May 2016

Monsters in the Dark: High Energy Signatures of Black Hole Formation with Multimessenger Astronomy

Alexander L. Urban
University of Wisconsin-Milwaukee

Monsters in the Dark: High Energy Signatures of Black Hole Formation with Multimessenger Astronomy

by

Alexander L. Urban

A Dissertation Submitted in
Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
in Physics

at

The University of Wisconsin–Milwaukee
May 2016

ABSTRACT

MONSTERS IN THE DARK:
GLIMPSING THE HIGH ENERGY SIGNATURES OF BLACK HOLE FORMATION
WITH MULTIMESSENGER ASTRONOMY

by

Alexander L. Urban

The University of Wisconsin–Milwaukee, 2016
Under the Supervision of Professor Patrick R. Brady

When two compact objects inspiral and violently merge it is a rare cosmic event, producing fantastically "luminous" gravitational wave emission. It is also fleeting, staying in the Laser Interferometer Gravitational-wave Observatory's (LIGO) sensitive band only for somewhere between tenths of a second and several tens of minutes. However, when there is at least one neutron star, disk formation during the merger may power a slew of potentially detectable electromagnetic counterparts, such as short γ-ray bursts (GRBs), afterglows, and kilonovae. These explosions span the full electromagnetic spectrum and are expected within seconds, hours or days of the merger event. To learn as much astrophysics as possible requires targeted observations at every stage of this process, demanding a coordinated worldwide effort across many facilities and multiple astronomical disciplines, all in nearly real-time. In this dissertation I outline some of the major obstacles facing the multimessenger astronomy effort, including computation, data analysis and sky localization for LIGO source candidates, as well as disseminating this information quickly to the astronomical community. I also report on the performance of some of these services during Advanced LIGO's first Observing Run, and on my experience at LIGO Livingston Observatory during the first Observing Run of LIGO's Advanced stage, during which the instruments directly detected gravitational waves for the very first time. (The transient source GW150914 was observed 14 September 2015, and is consistent with a binary black hole merger at redshift $z \approx 0.09$.) I also participate in time-domain optical astronomy with the intermediate Palomar Transient Factory (iPTF) collaboration, searching for orphaned afterglow candidates to better understand the nature of relativistic outbursts such as GRBs.

To the epsilons:

Nikolas, Jasmine, Ivy, Lilly, Rutger,

Matt-Matt, Savannah, Noah, Jaxon,

Kialin, Marley, Carson, Casey, Chelsea, McKenzie,

Nicole and Jamee

I was only holding this candle so I could try to show you the way.

"Kids should be allowed to break stuff more often. That's a consequence of exploration. Exploration is what you do when you don't know what you're doing."

Neil deGrasse Tyson

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AGN active galactic nucleus

aLIGO Advanced LIGO

ASKAP Australian Square Kilometer Array Pathfinder

BAT Burst Alert Telescope (instrument on *Swift*)

BAYESTAR BAYESian TriAngulation and Rapid localization

BBH binary black hole

BH black hole

BNS binary neutron star

CTIO Cerro Tololo Inter-American Observatory

CBC compact binary coalescence

cWB Coherent WaveBurst

DECam Dark Energy Camera (instrument on the Blanco 4-m telescope at CTIO)

EM electromagnetic

ER8 eighth engineering run

FAR false alarm rate

FOV field of view

GBM Gamma-ray Burst Monitor (instrument on *Fermi*)

GCN Gamma-ray Coordinates Network

GRB gamma-ray burst

GSC Gas Slit Camera

GW gravitational wave

IMR inspiral-merger-ringdown

IPN InterPlanetary Network

iPTF intermediate Palomar Transient Factory

ISS International Space Station

KN kilonova

LAT Large Area Telescope

LIB LALInference Burst

LIGO Laser Interferometer GW Observatory

LOFAR Low Frequency Array

LMC Large Magellanic Cloud

LT Liverpool Telescope

MBTA Multi-Band Template Analysis

ML maximum likelihood

MWA Murchison Widefield Array

NSBH neutron star–black hole

NSBH neutron star–black hole

NS neutron star

O1 Advanced LIGO's first observing run

oLIB Omicron+LALInference Burst

P48 Palomar 48 inch Oschin telescope

P200 Palomar 200 inch Hale telescope

PESSTO Public ESO Spectroscopic Survey of Transient Objects

PTF Palomar Transient Factory

RMS root mean square

SKA Square Kilometer Array

S/N signal-to-noise ratio

SN supernova

UVOT UV/Optical Telescope (instrument on *Swift*)

VISTA@ESO Visible and Infrared Survey Telescope

VLA Karl G. Jansky Very Large Array

VLT Very Large Telescope

VST@ESO VLT Survey Telescope

XRT X-ray Telescope (instrument on *Swift*)

ACKNOWLEDGEMENTS

In a letter to Robert Hook dating from 1676, Isaac Newton once said, "If I have seen farther it is only by standing on the shoulders of giants." Of course Newton's sentiment is true for all of us, and I am no different: the road from a little boy desperately frightened of the black hole hiding under his bed, to a grown man nervously defending a thesis about those very same monsters, was paved by a lot of hands. Since this dissertation is (perhaps ironically) the most deeply personal thing I've ever done, it's about time I gave them the salute they deserve.

First, the obvious. My mother, Mickey Carducci, gave me so many things I don't even think I can name them all, but here's an honest shot: Life, love and support; a roof to live under, a childhood of happy memories; endless patience when all I wanted to do was watch *Star Wars*; that time the honest-to-goodness real Batman came to my fifth birthday party; nearly three decades of Christmases, Easters, Halloweens and movie nights; and making sure I knew the sky was never the limit, but the thing I should be reaching for in the first place. Thanks, mom.

Pops, you're not out of the woods either. You couldn't have known it at the time (because nerding out over stuff is kind of our whole thing), but it was when we went to see the film version of Carl Sagan's *Contact* in the summer of '97 that I decided space would be my whole life. I can still taste the soda and popcorn in the arcade room outside the theater at Nickel's Plaza. Thanks for that, dawg.

Hey Grandmother-type-guy, my most favoritest person in the whole wide world. Yes, you. Thanks for all the folksy wisdom, the late night tea-times, the homemade costumes for school, and for helping me write letters to NASA. Also, while I've got your attention: remember when I went to college, and you were worried I'd gotten too big for you? Well, so far I'm still the same young man who would spend a whole Sunday cooking pasta just to make you feel better. I don't ever want to be anyone else.

Aunt Julie! I know how much a fan you are of coincidences, and you know how much I am of not ignoring them. So I'll humor you this one: It's more than a little remarkable, isn't it, that in recent years the universe seems very keen to put us together? I suppose it is remarkable – I just remarked on it. Thanks for the always open door, for frequently acting as my sounding board, and for being the Thelma to my Louise. (Or, wait, is it the other way around?) I can only hope having an astrophysicist in your

pocket has been at least as awesome.

There are bound to be people I'm forgetting, kindred souls I owe my life to but haven't the space to recognize properly. Thanks Josh, Amanda, Sheena and Vanessa for growing up with me. I couldn't be who I am today without a family like you. Eric Nelson, you jerk, thanks for being such a great friend and music companion all these years. And to Mary and Dave Nelson: I *literally* couldn't have made it this far without your wisdom, patience, and faith in my abilities. Thank-you so much.

To Justin Ellis, Maddie and Les Wade, Sydney Chamberlin, Tom and Whitney Linz, Matt Brinson, Joe Simon, Dan Murray, and Chris Pankow: Thanks for all the jazz shows, game nights, nerd rants, and purely philosophical debates. You can't possibly know how much they've kept me grounded during this great, ugly process. Branson Stephens and Kristina Islo, thank you both for understanding me on an emotional level and for being aggressively supportive. Emily Searle-White, my captain, there is one place I belong and that is by your side, always. Signorina Dr. Rachaglini E. Corsi, Esq., formerly of Meadville, you're just an all-around awesome dude. Bethy Winkopp, same comment, and also I suppose being acknowledged in a skyence paper now officially makes you the uncontestable queen of space. You're welcome!

Thanks to the members of 5,015 Unicorns – in particular Allison Fantetti, Blair Klostermeier, Lindsay Ross, and Olive Mash – for being the best trivia team around (and generally awesome dudes). To Mark Poe, hey broh, thanks for going full ham on that RAVEN stuff broh. To Lauren Furst and Heather Markovich, thanks for believing in my intelligence despite my constant protests. To Stefanie Pinnow, Jen Lassiter, Kendra Pilon, Heidi Matera, and Tonia Klein for smooth and painless travel; to Bill Tuma for mocking my various plights (in the most impossibly endearing way!); to Katherine Quinn for being the most amazing person in the world to nerd out over stuff with; and to Wendie Spengler. Thank you for making me feel worthwhile.

To Mara Nichols, the girl from a storybook and leader of the literati. Paraphrasing one of our mutual heroes: to me you are always *"the"* woman; people have seldom heard you mentioned under any other name.

And to every one of the greatest teachers and mentors I've ever had: Mrs. Mesnik, Mrs. Burgess, Mr. Naggle, Mrs. Glotzbecker, Mr. Huyck, Mr. Toth, Mr. Rinehardt, Mr. Cain, Dr. Lo Bello, Dr. Lakins, Dr. Lombardi, Dr. Willey, Dr. Rahman, and Patrick Brady. This work belongs to all of you as much as it does to me.

~

This research is part of the ongoing Laser Interferometer Gravitational-wave Observatory (LIGO) project, funded by the U.S. National Science Foundation, as well as the international LIGO-Virgo Collaboration. Data from the LIGO observatories in Hanford, WA, and Livingston Parish, LA, are used extensively, and three months of research was conducted directly at LIGO Livingston Observatory during the first nominal Observing Run in autumn, 2015. I thank the operators, staff, and run coordinator on-site – particularly Jess McIver, Janeen Romie, Duncan Macleod, Joe Betzweiser, and Shivaraj Kandhasamy – for their warm reception and patient guidance during this on-site stay. The research reported herein has made use of data, software and/or web tools obtained from the LIGO Open Science Center[1] (The LIGO Scientific Collaboration, 2016), a service of LIGO Laboratory and the LIGO Scientific Collaboration.

This work is also based in part on observations obtained with the P48 Oschin telescope as part of the Intermediate Palomar Transient Factory (iPTF) project, a scientific collaboration among Caltech, LANL, UW-Milwaukee, the Oskar Klein Center, the Weizmann Institute of Science, the TANGO Program of the University System of Taiwan, and the Kavli IPMU. I thank the staff of Caltech Optical Observatories, particularly Mansi Kasliwal, Yi Cao and Shri Kulkarni, for kindly hosting a visit to Caltech in November, 2013.

The National Energy Research Scientific Computing Center (NERSC) provides staff, computational resources, and data storage for the iPTF project. Some of the data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California, and NASA; the observatory was made possible by the generous financial support of the W. M. Keck Foundation.

This work made use of data supplied by the UK *Swift* Science Data Centre at the University of Leicester, as well as the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA. This research has also made use of the High Energy Astrophysics Science Archive Research Center (HEASARC) database[2], maintained at NASA Goddard Space Flight Center.

[1] https://losc.ligo.org
[2] http://heasarc.gsfc.nasa.gov

In the course of completing this project, I have been supported by teaching assistantships at the University of Wisconsin-Milwaukee, and under research assistantships funded by NSF grants PHY-0970074 and PHY-1307429 as well as the UWM Research Growth Initiative. I thank my advisor, Patrick Brady, and the faculty and staff of the Leonard E. Parker Center for Gravitation, Cosmology and Astrophysics (CGCA) for their financial and moral support.

Part I

Big Things Have Small Beginnings

Chapter 1

A Brief History of Things That Go Bump in the Night

"The Cosmos is all that is, or ever was, or ever will be. Our contemplations of the Cosmos stir us: there is a tingling in the spine, a catch in the voice, a faint sensation, as if a distant memory, of falling from a height. We know we are approaching the grandest of mysteries."

Carl Sagan, *Cosmos: A Personal Voyage*

IT WAS A DARK AND STORMY NIGHT.

Many a hackneyed horror story has begun this way. Originally the opening line of English novelist Edward Bulwer-Lytton's 1830 novel *Paul Clifford*, the phrase has since come to be derided in literary and horror fiction circles as the prime example of a florid, gravely melodramatic style of writing commonly referred to as "purple prose." (Some authors, such as the fantasy novelist Madeline L'Engle in *A Wrinkle in Time*, have even actively taken to playful parodies of this infamous literary trope.) To an astronomer, however, the patter evokes horror stories of a very different kind: we have come to love the dark, but so thoroughly dread the clouds.

This dissertation covers quite a breadth of topics – from the theoretical background for monstrous, fast, relativistic explosions, to the astounding feats of engineering necessary to build instrumentation that can detect them, to data analysis methods and results of searches for certain classes of transient high-energy astrophysical phenomena. Before launching into technical discussions, however, it is worth motivating this work by placing it in its larger historical context. There are many reasons I can think of for doing this, but perhaps the most imminently compelling is to answer the following

question: *What have you done that is different from anything that came before you?*

To that end, in the annals of scientific history there are two particular plot threads that weave together to lend context to the work I present in this dissertation. I will begin by telling each story in brief, focusing on the details that will be relevant again in later chapters; certain topics mentioned here will be fleshed out in more detail in chapter 2. (Whenever possible, I have embedded hyperlinks in this document to specific chapters, sections, figures, tables, external URLs and cited works where they are referenced.) I will then conclude this chapter by briefly laying out the organizational structure of the rest of this work.

The first order of business is a beautiful bit of expository narrative that engages a lingering scientific mystery, with a dash of Cold War-era spy thriller thrown in for good measure. It was a time when we as a species, frightened at our own destructive potential, for once decided to put down our weapons and reach for the stars instead.

And it begins, of course, on a dark and stormy night....

1.1 Gamma-Ray Bursts: A Scientific Mystery Story

Project *Vela* was a program designed by the U.S. military to monitor the Soviet Union's compliance with the Partial Nuclear Test Ban Treaty of 1963. As its name suggests, this treaty forbade either of the Superpowers from testing nuclear weapons under water, in the upper atmosphere, or in outer space for any reason, with the stated goal of slowing an arms race in full swing at the time and reducing excessive nuclear fallout in Earth's atmosphere. Named for the Spanish verb *velar* meaning "to watch over," the *Vela* program consisted of a constellation of twelve satellites placed into near-Earth orbit (just above the Van Allen radiation belts) in pairs throughout the 1960s. Each satellite worked by detecting hard X-rays, neutrons and γ-rays; the detonation of a nuclear weapon would have a signature in all three, and its γ-ray "light curve" (apparent brightness as a function of time) has a distinctive double peak shape. Triangulation via timing across at least two satellites would also allow a crude localization of any given signal.

As the story goes, a team at the Los Alamos National Laboratory led by Ray Klebesadel noticed a bizarre flash of γ-ray photons on 2 July 1967 around 14:19 UTC. The

signal they observed was quite unlike any known nuclear weapons signature, so the team quickly understood that this was likely *not* due to nuclear weapons and did not consider the matter particularly urgent. Nevertheless, it was a curious scientific oddity. As more *Vela*-series satellites were launched with better instrumentation, Klebesadel et al. continued to find similar signals that they could not easily explain in their data. By the early 1970s, sixteen anomalous events were observed in total. Through triangulation, the team was able to estimate the on-sky position where the signals originated from, and ruled out the Earth and the Sun as sources.

A publicly-available report on these anomalous signals was not published until several years later, in 1973 (Klebesadel et al., 1973). In the first unclassified article's abstract the authors provide a summary of the events that still serves perfectly to define them:

> Sixteen short bursts of photons in the energy range 0.2-1.5 MeV have been observed between 1969 July and 1972 July using widely separated spacecraft. Burst durations ranged from less than 0.1 s to $\sim$30 s, and time-integrated flux densities from $\sim 10^{-5}$ erg cm^{-2} to $\sim 2 \times 10^{-4}$ erg cm^{-2} in the energy range given. Significant time structure within bursts was observed. Directional information eliminates the Earth and Sun as sources.

As the events persisted, he also gave them a suitably imposing name that we still use to this day: the *γ-ray burst* (GRB).

What astrophysical maelstrom could be producing these bursts? It is an intellectually galvanizing question. For a long time, the true origin and nature of GRBs was deeply mysterious. Many in the astronomical community argued that the bursts must originate in the Milky Way galaxy, largely because if their source is any farther than that it would have to be inconceivably bright and involve a really rather calamitous and catastrophic energy realease.

The situation remained rather murky until NASA launched the Compton Gamma Ray Observatory (CGRO) in 1991. The CGRO mission had an on-board instrument called the Burst and Transient Source Explorer (BATSE; Fishman et al. 1992, Paciesas et al. 1999) which was the first space borne telescope specially designed and built to discover GRBs. Its detector array consisted of a set of eight sodium iodide (NaI) crystals, one at each corner of the spacecraft, which covered the photon energy range from 20 keV to 2 MeV. With BATSE, GRBs were detected at a rate of about 1 day^{-1}

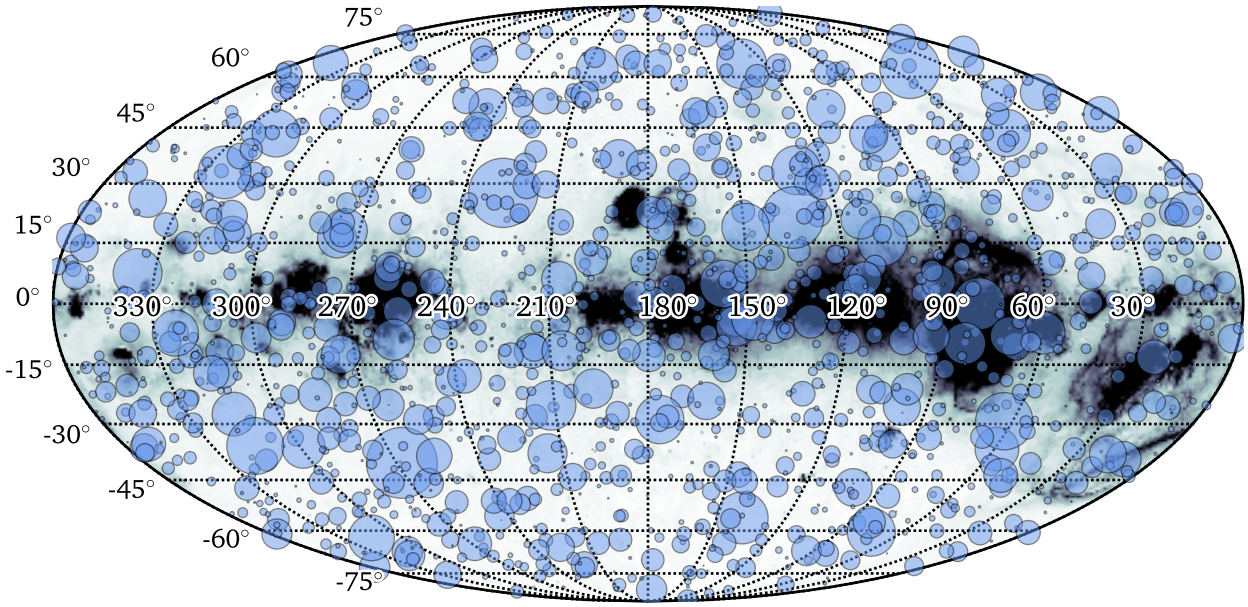

Figure 1.1: All-sky distribution of GRBs from the BATSE 4B catalogue (Paciesas et al., 1999). Sky locations are shown on a Mollweide projection in galactic coordinates and marked with a transparent, light blue circle, whose size is proportional to the positional error radius of the corresponding BATSE burst. A mosaic of the Milky Way in Hα, collected from the WHAM, VTSS, and SHASSA surveys[1], is also shown for reference; note that the BATSE bursts appear to have no preferred direction on the sky, and do not in general track the shape of the Milky Way.

with $\sim$1-10° sky resolution, a reasonable improvement on the capabilities of the *Vela* satellites. BATSE detected over 2000 GRBs during its 9-year mission lifetime and made vast headway toward resolving some long-standing mysteries about these enigmatic monsters in the night.

In particular, two simple observations about the population of BATSE bursts seem to disfavor a Milky Way origin for GRBs:

1. BATSE bursts are isotropically distributed across the entire sky, with no obvious preferred direction in either the north or south galactic hemispheres, and certainly do not trace the shape of the Milky Way (see Fig. 1.1.)

2. A cumulative histogram $N(> S)$ of the number of sources appearing brighter than a given flux threshold S does not result in the trend one would expect of sources uniformly distributed in a Euclidean volume (see Fig. 1.2.)

[1]Accessed from http://lambda.gsfc.nasa.gov/product/foreground/fg_halpha_get.cfm.

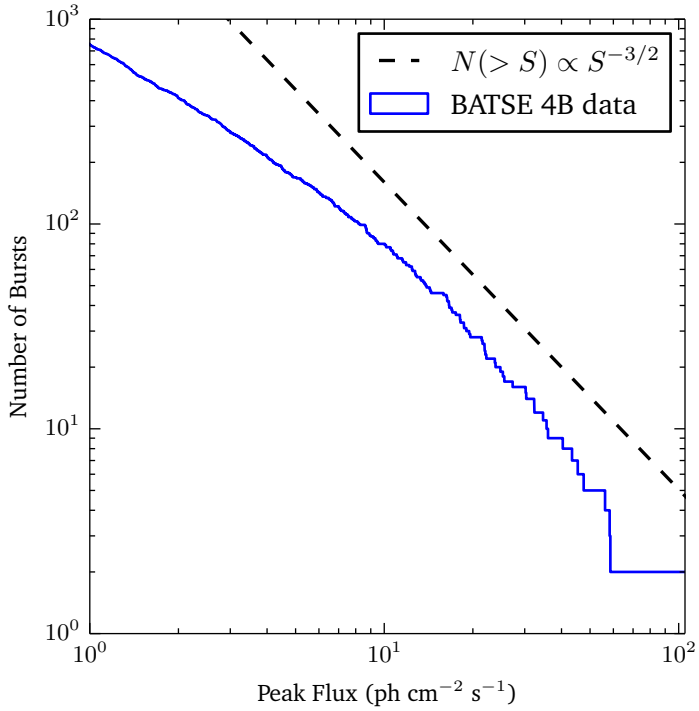

Figure 1.2: Cumulative number of sources versus peak flux (in 256 ms time bins) from the BATSE 4B catalogue (Paciesas et al., 1999). A number distribution with $N(> S) \propto S^{-3/2}$, as would be expected for standard candles uniformly distributed in a Euclidean volume, is shown as a dashed line for reference. Note that BATSE bursts roughly follow this power law trend at the bright end (*i.e.* at peak fluxes $\gtrsim 10$ photons cm^{-2} s^{-1}, with small-number fluctuations near 100 ph cm^{-2} s^{-1}) but deviate substantially at the faint end.

It should be stressed that neither of these features in the BATSE data point to anything conclusive, but taken together they are rather evocative. For instance, the first point by itself is consistent with the sources of GRBs being very nearby (with a distance from Earth smaller than the thickness of the Milky Way disk) or very far away and scattered across other galaxies. If GRBs were originating from the disk of the Milky Way, one would expect only to find them in the galactic plane; if they were originating from the Milky Way halo, one would still expect some anisotropy because the Sun is $\sim$8 kpc from the galactic center (and so more sources might have been seen *e.g.* in the direction of Sagittarius than that of Auriga).

The second point is especially suggestive. Consider a collection of homogeneously distributed standard candles (that is, a class of sources that all have the same intrinsic brightness) in a Euclidean universe. Such objects should be found in numbers that scale with distance as $N(< R) \propto R^3$ because they are found homogeneously throughout

space. Since the observed flux scales with distance as $S \propto R^{-2}$, the observed number distribution in flux ought to scale as $N(> S) \propto S^{-3/2}$. This is roughly consistent with BATSE GRBs at the bright end (peak fluxes $\gtrsim 10$ photons cm^{-2} s^{-1}) but not at all the case at the faint end, where we see a significant roll-off in the number of sources. (Note also that this argument still applies for non-standard candles, as long as the luminosity function is independent of distance.)

What might these observations imply about the astrophysical population of GRBs? At this point in the story we can only stimulate a conversation; they *could* mean a lot of things, and therefore do not tell us anything conclusive. For example, it may be that brighter GRBs could be localized with higher precision, or that a NaI scintillation detector's sensitivity diminishes as bursts become fainter, so that low-flux GRBs are not detected as efficiently. Presumably the brighter bursts tend to occur nearer to the solar system, so this would lead to a selection effect, potentially making both the roll-off in $N(> S)$ at faint fluxes and the apparent isotropy of burst locations an instrumental artifact. But at least one alternative possibility is tantalizing. It is well established that the inverse square law between intrinsic luminosity and apparent flux is not satisfied at distances comparable to the cosmic horizon (*i.e.* out into the Hubble flow) where the universe is not Euclidean and redshift effects become important. If GRBs actually tend to occur at cosmological distances, it would naturally explain why their observed population is both isotropic and evidently non-Euclidean – especially if the astrophysical burst population evolves over cosmic timescales, as the universe ages.

The GRB source distance debate was finally settled with style and panache on 8 May 1997. Four hours after the detection of GRB 970508 (so named for the date on which it occurred, *i.e.* GRB YYMMDD – a convention widely used for GRBs) a rapidly fading optical counterpart was discovered at the same location (the second time such an optical counterpart to a GRB was ever discovered; see Djorgovski et al. 1997, Metzger et al. 1997). Spectral analysis of this optical "afterglow" revealed atomic absorption features associated with doubly ionized magnesium (Mg II) and iron (Fe II) uniformly shifted from their known wavelengths as measured on Earth. Thus, the redshift of this burst was measured as $z = 0.835$. This measurement unequivocally places GRB 970508 at a radial comoving distance $D_M = 2.92$ Gpc – *well* beyond the Milky Way, far out into the Hubble flow.

A sizeable portion of GRBs since then have had associated afterglows in the X-ray,

optical and radio bandpasses (see Piran 2004 for a review of burst afterglows as of 2004). Where spectral analysis is possible, a significant redshift has been ovserved in almost every case – uambiguously demonstrating that γ-ray bursts originate from cosmological distances after all.

How bright is a γ-ray burst "up close?" Based on the measured flux, redshift and duration of an average burst, one finds that the total energy output in the GRB rest frame is on the order of 10^{52} erg when corrected for relativistic beaming (see below for a full explanation). That number is *enormous*. To put it in context, consider that the Sun is far and away the largest energy producer I am at all familiar with, because it has single-handedly sustained all the day-to-day operations on this planet for nearly 5 billion years. How long would we have to wait for the Sun to put out a comparable amount of energy? Assuming a fixed solar luminosity of 3.846×10^{33} erg s^{-1}, this comes out to $\sim 10^{19}$ seconds, or about *ten billion years*. GRBs release about as much energy in a matter of seconds (or minutes) as the Sun will have done over its *entire lifetime*. By a hefty margin, they are therefore the brightest known transient explosions in the universe, without question the best bangs since the Big one.[2]

Not at all a bad find from a group of people originally looking for unsanctioned nuclear weapons tests.

1.1.1 The Long and the Short of It

Now that their existence and relevance to cosmology have been established, let us see what deeper conclusions we can draw about GRBs by considering their observed properties.

We have already seen that typical GRBs range from lasting fractions of a second to several minutes. A good question to ask at this point might therefore be: How often and how extremely do burst durations fluctuate? To answer this question GRB data analysts look at a statistic called T_{90}, defined as the median timescale over which 90% of the total fluence (or integrated flux) was collected. A histogram of T_{90} immediately reveals a slappingly obvious double peak structure, suggesting the superposition of two underlying distributions[3] with a dividing line at roughly 2 seconds (see Fig. 1.3). For

[2] With apologies to Douglas Adams.

[3] It turns out that the *hardness ratio* of a burst – the ratio of high-energy photons to low-energy ones, effectively the "color" in γ-rays – strongly correlates with T_{90}, revealing another double peak distribution

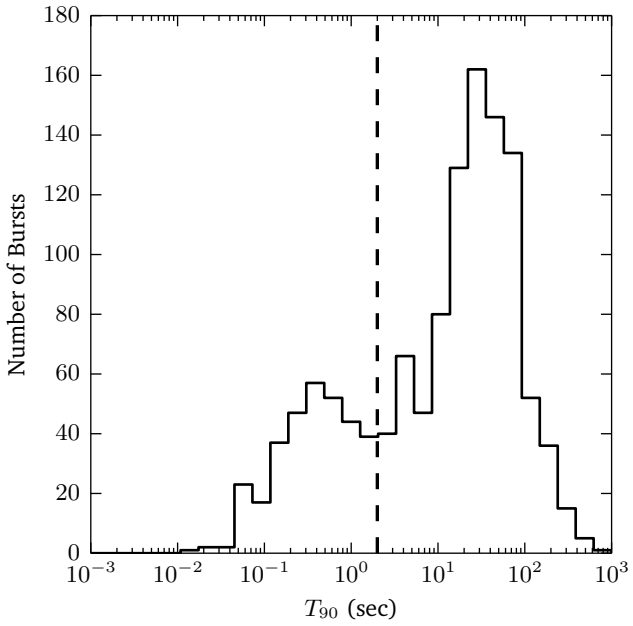

Figure 1.3: Histogram of the durations (in seconds) of GRBs from the BATSE 4B catalogue (Paciesas et al., 1999). The metric for duration shown here is a commonly-used statistic called T_{90}, defined as the median timescale over which 90% of a burst's total energy was collected based on its measured light curve. Note that there is a strong double-peak structure evident here, suggesting two different populations of progenitor sources separated at roughly 2 sec (marked with a dashed line).

this reason the T_{90} statistic turns out to be a good discriminator between two populations of sources. Any burst with $T_{90} \gtrsim 2$ seconds we call *long*. Creatively, any burst with $T_{90} \lesssim 2$ seconds is called *short*.

In both short and long GRBs, the observed high energy prompt spectrum is typically a power law $F_\nu \propto \nu^{-\alpha}$ with spectral index $\alpha \approx 1$, and thus clearly nonthermal in nature. In many cases, there are also wild fluctuations on timescales of $\delta t \sim 10^{-2}$ seconds seen in the γ-ray light curve (Fig. 1.4). Given that no source can vary coherently faster than the light travel time across its full length, this constrains the characteristic radius of the initial GRB explosion to be $R \lesssim c\,\delta t \sim 3000$ kilometers, or about the distance from New York City to Boulder, CO. On astronomical scales this is a *very* small distance, especially considering the intrinsic brightness of observed GRBs. There would need to be an extraordinarily high number of γ-rays contained in a very small source region during the burst, in which case the higher energy γ-rays would quickly and efficiently interact with low energy ones to produce e^+e^- pairs, resulting in a thermal spectrum

in the T_{90} – HR plane. Long GRB photons tend to be soft, while short GRB photons tend to be hard. See Piran (2004); D'Avanzo (2015).

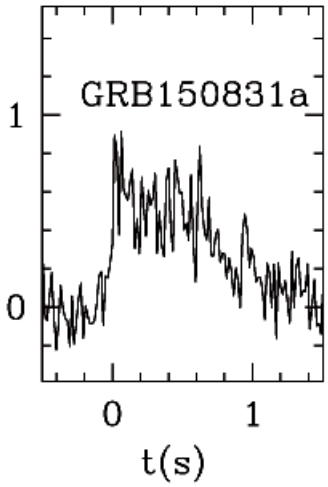

Figure 1.4: The γ-ray light curve of GRB 150831A, as measured by the *Swift* Burst Alert Telescope (BAT). Note rapid variations on timescales of 10^{-2} sec. For this GRB, $T_{90} = 1.15$ sec and $\alpha = 0.81$. (This figure is drawn from the *Swift* online source catalogue, and reproduced with permission from the *Swift* BAT team.)

glaringly inconsistent with what is observed. So how do the γ-rays overcome this and escape?

This puzzling inconsistency is often referred to as the *compactness problem*, and was once used to argue that GRBs must occur in the galactic neighborhood. As we have seen, direct measurements of afterglow redshifts completely rule out this possibility. If the initial GRB explosion expands with a huge bulk Lorentz factor of $\Gamma \equiv (1-v/c)^{-1/2} \sim 100$ or greater – that is, at over 99.995% the speed of light along the line of sight to Earth – then γ-rays are $\Gamma^{4+2\alpha} \sim 10^{12}$ times more likely to escape before succumbing to pair production, for two reasons:

1. Suppose we observe the explosion from an angle $\theta_{\rm obs}$ away from its bulk direction of motion. Due to time dilation, the observed timescale δt is shorter than its rest-frame value by a factor Γ; relativistic aberration reduces it further by a factor $\Gamma^{-1}[1-(v/c)\cos\theta_{\rm obs}]^{-1} \sim \Gamma$ if $\Gamma \gg 1$, $\theta_{\rm obs} \ll 1$, and $\Gamma \gg \theta_{\rm obs}\Gamma$. Thus, in the rest frame of the source, the characteristic emission radius $R \lesssim \Gamma^2 c\,\delta t$ can be a factor Γ^2 larger than it would be if the explosion were non-relativistic, and its cross-sectional area is larger by a factor Γ^4: the photons are *far* less densely packed.

2. These effects also blueshift the energy per individual photon by a factor Γ^2 relative to the observer. Because the measured spectral energy distribution is a power

law, if some fraction $f_{\rm obs} \propto \int_{E_1}^{E_2} E^{-\alpha}\, dE$ of photons are observed to be below the energy threshold necessary for pair production, then $f = \int_{E_1}^{E_2} (E/\Gamma^2)^{-\alpha}\, dE = \Gamma^{2\alpha} f_{\rm obs}$ is the fraction in the rest frame.

Hence, the nonthermal spectra of most GRBs are naturally explained if the explosion producing them is *ultra-relativistic* (Piran, 2004). Furthermore we can still roughly estimate the radius of the initial explosion, with $R \sim 10^{11}$ meters if $\Gamma \sim 100$, or approximately the size of the orbit of Mars.[4] This is roughly the size of the most massive stars, which as we will see is not a coincidence. Note that these inferences are made purely on the basis of relativistic kinematics.

All of this suggests a picture where both long and short GRBs involve a massive explosion, during which the ejecta receive a collossal amount of kinetic energy and are then accelerated to very near the speed of light. Most of this kinetic energy is then efficiently converted to the observed flux in γ-rays through some mechanism, and only a small fraction of kinetic energy is dissipated. This might be the case if, for instance, γ-ray emission is due to the ejecta sweeping up surrounding material that was initially more or less at rest, such as the interstellar medium (the very diffuse gas and dust between stars). We can constrain the ejecta's rest mass, $M_{\rm ej}$, by equating kinetic energy with observed energy and considering the efficiency ϵ_c of energy conversion during this process: $\epsilon_c(\Gamma - 1) M_{\rm ej} c^2 = 4\pi S D^2$. Here S is the total GRB fluence and D its measured distance from Earth. We find that

$$M_{\rm ej} \simeq \epsilon_c^{-1} \left(\frac{S}{10^{-6}\,{\rm erg\ cm^{-2}}}\right) \left(\frac{D}{3000\ {\rm Mpc}}\right)^2 \left(\frac{\Gamma - 1}{100}\right)^{-1} 10^{-5}\ {\rm M}_\odot \tag{1.1}$$

where factors appearing in parentheses are typically of order 1–100 (Rhoads, 1997; Sari et al., 1999; Fruchter et al., 1999). Remarkably, in order to efficiently ($\epsilon_c \sim 1$) convert such extreme kinetic energy into γ-rays, the ejecta must contain a comparitively tiny amount of baryonic matter, amounting to ∼0.0001% the mass of a typical O- or B-type star. This mass constraint is known as the *baryon loading* of the outflow.

A totally separate issue is the sheer scale of energy produced by a typical GRB. Extrapolating from the observed burst fluence and distance to source, the total observed energy $4\pi S D^2$ can exceed 10^{54} erg in some cases if the γ-ray emission is roughly

[4] Lorentz factors of $\Gamma \sim 100$ are consistent with measurements of the opening jet angles of most long GRBs; see *e.g.* Piran (2004).

isotropic. The problem is, this is so much energy on such a short timescale (remember: minutes, at most!) that we struggle to think of known astrophysical mechanisms capable of powering it. The rest mass-energy $E = Mc^2$ of a very massive star is comparable, but this would leave no energy budget for *e.g.* supernovae or gravitational wave emission, and we have already constrained the emitting region to be something like the size of a massive star.

Our way out this time is to consider again the effects of aberration. This effect can be seen in everyday experience (think of the way rain seems to fall harder against a car windshield as you barrel down the highway than it does when you're standing still, even on a windless day) but as usual its relativistic version is much more extreme. In an astronomical context, it results in narrow jet-like structures that often show up in sources such as quasars, within which matter flows outward at speeds $v/c \sim 1$. The GRB case is even more extreme. It is now understood that radiation from GRBs is beamed in narrow, roughly conical jets of angular width $\theta_{\rm jet} \sim 1$–10° in long GRBs (Frail et al., 2001; Lipunov et al., 2007) and $\theta_{\rm jet} \sim 10$–20° in short GRBs (Berger, 2014a; D'Avanzo, 2015). This adjusts the energy scale (and the ejecta mass limit, Eq. 1.1) by a beaming factor $f_b^{-1} \equiv 1 - \cos\theta_{\rm jet} \approx \theta_{\rm jet}^2/4\pi \sim 10^{-2}$ (see *e.g.* Cenko et al. 2015), making it consistent with the energy budget of most stellar core collapse or compact object merger scenarios. Relativistic aberration also provides a way to confirm the inferences made about GRBs in this section so far, and sheds some light on the central engine driving these monstrous explosions. Robust observational support, both for beaming and for ultra-relativistic bulk velocities in the initial explosion, is evident in the light curves and optical spectra of GRB afterglows (see Piran 2004 for a review).

A Long Time Ago in a Galaxy Far, Far Away....

Given the similarities between the two flavors of GRB, it is natural to wonder now about the nature of their differences. How can two presumably distinct progenitor populations produce such similar physics? To get a sense of this, we look for any hint of a relationship between GRBs and other astrophysical phenomena that are better understood. Consider long GRBs. These tend to be discovered in very active star forming regions, and in galaxies where the rate of star formation is still quite high (*e.g.* young spirals at high redshift; see Christensen et al. 2004; Jakobsson et al. 2011). The

cosmic rate density of long GRBs is also found to closely track the star formation rate history of the universe, well into redshifts $z \gtrsim 2$ (Wanderman & Piran, 2010).

In April 1998, the positional error box of GRB 980425 was found to contain an optically variable source later confirmed to be a Type Ic supernova, designated SN 1998bw (Woosley et al., 1999). This marked the first time a supernova was found in coincidence with a GRB, and was interpreted as a strong hint at a deeper connection. Later observations found that the optical light curves of several long GRB afterglows were well-described as the sum of a power law decay (as is typical of "classical" GRB afterglows) and a luminous Type Ib or Type Ic supernova, whose optical spectra are distinguished by a lack of H and Si I absorption lines. The connection between supernovae and long-duration GRBs was rendered unequivocal in March 2003 when SN 2003dh was identified with the afterglow of GRB 030329 (Berger et al., 2003), firmly establishing that these two classes of optical transient are closely associated. (It has since become clear that, while most GRBs with associated afterglows are discovered at high redshifts $z \gtrsim 1$, the supernova counterpart is bright enough to be visible only for the nearest bursts.)

It is now widely accepted that long GRBs generally occur very shortly after a new black hole is formed, explaining their association with stripped core-collapse supernovae (Piran, 2004). In this scenario, the central engine is powered by infalling material that *used* to compose the inner guts of a massive, metal-rich star. The progenitor star had shed its outermost hydrogen layer near the end of its life, and when the star's core finally gets hot enough to fuse the radioactive isotope ^{56}Ni, a catastrophic instability arises. The star can no longer sustain its weight with nuclear fusion, so its core begins to collapse. The stellar core implodes on itself, while the outer layers of the star rebound and expel outward, producing a supernova. In some cases, not even the degeneracy pressure of atomic neutrons could prevent complete core collapse.[5] In this scenario a black hole is formed, anywhere from 5–50 $M_\odot$ in girth. But the newborn monster is ravenous, devouring material left over from the collapse; this material radiates on its way in, pulverizing the surrounding environment with a psychotic amount of energy and powering a long-duration GRB when the ejecta quickly accelerate to $\Gamma \gtrsim 100$ in a narrow jet along the axis of rotation. The fact that the star shed its outer-

[5] *Neutron degeneracy pressure* will halt the collapse if there is not enough mass to overcome it, resulting in the *second* most compact object known to science: a neutron star.

most hydrogen layer before dying is crucial in producing a clean burst consistent with the spectra and energetics observed in most cases (see *e.g.* Cenko et al. 2013, 2015), because in this case the ejecta mass can be sufficiently small that most of its kinetic energy is converted to photons (Eq. 1.1).

Plenty of open questions still remain: Do we really understand beaming and energetics in long GRBs? What is the true cosmic event rate of long GRBs, unbiased by considering just the ones we can see from Earth? Does this rate density scale with the cosmic star formation history of the universe, even at high redshifts? Is the true astrophysical population of GRBs dominated by unseen darker bursts or "dirty fireballs" where the ejecta entrain more baryonic mass? These questions will be revisited vengeance in chapter 4.

Short GRBs, on the other hand, remain a bit more mysterious. Much of this owes to the difficulty in detecting them due to their short duration, but they also seem to be systematically less luminous than their long-duration counterparts, and tend to be discovered at smaller redshifts (D'Avanzo, 2015). There are also other clues. Perhaps the most noteworthy detail is that short GRBs lack associated supernovae and are far and away more likely to be found out in the halos of old elliptical galaxies, where star formation has long since ground to a halt (Berger, 2014a). This is the same sort of environment where old, massive binary star systems go after they die: supernovae explode asymmetrically, producing a "kick" that propels the supernova remnant (either a neutron star or black hole) out into the galactic halo. More than half the massive stars in our own galaxy are in binaries. It is known that some of these binaries survive through the supernovae of both stars, leaving behind compact binary systems with neutron stars and black holes. These compact binaries will steadily radiate orbital energy in the form of gravitational waves (see chapter 2), slowly spiraling toward each other over the course of hundreds of millions of years. Short GRBs may well be produced in the very last few seconds of this elegant dance of death, when the two compact objects – at least one of which must be a neutron star – eventually merge (Metzger & Berger, 2012).

At the time of writing, both binary neutron star (*e.g.* Hulse & Taylor 1975) and binary black hole (Abbott et al., 2016c) systems have been observed in nature, and glimpses at an association may have been observed (Connaughton et al., 2016a). The connection between compact binary mergers and short GRBs is a hypothesis we hope

to directly test with the LIGO experiment.

1.2 Listening to the Universe with Gravitational Waves

There is another chronicle from the history dossier of science that will help us understand the full astronomical impact of compact binary inspirals, and this time, it involves quite the cast of characters.

A decade after successfully explaining the bizarre nature of motion near the speed of light, Albert Einstein published his General Theory of Relativity in November 1915. The General Theory was born out of a need to unite the *Special* Theory of Relativity with a model of gravitation, and was in its own way the resolution to a mystery 300 years in the making. For three centuries, science was content with Isaac Newton's worldview: that the universe has a clockwork structure, and that gravity is a *force* causing attraction between any two masses. This way of thinking is not exactly wrong, as far as it goes;[6] but it is incomplete, and there are a number of problems. Chief among these is the fact that it violates special relativity's explicit decree that no information may pass faster than the speed of light, which is better thought of in this context as the "speed of causality." That the speed of light does not respect classical laws of relative velocity is an experimental fact, first established by Albert A. Michelson and Edward W. Morley at Case Western Reserve University in Cleveland, Ohio, in the spring and summer of 1887.

Einstein's diplomatic solution to this conflict between new and old was to completely rewrite our understanding of the nature of gravity. In the framework of General Relativity, gravitation is no longer regarded as a force between two or more bodies. Instead, it is now thought of as *curvature in the fabric of reality itself*: space and time are no longer separate things, but form a unified geometry, which is then bent and distorted by the presence of masses. The curvature of "spacetime" then communicates to massive objects how they should move, and this entire back-and-forth interaction takes place at the speed of light, a mere 299,792,458 meters per second.

As is discussed in more detail in many other works (*e.g.* Wald 1984; Hartle 2003), gravity as the curvature of spacetime explains much about the universe that simply does

[6] And certainly not naive: Newtonian mechanics explains *e.g.* the orbits of planets and the structure of the Sun, and we needed nothing more than Newton's equations to get to the Moon.

not fit into the clockwork Newtonian point of view. For instance, the planet Mercury precesses by 43 seconds of arc per century in its orbit around the Sun. Observations of background stars near the Sun, taken during a solar eclipse when most of the foreground light is blocked out, reveal that beams of light are bent as they pass around the Sun by an amount that agrees to high precision with General Relativity. Later tests, such as the Pound-Rebka experiment, verified that photons are redshifted as they lose energy climbing out of a gravitational well. In 2004, Gravity Probe B tested the prediction that Earth will drag spacetime along as it rotates; once again, General Relativity was upheld. It is even true that time passes slightly differently in low-Earth orbit than it does on the ground – a fact you might have relied on to get to work on time this morning, if you used your GPS to navigate traffic hazards.

But there is one prediction from Einstein that, prior to September 2015, had never been directly tested. If spacetime can be bent and distorted then it should also be prone to wiggles: any object whose quadrupole moment changes over time – that is, any change that is not spherically or cylindrically symmetric (see chapter 2) – should make ripples in the fabric of spacetime (see *e.g.* Creighton & Anderson 2011; Maggiore 2008). We call this phenomenon *gravitational waves* (GWs), and fully expect that GWs are produced by a wide variety of astrophysical phenomena. The 1993 Nobel Prize in Physics was even awarded to Russell A. Hulse and Joseph H. Taylor, Jr. (Hulse & Taylor, 1975; Taylor & Weisberg, 1982), whose observations of the binary pulsar system PSR J1913+16 show that the neutron stars in this system orbit at a rate whose decay over time is consistent with General Relativity, which explains the orbital decay in terms of the neutron stars gradually radiating away their energy in the form of gravitational waves as they spiral toward one another. (If this sounds familiar, recall it is the same mechanism thought to power short GRBs, when the neutron stars eventually merge.)

Even the most extreme astrophysical sources – *e.g.* two colliding stellar mass black holes at a distance of 400 Mpc – would only stretch and constrict the local spacetime near Earth by an amount on the order of 10^{-21} meters, or about one millionth the width of a proton. But there are bold new astronomical ventures that give us hope. The Laser Interferometer Gravitational-wave Observatory (LIGO: LIGO Scientific Collaboration et al. 2015) experiment is a set of kilometer-scale interferometers, one in Livingston, LA, and the other in Hanford, WA, designed to detect very weak gravitational waves by the effect they have on masses suspended very close to free fall (see chapter 2 and the

discussion therein). As will be discussed in chapters 2, 3, 5, 6, and 7, the LIGO Scientific Collaboration has announced the first confirmed direct detection of gravitational waves from the inspiral and merger of two stellar mass black holes. This is a brand new era of gravitational wave astronomy. It is analogous to revealing a new sense on the universe ("hearing" as well as "seeing") that will bring with it a wealth of new astrophysical information, including much that we cannot anticipate.

1.3 Goals of This Thesis

The primary goal of this work is to outline in explicit detail all the ways in which I helped. I begin in chapter 2 by briefly discussing the effect gravitational waves should have on freely falling test masses, then outlining the experimental design of the LIGO project. I will both motivate and close this section up with a discussion of "electromagnetic follow-up," the collaborative effort to train high-energy, optical and radio telescopes in space and around the world on gravitational wave transients discovered by LIGO. (Aspects of the discussion presented here were the result of a three month stay at LIGO Livingston Observatory in autumn 2015.) I will then close chapter 2 with an overview of two high-energy space satellites (the *Swift* Gamma-Ray Burst Explorer and the *Fermi* Gamma-Ray Space Telescope) and an optical facility (the Zwicky Transient Facility, ZTF) on Mt. Palomar in San Diego County, California.

Chapter 3 will contain work I have contributed to prepare for electromagnetic follow-up of LIGO detection candidates. Chapter 3 details a software pipeline designed to rapidly identify coincidences in time and sky location between LIGO detection candidates and GRBs discovered by *Swift* and *Fermi*, within a minute or so of the events having been observed. This will be very useful if a source is observed concurrently in gravitational waves and high-energy photons, because the search is more sensitive than a blind all-sky search for GWs alone, and many in the astronomical community will have enormous interest in locating the afterglows of these joint detections.

In chapter 4 I relay a series of observations made with the Palomar Transient Factory, a precursor to ZTF. With this facility, our team successfully made the first discovery of a long GRB (designated GRB 140226A) by its optical afterglow emission (the optical transient iPTF14yb). In the same chapter I then constrain the cosmic event rate of

on- and off-axis optical afterglows, both of GRBs and of failed GRB "dirty fireballs" in which the outgoing jet may entrain more baryonic mass (thus reducing the peak energy of the explosion). I also discuss the implications of these rate constraints for planned orphaned afterglow searches with wide-field optical facilities such as ZTF.

In chapters 5 and 6, I discuss work that was done on-site at the LIGO Livingston Observatory in Livingston Parish, LA, as part of the LIGO Scientific Collaboration's On-Site Fellows program. In particular, in chapter 5 I report on my role in the first direct observation of gravitational waves from a binary black hole merger and on the performance of various services during the first Observing Run of Advanced LIGO. Then, in chapter 6, I report on the high energy, optical and radio follow-up campaign surrounding the gravitational wave transient GW150914.

In the 7th and final chapter, I conclude with a further re-contextualization of this work, an optimistic look to the future of time domain multimessenger surveys of the sky. There may be little reason to fear the dark – after all, we could never see the stars without it.

Chapter 2

Multimessenger Astronomy in the Advanced LIGO Era[1]

"Sir Isaac Newton, renowned inventor of the milled-edge coin and the catflap!"
"The what?" said Richard.
"The catflap! A device of the utmost cunning, perspicuity and invention. It is a door within a door, you see, a ..."
"Yes," said Richard, "there was also the small matter of gravity."

Douglas Adams, *Dirk Gently's Holistic Detective Agency*

On 14 September 2015, at 9:50:45 UT, the two facilities of the Laser Interferometer Gravitational-wave Observatory (LIGO) in Livingston, LA and Hanford, WA separately observed a transient gravitational wave (GW) signal within $6.9^{+0.5}_{-0.4}$ ms of one another (Abbott et al., 2016c).[2] Designated GW150914 based on its calendar date, the observed signal is consistent with a binary black hole merger of initial component masses $m_1 = 36^{+5}_{-4}\,M_\odot$, $m_2 = 29^{+4}_{-4}\,M_\odot$, occurring at a redshift $z = 0.09^{+0.03}_{-0.04}$ (and a luminosity distance $D_L = 410^{+160}_{-180}$ Mpc, all as measured in the local Earth-based observer frame). After the observed GW signal swept up in frequency from 35 to 350 Hz over a period of roughly 0.2 seconds, the black holes violently merged, leaving behind a final black hole of mass $m = 62^{+4}_{-4}\,M_\odot$ and dimensionless spin $a = c|\mathbf{S}|/Gm^2 = 0.67^{+0.05}_{-0.07}$ (where S is the spin angular momentum), radiating $3.0^{+0.5}_{-0.5}\,M_\odot c^2$ of energy in the process (The LIGO Scientific Collaboration & the Virgo Collaboration, 2016b). This consitutes the

[1] Portions of this chapter were completed as part of the LSC Fellows Program at the LIGO Livingston Observatory in autumn 2015, concurrent with the first Observing Run of Advanced LIGO.

[2] For each of these parameters we report the median value and range of the 90% credible interval.

first direct detection of GWs and the first direct evidence for binary black hole systems occurring in nature. The final black hole is several times more massive than any other known stellar-mass black hole (Abbott et al., 2016a).

As discussed in chapter 1, prior to GW150914, the best evidence supporting gravitational waves as a real astrophysical phenomenon came from measurements of the decaying orbital period of the binary pulsar PSR J1913+16 (Hulse & Taylor, 1975; Taylor & Weisberg, 1982). The orbital decay of PSR J1913+16 is consistent with predictions of General Relativity to very high precision, and demonstrates without ambiguity that energy is being radiated from the system. While General Relativity explains this energy loss as having been carried away in the form of gravitational waves, the frequency of any radiation from this binary pulsar is too low for LIGO to observe directly (see below). Instead, GWs from compact binary inspirals will enter LIGO's sensitive frequency band only in the last moments before merger, as the two bodies come within tens or hundreds of kilometers of one another and attain orbital velocities that are an appreciable fraction of the speed of light.

The source of GW150914 was a binary black hole merger, which is not expected to contain or interact with very much matter.[3] However, in compact binary systems with at least one neutron star, we expect accretion disks to form on dynamical timescales of $\sim$0.01–0.1 sec as the two bodies coalesce or as one of the neutron stars is tidally disrupted (Metzger & Berger, 2012). As alluded to in chapter 1, this accretion is thought to power short-duration γ-ray bursts (GRBs) within a few seconds of merger, and X-ray ($\sim$minutes), optical ($\sim$hours to days) and radio ($\sim$months to years) afterglows on longer timescales (D'Avanzo, 2015). Furthermore, r-process nucleosynthesis in ejecta from the accretion disk may power so-called "kilonova" emission within a day or so of merger (Berger, 2014a). Clearly, although the accretion flow during the merger process is very short-lived, it is also remarkably explosive: the electromagnetic transients described here can attain isotropic equivalent luminosities $L_{\rm iso} \sim 10^{50}$ erg s^{-1}, approaching the brightest known classes of transient explosion in the universe.

In this chapter, we outline the methodology of the LIGO experiment and describe some of its scientific goals in the context of *multimessenger astronomy:* the effort to

[3]However, the *Fermi* Gamma-ray Burst Monitor (GBM) recorded a very weak γ-ray transient some 0.4 sec after the measured coalescence time of GW150914 (Connaughton et al., 2016a). Although this transient is of very low statistical significance, its temporal coincidence with GW150914 is perhaps a bit suggestive. We will discuss this transient in more detail in chapters 5 and 6.

study several astrophysical sources using gravitational wave, electromagnetic, and neutrino emission. We begin by using General Relativity to describe the effect that a gravitational plane wave will have as it passes a set of freely falling test masses, then detail the interferometric LIGO and Virgo detectors, including their major sources of noise. We also give a brief overview of common matched filter data analysis techniques used to search for signals originating from compact binary coalescence (CBC) in the LIGO data stream. Next, we outline a method for rapid sky localization of CBC signal candidates, and address the technical obstacles facing our effort to inform high-energy, optical, and radio astronomers in nearly real time when significant signal candidates are found. Finally, we conclude this chapter with a description of two space-based γ-ray observatories (the *Swift* and *Fermi* γ-ray space telescopes) and an optical observatory (the Palomar Transient Factory in California) which are used to perform follow-up observations of LIGO source candidates and to observe GRBs and their afterglows.

2.1 Gravitational Waves

In order to properly understand how interferometric gravitational wave detectors work, we must first understand some basic properties of how GWs propagate. In particular, we need to see how a passing gravitational wave will affect motion on, say, a ring of freely falling test masses. A simple toy model will help to illustrate.[4]

Recall from General Relativity that gravitation is understood as an effect of curvature in the surrounding spacetime. Mathematically, the metric tensor g_{ab} represents the differential separation between two points according to the line element

$$ds^2 = g_{ab}\left(\mathrm{d}x^a \otimes \mathrm{d}x^b\right) \tag{2.1}$$

(where $\mathrm{d}x^a$ is a 1-form on the spacetime manifold)[5]. Inertial observers move along geodesics of the spacetime; the mutual acceleration a^a felt by a pair of freely falling

[4]To save time, in this chapter we presume a familiarity with General Relativity. The relevant aspects of Einstein's theory will be reviewed only very briefly. A more careful treatment may be found in several places, *e.g.* Wald (1984), Hartle (2003), Creighton & Anderson (2011), and Maggiore (2008). Each of these works has influenced the presentation of this section.

[5]Throughout this section, we use the Einstein summation convention where repeated upper and lower indices denote a sum (*e.g.* ${g_\alpha}^\alpha = \sum_{\alpha=0}^{3} {g_\alpha}^\alpha$). We also use abstract index notation, where Greek indices $(\alpha, \beta, \ldots)$ denote the components of a tensor in some coordinate system, mid-Latin indices $(i, j, \ldots)$ denote only the spatial components, and early Latin indices $(a, b, \ldots)$ represent the rank of the general tensor.

test particles due to local curvature is described by the geodesic deviation equation

$$a^a = -R_{cbd}{}^a X^b T^c T^d, \tag{2.2}$$

where X^a is a separation vector, T^a the local tangent to the geodesic, and the Riemann tensor

$$R_{abc}{}^d = \partial_b \Gamma^d_{ac} - \partial_a \Gamma^d_{bc} + \Gamma^e_{ca}\Gamma^d_{be} - \Gamma^e_{cb}\Gamma^d_{ae} \tag{2.3}$$

encodes everything physical about spacetime curvature. It is constructed from derivatives of g_{ab} in the form of the Christoffel symbols

$$\Gamma^c_{ab} = \frac{1}{2} g^{cd} \left(\partial_a g_{bd} + \partial_b g_{ad} - \partial_d g_{ab}\right) \tag{2.4}$$

where ∂_a is used here as an ordinary partial derivative.

Consider the flat Minkowski spacetime of *Special* Relativity, η_{ab}, written in standard Cartesian coordinates (t, x, y, z) where the metric components are $\eta_{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1)$. Imagine a perturbation $h_{\mu\nu}$ to the Minkowski metric that is small when expressed in these coordinates, so that the full metric is

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \tag{2.5}$$

It is now straightforward, if tedious, to linearize Einstein's equations in

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} \tag{2.6}$$

where $R_{\mu\nu} = R_{\mu\alpha\nu}{}^\beta$ and $R = R_\alpha{}^\alpha$. With no small amount of effort it can also be shown that several gauge freedoms exist for $h_{\mu\nu}$, and the physics of the situation is most easily understood under the following gauge conditions:

$$h = h_\mu{}^\mu = 0 \qquad \text{(traceless)} \tag{2.7}$$

$$h_{\mu 0} = 0 \qquad \text{(purely spatial)} \tag{2.8}$$

$$\partial^\nu h_{\mu\nu} = 0 \qquad \text{(transverse)}. \tag{2.9}$$

Because of properties (2.7)–(2.9), this is often referred to as the *transverse-traceless* or TT gauge. Furthermore, since the metric components $h_{\mu\nu} = h_{\nu\mu}$, the TT gauge reveals that only two nonzero components remain. For reasons that will soon become obvious, we refer to these as $h_+(x^\mu)$ and $h_\times(x^\mu)$.

It can also be shown that the vacuum Einstein field equation reduces to a wave equation for $h_{\mu\nu}$ in the TT gauge,

$$\Box h_{\mu\nu} = \partial_\alpha \partial^\alpha h_{\mu\nu} = 0. \tag{2.10}$$

Its simplest solution is a plane wave propagating in, say, the $+z$-direction – hence the name "gravitational wave." In the TT gauge, this plane wave solution looks like

$$[h_{\mu\nu}] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.11}$$

where $h_+ = h_+(t - z/c)$ and $h_\times = h_\times(t - z/c)$ are each functions only of the retarded time.

These plane waves will propagate at the speed of light, and have two linearly independent polarization states described by h_+ and $h_\times$ respectively. What effect does a plane wave have as it passes? Consider again the relative acceleration between two nearby, freely falling particles, Eq. (2.2). Let $\vec{\xi} = \xi(\sin\theta\cos\varphi\,\hat{\mathbf{x}} + \sin\theta\sin\varphi\,\hat{\mathbf{y}} + \cos\theta\,\hat{\mathbf{z}})$ be the (purely spatial) separation vector of one of these particles relative to the other, and note that the relative acceleration is $a_i = -R_{0i0j}\xi^j$, so

$$a_1 = \frac{1}{2}\xi\sin\theta\left(\ddot{h}_+\cos\varphi + \ddot{h}_\times\sin\varphi\right) \tag{2.12}$$

$$a_2 = \frac{1}{2}\xi\sin\theta\left(-\ddot{h}_+\sin\varphi + \ddot{h}_\times\cos\varphi\right) \tag{2.13}$$

$$a_3 = 0. \tag{2.14}$$

The first thing to notice is that the acceleration is purely transverse, as one might expect. Our two particles will only be accelerated by the passing gravitational wave in the plane perpendicular to the wave's direction of travel. Second, the particles will be affected most when they are initially separated perpendicular to the passing wave ($\theta = \pi/2$) and they will be completely unaffected when they are separated in the direction of the passing wave ($\theta = 0, \pi$). Lastly, the angle φ rotates between the $+$ and $\times$ polarizations; the associated lines of force are visualized in Fig. 2.1, and we can start to understand why we gave the polarization states these monickers.

To visualize what happens to our freely falling test particles over time as the wave passes, note that the magnitude of the separation vector $\vec{\xi}$ gives a measure of the

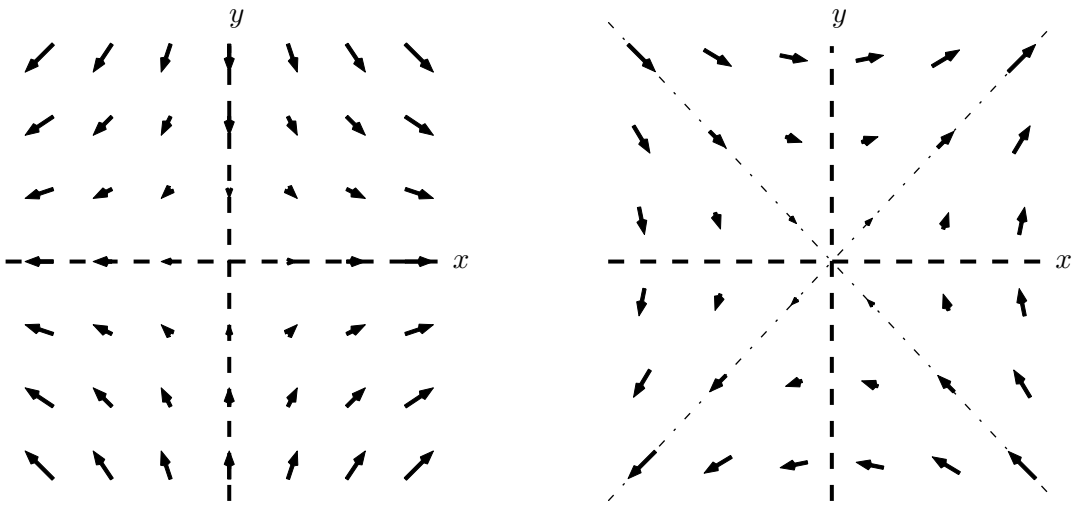

(a) Plus polarization (b) Cross polarization

Figure 2.1: Lines-of-force diagram for (a) a purely plus- and (b) a purely cross-polarized gravitational wave in the transverse plane at the start of a wave cycle; see Eqs. (2.12) and (2.13). The names "plus" and "cross" arise from the broad shapes of these lines of force.

distance between these particles at any given time. The component of acceleration in that direction is

$$\begin{aligned}\frac{d^2\xi}{dt^2} &= \mathbf{a}\cdot\frac{\vec{\xi}}{\xi}\\ &= a_1\sin\theta\cos\varphi + a_2\sin\theta\sin\varphi\\ &= \frac{1}{2}\xi\sin^2\theta\left(\ddot{h}_+\cos 2\varphi + \ddot{h}_\times\sin 2\varphi\right).\end{aligned}$$

If the particles were initially at rest with respect to one another, we can integrate directly and see that

$$\xi(t) = \xi(0)\left[1 + \frac{1}{2}\sin^2\theta\left(h_+\cos 2\varphi + h_\times\sin 2\varphi\right)\right]. \tag{2.15}$$

As expected, we see that when $\theta = 0$ or $\theta = \pi$ the separation vector does not change over time, and its change is most drastic when $\theta = \pi/2$. Furthermore, the rotation between h_+ and $h_\times$ depends only on 2φ, so these are quadrupolar fields – an important fact to consider when we ask: what kinds of things actually produce gravitational waves?

A quarter-cycle of a plane gravitational wave is illustrated in Fig. 2.2, where we imagine its effect on a ring of test masses. In a little while we will see how the LIGO

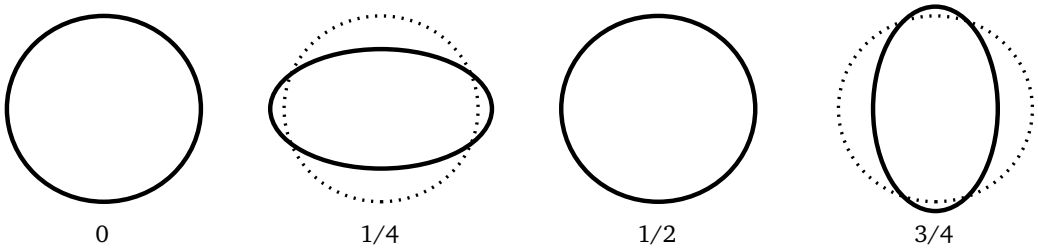

(a) Plus polarization

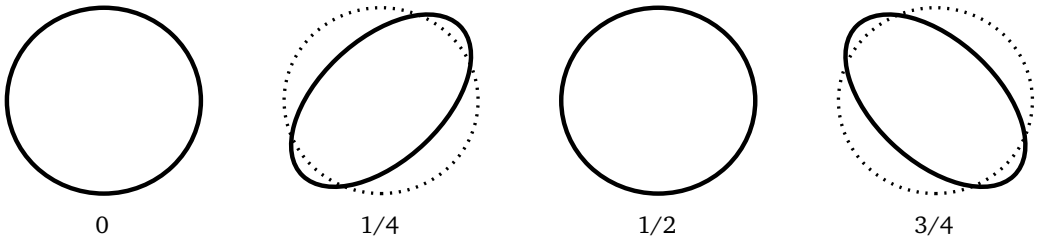

(b) Cross polarization

Figure 2.2: Distortion of a ring of test masses lying in the plane perpendicular to (a) a purely plus- and (b) a purely cross-polarized gravitational wave. In each case the initially circular ring is distorted into an ellipse whose orientation depends on the polarization of the wave. Over a full cycle of the passing wave, the ring will first distort to an ellipse (1/4 cycle), return to a circular shape (1/2 cycle), distort in the other direction (3/4 cycle), and finally return again to its initial shape.

experiment uses a set of suspended test masses to measure the effect of passing gravitational waves in the frequency range from $\sim$10–10^4 Hz, but first, we will examine an important astrophysical source of gravitational waves in this bandpass: the inspiral and merger of ultra-compact bodies.

2.1.1 Astrophysical Sources: Compact Binary Coalescence

To lend some clarity to this discussion, we will briefly take a step back and think more generally about metric perturbations in an arbitrary gauge. We write $\overline{h}_{\mu\nu}$ to distinguish this from the specific TT gauge.[6]

From the Einstein field equations and the conservation of energy,

$$\partial_\mu \overline{h}^{\mu\nu} = 0, \tag{2.16}$$

[6]Strictly speaking, the object $\overline{h}_{\mu\nu} = h_{\mu\nu} - (1/2)h\eta_{\mu\nu}$ is what we call the "trace-reversed" metric perturbation; in the TT gauge we have $\overline{h}_{\mu\nu} = h_{\mu\nu}$.

we see that $\overline{h}_{\mu\nu}$ satisfies a wave equation

$$\Box\overline{h}^{\mu\nu} = -\frac{16\pi G}{c^4}T^{\mu\nu} \tag{2.17}$$

with the stress-energy tensor $T^{\mu\nu}$ acting as a source term. Its general solution

$$\overline{h}^{\mu\nu}(t,\mathbf{x}) = \frac{4G}{c^4}\int \frac{T^{\mu\nu}(t-|\mathbf{x}-\mathbf{x}'|/c,\mathbf{x}')}{|\mathbf{x}-\mathbf{x}'|}\,d^3\mathbf{x}' \tag{2.18}$$

can be written down using the familiar Green's function in Minkowski spacetime. We seek an approximate solution in the radiation zone, where the field point r is much farther away than the GW wavelength λ is long, and λ itself is longer than the characteristic size R of the emitting region. (This is natural for the astrophysical sources LIGO is sensitive to, where r is *at least* on the order of parsecs and $\lambda \sim 10^4$–10^7 meters, constraining the sources themselves to be rather compact.) Under these conditions, $|\mathbf{x}-\mathbf{x}'| \approx r$ is roughly constant over the size of the source. Furthermore, if the source is slowly moving radially then $t-|\mathbf{x}-\mathbf{x}'|/c \approx t-r/c$ and the metric perturbation is

$$\overline{h}^{\mu\nu}(t,\mathbf{x}) = \frac{4G}{c^4 r}\int T^{\mu\nu}(t-r/c,\mathbf{x}')\,d^3\mathbf{x}'. \tag{2.19}$$

Ultimately, what we most want to understand is the behavior of the spatial components of $\overline{h}_{\mu\nu}$, since all the pure physics is contained in the TT gauge. To that end, we can use the conservation of mass-energy ($\partial_\mu T^{\mu\nu} = 0$) to show that

$$\overline{h}^{ij}(t,\mathbf{x}) = \frac{2G}{c^4 r}\frac{\partial^2}{\partial t^2}\int x'^i x'^j T^{00}(t-r/c,\mathbf{x}')\,d^3\mathbf{x}'. \tag{2.20}$$

This is the second time derivative of the familiar quadrupole tensor, which is defined with spatial components

$$I^{ij}(t) = \int x'^i x'^j T^{00}(t-r/c,\mathbf{x}')\,d^3\mathbf{x}', \tag{2.21}$$

in which case we see that

$$\overline{h}^{ij} = \frac{2G}{c^4 r}\ddot{I}^{ij}(t-r/c). \tag{2.22}$$

Finally, let $P_{ij} = \delta_{ij} - n_i n_j$ be a projection operator, where $n_i = x_i/r$ is a unit vector in the gravitational wave's direction of travel. If

$$\mathcal{I}_{ij} = P_{ik}I^{kl}P_{lj} - \frac{1}{2}P_{ij}P_{kl}I^{kl} \tag{2.23}$$

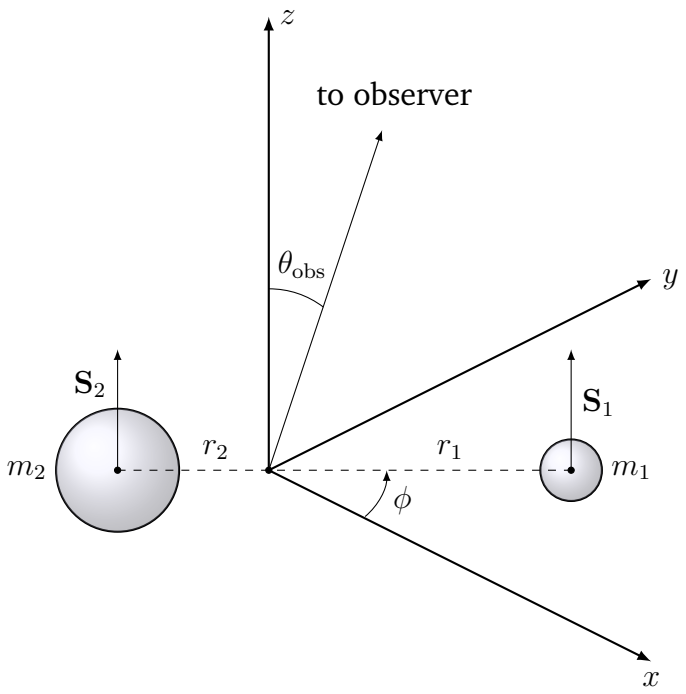

Figure 2.3: Diagram of a binary system orbiting in the x–y plane, with the origin located at the barycenter of the system. The masses (m_1, m_2) and orbital separation (r_1, r_2) are drawn to scale with $m_2 = 2m_1$. The orbital phase (ϕ) and relative viewing angle of a distant observer ($\theta_{\rm obs}$) are also visualized. For simplicity, objects in the binary have aligned spins ($\mathbf{S}_1$, $\mathbf{S}_2$), so the total angular momentum is oriented along the z-axis.

is our quadrupole tensor made traceless and projected into the transverse plane, then in the TT gauge we have

$$h_{ij}(t, \mathbf{x}) = \frac{2G}{c^4 r} \ddot{\mathcal{I}}_{ij}(t - r/c). \tag{2.24}$$

In this way, we see that gravitational waves are produced in the radiation zone by any source whose quadrupole moment changes non-linearly over time; that is, by accelerating masses (or massive fluids).

Now imagine a binary system orbiting in the x–y plane (Fig. 2.3). The system consists of two point masses m_1 and m_2, orbiting at a distance of r_1 and r_2 from the origin, respectively. The orbital separation between these two masses is $a = r_1 + r_2$ and the orbital angular momentum vector is oriented along the z axis. The total mass $M = m_1 + m_2$ and reduced mass $\mu = m_1 m_2 / M$, and we note that $r_1 = a m_2 / M$ and $r_2 = a m_1 / M$. The quadrupole moment of this system is trace-free, and its non-vanishing

components are

$$I_{11} = \frac{1}{2}\mu a^2 \left(1 + \cos 2\phi\right) \tag{2.25}$$

$$I_{22} = \frac{1}{2}\mu a^2 \left(1 - \cos 2\phi\right) \tag{2.26}$$

$$I_{12} = I_{21} = \frac{1}{2}\mu a^2 \sin 2\phi \tag{2.27}$$

where the orbital phase $\phi = \omega t$ uniformly increases with time at a rate set by the orbital angular frequency ω. To an observer a large distance r away on the z axis, we take two time derivatives and see that the metric perturbation (which is already both transverse and trace-free, *i.e.* $I_{ij} = \mathcal{I}_{ij}$) is

$$[h_{ij}] = -\frac{4G\mu a^2\omega^2}{c^4 r}\begin{bmatrix} \cos 2\phi & \sin 2\phi & 0 \\ \sin 2\phi & -\cos 2\phi & 0 \\ 0 & 0 & 0 \end{bmatrix}. \tag{2.28}$$

The two GW polarizations are therefore

$$h_+(t) = -\frac{4G\mu a^2\omega^2}{c^4 r}\cos 2\phi(t) \tag{2.29}$$

$$h_\times(t) = -\frac{4G\mu a^2\omega^2}{c^4 r}\sin 2\phi(t). \tag{2.30}$$

Notice that gravitational waves from this system are observed to be monochromatic at twice the orbital frequency, $f_{\rm GW} = 2f_{\rm orb} = \omega/\pi$, since the waves are quadrupolar. However, the "monochromatic" part of this statement isn't *quite* true – gravitational waves will carry energy away from the binary, its orbit will decay, and the bodies will steadily spiral toward one another until one day they violently merge. During this process, both the frequency and amplitude of radiation will sweep up as time goes on, yielding a waveform that is almost – but not quite – entirely unlike the sound of a bird chirping. (For this reason, we commonly refer to compact binary GW signals as "chirp" waveforms. See Fig. 2.4 for a cartoon of GW150914's characteristic chirp signal.) To correct our understanding of this process, we can introduce the parametric variable

$$v = (\pi G M f_{\rm GW})^{1/3} = \left(\frac{2\pi G M}{P_{\rm orb}}\right)^{1/3} = \sqrt{\frac{GM}{a}} \tag{2.31}$$

(where $P_{\rm orb} = 2\pi/\omega$ is the orbital period) as a surrogate for the GW frequency. In terms of v, and for any observer inclined at an angle ι with respect to the orbital plane, the

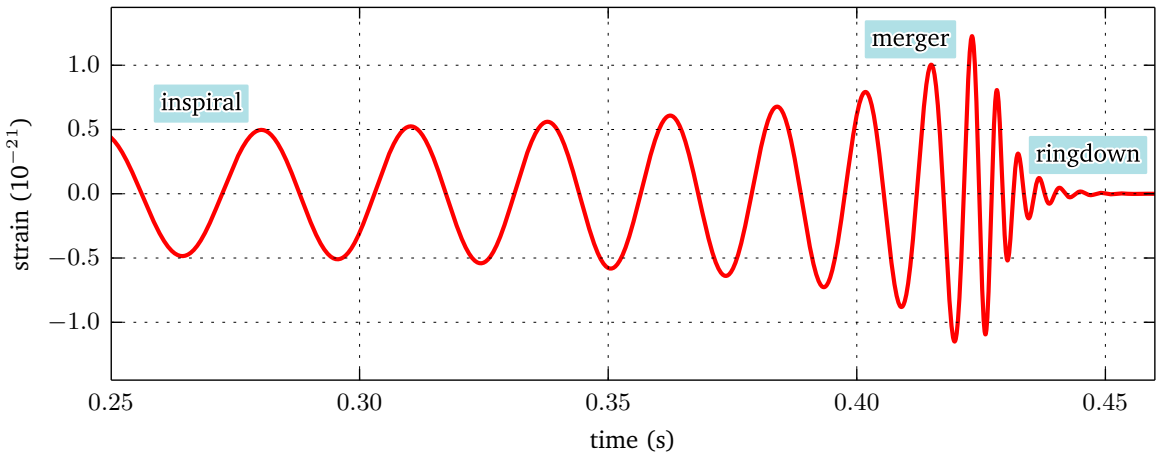

Figure 2.4: Estimated gravitational wave strain amplitude from transient source GW150914, projected onto the H1 detector (see section 2.2). Data shown are from numerical relativity models of the black hole horizons as the black holes coalesce, following a similar presentation in Abbott et al. (2016c). For illustration, we show the portions of the observed waveform corresponding to late inspiral and merger of the two black holes, followed by ringdown emission as the final black hole settles down. Note the duration ($\approx$0.2 seconds) of the signal in H1's sensitive band.

GW polarizations are

$$h_+(t(v)) = -\frac{2G\mu}{c^2 r}\left(1+\cos^2\iota\right)\left(\frac{v}{c}\right)^2\cos 2\phi(v) \tag{2.32}$$

$$h_\times(t(v)) = -\frac{4G\mu}{c^2 r}\cos\iota\left(\frac{v}{c}\right)^2\sin 2\phi(v). \tag{2.33}$$

Time, frequency, and phase evolution are then each determined by differential equations in the post-Newtonian parameter v, with frequency evolution in particular characterized by the *chirp mass*

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}} = \frac{c^3}{G}\left(\frac{5}{96}\pi^{-8/3} f^{-11/3}\dot{f}\right)^{3/5}. \tag{2.34}$$

In practice, these equations are expanded in powers of v/c and solved either numerically or analytically, and the resulting waveform accuracy presumably improves with every successive order kept. Moreover, because the distance to source r can be measured independently of the other amplitude coefficients, GW signals from compact binary coalescence are "standard sirens" in the same way that Cepheid variable stars are standard candles. Precise waveform models also allow sensitive matched-filter signal searches to be done with LIGO data, as will be discussed in section 2.2.3.

Assuming LIGO will observe binary systems that contain stellar-mass objects, we can use Kepler's third law, $GM = a^3\omega^2 = \pi^2 a^3 f_{\rm GW}^2$, to place constraints on the size

Electromagnetic Counterpart	Viewing Angle	Detectability Timescale	Bandpass
short GRB	$\theta_{\rm obs} \lesssim \theta_{\rm jet}$	$\sim$seconds	γ-ray/hard X-ray
X-ray afterglow	$\theta_{\rm obs} \lesssim \theta_{\rm jet}$	$\sim$minutes–hours	X-ray
optical afterglow	$\theta_{\rm obs} \lesssim 2\theta_{\rm jet}$	$\sim$minutes–days	optical
kilonova	isotropic	$\sim$days	optical/near-IR
radio afterglow	isotropic	$\sim$weeks–years	radio

Table 2.1: Summary of expected electromagnetic counterparts of NS-NS/NS-BH mergers as a function of the observing angle, $\theta_{\rm obs}$ (Metzger & Berger 2012; see Fig. 2.3). Within seconds after merger, a centrifugally supported accretion disk forms around the merger remnant (usually a BH). Dynamically rapid accretion lasting $\lesssim 1$ sec powers a collimated relativistic jet with half-opening angle $\theta_{\rm jet}$, yielding a short, hard γ-ray burst; this emission is only detectable to observers in the line of fire ($\theta_{\rm obs} \lesssim \theta_{\rm jet}$) because of relativistic beaming. Next, nonthermal "afterglow" emission is produced as the relativistic jet interacts with the surrounding interstellar medium. X-ray and optical afterglows will be detectable within $\sim$minutes from viewing angles $\theta_{\rm obs} \lesssim \theta_{\rm jet}$, while optical afterglows remain observable on timescales up to $\sim$days-weeks and from viewing angles $\theta_{\rm obs} \lesssim 2\theta_{\rm jet}$. As the jet slows and expands laterally, radio afterglows become roughly isotropic and remain visible on timescales of $\sim$weeks-months. They may also be produced on timescales of $\sim$years by sub-relativistic ejecta. Finally, short-lived, roughly isotropic optical counterparts called "kilonovae" can also accompany the merger. Powered by radioactive decay of heavy elements synthesized in the ejecta, kilonovae may last up to $\sim$a few days after merger. Note, the viewing angle $\theta_{\rm obs}$ is equivalent to the inclination angle ι of the orbital plane.

of the emitting region. Typically, if objects in the binary are each $1\,M_\odot$, the inspiral signal sweeps past 10 Hz at orbital separations of $a \sim 10^5$ m or less – the binary must be only at most several hundred kilometers across for $f_{\rm GW}$ to be in LIGO's sensitive range. Stellar mass objects this close together must be extraordinarily compact, which implies they must be neutron stars (NS) or black holes (BH). As of September 2015, both NS-NS and BH-BH systems have been observed in nature through radio and GW means, respectively. It is believed that NS-BH systems occur naturally as well, although to date there have been no direct observations. Each of these systems evolve across LIGO's sensitive frequency band over a span of $\sim$0.1–1000 seconds at the very end of their inspiral, when each body's orbital velocity has become a sizeable fraction of the speed of light. When the two objects merge at long last, they form a final black hole (or perhaps, in some cases, a hypermassive neutron star).

If at least one neutron star is present in the binary then we expect that either unstable collapse (NS-NS) or tidal deformation (NS-BH) during merger will instigate dynamically rapid disk formation on timescales of $\sim$0.1–1 sec (Metzger & Berger, 2012; D'Avanzo, 2015). In either case this disk quickly accretes onto the merger remnant, powering short-duration γ-ray bursts (-1–6 seconds after merger) and ultimately their X-ray ($\sim$minutes–hours), optical ($\sim$hours–days), and radio ($\sim$weeks–years) afterglows

(Table 2.1; see Metzger & Berger 2012). The short GRB explosion is both fleeting – lasting only about as long as the accretion timescale – and highly collimated (with opening jet angles $\theta_{\rm jet} \sim 20^\circ$) due to relativistic abberation (ejecta during this explosion are accelerated to Lorentz factors of $\Gamma_0 \sim 100$ or more; D'Avanzo 2015). However, as the ejecta slow down and interact with the interstellar medium, the afterglow emission expands laterally and can be observed from wider viewing angles. Furthermore, a roughly isotropic "kilonova" explosion powered by r-process nucleosynthesis in the disk ejecta is expected to become visible within about a day or so of the merger. Observing electromagnetic transients at every stage of this process would provide great insight into the binary's merger dynamics and give us more opportunities to test both General Relativity and specific short GRB emission models, among other things. A population of GW detections with electromagnetic counterparts will even make possible a local-universe measurement of the Hubble constant (H_0) by giving independent measurements of both luminosity distance and redshift (Berger, 2014a). Accordingly, astronomers around the world are heavily invested in triggering high-energy, optical and radio telescopes on real-time LIGO discoveries. Indeed, contributions to this effort will be the primary focus of this dissertation.

There are many other astrophysical bodies with nonlinear changes in quadrupole moment that we expect to radiate GWs over LIGO's sensitive frequency band. In this dissertation, however, we focus exclusively on transient signals from compact binary coalescence that do not recur and can only be observed once.

Now that we have understood something about LIGO's transient source population, we now turn to the Advanced LIGO instrumentation.

2.2 Interferometric Gravitational Wave Detectors

The two LIGO detectors, together with a third facility called Virgo in Cascina, Italy, initially conducted observations between 2002 and 2010 (Abbott et al., 2016c). No GW detections were made during this period. The LIGO and Virgo Collaboration subsequently began a series of upgrades to the network of detectors meant to make them an order of magnitude more sensitive to GW sources in the local universe. The LIGO facilities completed their first three-month Observing Run (O1) between 18 September

2015 and 13 January 2016, detecting GW150914 just prior to the official start of the Run, as operators and engineers were preparing the Hanford (H1) and Livingston (L1) detectors for stable data-gathering. At the time of writing, analysis of data from the whole of O1 is ongoing. The detectors are currently receiving another round of upgrades, and will begin gathering data again with yet more sensitivity in July 2016 in the frequency range from 25 Hz to 8 kHz.

In this section, I give an executive overview of the experimental methodology of Advanced LIGO. I begin by describing the basic setup of the project, then consider its most important sources of noise, which sets the distance sensitivity to compact binary sources and their rate of discovery. I briefly outline some common matched filter signal analysis techniques used to rapidly identify CBC signals in the LIGO data stream, and consider our ability to localize them on the sky. I then close this section by considering the technical obstacles facing the electromagnetic follow-up effort, and how they might be addressed.

2.2.1 Basic Experimental Setup

In an effort to distinguish GW signals from local instrumental and environmental disturbances and to provide reasonable sky localization, the three ground-based GW detectors H1, L1, and Virgo are widely separated geographically, with H1 and L1 straddling the contiguous United States and Virgo situated across the Atlantic in southern Europe. Each site operates a single GW detector (LIGO Scientific Collaboration et al., 2015; Acernese et al., 2015a), and each works on the same basic principle: a modified Michelson interferometer measures the differential change in length between two orthogonal L-shaped arms (Fig. 2.5). Each arm of the interferometer is formed by two mirrors which act as test masses (the *input* and *end* test masses, ITM and ETM) separated by a distance L_* in the absence of gravitational radiation. At the LIGO sites (H1 and L1) the arms are $L_* = 4$ km long, while at Virgo the arms are $L_* = 3$ km. The primary data stream then records the differential change in length between the x- and y-arms,

$$h(t) = \frac{L_x(t) - L_y(t)}{L_*}. \tag{2.35}$$

We call this quantity the *strain* due to gravitational waves. Differential length variations alter the phase difference between the laser beams as they return to the beam splitter,

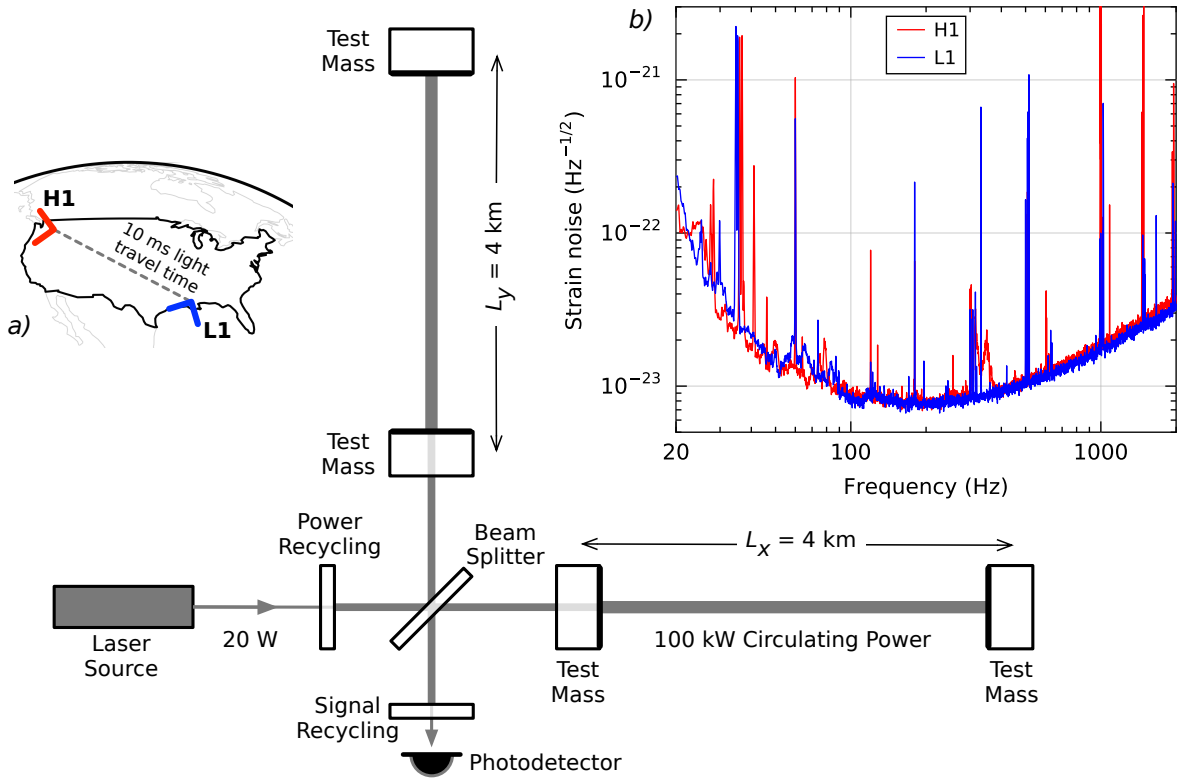

Figure 2.5: Simplified diagram of an Advanced LIGO detector (not to scale). A gravitational wave propagating orthogonally to the detector plane and linearly polarized parallel to the 4 km optical cavities will have the effect of lengthening one 4 km arm and shortening the other during one half-cycle of the wave (see section 2.1); these length changes are then reversed during the other half-cycle. An output photodetector records these differential cavity length variations as the wave passes. While a detector's directional response is maximal for this case, it is still significant for most other angles of incidence or polarizations (note, gravitational waves propagate freely through the Earth). Inset (a): Location and orientation of the LIGO detectors at Hanford, WA (H1) and Livingston, LA (L1). Inset (b): The instrument noise for each detector near the time GW150914 was observed; this is an amplitude spectral density, expressed in terms of equivalent gravitational-wave strain amplitude. The sensitivity is limited by photon shot noise at frequencies above 150 Hz, and by a superposition of other noise sources at lower frequencies (see section 2.2.2). Narrow-band features include calibration lines (3338, 330, and 1080 Hz), vibrational modes of suspension fibers (500 Hz and harmonics), and 60 Hz electric power grid harmonics. This figure is reproduced from Abbott et al. (2016c) with permission from the LIGO Open Science Center (see https://losc.ligo.org).

transmitting an optical signal to the output photodiode that is proportional to the GW strain $h(t)$.

How is the measured $h(t)$ related to h_+ and $h_\times$? Answering this question will involve writing down the Euler angles for a transformation from the source frame to the lab frame (see *e.g.* Nishizawa et al. 2009; Abbott et al. 2016d), accounting for the inclination angle ι of the detector with respect to the source's orbital plane, the polarization angle ψ describing the azimuthal orientation of the source frame, and the sky location of the source, which we can describe relative to the detector with spherical coordinates (θ, φ). (In these coordinates, we take the origin to be at the beamsplitter, and the x- and y-directions to be along the interferometer arms.) An arbitrary source emitting gravitational waves from a distance r will have plus and cross polarizations that depend on r, ι, and the emission physics, so the strain on a detector (absent any noise) is

$$h(t) = F_+(\theta, \varphi, \psi) h_+(t; r, \iota, \dots) + F_\times(\theta, \varphi, \psi) h_\times(t; r, \iota, \dots) \tag{2.36}$$

where $F_+, F_\times$ are the antenna response patterns for each polarization. After projecting onto the detector plane, one can show that, for the interferometric LIGO and Virgo detectors,

$$F_+(\theta, \varphi, \psi) = \frac{1}{2}\left(1 + \cos^2\theta\right)\cos 2\varphi \cos 2\psi + \cos\theta \sin 2\varphi \sin 2\psi \tag{2.37}$$

$$F_\times(\theta, \varphi, \psi) = -\frac{1}{2}\left(1 + \cos^2\theta\right)\cos 2\varphi \sin 2\psi + \cos\theta \sin 2\varphi \cos 2\psi. \tag{2.38}$$

Notice that this L-shaped configuration has antenna response patterns that support most of the sky at fixed ψ, so that at least one polarization will almost always contribute significantly to $h(t)$. Note, there are four directions in the detector plane from which a gravitational plane wave would displace the x and y arms equally (producing zero strain).

2.2.2 Sources of Noise

A $\mu = 1\, M_\odot$ binary at a distance $r = 100$ Mpc will produce strain with characteristic amplitude $4G\mu/c^2 r \sim 10^{-21}$. To achieve the sensitivity needed to measure gravitational waves this weak, the LIGO detectors include at least three significant improvements on the basic Michelson interferometer described above (LIGO Scientific Collaboration

et al., 2015; Abbott et al., 2016c). First, the test masses on each arm form a resonant optical cavity (called a *Fabry-Pérot cavity*) that multiplies any change in the phase of circulating laser light due to GWs by a factor of ∼300. Second, a partially transmissive power-recycling mirror installed at the input causes resonant buildup of laser light across the entire interferometer array. A 20 W laser input is increased to 700 W when it reaches the beam splitter, and is further increased to 100 kW as it circulates in each arm cavity. Finally, a partially transmissive signal-recycling mirror at the output then broadens the bandwidth of the arm cavities. The input beam is a 1064 nm wavelength Nd:YAG laser which is stabilized in amplitude, frequency, and beam geometry. The GW signal is then extracted at the output port. The nominal $h(t)$ sampling rate is 16 kHz, allowing signals to be analyzed (in principle) up to 8 kHz.

These optics and interferometry techniques are designed to simplify and enhance the process of converting incident GW strain to output optical signal, so that we can minimize the influence of photon shot noise (the dominant noise source above 500 Hz; see Fig. 2.5). Other actions are taken to mitigate strain noise at lower frequencies. In order to detect gravitational waves, the test masses must be kept very close to free-fall and must not be susceptible to displacement by ground motion in the environment near the detector. We achieve this by isolating the ITMs and ETMs from seismic activity, minimizing noise at frequencies ∼10–100 Hz. Each test mass is installed as the final stage of a quadruple-pendulum suspension system (with a frequency response that goes as f^{-8}), which is then further supported by an active seismic isolation platform. Taken together, these systems provide more than 10 orders of magnitude of isolation from ground motion for frequencies above 10 Hz. Thermal noise at ∼100–500 Hz is mitigated by using materials with low mechanical loss in the test masses and their suspensions to reduce the impact of Brownian motion. The test masses themselves are 40 kg fused silica substrates with low-loss dielectric optical coatings, and they are suspended with fused silica fibers from the stage above.

To minimize any additional noise sources, all components of the interferometer other than the laser source are mounted on vibration isolation stages in ultra-high vacuum. Optical phase fluctuations caused by Rayleigh scattering are reduced by maintaining the pressure in the tubes containing the arm-cavity beams (1.2 m in diameter) strictly below 1 μPa.

How are the data calibrated? At the LIGO sites, servo controls hold the 4 km arm

cavities in resonance and keep the optical components precisely aligned. Strain data from the output photodiode is calibrated by measuring its response to lateral test mass motion, which is controlled by photon pressure from an additional modulated laser beam. For GW150914, the calibration was established to a 1σ uncertainty of less than 10% in amplitude and 10° in phase. It was also monitored continuously with laser excitations at specific frequencies (refer again to Fig. 2.5). Two alternative methods are used to validate the absolute calibration: one involving the main laser wavelength, and the other, a radio-frequency oscillator. (Furthermore, the detector response to gravitational waves is also tested by injecting simulated waveforms straight into the detector hardware with the calibration laser.)

To monitor environmental disturbances – things such as microseismic ground motion induced by weather, wind motion, trains and large trucks near the detector – and their influence on the interferometers, each LIGO observatory site is equipped with a vast array of sensors. These include seismometers, accelerometers, microphones, magnetometers, radio receivers, weather sensors, AC-power line monitors, and a cosmic-ray detector. Another $\sim 10^5$ channels record a breathtaking wealth of information about the interferometer's operation status and the state of its control systems. Data gathering is synchronized to Global Positioning System (GPS) time to better than 10 μs. Timing accuracy is verified with both an atomic clock and a secondary GPS receiver at each observatory site. This timing is critical for triangulating signal candidates on the sky, as we will see in section 2.2.4.

In their most sensitive band, 100-300 Hz, the O1-era LIGO detectors were 3 to 5 times more sensitive in strain than they were during Initial LIGO; at lower frequencies the improvement is even better, with over ten times better sensitivity below 60 Hz. Because the detector response, Eq. (2.36), is proportional to GW amplitude, at low redshift the (approximately Euclidean) volume of space to which they are sensitive increases with the cube of strain sensitivity. For binary black holes with masses similar to GW150914, the space-time volume surveyed by all observations reported to date improves over previous ones by more than an order of magnitude.

2.2.3 Detection Methods

Identification of binary signals in the $h(t)$ data stream boils down to separating sufficiently strong (or "loud") signals from random noise by assigning some measure of their statistical significance (see Creighton & Anderson 2011). In general, the ith detector in the network will have a calibrated data stream, $s_i(t)$, that is populated both by that detector's response to a passing GW signal, $h_i(t)$, and by local noise disturbances, $n_i(t)$. These will combine linearly as the time series

$$s_i(t) = h_i(t) + n_i(t). \tag{2.39}$$

For a network of N detectors, the combined likelihood ratio

$$\Lambda_{\rm net} = \prod_{i=1}^{N} \Lambda_i = \exp\left(\sum_{i=1}^{N} \left[(s_i|h)_i - \frac{1}{2}(h|h)_i \right] \right) \tag{2.40}$$

is the optimal ranking statistic if $h(t)$ is known and if $n_i(t)$ result from a Gaussian random process, where

$$(f|g)_i = 4\,\mathrm{Re} \int_0^\infty \frac{\tilde{f}^*(f)\,\tilde{g}(f)}{S_i(f)}\,df \tag{2.41}$$

is the noise-weighted inner product of two functions f, g in the ith detector, whose power spectral density is $S_i(f)$. (A tilde denotes the frequency domain representation of a time series.) In the ith detector, the signal-to-noise ratio

$$\rho_i^2 = \frac{(s_i|h)_i^2}{(h|h)_i} \tag{2.42}$$

and the max log-likelihood, $\ln \Lambda_{i,\max} \sim (s_i|h)$. In Gaussian noise, this quantity is Gaussian-distributed with zero mean and unit variance when no signal is present; if a GW is passing by, then the mean becomes $(h|h)_i^{1/2}$. The network SNR is simply the quadrature sum of single-detector SNRs,

$$\rho_{\rm net}^2 = \sum_{i=1}^{N} \rho_i^2 = \sum_{i=1}^{N} \frac{(s_i|h)_i^2}{(h|h)_i}. \tag{2.43}$$

This suggests a method of measuring significance: since compact binary signals are known to high precision, data streams may be filtered with template waveforms parametrized by masses (m_1, m_2) and spins $(\mathbf{S}_1, \mathbf{S}_2)$; we then keep the template waveform that maximizes $\rho_{\rm net}$ with consistent template parameters and times of arrival in each detector.

System Type	Volumetric Rate (yr^{-1} Gpc^{-3})	Max Distance (Mpc)	Detection Rate (yr^{-1})
NS-NS	$\sim$10–10^4	445	0.4–400
NS-BH	$\sim$10–10^3	927	0.2–300
BH-BH	$\sim10^{-1}$–10^2	2187	0.4–1000

Table 2.2: Anticipated discovery rate of compact binary mergers at Advanced LIGO design sensitivity, according to Abadie et al. (2010). The measured uncertainty in cosmic volumetric rate for each type of system is uncertain by orders of magnitude, so we give only order-of-magnitude estimates. The distance quoted is the maximum distance at which Advanced LIGO would observe an optimal source at 5σ confidence.

After accounting for the large number of templates used (Lyons, 2008), this gives a way to compute a p-value for the null hypothesis (no signal is present).

In practice, there are several factors complicating this procedure. First, seismic activity and other disturbances mean that the noise in LIGO and Virgo detectors is both non-stationary and non-Gaussian (Nuttall et al., 2015; The LIGO Scientific Collaboration & the Virgo Collaboration, 2016a). We get around the former problem by demanding that any candidate signal is observed at coincident times in at least two detectors, dramatically reducing the rate of false positives. We also reject non-stationary noise fluctuations by performing χ^2 tests across the detector network as a signal consistency check (Allen, 2005), and by vetoing chunks of analysis time over which well-understood transient noise glitches occur. (Re-weighting the single-detector SNRs by the measured χ^2 value also makes the search behave closer to a Gaussian; see Allen et al. 2012.) Because the background distribution is not known *a priori*, we must measure it empirically with histograms of the ranking statistic ρ_{net} in off-source chunks of analysis time. A few methods exist for doing this, with some that are computationally fast but less sensitive (Cannon et al., 2012a) and another that is computationally expensive but more sensitive (Usman et al., 2015). The faster methods are used to perform real-time searches for compact binary signals that deliver source candidates within 30-60 seconds of merger, making it possible to send early warning to astronomers so that they can target the littany of electromagnetic counterparts (see sections 2.2.4 and 2.2.5). The more sensitive method is then used to assign a signal's final measured significance.

How often will LIGO and Virgo observe compact binary merger transients? Notice from Eqs. (2.32) and (2.43) that $\rho_{\mathrm{net}} \propto 1/r$, where r is the distance to source. Thus,

$f_{\rm low}$ (Hz)	Time for 1.4-1.4 $M_\odot$ (sec)	Time for 1.4-10 $M_\odot$ (sec)	Time for 30-30 $M_\odot$ (sec)
30 Hz	54.7	12.0	0.278
25 Hz	88.8	19.6	0.473
10 Hz	1020	227	6.07

Table 2.3: Predicted duration of compact binary merger signals across the Advanced LIGO frequency range at O1 ($f_{\rm low} = 35$ Hz), O2 ($f_{\rm low} = 25$ Hz), and design ($f_{\rm low} = 10$ Hz) sensitivity. The time quoted is the time taken for the binary to cross from $f_{\rm low}$ to either 8 kHz or its innermost stable circular orbit (whichever is lower). From left to right, numbers are quoted for binaries of typical masses for NS-NS, NS-BH, and BH-BH systems, respectively. These signal durations are computed using the LALInference software library (Veitch et al., 2015).

the power spectral density $S_i(f)$ essentially sets the ith detector's distance sensitivity to these sources; at the time of GW150914, both H1 and L1 were operating with a 5σ angle- and orientation-averaged sensitivity to NS-NS sources of 70-80 Mpc.[7] Table 2.2 lists the measured rate of compact binary mergers of various types, and their expected rate of discovery given the Advanced LIGO commissioning schedule. Note that, at design sensitivity, we expect some 40 NS-NS detections per year (Abadie et al., 2010) with up to 30-40 minutes' early warning for astronomers (Table 2.3).

2.2.4 Sky Localization

For a single LIGO detector, all-sky amplitude sensitivity is quantified by the RMS antenna pattern

$$F^2_{\rm rms}(\theta,\varphi) = \frac{1}{2}\left[F^2_+(\theta,\varphi,\psi) + F^2_\times(\theta,\varphi,\psi)\right] \tag{2.44}$$

which is independent of polarization angle ψ (see Eq. 2.36). The LIGO and Virgo detectors are all ground-based, so sky locations should really be expressed relative to some geocentred coordinate system like the standard geographic coordinates (latitude Θ and longitude Φ) or equatorial coordinates (right ascension α and declination δ in the J2000.0 epoch). The latter system is far more appropriate for astrophysical sources since (α, δ) are fixed on the sky, but it is easier to visualize network effects in geographic coordinates, so for the moment we will think in terms of them. Consider a network of N GW detectors. If individual instruments are indexed by H, then the network RMS

[7] NS-NS sources are the standard metric for distance sensitivity because they remain in band for hundreds of cycles and emit across LIGO's entire frequency range.

antenna pattern is

$$F^2_{\mathrm{rms,net}}(\Theta,\Phi) = \frac{1}{N}\sum_H F^2_{\mathrm{rms},H}(\Theta,\Phi). \tag{2.45}$$

The H1, L1 detectors are not quite co-located (Abbott et al., 2016d), so Earth's curvature causes them to be slightly mis-aligned, and they also have different orientations of the x and y arms (Fig. 2.5). Mis-alignment guarantees that at least one detector will always have some strain response, even if it is very small. But because they lie only 3000 km apart on the same continent, the $\{\mathrm{H1},\mathrm{L1}\}$ antenna response patterns are similar enough that the maximal network coverage is strongly degenerate. This degeneracy in F_{rms} for the H1-L1 detector network is illustrated in Fig. 2.6. (Note that, over the course of a day, Earth's rotation will slide the network's high-sensitivity lobes across every part of the sky.) During O1, H1 and L1 were the only instruments observing, but there are plans for Virgo to come online during late O2 or early O3. Virgo's influence on the network will improve all-sky coverage and break the H1-L1 degeneracy for the nearest sources, but its lower strain sensitivity means that distant signals may be much quieter in the Virgo data stream, so for these sources the degeneracy will remain (Singer et al., 2014).

The network antenna pattern, Eq. (2.45), gives the coherent all-sky sensitivity of detectors to $h(t)$ signals at fixed amplitude. With a known waveform model (as is the case for compact binary signals) and with measurements of ι and ψ, this allows us to reconstruct the on-sky position of an observed source. Additional constraints on sky location come from enforcing phase consistency across detectors and from triangulating the signal based on differences in its observed arrival time at each detector. Regarding the latter method, let D_{IJ} be the distance between detectors I and J. Since GWs are predicted to propagate at the speed of light, the maximum arrival time difference is the light-travel time, D_{IJ}/c, and this will only be measured when the GW travels along the axis joining each detector. When the wave passes from an angle ϑ_{IJ} away from this axis, the measured time delay is

$$\Delta t_{IJ} = \frac{D_{IJ}}{c}\cos\vartheta_{IJ}. \tag{2.46}$$

Thus, the locus of constant time delay is a ring on the sky; timing uncertainty (see section 2.2.2) turns this ring into an annulus. For networks with three or more detectors, the annuli from each separate detector pair will overlap in the true direction of the source.

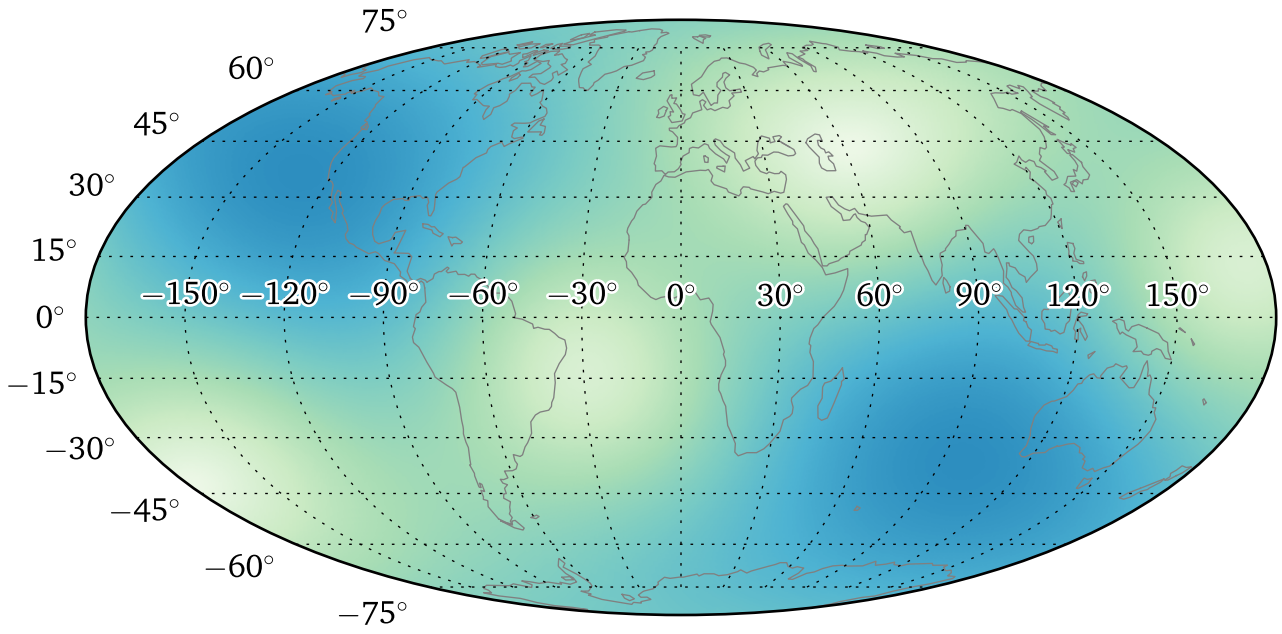

Figure 2.6: Coherent sensitivity of the H1-L1 GW detector network. This map is a Mollweide projection in geographic coordinates, emphasizing spatial relationships with respect to the Earth-fixed GW detector network, as well as possible ground-based telescope sites. Following a similar presentation in Singer et al. (2014), shading indicates the RMS network antenna pattern, Eq. (2.45), with darker blue areas corresponding to high sensitivity and white corresponding to null sensitivity.

In practice, for compact binary sources these methods are combined using Bayesian statistics in the rapid BAYESTAR sky localization algorithm (Singer & Price, 2016). This algorithm begins with the recovered signal template and marginalizes over inclination, polarization, orbital phase and distance to source, demanding phase consistency across detectors and using both amplitude sensitivity and signal timing to reconstruct the observed signal's sky location. Because BAYESTAR uses the recovered signal template, its sky localization ability will improve with the network SNR. Since the pre-O1 engineering period began, BAYESTAR has been run on all CBC signal candidates from specialized computing hardware hosted at the California Institute of Technology. This machinery uses Intel Haswell (Xeon Phi) processors with 16 physical cores (and 32 threads), bringing the time taken for BAYESTAR localization down to < 60 seconds.

A typical sky map showing BAYESTAR's reconstruction of a simulated NS-NS signal, found in simulated Gaussian H1-L1 strain noise modeled on the O1-era Advanced LIGO commissioning schedule, is shown in Fig. 2.7. Notice that the H1-L1 network degen-

eracy effectively breaks the timing annulus up into two lobes straddling the north and south geographic hemispheres. Compared to the full parameter estimation software, which uses Markov Chain Monte Carlo (MCMC) methods to measure both intrinsic and extrinsic CBC model parameters and can take ∼several hours per signal candidate to compute, BAYESTAR does comparably well in the H1-L1 network. However, when Virgo is added to the network with much lower sensitivity than H1/L1's O2-era strain spectra, full parameter estimation can improve on BAYESTAR's localization area by factors of 5 or more. As a metric for rapid sky localization performance we use the *searched area*, that is, the sky area contained in the smallest contour around both BAYESTAR's computed maximum *a posteriori* probability and the associated true location of a simulated NS-NS source. The median searched area in O1(O2)-era strain spectra is ≈170 deg^2 (120 deg^2), as reported in Singer et al. (2014).

2.2.5 Electromagnetic Follow-up

In the interest of maximizing the science output of Advanced LIGO and Virgo, it is important to share results as rapidly we can with astronomer partners at optical, radio, high energy and neutrino facilities around the world. At the same time, we want to keep our analysis as clean and unbiased as possible, especially in the era of first detections. To this end, we have worked to communicate results to a privately-subscribed subset of the wider astronomical community via the Gamma-ray Coordination Network (GCN[8]), maintained at NASA Goddard Spaceflight Center, during O1 (see chapter 6). We also put together a collection of volunteers to act as "EM follow-up advocates" in the event of an interesting GW detection candidate. Together with an on-call rapid response team (RRT) and the operators and staff on control room shifts at the two LIGO sites, this team would meet collectively to decide whether a GW detection candidate was worthy of astronomical follow-up within ∼10–100 minutes of its appearance in the Gravitational-wave Candidate Event Database (GraceDB[9]).

This rapid response team included experts in both instrumentation and data analysis from the LIGO project. The presence of detector characterization (DetChar) experts proved crucial, largely because these are the people extensively knowledgeable

[8] http://gcn.gsfc.nasa.gov
[9] https://gracedb.ligo.org

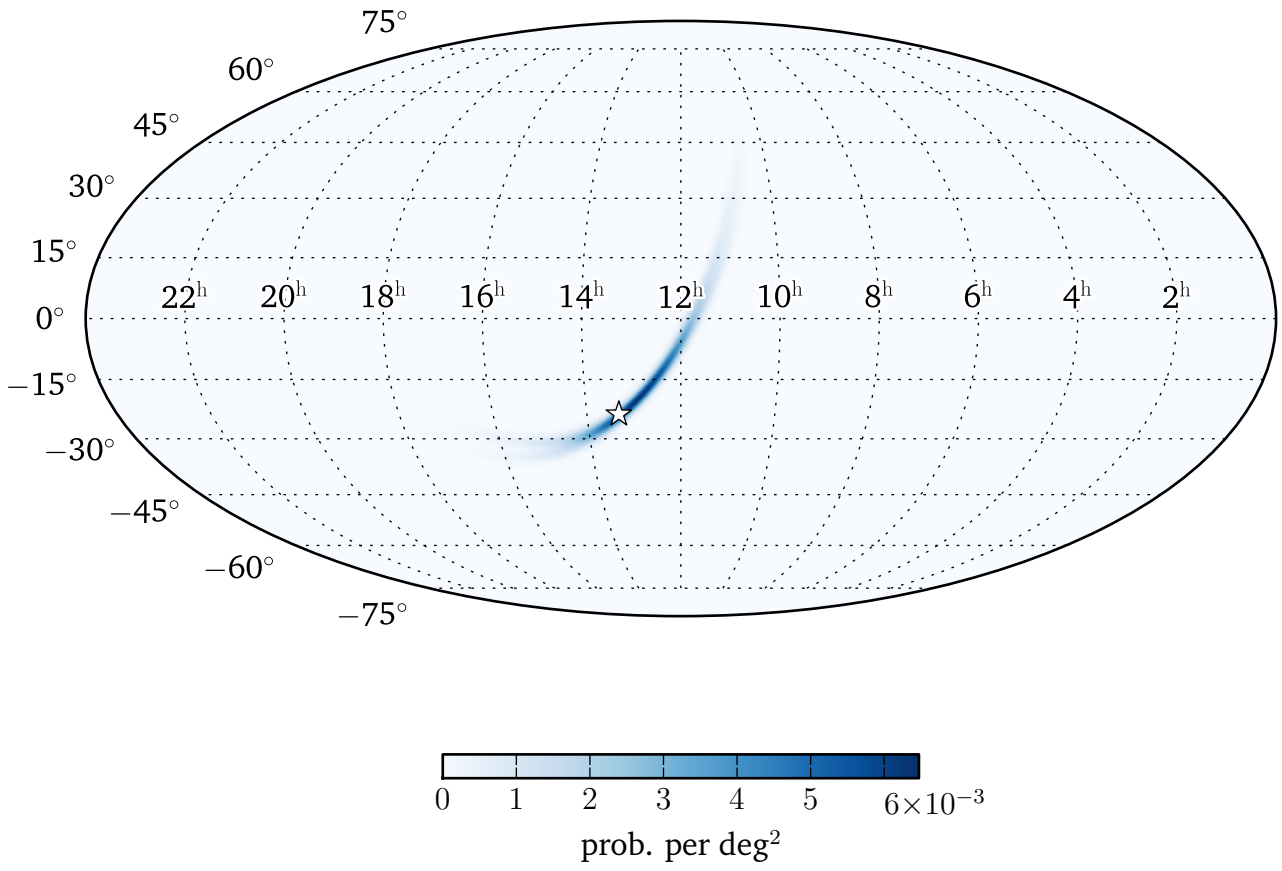

Figure 2.7: Localization of a typical simulated binary neutron star signal in simulated noise based on the circa 2015 Advanced LIGO commissioning schedule (Singer et al., 2014). The true location of the simulated signal is marked with a white star. This map is a Mollweide projection in equatorial (J2000.0) coordinates, with shading proportional to the reconstructed posterior probability per deg^2. The signal was recovered near the 5σ significance threshold with $\rho_{\rm net}$ = 12.7, and its 90% confidence area of 530 deg^2 is near the median of sky maps in this study. Nearly all of the posterior probability density is concentrated in a single, long, thin "island" over the southern H1-L1 network antenna pattern maxima at the time of coalescence (*cf.* Fig. 2.6). The shape and width of this island is due to relative timing uncertainty between the two detectors, and its edge is forked like a snakes tongue, with one fork corresponding to the binary having face-on inclination ($\iota \approx 0°$) and the other fork corresponding to face-off ($\iota \approx 180°$). In this study, approximately 50% of recovered sky posteriors are unimodal, with 35% having two degenerate modes and the remaining 15% having three or more.

in identifying sources of noise that affect the astrophysical searches. This expertise was necessary in order to reliably classify transient noise phenomena which may conspire to resemble a GW signal in two or more detectors, and only a seasoned DetChar representative would have the experience needed to determine whether a given candidate is likely to be due to *e.g.* some noise or environmental artifact rather than a *bona fide* gravitational wave. In order to train the EM follow-up program, during O1, we agreed to broadcast any signal candidate over the private GCN feed that was not likely to be caused by transient noise, not associated with a planned hardware injection, not a test

of the data analysis pipelines or computing infrastructure, and had an estimated false alarm rate (FAR) below 3.8×10^{-7} Hz (≈ 1 month^{-1}). This program was fully implemented for the genuine signal GW150914, as will be described in chapter 6. While we expect the nature of specific noise transients to change in future engineering runs as the detectors' strain spectra evolve, the global EM follow-up procedure is expected to stand basically in place, with alerts broadcast over the GCN becoming public after LIGO announces its first four transient GW discoveries.

The primary data quality question, as it relates to electromagnetic follow-up, is to whether the GW detection candidate is likely to have been caused by glitches in one or more detectors. Answering this question quickly is of course not always straightforward (or indeed, even possible), and at this time cannot naturally be done in an automated way. For this reason we advocated a policy of "shoot first, ask questions later" during O1: unless something was obviously wrong with the calibrated $h(t)$ data streaming from the sites at the event time, the rapid response team would make the call to distribute an alert to astronomers. EM follow-up advocates would continue to look into the situation with the rapid response team, keeping a close real-time communication channel open during this process, and would issue a retraction within an hour or two if the DetChar expert(s) felt it was necessary. Given that high-energy EM counterparts of GW transients are expected to reach peak intensity within hours of the event, we found this made rapid observations possible and places the burden of deciding to follow up on individual astronomers. Note that this has been the model used in the GRB community for over ten years.

In the GCN system, "Notices" are automated machine-readable alerts sent within tens of seconds or minutes of an event; this timescale is consistent with LIGO's low-latency data products, while the full parameter estimation usually finishes in time for a more detailed, human readable GCN "Circular" once the relevant experts have had time to process the detection candidate (see Figure 2.8). Astronomers generally expect that some non-negligible fraction of Notices will be triggered on noise events and other spurious artifacts, while the Circulars will contain more reliable information about event candidates that survive an initial data quality vetting process. Before the start of O1, the expectation among astronomers was that LIGO would distribute $\sim$1 false alarm event per calendar month, in order to provide the astronomers with a population of events to train their follow-up capabilities on.

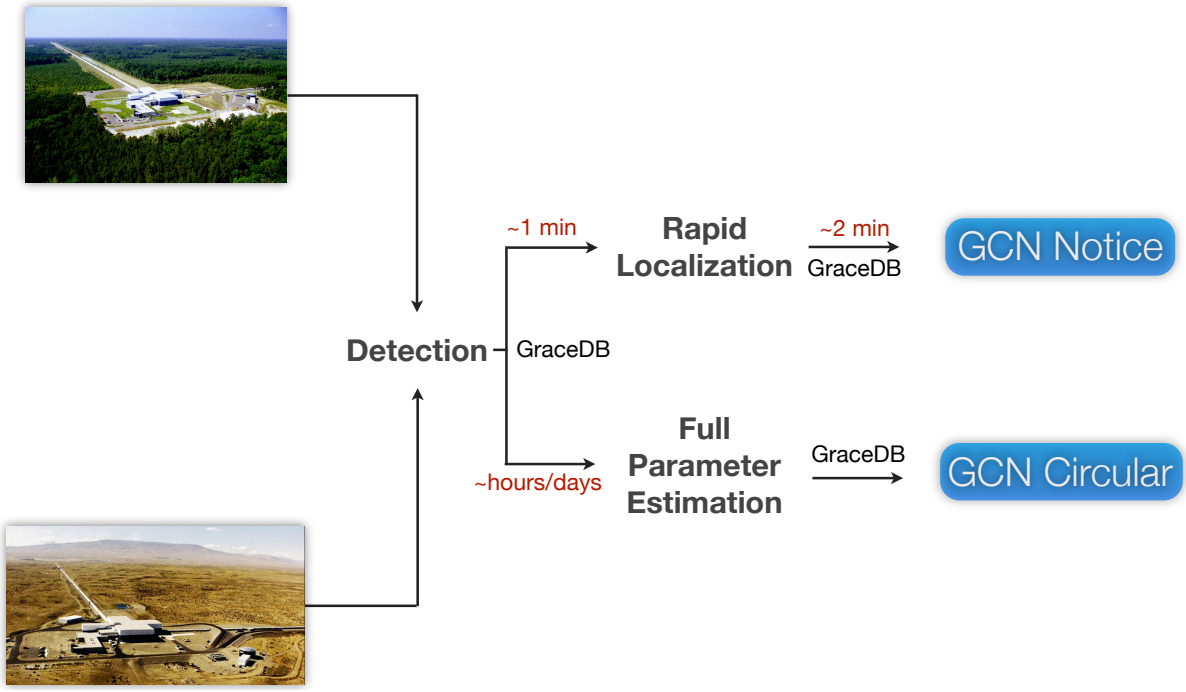

Figure 2.8: Flow chart illustrating the stream of data products on various timescales, and of the dissemination of GCN Notices (machine-readable alerts in low latency) and Circulars (more carefully vetted human-readable alerts ∼2–36 hours later). This figure is reproduced from Urban et al. (2015).

The remaining question is, what constitutes an "obvious" noise transient? The DetChar representative to the rapid response team will advise the EM follow-up advocate on the quality of the data at the time of a candidate event based on the following general guideline criteria:

1. Are any of the common classes of glitch (*e.g.* RF45 glitches, whistles, or "blip" gliches) in occurrence near the time of the event? Is the event itself contaminated or likely to be caused by such a glitch? (Refer to Smith 2015 for a summary of the commonly understood glitch classes.)

2. Is there significant coupling between $h(t)$ and any auxiliary instrument status channels?

3. Is the local environment in a state known or likely to produce noise couplings? (That is, are there wind gusts exceeding 10 MPH? Have there been any nearby or especially powerful earthquakes? Has there been a recent dip in temperature?)

4. Is there anything else wrong with the state of the interferometers not accounted

for by operators in the control room? For instance, is the calibration solid near the time of the event?

A number of tools exist to aid in addressing these questions, including (but not limited to) time-frequency scans of the important auxiliary channels, as well as the iDQ data quality infrastructure (Essick, 2015). The consulting DetChar expert will of course be familiar with these and know the proper course of action to take in investigating.

In the Event that an Alert is Vetoed

If detector characterization experts determine that the $h(t)$ stream in one or more detectors near the time of an interesting candidate is suffering from bad calibration or saturated with obvious, well-understood glitches, they may recommend rejecting the candidate from astronomical follow-up. In this case, a retraction GCN Notice would be issued. (If the decision is made early enough, for example if an operator on shift declares one of the instruments to not be in a good state, the Preliminary GCN Notice may be preempted altogether.) The detector characterization expert would then provide a brief but detailed summary explanation justifying this veto, carefully noting any environmental concerns or glitch artifacts that cast suspicions on the event. For example, if a candidate is rejected based on local seismic activity, the explanation might read as follows:

> "There had been significant ground motion within a second or so of the event time, as recorded in beam splitter optical lever channels and various accelerometers. Though these features were seen at low frequencies, similar disturbances have commonly been observed to upconvert to higher frequencies, where low latency searches generally look for signal candidates."[10]

Justification for the veto would be logged manually in the event's GraceDB entry.

In rarer cases, the retraction may be issued after a GCN Circular has been posted; in this case, it is up to the EM follow-up advocate to issue a retraction Circular.

In the Event that an Alert is Approved

On the other hand, if DetChar, the rapid response team, and the on-duty EM follow-up advocate all agree that the $h(t)$ streams are clean (or at least free of any obvious

[10]This example is paraphrased from a log entry by Peter Shawhan in GraceDB event G195945, which occurred in the second month of O1 on 28 October 2015.

glitch artifacts and other common noise sources), then the EM follow-up advocate will issue an Initial GCN Notice (which contains a link to a low-latency sky map) and begin assembling the human readable GCN Circular.[11] Once information has been broadcast over the GCN, astronomers at various facilities will then decide whether or not to conduct observations.

2.3 Space-based γ-Ray Observatories

Because we expect electromagnetic transients that occur as a result of compact binary merger, jointly observing them can greatly enhance our understanding of both the binaries' merger dynamics and their role in the larger cosmological context. For most neutron star binaries whose orbital plane is inclined at an angle $\iota \leq 20°$, the first EM counterpart we expect to see are short-duration GRBs. However, $\sim$1-minute latencies in data analysis (see sections 2.2.3 and 2.2.4) mean that, for the moment, it is impractical to trigger γ-ray observations on LIGO discoveries, since the GRBs will presumably occur during or immediately after merger. In chapter 3 we will explore a method by which independently-discovered GRBs may be rapidly correlated with LIGO signal candidates within 60-120 seconds of merger. But doing this will of course require instrumentation, and most reliable high energy obvservatories are space-based to overcome the effects of Earth's atmosphere, which raises the possibility of data latency issues.

Fortunately, there are two satellite observatories[12] – the *Swift* and *Fermi* missions – with instruments specially designed to detect GRBs and which report confident discoveries over the GCN system on timescales similar to LIGO's data analysis pipelines. Both missions are led by the US-based National Aeronautics and Space Administration (NASA), and we use their rapid data products to search for coincidences with CBC candidates in the LIGO data stream. Leading up to O1, an automated system was developed and installed on the same hardware used for LIGO sky localization. This system ingested GCN Notices from the *Swift* and *Fermi* satellites, filtered them for GRB-like objects likely to be astrophysical in origin, and uploaded them to GraceDb. In this

[11] For full notes on this process, please see the EM follow-up advocate wiki page.

[12] An InterPlanetary Network (IPN) of high energy observatories also exists, but since localization can take several hours, we cannot use this in near-real time.

section, we briefly examine the sensitivity of each of these satellite instruments.

2.3.1 *Swift* Burst Alert Telescope (BAT)

Named for a family of bird that ranks among the fastest animals on Earth, the *Swift* satellite (Gehrels et al., 2004a) is a high energy observatory with rapid autonomous slew capability. The spacecraft uses reaction wheels to trigger its narrow-field X-Ray Telescope (XRT) and Ultraviolet/Optical Telescope (UVOT) on GRBs discovered by its wide-field Burst Alert Telescope (BAT). The BAT instrument consists of a 5200 cm^2 plane of CdZnTe tiles (Fig. 2.9), which is photoconductive to hard X-ray photons in the spectral range 15-150 keV. Suspended above the detector array is a coded aperture mask with 52,000 randomly placed 5-mm lead tiles, which is used to localize GRBs with 1-4 arcminute accuracy. *Swift* BAT has a fully-coded instantaneous field of view (FOV) of 1.4 sr (4600 deg^2) and a partially-coded FOV of just under 3.0 sr (9800 deg^2). Onboard software computes the GRB direction within 15 seconds and transmits to the ground by way of NASA's Tracking and Data Relay Satellite System (TDRSS), at which point a Notice is broadcast over the GCN. Because the onboard *Swift* software is efficient in identifying genuine bursts and their sky location, we upload all GCN Notices supplied by the *Swift* BAT instrument to GraceDb.

2.3.2 *Fermi* Gamma-ray Burst Monitor (GBM)

The *Fermi* Gamma-ray Space Telescope is another high energy astronomy platform in low-Earth orbit. While its main instrument, the Large Area Telescope (LAT: Abdo et al. 2009), is used to study a diverse set of astrophysical phenomena, *Fermi* also has a secondary instrument called the Gamma-ray Burst Monitor (GBM: Meegan et al. 2009a; Fig. 2.9) which is used to discover GRBs. This instrument is a collection of scintillation detectors with $\sim$1-10° on-sky precision. It consists of 12 NaI crystals, in sets of three mutually orthogonal mounts, and 2 BGO crystals; the total spectral range is from 8 keV to 40 MeV. Sky localization is done with triangulation, and the instantaneous field of view is $\sim 3\pi$ (31,000 deg^2), or essentially every part of the sky not occluded by the Earth. Upon detecting a potential GRB, on-board software will compute a crude sky location and transmit to the ground via TDRSS within 30 seconds. However, many

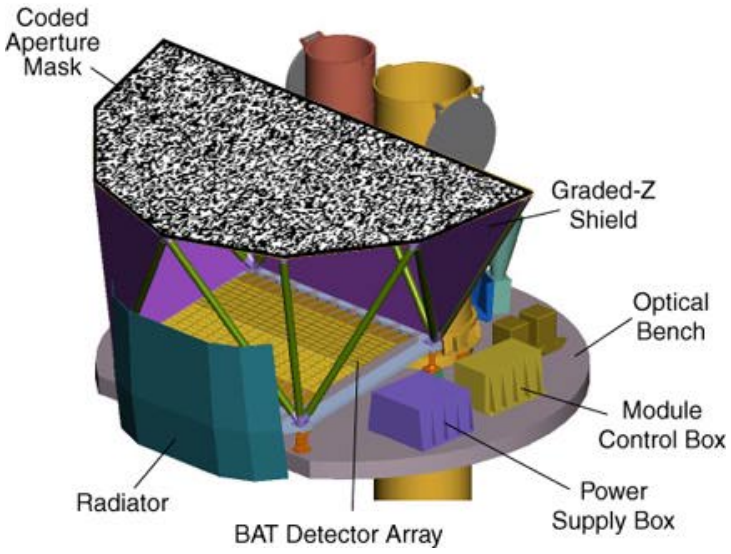

(a) *Swift* BAT (*Image source:* http://swift.gsfc.nasa.gov)

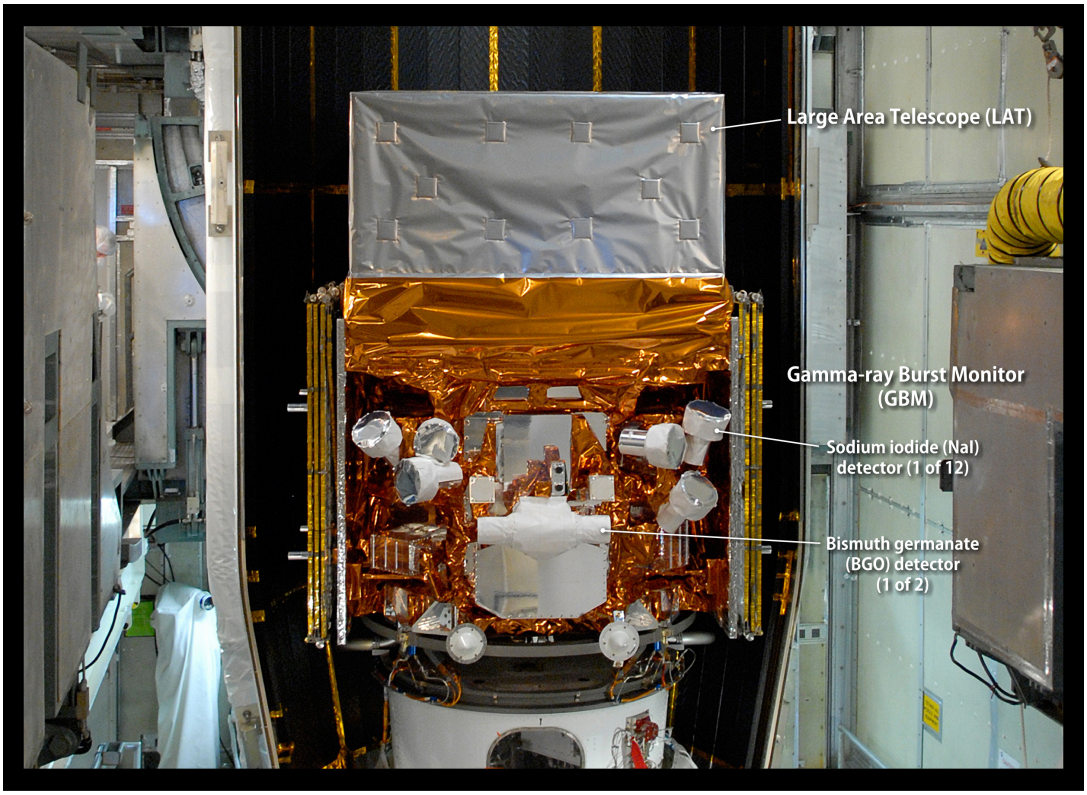

(b) *Fermi* GBM (*Image source:* http://gammaray.nsstc.nasa.gov)

Figure 2.9: Diagram of the *Swift* and *Fermi* γ-ray observatories. *Swift* BAT has much better on-sky precision than *Fermi* GBM, but the latter instrument has better sensitivity across the harder spectral range where short GRBs are typically found. For this reason *Fermi* supplies short GRBs at a rate of $\sim$1 week^{-1}. These diagrams are reproduced from NASA public webpages with permission from the *Swift* and *Fermi* instrument teams.

spurious non-astrophysical sources survive this process, and more careful sky localization computed on the ground will account for a systematic bias of 3° or so due to lateral motion of the spacecraft in its orbit. This more careful sky localization excludes more junk artifacts, and is broadcast over GCN typically within ∼1 minute of a GRB. The ground-computed localization Notices are the ones we ingest into GraceDb for LIGO searches.

2.4 Time Domain Optical Astronomy

As X-ray afterglows become detectable in the minutes and hours following a compact binary merger, it will be very useful to trigger instruments such as *Swift*'s XRT, which has to date been used extensively to study GRB afterglows and supernovae. A major obstacle in this regard is that *Swift* XRT has a very narrow FOV, making it difficult to search the hundreds of deg^2 supported by a typical LIGO sky map. Wide-field optical telescopes will be crucial for this endeavor, especially ones that can take deep images and slew very quickly, covering a large sky area with the sensitivity needed to catch off-axis optical afterglows. In this section we examine one such facility on Mt. Palomar outside of San Diego, CA, which has been used extensively during observations presented in chapters 4 and 6.

2.4.1 Palomar Transient Factory

The intermediate Palomar Transient Factory (iPTF: Law et al. 2004) consists of three optical telescopes used for discovery (P48), followup (P60) and spectroscopy (P200). The P48 has a 7.1 deg^2 field of view and a 5σ optical sensitivity of ≈20 mag, with an integration time of 60 sec, followed by 30 sec for slewing. During the initial phase of the project, from 2009-2013, PTF spent most of its observing time at a 3-5 day cadence searching for supernovae and for variable sources in the Milky Way, but in its Intermeidate phase (late 2013-present) there have also been rapid searches at ∼2-hour cadences for fast relativistic transients. Two such optical transients have been discovered to date: PTF11agg, a good candidate "dirty fireball" explosion that may have resulted from a GRB progenitor with too much baryonic mass; and iPTF14yb, the first confirmed independent discovery of an on-axis long GRB by its optical afterglow

emission. iPTF14yb will be discussed in great detail in chapter 4.

Work has also been done by following up significant short GRBs from *Fermi* GBM with the iPTF facilities, in part as a "dress rehearsal" for follow-up of LIGO signal candidates (Singer et al., 2015). Ultimately, iPTF will give way to the Zwicky Transient Facility (ZTF), a survey whose camera will be sensitive in the $ugrizy$ bands down to 20.5 mag, but whose instantaneous FOV will total some 47 deg^2. With improvements in imaging and slew capability, the ZTF camera will be one of the best-equipped to quickly scan LIGO sky localization areas for electromagnetic counterparts.

2.5 The Big Picture

We have seen that major engineering breakthroughs in instrumentation will allow compact binary mergers to be observed, first in gravitational waves and then with high energy and optical instruments. The breadth and depth of coverage this makes possible is truly unprecedented in the history of astronomy; for the first time, we have the instrumentation and computing power to learn as much as possible about the most energetic transient explosions in the universe. With this in mind, we turn now to describe efforts at doing just that.

Part II

Never Ignore a Coincidence

(...unless you're busy, in which case, always ignore a coincidence)

Chapter 3

Rapid Identification of Electromagnetic Counterparts to Gravitational-wave Transients

Open here I flung the shutter, when, with many a flirt and flutter,
In there stepped a stately raven of the saintly days of yore.
Not the least obeisance made he; not a minute stopped or stayed he;
But, with mien of lord or lady, perched above my chamber door–
Perched upon a bust of Pallas just above my chamber door–
Perched, and sat, and nothing more.

Then this ebony bird beguiling my sad fancy into smiling,
By the grave and stern decorum of the countenance it wore,
'Though thy crest be shorn and shaven, thou,' I said, 'art sure no craven.
Ghastly grim and ancient raven wandering from the nightly shore–
Tell me what thy lordly name is on the Night's Plutonian shore!'
Quoth the raven, 'Nevermore.'

Edgar A. Poe, *The Raven*

Short-duration γ-ray bursts (GRBs) are incredibly brief ($\sim$0.01–2), ultra-luminous ($\sim 10^{47}$ erg s^{-1}) explosions in distant galaxies (redshifts $0.05 \lesssim z \lesssim 1$) that emit high-energy photons over a spectral range spanning from tens of keV to several MeV. There is some evidence that this population of bursts is associated with old elliptical galaxies, where star formation is no longer active and dead remnants of supernovae abound (see chapter 1). Properties in common with long-duration GRBs suggest a progenitor system that involves accretion of disrupted matter onto a newly formed black hole as the central engine driving the explosion, while the timescale, energetics and galaxy

demographics are suggestive of neutron star binary merger (Berger, 2014a; D'Avanzo, 2015). Like their long-duration analogues, most short GRBs are followed by bright, non-thermal "afterglows" in the X-ray ($\Delta t \sim$minutes), optical ($\Delta t \sim$hours to days), and radio ($\Delta t \sim$months to years) bandpasses.

If short GRBs are indeed caused by rapid accretion of displaced material during neutron star binary merger events, then there are several open questions as to the efficiency and dynamics of this process. Simultaneous observation of merger events in both gravitational waves (GWs) and high energy photons will do much to clarify these issues, especially if the afterglow emission can also be observed (Metzger & Berger, 2012). However, short GRBs are highly collimated: the ejecta move at initial Lorentz factors $\Gamma_0 \sim 10^2$ or more, so γ-ray emission is beamed in relativistic jets of half-opening angle $\theta_{\rm jet} \sim 20^\circ$, and the line of sight to Earth must be within this angle for emission to be viewable. For this reason the true cosmic event rate of short GRBs may be a factor $f_b \equiv (1 - \cos\langle\theta_{\rm jet}\rangle)^{-1} \sim$ 10–100 times higher than the *observed* rate. Afterglow emission is predicted to be visible from wider viewing angles as the ejecta slow down and expand laterally after interacting with the interstellar medium (Rhoads, 1999; Sari et al., 1999). The merger may also be accompanied by bright "kilonova" emission powered by r-process neucleosynthesis in the explosive ejecta within a few days of the initial burst (Metzger & Berger, 2012).

The Laser Interferometer Gravitational-wave Observatory (LIGO: LIGO Scientific Collaboration et al. 2015) network of detectors recently discovered transient source GW150914, a binary black hole merger at redshift $z = 0.09^{+0.03}_{-0.04}$ (Abbott et al., 2016c; The LIGO Scientific Collaboration & the Virgo Collaboration, 2016b), by directly observing GWs for the first time, confirming both that compact binary mergers do occur and that their merger can be observed with the LIGO instruments. (While short GRB emission is not expected from a system with two black holes, NASA's *Fermi* Gamma-ray Space Telescope did record a weak γ-ray transient some 0.4 seconds after merger; see Connaughton et al. 2016a.) In the first observing phase of its Advanced configuration, between 18 Sept 2015 and 13 Jan 2016, the two LIGO detectors in Hanford, WA (H1) and Livingston, LA (L1) were sensitive to NS-NS binary mergers out to 60–80 Mpc in the frequency range from 35 Hz to 8 kHz (chapter 2). Analysis of data from this first Observing Run (O1) of Advanced LIGO is ongoing (chapter 5). The detectors are currently offline, receiving another round of upgrades designed to improve their low-

frequency sensitivity down to ≈25 Hz and their distance sensitivity by 1–5%. A second Observing Run is planned for late July, with the Virgo detector (Acernese et al., 2015a) in Cascina, Italy joining perhaps as early as Dec 2016.

During O1, systems were in place to rapidly identify GW transients in the LIGO data stream. Separate efforts use both a matched filter technique with precise models of compact binary waveforms (*e.g.* Cannon et al. 2012a), and unmodelled techniques making minimal assumptions about waveform structure (Klimenko et al., 2008, 2016; Lynch et al., 2015). The matched filter searches identify both neutron star (NS) and black hole (BH) binary mergers and report signal candidates within 30-60 seconds of the measured merger time; while the unmodelled searches identify generic GW transients (including Galactic supernovae) within several minutes. Each of these searches is done "blindly," that is, across the entire sky and without being triggered by any other process. Fully coherent searches for GW signals triggered by specific GRBs are also done (Sutton et al., 2010; Williamson et al., 2014), but these are computationally expensive and may take several hours to estimate the significance of signal candidates.

In this chapter we present a method to rapidly identify associations between real-time GW signals and GRB detections. We rely on two high energy observatories in low-Earth orbit: the Burst Alert Telescope (BAT: Barthelmy et al. 2005) instrument aboard the *Swift* Gamma-Ray Burst Explorer and the Gamma-ray Burst Monitor (GBM: Meegan et al. 2009a) instrument aboard the *Fermi* Gamma-Ray Space Telescope. Each of these instruments publicly reports GRB discoveries over the Gamma-ray Coordination Network (GCN[1]) within 15–60 seconds of detecting a burst, making possible a rapid search for time and sky coincidence with LIGO signal candidates. In section 3.1 we lay out the scheme this search uses to assign significance to coincident events, describing its implementation and background estimation in section 3.2. We report the results of a mock data challenge in section 3.3, describing the sensitivity of this search in Gaussian noise compared to both an untriggered all-sky matched filter search and to the fully coherent method. Finally, in section 3.4 we offer concluding remarks and a look at the prospects of coincident detections in future observing runs.

[1] http://gcn.gsfc.nasa.gov

3.1 Statistical Framework

It is most natural to treat coincident observations of GRB and GW signals with a unified statistical formalism. Let generic data from interferometric GW detectors be represented by $\mathbf{g}$, and generic data from EM observations by $\mathbf{I}$. Since the only relevant information available in low latency is the time of arrival and the sky error region for each signal, we consider $\mathbf{g}$ and $\mathbf{I}$ to be *coincident* if one is observed to occur within a given window of time of the other, whence a more detailed test for coincidence in location on the sky is triggered. For short hard GRBs we accept time coincidences with a uniform prior if $\Delta t \in [-1\text{ s}, +5\text{ s}]$, where $\Delta t = t_{\mathbf{I}} - t_{\mathbf{g}}$ is the difference in estimated arrival times between the two signals.

We then define the joint likelihood ratio

$$\Lambda(\mathbf{g}, \mathbf{I}) = \frac{p(\mathbf{g}, \mathbf{I} \,|\, 1)}{p(\mathbf{g}, \mathbf{I} \,|\, 0)} \tag{3.1}$$

as a statistic comparing the probability density of observing coincident signals with data $(\mathbf{g}, \mathbf{I})$ assuming the signals have a common source (signified here by the binary value 1) to that assuming the apparent association is actually spurious (signified by 0). It is crucial to consider and interpret each of these scenarios very carefully. In the latter case, note that because any "false alarm" reported by this search amounts to an accidental association of two independent, unrelated events, there are essentially three possibilities:

1. The GW and the GRB are unrelated – but otherwise genuine – signals that, by pure happenstance, occur within six seconds of each other.

2. One of the signals is actually a background artifact, but the other is astrophysical.

3. Both signals are due to background noise and their coincidence in time is purely serendipitous.

Since NS-NS mergers occur at most only a few times per month out to 200 Mpc (Abadie et al., 2010) but GRBs (of either the long or short variety) and similar transients are detected ∼once a day or so (Piran 2004; Berger 2014a; see below), the first of these possibilities is exceedingly rare compared to the other two. Moreover, *Swift* BAT and *Fermi* GBM routinely report GRBs within ∼1 minute of the burst, and punctually retract occasional spurious GRB candidates within a few hours, so any association between some

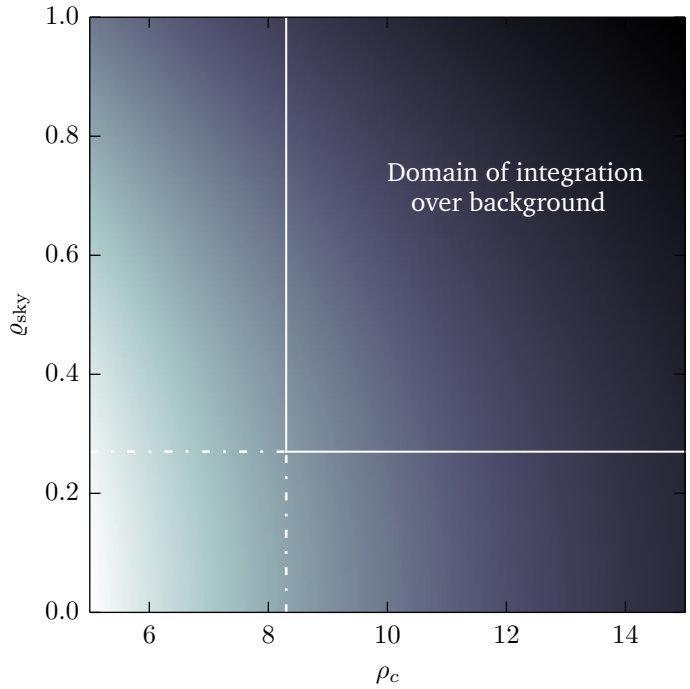

Figure 3.1: Toy model illustrating the integration problem of the search described in this paper. A simulated background distribution describing the fraction of noise events occurring within GRB on-source windows and binned in the network SNR ρ_c and sky overlap factor $\varrho_{\rm sky}$ is shaded in grayscale, in descending value from light to dark. The assumption is that time and sky coincidence are statistically independent statements, so $p(\rho_c, \varrho_{\rm sky} \,|\, n)$ is separable and the integration over background for a specific event is carried out over the rectangular domain indicated here.

GW signal candidate and a non-GRB event may be retracted on the same timescale. Therefore, the dominant region of parameter space under the null hypothesis is the one that asserts a false association between an astrophysical GRB and a background GW event, so that

$$p(\mathbf{g}, \mathbf{I} \,|\, 0) \simeq p(\mathbf{g} \,|\, n)\, p(\mathbf{I} \,|\, s) \tag{3.2}$$

where n and s refer to the noise and signal hypotheses in their respective contexts.

Now consider the numerator in Eq. (3.1). Assuming the signals are genuine and have a common astrophysical source, their location on the sky will (of course) be strongly correlated. The probability density of jointly observing two transients $\mathbf{g}, \mathbf{I}$ under this hypothesis is got by marginalizing over sky location, so that

$$p(\mathbf{g}, \mathbf{I} \,|\, 1) = \int_{S^2} p(\mathbf{g} \,|\, \hat{\Omega}, s)\, p(\mathbf{I} \,|\, \hat{\Omega}, s)\, p(\hat{\Omega} \,|\, 1)\, d^2\hat{\Omega} \tag{3.3}$$

where s refers to the astrophysical signal hypothesis separately for $\mathbf{g}$ and $\mathbf{I}$; $p(\hat{\Omega} \,|\, 1)$ is

a prior distribution for the sky location of the common signal source; and $p(\mathbf{g}\,|\,\hat{\Omega}, s)$, $p(\mathbf{I}\,|\,\hat{\Omega}, s)$ are independently measured by the two experiments. These functions are normalized to the parameter space probability densities $p(\mathbf{g}\,|\,s)$ and $p(\mathbf{I}\,|\,s)$ respectively. Normalizing instead to unity would result in the Bayesian posterior all-sky maps for each signal,

$$\mu_{\mathbf{g}}(\hat{\Omega}) = p(\hat{\Omega}\,|\,\mathbf{g}, s), \qquad \mu_{\mathbf{I}}(\hat{\Omega}) = p(\hat{\Omega}\,|\,\mathbf{I}, s) \tag{3.4}$$

in terms of which the probability density may be written

$$p(\mathbf{g}, \mathbf{I}\,|\,1) = p(\mathbf{g}\,|\,s)p(\mathbf{I}\,|\,s)\,\varrho_{\mathrm{sky}}(\mathbf{g}, \mathbf{I}) \tag{3.5}$$

where the quantity

$$\varrho_{\mathrm{sky}}(\mathbf{g}, \mathbf{I}) = (4\pi)^2 \int_{S^2} \mu_{\mathbf{g}}(\hat{\Omega})\,\mu_{\mathbf{I}}(\hat{\Omega})\,p(\hat{\Omega}\,|\,1)\,d^2\hat{\Omega} \tag{3.6}$$

is understood in a Bayesian way as the evidence factor for GW and EM signals having the same sky location. We use a uniform prior on sky location (that is, $p(\hat{\Omega}\,|\,1) = \mathrm{const.}$). In terms of quantities that are measured empirically, the joint likelihood ratio becomes

$$\Lambda(\mathbf{g}, \mathbf{I}) \simeq \Lambda_{\mathbf{g}}\,\varrho_{\mathrm{sky}}(\mathbf{g}, \mathbf{I}) \tag{3.7}$$

where

$$\Lambda_{\mathbf{g}} = \frac{p(\mathbf{g}\,|\,s)}{p(\mathbf{g}\,|\,n)} \tag{3.8}$$

is the likelihood ratio associated with the GW signal independently.

3.1.1 Assigning and Interpreting Significance

In any given search for signals buried in noise, one seeks a way to distinguish those signals from the background. In the context of the joint GW-EM coincidence search, the background is characterized by false associations between astrophysical GRBs, which are independently reported with high significance from a separate experiment, and noise artifacts from an untriggered LIGO-Virgo search. The proposed search is therefore hierarchical, and the joint likelihood ratio, Eq. (3.7), essentially re-ranks signal candidates based on their "loudness" (in terms of network SNR) as well as their coincidence in both time and sky position with the GRB, folding all of this information into a

single ranking statistic. Significance is then assigned based on the distribution $p(\Lambda \,|\, 0)$ of Λ for the background. In particular, the new p-value[2] (or "false alarm probability")

$$p = P(\Lambda \geq \Lambda^* \,|\, 0) = \int_{\Lambda^*}^{\infty} p(\Lambda \,|\, 0)\, d\Lambda \tag{3.9}$$

is the fraction of background events expected with joint likelihood ratio $\Lambda \geq \Lambda^*$, where Λ^* is the ranking statistic assigned to a specific pair of events. A detection threshold on the p-value is set prior to running the experiment in such a way that only some small fraction of events is likely to raise a false alarm on average (see *e.g.* Allen et al. 2012).

In previous LIGO and Virgo searches, the preferred way to report significance is not through the p-value directly, but through a statistic derived from it called the *false alarm rate* (FAR), which we write[3] as R. Assuming the background is Poisson-distributed with rate parameter

$$\lambda(\Lambda^*) = -\ln[1 - P(\Lambda \geq \Lambda^* \,|\, 0)], \tag{3.10}$$

one equates the measured p-value to the probability of a nonzero number of background events occuring with $\Lambda \geq \Lambda^*$:

$$P(\Lambda \geq \Lambda^* \,|\, 0) = \sum_{k=1}^{\infty} \frac{\lambda^k(\Lambda^*)}{k!} e^{-\lambda(\Lambda^*)} = 1 - e^{-\lambda(\Lambda^*)}. \tag{3.11}$$

The FAR statistic is then defined by

$$\lambda(\Lambda^*) = RT \tag{3.12}$$

where T is the total analyzed time in the search; that is, the total amount of time during which at least two interferometers in the LIGO/Virgo network were online and searching for signals.

In the joint coincidence search, we estimate the average number of accidental associations made between any single GRB and a noise glitch within some on-source time window τ according to

$$N_{\mathrm{acc}} \simeq \tau R_{\mathrm{GW}} \times \mathcal{C}_{\mathrm{acc}}(\varrho_{\mathrm{sky}}) \tag{3.13}$$

where R_{GW} is the false alarm rate reported by an untriggered search for gravitational waves alone and the quantity

$$\mathcal{C}_{\mathrm{acc}}(\varrho^*_{\mathrm{sky}}) = 1 - \mathrm{CDF}(\varrho^*_{\mathrm{sky}} \,|\, 0) \tag{3.14}$$

[2] In this section we must be careful to distinguish the p-value, which is just a number, from any probability distributions $p(x)$. The double-duty notation is unfortunate, but we are stuck with it.

[3] In some cases, it is useful to think in terms of the *inverse false alarm rate* (IFAR), which in this work we write as $\mathcal{I}$.

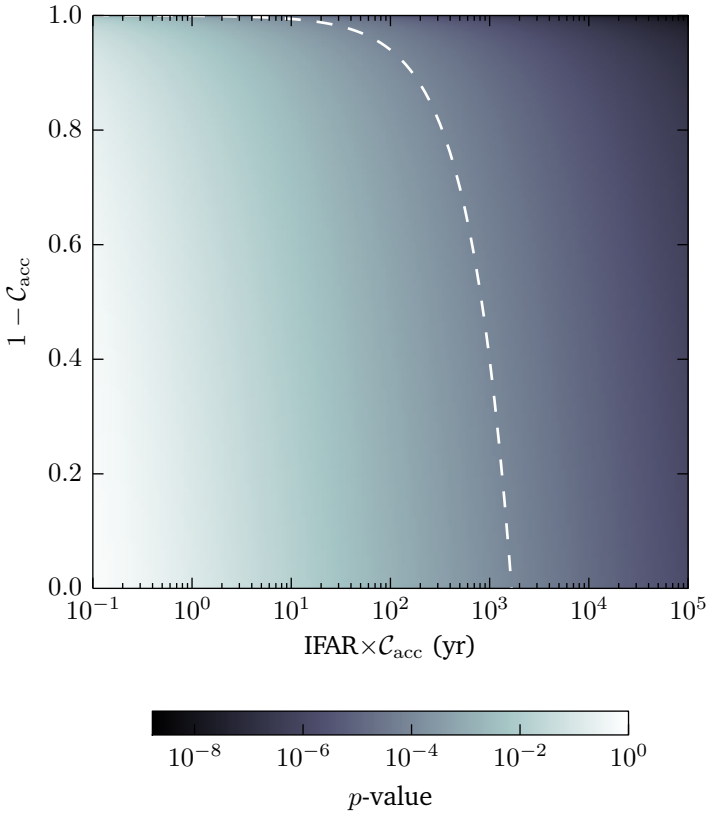

Figure 3.2: Illustration of the dependence of combined false alarm probability (p-value) on coincidence in time (quantified by the inverse coincident false alarm rate with sky factor removed, $\mathcal{IC}_{\rm acc}$) and coincidence on the sky (quantified by the accidental coincidence fraction $\mathcal{C}_{\rm acc}$). The p-value decreases toward the upper right-hand side of the plot, with a contour at $p = 10^{-4}$ shown as a dashed curve. Note that if a GW candidate is reported with high $R_{\rm GW}$ but precisely localized, then $\mathcal{C}_{\rm acc}$ can act to push the hierarchical FAR down dramatically.

is the cumulative fraction of accidental associations with $\varrho_{\rm sky} \geq \varrho^*_{\rm sky}$. Since the probability P of reporting a false association with a given GRB in this amount of time is

$$P(N_{\rm acc} \,|\, 0) = 1 - e^{-N_{\rm acc}}, \tag{3.15}$$

the significance we assign to each pair of events is the *coincident false alarm rate*,

$$R = \dot{N}_{\rm GRB}\, N_{\rm acc}. \tag{3.16}$$

Here $\dot{N}_{\rm GRB}$ corrects the trials factor (Lyons, 2008) for the annual rate of discovery of GRBs, so that if $n = \dot{N}_{\rm GRB} T$ GRBs are expected in the amount of time LIGO and Virgo

are searching for GW signals then

$$\begin{aligned} P(\Lambda \geq \Lambda^* \,|\, 0) &\simeq 1 - e^{-\lambda} \\ &= 1 - \left[1 - P(N_{\rm acc} \,|\, 0)\right]^n \\ &\simeq nP(N_{\rm acc} \,|\, 0) + \mathcal{O}\left[P^2(N_{\rm acc} \,|\, 0)\right] \end{aligned} \tag{3.17}$$

with $\lambda = RT$. In terms of reporting associations to astronomers, the coincident false alarm rate is the favored ranking statistic because it provides a quantitative way to mitigate spurious artifacts.

3.1.2 The Rate of *Swift* and *Fermi* GRBs

To accurately estimate the background of the coincidence search, one must measure the rate $\dot{N}_{\rm GRB}$ at which both long and short GRBs are discovered and reported in real-time. While the *Swift* satellite launched in late 2004 and has been steadily reporting GRBs ever since, the true, current, unbiased rate will also include those discovered by the *Fermi* spacecraft, which launched in 2008. Using the High Energy Astrophysics Science Archive Research Center (HEASARC[4]) database, we have obtained a complete list of all GRBs discovered by either observatory between 2008 July 14 UT (the day *Fermi* GBM discovered its first GRB) and 2015 September 1. There were some 444 GRBs reported by *Swift* BAT during this period, and 2102 reported independently between both instruments. A histogram of the time difference, Δt, between GRBs in this population is shown in Fig. 3.3. Clearly, the discovery of GRBs by either of these satellites is well-modelled as a Poisson process, with a measured rate parameter $\dot{N}_{\rm GRB} = 0.807 \pm 0.018$ day^{-1}.

3.2 Implementation of the Search

In practice, then, the proposed real-time coincidence search pipeline proceeds as follows. For compact binary signals, the two factors appearing in Eq. (3.7) may each be calculated within a minute or so of the merger event if communication is open between LIGO-Virgo and electromagnetic (EM) astronomers. EM transients will have

[4] https://heasarc.gsfc.nasa.gov

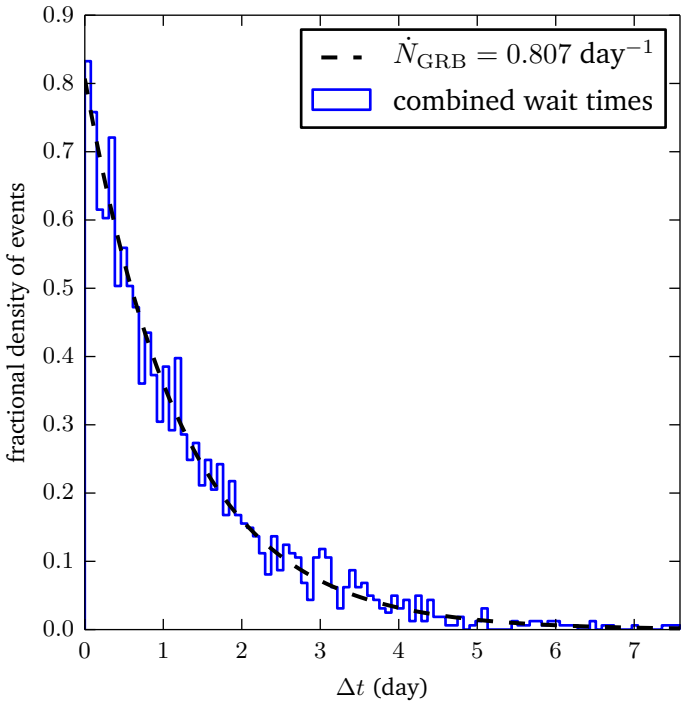

Figure 3.3: Histogram of wait times between GRBs discovered by *Swift* and *Fermi* between 2008 July 14 and 2015 September 1. The fit is to an Erlang distribution with rate parameter $\dot{N}_{\rm GRB} = 0.807$ day^{-1}, the measured rate of independently-discovered GRBs by either spacecraft. GRBs which were jointly discovered by both *Swift* and *Fermi* are counted only once.

been detected by either of two space-based observatories (the *Swift* BAT or the *Fermi* GBM instruments), identified as likely GRBs, and finally transmitted along with information on sky location and detection significance over the GCN system. These GCN Notices are then automatically filtered and entered into a centralized database (the Gravitational-wave Candidate Event Database, GraceDb[5]; see chapter 2), where GW candidates identified by the LIGO and Virgo collaborations' own online data analysis pipelines are also stored. Whenever either type of trigger lands in GraceDb – the GRBs from *Swift* and *Fermi* or the GW candidates identified by LIGO – an automated monitor responds immediately, looking for coincident events of the opposite type within the 6-second window discussed above (see Fig. 3.4). We refer to this algorithm and the software that implements it as the Rapid, on-source VOEvent[6] Coincidence Monitor (RAVEN).

[5] https://gracedb.ligo.org

[6] GCN Notices are broadcast as machine-readable XML files known as VOEvents. This format is based on the International Virtual Observatory Alliance; see http://www.ivoa.net for more information.

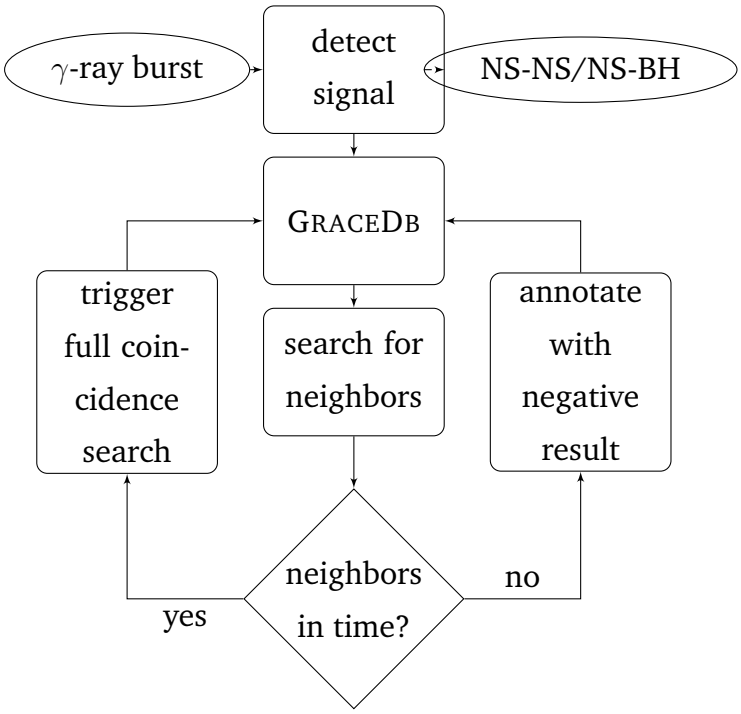

Figure 3.4: Organizational flow chart of the proposed low-latency coincidence pipeline. Detections of the EM and GW signals are each uploaded to a central database (GRACEDB), at which point the coincidence search is triggered automatically. When the pipeline finishes running, every relevant entry in the database is annotated with its results.

Note that the procedure described here can easily be generalized to find associations between arbitrary GW and EM transients, as long as some time window τ is set beforehand. For instance, the LIGO and Virgo interferometers are also sensitive to supernovae in the Milky Way galaxy, and these may produce long-duration GRBs under the right conditions (see *e.g.* Piran 2004). Real-time data analysis pipelines searching for generic GW transients will identify supernova signatures within several minutes (Klimenko et al., 2008, 2016; Lynch et al., 2015) and RAVEN is fully capable of identifying these with both long GRBs and with neutrino emission reported by sites in the SuperNova Early Warning System (SNEWS[7]). In fact, all of this was done as part of regular monitoring during O1. In this work, however, we focus only on the special case of compact binary mergers.

[7] http://snews.bnl.gov

3.2.1 Calculation of $\varrho_{\rm sky}$

To implement the RAVEN pipeline, one must inevitably approximate the integral in Eq. (3.6) numerically. For the purposes of this paper and during online searches for compact binary coalescence we use the BAYESTAR rapid sky localization algorithm (Singer & Price, 2016), which employs Bayesian statistics to localize GW signal candidates on the sky from information on time-of-arrival, phase-on-arrival and network SNR. This code stores posterior distributions in a discretized format using the Flexible Image Transport System (FITS[8]), with the unit sphere pixelated according to the Hierarchical Equal Area isoLatitude Pixelization (HEALPix: Górski et al. 2005) scheme. Each pixel in the map is assigned a probability mass approximating the integral of $\mu_{\mathbf{g}}$ over the area of that pixel. In other words, the GW sky map $\mu_{\mathbf{g}}(\hat{\Omega})$ is approximated by a discrete list $\mu_{\mathbf{g},i}$ of probability masses corresponding to the pixels making up an image of the map, each of which have the same area in the HEALPix projection. The resolution (*i.e.*, the pixel size) of each GW sky map scales inversely with the SNR of its corresponding trigger.

For the sake of consistency and convenience, we use a similar procedure to generate the GRB all-sky map, approximating the continuous function $\mu_{\mathbf{I}}(\hat{\Omega})$ by a discrete list of probability masses $\mu_{\mathbf{I},i}$ with the same angular resolution as the GW maps. We always have the freedom to control this resolution because of the algorithm used to construct the sky map: when a GRB is detected, an inferred sky location $\hat{\alpha}$ and error radius σ are reported by the γ-ray instrument and known to astronomers as soon as news arrives of the event. In principle, the appropriate thing to do is to fit this to a wrapped Gaussian whose mean direction is that of the unit vector $\hat{\alpha}$, which we approximate by the 2-dimensional von Mises-Fisher distribution (see *e.g.* Kanti & Jupp 1999)

$$\mu_{\mathbf{I}}(\hat{\Omega};\hat{\alpha},\kappa) = \frac{\kappa}{4\pi\sinh\kappa}\exp\left(\kappa\hat{\alpha}\cdot\hat{\Omega}\right). \tag{3.18}$$

Here κ is a "concentration parameter" directly related to the error radius by

$$\frac{I_{3/2}(\kappa)}{I_{1/2}(\kappa)} = \coth\kappa - \kappa^{-1} = \exp\left(-\sigma^2/2\right) \tag{3.19}$$

where I_α is the modified Bessel function of the first kind with index α. In the limit $\sigma \to 0$, which is appropriate for most GRBs, the concentration parameter $\kappa \to \infty$ and Eq. (3.18) reduces to a multivariate Gaussian with variance $\sigma^2 \simeq 1/\kappa$.

[8] http://fits.gsfc.nasa.gov

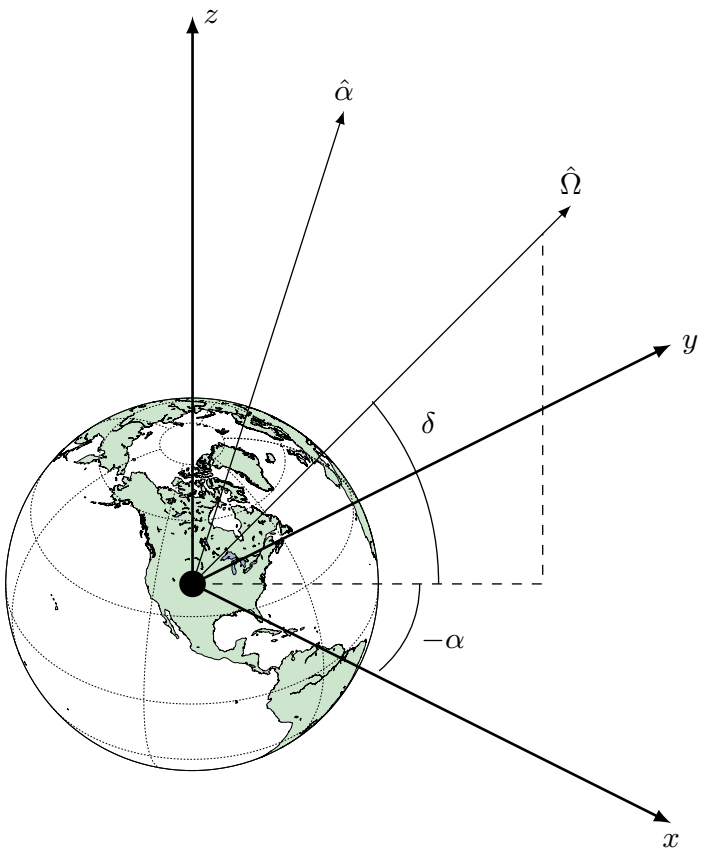

Figure 3.5: Diagram of the geometry of directional statistical quantities. When a GRB is reported it is accompanied by an estimate of the measurement uncertainty (encoded by the error radius σ) and an inferred sky location, which we take to be in the direction of the unit vector $\hat{\alpha}$. The unit vector $\hat{\Omega}$ is the random variable described by either of the sky maps $\mu_{\mathbf{g}}$ or $\mu_{\mathbf{I}}$. (Up to selection of an arbitrary prime meridian, the usual right ascension and declination coordinates (α, δ) are shown for comparison).

There is a minor but computationally significant problem: in the von Mises-Fisher distribution, Eq. (3.18), the maximum value of the PDF scales like $\kappa/2\pi$ as $\kappa \to \infty$ and the direct operation $\exp(\kappa)$ becomes intractable (that is, as the support of the PDF shrinks to the size of one element on a finely pixelated sphere, it more closely resembles a delta function and $\exp(\kappa)$ results in a value of `inf` or `NaN`). But there is a simple workaround. If the center of the ith pixel lies in the direction $\hat{\Omega}_i$, then we avoid overflow errors by first calculating

$$\xi_i = \kappa \left(\hat{\alpha} \cdot \hat{\Omega}_i - 1 \right) \tag{3.20}$$

and then using $\{\xi_i\}$ to generate the normalized sky map

$$\mu_{\mathbf{I},i} = \frac{N}{4\pi} \frac{\exp\left(\xi_i\right)}{\sum_j \exp\left(\xi_j\right)} \tag{3.21}$$

where N is the total number of pixels on the sky (each having area $4\pi/N$). This series of calculations is valid because the discrete distribution is computed up to a constant factor in the first step, Eq. (3.20), and normalized in the second step, Eq. (3.21). With a uniform prior

$$p(\hat{\Omega} \,|\, 1) = 1/4\pi$$

and dropping unnecessary factors of 4π one may then approximate $\varrho_{\rm sky}$ numerically by

$$\varrho_{\rm sky} \simeq \frac{4\pi}{N} \sum_{i=1}^{N} \mu_{\mathbf{g},i}\, \mu_{\mathbf{I},i}. \tag{3.22}$$

Empirically, *Swift* GRBs tend to have narrow localization regions whose support covers less than the area of a single pixel in even the most finely pixelated GW sky maps, while *Fermi* GRBs have error boxes that can span tens to hundreds of square degrees. This is due to the nature of the localization technique used on each satellite: *Swift* BAT employs a coded aperture mask over a plane of CdZnTe hard X-ray detector tiles (Barthelmy et al., 2005), inferring the direction to source from the "shadow" it casts over the detector plane; while *Fermi* GBM uses four sets of triplet NaI scintillators and two BGO crystals latched to the side of the spacecraft to cover three orthogonal directions, using triangulation to infer the direction to source (Meegan et al., 2009a). In the case of *Swift* BAT, localization is excellent but short GRB detection rates are low because they occur on timescales shorter than the spacecraft can slew. The BAT instrument is also relatively insensitive over the higher energy bandpass where short GRBs typically emit most of their photons. Localization by *Fermi* GBM is less precise and varies with the SNR of the signal, but sky coverage is nearly complete (GBM has access to nearly every part of the sky not occluded by Earth) and the instrument is much more sensitive to the systematically harder photons from short GRBs. *Fermi* GBM detects short GRBs at a rate of $\sim$0.5–1 week^{-1} on average.

A typical example showing the localization abilities of each spacecraft is GRB 131105A, a long-duration burst detected by both *Swift* BAT (Cummings et al., 2013) and *Fermi* GBM (Fitzpatrick & Jenke, 2013). Each of these detections was reported as a GCN Notice within tens of seconds, with a significance of 10.8σ reported by *Swift* and 11.7σ reported by *Fermi*. The *Swift* BAT error circle is $6'$ wide, while the *Fermi* GBM error radius was reported as $6.18°$.

3.2.2 Background Estimation

The dominant computational cost in doing a search for time- and sky-coincident events is in performing background estimation. By construction, the time-coincident piece of the background is already contained in the rate of all-sky false alarm events R_{GW}, the on-source time window τ and the rate $\dot{N}_{\mathrm{GRB}}$ of discovery of GRBs. Therefore the only open question is to how accidental associations between GRBs and background events compare on the sky. This question is best addressed by estimating the fraction $\mathcal{C}_{\mathrm{acc}}(\varrho_{\mathrm{sky}})$ of accidental association events whose sky maps overlap to a degree better than ϱ_{sky}, Eq. (3.14).

Since any pair of unrelated events will have statistically independent sky maps, we compare ϱ_{sky} to the value it would have taken if the GRB had been observed to come from any random direction on the sky. This immediately suggests performing a convolution, although because of the directional nature of the problem, we must be scrupulously careful about defining the domain of integration. What we wish to compute is ϱ_{sky} as a function of $\hat{\alpha}$,

$$\varrho_{\mathrm{sky}}(\hat{\alpha}) = 4\pi \int_{S^2} \mu_{\mathbf{I}}(\hat{\Omega}; \hat{\alpha}, \kappa)\, \mu_{\mathbf{g}}(\hat{\Omega})\, d^2\hat{\Omega}, \tag{3.23}$$

at fixed concentration parameter κ (assuming for simplicity that the GRB would have been localized with identical precision regardless of what direction it came from). Then, for points uniformly sampled on S^2, the 1-dimensional distribution of values of ϱ_{sky} is

$$\mathcal{C}_{\mathrm{acc}}(\varrho_{\mathrm{sky}}) = \frac{1}{4\pi} \int_{D(\varrho_{\mathrm{sky}})} d^2\hat{\alpha} \tag{3.24}$$

where the (not necessarily simply connected) domain

$$D(\varrho^*_{\mathrm{sky}}) = \left\{ \hat{\alpha} \,|\, \varrho_{\mathrm{sky}}(\hat{\alpha}) \geq \varrho^*_{\mathrm{sky}} \right\}$$

is a subset of S^2 bounded by level curves of $\varrho_{\mathrm{sky}}(\hat{\alpha})$. Put another way, after convolving both sky maps we first compute $D(\varrho_{\mathrm{sky}})$ as the smallest contour surrounding both the convolution's maximum value and the reported location $\hat{\alpha}$ of the GRB – the integral in Eq. (3.24) – and then we measure $\mathcal{C}_{\mathrm{acc}}$ geometrically as the ratio between the area of this contour and the area of the entire sky.

As mentioned, taking the convolution (3.23) can be an expensive step if done by sheer brute force. However, because the von Mises-Fisher PDF (3.18) is a function only

of $\hat{\alpha} \cdot \hat{\Omega}$ at fixed κ, there is guaranteed to be a rotation of S^2 that sets $\hat{\alpha}$ to be the North Pole (for example) and makes $\mu_{\mathbf{I}}$ azimuthally symmetric. In terms of a decomposition into the standard spherical harmonic modes $Y_{lm}(\hat{\Omega})$, this means there always exists a coordinate chart in which $\mu_{\mathbf{I}}$ has contributions to its spectrum only from the zonal modes where $m = 0$:

$$\mu_{\mathbf{I}}(\hat{\Omega}; \hat{\alpha}, \kappa) = \sum_{l=0}^{\infty} k_l Y_{l0}(\hat{\Omega})$$

for expansion coefficients k_l. If $\mu_{\mathbf{g}}$ is expanded in this chart (with spectral contributions generally distributed across all modes) as

$$\mu_{\mathbf{g}}(\hat{\Omega}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} a_{lm} Y_{lm}(\hat{\Omega})$$

for coefficients a_{lm} then the 2D convolution is given by the Funk-Hecke theorem (Basri & Jacobs, 2003) as

$$\varrho_{\text{sky}}(\hat{\alpha}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sqrt{\frac{(4\pi)^3}{2l+1}} k_l a_{lm} Y_{lm}(\hat{\alpha}). \tag{3.25}$$

This statement is analogous to the familiar convolution theorem on $\mathbb{R}$, and has a similar numerical implementation – reducing the number of calculations needed from $\mathcal{O}[N^2]$ to $\mathcal{O}[N \log N]$, where N is the number of pixels in each sky map.

Several remarks are worth making at this stage. First, the von Mises-Fisher expansion coefficients k_l are found to scale as

$$k_l \simeq \exp\left[-\frac{l(l+1)}{2\kappa}\right]$$

for $l \gg 1$, which is almost always the case for GRB sky maps. Thus, the von Mises-Fisher distribution acts as a low-pass filter on the spectral modes of $\mu_{\mathbf{g}}$ (analogous to frequency domain filtering with a 1D Gaussian), and we are able to truncate the sum at some $l_{\max}$ to get

$$\varrho_{\text{sky}}(\hat{\alpha}_i) \simeq \sum_{l=0}^{l_{\max}} \sum_{m=-l}^{l} \sqrt{\frac{(4\pi)^3}{2l+1}} k_l a_{lm} Y_{lm}(\hat{\alpha}_i) \tag{3.26}$$

for the mean direction $\hat{\alpha}_i$ projected onto the ith pixel. In practice, we choose $l_{\max}$ to scale with the number N of pixels in the discretized unit sphere, then use the cumulative histogram of the $\varrho_{\text{sky}}(\hat{\alpha}_i)$ over all N pixels to measure $\mathcal{C}_{\text{acc}}(\varrho_{\text{sky}})$. Second, because of the azimuthal symmetry of the von Mises-Fisher PDF, it also acts as a smoothing

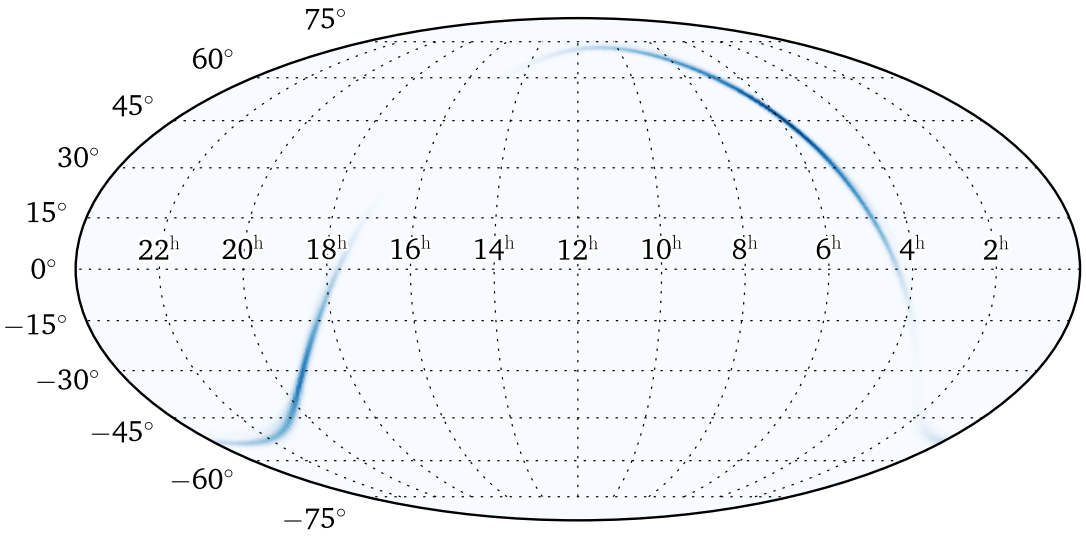

(a) Gravitational wave sky map

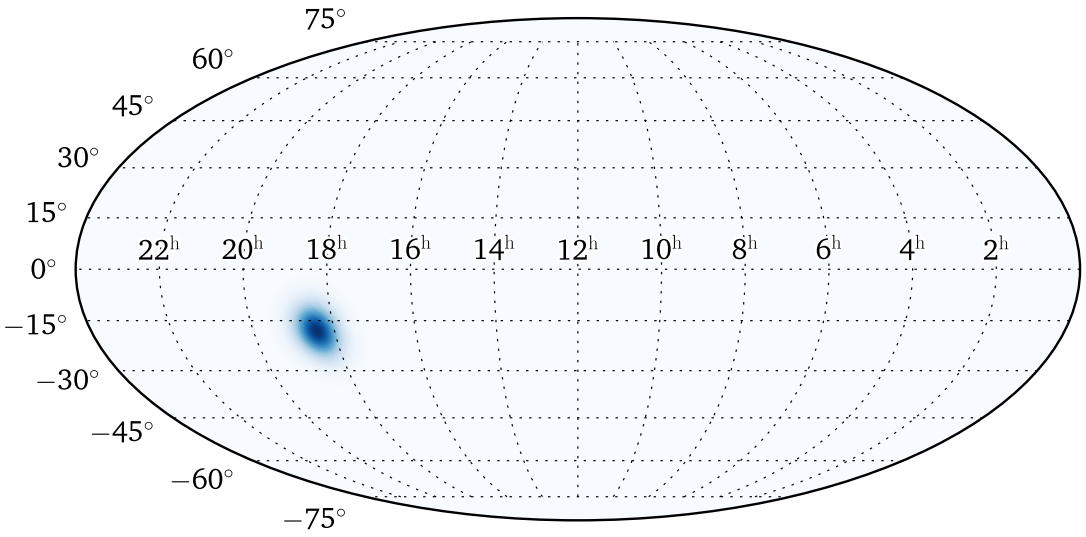

(b) GRB sky map (simulated with $\sigma = 5°$)

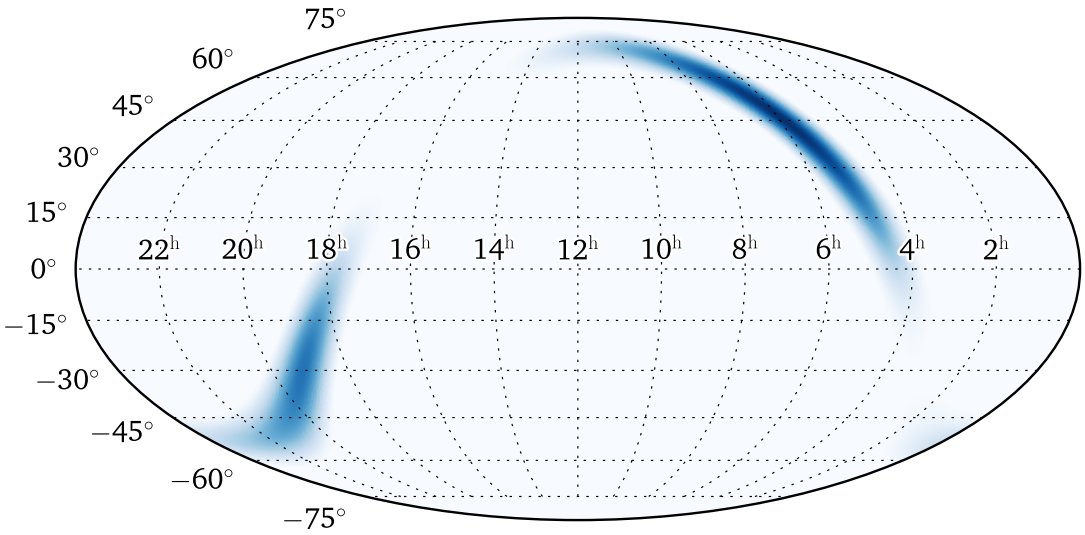

(c) Convolution of (a) with (b)

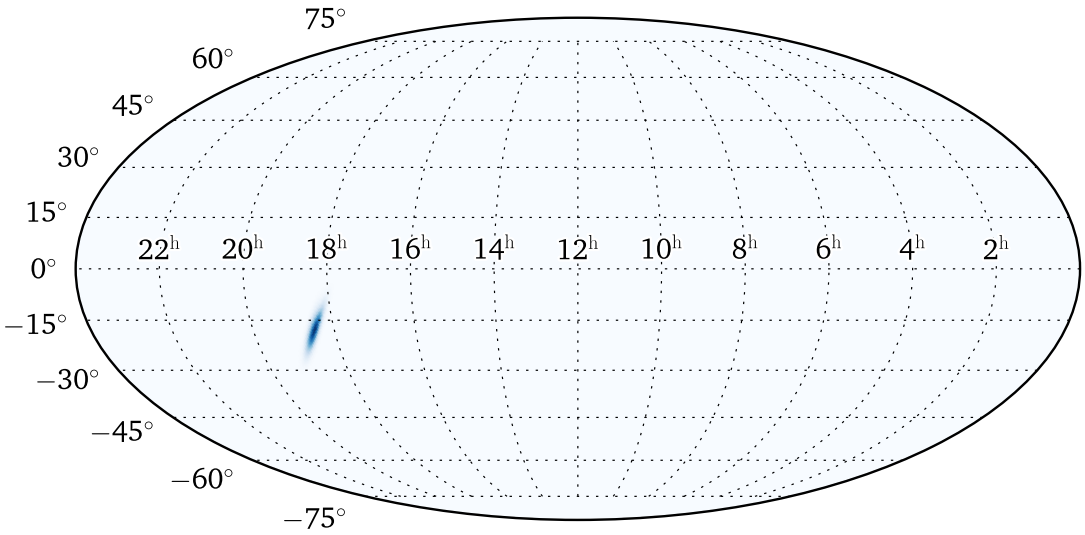

(d) Combined sky map, $p(\hat{\Omega} \,|\, \mathbf{g}, \mathbf{I}, s)$

Figure 3.6: Sky location posterior for (a) a simulated GW signal recovered with network SNR $\simeq 12.5$ in Gaussian noise with 2015-era sensitivity (computed with BAYESTAR); (b) A simulated GRB accompanying this signal with on-sky precision $\sigma = 5°$, typical for *Fermi* GBM; (c) the convolution of (a) with (b); and (d) the joint sky map, all in equatorial (J2000.0) coordinates. The GW posterior has two long "islands" forming a broken triangulation ring, typifical of $\approx 50\%$ of GW sky maps (see Singer et al. 2014). The *Fermi* GBM-like localization provides a joint map significantly smaller than either of the individual ones.

kernel on $\mu_{\mathbf{g}}$ (not unlike the 2D Guassian smoothing commonly used in image processing), with the degree of smoothing set by the concentration parameter κ. (As $\kappa \to \infty$ the kernel is a δ-function that returns $\mu_{\mathbf{g}}$ exactly, while for $\kappa \to 0$ the kernel smears $\mu_{\mathbf{g}}$ evenly across the entire sky – see Fig. 3.6.) Finally, we see that the statistic ϱ_{sky} is proportional to the evidence factor in a joint Bayesian sky map combining data from each experiment:

$$p(\hat{\Omega} \,|\, \mathbf{g}, \mathbf{I}, s) = \frac{\mu_{\mathbf{g}}(\hat{\Omega})\,\mu_{\mathbf{I}}(\hat{\Omega}; \hat{\alpha}, \kappa)}{\varrho_{\mathrm{sky}}(\hat{\alpha})/4\pi}. \tag{3.27}$$

Because $\mu_{\mathbf{g}}$ and $\mu_{\mathbf{I}}$ have radically different shapes, the joint sky map typically provides a quite dramatic improvement over one or both of them, and can be made available as a data product of the coincidence search.

All pieces are now in place to compute the ranking statistic (3.16) for each coincident candidate, and provide an improved inference of sky location. In case of a single event having multiple coincident neighbors (as would happen if the same GW signal candidate is caught by multiple LIGO/Virgo data analysis pipelines, or if two satellites separately observe the same GRB), the RAVEN coincidence pipeline will run on all possible associations, and action may be taken appropriately. The results of the search are to be made available to the wider astronomical community through LIGO and Virgo's connection to the GCN broadcasting service. The total estimated latency for this coincidence pipeline has been on the order of 1–2 minutes after merger during LIGO/Virgo engineering periods and during O1. All of this bodes very well for directed X-ray and optical followup programs. We now turn to the question of this pipeline's performance compared to a blind search for GW signals alone.

3.3 Simulation

To gauge the efficacy and efficiency of the proposed real-time coincidence pipeline, we have arranged a simulation echoing the procedure used during a recent large-scale mock data challenge (Singer et al., 2014). Two simulated observing runs were performed using Gaussian noise modelled on the 2015- and 2016-era LIGO/Virgo commissioning schedule (LIGO Scientific Collaboration et al., 2015; Acernese et al., 2015a). In the 2015 case, the GW detector network consists of Hanford (H1) and Livingston (L1), while the 2016 scenario includes H1, L1 and Virgo. This setup is meant to re-

flect and anticipate electromagnetic followup of gravitational wave signals in the early advanced detector era, with two important differences expected in 2016 compared to 2015: first, an increase in sensitivity in the H1-L1 detector network (so that a greater number of cycles may be observed during compact binary inspirals), and second, the addition of the Virgo detector with roughly half the relative strain sensitivity of H1 and L1. In both scenarios, a 60-day segment of data containing only a realization of Gaussian noise was prepared from the power spectral densities expected in each detector in the relevant era (Abbott et al. 2016d; see Fig. 3.7). Then, a population of some ∼50,000 NS-NS binary inspiral signals is injected into these data (with separate signal populations for the 2015 and 2016 scenarios). Signal parameters are drawn from astrophysically motivated mass (m_1, m_2) and spin (χ_1, χ_2) distributions and with polarization angle (ψ), coalescence phase (ϕ_c), time of arrival at geocenter, and sky location (α, δ) drawn from the appropriate uniform distributions. Because we are interested in short GRBs, in this study, the inclination angle ι is drawn from a uniform distribution between $0° \leq \iota \leq 30°$. The (luminosity) distance to source is drawn from a distribution that scales as r^2 so that the signal population is uniform in 3-dimensional volume. (In 2015, the maximum injected distance is 220 Mpc; in 2016, it is 440 Mpc.) We then mean to characterize the population of injected signals against the population of background events (*e.g.* stationary and Gaussian noise artifacts in the detectors).

The real-time data analysis pipeline GSTLAL (Cannon et al., 2012a) was run on the full data stream, with and without injected signals, in both the 2015 and 2016 scenarios. In its online mode, this pipeline is designed to rapidly ($\lesssim$30–60 sec) identify signal candidates with the lowest combined false alarm rate ($R_{\rm GW}$) and highest network SNR out of a cluster of events coincident within the light-travel time of every pair of GW detectors in the network. The "time of arrival" reported by GSTLAL is estimated by analyzing the times at which coincident signals in two or more detectors have been identified as having the highest single-detector SNR, with χ^2 tests done as a signal consistency check (Allen, 2005). Candidates are then localized on the sky with BAYESTAR; the properties and morphologies of the resulting posterior sky maps were analyzed, along with a discussion of strategies for astronomical followup based on these results, in Singer et al. (2014). However, this study focused exclusively on signal candidates with p-values above the 5σ detection threshold. In an online search BAYESTAR produces posteriors for *all* signal candidates reported to GRACEDB, and the

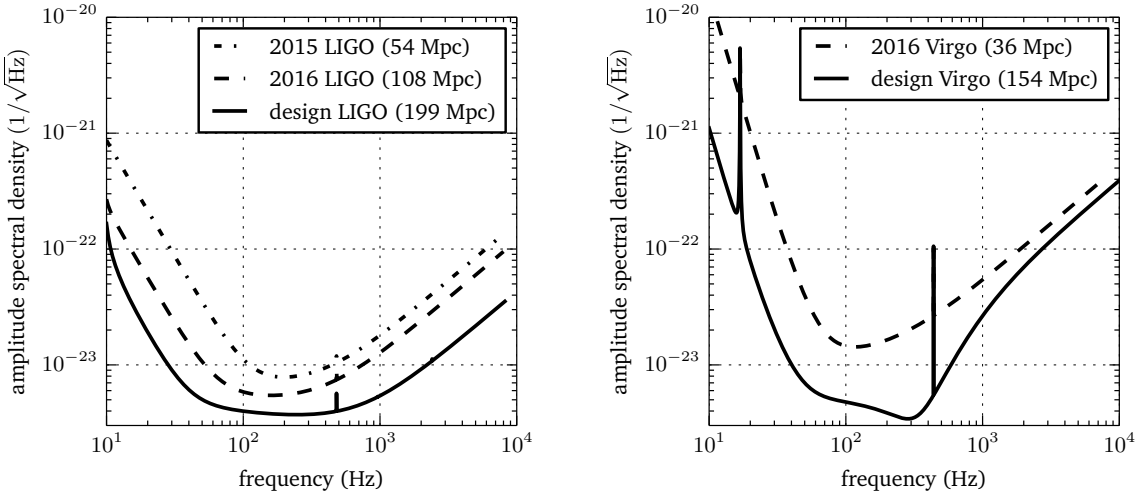

(a) 2015 (b) 2016

Figure 3.7: Model detector noise amplitude spectral density curves. The LIGO 2015, 2016, and final design noise curves are shown in the left panel and the Virgo 2016 and final design noise curves in the right panel. The angle- and orientation-averaged $\rho_c = 8$ sensitive distance range for (1.4, 1.4) $M_\odot$ NS-NS mergers is given for each detector; note that in this simulation, the sensitive distance will be enhanced by the fact that every simulated binary is face-on ($\iota \leq 30°$). These amplitude spectral density curves are identical to those used in the study reported in Singer et al. (2014).

vast preponderance of these have very low significance (and low network SNR). It remains to be seen how the sky maps of lower-significance candidates and background events compare with those of "gold-plated" detections.

We endeavor to address part of this question by simulating a population of *Swift* short GRBs matching the population of injected GW signals. For the *Swift* BAT detection time, we select a random value within the on-source window $[-1\text{ s}, +5\text{ s}]$ around the geocentric merger time of the every GW signal. Since *Swift* BAT typically reports GRBs with an on-sky error radius on the scale of arcminutes (Barthelmy et al., 2005), we also select the true injected sky position with a very small error radius of 0.05 degrees, so that the reconstructed GRB sky map $\mu_{\mathbf{I}}(\hat{\Omega})$ behaves like a δ-function picking out the true location of the source. With correlated populations of GW signals and GRBs generated from the same list of injections (Fig. 3.8), and assuming *Swift* BAT has 100% efficiency to GRBs that occur within LIGO/Virgo's sensitive volume, we then run the RAVEN coincidence search in batch mode, looking for time coincidences between GWs and GRBs and triggering a significance calculation with timing and sky information if a neighbor is found. (An example sky posterior for a NS-NS signal injection recovered

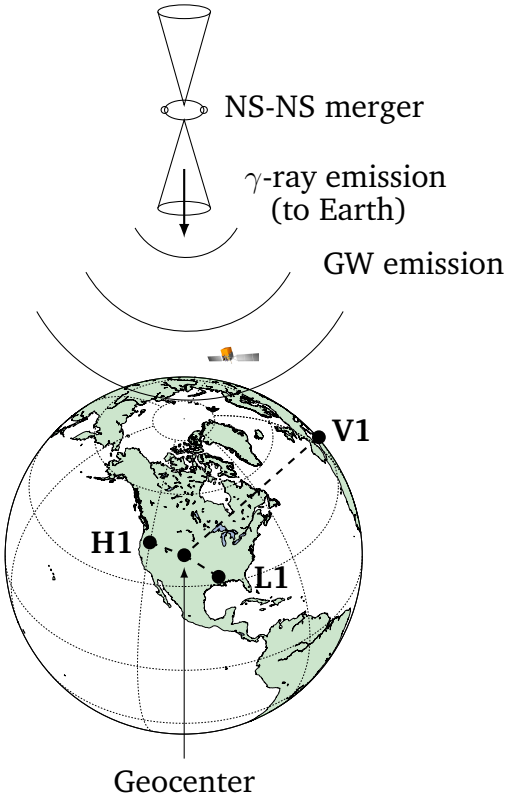

Figure 3.8: Schematic diagram of the joint GRB-GW detection geometry. A coalescing NS-NS binary produces a γ-ray burst visible to the *Swift* BAT instrument, and a corresponding GW signal is visible to LIGO and Virgo. The actual time of coalescence of the binary is geocentered – that is, interpreted as a measurement in the rest frame of the Earth at the moment the GW signal passes through its center – and then injected into the data stream of each interferometer, accounting for instrument response and for light-travel time. The γ-ray emission is then simulated as a GRB reported by *Swift* BAT at a time randomly selected within a window of $[-1, +5]$ seconds around the actual geocentric merger time of the injected GW signal. The RAVEN coincidence pipeline responds according to the workflow in Fig. 3.4.

with low SNR is shown in Fig. 3.6(a) with the GRB location overlaid)

To get a sufficiently large sample size ($\sim$50,000 in 60 days of analysis time), we have prepared NS-NS inspiral signals at an unrealistically high Poisson event rate. This is justified because (1) the point of this study is to quantify statistical behavior of the noise background compared to astrophysical sources in RAVEN's search method; and (2) this is the most computationally efficient method of implementing the search with a controlled source population. Our simulation methodology is equivalent to repeating the same experiment many times (each with identical noise realizations) over an ensemble average of astrophysical events. However, not all injected signals are recovered by the analysis because GW signal candidates are only saved if they have a network SNR of 4 or greater in each detector within some small time window. Some low-SNR candi-

dates may also suffer correspondingly poor timing estimation by the GSTLAL pipeline; thus, although all short GRBs are generated from injected GW signals, only a subset will have counterpart GW candidates identified by GSTLAL.

3.3.1 Results

Because our simulation methodology is designed so that GSTLAL has access to a data stream wholly devoid of astrophysical signals, we can cleanly characterize the noise of this search without the contaminating effects of astrophysics. In principle, one would expect that an untriggered GW experiment whose time duration is T ought to observe

$$\bar{N}(R_{\rm GW}) = TR_{\rm GW} = \frac{T}{\mathcal{I}_{\rm GW}} \tag{3.28}$$

noise transients with false alarm rate $R_{\rm GW}$ (inverse FAR $\mathcal{I}_{\rm GW}$) or lower (higher). Fig. 3.9 compares the number of background events at or above a given $\mathcal{I}_{\rm GW}$ in both the 2015 and 2016 scenarios. We find, unsurprisingly, that both searches report background events that follow the power law distribution, Eq. (3.28), up to small-number statistics at high $\mathcal{I}_{\rm GW}$.

As noted in section 3.1.2, to very good approximation the process of jointly discovering GRBs by *Swift* BAT and *Fermi* GBM is a Poisson process with rate $\dot{N}_{\rm GRB} = 0.807$ day^{-1}. To simulate the RAVEN background, we need only model a population of *Swift/Fermi* GRBs wholly unrelated to any noise transients in the LIGO/Virgo data streams. We accomplish this by drawing $N_{\rm GRB}$ random event times during the analysis period, where $N_{\rm GRB}$ itself is drawn from a Poisson distribution with rate parameter $\lambda = \dot{N}_{\rm GRB}T$. For each "GRB," we then select the most significant noise transient within the 6-second time window described above. To accumulate a suitably large population of coincident noise events, we repeat this process a large number (∼20,000) of times, normalizing to the total number of noise transients reported by GSTLAL over the entire analysis period. We find the result of this simulation is consistent with the expected fraction of coincident background events,

$$\bar{N}(R) \propto R = \frac{1}{\mathcal{I}} = \frac{\tau\dot{N}_{\rm GRB}}{\mathcal{I}_{\rm GW}}, \tag{3.29}$$

in both 2015 and 2016.

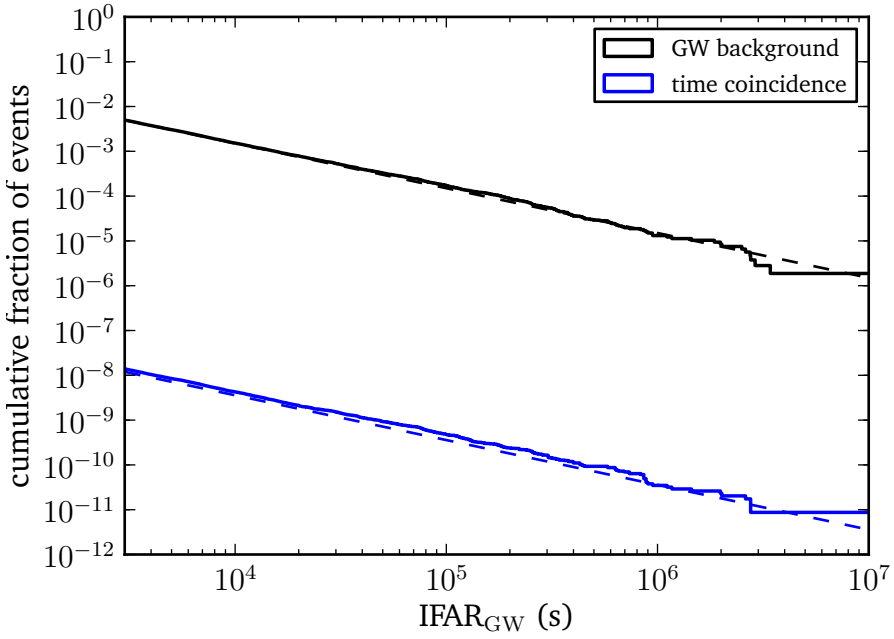

(a) 2015

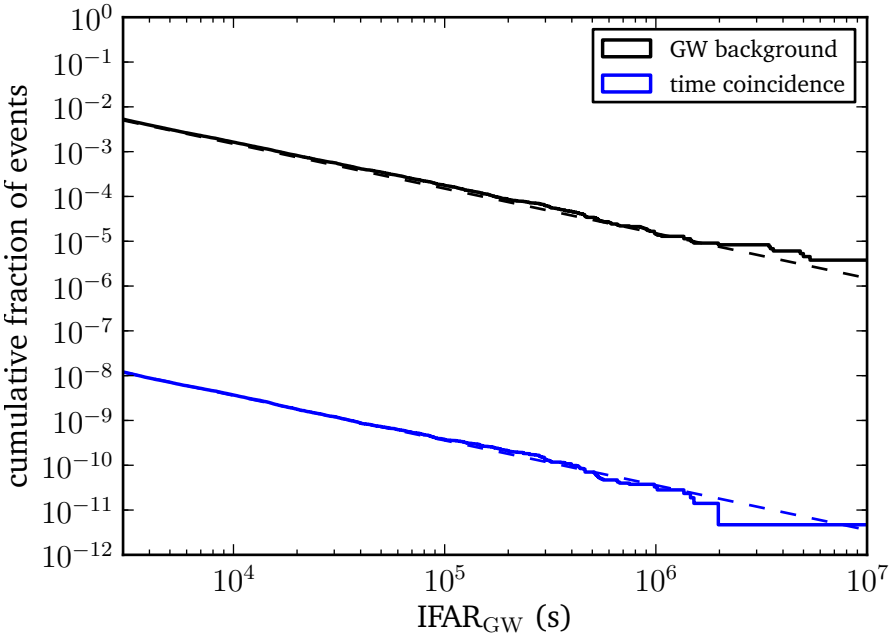

(b) 2016

Figure 3.9: Signal background (*i.e.* no injections present) for both the untriggered matched filter GW search (GSTLAL) and the RAVEN search using only time coincidence in (a) the 2015 and (b) the 2016 scenarios. In each case we show the "inverse false alarm rate" (IFAR) ranking statistic, $\mathcal{I}_{\rm GW}$, equivalent to $1/R_{\rm GW}$ (so that significance increases down and to the right).

To characterize the sky coincident background, we have run BAYESTAR on all GST-

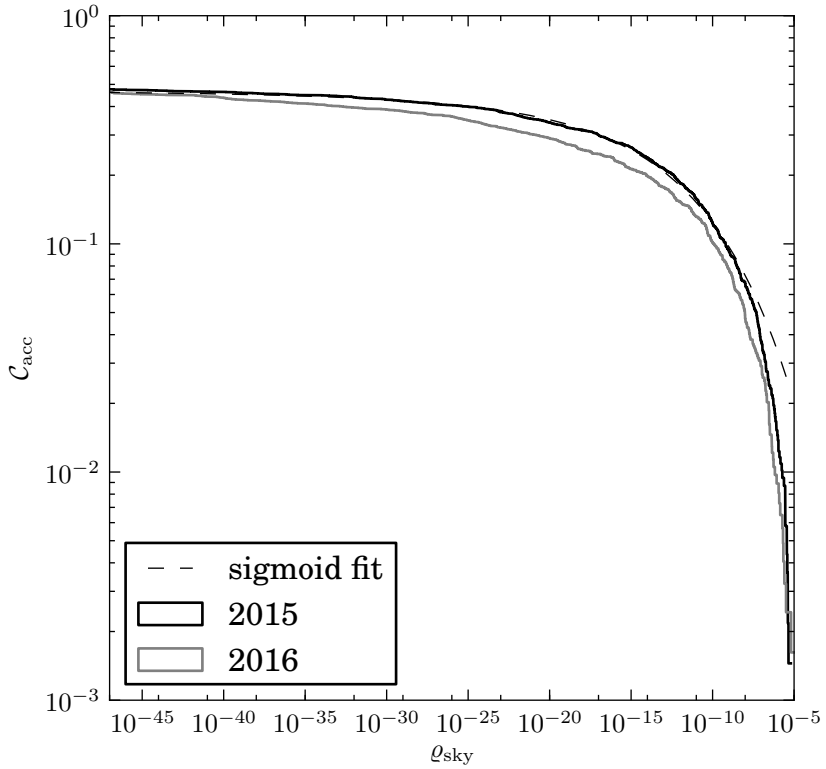

Figure 3.10: The cumulative fraction $\mathcal{C}_{\rm acc}(\varrho^*_{\rm sky})$ of accidental associations with $\varrho_{\rm sky} \geq \varrho^*_{\rm sky}$. The solid black and grey curves are the empirical histograms in the 2015 and 2016 scenarios, respectively, measured by localizing background events and modeling a *Swift* short GRB at a randomly selected point on the sky. (It is found that roughly 75% of such associations have $\varrho_{\rm sky} = 0$.) The black dashed curve is a fit of the 2015 empirical histogram to a rational sigmoid function that overestimates the measured value at large $\varrho_{\rm sky}$.

LAL background events with FAR $R_{\rm GW} \leq 10^{-4}$ Hz. Noting that *Swift* GRBs could come from any location on the sky but are localized to arcminute precision, we then select sky points at random and compute $\varrho_{\rm sky}$ for each background event's sky map. The cumulative distribution of $\varrho_{\rm sky}$ got from this procedure is plotted in Fig. 3.10; note that $\varrho_{\rm sky} = 0$ (and $\mathcal{C}_{\rm acc} = 1$) roughly 75% of the time in both 2015 and 2016. This is explained by the fact that BAYESTAR's sky maps tend to support only at most ≈25% of the sky on average, so random sky points will only coincide with BAYESTAR sky map support roughly this often. In 2016, the noise population has systematically fewer coincident "sky localizations" at or above a fixed $\varrho_{\rm sky}$ value; this is because localization in 2016 has improved with the sensitivity of H1 and L1.

Given that the signal background is well-described as a Poisson process in the ex-

pected way, we can now examine the effect of RAVEN's hierarchical ranking scheme on revealing more astrophysical signals than is possible with the untriggered GSTLAL search. To begin with, the most powerful effect of time coincidence alone is to reduce the fraction of contaminating background events by a factor $\tau\dot{N}_{\mathrm{GRB}} = 5.6 \times 10^{-5}$. In the 2015 scenario, this amounts to a relative improvement of $(4154 - 3105)/3105 = 33.8\%$ in the number of injected signals recovered by RAVEN with greater than 5σ significance compared to GSTLAL; in 2016, the relative improvement is similar at $(2941 - 2208)/2208 = 33.2\%$. There were slightly fewer surviving events in the 2016 configuration than in the 2015 one; this is because adding a third detector required us to apportion the two months of Gaussian noise to different combinations of detectors. In the 2015 simulation, all 60 days of data were allocated to the H1-L1 network. In 2016 about 43 days were devoted to the H1-L1-Virgo and H1-L1 networks, and the remaining 17 days with only H1-Virgo and L1-Virgo coincident time contributed relatively few detections due to Virgo's reduced sensitivity.

Note that this GRB study is selecting for face-on binaries ($\iota \leq 30°$) whose GW emission is linearly polarized and has maximum amplitude. Relative to *e.g.* the 2015-era results reported in Singer et al. (2014), which draw injected signals from an astrophysical distribution uniform in $\cos\iota$ but use the exact same noise realization (and earlier versions of the same software) as this study, RAVEN reports over 450% more signals recovered with $>5\sigma$ confidence. However, it is worth noting that while RAVEN is more sensitive to signals with GRB counterparts, these on-axis binaries occur more rarely than the general NS-NS population by a factor $f_b = (1 - \cos\iota_{\mathrm{max}})^{-1} \sim 10$. All of these issues will be discussed in the next section.

For GW events recovered with FAR $R_{\mathrm{GW}} \leq 10^{-6}$ Hz by GSTLAL, sky coincidence with *Swift* GRBs provides only a median factor of 10^{-3}–10^{-2} improvement in false alarm rate in both 2015 and 2016. When the GRB localization is this precise, our method of computing $\mathcal{C}_{\mathrm{acc}}$, Eq. (3.24), amounts to what is referred to in Singer et al. (2014) as the "searched area" statistic: the area of the smallest credible region containing the true location of the source. (In the case of *Swift* BAT localization, the quantity $\mathcal{C}_{\mathrm{acc}}$ is effectively the ratio of this area to that of the entire sky, 41253 deg^2.) Cumulative histograms of the searched area (and $\mathcal{C}_{\mathrm{acc}}$) for the 2015 and 2016 populations of recovered signal injections is shown in Fig. 3.11. Note that 2016-era localizations of signal candidates with significance this low have systematically larger searched area in 2016

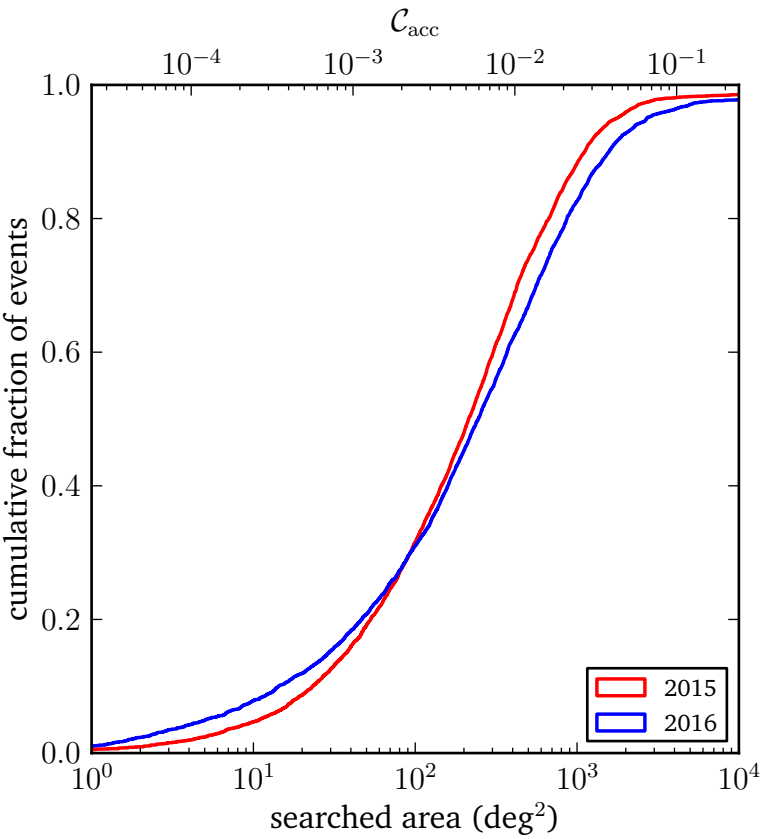

Figure 3.11: Cumulative histogram of sky localization areas in the simulated 2015 (H1-L1; red) and 2016 (H1-L1-Virgo; blue) scenarios. The metric used is the "searched area," *i.e.* the area of the smallest confidence region containing the true location of the source. The top axis shows the estimated cumulative fraction $\mathcal{C}_{\rm acc} = A/(41253\ \mathrm{deg}^2)$ of background events whose measured on-sky overlap is better than one whose searched area is A, assuming any coincident GRB was discovered by *Swift* BAT (and therefore has arcminute-scale precision).

than in 2015; this happens because low-SNR signals are buried in the reduced sensitivity of Virgo, and because the H1-L1 timing uncertainty goes from 131 μs in 2015 to 158 μs in 2016, so the sky error is larger by a factor $(158/131)^2 = 1.45$ at fixed SNR.[9]

The statistic space of recovered injections is shown in Fig. 3.12, illustrating the effects of time and sky coincidence separately. As we might expect, the *Swift* BAT sky precision makes a noticeable difference only right at the cusp of a given significance threshold, but provides 15–18% more recovered signals than time coincidence alone. We also note that, were the GRB localization less precise (as is typically the case for bursts reported by *Fermi* GBM), there is still significant benefit in analyzing sky coincidence because the shapes of GBM error boxes can break the H1-L1 degeneracy (see

[9]However, clearly the 2016 network is more sensitive (and will produce better parameter constraints) at fixed *distance*.

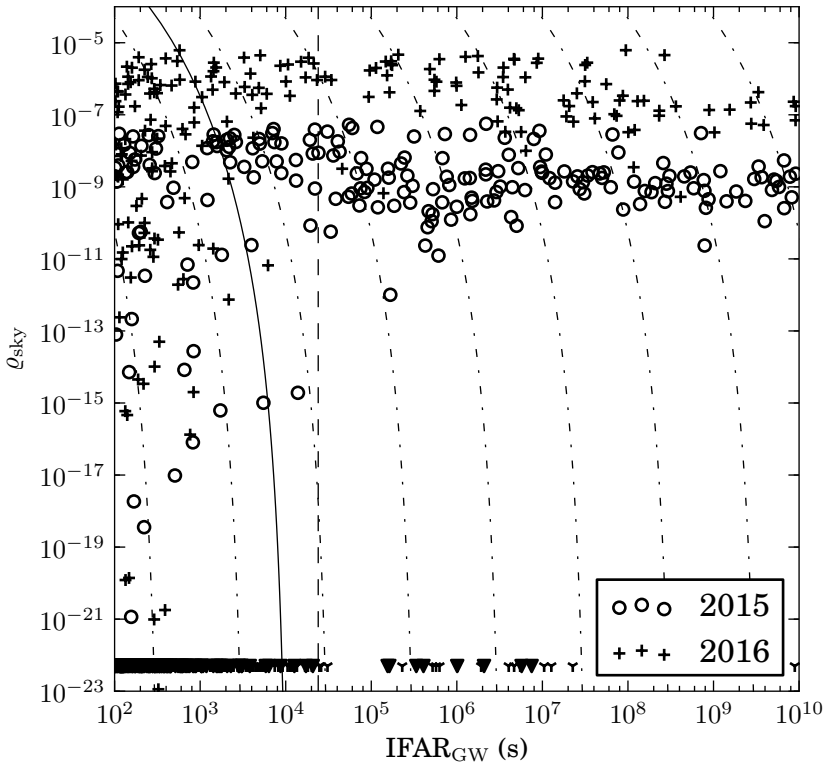

Figure 3.12: Contours of constant FAR in the $\mathcal{I}_{\rm GW}$-$\varrho_{\rm sky}$ plane. The dashed-dotted lines are the contours calculated with the rational sigmoid fit to $\mathcal{C}_{\rm acc}(\varrho^*_{\rm sky})$ in 2015, increasing to the right, with the black solid curve indicating the contour corresponding to a FAR of 1 century^{-1}. The black dashed line indicates the same value of FAR neglecting information on sky coincidence. A set of recovered 2015 and 2016 injections with corresponding GRBs has also been plotted, illustrating the improvement in sensitivity one gets from these two pieces of information if a FAR detection threshold of 1 century^{-1} is applied (with candidates to the left of either curve being considered detections).

chapter 2) and dramatically reduce the on-sky area a telescope needs to scan before finding the true source location. We defer a full analysis of the effect of *Fermi* GBM sky localization on RAVEN searches to a future work.

3.3.1.1 Comparison with Coherent Search

For a network of N detectors with well-measured noise, the combined likelihood ratio

$$\Lambda_{\mathbf{g}} = \prod_{i=1}^{N} \Lambda_i = \exp\left(\sum_{i=1}^{N}\left[(s_i|h)_i - \frac{1}{2}(h|h)_i\right]\right) \tag{3.30}$$

is the optimal matched filter ranking statistic with signal template $h(t)$ in Gaussian noise (Harry & Fairhurst, 2011), where s_i is the calibrated data stream and

$$(f|g)_i = 2\,\mathrm{Re}\int_{f_1}^{f_2} \frac{\tilde{f}^*(f)\,\tilde{g}(f)}{S_i(f)}\,df \tag{3.31}$$

the noise-weighted inner product of two functions f, g in the ith detector, whose power spectral density is $S_i(f)$. (A tilde denotes the frequency domain representation of a time series.) In terms of the coincidence search performed by pipelines such as GstLAL, the network SNR is simply the quadrature sum of single-detector SNRs,

$$\rho_c^2 = \sum_{i=1}^{N} \rho_i^2 = \sum_{i=1}^{N} \frac{(s_i|h)_i^2}{(h|h)_i}, \tag{3.32}$$

since the ρ_i^2 are individually proportional to the maximum log-likelihoods $\ln \Lambda_{i,\mathrm{max}}$. Typically, a basic consistency of templates and arrival times are required across each of the N detectors, with χ^2 tests further checking signal consistency as noted above.

Coincidence searches making use of ρ_c may be done in nearly real time, as is the case during online searches for compact binary coalescence by pipelines such as GstLAL. However, ρ_c^2 is not in general proportional to $\ln \Lambda_{\mathbf{g},\mathrm{max}}$ because it does not account for correlations between detectors. The network SNR is *not* an optimal statistic, even in Gaussian noise – for that, one must compute the *coherent* SNR, which is formed from a coherent sum across all N data streams before carrying out a search. The coherent SNR ρ_{coh} may be thought of as the effective network SNR of a hypothetical, completely uncorrelated, synthetic 2-detector network, with each synthetic detector corresponding to one of the allowed polarization states in General Relativity. The synthetic plus and cross data streams are given by

$$\tilde{s}_{+,\times}(f) = \sum_{i=1}^{N} \frac{F^i_{+,\times}\tilde{s}_i(f)}{S_i(f)} \left(\sum_{j=1}^{N} \frac{(F^j_{+,\times})^2}{S_j(f)}\right)^{-1} \tag{3.33}$$

where $F^i_{+,\times}(\alpha, \delta, \psi)$ are the antenna response functions of the ith detector in the network (see chapter 2), which depend on sky location (RA, dec) and polarization angle of the source. The plus and cross power spectral densities are

$$\frac{1}{S_{+,\times}(f)} = \left(\sum_{i=1}^{N} \frac{(F^i_{+,\times})^2}{S_i(f)}\right)^{-1} \tag{3.34}$$

and in terms of this synthetic network the coherent SNR is

$$\rho^2_{\rm coh} = \frac{(s_+|h)^2_+}{(h|h)_+} + \frac{(s_\times|h)^2_\times}{(h|h)_\times} \tag{3.35}$$

where the indices $+, \times$ refer to data streams and inner products in the signal space of the corresponding plus and cross synthetic detectors, respectively.

In Gaussian noise, the single-detector SNRs are Gaussian-distributed since they are linear operations on the data stream $s_i(t) = n_i(t) + h(t)$. However, specializing to the case of NS-NS binary inspiral, the waveform $h(t)$ depends on twice the orbital phase, does not precess, and radiates most of its power during the inspiral. Therefore, there are two linearly independent contributions to the waveform (call them h_0 and $h_{\pi/2}$) and hence to the SNR. The background (trigger events with no signal present) in any matched filter search for these sources in the N-detector network will measure a χ^2 distribution with $2N$ degrees of freedom in the case of ρ_c^2, and strictly four degrees of freedom (regardless of N) in the case of $\rho^2_{\rm coh}$. When a signal is present the SNRs will still follow a noncentral χ^2 distribution, with the same number of degrees of freedom but with non-centrality parameters $\lambda = \sum_{i=1}^N (h|h)_i$ for ρ_c^2 and $\lambda = (h|h)_+ + (h|h)_\times$ for $\rho^2_{\rm coh}$. For searches triggered by a GRB, a fully coherent matched filter search will improve upon any coincidence search because the detector response to GWs is modelled and the background has fewer degrees of freedom (in a network with three or more detectors).

While there is a fully coherent matched filter pipeline designed to search through LIGO data (Williamson et al., 2014), it is impractical to use it in this study due to the computational cost. Therefore, we simulate the output of an ideal coherent search by first computing the noncentrality parameter $\lambda = (h|h)_+ + (h|h)_\times$ for every NS-NS source in our simulated population given the strain power spectral densities in 2015 and 2016 (Fig. 3.7). The distribution of coherent SNRs measured for an individual source at a distance D and other parameters $\bar\theta = (\alpha, \delta, \iota, \psi, \phi_c, m_1, m_2, \chi_1, \chi_2)$ across a sequence of Gaussian noise realizations $\mathbf{n}$ in each detector is

$$p(\rho^2_{\rm coh} \,|\, D, \bar\theta, \mathbf{n}) = \mathrm{PDF}_{\chi^2}(\rho^2_{\rm coh}; \lambda, 4). \tag{3.36}$$

We are interested in the distance sensitivity of a coherent search, so we marginalize over the parameters $\bar\theta$:

$$p(\rho^2_{\rm coh} \,|\, D, \mathbf{n}) = \int p(\rho^2_{\rm coh} \,|\, D, \bar\theta, \mathbf{n})\, p(\bar\theta)\, d\bar\theta \tag{3.37}$$

where the Bayesian prior distribution on each parameter is encoded by

$$p(\bar{\theta})\,d\bar{\theta} = \prod_{i=1}^{M} p(\theta_i)\,d\theta_i. \tag{3.38}$$

The distance sensitivity is then given by the fraction of events whose ranking statistic would be larger than some threshold $\bar{\rho}^2$, *i.e.*

$$P(\rho^2_{\rm coh} \geq \bar{\rho}^2 \,|\, D, \mathbf{n}) = \int_{\bar{\rho}^2}^{\infty} p(\rho^2_{\rm coh} \,|\, D, \mathbf{n})\,d\rho^2_{\rm coh}. \tag{3.39}$$

In practice we approximate this integral by first binning up in distance, D. For every signal in a given bin, we calculate the CDF of the SNR distribution, Eq. (3.36), at the threshold $\bar{\rho}^2$; we then sum this over every source in the bin (and weight this sum by the number of sources in the bin). Because the signals were drawn from astrophysical parameter distributions $p(\bar{\theta})$, this approximation to Eq. (3.39) improves with the number of sources in a given bin.[10]

To compare the sensitivity of RAVEN with that of the idealized coherent search, we use a similar histogram procedure, binning up the distance to source and computing the fraction of simulated signals in each bin which were recovered with a significance of 5σ or better. (Strictly speaking, this answers a slightly different question than the one posed for the coherent search because here we have access to only one noise realization.) Histograms for the raw GSTLAL search, RAVEN with time coincidence only, RAVEN with time and sky coincidence, and the idealized coherent search are all shown for both 2015 and 2016 scenarios in Fig. 3.13. While the behavior of each of these histograms doubtless suffers from small-number statistics in bins at low D, the measurement of efficiency is more precise at farther distances and clear patterns do emerge. In every case, we find that RAVEN's search sensitivity improves with each coincidence constraint added: time coincidence improves distance efficiency by about 10%, while sky coincidence adds a further 5% in distance. Volumetrically, this explains the overall factor of $(1.15)^3 \approx 52\%$ increase in number of recovered signals as reported in section 3.3. Moreover, decreasing the significance threshold to only 3σ (or a p-value of 2.7×10^{-3}) in a time-and-sky coincidence search more closely mathces the 5σ sensitivity of the coherent search at larger distances. This improves over GSTLAL by about 28% in

[10] Because the sources are uniform in volume, the distance distribution is $p(D) \propto D^2$. Unfortunately this means the approximation is poor at small distances, where there are far fewer simulated signals.

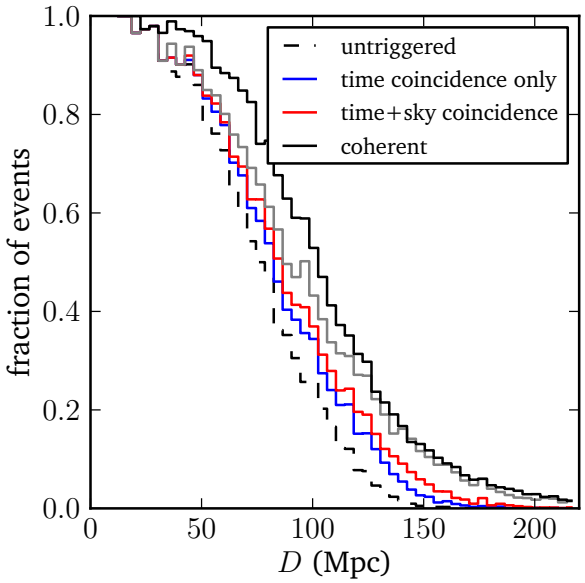

(a) 2015

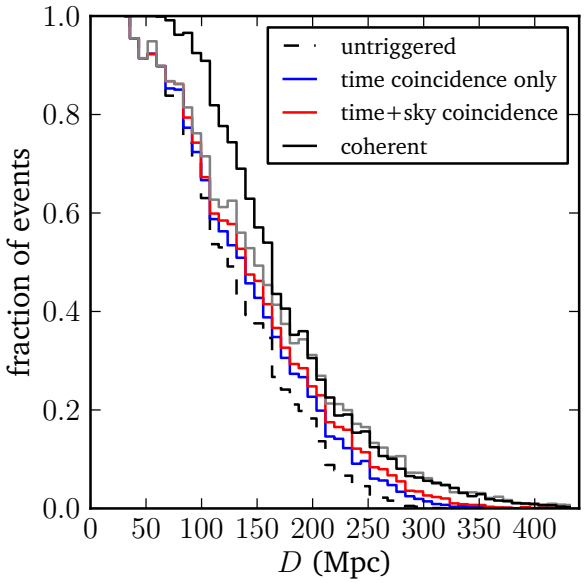

(b) 2016

Figure 3.13: Detection efficiency as a function of distance in (a) the 2015 and (b) the 2016 scenario, with an applied significance threshold of 5σ (amounting to a threshold on p-value of 5.7×10^{-7}). For full comparison, efficiency curves are shown for an untriggered all-sky search (dashed-dotted), a GRB search with time coincidence only (red), a GRB search accounting for both time and sky coincidence (blue), and a fully coherent GRB-triggered search (solid black). The grey curve illustrates a threshold of 3σ in p-value applied to the time- and sky-coincident search, which better matches the sensitivity of the coherent search at large distances.

distance in both 2015 and 2016, giving more than a factor of 2 improvement in volume sensitivity.

3.4 Summary and Discussion

In Gaussian noise, one expects the FAR statistic to obey a scaling law

$$R_{GW} \propto \exp\left(-\frac{\rho_c^2}{2}\right).$$

Since GW signals from compact binary coalescence are predicted to decay in amplitude like $1/D$ as they expand through the universe, the network SNR $\rho_c \propto 1/D$. Put another way, the maximum distance at which an optimally-oriented source with total mass M and reduced mass μ would be observed with single-detector SNR ρ is given by the scaling law (Maggiore, 2008)

$$D_H = 72.5\ \mathrm{Mpc}\ \left(\frac{M}{M_\odot}\right)^{1/3}\left(\frac{\mu}{M_\odot}\right)^{1/2}\left(\frac{1}{\rho}\right) \times \left[\int_{f_1/\mathrm{Hz}}^{f_2/\mathrm{Hz}}\left(\frac{f}{100\mathrm{Hz}}\right)^{-7/3}\left(\frac{10^{-46}\ \mathrm{Hz}^{-1}}{S(f)}\right) d\left(\frac{f}{\mathrm{Hz}}\right)\right]^{1/2}$$

where $f^{-7/3}$ is the approximate power spectral density behavior of a NS-NS inspiral signal. This quantity is called the *horizon distance* and is often used as a metric for detector sensitivity with $\rho = 8$ (since this corresponds in Gaussian noise to a network SNR of nearly 12, and is just on the cusp of being detectable). In the 2015 configuration, both H1 and L1 have $D_H = 54$ Mpc; in 2016, H1 and L1 have $D_H = 108$ Mpc and Virgo has $D_H = 36$ Mpc. This is borne out by the simulated population of signals with a full range of ι analyzed in Singer et al. (2014).

In the case of a RAVEN search, we are selecting for face-on compact binaries so that the horizon distance is improved by roughly the average "beaming factor" $f_b^{-1} \sim 10\%$. (Of course we also rely on the presence of short GRB emission, and the true beaming factor will depend on things such as the efficiency of producing such emission and the true beaming angle of the burst.) The results of the previous section verify that this improvement in horizon distance for GRBs is a real phenomenon in the time-coincident search; sky coincidence then changes the scaling and improves distance sensitivity yet further. Finally, dropping the RAVEN significance threshold to only 3σ lowers the SNR

threshold on "detectable" signals, boosting the horizon distance by a factor of $\sim$2. We believe a 3σ threshold is justified because the RAVEN search is not in place to claim detections; rather, it is a means of identifying potentially interesting signal candidates that would otherwise be considered "sub-threshold" and lost in the shuffle of background statistics. Jointly observing GRBs and their afterglows coincident in time and sky location with a GW candidate ranks among one of the highest priorities in multimessenger astronomy.

While RAVEN is sensitive at 3σ to a horizon distance twice as large as the untriggered GstLAL search (and therefore to a horizon volume nearly ten times as large), the beaming factor f_b means that the events RAVEN is searching for are intrinsically more rare. In terms of number of detections, if the average short GRB beaming angle is 30° then the beaming factor and expanded search volume produce a zero-sum game: if *Swift* and *Fermi* are highly efficient at discovering local short GRBs, then there will be (on average) $\sim$1 coincident sub-threshold short GRB detection for every "gold-plated" NS-NS inspiral discovered without the help of EM counterparts. During O1, the distance sensitivity of H1 and L1 varied between 60 and 80 Mpc; during O2, the horizon distance may be 5–10% better still.

Chapter 4

Constraints on the Cosmic Event Rate of Fast Relativistic Transients

Alex, I feel like you've made a lot of that up.

Jamee Cremeans, personal correspondence

Captain Picard: *"I understand what you've done here, Q. But I think the lesson could have been learned without the loss of 18 members of my crew."*

Q: *"If you can't take a little bloody nose, maybe you ought to go back home and crawl under your bed. It's not safe out here. It's wondrous, with treasures to satiate desires both subtle and gross. But it's not for the timid."*

Star Trek: The Next Generation (episode *Q Who?*)

Recall from chapter 1 that long-duration γ-ray bursts (GRBs) are brief (~ 1 min), ultra-luminous ($\sim 10^{50}$ erg s^{-1}) explosions in distant galaxies (redshifts $0.1 \lesssim z \lesssim 10$) that emit high-energy photons over a spectral range spanning from tens of keV to several MeV. There is strong evidence that these bursts are associated with stripped core-collapse supernovae, suggesting a progenitor system that involves accretion of disrupted matter onto a newly formed black hole as the central engine driving the explosion (Piran, 2004). Most bursts are followed by bright, non-thermal "afterglows" in the X-ray ($\Delta t \sim$ minutes), optical ($\Delta t \sim$ hours to days), and radio ($\Delta t \sim$ months to years) bandpasses, where Δt is the observed time since the initial burst.

The standard model of long-duration GRBs makes two important requirements. The explosion must be *ultra-relativistic*, expanding initially with Lorentz factors $\Gamma_0 \gtrsim 100$, since the prompt emission is observed to be non-thermal (*i.e.* the ejecta should be

optically thin to pair production at typical peak spectral energies of several hundred keV)[1]. The explosion must also be *tightly collimated*, producing radiation in roughly conical jets with half-opening angle $\theta_{\rm jet} \sim 1$–10°, due to simple energy conservation: An isotropic explosion would have luminosities $L_{\rm iso} \gtrsim 10^{54}$ erg s^{-1}, and this is difficult to explain in terms of known physical mechanisms. (Even compact object scenarios limit $L \lesssim 10^{53}$–10^{54} erg s^{-1}.) On the other hand, a narrowly beamed burst is consistent with the energy budget of a core-collapse scenario, and would release energies comparable to e.g. type Ibc supernovae. Robust observational support for beaming is offered by achromatic "jet breaks" visible in the light curves of several optical afterglows; these are explained by the jet suddenly widening as the ejecta slow down due to interactions with the circum-burst medium (*cf.* Rhoads 1997, Sari et al. 1999, Fruchter et al. 1999).

If ejecta are to be accelerated to such extreme velocities, then the outgoing jet must entrain only a very small mass $M_{\rm ej} \sim 10^{-5}\,M_\odot$, else too much kinetic energy is dissipated. This is referred to as the "baryon loading" problem. The vast preponderance of observed GRBs have prompt spectra with peak energies of a few hundred keV, requiring very "clean" outflows likely originating from the core collapse of a massive star progenitor that has shed its outer hydrogen layer. But there is mounting evidence that the astrophysical population of long GRBs may be dominated by sub-luminous or "dark" bursts, whose peak spectral energies are below the γ-ray bandpass (see e.g. Butler et al. 2010). Thus, some outstanding questions remain. Could these less energetic, fainter bursts (e.g., X-ray flashes; Heise et al. 2001) result from an outflow with more entrained mass (a so-called "dirty fireball"; Dermer et al. 2000; Huang et al. 2002)? Or can other properties, such as viewing angle (Granot et al., 2005) or the nature of the remnant (Mazzali et al., 2006), explain them?

A separate issue is that the tight collimation immediately implies most astrophysical GRBs are beamed away from Earth. If a typical opening angle $\theta_{\rm jet} \approx 10^\circ$, then only a fraction $f_b^{-1} \equiv 1 - \cos\langle\theta_{\rm jet}\rangle \sim 10^{-2}$ of GRBs are visible to Earth-based detectors (Guetta & Della Valle, 2007). However, as the outflow slows to a Lorentz factor $\Gamma \sim \theta_{\rm jet}^{-1}$ it begins to expand laterally, illuminating an increasing fraction of the sky with diminishing intensity (Rhoads, 1999; Sari et al., 1999). The *afterglows* of most bursts

[1]A possible photospheric component has been identified in the prompt high-energy spectra of several GRBs (e.g., Ryde et al. 2010; Guiriec et al. 2011), but this does not dramatically ease the requirement of ultra-relativistic expansion.

should therefore become visible at late times $t \gg \Delta t_{\rm GRB}$, at the cost of much weaker apparent brightness. But despite dedicated efforts to uncover these "orphan" afterglows in the X-ray (Nakar & Piran, 2003; Greiner et al., 2000), optical (Becker et al., 2004; Rykoff et al., 2005; Rau et al., 2006), and radio (Gal-Yam et al., 2006) bandpasses, no bona fide off-axis candidate has been discovered to date. This begs the question: Do we really understand beaming in GRBs?

Each of these issues can be addressed by sensitive, wide-field surveys that target relativistic explosions independent of a high-energy trigger. In the coming years, instruments such as the Zwicky Transient Facility (ZTF: Bellm 2014) and the Large Synoptic Survey Telescope (LSST: Ivezic et al. 2008) will have enough optical coverage at $\sim$hour-long cadences to make fast transient surveys viable (Fig. 4.1), and important technical milestones have already been achieved. Recently, the Intermediate Palomar Transient Factory (iPTF: Law et al. 2009a; Rau et al. 2009) reported discovery of iPTF14yb (Cenko et al. 2015; see section 4.1), the first unambiguous example of a long-duration GRB discovered by its optical afterglow – a high-energy counterpart, GRB 140226A, was later identified in archival data from the Inter-Planetary Network (IPN: Hurley et al. 2010a) of satellite detectors. Similarly, the "afterglow-like" transient PTF11agg was discovered during the earlier PTF survey, although it was never reliably classified due to a lack of optical spectra (Cenko et al., 2013).

In this chapter we attempt a comprehensive discussion of future relativistic transient surveys, as informed by the astrophysical event rates inferred from iPTF14yb and PTF11agg. We begin in section 4.1 by recounting the discovery of transient source iPTF14yb. In section 4.2 we proceed to outline a somewhat novel rate estimation scheme based on Bayesian statistics. Then, combining data from iPTF with those of other surveys, in section 4.3 we place significant constraints on the rates of three distinct but related types of optical transient: (1) typical, on-axis GRB afterglows discovered independent of the high-energy counterpart; (2) the off-axis afterglows of otherwise typical GRBs beamed away from Earth; and (3) "dirty fireballs" resulting from GRB-like explosions that entrain too much baryonic mass to produce γ-ray emission. Finally, in section 4.4 we conclude by clarifying a set of optimal search strategies for fast, explosive relativistic transients in the ZTF/LSST era.

Throughout this work, we adopt a standard ΛCDM cosmology with $H_0 = 67\,{\rm km\ s^{-1}\ Mpc^{-1}}$, $\Omega_{\rm m} = 0.32$, and $\Omega_\Lambda = 1 - \Omega_{\rm m} = 0.68$ (Planck Collaboration et al., 2013). All

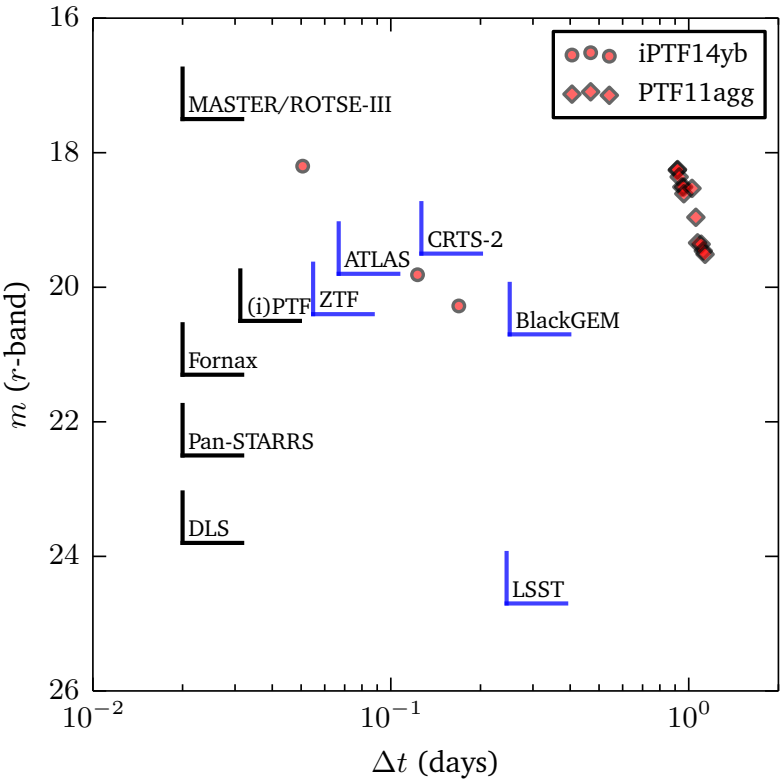

Figure 4.1: Limiting r-band magnitude against cadence for several ongoing (black) and planned (blue) synoptic all-sky surveys. For the ongoing surveys a timescale of 0.5 hr is shown (with a more realistic cadence of 1.0 hr for (i)PTF), while for the planned surveys the average cadence is estimated as the time $\delta t \times 0.1 \times 41253 \ \mathrm{deg}^2/\Omega_{\mathrm{FOV}}$ taken to image 10% of the sky and return to the same field, where δt is the total average imaging time (exposure plus readout and slew; see Tables 4.1 and 4.3 for references). The light curves of iPTF14yb and PTF11agg, as taken in the r-band with the P48, are shown for comparison.

quoted uncertainties are 1σ (68%) credible intervals unless otherwise noted, and UT times are used throughout. Any optical magnitudes discussed are measured in the AB system.

4.1 Discovery of Transient Source iPTF14yb[2]

As part of regular monitoring observations with the Palomar 48 inch Oschin Schmidt telescope (P48)[3], we discovered a new transient source, designated iPTF14yb, at J2000.0 location $\alpha = 14^{\rm h}45^{\rm m}58.01^{\rm s}$, $\delta = +14^\circ59'35\farcs1$ (estimated uncertainty of 80 mas in each coordinate). iPTF14yb was first detected in a 60 s image beginning at 10:17:37 on 2014 Feb 26, with apparent magnitude $r' = 18.16 \pm 0.03$. Subsequent P48 monitoring revealed rapid same-night fading from the source (see Fig. 4.1).

Nothing was detected at the location of iPTF14yb in a P48 image beginning at 09:04:46 on 2014 Feb 26 (i.e., 1.21 hr before the first detection) to a limit of $r' > 21.16$ mag. A coaddition of all existing iPTF P48 images of this location, spanning the time range from 2009 May 28 to 2014 Feb 24, also reveals no quiescent counterpart to $r' > 22.9$ mag.

We obtained target-of-opportunity X-ray observations with the *Swift* satellite (Gehrels et al. 2004a; see chapter 2) beginning at 17:11 on 2014 Feb. 26 A bright counterpart was identified in the X-Ray Telescope (XRT; Burrows et al. 2005a) images at the location of iPTF14yb. The X-ray spectrum is well described by a power law with a photon index $\Gamma = 2.1^{+0.5}_{-0.3}$.

Finally, we obtained a CCD spectrum of iPTF14yb with the Keck LRIS instrument beginning at 15:26 on 2014 Feb. 26. Superimposed over a relatively flat continuum ($f_\lambda \propto \lambda^{-1.3\pm0.1}$), we identify strong metal absorption lines from Mg 2, Fe 2, Al 2, C 4, Si 2, Si 4, C 2, and O 1 at $z = 1.9733 \pm 0.0003$. A damped Lyman α (DLA) system with $\log(N_{\rm H\,I}/{\rm cm}^2) = 20.7 \pm 0.2$ is also observed at this redshift, and the onset of the Lyman α forest blueward of H 1 implies that this is the redshift of iPTF14yb.

[2]The content of this section is based upon Cenko et al. (2015). On this work, I am listed as 2nd author; my direct contributions were primarily to the calculation of the event rate and some data analysis, as well as a portion of the text.

[3]P48 data processing is described by Laher et al. (2014), while photometric calibration of iPTF data is discussed by Ofek et al. (2012).

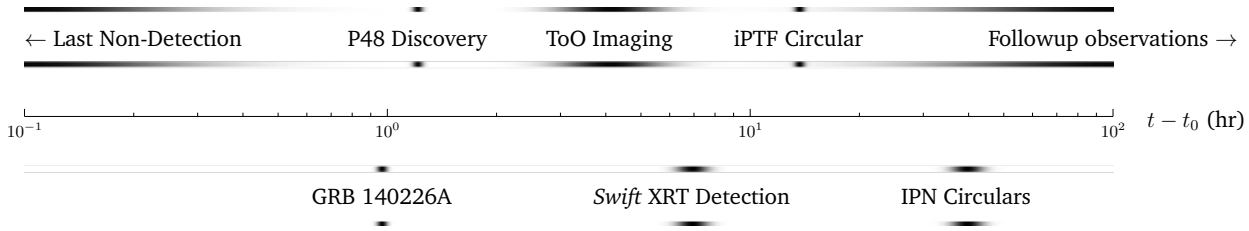

Figure 4.2: Timeline of discovery and announcements related to iPTF14yb and GRB 140226A. Note the optical transient iPTF14yb was discovered independently by the P48 instrument, nearly two full days before any announcement from IPN about GRB 140226A.

4.1.1 Association with GRB 140226A

After receiving notification of our discovery of iPTF14yb, the Inter-Planetary Network of high energy detectors (IPN: Hurley et al. 2010a) reported discovery of GRB 140226A, a possible counterpart of iPTF14yb (Hurley et al., 2014). GRB 140226A was detected by the *Odyssey*, *INTEGRAL*, and *Konus* satellites at 10:02:57 on 2014 Feb 26; that is, 14.7 min before the midpoint of our P48 discovery image, and 58.2 min after our last P48 nondetection. (Fig. 6.1 shows a full timeline of observations and announcements.) The *Konus* light curve[4] shows a single pulse with a duration of 15 s (i.e., a long-duration GRB), and a 20 keV – 10 MeV γ-ray fluence of $(5.6 \pm 1.1) \times 10^{-6}$ erg cm^{-2} (Golenetskii et al., 2014). The time-averaged spectrum is well fit by a cut-off exponential model with $\alpha = -1.1 \pm 0.1$ and $E_{\rm pk} = 414 \pm 79$ keV (Golenetskii et al. 2014; see also the discussion in chapter 1). At this time, the location of iPTF14yb was below the horizon for the Burst Alert Telescope (BAT) onboard *Swift*, while the Gamma-Ray Burst Monitor (GBM) on *Fermi* was turned off because it was passing through the South Atlantic Anomaly (Hurley et al., 2014).

We can estimate the *a posteriori* probability of chance coincidence, both spatially and temporally, with a procedure similar to the one outlined in chapter 3. The IPN localized GRB 140226A to a timing annulus with area of 210 deg^2 (Hurley et al., 2014). Thus, the likelihood of chance spatial association is ~ 0.005. Similarly, since 2010 Jan 1, the IPN has been detecting GRBs at a rate of ~ 0.88 day^{-1}. Therefore, the likelihood of an unrelated IPN GRB being detected within the 73 min period between the last P48 upper limit and the first detection of iPTF14yb is ~ 0.044. Hence, the joint probability

[4]See http://www.ioffe.rssi.ru/LEA/GRBs/GRB140226A.

of chance coincidence is quite small, $P \sim 2 \times 10^{-4}$. We conclude that iPTF14yb is very likely associated with GRB 140226A and we will proceed with this assumption for the remainder of this chapter.

4.1.2 iPTF14yb in the Long GRB Context

We now compare the observed properties of iPTF14yb and its host galaxy with the known population of long GRBs as a final consistency check. We fit the X-ray light curve to a power law of the form $f_\nu \propto t^{-\alpha}$, finding $\alpha_{\rm X} = 1.54 \pm 0.11$ ($\chi^2 = 0.46$ for 2 degrees of freedom). At late times ($\Delta t \gtrsim 10$ days), the observed optical decay flattens, and in our last DEIMOS image the emission at the transient location is clearly spatially resolved. We interpret this to result from the emergence of an underlying host galaxy with $R \gtrsim 24.6$ mag. Neglecting the first point in the R/r-band light curve (where the decay appears shallower), we find an optical decay index of $\alpha_{\rm O} = 1.02 \pm 0.02$ ($\chi^2 = 61.3$ for 10 degrees of freedom). The simultaneous GROND optical/near-infrared spectral energy distribution (SED) at $\Delta t = 1.0$ days is well fit by a power law with index $\beta_{\rm O} = 0.63 \pm 0.10$ with no evidence for extinction in the host galaxy. All of this is broadly consistent with standard afterglow models (e.g., Granot & Sari 2002; see chapter 1) for expansion into a constant-density circumburst medium with electron index $p \approx 2.5$ and a cooling break between the X-ray and optical bands. Furthermore, these properties are typical of the behavior of early X-ray (e.g., Evans et al. 2009) and optical (e.g., Cenko et al. 2009) afterglow light curves.

The temporal decay indices observed in the X-ray and (especially) the optical are difficult to reconcile with evolution after the jet break (e.g., Sari et al. 1999), as would be expected for an off-axis orphan afterglow ($\alpha_{\rm orphan} \gtrsim 2$).[5] As with the rapid rise from our P48 non-detection 1.2 hr before discovery, this further reinforces the association with GRB 140226A, as it suggests iPTF14yb was initially viewed from within the jet opening angle. Unfortunately, emergence of the host galaxy in optical observations at $\Delta t \approx 10$ days greatly complicates our ability to detect any jet break feature in the afterglow light curve, which would have offered robust support for such a geometry.

The observed optical spectrum is typical of low-resolution spectra taken of long GRB

[5] Even if we allow the outburst time to vary freely in our power-law fits, the best-fit temporal indices in the X-ray and optical are still $\lesssim 2.0$.

Survey Camera	E_A (deg^2 day)	$T_{\rm all-sky}$ (day)	R (yr^{-1})	Reference
PTF (high cadence)	...	0.031ϵ	$1.2^{+3.4}_{-1.1}\epsilon^{-1}_{\rm PTF} \times 10^4$	(this chapter)
PTF (SN cadence)	...	0.749ϵ	$487^{+726}_{-356}\epsilon^{-1}_{\rm PTF}$	(this chapter)
iPTF	24637	0.597ϵ	$611^{+910}_{-446}\epsilon^{-1}_{\rm iPTF}$	Cenko et al. (2015) / this chapter
PS1/MDS	40.4	6.1×10^{-4}	$\lesssim 1.8 \times 10^6$	Berger et al. (2013b)
DLS	1.1	1.1×10^{-5}	$\lesssim 9.8 \times 10^7$	Becker et al. (2004)
Fornax	1.9	2.2×10^{-5}	$\lesssim 5.0 \times 10^7$	Rau et al. (2008)
ROTSE-III	635	1.5×10^{-2}	$\lesssim 7.5 \times 10^4$	Rykoff et al. (2005)
MASTER	...	2.4×10^{-2}	$\lesssim 4.5 \times 10^4$	Lipunov et al. (2007)
Combined	25315	0.64	545^{+812}_{-398}	(this chapter)

Table 4.1: Constraints on the all-sky rate of fast optical transients. Here E_A is the total areal exposure, $T_{\rm all-sky}$ the equivalent all-sky survey time and R the constrained all-sky rate of fast optical transients for a given survey.

afterglows (e.g., Fynbo et al. 2009), with DLA absorption and strong features from both low- and high-ionization-state metal transitions. Unlike many other GRB afterglows, however, no fine structure lines are apparent in the spectrum of iPTF14yb. This may be caused either by absorbing material that is more distant from the explosion, or simply a lack of spectral resolution.

4.2 Bayesian Rate Estimation Scheme

A robust scheme for constraining the generic event rate R of a certain class of transients is afforded by Bayes' theorem in the following way. Assuming the true rate is roughly fixed over time, suppose a survey S discovers[6] n_S such transients after searching a total "volume" V_S with efficiency ϵ_S in the r-band. The volume referred to here need not be a physical volume; for example, if one is interested in the event rate per field of view per unit time, then the survey volume searched is the total areal exposure E_A. The efficiency factor ϵ_S, on the other hand, is always understood to mean the fraction of interesting astrophysical sources that are identified by transient detection pipelines when they appear in survey images.

According to Bayes' theorem, a survey's measurement of R is fully described by the

[6]Here, as in Cenko et al. (2013), it is critical to distinguish between *detection* (a bright new source has appeared in a survey image) and *discovery* (the source has been flagged as astrophysically interesting). In particular, due to the large number of uncatalogued asteroids near e.g. the IPTF camera's limiting sensitivity, discovery requires at least two independent detections.

posterior probability density

$$p(R|n_S, \epsilon_S, V_S) = \frac{\mathcal{L}(n_S|R, \epsilon_S, V_S)}{P(n_S|\epsilon_S, V_S)}\, p(R|\epsilon_S, V_S). \tag{4.1}$$

Heuristically, the various pieces of Eq. (4.1) are interpreted in the usual way under Bayesian statistics. The *prior probability* $p(R|\epsilon_S, V_S)$ encodes any prior knowledge about the rate, and must be chosen with great care to minimize bias. Likewise, at fixed R the game of discovering transients essentially becomes a counting experiment, so the *likelihood function*

$$\mathcal{L}(n_S|R, \epsilon_S, V_S) = \frac{(\epsilon_S V_S R)^{n_S}}{n_S!} \exp(-\epsilon_s V_S R) \tag{4.2}$$

is just the Poisson distribution with rate parameter $\epsilon_S V_S R$ (assuming R is uniform over the volume V_S). The remaining term is the *evidence factor*

$$\mathcal{E} \equiv P(n_S|\epsilon_S, V_S) = \int_0^\infty \mathcal{L}(n_S|R, \epsilon_S, V_S)\, p(R|\epsilon_S, V_S)\, dR \tag{4.3}$$

and is understood to serve as a normalization constant.

We note that a rate estimation method like this one is most useful when little is known about the transients of interest. If only a smattering of events has been observed, then the details of e.g. their true redshift distribution and luminosity function remain unclear, but a modest estimate of the event rate is helpful in designing better experiments. However, in light of the small-number statistics at hand, any measurement of the rate is going to be highly uncertain. We therefore combine data from multiple surveys with the following rule: Every time an interesting transient is discovered, we use the measured posterior from the last one as its prior, assigning to the first discovery a uniformly flat prior $p(R|\epsilon_S, V_S) \propto 1$. Since each discovery is statistically independent, after N discoveries the measured posterior is

$$\begin{aligned} p(R|\{n_i\}, \{\epsilon_i\}, \{V_i\}) &= \frac{1}{\mathcal{E}} \prod_{i=1}^{N} p(R|n_i, \epsilon_i, V_i) \\ &= \frac{V^{n+1} R^n}{n!} \exp(-VR) \end{aligned} \tag{4.4}$$

where $n = \sum_{i=1}^{N} n_i$ and $V = \sum_{i=1}^{N} \epsilon_i V_i$ respectively. (In other words, the result inferred from several surveys separately discovering N transients in respective volumes $\{\epsilon_i V_i\}$ is as though a single survey discovered them in a volume V.)

4.2.1 All-sky Rate of Fast Relativistic Transients

To illustrate, we will now compute the annual all-sky rate of fast relativistic transients. By this we mean extragalactic sources that show rapid fading in the optical (falling by at least one magnitude over the course of a night) and a non-thermal spectrum indicating relativistic ($\Gamma \gtrsim 10$) or ultra-relativistic ($\Gamma \gtrsim 100$) outflow, but which are discovered independent of any high-energy trigger. Because the quantity of interest is the per-field rate integrated over the whole sky, the relevant survey volume is the effective search time

$$T = \frac{\epsilon E_A}{41253\,\mathrm{deg}^2}. \tag{4.5}$$

In the iPTF era (2013 Jan 1 through 2014 March 1), only one event (iPTF14yb) is a confirmed ultra-relativistic outburst out of $T_{\mathrm{iPTF}} = 0.597\epsilon$ effective all-sky days of searching (since the total areal exposure $E_A = 24637$ deg^2; see Cenko et al. 2015). Hence our Bayesian rate estimation scheme (Eq. 4.4) produces exactly the most likely inferred rate of $611\epsilon^{-1}$ yr^{-1} derived in Cenko et al. when only data in this era and from this survey are considered. To report a 1σ error range on the rate measurement (and all others[7] with $n \geq 1$), we choose a Bayesian credible interval $[R_L, R_H]$ using a "water-lowering" algorithm that satisfies

$$\int_{R_L}^{R_H} p(R|n,T)\,dR = 0.68$$
$$p(R_L|n,T) = p(R_H|n,T) \tag{4.6}$$

resulting in a credible interval of (165–1521)$\epsilon_{\mathrm{iPTF}}^{-1}$ yr^{-1}. Non-detections from other high cadence, wide-field surveys (e.g. Pan-STARRS1, ROTSE-III and MASTER; see Berger et al. (2013b); Rykoff et al. (2005); Lipunov et al. (2007)) may also be included, but due to the comparatively low estimated total areal exposure of these surveys they are not found to further significantly constrain the all-sky rate (see Table 4.1).

However, the transient source PTF11agg was discovered in the initial PTF era between 2009 and 2012 (Cenko et al. (2013)), and its discovery certainly has some bearing on the all-sky rate. This transient exhibited all the hallmarks of a relativistic outburst, with rapid fading in the optical (1.2 mag in 5.3 hr), a light curve reminiscent

[7] In this section and in what follows, for surveys that produce zero discoveries, we report an upper limit with 95% confidence (corresponding to ≈ 3 events).

of on-axis GRB afterglows, and a long-lasting (∼years), scintillating radio counterpart. Deeper Keck/LRIS imaging at later times also revealed a faint ($R = 26.2$) blue ($g' - R = 0.17$) quiescent counterpart spatially coincident with PTF11agg, which is likely to be its host. Because this transient's light curve failed to demonstrate even the most rudimentary properties of off-axis afterglow models, and because the chance of serendipitously discovering an on-axis afterglow after missing the high-energy trigger due to lack of *Swift* or *Fermi* coverage was calculated as 2.6%, Cenko et al. present it as a candidate dirty fireball (*i.e.* a baryon-loaded relativistic outburst). Though the faintness of the host and absence of any Lyman break constrain the redshift to be $0.5 \lesssim z \lesssim 3$, lack of a definite spectroscopic redshift renders this interpretation fairly uncertain. Indeed, other authors have proffered alternative explanations in the context of compact binary coalescence (*cf.* Wang & Dai (2013), Wu et al. (2014)).

Also uncertain for PTF data is the total areal exposure to hours-long transients, since observing strategies at the time did not provide accurate estimation of fast transient rates (e.g., software pipelines for detection and rapid spectroscopic followup were not fully in place at the Palomar site, and target-of-opportunity programs were in their infancy). Further complicating matters is the fact that PTF11agg was discovered in an unusually high-cadence field, one with an order of magnitude more pointings on a single night than the slower-cadence fields typifying SN searches of the era. It is possible this discovery in a high-cadence field is purely serendipitous, but it is also possible the PTF pipelines would be unable to respond to a PTF11agg-like event at the lower cadence. To estimate the effective all-sky search time $T_{\rm PTF}$, we compare two lists of PTF fields: the first were imaged at least twice on the same night with reasonably good photometry (SN cadence) while the second were imaged at least ten times on a single night (high cadence). We then stack the differences τ_i in their image times, weighted by the 7.1 deg^2 field of view of the P48 camera and neglecting the overhead time δt (which includes exposure, readout and slew, and is in general negligible compared to the τ_i). We find that

$$\begin{aligned} T_{\rm PTF} &\approx \epsilon_{\rm PTF} \times \frac{7.1\,{\rm deg}^2}{41253\,{\rm deg}^2} \sum_i \tau_i \\ &= \begin{cases} 0.031 \epsilon_{\rm PTF}\ {\rm days} & ({\rm high\ cadence}) \\ 0.749 \epsilon_{\rm PTF}\ {\rm days} & ({\rm SN\ cadence}) \end{cases} \end{aligned} \tag{4.7}$$

where $\epsilon_{\rm PTF}$ is the unknown efficiency of the PTF transient discovery pipeline. Clearly

Designation	RA (J2000.0)	Dec (J2000.0)	UT Date	$m_{\rm peak}$ (P48 r-band)	Class
PTF09fso	$22^{\rm h}13^{\rm m}33.58^{\rm s}$	$-22°30'51''.9$	2013 October 22	17.29 ± 0.04	Blazar
PTF10lfj	$13^{\rm h}42^{\rm m}03.46^{\rm s}$	$+28°18'04''.0$	2014 February 22	17.07 ± 0.01	RR Lyrae
iPTF13hb	$13^{\rm h}00^{\rm m}37.65^{\rm s}$	$+27°30'24''.9$	2013 February 18	17.79 ± 0.02	M Dwarf
iPTF13bde	$16^{\rm h}30^{\rm m}25.03^{\rm s}$	$+39°44'25''.6$	2013 May 15	18.18 ± 0.02	M Dwarf
iPTF13bku	$13^{\rm h}27^{\rm m}11.01^{\rm s}$	$+12°13'05''.3$	2013 June 1	18.49 ± 0.04	M Dwarf
iPTF13cro	$23^{\rm h}18^{\rm m}07.82^{\rm s}$	$+11°49'37''.3$	2013 August 10	18.63 ± 0.03	M Dwarf
iPTF13dkt	$23^{\rm h}27^{\rm m}47.68^{\rm s}$	$+25°15'33''.7$	2013 September 13	18.28 ± 0.04	M Dwarf
iPTF14yb	$14^{\rm h}45^{\rm m}58.01^{\rm s}$	$+14°59'35''.1$	2014 February 26	18.16 ± 0.03	GRB Afterglow
iPTF14cvk	$01^{\rm h}29^{\rm m}52.03^{\rm s}$	$+30°55'36''.0$	2013 September 29	17.80 ± 0.02	M Dwarf
iPTF14cvl	$02^{\rm h}43^{\rm m}36.84^{\rm s}$	$+17°42'40''.6$	2013 October 1	18.05 ± 0.02	M Dwarf
iPTF14cvm	$05^{\rm h}41^{\rm m}53.90^{\rm s}$	$+60°44'55''.9$	2013 December 23	17.52 ± 0.08	M Dwarf
iPTF14cxm	$08^{\rm h}00^{\rm m}21.99^{\rm s}$	$+47°28'10''.4$	2014 January 23	18.86 ± 0.05	M Dwarf
iPTF14cxn	$06^{\rm h}46^{\rm m}27.10^{\rm s}$	$+41°03'51''.4$	2014 March 26	18.55 ± 0.06	M Dwarf
iPTF14czp	$20^{\rm h}57^{\rm m}16.30^{\rm s}$	$+44°16'43''.7$	2013 September 11	16.50 ± 0.01	M Dwarf

Table 4.2: Results of an archival search for fast optical transients in iPTF data.

there is significant difference between these two possibilities, with an order of magnitude less coverage in the high-cadence fields than in the ones with SN cadence. Discovery of one transient in this data means the most likely inferred rate is $487\epsilon_{\rm PTF}^{-1}$ yr^{-1} if its presence in a high-cadence field is serendipitous, with an estimated 1σ uncertainty ranging from (131–1213)$\epsilon_{\rm PTF}^{-1}$ yr^{-1}. On the other hand, in the event PTF11agg could *only* have been discovered in such a field, the most likely inferred rate is $1.2\epsilon_{\rm PTF}^{-1} \times 10^4$ yr^{-1} with a 1σ credible interval of (0.1–4.6)$\epsilon_{\rm PTF}^{-1} \times 10^4$ yr^{-1}. While the latter rate constraint is *very* large compared to that inferred from iPTF14yb and the former is more or less on par (Fig. 4.3), each rate constraint is within 1–2σ overlap with the posterior mode from iPTF14yb and therefore the higher rate cannot conclusively be ruled out. We list separately in Table 4.1 each of these inferred rates, noting that future fast transient surveys are needed to resolve the issue. We will continue to revisit PTF11agg in section 4.3.3.

4.2.2 Efficiency of iPTF

In order to properly interpret any statement about event rates we must understand the efficiency parameter ϵ. Broadly speaking, this number represents the fraction of *detectable* interesting transients on-sky that are correctly identified by the image subtraction pipeline, where "interesting transient" in this context refers to a GRB-like rel-

ativistic outflow. To estimate the efficiency of the iPTF survey to rapidly fading optical transients, we have performed an archival search over data between 1 February 2013 and 24 June 2014 for sources whose light curve exhibits a change $\Delta m \gtrsim 1$ mag over a timescale $\Delta t \lesssim 3$ hr, filtering out artifacts of image subtraction, poor photometry, etc. (A threshold on iPTF's own image ranking statistic was also applied to weed out innumerable junk events.) Of the surviving list of 288 candidates, we find only two relativistic outflow events (iPTF14yb itself and PTF09fso, a known blazar) as well as a total of twelve stellar flares and a smattering of Solar System asteroids, which we ignore for the sake of brevity. Of the stellar flares, eleven are M dwafs consistent with locations of known stars from the Wide-field Infrared Survey Explorer (WISE: Wright et al. 2010) and Sloan Digital Sky Survey (SDSS: Ahn et al. 2014) catalogues, and one is an RR Lyrae star. The detailed result of this search is shown in Table 4.2. Note that repeating this exercise for the PTF era between 2009 and 1 February 2013 is much more difficult, owing to the problems with observing strategies for same-night events mentioned above (see Cenko et al. (2013) for a more detailed discussion). For this reason we forego an estimate of the PTF-era efficiency here.

Based on the results of this search, we argue the efficiency parameter $\epsilon_{\rm iPTF} \sim 1$ because, failing one preexisting PTF-era blazar, no relativistic outburst at a cosmological distance was missed by the automated iPTF image subtraction pipeline. Following Berger et al. (2013b), to be conservative we estimate $\epsilon_{\rm iPTF} = 0.8$. Therefore the finall all-sky rate constrained from iPTF is 764 yr^{-1}, with a 68% credible interval of 206–1900 yr^{-1}. After combining all available surveys, the rate constraint is improved slightly to 706 yr^{-1}, with a 68% credible interval of 190–1760 yr^{-1} (see Table 4.1).

4.3 The Rate of Relativistic Transients

4.3.1 On-Axis Afterglows

In lieu of the results of the previous section we argue the efficiency parameter $\epsilon_{\rm iPTF} \sim 1$ because, failing one preexisting PTF-era blazar, no relativistic outburst at a cosmological distance was missed by the automated iPTF image subtraction pipeline. Therefore we now move on to constrain the volumetric rates of three categories of relativistic outburst: "classical" on-axis GRB afterglows; off-axis orphaned afterglows,

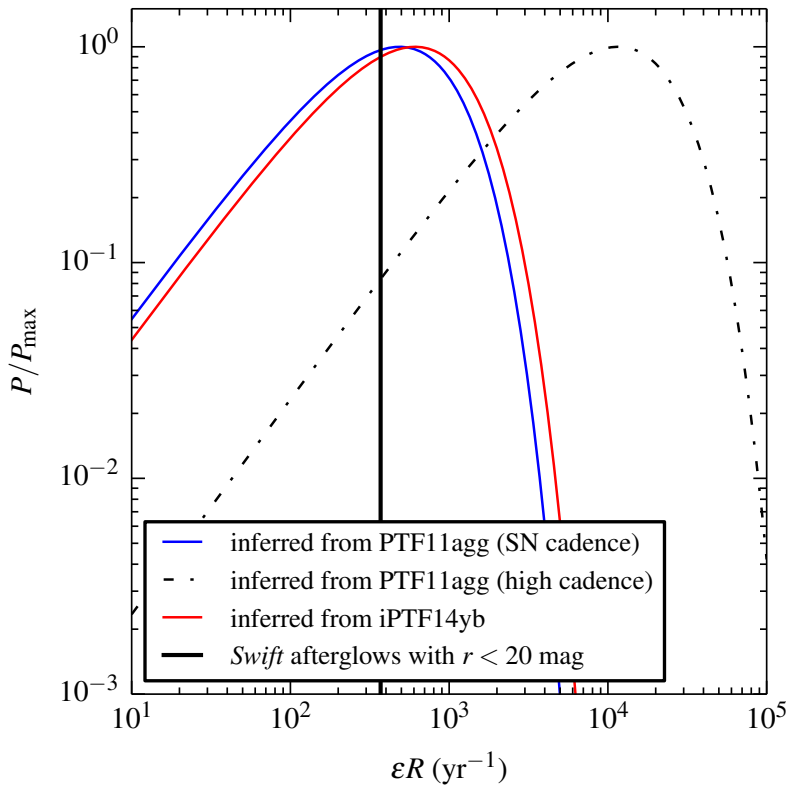

Figure 4.3: Posterior distributions for the all-sky rate of fast optical transients (down to $m \approx 20$ mag) as inferred from iPTF14yb (solid red) and PTF11agg in the case of cadences suited to SN searches (solid blue) and strictly high-cadence fields (dashed-dotted black). For comparison, the solid black line is the all-sky rate of *Swift* afterglows brighter than $r < 20$ mag (Lien et al., 2014).

where the high-energy counterpart is missed due to viewing angle effects between the source and Earth; and dirty fireballs, where the high-energy prompt emission is either highly suppressed due to baryon loading or absent altogether. We begin by framing the iPTF14yb rate in the context of long-duration GRBs.

For cosmological sources that occur with volumetric rate density $\mathcal{R}(z)$ in the rest frame of the source, the number density observable between redshift z and $z + dz$ per unit time in the observer rest frame is

$$n(z)\,dz = \frac{\mathcal{R}(z)}{1+z}\frac{dV}{dz}\,dz \tag{4.8}$$

where dV/dz is the comoving volume element. Assuming $\mathcal{R}(z)$ is slowly varying on yearly timescales, a transient survey that searches for an equivalent all-sky time T_S

with sensitivity[8] $\varepsilon(z)$ will discover

$$\lambda = T_S \int \varepsilon(z) n(z)\, dz \tag{4.9}$$

such sources on average. As with the all-sky rate constraint in section 4.2.1, because only a single event has been discovered to date there is precious little new information that can truly be gleaned from it. In particular, the true rate density $\mathcal{R}(z)$ of iPTF14yb-like events is a profound open question, one that has obvious bearing on the long GRB rate. A sizeable sample of optical GRB afterglows like iPTF14yb, unbiased by the usual selection effects that plague γ-ray telescopes, would go a long way toward answering this question in earnest. For the moment, suppose the rate density $\mathcal{R}$ is a constant, at least out to redshifts accessible to the current generation of optical telescopes. A Bayesian framework similar to the one used earlier to constrain the all-sky rate R directly applies, whence the volumetric likelihood function is

$$\mathcal{L}(n|\mathcal{R}, \epsilon, E_A, m_{\rm lim}) \propto \lambda^n e^{-\lambda} \tag{4.10}$$

and the posterior is computed as in Eq. 4.1.

After applying a standard K-correction (Hogg et al., 2002) at redshift $z_0 = 1.98$ and spectral flux density $f_\lambda \propto \lambda^{-1.3}$ in wavelength units, iPTF14yb is seen to have a peak absolute r-band magnitude $M_{\rm peak} \approx -27.5$ nearly $\sim 10^3$ s after the burst in its rest frame. This is fairly typical of *Swift* afterglows (e.g. Cenko et al. 2009), so in what follows, we model the survey sensitivity on the apparent r-band magnitude

$$m(z) = M_{\rm peak} + 5 \log_{10} \frac{D_L(z)}{10\,{\rm pc}} + K(z) \tag{4.11}$$

of a standard candle with peak luminosity $M_{\rm peak}$ at redshft z, where $D_L(z)$ is the luminosity distance and $K(z)$ the K-correction for iPTF14yb. The sensitivity function is

$$\varepsilon(z) = \Theta\left[m_{\rm lim} - 1 - m(z)\right] \tag{4.12}$$

(note that we require the transient to visibly fade by at least 1 mag in order to be flagged as astrophysically interesting). In practice, the sensitivity function (Eq. 4.12) simply cuts off the integral in Eq. 4.9 at some maximum redshift $z_{\rm max}$. Constraints on

[8] In this context, *sensitivity* is intended to mean the fraction of sources occurring at redshift z that are discovered by the survey. In principle, this will depend on the luminosity function of the source and on the limiting magnitude of the survey camera.

$\mathcal{R}$ have been tabulated in Table 4.3; based on the discovery of iPTF14yb, we measure the rate to be 1.2 Gpc^{-3} yr^{-1} with a 1σ credible interval of (0.3–3.1) Gpc^{-3} yr^{-1}. Wanderman & Piran (2010) find the local universe rate of long-duration *Swift* GRBs to be 1.3 Gpc^{-3} yr^{-1} without accounting for beaming; consistency of these two rates makes the association of iPTF14yb with GRB 140226A all the more robust.

4.3.2 Constraints on Off-Axis Events

Based on non-detections in the iPTF data, we may also further constrain the rate of off-axis "orphaned" afterglows (*i.e.* those viewed from an observing angle $\theta_{\rm obs}$ greater than the opening jet angle $\theta_{\rm jet}$ of high-energy emission at the source). The most constraining prior limit on this rate comes from the Pan-STARRS1 survey (Berger et al., 2013b) at $\mathcal{R}_{\rm orph} \lesssim 1 \times 10^3$ Gpc^{-3} yr^{-1} assuming standard candle sources with $M_{\rm peak} = -24$ and $f_\lambda \propto \lambda^{-1}$. In our case, using the same method as in section 4.3.1, we may set an upper limit with 95% confidence of

$$\mathcal{R}_{\rm orph} \lesssim \frac{3}{T_{\rm iPTF} V} = 26 \text{ Gpc}^{-3} \text{ yr}^{-1} \tag{4.13}$$

from an effective areal coverage of $E_A = 24637$ deg^2 days at $M_{\rm peak} = -24$. At $M_{\rm peak} = -20$, the rate constraint is $\mathcal{R}_{\rm orph} \lesssim 1.2 \times 10^3$ Gpc^{-3} yr^{-1}, while orphaned afterglows less luminous than this will presumably be lost on the P48 camera.

4.3.3 "Dirty Fireball" Events and Revisiting PTF11agg

The optical transient PTF11agg (Cenko et al., 2013) was discovered in the initial phase of PTF as part of regular monitoring during a high-cadence experiment with the P48 camera. The transient was never spectroscopically classified, leading to some ambiguity as to its true nature: its light curve is consistent with the population of *Swift* long GRB afterglows (Cenko et al., 2009), but no high-energy counterpart was observed by *Swift* or *Fermi*. A faint, blue quiescent counterpart was observed some in the same location several weeks later, suggesting that PTF11agg occurred somewhere in the redshift range $0.5 \lesssim z \lesssim 3$ (assuming this counterpart is the transient's host). In Fig. 4.4 we show the rate constraint (using the Bayesian scheme described in section 4.2) as a function of redshift; the cosmic rate constraint is consistent with both iPTF14yb and the rate of *Swift* bursts in this range. However, we strongly caution that

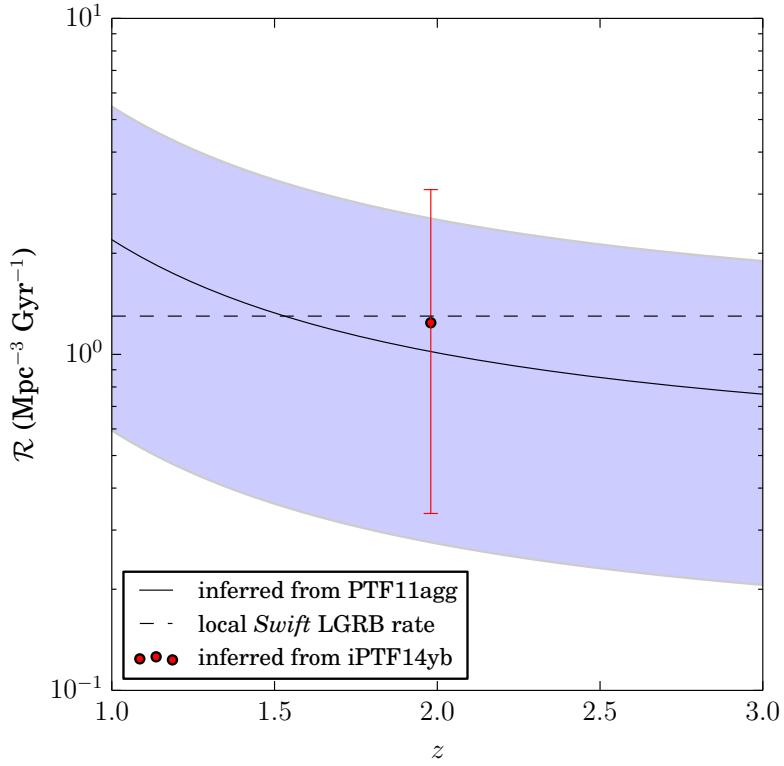

Figure 4.4: Volumetric rate of fast optical transients as inferred from iPTF14yb and PTF11agg. In the latter case, the rate is computed as a function of its unknown redshift, with the shaded region visualizing the 1σ uncertainty interval at fixed z. For comparison, the local long GRB rate (Wanderman & Piran, 2010) is indicated as a dashed line.

the question of efficiency, ϵ, to rapidly fading optical transients is much more uncertain in the PTF era than for iPTF.

4.4 Discussion and Conclusions

The most immediately tangible scientific statement that may now be made is an answer to the following question: in searching for fast optical transients such as iPTF14yb, is it more beneficial to increase optical sensitivity ($m_{\rm lim}$) or to maximize sky coverage (E_A)? In the coming years, as more large-scale synoptic sky surveys see first light, settling this question will be among the primary science drivers. Next-generation optical survey cameras, led by LSST, will be very sensitive at $m_{\rm lim} = 24.7$ mag with modest potential for rapid sky coverage ($\Omega_{\rm FOV} = 9.6$ deg^2 and average cadence ≈ 3 days). Others such as ZTF will be less sensitive at $m_{\rm lim} = 20.4$ but will also have much less compe-

Survey Camera	$m_{\rm lim}$ (5σ)	$\Omega_{\rm FOV}$ (deg^2)	$N_{\rm obs}$	VT (Gpc^3 yr)	$\mathcal{R}$ (Gpc^{-3} yr^{-1})	n
iPTF	20.7	7.1	19	0.81	$1.2^{+1.9}_{-0.9}$	$1.0^{+1.5}_{-0.7}$
PS1/MDS	22.5	7.0	3.1×10^{-2}	7.82×10^{-3}	$\lesssim 3.84 \times 10^{2}$	$1.1^{+1.7}_{-0.7} \times 10^{-2}$
DLS	23.8	20	8.5×10^{-4}	3.09×10^{-4}	$\lesssim 9.71 \times 10^{3}$	$4.3^{+6.8}_{-3.1} \times 10^{-4}$
Fornax	21.3	...	1.5×10^{-3}	2.35×10^{-4}	$\lesssim 1.28 \times 10^{4}$	$3.3^{+5.2}_{-2.4} \times 10^{-4}$
ROTSE-III	17.5	2.64	4.9×10^{-1}	7.73×10^{-3}	$\lesssim 3.88 \times 10^{2}$	$1.1^{+1.7}_{-0.8} \times 10^{-2}$
ATLAS*	19.8	30	473	15.3	—	18^{+29}_{-13}
ZTF*	20.4	47	288	11.8	—	14^{+21}_{-10}
LSST*	24.7	9.6	54	4.48	—	$5.5^{+8.3}_{-4.0}$
CRTS-2*	19.5	19	125	3.55	—	$4.3^{+6.6}_{-3.1}$
BlackGEM*	20.7	18.5	63	2.85	—	$3.5^{+5.2}_{-2.6}$

Table 4.3: The volumetric rate of fast optical transients at $M = -27.5$ mag. For each survey camera we list the limiting magnitude $m_{\rm lim}$; the instantaneous field of view $\Omega_{\rm FOV}$; the average number $N_{\rm obs}$ of single-image observations per field in an equivalent 3π survey averaging 6.5 hours observing per night; the total spacetime volume VT probed for transients at $M = -27.5$; and finally the expected number of discoveries n using the volumetric rate constrained from iPTF14yb. The first five surveys are ongoing as of 2014 March 1, and the rate $\mathcal{R}$ listed is that constrained by the given survey. Performance for the five future surveys (indicated with an asterisk) is estimated according to Ivezic et al. (2008); Bellm (2014). The stated $N_{\rm obs}$ and n values are *per annum*.

tition for telescope time and a much larger instantaneous field of view at $\Omega_{\rm FOV} = 47$ deg^2. Assuming access to a total solid angle of 3π per year for a sample of commissioned surveys, we have computed the expected number n of iPTF14yb-like discoveries per year at a cosmological rate of 1.2 Gpc^{-3} yr^{-1} according to

$$n = \frac{3 N_{\rm obs} \Delta t}{4} \int_0^{z_{\rm max}} \frac{\mathcal{R}}{1+z} \frac{dV}{dz} \, dz \tag{4.14}$$

where $N_{\rm obs}$ is the number of single-image observations per field per year covering an average timescale of 6.5 hours in a single night, with successive images staggered at intervals of $\Delta t \simeq 1$ hour (Table 4.3) and $z_{\rm max}$ is the maximum redshift at which an iPTF14yb-like transient remains discoverable to a survey with limiting magnitude $m_{\rm lim}$. These results are also listed in Table 4.3.

Our predicted discovery rates are suggestive and encouraging. From our constraint on $\mathcal{R}$, we find on the order of 40 untriggered GRB afterglows are likely to be discovered per year if these surveys perform as expected. This will doubtless be a boon to the community of GRB astronomers, for whom a rich, unbiased population of aftergows will provide unprecedented insight into their true rate density and luminosity function. Follow-up spectroscopic studies of their host galaxies may also constrain demographics

in terms of metalicity and star formation history, making untriggered searches for "run-of-the-mill" afterglows their own reward. But perhaps more importantly, this exercise clarifies the optimal search strategy, at least for on-axis afterglows: to maximize discovery yield, coverage on-sky appears to be far more important than optical sensitivity. Surveys that sacrifice deep limiting magnitude in favor of wide instantaneous field of view and rapid cadence, such as ZTF and ATLAS, are likely to discover at least a factor of 3-6 more events than those such as LSST with more powerful survey instruments but fewer total exposures at $\sim$hours-long timescales (Fig. 4.1). Though the details of rate history and luminosity function are left to future experiments, the reason for this is already fairly clear: most on-axis afterglows are extremely luminous (with $M \simeq -28$ on average) and a telescope with $m_{\rm lim} \sim 20$ is already sufficient to probe out to peak star formation at redshift $z \sim 2$-3 for these transients (while only comparatively few afterglows per year originate from redshifts at $z \gtrsim 4$). It should be noted, however, that a survey like LSST is suited to probe high redshifts for *all* afterglows – not merely the most intrinsically luminous ones – and this presents a unique opportunity to see how well GRBs trace star formation history in this regime.

Part III

Now, Bring Me That Horizon

Chapter 5

Results from Advanced LIGO's First Observing Run[1]

All along the watchtower, princes kept the view
While all the women came and went; barefoot servants, too.
And outside, in the distance, a wildcat did growl
Two riders were approaching. The wind began to howl.

Bob Dylan, *All Along the Watchtower*

As discussed in chapters 1 and 2, prior to GW150914, the best evidence supporting gravitational waves (GWs) as a real astrophysical phenomenon came from measurements of the decaying orbital period of the binary pulsar PSR J1913+16 (Hulse & Taylor, 1975; Taylor & Weisberg, 1982). The orbital decay of PSR J1913+16 is consistent with predictions of General Relativity to very high precision, and demonstrates without ambiguity that energy is being radiated from the system. While General Relativity explains this energy loss as having been carried away in the form of gravitational waves, the frequency of any radiation from this binary pulsar will be too low for the Laser Interferometer Gravitational-wave Observatory (LIGO: LIGO Scientific Collaboration et al. 2015) to observe directly. Instead, GWs from compact binary inspirals will enter LIGO's sensitive frequency band only in the last moments before merger, as the two bodies come within hundreds of kilometers of one another and attain orbital velocities that are an appreciable fraction of the speed of light.

[1]Portions of this chapter were completed as part of the LSC Fellows Program at the LIGO Livingston Observatory in autumn 2015, concurrent with the first Observing Run of Advanced LIGO.

The two LIGO detectors, together with a third facility called Virgo in Cascina, Italy (Acernese et al., 2015a), initially conducted observations between 2002 and 2010 (Abbott et al., 2016c). No GW detections were made during this period, so the LIGO and Virgo Collaboration commissioned a series of upgrades to the network of detectors meant to make them an order of magnitude more sensitive to the local universe. The LIGO facilities completed their first three-month Observing Run (O1) in an early "Advanced" configuration between 18 September 2015 and 13 January 2016, detecting GW150914 during an engineering period just prior to the official start of the Run, as operators and engineers were preparing the Hanford, WA (H1) and Livingston, LA (L1) detectors for stable data-gathering. At the time of writing, analysis of data from the whole of O1 is ongoing. The detectors are currently receiving another round of upgrades, and will begin gathering data again with yet more sensitivity in July 2016 in the frequency range from 25 Hz to 8 kHz.

This chapter will focus exclusively on material that has either already entered the public domain or is related to the performance of certain computing services. In section 5.1 we describe aspects of the discovery of the gravitational wave transient GW150914, including sky localization, electromagnetic (EM) follow-up alerts, and data quality checks. Section 5.2 will report on the behavior of three automated critical services: rapid sky localization with BAYESTAR (Singer & Price, 2016), filtering and storage of real-time γ-ray burst (GRB) alerts received over the Gamma-ray Coordinates Network (GCN[2]), and the automated RAVEN GRB coincidence pipeline. Section 5.3 estimates constraints on the properties of short GRBs from non-detections reported by RAVEN. Finally, section 5.4 offers a summary and some concluding remarks.

5.1 Role in the Discovery of GW150914

On 14 September 2015, at 9:50:45 UT, the two LIGO facilities in Livingston, LA and Hanford, WA separately observed a transient gravitational wave (GW) signal within $6.9^{+0.5}_{-0.4}$ ms of one another (Abbott et al., 2016c).[3] Designated GW150914 based on its calendar date, the observed signal is consistent with a binary black hole merger of initial component masses $m_1 = 36^{+5}_{-4}\,M_\odot$, $m_2 = 29^{+4}_{-4}\,M_\odot$, occurring at a redshift

[2] http://gcn.gsfc.nasa.gov

[3] For each of these parameters we report the median value and range of the 90% credible interval.

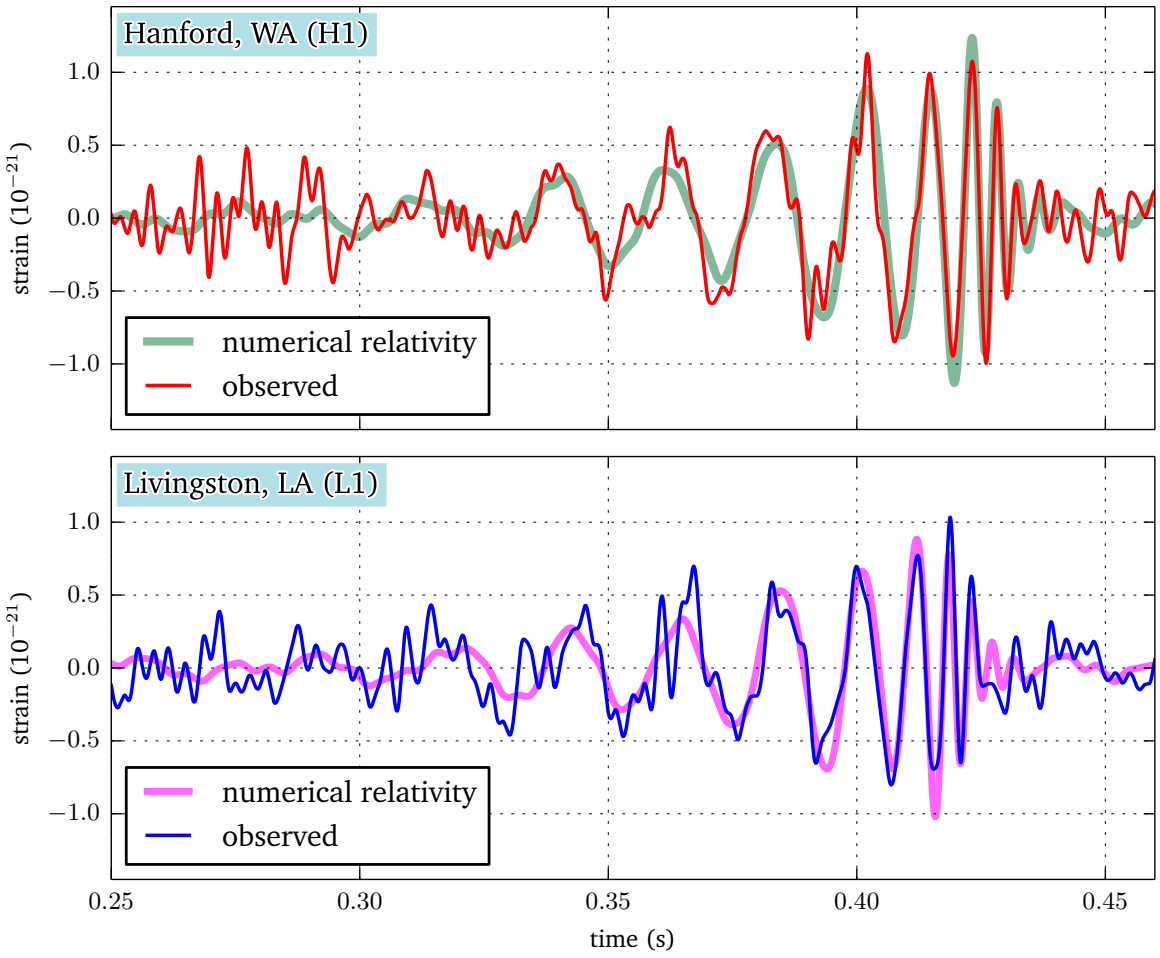

Figure 5.1: The transient gravitational wave source GW150914 as observed by the LIGO Hanford (H1) and Livingston (L1) detectors. Times are shown relative to 14 September 2015 at 09:50:45 UT. Following Abbott et al. (2016c), for visualization, all time series are filtered with a 35–350 Hz bandpass filter to suppress large fluctuations outside the detectors' most sensitive frequency band, and band-reject filters to remove the strong instrumental spectral lines. *Top:* H1 strain. *Bottom:* L1 strain. (Note, GW150914 arrived first at L1 and 6.9 ± 0.5 ms later at H1.) Thick, lightly-colored lines show numerical relativity simulations of a system whose parameters are consistent with those recovered from GW150914, filtered with the same 35–350 Hz bandpass filter and projected onto the H1 and L1 detectors.

$z = 0.09^{+0.03}_{-0.04}$ (and a luminosity distance $D_L = 410^{+160}_{-180}$ Mpc from Earth). After their GW signal swept up in frequency from 35 to 350 Hz over a period of roughly 0.2 seconds, the black holes violently merged, leaving behind a final black hole of mass $m = 62^{+4}_{-4}\,M_\odot$ and dimensionless spin $a = c|\mathbf{S}|/Gm^2 = 0.67^{+0.05}_{-0.07}$ (where $\mathbf{S}$ is the spin angular momentum), radiating $3.0^{+0.5}_{-0.5}\,M_\odot c^2$ of energy in the process (The LIGO Scientific Collaboration & the Virgo Collaboration, 2016b). This consitutes the first confirmed direct detection of GWs and the first direct evidence for binary black hole systems occurring in nature. The final black hole is several times more massive than any other known stellar-mass black hole (Abbott et al., 2016a).

GW150914 was first reported by Coherent WaveBurst (cWB: Klimenko et al. 2008, 2016), a data analysis pipeline designed to identify generic GW transients in the LIGO data stream. The characteristic "chirp" waveform (see chapter 2) is visible by eye in the data stream after bandpassing from 35–350 Hz, and from the time-frequency structure of the signal (Fig. 5.1) one can crudely estimate m_1, m_2 using Kepler's law. A sky map produced by the cWB pipeline became available within 15 minutes of merger, but the first human-readable GCN Circular was not distributed to astronomers until nearly two days later, and no information was given about masses, significance, or the nature of the signal (Abbott et al., 2016b). More refined searches using matched-filter techniques (Cannon et al., 2012a; Usman et al., 2015) later identified the transient with higher significance, and the LALInference compact binary parameter estimation architecture (Veitch et al., 2015) gave more precise measurements of the masses and sky location within several weeks. The full sequence of data products communicated with astronomers, and follow-up observations that were conducted, is described in chapter 6.

As a data quality check on GW150914, we[4] have investigated the x- and y-arm photon calibration readback channels at both LIGO sites (H1, L1) in $(-5, +5)$-second windows around 2015-09-14 09:50:54 UT (see Figs. 5.2 and 5.3). We find no evidence of misbehavior in these channels. There is structure present in the spectrogram, but these are spectral lines due only to measurements of the arm lengths at those frequencies. This implies that the photon calibration laser is not responsible for any transient actuation near the time of GW150914.

[4] I thank J. McIver, A. Lundgren, and S. Kandhasamy for guidance in completing this investigation.

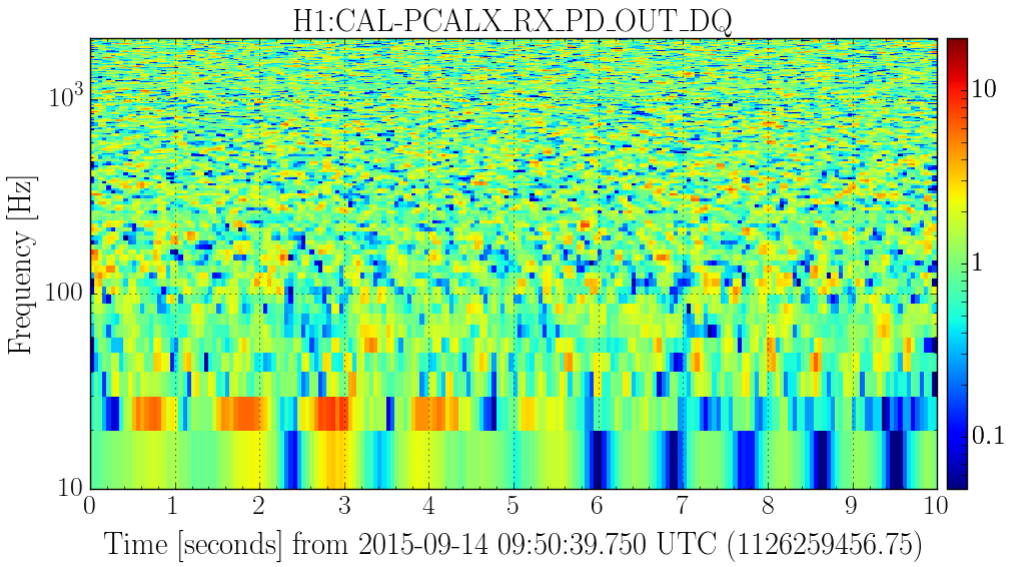

(a) H1

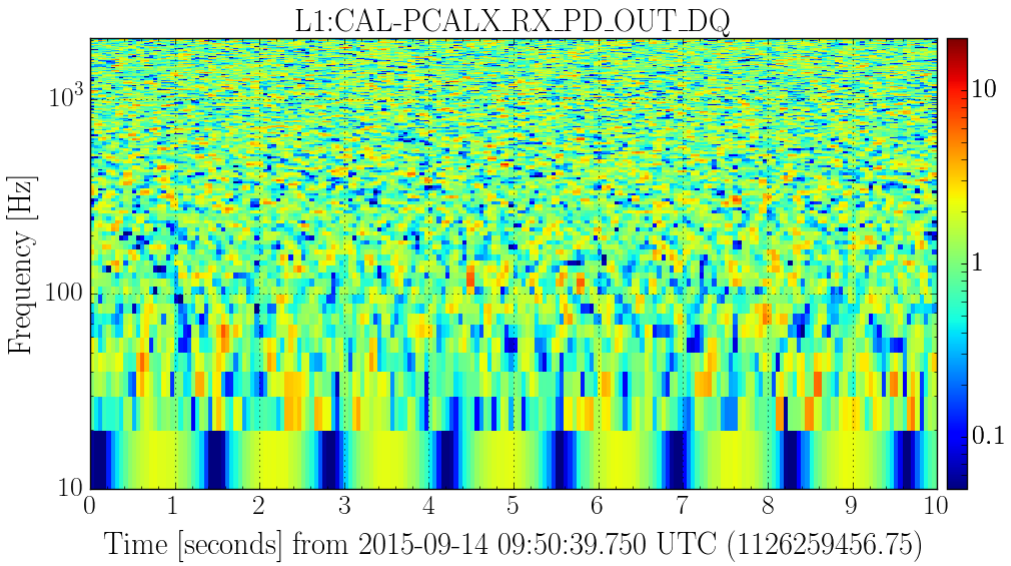

(b) L1

Figure 5.2: Spectrogram of the x-arm photon calibration in (a) H1 and (b) L1 near the time of GW150914. Note there is no obvious time-frequency structure that would indicate behavior out of the ordinary. (For reference, GW150914 occured 5 seconds after the start time of these plots.)

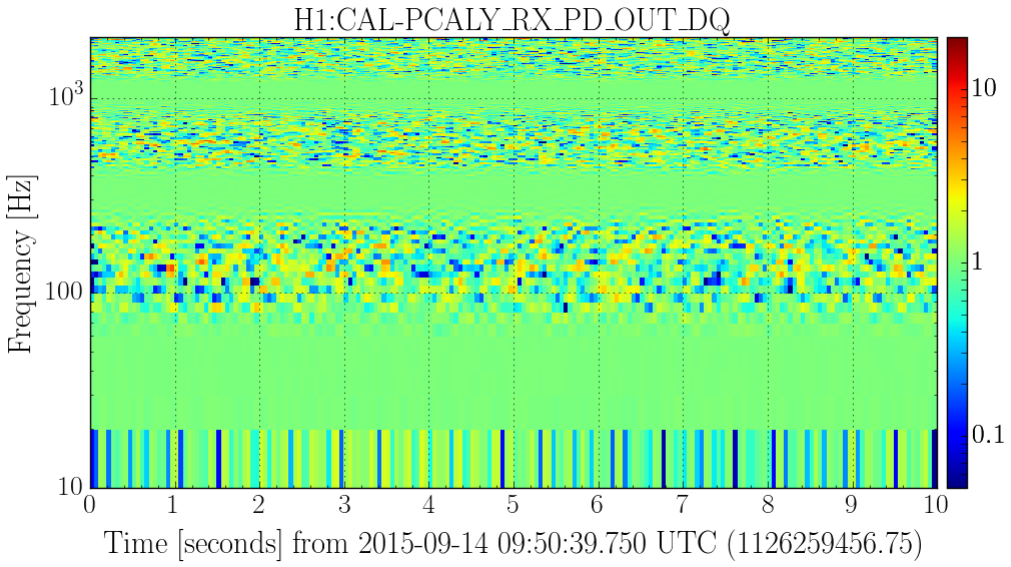

(a) H1

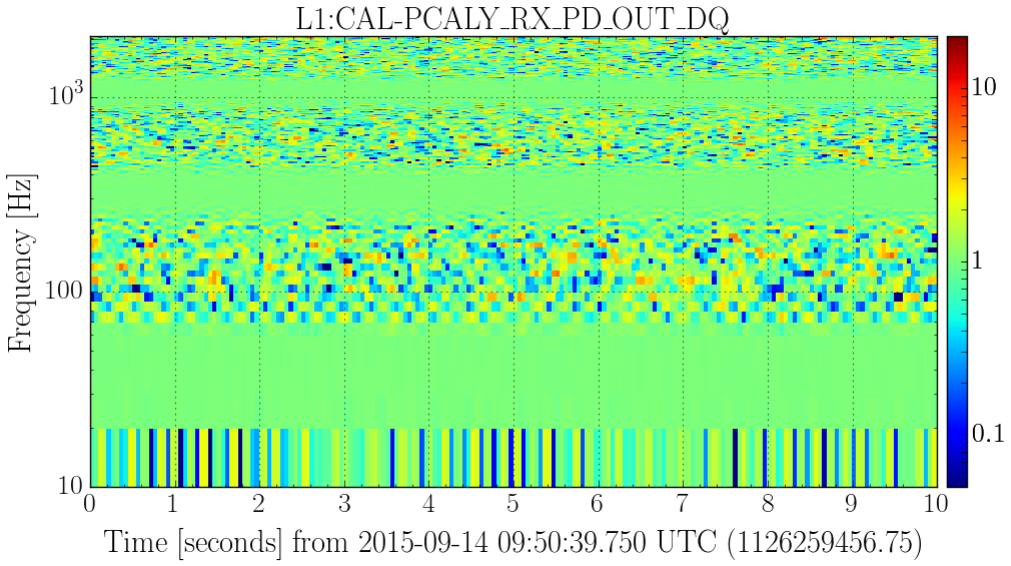

(b) L1

Figure 5.3: Spectrogram of the y-arm photon calibration in (a) H1 and (b) L1 near the time of GW150914. Note the spectral line structure present in both instruments, which is used to calibrate the arm lengths at specific frequencies. There is no obvious time-frequency structure that would indicate behavior out of the ordinary. (For reference, GW150914 occured 5 seconds after the start time of these plots.)

5.2 Behavior of Critical Services

In the pre-Run engineering phase and throughout the whole of O1, all real-time compact binary GW signal candidates were localized with the BAYESTAR rapid sky localization algorithm (Singer & Price, 2016). This process is automated, with BAYESTAR triggered when (and only when) one of the real-time matched filter pipelines uploads a signal candidate to the Gravitational-wave Candidate Event Database (GRACEDB[5]). During O1, the pipelines uploaded only those candidates with a false alarm rate (FAR) lower than 10^{-4} Hz ($\sim$10 day^{-1}). This threshold is *much* higher than the 5σ confidence threshold, which would amount to $\sim 10^{-13}$ Hz (or $\sim$1/10,000 yr) if the signal background is a Poisson process.[6] In terms of GRACEDB submission time, the average latency of the GSTLAL online matched filter search (Cannon et al., 2012a) was below 52 seconds after the measured coalescence time for NS-NS, NS-BH, and BH-BH binary inspirals. BAYESTAR was run on dedicated hardware hosted at the California Institute of Technology which consisted of Intel Haswell (Xeon Phi) processors with 16 physical cores (32 threads). Outside of maintenance periods on the LIGO Caltech computer cluster and an isolated incident with unusually high volume from GRACEDB, BAYESTAR sky maps were successfully computed for all real-time signal candidates within 20–60 seconds of their arrival in the database (and 70–110 seconds after merger).

5.2.1 Electromagnetic Follow-up Program

As discussed in chapter 2, prior to the start of O1 the LIGO Scientific Collaboration had determined that any signal candidate with FAR $\leq\sim$ 1 month^{-1} should be disseminated to a privately-subscribed group of 74 astronomy collaborations. In total, 63 of these groups were active and had the instrument time to observe LIGO source candidates during O1. The relatively high threshold in FAR was designed to guarantee that, even in the absence of "gold-plated" detections, our astronomer partners would still have candidates with typical sky localizations on which to train their follow-up programs. Signal candidates meeting this criterion would be disseminated as machine-readable Notices across the GCN system; however, during O1, GCN Notices

[5] https://gracedb.ligo.org

[6] Nevertheless, low-significance signal candidates are interesting for other reasons, including pipeline diagnostics, data quality testbeds, and, crucially, GW triggers for the RAVEN pipeline.

were only broadcast to LIGO members internally. After more careful vetting on ∼hour-long timescales, we then follow these with human-readable, prose-form GCN Circulars; these Circulars are then broadcast to the private group of astronomers.

Given that O1 lasted some four calendar months, the expected number of signals at or below this FAR threshold is four per data analysis pipeline, according to Poisson statistics. However, the pipelines are not independent (*i.e.* multiple pipelines may respond to the same event) and they do not normalize their FARs equivalently (some normalize to the total analyzed time, while others adjust for the duty cycle of each detector and normalize to the "wall clock" time) so the total number of expected candidates is still around 4–6.

During O1 there were 6 events with FARs below the stated 1/month threshold, identified by both matched-filter and generic transient searches. Three of these signal candidates were identified as transient noise artifacts within an hour or so of the initial event. For these three events, no GCN Circular was ever broadcast. The remaining three events were disseminated to our astronomer colleagues within a day or two of the observed signal. Of these, one was ultimately retracted several weeks later when follow-up estimates found that its significance was much lower than initially thought. Another was GW150914, and we will look closely at this as a flagship example of the full EM follow-up program in chapter 6.

5.3 Population of GRBs Analyzed by RAVEN

In the pre-engineering phase and during the whole of O1, GRBs discovered by the *Swift* Burst Alert Telescope (BAT: Barthelmy et al. 2005) and the *Fermi* Gamma-ray Burst Monitor (GBM: Meegan et al. 2009a) and reported in real time as GCN Notices were ingested by an automated monitoring service and stored in GraceDb. During O1 there were 76 electromagnetic transients identified by these satellites which were later confirmed as GRBs; their designations, T_{90} statistics, event times, and GraceDb links are shown in Tables[7] 5.1 and 5.2. The average latency of *Swift* GRBs is between 15-26 seconds, while that of *Fermi* bursts is 56-70 seconds (see chapter 2).

Of these 76 GRBs, 13 are classified as "short" based on their T_{90} and hardness ratio

[7]Many thanks to R. Coyne and D. Talukder for their assistance in producing this table.

(*i.e.*, the ratio of highest-energy γ-rays to lower-energy ones). Another 55 were classified as "long," and the remaining 8 bursts are described as "ambiguous" because it is not clear which astrophysical population they belong to. For each GRB, RAVEN performs an automated search for coincident GW candidates identified by any of LIGO/Virgo's real-time analysis pipelines, including both matched-filter searches for compact binary coalescence and more general searches for generic transients. Two coincidence windows are used for every burst: one, a $[-5, +1]$ second window based on models of compact binary merger (Metzger & Berger, 2012) and the other a $[-600, +60]$ second window based on models of supernova emission. (The wider coincidence window will of course capture more accidental associations, or background events; see chapter 3.) During O1, RAVEN did not identify any GW candidates within the shorter 6-second window. The wider, 10-minute window captured only accidental associations between GRBs and known "hardware injections" performed at the LIGO sites by applying radiation pressure against the test mass mirrors using the calibration laser. The average latency of search results from RAVEN is 30–80 seconds after a burst (when the trigger is a GRB) or 70–120 seconds after merger (when the trigger is a GW candidate).

5.4 Summary and Discussion

During O1, the horizon distance (see chapter 3) to NS-NS mergers in L1 measured around 60–70 Mpc, while in H1 the horizon distance got as high as 80 Mpc. The detectors are currently receiving another round of upgrades meant to make them more sensitive to transients in the local universe by an amount 5–10% in distance. Given that at least one binary black hole system has been observed with the LIGO experiment to date, and given that there may be a very weak γ-ray transient associated with it (Connaughton et al., 2016a), this is an exciting time to pursue multimessenger astronomy. In the next chapter, we turn to a detailed working example: the EM follow-up campaign surrounding GW150914.

Table 5.1: O1 GRBs analyzed by RAVEN.

GRB Designation	GRACEDB ID	Satellite	Classification (Long/Short)	Time (UT)	T_{90} (sec)	z
GRB150919606	E185587	*Fermi*	Long	14:33:18.900	6.656 ± 0.923	–
GRB150922A	E186292	*Fermi*	Short	05:37:29.076	0.144 ± 0.036	–
GRB150922718	E186424	*Fermi*	Long	17:13:42.887	15.872 ± 5.838	–
GRB150922883	E186470	*Fermi*	Ambiguous	21:11:32.732	2.816 ± 0.572	–
GRB150923297	E186583	*Fermi*	Short	07:07:36.184	0.192 ± 1.462	–
GRB150923429	E186620	*Fermi*	Short	10:18:17.924	0.192 ± 0.143	–
GRB150923864	E186747	*Fermi*	Short	20:44:23.646	1.792 ± 0.091	–
GRB150923995	E186783	*Fermi*	Long	23:52:52.541	17.152 ± 6.676	–
GRB150925A	E187124	*Swift*	Long	04:09:28.81	121.0 ± 27.0	–
GRB150928359	E188005	*Fermi*	Long	08:37:19.023	53.504 ± 4.615	–
GRB151001348	E188814	*Fermi*	Long	08:20:35.169	377.862 ± 8.083	–
GRB151001A	E188893	*Swift*	Long	15:04:22.93	8.94 ± 1.07	–
GRB151001B	E188928	*Swift*	Long	18:29:36.78	109.0 ± 17.0	–
GRB151003729	E189473	*Fermi*	Long	17:29:59.933	44.032 ± 5.043	–
GRB151004A	E189755	*Swift*	Long	18:09:04.77	128.4 ± 17.77	–
GRB151006A	E190216	*Swift*	Long	09:55:01.88	203.9 ± 41.6	–
GRB151009949	E191188	*Fermi*	Long	22:47:03.449	18.944 ± 2.429	–
GRB151011136	E191562	*Fermi*	Long	03:15:27.254	25.344 ± 2.996	–
GRB151014592	E192530	*Fermi*	Long	14:13:03.581	34.304 ± 1.448	–
GRB151021A	E194197	*Swift*	Long	01:29:12.57	110.2 ± 3.7	2.33
GRB151021B	E194374	*Fermi*	Long	18:59:28.923	7.229 ± 0.602	–
GRB151022577	E194592	*Fermi*	Short	13:51:02.089	0.32 ± 0.648	–
GRB151022A	E194587	*Swift*	Long	14:06:32.11	116.7 ± 23.0	–
GRB151023104	E194819	*Fermi*	Long	02:29:25.137	10.24 ± 3.874	–
GRB151023A	E194829	*Swift*	Long	13:43:04.81	10.66 ± 4.41	–
GRB151024179	E194977	*Fermi*	Ambiguous	04:17:53.560	4.608 ± 2.36	–
GRB151026169	E195464	*Fermi*	Long	04:03:06.559	53.248 ± 2.172	–
GRB151026523	E195558	*Fermi*	Long	12:32:38.939	63.232 ± 7.952	–
GRB151027A	E195704	*Swift*	Long	03:58:24.15	129.69 ± 5.55	0.81
GRB151027B	E195894	*Swift*	Long	22:40:40.66	80.0 ± 35.78	4.063
GRB151030A	E196738	*Fermi*	Long	23:58:22.637	116.482 ± 0.923	–
GRB151031A	E196797	*Swift*	Ambiguous	05:50:30.34	5.0 ± 2.24	1.167
GRB151107A	E198744	*Swift*	Ambiguous	17:19:36.61	–	–
GRB151107B	E198774	*Fermi*	Long	20:24:52.297	139.01 ± 6.446	–
GRB151111A	E199730	*Swift*	Long	08:33:23.41	76.93 ± 12.59	3.5
GRB151112A	E200054	*Swift*	Ambiguous	13:44:48.08	19.32 ± 31.24	4.1
GRB151114A	E200499	*Swift*	Ambiguous	09:59:34.15	4.86 ± 0.98	–
GRB151114645	E200556	*Fermi*	Long	15:28:24.519	34.816 ± 1.95	–
GRB151117442	E201325	*Fermi*	Long	10:36:59.778	58.561 ± 4.672	–
GRB151118A	E201488	*Swift*	Long	03:06:30.01	23.4 ± 10.5	–
GRB151118554	E201590	*Fermi*	Long	13:18:05.276	40.897 ± 10.555	–

Table 5.2: O1 GRBs analyzed by RAVEN (cont.)

GRB Designation	GRACEDB ID	Satellite	Classification (Long/Short)	Time (UT)	T_{90} (sec)	z
GRB151120A	E202051	*Fermi*	Long	08:22:53.704	28.416 ± 4.615	–
GRB151122A	E202680	*Fermi*	Long	17:00:45.006	51.2 ± 15.496	–
GRB151126293	E203647	*Fermi*	Long	07:01:17.239	8.448 ± 3.556	–
GRB151127A	E203926	*Swift*	Short	09:08:49.80	0.19 ± 0.04	–
GRB151129333	E204457	*Fermi*	Long	08:00:06.085	52.224 ± 7.455	–
GRB151130160	E204759	*Fermi*	Long	03:50:50.019	20.224 ± 4.36	–
GRB151202565	E205336	*Fermi*	Short	13:33:49.808	0.704 ± 1.159	–
GRB151205A	E206159	*Swift*	Long	15:46:00.93	62.8 ± 12.3	–
GRB151205B	E206220	*Swift*	Short	21:43:14.53	1.4 ± 0.2	–
GRB151210041	E207067	*Fermi*	Long	00:59:16.643	37.633 ± 1.717	–
GRB151210A	E207108	*Swift*	Long	03:12:56.46	94.9 ± 10.4	–
GRB151211672	E207484	*Fermi*	Long	16:07:34.520	40.897 ± 1.493	–
GRB151212030	E207571	*Fermi*	Long	00:42:58.448	22.272 ± 4.615	–
GRB151212064	E207580	*Fermi*	Long	01:32:04.972	13.312 ± 1.557	–
GRB151215A	E208363	*Swift*	Long	03:01:28.95	17.8 ± 1.0	2.59
GRB151218857	E209019	*Fermi*	Ambiguous	20:33:31.902	3.328 ± 2.064	–
GRB151219567	E209201	*Fermi*	Long	13:36:22.836	62.72 ± 6.676	–
GRB151222A	E210017	*Fermi*	Short	08:10:13.624	0.768 ± 0.362	–
GRB151227A	E211350	*Fermi*	Ambiguous	01:44:07.692	3.389 ± 0.602	–
GRB151227B	E211389	*Fermi*	Long	05:13:48.856	43.008 ± 0.362	–
GRB151228A	E211622	*Swift*	Short	03:05:12.38	0.27 ± 0.01	–
GRB151228B	E211819	*Swift*	Long	22:47:14.94	48.0 ± 16.0	–
GRB151229A	E211923	*Swift*	Short	06:50:27.96	1.78 ± 0.44	–
GRB151229486	E211974	*Fermi*	Short	11:40:06.468	0.16 ± 0.202	–
GRB151231A	E212513	*Fermi*	Long	10:37:47.522	71.425 ± 0.724	–
GRB151231B	E212538	*Fermi*	Short	13:38:08.165	0.832 ± 0.405	–
GRB160101A	E212652	*Fermi*	Long	00:43:53.610	4.669 ± 0.602	–
GRB160101B	E212695	*Fermi*	Long	05:10:12.860	22.013 ± 1.619	–
GRB160102500	E213004	*Fermi*	Long	11:59:22.628	25.344 ± 1.493	–
GRB160102936	E213119	*Fermi*	Long	22:28:16.995	10.496 ± 1.145	–
GRB160104A	E213522	*Swift*	Long	11:24:10.84	16.2 ± 2.5	–
GRB160104918	E213637	*Fermi*	Long	22:01:26.729	44.288 ± 1.493	–
GRB160106948	E214439	*Fermi*	Long	22:45:30.929	39.425 ± 0.724	–
GRB160107A	E214740	*Fermi*	Long	22:20:41.502	113.922 ± 17.755	–
GRB160111115	E215644	*Fermi*	Long	02:45:03.291	26.88 ± 6.676	–

Chapter 6

Localization and Broadband Follow-up of the Gravitatonal Wave Transient GW150914[1]

"The night is dark and full of terrors!"
Melissandre, *Game of Thrones*

In chapters 2 and 5 I reported on the electromagnetic (EM) follow-up effort planned for Advanced LIGO's first Observing Run (O1). This chapter will detail an example follow-up campaign conducted after the first direct detection of a gravitational wave (GW) transient.

A new generation of GW detectors is making deeper searches possible for GW signal events, with the Laser Interferometer Gravitational-wave Observatory (LIGO; LIGO Scientific Collaboration et al. 2015) becoming operational and Virgo (Acernese et al., 2015b) expected to join soon (Abbott et al., 2016f). As discussed in chapter 2, some of the most promising astrophysical sources of GW signals are also expected to produce broadband electromagnetic (EM) emission (as well as neutrinos). This has created new and exciting opportunities for multimessenger astronomy.

[1]The text of this chapter is adapted from Abbott et al. (2016b). My direct contributions to this work were primarily to act as editor for section 7.1, to compose and edit Fig. 7.1, to track down data for Table 7.2, and to assist in communicating results with the wider community of astronomer partners as detailed in chapters 2 and 5. I was also actively involved in discussions that led to the content presented in all sections.

In a compact binary coalescence (CBC) event, a close binary featuring two neutron stars (NSs), two black holes (BHs), or a NS and a BH experiences orbital decay as the binary emits gravitational radiation. In a NS binary—a binary neutron star (BNS) or neutron star–black hole (NSBH) system—we expect EM signatures due to energetic outflows at different timescales and wavelengths. If a relativistic jet forms, we may observe a short gamma-ray burst (GRB) lasting on the order of one second or less, followed by X-ray, optical and radio afterglows of hours–days duration (e.g., Eichler et al. 1989; Narayan et al. 1992; Nakar 2007; Berger 2014b; Fong et al. 2015). Rapid neutron capture in the sub-relativistic ejecta (e.g., Lattimer & Schramm 1976) is hypothesized to produce a kilonova or macronova, an optical–near infrared signal lasting hours–weeks (e.g., Li & Paczyński 1998). Eventually, we may observe a radio blast wave from this sub-relativistic outflow, detectable for months–years (e.g., Nakar & Piran 2011). Furthermore, several seconds prior to or tens of minutes after merger, we may see a coherent radio burst lasting milliseconds (e.g., Hansen & Lyutikov 2001; Zhang 2014). In short, a NS binary can produce EM radiation over a wide range of wavelengths and time scales. On the other hand, in the case of a stellar-mass binary black hole (BBH), the current consensus is that no significant EM counterpart emission is expected, except for those in highly improbable environments pervaded by large ambient magnetic fields or baryon densities.

6.0.1 Past Follow-up Efforts

The first gravitational-wave-triggered EM observations were carried out during the 2009–2010 science run of the initial LIGO and Virgo detectors (Abadie et al., 2012), featuring real-time searches for un-modeled GW bursts and CBCs. GW candidates were identified—typically within 30 minutes—and their inferred sky locations were used to plan follow-up observations with over a dozen optical and radio telescopes on the ground plus the *Swift* satellite (Gehrels et al., 2004b). Tiles were assigned to individual facilities to target known galaxies that were consistent with the GW localizations and that were within the 50 Mpc nominal BNS detectability horizon. Eight GW candidates were followed up. Though none of the GW candidates were significant enough to constitute detections and the EM candidates found were judged to be merely serendipitous sources (Evans et al., 2012; Aasi et al., 2014), the program demonstrated the feasibil-

ity of searching in real-time for GW transients, triggering follow-up, and analyzing GW and EM observations jointly.

In preparing for Advanced detector operations, the LIGO and Virgo collaborations worked with the broader astronomy community to set up an evolved and greatly expanded EM follow-up program.[2] Seventy-four groups with access to ground- and space-based facilities joined, of which 63 were operational during Advanced LIGO's first observing run (O1). Instead of centrally planning the assignment of tiles to facilities, we have set up a common EM bulletin board for facilities and observers to announce, coordinate, and visualize the footprints and wavelength coverage of their observations. The new program builds on the Gamma-ray Coordinates Network (GCN)[3] system that has long been established for broadband follow-up of GRBs. We distribute times and sky positions of event candidates via machine-readable Notices, and participating facilities communicate the results of observations via short bulletins, GCN Circulars. A key difference is that GRB Notices and Circulars are instantly public, whereas GW alert Notices and follow-up Circulars currently are restricted to participating groups until the event candidate in question has been published. After four high-confidence GW events have been published, further high-confidence GW event candidates will be promptly released to the public.

After years of construction and commissioning, the Advanced LIGO detectors at Livingston, Louisiana, and Hanford, Washington, began observing in 2015 September with about 3.5 times the distance reach (> 40 times the sensitive volume) of the earlier detectors. A strong GW event was identified shortly after the pre-run calibration process was completed. Deep analysis of this event, initially called G184098 and later given the name GW150914, is discussed in detail in Abbott et al. (2016a) and companion papers referenced therein. In this paper we describe the initial low-latency analysis and event candidate selection (section 6.1), rapid determination of likely sky localization (section 6.2), and the follow-up EM observations carried out by partner facilities (section 6.3, section 6.4). For analyses of those observations, we refer the reader to the now-public GCN Circulars[4] and to a number of recent papers. We end with a brief discussion of EM counterpart detection prospects for future events.

[2] See program description at http://www.ligo.org/scientists/GWEMalerts.php.

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

[4] All Circulars related to GW150914 are collected at http://gcn.gsfc.nasa.gov/other/GW150914.gcn3

6.1 Data Analysis and Discovery

As configured for O1, four low-latency pipelines continually search for transient signals that are coincident in the two detectors within the 10 ms light travel time separating them. Coherent WaveBurst (cWB; Klimenko et al. 2015) and Omicron+LALInference Burst (oLIB; Lynch et al. 2015) both search for un-modeled GW bursts (Abbott et al., 2016d) and produce sky position probability maps. At the start of O1, the pipelines GSTLAL (its name derived from GStreamer and LAL, the LIGO Algorithm Library; Cannon et al. 2012b; Messick et al. 2016) and Multi-Band Template Analysis (MBTA; Adams et al. 2015) search for compact binary merger transients using matched filter techniques. At the start of O1, both the online GSTLAL and MBTA were configured to search specifically for NS binary mergers, but their template banks were expanded to include binary black hole systems after the discovery of GW150914. Since CBC waveforms can be precisely computed from general relativity, GSTLAL and MBTA are more sensitive to CBC signals than the burst search pipelines are. The BAYESTAR algorithm (Singer & Price, 2016) calculates sky maps for all CBC candidates. All four detection pipelines report candidates within a few minutes of data acquisition.

LIGO conducted a series of engineering runs throughout Advanced LIGO's construction and commissioning to prepare to collect and analyze data in a stable configuration. The eighth engineering run (ER8) began on 2015 August 17 at 15:00[5] and critical software was frozen by August 30. The rest of ER8 was to be used to calibrate the detectors, to carry out diagnostic studies, to practice maintaining a high coincident duty cycle, and to train and tune the data analysis pipelines. Calibration was complete by September 12 and O1 was scheduled to begin on September 18. On 2015 September 14, cWB reported a burst candidate to have occurred at 09:50:45 with a network signal-to-noise ratio (S/N) of 23.45 and an estimated false alarm rate (FAR) < 0.371 yr^{-1} based on the available (limited at that time) data statistics. Also, oLIB reported a candidate with consistent timing and S/N. No candidates were reported at this time by the low-latency GSTLAL and MBTA pipelines, ruling out a BNS or NSBH merger.

Although GW150914 was detected before O1 officially began, the LIGO and Virgo collaborations decided to send an alert to partner facilites because the preliminary FAR estimate satisfied our planned alert threshold of 1 month^{-1}, indicating a potentially in-

[5] All dates and times are in UT.

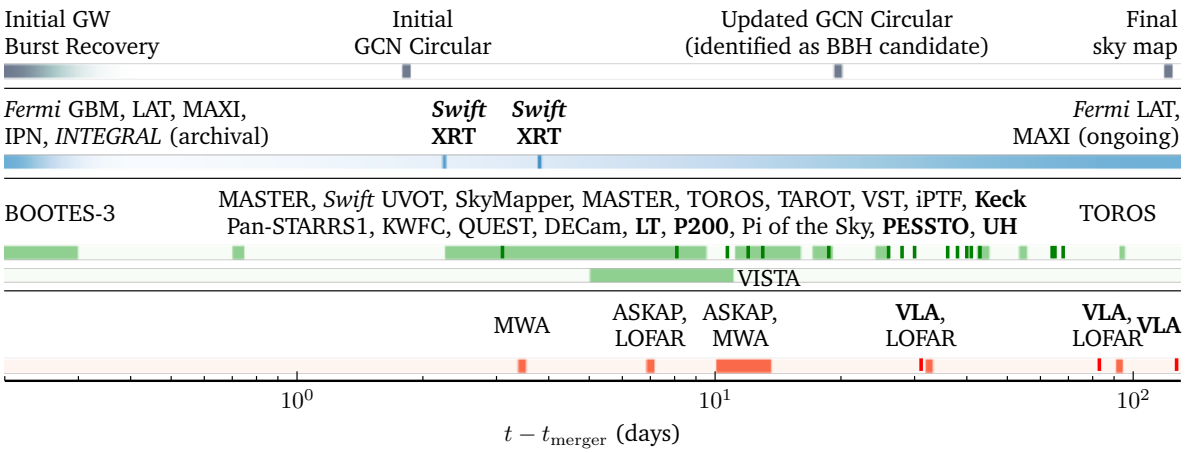

Figure 6.1: Timeline of observations of GW150914, separated by band and relative to the time of the GW trigger. The top row shows GW information releases. The bottom four rows show high-energy, optical, near-infrared, and radio observations respectively. Optical spectroscopy and narrow-field radio observations are indicated with darker tick marks and boldface text.

teresting source (see chapter 2). Though we had not planned to disseminate real-time GCN Notices before the formal start of O1, most of the computing infrastructure was already in place. Basic data quality checks were done within hours of GW150914; both interferometers were stable and the data stream was free of artifacts (Abbott et al., 2016b). A cWB sky map was available 17 min after the data were recorded, and a LALInference Burst (LIB) sky map after 14 hr. After extra data integrity checks and an update to the GCN server software, these two sky maps were communicated to observing partners in a GCN Circular nearly two days after the event occurred (GCN 18330). Mass estimates were not released in this initial Circular, and observers may have assumed the event was associated with a BNS system or a GW burst (e.g., from a nearby core-collapse supernova). The knowledge that GW150914 was consistent with a BBH inspiral and merger was only shared later, on October 3 (GCN 18388). Figure 6.1 shows the chronology of the GW detection alerts and follow-up observations.

The data were subsequently re-analyzed offline with two independent matched-filter searches, both using a template bank that includes both NS binary and BBH mergers. The waveform was confirmed to be consistent with a BBH merger and this information was shared with observers about 3 weeks after the event (GCN 18388). The FAR was evaluated with the data collected through 20 October, reported to be less than 1 in 100 years (GCN 18851; Abbott et al. 2016c), and ultimately determined to be

much lower. The final results of the offline search are reported in Abbott et al. (2016a).

6.2 Sky Maps

We produce and disseminate all-sky probability maps using a sequence of algorithms with increasing accuracy and computational cost. Here, we compare four location estimates: the prompt cWB and LIB localizations that were initially shared with observing partners, and the rapid BAYESTAR localization and the final localization from LALInference. All four are shown separately shown in Fig. 6.2, and their overlap is illustrated in Fig. 6.3.

cWB performs a constrained maximum likelihood estimate of the reconstructed signal on a sky grid (Klimenko et al., 2015) weighted by the detectors' antenna patterns (Essick et al., 2015) and makes minimal assumptions about the waveform morphology. With two detectors, this amounts to restricting the signal to only one of two orthogonal GW polarizations throughout most of the sky. LIB performs Bayesian inference assuming the signal is a sinusoidally modulated Gaussian (Lynch et al., 2015). While this assumption may not perfectly match the data, it is flexible enough to produce reliable localizations for a wide variety of waveforms, including BBH inspiral-merger-ringdown signals (Essick et al., 2015). BAYESTAR produces sky maps by triangulating the times, amplitudes, and phases on arrival supplied by the CBC pipelines (Singer & Price, 2016). BAYESTAR was not available promptly because the low-latency CBC searches were not configured for BBHs; the localization presented here is derived from the offline CBC search. LALInference performs full forward modeling of the data using a parameterized CBC waveform which allows for BH spins and detector calibration uncertainties (Veitch et al., 2015). It is the most accurate method for CBC signals but takes the most time due to the high dimensionality. We present the same LALInference map as Abbott et al. (2016e), with a spline interpolation procedure to include the potential effects of calibration uncertainties. The BAYESTAR and LALInference maps were shared with observers on 2016 January 13 (GCN 18858), at the conclusion of the O1 run. Since GW150914 is a CBC event, we consider the LALInference map to be the most accurate, authoritative, and *final* localization for this event. This map has a 90% credible region with area 620 deg^2.

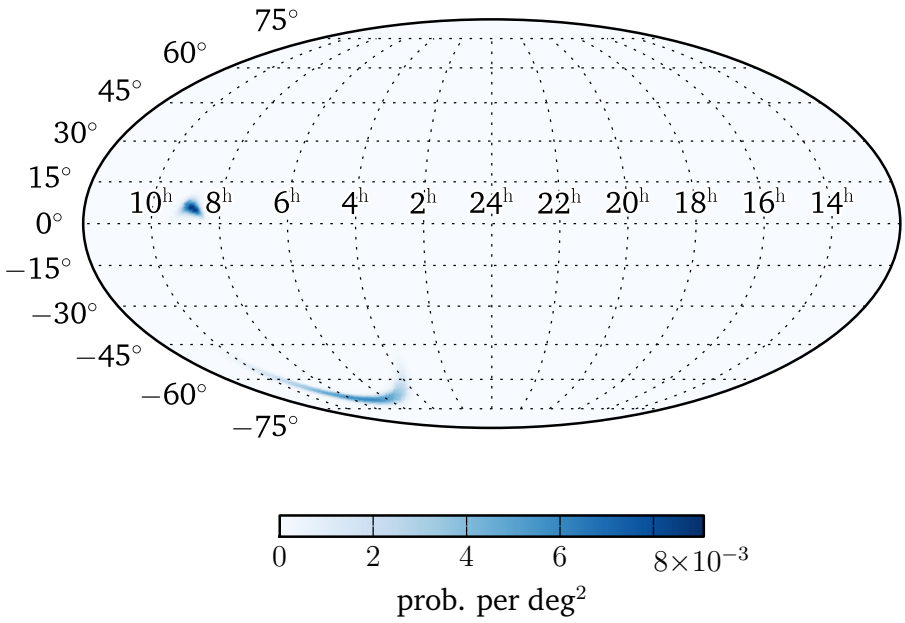

(a) cWB

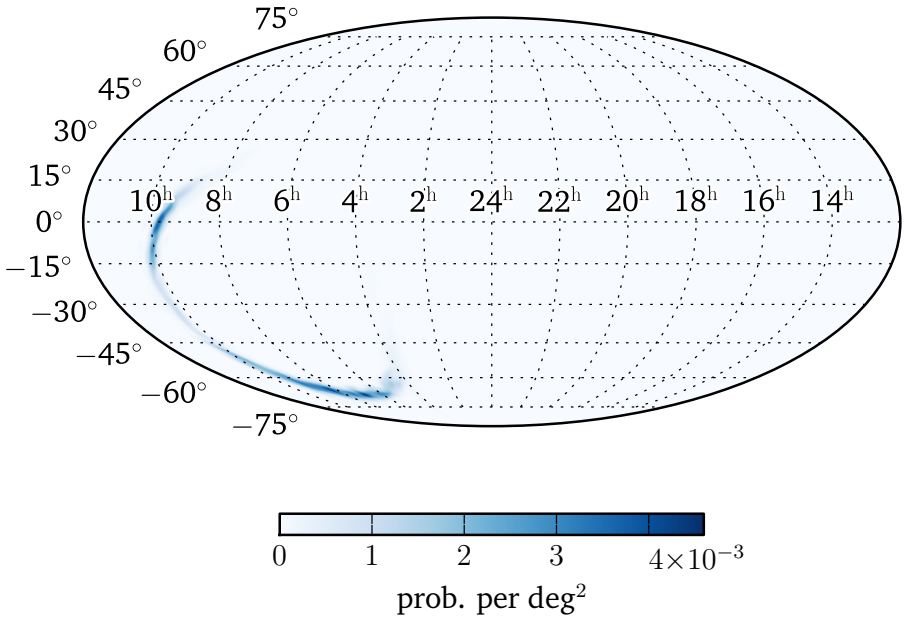

(b) LIB

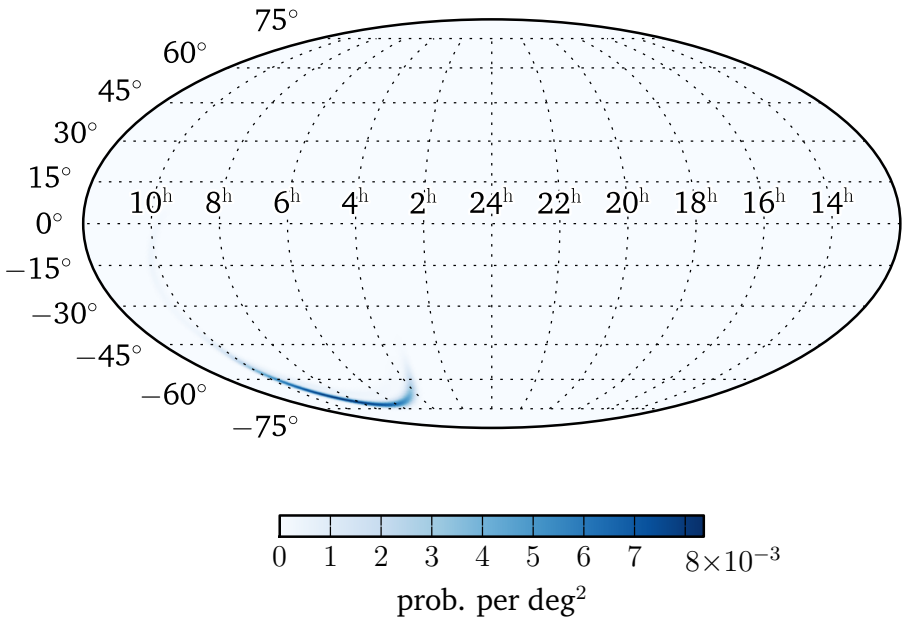

(c) BAYESTAR

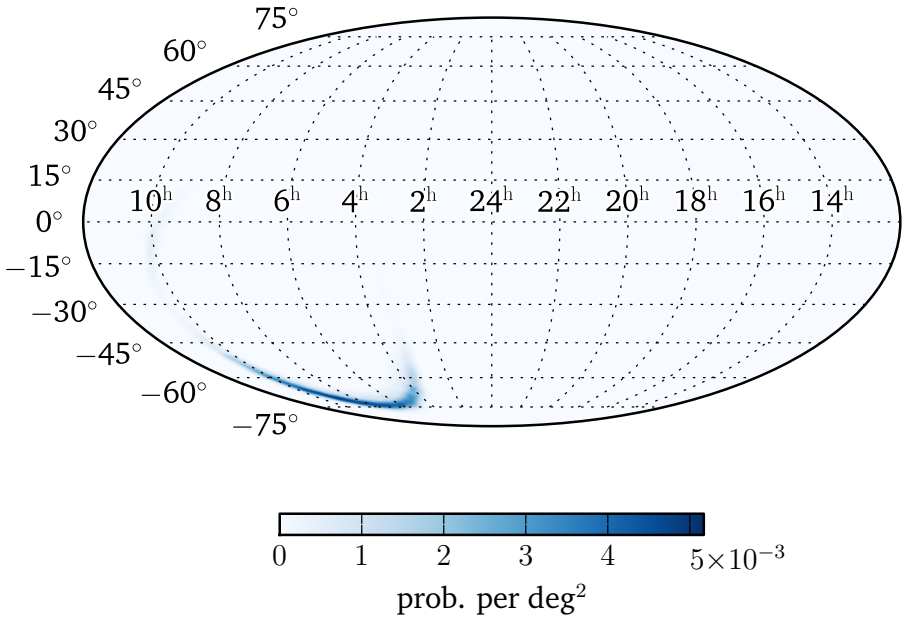

(d) LALInference

Figure 6.2: Sky location posterior for transient source GW150914, as computed by (a) cWB, (b) LIB, (c) BAYESTAR, and (d) LALInference, all in equatorial (J2000.0) coordinates. Each of these distributions forms a segment of an annulus consistent with the relative time-of-arrival of GW150914 in the H1 and L1 detectors; see Fig. 6.3.

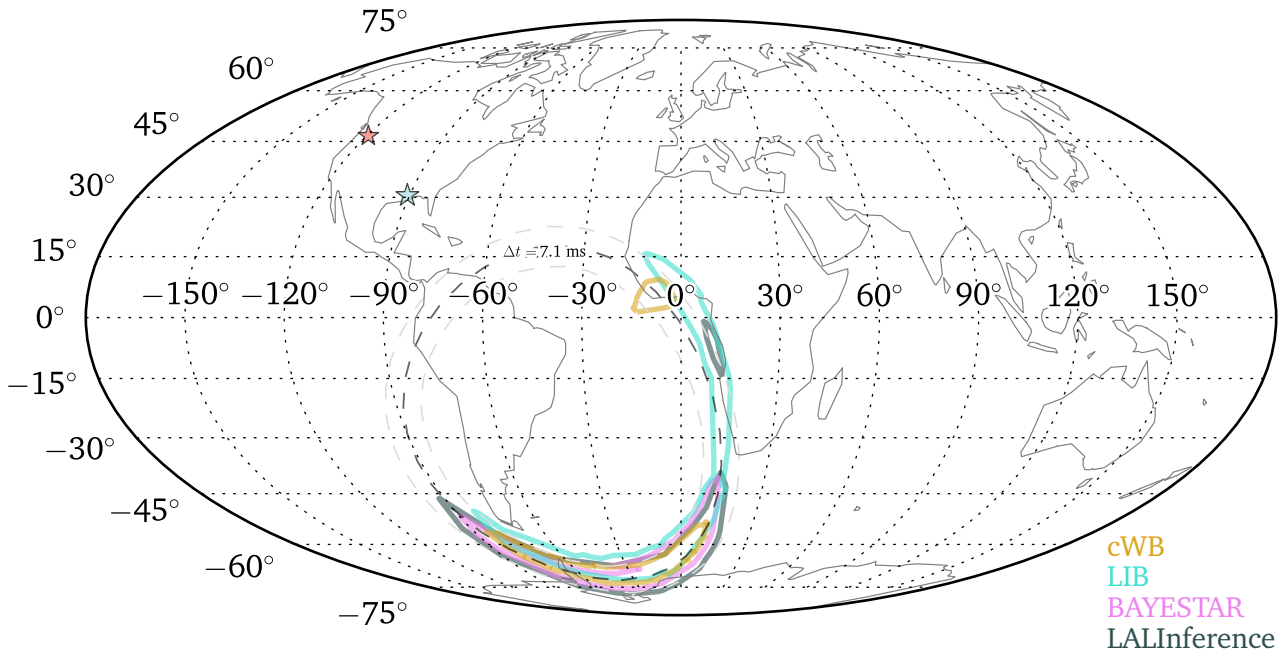

Figure 6.3: Comparison of different GW sky maps, showing the 90% credible level contours for each algorithm. This is a Mollweide projection in geographic coordinates. Locations of the Hanford (H1, red) and Livingston (L1, blue) detectors are each marked with a star, and the dashed annulus corresponding to a relative time-of-arrival difference $\Delta t = t_{\mathrm{L1}} - t_{\mathrm{H1}} = 7.1 \pm 0.5$ ms is shown for comparison.

All of the sky maps agree qualitatively, favoring a broad, long section of arc in the Southern hemisphere, and to some extent a shorter section of almost the same arc located near the equator. While the majority of LIB's probability is concentrated in the Southern hemisphere, a significant portion of the 90% confidence region extends into the Northern hemisphere. The LALInference sky map shows much less support in the Northern hemisphere which is likely associated with the stronger constraints provided by full and precise CBC waveforms. Finally the cWB localization algorithm also supports an isolated hot spot near $\alpha \sim 9^{\mathrm{h}}, \delta \sim 5^{\circ}$. While all algorithms assume elliptical polarization throughout most of the sky, cWB's assumptions are relaxed near this island, where the detector responses make it possible to distinguish other polarizations (see Fig. 6.3).

The dominant feature in all four sky maps is an annulus with polar angle $\theta_{\mathrm{HL}} \approx 45^{\circ}$, determined via triangulation by the time-of-arrival difference Δt_{HL} between the Hanford and Livingston detectors. However, refinements are possible by demanding amplitude and phase consistency across the detector network, and from the mildly directional antenna patterns of the LIGO detectors (Kasliwal & Nissanke 2014; Singer et al. 2014; see chapter 2. In particular, the detectors' antenna patterns dominate the modulation around the ring for unmodelled reconstructions through a correlation with the inferred distance to source (Essick et al., 2015). As shown in Fig. 6.3, the algorithms all infer annuli with polar angles that are consistent at the 1σ level.

6.2.1 Comparison of Gravitational Wave Sky Maps

For this signal event, cWB produces smaller confidence regions than the other algorithms. While cWB produces reasonably accurate maps for typical BBH signals, it can systematically misestimate the sizes of large confidence regions (Essick et al., 2015). The other algorithms are more self-consistent in this regime. Only the LALInference results account for calibration uncertainty (systematic errors in the conversion of the photocurrent into the GW strain signal). Because systematic errors in the calibration phase affect the measured arrival times at the detectors, the main effect is to broaden the position uncertainty relative to the other sky maps.

To quantify sky map comparisons, we can use the intersections of the 90% confidence regions as well as the fidelity $F(p,q) = \int \sqrt{pq}\,\mathrm{d}\Omega \in [0,1]$ between two maps p

and q. All these measures show that the sky maps are similar but not identical. Typically, this level of quantitative disagreement is distinguishable by eye and has been observed in large simulation campaigns (Singer et al., 2014; Berry et al., 2015; Essick et al., 2015) for approximately 10%–20% of the simulated signals. This even includes the bi-modality of LIB's $\theta_{\rm HL}$ distribution (see inset of Fig. 6.3), which is associated with a degeneracy in the signal's chirality (or equivalently the binary's inclination) at different points around the ring. Similar features were noted for BNS systems as well (Singer et al., 2014).

6.3 Follow-up Observations

Twenty-five participating teams of observers responded to the GW alert to mobilize satellites and ground-based telescopes spanning 19 orders of magnitude in electromagnetic wavelength. Fig. 6.1 shows a timeline of the observations. Observations and archival analysis started shortly after the candidate was announced, two days after the event was recorded. Most facilities followed tiling strategies based on the cWB and LIB sky maps. Some groups, considering the possibility of a NS merger or core-collapse supernova, selected fields based on the areal density of nearby galaxies or targeted the Large Magellanic Cloud (LMC) (e.g., Annis et al. 2016). Had the BBH nature of the signal been promptly available, most groups would not have favored local galaxies because LIGO's range for BBH mergers is many times larger than for BNSs. The campaign is summarized in detail below.

6.3.1 Gamma-ray and X-ray Observations

The *Fermi* Gamma-ray Burst Monitor (GBM; Meegan et al. 2009b), *INTEGRAL* (Winkler et al., 2003), and the InterPlanetary Network (IPN; Hurley et al. 2010b) searched for prompt high-energy emission temporally coincident with the GW event. Although no GRB in coincidence with GW150914 was reported, an off-line analysis of the Fermi GBM (8 keV–40MeV) data revealed a weak transient with duration of $\sim 1\,$s (Connaughton et al., 2016b). A similar analysis was performed for the instruments onboard *INTEGRAL* (Winkler et al., 2003), particularly the spectrometer's anticoinci-

dence shield (SPI–ACS, von Kienlin et al., 2003, 75 keV–1 MeV) [6]. No significant signals were detected, setting upper limits on the hard X-ray fluence at the time of the event (Savchenko et al., 2016). (See Savchenko et al. (2016) and Connaughton et al. (2016b) for a detailed comparison of the *INTEGRAL* SPI-ACS upper limit and the *Fermi* GBM transient.) Data from the six-spacecraft, all-sky, full-time monitor IPN, (*Odyssey*–HEND, *Wind*–Konus, *RHESSI*, *INTEGRAL*–SPI-ACS, and *Swift*–BAT[7]) revealed no bursts around the time of GW150914 apart from the weak GBM signal (Hurley et al., in preparation).

The *Fermi* Large Area Telescope (LAT), MAXI and *Swift* searched for high-energy afterglow emission. The LIGO localization first entered the *Fermi* LAT field of view (FOV) at 4200 s after the GW trigger and was subsequently observed in its entirety over the next 3 hr and every 3 hr thereafter at GeV energies (Fermi-LAT collaboration, 2016). The entire region was also imaged in the 2–20 keV X-ray band by the MAXI Gas Slit Camera (GSC; Matsuoka et al. 2009) aboard the International Space Station (ISS) from 86 to 77 min before the GW trigger and was re-observed during each subsequent ~ 92 min orbit. The *Swift* X-ray Telescope (XRT; Burrows et al. 2005b) followed up the GW event starting 2.25 days after the GW event, and covered 5 tiles containing 8 nearby galaxies for a total $\sim$0.3 deg^2 area in the 0.3–10 keV energy range. A 37-point tiled observation of the Large Magellanic Cloud was executed a day later. *Swift* UV/Optical Telescope (UVOT) provided simultaneous ultraviolet and optical observations, giving a broadband coverage of 80% of the *Swift* XRT FOV. Details of these observations are given in Evans et al. (2016).

6.3.2 Optical and Near-IR Observations

The optical and near-infrared observations fell into roughly two stages. During the first week, wide FOV (1–10 deg^2) telescopes tiled large areas to identify transient candidates, and then larger but narrower FOV telescopes obtained classification spectroscopy and further photometry. The wide FOV instruments included DECam on the CTIO Blanco telescope (Flaugher et al., 2015; Dark Energy Survey Collaboration et al., 2016), the Kiso Wide Field Camera (KWFC, J-GEM; Sako et al., 2012), La Silla QUEST (Baltay

[6] *INTEGRAL*'s coded-mask imager (IBIS, Ubertini et al., 2003, 20–200 keV) was pointed far outside the GW localization region.

[7] *Swift* Burst Alert Telescope did not intersect the GW localization at the time of the trigger

et al., 2007), MASTER-SAAO twin robotic telescope of the Global MASTER Robotic Net (Lipunov et al., 2010), the Palomar 48 inch Oschin telescope (P48) as part of the intermediate Palomar Transient Factory (iPTF; Law et al. 2009b), Pan-STARRS1 (Kaiser et al., 2010), SkyMapper (Keller et al., 2007), TAROT-La Silla (Boër et al., 1999, node of the TAROT-Zadko-Algerian National Observatory-C2PU collaboration), and the VLT Survey Telescope (VST@ESO; Capaccioli & Schipani, 2011, GRAvitational Wave Inaf TeAm, Brocato et al. 2016 in preparation)[8] in the optical band, and the Visible and Infrared Survey Telescope (VISTA@ESO; Emerson et al. 2006)[9] in the near infrared. They represent different classes of instruments ranging in diameter from 0.25 to 4 m and reaching apparent magnitudes from 18 to 22.5. About one third of these facilities followed a galaxy-targeted observational strategy, while the others tiled portions of the GW sky maps covering 70–590 deg^2. A narrow (arcminute) FOV facility, the 1.5 m EABA telescope in Bosque Alegre operated by the TOROS collaboration (M. Diaz et al. 2016, in prep.), also participated in the optical coverage of the GW sky maps. *Swift* UVOT observed simultaneously with XRT, giving a broadband coverage of 80% of the *Swift* XRT FOV.

A few tens of transient candidates identified by the wide-field telescopes were followed on the 10 m Keck II telescope (DEIMOS; Faber et al., 2003), the 2 m Liverpool Telescope (LT; Steele et al. 2004), the Palomar 200 inch Hale telescope (P200; Bracher 1998), the 3.6 m ESO New Technology Telescope (within the Public ESO Spectroscopic Survey of Transient Objects, PESSTO; Smartt et al. 2015), and the University of Hawaii 2.2 m telescope (SuperNovae Integral Field Spectrograph, SNIFS).

An archival search for bright optical transients was conducted in the CASANDRA-3 all-sky camera database of BOOTES-3 (Castro-Tirado et al., 2012) and the all-sky survey of the Pi of the Sky telescope (Mankiewicz et al., 2014), both covering the entire southern sky map. The BOOTES-3 images are the only observations simultaneous to GW150914 available to search for prompt/early optical emission. They reached a limiting magnitude of 5 due to poor weather conditions (GCN 19022). The Pi of the Sky telescope images were taken 12 days after GW150914 and searched to find slowly fading transients brighter than $R < 11.5$ mag (GCN 19034).

[8]ESO proposal ID:095.D-0195,095.D-0079
[9]ESO proposal ID:095.D-0771

6.3.3 Radio Observations

The radio telescopes involved in the EM follow-up program have the capability to observe a wide range of frequencies with different levels of sensitivity, and a range of FOVs covering both the northern and southern skies. The Low Frequency Array (LOFAR; van Haarlem et al. 2013) and the Murchison Widefield Array (MWA; Tingay et al. 2013) are phased array dipole antennas sensitive to meter wavelengths with large FOVs ($\approx 50\,\mathrm{deg}^2$ with uniform sensitivity for the LOFAR observations carried out as part of this follow-up program; and up to $1200\,\mathrm{deg}^2$ for Murchison Widefield Array). The Australian Square Kilometer Array Pathfinder (ASKAP; Schinckel et al. 2012) is an interferometric array composed of thirty-six 12 m-diameter dish antennas. The Karl G. Jansky Very Large Array (VLA; Perley et al. 2009) is a twenty-seven antenna array, with dishes of 25 m diameter. Both Australian Square Kilometer Array Pathfinder and the VLA are sensitive from centimeter to decimeter wavelengths.

MWA started observing $3\,\mathrm{d}$ after the GW trigger with a 30 MHz bandwidth around a central frequency of 118 MHz and reached a root mean square (RMS) noise level of about 40 mJy/beam. The ASKAP observations used the five-element Boolardy Engineering Test Array (BETA; Hotan et al. 2014), which has a FOV of $\approx 25\,\mathrm{deg}^2$ and FWHM synthesized beam of $1' - 3'$. These observations were performed with a 300 MHz bandwidth around a central frequency of 863.5 MHz, from ≈ 7 to $\approx 14\,\mathrm{d}$ after the GW trigger, reaching RMS sensitivities of $1-3$ mJy/beam. LOFAR conducted three observations from ≈ 7 d to ≈ 3 months following the GW trigger, reaching a RMS sensitivity of ≈ 2.5 mJy/beam at 145 MHz, with a bandwidth of 11.9 MHz and a spatial resolution of $\approx 50''$. ASKAP, LOFAR, and MWA all performed tiled observations aimed at covering a large area of the GW region.

The VLA performed follow-up observations of GW150914 from ≈ 1 month to ≈ 4 months after the GW trigger[10], and targeted selected candidate optical counterparts detected by iPTF. VLA observations were carried out in the most compact array configuration (D configuration) at a central frequency of $\approx 6\,\mathrm{GHz}$ (primary beam FWHP of $\approx 9'$, and synthesized beam FWHP of $\approx 12''$). The RMS sensitivity of these VLA observations was $\approx 8 - 10\mu$Jy/beam.

Instrument acknowledgements for the observations reported here appear in ap-

[10]VLA/15A-339, PI: A. Corsi

pendix A.

6.4 Coverage

Using the GW data by itself, we can only constrain the position of the source on the sky to an area of 590 deg^2 (90% confidence). The inferred redshift is $z = 0.09^{+0.03}_{-0.04}$, corresponding to a luminosity distance of 410^{+160}_{-180} Mpc (Abbott et al., 2016e).

By far the most complete coverage of the area is at the highest energies. The *INTEGRAL* SPI–ACS provided the largest effective area in the 75 keV–1 MeV range, albeit with significantly varying detection efficiency. Owing to its nearly omnidirectional response, it had a full coverage of the GW probability map (GCN 18354; Savchenko et al. 2016). *Fermi* GBM captured 75% of the localization at the time of the GW trigger and the entire area by 25 min after (GCN 18339). *Fermi* LAT observations started 4200 s after the trigger and the entire localization continued to be observed every three hours.

Coverage in X-rays is complete down to 10^{-9} erg cm^{-2} s^{-1} with the MAXI observations, but relatively sparse at fainter flux, with the *Swift* XRT tiles spanning about 5 deg^2 and enclosing a probability of $\sim$0.3% in the energy range 0.3–10 keV to a depth of 10^{-13}–10^{-11} erg cm^{-2} s^{-1} (GCNs 18331, 18346).

Optical facilities together tiled about 870 deg^2 and captured a containment probability of 57% of the initial LIB sky map, though only 36% of the final LALInference sky map that was available after the observations were completed. The depth varies widely among these facilities. MASTER and Pan-STARRS1 covered the most area with their observations, while large areas also were covered by the iPTF, Dark Energy Camera (DECam), VST@ESO and La Silla–QUEST. The contained probability of the initial sky maps is dominated by MASTER, DECam, Pan-STARRS1, La Silla–QUEST and VST@ESO, while the final sky map is contained best by MASTER, DECam and VST@ESO. Relatively small area and contained probability were covered by facilities that targeted nearby galaxies. The only near-infrared facility, VISTA@ESO, covered 70 deg^2 and captured a containment probability of 8% of the final LALInference sky map.

The radio coverage is also extensive, with a contained probability of 86%, dominated by MWA in the 118 MHz band (GCN 18345).

Deep photometry, broadband observations and spectroscopy identified the majority of the candidates to be normal population type Ia and type II supernova, with a few dwarf novae and active galactic nuclei (AGNs), all very likely unrelated to GW150914. Candidate classification, comparison of redshift with the GW distance, and use of source age are crucial constraints to rule candidates in and out. Detailed discussions of candidate selection, spectroscopic and broadband follow-up are presented in survey-specific publications about iPTF candidates (Kasliwal et al., 2016) and about PESSTO follow-up of Pan-STARRS1 candidates (Smartt et al., 2016).

6.5 Sensitivity

Since the follow-up program was primarily designed to search for counterparts to BNS and NSBH systems, it is interesting to note that the observational campaign would have provided powerful constraints. If GW150914 had an associated short GRB, it would easily have been detectable within the actual distance of GW150914 (Berger, 2014c). A BNS coalescence at ~ 70 Mpc, the average distance at which it could have been detected during O1 (Martynov et al., 2016), might produce a short gamma-ray burst X-ray afterglow 11 hours after the burst with isotropic-equivalent flux of 2×10^{-11} to 6×10^{-8} erg cm^{-2} s^{-1} (Berger, 2014c). A BNS at that distance might also produce a kilonova (e.g., Metzger et al. 2010; Barnes & Kasen 2013; Tanaka & Hotokezaka 2013; Grossman et al. 2014) within a few days after merger with apparent magnitude in the range 17–24. This range lies within the depth reached in the optical band but also in the near IR where observations (Tanvir et al., 2013; Berger et al., 2013a) suggest the bulk of the emission. Finally, this BNS system might have produced radio afterglows in the range of 0.1–15 mJy (e.g., Hotokezaka & Piran 2015). We note that many of these possible counterparts could have been detected by the EM follow-up effort associated with GW150914. Radio observations from wide field facilities were sensitive to the bright counterparts at low frequencies while the VLA to fainter counterparts at frequencies above a few GHz.

6.6 Conclusions

Transient source GW150914 is consistent with the inspiral and merger of two BHs of masses 36^{+5}_{-4} and $29^{+4}_{-4}\,\mathrm{M}_{\odot}$ respectively, resulting in the formation of a final BH of mass $62^{+4}_{-4}\,\mathrm{M}_{\odot}$ (Abbott et al., 2016a). In classical general relativity, a *vacuum* BBH merger does not produce any EM or particle emission whatsoever. Whereas supermassive BBHs in galactic centers may appear as dual AGNs or have other distinctive EM signatures due to interactions with gas or magnetic fields, stellar BBH systems are not expected to possess detectable EM counterparts. The background gas densities and magnetic field strengths should therefore be typical of the interstellar medium, which are many orders of magnitude smaller than the environments of EM bright supermassive BBHs. Although GW150914 is loud in GWs and expected to be absent in all EM bands, thorough follow-up observations were pursued to check for EM emission. Future EM follow-ups of GW sources will shed light on the presence or absence of firm EM counterparts and astrophysical processes that may trigger EM emission from these systems.

The EM campaign following GW150914 successfully demonstrates the capability of the observing partners to cover large swaths of the sky localization area, to identify candidates, and to activate larger telescopes for photometric and spectroscopic characterization within a few days of an event. We note that the information about the source's BBH nature and updated sky maps were sent out twenty days and four months after the event, respectively. This resulted in some instruments covering much less of the probability region or to the required depth of GW150914 than they may have planned for. We expect future alerts to be issued within tens of minutes with more information about the signal type, and more rapid updates of the maps. The follow-up efforts would have been sensitive to a wide range of emission expected from BNS or NSBH mergers. However, the widely variable sensitivity reached across the sky localization area continues to be a challenge for an EM counterpart search.

The number of galaxies (with luminosities $L \geq 0.1L^{\star}$; Blanton et al., 2003) within the comoving volume of 10^{-2} Gpc^3 corresponding to the 90% credible area of the LALInference sky map and within the 90% confidence interval distance is $\sim 10^5$. Such a number makes it impossible to identify the host galaxy in the absence of an EM counterpart detection. The presence of a third GW detector such as Virgo would have

improved the sky localization of GW150914 to a few tens of square degrees both for the un-modeled and CBC searches. The future addition of more GW detectors to the global network (Abbott et al., 2016f) will significantly improve the efficiency of searches for EM counterparts.

In summary, we have described the EM follow-up program carried out for the first GW source detected by Advanced LIGO. Within two days of the initial tentative detection of GW150914 a GCN circular was sent to EM follow-up partners alerting them to the event and providing them with initial sky maps. Twenty-five EM observing teams mobilized their resources, and over the ensuing three months observations were performed with a diverse array of facilities over a broad wavelength range (from radio to γ-ray). Findings from those observations are being disseminated in other papers. The localization and broadband follow-up of this GW event constitutes an important first step in a new era of gravitational wave multi-messenger astronomy.

This chapter is based on LIGO document LIGO-P1500227-v11.

Part IV

"And Now His Watch Has Ended"

Chapter 7

Unafraid of the Dark

Though my soul may set in darkness, it will rise in perfect light;
I have loved the stars too fondly to be fearful of the night.

Sarah Williams, *The Old Astronomer (To His Pupil)*

"Let me tell you about 'scared.' Your heart is beating so hard, I can feel it through your hands. There's so much blood and oxygen pumping through your brain, it's like rocket fuel. Right now you could run faster and you can fight harder, you can jump higher than ever in your life, and you are so alert it's like you can slow down time. What's wrong with scared? Scared is a superpower! Your superpower! There **is** *danger in this room, and guess what? It's you!"*

The 12th Doctor, addressing a frightened child, *Doctor Who* (episode *Listen*)

7.1 Looking Back

At last we begin the end with a bit of an author's note. When I was very young, I was terrified of black holes. The very concept, that you could never escape, no matter how hard you try, made them the perfect monster under my bed. My grandmother gave me a book at age seven that was my first exposure to the concept of a black hole; it gave me nightmares for weeks. Twenty-one years later, black holes have played a direct role in a major scientific milestone.

It is an intellectually galvanizing time in the field of time domain transient astronomy. With the escalation of instrument sensitivity and computational power afforded by Moore's Law, "traditional" observatories around the world, in Earth orbit, and throughout the solar system are now able to cover the full electromagnetic spectrum on timescales ranging from fractions of a second to several days, yielding a rich population of totally new astrophysical sources (*e.g.* relativistic tidal disruption flares,

fast radio bursts) and a deeper understanding of old, well-trodden favorites (*e.g.* supernovae). Moreover, the massive Laser Interferometer Gravitational-wave Observatory (LIGO: LIGO Scientific Collaboration et al. 2015) experiment has made the first confirmed direct detection of gravitational waves (GW) in the form of transient source GW150914 (Abbott et al., 2016c). Many of the astrophysical sources expected to populate LIGO's frequency band, particularly those involving the coalescence of ultra-compact objects (*e.g.* a neutron star and a black hole), are also expected to produce fantastically bright electromagnetic emission that will be visible to a battalion of high energy, optical, and radio observatories both in space and on the ground.

All of this comes at a deeply fortuitous time when excitement in the community runs astronomically high. Rapid, wide-field optical facilities such as the Palomar Transient Factory (PTF: Law et al. 2009a) have made significant headway in resolving a number of outstanding astrophysical mysteries in recent years, while concurrent observations of transients discovered by LIGO will give unprecedented insight into the nature of some of the most energetic explosive events in the universe. With next-generation instruments like the Zwicky Transient Facility (ZTF: Bellm 2014) and the Large Synoptic Survey Telescope (LSST: Ivezic et al. 2008), astronomers will have high-cadence access to the transient sky on a scale never before attained. The confluence of all these experiments will allow us to constrain the physics of fast relativistic transient phenomena, and give us a new sense on the local universe, doubtless revealing much that cannot be anticipated.

7.2 Looking Ahead

Analysis of data from the whole of O1 is ongoing. Future studies will contain constraints on the properties of short, hard γ-ray bursts assuming the progenitor system involves inspiraling compact binary systems; this will involve a mock population of sources, similar to chapter 3, but using actual O1 strain spectra from both H1 and L1. Based on the weak γ-ray transient observed 0.4 seconds after GW150914 (Connaughton et al., 2016a), it will also be necessary to compare the population of 76 GRBs reported during O1 (chapter 5) to GW signal candidates uncovered by matched filter searches with template banks that include binary black hole systems.

In the coming years, during O2 and beyond, the LIGO and Virgo (Acernese et al., 2015a) detectors will be online for about one full observing run per year on average, slated to last between 3–6 months each. When the LIGO detectors are in observing mode, the electromagnetic followup effort will be heavily invested in using high-energy instruments such as *Swift* and *Fermi* and optical facilities such as ZTF. At a source rate density of $\sim 10^{-4}$ Mpc^{-3} yr^{-1}, Advanced LIGO expects to discover $\sim$40 binary neutron star merger systems per operational year at design sensitivity (Abadie et al., 2010), with a rich opportunity for regular followup of these sources in the X-ray and optical bands. While scanning for counterparts and during LIGO instruments' downtime, we could use optical facilities such as ZTF to search for generic fast relativistic transients (orphan afterglows, tidal disruption flares, etc.) and to follow these up with *Swift* XRT. At a detection rate of $\sim$1–10 yr$^{-}1$, this will give a substantial population of sources.

Much more distant sources (i.e., out into the Hubble flow at redshifts $z \gtrsim 0.5$–5) will be far too quiet for LIGO to detect, but interesting tests can still be done with optical astronomy. The wide field-of-view of PTF/ZTF's survey camera, and its rapid imaging and slew capabilities, have allowed fast-cadence searches to be done at optical sensitivities of $m \approx 20$ mag. This has made it feasible to search for optical transients that fade on $\sim$hour-long timescales, such as PTF11agg and iPTF14yb. Since most GRBs (both long and short) are beamed away from Earth, their intrinsic rate density may be a factor $\sim$100 times higher than the *observed* rate density. The standard model of GRB emission predicts that as ejecta from the central explosion decelerate substantially, the jet will expand laterally, so that the afterglow illuminates an increasing fraction of the sky at diminishing intensity (Rhoads, 1999). The afterglows of most bursts should therefore be visible several hours (or days) after the initial explosion, even though the high-energy trigger is beamed away from Earth. Discovery of such an "orphan afterglow" would instantly confirm our broad picture of relativistic beaming in GRBs. Observing the off-axis afterglows of nearby compact binary mergers will also open a wealth of new astrophysics as has been discussed at great length in chapters 2 and 3.

The woods are lovely, dark and deep. But I have promises to keep, and miles to go before I sleep. — ROBERT FROST

Bibliography

Aasi, J., et al. 2014, ApJS, 211, 7

Abadie, J., et al. 2012, Astronomy and Astrophysics, 539, A124

—. 2010, Classical and Quantum Gravity, 27, 173001

Abbott, B., et al. 2016a, Phys. Rev. Lett., 116, 061102

—. 2016b, Class. Quantum Grav., submitted, arXiv:1602:03844

—. 2016c, Phys. Rev. D, submitted, arXiv:1602.03839

—. 2016d, Phys. Rev. D, submitted, arXiv:1602.03843

—. 2016e, Phys. Rev. Lett., submitted, arXiv:1602.03840

—. 2016f, Living Reviews in Relativity, 19, 1

Abbott, B. P., et al. 2016a, ApJL, 818, L22

—. 2016b, ArXiv e-prints

—. 2016c, Physical Review Letters, 116, 061102

—. 2016d, Living Reviews in Relativity, 19

Abdo, A. A., et al. 2009, ApJL, 706, L138

Acernese, F., et al. 2015a, Classical and Quantum Gravity, 32, 024001

—. 2015b, Classical and Quantum Gravity, 32, 024001

Adams, T., et al. 2015, ArXiv e-prints

Ahn, C. P., et al. 2014, ApJS, 211, 17

Allen, B. 2005, Phys. Rev. D, 71, 062001

Allen, B., Anderson, W. G., Brady, P. R., Brown, D. A., and Creighton, J. D. E. 2012, Phys. Rev. D, 85, 122006

Annis, J., Soares-Santos, M., Berger, E.,, et al. 2016, ArXiv e-prints

Baltay, C., et al. 2007, PASP, 119, 1278

Barnes, J. and Kasen, D. 2013, ApJ, 775, 18

Barthelmy, S. D., et al. 2005, Space Science Reviews, 120, 143

Basri, R. and Jacobs, D. W. 2003, IEEE Trans. Pattern Anal. Mach. Intell., 25, 218

Becker, A. C., et al. 2004, ApJ, 611, 418

Bellm, E. C. 2014, ArXiv e-prints

Berger, E. 2014a, Annual Review of Astronomy and Astrophysics, 52, 43

—. 2014b, Annual Reviews of Astronomy and Astrophysics, 52, 43

—. 2014c, Annual Reviews of Astronomy and Astrophysics, 52, 43

Berger, E., Fong, W., and Chornock, R. 2013a, ApJL, 774, L23

Berger, E., et al. 2003, Nature, 426, 154

—. 2013b, ApJ, 779, 18

Berry, C. P. L., et al. 2015, ApJ, 804, 114

Blackburn, L., Briggs, M. S., Burns, E.,, et al. 2015, GCN, 18339, 1

Blanton, M. R., et al. 2003, ApJ, 592, 819

Boër, M., et al. 1999, Astronomy and Astrophysics, Supplement, 138, 579

Bracher, K. 1998, Mercury, 27, 4

Burrows, D. N., et al. 2005a, Space Science Reviews, 120, 165

—. 2005b, Space Science Reviews, 120, 165

Butler, N. R., Bloom, J. S., and Poznanski, D. 2010, ApJ, 711, 495

Cannon, K., et al. 2012a, ApJ, 748, 136

—. 2012b, ApJ, 748, 136

Capaccioli, M. and Schipani, P. 2011, The Messenger, 146, 2

Castro-Tirado, A. J., et al. 2012, in Astronomical Society of India Conference Series, Vol. 7, Astronomical Society of India Conference Series, 313–320

Castro-Tirado, A. J., Tello, J. C., Cunniffe, R.,, et al. 2016, GCN, 19022, 1

Cenko, S. B., et al. 2009, ApJ, 693, 1484

Cenko, S. B., et al. 2013, ApJ, 769, 130

—. 2015, ApJL, 803, L24

Christensen, L., Hjorth, J., and Gorosabel, J. 2004, Astronomy and Astrophysics, 425, 913

Connaughton, V., et al. 2016a, ArXiv e-prints

—. 2016b, ArXiv e-prints

Creighton, J. and Anderson, W. 2011, Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis. (Wiley-VCH)

Cummings, J. R., et al. 2013, GRB Coordinates Network, 15446

Cwiek, A., Zarnecki, A. F., Mankiewicz, L.,, et al. 2016, GCN, 19034, 1

Dark Energy Survey Collaboration, et al. 2016, ArXiv e-prints

D'Avanzo, P. 2015, Journal of High Energy Astrophysics, 7, 73

Dermer, C. D., Chiang, J., and Mitman, K. E. 2000, ApJ, 537, 785

Djorgovski, S. G., et al. 1997, Nature, 387, 876

Eichler, D., Livio, M., Piran, T., and Schramm, D. N. 1989, Nature, 340, 126

Emerson, J., McPherson, A., and Sutherland, W. 2006, The Messenger, 126, 41

Essick, R. 2015, LIGO-T1300567

Essick, R., Vitale, S., Katsavounidis, E., Vedovato, G., and Klimenko, S. 2015, ApJ, 800, 81

Evans, P. A., et al. 2009, MNRAS, 397, 1177

—. 2012, ApJS, 203, 28

Evans, P. A., Kennea, J. A., Barthelmy, S.,, et al. 2015a, GCN, 18331, 1

—. 2015b, GCN, 18346, 1

Evans, P. A., et al. 2016, ArXiv e-prints

Faber, S. M., et al. 2003, in Proceedings of the SPIE, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood, 1657–1669

Fermi-LAT collaboration. 2016, ArXiv e-prints

Ferrigno, C., Savchenko, V., Mereghetti, S.,, et al. 2015, GCN, 18354, 1

Fishman, G. J., Meegan, C. A., Wilson, R. B., Paciesas, W. S., and Pendleton, G. N. 1992, in NASA Conference Publication, Vol. 3137, NASA Conference Publication, ed. C. R. Shrader, N. Gehrels, & B. Dennis, 26–34

Fitzpatrick, G. and Jenke, P. 2013, GRB Coordinates Network, 15455

Flaugher, B., et al. 2015, Astronomical Journal, 150, 150

Fong, W., Berger, E., Margutti, R., and Zauderer, B. A. 2015, ApJ, 815, 102

Frail, D. A., et al. 2001, ApJL, 562, L55

Fruchter, A. S., et al. 1999, ApJL, 519, L13

Fynbo, J. P. U., et al. 2009, ApJS, 185, 526

Gal-Yam, A., et al. 2006, ApJ, 639, 331

Gehrels, N., et al. 2004a, ApJ, 611, 1005

—. 2004b, ApJ, 611, 1005

Golenetskii, S., Aptekar, R., Frederiks, D., Pal'Shin, V., Oleynik, P., Ulanov, M., Svinkin, D., and Cline, T. 2014, GRB Coordinates Network, 15889, 1

Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. 2005, ApJ, 622, 759

Granot, J., Ramirez-Ruiz, E., and Perna, R. 2005, ApJ, 630, 1003

Granot, J. and Sari, R. 2002, ApJ, 568, 820

Greiner, J., Hartmann, D. H., Voges, W., Boller, T., Schwarz, R., and Zharikov, S. V. 2000, Astronomy and Astrophysics, 353, 998

Grossman, D., Korobkin, O., Rosswog, S., and Piran, T. 2014, MNRAS, 439, 757

Guetta, D. and Della Valle, M. 2007, ApJL, 657, L73

Guiriec, S., et al. 2011, ApJL, 727, L33

Hansen, B. M. S. and Lyutikov, M. 2001, MNRAS, 322, 695

Harry, I. W. and Fairhurst, S. 2011, Phys. Rev., D83, 084002

Hartle, J. B. 2003, Gravity : an introduction to Einstein's general relativity (Addison Wesley)

Heise, J., Zand, J. I., Kippen, R. M., and Woods, P. M. 2001, in Gamma-ray Bursts in the Afterglow Era, ed. E. Costa, F. Frontera, & J. Hjorth, 16

Hogg, D. W., Baldry, I. K., Blanton, M. R., and Eisenstein, D. J. 2002, ArXiv Astrophysics e-prints

Hotan, A. W., et al. 2014, Publications of the Astron. Soc. of Australia, 31, e041

Hotokezaka, K. and Piran, T. 2015, MNRAS, 450, 1430

Huang, Y. F., Dai, Z. G., and Lu, T. 2002, MNRAS, 332, 735

Hulse, R. A. and Taylor, J. H. 1975, ApJL, 195, L51

Hurley, K., et al. 2010a, in American Institute of Physics Conference Series, Vol. 1279, American Institute of Physics Conference Series, ed. N. Kawai & S. Nagataki, 330–333

Hurley, K., et al. 2010b, in American Institute of Physics Conference Series, Vol. 1279, American Institute of Physics Conference Series, ed. N. Kawai & S. Nagataki, 330–333

Hurley, K., et al. 2014, GRB Coordinates Network, 15888, 1

Ivezic, Z., et al. 2008, ArXiv e-prints

Jakobsson, P., Malesani, D., Hjorth, J., Fynbo, J. P. U., and Milvang-Jensen, B. 2011, Advances in Space Research, 47, 1416

Kaiser, N., et al. 2010, in Proceedings of the SPIE, Vol. 7733, Ground-based and Airborne Telescopes III, 77330E

Kanti, M. and Jupp, P. E. 1999, Directional Statistics (Wiley)

Kaplan, D., Croft, S., Bannister, K.,, et al. 2015, GCN, 18345, 1

Kasliwal, M. M., et al. 2016, ArXiv e-prints

Kasliwal, M. M. and Nissanke, S. 2014, ApJL, 789, L5

Keller, S. C., et al. 2007, Publications of the Astron. Soc. of Australia, 24, 1

Klebesadel, R. W., Strong, I. B., and Olson, R. A. 1973, ApJL, 182, L85

Klimenko, S., et al. 2016, Phys. Rev. D, 93, 042004

—. 2015, arXiv:1511.05999

Klimenko, S., Yakushin, I., Mercer, A., and Mitselmakher, G. 2008, Classical and Quantum Gravity, 25, 114029

Laher, R. R., et al. 2014, PASP, 126, 674

Lattimer, J. M. and Schramm, D. N. 1976, ApJ, 210, 549

Law, N. M., et al. 2009a, PASP, 121, 1395

—. 2009b, PASP, 121, 1395

Law, N. M., Rutledge, R. E., and Kulkarni, S. R. 2004, MNRAS, 350, 1079

Li, L.-X. and Paczyński, B. 1998, ApJL, 507, L59

Lien, A., Sakamoto, T., Gehrels, N., Palmer, D. M., Barthelmy, S. D., Graziani, C., and Cannizzo, J. K. 2014, ApJ, 783, 24

LIGO Scientific Collaboration, et al. 2015, Classical and Quantum Gravity, 32, 074001

LIGO Scientific Collaboration and Virgo. 2015a, GCN, 18330, 1

—. 2015b, GCN, 18388, 1

—. 2016a, GCN, 18858, 1

—. 2016b, GCN, 18851, 1

Lipunov, V., et al. 2010, Advances in Astronomy, 2010, 30

Lipunov, V. M., et al. 2007, Astronomy Reports, 51, 1004

Lynch, R., Vitale, S., Essick, R., Katsavounidis, E., and Robinet, F. 2015, ArXiv e-prints

Lyons, L. 2008, Ann. Appl. Stat., 3, 887

Maggiore, M. 2008, Gravitational Waves, Volume 1: Theory and Experiments (Oxford University Press)

Mankiewicz, L., et al. 2014, in Revista Mexicana de Astronomia y Astrofisica, vol. 27, Vol. 45, Revista Mexicana de Astronomia y Astrofisica Conference Series, 7–11

Martynov, D. V., et al. 2016, ArXiv e-prints

Matsuoka, M., et al. 2009, Publications of the ASJ, 61, 999

Mazzali, P. A., et al. 2006, Nature, 442, 1018

Meegan, C., et al. 2009a, ApJ, 702, 791

—. 2009b, ApJ, 702, 791

Messick, C., et al. 2016, https://dcc.ligo.org/LIGO-P1600009

Metzger, B. D. and Berger, E. 2012, ApJ, 746, 48

Metzger, B. D., et al. 2010, MNRAS, 406, 2650

Metzger, M. R., Djorgovski, S. G., Kulkarni, S. R., Steidel, C. C., Adelberger, K. L., Frail, D. A., Costa, E., and Frontera, F. 1997, Nature, 387, 878

Nakar, E. 2007, Physics Reports, 442, 166

Nakar, E. and Piran, T. 2003, New Astronomy, 8, 141

—. 2011, Nature, 478, 82

Narayan, R., Paczynski, B., and Piran, T. 1992, ApJL, 395, L83

Nishizawa, A., Taruya, A., Hayama, K., Kawamura, S., and Sakagami, M.-A. 2009, Phys. Rev. D, 79, 082002

Nuttall, L. K., et al. 2015, Classical and Quantum Gravity, 32, 245005

Ofek, E. O., et al. 2012, PASP, 124, 62

Paciesas, W. S., et al. 1999, ApJS, 122, 465

Perley, R., et al. 2009, IEEE Proceedings, 97, 1448

Piran, T. 2004, Reviews of Modern Physics, 76, 1143

Planck Collaboration, et al. 2013, ArXiv e-prints

Rau, A., Greiner, J., and Schwarz, R. 2006, Astronomy and Astrophysics, 449, 79

Rau, A., et al. 2009, PASP, 121, 1334

Rau, A., Ofek, E. O., Kulkarni, S. R., Madore, B. F., Pevunova, O., and Ajello, M. 2008, ApJ, 682, 1205

Rhoads, J. E. 1997, ApJL, 487, L1

—. 1999, ApJ, 525, 737

Ryde, F., et al. 2010, ApJL, 709, L172

Rykoff, E. S., et al. 2005, ApJ, 631, 1032

Sako, S., et al. 2012, in Proceedings of the SPIE, Vol. 8446, Proceedings of the SPIE, 6

Sari, R., Piran, T., and Halpern, J. P. 1999, ApJL, 519, L17

Savchenko, V., et al. 2016, ApJL, 820, L36

Schinckel, A. E., Bunton, J. D., Cornwell, T. J., Feain, I., and Hay, S. G. 2012, 8444, 84442A

Singer, L. P., et al. 2015, ApJ, 806, 52

Singer, L. P. and Price, L. R. 2016, Phys. Rev. D, 93, 024013

Singer, L. P., et al. 2014, ApJ, 795, 105

Singer, L. P., et al. 2014, ApJ, 795, 105

Smartt, S. J., et al. 2016, ArXiv e-prints

—. 2015, Astronomy and Astrophysics, 579, A40

Smith, J. 2015, LIGO-G1500642

Steele, I. A., et al. 2004, in Proceedings of the SPIE, Vol. 5489, Ground-based Telescopes, ed. J. M. Oschmann, Jr., 679–692

Sutton, P. J., et al. 2010, New J. Phys., 12, 053034

Tanaka, M. and Hotokezaka, K. 2013, ApJ, 775, 113

Tanvir, N. R., Levan, A. J., Fruchter, A. S., Hjorth, J., Hounsell, R. A., Wiersema, K., and Tunnicliffe, R. L. 2013, Nature, 500, 547

Taylor, J. H. and Weisberg, J. M. 1982, ApJ, 253, 908

The LIGO Scientific Collaboration. 2016, LIGO Open Science Center release of GW150914, https://losc.ligo.org/events/GW150914/, DOI 10.7935/K5MW2F23

The LIGO Scientific Collaboration and the Virgo Collaboration. 2016a, ArXiv e-prints

—. 2016b, ArXiv e-prints

Tingay, S. J., et al. 2013, Publications of the Astron. Soc. of Australia, 30, e007

Ubertini, P., et al. 2003, Astronomy and Astrophysics, 411, L131

Urban, A. L., Stephens, B., Romie, J., and Giardina, D. 2015, LIGO-G1501296

Usman, S. A., et al. 2015, ArXiv e-prints

van Haarlem, M. P., et al. 2013, Astronomy and Astrophysics, 556, A2

Veitch, J., et al. 2015, Phys. Rev. D, 91, 042003

von Kienlin, A., et al. 2003, Astronomy and Astrophysics, 411, L299

Wald, R. M. 1984, General relativity (University of Chicago Press)

Wanderman, D. and Piran, T. 2010, MNRAS, 406, 1944

Wang, L. J. and Dai, Z. G. 2013, ApJL, 774, L33

Williamson, A., Biwer, C., Fairhurst, S., Harry, I., Macdonald, E., Macleod, D., and Predoi, V. 2014, Phys. Rev., D90, 122004

Winkler, C., et al. 2003, Astronomy and Astrophysics, 411, L1

Woosley, S. E., Eastman, R. G., and Schmidt, B. P. 1999, ApJ, 516, 788

Wright, E. L., et al. 2010, Astronomical Journal, 140, 1868

Wu, X. F., Gao, H., Ding, X., Zhang, B., Dai, Z. G., and Wei, J. Y. 2014, ApJL, 781, L10

Zhang, B. 2014, ApJL, 780, L21

Appendix A

Acknowledgements: Localization and Broadband Follow-up of the Gravitational Wave Transient GW150914[1]

The authors gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO 600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research, for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council of Scientific and Industrial Research of India, Department of Science and Technology, India, Science & Engineering Research Board

[1]This appendix is adapted from a supplement to Abbott et al. (2016b), and is designed to capture acknowledgements for all the various instruments used to conduct observations reported in that work.

(SERB), India, Ministry of Human Resource Development, India, the Spanish Ministerio de Economía y Competitividad, the Conselleria d'Economia i Competitivitat and Conselleria d'Educació, Cultura i Universitats of the Govern de les Illes Balears, the National Science Centre of Poland, the European Commission, the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, the Hungarian Scientific Research Fund (OTKA), the Lyon Institute of Origins (LIO), the National Research Foundation of Korea, Industry Canada and the Province of Ontario through the Ministry of Economic Development and Innovation, the National Science and Engineering Research Council Canada, Canadian Institute for Advanced Research, the Brazilian Ministry of Science, Technology, and Innovation, Russian Foundation for Basic Research, the Leverhulme Trust, the Research Corporation, Ministry of Science and Technology (MOST), Taiwan and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, MPS, INFN, CNRS and the State of Niedersachsen/Germany for provision of computational resources.

The Australian SKA Pathfinder is part of the Australia Telescope National Facility which is managed by CSIRO. Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. Establishment of the Murchison Radio-astronomy Observatory was funded by the Australian Government and the Government of Western Australia. ASKAP uses advanced supercomputing resources at the Pawsey Supercomputing Centre. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site.

AJCT acknowledges support from the Junta de Andalucía (Project P07-TIC-03094) and Univ. of Auckland and NIWA for installing of the Spanish BOOTES-3 station in New Zealand, and support from the Spanish Ministry Projects AYA2012-39727-C03-01 and 2015-71718R.

Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and

Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft and the Collaborating Institutions in the Dark Energy Survey.

The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, the National Optical Astronomy Observatory, the University of Nottingham, The Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University.

The DES data management system is supported by the National Science Foundation under Grant Number AST-1138766. The DES participants from Spanish institutions are partially supported by MINECO under grants AYA2012-39559, ESP2013-48274, FPA2013-47986, and Centro de Excelencia Severo Ochoa SEV-2012-0234. Research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478.

The *Fermi* LAT Collaboration acknowledges support for LAT development, operation and data analysis from NASA and DOE (United States), CEA/Irfu and IN2P3/CNRS (France), ASI and INFN (Italy), MEXT, KEK, and JAXA (Japan), and the K.A. Wallenberg Foundation, the Swedish Research Council and the National Space Board (Sweden). Science analysis support in the operations phase from INAF (Italy) and CNES (France) is also gratefully acknowledged. The *Fermi* GBM Collaboration acknowledges the support of NASA in the United States and DRL in Germany.

GRAWITA acknowledges the support of INAF for the project "Gravitational Wave

Astronomy with the first detections of adLIGO and adVIRGO experiments."

This works exploited data by INTEGRAL, an ESA project with instruments and science data centre funded by ESA member states (especially the PI countries: Denmark, France, Germany, Italy, Switzerland, Spain), and with the participation of Russia and the USA. The SPI ACS detector system has been provided by MPE Garching/Germany. We acknowledge the German INTEGRAL support through DLR grant 50 OG 1101.

IPN work is supported in the U.S. under NASA Grant NNX15AU74G.

This work is partly based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Intermediate Palomar Transient Factory (iPTF) project, a scientific collaboration among the California Institute of Technology, Los Alamos National Laboratory, the University of Wisconsin, Milwaukee, the Oskar Klein Center, the Weizmann Institute of Science, the TANGO Program of the University System of Taiwan, and the Kavli Institute for the Physics and Mathematics of the Universe. MMK and YC acknowledge funding from the National Science Foundation PIRE program grant 1545949. AAM acknowledges support from the Hubble Fellowship HST-HF-51325.01. Part of the research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

J-GEM is financially supported by KAKENHI Grant No. 24103003, 15H00774 and 15H00788 of MEXT Japan, 15H02069 and 15H02075 of JSPS, and "Optical and Near-Infrared Astronomy Inter-University Cooperation Program" supported by MEXT.

The Liverpool Telescope is operated on the island of La Palma by Liverpool John Moores University in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofisica de Canarias with financial support from the UK Science and Technology Facilities Council.

LOFAR, the Low Frequency Array designed and constructed by ASTRON, has facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. R. Fender acknowledges support from ERC Advanced Investigator Grant 267697.

MASTER Global Robotic Net is supported in parts by Lomonosov Moscow State Uni-

versity Development programm, Moscow Union OPTICA , Russian Science Foundation 16-12-00085, RFBR15-02-07875, National Research Foundation of South Africa.

We thank JAXA and RIKEN for providing MAXI data, KAKENHI Grant No. 24103002, 24540239, 24740186, and 23000004 of MEXT, Japan.

This work uses the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government Department of Industry and Science and Department of Education (National Collaborative Research Infrastructure Strategy: NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. The MWA acknowledges the iVEC Petabyte Data Store and the Initiative in Innovative Computing and the CUDA Center for Excellence sponsored by NVIDIA at Harvard University.

Pan-STARRS is supported by the University of Hawaii and the National Aeronautics and Space Administration's Planetary Defense Office under Grant No. NNX14AM74G. The Pan-STARRS-LIGO effort is in collaboration with the LIGO Consortium and supported by Queen's University Belfast. The Pan-STARRS1 Sky Surveys have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, and the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation Grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), and the Los Alamos National Laboratory. This work is based (in part) on observations collected at the European Organisation for Astronomical Research in the Southern Hemisphere, Chile as part of PESSTO, (the Public ESO Spectroscopic Survey for Transient Objects Survey) ESO programs188.D-3003, 191.D-0935.

Some of the data presented herein were obtained at the Palomar Observatory, Cali-

fornia Institute of Technology.

SJS acknowledges funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement n^{o} [291222] and STFC grants ST/I001123/1 and ST/L000709/1. MF is supported by the European Union FP7 programme through ERC grant number 320360. KM acknowledges support from the STFC through an Ernest Rutherford Fellowship

FOE acknowledges support from FONDECYT through postdoctoral grant 3140326.

Parts of this research were conducted by the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020.

Funding for *Swift* is provided by NASA in the US, by the UK Space Agency in the UK, and by the Agenzia Spaziale Italiana (ASI) in Italy. This work made use of data supplied by the UK *Swift* Science Data Centre at the University of Leicester. We acknowledge the use of public data from the *Swift* data archive.

The TOROS collaboration acknowledges support from Ministerio de Ciencia y Tecnología (MinCyT) and Consejo Nacional de Investigaciones Cientificas y Tecnológicas (CONICET) from Argentina and grants from the USA NSF PHYS 1156600 and NSF HRD 1242090.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

VST and VISTA observations were performed at the European Southern Observatory, Paranal, Chile. We acknowledge ESO personnel for their assistance during the observing runs.

This appendix is based on LIGO document LIGO-P1500227-v11.

Alexander L. Urban

Milwaukee, WI 53211, USA

EDUCATION

Aug. 2010 – May 2016 **Ph.D., Physics**

Leonard E. Parker Center for Gravitation,
Cosmology and Astrophysics,
University of Wisconsin–Milwaukee
3135 N. Maryland Ave., Milwaukee, WI 53211
Advisor: Dr. Patrick R. Brady

Thesis title: Monsters in the Dark:
High Energy Signatures of Black
Hole Formation with Multimessenger Astronomy

Aug. 2006 – May 2010 **B.S., Physics**

Allegheny College
520 N. Main St., Meadville, PA 16335
(Mathematics, Astronomy, Political Science minors)
Advisors: Dr. James Lombardi, Jr., Dr. Shafiqur Rahman

RESEARCH EXPERIENCE

Jan. 2011 – present **Graduate Research Assistant**

Leonard E. Parker Center for Gravitation,
Cosmology and Astrophysics,
University of Wisconsin–Milwaukee, Milwaukee, WI
Supervisor: Dr. Patrick R. Brady
Project(s): See *Statement of Research Interests* below.

June 2009 – Aug. 2010	**Undergaduate Research Assistant** Allegheny College, Meadville, PA *Supervisor*: Dr. Daniel Willey *Project*: Numerical and laboratory modeling of maser emission by He-OCS collisions within interstellar molecular clouds
Relevant skills and qualifications:	Experience with Python, HTML, Matlab, C and Fortran Experience with various data analysis software packages Designing and maintaining real-time code, optimized hardware Outstanding public speaking and communication

STATEMENT OF RESEARCH INTERESTS

Gravitational-wave (GW) detection with the ground-based LIGO and Virgo interferometers, as well as time-domain optical astronomy with the intermediate Palomar Transient Factory (iPTF) collaboration. Broadly, I am interested in high-energy transient phenomena and their electromagnetic counterparts, *e.g.* compact binary coalescence, γ-ray bursts (GRBs), optical afterglows, supernovae, kilonovae, and tidal disruption flares. I have contributed a real-time search for coincidence between GW candidates and GRBs from the *Swift* and *Fermi* satellites, and manage infrastruture that produces sky location estimates for the GW candidates within 2 minutes of coalescence time. I have also worked as part of a team that achieved the first unambiguous discovery of a long GRB by its optical afterglow, using this to prepare for "orphaned" afterglow searches in the upcoming era of fast, wide-field optical surveys such as the Zwicky Transient Facility (ZTF) and the Large Synoptic Survey Telescope (LSST).

PUBLICATIONS

(An asterisk (*) indicates this publication is peer-reviewed)

* B. P. Abbot, et al. (the LIGO Scientific Collaboration, the Virgo Collaboration, and electromagnetic astronomer partners), *Localization and Broadband Follow-up of the Gravitational Wave Transient GW150914*, arXiv:1602.08492 (2016; ADS entry).

* B. P. Abbot, et al. (the LIGO Scientific Collaboration and the Virgo Collaboration), *Observation of Gravitational Waves from a Binary Black Hole Merger*, PRL **116**, 6 (2016; ADS entry).

* S. B. Cenko, **A. L. Urban** et al., *iPTF14yb: The first discovery of a gamma-ray burst afterglow independent of a high-energy trigger*, ApJL **803**, L24 (2015; ADS entry).

* C. P. L. Berry, I. Mandel, H. Middleton, L. P. Singer, **A. L. Urban** et al., *Parameter estimation for binary neutron-star coalescences with realistic noise during the Advanced LIGO era*, ApJ **804**, 114 (2015; ADS entry).

* L. P. Singer, L. R. Price, B. Farr, **A. L. Urban** et al., *The first two years of electromagnetic follow-up with advanced LIGO and Virgo*, ApJ **795**, 105 (2014; ADS entry).

A. L. Urban et al., *Constraints on the cosmic event rate of fast relativistic transients from the iPTF survey: Orphan afterglows and dirty fireball explosions*, in preparation.

A. L. Urban et al., *Rapid identification of gravitational-wave counterparts to electromagnetic transients: Compact binary inspirals and short γ-ray bursts*, in preparation.

B. Farr et al. (including **A. L. Urban**), *Parameter estimation on gravitational waves from neutron star binaries with spinning components*, arXiv:1508.05336.

AWARDS & HONORS

2015	LSC (LIGO Scientific Collaboration) Fellow at the Livingston Observatory
2012–2015	UWM Physics Research Excellence Award
2010–2015	UWM Chancellor's Graduate Student Award
2014	Outstanding Poster Presentation by a Graduate Student, 2014 LIGO-Virgo Collaboration Meeting, Nice, France
2010	Best Senior Comprehensive Project Thesis (Allegheny College)
2009–2010	Outstanding Physics Senior Award (Allegheny College)
2006–2010	Allegheny College Trustee Scholarship

TEACHING EXPERIENCE

Spring 2015	**Undergraduate Physics Tutor** Dept. of Physics, University of Wisconsin–Milwaukee
Spring 2015	**Teaching Assistant** Physics 108 (introductory physics lab without calculus) Dept. of Physics, University of Wisconsin–Milwaukee
Spring 2011–2013	**Teaching Assistant** Physics 215 (calculus-based introductory EM lab) Dept. of Physics, University of Wisconsin–Milwaukee
Autumn 2012	**1-Day Guest Lecturer** Physics 717 (General Relativity & Gravitation) Dept. of Physics, University of Wisconsin–Milwaukee
Autumn 2010–2011	**Teaching Assistant** Physics 214 (calculus-based introductory mechanics lab) Dept. of Physics, University of Wisconsin–Milwaukee
Autumn 2011	**Teaching Assistant** Physics 109 (introductory mechanics discussion without calculus) Dept. of Physics, University of Wisconsin–Milwaukee