chapter 5).

Event Rate This column displayed the FAR for the trigger, as estimated by the reporting pipeline. For the burst trigger generators, this was the FAR as estimated in the longer span of background data. See Section 9.2 for details.

Frequency The frequency column, for burst events, displayed the estimated central frequency for the GW signal.

Status The status column displayed "processed" for most triggers, or "alert" for triggers that were identified as event candidates. The status column status could also display "overload", indicating that LUMIN did not launch the Matlab tiling jobs for a particular trigger because of an excess of triggers near that time.

Scopes The scopes column held links to see the skymap plots and tile files for each telescope in the network. See Section 9.4 for descriptions of these files.

View Times The view times column provided a link to a plot of a timeline, showing when the top selected tile for each telescope was viewable. See Section 9.4 for details about calculating view times.

ETG The ETG column listed which trigger generator, Omega, cWB unconstrained cWB linear, or MBTA, generated the trigger. It also provided a link to the online web pages generated by that particular trigger generator.

Checklist The final column in the table was only used for triggers labeled in green with the "alert" status, i.e. event candidates. The last column provided a link to the checklist of things the on-call team would verify for each event. Details of the checklist are described in Section 9.6.

Typically, the on-call team would check the LUMIN Events Page immediately after receiving a new alert, and then use the information and links provided while in discussion to evaluate the trigger. A number of mechanisms were provided to inform the on-call team of a new event candidate:

E-mail and Text Messages When a trigger passed all of the automated cuts, LUMIN used the GraCEDb client to pass a message, known as the LUMIN_GO flag, into GraCEDb, Through its own web interface, GraCEDb provided tools for any member of the collaboration to sign up for e-mail or SMS text messages on mobile devices. Users could specify which subset of triggers they preferred to receive. For example, on-call monitors were encouraged to request alerts to their cell phones for all triggers with LUMIN_GO flags.

Control Room Alarms LUMIN maintained a web-accessible text file that contained only a single string, the GraCEDb ID of the last event candidate to pass all automated cuts. In each of the LIGO control rooms, a program regularly polled

this text file, looking for new events. An audio alarm sounded in the control room when a new event was found, alerting the SciMons.

LUMIN Status Page An alternative form of audio and visual alarm was available both inside and outside the control room through the web-based LUMIN Status Page. This web page used an auto-refresh feature every 60 seconds to stay current. The page displayed only events with a FAR less than five events per day, that is, only the most significant events of the day. Most of the time, the page was dominated by a black background. However, at times with an event candidate with a time stamp less than thirty minutes old, the page turned to a bright green color, and sounded a distinctive audio alarm every two minutes. Leaving this web page open on a browser, with a reasonable audio system, effectively alerted a user to each new event candidate.

9.6 Manual Validation

9.6.1 Purpose of Manual Validation

The decision to manually validate each trigger created many of the challenges, as well as much of the excitement, that came along with the online analysis. The purpose of manual validation was to remove triggers that were "obviously" not real events before they were passed to our astronomer partners, and so save unnecessary use of observing time. In principle, most of the checks that the on-call team was asked to make could have been automated. In fact, some of the "manual" checks

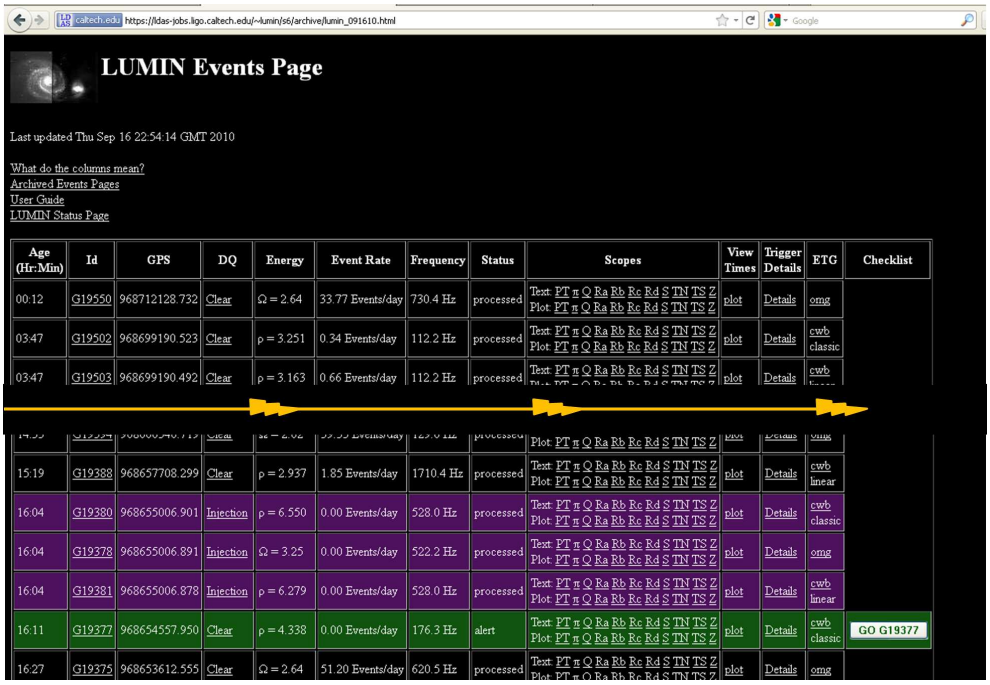

Figure 9.2: A sample screen shot of the LUMIN Events page. Each line represents a trigger from one of the three online trigger generators. Green lines indicate event candidates that require attention from the on-call team, and purple lines indicate scheduled hardware injections.

were actually performed by a script. The true value of the manual validation rested in checking for problems that were not anticipated in advance. By including control room personnel in the manual validation process, there was a good chance that any major problems with the detectors, even unexpected problems, would be brought to the attention of the team, and could veto a trigger if needed. Moreover, any software system, no matter how carefully designed and checked, may contain bugs. The online system was designed so that no telescope ever received a trigger without a manually executed request. This was a safeguard against the fear that the system may have some failure mode where spurious alerts would be repeatedly sent, leading to multiple, extraneous observations.

The manual validation was not intended to eliminate every trigger caused by noise in the detectors. The best way (arguably, the only way) to distinguish between noise events and true gravitational wave signals is with a statistical comparison of zero-lag and time-shifted data. In fact, the manual validation was expected to pass along some noise fluctuations as event candidates. The manual validation was designed to veto a trigger only if it coincided with some detector or software abnormality.

9.6.2 The EVO Conference

When an alert about a new event candidate went out to the LIGO/Virgo collaboration (see Section 9.5), the on-call team quickly assembled for an EVO

conference. EVO (http://evo.caltech.edu) is an application that allows voice, text, and video conferencing, similar to commercial services such as Google Chat and Skype. To allow the relevant parties to quickly get in touch, an EVO "meeting" was permanently maintained only for this purpose. The on-call team needed only to login to this EVO meeting to make contact with each other. In addition, the on-shift pipeline expert typically remained logged-in during her shift. This facilitated a hand-off between shifts, and helped ensure that at least one pipeline expert was always available. In addition to the on-call team, the EVO discussion of a trigger often included additional collaboration members. The assembled team discussed the current trigger, and performed each of the checks in the checklist described in the next section. Typically, the pipeline expert would take the lead in keeping the call organized and efficient. The SciMons, through discussions with the local operators and their own observations, were often able to provide insight into the current detector state and possible sources of excess noise.

9.6.3 The Checklist

From the LUMIN Events Page, a member of the LIGO/Virgo collaboration could click on the button in the "Checklist" column to view the LUMIN Event Checklist Page (Figure 9.3). This page displayed the trigger information listed on the LUMIN Events Page, along with a checklist of actions that the on-call team should perform before approving an event for follow-up observation. More detailed

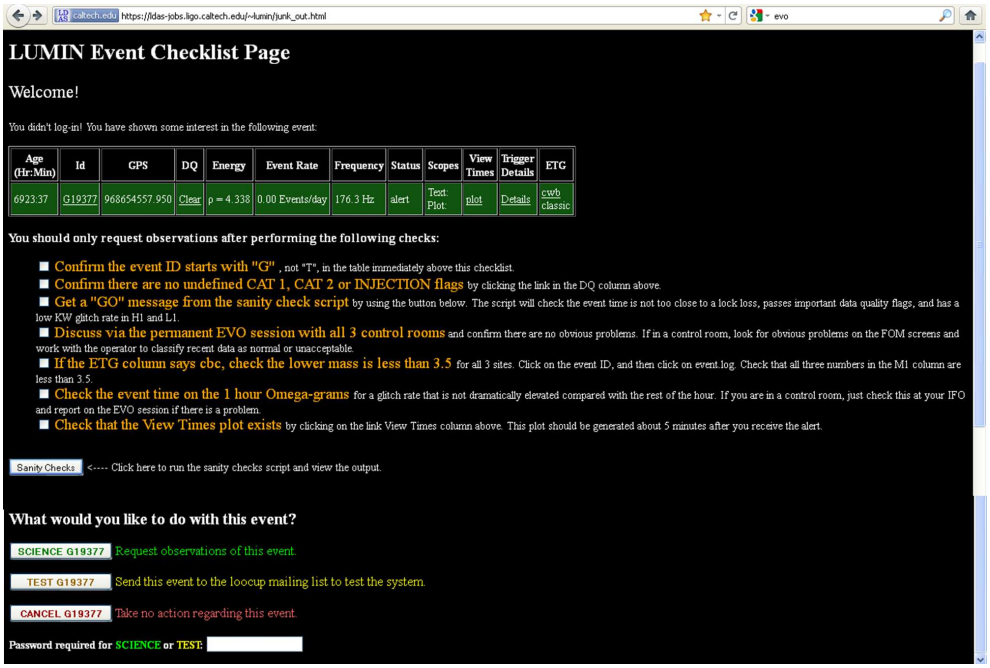

Figure 9.3: A sample screen shot of the LUMIN checklist page. The on-call team uses this list to guide their discussion about an event candidate. If the trigger passes all of the checks, a member of the team uses the SCIENCE button at the bottom of the page to send the trigger's estimated sky position to astronomers and robotic telescopes for observation.

instructions about how to perform each check, and about expectations for the EVO phone conference, were available to the collaboration through the LUMIN User Guide, which was linked from the LUMIN Events Page. Most of the checks were straightforward, and were quickly completed by inspecting one file or another for problems. However, the more subjective checks often led to nuanced discussions, and were the dominant source of latency in the pipeline.

One such check was the instruction to "Check the event time on the 1 hour Omega-grams for a glitch rate that is not dramatically elevated compared with the rest of the hour." The so-called Omega-grams are a common control room tool, and display the frequency and SNR of glitches, or short bursts of power in the GW data stream of each detector. Of course, any single glitch, if found to be in coincidence across several detectors, could be a GW signal. However, an unusually high rate of glitches in a single detector typically indicates a detector problem. Deciding exactly what was and what was not a dramatically elevated glitch rate is a matter of judgment, and this led to time-consuming discussions. Another source of subjectivity in the checklist was the request to check for "obvious problems". Of course, what is obviously a problem to one SciMon may seem subtle or unimportant to another. Typical gray areas included elevated levels of seismic noise around the event time, and the presence of a truck or a train in the vicinity of a detector around the time of an event. While a SciMon may be certain that a train passed near the detector site within minutes of an event, it may not be clear if that train caused the

particular trigger in question.

The LUMIN Event Checklist Page also contained a button to run the "sanity check script". The sanity check script is a Matlab script that read files at each detector site to determine if there were any lock losses within two minutes of the event time. The script also checked the output files from another glitch finding software, KleineWelle, to check for unusually high glitch rates at the two LIGO sites. The sanity check script could be re-run as needed, and the output was viewable on a dedicated web page that was displayed after clicking the "Sanity Checks" button on the LUMIN Event Checklist Page.

Finally, the checklist page contained buttons labeled "SCIENCE", "TEST", and "CANCEL". The green SCIENCE button, when pressed, launched the telescope communication scripts (Section 9.7), and sent the coordinates of the tiles LUMIN selected to the telescope network. The TEST button communicated with LUMIN by the same channel as the SCIENCE button, but only sent the event coordinates to e-mail lists within the LIGO/Virgo collaboration, or in some cases, pinged the servers controlling robotic telescopes. Finally, the CANCEL button did not send any message to LUMIN, but only displayed the "LUMIN Confirmation Page", as did the other two buttons.

For any event, approved or vetoed, the on-call team accessed the LUMIN Confirmation Page at the conclusion of the EVO discussion. The LUMIN Confirmation Page included a link that, when clicked, generated a template "Event Report" for

the event (Figure 9.4). The Event Reports were wiki pages that served as the written record of the EVO conference. The report recorded who was on the call, if the team decided to approve or veto the event, how quickly a decision was reached, and what factors were considered in the decision. Typically, the reports were filled out primarily by the pipeline expert on the call.

9.6.4 Training, Documentation, and Testing

For the autumn observing period, a team of about twenty individuals were trained as pipeline experts to take shifts. Each expert scheduled herself for eight hour shifts so that one pipeline expert was on-call during any moment in the run. The pipeline experts guided the EVO phone conferences, filled out the Event Reports, launched the telescope communication scripts for validated events, and responded to problems in the online system. Training for the experts was provided in several forms. A dedicated EVO phone conference, which was recorded and posted on the web, walked the experts through the procedures. Several instructional wiki pages, the online "LUMIN User Guide", as well as an Expert Manual that was archived as a LIGO technical document, were all made available to the Pipeline Experts. In addition, the entire online system, including the on-call team, was exercised for nearly three weeks prior to the beginning of the autumn run. During this testing period, the on-call team responded to live triggers in the detectors. Triggers that passed the checklist were forwarded to the LIGO/Virgo community via the "TEST"

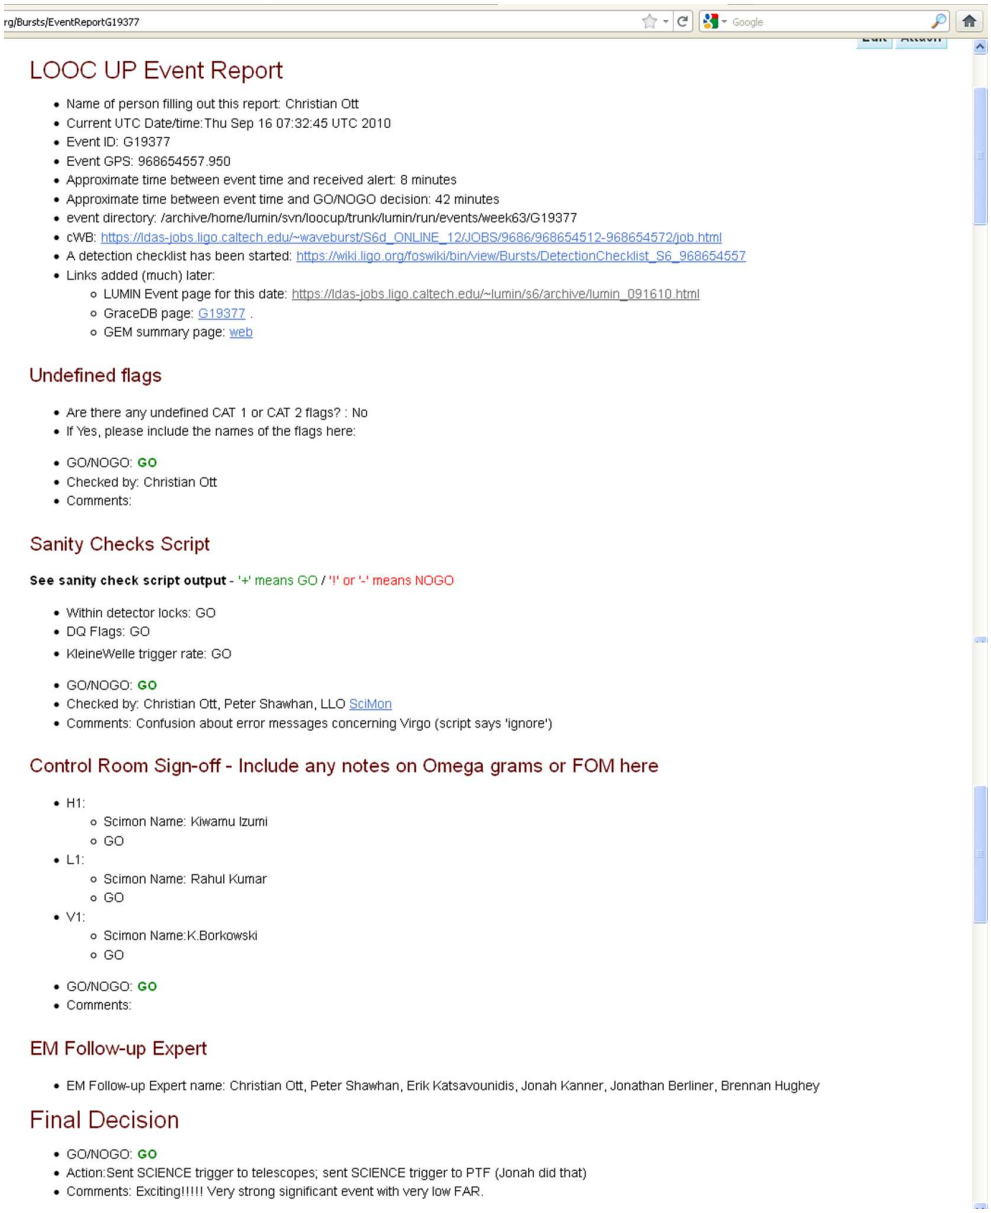

rg/Bursts/EventReportG19377

LOOC UP Event Report

- Name of person filling out this report: Christian Ott
- Current UTC Date/time:Thu Sep 16 07:32:45 UTC 2010
- Event ID: G19377
- Event GPS: 968654557.950
- Approximate time between event time and received alert: 8 minutes
- Approximate time between event time and GO/NOGO decision: 42 minutes
- event directory: /archive/home/lumin/svn/loocup/trunk/lumin/run/events/week63/G19377
- cWB: https://ldas-jobs.ligo.caltech.edu/~waveburst/S6d_ONLINE_12/JOBS/9686/968654512-968654572/job.html
- A detection checklist has been started: https://wiki.ligo.org/foswiki/bin/view/Bursts/DetectionChecklist_S6_968654557
- Links added (much) later:
  - LUMIN Event page for this date: https://ldas-jobs.ligo.caltech.edu/~lumin/s6/archive/lumin_091610.html
  - GraceDB page: G19377 .
  - GEM summary page: web

Undefined flags

- Are there any undefined CAT 1 or CAT 2 flags? : No
- If Yes, please include the names of the flags here:

- GO/NOGO: **GO**
- Checked by: Christian Ott
- Comments:

Sanity Checks Script

See sanity check script output - '+' means GO / '!' or '-' means NOGO

- Within detector locks: GO
- DQ Flags: GO
- KleineWelle trigger rate: GO

- GO/NOGO: **GO**
- Checked by: Christian Ott, Peter Shawhan, LLO SciMon
- Comments: Confusion about error messages concerning Virgo (script says 'ignore')

Control Room Sign-off - Include any notes on Omega grams or FOM here

- H1:
  - Scimon Name: Kiwamu Izumi
  - GO
- L1:
  - Scimon Name: Rahul Kumar
  - GO
- V1:
  - Scimon Name:K.Borkowski
  - GO

- GO/NOGO: **GO**
- Comments:

EM Follow-up Expert

- EM Follow-up Expert name: Christian Ott, Peter Shawhan, Erik Katsavounidis, Jonah Kanner, Jonathan Berliner, Brennan Hughey

Final Decision

- GO/NOGO: **GO**
- Action:Sent SCIENCE trigger to telescopes; sent SCIENCE trigger to PTF (Jonah did that)
- Comments: Exciting!!!!! Very strong significant event with very low FAR.

Figure 9.4: A sample screen shot of a LUMIN Event Report. The on-call team filled out a report for each event candidate. A report template was generated automatically by links within the LUMIN web pages. The report served as a record of the EVO phone call, and documented what factors the team considered when deciding to approve or veto an event.

button on the LUMIN Event Checklist Page. Moreover, unlisted hardware injections were added to the data during this period, so that the on-call team could respond, without warning, to simulated signals. In addition to the documentation provided for the pipeline experts, documentation designed specifically for SciMons was also included in the LUMIN User Guide. These instructions were linked both from the LUMIN Events Page, and from the main source of SciMon instructions.

9.7 Alerts to Telescopes

9.7.1 Alerts through the LUMIN web interface

When the SCIENCE button on the LUMIN Checklist Page was pressed, a socket connection passed a request to LUMIN to make an observing request. LUMIN would then make a series of automated checks. For LUMIN to pass the coordinates to a particular telescope, it required the following:

- The user had entered the correct password on the LUMIN Checklist Page.
- The observation file for the event existed.
- The tile file for the event existed.
- No previous observing requests existed for the event.
- The observation file showed that at least one of the selected tiles was observable for that particular scope for some period of time.

Notice that some of these requirements were intended to guard against spurious requests due to user error. No more than one observing request could go out for each event; this was especially important since multiple users could view the LUMIN Checklist Page, and a miscommunication on the EVO phone call could have led to more than one person pressing the SCIENCE button. The password restriction was intended to limit spurious pressing of the SCIENCE button by members of the LIGO/Virgo collaboration who were not directly connected with the EM follow-up project. The LIGO/Virgo collaboration contains hundreds of individuals, all of whom have access to the LUMIN web pages. However, the password for the SCIENCE button was only shared with the roughly twenty low-latency pipeline experts. Another design feature used sociological pressures to encourage scrupulous use of the SCIENCE button. LUMIN recorded the login name of the user who pressed the SCIENCE button, and shared this information over a LIGO/Virgo e-mail list, along with the trigger details. These safe-guards were found to be sufficient; during the seven week autumn observing period, no observing request was sent without first getting approval on the EVO phone conference.

Each time the SCIENCE button was pressed, LUMIN checked if an observing request was possible for QUEST, TAROT, Zadko, Pi of the Sky, SkyMapper, ROTSE III, the Liverpool Telescope, and LOFAR. The ROTSE III system includes four robotic telescopes, and TAROT includes two. So, during the autumn run, as many as twelve telescopes could have been alerted to the presence of an event

through this interface. Most of these alerts communicated with servers controlling robotic telescopes, and were designed to allow a rapid, automated response. The exception to this practice of automated observing was the LOFAR telescope; alerts to LOFAR were e-mailed to a member of the LOFAR team by a Python script. The Pi of the Sky prototype telescope was undergoing maintenance during the observing run, so in practice, communication with Pi of the Sky was also mediated by a member of that collaboration.

The communication scripts used by LUMIN to pass alerts to TAROT, SkyMapper, Pi of the Sky, and Zadko were all very similar, and were based on Tcl communication scripts created by the TAROT team to make use of a protocol used by the GCN. The ROTSE team also made use of the GCN protocol, but communication was handled by a C code library created by the ROTSE team. The QUEST team used a protocol designed just for this purpose. The server that controls the QUEST telescope routinely polled a unique URL for each night, located on the CalTech server. LUMIN posted the trigger information to the URL for each alert, which was discovered by the QUEST server within a few minutes.

The Liverpool Telescope received alerts through a purpose built Python script. This script was unique within LUMIN in that it included logic for checking if the event was currently viewable by the Liverpool Telescope. The other communication scripts expected computers local to the telescope determine whether the coordinates were viewable and to protect the telescope from pointing in unfavorable locations,

for example, directly at the moon.

9.7.2 Telescopes with Low FAR

All of the observatories that received alerts through the LUMIN web interface, agreed to accept alerts with a target rate of between one per week and a few per week. The FAR thresholds applied by LUMIN, described in Section 9.2, were designed to accommodate this target. However, two of the collaborations in the network, Swift and PTF, preferred to receive a more limited selection of triggers. Communication with each of these observatories was handled as a special case, described below.

The Swift observatory participated in the EM follow-up program through a target-of-opportunity (TOO) agreement for around one set of observations for each observing period. This amounted to a target rate of roughly one GW trigger each month, rather than one each week. In an effort to more selectively choose targets, GEM imposed a threshold on the total P statistic that was captured in five pointings of the Swift X-ray Telescope, in addition to a threshold on FAR.

The communication with the Swift team was achieved through a Swift TOO internet site, that allows users to enter observing requests through a web form. Following the EVO phone conference, if a trigger was approved for observations and met the special GEM criteria, then a designated team member would pass the tiles selected by GEM to the Swift mission controllers.

PTF participated only in the autumn observing period, with a target trigger

rate of one event per month. Experience taught us that the target FAR should be set higher than the desired trigger rate, because FAR was expressed in livetime instead of calendar time, some triggers were vetoed by the on-call team, and some estimated sky positions are not in observable places on the sky. A FAR threshold of 0.15 events per day was selected for PTF. As part of the checks performed by the on-call team, pipeline experts monitored for triggers with a FAR lower than 0.15 events per day. The tile selection for PTF was handled by LUMIN, in a manner similar to the other telescopes. The script to pass events into the PTF scheduler was a custom built Python script that made use of the VOEvent format for passing trigger information. When the on-call team approved a trigger for observing with PTF, the pipeline expert executed a command on the CalTech computers that passed a PTF observing request into LUMIN. LUMIN then treated the observation request in a manner similar to observing requests made through the web interface.

Chapter 10

Data Set and Analysis

The full low-latency pipeline was used during two data collection periods. The "winter run" collected data between December 17, 2009 and January 8, 2010, and the "autumn run" collected data from September 2 to October 20, 2010. This chapter presents trigger properties of the event candidates and an analysis of the significance of the triggers. Also included is an overview of the analyses used to search for EM transients in the optical data.

10.1 Gravitational Wave Event Candidates

10.1.1 Winter Run

During S6, the Virgo detector stopped taking data for a commissioning break between January and August, 2011. This led to two distinct periods of three site observations. During the first period, in winter 2009/2010, the QUEST and TAROT telescopes received triggers from the LIGO/Virgo network.

The main threshold in LUMIN for the winter run was a demand that triggers have a FAR of less than 1.0 events per day. Near the end of the run, this threshold was relaxed to 2.0 events per day. The run yielded nine triggers that passed the

Table 10.1: Table showing the event candidates in the winter run. The FAR is shown in units of events per day. Only observations by optical band instruments are noted in the last column.

Date	Pipeline	FAR	Notes
Dec 20, 2009	Omega	0.7	Request sent; no observations
Dec 20, 2009	Omega	0.0	Vetoed because it was the 2nd event of the day
Dec 21, 2009	cWB	0.9	Vetoed because of weak localization
Dec 26, 2009	Omega	0.3	Request sent; no observations
Dec 29, 2009	Omega	0.7	Collected 12 images with QUEST telescope
Jan 2, 2010	cWB	0.5	Request sent; no observations
Jan 3, 2010	Omega	0.8	Vetoed because of weak localization and 2nd event in 24 hours
Jan 3, 2010	cWB	1.3	Request sent on sub-threshold event on January 7 to collect extra data (6 images with TAROT)
Jan 6, 2010	cWB	0.4	Request sent; no observations
Jan 6, 2010	Omega	4.5	Request sent on sub-threshold event, only to gather more data (9 QUEST Images)
Jan 7, 2010	cWB	1.6	Collected 12 images with QUEST telescope

automated cuts, six of which were cleared by the manual validation. In addition, two weaker triggers were used for EM observations, to increase the amount of optical data. Of the eight observing requests, three resulted in observations with the QUEST telescope, and 1 trigger led to observations with the TAROT telescope. Some of the properties of the winter run event candidates may be seen in Table 10.1.

The number of triggers may be compared with the livetime of the network. The network livetime between midnight GMT on December 17 2009 and midnight GMT on January 8 2010, totals to 9.5 days of triple coincidence livetime after applying

Category 2 vetoes. The target FAR of 1.0 events per day implies we should expect around 9 triggers from each pipeline. In fact, there were five Omega triggers with a FAR of less than 1 event per day, and only three cWB triggers more significant than 1 event per day. It is unclear why there seem to be less triggers than expected. However, the rate of zero-lag triggers agreed with the FAR prediction up to an order unity factor, and this allowed enough control over the observing rate to honor our agreements with the partner observatories.

Monitors in this run were instructed to reject events with weak localization, using the criteria that the sum of the likelihoods in all three QUEST fields contained $P > 0.50$. In addition, to place a hard limit on the rate of requests to observatories, monitors were told to reject any trigger if there had been an observing request within the past 24 hours. These two criteria were used by the monitors to veto 3 of the 9 triggers that lead to LUMIN alerts. For the autumn run, both of these criteria were intentionally dropped, but other criteria became more important.

10.1.2 Autumn Run

There were several significant changes to the online system from the winter run to the autumn run. The number of persons volunteering for "on-call" shifts increased from a small handful, around five, to roughly twenty. Omega Pipeline began using the Ω statistic as its detection statistic, replacing probSignal. cWB added a second mode of running, which imposed a linear polarization constraint on

the signal model, and so effectively ran two separate ETGs during the autumn run. In addition, the MBTA pipeline was added, bringing the total count of ETGs to four. To accommodate the larger number of ETGs, LUMIN's threshold for an alert was tightened to 0.25 events per day. Finally, the telescopes Pi of the Sky, ROTSE, PTF, SkyMapper, Zadko, LOFAR, and Liverpool Telescope joined the project for the autumn run.

Table 10.2 shows some of the properties of the event candidates in the autumn run. The triple coincidence livetime for this run, after applying Category 2 data quality, is 9.0 days, as measured between midnight GMT between September 4 and October 14, 2010. Formally, the run extended to October 20, but some commissioning work began on the LIGO interferometers October 15, so very little triple coincidence livetime was available after this date. So, for each pipeline, we expect $9/4 = 2.25$ events. We see 3 events from linearly polarized cWB, 4 from unmodelled cWB, and 3 Omega events. Of the 8 MBTA triggers, only 2 passed the cut that demanded at least one of the objects had a mass consistent with a neutron star. Of the 3 Omega events, two were spurious triggers caused by a failure mode of the background estimate. Omega Pipeline used data from the twenty four hours of calendar time preceding the trigger to estimate the background rates. If no or very little livetime occurred in that range of data, then the algorithm could fail, and erroneously marked triggers as having an inflated significance. Rejecting these two triggers leaves only one Omega trigger that was a candidate for optical follow up.

Table 10.2: Table showing the event candidates in the autumn run. The FAR is shown in units of events per day. Only observations by optical band instruments are noted in the last column.

Date	Pipeline	Detection Statistic	FAR	Notes
Sept 4	Omega	$\Omega = 2.78$	0.00	Vetoed due to invalid FAR estimate
Sept 4	Omega	$\Omega = 2.89$	0.00	Vetoed due to invalid FAR estimate
Sept 7	cWB (both)	$\eta = 3.45$	0.13	Observed by ROTSE IIIc
Sept 16	cWB (U)	$\eta = 4.34$	0.00	Observed by ROTSE, TAROT, Zadko, SkyMapper
Sept 18	MBTA	$\rho_c = 11.5$	0.09	Vetoed because mass estimate exceeded limit
Sept 19	MBTA	$\rho_c = 10.0$	0.16	Observed by LOFAR, TAROT, Zadko, SkyMapper, QUEST, and ROTSE
Sept 25	MBTA	$\rho_c = 11.3$	0.01	Veoted because mass estimate exceeded limit
Sept 26	cWB (both)	$\eta = 3.73$	0.00	Observed by PTF and ROTSE
Sept 28	MBTA	$\rho_c = 11.9$	0.07	Vetoed because mass estimate exceeded limit
Sept 29	MBTA	$\rho_c = 11.8$	0.02	Vetoed because mass estimate exceeded limit
Oct 2	MBTA	$\rho_c = 10.2$	0.20	Vetoed because mass estimate exceeded limit
Oct 2	MBTA	$\rho_c = 15.4$	0.01	Vetoed because mass estimate exceeded limit
Oct 3	Omega	$\Omega = 3.06$	0.21	Observed by LOFAR, Pi of the Sky, SkyMapper, Liverpool, QUEST, ROTSE
Oct 6	MBTA	$\rho_c = 10.1$	0.18	Unobservable in optical due to proximity to sun
Oct 12	cWB (both)	$\eta = 3.3$	0.22	Vetoed due to high winds and glitch rate at Virgo

In total, the nine days of livetime led to 15 triggers alerting the on-call team. After rejecting the spurious Omega triggers and the MBTA triggers without a NS template, there were seven remaining triggers for the run. Compared with the target of one trigger for each day of livetime, we conclude that the FAR system has, in this case, been successful in controlling the trigger rate.

Looking at the manual validation process, only one event was rejected due to worrisome conditions at the interferometers. The Omega triggers rejected due to bad background estimate demonstrate the type of problem that the manual validation was intended to protect against. The rejection of MBTA triggers with high mass templates was a planned part of the procedure. In short, the manual validation seems to have been reasonably successful in rejecting only obviously bad events.

Of note is the September 16 event, discovered by cWB with an unusually low FAR. The online background estimate calculated this event to occur once every 191 days of livetime, that is, to have had only a 5% chance of arising from background in the course of the autumn run. This is discussed further in Section 10.3.

10.2 Optical Data

In the winter run, four of the eight imaging requests resulted in some follow-up data collection, each with a single telescope.

The greater variety of telescopes in the autumn run led to a more thorough response to triggers. Observing requests were made for six GW triggers. One of

these triggers had a reconstructed position within a few degrees of the sun, and so was unobservable with optical instruments. Optical follow-up data was collected for all five of the remaining triggers. In Figure 10.1, timelines show the response of five telescope systems to these triggers. In the figure, each horizontal line corresponds to one GW trigger. Marker positions indicate the time between the GW time stamp and an associated follow-up image.

The lowest latency responses recorded images with Zadko 36 minutes GW timestamp, and with TAROT 43 minutes after the GW timestamp. In these cases, nearly all of the latency was in the alert generation pipeline, with the robotic telescopes responding promptly to observing requests. In other cases, robotic schedulers needed to wait for dark skies or for a target to rise above the horizon, leading to wait times of a few hours up to twenty hours. In many cases, follow-up images continued for several days, up to a month after the trigger time. For all five GW triggers, at least one image was captured within the first twenty-four hours following the alert.

10.3 Analysis of Gravitational Wave Event Candidates

The first question to ask about the event candidates in Tables 10.1 and 10.2 is if any of them are inconsistent with the background, and so are promising candidates. Each run spanned livetime totaling roughly nine days. As a somewhat arbitrary threshold, we may seek candidates which are marginally inconsistent with noise as those that have less than a 10% chance of arising from the background, and so look

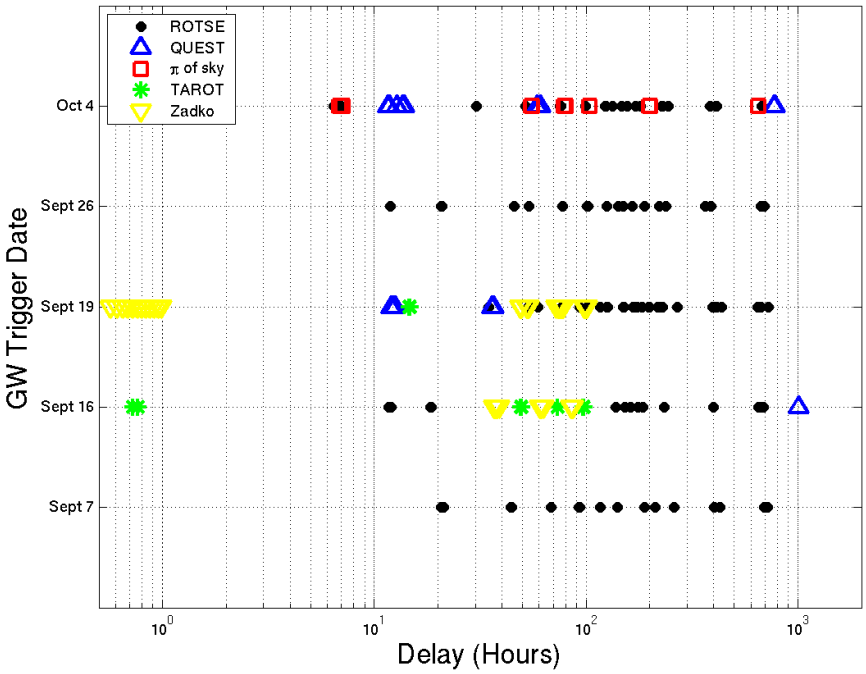

Figure 10.1: Figure showing the delay between the trigger time and each observation for five different telescope systems in the autumn run. The dates on the y-axis indicate GW triggers in the autumn run. Each marker on the plot shows the number of hours between the GW trigger and a follow-up observation.

for triggers with a FAR of less than one event per 90 days, or 0.01 events per day. In the winter run, one of the December 20 Omega triggers seems to meet this criteria. However, a closer look at the background estimate from this day reveals that the apparent high significance comes from a lack of background data available at run time. Using an expanded background sample shows that this event actually has a measured FAR of 1 event per 88 days, and so is very close to the selected threshold. At the time, however, our procedure was to limit observing requests to one event in a twenty four hour period, and this caused us to reject this trigger for follow-up study. In retrospect, this was unfortunate, and this procedure was modified for the autumn run. With the evidence available, there is no reason to believe this event represents a true signal.

Looking to the autumn run, we first eliminate the Omega events with faulty background estimates, and the MBTA events with high masses that fall outside of this analysis. This does leave two triggers of note. First, the cWB trigger event of September 26 has a FAR that was measured to be around 0.015 events per day, or one background event for every 68 days of livetime. This does not quite meet our criteria of less than a 10% chance of arising from background, but it is certainly close. However, this trigger occurred in an interval of data marked as having a train in the vicinity of the Livingston interferometer. Trains are known to cause excess seismic noise, and so excess triggers. Properly evaluating the FAR of this trigger would mean comparing with other times of elevated noise, and would likely degrade

the significance of the trigger further. In short, it would be very difficult to make a case for this trigger as a true gravitational wave event.

The Sept 16 cWB trigger was seen to have a FAR of roughly 1 event per 191 days, or a 5% chance of arising from noise in a single pipeline in the autumn run. At this significance, it is possible that a rare EM transient found in association with the trigger could make a convincing case for a true detection. Moreover, the trigger spectrogram was immediately recognized by the on-call team as exhibiting the "chirp" pattern that is characteristic of inspiraling compact objects. This chirp pattern is known to be highly unusual in noise events. Notably, the event report, posted by the on-call team one hour after the trigger time-stamp, includes the comment: "Exciting!!!!! Very strong significant event with very low FAR."

The skymap produced by cWB may be seen in Figure 10.2. Marked as a rectangle is the location and approximate size of the field imaged by ROTSE and TAROT. A closer view of this region is shown in Figure 10.3. In this Figure, the locations of galaxies in the GWGC are marked as circles. The five galaxies observed by the Zadko telescope in response to the GW trigger are marked as filled circles.

If the trigger was caused by a compact object merger, then we would expect template-based searches, such as MBTA, to distinguish the event from noise better than the Burst pipelines. In fact, MBTA did find triggers in the Hanford and Livingston data at the time of this event, but failed to find a trigger in Virgo. Since MBTA demanded triple coincident triggers, a requirement mainly imposed

by the desire to perform sky localization, it reported no event candidate at this time. However, two weeks after the event was discovered by the online analysis, the flagship inspiral analysis pipeline was run over the data, and did find an inspiral event, with the remarkable FAR of 1 in 7000 years. In fact, the high significance of the event seemed convincing enough to label it the first GW detection based on the GW data alone.

The unusual nature of this event spurred an eager search for a counterpart in the associated optical data, described in the next section. In March, 2011, the trigger was revealed to have been caused by an unmarked, simulated signal, injected into the interferometers as part of an excercise known as the blind injection challenge [1]. The immediate identification of this injection as a significant event, and successful acquisition of associated optical data, were successes of the EM counterpart search implementation.

10.4 September 16 Image Analysis

Finding transients in optical images is a complex endeavor. To highlight some possible approaches and common issues, we present the results of a preliminary analysis completed for images associated with the September 16, 2010 trigger.

[1] http://www.ligo.org/news/blind-injection.php

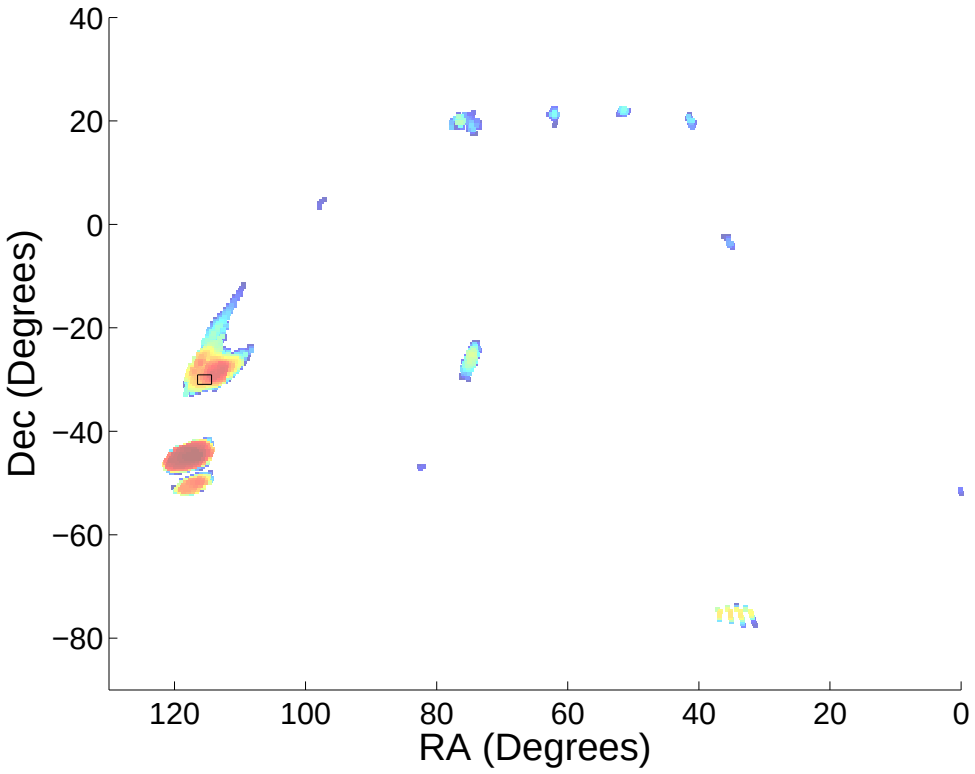

Figure 10.2: Most of the September 16 skymap, produced by cWB. The cWB skymap is shown with the highest likelihood regions in red. The field that was chosen for ROTSE and TAROT is marked as a rectangle.

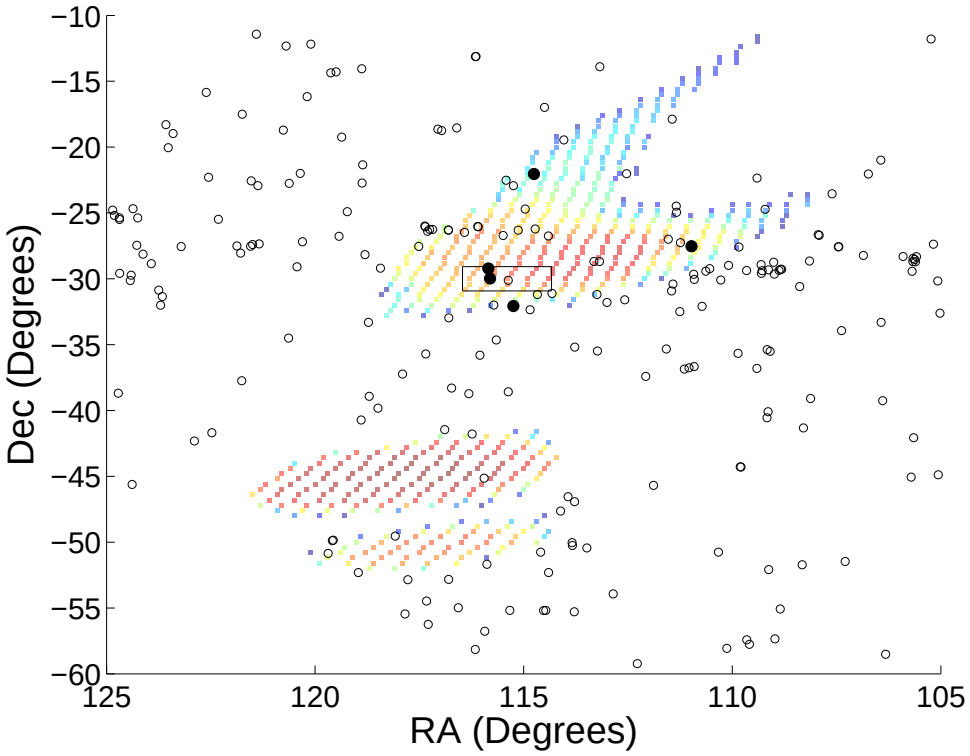

Figure 10.3: Close up of the region of interest for the September 16 skymap. Positions of galaxies are marked as circles, and the cWB skymap is shown with the highest likelihood regions in red. The field that was chosen for ROTSE and TAROT is marked as a rectangle. The five galaxies observed by Zadko are shown as solid circles.

10.4.1 ROTSE

The ROTSE telescopes in Australia (ROTSE IIIa) and Namibia (ROTSE IIIc) collected a total of 163 images for this trigger, with 102 of them containing usable data. The first set of thirty images was collected starting 11 hours and 44 minutes after the trigger time.

In order to search for transients with various time-scales, both a long timescale and a short timescale analysis was performed. The short timescale analysis assumed the transient would only be visible the first night, and sought objects with fading light curves in the thirty images collected on September 16. The long timescale analysis assumed that the object would fade over the course of several days, and so coadded the images for each night, and then looked for a light curve that faded over several nights.

In both cases, the analysis made use of an image subtraction algorithm developed by the ROTSE team [101]. The algorithm attempts to correct for differences in point-spread functions between two images, and then subtracts pixel by pixel. The resulting difference image should clearly reveal any object that has significantly changed brightness in the time between the two images. After image subtraction, the pipeline used the package SExtractor to identify point objects in the resulting difference images. Any object that was identified in two difference images was labeled as a candidate transient. This resulted in 65 candidates for the short timescale analysis, and 36 candidates in the long timescale analysis.

Following this, cuts were applied to the location and morphology of the candidates. Candidates were visually inspected for signs of faulty subtraction, and rejected if they seemed to be a result of this, and they were checked against catalogs of known objects to ensure they were really new. After these cuts, 16 short timescale, and 14 long timescale, candidates remained.

Light curves were generated for the remaining candidates. The light curves were inspected for a dimming signature. None of the light curves showed consistent dimming, and so no potential transients were identified.

10.4.2 TAROT and Zadko

The TAROT telescope took a set of two images beginning 44 minutes after the trigger, and more images over the following four nights for a total of twenty images. The Zadko telescope took a series of 30 images 36 hours after the trigger time, but then was unable to observe due to a mechanical problem. Zadko took some additional fields 160 days after the trigger time, for a total of 63 usable images.

The TAROT and Zadko images were searched using the same pipeline, with different choices of pipeline parameter values. Unlike the analysis used for ROTSE and SkyMapper, this pipeline did not employ image subtraction at any stage. Instead, after data reduction and photometric calibration, the first step used SExtractor to create a catalog of objects in each image. A series of automated cuts was applied to these lists of objects, typically resulting in a small number of candidates

to be visually inspected.

For the TAROT images, the pipeline first looked for objects that could be found in at least seven of the images. It then kept only objects that were brighter than 12th magnitude at some point in the sequence. Objects close to the edge of images were rejected. To be considered a candidate transient, the object must fade by more than 1.5 magnitudes over the course of observations. A second, similar analysis was run on the same images, with modified parameters to search broadly for objects brighter than 10th magnitude.

The Zadko analysis was run with slightly different parameters for each of the five galaxies observed. The Zadko images were searched for transients which had a peak brightness greater than magnitude 14.5, and which faded by more than 4 magnitudes in the 160 days between observations.

Each analysis resulted in zero to three candidates. Each candidate was visually inspected, and found to be due to some kind of instrument artifact, such as "wings" on saturated CCDs.

10.4.3 SkyMapper

SkyMapper took 21 images in a mosaic seven days after the September 16 trigger. An image subtraction based pipeline was employed to search for objects in these images.

The pipeline used SExtractor to identify objects in the difference images, and

then employed a series of cuts based on the object's shape parameters to reject artifacts and uninteresting objects. Cuts to reject artifacts included rejection of non-Gaussian shapes, rejection of misalignment based on a pattern of bright pixels next to dark pixels, and a check against lists of known variable stars. The pipeline found no candidate transients to 20th magnitude.

10.5 Image Analysis Discussion

The various approaches to searching for transients in images reveal some of the general issues that arise in such a search. The specific type and tuning of the analysis has a major impact on the number of false transients that pass the cuts. The SkyMapper analysis reveals that cutting artifacts based on shape parameters can be a very powerful tool for removing them. In contrast, the superficially similar ROTSE analysis, without cuts on shape parameters, yielded more false positives, and so relied more heavily on manual inspection of images. The TAROT and Zadko analyses had few false positives, but to achieve this, made significant demands on the brightness and repeated observing of objects.

The fact that the EM transient FAR depends strongly on pipeline choice has serious implications for the potential to use transient discovery as a method to boost confidence in a GW candidate signal. Any increase in the statistical confidence of a GW candidate will depend on having an estimate of the FAR of the associated EM transient. Since the false positive rate is pipeline dependent, this transient FAR will

need to be estimated using the same pipeline that is employed to search the GW follow-up images. To measure such a FAR methodically means developing cuts that do not rely on manual inspection, a potentially challenging issue.

In addition to the types of features used to eliminate background, searches for optical transients may differ in the properties of the transients they seek. For example, the searches with TAROT, Zadko, and ROTSE all made assumptions about how much, and how quickly, an optical transient should fade. So, a search may only be as successful as the model of the transient it is seeking. Should the transient fade over the course of an hour, a day, or a month? Does the transient only fade, or is a brightening time observable? Searching for broader classes of transients will generally lead to more false positives and more identifications of objects not linked with GWs, such as variable stars. Thus, a trade-off exists between searching for broad classes of transients and searching sensitively for specific transients with a model of the expected light curve.

Even once the parameters of a pipeline are set, measuring the FAR is not a simple matter. Other dependencies may be important when determining the FAR. For example, a data set in the galactic plane will have more stars, and so more false positives, than a data set outside the galactic plane. Similar reasoning applies to data sets with very nearby galaxies, which are extended objects, and so may cause confusion for catalog creation programs such as SExtractor.

A consideration of the rates of some common classes of EM transients shows

that the problem of unassociated transients is extremely important. In fact, within the volume where NS mergers are observable with Advanced LIGO, EM transients that are *not* associated with the mergers will be much more common than those that are. After image subtraction, the greatest number of remaining objects are typically detector artifacts or solar system objects, such as asteroids. Demanding spatial coincidence with a galaxy, demanding the point-spread function is consistent with a point object source, and requiring that the object exhibits no observable proper motion are important handles for removing these types of search contaminants. At this stage, a large area suvey will still contain a number of astrophysical transients that are unlikely to be associated with observable GW emission. A study of the rates of these unassociated transients is presented by Shri Kulkarni and Mansi Kasliwal in [54]. The authors find the dominant sources of astrophysical contaminants to be type Ia supernovae and M-dwarf flare stars. They estimate that, even after demanding coincidence with a galaxy, between 1% and 10% of GW triggers from the advanced detector network will have sky regions that will contain one of these objects. A false coincidence rate of $\sim 10\%$ would mean that any associations between EM transients and GW triggers could not be made with confidence. Finding the EM counterparts to GW triggers, which should be a more rare class of transient, is a "needle in a haystack" problem. At a recent conference, Josh Bloom suggested the solution to this problem is to "learn to recognize hay." In other words, the real data analysis problem of searching for EM transients is not only a question of finding transients –

there are lots of unrelated transients to find – but a question of *classifying* transients. Making a high confidence association between an Advanced LIGO trigger and an EM counterpart will require correctly classifying the familiar, unrelated transients, such as distant supernovae and M-dwarf flares.

This all suggests that, while the identification of an optical transient will surely boost confidence in gravitational wave detection, a quantitative statement of the increased confidence will take a great deal of effort to achieve. It also suggests that care to develop and test a sophisticated transient identification pipeline would be well worth the investment. It is worth noting here that some synergies in this direction certainly exist. Large optical surveys, such as PTF, Pan-STARRS, and the future LSST, are motivating a great deal of work in developing fast, automated, transient finding pipelines (See, for example, [87, 25]).

Chapter 11

Summary

During the 2009/2010 science runs of LIGO and Virgo, data from gravitational wave observatories was used to direct follow-up observations from robotic telescopes for the first time. A large number of obstacles, both technical and sociological, needed to be overcome to make this possible.

Perhaps most daunting was the task of resolving the discrepancy between the large uncertainty in GW position reconstruction, which is of order tens of square degrees, and the size of the FOV of current astronomical equipment, which is typically no more than a few square degrees. The limited range of the 2009/2010 network of GW observatories provided a key to partially resolving this discrepancy. Focusing the search on likely host galaxies in the nearby universe made position reconstruction possible, at least for a fraction of potential sources. Near the detection threshold, signal source positions in simulations were recovered 30-50% of the time, with three fields of a 1.85° FOV telescope.

Another new feature of the search was the completeness of the analysis in low-latency. While many control room tools have run in low-latency in the past, this was the first time that a complete search for GW transients, including coincidence and/or coherence between sites, background estimation, and data quality checks, had been

performed in low latency. Moreover, an on-call team was immediately alerted to the presence of significant triggers, opening the door to follow-up studies in the minutes and hours following an event. To make this run effectively, a collection of software programs had to communicate smoothly and continuously. In addition, the protocol for the on-call team, including schedules and instructions for handling events, had to be constructed, documented, and communicated to team members. An interactive, web based system was created for the team to collect, view, and document information on events, and to execute software to pass coordinate information to robotic telescopes.

Finally, the arrangements with astronomer colleagues demanded some care. Details of events communication, observing schedules, data access, and even data analysis all had to be arranged. The large, heterogeneous network made this a demanding task, with each astronomical collaboration settling on different procedures in all these areas.

Though this study demonstrates significant progress, some work remains to fully prepare for EM follow-up observations in the era of advanced GW detectors. Choosing fields to observe in 2009/2010 relied on the limited range of the GW network. As LIGO and Virgo move from a binary NS inspiral horizon of roughly 35 Mpc to 450 Mpc, the number of potential host galaxies within that horizon will increase enormously, as the search volume will increase by a factor of a thousand. It is unclear if the galaxy weighting procedure used in the last run will be effective

when applied to so many potential host galaxies. In the next few years, work will be needed to find observing strategies and source catalogs that can accommodate this expanding range. In fact, at least one effort is currently underway to construct a complete catalog of the Northern sky out to the sky *averaged* advanced detector range of 200 Mpc [44]. Even with this complete, it seems unlikely that weighting the several hundreds of galaxies associated with a typical uncertainty region will reduce the needed amount of observing time. However, these options should be considered more carefully.

Some work will be needed to construct the data products and delivery software that will be available to astronomers working with advanced detectors. The 2009/2010 model, where gravitational wave scientists made decisions regarding which observatories will receive which triggers, will likely not be sustainable. This means triggers will need to be delivered with more complete information, including significance evaluation, and possibly a complete skymap. Some technical questions need to be addressed, such as constructing a file format for the $\sim$ 1000 pixel skymaps, as well as more basic questions, such as how to sensibly communicate data quality information about triggers. Additionally, it may be desirable to reduce latency by eliminating the manual validation of each trigger, and this would place further demands on the online software system.

Moreover, to gain benefits from the follow-up observation, one must be able to reliably identify significant transients in the images. Development, testing, and

characterization of transient classification pipelines will be needed for this to be successful in the future. Fortunately, the GW community is not alone in seeking to improve this technology. Current and future wide area optical surveys, including Pan-STARRS, SkyMapper, PTF, and the future LSST, are placing increasing demands on the accuracy of transient identification software. Hopefully, progress generated within these projects will benefit the GW community as well.

Despite the challenges that lie ahead, the beginning of low-latency GW analysis and EM follow-up observations has had significant impacts on the gravitational wave community, as well as the greater astronomical community. Several recent conferences have included EM follow-up observations to GW signals as a central discussion point [1] [2] [3]. The demonstration of a complete, working, low-latency system with position reconstruction, has set a precedent for the next generation of LIGO and Virgo, expected to be operational around 2015. Discussions about low-latency triggers from advanced detectors now center around questions such as "What information will be available?" and "How should effective observing strategies be designed?" The questions of *if* low-latency triggers, or *if* estimated source positions, will be available, are no longer points of discussion – the existence of the needed tools is now a fact.

[1] http://www.gravity.phys.uwm.edu/conferences/gwpaw/

[2] http://www.ligo.caltech.edu/gwodw2011/program.shtml

[3] http://www.confcon.com/head2011/

Bibliography

[1] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, and et al. Implementation and testing of the first prompt search for electromagnetic counterparts to gravitational wave transients. *ArXiv e-prints*, September 2011.

[2] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, C. Adams, R. Adhikari, P. Ajith, B. Allen, G. Allen, E. Amador Ceron, and et al. Calibration of the LIGO gravitational wave detectors in the fifth science run. *Nuclear Instruments and Methods in Physics Research A*, 624:223–240, December 2010.

[3] J. Abadie et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys. Rev. D*, 81(10):102001, May 2010.

[4] J Abadie et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27:173001, 2010.

[5] J Abadie et al. Search for gravitational waves from compact binary coalescence in LIGO and Virgo data from S5 and VSR1. *Phys. Rev. D*, 82(10):102001, 2010.

[6] B. Abbott, R. Abbott, R. Adhikari, A. Ageev, B. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. Araya, H. Armandula, and et al. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Phys. Rev. D*, 72(8):082001–+, October 2005.

[7] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, and et al. All-Sky LIGO Search for Periodic Gravitational Waves in the Early Fifth-Science-Run Data. *Phys. Rev. Letters*, 102(11):111102–+, March 2009.

[8] B. P. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev. D*, 80(10):102001, Nov 2009.

[9] B P Abbott et al. Search for gravitational-wave bursts associated with gamma-ray bursts using data from LIGO science run 5 and Virgo science run 1. *The Astrophysical Journal*, 715:1438, 2010.

[10] C W Akerlof, R L Kehoe, T A McKay, et al. The ROTSE-III robotic telescope system. *Publications of the Astronomical Society of the Pacific*, 115(803):132–140, 2003.

[11] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton. FINDCHIRP: an algorithm for detection of gravitational waves from inspiraling compact binaries. *ArXiv General Relativity and Quantum Cosmology e-prints*, September 2005.

[12] N Andersson, V Ferrari, D I Jones, K D Kokkotas, B Krishnan, J S Read, L Rezzol la, and B Zink. Gravitational waves from neutron stars: promises and challenges. *General Relativity and Gravitation*, 43(2):409–436, 2011.

[13] C Baltay, D Rabinowitz, P Andrews, A Bauer, N Ellman, et al. The QUEST large area CCD camera. *Publications of the Astronomical Society of the Pacific*, 119:1278–1294, 2007.

[14] J. F. Beacom and P. Vogel. Can a supernova be located by its neutrinos? *Phys. Rev. D*, 60(3):033007–+, August 1999.

[15] E Berger. A short GRB no-host problem? Investigating large progenitor offsets for short GRBs with optical afterglows. *The Astrophysical Journal*, 722:1946, 2010.

[16] L. Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Reviews in Relativity*, 9:4–+, June 2006.

[17] J. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter. Black-hole binaries, gravitational waves, and numerical relativity. *Reviews of Modern Physics*, 82:3069–3119, October 2010.

[18] S. Chatterji, A. Lazzarini, L. Stein, P. J. Sutton, A. Searle, and M. Tinto. Coherent network analysis technique for discriminating gravitational-wave bursts from instrumental noise. *Phys. Rev. D*, 74(8):082005–+, October 2006.

[19] Alessandra Corsi and Peter Mészáros. Gamma-ray burst afterglow plateaus and gravitational waves: Multi-messenger signature of a millisecond magnetar? *The Astrophysical Journal*, 702, 2009.

[20] D M Coward, M Todd, T P Vaalsta, M. Laas-Bourez, A Klotz, A Imerito, L. Yan, P Luckas, A B Fletcher, M G Zadnik, R R Burman, D G Blair, J Zadko, M Boer, P Thierry, E J Howell, S Gordon, A Ahmat, J Moore, and K Frost. The Zadko telescope: A southern hemisphere telescope for optical transient searches, multi-messenger astronomy and education. *Preprint*, 2010.

[21] Curt Cutler, Theocharis A. Apostolatos, Lars Bildsten, Lee Samuel Finn, Eanna E. Flanagan, Daniel Kennefick, Dragoljubov M. Markovic, Amos Ori, Eric Poisson, Gerald Jay Sussman, and Kip S. Thorne. The last three minutes: Issues in gravitational wave measurements of coalescing compact binaries. *Phys. Rev. Letters*, 70:2984–2987, 1993.

[22] N Dalal, D E Holz, S Hughes, and B Jain. Short GRB and binary black hole standard sirens as a probe of dark energy. *Phys. Rev. D*, 74:063006, 2006.

[23] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys. Rev. Letters*, 85:3761–64, 2000.

[24] S Desai, E O Kahya, and R Woodard. Reduced time delay for gravitational waves with dark matter emulators. *Phys. Rev. D*, 77:124041, 2008.

[25] C. Donalek, A. Mahabal, B. Moghaddam, S. G. Djorgovski, A. J. Drake, M. J. Graham, and R. Williams. Real Time Classification of Transient Events in Synoptic Sky Surveys. In *American Astronomical Society Meeting Abstracts #215*, volume 42 of *Bulletin of the American Astronomical Society*, pages 477.03–+, January 2010.

[26] S Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New Journal of Physics*, 11:123006, 2009.

[27] Rob Fender, Ralph Wijers, Ben Stappers, et al. The LOFAR transients key project. *Proceedings of Science*, MQW6:104, 2006.

[28] V. Ferrari, S. Matarrese, and R. Schneider. Gravitational wave background from a cosmological population of core-collapse supernovae. *MNRAS*, 303:247–257, February 1999.

[29] V. Ferrari, S. Matarrese, and R. Schneider. Stochastic background of gravitational waves generated by a cosmological population of young, rapidly rotating neutron stars. *MNRAS*, 303:258–264, February 1999.

[30] Eanna E. Flanagan and Tanja Hinderer. Constraining neutron star tidal Love numbers with gravitational wave detectors. *Phys. Rev. D*, 77:021502, 2008.

[31] D B Fox et al. The afterglow of grb 050709 and the nature of the short-hard gamma-ray bursts. *Nature*, 437:845, 2005.

[32] C L Fryer, D E Holz, and S A Hughes. Gravitational wave emission from core collapse of massive stars. *The Astrophysical Journal*, 565:430–446, 2002.

[33] N Gehrels, J K Cannizzo, and J P Norris. Gamma-ray bursts in the swift era. *New Journal of Physics*, 9:37, 2007.

[34] N Gehrels et al. The swift gamma-ray burst mission. *The Astrophysical Journal*, 611:1005, 2004.

[35] Jonathan Granot, Alin Panaitescu, Pawan Kumar, and Stan E Woosley. Off-axis afterglows emission from jetted gamma-ray bursts. *The Astrophysical Journal Letters*, 570:L61 – L64, 2002.

[36] F. Halzen and G. G. Raffelt. Reconstructing the supernova bounce time with neutrinos in IceCube. *Phys. Rev. D*, 80(8):087301–+, October 2009.

[37] W E Harris. *The Astronomical Journal*, 122:1487, 1996.

[38] B. A. Jacoby, P. B. Cameron, F. A. Jenet, S. B. Anderson, R. N. Murty, and S. R. Kulkarni. Measurement of Orbital Decay in the Double Neutron Star Binary PSR B2127+11C. *The Astrophysical Journal Letters*, 644:L113–L116, June 2006.

[39] Piotr Jaranowski and Andrzej Krolak. Optimal solution to the inverse problem for the gravitational wave signal of a coalescing compact binary. *Phys. Rev. D*, 49:1723–39, 1994.

[40] D A Kann. The afterglows of Swift-era gamma-ray burst. 1. comparing pre-Swift and Swift-era long/soft (type ii) GRB optical afterglows. *The Astrophysical Journal*, 720:1513, 2010.

[41] D. A. Kann, S. Klose, B. Zhang, S. Covino, N. R. Butler, D. Malesani, E. Nakar, A. C. Wilson, L. A. Antonelli, G. Chincarini, B. E. Cobb, P. D'Avanzo, V. D'Elia, M. Della Valle, P. Ferrero, D. Fugazza, J. Gorosabel, G. L. Israel, F. Mannucci, S. Piranomonte, S. Schulze, L. Stella, G. Tagliaferri, and K. Wiersema. The Afterglows of Swift-era Gamma-Ray Bursts. II. Type I GRB versus Type II GRB Optical Afterglows. *The Astrophysical Journal*, 734:96–+, June 2011.

[42] Jonah Kanner, Tracy L Huard, Szabolcs Marka, David C Murphy, Jennifer Piscionere, Molly Reed, and Peter Shawhan. LOOC UP: locating and observing optical counterparts to gravitational wave bursts. *Classical and Quantum Gravity*, 25(18):184034, 2008.

[43] I D Karachentsev, V E Karachentseva, W K Huchtmeie, and D I Makarov. *The Astronomical Journal*, 127:2031, 2004.

[44] M. Kasliwal. *Private communication.*

[45] Luke Zoltan Kelley, Enrico Ramirez-Ruiz, Marcel Zemp, Jurg Diemand, and Ilya Mandel. The distribution of coalescing compact binaries in the local universe: Prospects for gravitational-wave observations. *The Astrophysical JournalL*, 725(1):L91, 2010.

[46] K. Kiuchi, M. Shibata, P. J. Montero, and J. A. Font. Gravitational Waves from the Papaloizou-Pringle Instability in Black-Hole-Torus Systems. *Phys. Rev. Letters*, 106(25):251102–+, June 2011.

[47] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev. D*, 72(12):122002–+, December 2005.

[48] S Klimenko, G Vedovato, M Drago, G Mazzolo, G Mitselmakher, C Pankow, G Prodi, V Re, F Salemi, and I Yakusin. Localization of gravitational wave sources with networks of advanced detectors. *submitted to Phys. Rev. D*, 2011.

[49] A Klotz, M Bor, J L Atteia, and B Gendre. Early optical observations of gamma-ray bursts by the TAROT telescopes: Period 2001-2008. *The Astronomical Journal*, 137:4100, 2009.

[50] S Kobayashi and P Meszaros. Gravitational radiation from gamma-ray burst progenitors. *The Astrophysical Journal*, 589:861–870, 2003.

[51] Shiho Kobayashi and Peter Mészáros. Polarized gravitational waves from gamma-ray bursts. *The Astrophysical Journal*, 585:L89–L92, 2003.

[52] O. Korobkin, E. B. Abdikamalov, E. Schnetter, N. Stergioulas, and B. Zink. Stability of general-relativistic accretion disks. *Phys. Rev. D*, 83(4):043007–+, February 2011.

[53] M. Kramer, I. H. Stairs, R. N. Manchester, M. A. McLaughlin, A. G. Lyne, R. D. Ferdman, M. Burgay, D. R. Lorimer, A. Possenti, N. D'Amico, J. M. Sarkissian, G. B. Hobbs, J. E. Reynolds, P. C. C. Freire, and F. Camilo. Tests of General Relativity from Timing the Double Pulsar. *Science*, 314:97–102, October 2006.

[54] S R Kulkarni and M M Kasliwal. Transients in the local universe. *Preprint*, arXiv:0903.0218, 2009.

[55] Nicholas M. Law, Shrinivas R. Kulkarni, and Richard G. Dekany an d others. Palomar transient factory: System overview, performance, and first results. *PASP*, 121(886):1395–1408, 2009.

[56] I. Leonor, L. Cadonati, E. Coccia, S. D'Antonio, A. Di Credico, V. Fafone, R. Frey, W. Fulgione, E. Katsavounidis, C. D. Ott, G. Pagliaroli, K. Scholberg, E. Thrane, and F. Vissani. Searching for prompt signatures of nearby core-collapse supernovae by a joint analysis of neutrino and gravitational wave data. *Classical and Quantum Gravity*, 27(8):084019–+, April 2010.

[57] L.-X. Li and B. Paczyński. Transient Events from Neutron Star Mergers. *The Astrophysical Journal Letters*, 507:L59–L62, November 1998.

[58] D. R. Lorimer. Binary and Millisecond Pulsars. *Living Reviews in Relativity*, 11:8–+, November 2008.

[59] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford. A Bright Millisecond Radio Burst of Extragalactic Origin. *Science*, 318:777–, November 2007.

[60] A. I. MacFadyen and S. E. Woosley. Collapsars: Gamma-Ray Bursts and Explosions in "Failed Supernovae". *The Astrophysical Journal*, 524:262–289, October 1999.

[61] M. Maggiore. Gravitational wave experiments and early universe cosmology. *Physics Reports*, 331:283–367, July 2000.

[62] Katarzyna Malek, Tadeusz Batsch, Mikolaj Cwiok, et al. General overview of the Pi of the Sky system. *Proceedings of SPIE*, 7502:75020D, 2009.

[63] F Marion. Multiband search of coalescing binaries applied to virgo citf data. *Proceedings of Rencontres de Moriond on Gravitational Waves and Experimental Gravity 2003*, 2004.

[64] Dragoljub Markovic. On the possibility of determining cosmological parameters from measurements of gravitational waves emitted by coalescing, compact binaries. *Phys. Rev. D*, 48:4738–4756, 1993.

[65] P. Mészáros. Gamma-ray bursts. *Reports on Progress in Physics*, 69:2259–2321, August 2006.

[66] B. D. Metzger, D. Giannios, T. A. Thompson, N. Bucciantini, and E. Quataert. The protomagnetar model for gamma-ray bursts. *MNRAS*, 413:2031–2056, May 2011.

[67] B D Metzger, G Martnez-Pinedo, S. Darbha, E Quataert, A Arcones, D Kasen, R Thomas, P Nugent, I V Panov, and N T Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *MNRAS*.

[68] S Murphy, S Keller S, B Schmidt, P Tisserand, M Bessell, P Francis, and G Da Costa. SkyMapper and the southern sky survey: A valuable resource for stellar astrophysics. *Astronomical Society of the Pacific Conference Series*, 404:356, 2009.

[69] E Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442:166, 2007.

[70] Samaya Nissanke, Daniel E Holz, Scott A Hughes, Neal Dalal, and Jonathan L Sievers. Exploring short hard gamma-ray bursts as gravitational-wave standard sirens. *The Astrophysical Journal*, 725(1):496, 2010.

[71] Laura K Nuttall and Patrick J Sutton. Identifying the host galaxy of gravitational wave signals. *Phys. Rev. D*, 82:102002, 2010.

[72] Ryan M O'Leary, Richard O'Shaughnessy, and Frederic A Rasio. Dynamical interactions and the black-hole merger rate of the universe. *Phys. Rev. D*, 76(6):061504, 2007.

[73] Christian D Ott. The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, 26:063001, 2009.

[74] G Paturel, P Fouque, L Bottinelli, and L Gouguenheim. *Astronomy and Astrophysics Supplement Series*, 80:299, 1989.

[75] G. Paturel, C. Petit, P. Prugniel, G. Theureau, J. Rousseau, M. Brouty, P. Dubois, and L. Cambrésy. HYPERLEDA. I. Identification and designation of galaxies. *Astronomy and Astrophysics*, 412:45–55, December 2003.

[76] P. C. Peters and J. Mathews. Gravitational Radiation from Point Masses in a Keplerian Orbit. *Physical Review*, 131:435–440, July 1963.

[77] E. S. Phinney. The rate of neutron star binary mergers in the universe - Minimal predictions for gravity wave detectors. *The Astrophysical Journal Letters*, 380:L17–L21, October 1991.

[78] T. Piran. The physics of gamma-ray bursts. *Reviews of Modern Physics*, 76:1143–1210, October 2004.

[79] A. L. Piro and E. Pfahl. Fragmentation of Collapsar Disks and the Production of Gravitational Waves. *The Astrophysical Journal*, 658:1173–1176, April 2007.

[80] Arne Rau, Shrinivas R Kulkarni, Nicholas M Law, et al. Exploring the optical transient sky with the Palomar Transient Factory. *Publications of the Astronomical Society of the Pacific*, 121(886):1334–1351, 2009.

[81] James E Rhoads. Dirty fireballs and orphan afterglows: A tale of two transients. *The Astrophysical Journal*, 591:1097–1103, 2003.

[82] P. R. Saulson. *Fundamentals of interferometric gravitational wave detectors*. 1994.

[83] B F Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310–311, Sep 1986.

[84] B. F. Schutz. Networks of gravitational wave detectors and three figures of merit. *Classical and Quantum Gravity*, 28(12):125023–+, June 2011.

[85] Naoki Seto. Prospects of LIGO for constraining inclination of merging compact binaries associated with three-dimensionally localized short-hard GRBs. *Phys. Rev. D*, 75:024016, 2007.

[86] Xavier Siemens et al. Gravitational wave burst form cosmic (super)strings: Quantitative analysis and constraints. *Phys. Rev. D*, 73:105001, 2006.

[87] D. L. Starr, J. S. Bloom, J. M. Brewer, N. R. Butler, D. Poznanski, M. Rischard, and C. Klein. The Berkeley Transient Classification Pipeline: Deriving Real-time Knowledge from Time-domain Surveys. In D. A. Bohlender, D. Durand, & P. Dowler, editor, *Astronomical Data Analysis Software and Systems XVIII*, volume 411 of *Astronomical Society of the Pacific Conference Series*, pages 493–+, September 2009.

[88] Iain A Steele, Robert J Smith, Paul C Rees, et al. The Liverpool Telescope: performance and first results. *Proceedings of the SPIE*, 5489:679–692, 2004.

[89] R B Tully. *The Astrophysical Journal*, 321:280, 1987.

[90] R B Tully, L Rizzi, E J Shaya, H M Courtois, D I Makarov, and B A Jacobs. *AJ*, 138:323, 2009.

[91] G. Ushomirsky, C. Cutler, and L. Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *MNRAS*, 319:902–932, December 2000.

[92] Tanmay Vachaspati. Cosmic sparks from superconducting strings. *Phys. Rev. Letters*, 101:141301, 2008.

[93] Michele Vallisneri. Prospects for gravitational-wave observations of neutron-star tidal disruption in neutron-star-black-hole binaries. *Phys. Rev. Letters*, 84:3519–3522, 2000.

[94] J. Weber. Gravitational-Wave-Detector Events. *Physical Review Letters*, 20:1307–1308, June 1968.

[95] J. M. Weisberg, D. J. Nice, and J. H. Taylor. Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16. *The Astrophysical Journal*, 722:1030–1034, October 2010.

[96] J. M. Weisberg and J. H. Taylor. The Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis. In F. A. Rasio & I. H. Stairs, editor, *Binary Radio Pulsars*, volume 328 of *Astronomical Society of the Pacific Conference Series*, pages 25–+, July 2005.

[97] D. J. White, E. J. Daw, and V. S. Dhillon. A list of galaxies for gravitational wave searches. *Classical and Quantum Gravity*, 28(8):085016–+, April 2011.

[98] S. E. Woosley. Gamma-ray bursts from stellar mass accretion disks around black holes. *The Astrophysical Journal*, 405:273–277, March 1993.

[99] S. E. Woosley and J. S. Bloom. The Supernova Gamma-Ray Burst Connection. *Annual Review of Astronomy and Astrophysics*, 44:507–556, September 2006.

[100] K. Yamamoto, T. Uchiyama, S. Miyoki, M. Ohashi, K. Kuroda, H. Ishitsuka, T. Akutsu, S. Telada, T. Tomaru, T. Suzuki, N. Sato, Y. Saito, Y. Higashi, T. Haruyama, A. Yamamoto, T. Shintomi, D. Tatsumi, M. Ando, H. Tagoshi, N. Kanda, N. Awaya, S. Yamagishi, H. Takahashi, A. Araya, A. Takamori, S. Takemoto, T. Higashi, H. Hayakawa, W. Morii, and J. Akamatsu. Current status of the CLIO project. *Journal of Physics Conference Series*, 122(1):012002–+, July 2008.

[101] F. Yuan and C. W. Akerlof. Astronomical Image Subtraction by Cross-Convolution. *The Astrophysical Journal*, 677:808–812, April 2008.