

**The needle in the 100 deg² haystack:
The hunt for binary neutron star mergers with LIGO and
Palomar Transient Factory**

Thesis by
Leo P. Singer

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2015
(Defended November 24, 2014)

© 2015

Leo P. Singer

All Rights Reserved

To the love of my life, my wife Kristin, and our precious son Isaac.

Bacon in his instruction tells us that the scientific student ought not to be as the ant, who gathers merely, nor as the spider who spins from her own bowels, but rather as the bee who both gathers and produces. All this is true of the teaching afforded by any part of physical science. Electricity is often called wonderful, beautiful; but it is so only in common with the other forces of nature. The beauty of electricity or of any other force is not that the power is mysterious, and unexpected, touching every sense at unawares in turn, but that it is under law, and that the taught intellect can even now govern it largely. The human mind is placed above, and not beneath it, and it is in such a point of view that the mental education afforded by science is rendered super-eminent in dignity, in practical application and utility; for by enabling the mind to apply the natural power through law, it conveys the gifts of God to man.

Michael Faraday, Lecture notes of 1858, quoted in *The Life and Letters of Faraday* (1870) by Bence Jones, Vol. 2, p. 404

Acknowledgments

This is LIGO Document Number LIGO-P1400223-v10. I carried out the work presented in this thesis within the LIGO Scientific Collaboration (LSC) and the Intermediate Palomar Transient Factory (iPTF) collaboration. The methods and results I present are under review and are potentially subject to change. The opinions expressed here are my own and not necessarily those of the LSC or iPTF.

I gratefully acknowledge funding from the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory, which provided support for this work. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the NSF and operates under cooperative agreement PHY-0107417. I thank the NSF for supporting my research directly through a Graduate Research Fellowship. This work is based on observations obtained with the Palomar 48-inch Oschin telescope and the robotic Palomar 60-inch telescope at the Palomar Observatory as part of the Intermediate Palomar Transient Factory 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. The work in this thesis is partly funded by *Swift* Guest Investigator Program Cycle 9 award 10522 (NASA grant NNX14AC24G) and Cycle 10 award 10553 (NASA grant NNX14AI99G).

Thank you, Mom, thank you Dad, for an upbringing full of love, learning, and love of learning.

Thank you, my wife Kristin, thank you, my son Isaac, for your love and for your patience with me.

Thank you, Susan Bates, for your tutoring in problem solving that has resonated with me from elementary school through every day of my scientific career.

Thank you, John Jacobson, Amanda Vehslage, Tandra Walker, and Dr. Philip Terry-Smith, for the most inspiring courses in my high school education, and for molding me into a responsible and well-rounded individual.

Thank you, Profs. Luis Orozco and Betsy Beise, for your mentoring and friendship as well as the University of Maryland undergraduate physics courses that I enjoyed so much. Thank you for initiating me into physics research, and for sending me to graduate school so well prepared.

Thank you, Prof. Alan Weinstein, for being an outstanding (and, when necessary, forbearing) thesis advisor, for engineering the many wonderful collaborations that I have been a part of at Caltech, and for showing me how to thrive within a Big Science experiment.

Thank you, Prof. Shri Kulkarni, for recruiting me into PTF, for engineering a totally original cross-disciplinary research opportunity in physics and astronomy, and for placing trust in me. I am continually in awe of how that trust has paid off. I thank my colleagues in PTF for welcoming me into their highly capable and exciting team.

Thank you, Prof. Christian Ott, for teaching me two formative courses. I was able to write BAYESTAR, my greatest contribution so far to Advanced LIGO, only because the latter of these courses (Ay 190: Computational Astrophysics) was fresh in my head.

Thank you, Prof. David Reitze, for making me feel like the success of Advanced LIGO depends upon me. (I think that you inspire that same feeling in everyone at LIGO Laboratory.)

Thank you, Rory Smith, my officemate, for ducking good-naturedly whenever I wanted to chuck a chair out the window of 257 West Bridge. (Despite many strong oaths, no chairs were actually chucked during the writing of this thesis.)

Thank you, Nick Fotopoulos, Larry Price, Brad Cenko, and Mansi Kasliwal, for your collaboration and friendship throughout my studies, friendships that I hope to keep and to nurture.

Abstract

The Advanced LIGO and Virgo experiments are poised to detect gravitational waves (GWs) directly for the first time this decade. The ultimate prize will be joint observation of a compact binary merger in both gravitational and electromagnetic channels. However, GW sky locations that are uncertain by hundreds of square degrees will pose a challenge. I describe a real-time detection pipeline and a rapid Bayesian parameter estimation code that will make it possible to search promptly for optical counterparts in Advanced LIGO. Having analyzed a comprehensive population of simulated GW sources, we describe the sky localization accuracy that the GW detector network will achieve as each detector comes online and progresses toward design sensitivity. Next, in preparation for the optical search with the iPTF, we have developed a unique capability to detect optical afterglows of gamma-ray bursts (GRBs) detected by the *Fermi* Gamma-ray Burst Monitor (GBM). Its comparable error regions offer a close parallel to the Advanced LIGO problem, but *Fermi*'s unique access to MeV–GeV photons and its near all-sky coverage may allow us to look at optical afterglows in a relatively unexplored part of the GRB parameter space. We present the discovery and broadband follow-up observations (X-ray, UV, optical, millimeter, and radio) of eight GBM–iPTF afterglows. Two of the bursts (GRB 130702A / iPTF13bxi and GRB 140606B / iPTF14bfu) are at low redshift ($z = 0.145$ and $z = 0.384$, respectively), are sub-luminous with respect to “standard” cosmological bursts, and have spectroscopically confirmed broad-line type Ic supernovae. These two bursts are possibly consistent with mildly relativistic shocks breaking out from the progenitor envelopes rather than the standard mechanism of internal shocks within an ultra-relativistic jet. On a technical level, the GBM–iPTF effort is a prototype for locating and observing optical counterparts of GW events in Advanced LIGO with the Zwicky Transient Facility.

Contents

Acknowledgments	iv
Abstract	vi
1 The road to Advanced LIGO	1
1.1 Challenges	4
1.2 Aims of this thesis	8
2 Range and sky resolution of GW detector networks	14
2.1 Basic matched filter search	14
2.2 Measures of detector sensitivity	19
2.3 Fisher information matrix: single detector	20
2.4 Independence of intrinsic and extrinsic errors	22
2.5 Interpretation of phase and time errors	25
2.6 Position resolution	26
2.6.1 Marginalization over nuisance parameters	29
2.6.2 Spatial interpretation	29
2.6.3 Outline of calculation	31
2.6.4 Example calculation for HLV network	31
2.6.5 Improvement in localization due to coherence	32
2.6.6 Revision to LIGO observing scenarios document	37
2.7 Summary	40

3	Early warning GW detection	41
3.1	Prospects for early-warning detection and EM follow-up	45
3.2	Novel real-time algorithm for CBC detection	49
3.2.1	Discrete-time representation of a matched filter	49
3.2.2	The LLOID method	52
3.2.2.1	Selectively reducing the sample rate of the data and templates . . .	53
3.2.2.2	Reducing the number of filters with the SVD	54
3.2.2.3	Early-warning output	55
3.2.3	Comparison of computational costs	57
3.2.3.1	Conventional TD method	57
3.2.3.2	Conventional FD method	57
3.2.3.3	LLOID method	58
3.2.3.4	Speedup of LLOID relative to TD method	59
3.3	Implementation	60
3.3.1	Planning stage	60
3.3.2	Filtering stage	61
3.4	Results	62
3.4.1	Setup	62
3.4.2	Measured SNR loss	66
3.4.3	Other potential sources of SNR loss	67
3.4.4	Lower bounds on computational cost and latency compared to other methods	70
3.4.5	Extrapolation of computational cost to an Advanced LIGO search	70
3.4.6	Measured latency and overhead	72
3.5	Conclusions	72
4	BAYESTAR: Rapid Bayesian sky localization of BNS mergers	76
4.1	Bayesian probability and parameter estimation	78
4.2	The BAYESTAR likelihood	80
4.3	Properties	81
4.4	Prior and problem setup	84
4.5	Marginal posterior	86

4.5.1	Integral over angles and time	86
4.5.2	Integral over distance	87
4.5.2.1	Adaptive Gaussian quadrature method	87
4.5.2.2	Fixed order Gaussian quadrature method	90
4.6	Adaptive mesh refinement	92
4.7	Run time	93
4.8	Case study	94
5	The first two years of EM follow-up with Advanced LIGO and Virgo	96
5.1	Sources and sensitivity	98
5.1.1	Observing scenarios	99
5.1.2	Simulated waveforms	99
5.1.3	Sensitivity to assumptions	101
5.1.4	Source locations	102
5.1.5	Duty cycle	102
5.2	Detection and position reconstruction	102
5.2.1	Template waveforms	103
5.2.2	Detection threshold	103
5.2.3	Sky localization and parameter estimation	104
5.3	Results	105
5.3.1	2015	108
5.3.2	2016	116
5.4	Discussion	120
5.4.1	Caveats	120
5.4.2	Detection scenarios	121
5.4.3	Comparison with other studies	122
5.4.4	Conclusion	124
6	Discovery and redshift of an optical afterglow in 71 square degrees: iPTF13bxi and GRB 130702A	127
6.1	Introduction	128
6.2	Discovery	130

6.3	Broadband photometric follow-up	132
6.4	Optical spectroscopy and host galaxy environment	135
6.5	GRB 130702A in context	137
6.6	Conclusion	138
7	<i>Fermi</i>, iPTF, and the GRB–supernova connection	141
7.1	Introduction	142
7.2	Search methodology	145
7.2.1	Automated TOO Marshal: alerts and tiling	145
7.2.2	Triggering the P48	146
7.2.3	Automated candidate selection	148
7.2.4	Visual scanning in Treasures Portal	149
7.2.5	Archival vetting in the Transient Marshal	151
7.2.6	Photometric, spectroscopic, and broad-band follow-up	152
7.2.7	Long-term monitoring and data reduction	154
7.3	The GBM–iPTF bursts	155
7.3.1	GRB 130702A / iPTF13bxl	159
7.3.2	GRB 131011A / iPTF13dsw	162
7.3.3	GRB 131231A / iPTF13ekl	163
7.3.4	GRB 140508A / iPTF14aue	165
7.3.5	GRB 140606B / iPTF14bfu	166
7.3.6	GRB 140620A / iPTF14cva	168
7.3.7	GRB 140623A / iPTF14cyb	170
7.3.8	GRB 140808A / iPTF14eag	171
7.4	The population in context	173
7.4.1	Selection effects	173
7.4.2	GRBs as standard candles?	175
7.4.3	Shock breakout	177
7.5	Looking forward	181
8	Conclusion	194
8.1	Next steps	195

8.1.1	Spin and NSBH mergers	195
8.1.2	Sub-threshold signals in rapid localization	196
8.1.3	Distance-resolved rapid localizations	196
8.1.4	GRBs beyond the <i>Fermi</i> bandpass	197
8.1.5	<i>Fermi</i> GBM and iPTF as a short GRB factory	197
8.2	Future directions	199
8.2.1	Early warning and dynamically tuned squeezing	199
8.2.2	LIGO as a short GRB early-warning system	201
8.2.3	Optical counterpart search with ZTF	202
8.3	Conclusion	205
A	Computer codes	206
A.1	Astropy	206
A.2	GSTLAL	206
A.3	GStreamer	206
A.4	HEALPix	207
A.5	LALSuite	207
A.6	Code listing: sky resolution from Fisher matrices	207
B	Low frequency cutoff for inspiral searches	212
C	Sky map file format	216
C.1	Mandatory	216
C.2	Optional	217
C.3	Example code in Python	218
C.4	Specimen	218
D	“First Two Years” data release	220
E	Observations of iPTF/GBM afterglows	224
	Acronyms	233

List of Figures

1.1	Effect of a GW on a ring of free-falling test particles	2
1.2	A basic BNS inspiral waveform	3
1.3	Aerial photographs of ground-based GW detectors	4
1.4	Typical light curves of Advanced LIGO optical counterparts	5
1.5	PTF and ZTF cameras	8
1.6	Example GW localization, comparing Initial LIGO method to BAYESTAR	10
1.7	Discovery of GRB 130702A / iPTF13bxi	12
2.1	Antenna patterns	18
2.2	Position resolution as a function of sky position for a three detector network	33
2.3	Log ratios of SNRs in pairs of detectors	34
2.4	Improvement in sky resolution due to coherence	35
2.5	Fisher matrix error ellipses of a random sample of GW events	39
3.1	Expected number of BNS sources detectable before coalescence	45
3.2	Localization area as a function of time before coalescence	47
3.3	Sky resolution versus time before merger for a random sample of events	50
3.4	Schematic of LLOID pipeline	63
3.5	Mass parameters of sub-bank	64
3.6	Mismatch between LLOID filter response and nominal templates	75
4.1	Timeline of CBC EM counterparts and the Advanced LIGO analysis	78
4.2	Cramér–Rao lower bound on time and phase accuracy	84
4.3	Ratio of Fisher matrix elements between autocorrelation likelihood and full GW data	85
4.4	Initial subdivisions for radial integral	88

4.5	BAYESTAR run time	94
5.1	Amplitude spectral densities for early Advanced LIGO configurations	100
5.2	BAYESTAR P - P plots	107
5.3	Cumulative histograms of sky areas	110
5.4	Cumulative histograms of angle offsets	111
5.5	The HL degeneracy	113
5.6	A typical bimodal localization, circa 2015	114
5.7	A typical unimodal localization, circa 2015	115
5.8	Frequency of localizations with one, two, or more modes	115
5.9	Breakdown of 2016 scenario by detector network	116
5.10	Typical HLV localization, circa 2016	118
5.11	Typical HV localization, circa 2016	119
6.1	Discovery of GRB 130702A / iPTF13bxl	129
6.2	Calibration of <i>Fermi</i> systematic errors	131
6.3	Light curve of GRB 130702A / iPTF13bxl	134
6.4	Broadband SED of GRB 130702A / iPTF13bxl	136
6.5	Optical spectra	137
7.1	Screen shot of iPTF TOO Marshal	147
7.2	Screen shot of iPTF Treasures portal	151
7.3	Screen shot of iPTF Transient Marshal	153
7.4	Optical light curves of long GRBs	156
7.5	Containment probability and age	158
7.6	Afterglow spectra	159
7.7	Discovery of GRB 130702A / iPTF13bxl	163
7.8	Discovery of GRB 131011A / iPTF13dsw	164
7.9	Discovery of GRB 131231A / iPTF13ekl	165
7.10	Discovery of GRB 140508A / iPTF14aue	166
7.11	Discovery of GRB 140606B / iPTF14bfu	168
7.12	Discovery of GRB 140620A / iPTF14cva	170

7.13	Discovery of GRB 140623A / iPTF14cyb	171
7.14	Discovery of GRB 140808A / iPTF14eag	172
7.15	Fluence and error radius	174
7.16	Cumulative redshift distribution	175
7.17	Amati relation	176
7.18	Radiative efficiency of prompt phase	181
7.19	Light curve and SED of GRB 130702A / iPTF13bxl.	186
7.20	Light curve of GRB 131011A / iPTF13dsw.	187
7.21	Light curve and SED of GRB 131231A / iPTF13ekl.	188
7.22	Light curve and SED of GRB 140508A / iPTF14aue.	189
7.23	Light curve and SED of GRB 140606B / iPTF14bfu.	190
7.24	Light curve and SED of GRB 140620A / iPTF14cva.	191
7.25	Light curve of GRB 140623A / iPTF14cyb.	192
7.26	Light curve and SED of GRB 140808A / iPTF14eag.	193
8.1	Sensitivity curves for various squeezing configurations	200
8.2	The ZTF camera	202
8.3	Kilonova phase space accessible with various optical instruments	204
B.1	Accumulated SNR as a function of low frequency cutoff	215

List of Tables

1.1	Comparison of the survey speeds of the PTF and ZTF cameras. Reproduced from a presentation by E. Bellm.	9
2.1	Revised sky resolution predictions for Advanced LIGO observing scenarios	38
3.1	Number of sources that are localizable before merger	48
3.2	Notation used to describe filters	57
3.3	Filter design for a sub-bank of 1314 templates	65
3.4	Computational cost of LLOID versus conventional matched filter methods	70
5.1	Detection rate and sky localization accuracy	109
7.1	Number of optical transient candidates surviving each vetting stage	149
7.2	GBM-iPTF detections.	155
7.3	Log of P48 tilings for <i>Fermi</i> GBM bursts.	157
7.4	Log of spectroscopic observations	160
B.1	Fractional accumulated SNR for selected low frequency cutoffs	214
D.1	Simulated BNS signals for 2015 scenario	222
D.2	Detections and sky localization areas for 2015 scenario	222
D.3	Simulated BNS signals for 2016 scenario	223
D.4	Detections and sky localization areas for 2016 scenario	223
E.1	Optical observations of GBM-iPTF afterglows	225
E.2	Radio observations of GBM-iPTF afterglows	231

Chapter 1

The road to Advanced LIGO

Einstein's general theory of relativity holds that the laws of motion play out in a curved space-time, with curvature caused by the presence of matter and energy. This strange statement has some even stranger consequences. One of the earliest solutions of Einstein's equation predicted black holes (BHs), stars made of pure space-time curvature, whose gravitational wells are so deep that nothing, not even light, can escape. We now know that when a massive star exhausts the last of its fuel, it can collapse to form a neutron star (NS)—the densest possible stable arrangement of matter, something akin to a gigantic nucleus with atomic number 10^{57} —or a stellar-mass BH. This gravitational collapse can be messy and loud. It may produce a relativistic shock wave that powers a long GRB, and it may drive a supernova explosion that outshines the late star's host galaxy in visible light for several weeks. Long afterward, the strong gravitational field of a compact object can have other interesting consequences. If the star has a binary companion from which it can accrete matter, it can power a wide range of high-energy transient phenomena. However, all of these processes occur in basically static (but strongly curved) space-time.

In the dynamical regime, Einstein's theory predicts GWs that transmit energy via propagating disturbances in space-time, much as the dynamical solutions of Maxwell's equations carry energy as light. Operationally, the effect of a passing GW is to slightly change the separation between free-falling objects (see Figure 1.1). The brightest source of gravitational waves that we think nature can make is a binary system of two compact objects (NSs and/or BHs). If a compact binary is in a tight enough orbit, gravitational radiation can efficiently carry away energy and angular momentum. This orbital decay was famously observed in the binary pulsar PSR 1913+16 (Hulse & Taylor, 1975; Taylor & Weisberg, 1982), for which Hulse and Taylor received the Nobel




Figure 1.1 Effect of a GW on a ring of free-falling test particles. Left: a ‘+’ polarized GW, causing the test particles to be alternately squeezed or stretched in two orthogonal directions. Right: a ‘x’ polarized GW, causing a stretching and squeezing in a sense that is rotated 45° relative to the ‘+’ polarization.

Prize in Physics in 1993. The energy loss eventually will become a runaway process, as the orbital separation decreases and the system radiates even more gravitational waves. See Figure 1.2 for an illustration of the basic GW “inspiral” waveform due to a compact binary coalescence (CBC). Ultimately, the two compact objects will coalesce: they will become a single perturbed BH, which will ring down as it settles into rotationally symmetric stationary state.

If one or both of the binary companions is a NS, the merger process itself can also be messy and loud. The immense tidal forces can tear apart the NS before it takes the final plunge. The resultant hot, highly magnetized accretion flow may create the conditions necessary for a highly relativistic jet (Rezzolla et al., 2011). This process is thought to power short GRBs (Paczynski, 1986; Eichler et al., 1989; Narayan et al., 1992; Rezzolla et al., 2011).

The Laser Interferometer GW Observatory (LIGO) and Virgo have been constructed with the aim of directly detecting GWs from CBCs of binary neutron stars (BNSs), among other potential sources. This will provide a singularly dramatic confirmation of Einstein’s relativity in the otherwise largely untested strong-field dynamical regime. GW observations could even test alternative theories of gravity (Section 6, Will, 2006; Del Pozzo et al., 2013) or constrain the NS equation of state (Read et al., 2009). A temporal coincidence between a CBC event and a short GRB would also settle the question of the progenitors of at least some of these elusive explosions. No GW events were detected in LIGO–Virgo observing runs at an initial sensitivity (Abadie




Figure 1.2 A basic CBC “inspiral” waveform. The red and blue traces correspond to two orthogonal GW polarizations (see Chapter 2). At the lowest post-Newtonian order, the signal is shrunk or dilated in time by a single mass parameter: the chirp mass, \mathcal{M} , a combination of the component masses defined in Section 3.1.

et al., 2012c). However, LIGO is just now finishing its transformation into Advanced LIGO, with Advanced Virgo soon to follow suit. Both designed to be ultimately ten times more sensitive than their predecessors, they will be able to monitor a thousand times more volume within the local Universe. The first detections are expected over the next few years (Abadie et al., 2010b).

A perhaps even greater prize would be detecting both the GW signal and an optical transient resulting from the same BNS merger event. An optical afterglow (van Eerten & MacFadyen, 2011) would aid in the understanding of the physics of the relativistic jet (for the “classic” model, see Sari et al. 1998), and a bright on-axis afterglow would be the most obvious signpost by which to locate the host galaxy. These signatures, however, are expected to be rare because, like the short GRB itself, we have to be inside the collimated cone of the jet to see them. Perhaps a more promising optical signature (Metzger & Berger, 2012) would be that of a roughly omnidirectional “kilonova” powered by the radioactive decay of the hot r -process ejecta (Li & Paczyński, 1998; Barnes & Kasen, 2013a) or a “kilonova precursor” powered by free neutrons in the fast-moving outer layers of the ejecta (Metzger et al., 2015). A kilonova could inform us about the nature and distribution of the ejecta, and tell us whether the merged compact object collapsed directly to a black hole or went through a brief phase as a hyper-massive neutron star (Metzger & Fernández,



Figure 1.3 From left to right: aerial views of LIGO Hanford Observatory (reproduced from <http://ligo.org>), Virgo (Reproduced from <http://virgo.lal.in2p3.fr>), and LIGO Livingston Observatory (reproduced from <http://ligo.org>).

2014). See Figure 1.4 for typical r -band light curves of these optical signatures. If we could detect GW and electromagnetic (EM) emission from a sufficiently large number of CBCs, then we could simultaneously measure their luminosity distances and redshifts, thereby adding an almost calibration-free “standard siren” to the cosmological distance ladder (Schutz, 1986; Holz & Hughes, 2005; Dalal et al., 2006; Nissanke et al., 2010).¹

1.1 Challenges

Some first steps toward multimessenger observations were taken in the last LIGO–Virgo science run. The first low-latency CBC search was deployed, including an online matched filter analysis for detection (Multi-Band Template Analysis; MBTA), a fast but ad hoc algorithm for sky localization (Abadie et al., 2012a), and a system for sending alerts to optical facilities (Kanner et al., 2008). The whole process from data acquisition to alerts took about half an hour (dominated by a final human-in-the-loop check; see Chapter 3 for a full timing budget). This period also saw the development of the first practical Bayesian parameter estimation codes, which at the time took a few weeks to thoroughly map the parameter space of any detection candidate (Aasi et al., 2013b). A consortium of X-ray, optical, and radio telescopes participated in searching for EM counterparts (Abadie et al., 2012b; Aasi et al., 2014).

Although these were important proofs of concept, the increased sensitivity of the Advanced LIGO detectors will force us toward more sophisticated approaches at each stage of the process

¹Third-generation detectors such as the proposed Einstein Telescope will be able to simultaneously determine distances and redshifts of BNS mergers *through GWs observations alone* by measuring the orbital frequency at which tidal disruption occurs; see Messenger & Read (2012).




Figure 1.4 Light curves of short GRB afterglows, scaled to an Advanced LIGO range of 200 Mpc. Thin gray lines are afterglows of *Swift* short GRBs that have known redshifts. Thick black lines are synthetic on-axis afterglows from [van Eerten & MacFadyen \(2011\)](#) with jet half-opening angles of 0.2 rad and observer angles of 0 rad. Jet energies E_j , in units of ergs, and circumburst densities n , in units of cm^{-3} , are labeled on the plot. The solid, deep red line is the r -band neutron-powered kilonova precursor model from [Metzger et al. \(2015\)](#) with opacity $\kappa_r = 30 \text{ cm}^2 \text{ g}^{-1}$, free neutron mass $m_n = 10^{-4} M_\odot$, and electron fraction $Y_e = 0.05$. The solid, light red lines represent kilonova models from [Barnes & Kasen \(2013a\)](#) with ejected masses of $m_{\text{ej}} = 10^{-3}, 10^{-2}$, or $10^{-1} M_\odot$ and characteristic velocities $\beta = v/c = 0.1, 0.2$, or 0.3 . The [Metzger et al. \(2015\)](#) kilonova precursors are blue: they peak at about 0.1 mag brighter in the g band than in r . The [Barnes & Kasen \(2013a\)](#) kilonova models are red: they are about 1 mag brighter in the i band than in r .

(detection, sky localization, and EM follow-up). The first challenge is longer GW signals. A significant part of Advanced LIGO’s expanded detection volume comes from better sensitivity at low frequencies, moving the seismic noise cutoff from ~ 40 Hz to ~ 10 Hz (see Appendix B for a detailed discussion of the sensitivity as a function of low frequency cutoff). CBC signals are chirps, ramping from low to high frequency as $f \propto t^{-3/8}$. Although a typical BNS merger signal would remain in band for Initial LIGO for about 25 s, it would be detectable by Advanced LIGO for as long as 1000 s (see Equation (3.1) in Chapter 3). A second consequence of longer signals is that the signal can accumulate more power and a larger total phase shift while in band, improving the ability to measure the mass of the binary but dramatically increasing the number of GW templates required to adequately tile the parameter space. A third problem is that we cannot assume that the detector and the instrument noise are in a stationary state for the durations of these long signals; we must adaptively condition or *whiten* the data as the noise level rises or falls, and we must be able to carry on integrating the signal over gaps or glitches. These are all formidable problems for traditional fast Fourier transform (FFT)-based matched filter pipelines, which have inflexible data handling, whose latency grows with the length of the signal, and whose computational requirements increase with both the length and number of template signals. To effectively search for these signals in real time we need a detection pipeline whose latency and computational demands do not scale much with the duration of the GW signal.

The second challenge is that the sky localization must be both fast *and* accurate. The original rapid sky localization and full Bayesian parameter estimation codes entailed an undesirable tradeoff of response time and accuracy: the former took only minutes, but produced sky areas that were 20 times larger than the latter, which could take days (Sidery et al., 2014). This compromise was somewhat acceptable in Initial LIGO because, given the small number of galaxies within the detectable volume, one could significantly reduce the area to be searched by selecting fields that contained nearby galaxies (Kopparapu et al., 2008; White et al., 2011). With the expanded range of Advanced LIGO enclosing many more galaxies, this will still be a valuable strategy, but will be somewhat less effective (Nissanke et al., 2013). Given that the predicted optical signatures of BNS mergers are faint (with kilonovae predicted to be fainter than $R > 22$ mag) and may peak in under a day, it is essential that the rapid localization be as accurate as possible. Ideally, it should be just as accurate as the localization from the full Bayesian parameter estimation.

Third and most importantly, we have to build the instruments, software, collaborations,

and observational discipline to search through areas of hundreds of deg^2 for the faint, rapidly fading optical counterparts. We need deep, wide-field optical survey telescopes to scan the GW localizations and detect new transient or variable sources, robotic follow-up telescopes to track photometric evolution and obtain color information, a network of 5-m class and larger telescopes to secure spectroscopic classifications, as well as X-ray and radio telescopes that can act on targets of opportunity (TOOs). To identify which among tens or hundreds of thousands of optical transients to follow up we need real-time image subtraction, machine learning, and integration with archival survey data, not to mention a team of human observers in the loop and executing deep spectroscopic observations on large-aperture telescopes.

There is a fortunate convergence between the construction of Advanced LIGO and Virgo and the deployment of deep, high-cadence, synoptic, optical transient surveys. Experiments like the Palomar Transient Factory (PTF; [Rau et al. 2009](#); [Law et al. 2009](#)) have focused on discovering rare or rapidly rising optical transients, but should also be well suited to searching for optical counterparts of GW sources. The key instrument in PTF is the Canada–France–Hawaii $12\,288 \times 8\,192$ pixel CCD mosaic (CFH12k) camera ([Rahmer et al., 2008](#)) on the Palomar 48-inch Oschin telescope (P48), capable of reaching limits of $R \approx 20.6$ mag in 60 s over a wide, 7.1 deg^2 field of view (FOV). In its planned successor, the Zwicky Transient Facility (ZTF), this will be replaced by a new 47 deg^2 camera. With a larger FOV and faster readout electronics, ZTF will achieve an order of magnitude faster volumetric survey rate (see Figure 1.5 for an illustration of the PTF and ZTF cameras and Table 1.1 for a comparison of survey speeds). A real-time image subtraction and machine learning pipeline supplies a stream of new optical transient candidates from which a team of human observers selects the most interesting targets for multicolor photometry on the robotic Palomar 60-inch telescope (P60) and spectroscopic classification on the Palomar 200-inch Hale telescope (P200) and other large telescopes. Lessons learned by PTF will inform the planning and operation of future optical transient surveys such as BlackGEM² (which will be dedicated to following up GW sources) and the Large Synoptic Survey Telescope (LSST), as relates to both blind transient searches and targeted searches for optical counterparts of GW candidates.

²<https://www.astro.ru.nl/wiki/research/blackgemarray>




Figure 1.5 The PTF (left) and ZTF (right) cameras. Reproduced from a presentation by E. Bellm.

1.2 Aims of this thesis

The aim of my thesis is to deliver the major, fully and realistically characterized and tested, pieces of the search for optical counterparts of BNS mergers, including detection and parameter estimation as well as the optical transient search itself. Here is a chapter-by-chapter summary of the content of this thesis.

Chapter 2 introduces the basic principles of a matched filter bank GW search. We describe the range of a GW detector in terms of its directional sensitivity or antenna pattern, its noise power spectral density (PSD), and the signal-to-noise ratio (SNR). We then apply the Fisher information matrix formalism to compute the approximate sky resolution of a network of GW detectors. There is a great deal of prior literature on this topic that considers GW sky localization in terms of timing triangulation (see, for instance, [Fairhurst 2009](#)). Our calculation captures the additional contributions of the phases and amplitudes on arrival at the detectors, which we show to be significant, especially near the plane of the detectors where timing triangulation is formally

Table 1.1. Comparison of the survey speeds of the PTF and ZTF cameras. Reproduced from a presentation by E. Bellm.

	PTF	ZTF
Active area	7.26 deg ²	47 deg ²
Overhead time	46 s	< 15 s
Optimal exposure time	60 s	30 s
Relative areal survey rate	1x	15.0x
Relative volumetric survey rate	1x	12.3x

degenerate. Our derivation is extremely compact, and evaluating it is only marginally more complicated than the timing triangulation approach. We discuss the sky localization accuracy as a function of direction in the sky, and build some intuition that we will rely upon in future chapters. This chapter is in preparation as a separate paper and as a proposed update to a living document describing the GW detector commissioning and observing schedule ([Aasi et al., 2013c](#)).

Chapter 3 describes a novel matched filtering algorithm that is capable of detecting a GW signal within seconds after the merger, or even seconds *before*. This algorithm, called LLOID, uses orthogonal decomposition and multirate signal processing to bring the computational demands of an online BNS search within the scope of current resources. My contributions to LLOID include: working on the pipeline to drive the latency to ~ 10 s, improving data handling to be able to skip over glitches in the data efficiently without unduly sacrificing SNR, studying the signal processing and computational aspects of the algorithm, improving the time and phase accuracy of triggers, and preparing the first complete description of it for the literature ([Cannon et al., 2012](#)). LLOID has been extensively tested both offline and in real time with simulated and commissioning data in a series of “engineering runs,” and will serve as the flagship low-latency BNS detection pipeline in Advanced LIGO. This chapter is in preparation as a standalone technical paper.

Chapter 4 develops a new rapid sky localization algorithm, BAYESTAR, that takes just tens of seconds, but achieves about the same accuracy (see Figure 1.6) as the full parameter estimation. It owes its speed to three innovations. First, like the ad hoc Initial LIGO code, it takes as input the matched filter parameter estimates from the detection pipeline rather than the full GW time

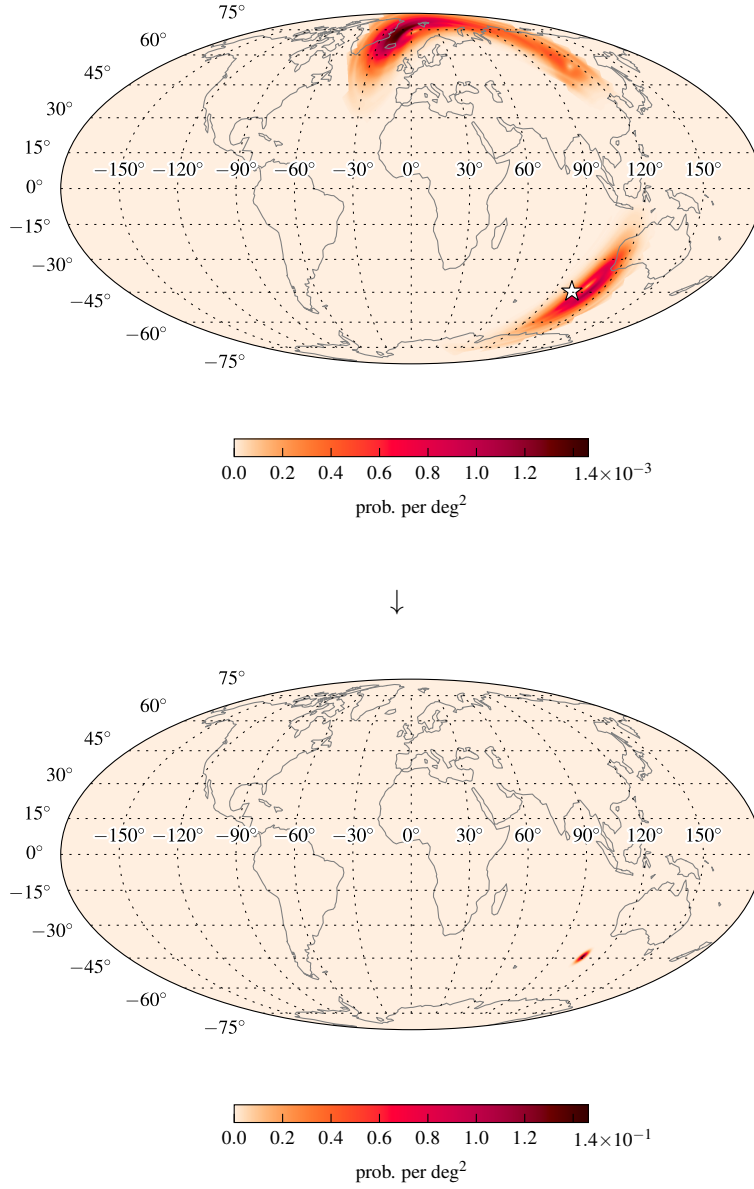


Figure 1.6 Probability sky maps for a simulated three-detector, $(1.26, 1.49) M_{\odot}$ merger event from the third engineering run. Above is the sky map from the Initial LIGO rapid localization code. Below is the sky map from BAYESTAR. (This is event G71031.)

series. Second, using a result from Chapter 2, it concerns itself with sky location only and not the masses of the signal, exploiting the fact that the errors in the intrinsic and extrinsic parameters of a BNS signal are approximately uncorrelated. Third, though fully Bayesian, unlike the full parameter estimation it does not use Markov chain Monte Carlo (MCMC) sampling; instead it uses an adaptive sampling grid and low-order Gaussian quadrature. The result is both inherently fast and also highly parallelizable. Like LLOID, it has been tested with both extensive offline simulations and in online engineering runs.

Having assembled the Advanced LIGO real-time BNS pipeline in Chapters 3 and 4, in **Chapter 5** we provide a detailed description of what the first Advanced LIGO detections and sky localizations may look like. Because our new rapid sky localization algorithm is orders of magnitude faster than the full parameter estimation, for the first time we can perform end-to-end analyses of thousands of events, thereby providing a statistically meaningful and comprehensive description of the areas and morphologies that will arise in the early Advanced LIGO configurations. The first 2015 observing run is expected to involve only the two LIGO detectors in Hanford, Washington, and Livingston, Louisiana, and not the Virgo detector in Cascina, Italy. [Aasi et al. \(2013c\)](#) assumed, based on timing triangulation considerations, that two detector networks would always produce localizations that consist of degenerate annuli spanning many thousands of deg^2 . We find that the interplay between the phase and amplitude on arrival (i.e., the GW polarization) and prior distribution powerfully break this degeneracy (see also [Raymond et al. 2009](#); [Kasliwal & Nissanke 2014](#)), limiting *almost all* areas to below 1000 deg^2 , with a median of about 600 deg^2 . We elucidate one curious degeneracy that survives, that causes most source localizations to equally favor the true position of the source as well as a position at the polar opposite. We then model the 2016 observing run, which has the LIGO detectors operating with somewhat deeper sensitivity and has Advanced Virgo online. Even with Virgo’s sensitivity delayed with the staggered commissioning, adding the third detector shrinks areas to a median of $\sim 200 \text{ deg}^2$. As the detectors continue approaching final design sensitivity and as more detectors come online, areas will continue to shrink to $\sim 10 \text{ deg}^2$ and below. This chapter is published as [Singer et al. \(2014\)](#). The supplementary data described in Appendix D contains a browsable catalog of simulated GW sky maps, in the format that will be used for sending GW alerts (which is described in Appendix C).




Figure 1.7 P48 imaging of GRB 130702A and discovery of iPTF13bxi. The left panel illustrates the γ -ray localizations (red circle: 1σ GBM; green circle: Large Area Telescope (LAT); blue lines: 3σ InterPlanetary Network (IPN)) and the 10 P48 reference fields that were imaged (light gray rectangles). For each P48 pointing, the locations of the 11 CCD chips are indicated with smaller rectangles (one CCD in the camera is not currently operable). The small black diamond is the location of iPTF13bxi. The right panels show P48 images of the location of iPTF13bxi, both prior to (top) and immediately following (bottom) discovery. Reproduced from [Singer et al. \(2013\)](#).

Chapter 6 confronts the search for optical transients in large error regions with the Intermediate Palomar Transient Factory; (iPTF; [Kulkarni 2013](#)) and its planned successor, the Zwicky Transient Facility (ZTF; [Kulkarni 2012](#); [Bellm 2014](#); [Smith et al. 2014b](#)). The surest way to convince ourselves that the search for optical counterparts of GW transients will be effective is to try it out and discover something. As a model problem, for the past year we have searched for optical counterparts of GRBs detected by the Gamma-ray Burst Monitor (GBM; [Meegan et al. 2009a](#)) instrument onboard the *Fermi* satellite. Like LIGO, *Fermi* GBM produces coarse localizations that are uncertain by $\sim 100 \text{ deg}^2$, and though afterglows of long GRBs are much brighter than anticipated LIGO optical counterparts, the important timescales for follow-up observations are similar. *Fermi* GBM bursts are also interesting in their own right. *Fermi* GBM and the *Swift* Burst Alert Telescope (BAT; [Barthelmy et al. 2005](#)) have highly complementary strengths: fields of view of 70% and 10% of the sky respectively, and energy bandpasses of $\sim \text{few keV}$ to 300 GeV

(when including the *Fermi* Large Area Telescope or LAT; [Atwood et al. 2009](#)) and 15–150 keV respectively. However, with the GBM’s coarse localization, very few *Fermi* bursts have been studied outside the gamma-ray band (the exception being bursts that are coincidentally also detected by *Fermi* LAT or *Swift* BAT). In this chapter, we relate the discovery, redshift, and broadband observations of GRB 130702A and its optical afterglow, iPTF13bxi [Singer et al. \(2013\)](#). This is the first discovery of an optical afterglow based solely on a *Fermi* GBM localization (see discovery image in Figure 1.7). This is a notable event in and of itself for several other reasons. First, its redshift places it among the nearest GRBs ever recorded. Second, its prompt energy release in gamma rays is intermediate between bright, cosmologically distant, “standard” bursts, and nearby low-luminosity GRBs (llGRBs) which comprise many of the well-studied GRB–supernova (SN). Finally, because of its low redshift we were able to spectroscopically detect its associated broad-line type Ic SN, establishing it as a test for the GRB–SN connection.

Chapter 7 reports on the total of eight GBM–iPTF afterglows that we have discovered in one year of this experiment. In this chapter, we present our broadband follow-up including spectroscopy as well as X-ray, UV, optical, sub-millimeter, millimeter, and radio observations. We study possible selection effects in the context of the total *Fermi* and *Swift* GRB samples. We identify one new outlier on the Amati relation, challenging its application to standardize GRB luminosities. We find that two bursts are consistent with a mildly mildly relativistic shock breaking out from the progenitor star, rather than the ultra-relativistic internal shock mechanism that powers standard cosmological bursts. Finally, in the context of the Zwicky Transient Facility (ZTF), we discuss how we will continue to expand this effort to find optical counterparts of binary neutron star mergers that should soon be detected Advanced LIGO and Virgo.

Chapter 2

Range and sky resolution of GW detector networks

This chapter is reproduced from a work in preparation, of which I will be the sole author. Section 2.2 is reproduced from [Singer et al. \(2014\)](#), copyright © 2014 The American Astronomical Society.

In this chapter, we will use a basic description of the signal and noise received by a GW detector network to derive a matched filter bank, the prevailing technique used to search for well-modeled CBC signals in LIGO data. This model will allow us to calculate the range and angular resolution of a network of detectors.

2.1 Basic matched filter search

With interferometric detectors like LIGO and Virgo, the astrophysical signal is embedded in a time series measurement, the strain or the differential change in the lengths of the detectors' two arms. Many noise sources enter the detector in different subsystems, get filtered by the detector's response, and add to the measured strain. There are "fundamental" noise sources, such as quantum fluctuations in the laser field that result in shot noise at low frequency and radiation pressure noise at high frequency. Other noise sources are "technical," meaning that they arise from the implementation of the detector as a realizable non-ideal system; examples include glitches due to scattered light, laser frequency fluctuations, cross-coupling between length degrees of freedom,

coupling between angular and length degrees of freedom, and time-varying alignment drifts. Other noise sources are “environmental,” such seismic or anthropogenic ground motion noise.

For the purpose of GW data analysis, the most important division is between quasi-stationary Gaussian-like noise and transient noise sources (“glitches”). Extracting astrophysical signals from the data requires frequency domain (FD) techniques (whitening, matched filtering) to suppress the former and time domain (TD) approaches (coincidence, candidate ranking, time slides) to deal with the latter.

CBC searches are greatly aided by the fact that their GW signals can (at least in principle) be predicted with exquisite precision throughout LIGO’s sensitive band. Therefore, a standard approach to CBC detection is matched filtering; a representative set of model waveforms is assembled into a template bank with which the data is convolved.

In the TD, the strain observed by a single GW interferometer is

$$y_i(t) = x_i(t; \boldsymbol{\theta}) + n_i(t). \quad (2.1)$$

In the FD,

$$Y_i(\omega) = \int_{-\infty}^{\infty} y(t) e^{-i\omega t} dt = X_i(\omega; \boldsymbol{\theta}) + N_i(\omega), \quad (2.2)$$

where $X_i(\omega; \boldsymbol{\theta})$ is the GW signal given a parameter vector $\boldsymbol{\theta}$ that describes the GW source, and $N_i(\omega)$ is that detector’s Gaussian noise with one-sided PSD $S_i(\omega) = E[|n_i(\omega)|^2] + E[|n_i(-\omega)|^2] = 2E[|n_i(\omega)|^2]$. We shall denote the combined observation from a network of detectors as $\mathbf{Y}(\omega) \equiv \{Y_i(\omega)\}_i$.

Under the assumptions that the detector noise is Gaussian and that the noise from different detectors are uncorrelated, the likelihood of the observation, \mathbf{y} , conditioned on the value of $\boldsymbol{\theta}$, is a product of Gaussian distributions:

$$\mathcal{L}(\mathbf{Y}; \boldsymbol{\theta}) = \prod_i p(Y_i | \boldsymbol{\theta}) \propto \exp \left[-\frac{1}{2} \sum_i \int_0^{\infty} \frac{|Y_i(\omega) - X_i(\omega; \boldsymbol{\theta})|^2}{S_i(\omega)} d\omega \right]. \quad (2.3)$$

A CBC source is specified by a vector of extrinsic parameters describing its position and

orientation, and intrinsic parameters describing the physical properties of the binary components:

$$\theta = \left[\begin{array}{ll} \alpha & \text{right ascension} \\ \delta & \text{declination} \\ r & \text{distance} \\ t_{\oplus} & \text{arrival time at geocenter} \\ \iota & \text{inclination angle} \\ \psi & \text{polarization angle} \\ \phi_c & \text{coalescence phase} \\ \hline m_1 & \text{first component's mass} \\ m_2 & \text{second component's mass} \\ \mathbf{S}_1 & \text{first component's spin} \\ \mathbf{S}_2 & \text{second component's spin} \end{array} \right] \left. \begin{array}{l} \\ \\ \\ \\ \\ \\ \\ \text{extrinsic} \\ \text{parameters,} \\ \theta_{\text{ex}} \\ \\ \text{intrinsic} \\ \text{parameters,} \\ \theta_{\text{in.}} \end{array} \right\} \quad (2.4)$$

This list of parameters involves some simplifying assumptions. Eccentricity is omitted: although it does play a major role in the evolution and waveforms of neutron star–black hole (NSBH) and binary black hole (BBH) sources formed by dynamical capture (East et al., 2013), BNS systems formed by binary stellar evolution should almost always circularize due to tidal interaction (Belczynski et al., 2002) and later GW emission (Peters, 1964) long before the inspiral enters LIGO’s frequency range of $\sim 10\text{--}1000$ kHz. Tidal deformabilities of the NSs are omitted because the signal imprinted by the companions’ material properties is so small that it will only be detectable by an Einstein Telescope-class GW observatory (Hinderer et al., 2010). Furthermore, in GW detection efforts, especially those focused on BNS systems, the component spins \mathbf{S}_1 and \mathbf{S}_2 are often assumed to be nonprecessing and aligned with the system’s total angular momentum and condensed to a single scalar parameter χ , or even neglected entirely: $\mathbf{S}_1 = \mathbf{S}_2 = 0$.

Assuming circular orbits and no spin precession, we can write the GW signal in each detector as a linear combination of two basis waveforms, H_0 and $H_{\pi/2}$. For nonprecessing systems, H_0 and $H_{\pi/2}$ are approximately “in quadrature” in the same sense as the sine and cosine functions, being nearly orthogonal and out of phase by $\pi/2$ at all frequencies. If H_0 and $H_{\pi/2}$ are Fourier transforms of real functions, then $H_0(\omega) = H_0^*(-\omega)$ and $H_{\pi/2}(\omega) = H_{\pi/2}^*(-\omega)$, and we can

write (assuming an arbitrary phase convention)

$$H_{\pi/2}(\omega) = H_0(\omega) \cdot \begin{cases} -i & \text{if } \omega \geq 0 \\ i & \text{if } \omega < 0 \end{cases}. \quad (2.5)$$

For brevity, we define $H \equiv H_0$ and write all subsequent equations in terms of the H basis vector alone. Then, we can write the signal model in a way that isolates all dependence on the extrinsic parameters, θ_{ex} , into the coefficients and all dependence on the intrinsic parameters, θ_{in} , into the basis waveform, by taking the Fourier transform of Equation (2.8) of [Harry & Fairhurst \(2011b\)](#):

$$X_i(\omega; \theta) = e^{-i\omega(t_{\oplus} - \mathbf{d}_i \cdot \mathbf{n})} \frac{r_{1,i}}{r} e^{2i\phi_c} \left[\frac{1}{2} (1 + \cos^2 \iota) \Re\{\zeta\} - i(\cos \iota) \Im\{\zeta\} \right] H(\omega; \theta_{\text{in}}) \quad (2.6)$$

for $\omega \geq 0$, where

$$\zeta = e^{-2i\psi} (F_{+,i}(\alpha, \delta, t_{\oplus}) + iF_{\times,i}(\alpha, \delta, t_{\oplus})). \quad (2.7)$$

The quantities $F_{+,i}$ and $F_{\times,i}$ are the dimensionless detector antenna factors, defined such that $0 \leq F_{+,i}^2 + F_{\times,i}^2 \leq 1$. They depend on the orientation of detector i as well as the sky location (as depicted in Figure 2.1) and sidereal time of the event and are presented in [Anderson et al. \(2001\)](#). In a coordinate system with the x and y axes aligned with the arms of a detector, its antenna pattern is given in spherical polar coordinates as

$$F_+ = -\frac{1}{2}(1 + \cos^2 \theta) \cos 2\phi, \quad (2.8)$$

$$F_{\times} = -\cos \theta \sin 2\phi. \quad (2.9)$$

The unit vector \mathbf{d}_i represents the position of detector i in units of light travel time.¹ The vector \mathbf{n} is the direction of the source. The negative sign in the dot product $-\mathbf{d}_i \cdot \mathbf{n}$ is present because the direction of travel of the GW signal is opposite to that of its sky location. The quantity $r_{1,i}$ is a fiducial distance at which detector i would register SNR=1 for an optimally oriented binary

¹When considering transient GW sources such as those that we are concerned with in this thesis, the origin of the coordinate system is usually taken to be the geocenter. For long-duration signals such as those from statically deformed neutron stars, the solar system barycenter is a more natural choice.




Figure 2.1 The directional dependence of the $+$, \times , and root mean square (RMS) antenna patterns of a LIGO-style GW detector. The detector is at the center of the light box, with its two arms parallel to the horizontal edges.

(face-on, and in a direction perpendicular to the interferometer's arms):

$$r_{1,i} = 1/\sigma_i, \quad \sigma_i^2 = \int_0^\infty \frac{|H(\omega; \theta_{\text{in}})|^2}{S_i(\omega)} d\omega. \quad (2.10)$$

More succinctly, we can write the signal received by detector i in terms of observable extrinsic parameters $\theta_i = (\rho_i, \gamma_i, \tau_i)$, the amplitude ρ_i , phase γ_i , and time delay τ_i on arrival at detector i :

$$X_i(\omega; \theta_i, \theta_{\text{in}}) = X_i(\omega; \rho_i, \gamma_i, \tau_i, \theta_{\text{in}}) = \frac{\rho_i}{\sigma_i} e^{i(\gamma_i - \omega\tau_i)} H(\omega; \theta_{\text{in}}). \quad (2.11)$$

The prevailing technique for detection of GWs from CBCs is to realize a maximum likelihood (ML) estimator (MLE) from the likelihood in Equation (2.3) and the signal model in Equation (2.11). Concretely, this results in a bank of matched filters, or the cross-correlation between a template waveform and the incoming data stream,

$$z_i(\tau_i; \theta_{\text{in}}) = \frac{1}{\sigma_i(\theta_{\text{in}})} \int_0^\infty \frac{H^*(\omega; \theta_{\text{in}}) Y_i(\omega) e^{i\omega\tau_i}}{S_i(\omega)} d\omega. \quad (2.12)$$

The ML point estimates of the signal parameters, $\text{MLE}(\mathbf{y}) = \{\{\hat{\theta}_i\}_i, \hat{\theta}_{\text{in}}\} = \{\{\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i\}_i, \hat{\theta}_{\text{in}}\}$, are

given by

$$\hat{\boldsymbol{\theta}}_{\text{in}}, \{\hat{\tau}_i\}_i = \underset{\boldsymbol{\theta}_{\text{in}}, \{\hat{\tau}_i\}_i}{\operatorname{argmax}} \sum_i |z_i(\tau_i; \boldsymbol{\theta}_{\text{in}})|^2, \quad (2.13)$$

$$\hat{\rho}_i = |z_i(\hat{\tau}_i; \hat{\boldsymbol{\theta}}_{\text{in}})|, \quad (2.14)$$

$$\hat{\gamma}_i = \arg z_i(\hat{\tau}_i; \hat{\boldsymbol{\theta}}_{\text{in}}). \quad (2.15)$$

A detection candidate consists of $\{\{\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i\}_i, \hat{\boldsymbol{\theta}}_{\text{in}}\}$. There are various ways to characterize the significance of a detection candidate. In Gaussian noise, the maximum likelihood for the network is obtained by maximizing the network SNR, ρ_{net} ,

$$\hat{\rho}_{\text{net}} = \max_{\boldsymbol{\theta}} \sum_i |z_i(\boldsymbol{\theta})|^2 = \sqrt{\sum_i \hat{\rho}_i^2}; \quad (2.16)$$

this, therefore, is the simplest useful candidate ranking statistic.

2.2 Measures of detector sensitivity

The sensitivity of a single GW detector is customarily described by the horizon distance, or the maximum distance at which a particular source would create a signal with a maximum fiducial single-detector SNR, ρ . It is given by

$$d_{\text{H}} \approx \frac{G^{5/6} M^{1/3} \mu^{1/2}}{c^{3/2} \pi^{2/3} \rho} \sqrt{\frac{5}{6} \int_{f_1}^{f_2} \frac{f^{-7/3}}{S(f)} df}, \quad (2.17)$$

where G is Newton's gravitational constant, c is the speed of light, M the sum of the component masses, μ the reduced mass, $f^{-7/3}$ the approximate power spectral density (PSD) of the inspiral signal, and $S(f)$ the PSD of the detector's noise. The lower integration limit f_1 is the low-frequency extent of the detector's sensitive band. For the Advanced LIGO and Virgo detectors, ultimately limited at low frequency by ground motion ([Adhikari, 2014](#)), we take $f_1 = 10$ Hz. Using a typical value of the detector sensitivity $S(100 \text{ Hz}) = 10^{-46} \text{ Hz}^{-1}$, we can write Equation (2.17) as a scaling

law:

$$d_H \approx 72.5 \text{ Mpc} \left(\frac{M}{M_\odot} \right)^{1/3} \left(\frac{\mu}{M_\odot} \right)^{1/2} \left(\frac{1}{\rho} \right) \cdot \left[\int_{f_1/\text{Hz}}^{f_2/\text{Hz}} \left(\frac{f}{100 \text{ Hz}} \right)^{-7/3} \left(\frac{10^{-46} \text{ Hz}^{-1}}{S(f)} \right) d \left(\frac{f}{\text{Hz}} \right) \right]^{1/2}. \quad (2.18)$$

For BNS masses, the inspiral ends with a merger and black hole ring down well outside LIGO's most sensitive band. A reasonable approximation is to simply truncate the SNR integration at the last stable orbit of a Schwarzschild black hole with the same total mass (Maggiore, 2008),

$$f_2 = \frac{c^3}{6\sqrt{6}\pi G M} \approx (4400 \text{ Hz}) \frac{M_\odot}{M}. \quad (2.19)$$

Usually, $\rho = 8$ is assumed because $\rho = 8$ signals in two detectors (for a root-sum-squared network SNR of $\rho_{\text{net}} = 8\sqrt{2} = 11.3$) is nearly adequate for a confident detection (see discussion of detection thresholds in Section 5.2). Another measure of sensitivity is the BNS range d_R , the volume-, direction-, and orientation-averaged distance of a source with $\rho \geq 8$, drawn from a homogeneous population. Due to the directional sensitivity or antenna pattern of interferometric detectors, the range is a factor of 2.26 smaller than the horizon distance for the same SNR threshold. See also Allen et al. (2012); Abadie et al. (2012d).²

2.3 Fisher information matrix: single detector

We can predict the uncertainty in the ML estimates without working out its full distribution. The Cramér–Rao lower bound (CRLB) gives its covariance in the asymptotic limit of high SNR. The CRLB has been widely applied in GW data analysis to estimate parameter estimation uncertainty (for example, Balasubramanian et al. 1996; Fairhurst 2009; Ajith & Bose 2009; Wen & Chen 2010; Aasi et al. 2013c; Fairhurst 2014)³. We will momentarily consider the likelihood for a single

²Even at its final design sensitivity, Advanced LIGO's range for BNS mergers is only 200 Mpc or $z = 0.045$ (assuming the WMAP nine-year Λ CDM cosmology; Hinshaw et al., 2013). The horizon distance, 452 Mpc or $z = 0.097$, is only modestly cosmic. Because of the small distances considered in this study, we do not distinguish between different distance measures, nor do our gravitational waveforms contain any factors of $(1+z)$.

³The Fisher matrix is also used in construction of CBC matched filter banks. The common procedure is to place templates uniformly according to the determinant of the signal space metric, which is the Fisher matrix. This is equivalent to uniformly sampling the Jeffreys prior. In practice, this is done either by constructing a hexagonal lattice (Cokelaer,

detector:

$$\mathcal{L}(Y_i; \rho_i, \gamma_i, \tau_i, \theta_{\text{in}}) \propto \exp \left[-\frac{1}{2} \int_0^\infty \frac{|Y_i(\omega) - X_i(\omega; \rho_i, \gamma_i, \tau_i, \theta_{\text{in}})|^2}{S_i(\omega)} d\omega \right], \quad (2.20)$$

with $X_i(\omega; \rho_i, \gamma_i, \tau_i, \theta_{\text{in}})$ given by Equation (2.11).

The Fisher information matrix for a measurement y described by the unknown parameter vector θ is the conditional expectation value

$$\mathcal{I}_{jk} = \mathbb{E} \left[\left(\frac{\partial \log \mathcal{L}(Y_i; \theta)}{\partial \theta_j} \right) \left(\frac{\partial \log \mathcal{L}(Y_i; \theta)}{\partial \theta_k} \right) \middle| \theta \right]. \quad (2.21)$$

Note that if $\log \mathcal{L}$ is twice differentiable in terms of θ , then the Fisher matrix can also be written in terms of second derivatives as

$$\mathcal{I}_{jk} = \mathbb{E} \left[-\frac{\partial^2 \log \mathcal{L}(Y_i; \theta)}{\partial \theta_j \partial \theta_k} \middle| \theta \right]. \quad (2.22)$$

In this form, we can recognize the Fisher matrix as the expectation value, conditioned on the true parameter values, of the Hessian matrix of the log likelihood. It describes how strongly the likelihood depends, on average, on the parameters. If $\hat{\theta}$ is an unbiased estimator of θ , $\tilde{\theta} = \hat{\theta} - \theta$ is the measurement error, and $\Sigma = \mathbb{E}[\tilde{\theta}\tilde{\theta}^\top]$ is the covariance of the measurement error, then the CRLB says that $\Sigma \geq \mathcal{I}^{-1}$, in the sense that $(\Sigma - \mathcal{I}^{-1})$ is positive semi-definite.

When (as in our assumptions) the likelihood is Gaussian, Equation (2.21) simplifies to

$$\mathcal{I}_{jk} = \int_0^\infty \Re \left[\left(\frac{\partial X_i}{\partial \theta_j} \right)^* \left(\frac{\partial X_i}{\partial \theta_k} \right) \right] \frac{1}{S_i(\omega)} d\omega. \quad (2.23)$$

This form is useful because it involves manipulating the signal $X_i(\omega)$ rather than the entire observation $Y(\omega)$. In terms of the k th SNR-weighted moment of angular frequency,

$$\overline{\omega^k}_i = \left[\int_0^\infty \frac{|h(\omega)|^2}{S_i(\omega)} \omega^k d\omega \right] \left[\int_0^\infty \frac{|h(\omega)|^2}{S_i(\omega)} d\omega \right]^{-1}, \quad (2.24)$$

2007) or sampling stochastically (Harry et al., 2009; van den Broeck et al., 2009; Ajith et al., 2014; Manca & Vallisneri, 2010; Privitera et al., 2014)

the Fisher matrix for the signal in the i th detector is

$$\mathcal{I}_i = \begin{pmatrix} \mathcal{I}_{\theta_i, \theta_i} & \mathcal{I}_{\theta_i, \theta_{\text{in}}} \\ \mathcal{I}_{\theta_i, \theta_{\text{in}}}^\top & \rho_i^2 \mathcal{I}_{\theta_{\text{in}}, \theta_{\text{in}}} \end{pmatrix} \quad (2.25)$$

where

$$\mathcal{I}_{\theta_i, \theta_i} = \begin{matrix} & \rho_i & \gamma_i & \tau_i \\ \begin{matrix} \rho_i \\ \gamma_i \\ \tau_i \end{matrix} & \begin{pmatrix} 1 & 0 & 0 \\ 0 & \rho_i^2 & -\rho_i^2 \bar{\omega}_i \\ 0 & -\rho_i^2 \bar{\omega}_i & \rho_i^2 \bar{\omega}_i^2 \end{pmatrix} \end{matrix}. \quad (2.26)$$

(This is equivalent to an expression given in [Grover et al. 2014](#).) The information matrix elements that relate to the intrinsic parameters can also be expressed as linear combinations of the angular frequency moments. However, as we will see in the next section, we need not compute these matrix elements if we are only interested in sky localization accuracy.

2.4 Independence of intrinsic and extrinsic errors

If all of the detectors have the same noise PSDs up to multiplicative factors, $c_1 S_1(\omega) = c_2 S_2(\omega) = \dots = c_n S_n(\omega) \equiv S(\omega)$, then we can show that the errors in the intrinsic parameters (masses) are not correlated with sky position errors. This is because we can change variables from amplitudes, phases, and times to amplitude ratios, phase differences, and time differences. With N detectors, we can form a single average amplitude, time, and phase, plus $N - 1$ linearly independent differences. The averages are correlated with the intrinsic parameters, but neither are correlated with the differences. Since only the differences inform sky location, this gives us license to neglect uncertainty in masses when we are computing sky resolution.

This is easiest to see if we make the temporary change of variables $\rho \rightarrow \varsigma = \log \rho$. This allows us to factor out the SNR dependence from the single-detector Fisher matrix. The extrinsic part becomes

$$\mathcal{I}_{\theta_i, \theta_i} = \begin{matrix} & \varsigma_i & \gamma_i & \tau_i \\ \begin{matrix} \varsigma_i \\ \gamma_i \\ \tau_i \end{matrix} & \begin{pmatrix} \rho_i^2 & 0 & 0 \\ 0 & \rho_i^2 & -\rho_i^2 \bar{\omega}_i \\ 0 & -\rho_i^2 \bar{\omega}_i & \rho_i^2 \bar{\omega}_i^2 \end{pmatrix} \end{matrix} = \rho_i^2 \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -\bar{\omega}_i \\ 0 & -\bar{\omega}_i & \bar{\omega}_i^2 \end{pmatrix}. \quad (2.27)$$

Due to our assumption that the detectors' PSDs are proportional to each other, the noise moments are the same for all detectors, $\overline{\omega^k_i} \equiv \overline{\omega^k}$. Then we can write the single-detector Fisher matrix as

$$\mathcal{I}_i = \rho_i^2 \begin{pmatrix} A & B \\ B^\top & C \end{pmatrix}, \quad (2.28)$$

with the top-left block A comprising the extrinsic parameters and the bottom-right block C the intrinsic parameters.

Information is additive, so the Fisher matrix for the whole detector network is

$$\mathcal{I}_{\text{net}} = \begin{pmatrix} \rho_1^2 A & 0 & \cdots & 0 & \rho_1^2 B \\ 0 & \rho_2^2 A & & \vdots & \rho_1^2 B \\ \vdots & & \ddots & 0 & \vdots \\ 0 & 0 & \cdots & \rho_N^2 A & \rho_N^2 B \\ \rho_1^2 B^\top & \rho_2^2 B^\top & \cdots & \rho_N^2 B^\top & \rho_{\text{net}}^2 C \end{pmatrix}. \quad (2.29)$$

Now we introduce the change of variables that sacrifices the N th detector's extrinsic parameters for the network averages,

$$\begin{aligned} \zeta_N &\rightarrow \bar{\zeta} = (\sum_i \rho_i^2 \zeta_i) / \rho_{\text{net}}^2, \\ \gamma_N &\rightarrow \bar{\gamma} = (\sum_i \rho_i^2 \gamma_i) / \rho_{\text{net}}^2, \\ \tau_N &\rightarrow \bar{\tau} = (\sum_i \rho_i^2 \tau_i) / \rho_{\text{net}}^2, \end{aligned} \quad (2.30)$$

and replaces the first $N - 1$ detectors' extrinsic parameters with differences,

$$\left. \begin{aligned} \zeta_i &\rightarrow \delta\zeta_i = \zeta_i - \bar{\zeta} \\ \gamma_i &\rightarrow \delta\gamma_i = \gamma_i - \bar{\gamma} \\ \tau_i &\rightarrow \delta\tau_i = \tau_i - \bar{\tau} \end{aligned} \right\} \text{ for } i = 1, \dots, N - 1. \quad (2.31)$$

The Jacobian matrix that describes this change of variables is

$$J = \begin{pmatrix} 1 & 0 & \cdots & 0 & 1 & 0 \\ 0 & 1 & & 0 & 1 & 0 \\ \vdots & & \ddots & & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 1 & 0 \\ \frac{-\rho_1^2}{\rho_N^2} & \frac{-\rho_2^2}{\rho_N^2} & \cdots & \frac{-\rho_{N-1}^2}{\rho_N^2} & 1 & 0 \\ 0 & 0 & \cdots & 0 & 0 & 1 \end{pmatrix}. \quad (2.32)$$

The transformed network Fisher matrix is block diagonal,

$$\mathcal{I}_{\text{net}} \rightarrow J^\top \mathcal{I}_{\text{net}} J = \begin{pmatrix} \rho_1^2(1 + \frac{1}{\rho_1^4})A & \frac{\rho_1^2\rho_2^2}{\rho_N^2}A & \cdots & \frac{\rho_1^2\rho_{N-1}^2}{\rho_N^2}A & 0 & 0 \\ \frac{\rho_1^2\rho_2^2}{\rho_N^2}A & \rho_2^2(1 + \frac{1}{\rho_1^4})A & & \frac{\rho_2^2\rho_{N-1}^2}{\rho_N^2}A & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots & \vdots \\ \frac{\rho_1^2\rho_{N-1}^2}{\rho_N^2}A & \frac{\rho_2^2\rho_{N-1}^2}{\rho_N^2}A & \cdots & \rho_{N-1}^2(1 + \frac{1}{\rho_1^4})A & 0 & 0 \\ 0 & 0 & \cdots & 0 & \rho_{\text{net}}^2A & \rho_{\text{net}}^2B \\ 0 & 0 & \cdots & 0 & \rho_{\text{net}}^2B^\top & \rho_{\text{net}}^2C \end{pmatrix}. \quad (2.33)$$

The top-left block contains $N - 1$ relative amplitudes, phases, and times on arrival, all potentially correlated with each other. The bottom-right block contains the average amplitudes, phases, and times, as well as the masses. The averages and the masses are correlated with each other, but are not correlated with the differences. Because only the differences are informative for sky localization, we drop the intrinsic parameters from the rest of the Fisher matrix calculations in this chapter.

2.5 Interpretation of phase and time errors

We take a brief digression to discuss the physical interpretation of the time and amplitude errors.

For our likelihood, the CRLB implies that

$$\text{cov} \begin{pmatrix} \tilde{\rho}_i \\ \tilde{\gamma}_i \\ \tilde{\tau}_i \end{pmatrix} \geq \mathcal{I}^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \rho_i^2 \bar{\omega}_i^2 / \omega_{\text{rms},i}^2 & \rho_i^2 \bar{\omega}_i / \omega_{\text{rms},i}^2 \\ 0 & \rho_i^2 \bar{\omega}_i / \omega_{\text{rms},i}^2 & \rho_i^2 / \omega_{\text{rms},i}^2 \end{pmatrix}, \quad (2.34)$$

where $\omega_{\text{rms},i}^2 = \bar{\omega}_i^2 - \bar{\omega}_i^2$. Reading off the $\tau\tau$ element of the covariance matrix reproduces the timing accuracy in Equation (24) of [Fairhurst \(2009\)](#),

$$\text{std}(\hat{\tau}_i - \tau_i) \geq \sqrt{(\mathcal{I}^{-1})_{\tau\tau}} = \frac{\rho_i}{\omega_{\text{rms},i}}. \quad (2.35)$$

The Fisher matrix in Equation (2.27) is block diagonal, which implies that estimation errors in the signal amplitude ρ are uncorrelated with the phase γ and time τ . A sequence of two changes of variables lends some physical interpretation to the nature of the coupled estimation errors in γ and τ .

First, we put the phase and time on the same footing by measuring the time in units of $1/\sqrt{\omega^2}$ with a change of variables from τ to $\gamma_\tau = \sqrt{\omega^2}\tau$:

$$\mathcal{I}' = \begin{matrix} & \begin{matrix} \rho_i & \gamma_i & \gamma_{\tau,i} \end{matrix} \\ \begin{matrix} \rho_i \\ \gamma_i \\ \gamma_{\tau,i} \end{matrix} & \begin{pmatrix} 1 & 0 & 0 \\ 0 & \rho_i^2 & -\rho_i^2 \frac{\bar{\omega}_i}{\sqrt{\omega_i^2}} \\ 0 & -\rho_i^2 \frac{\bar{\omega}_i}{\sqrt{\omega_i^2}} & \rho_i^2 \end{pmatrix} \end{matrix}. \quad (2.36)$$

The second change of variables, from γ and γ_τ to $\gamma_\pm = \frac{1}{\sqrt{2}}(\gamma \pm \gamma_\tau)$, diagonalizes the Fisher matrix:

$$\mathcal{I}'' = \begin{matrix} & \begin{matrix} \rho_i & \gamma_{+,i} & \gamma_{-,i} \end{matrix} \\ \begin{matrix} \rho_i \\ \gamma_{+,i} \\ \gamma_{-,i} \end{matrix} & \begin{pmatrix} 1 & 0 & 0 \\ 0 & \left(1 - \frac{\bar{\omega}_i}{\sqrt{\omega_i^2}}\right) \rho_i^2 & 0 \\ 0 & 0 & \left(1 + \frac{\bar{\omega}_i}{\sqrt{\omega_i^2}}\right) \rho_i^2 \end{pmatrix} \end{matrix}. \quad (2.37)$$

Thus, in the appropriate time units, the *sum and difference* of the phase and time of the signal are measured independently.

2.6 Position resolution

Finally, we will calculate the position resolution of a network of GW detectors. We could launch directly into computing derivatives of the full signal model from Equation (2.6) with respect to all of the parameters, but this would result in a very complicated expression. Fortunately, we can take two shortcuts. First, since we showed in Section 2.4 that the intrinsic parameters are correlated only with an overall nuisance average arrival time, amplitude, and phase, we need not consider the derivatives with respect to mass at all. Second, we can reuse the extrinsic part of the single detector Fisher matrix from Equation (2.27) by computing the much simpler Jacobian matrix to transform from the time, amplitude, and phase on arrival, to the parameters of interest.

We begin by transforming the single-detector Fisher matrix from a polar to a rectangular representation of the complex amplitude given in Equations (2.14, 2.13), $\rho_i, \gamma_i \rightarrow \Re[z_i] = \rho_i \cos \gamma_i, \Im[z_i] = \rho_i \sin \gamma_i$:

$$\mathcal{I}_i = \begin{matrix} & \Re[z_i] & \Im[z_i] & \tau_i \\ \begin{matrix} \Re[z_i] \\ \Im[z_i] \\ \tau_i \end{matrix} & \begin{pmatrix} 1 & 0 & \bar{\omega}_i b_i \\ 0 & 1 & -\bar{\omega}_i b_i \\ \bar{\omega}_i b_i & -\bar{\omega}_i b_i & \rho_i^2 \bar{\omega}_i^2 \end{pmatrix} \end{matrix}. \quad (2.38)$$

Consider a source in a “standard” orientation with the direction of propagation along the $+z$ axis, such that the GW polarization tensor may be written in Cartesian coordinates as

$$H = \frac{1}{r} e^{2i\phi_c} \begin{pmatrix} \frac{1}{2}(1 + \cos^2 \iota) & i \cos \iota & 0 \\ i \cos \iota & -\frac{1}{2}(1 + \cos^2 \iota) & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (2.39)$$

Now introduce a rotation matrix R that actively transforms this source to the Earth-relative polar coordinates θ, ϕ , and gives the source a polarization angle ψ (adopting temporarily the notation

$c_\theta = \cos \theta$, $s_\theta = \sin \theta$):

$$R = R_z(\phi) R_y(\theta) R_z(\psi) R_y(\pi) \quad (2.40)$$

$$= \begin{pmatrix} c_\phi & -s_\phi & 0 \\ s_\phi & -c_\phi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} c_\theta & 0 & s_\theta \\ 0 & 1 & 0 \\ -s_\theta & 0 & c_\theta \end{pmatrix} \begin{pmatrix} c_\psi & -s_\psi & 0 \\ s_\psi & -c_\psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \quad (2.41)$$

(The rightmost rotation reverses the propagation direction so that the wave is traveling *from* the sky position θ, ϕ .) With the (symmetric) detector response tensor D_i , we can write the received amplitude and arrival time as

$$z_i = r_{1,i} \text{Tr} \left[D_i R H R^\top \right], \quad (2.42)$$

$$\tau_i = t_\oplus + \mathbf{d}_i^\top R \mathbf{k}. \quad (2.43)$$

Equivalently, we can absorb the rotation R and the horizon distance $r_{1,i}$ into the polarization tensor, detector response tensors, and positions,

$$H \rightarrow H' = R_z(\psi) R_y(\pi) H R_y(\pi)^\top R_z(\psi)^\top, \quad (2.44)$$

$$D_i \rightarrow D'_i = r_{1,i} R_y(\theta)^\top R_z(\phi)^\top D_i R_z(\phi) R_y(\theta), \quad (2.45)$$

$$\mathbf{d}_i \rightarrow \mathbf{d}'_i = R_y(\theta)^\top R_z(\phi)^\top \mathbf{d}_i, \quad (2.46)$$

$$\mathbf{k} \rightarrow \mathbf{k}' = (0, 0, -1). \quad (2.47)$$

Now the model becomes

$$H' = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (2.48)$$

$$z_i = \text{Tr} [D'_i H'] = h_+(D'_{00} - D'_{11}) + 2h_\times D'_{01}, \quad (2.49)$$

$$\tau_i = t_\oplus + (\mathbf{d}'_i) \cdot \mathbf{k}, \quad (2.50)$$

where

$$h_+ = \frac{1}{r} e^{2i\phi_c} \left[\frac{1}{2} (1 + \cos^2 \iota) \cos 2\psi + i \cos \iota \sin 2\psi \right], \quad (2.51)$$

$$h_\times = \frac{1}{r} e^{2i\phi_c} \left[\frac{1}{2} (1 + \cos^2 \iota) \sin 2\psi - i \cos \iota \cos 2\psi \right]. \quad (2.52)$$

We insert an infinitesimal rotation δR to perturb the source's orientation from the true value:

$$z_i = \text{Tr} \left[D'_i(\delta R) H'(\delta R)^\top \right], \quad (2.53)$$

$$\tau_i = t_\oplus + (\mathbf{d}'_i)^\top (\delta R) \mathbf{k}'. \quad (2.54)$$

We only need a first order expression for δR , because we will be taking products of first derivatives of it⁴:

$$\delta R = \begin{pmatrix} 1 & 0 & \delta\theta \\ 0 & 1 & \delta\phi \\ -\delta\theta & -\delta\phi & 1 \end{pmatrix}. \quad (2.55)$$

We construct a Jacobian matrix J_i to transform from the single-detector observables $(\Re[z_i], \Im[z_i], \tau_i)$ to the position perturbations, polarization components, and geocentered arrival time $(\delta\theta, \delta\phi, \Re[h_+], \Im[h_+], \Re[h_\times], \Im[h_\times], t_\oplus)$:

$$J_i^\top = \begin{matrix} & \Re[z_i] & \Im[z_i] & \tau_i \\ \begin{matrix} \delta\theta \\ \delta\phi \\ \Re[h_+] \\ \Im[h_+] \\ \Re[h_\times] \\ \Im[h_\times] \\ t_\oplus \end{matrix} & \begin{pmatrix} -2\Re[h_+]D'_{02} - 2\Re[h_\times]D'_{12} & -2\Im[h_+]D'_{02} - 2\Im[h_\times]D'_{12} & -d'_0 \\ -2\Re[h_\times]D'_{02} + 2\Re[h_+]D'_{12} & -2\Im[h_\times]D'_{02} + 2\Im[h_+]D'_{12} & -d'_1 \\ D'_{00} - D'_{11} & 0 & 0 \\ 0 & D'_{00} - D'_{11} & 0 \\ 2D'_{01} & 0 & 0 \\ 0 & 2D'_{01} & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{matrix}. \quad (2.56)$$

⁴Caution: the angles $\delta\theta$ and $\delta\phi$ represent displacements in two orthogonal directions, but are *not* necessarily simply related to θ and ϕ .

We transform and sum the information from each detector,

$$\mathcal{I}_{\text{net}} = \sum_i J_i^\top \mathcal{I}_i J_i. \quad (2.57)$$

2.6.1 Marginalization over nuisance parameters

To extract an area from the Fisher matrix, we must first marginalize or discard the nuisance parameters. Note that marginalizing parameters of a multivariate Gaussian distribution amounts to simply dropping the relevant entries in the mean vector and covariance matrix. Since the information is the inverse of the covariance matrix, we need to invert the Fisher matrix, drop all but the first two rows and columns, and then invert again.

This procedure has a shortcut called the Schur complement (see, for example, [Press et al. 2007a](#)). Consider a partitioned square matrix M and its inverse:

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \quad M^{-1} = \begin{pmatrix} \tilde{A} & \tilde{B} \\ \tilde{C} & \tilde{D} \end{pmatrix}. \quad (2.58)$$

If A and B are square matrices, then the upper-left block of the inverse can be written as

$$\tilde{A}^{-1} = A - BD^{-1}C. \quad (2.59)$$

If we partition the \mathcal{I}_{net} similarly, the A block consists of the first two rows and columns and D is the lower right block that describes all other parameters. Because the Fisher matrix is symmetric, the off-diagonal blocks satisfy $C = B^\top$. Then the Schur complement

$$\mathcal{I}_{\text{marg}} = A - BD^{-1}B^\top \quad (2.60)$$

gives us the information matrix marginalized over all parameters but $\delta\theta$ and $\delta\phi$.

2.6.2 Spatial interpretation

How do we extract the dimensions of the localization from the Fisher matrix? If there are $N \leq 2$ detectors, then the Fisher matrix must be degenerate, because there are $3N$ measurements and 7

parameters:

$$\left\{ \begin{array}{c} \delta\theta \\ \delta\phi \\ \Re[h_+] \\ \Im[h_+] \\ \Re[h_\times] \\ \Im[h_\times] \\ t_\oplus \end{array} \right\} = 7 \text{ parameters} \quad \longleftrightarrow \quad \left\{ \begin{array}{c} \Re[z_i] \\ \Im[z_i] \\ \tau_i \end{array} \right\} \times N = 3N \text{ observables.}$$

Therefore, for $N = 2$ detectors, the marginalized Fisher matrix $\mathcal{I}_{\text{marg}}$ is singular. Its only nonzero eigenvalue λ describes the width of an annulus on the sky. The width of the annulus that contains probability p is given by

$$L_p = 2\sqrt{2} \operatorname{erf}^{-1}(p) / \sqrt{\lambda}. \quad (2.61)$$

The prefactor $2\sqrt{2}/\operatorname{erf}^{-1}(p)$ is the central interval of a normal distribution that contains a probability p , and is ≈ 3.3 for $p = 0.9$. *Caution:* for two-detector networks, priors play an important role in practical parameter estimation and areas can be much smaller than one would predict from the Fisher matrix (see Chapter 5 for more discussion).

For $N \geq 3$ detectors, the parameters are over-constrained by the data and the Fisher matrix describes the dimensions of an ellipse. Within a circle of radius r centered on the origin, the enclosed probability p is

$$p = \int_0^{2\pi} \int_0^r \frac{1}{2\pi} e^{-s^2/2} s \, ds \, d\phi = 1 - e^{-r^2/2}. \quad (2.62)$$

Therefore the radius r of the circle that contains a probability p is

$$r = \sqrt{-2 \ln(1 - p)}. \quad (2.63)$$

Suppose that the eigenvalues of the Fisher matrix are λ_1 and λ_2 . This describes a 1σ uncertainty ellipse that has major and minor radii $\lambda_1^{-1/2}$, $\lambda_2^{-1/2}$, and area $A_{1\sigma} = \pi/\sqrt{\lambda_1\lambda_2} = \pi/\sqrt{\det \mathcal{I}}$.

Then the area of an ellipse containing probability p is

$$A_p = -2\pi \ln(1 - p) / \sqrt{\det \mathcal{I}}, \quad (2.64)$$

or, more memorably for the 90th percentile, $A_{0.9} = 2\pi \ln(10) / \sqrt{\det \mathcal{I}}$.

2.6.3 Outline of calculation

Using the above derivation, we arrive at a prediction for the sky resolution of a GW detector network. We took some shortcuts that allowed us to avoid directly evaluating the complicated derivatives of the signal itself with respect to sky location. As a result, the expressions involved in each step are simple enough to be manually entered into a computer program. However, because the procedure involves several steps, we outline it once again below.

1. Compute, for each detector, the horizon distance $r_{1,i}$, the angular frequency moments $\bar{\omega}_i$ and $\bar{\omega}_i^2$, and (h_+, h_\times) from Equations (2.51, 2.52). (These can be reused for multiple source positions as long as the masses and the detector noise PSDs are the same.)
2. For a given ϕ, θ, ψ , compute the complex received amplitude z_i from Equations (2.48, 2.49), the extrinsic Fisher matrix from Equation (2.27), and the Jacobian from Equation (2.56).
3. Sum the information from all detectors using Equation 2.57.
4. Compute the marginalized Fisher matrix from the Schur complement using Equation (2.60).
5. If there are two detectors, find the width L_p of the ring describing the p th quantile using Equation (2.61). If there are three or more detectors, find the area A_p of the p th quantile using Equation (2.64).
6. (Optionally, convert from (steradians) to (square) degrees.)

See code listing in Appendix A.6.

2.6.4 Example calculation for HLV network

As an example, we calculate the 90% area as a function of sky position. We consider a three-detector, Hanford–Livingston–Virgo (HLV) network at final design sensitivity. Our source is a 1.4–1.4 M_\odot

BNS merger at a fixed distance of 180 Mpc and a fixed inclination angle of $\iota = 30^\circ$. The signal model is a stationary phase approximation waveform accurate to an order of 3.5 PN⁵. Aside from sky location, the polarization angle ψ remains a free parameter; for the purpose of this example calculation, at each sky location we pick the value of ψ that minimizes the area. The result is shown in Figure 2.2(a). The plot is shown in geographic coordinates to preserve spatial relationships to the Earth-fixed detector locations and antenna patterns.

For these sources, the area ranges from about 3 to 200 deg². It is smaller than 10 deg² across much of the sky, but much larger than 10 deg² in a broad ring that is oblique to but near the equator. This is the great circle that is parallel to the plane of the detectors. Here, the times of arrival are more sensitive than anywhere else to the azimuthal angle around the ring, but are constant to first order in the elevation angle relative to this plane. In Figure 2.2(b), we show the sky resolution calculated using times of arrival only. Without the coherence between detectors, the Fisher matrix becomes singular everywhere along this ring.

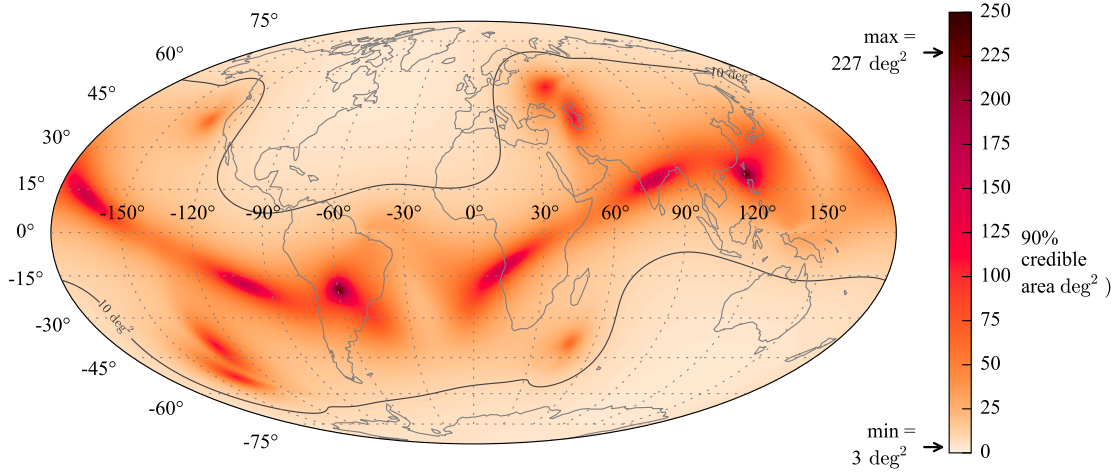
Other prominent features in Figure 2.2 include four small knots of coarse resolution at large angles to the great circle described above. These are positions where the *ratio* of the antenna patterns of any two detectors is large—for instance, positions where the Hanford detector is very sensitive but the Livingston detector is not. In Figure 2.3, we have plotted a heatmap of the logarithm of the ratios of the SNRs in each of the three pairs of detectors; these plots exhibit hotspots at the same locations as those isolated spots of large area uncertainty in Figure 2.2.

2.6.5 Improvement in localization due to coherence

Much of the current expectations for GW sky localization accuracy are based on Fisher matrix calculations that use time of arrival information alone (Fairhurst, 2009; Wen & Chen, 2010; Aasi et al., 2013c; Fairhurst, 2014). There is a question as to whether these estimates can be straightforwardly ‘calibrated’ with a scale factor to give realistic areas that account for coherence. Grover et al. (2014) studied the relationship between the areas predicted this way versus areas found using the full Fisher matrix or by actually performing MCMC parameter estimation. They postulated two possible expressions for the ratio between the area computed with coherence

⁵The “TaylorF2” waveform.

(a) Coherent



(b) Time of arrival only

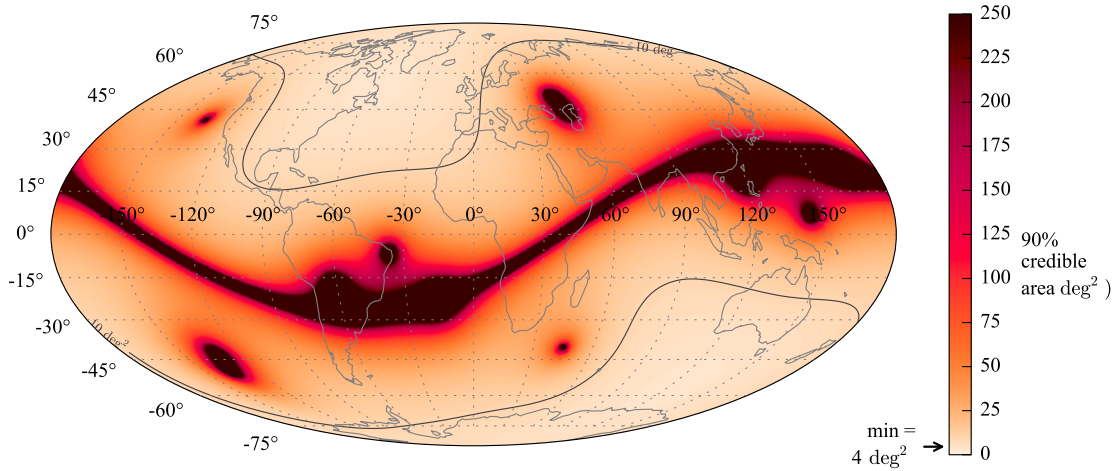


Figure 2.2 Position resolution as a function of sky position for a three-detector (HLV) network at final design sensitivity. The distance of the source is set to a constant value of 180 Mpc and the inclination angle is fixed to $\iota = 30^\circ$. At each sky location, the polarization angle ψ is varied to minimize the position uncertainty. Panel (a) shows the area of the 90% credible region as a function of geographic coordinates. The minimum and maximum areas are marked on the color bar. The contour within which sources are localized to 10 deg^2 is shown as a black curve. Panel (b) shows the area that would be found using time of arrival information only (but assuming the same polarization angle). Note that in (b), the maximum area is off the scale; black regions of the plot are localized to areas worse than 250 deg^2 .

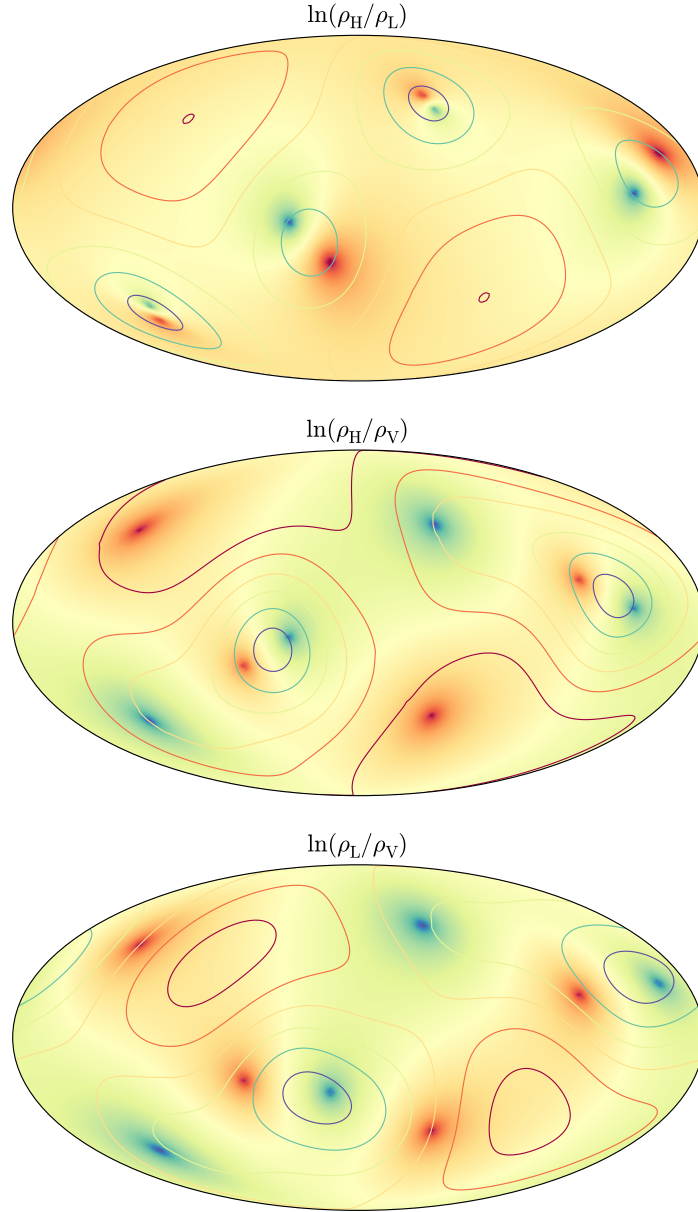


Figure 2.3 Log ratios of SNRs in pairs of detectors, $\ln(\rho_i/\rho_j)$ for detectors i and j . Contour lines show the network SNR for the pair of detectors, $\sqrt{\rho_i^2 + \rho_j^2}$.




Figure 2.4 Improvement in localization by including coherence. The color represents the ratio between the area as calculated assuming amplitude and phase consistency between detectors, and the area as calculated from just time delays on arrival. The median value and the two predictions from (Grover et al., 2014) are marked on the color bar.

versus the area computed from timing only:

$$\frac{A \text{ (coherent)}}{A \text{ (timing only)}} = \frac{\overline{\omega^2} - \bar{\omega}^2}{\bar{\omega}^2} \quad (2.65)$$

or

$$\frac{A \text{ (coherent)}}{A \text{ (timing only)}} = \sqrt{\frac{\overline{\omega^2} - \bar{\omega}^2}{\bar{\omega}^2}}. \quad (2.66)$$

This is motivated by the fact that the linear dimensions of the uncertainty ellipse, using only timing, are proportional to $1/\omega_{\text{rms}} = 1/\sqrt{\overline{\omega^2} - \bar{\omega}^2}$, whereas if the phase was known exactly (and removed altogether from the Fisher matrix, without marginalization), the linear dimensions would be proportional to $1/\sqrt{\bar{\omega}^2}$.

Note that in a two-detector network, the *area* scales as $A \propto 1/\sqrt{\omega_{\text{rms}}}$, because the localization regions are annuli whose areas are determined by a single linear dimension. For a network of three or more detectors, the area scales as $A \propto 1/\omega_{\text{rms}}$, because the area depends on two linear dimensions. (Similar scaling relations would hold for *phase-only* localization). Equation (2.65), therefore, may apply to three, but not two, detectors; conversely, we should expect that Equation (2.65) holds for two but not three detectors.

However, there are other problems with Equations (2.65) and (2.66). First, for two detector networks there is *no* improvement in sky localization due to adding coherence; the marginal Fisher matrix is *identical* whether computed with timing only or with the full signal model. When combined with a *distance prior*, coherence can partially break degeneracies and improve localization relative to triangulation, but this effect is not represented by the Fisher matrix calculation. Second, they do not account for the fact that in networks of three or more detectors, coherence *itself* can inform sky localization, even *without* time of arrival information; the effect of phase measurement goes beyond the correlations with time of arrival measurement errors. Third, wherever coherence breaks degeneracies in the timing analysis, the ratio in areas can be arbitrarily extreme because the area computed from timing can be almost arbitrarily large. Fourth, in configurations where the posterior probability distributions have multiple modes, coherence and priors can reduce areas significantly (see Chapter 5) in a way that is not captured by the Fisher matrix. This is a problem with the Fisher matrix itself, inherited by Equations (2.65) and (2.66).

In Figure 2.4 we show the ratio between the area computed using the coherent Fisher matrix versus the area computed using timing alone. The median ratio of 0.69 is marked on the color bar, along with the values of the two Grover et al. (2014) formulas, Equation (2.65) ≈ 0.56 and Equation (2.66) ≈ 0.75 . The median is a little closer to Equation (2.66), contrary to what would be expected for a three-detector network. There are small pockets where the ratio is almost 1, i.e., no improvement. There is also a thin strip and two broad regions near the detector plane where the ratio is almost arbitrarily small (because the triangulation localization has a degeneracy that is broken by coherence). Neither extreme is well described by the Grover et al. (2014) formulas, and the median is not particularly close to one formula or the other. The improvement in area seems to be highly sensitive to the detector configuration and the sky position.

2.6.6 Revision to LIGO observing scenarios document

Recently, the LIGO/Virgo collaboration published a document outlining possible commissioning and observing timetables from 2015 through 2022 (Aasi et al., 2013c). This document estimates the detection rate and sky localization accuracy as the detector network evolves. Sky areas are estimated using triangulation considerations, as in Fairhurst (2009). As we have now shown, this can modestly overestimate the true uncertainties for three-detector networks.

Here, we characterize the resolution of an HLV network assuming full coherence. First, we generated a sample of detectable BNS signals by drawing samples from a spatially uniform, isotropic distribution and checking whether the SNR was ≥ 4 in at least two detectors and the network SNR in those detectors, ρ_{net} , was at least 12.⁶ Each detector has a random and independent duty cycle of 80%. For all surviving sources, we computed the area from the Fisher matrix, assuming full coherence and also using only timing information. The 90% confidence ellipses for all of these sources are shown in Figure 2.5.

These assumptions are similar to, but not exactly the same as, the “2019” scenario from Aasi et al. (2013c). They assume that Virgo is slightly less sensitive due to differing commissioning timetables. We assume a Virgo BNS range of 154 Mpc while they assume 65-130 Mpc. They assumed a single-detector threshold of 5, and calculated the network SNR from all operating detectors, rather than just those detectors with SNRs above the single-detector threshold. Fur-

⁶This is not to be confused with the different simulated detection population that we introduce in Chapter 5, which arises from HL and HLV detector networks at *early* Advanced LIGO/Virgo sensitivity.

Table 2.1. Revised sky resolution predictions for Advanced LIGO observing scenarios

Epoch	Range (Mpc)		% Localized within		Median area (deg ²)
	LIGO	Virgo	5 deg ²	20 deg ²	
2019+ (Aasi et al., 2013c)	200	65–130	3–8	8–28	N/A
HLV (coherent)	199	154	5.5	36	43
HLV (timing)	199	154	3.4	27	73

Note. — Revised predictions of sky resolution at final Advanced LIGO+Virgo design sensitivity, compared with [Aasi et al. \(2013c\)](#). (The median value is omitted from the first row because [Aasi et al. \(2013c\)](#) does not specify one.)

thermore, [Aasi et al. \(2013c\)](#) neglected the contribution to sky localization from detectors with $\rho < 3$.

Summary statistics compared against [Aasi et al. \(2013c\)](#) are shown in Table 2.1. Percentiles were computed by dividing the number of three-detector events localized within a given area by the total number of events (i.e., the denominator included both two- and three-detector events; equivalent to treating the areas for two-detector events as infinite). For timing only, we find that 4.5% of sources are localized within 5 deg², consistent with the range of 3–8% in [Aasi et al. \(2013c\)](#). However, we find that 30% of sources are localized within 20 deg², at the high end of the claimed range of 8–28%. This moderate disagreement is probably due to differing Virgo sensitivities and slightly different detection criteria. Adding full coherence, we find that 7.4% of sources are found within 5 deg² and 38% within 20 deg². The median decreases from 83 to 46 deg², shrinking in area by a factor of 0.55.

Assuming statistically independent single-detector duty cycles⁷ of 0.8, all three detectors are operational about half the time; most of the rest of the time only a pair of detectors is observing. Including the localization of two-detector events would modestly improve the median area.

⁷In previous LIGO/Virgo science runs, the duty cycles of distinct detectors were somewhat correlated. The detectors tended to remain locked longer at night due to reduced local anthropogenic ground motion. Advanced detectors with improved seismic isolation systems may be more resistant to seismically triggered lock loss and have duty cycles that are less correlated.




Figure 2.5 90% confidence ellipses of a random sample of sources that would be detectable by a three-detector, HLV network at final design sensitivity.

2.7 Summary

We have presented an idealized model of the noise and signal content in GW observations with a network of laser interferometric detectors. This model led us to the matched filter, the maximum likelihood estimator, and the network SNR as the simplest possible statistic for discriminating between the presence or absence of an astrophysical signal. Having defined SNR, we worked out the horizon distance of a GW detector. Next, we used the Fisher matrix formalism to calculate the sky resolution of a GW detector network. Previous approaches ([Fairhurst 2009](#), etc.) have considered only arrival time measurements, but our computation also accounts for measurements of phase and amplitude on arrival. Despite this, our method as outlined in [Section 2.6.3](#) is only slightly more complicated than the timing-only calculation. Though the Fisher matrix analysis cannot be applied to two-detector networks, we do expect it to accurately predict the sky localization accuracy of signals that are confidently detected by networks of three or more detectors of comparable sensitivity. We endorse it as the most sophisticated analysis worth carrying out, short of performing full Bayesian parameter estimation on a population of simulated signals. We advocate using it to revise the overly pessimistic sky resolution predictions for the 2019 and later scenarios in [Aasi et al. \(2013c\)](#).

The horizon distance, the observables and parameters in GW observations of compact binaries, the Fisher matrix, and estimates of sky resolution will recur in later chapters. In the next chapter, we will use all of these results to study the prospects for detecting and localizing BNS signals in near real-time, from hundreds of seconds before to seconds after merger.

Chapter 3

Early warning G W detection

KIPP CANNON¹, ROMAIN CARIOU², ADRIAN CHAPMAN³, MIREIA CRISPIN-ORTUZAR⁴,
NICKOLAS FOTOPOULOS³, MELISSA FREI^{5,6}, CHAD HANNA⁷, ERIN KARA⁸, DREW KEPPEL^{9,10},
LAURA LIAO¹¹, STEPHEN PRIVITERA³, ANTONY SEARLE³, LEO SINGER³, AND ALAN WEINSTEIN³

¹Canadian Institute for Theoretical Astrophysics, Toronto, ON, Canada

²Département de Physique, École Normale Supérieure de Cachan, Cachan, France

³LIGO Laboratory, California Institute of Technology, MC 100-36, 1200 E. California Blvd., Pasadena, CA, USA

⁴Facultat de Física, Universitat de València, Burjassot, Spain

⁵Department of Physics, University of Texas at Austin, Austin, TX, USA

⁶Center for Computational Relativity and Gravitation and School of Mathematical Sciences, Rochester Institute of Technology, Rochester, NY, USA

⁷Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada

⁸Department of Physics and Astronomy, Barnard College, Columbia University, New York, NY, USA

⁹Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, Hannover, Germany

¹⁰Institut für Gravitationsphysik, Leibniz Universität Hannover, Hannover, Germany

¹¹Department of Chemistry and Biology, Ryerson University, Toronto, ON, Canada

This chapter is reproduced in part from Cannon et al. (2012), which was published in The Astrophysical Journal as “Toward early-warning detection of gravitational waves from compact binary coalescence,” copyright © 2012 The American Astronomical Society. Note that the author list is alphabetical because of the large number of contributors to this algorithm and codebase, initiated and organized by K.C., C.H., and D.K. My contributions to that code are related to the interpolation, decimation, and triggering stages. My contributions to the project include all of the calculations related to “early warning” detection and localization, all of the accounting of the computational cost, and the measurement of the mismatch of the template bank. I prepared all of the figures and tables in this publication and about 80% of the text, and was the corresponding author.

In the first generation of ground-based laser interferometers, the GW community initiated a project to send alerts when potential GW transients were observed in order to trigger follow-up observations by EM telescopes. The typical latencies were 30 minutes (Hughey, 2011). This was an important achievement, but too late to catch any prompt (i.e., simultaneous with gamma-ray emission) optical flash and later than would be desirable to search for an on-axis optical afterglow (which fades rapidly as a power law in time; see for example Metzger & Berger 2012). Since the GW signal is in principle detectable even before the tidal disruption, one might have the ambition of reporting GW candidates not minutes after the merger, but seconds before. We explore one essential ingredient of this problem, a computationally inexpensive real-time filtering algorithm for detecting inspiral signals in GW data. We also consider the prospects for advanced GW detectors and discuss other areas of work that would be required for rapid analysis.

In October 2010, LIGO completed its sixth science run (S6) and Virgo completed its third science run (VSR3). While both LIGO detectors and Virgo were operating, several all-sky detection pipelines operated in a low-latency configuration to send astronomical alerts, namely, Coherent WaveBurst (cWB), Omega, and Multi-Band Template Analysis (MBTA; Abadie et al., 2012a,b). cWB and Omega are both unmodeled searches for bursts based on time-frequency decomposition of the GW data. MBTA is a novel kind of template-based inspiral search that was purpose-built for low latency operation. MBTA achieved the best GW trigger-generation latencies, of 2–5 minutes. Alerts were sent with latencies of 30–60 minutes, dominated by human vetting. Candidates were sent for EM follow-up to several telescopes; *Swift*, LOFAR, ROTSE, TAROT, QUEST, SkyMapper, Liverpool Telescope, Pi of the Sky, Zadko, and Palomar Transient Factory imaged some of the most likely sky locations (Abadie et al., 2012b; Evans et al., 2012; Aasi et al., 2014).

There were a number of sources of latency associated with the search for CBC signals in S6/VSR3 (Hughey, 2011), listed here.

Data acquisition and aggregation ($\gtrsim 100$ ms) The LIGO data distribution system collects data in real time, but distributes it to computers in the control rooms 16 times a second, and archives it for immediate offsite replication in blocks of 16 s (Bork et al., 2001). Data are copied from all of the GW observatories to the analysis clusters over the Internet, which is capable of high

bandwidth but only modest latency. Altogether, it takes about 16 s to transmit the data to the analysis clusters, but with moderate changes in infrastructure could be reduced to ~ 100 ms if the data were streamed to the computing clusters in real-time without blocking it into 16 s chunks.

Data conditioning (~ 1 min) Science data must be calibrated using the detector’s frequency response to gravitational radiation. Currently, data are calibrated in blocks of 16 s. Within ~ 1 minute, data quality is assessed in order to create veto flags. These are both technical sources of latency that might be addressed with improved calibration and data quality software for advanced detectors.

Trigger generation (2–5 min) Low-latency data analysis pipelines deployed in S6/VSR3 achieved an impressive latency of minutes. However, second to the human vetting process, this dominated the latency of the entire EM follow-up process. Even if no other sources of latency existed, this trigger generation latency is too long to catch prompt or even extended emission. Low-latency trigger generation will become more challenging with advanced detectors because inspiral signals will stay in band up to 10 times longer. In this work, we will focus on reducing this source of latency.

Alert generation (2–3 min) S6/VSR3 saw the introduction of low-latency astronomical alerts, which required gathering event parameters and sky localization from the various online analyses, downselecting the events, and calculating telescope pointings. If other sources of latency improve, the technical latency associated with this infrastructure could dominate, so work should be done to improve it.

Human validation (10–20 min) Because the new alert system was commissioned during S6/VSR3, all alerts were subjected to quality control checks by human operators before they were disseminated. This was by far the largest source of latency during S6/VSR3. Hopefully, confidence in the system will grow to the point where no human intervention is necessary before alerts are sent, so we give it no further consideration here.

This chapter will focus on reducing the latency of trigger production. Data analysis strategies for advance detection of CBCs will have to strike a balance between latency and throughput.

CBC searches consist of banks of matched filters, or cross-correlations between the data stream and a bank of nominal “template” signals. There are many different implementations of matched filters, but most have high throughput at the cost of high latency, or low latency at the cost of low throughput. The former are epitomized by the overlap-save algorithm for FD convolution, currently the preferred method in GW searches. The most obvious example of the latter is direct TD convolution, which can be done in real-time. However, its cost in floating point operations per second is linear in the length of the templates, so it is prohibitively expensive for long templates. The computational challenges of low-latency CBC searches are still more daunting for advanced detectors for which the inspiral signal remains in band for a large fraction of an hour (see Appendix B).

Fortunately, the morphology of inspiral signals can be exploited to offset some of the computational complexity of known low-latency algorithms. First, the signals evolve slowly in frequency, so that they can be broken into contiguous band-limited time intervals and processed at possibly lower sample rates. Second, inspiral filter banks consist of highly similar templates, admitting methods such as the singular value decomposition (SVD) (Cannon et al., 2010) or the Gram-Schmidt process (Field et al., 2011) to reduce the number of templates.

Several efforts that exploit one or both of these properties are under way to develop low-latency CBC search pipelines with tractable computing requirements. One example is MBTA (Marion & the Virgo Collaboration, 2003; Buskulic et al., 2010), which was deployed in S6/VSR3. MBTA consists of multiple, usually two, template banks for different frequency bands, one which is matched to the early inspiral and the other which is matched to the late inspiral. An excursion in the output of any filter bank triggers coherent reconstruction of the full matched filtered output. Final triggers are built from the reconstructed matched filter output. Another novel approach using networks of parallel, second-order infinite impulse response (IIR) filters is being explored by Hooper et al. (2010) and Luan et al. (2012).

We will use both properties to demonstrate that a very low latency detection statistic is possible with current computing resources. Assuming the other technical sources of latency can be reduced significantly, this could make it possible to send prompt (< 1 minute) alerts to the astronomical community.

The chapter is organized as follows. First, we describe the standard, offline CBC detection process. Using a simple model of the detection and sky localization accuracy of this search, we




Figure 3.1 Expected number of NS–NS sources that could be detectable by Advanced LIGO a given number of seconds before coalescence. The heavy solid line corresponds to the most probable yearly rate estimate from [Abadie et al. \(2010b\)](#). The shaded region represents the 5%–95% confidence interval arising from substantial uncertainty in predicted event rates.

study the prospects for early-warning detection. Then, we provide an overview of our novel method for detecting CBC signals near real-time. We then describe a prototype implementation using open source signal processing software. To validate our approach we present a case study focusing on a particular subset of the NS–NS parameter space. We conclude with some remarks on what remains to be prepared for the advanced detector era.

3.1 Prospects for early-warning detection and EM follow-up

Before the GW signal leaves the detection band, we can imagine examining the SNR accumulated up to that point and, if it is already significant, releasing an alert immediately, trading SNR and sky localization accuracy for pre-merger detection.

In the quadrupole approximation, the instantaneous frequency of the GW inspiral signal is related to the time t relative to coalescence (Section 5.1 of [Sathyaprakash & Schutz, 2009](#)) through

$$f(t) = \frac{1}{\pi \mathcal{M}_t} \left[\frac{5}{256} \frac{\mathcal{M}_t}{t} \right]^{3/8}, \quad (3.1)$$

where $\mathcal{M} = M^{2/5}\mu^{3/5}$ is the chirp mass of the binary, $\mathcal{M}_t = GM/c^3$ is the chirp mass in units of time, M is the total mass, and μ is the reduced mass. The expected value of the single-detector SNR for an optimally oriented (source at detector's zenith or nadir, orbital plane face-on) inspiral source is (Abadie et al., 2010b)

$$\rho = \frac{\mathcal{M}_t^{5/6} c}{\pi^{2/3} D} \sqrt{\frac{5}{6} \int_{f_{\text{low}}}^{f_{\text{high}}} \frac{f^{-7/3}}{S(f)} df}, \quad (3.2)$$

where D is the luminosity distance and $S(f)$ is the one-sided power spectral density of the detector noise. f_{low} and f_{high} are low- and high- frequency limits of integration which may be chosen to extend across the entire bandwidth of the detector. If we want to trigger at a time t before merger, then we must cut off the SNR integration at $f_{\text{high}} = f(t)$, with $f(t)$ given by Equation (3.1) above.

Figure 3.1 shows projected early detectability rates for NS–NS binaries in Advanced LIGO assuming the anticipated detector sensitivity for the ‘zero detuning, high power’ configuration described in Shoemaker (2010) and NS–NS merger rates estimated in Abadie et al. (2010b). The merger rates have substantial measurement uncertainty due to the small sample of known double pulsar systems that will merge within a Hubble time; they also have systematic uncertainty due to sensitive dependence on the pulsar luminosity distribution function (Kalogera et al., 2004). The most probable estimates indicate that at a single-detector SNR threshold of 8 we will observe a total of 40 events yr^{-1} ; $\sim 10 \text{ yr}^{-1}$ will be detectable within 10 s of merger and $\sim 5 \text{ yr}^{-1}$ will be detectable within 25 s of merger if analysis can proceed with near zero latency.

We emphasize that any practical GW search will include technical delays due to light travel time between the detectors, detector infrastructure, and the selected data analysis strategy. Figure 3.1 must be understood in the context of all of the potential sources of latency, some of which are avoidable and some of which are not.

EM follow-up requires estimating the location of the GW source. The localization uncertainty can be estimated from the uncertainty in the time of arrival of the GWs, which is determined by the signal's effective bandwidth and SNR (Fairhurst, 2009). Table 3.1 and Figure 3.2 show the estimated 90% confidence area versus time of the loudest coalescence events detectable by Advanced LIGO and Advanced Virgo. This is the *minimum* area; localization is best at high elevation from the plane containing the detectors, and worst at zero elevation. Fairhurst also cautions that his Fisher matrix calculation fails to capture disconnected patches of probability,




Figure 3.2 Area of the 90% confidence region as a function of time before coalescence for sources with anticipated detectability rates of 40, 10, 1, and 0.1 yr^{-1} . The heavy dot indicates the time at which the accumulated SNR exceeds a single-detector threshold of 8.

which occur prominently in networks of three detectors where there are generally two local maxima on opposite sides of the plane of the detectors. Aside from the mirror degeneracy, characterizing the uncertainty region by the Fisher matrix alone tends to overestimate, rather than underestimate, the area for low-SNR events, but this effect is generally more than compensated for by the source being in an unfavorable sky location. For these reasons, the localization uncertainty estimated from timing is highly optimistic and will only suffice for an order-of-magnitude estimate. Once per year, we expect to observe an event with a final single-detector SNR of ≈ 27 whose location can be constrained to about 1300 deg^2 (3.1% of the sky) within 25 s of merger, 260 deg^2 (0.63% of the sky) within 10 s of merger, and 0.82 deg^2 (0.0020% of the sky) at merger.

The picture is qualitatively similar when we track sky localization area versus time for individual events using the coherent Fisher matrix approach described in Section 2.6.3. In Figure 3.3, we plot the 90% confidence area as a function of time before coalescence for the event sample described in Section 2.6.6. The track begins at the earliest time that the source is “confidently” detectable (single-detector SNR threshold of $\rho \geq 4$ in at least two detectors, and the network SNR from all of the detectors that are above threshold together yield a network SNR $\rho_{\text{net}} \geq 12$). We find that about 1% of events are detectable 100 s before merger, with areas of $\sim 10^3$ – 10^4 deg^2 . About 20% of sources are detectable 10 s before merger, with areas of $\sim 10^2$ – 10^3 deg^2 . By the time

Table 3.1 Horizon distance, SNR at merger, and area of 90% confidence at selected times before merger for sources with expected detectability rates of 40, 10, 1, and 0.1 yr⁻¹.

Rate yr ⁻¹	Horizon (Mpc)	Final SNR	A(90%) (deg ²)			
			25 s	10 s	1 s	0 s
40	445	8.0	—	—	—	9.6
10	280	12.7	—	1200	78	3.8
1	130	27.4	1300	260	17	0.8
0.1	60	58.9	280	56	3.6	0.2

Note. — A dash (—) signifies that the confidence area is omitted because at the indicated time the SNR would not have crossed the detection threshold of 8.

that the full signal has been acquired, the areas shrink to $\sim 10\text{--}100\text{ deg}^2$.

After merger, typical short GRB optical afterglows, scaled to Advanced LIGO distances, should be within reach of moderately deep optical transient experiments such as the Palomar Transient Factory (PTF; [Rau et al. 2009](#); [Law et al. 2009](#)), its successor ZTF ([Kulkarni, 2012](#); [Bellm, 2014](#); [Smith et al., 2014b](#)), the eagerly awaited Large Synoptic Survey Telescope (LSST; [Ivezic et al. 2008](#)), and also the BlackGEM array (dedicated to GW candidate follow-up). However, the exposure and overhead times of these instruments are too long to search the pre-merger error boxes. It is possible to reduce the localization uncertainty by only looking at galaxies from a catalog that lie near the sky location and luminosity distance estimate from the GW signal ([Nuttall & Sutton, 2010](#)) as was done in S6/VSR3. Within the expected Advanced LIGO NS–NS horizon distance, the number of galaxies that can produce a given signal amplitude is much larger than in Initial LIGO and thus the catalog will not be as useful for downselecting pointings for most events. However, exceptional GW sources will necessarily be extremely close. Within this reduced volume there will be fewer galaxies to consider for a given candidate and catalog completeness will be less of a concern. For some nearby events, using a galaxy catalog in connection with *pre-merger* triggers could make it possible to search for optical flashes with rapidly slewing telescopes such as TAROT. Last, if *Swift* had a fully automated TOO mode, then for exceptional sources that are detectable ~ 100 s before merger, one could slew *Swift* so as to use BAT to monitor the area of $\sim 1000\text{ deg}^2$ for the GRB itself, building a virtual all-sky GRB monitor out of LIGO and *Swift*.

combined (Sathyaprakash et al., in preparation).

3.2 Novel real-time algorithm for CBC detection

In this section, we describe a decomposition of the CBC signal space that reduces TD filtering cost sufficiently to allow for the possibility of early-warning detection with modest computing requirements. We expand on the ideas of [Marion & the Virgo Collaboration \(2003\)](#) and [Buskulic et al. \(2010\)](#) that describe a multi-band decomposition of the compact binary signal space that resulted in a search with minutes' latency during S6/VSR3 ([Hughey, 2011](#)). We combine this with the SVD rank-reduction method of [Cannon et al. \(2010\)](#) that exploits the redundancy of the template banks.

3.2.1 Discrete-time representation of a matched filter

Here, we translate the basic matched filter bank search described in Section 2.1 to a discrete time representation, suitable for studying the filter design. Suppose that the observed data $y[k]$ consists of a known, nominal signal $x[k]$, and additive, zero-mean noise $n[k]$:

$$y[k] = x[k] + n[k].$$

A matched filter is a linear filter, defined as

$$z[k] = \sum_{n=0}^{N-1} h[n] y[k-n] = z_s[k] + z_n[k],$$

where z_s is the response of the filter to the signal alone and z_n is the response of the signal to noise alone. The matched filter's coefficients maximize the ratio of the expectation of the filter's instantaneous response to the variance in the filter's output:

$$(\text{signal to noise})^2 = \frac{\mathbb{E}[z[0]]^2}{\text{var}[z[k]]} = \frac{z_s[0]^2}{\text{var}[z_n[k]]}.$$




Figure 3.3 Trajectories of 90% localization area versus time before coalescence for a random sample of BNS signals. Each track starts when the signal crosses a single-detector SNR threshold of $\rho \geq 4$ in at least two detectors, and when all of the detectors that are above threshold together yield a network SNR $\rho_{\text{net}} \geq 12$. The bottom panel shows the cumulative fraction of events that are detectable a given time before coalescence.

It is well known (see, for example, [Turin, 1960](#)) that if $n[k]$ is Gaussian and wide-sense stationary, then the optimum is obtained when

$$\tilde{h}[n] = \tilde{x}^*[n] \tilde{S}^{-1}[n],$$

up to an arbitrary multiplicative constant. Here, $\tilde{h}[n]$, $\tilde{x}[n]$, and $\tilde{y}[n]$ are the discrete Fourier transforms (DFTs) of $h[k]$, $x[k]$, and $y[k]$, respectively; $\tilde{S}[n] = \text{E} [\tilde{n}[n] \tilde{n}^*[n]]$ is the folded, two-sided, discrete power spectrum of $n[k]$. It is related to the continuous, one-sided power spectral density $S(f)$ through

$$\tilde{S}[n] = \begin{cases} S(n) & \text{if } n = 0 \text{ or } n = N/2 \\ S(nf^0/2N)/2 & \text{if } 0 < n < N/2 \\ \tilde{S}[N - n] & \text{otherwise,} \end{cases}$$

where N is the length of the filter and f^0 is the sample rate. (In order to satisfy the Nyquist-Shannon sampling criterion, it is assumed that the detector's continuous noise power spectral density $S(f)$ vanishes for all $f > f^0/2$, or alternatively, that the data are low-pass filtered prior to matched filtering.) The DFT of the output is

$$\tilde{z}[n] = \tilde{x}^*[n] \tilde{S}^{-1}[n] \tilde{y}[n] \equiv \left(\tilde{S}^{-1/2}[n] \tilde{x}[n] \right)^* \left(\tilde{S}^{-1/2}[n] \tilde{y}[n] \right). \quad (3.3)$$

The placement of parentheses in Equation (3.3) emphasizes that the matched filter can be thought of as a cross-correlation of a whitened version of the data with a whitened version of the nominal signal. In this chapter, we shall not describe the exact process by which the detector's noise power spectrum is estimated and deconvolved from the data; for the remainder of this chapter we shall define $y[k]$ as the *whitened* data stream. Correspondingly, from this point on we shall use $h[k]$ to describe the *whitened* templates, being the inverse DFT of $\left(\tilde{S}^{-1/2}[n] \tilde{x}[n] \right)^*$.

Inspiral signals are continuously parameterized by a set of intrinsic source parameters θ that determine the amplitude and phase evolution of the GW strain. For systems in which the effects of spin can be ignored, the intrinsic source parameters are just the component masses of the binary, $\theta = (m_1, m_2)$. For a given source, the strain observed by the detector is a linear combination of two waveforms corresponding to the '+' and '×' GW polarizations. Thus, we must design two filters for each θ .

The coefficients for the M filters are known as templates, and are formed by discretizing and time reversing the waveforms and weighting them by the inverse amplitude spectral density of the detector's noise. To construct a template bank, templates are chosen with $M/2$ discrete signal parameters $\theta_0, \theta_1, \dots, \theta_{M/2-1}$. These are chosen such that any possible signal will have an inner product ≥ 0.97 with at least one template. Such a template bank is said to have a *minimal match* of 0.97 (Owen & Sathyaprakash, 1999).

Filtering the detector data involves a convolution of the data with the templates. For a unit-normalized template $h_i[k]$ and whitened detector data $y[k]$, both sampled at a rate f^0 , the result can be interpreted as the SNR, $z_i[k]$, defined as

$$z_i[k] = \sum_{n=0}^{N-1} h_i[n] y[k-n]. \quad (3.4)$$

This results in M SNR time series. Local peak-finding across time and template indices results in single-detector triggers. Coincidences are sought between triggers in different GW detectors in order to form detection candidates.

Equation (3.4) can be implemented in the TD as a bank of finite impulse response (FIR) filters, requiring $\mathcal{O}(MN)$ floating point operations per sample. However, it is typically much more computationally efficient to use the convolution theorem and the FFT to implement fast convolution in the FD, requiring only $\mathcal{O}(M \lg N)$ operations per sample but incurring a latency of $\mathcal{O}(N)$ samples.

3.2.2 The LLOID method

Here we describe a method for reducing the computational cost of a TD search for CBCs. We give a zero latency, real-time algorithm that competes in terms of floating point operations per second with the conventional overlap-save FD method, which by contrast requires a significant latency due to the inherent acausality of the Fourier transform. Our method, called Low Latency Online Inspiral Detection (LLOID), involves two transformations of the templates that produce a network of orthogonal filters that is far more computationally efficient than the original bank of matched filters.

The first transformation is to chop the templates into disjointly supported intervals, or *time slices*. Since the time slices of a given template are disjoint in time, they are orthogonal with respect

to time. Given the chirp-like structure of the templates, the “early” (lowest frequency) time slices have significantly lower bandwidth and can be safely downsampled. Downsampling reduces the total number of filter coefficients by a factor of ~ 100 by treating the earliest part of the waveform at $\sim 1/100$ of the full sample rate. Together, the factor of 100 reduction in the number of filter coefficients and the factor of 100 reduction in the sample rate during the early inspiral save a factor of $\sim 10^4$ floating point operations per second (FLOPS) over the original (full sample rate) templates.

However, the resulting filters are still not orthogonal across the parameter space and are in fact highly redundant. We use the SVD to approximate the template bank by a set of orthogonal *basis filters* (Cannon et al., 2010). We find that this approximation reduces the number of filters needed by another factor of ~ 100 . These two transformations combined reduce the number of floating point operations to a level that is competitive with the conventional high-latency FD-matched filter approach. In the remainder of this section we describe the LLOID algorithm in detail and provide some basic computational cost scaling.

3.2.2.1 Selectively reducing the sample rate of the data and templates

The first step of our proposed method is to divide the templates into time slices in a TD analog to the FD decomposition employed by MBTA (Marion & the Virgo Collaboration, 2003; Buskulic et al., 2010). The application to GW data analysis is foreshadowed by an earlier FD convolution algorithm, proposed by Gardner (1995), based on splitting the impulse response of a filter into smaller blocks. We decompose each template $h_i[k]$ into a sum of S non-overlapping templates,

$$h_i[k] = \sum_{s=0}^{S-1} \begin{cases} h_i^s[k] & \text{if } t^s \leq k/f^0 < t^{s+1} \\ 0 & \text{otherwise,} \end{cases} \quad (3.5)$$

for S integers $\{f^0 t^s\}$ such that $0 = f^0 t^0 < f^0 t^1 < \dots < f^0 t^S = N$. The outputs of these new time-sliced filters form an ensemble of partial SNR streams. By linearity of the filtering process, these partial SNR streams can be summed to reproduce the SNR of the full template.

Since waveforms with neighboring intrinsic source parameters θ have similar time-frequency evolution, it is possible to design computationally efficient time slices for an extended region of parameter space rather than to design different time slices for each template.

For concreteness and simplicity, consider an inspiral waveform in the quadrupole approximation, for which the time-frequency relation is given by Equation (3.1). This monotonic time-frequency relationship allows us to choose time slice boundaries that require substantially less bandwidth at early times in the inspiral.

An inspiral signal will enter the detection band with some low frequency f_{low} at time t_{low} before merger. Usually the template is truncated at some prescribed time t^0 , or equivalent frequency f_{high} , often chosen to correspond to the last stable orbit (LSO). The beginning of the template is critically sampled at $2f_{\text{low}}$, but the end of the template is critically sampled at a rate of $2f_{\text{high}}$. In any time interval smaller than the duration of the template, the bandwidth of the filters across the entire template bank can be significantly less than the full sample rate at which data are acquired.

Our goal is to reduce the filtering cost of a large fraction of the waveform by computing part of the convolution at a lower sample rate. Specifically we consider here time slice boundaries with the smallest power-of-two sample rates that sub-critically sample the time-sliced templates. The time slices consist of the S intervals $[t^0, t^1), [t^1, t^2), \dots, [t^{S-1}, t^S)$, sampled at frequencies f^0, f^1, \dots, f^{S-1} , where f^s is at least twice the highest nonzero frequency component of any filter in the bank for the s th time slice.

The time-sliced templates can then be downsampled in each interval without aliasing, so we define them as

$$h_i^s[k] \equiv \begin{cases} h_i \left[k \frac{f}{f^s} \right] & \text{if } t^s \leq k/f^s < t^{s+1} \\ 0 & \text{otherwise.} \end{cases} \quad (3.6)$$

We note that the time slice decomposition in Equation (3.5) is manifestly orthogonal since the time slices are disjoint in time. In the next section, we examine how to reduce the number of filters within each time slice via SVD of the time-sliced templates.

3.2.2.2 Reducing the number of filters with the SVD

As noted previously, the template banks used in inspiral searches are by design highly correlated. Cannon et al. (2010) showed that applying the SVD to inspiral template banks greatly reduces the number of filters required to achieve a particular minimal match. A similar technique can be applied to the time-sliced templates as defined in Equation (3.6) above. The SVD is a matrix

factorization that takes the form

$$h_i^s[k] = \sum_{l=0}^{M-1} v_{il}^s \sigma_l^s u_l^s[k] \approx \sum_{l=0}^{L^s-1} v_{il}^s \sigma_l^s u_l^s[k], \quad (3.7)$$

where $u_l^s[k]$ are orthonormal *basis templates* related to the original time-sliced templates through the *reconstruction matrix*, $v_{il}^s \sigma_l^s$. The expectation value of the fractional loss in SNR is the SVD tolerance, given by

$$\left[\sum_{l=0}^{L^s-1} (\sigma_l^s)^2 \right] \left[\sum_{l=0}^{M-1} (\sigma_l^s)^2 \right]^{-1},$$

determined by the number L^s of basis templates that are kept in the approximation. Cannon et al. (2010) showed that highly accurate approximations of inspiral template banks could be achieved with few basis templates. We find that when combined with the time slice decomposition, the number of basis templates L^s is much smaller than the original number of templates M and improves on the rank reduction demonstrated in Cannon et al. (2010) by nearly an order of magnitude.

Because the sets of filters from each time slice form orthogonal subspaces, and the basis filters within a given time slice are mutually orthogonal, the set of all basis filters from all time slices forms an orthogonal basis spanning the original templates.

In the next section, we describe how we form our early-warning detection statistic using the time slice decomposition and the SVD.

3.2.2.3 Early-warning output

In the previous two sections, we described two transformations that greatly reduce the computational burden of TD filtering. We are now prepared to define our detection statistic, the early-warning output, and to comment on the computational cost of evaluating it.

First, the sample rate of the detector data must be decimated to match sample rates with each of the time slices. We will denote the decimated detector data streams using a superscript “ s ” to indicate the time slices to which they correspond. The operator H^\downarrow will represent the appropriate decimation filter that converts between the base sample rate f^0 and the reduced sample rate f^s :

$$y^s[k] = \left(H^\downarrow y^0 \right) [k].$$

We shall use the symbol H^\dagger to represent an interpolation filter that converts between sample rates f^{s+1} and f^s of adjacent time slices,

$$y^s[k] = \left(H^\dagger y^{s+1} \right) [k].$$

From the combination of the time slice decomposition in Equation (3.6) and the SVD defined in Equation (3.7), we define the early-warning output accumulated up to time slice s using the recurrence relation,

$$z_i^s[k] = \overbrace{\left(H^\dagger \rho_i^{s+1} \right) [k]}^{\text{S/N from previous time slices}} + \underbrace{\sum_{l=0}^{L^s-1} v_{il}^s \sigma_l^s}_{\text{reconstruction}} \overbrace{\sum_{n=0}^{N^s-1} u_l^s [n] y^s[k-n]}^{\text{orthogonal FIR filters}}. \quad (3.8)$$

Observe that the early-warning output for time slice 0, $z_i^0[k]$, approximates the SNR of the original templates. The signal flow diagram in Figure 3.4 illustrates this recursion relation as a multirate filter network with a number of early-warning outputs.

Ultimately, the latency of the entire LLOID algorithm is set by the decimation and interpolation filters because they are generally time symmetric and slightly acausal. Fortunately, as long as the latency introduced by the decimation and interpolation filters for any time slice s is less than that time slice's delay t^s , the total latency of the LLOID algorithm will be zero. To be concrete, suppose that the first time slice, sampled at a rate $f^0 = 4096$ Hz, spans times $[t^0, t^1) = [0 \text{ s}, 0.5 \text{ s})$, and the second time slice, sampled at $f^1 = 512$ Hz, spans $[t^1, t^2) = [0.5 \text{ s}, 4.5 \text{ s})$. Then the second time slice's output, $z_i^1[k]$, will *lead* the first time slice's output, $z_i^0[k]$, by 0.5 s. A decimation filter will be necessary to convert the 4096 Hz input signal $y[k] \equiv y^0[k]$ to the 512 Hz input $y^1[k]$, and an interpolation filter will be necessary to match the sample rates of the two early-warning outputs. In this example, as long as the decimation and interpolation filters are together acausal by less than $t^1 = 0.5$ s, the total SNR $z_i^0[k]$ will be available with a latency of zero samples. When zero latency is important, we may take this as a requirement for the decimation and interpolation filter kernels.

In the next section, we compute the expected computational cost scaling of this decomposition and compare it with the direct TD implementation of Equation (3.4) and higher latency blockwise

Table 3.2 Notation used to describe filters.

	Definition
f^s	Sample rate in time slice s
M	Number of templates
N	Number of samples per template
S	Number of time slices
L^s	Number of basis templates in time slice s
N^s	Number of samples in decimated time slice s
N^\downarrow	Length of decimation filter
N^\uparrow	Length of interpolation filter

FD methods.

3.2.3 Comparison of computational costs

We now examine the computational cost scaling of the conventional TD or FD matched filter procedure as compared with LLOID. For convenience, Table 3.2 provides a review of the notation that we will need in this section.

3.2.3.1 Conventional TD method

The conventional, direct TD method consists of a bank of FIR filters, or sliding-window dot products. If there are M templates, each N samples in length, then each filter requires MN multiplications and additions per sample, or, at a sample rate f^0 ,

$$2MNf^0 \text{ FLOPS.} \quad (3.9)$$

3.2.3.2 Conventional FD method

The most common FD method is known as the overlap-save algorithm, described in [Press et al. \(2007b\)](#). It entails splitting the input into blocks of D samples, $D > N$, each block overlapping the previous one by $D - N$ samples. For each block, the algorithm computes the forward FFT of the data and each of the templates, multiplies them, and then computes the reverse FFT.

Modern implementations of the FFT, such as the ubiquitous `fftw`, require about $2D \lg D$ operations to evaluate a real transform of size D ([Johnson & Frigo, 2007](#)). Including the forward

transform of the data and M reverse transforms for each of the templates, the FFT costs $2(M + 1)D \lg D$ operations per block. The multiplication of the transforms adds a further $2MD$ operations per block. Since each block produces $D - N$ usable samples of output, the overlap-save method requires

$$f^0 \cdot \frac{2(M + 1) \lg D + 2M}{1 - N/D} \text{ FLOPS.} \quad (3.10)$$

In the limit of many templates, $M \gg 1$, we can neglect the cost of the forward transform of the data and of the multiplication of the transforms. The computational cost will reach an optimum at some large but finite FFT block size $D \gg N$. In this limit, the FD method costs $\approx 2f^0 M \lg D$ FLOPS.

By adjusting the FFT block size, it is possible to achieve low latency with FD convolution, but the computational cost grows rapidly as the latency in samples ($D - N$) decreases. It is easy to show that in the limit of many templates and long templates, $M, \lg N \gg 1$, the computational cost scales as

$$\left(1 + \frac{\text{template length}}{\text{latency}}\right) (2f^0 M \lg N).$$

3.2.3.3 LLOID method

For time slice s , the LLOID method requires $2N^s L^s f^s$ FLOPS to evaluate the orthogonal filters, $2ML^s f^s$ FLOPS to apply the linear transformation from the L^s basis templates to the M time-sliced templates, and Mf^s FLOPS to add the resultant partial SNR stream.

The computational cost of the decimation of the detector data is a little bit more subtle. Decimation is achieved by applying an FIR anti-aliasing filter and then downsampling, or deleting samples in order to reduce the sample rate from f^{s-1} to f^s . Naively, an anti-aliasing filter with $(f^{s-1}/f^s)N^\downarrow$ coefficients should demand $2N^\downarrow (f^{s-1})^2/f^s$ FLOPS. However, it is necessary to evaluate the anti-aliasing filter only for the fraction f^s/f^{s-1} of the samples that will not be deleted. Consequently, an efficient decimator requires only $2N^\downarrow f^{s-1}$ FLOPS. (One common realization is an ingenious structure called a *polyphase decimator*, described in Chapter 1 of [Jovanovic-Dolecek \(2002\)](#).)

The story is similar for the interpolation filters used to match the sample rates of the partial SNR streams. Interpolation of a data stream from a sample rate f^s to f^{s-1} consists of inserting zeros between the samples of the original stream, and then applying a low-pass filter with

$(f^{s-1}/f^s)N^\uparrow$ coefficients. The low-pass filter requires $2MN^\uparrow(f^{s-1})^2/f^s$ FLOPS. However, by taking advantage of the fact that by construction a fraction, f^s/f^{s-1} , of the samples are zero, it is possible to build an efficient interpolator that requires only $2MN^\uparrow f^{s-1}$ FLOPS. (Again, see [Jovanovic-Dolecek \(2002\)](#) for a discussion of *polyphase interpolation*.)

Taking into account the decimation of the detector data, the orthogonal FIR filters, the reconstruction of the time-sliced templates, the interpolation of SNR from previous time slices, and the accumulation of SNR, in total the LLOID algorithm requires

$$\sum_{s=0}^{S-1} (2N^s L^s + 2ML^s + M) f^s + 2 \sum_{f^s \in \{f^k : 0 < k < S\}} (N^\downarrow f^0 + MN^\uparrow f^{s-1}) \quad (3.11)$$

FLOPS. The second sum is carried out over the set of distinct sample rates (except for the base sample rate) rather than over the time slices themselves, as we have found that it is sometimes desirable to place multiple adjacent time slices at the same sample rate in order to keep the size of the matrices that enter the SVD manageable. Here we have assumed that the decimation filters are connected in parallel, converting from the base sample rate f^0 to each of the time slice sample rates f^1, f^2, \dots , and that the interpolation filters are connected in cascade fashion with each interpolation filter stepping from the sample rate of one time slice to the next.

We can simplify this expression quite a bit by taking some limits that arise from sensible filter design. In the limit of many templates, the cost of the decimation filters is negligible as compared to the cost of the interpolation filters. Typically, we will design the interpolation filters with $N^\uparrow \lesssim L^s$ so that the interpolation cost itself is negligible compared with the reconstruction cost. Finally, if the number of basis templates per time slice L^s is not too small, the reconstruction cost dominates over the cost of accumulating the partial SNR. In these limits, the cost of LLOID is dominated by the basis filters themselves and the reconstruction, totaling $2 \sum_{s=0}^{S-1} f^s L^s (N^s + M)$ FLOPS.

3.2.3.4 Speedup of LLOID relative to TD method

If the cost of the basis filters dominates, and the frequency of the templates evolves slowly enough in time, then we can use the time-frequency relationship of Equation (3.1) to estimate the speedup

relative to the conventional, direct TD method. The reduction in FLOPS is approximately

$$\frac{2 \sum_{s=0}^{S-1} f^s L^s N^s}{2MNf^0} \approx \frac{\alpha}{(t_{\text{low}} - t_{\text{high}})(f^0)^2} \int_{t_{\text{low}}}^{t_{\text{high}}} (2f(t))^2 dt = \frac{16\alpha (t_{\text{low}}f^2(t_{\text{low}}) - t_{\text{high}}f^2(t_{\text{high}}))}{(f^0)^2 (t_{\text{low}} - t_{\text{high}})}, \quad (3.12)$$

where $\alpha \approx L^s/M$ is the rank reduction factor, or ratio between the number of basis templates and the number of templates. This approximation assumes that the frequency of the signal is evolving very slowly so that we can approximate the time slice sample rate as twice the instantaneous GW frequency, $f^s \approx 2f(t)$, and the number of samples in the decimated time slice as the sample rate times an infinitesimally short time interval, $N^s \approx 2f(t) dt$. The integral is evaluated using the power-law form of $f(t)$ from Equation (3.1). Substituting approximate values for a template bank designed for component masses around $(1.4, 1.4) M_\odot$, $\alpha \approx 10^{-2}$, $t_{\text{low}} = 10^3$ s, $f_{\text{low}} = 10^1$ Hz, $f_{\text{high}} = f_{\text{ISCO}} \approx 1570$ Hz, $f^0 = 2f_{\text{ISCO}}$, and $t_{\text{high}} = f_{\text{ISCO}}^{-1}$, we find from Equation (3.12) that the LLOID method requires only $\sim 10^{-6}$ times as many FLOPS as the conventional TD method.

3.3 Implementation

In this section we describe an implementation of the LLOID method described in Section 3.2 suitable for rapid GW searches for CBCs. The LLOID method requires several computations that can be completed before the analysis is underway. Thus, we divide the procedure into an offline planning stage and an online, low-latency filtering stage. The offline stage can be done before the analysis is started and updated asynchronously, whereas the online stage must keep up with the detector output and produce search results as rapidly as possible. In the next two subsections we describe what these stages entail.

3.3.1 Planning stage

The planning stage begins with choosing templates that cover the space of source parameters with a hexagonal grid (Cokelaer, 2007) in order to satisfy a minimal match criterion. This assures a prescribed maximum loss in SNR for signals whose parameters do not lie on the hexagonal grid. Next, the grid is partitioned into groups of neighbors called *sub-banks* that are appropriately sized so that each sub-bank can be efficiently handled by a single computer. Each sub-bank contains

templates of comparable chirp mass, and therefore similar time-frequency evolution. Dividing the source parameter space into smaller sub-banks also reduces the offline cost of the SVD and is the approach considered in Cannon et al. (2010). Next, we choose time slice boundaries as in Equation (3.6) such that all of the templates within a sub-bank are sub-critically sampled at progressively lower sample rates. For each time slice, the templates are downsampled to the appropriate sample rate. Finally, the SVD is applied to each time slice in the sub-bank in order to produce a set of orthonormal basis templates and a reconstruction matrix that maps them back to the original templates as described in Equation (3.7). The downsampled basis templates, the reconstruction matrix, and the time slice boundaries are all saved to disk.

3.3.2 Filtering stage

The LLOID algorithm is amenable to latency-free, real-time implementation. However, a real-time search pipeline would require integration directly into the data acquisition and storage systems of the LIGO observatories. A slightly more modest goal is to leverage existing low latency, but not real-time, signal processing software in order to implement the LLOID algorithm.

We have implemented a prototype of the low-latency filtering stage using an open-source signal processing environment called GStreamer¹ (version 0.10.33). GStreamer is a vital component of many Linux systems, providing media playback, authoring, and streaming on devices from cell phones to desktop computers to streaming media servers. Given the similarities of GW detector data to audio data it is not surprising that GStreamer is useful for our purpose. GStreamer also provides some useful stock signal processing elements such as resamplers and filters. We have extended the GStreamer framework by developing a library called `gstlal`² that provides elements for GW data analysis.

GStreamer pipelines typically operate with very low (in some consumer applications, imperceptibly low) latency rather than in true real time because signals are partitioned into blocks of samples, or *buffers*. This affords a number of advantages, including amortizing the overhead of passing signals between elements and grouping together sequences of similar operations. However, buffering a signal incurs a latency of up to one buffer length. This latency can be made small at the cost of some additional overhead by making the buffers sufficiently small. In any case,

¹<http://gstreamer.net/>

²<https://www.lsc-group.phys.uwm.edu/daswg/projects/gstlal.html>

buffering is a reasonable strategy for low-latency LIGO data analysis because, as we previously remarked, the LIGO data acquisition system has a granularity of 1/16 s.

3.4 Results

In this section we evaluate the accuracy of the LLOID algorithm using our GStreamer-based implementation described in the previous section. We calculate the measured SNR loss due to the approximations of the LLOID method and our implementation of it. Using a configuration that gives acceptable SNR loss for our chosen set of source parameters, we then compare the computational cost in FLOPS for the direct TD method, the overlap-save FD method, and LLOID.

3.4.1 Setup

We examine the performance of the LLOID algorithm on a small region of compact binary parameter space centered on typical NS–NS masses. We begin by constructing a template bank that spans component masses from 1 to 3 M_{\odot} using a simulated Advanced LIGO noise power spectrum (Shoemaker, 2010)³. Waveforms are generated in the frequency domain in the stationary phase approximation at (post)^{3.5}-Newtonian order in phase and Newtonian order in amplitude (the TaylorF2 waveforms described in Buonanno et al., 2009). Templates are truncated at 10 Hz, where the projected sensitivity of Advanced LIGO is interrupted by the “seismic wall.” This results in a grid of 98 544 points, or $2 \times 98\,544 = 197\,088$ templates. Then we create sub-banks by partitioning the parameter space by chirp mass. Figure 3.5 illustrates this procedure. We concentrate on a sub-bank with 657 points with chirp masses between 1.1955 and 1.2045 M_{\odot} , or $2 \times 657 = 1314$ templates. With this sub-bank we are able to construct an efficient time slice decomposition that consists of 11 time slices with sample rates of 32–4096 Hz, summarized in Table 3.3. We use this sub-bank and decomposition for the remainder of this section.

³<http://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T0900288&version=3>

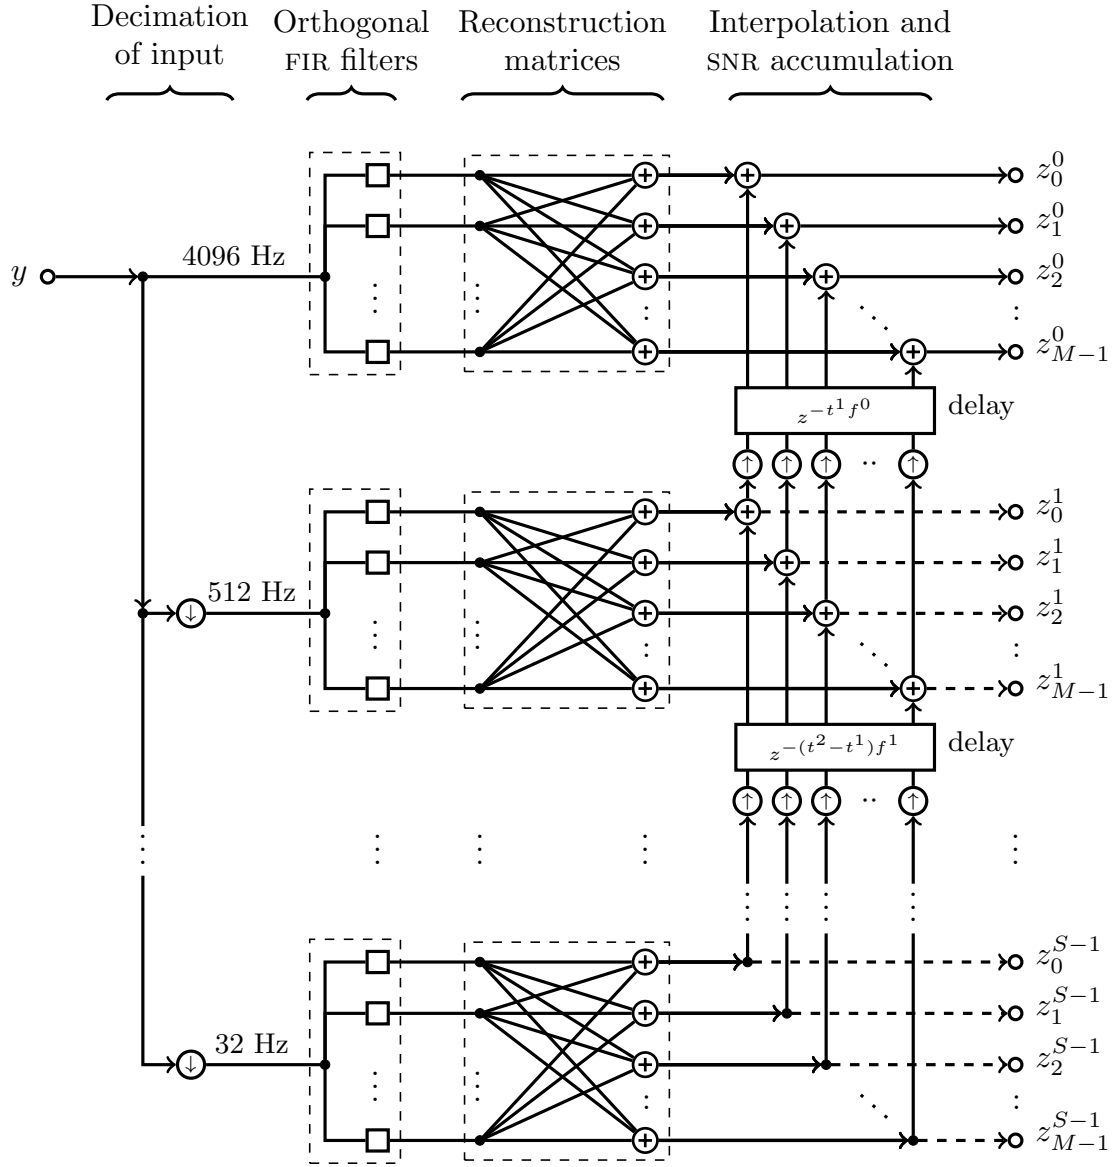


Figure 3.4 Schematic of LLOID pipeline illustrating signal flow. Circles with arrows represent interpolation \Uparrow or decimation \Downarrow . Circles with plus signs represent summing junctions \oplus . Squares \square stand for FIR filters. Sample rate decreases from the top of the diagram to the bottom. In this diagram, each time slice contains three FIR filters that are linearly combined to produce four output channels. In a typical pipeline, the number of FIR filters is much less than the number of output channels.




Figure 3.5 Source parameters selected for sub-bank used in this case study, consisting of component masses m_1 and m_2 , between 1 and 3 M_\odot , and chirp masses \mathcal{M} between 1.1955 and 1.2045 M_\odot .

Table 3.3 Filter design for a sub-bank of 1314 templates.

ableline		f^s	$[t^s, t^{s+1})$	N^s	$-\log_{10}(1-\text{SVD tolerance})$					
		(Hz)	(s)		1	2	3	4	5	6
		4096	$[0, 0.5)$	2048	1	4	6	8	10	14
		512	$[0.5, 4.5)$	2048	2	6	8	10	12	16
		256	$[4.5, 12.5)$	2048	2	6	8	10	12	15
		128	$[12.5, 76.5)$	8192	6	20	25	28	30	32
		64	$[76.5, 140.5)$	4096	1	8	15	18	20	22
		64	$[140.5, 268.5)$	8192	1	7	21	25	28	30
		64	$[268.5, 396.5)$	8192	1	1	15	20	23	25
		32	$[396.5, 460.5)$	2048	1	1	3	9	12	14
		32	$[460.5, 588.5)$	4096	1	1	7	16	18	21
		32	$[588.5, 844.5)$	8192	1	1	8	26	30	33
		32	$[844.5, 1100.5)$	8192	1	1	1	12	20	23

Note. — From left to right, this table shows the sample rate, time interval, number of samples, and number of orthogonal templates for each time slice. We vary SVD tolerance from $(1 - 10^{-1})$ to $(1 - 10^{-6})$.

3.4.2 Measured SNR loss

The SNR loss is to be compared with the mismatch of 0.03 that arises from the discreteness of the template bank designed for a minimal match of 0.97. We will consider an acceptable target SNR loss to be a factor of 10 smaller than this, that is, no more than 0.003.

We expect two main contributions to the SNR loss to arise in our implementation of the LLOID algorithm. The first is the SNR loss due to the truncation of the SVD at $L^s < M$ basis templates. As remarked upon in Cannon et al. (2010) and Section 3.2.2.2, this effect is measured by the SVD tolerance. The second comes from the limited bandwidth of the interpolation filters used to match the sample rates of the partial SNR streams. The maximum possible bandwidth is determined by the length of the filter, N^\uparrow . SNR loss could also arise if the combination of both the decimation filters and the interpolation filters reduces their bandwidth measurably, if the decimation and interpolation filters do not have perfectly uniform phase response, or if there is an unintended subsample time delay at any stage.

To measure the accuracy of our GStreamer implementation of LLOID including all of the above potential sources of SNR loss, we conducted impulse response tests. The GStreamer pipeline was presented with an input consisting of a unit impulse. By recording the outputs, we can effectively “play back” the templates. These impulse responses will be similar, but not identical, to the original, nominal templates. By taking the inner product between the impulse responses for each output channel with the corresponding nominal template, we can gauge exactly how much SNR is lost due to the approximations in the LLOID algorithm and any of the technical imperfections mentioned above. We call one minus this dot product the *mismatch* relative to the nominal template.

The two adjustable parameters that affect performance and mismatch the most are the SVD tolerance and the length of the interpolation filter. The length of the decimation filter affects mismatch as well, but has very little impact on performance.

Effect of SVD tolerance We studied how the SVD tolerance affected SNR loss by holding $N^\downarrow = N^\uparrow = 192$ fixed as we varied the SVD tolerance from $(1 - 10^{-1})$ to $(1 - 10^{-6})$. The minima, maxima, and median mismatches are shown as functions of SVD tolerance in Figure 3.6(a). As the SVD tolerance increases toward 1, the SVD becomes an exact matrix factorization, but the computational cost increases as the number of basis filters increases. The conditions presented

here are more complicated than in the original work (Cannon et al., 2010) due to the inclusion of the time-sliced templates and interpolation, though we still see that the average mismatch is approximately proportional to the SVD tolerance down to $(1 - 10^{-4})$. However, as the SVD tolerance becomes even higher, the median mismatch seems to saturate at around 2×10^{-4} . This could be the effect of the interpolation, or an unintended technical imperfection that we did not model or expect. However, this is still an order of magnitude below our target mismatch of 0.003. We find that an SVD tolerance of $(1 - 10^{-4})$ is adequate to achieve our target SNR loss.

Effect of interpolation filter length Next, keeping the SVD tolerance fixed at $(1 - 10^{-6})$ and the length of the decimation filter fixed at $N^\downarrow = 192$, we studied the impact of the length N^\uparrow of the interpolation filter on mismatch. We use GStreamer’s stock `audioresample` element, which provides an FIR decimation filter with a Kaiser-windowed sinc function kernel. The mismatch as a function of N^\uparrow is shown in Figure 3.6(b). The mismatch saturates at $\sim 2 \times 10^{-4}$ with $N^\uparrow = 64$. We find that a filter length of 16 is sufficient to meet our target mismatch of 0.003.

Having selected an SVD tolerance of $(1 - 10^{-4})$ and $N^\uparrow = 16$, we found that we could reduce N^\downarrow to 48 without exceeding a median mismatch of 0.003.

We found that careful design of the decimation and interpolation stages made a crucial difference in terms of computational overhead. Connecting the interpolation filters in cascade fashion rather than in parallel resulted in a significant speedup. Also, only the shortest interpolation filters that met our maximum mismatch constraint resulted in a sub-dominant contribution to the overall cost. There is possibly further room for optimization beyond minimizing N^\uparrow . We could design custom decimation and interpolation filters, or we could tune these filters separately for each time slice.

3.4.3 Other potential sources of SNR loss

One possible source of SNR loss for which we have not accounted is the leakage of sharp spectral features in the detector’s noise spectrum due to the short durations of the time slices. In the LLOID algorithm, as with many other GW search methods, whitening is treated as an entirely separate data conditioning stage. In this chapter, we assume that the input to the filter bank is already whitened, having been passed through a filter that flattens and normalizes its spectrum. We elected to omit a detailed description of the whitening procedure since the focus here is on the

implementation of a scalable inspiral filter bank.

However, the inspiral templates themselves consist of the GW time series convolved with the impulse response of the whitening filter. As a consequence, the LLOID algorithm must faithfully replicate the effect of the whitening filter. Since in practice the noise spectra of ground-based GW detectors contain both high- Q lines at mechanical, electronic, and control resonances and a very sharp rolloff at the seismic wall, the frequency response of the LLOID filter bank must contain both high- Q notches and a very abrupt high-pass filter. FIR filters with rapidly varying frequency responses tend to have long impulse responses and many coefficients. Since the LLOID basis filters have, by design, short impulse responses and very few coefficients, one might be concerned about spectral leakage contaminating the frequency response of the LLOID filter bank.

The usual statement of the famous Nyquist–Shannon theorem, stated below as Theorem 1, has a natural dual, Theorem 2, that addresses the frequency resolution that can be achieved with an FIR filter of a given length.

Theorem 1. (After [Oppenheim et al., 1997](#), p. 518) *Let $x(t)$ be a band-limited signal with continuous Fourier transform $\tilde{x}(f)$ such that $\tilde{x}(f') = 0 \forall f' : |f'| > f_M$. Then, $x(t)$ is uniquely determined by its discrete samples $x(n/f^0)$, $n = 0, \pm 1, \pm 2, \dots$, if $f^0 > 2f_M$.*

Theorem 2. *Let $x(t)$ be a compactly supported signal such that $x(t') = 0 \forall t' : |t'| > t_M$. Then its continuous Fourier transform $\tilde{x}(f)$ is uniquely determined by the discrete frequency components $\tilde{x}(n \Delta f)$, $n = 0, \pm 1, \pm 2, \dots$, if $\Delta f < 1/(2t_M)$.*

Another way of stating Theorem 2 is that, provided $x(t)$ is nonzero only for $|t| < 1/(2\Delta f)$, the continuous Fourier transform can be reconstructed at any frequency f from a weighted sum of sinc functions centered at each of the discrete frequency components, namely,

$$\tilde{x}(f) \propto \sum_{n=-\infty}^{\infty} \tilde{x}(n \Delta f) \operatorname{sinc} [\pi(f - n \Delta f)/\Delta f].$$

Failure to meet the conditions of this dual of the sampling theorem results in spectral leakage. For a TD signal to capture spectral features that are the size of the central lobe of the sinc function, the signal must have a duration greater than $1/\Delta f$. If the signal $x(t)$ is truncated by sampling it for a shorter duration, then its Fourier transform becomes smeared out; conceptually, power “leaks” out into the side lobes of the sinc functions and washes away sharp spectral features. In

the GW data analysis literature, the synthesis of inspiral matched filters involves a step called *inverse spectrum truncation* (see [Fairhurst, 2009](#), Section VII) that fixes the number of coefficients based on the desired frequency resolution.

In order to effectively flatten a line in the detector’s noise power spectrum, the timescale of the templates must be at least as long as the damping time τ of the line, $\tau = 2Q/\omega_0$, where Q is the quality factor of the line and ω_0 is the central angular frequency. To put this into the context of the sampling theorem, in order to resolve a notch with a particular Q and f_0 , an FIR filter must achieve a frequency resolution of $\Delta f \gtrsim \pi f_0/Q$ and therefore its impulse response must last for at least a time $1/\Delta f = Q/\pi f_0$. For example, in the S6 detector configuration known as “Enhanced LIGO,” the violin modes ([Penn et al., 2007](#)) had $Q \sim 10^5$ and $\omega_0 \sim (2\pi)340 \text{ rad s}^{-1}$, for a coherence time $\tau \sim 10^2 \text{ s}$.

In our example template bank, many of the time slices are much shorter than this. However, in summation the time slices have the same duration as the full templates themselves, and the full templates are much longer than many coherence times of the violin mode. For this reason, we speculate that LLOID should be just as robust to sharp line features as traditional FFT-based searches currently employed in the GW field. Future works must verify this reasonable supposition with numerical experiments, including impulse response studies similar to the ones presented here but with detector noise power spectra containing lines with realistically high quality factors.

There could, in principle, be lines with coherence times many times longer than the template duration. For example, the Q of the violin modes may increase by orders of magnitude in Advanced LIGO ([Strain & Cagnoli, 2006](#)). Also, there are certainly narrow lines that are non-stationary. Both of these cases can be dealt with by preprocessing the data with bandstop filters that attenuate the lines themselves and also conservatively large neighborhoods around them. If such bandstops were implemented as an FIR filter, they could be built into the time slices without any difficulty.

Another way to deal with line features with coherence times much longer than the templates would be to entirely ‘factor’ the whitening out of the LLOID filter bank. Any line features could be notched out in the whitening stage with IIR filters, which can achieve infinitely high Q at just second order. If the detector data were passed through the whitening filter twice, then time-sliced filters need not depend on the detector’s noise power spectral density at all. In such a variation

Table 3.4 Computational cost in FLOPS and latency in seconds of the direct TD method, the overlap-save FD method, and LLOID.

Method	FLOPS (Sub-bank)	Latency (s)	FLOPS (NS-NS)	number of Machines
Direct (TD)	4.9×10^{13}	0	3.8×10^{15}	$\sim 3.8 \times 10^5$
Overlap-save (FD)	5.2×10^8	2×10^3	5.9×10^{10}	~ 5.9
LLOID (theory)	6.6×10^8	0	1.1×10^{11}	~ 11
LLOID (prototype)	(0.9 cores)	0.5	—————	$\gtrsim 10$

Note. — Cost is given for both the sub-bank described in Section 3.4.1 and a full 1–3 M_\odot NS–NS search. The last column gives the approximate number of machines per detector required for a full Advanced LIGO NS–NS search.

on the LLOID method, the basis filters could be calculated from the *weighted* SVD (Gabriel & Zamir, 1979; Jackson, 2003, Chapter 3.6) of the time-sliced templates, using the covariance of the detector noise as a weight matrix.

3.4.4 Lower bounds on computational cost and latency compared to other methods

We are now prepared to offer the estimated computational cost of filtering this sub-bank of templates compared to other methods. We used the results of the previous subsections to set the SVD tolerance to $(1 - 10^{-4})$, the interpolation filter length to 16, and the decimation filter length to 48. Table 3.4 shows the computational cost in FLOPS for the sub-bank we described above. For the overlap-save FD method, an FFT block size of $D = 2N$ is assumed, resulting in a latency of (N/f^0) seconds. Both the FD method and LLOID are five orders of magnitude faster than the conventional, direct TD method. However, the FD method has a latency of over half of an hour, whereas the LLOID method, with suitable design of the decimation and interpolation filters, has no more latency than the direct TD method.

3.4.5 Extrapolation of computational cost to an Advanced LIGO search

Table 3.4 shows that the LLOID method requires 6.6×10^8 FLOPS to cover a sub-bank comprising 657 out of the total 98 544 mass pairs. Assuming that other regions of the parameter space have

similar computational scaling, an entire single-detector search for NS–NS signals in the 1–3 M_{\odot} component mass range could be implemented at $(98\,544/657) \approx 150$ times the cost, or 9.9×10^{10} FLOPS.

We computed the computational cost of a full Advanced LIGO NS–NS search a second way by dividing the entire 1–3 M_{\odot} parameter space into sub-banks of 657 points apiece, performing time slices and SVDs for each sub-bank, and tabulating the number of floating point operations using Equation (3.11). This should be a much more accurate measure because template length varies over the parameter space. Lower chirp mass templates sweep through frequency more slowly and require more computations while higher chirp mass templates are shorter and require fewer computations. Despite these subtleties, this estimate gave us 1.1×10^{11} FLOPS, agreeing with the simple scaling argument above.

Modern (circa 2011) workstations can achieve peak computation rates of up to $\sim 10^{11}$ FLOPS. In practice, we expect that a software implementation of LLOID will reach average computation rates that are perhaps a factor 10 less than this, $\sim 10^{10}$ FLOPS per machine, due to non-floating point tasks including bookkeeping and thread synchronization. Given these considerations, we estimate that a full Advanced LIGO, single-detector, NS–NS search with LLOID will require ~ 10 machines.

By comparison, using the conventional TD method to achieve the same latency costs 4.9×10^{13} FLOPS for this particular sub-bank, and so simply scaling up by the factor of 150 suggests that it would require 7.4×10^{15} FLOPS to search the full parameter space. To account for the varying sample rate and template duration across the parameter space, we can also directly calculate the cost for the full TD method search using Equation (3.9), resulting in 3.8×10^{15} FLOPS, agreeing within an order of magnitude. This would require $\gtrsim 10^5$ current-day machines. Presently, the LIGO Data Grid⁴ consists of only $\sim 10^4$ machines, so direct TD convolution is clearly impractical.

The overlap-save FD method is slightly more efficient than LLOID for this particular sub-bank, requiring 5.2×10^8 FLOPS. The scaling argument projects that a full FD search would require 7.8×10^{10} FLOPS. The direct calculation from Equation (3.10) gives 5.9×10^{10} FLOPS, in order-of-magnitude agreement. In this application, the conventional FD search is scarcely a factor of two faster than LLOID while gaining only 0.3% in SNR, but only at the price of thousands of seconds of latency.

⁴<https://www.lsc-group.phys.uwm.edu/lscdatagrid/>

3.4.6 Measured latency and overhead

Our GStreamer pipeline for measuring impulse responses contained instrumentation that would not be necessary for an actual search, including additional interpolation filters to bring the early-warning outputs back to the full sample rate and additional outputs for recording signals to disk.

We wrote a second, stripped pipeline to evaluate the actual latency and computational overhead. We executed this pipeline on one of the submit machines of the LIGO–Caltech cluster, a Sun Microsystems Sun Fire™ X4600 M2 server with eight quad-core 2.7 GHz AMD Opteron™ 8384 processors. This test consumed $\sim 90\%$ of the capacity of just one out of the 32 cores, maintaining a constant latency of ~ 0.5 s.

The measured overhead is consistent to within an order of magnitude with the lower bound from the FLOPS budget. Additional overhead is possibly dominated by thread synchronization. A carefully optimized GStreamer pipeline or a hand-tuned C implementation of the pipeline might reduce overhead further.

The 0.5 s latency is probably due to buffering and synchronization. The latency might be reduced by carefully tuning buffer lengths at every stage in the pipeline. Even without further refinements, our implementation of the LLOID algorithm has achieved latencies comparable to the LIGO data acquisition system itself.

3.5 Conclusions

We have demonstrated a computationally feasible filtering algorithm for the rapid and even early-warning detection of GWs emitted during the coalescence of NSs and stellar-mass BHs. It is one part of a complicated analysis and observation strategy that will unfortunately have other sources of latency. However, we hope that it will motivate further work to reduce such technical sources of GW observation latency and encourage the possibility of even more rapid EM follow-up observations to catch prompt emission in the advanced detector era.

CBC events may be the progenitors of some short hard GRBs and are expected to be accompanied by a broad spectrum of EM signals. Rapid alerts to the wider astronomical community will improve the chances of detecting an EM counterpart in bands from gamma-rays down to radio. In the Advanced LIGO era, it appears possible to usefully localize a few rare

events prior to the GRB, allowing multi-wavelength observations of prompt emission. More frequently, low-latency alerts will be released after merger but may still yield extended X-ray tails and early on-axis afterglows.

The LLOID method is as fast as conventional FFT-based, FD convolution but allows for latency free, real-time operation. We anticipate requiring $\gtrsim 40$ modern multi-core computers to search for binary NSs using coincident GW data from a four-detector network. In the future, additional computational savings could be achieved by conditionally reconstructing the SNR time series only during times when a composite detection statistic crosses a threshold (Cannon et al., 2011). However, the anticipated required number of computers is well within the current computing capabilities of the LIGO Data Grid.

We have shown a prototype implementation of the LLOID algorithm using GStreamer, an open-source signal processing platform. Although our prototype already achieves latencies of less than one second, further fine tuning may reduce the latency even further. Ultimately the best possible latency would be achieved by tighter integration between data acquisition and analysis with dedicated hardware and software. This could be considered for third-generation detector design. Also possible for third-generation instruments, the LLOID method could provide the input for a dynamic tuning of detector response via the signal recycling mirror to match the frequency of maximum sensitivity to the instantaneous frequency of the GW waveform. This is a challenging technique, but it has the potential for substantial gains in SNR and timing accuracy (Meers et al., 1993).

Although we have demonstrated a computationally feasible statistic for advance detection, we have not yet explored data calibration and whitening, triggering, coincidence, and ranking of GW candidates in a framework that supports early EM follow-up. One might explore these, and also use the time slice decomposition and the SVD to form low-latency signal-based vetoes (e.g., χ^2 statistics) that have been essential for glitch rejection used in previous GW CBC searches. These additional stages may incur some extra overhead, so computing requirements will likely be somewhat higher than our estimates.

Future work must more deeply address sky localization accuracy in a realistic setting as well as observing strategies. Here, we have followed Fairhurst (2009) in estimating the area of 90% localization confidence in terms of timing uncertainties alone, but it would be advantageous to use a galaxy catalog to inform the telescope tiling (Nuttall & Sutton, 2010). Because early detections

will arise from nearby sources, the galaxy catalog technique might be an important ingredient in reducing the fraction of sky that must be imaged. Extensive simulation campaigns incorporating realistic binary merger rates and detector networks will be necessary in order to fully understand the prospects for early-warning detection, localization, and EM follow-up using the techniques we have described.

Acknowledgements

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the NSF and operates under cooperative agreement PHY-0107417. C.H. thanks Ilya Mandel for many discussions about rate estimates and the prospects of early detection, Patrick Brady for countless fruitful conversations about low-latency analysis methods, and John Zweizig for discussions about LIGO data acquisition. N.F. thanks Alessandra Corsi and Larry Price for illuminating discussions on astronomical motivations. L.S. thanks Shaun Hooper for productive conversations on signal processing. This research is supported by the NSF through a Graduate Research Fellowship to L.S. and by the Perimeter Institute for Theoretical Physics through a fellowship to C.H. D.K. is supported from the Max Planck Gesellschaft. M.A.F. is supported by NSF Grant PHY-0855494.

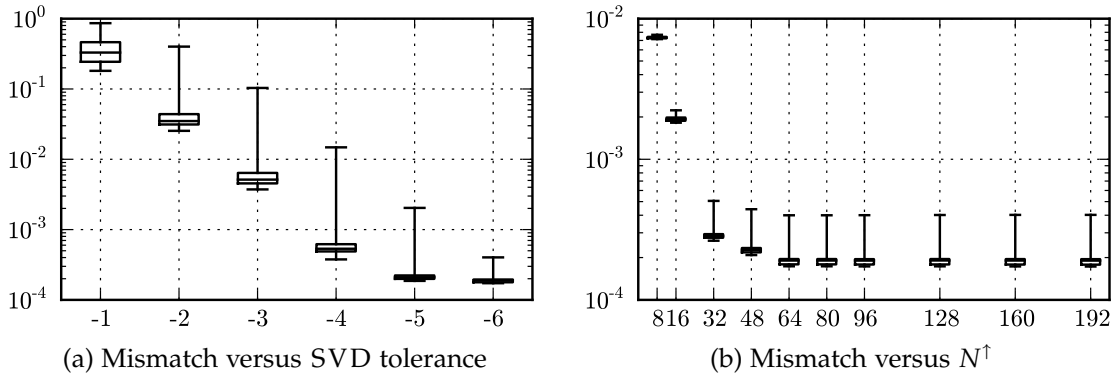


Figure 3.6 Box-and-whisker plot of mismatch between nominal template bank and LLOID measured impulse responses. The upper and lower boundaries of the boxes show the upper and lower quartiles; the lines in the center denote the medians. The whiskers represent the minimum and maximum mismatches over all templates. In (a) the interpolation filter length is held fixed at $N^\dagger = 192$, while the SVD tolerance is varied from $(1 - 10^{-1})$ to $(1 - 10^{-6})$. In (b), the SVD tolerance is fixed at $(1 - 10^{-6})$ while N^\dagger is varied from 8 to 192.

Chapter 4

BAYESTAR: Rapid Bayesian sky localization of BNS mergers

This chapter is reproduced from a work in preparation for Physical Review D. The authors will be Leo P. Singer and Larry R. Price. The introduction of this chapter is reproduced in part from [Singer et al. \(2014\)](#), copyright © 2014 The American Astronomical Society.

We expect this decade to bring the first direct detection of GWs from compact objects. The LIGO and Virgo detectors are being rebuilt with redesigned mirror suspensions, bigger optics, novel optical coatings, and higher laser power ([Harry, 2010](#); [Acernese et al., 2013](#)). In their final configuration, Advanced LIGO and Virgo are expected to reach ~ 10 times further into the local universe than their initial configurations did. The best-understood sources for LIGO and Virgo are BNS mergers. They also offer a multitude of plausible EM counterparts ([Metzger & Berger, 2012](#)) including collimated short-hard gamma-ray bursts (short GRBs; see for example [Paczynski 1986](#); [Eichler et al. 1989](#); [Narayan et al. 1992](#); [Rezzolla et al. 2011](#)) and X-ray/optical afterglows, near-infrared kilonovae (viewable from all angles; [Li & Paczyński, 1998](#); [Barnes & Kasen, 2013b](#), etc.), and late-time radio emission ([Nakar & Piran, 2011](#); [Piran et al., 2013](#)). Yet, typically poor GW localizations of $\gtrsim 100 \text{ deg}^2$ will present formidable challenges to observers hunting for their EM counterparts.

The final Initial LIGO–Virgo observing run pioneered the first accurate, practical parameter estimation and position reconstruction methods for BNS signals. This included a prompt,

semi-coherent, ad hoc analysis (Timing++, designed to work with MBTA; [Abadie et al. 2012a](#)), and the first version of a rigorous Bayesian MCMC analysis (LALINFERENCE; [Aasi et al. 2013b](#)). Though important milestones, there was an undesirable compromise made between accuracy and speed: though the former analysis took only minutes, it produced areas that were typically 20 times larger than the latter, which could take days ([Siderly et al., 2014](#)). To increase the odds of finding a relatively bright but rapidly fading afterglow, one wants localizations that are both prompt *and* accurate, to begin optical searches within minutes to hours of a GW detection. To increase the odds of finding a kilonova, one wants localizations that are reliably available in under one day, to allow as much time as possible for multiple deep exposures. See Figure 4.1 for a timeline of the most promising EM counterparts as compared to the response times of the various steps in the GW analysis.

To that end, in this chapter we develop a rapid and accurate Bayesian sky localization method that takes mere seconds but achieves approximately the same accuracy as the full MCMC analysis. We call this algorithm BAYESian TriAngulation and Rapid localization (BAYESTAR)¹. It differs from existing techniques in several important ways. In the first place, we treat the matched-filter detection pipeline as a measurement system in and of itself, treating the point-parameter estimates that it provides as the experimental input data rather than the full GW time series that is used by the MCMC analysis. This drastically reduces the dimensionality of both the data and the signal model. It also permits us to avoid directly computing the expensive post-Newtonian model waveforms, making the likelihood itself much faster to evaluate. Finally, instead of using MCMC or some similar method for statistical sampling, we make use of a deterministic quadrature scheme. This algorithm is unique in that it bridges the detection and parameter estimation of GW signals, two tasks that have until now involved very different numerical methods and time scales. We expect that BAYESTAR will take on a key role in observing CBC events in both GW and optical channels during the Advanced LIGO era.

¹A pun on the Cylon battleships in the American television series *Battlestar Galactica*. The defining characteristic of the Cylons is that they repeatedly defeat humanity by using their superhuman information-gathering ability to coordinate overwhelming forces. The name also suggests that, like the Cylons, GW detectors may some day rise against us humans.

We do not like to mention the final 'L' in the acronym, because then it would be pronounced BAYESTARL, which sounds stupid.




Figure 4.1 Rough timeline of compact binary merger electromagnetic emissions in relation to the timescale of the Advanced LIGO/Virgo analysis described in this thesis. The time axis measures seconds after the merger.

4.1 Bayesian probability and parameter estimation

In the Bayesian framework, parameters are inferred from the GW data by forming the posterior distribution, $p(\theta|\mathbf{Y})$, which describes the probability of the parameters given the observations. Bayes' rule relates the likelihood $p(\mathbf{Y}|\theta)$ to the posterior $p(\theta|\mathbf{Y})$,

$$p(\theta|\mathbf{Y}) = \frac{p(\mathbf{Y}|\theta)p(\theta)}{p(\mathbf{Y})}, \quad (4.1)$$

introducing the prior distribution $p(\theta)$ which encapsulates previous information about the parameters (for example, arising from earlier observations or from known physical bounds on the parameters), and the evidence $p(\mathbf{Y})$. In parameter estimation problems such as those we are concerned with in this chapter, the evidence serves as nothing more than a normalization factor. However, when comparing two models with different numbers of parameters, the ratio of the evidences quantifies the relative parsimony of the two models, serving as a precise form of Occam's razor.

The parameters that describe a BNS merger signal are listed in Equation (2.4). The choice of prior is determined by one's astrophysical assumptions, but during LIGO's sixth science run (S6) when LIGO's Bayesian CBC parameter estimation pipelines were pioneered, the prior was taken to be isotropic in sky location and binary orientation, and uniform in volume, arrival time, and the component masses (Aasi et al., 2013b).

In Bayesian inference, although it is often easy to write down the likelihood or even the full posterior in closed form, usually one is interested in only a subset β of all of the model's parameters, the others, λ , being nuisance parameters. In this case, we integrate away the nuisance

parameters, forming the marginal posterior

$$p(\boldsymbol{\beta}|\mathbf{Y}) = \int \frac{p(\mathbf{Y}|\boldsymbol{\beta}, \boldsymbol{\lambda})p(\boldsymbol{\beta}, \boldsymbol{\lambda})}{p(\mathbf{Y})} d\boldsymbol{\lambda}, \quad (4.2)$$

with $\boldsymbol{\theta} = (\boldsymbol{\beta}, \boldsymbol{\lambda})$. For instance, for the purpose of locating a GW source in the sky, all parameters (distance, time, orientation, masses, and spins) except for (α, δ) are nuisance parameters.

Bayesian parameter estimation has many advantages, including broad generality and the ability to make probabilistically meaningful statements even with very low SNR measurements. However, in problems of even modest complexity, the marginalization step involves many-dimensional, ill-behaved integrals. The powerful MCMC integration technique has become almost synonymous with Bayesian inference. Though powerful, MCMC is inherently non-deterministic and resistant to parallelization, as well as (at least historically) slow. With the most sophisticated CBC parameter estimation codes, it still takes days to process a single event. This delay is undesirable for planning targeted EM follow-up searches of LIGO events.

In what follows, we describe a complementary rapid parameter estimation scheme that can produce reliable positions estimates within minutes of a detection. We can even use our scheme to speed up the full MCMC analysis and make the refined parameter estimates available more quickly. The key difference is that we start not from the GW signal itself, but from the point parameter estimates from the detection. By harnessing the detection pipeline in this manner, we arrive at a simpler Bayesian problem that is amenable to straightforward, deterministic, numerical quadrature.

There are many practical advantages of doing so. For one, there are difficulties in synchronously gathering together the calibrated GW strain data, auxiliary instrument channels, and data quality vetoes from all of the sites. The data consumed by the real-time detection pipeline are not necessarily final. Longer-running follow-up analyses can benefit from offline calibration, whereas the rapid sky localization need not re-analyze the online data. Moreover, the dimensionality of the problem is greatly reduced, and the problem becomes computationally easier. Finally, by breaking free of the MCMC framework, the results are much easier to use for pointing telescopes. With MCMC algorithms, it is often desirable to bin or interpolate the cloud of sample points to provide a smooth, high-resolution representation of the probability distribution. Reliable post-processing of MCMC chains often relies on clustering and kernel density estimation, both of which prohibit

very large numbers of samples due to rapid growth of computational cost. With BAYESTAR, the natural form of the result is an adaptively sampled mesh with high resolution only where it is needed. The output from BAYESTAR is therefore extremely convenient for packaging into a Flexible Image Transport System (FITS) file for transmission with GW alerts (see Appendix C for details).

4.2 The BAYESTAR likelihood

For the purpose of rapid sky localization, we assume that we do not have access to the GW data \mathbf{Y} itself, and that our only contact with it is through the ML parameter estimates $\{\{\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i\}_i, \hat{\theta}_{\text{in}}\}$. Although this is a significant departure from conventional GW parameter estimation techniques, we can still apply the full Bayesian machinery of Equation (4.2) to compute a posterior distribution for the sky location.

The relevant likelihood is now the probability of the ML estimates, conditioned on the true parameter values, and marginalized over all possible GW observations:

$$p(\{\hat{\theta}_i\}_i, \hat{\theta}_{\text{in}} | \theta) \propto \int_{\mathbf{Y} | \text{MLE}(\mathbf{Y}) = \{\{\hat{\theta}_i\}_i, \hat{\theta}_{\text{in}}\}} p(\mathbf{Y} | \theta) p(\theta) d\mathbf{Y}. \quad (4.3)$$

Although we may not be able to evaluate this equation directly, with some educated guesses we can create a likelihood that has many properties in common with it. Any valid approximate likelihood must have the same Fisher matrix as shown in Equation (2.37). It must also have the same limiting behavior: it should be periodic in the phase error $\hat{\gamma}_i$ and go to zero as $\hat{\tau}_i \rightarrow \pm\infty$, $\hat{\rho}_i \rightarrow 0$, or $\hat{\rho}_i \rightarrow \infty$. Additionally, when $\hat{\tau}_i = 0$, the distribution of $\hat{\rho}_i^2$ should reduce to a noncentral χ^2 distribution with two degrees of freedom, centered about ρ_i^2 , because the complex matched filter time series $z_i(t)$ is Gaussian (under the ideal assumption the GW strain time series is Gaussian).

These conditions could be satisfied by realizing a multivariate Gaussian distribution with covariance matrix $\Sigma = \mathcal{I}^\top$, and then replacing individual quadratic terms in the exponent of the form $-\tilde{\theta}^2/2$ with $\cos \tilde{\theta}$.

Another way is to plug the signal model from Equation (2.11) *evaluated at the ML parameter*

estimates into the single-detector likelihood in Equation (2.20):

$$p(\hat{\theta}_i | \theta) := p(Y_i(\omega) = X_i(\omega; \hat{\theta}_i) | \theta) \\ \propto \exp \left[-\frac{1}{2} \int_0^\infty \frac{\left| \frac{\hat{\rho}_i}{\sigma_i(\hat{\theta}_{\text{in}})} e^{i(\hat{\gamma}_i - \omega \hat{\tau}_i)} H(\omega; \hat{\theta}_{\text{in}}) - \frac{\rho_i}{\sigma_i(\theta_{\text{in}})} e^{i(\gamma_i - \omega \tau_i)} H(\omega; \theta_{\text{in}}) \right|^2}{S_i(\omega)} d\omega \right].$$

If we assume that $\hat{\theta}_{\text{in}} = \theta_{\text{in}}$, then this reduces to

$$p(\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i | \rho_i, \gamma_i, \tau_i) \propto \exp \left[-\frac{1}{2} \hat{\rho}_i^2 - \frac{1}{2} \rho_i^2 + \hat{\rho}_i \rho_i \Re \left\{ e^{i\tilde{\gamma}_i} a_i^*(\tilde{\tau}_i) \right\} \right], \quad (4.4)$$

with $\tilde{\gamma}_i = \hat{\gamma}_i - \gamma_i$, $\tilde{\tau}_i = \hat{\tau}_i - \tau_i$, and the template's autocorrelation function $a_i(t; \theta_{\text{in}})$ defined as

$$a_i(t; \theta_{\text{in}}) := \frac{1}{\sigma_i^2(\theta_{\text{in}})} \int_0^\infty \frac{|H(\omega; \theta_{\text{in}})|^2}{S_i(\omega)} e^{i\omega t} d\omega. \quad (4.5)$$

To assemble the joint likelihood for the whole network, we just form the product from the individual detectors:

$$p(\{\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i\}_i | \{\rho_i, \gamma_i, \tau_i\}_i) \propto \exp \left[-\frac{1}{2} \sum_i \hat{\rho}_i^2 - \frac{1}{2} \sum_i \rho_i^2 + \sum_i \hat{\rho}_i \rho_i \Re \left\{ e^{i\tilde{\gamma}_i} a_i^*(\tilde{\tau}_i) \right\} \right]. \quad (4.6)$$

4.3 Properties

First, observe that at the true parameter values, $\hat{\theta}_i = \theta_i$, the logarithms of Equation (4.4) and Equation (2.20) have the same Jacobian. This is because the derivatives of the autocorrelation function are

$$a^{(n)}(t) = i^n \overline{\omega^n},$$

with $\overline{\omega^n}$ defined in Equation (2.24). For example, the first few derivatives are

$$a(0) = 1, \quad \dot{a}(0) = i\overline{\omega}, \quad \ddot{a}(0) = -\overline{\omega^2}.$$

Using Equation (2.22), we can compute the Fisher matrix elements for the autocorrelation

likelihood given by Equation (4.4), with detector subscript suppressed:

$$\mathcal{I}_{\rho\rho} = 1,$$

$$\mathcal{I}_{\rho\gamma} = 0,$$

$$\mathcal{I}_{\rho\tau} = 0,$$

$$\mathcal{I}_{\gamma\gamma} = \rho^2 \int_{-T}^T |a(t)|^2 w(t; \rho) dt, \quad (4.7)$$

$$\mathcal{I}_{\tau\tau} = -\rho^2 \int_{-T}^T \Re[a^*(t)\ddot{a}(t)] w(t; \rho) dt, \quad (4.8)$$

$$\mathcal{I}_{\gamma\tau} = -\rho^2 \int_{-T}^T \Im[a^*(t)\dot{a}(t)] w(t; \rho) dt, \quad (4.9)$$

where

$$w(t; \rho) = \frac{\exp\left[\frac{\rho^2}{4}|a(t)|^2\right] \left(I_0\left[\frac{\rho^2}{4}|a(t)|^2\right] + I_1\left[\frac{\rho^2}{4}|a(t)|^2\right]\right)}{2 \int_{-T}^T \exp\left[\frac{\rho^2}{4}|a(t')|^2\right] I_0\left[\frac{\rho^2}{4}|a(t')|^2\right] dt'}. \quad (4.10)$$

The notation I_k denotes a modified Bessel function of the first kind. Matrix elements that are not listed have values that are implied by the symmetry of the Fisher matrix. Note that the minus signs are correct but a little confusing: despite them, $\mathcal{I}_{\gamma\gamma}, \mathcal{I}_{\tau\tau} \geq 0$ and $\mathcal{I}_{\gamma\tau} \leq 0$. The time integration limits $[-T, T]$ correspond to a flat prior on arrival time, or a time coincidence window between detectors.

We can show that the weighting function $w(t; \rho)$ approaches a Dirac delta function as $\rho \rightarrow \infty$, so that the Fisher matrix for the autocorrelation likelihood approaches the Fisher matrix for the full GW data, Equation (2.27), as $\rho \rightarrow \infty$. The Bessel functions asymptotically approach:

$$I_0(x), I_1(x) \rightarrow \frac{e^x}{\sqrt{2\pi x}} \text{ as } x \rightarrow \infty.$$

For large ρ , the exponential dominates and we can write:

$$w(t; \rho) \rightarrow \frac{\exp\left[\frac{\rho^2}{2}|a(t)|^2\right]}{\int_{-T}^T \exp\left[\frac{\rho^2}{2}|a(t')|^2\right] dt'} \text{ as } \rho \rightarrow \infty.$$

The Taylor expansion of $|a(t)|^2$ is

$$|a(t)|^2 = 1 + \frac{1}{2} \left(\frac{\partial^2}{\partial t^2} |a(t)|^2 \Big|_{t=0} \right) t^2 + \mathcal{O}(t^4) = 1 - \omega_{\text{rms}}^2 t^2 + \mathcal{O}(t^4).$$

Substituting, we find that $w(t; \rho)$ approaches a normalized Gaussian distribution:

$$w(t; \rho) \approx \frac{\exp \left[-\frac{1}{2} \rho^2 \omega_{\text{rms}}^2 t^2 \right]}{\int_{-T}^T \exp \left[-\frac{1}{2} \rho^2 \omega_{\text{rms}}^2 (t')^2 \right] dt'}.$$

And finally, because the Dirac delta function may be defined as the limit of a Gaussian, $w(t; \rho) \rightarrow \delta(t)$ as $\rho \rightarrow \infty$.

We can now write the Fisher matrix for the autocorrelation likelihood in a way that makes a comparison to the full signal model explicit. Define:

$$\begin{aligned} \mathcal{I}_{\gamma\gamma} &= \rho^2 \cdot \mathfrak{H}_{\gamma\gamma}(\rho), \\ \mathcal{I}_{\tau\tau} &= \rho^2 \bar{\omega}^2 \cdot \mathfrak{H}_{\tau\tau}(\rho), \\ \mathcal{I}_{\gamma\tau} &= -\rho^2 \bar{\omega} \cdot \mathfrak{H}_{\gamma\tau}(\rho). \end{aligned}$$

Now, the \mathfrak{H}_{ij} ² contain the integrals from Equations (4.7, 4.8, 4.9) and encode the departure of the autocorrelation likelihood from the likelihood of the full data at low SNR. All of the $\mathfrak{H}_{ij}(\rho)$ are sigmoid-type functions that asymptotically approach 1 as $\rho \rightarrow \infty$ (see Figures 4.2 and 4.3). The transition SNR ρ_{crit} is largely the same for all three nontrivial matrix elements, and is determined by the time coincidence window T and the signal bandwidth ω_{rms} .

In the limit of large SNR, our interpretation is that the point estimates $(\hat{\rho}, \hat{\gamma}, \hat{\tau})$ contain all of the information about the underlying extrinsic parameters.

On the other hand, in the low SNR limit, the diminishing value of $\mathfrak{H}_{ij}(\rho)$ reflects the fact that some information is lost when the full data \mathbf{x} is discarded. Concretely, as the prior interval T becomes large compared to $1/\rho\omega_{\text{rms}}$, the ML estimator becomes more and more prone to picking up spurious noise fluctuations far from the true signal. Clearly, when the coincidence window T is kept small as possible, more information is retained in the ML point estimates. Put another

²The Fish(er) factor.




Figure 4.2 CRLB on root-mean square timing uncertainty and phase error, using the likelihood for the full GW data (Equation 2.20; dashed diagonal line) or the autocorrelation likelihood (Equation 4.4; solid lines) with a selection of arrival time priors.

way, if T is small, then the transition SNR ρ_{crit} is also small and fainter signals become useful for parameter estimation. In this way, the BAYESTAR likelihood exhibits the *threshold effect* that is well-known in communication and radar applications (Barankin, 1949; McAulay & Seidman, 1969; McAulay & Hofstetter, 1971).

In the following sections, we describe the numerical implementation of BAYESTAR.

4.4 Prior and problem setup

The detection pipeline supplies a candidate, $\{\{\hat{\rho}_i, \hat{\gamma}_i, \hat{\tau}_i\}_i, \hat{\theta}_{\text{in}}\}$, and discretely sampled noise PSDs, $S_i(\omega_j)$, for all detectors. We compute the GW signal for a source with intrinsic parameters equal to the detection pipeline's estimate, $H(\omega; \hat{\theta}_{\text{in}})$. Then we find the SNR=1 horizon distance $r_{1,i}$ for each detector by numerically integrating Equation (2.10).

We have no explicit prior on the intrinsic parameters; in our analysis they are fixed at their ML estimates, $\hat{\theta}_{\text{in}}$.

The arrival time prior is connected to the origin of the detector coordinate system. Given the Earth-fixed coordinates of the detectors \mathbf{n}_i and the arrival times τ_i , we compute their averages




Figure 4.3 Ratio between Fisher matrix elements (solid: $\hat{\Omega}_{\gamma\gamma}$, dashed: $\hat{\Omega}_{\gamma\tau}$, dotted: $\hat{\Omega}_{\tau\tau}$) for the autocorrelation likelihood and the full GW data. Colors correspond to different arrival time priors as in Figure 4.2.

weighted by the timing uncertainty formula:

$$\langle \mathbf{n} \rangle = \frac{\sum_i \frac{\mathbf{n}_i}{(\hat{\rho}_i \omega_{\text{rms},i})^2}}{\sum_i \frac{1}{(\hat{\rho}_i \omega_{\text{rms},i})^2}}, \quad \langle \hat{\tau} \rangle = \frac{\sum_i \frac{\hat{\tau}_i}{(\hat{\rho}_i \omega_{\text{rms},i})^2}}{\sum_i \frac{1}{(\hat{\rho}_i \omega_{\text{rms},i})^2}}.$$

Then, we subtract these means:

$$\mathbf{n}_i \leftarrow \mathbf{n}_i - \langle \mathbf{n} \rangle, \quad \hat{\tau}_i \leftarrow \hat{\tau}_i - \langle \hat{\tau} \rangle.$$

In these coordinates, now relative to the weighted detector array barycenter, the arrival time prior is uniform in $-T \leq t \leq T$, with $T = \max_i |\mathbf{n}_i|/c + 5$ ms.

The distance prior is a user-selected power of distance,

$$p(r) \propto \begin{cases} r^m & \text{if } r_{\min} < r < r_{\max} \\ 0 & \text{otherwise,} \end{cases}$$

where $m = 2$ for a prior that is uniform in volume, and $m = -1$ for a prior that is uniform in the

logarithm of the distance. If a distance prior is not specified, the default is uniform in volume out to the maximum SNR=4 horizon distance:

$$m = 2, \quad r_{\min} = 0, \quad r_{\max} = \frac{1}{4} \max_i r_{1,i}.$$

Finally, the prior is uniform in $-1 \leq \cos \iota \leq 1$ and $0 \leq \psi < 2\pi$.

We compute the autocorrelation function for each detector from $t = 0$ to $t = T$ at intervals of $\Delta t = 1/f_s$, where f_s is the smallest power of two that is greater than or equal to the Nyquist rate. Because BNS signals typically terminate at about 1500 Hz, a typical value for Δt is $(4096 \text{ Hz})^{-1}$. We use a pruned FFT because for BNS systems, the GW signal remains in LIGO's sensitive band for $\sim 100\text{--}1000$ s, whereas $T \sim 10$ ms.³

4.5 Marginal posterior

The marginal posterior as a function of sky location is

$$f(\alpha, \delta) \propto \int_0^\pi \int_{-1}^1 \int_{-T}^T \int_{r_{\min}}^{r_{\max}} \int_0^{2\pi} \exp \left[-\frac{1}{2} \sum_i \rho_i^2 + \sum_i \hat{\rho}_i \rho_i \Re \left\{ e^{i\tilde{\gamma}_i} a_i^*(\tilde{\tau}_i) \right\} \right] r^m d\phi_c dr dt_\oplus d\cos \iota d\psi. \quad (4.11)$$

To marginalize over the coalescence phase, we can write $\tilde{\gamma}_i = \tilde{\gamma}'_i + 2\phi_c$. Then integrating over ϕ_c and suppressing normalization factors, we get

$$f(\alpha, \delta) \rightarrow \int_0^\pi \int_{-1}^1 \int_{-T}^T \int_{r_{\min}}^{r_{\max}} \exp \left[-\frac{1}{2} \sum_i \rho_i^2 \right] I_0 \left[\left| \sum_i \hat{\rho}_i \rho_i e^{i\tilde{\gamma}'_i} a_i^*(\tilde{\tau}_i) \right| \right] r^m dr dt_\oplus d\cos \iota d\psi. \quad (4.12)$$

In the above equation, we need not distinguish between $\tilde{\gamma}_i$ and $\tilde{\gamma}'_i$ because the likelihood is now invariant under arbitrary phase shifts of all of the detectors' signals.

4.5.1 Integral over angles and time

The integrand is periodic in ψ , so simple Newton–Cotes quadrature over ψ exhibits extremely rapid convergence. We therefore sample the posterior on a regular grid of 10 points from 0 to π .

³See <http://www.fftw.org/pruned.html> for a discussion of methods for computing the first K samples of an FFT of length N .

The integral over $\cos \iota$ converges with just as rapidly with Gauss–Legendre quadrature, so we use a 10-point Gauss–Legendre rule for integration over $\cos \iota$.

We sample t_{\oplus} regularly from $-T$ to T at intervals of Δt . This is typically $\sim 2(10 \text{ ms})(4096 \text{ Hz}) \approx 80$ samples. We use Catmull–Rom cubic splines to interpolate the real and imaginary parts of the autocorrelation functions between samples.

4.5.2 Integral over distance

The innermost integral over r is a little bit more challenging. The distance integrand has both a sharp maximum-likelihood peak and an extended power-law prior-dominated tail. At any values of the other parameters, either the peak or the tail might dominate. If the peak dominates, it may not be successfully resolved by any general-purpose fixed-order quadrature scheme. Any general adaptive integrator may take a long time to find the peak.


We tried two ways to evaluate the distance integral, first with adaptive Gaussian quadrature and second with fixed-order Gaussian quadrature. The second was ultimately faster and just as accurate for the example events that we tested, but we will describe both methods because they were instructive.

4.5.2.1 Adaptive Gaussian quadrature method

The first method was to use adaptive Gaussian quadrature, but with an educated choice of initial subdivisions. The distance integral can be written as

$$\begin{aligned} \mathcal{F} &= \int_{r_{\min}}^{r_{\max}} \exp\left[-\frac{p^2}{r^2}\right] I_0\left[\frac{b}{r}\right] r^m dr \\ &= \int_{r_{\min}}^{r_{\max}} \exp\left[-\frac{p^2}{r^2} + \frac{b}{r}\right] \bar{I}_0\left[\frac{b}{r}\right] r^m dr \end{aligned} \quad (4.13)$$

Figure 4.4 Illustration of initial subdivisions for distance integration scheme. Distance increases from left to right. In the color version, the left-hand tail, the left- and right-hand sides of the maximum likelihood peak, and the right-hand tail, are colored cyan, red, green, and blue, respectively.



where

$$p^2 = \frac{1}{2} r^2 \sum_i \rho_i^2 \quad (4.14)$$

$$b = r \left| \sum_i \hat{\rho}_i \rho_i e^{i\tilde{\gamma}_i} a_i^*(\tilde{\tau}_i) \right| \quad (4.15)$$

$$\bar{I}_0(x) = \exp(-|x|) I_0(x). \quad (4.16)$$

The coefficients p^2 and b are nonnegative and independent of distance. The symbol \bar{I}_0 denotes an exponentially scaled Bessel function. In the limit of large argument, $I_0(|x|) \sim \exp(|x|)/\sqrt{2\pi|x|}$ (Olver et al., 2010; DLMF, 2014)⁴. The scaled Bessel is useful for evaluation on a computer because it has a relatively small range $(0, 1]$ and varies slowly in proportion to $x^{1/2}$. If we neglect both the Bessel function and the r^m prior, then the approximate likelihood $\exp(-p^2/r^2 + b/r)$ is maximized when

$$r = r_0 \equiv 2p^2/b, \quad (4.17)$$

with p and b defined above in Equations (4.14, 4.15). The likelihood takes on a factor η (say, $\eta = 0.01$) of its maximum value when

$$r = r_{\pm} = \left(\frac{1}{r_0} \mp \frac{\sqrt{-\log \eta}}{p} \right)^{-1}. \quad (4.18)$$

We have now identified up to five breakpoints that partition the distance integrand into up to four intervals with quantitatively distinct behavior. These intervals are depicted in Figure 4.4 with distance, r , increasing from left to right. There is a left-hand or small distance tail in which the integrand is small and monotonically increasing, a left- and right-hand side of the maximum likelihood peak, and a right-hand tail in which the integrand is small and monotonically

⁴<http://dlmf.nist.gov/10.40.E1>

decreasing. These breakpoints are:

$$r_{\text{break}} = \left\{ r \in \begin{Bmatrix} r_{\min} \\ r_- \\ r_0 \\ r_+ \\ r_{\max} \end{Bmatrix} : r_{\min} \leq r \leq r_{\max} \right\}. \quad (4.19)$$

We used these breakpoints as initial subdivisions in an adaptive Gaussian quadrature algorithm⁵. This function estimates the integral over each subdivision and each interval's contribution to the total error, then subdivides the interval whose error contribution is largest. Subdivisions continue until a fixed total fractional error is reached. In this way, most integrand evaluations are expended on the most important distance interval, whether that happens to be the tails (when the posterior is dominated by the prior) or the peak (when the posterior is dominated by the observations).

4.5.2.2 Fixed order Gaussian quadrature method

Let's say that we want to evaluate an integral

$$\mathcal{F}(x_1, x_2) = \int_{x_1}^{x_2} f(x') dx'. \quad (4.20)$$

Suppose that this integral is resistant to direct application of standard quadrature techniques, but we have another function $g(x)$ that we can integrate analytically or at least numerically with relative ease,

$$\mathcal{G}(x) = G(x) - G(x_0) = \int_{x_0}^x g(x') dx'. \quad (4.21)$$

There is an *ultimate*, though tautological, change of variables that makes it trivial to evaluate $\mathcal{G}(x)$,

$$\mathcal{G}(x) = \int_0^{\mathcal{G}(x)} dG. \quad (4.22)$$

⁵for instance, GNU Scientific Library (GSL)'s `gsl_integrate_qagp` function, http://www.gnu.org/software/gsl/manual/html_node/QAGP-adaptive-integration-with-known-singular-points.html

If we apply the same change of variables to Equation (4.20), we get

$$\mathcal{F}(x_1, x_2) = \int_{\mathcal{G}(x_1)}^{\mathcal{G}(x_2)} \frac{f(\mathcal{G}^{-1}(G))}{g(\mathcal{G}^{-1}(G))} dG. \quad (4.23)$$

If $f(x)$ and $g(x)$ have sufficiently similar behavior, then their quotient f/g will be better behaved than f itself, and more amenable to any given quadrature technique. When the transformed integral Equation (4.23) is evaluated using Monte Carlo integration, this technique is referred to as importance sampling.

We use the same kind of change of variables with fixed-order Gaussian quadrature. We first find the approximate ML distance r_0 using Equation (4.17). If $r_0 < r_{\min}$ or $r_0 > r_{\max}$, then the ML peak lies outside of the limits of integration and the original integrand behaves like a low-order power law. In this case, we use 10-point Gaussian quadrature to evaluate Equation (4.13).

If, on the other hand, $r_{\min} \leq r_0 \leq r_{\max}$, then the integrand contains a sharp peak that we can smooth out with importance sampling. We first re-scale the distance integral Equation (4.13) so that the peak value of the integrand is ~ 1 :

$$\mathcal{F} = \exp\left(\frac{b^2}{4p^2}\right) r_0^m \int_{r_{\min}}^{r_{\max}} \exp\left[-\left(\frac{p}{r} - \frac{b}{2p}\right)^2\right] \left[\frac{r}{r_0}\right]^m dr. \quad (4.24)$$

The importance sampling integrand $g(r)$ is

$$g(r) = \exp\left[-\left(\frac{p}{r} - \frac{b}{2p}\right)^2\right] \frac{1}{r^2}, \quad (4.25)$$

with definite integral

$$\mathcal{G}(r) = \frac{\sqrt{\pi}}{p} Q\left[\sqrt{2}\left(\frac{p}{r} - \frac{b}{2p}\right)\right], \quad (4.26)$$

where Q is the cumulative distribution function for the upper tail of the standard normal distribution,

$$Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{x'^2}{2}\right] dx'. \quad (4.27)$$

Its inverse is

$$\mathcal{G}^{-1}(G) = \left[\frac{1}{r_0} + \frac{1}{\sqrt{2p}} Q^{-1}\left(\frac{pG}{\sqrt{\pi}}\right)\right]^{-1}. \quad (4.28)$$

The transformed integral is

$$\mathcal{F} = \exp\left(\frac{b^2}{4p^2}\right) r_0^m \int_{\mathcal{G}(r_{\min})}^{\mathcal{G}(r_{\max})} \bar{I}_0 \left[\frac{b}{r}\right] \left[\frac{r}{r_0}\right]^m r^2 \bigg|_{r=\mathcal{G}^{-1}(G)} dG. \quad (4.29)$$

This integral is evaluated numerically with a 10-point Gauss–Legendre rule.

Last, there are a few special cases where Equation (4.13) can be evaluated exactly. If $p^2 = 0$ (which implies $b = 0$), then the integral may be expressed in terms of logarithms. If $b = 0$ but $p^2 \neq 0$, then the integral may be expressed in terms of the exponential integral function,

$$E_n(x) = \int_1^\infty \frac{\exp(-xt)}{t^n} dt. \quad (4.30)$$

The distance integral, therefore, takes either one or ten function evaluations.

4.6 Adaptive mesh refinement

We have explained how we evaluate the marginal posterior at a given sky location. Now we must specify where we choose to evaluate it.

Our sampling of the sky relies completely on the Hierarchical Equal Area isoLatitude Pixelization (HEALPix), a special data structure designed for all-sky maps. Hierarchical Equal Area isoLatitude Pixelization (HEALPix) divides the sky into equal-area pixels. There is a hierarchy of HEALPix resolutions. A HEALPix resolution may be designated by its order N . The $N = 0$ th order or base tiling has 12 pixels. At every successive order, each tile is subdivided into four new tiles. A resolution may also be referred to by the number of subdivisions along each side of the base tiles, $N_{\text{side}} = 2^N$. There are $N^{\text{pix}} = 12N_{\text{side}}^2$ pixels at any given resolution. The HEALPix projection uniquely specifies the coordinates of the center of each pixel by providing a mapping from the resolution and pixel index $(N_{\text{side}}, i_{\text{pix}})$ to right ascension and declination (α, δ) .

We begin by evaluating the posterior probability density at the center of each of the $N_{\text{pix},0} = 3072$ pixels of an $N_{\text{side},0} = 16$ HEALPix grid. At this resolution, each pixel has an area of 13.4 deg^2 . We then rank the pixels by contained probability (assuming constant probability density within a pixel), and subdivide the most probable $N_{\text{pix},0}/4$ pixels into $N_{\text{pix},0}$ new pixels. We then evaluate the posterior again at the pixels that we subdivided, sort again, and repeat seven times,

so that we have evaluated the posterior probability density a total of $8N_{\text{pix},0}$ times. On *most* iterations, we descend one level deeper in HEALPix resolution. The resulting map is a tree structure that consists of a mix of several resolutions. We traverse the tree and flatten it into the highest resolution represented. The highest possible resolution is $N_{\text{side}} = 2^{11}$, with an area of $\approx 10^{-3} \text{ deg}^2$ per pixel.⁶

Within each iteration, all of the marginalization integrals may be evaluated in parallel. In our C language implementation, the inner loop over pixels is parallelized with OpenMP⁷. See Appendix A.5 for source code for BAYESTAR.

4.7 Run time

Since BAYESTAR is designed as one of the final steps in the real-time BNS search, it is important to characterize how long it takes to calculate a sky map. We compiled BAYESTAR with the Intel C Compiler (icc) at the highest architecture-specific optimization setting (-xhost). We timed it under Scientific Linux 6 on a Supermicro 1027GR-TRFT system with dual Intel Xeon E5-2670 CPUs clocked at 2.60GHz, providing a total of 16 cores and capable of executing 32 threads simultaneously (with hyperthreading). In Figure 4.5, we show how long it took to calculate a localization with BAYESTAR as the number of OpenMP threads was varied from 1 to 32. This is a violin plot, a smoothed vertical histogram. The red surface shows run times for a two-detector network (HL) modeled on the first scheduled Advanced LIGO observing run in 2015, and the blue surface shows run times for a three-detector network (HLV) based on the second planned observing run in 2016. These are the two observing scenarios that will be discussed at length in the following chapter.

Several features are apparent. First, at any number of threads, the two configurations have similar run times, although the 2016 events contain a subpopulation of outliers that take about 2.5 times as long as the 2015 events. These are probably due to taking one of the more expensive code branches in the distance integral. Second, the run times decrease proportionally to the number of threads. A slight flattening at 32 threads may be due to increasing use of hyperthreading. Based on experiences running BAYESTAR on the 32-core (64 threads with hyperthreading) cluster login

⁶Although the resulting sky map contains $N_{\text{pix}} \approx 5 \times 10^6$ pixels, at most $\approx 2 \times 10^4$ pixels have distinct values. For the purpose of delivery to observers, therefore, the output is always gzip-compressed with a ratio of $\approx 250 : 1$.

⁷<http://openmp.org/>




Figure 4.5 Violin plot of BAYESTAR run times as the number of OpenMP threads is varied from 1 to 32. The 2015 scenario is shown in red and the 2016 scenario in blue.

machine, we expect the almost ideal parallel speedup to continue on machines with more threads. With just one thread, the BAYESTAR analysis takes 500–1500 s, already orders of magnitude faster than the full parameter estimation. With 32 threads, BAYESTAR takes just 20–60 s. This latency is comparable to the other stages (data aggregation, trigger generation, alert distribution) in the real-time BNS analysis. The 32-thread configuration is representative of how BAYESTAR might be deployed in early Advanced LIGO.⁸

4.8 Case study

We have completed our description of the BAYESTAR algorithm. In the following chapter, we will describe the first results with BAYESTAR, a detailed case study of the sky localization capabilities of the earliest configurations of Advanced LIGO and Virgo. Because it is typically three orders of magnitude faster than the conventional MCMC analysis, for the first time we are able to generate detailed sky maps for thousands of simulated GW sources. With such a large

⁸BAYESTAR has been successfully ported to the Intel’s Many Integrated Core (MIC) architecture and has been tested in a 500 thread configuration on a system with dual Intel Xeon Phi 5110P coprocessors.

sample, we can for the first time unambiguously describe areas and morphologies that should be typical of networks of just two detectors. This case study will be the subject of the next chapter.

Chapter 5

The first two years of EM follow-up with Advanced LIGO and Virgo

LEO P. SINGER¹, LARRY R. PRICE¹, BEN FARR^{2,3}, ALEX L. URBAN⁴, CHRIS PANKOW⁴,
SALVATORE VITALE⁵, JOHN VEITCH^{6,3}, WILL M. FARR³, CHAD HANNA^{7,8}, KIPP CANNON⁹,
TOM DOWNES⁴, PHILIP GRAFF¹⁰, CARL-JOHAN HASTER³, ILYA MANDEL³, TREVOR SIDERY³,
AND ALBERTO VECCHIO³

¹LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

²Department of Physics and Astronomy & Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, USA

³School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT, UK

⁴Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA

⁵Massachusetts Institute of Technology, 185 Albany Street, Cambridge, MA 02139, USA

⁶Nikhef, Science Park 105, Amsterdam 1098XG, The Netherlands

⁷Perimeter Institute for Theoretical Physics, Ontario N2L 2Y5, Canada

⁸The Pennsylvania State University, University Park, PA 16802, USA

⁹Canadian Institute for Theoretical Astrophysics, University of Toronto, Toronto, Ontario, M5S 3H8, Canada

¹⁰NASA Goddard Space Flight Center, Greenbelt, MD, USA

This chapter is reproduced in part from [Singer et al. \(2014\)](#), which was published under the same title in The Astrophysical Journal, copyright © 2014 The American Astronomical Society. My contributions included developing the rapid sky localization code, planning the simulation campaign, running the detection pipeline for the 2016 scenario, running the rapid sky localization code for both scenarios, reducing and analyzing the results, generating most of the tables and figures as well as the data release, and writing about 95% of the text. A.U. developed and ran the pipeline for the 2015 scenario, contributed to the data reduction, and created Figure 5.8. B.F., S.V., J.V., and P.G. ran the stochastic samplers. The other authors contributed to planning, code development, and editing of the paper.

Several planned optical astronomy projects with a range of fields of view and apertures are preparing to pursue optical counterparts of BNS mergers detected by Advanced LIGO and Virgo. These include ZTF (Kulkarni, 2012; Bellm, 2014; Smith et al., 2014b), Pan-STARRS¹, BlackGEM², and LSST (Ivezic et al., 2008), to name a few. Advanced LIGO is scheduled to start taking data in 2015 (Aasi et al., 2013c). Preparations for joint EM and GW observations require a complete understanding of when and how well localized the first GW detections will be. Plausible scenarios for the evolution of the configuration and sensitivity of the worldwide GW detector network as it evolves from 2015 through 2022, as well as rough estimates of sky localization area, are outlined in Aasi et al. (2013c).

To provide a more realistic and complete picture, we have conducted Monte Carlo simulations of the 2015 and 2016 detector network configurations, probing the transition from two to three detectors as Advanced Virgo is scheduled to begin science operation. Prior work has focused on various aspects of position reconstruction with advanced GW detectors (Fairhurst, 2009; Wen & Chen, 2010; Fairhurst, 2011; Vitale & Zanolin, 2011; Rodriguez et al., 2014; Nissanke et al., 2011; Nissanke et al., 2013; Kasliwal & Nissanke, 2014; Grover et al., 2014; Sidery et al., 2014), but ours is the first to bring together a large astrophysically motivated population, an educated guess about the detector commissioning timetable, a realistic SNR distribution, and the Advanced LIGO/Virgo data analysis pipeline itself.

We have simulated hundreds of GW events, recovered them with a real-time detection pipeline, and generated sky maps using both real-time and thorough off-line parameter estimation codes that will be operating in 2015 and beyond. This study contains some of the first results with BAYESTAR, a rapid Bayesian position reconstruction code that will produce accurate sky maps less than a minute after any BNS merger detection. The LALINFERENCE_MCMC (van der Sluys et al., 2008b; Raymond et al., 2009), LALINFERENCE_NEST (Veitch & Vecchio, 2010), and LALINFERENCE_BAMBI (Graff et al., 2012, 2014) stochastic samplers were also used to follow up a subset of detected GW events. Though these analyses are significantly more computationally costly than BAYESTAR, taking hours to days, they can provide improved sky location estimates when the GW signal is very weak in one detector, and also yield not just sky localization but

¹<http://pan-starrs.ifa.hawaii.edu/public/>

²<https://www.astro.ru.nl/wiki/research/blackgemarray>

the full multidimensional probability distribution describing the parameters of a circularized compact binary merger. All four algorithms are part of the LALINFERENCE library ([Aasi et al., 2013b](#)), developed specifically for estimating the parameters of GW sources from ground-based detectors. Together, these analyses will be able to provide sky localizations on time scales that enable searching for all expected electromagnetic counterparts of compact binary mergers (except the GRB itself).

With the benefit of a much larger sample size, important features of the 2015 and 2016 configurations come into focus. First, we find that, even in 2015 when only the two LIGO detectors are operating (or in 2016 during periods when the Virgo detector is not in science mode), there is at least a 60% chance of encountering the source upon searching an area of about 200 deg^2 . Second, many of these two-detector events will not be localized to a single simply connected region in the sky. We elucidate two nearly degenerate sky locations, separated by 180° , that arise when only the two LIGO detectors are operating. When a GW source falls within this degeneracy, its sky map will consist of two diametrically opposed islands of probability. Third, in our simulations, we add a third detector, Advanced Virgo, in 2016. Even though, at that time, Virgo is anticipated to be only one-third as sensitive as the other two detectors due to differing LIGO and Virgo commissioning timetables, we find that coherence with the signal in Virgo generally breaks the previously mentioned degeneracy and shrinks areas to a third of what they were with two detectors. Fourth and most importantly, a picture of a typical early Advanced LIGO event emerges, with most occurring in a limited range of Earth-fixed locations, and most sky maps broadly fitting a small number of specific morphologies.

5.1 Sources and sensitivity

BNS systems are the most promising and best understood targets for joint GW and EM detection. Though rate estimates remain uncertain, ranging from 0.01 to $10 \text{ Mpc}^{-3} \text{ Myr}^{-1}$, we choose to work with the “realistic” rate obtained from [Abadie et al. \(2010b\)](#) of $1 \text{ Mpc}^{-3} \text{ Myr}^{-1}$. This rate leads to a GW detection rate of 40 yr^{-1} at final Advanced LIGO design sensitivity. Some mergers of NSBHs are also promising sources of GW and EM emission. Two Galactic high-mass X-ray binaries (HMXBs) have been identified as possible NSBH progenitors ([Belczynski et al., 2011, 2013](#)). From these can be extrapolated a lower bound on the GW detection rate of at least 0.1 yr^{-1}

at Advanced LIGO’s final design sensitivity, although rates comparable to BNS detections are empirically plausible. Black holes in binaries may possess large spins, causing precession during the inspiral. Precession-altered phase evolution can aid in parameter estimation ([van der Sluys et al., 2008a,b](#); [Harry et al., 2014](#); [Nitz et al., 2013](#); [Raymond et al., 2009](#)), but models of waveforms suitable for rapid detection and parameter estimation are still under active development ([Blackman et al., 2014](#); [Hannam et al., 2014](#); [Taracchini et al., 2014](#)). As for the binary black hole mergers detectable by Advanced LIGO and Virgo, there are currently no compelling mechanisms for electromagnetic counterparts associated with them. We therefore restrict our attention to BNS mergers, because they have the best understood rates, GW signal models, and data analysis methods.

5.1.1 Observing scenarios

[Aasi et al. \(2013c\)](#) outline five observing scenarios representing the evolving configuration and capability of the Advanced GW detector array, from the first observing run in 2015, to achieving final design sensitivity in 2019, to adding a fourth detector at design sensitivity by 2022. In this study, we focus on the first two epochs. The first, in 2015, is envisioned as a three-month science run. LIGO Hanford (H) and LIGO Livingston (L) Observatories are operating with an averaged $(1.4, 1.4) M_{\odot}$ BNS range between 40 and 80 Mpc. The second, in 2016–2017, is a six-month run with H and L operating between 80 and 120 Mpc and the addition of Advanced Virgo (V) with a range between 20 and 60 Mpc. For each configuration, we used model noise PSD curves in the middle of the ranges in [Aasi et al. \(2013c\)](#), plotted in Figure 5.1. For H and L, we used the “early” and “mid” noise curves from [Barsotti & Fritschel \(2012\)](#) for the 2015 and 2016 scenarios respectively. For V in 2016, we used the geometric mean of the high and low curves of [Aasi et al. \(2013c\)](#). Final LIGO and Virgo design sensitivity is several steps further in the commissioning schedule than that which we consider in this paper.

5.1.2 Simulated waveforms

For each of the two scenarios we made synthetic detector streams by placing post-Newtonian inspiral signals into two months of colored Gaussian noise. We used “SpinTaylorT4” waveforms, employing the TaylorT4 approximant and accurate to 3.5PN order in phase and 1.5PN order in




Figure 5.1 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 averaged $\rho = 8$ range d_R for $(1.4, 1.4) M_\odot$ BNS mergers is given for each detector.

amplitude (Buonanno et al., 2003, 2006; Buonanno et al., 2009).³ There was an average waiting time of ≈ 100 s between coalescences. At any given time, one BNS inspiral signal was entering LIGO’s sensitive band while another binary was merging, but both signals were cleanly separated due to their extreme narrowness in time-frequency space. The PSD estimation used enough averaging that it was unaffected by the overlapping signals. Component masses were distributed uniformly between 1.2 and $1.6 M_\odot$, bracketing measured masses for components of known BNS systems as well as the $1\text{-}\sigma$ intervals of the intrinsic mass distributions inferred for a variety of NS formation channels (Pejcha et al., 2012; Özel et al., 2012).

We gave each NS a randomly oriented spin with a maximum magnitude of $\chi = c |\mathbf{S}| / Gm^2 \leq 0.05$, where \mathbf{S} is the star’s spin angular momentum and m is its mass. This range includes the most rapidly rotating pulsar that has been found in a binary, PSR J0737-3039A (Burgay et al., 2003; Brown et al., 2012). However, the fastest-spinning millisecond pulsar, PSR J1748-2446ad (Hessels et al., 2006), has a dimensionless spin parameter of ~ 0.4 , and the theoretical evolutionary limits on NS spin-up in BNS systems are uncertain.

³There is a C language implementation as the function `XLALSimInspiralSpinTaylorT4` in `lalsimulation`. See Acknowledgments and Appendix D.

5.1.3 Sensitivity to assumptions

The total detection rate depends on some of these assumptions, and in particular is sensitive to the assumed NS mass distribution. As can be seen from Equation (2.17), binaries with the greatest and most symmetric component masses can be detected at the farthest distance. According to Equation (2.19), for BNS systems the merger always occurs at kHz frequencies, on the upward $S(f) \propto f^2$ slope of the noise curves in Figure 5.1 in the regime dominated by photon shot noise (Buonanno & Chen, 2001; Adhikari, 2014). As a result, the integral in Equation (2.17) depends only weakly on masses. For equal component masses, the horizon distance scales as $d_H \propto m_{\text{NS}}^{5/6}$, so the detection rate scales rapidly with mass as $\dot{N} \propto d_H^3 \propto m_{\text{NS}}^{2.5}$.

The normalized distribution of sky localization areas depends only weakly on the distribution of NS masses. Fairhurst (2009) computes the approximate scaling of sky localization area by considering the Fisher information associated with time of arrival measurement. Valid for moderately high SNR, the RMS uncertainty in the time of arrival in a given detector is

$$\sigma_t = \frac{1}{2\pi\rho\sqrt{\overline{f^2} - \overline{f}^2}}, \quad (5.1)$$

where $\overline{f} = \overline{f^1}$, and $\overline{f^k}$ is the k th moment of frequency, weighted by the signal to noise per unit frequency,

$$\overline{f^k} \approx \left[\int_{f_1}^{f_2} \frac{|h(f)|^2 f^k}{S(f)} df \right] \left[\int_{f_1}^{f_2} \frac{|h(f)|^2}{S(f)} df \right]^{-1}. \quad (5.2)$$

As in Equation (2.17), we can substitute the approximate inspiral signal spectrum $|h(f)|^2 \propto f^{-7/3}$. The areas then scale as the product of the timing uncertainty in individual detectors, or as simply the square of Equation (5.1) for a network of detectors with similar (up to proportionality) noise PSDs. As m_{NS} varies from 1 to 2 M_\odot , the upper limit of integration f_2 given by Equation (2.19) changes somewhat, but areas change by a factor of less than 1.5. (See also Grover et al. 2014 for scaling of sky localization area with mass).

Introducing faster NS spins would result in smaller sky localization areas, since orbital precession can aid in breaking GW parameter estimation degeneracies (Raymond et al., 2009). However, rapid spins could require more exotic BNS formation channels, and certainly would require using more sophisticated and more computationally expensive GW waveforms for parameter

estimation.

5.1.4 Source locations

Source locations were random and isotropic, and uniform in distance cubed. The source distribution was cut off at the $\rho = 5$, $(1.6, 1.6) M_{\odot}$ horizon distance, far enough away that the selection of detected binaries was determined primarily by the sensitivity of the instruments and the detection pipeline, not by the artificial distance boundary.

5.1.5 Duty cycle

Following [Aasi et al. \(2013c\)](#), we assumed that each detector had an independent and random $d = 80\%$ duty cycle. In the 2015, HL configuration, this implies that both detectors are in operation $d^2 = 64\%$ of the time. In 2016, there are three detectors operating $d^3 = 51.2\%$ of the time and each of three pairs operating $d^2(1 - d) = 12.8\%$ of the time. We do not simulate any of the time when there are one or zero detectors operating, but instead fold this into conversion factors from our Monte Carlo counts to detection rates.

5.2 Detection and position reconstruction

Searches for GWs from compact binaries ([Allen et al., 2012](#); [Babak et al., 2013](#)) employ banks of matched filters, in which the data from all of the detectors are convolved with an array of template waveforms. The output of each filter is the instantaneous SNR with respect to that template in that detector. An excursion above a threshold SNR in two or more detectors with exactly the same binary parameters and within approximately one light-travel time between detectors is considered a coincidence. Coincidences may be accidental, due to chance noise fluctuations or, in real GW data streams, environmental disturbances and instrument glitches. Coincidences with sufficiently high ρ_{net} (root-sum-square of the SNR in the individual detectors) are considered detection candidates. A χ^2 statistic is used to aid in separating the true, astrophysical signals from accidental coincidences or false positives ([Allen, 2005](#); [Hanna, 2008](#); [Cannon et al., 2013](#)).

Offline inspiral searches used in past LIGO/Virgo science runs will be computationally strained in Advanced LIGO/Virgo due to denser template banks and BNS signals that remain

in band for up to $\sim 10^3$ s. To address these issues and achieve latencies of $\lesssim 1$ minute, a rapid matched-filter detection pipeline called GSTLALINSPIRAL (Cannon et al., 2012) has been developed. To mimic Advanced LIGO/Virgo observations as closely as possible, we used GSTLALINSPIRAL to extract simulated detection candidates from our two-month data streams.

5.2.1 Template waveforms

The templates were constructed from a frequency domain, post-Newtonian model describing the inspiral of two compact objects, accurate to 3.5 post-Newtonian order in phase and Newtonian order in amplitude (Buonanno et al., 2009).⁴ These waveforms neglect spins entirely. This is known to have a minimal impact on detection efficiency for BNS sources with low spins (Brown et al., 2012). These waveforms are adequate for recovering the weakly spinning simulated signals that we placed into the data stream.

5.2.2 Detection threshold

In our study, we imposed a single-detector threshold SNR of 4. A simulated signal is then considered to be detected by GSTLALINSPIRAL if it gives rise to a coincidence with sufficiently low false alarm probability as estimated from the SNR and χ^2 values. We follow the lead of Aasi et al. (2013c) in adopting a false alarm rate (FAR) threshold of $\text{FAR} \leq 10^{-2} \text{ yr}^{-1}$. Aasi et al. (2013c) claim that in data of similar quality to previous LIGO/Virgo science runs, this FAR threshold corresponds to a network SNR threshold of $\rho_{\text{net}} \geq 12$. Since our data is Gaussian and perfectly free of glitches, to obtain easily reproducible results we imposed a second explicit detection cut of $\rho_{\text{net}} \geq 12$. We find that our joint threshold on FAR and SNR differs negligibly from a threshold on SNR alone. Because any given simulated signal will cause multiple coincidences at slightly different masses and arrival times, for each simulated signal we keep only the matching candidate with the lowest SNR.

⁴These are in lalsimulation as the function `XLALSimInspiralTaylorF2`. See acknowledgements and Appendix D.

5.2.3 Sky localization and parameter estimation

All detection candidates are followed up with rapid sky localization by BAYESTAR and a subset were followed up with full parameter estimation by the LALINFERENCE/MCMC/NEST/BAMBI stochastic samplers. The three different stochastic samplers all use the same likelihood, but serve as useful cross-verification. Both BAYESTAR and the three stochastic samplers are coherent (exploiting the phase consistency across all detectors) and Bayesian (making use of both the GW observations and prior distributions over the source parameters). They differ primarily in their input data.

BAYESTAR’s likelihood function depends on only the information associated with the triggers comprising a coincidence: the times, phases, and amplitudes on arrival at each of the detectors. BAYESTAR exploits the leading-order independence of errors in the extrinsic and intrinsic parameters by holding the masses fixed at the values estimated by the detection pipeline. Marginalized posterior distributions for the sky positions of sources are produced by numerically integrating the posterior in each pixel of the sky map. Because BAYESTAR’s analysis explores only a small sector of the full parameter space, never performs costly evaluations of the post-Newtonian GW waveforms, and uses highly tuned standard numerical quadrature techniques, it takes well under a minute (see Figure 4.1).

On the other hand, the likelihood function used for the stochastic samplers depends on the full GW data, and is the combination of independent Gaussian distributions for the residual in each frequency bin after model subtraction. Posterior distributions for the sky position are produced by sampling the full parameter space of the signal model, then marginalizing over all parameters but the sky location. This method requires the generation of a model waveform for each sample in parameter space, making it far more expensive than the BAYESTAR approach, but independent of the methods and models used for detection. Most importantly, intrinsic parameters (including spins) can be estimated using these higher-latency methods. For the purposes of this study, parameter estimation used the same frequency-domain, non-spinning waveform approximant as the detection pipeline. Analyses that account for the spin of the compact objects are more costly, currently taking weeks instead of days to complete, and will be the subject of a future study.

5.3 Results

Of $\sim 100,000$ simulated sources spread across the 2015 and 2016 scenarios, ≈ 1000 events survived as confident GW detections.⁵ No false alarms due to chance noise excursions survived our detection threshold; all events which should have been detectable were detected. We constructed probability sky maps using BAYESTAR for all events and using LALINFERENCE_NEST/MCMC for a randomly selected subsample of 250 events from each scenario.⁶⁷ The top four panels (a, b, c, d) of Figure 5.3 show cumulative histograms of the areas in deg^2 inside of the smallest 50% and 90% confidence regions for each event, for both sky localization methods. These contours were constructed using a ‘water-filling’ algorithm: we sampled the sky maps using equal-area HEALPix (**Hierarchical Equal Area isoLatitude Pixelization!**; Górski et al., 2005) pixels, ranked the pixels from most probable to least, and finally counted how many pixels summed to a given total probability. In the bottom two panels (e), (f) of Figure 5.3, we also show a histogram of the smallest such constructed region that happened to contain the true location of each simulated source. We call this the searched area.

Panels (a–d) and (e), (f) may be thought of as measuring precision and accuracy respectively. The former measure how dispersed or concentrated each individual sky map is, while the latter describe how far the localization is from the true sky position. The 90% area histograms and the searched area histograms also answer different but complementary questions that relate to two different strategies for following up LIGO/Virgo events. One might decide in 2015 to search for optical counterparts of all GW events whose 90% areas are smaller than, for example, 200 deg^2 . By finding 200 deg^2 on the horizontal axis of the 90% area histogram, one would find that this corresponds to following up 10% of all GW detections. On the other hand, one might decide to always search the most probable 200 deg^2 area for every GW event, corresponding to a different confidence level for every event. In this case, one would find 200 deg^2 on the horizontal axis of the searched area histogram, and find that this strategy would enclose the true location of the GW source 64% of the time.

⁵There were slightly fewer surviving events in the 2016 configuration than in the 2015 configuration. 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 two months of data were allocated to the HL network. In 2016 about 43 days were devoted to the HLV and HL configurations, with the remaining 17 days of HV and LV mode contributing few detections.

⁶Results from LALINFERENCE_BAMBI are not shown in our plots because this sampler was run for only 30 events, and the sampling error bars would overwhelm the plots.

⁷The three stochastic samplers LALINFERENCE_NEST/MCMC/BAMBI were interchangeable to the extent that they used the same likelihood and produced sky maps that agreed with each other.

One might naively expect that self-consistency would require the 90% confidence area and searched area histograms to intersect at 90% of detections, but this is not generally required because the posteriors of different events have widely different dimensions. However, it is true that 90% of sources should be found within their respective 90% confidence contours. This can be formalized into a graphical self-consistency test called a probability–probability (P – P) plot (see [Sidery et al. 2014](#) for applications in GW parameter estimation). For each event, one follows the water-filling procedure to walk outward from the most probable pixel until one reaches the pixel that contains the true sky location, noting the total probability p that has been traversed. A cumulative histogram of p should be diagonal, within a binomial confidence band. It is already well established that LALINFERENCE_NEST/MCMC/BAMBI satisfy the P – P plot test when deployed with accurate templates and reasonable priors. We found at first that BAYESTAR’s P – P plots tended to sag below the diagonal, indicating that though the accuracy (i.e., searched area) was comparable to the stochastic samplers, the *precision* was overstated with confidence intervals that were only about 70% of the correct area. This was rectified by pre-scaling the SNRs from GSTLAL_INSPIRAL by a factor of 0.83 prior to running BAYESTAR. This correction factor suggests that, for example, an SNR 10 trigger from GSTLAL_INSPIRAL has the *effective information content* of an SNR 8.3 signal. The missing information may be due to losses from the discreteness of the template bank, from the SVD, from mismatch between the matched-filter templates and the simulated signals, from the small but nonzero correlations between masses and intrinsic parameters, or from elsewhere within the detection pipeline. The correction is hard-coded into the rapid localization. With it, the P – P plots are diagonalized without negatively affecting the searched area (see Figure 5.2).

The left-hand axes of all four panels of Figure 5.3 show the expected cumulative number of detections, assuming the ‘realistic’ BNS merger rates from [Abadie et al. \(2010b\)](#). We stress that the absolute detection rate might be two orders of magnitude smaller or one order of magnitude higher due to the large systematic uncertainty in the volumetric rate of BNS mergers, estimated from population synthesis and the small sample of Galactic binary pulsars ([Abadie et al., 2010b](#)). An additional source of uncertainty in the detection rates is the Advanced LIGO/Virgo commissioning schedule given in [Aasi et al. \(2013c\)](#). The proposed sensitivity in the 2016 scenario may be considered a plausible upper bound on the performance of the GW detector network in 2015, if commissioning occurs faster than anticipated. Likewise, the quoted sensitivity in the




Figure 5.2 P – P plots for BAYESTAR localizations in the 2015 and 2016 configurations. The gray lozenge around the diagonal is a target 95% confidence band derived from a binomial distribution.

2015 scenario is a plausible lower bound on the performance in 2016. The right-hand axes show the cumulative percentage of all detected sources. These percentages depend only on the gross features of the detector configuration and not on the astrophysical rates, so are relatively immune to the systematics described above.

Table 5.1 summarizes these results.

5.3.1 2015

Our 2015 scenario assumes two detectors (HL) operating at an anticipated range of 54 Mpc. About 0.1 detectable BNS mergers are expected, though there are nearly two orders of magnitude systematic uncertainty in this number due to the uncertain astrophysical rates. A detection in 2015 is possible, but likely only if the BNS merger rates or the detectors' sensitivity are on the modestly optimistic side. A typical or median event (with a localization area in the 50th percentile of all detectable events) would be localized to a 90% confidence area of $\sim 500 \text{ deg}^2$.

We find that the area histograms arising from the BAYESTAR rapid sky localization and the full parameter estimation agree within sampling errors, and that the sky maps resulting from the two analyses are comparable for any individual event. Put differently, the rapid sky localization contains nearly all of the information about sky localization for these events, with the full probability distributions over masses and spins becoming available when the stochastic samplers finish on a timescale of a day.

Figure 5.4(a) shows a histogram of the cosine of the angular separation between the true location of the simulated GW source and the maximum a posteriori (MAP) estimate (the mode of the sky map, or the most probable location). The main feature is a peak at low separation. However, there is a second peak at the polar opposite of the true location, 180° away; about 15% of events are recovered between 100 and 180° away from the true location.

Correspondingly, for any one event, it is common to find the probability distributed across two antipodal islands on opposite sides of the mean detector plane. We define this plane by finding the average of the two vectors that are normal to the planes of the two detectors' arms, and then taking the plane to which this vector is normal. This plane partitions the sky into two hemispheres. We find that one hemisphere is favored over the other by less than 20% (which is to say that the odds favoring one hemisphere over the other are as even as 60%/40%) for 20% of events.

The second peak admits a simple explanation as an unavoidable degeneracy due to the relative

Table 5.1. Detection rate and sky localization accuracy

		2015		2016	
Detectors LIGO (HL) BNS range Run duration No. detections		HL 54 Mpc 3 months 0.091		HLV 108 Mpc 6 months 1.5	
		rapid	full PE	rapid	full PE
Fraction 50% CR Smaller than	5 deg ²	—	—	9%	14%
	20 deg ²	2%	3%	15%	35%
	100 deg ²	30%	37%	32%	72%
	200 deg ²	74%	80%	62%	90%
	500 deg ²	100%	100%	100%	100%
Fraction 90% CR Smaller than	5 deg ²	—	—	2%	2%
	20 deg ²	—	—	8%	14%
	100 deg ²	3%	4%	15%	31%
	200 deg ²	10%	13%	19%	45%
	500 deg ²	44%	48%	39%	71%
Fraction Searched Area Smaller than	5 deg ²	3%	4%	11%	20%
	20 deg ²	14%	19%	23%	44%
	100 deg ²	45%	54%	47%	71%
	200 deg ²	64%	70%	62%	81%
	500 deg ²	87%	89%	83%	93%
Median Area {	50% CR	138 deg ²	124 deg ²	162 deg ²	43 deg ²
	90% CR	545 deg ²	529 deg ²	621 deg ²	235 deg ²
	searched	123 deg ²	88 deg ²	118 deg ²	29 deg ²

Note. — This table is a summary of the 2015 and 2016 scenarios, listing the participating detectors, BNS horizon distance, run duration, and fractions of events localized within 5, 20, 100, 200, or 500 deg². A dash (—) represents less than 1% of detections.

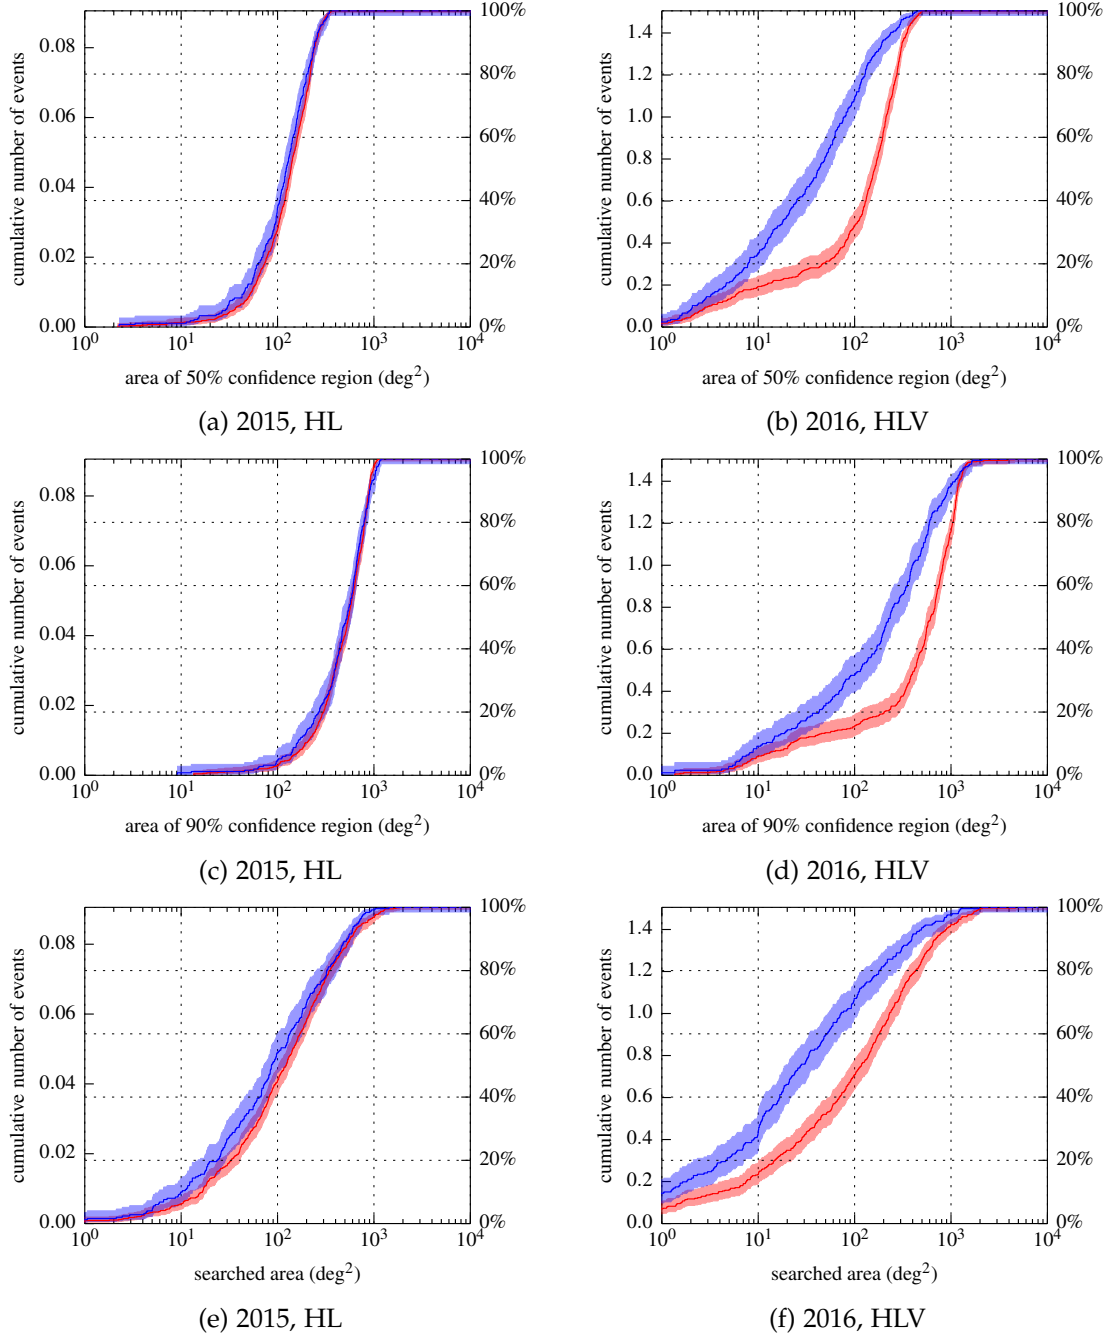


Figure 5.3 Cumulative histogram of sky localization areas in the 2015 (HL) and 2016 (HLV) scenarios. Plots in the left column (a, c, e) refer to the 2015 configuration and in the right column (b, d, f) to the 2016 configuration. The first row (a, b) shows the area of the 50% confidence region, the second row (c, d) shows the 90% confidence region, and the third row (e, f) shows the “searched area,” the area of the smallest confidence region containing the true location of the source. The red lines comprise all detections and their sky maps produced with BAYESTAR, and the blue lines represent sky maps for the random subsample of 250 detections analyzed with LALINFERENCE_NEST/MCMC. The light shaded region encloses a 95% confidence interval accounting for sampling errors (computed from the quantiles of a beta distribution; [Cameron, 2011](#)). The left axes show the number of detections localized within a given area assuming the “realistic” BNS merger rates from ([Abadie et al., 2010b](#)). The right axes show the percentage out of all detected events.




Figure 5.4 Normalized histogram of the cosine angular separation between the location of the simulated GW source and the MAP location estimate, for (a) the 2015 configuration and (b) the 2016 configuration. The red line encompasses all detections and their BAYESTAR localizations, and the blue line the subsample of 250 events analyzed by LALINFERENCE.NEST/MCMC. The inset shows the distribution of angle offsets for angles less than 60° .

positions of the H and L interferometers. Before the Hanford and Livingston sites were selected, it was decided that the detectors' arms would be as closely aligned as possible (Vogt, 1989, Section V-C-2). Significant misalignment would have created patches of the sky that were accessible to one detector but in a null of the other detector's antenna pattern, useless for a coincidence test.

The near alignment maximized the range of the detectors in coincidence, though at a certain expense of parameter estimation. Observe that the sensitivity of an interferometric GW detector is identical at antipodal points (i.e., symmetric under all rotations by 180°). Therefore, any source that lies in the plane of zero time delay between the detectors is always localized to two opposite patches. Because the HL detectors were placed nearby (at continental rather than intercontinental distances) on the surface of the Earth so as to keep their arms nearly coplanar, their combined network antenna pattern has two maxima that lie on opposite sides of that great circle. As a consequence, a large fraction of sources are localized to two islands of probability that cannot be distinguished based on time or amplitude on arrival. See Figure 5.5 for an illustration of these two degenerate patches.

A second undesirable side effect of the aligned antenna patterns is that GW polarization, observed as the phase difference on arrival at these two detectors, is of limited help for parameter estimation.

A fairly typical sky map morphology, even at modestly high SNR_{net} , will consist of two extended arc-shaped modes, one over North America and a mirror image on the opposite side of the Earth. See Figure 5.6 for a typical event exhibiting this degeneracy. In this example, it is also possible to discern two narrow stripes resembling the forked tongue of a snake. This is a reflection of the HL network's limited polarization sensitivity. It occurs when the GW phases on arrival support two different binary inclination angles, with the orbital plane nearly facing the observer but with opposite handedness (usually peaked at $\iota \approx 30^\circ$ and $\iota \approx 150^\circ$; see Schutz 2011). The two forks cross at a sky location where the GW data cannot distinguish between a clockwise or counterclockwise orbit.

The HL degeneracy is even apparent in earlier works on localization of GW bursts with networks of four or more detectors: Klimenko et al. (2011) drew a connection between accurate position reconstruction and sensitivity to both the '+' and 'x' GW polarizations, and noted that the close alignment of the HL detector network adversely affects position reconstruction. (They did not, however, point out the common occurrence of nearly 180° errors, or note that the worst

Figure 5.5 HL degeneracy. This, like all sky plots in this paper, is a Mollweide projection in geographic coordinates to emphasize spatial relationships with respect to the Earth-fixed GW detector network as well as possible ground-based telescope sites. Pluses denote the locations of signals whose best-estimate locations are offset by $\geq 100^\circ$, comprising the large-offset peak that is evident in Figure 5.4(a). The locations of zero time delay (simultaneous arrival at the H and L detectors) is shown as a thick black line. Shading indicates the RMS network antenna pattern, with darker areas corresponding to high sensitivity and white corresponding to null sensitivity.





Figure 5.6 Localization of a typical circa 2015 GW detection. This is a Mollweide projection in geographic coordinates. Shading is proportional to posterior probability per deg^2 . This is a moderately loud event with $\rho_{\text{net}} = 15.0$, but its 90% confidence area of 630 deg^2 is fairly typical, in the 60th percentile of all detections. The sky map is bimodal with two long, thin islands of probability over the northern and southern antenna pattern maxima. Neither mode is strongly favored over the other. Each island is forked like a snake’s tongue, with one fork corresponding to the binary having face-on inclination ($\iota \approx 0^\circ$) and the other fork corresponding to face-off ($\iota \approx 180^\circ$). This is event ID 18951 from Tables D.1 and D.2 and the online material (see Appendix D for more details).



GW localizations paradoxically occur where the HL network’s sensitivity is the greatest.)

The HL degeneracy affects most events that occur $\lesssim 30^\circ$ from one of the antenna pattern maxima. Most events that are $\gtrsim 50^\circ$ away have localizations that consist of a single long, thin arc or ring. See Figure 5.7 for an example.

In Figure 5.8, we have plotted a histogram of the number of disconnected modes comprising the 50% and 90% confidence regions and the searched area, for the rapid localizations in the 2015 configuration. The ratios of events having one, two, or three or more modes depend weakly on the selected confidence level. In 2015, using either the 50% contour or the searched area, we find that about half of the events are unimodal and about a third are bimodal, the rest comprising three or more modes. Using the 90% contour, we find that about a third of the events are unimodal and

Figure 5.7 Localization of a typical circa 2015 GW detection. This is a Mollweide projection in geographic coordinates. Shading is proportional to posterior probability per deg^2 . This event's $\rho_{\text{net}} = 12.7$ is near the threshold, but its 90% confidence area of 530 deg^2 near the median. The sky map consists of a single, long, thin island exhibiting the forked-tongue morphology. *This is event ID 20342 from Tables D.1 and D.2 and the online material (see Appendix D for more details).*





Figure 5.8 Frequency with which GW sky maps have one, two, or more disconnected modes during 2015. From top to bottom are the number of modes contained within the smallest confidence contour containing each simulated signal, the smallest 90% contour, and the smallest 50% contour. In 2015, roughly half of the sky maps will be unimodal, with most of the remainder being bimodal.




Figure 5.9 Breakdown of 2016 scenario by detector network. The top row shows the duty fraction of each subset of the detector network, the fraction of time when all three detectors (HLV) are observing, when any pair of detectors are observing (HL, LV, or HV), or when zero or one detectors are observing (—). The second row shows the fraction of coincident detections that occur under any given network (HLV, HL, LV, or HV). The last row shows the fraction of coincident detections for which the given detectors have signals above the single-detector threshold of $\rho = 4$.

about half are bimodal.

5.3.2 2016

In our 2016 scenario, the HL detectors double in range to 108 Mpc and the V detector begins observations with a range of 36 Mpc. Over this six-month science run we expect ~ 1.5 detections, assuming a BNS merger rate of $1 \text{ Mpc}^{-3} \text{ Myr}^{-1}$. Figure 5.9 shows how livetime and duty cycle breaks down according to detector network (HLV, HL, LV, or HV). About half of the time all three detectors are online, with the remaining time divided into four almost equal parts among the three pairs of detectors or ≤ 1 detector. However, the HLV network accounts for about three-quarters of detections and the HL network for most of the rest.

When all three detectors (HLV) are operating, most detections are comprised of H and L triggers, lacking a trigger from V because the signal is below the single-detector threshold of $\rho = 4$. Slightly more than half (57%) of all detections have a signal below threshold in one operating detector (almost always V) while slightly less than half (43%) consist of triggers from all operating detectors.

The first half consists mainly of HLV events that are detected by HL but not Virgo. For these events, the stochastic samplers provide a marked improvement in sky localization; their 90% confidence regions have about one-third as much area as their rapid localizations. This is because the rapid localization makes use of only the triggers provided by the detection pipeline, lacking information about the signal in Virgo if its SNR is < 4 . The stochastic samplers, on the other hand,


can use data from all operating detectors, regardless of SNR. Therefore, in the present analysis, an improved sky localization would be available for half of the detections on a timescale of a day. Fortunately, for BNS sources, it is immediately known whether an improved localization is possible, since this statement only depends on what detectors were online and which contributed triggers. On the other hand, it may be possible to provide prompt sky localizations for all events by simply lowering the single-detector threshold. If the single-detector threshold was dropped to unity, essentially no event would lack a Virgo trigger. There are also efforts to do away with the single-detector threshold entirely (Keppel, 2012, 2013). Simultaneously, there is promising work under way in speeding up the full parameter estimation using reduced order quadratures (Canizares et al., 2013), interpolation (Smith et al., 2014a), jump proposals that harness knowledge of the multimodal structure of the posterior (Farr et al., 2014), hierarchical techniques (Farr & Kalogera, 2013), and machine learning techniques to accelerate likelihood evaluation (Graff et al., 2012, 2014). It seems possible that the delayed improvement in sky localization may be a temporary limitation that can be overcome well before 2016.

The second half consists of HLV events with triggers from all three detectors and events that occur when only HL, HV, or LV are operating. For these, the BAYESTAR analysis and the full stochastic samplers produce comparable sky maps.


For nearby loud sources ($\rho_{\text{net}} \gtrsim 20$), the HLV network frequently produces compact sky localizations concentrated in a single island of probability. However at low SNR ($\rho_{\text{net}} \lesssim 20$), and especially for the events that are detected as only double coincidence (HL), the refined localizations from the full stochastic samplers often identify many smaller modes. A $\rho_{\text{net}} = 13.4$ example is shown in Figure 5.10. In this event, the rapid sky localization has two modes and has a morphology that is well-described by the HL degeneracy explained in Section 5.3.1. However, the refined, full parameter estimation breaks this into at least four smaller modes.

Of the remaining events, most occur when only the two HL detectors are operating. These look qualitatively the same as those in the 2015 case; their sky maps generally exhibit one or two slender islands of probability. However, percentage-wise, two-detector events are less well localized in the 2016 scenario than in the 2015 scenario. This unusual result is easily explained. Though the LIGO detectors improve in sensitivity at every frequency, with the particular noise curves that we assumed, the signal bandwidth is actually slightly lower with the 2016 sensitivity compared to 2015. This is because of improved sensitivity at low frequency. Applying Equation (5.1), we

Figure 5.10 Localization of a typical circa 2016 GW detection in the HLV network configuration. This is a Mollweide projection in geographic coordinates. This event consists of triggers in H and L and has $\rho_{\text{net}} = 13.4$. The rapid sky localization gives a 90% confidence region with an area of 1100 deg^2 and the full stochastic sampler gives 515 deg^2 . This is event ID 821759 from Tables D.3 and D.4 and the online material (see Appendix D for more details).




(a) BAYESTAR



(b) LALINFERENCE_MCMC

Figure 5.11 Rapid localization of a typical circa 2016 GW detection in the HV network configuration. This is a Mollweide projection in geographic coordinates. This event's $\rho_{\text{net}} = 12.2$ is near the detection threshold. Its 90% confidence area is 4600 deg^2 , but the true position of the source (marked with the white pentagram) is found after searching 65 deg^2 . This is event ID 655803 from Tables D.3 and D.4 and the online material (see Appendix D for more details).



find that for a $(1.4, 1.4) M_{\odot}$ binary at $\rho = 10$, one of the 2015 LIGO detectors has an RMS timing uncertainty of $131 \mu\text{s}$, whereas one of the 2016 LIGO detectors has a timing uncertainty of $158 \mu\text{s}$. Clearly, the 2016 detectors will produce more constraining parameter estimates for sources at any fixed distance as the SNR improves. However, for constant SNR the 2016 LIGO detectors should find areas that are $(158/131)^2 = 1.45$ times larger than events at the same SNR in 2015. This is indeed what we find.

Two-detector events involving Virgo (HV and LV) are rare, accounting for only about 6% of detections. Sky maps for these events sometimes exhibit multiple fringes spread over a quadrant of the sky. These are in part due to the increased importance of phase-on-arrival due to the oblique alignment of the LIGO and Virgo antenna patterns, which gives the network a limited ability to measure GW polarization. Occasionally there are also diffuse clouds of probability near the participating LIGO detector's two antenna pattern maxima, which may be a vestige of the antenna pattern. A typical HV event that exhibits both features is shown in Figure 5.11.

5.4 Discussion

5.4.1 Caveats

We reiterate that the scenarios we have described make assumptions about the astrophysical rate of BNS mergers and the Advanced LIGO/Virgo sensitivity as a function of time. The former is subject to orders of magnitude uncertainty due to the small sample of known galactic binary pulsars as well as model dependence in population synthesis (Abadie et al., 2010b). The latter could deviate from Aasi et al. (2013c) depending on actual Advanced LIGO/Virgo commissioning progress. However, the fractions of events localized within a given area are robust with respect to both of these effects.

We have dealt only with BNS mergers. NSBH mergers are also promising sources for closely related GW signals and EM transients. A similar, comprehensive investigation of GW sky localization accuracy for NSBH signals is warranted.

In this simulation, we have used ideal Gaussian noise, but selected a detection threshold that is designed to reproduce expected performance in detectors with realistically wide tails due to instrumental and environmental glitches. If Advanced LIGO's and Virgo's improved seismic isolation and control systems are even more effective at suppressing such glitches than their initial counterparts were, then the ρ_{net} threshold for confident detection would decrease, yielding discoveries earlier but with larger typical sky localization areas.

We remind the reader that the events comprising this study would be regarded as confident detections, with $\text{FAR} \lesssim 10^{-2} \text{ yr}^{-1}$, based on GW observations alone. In practice, some observers may choose to follow up more marginal detection candidates. For instance, a group with enough resources and telescope time to follow up one candidate per month might filter events with $\text{FAR} \leq 12 \text{ yr}^{-1}$. A high false alarm rate threshold will admit correspondingly lower ρ_{net} candidates with coarser localizations than those which we have presented here.

Finally, on a positive note, the number of detections is expected to increase considerably as commissioning proceeds toward final design sensitivity. Furthermore, sky localization will improve radically as the HLV detectors approach comparable sensitivity. The addition of two more planned ground-based GW detectors, LIGO-India and KAGRA, would likewise increase rates and improve sky localizations dramatically (Schutz, 2011; Veitch et al., 2012; Fairhurst, 2014; Nissanke et al., 2013; Aasi et al., 2013c).

5.4.2 Detection scenarios

From our representative sample of hundreds of early Advanced LIGO/Virgo events emerge a few common morphologies and several possible scenarios for the early detections of GWs from a BNS merger.

We find that in both 2015 and 2016, the detection rate is highly anisotropic and proportional to the cube of the network antenna pattern, with a strong excess above North America and the Indian Ocean and deficits in four spots over the south Pacific, south Atlantic, Caucasus, and north Pacific.

1. *HL event, single arc*—This scenario is relevant for the HL network configuration and applies to both 2015 and 2016. Figure 5.7 shows a typical sky map for a near-threshold detection with $\rho_{\text{net}} = 12.7$, exhibiting a single long, extended arc spanning $\sim 500 \text{ deg}^2$.
2. *HL event, two degenerate arcs*—This scenario also applies to 2015 or to HL livetime in 2016. Figure 5.6 shows a typical sky map with a moderately high $\rho_{\text{net}} = 15.0$, localized to $\sim 600 \text{ deg}^2$. Its localization embodies the HL degeneracy, with two strong, long, thin modes over North America and the Indian Ocean, separated by nearly 180° and therefore 12 hr apart in hour angle. Inevitably, one of these two modes will be nearly Sun-facing and inaccessible to optical facilities. Because of the bimodality, these sky maps can span slightly larger areas than case 1. After taking an inevitable 50% hit in visibility, such events resemble the single arc scenario.

Whether a given source falls into scenario 1 or 2 is largely determined by its sky location relative to the network antenna pattern. The transition occurs between $\sim 30^\circ$ and $\sim 50^\circ$ away from the two points of maximum sensitivity.

3. *HLV event, degeneracy broken by Virgo*—This scenario applies only to the 2016 configuration, while all three instruments are online. The rapid sky localization looks similar to the previous scenario, a pair of long, thin rings over the northern and southern hemispheres, but the full parameter estimation cuts this down to a handful of islands of probability covering as little as half to a third of the area, $\sim 200 \text{ deg}^2$. For such an event, the refined localization could be used to guide \sim day-cadence kilonova-hunting observations or to re-target the vetting of afterglow candidates arising from early-time observations. Several wide-field facilities could

be employed to monitor modes that lie in different hemispheres.

4. *HLV event, compactly localized*—This is another 2016, three-detector scenario. It describes many of the events that are detected with triggers in all three instruments. These are many of the best-localized events, with 90% confidence regions only a few times 10 deg^2 in area. These events are generally localized to one simply connected region and exhibit a less pronounced preference for particular sky locations. In this scenario, it is most likely that the rapid sky localization and the full parameter estimation will be similar. This is observationally the simplest scenario: just one of the several wide-field optical searches (for instance, ZTF or BlackGEM) would be able to scan the whole error region at a daily cadence.

5.4.3 Comparison with other studies

This is the first study so far to combine an astrophysically motivated source population, realistic sensitivity and detector network configurations, event selection effects arising from a genuine detection pipeline instead of an ad hoc threshold, and parameter estimation algorithms that will be deployed for GW data analysis. This study also has a much larger sample size and lower statistical uncertainty than most of the prior work. It is therefore somewhat difficult to compare results to other studies, which each have some but not all of these virtues.

To the best of the authors' knowledge, [Raymond et al. \(2009\)](#) were the first to point out the power of Bayesian priors for breaking sky degeneracies in two-detector networks, challenging a prevailing assumption at the time that two detectors could only constrain the sky location of a compact binary signal to a degenerate annulus. [Aasi et al. \(2013c\)](#) speculated that two-detector, 2015, HL configuration sky maps would be rings spanning “hundreds to thousands” of deg^2 and that coherence and amplitude consistency would “sometimes” resolve the localizations to shorter arcs. With our simulations, we would only revise that statement to read “hundreds to a thousand” deg^2 and change “sometimes” to “always.” [Kasliwal & Nissanke \(2014\)](#) recently argued for the feasibility of optical transient searches (in the context of kilonovae) with two-detector GW localizations.

[Aasi et al. \(2013c\)](#) used time-of-arrival triangulation ([Fairhurst, 2009](#)) to estimate the fraction of sources with 90% confidence regions smaller than 20 deg^2 , finding a range of 5–12% for 2016.

We find 14%. Our values are more optimistic, but perhaps also more realistic for the assumed detector sensitivity. Our sky localization takes into account phase and amplitude information, which [Grover et al. \(2014\)](#) points out can produce ≈ 0.4 times smaller areas compared to timing alone. However, it is clear from both [Aasi et al. \(2013c\)](#) and our study that such well-localized events will comprise an exceedingly small fraction of GW detections until the end of the decade. We therefore echo [Kasliwal & Nissanke \(2014\)](#) in stressing the importance of preparing to deal with areas of hundreds of deg^2 in the early years of Advanced LIGO and Virgo.

[Nissanke et al. \(2013\)](#) used an astrophysical distance distribution, drawing source positions uniformly from comoving volume for distances $d_L > 200$ Mpc and from a B -band luminosity-weighted galaxy catalog for distances $d_L \leq 200$ Mpc. They generated sky maps using their own MCMC code. In a similar manner to our present study, they imposed a threshold of $\rho_{\text{net}} > 12$. They explored several different GW detector network configurations. The most similar to our 2016 scenario was an HLV network at final design sensitivity. They found a median 95% confidence region area of $\sim 20 \text{ deg}^2$. In comparison, we find a 95 deg^2 confidence area of 374 deg^2 . Our much larger number is explained by several factors. First, we did not draw nearby sources from a galaxy catalog, so we have fewer loud, nearby sources. Second, since we accounted for duty cycle, poorly localized two-detector events account for a quarter of our sample. Third, and most important, we assumed Advanced Virgo’s anticipated initial sensitivity rather than its final design sensitivity.

[Rodríguez et al. \(2014\)](#) also studied an HLV network at final design sensitivity. Their simulated signals had identically zero noise content, the average noise contribution among all realizations of zero-mean Gaussian noise. All of their simulated events had a relatively high $\rho_{\text{net}} = 20$. They found a median 95% confidence area of 11.2 deg^2 . If we consider all of our 2016 scenario HLV events with $19.5 \leq \rho_{\text{net}} \leq 20.5$, we find a median area of 126 deg^2 . Our significantly larger number is once again partly explained by our less sensitive Virgo detector, which introduces several multimodal events even at this high ρ_{net} .

Similarly, [Grover et al. \(2014\)](#) and [Sidery et al. \(2014\)](#) studied a three-detector network, but at Initial LIGO design sensitivity. These studies were primarily concerned with evaluating Bayesian parameter estimation techniques with respect to triangulation methods. They found much smaller areas, with a median of about 3 deg^2 . Both papers used a source population that consisted mainly of very high-SNR signals with binary black hole masses, with distances distributed logarithmically. All of these effects contribute to unrealistically small areas.

Finally, [Kasliwal & Nissanke \(2014\)](#) made the first small-scale systematic study of localizability with two LIGO detectors, albeit at final Advanced LIGO design sensitivity. For these noise curves and a $(1.4, 1.4) M_{\odot}$ binary with single-detector $\rho = 10$, Equation (5.1) gives a timing uncertainty of $142 \mu\text{s}$. Their different choice of noise curves should result in areas that are $(131/142)^2 \approx 0.85$ times smaller than ours, at a given ρ_{net} . As we have, they imposed a network SNR threshold of $\rho_{\text{net}} \geq 12$ on all detections⁸. They found a median 95% confidence area of $\sim 250 \text{ deg}^2$ from a catalog of 17 events. From our 2015 scenario, we find a median 90% confidence area that is almost twice as large, $\sim 500 \text{ deg}^2$. Though we cannot directly compare our 90% area to their 95% area, our 95% area would be even larger. Several factors could account for this difference, including the smaller sample size in [Kasliwal & Nissanke \(2014\)](#). Also, [Kasliwal & Nissanke \(2014\)](#), like [Nissanke et al. \(2013\)](#), drew nearby sources from a galaxy catalog to account for clustering, so their population may contain more nearby, well-localized events than ours. Another difference is that [Kasliwal & Nissanke \(2014\)](#) do not report any multimodal localizations or the 180° degeneracy that we described in Section 5.3.1.

5.4.4 Conclusion

Many previous sky localization studies have found that networks of three or more advanced GW detectors will localize BNS mergers to tens of deg^2 . However, given realistic commissioning schedules, areas of hundreds of deg^2 will be typical in the early years of Advanced LIGO and Virgo.

We caution that multimodality and long, extended arcs will be a common and persistent feature of Advanced LIGO/Virgo detections. We caution that existing robotic follow-up infrastructure designed for GRBs, whose localizations are typically nearly Gaussian and unimodal, will need to be adapted to cope with more complicated geometry. In particular, we advise optical facilities to evaluate the whole GW sky map when determining if and when a given event is visible from a particular site.

We have elucidated a degeneracy caused by the relative orientations of the two LIGO detectors,

⁸[Nissanke et al. \(2013\)](#) and [Kasliwal & Nissanke \(2014\)](#) present a parallel set of results for a threshold $\rho_{\text{net}} > 8.5$, relevant for a coherent GW search described by [Harry & Fairhurst \(2011a\)](#). However, the coherent detection statistic described by [Harry & Fairhurst \(2011a\)](#) is designed for targeted searches at a known sky location (for instance, in response to a GRB). Thus the $\rho_{\text{net}} > 8.5$ threshold is not relevant for optical follow-up triggered by a detection from an all-sky GW search. Furthermore, this reduced threshold is not relevant to the HL configuration because the coherent detection statistic reduces to the network SNR for networks of two detectors or fewer.

such that position reconstructions will often consist of two islands of probability separated by 180° . We have shown that this degeneracy is largely broken by adding Virgo as a third detector, even with its significantly lower sensitivity. We have shown that sub-threshold GW observations are important for sky localization and parameter estimation.

We have demonstrated a real-time detection, sky localization, and parameter estimation infrastructure that is ready to deliver Advanced LIGO/Virgo science. The current analysis has some limitations for the three-detector network, an undesirable trade-off of sky localization accuracy and timescale. Work is ongoing to lift these limitations by providing the rapid sky localization with information below the present single-detector threshold and by speeding up the full parameter estimation by a variety of methods ([Canizares et al., 2013](#); [Smith et al., 2014a](#); [Farr et al., 2014](#); [Farr & Kalogera, 2013](#); [Graff et al., 2012](#)).

We have exhibited an approach that involves three tiers of analysis, which will likely map onto a sequence of three automated alerts with progressively more information on longer timescales, much as the way in which observers in the GRB community are used to receiving a sequence of Gamma-ray Coordinates Network (GCN) notices about a high-energy event.

The maximum timescale of the online GW analysis, about a day, is appropriate for searching for kilonova emission. However, the availability of BAYESTAR’s rapid localizations within minutes of a merger makes it possible to search for X-ray and optical emission. Due to jet collimation, these early-time signatures are expected to accompany only a small fraction of LIGO/Virgo events. However, the comparative brightness and distinctively short timescale of the optical afterglow makes it an attractive target. PTF has recently proved the practicality of wide-field afterglow searches through the blind discovery of afterglow-like optical transients ([Cenko et al., 2013b](#); [Cenko et al., 2014](#)) and the detection of optical afterglows of *Fermi* GBM bursts ([Singer et al., 2013](#)). We encourage optical transient experiments such as ZTF and BlackGEM to begin searching for EM counterparts promptly, based on the rapid GW localization. In the most common situation of no afterglow detection, the early observations may be used as reference images for longer-cadence kilonova searches.

Acknowledgements

L.P.S. and B.F. thank generous support from the NSF in the form of Graduate Research Fellowships. B.F. acknowledges support through NSF grants DGE-0824162 and PHY-0969820. A.L.U. and C.P. gratefully acknowledge NSF support under grants PHY-0970074 and PHY-1307429 at the University of Wisconsin–Milwaukee (UWM) Research Growth Initiative. J.V. was supported by the research programme of the Foundation for Fundamental Research on Matter, which is partially supported by the Netherlands Organisation for Scientific Research, and by Science and Technology Facilities Council grant ST/K005014/1. P.G. is supported by a NASA Postdoctoral Fellowship administered by the Oak Ridge Associated Universities.

GSTLAL analyses were produced on the NEMO computing cluster operated by the Center for Gravitation and Cosmology at UWM under NSF Grants PHY-0923409 and PHY-0600953. BAYESTAR analyses were performed on the LIGO–Caltech computing cluster. The MCMC computations were performed on Northwestern’s CIERA High-Performance Computing cluster GRAIL⁹.

We thank Patrick Brady, Vicky Kalogera, Erik Katsavounidis, Richard O’Shaughnessy, Ruslan Vaulin, and Alan Weinstein for helpful discussions.

This research made use of Astropy¹⁰ (Robitaille et al., 2013), a community-developed core Python package for Astronomy. Some of the results in this paper have been derived using HEALPix (Górski et al., 2005). Public-domain cartographic data is courtesy of Natural Earth¹¹ and processed with MapShaper¹².

BAYESTAR, LALINFERENCE_NEST, and LALINFERENCE_MCMC are part of the LIGO Algorithm Library Suite¹³ and the LIGO parameter estimation package, LALINFERENCE. Source code for GSTLAL¹⁴ and LALINFERENCE¹⁵ are available online under the terms of the GNU General Public License.

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the NSF and operates under cooperative agreement PHY-0757058.

⁹<http://ciera.northwestern.edu/Research/GraIl.Cluster.php>

¹⁰<http://www.astropy.org>

¹¹<http://www.naturalearthdata.com>

¹²<http://www.mapshaper.org>

¹³<http://www.lsc-group.phys.uwm.edu/cgit/lalsuite/tree>

¹⁴<http://www.lsc-group.phys.uwm.edu/cgit/gstlal/tree/>

¹⁵<http://www.lsc-group.phys.uwm.edu/cgit/lalsuite/tree/lalinference>

Chapter 6

Discovery and redshift of an optical afterglow in 71 square degrees: iPTF13bxl and GRB 130702A

LEO P. SINGER¹, S. BRADLEY CENKO², MANSI M. KASLIWAL^{3,13}, DANIEL A. PERLEY^{13,4},
ERAN O. OFEK⁵, DUNCAN A. BROWN^{1,6}, PETER E. NUGENT^{7,8}, S. R. KULKARNI⁴, ALESSANDRA CORSI⁹,
DALE A. FRAIL¹⁰, ERIC BELL⁴, JOHN MULCHAEY³, IAIR ARCAVI⁵, TOM BARLOW⁴,
JOSHUA S. BLOOM^{7,8}, YI CAO⁴, NEIL GEHRELS², ASSAF HORESH⁴, FRANK J. MASCI¹¹,
JULIE MCENERY², ARNE RAU¹², JASON A. SURACE¹¹, AND OFER YARON⁵

¹LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

²Astrophysics Science Division, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA

³Observatories of the Carnegie Institution for Science, 813 Santa Barbara St, Pasadena CA 91101, USA

⁴Cahill Center for Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

⁵Benoziyo Center for Astrophysics, The Weizmann Institute of Science, Rehovot 76100, Israel

⁶Department of Physics, Syracuse University, Syracuse, NY 13244, USA

⁷Department of Astronomy, University of California Berkeley, B-20 Hearst Field Annex # 3411, Berkeley, CA, 94720-3411

⁸Physics Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road MS 50B-4206, Berkeley, CA 94720, USA

⁹George Washington University, Corcoran Hall, Washington, DC 20052, USA

¹⁰National Radio Astronomy Observatory, P.O. Box O, Socorro, NM 87801, USA

¹¹Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA

¹²Max-Planck-Institut für extraterrestrische Physik, Giessenbachstrasse 1, 85748 Garching, Germany

¹³Hubble Fellow

This chapter is reproduced from [Singer et al. \(2013\)](#), which was published under the same title in The Astrophysical Journal Letters, copyright © 2014 The American Astronomical Society. My contributions include developing the P48 alert and tiling code, reducing the P48 photometry, assigning candidates for follow-up, analyzing the data, preparing

all of the figures, and writing about 70% of the text. S.B.C. and D.A.B. also contributed to candidate vetting. S.B.C. calculated the broadband afterglow models shown in Figure 6.4. D.A.P., A.C., and A.H. executed and reduced the radio observations. D.A.P. reduced the P60 photometry and wrote much of Section 6.5. The Magellan Baade spectrum was acquired by M.M.K. and reduced by J.M. The P200 spectrum was observed and reduced by E.B. Other authors made substantial contributions to the iPTF hardware and software.

6.1 Introduction

Our understanding of gamma-ray bursts (GRBs) has been propelled by our ability to localize these rare and energetic cosmic events precisely. *Compton Gamma-ray Observatory*/BATSE’s coarse localizations robustly demonstrated that GRBs were distributed isotropically on the sky and suggested that GRBs originate at cosmological distances (Meegan et al., 1992). Prompt arcminute localizations provided by *BeppoSAX* directly enabled the discovery of the first afterglows of long-duration GRBs (Costa et al., 1997; van Paradijs et al., 1997; Frail et al., 1997). Currently, the prompt slewing capabilities of the *Swift* satellite (Gehrels et al., 2004) enable the on-board narrow-field instruments to provide arcsecond localizations for ≈ 90 GRBs yr^{-1} within ≈ 100 s of the burst trigger.

With seven decades of simultaneous energy coverage, *Fermi* has opened a new window into the GRB phenomenon, the MeV to GeV regime. However, *Fermi* remains fundamentally limited by its localization capabilities. The Large Area Telescope (LAT; 20 MeV–300 GeV; 16% of all-sky; Atwood et al., 2009) can localize events with GeV photons to radii as small as $\sim 10'$. But the LAT only localizes a handful of GRBs each year. The Gamma-ray Burst Monitor (GBM; few keV–30 MeV; 70% of all-sky; Meegan et al. 2009b), on the other hand, detects GRBs at a rate of $\approx 250 \text{ yr}^{-1}$. However, typical GBM GRBs have localizations of many tens of square degrees (random plus systematic uncertainties). Consequently, no afterglows have been identified based solely on a GBM localization until this work¹.

The Palomar Transient Factory (PTF; Law et al. 2009) is developing the necessary instrumentation, algorithms, and observational discipline to detect optical counterparts to GBM GRBs. The wide 7.1 deg^2 field of view and sensitivity ($R \approx 20.6$ mag in 60 s) of the P48 and CFH12k cam-

¹The only comparable discovery was the afterglow of GRB 120716A in the $\approx 2 \text{ deg}^2$ error box from the IPN by Cenko et al. (2012).




Figure 6.1 P48 imaging of GRB 130702A and discovery of iPTF13bxi. The left panel illustrates the γ -ray localizations (red circle: 1σ GBM; green circle: LAT; blue lines: 3σ IPN) and the 10 P48 reference fields that were imaged (light gray rectangles). For each P48 pointing, the location of the 11 chips are indicated with smaller rectangles (one CCD in the camera is not currently operable). Our tiling algorithm places a large weight on the existence of deep P48 pre-explosion imaging (a necessity for high-quality subtraction images); the large gaps inside the GBM localization are fields without these reference images. The small black diamond is the location of iPTF13bxi. The right panels show P48 images of the location of iPTF13bxi, both prior to (top) and immediately following (bottom) discovery. We note that the LAT and IPN localizations were published *after* our discovery announcement (Singer et al., 2013).

era (Rahmer et al., 2008) are well-suited to identifying long-duration GRB afterglow candidates. The real-time software pipeline (Nugent et al., in prep.) enables rapid panchromatic follow-up with an arsenal of telescopes (e.g. Gal-Yam et al. 2011), essential to distinguish the true afterglow from background and foreground contaminants. Here, we present our discovery of iPTF13bxl, the afterglow of the *Fermi* GBM GRB 130702A, found by searching a sky area of 71 deg^2 with the Intermediate Palomar Transient Factory (iPTF).

6.2 Discovery

On 2013 July 2 at 00:05:23.079 UT, the *Fermi* GBM detected trigger 394416326. The refined human-generated (i.e., ground-based) localization, centered on $\alpha = 14^{\text{h}}35^{\text{m}}14^{\text{s}}$, $\delta = 12^{\circ}15'00''$ (J2000.0), with a quoted 68% containment radius of 4.0 (statistical uncertainty only), was disseminated less than an hour after the burst (Figure 6.1).

Fermi-GBM GRB positions are known to suffer from significant systematic uncertainties, currently estimated to be $\approx 2^{\circ}\text{--}3^{\circ}$. To characterize the full radial profile of the localization uncertainty, our GBM-iPTF pipeline automatically computed a probability map for the event, modeled on previous *Fermi*/*Swift* coincidences from 2010 March 30 through 2013 April 4. We fit a sigmoid function:

$$P(r) = \frac{1}{1 + (c_0 r)^{c_1}}, \quad (6.1)$$

where r is the angular distance to the *Swift* location, normalized by the in-flight or ground-based error radius for that burst. We find $c_0 = 1.35$, $c_1 = -2.11$ for in-flight GBM localizations and $c_0 = 0.81$, $c_1 = -2.47$ for ground-based GBM localizations (Figure 6.2).

Image subtraction within iPTF is greatly simplified by observing only pre-defined fields on the plane of the sky; this ensures that sources will fall on approximately the same detector location from night to night, minimizing a possible source of systematic uncertainty. Using a Hierarchical Equal Area isoLatitude Pixelization (HEALPix; Górski et al., 2005) bitmap representation of the probability distribution of the trigger location, we chose ten of these pre-defined fields to maximize the probability of enclosing the true (but still unknown) location of the burst (Figure 6.1). In this particular case, the ten selected fields did not include the center of the GBM localization because we lacked previous reference images there. Nonetheless, we estimated that these ten fields had a




Figure 6.2 Localization accuracy of *Fermi* GBM positions, generated by searching for coincidences with GRBs detected by the *Swift* satellite. The left panel shows the fraction of bursts contained within a given distance from the *Swift* position, both for in-flight and refined ground-based localizations. Ground-based localizations are on average about half as far from the true GRB positions as the in-flight localizations. The right panel shows a cumulative histogram of the *Fermi*–*Swift* distance, normalized by each trigger’s nominal 1σ radius (either ground-based or in-flight). Although the ground-based localizations are more accurate, the nominal ground-based error radii must be interpreted as describing a different confidence level than the in-flight error radius. The thick gray lines are fits to the logistic-like function in Equation 6.1.

38% chance of containing this GRB’s location. Given the youth, sky location, and probability of containment, we let our software trigger follow-up with the P48.

Starting at 04:17:23 UT ($\Delta t \equiv t - t_{\text{GBM}} = 4.2$ hr), we imaged each of these ten fields twice in 60 s exposures with the Mould *R* filter. These fields were then subjected to the standard iPTF transient search: image subtraction, source detection, and “real/bogus” machine ranking (Bloom et al., 2012; Brink et al., 2013) into likely astrophysical transient sources (“real”, or 1) or likely artifacts (“bogus”, or 0).

The iPTF real-time analysis found 27,004 transient/variable candidates in these twenty individual subtracted images. Of these, 44 were eliminated because they were determined to be known asteroids in the Minor Planet Checker database² using PyMPCChecker.³ Demanding a real/bogus rank greater than 0.1 reduced the list to 4214. Rejecting candidates that coincided with point sources in Sloan Digital Sky Survey (SDSS) brighter than $r' = 21$ narrowed this to 2470. Further, requiring detection in both P48 visits and imposing CCD-wide data quality cuts (e.g.,

²<http://www.minorplanetcenter.org/iau/mpc.html>

³<http://dotastro.org/PyMPC/PyMPC/>

bad pixels) eliminated all but 43 candidates. Following human inspection, seven sources were saved as promising transients in the iPTF database.

Two candidates, iPTF13bxh and iPTF13bxu, were near the cores of bright, nearby galaxies, environments that are inherently variable and also present a challenge to image subtraction. A third, iPTF13bxr, was coincident with a galaxy in SDSS with a quasar spectrum (SDSS J145359.72+091543.3). iPTF13bxt was close to a star in SDSS, and so was deemed a likely variable star. We did not consider these further for the afterglow search. The final three sources, iPTF13bxj (real-bogus score = 0.86), iPTF13bxk (real-bogus score = 0.49), and iPTF13bxl (real-bogus score = 0.83), remained as potential counterparts and were scheduled for $g'r'i'$ photometry with the robotic Palomar 60-inch telescope (P60; Cenko et al., 2006) and spectroscopic classification on the P200. iPTF13bxl ($\alpha = 14^{\text{h}}29^{\text{m}}14^{\text{s}}.78, \delta = +15^{\circ}46', 26''.4$) was immediately identified as the most promising candidate because it showed a significant intra-night decline. Our panchromatic follow-up (Sections 6.3 and 6.4) confirmed iPTF13bxl was indeed the afterglow. Subsequent spectroscopy revealed iPTF13bxj to be a Type II supernova at $z = 0.06$ and iPTF13bxk a quasar at $z = 2.4$.

Following our discovery announcement (Singer et al., 2013), the *Fermi* LAT and GBM teams published GCN circulars announcing the detection of GRB 130702A (Cheung et al., 2013; Collazzi & Connaughton, 2013). As seen by the GBM, GRB 130702A had a duration of $t_{90} \approx 59$ s and a 10 keV–1 MeV fluence of $f_{\gamma} = (6.3 \pm 2.0) \times 10^{-6}$ erg cm $^{-2}$. The best-fit power-law spectrum may suggest a classification as an X-ray flash. The LAT location was $0^{\circ}.9$ from iPTF13bxl, with a 90% statistical error radius of $0^{\circ}.5$. An IPN triangulation (Hurley et al., 2013) yielded a $3\text{-}\sigma$ annulus that was $0^{\circ}.46$ wide from its center to its edges. iPTF13bxl was $0^{\circ}.16$ from the annulus' centerline (Figure 6.1).

6.3 Broadband photometric follow-up

On 2013 July 3 at 4:10 UT, ($\Delta t = 28.1$ hr), the P60 obtained two sequences of Sloan $g'r'i'$ observations of the field of iPTF13bxl. P60 observations were calibrated relative to 20 reference stars in the SDSS (AB) system. Final reduction of the P48 observations was performed automatically at the Infrared Processing and Analysis Center (Ofek et al., 2012). We corrected the P48 and P60 photometry for Galactic extinction using maps from Schlafly & Finkbeiner (2011,

$$E(B - V) = 0.024 \text{ mag}.$$

The optical light curve is shown in Figure 6.3. We fit an achromatic broken power law to all filters and all times up to $\Delta t = 5$ days after the burst. A spectral slope of $\beta_{\text{O}} = 0.7 \pm 0.1$ is sufficient to characterize the post-break color, illustrated in the inset of Figure 6.4. We note that the optical decay ceased at $r' \approx 20$ mag after $\Delta t \approx 5$ days when the accompanying supernova started to dominate (Schulze et al., 2013). This supernova will be the subject of a future work.

Following our discovery of iPTF13bxl, we triggered TOO observations with the *Swift* XRT (Burrows et al., 2005) beginning at 00:50 UT on 2013 July 3 ($\Delta t = 1.03$ days). We downloaded the data products from the *Swift* XRT repository (Evans et al., 2007). The resulting 0.3–10 keV light curve is plotted in Figure 6.3. The spectrum is well fit by a power law with photon index $\Gamma = 2.0^{+0.14}_{-0.13}$, while the light curve fades in time with a power-law slope of $\alpha_{\text{X}} = 1.06 \pm 0.02$, in excellent agreement with the post-break optical decay.

After the discovery of the optical counterpart to GRB 130702A, we began observations with the Combined Array for Research in Millimeter-wave Astronomy (CARMA). All observations were carried out in single-polarization mode with the 3 mm receivers tuned to a frequency of 93 GHz, and were reduced using MIRIAD. We flux-calibrated the data using observations of MWC349 and 3C273. The afterglow was well-detected in both epochs, and we measured flux densities of 1.58 ± 0.33 mJy and 1.85 ± 0.30 mJy on July 4.13 and 5.17, respectively.

The position of iPTF13bxl was observed with the Karl G. Jansky Very Large Array (VLA) in C-band beginning at 6:20 UT on 2013 July 4 ($\Delta t = 2.29$ days). The observations were conducted using the standard WIDAR correlator setting. Data were reduced using the Astronomical Image Processing System package following standard practice. 3C286 was used for bandpass and flux calibration; J1415+1320 was used for gain and phase calibration. We detected a radio source with flux density of 1.49 ± 0.08 mJy at 5.1 GHz and 1.60 ± 0.08 mJy at 7.1 GHz. Errors on the measured flux were calculated as the quadrature sum of the map root-mean square and a fractional systematic error (of the order of 5%) to account for uncertainty in the flux density calibration.

The broadband spectral energy distribution (SED) is shown in Figure 6.4. We interpolated both the optical and X-ray observations to the mean time of the VLA and CARMA observations. In the context of the standard synchrotron afterglow model (Sari et al., 1998), the comparable X-ray and optical spectral and temporal indices at this time suggest both bandpasses lie in the same spectral regime, likely with $\nu > \nu_c$. This would imply a relatively hard electron spectral




Figure 6.3 P48, P60, and X-ray Telescope (XRT) light curves of iPTF13bxi. The broken power-law fit is shown up to $\Delta t = 5$ days. The XRT observations, re-binned to improve presentation, are shown in gray as $m(\text{AB}) - 6.5$ at 1 keV. A timeline in the top right puts the P48 and P60 observations in the context of the XRT follow-up, PTF's discovery GCN (Singer et al., 2013), the announcement of the LAT (Cheung et al., 2013) and IPN (Hurley et al., 2013) localizations, and the radio observations.

energy index ($N(\gamma_e) \propto \gamma_e^{-p}$) $p \approx 2$, possibly requiring a high-energy cut-off.

Also in Figure 6.4 we plot three broadband SED models synthesized using techniques similar to Perley et al. (2014c). Although these models are not formal fits to our highly under-constrained observations, they demonstrate overall consistency with standard synchrotron afterglow theory. Model “A” (dashed line; $\chi^2 = 126$) represents a constant-density (ISM) circumburst medium with $p = 2.1$, $\epsilon_B = 0.48$, $\epsilon_e = 0.41$, $E = 3 \times 10^{51}$ erg, $n = 1.2 \times 10^{-3}$ cm $^{-3}$. This model under-predicts the VLA bands, but this deviation could be due to scintillation or reverse shock emission at low frequencies. Model “B” (dotted line; $\chi^2 = 7$) is in a wind environment ($\rho \propto r^{-2}$) with $p = 2.1$, $\epsilon_B = 0.32$, $\epsilon_e = 0.43$, $E = 1.4 \times 10^{51}$ erg, $A_* = 4.8 \times 10^{-3}$ g cm $^{-1}$. This fits the data well except for a small discrepancy with the optical spectral slope. Model “C” (dotted-dashed line; $\chi^2 = 6$) is a similar wind model but with $p = 1.55$. Of the three, “C” fits the data best, but seems non-physical (high-energy cutoff required). Accurate determination of the underlying physical parameters would require tracing the evolution of the SED with time.

6.4 Optical spectroscopy and host galaxy environment

We obtained optical spectra of iPTF13bxi with the Double Spectrograph (DBSP) mounted on the P200 on 2013 July 3.17 and the Inamori-Magellan Areal Camera & Spectrograph (IMACS; Dressler et al., 2011) mounted on the 6 m Magellan Baade telescope on 2013 July 3.97 ($\Delta t = 1.2$ and 2.0 days, respectively). The resulting spectra are plotted in Figure 7.6.

Our initial DBSP spectrum exhibits a largely featureless, blue continuum. The higher SNR of the IMACS spectrum further reveals faint, narrow emission lines corresponding to [O III] and H α at a common redshift of $z = 0.145 \pm 0.001$ (luminosity distance $d_L = 680$ Mpc), which we adopt as the distance to GRB 130702A. The continuum of both spectra exhibit deviations from a single power-law, with excess flux (when compared with the late-time photometric spectral index of $\beta_O = 0.7$) visible at shorter wavelengths. This may suggest some contribution from either shock breakout or the emerging supernova at very early times post-explosion.

Three galaxies are visible in the immediate environment of iPTF13bxi in pre-explosion imaging (labeled “G1”, “G2”, and “G3” in Figure 6.1). Presumably the emission lines observed from iPTF13bxi arise in G1, given the small spatial offset ($0''.6$) and slit orientation (PA = 90). However, our spectra also reveal that galaxies G2 and G3 both lie at redshifts consistent with iPTF13bxi (e.g.,

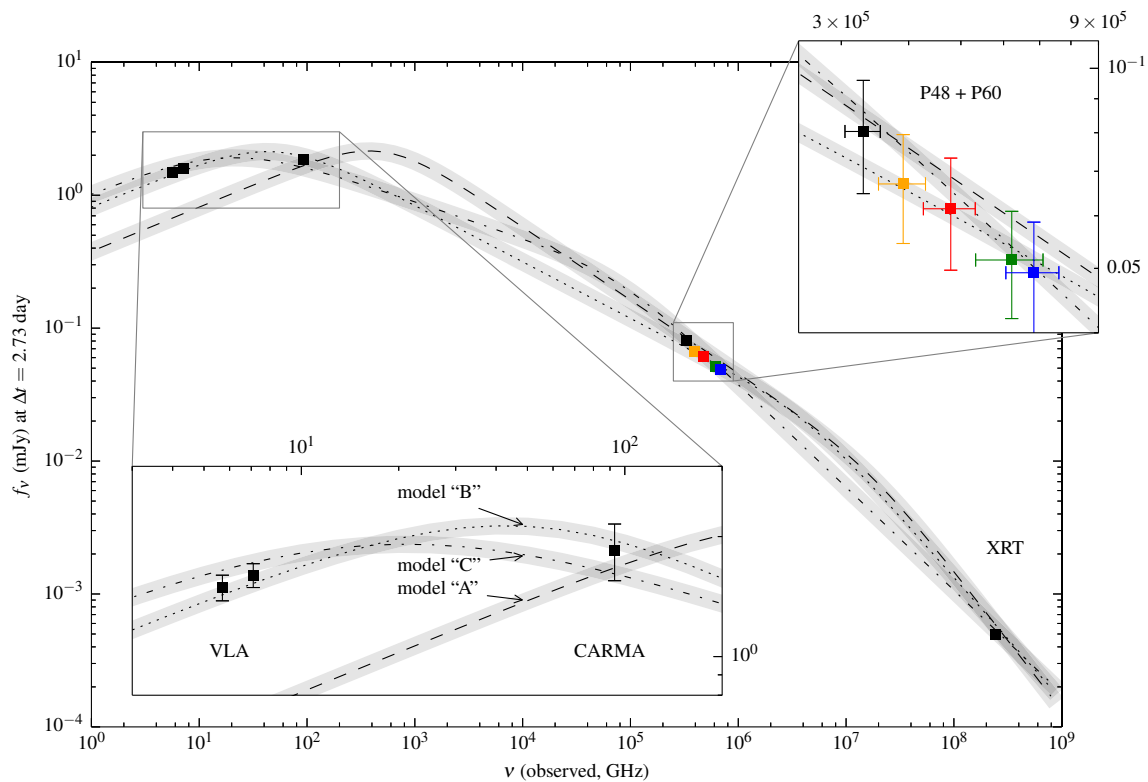


Figure 6.4 Broadband SED of iPTF13bxi. Two insets show details of the radio and optical observations respectively. The XRT and optical observations have been interpolated to the mean time of the radio observations ($\Delta t = 2.6$ days).




Figure 6.5 Optical spectra of iPTF13bxi and the nearby galaxy SDSS J142914.57+154619.3 (“G2”). Spectra in the left panel have been smoothed with a Savitzky-Golay filter. Our initial P200 spectrum of the afterglow (left panel, blue) exhibits a largely featureless blue continuum. A higher SNR spectrum taken the following night with IMACS (left panel, green) revealed faint emission features corresponding to [O III] and H α at $z = 0.145$ (top right panel). The bottom right panel shows a spectrum of the nearby galaxy G2, which has the same redshift as iPTF13bxi.

Figure 7.6). Observations with DEIMOS on the Keck II telescope reveal two more galaxies at the same redshift at separations of $1/2$ (SDSS J142910.29+154552.2) and $2/7$ (SDSS J142917.67+154352.2) from the transient. The explanation most consistent with past observations of long-duration GRB host galaxies (e.g., Savaglio et al. 2009) is that GRB 130702A exploded in a dwarf ($M_r \approx -16$ mag) member of this association or group, a relatively unusual environment (see also Kelly et al., 2013).

6.5 GRB 130702A in context

The prompt γ -ray isotropic energy release ($E_{\gamma, \text{iso}}$) of GRBs spans a range of six orders of magnitude, from $\sim 10^{48}$ – 10^{54} erg. At $z = 0.145$, the prompt emission from GRB 130702A constrains $E_{\gamma, \text{iso}} \lesssim (6.5 \pm 0.1) \times 10^{50}$ erg (90% upper limit; Amati et al., 2013). This value is significantly smaller than typical cosmological GRBs ($E_{\gamma, \text{iso}} \sim 10^{52}$ – 10^{54} erg; Amati 2006; Butler et al. 2007). Yet GRB 130702A greatly outshines the most nearby, sub-luminous events with well-studied supernovae, such as GRB 980425 ($E_{\gamma, \text{iso}} = 1.0 \times 10^{48}$ erg; Pian et al. 2000) and GRB 060218 ($E_{\gamma, \text{iso}} = 6.2 \times 10^{49}$ erg; Campana et al. 2006).

At all wavelengths, the counterpart behaves like a typical GRB afterglow scaled down in luminosity by a factor of ~ 10 compared to a “typical” *Swift* burst (or ~ 100 lower than a luminous pre-*Swift* burst) as observed at the same time. This is intuitively explained by an equivalent scaling down of the overall energy (per solid angle) of the burst and shockwave relative to more typical, high-luminosity bursts. It is not yet clear whether this energy difference is due primarily to the release of less relativistic ejecta by the burst overall, a wider jet, or a partially off-axis view of a structured jet. Late-time radio follow-up should help distinguish these models: an intrinsically low-energy GRB should produce a much earlier jet break than a widely-beamed burst, while a structured jet will actually produce an *increase* in flux at late times as the jet core spreads and its radiation enters our sightline.

Events with similar energetics have been found by *Swift*, e.g., GRB 050826 at $z = 0.30$ and GRB 120422A at $z = 0.28$ (Mirabal et al., 2007; Zhang et al., 2012). However, given their low intrinsic luminosities and higher redshift, the afterglows were too faint to identify late-time breaks and establish their shock energies E_K , making them difficult to physically interpret. GRB 130702A’s proximity avoids both these problems. Our observations suggest—and further observations should confirm—that its γ -ray and afterglow energetics are intermediate between these two previously quite-disparate classes of GRBs, helping to fill in the “gap” between the well-studied cosmological population and the class of less-luminous local GRBs and relativistic Type Ic supernovae (e.g., Soderberg et al. 2004, 2010).

6.6 Conclusion

Using the infrastructure outlined above, we estimate that a dedicated iPTF program would recover ~ 10 GBM afterglows each year. The addition of other surveys with comparably wide fields of view and apertures (e.g., Pan-STARRS, SkyMapper, CRTS) could increase this number, assuming they had similar real-time transient detection and follow-up programs in place. Since GBM detects GRBs in the classical γ -ray band, their optical counterparts should more closely resemble the pre-*Swift* population (≈ 1 mag brighter at a fixed time; Kann et al. 2010). Even if only a single event per year as nearby as GRB 130702A were uncovered, this would still represent a remarkable advance in our understanding of the GRB–supernova connection.

Furthermore, this work sets the stage for more discoveries in ongoing and future physics

experiments that are limited by similarly coarse position reconstruction. Later this decade, a network of advanced GW detectors including the Laser Interferometer GW Observatory (LIGO) and Virgo is expected to detect $\sim 0.4\text{--}400$ binary neutron star mergers per year (Abadie et al., 2010b), but with positions uncertain to tens to hundreds of deg^2 (Fairhurst, 2009; Nissanke et al., 2011; Aasi et al., 2013c).

Optical counterparts to GW sources will rarely (due to jet collimation) include bright, on-axis short-hard burst afterglows. Fainter r -process-fueled kilonovae (Li & Paczyński, 1998) or yet fainter off-axis afterglows (Rhoads, 1997) are expected to accompany binary neutron star mergers. Both of these signatures are predicted to be several magnitudes fainter than iPTF13bxl. Optical searches will be inundated with astrophysical false positives (Nissanke et al., 2013). This problem will only be exacerbated for future surveys covering larger areas (e.g., Zwicky Transient Facility; Kulkarni 2012) and/or with larger apertures (e.g., Large Synoptic Survey Telescope; Ivezić et al. 2008). However, a breathtakingly complete astrophysical picture could reward us: masses and spins measured in GWs; host galaxy and disruption ejecta in optical; circumstellar environment in radio. The case of GRB 130702A demonstrates for the first time that optical transients can be recovered from localization areas of $\sim 100\text{ deg}^2$, reaching a crucial milestone on the road to Advanced LIGO.

Acknowledgements

Optical photometry and spectroscopy referred to in this work will be made available via WISeREP⁴ (Yaron & Gal-Yam, 2012).

We acknowledge A. Weinstein, A. Gal-Yam, R. Quimby, V. Connaughton, and the *Fermi*-GBM team for valuable discussions, S. Caudill, S. Tinyanont, D. Khatami for P200 observing, and the developers of the COSMOS package for Magellan data reduction.

This research is supported by the NSF through a Graduate Research Fellowship for L.P.S., award PHY-0847611 for D.A.B., and NSF-CDI grant 0941742 for J.S.B. M.M.K. acknowledges generous support from the Carnegie-Princeton Fellowship. M.M.K. and D.A.P. are supported by NASA through the Hubble Fellowship grants HST-HF-51293.01 and HST-HF-51296.01-A, awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for

⁴<http://www.weizmann.ac.il/astrophysics/wiserep/>

Research in Astronomy, Inc., for NASA, under contract NAS 5-26555. E.O.O. is the incumbent of the Arye Dissentshik career development chair and is supported by grants from the Israeli Ministry of Science and the I-CORE Program. D.A.B. is further supported by an RCSA Cottrell Scholar award.

This research made use of Astropy⁵ (Robitaille et al., 2013), a community-developed core Python package for Astronomy. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

⁵<http://www.astropy.org>

Chapter 7

Fermi, iPTF, and the GRB–supernova connection

LEO P. SINGER¹, MANSI M. KASLIWAL², S. BRADLEY CENKO^{3,4}, DANIEL A. PERLEY^{27,5},
GEMMA E. ANDERSON^{6,7}, G. C. ANUPAMA⁸, IAIR ARCAVI^{9,10}, VARUN BHALERAO¹¹, BRIAN D. BUE¹²,
YI CAO⁵, VALERIE CONNAUGHTON¹³, ALESSANDRA CORSI¹⁴, ANTONINO CUCCHIARA^{28,3},
ROB P. FENDER^{6,7}, NEIL GEHRELS³, ADAM GOLDSTEIN^{28,15}, ASSAF HORESH¹⁶, KEVIN HURLEY¹⁷,
JOEL JOHANSSON¹⁸, D. A. KANN^{19,20}, CHRYSA KOUVELIOTOU¹⁵, KUIYUN HUANG²¹, S. R. KULKARNI⁵,
FRANK MASCI²², PETER NUGENT^{23,24}, ARNE RAU²⁰, UMAA D. REBBAPRAGADA¹², TIM D. STALEY^{6,7},
DMITRY SVINKIN²⁵, YUJI URATA²⁶, AND ALAN WEINSTEIN¹

¹LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

²Observatories of the Carnegie Institution for Science, 813 Santa Barbara St, Pasadena CA 91101, USA

³Astrophysics Science Division, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA

⁴Joint Space-Science Institute, University of Maryland, College Park, MD 20742, USA

⁵Cahill Center for Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

⁶Astrophysics, Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

⁷Physics & Astronomy, University of Southampton, Southampton SO17 1BJ, UK

⁸Indian Institute of Astrophysics, Koramangala, Bangalore 560 034, India

⁹Las Cumbres Observatory Global Telescope Network, 6740 Cortona Dr., Suite 102, Goleta, CA 93117, USA

¹⁰Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, USA

¹¹Inter-University Centre for Astronomy and Astrophysics (IUCAA), Post Bag 4, Ganeshkhind, Pune 411007, India

¹²Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

¹³CSPAR and Physics Department, University of Alabama in Huntsville, 320 Sparkman Dr., Huntsville, AL 35899, USA

¹⁴Texas Tech University, Physics Department, Lubbock, TX 79409-1051

¹⁵Astrophysics Office, ZP12, NASA Marshall Space Flight Center, Huntsville, AL 35812, USA

¹⁶Benoziyo Center for Astrophysics, Weizmann Institute of Science, 76100 Rehovot, Israel

¹⁷Space Sciences Laboratory, University of California-Berkeley, Berkeley, CA 94720, USA

¹⁸The Oskar Klein Centre, Department of Physics, Stockholm University, SE 106 91 Stockholm, Sweden

¹⁹Thüringer Landessternwarte Tautenburg, Sternwarte 5, 07778 Tautenburg, Germany

²⁰Max-Planck Institut für Extraterrestrische Physik, Giessenbachstr. 1, 85748 Garching, Germany

²¹Department of Mathematics and Science, National Taiwan Normal University, Lin-kou District, New Taipei City 24449, Taiwan

²²Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA

²³Department of Astronomy, University of California, Berkeley, CA 94720-3411, USA

²⁴Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

²⁵Ioffe Physical-Technical Institute, Politekhnicheskaya 26, St Petersburg 194021, Russia

²⁶Institute of Astronomy, National Central University, Chung-Li 32054, Taiwan

²⁷Hubble Fellow

²⁸NASA Postdoctoral Fellow

This chapter is reproduced from a paper to be titled “The Needle in the 100 deg² Haystack: Uncovering Afterglows of Fermi GRBs with the Palomar Transient Factory,” in preparation for The Astrophysical Journal.

7.1 Introduction

Deep synoptic optical surveys including the Palomar Transient Factory (PTF; [Rau et al. 2009](#); [Law et al. 2009](#)) and Pan-STARRS ([Kaiser et al., 2010](#)) have revealed a wealth of new transient and variable phenomena across a wide range of characteristic luminosities and time scales ([Kasliwal, 2011](#)). With a wide (7 deg²) instantaneous field of view, moderately deep sensitivity (reaching $R = 20.6$ mag in 60 s), a consortium of follow-up telescopes, sophisticated image subtraction and machine learning pipelines, and an international team of human-in-the-loop observers, PTF has been a wellspring of new or rare kinds of explosive transients (for instance, [Kasliwal et al. 2012](#); [Quimby et al. 2011](#)) and early-time observations of SNe or their progenitors (see, for example, [Nugent et al. 2011](#); [Corsi et al. 2012](#); [Ofek et al. 2013](#); [Gal-Yam et al. 2014](#)). PTF has even blindly detected the optical emission ([Cenko et al. 2014](#); Cenko et al., in preparation) from the rarest, brightest, and briefest of all known cosmic explosions, GRBs, hitherto only discoverable with the aid of precise localizations from space-based gamma-ray observatories. PTF has also detected explosions that optically resemble GRB afterglows but may entirely lack gamma-ray emission ([Cenko et al., 2013b](#)).

GRBs and their broadband afterglows are notoriously challenging to capture. They naturally evolve from bright to faint, and from high (gamma- and hard X-ray) to low (optical and radio)

photon energies, with information encoded on energy scales from 1 to 10^{16} GHz (Perley et al., 2014c) and time scales from 10^{-3} to 10^7 s. Only with a rapid sequence of handoffs between facilities graded by energy passband, field of view, and position accuracy, have we been able to find them, pinpoint their host galaxies, and constrain their physics. The *Swift* mission (Gehrels et al., 2004), with its 1.4 sr-wide (50% coded) Burst Alert Telescope (BAT; Barthelmy et al. 2005) and its ability to slew and train its onboard X-ray Telescope (XRT; Burrows et al. 2005) and UV/Optical Telescope (UVOT; Roming et al. 2005) on the location of a new burst within 100 s, has triumphed here: in 9 years of operation, it has tracked down ≈ 700 X-ray afterglows and enabled extensive panchromatic observations by a worldwide collaboration of ground-based optical and radio facilities.

Meanwhile, the *Fermi* satellite has opened up a new energy regime extending up to 300 GeV, with the Large Area Telescope (LAT; (Atwood et al., 2009)) detecting high-energy photons for about a dozen bursts per year. The Gamma-ray Burst Monitor (GBM; Meegan et al. 2009a), an all-sky instrument sensitive from 8 keV to 40 MeV, detects GRBs prolifically at a rate of $\approx 250 \text{ yr}^{-1}$, with a large number (about 44 yr^{-1}) belonging to the rarer short, hard bursts (Paciesas et al., 2012a). Although LAT can provide localizations that are as accurate as $\sim 10'$, *Fermi* GBM produces error circles that are several degrees across. Consequently, most *Fermi* bursts do not receive deep, broadband follow-up, with the properties of their afterglows largely unknown.

As part of the iPTF, over the past year we have developed the ability to rapidly tile these $\sim 100 \text{ deg}^2$ GBM error circles and pinpoint the afterglows. This TOO capability uses and briefly redirects the infrastructure of the ongoing synoptic survey, notably the machine learning software and the instrumental pipeline composed of the Palomar 48-inch Oschin telescope (P48; Rahmer et al. 2008), the robotic Palomar 60-inch telescope (P60; Cenko et al. 2006), and associated spectroscopic resources including the Palomar 200-inch Hale telescope (P200).

In Singer et al. (2013), we announced the first discovery of an optical afterglow based solely on a *Fermi* GBM localization.¹ That explosion, GRB 130702A / iPTF13bx1, was noteworthy for several reasons. First, it was detected by *Fermi* LAT. Second, it was at moderately low redshift, $z = 0.145$, yet had prompt energetics that bridged the gap between “standard”, bright cosmically

¹There are two earlier related cases. The optical afterglow of GRB 090902B was detected ex post facto in tiled observations with Robotic Optical Transient Search (ROTSE) about 80 min after the burst, but the afterglow was initially discovered with the help of an X-ray detection in *Swift* observations of the LAT error circle. GRB 120716A was identified by iPTF by searching a $\approx 2 \text{ deg}^2$ IPN error box (Cenko et al., 2012).

distant bursts and nearby sub-luminous bursts and X-ray flashes (XRFs). Third, due to its low redshift, the accompanying SN was spectroscopically detectable.

In this work, we begin with a detailed description of the operation of the iPTF GRB afterglow search. We then present seven more GBM–iPTF afterglows from the first 13 months of this project. In all eight cases, the association between the optical transient and the GRB was proven by the presence of high-redshift absorption lines in the optical spectra and the coincident detection of a rapidly fading X-ray source with *Swift* XRT. In two cases, the positions were further corroborated by accurate *Fermi* LAT error circles, and in four cases by accurate IPN triangulations involving distant spacecraft. In one case (GRB 140508A), the IPN triangulation was performed rapidly and was instrumental in selecting which optical transient candidates to follow up. In six cases, radio afterglows were detected. Our discovery rate of eight out of 35 events is consistent with the ages and searched areas of the GBM bursts, combined with the luminosity function of optical afterglows. Consequently, by tiling larger areas and/or stacking exposures, the iPTF afterglow search should be able to scale to more coarse localizations, such as those associated with short GRBs.

Next, we present extensive follow-up observations, including *R*-band photometry from the P48, multicolor photometry from the P60, spectroscopy (acquired with the P200, Keck, Gemini, APO, Magellan, and VLT), and radio observations with the Karl G. Jansky Very Large Array² (VLA), the Combined Array for Research in Millimeter-wave Astronomy (CARMA; [Bock et al. 2006](#); [Corder et al. 2010](#)), the Australia Telescope Compact Array (ATCA; [Frater et al. 1992](#)), and the Arcminute Microkelvin Imager (AMI; [Zwart et al. 2008](#)). We provide basic physical interpretations of the broadband SEDs of these afterglows. We find that seven of the events are consistent with the standard model of synchrotron cooling of electrons that have been accelerated by a single forward shock encountering either the constant-density circumburst interstellar medium (ISM; broadband behavior predicted in [Sari et al. 1998](#)) or a stellar (i.e., Wolf-Rayet) wind environment ([Chevalier & Li, 1999](#)). The exception, GRB 140620A / iPTF14cva, may be explained by an additional reverse shock or an inverse Compton component.

Two of the afterglows (GRB 130702A / iPTF13bxi and GRB 140606B / iPTF14bfu) faded away to reveal spectroscopically detected broad-line type Ic SNe (SNe Ic-BL). Despite the abundant photometric evidence for SNe in afterglow light curves (see [Li & Hjorth 2014](#) and references

²<http://www.vla.nrao.edu>

therein), the distinction of SN spectroscopy has been shared by scarcely tens³ out of ≈ 800 long *Swift* bursts in nine years of operation.

We estimate the kinetic energies of the relativistic blast waves for each burst from their X-ray afterglows (Freedman & Waxman, 2001). We find that although the gamma-ray energetics of these eight bursts are broadly similar to the *Swift* sample, two low-luminosity bursts (GRBs 130702A and 140606B) have significantly lower kinetic energies. We discuss the possibility that these two bursts arise not from a standard ultra-relativistic internal shock, but from a mildly relativistic shock as it breaks out from the progenitor star (see, for example, Nakar & Sari 2012).

We conclude by discussing prospects for targeted optical transient searches in wide areas. This is especially relevant for optical counterparts of gravitational wave events. We illustrate that optical afterglows of short bursts, which are intimately linked to the prime sources for the Advanced Laser Interferometer GW Observatory (LIGO) and Virgo, should be well within the reach of a similar approach using ZTF (Kulkarni, 2012; Bellm, 2014; Smith et al., 2014b).

7.2 Search methodology

7.2.1 Automated TOO Marshal: alerts and tiling

A program called the iPTF TOO Marshal monitors the stream of GCN notices⁴ from the three redundant, anonymous NASA/GSFC VOEvent servers. It listens for notices of type FERMI_GBM_GND.POS, sent by GBM’s automated on-ground localization, or FERMI_GBM_FIN.POS, sent by the GBM burst advocate’s human-in-the-loop localization.⁵

Upon receiving either kind of notice, the TOO Marshal determines if the best-estimate sky position is observable from Palomar at any time within the 24 hours after the trigger. The criterion for observability is that the position is at an altitude $> 23.5^\circ$ (i.e. airmass $\lesssim 2.5$), at least 20° from the center of the moon, at an hour angle between $\pm 6.5^\circ$, and that the Sun is at least 12° below the horizon at Palomar.

If the position is observable and the $1\text{-}\sigma$ statistical error radius r_{stat} reported in the GCN notice

³Between photometric, late-time red bumps and unambiguous spectral identifications, there are also GRB-SNe that have some SN-associated spectral features. The number of GRBs with spectroscopic SNe is, therefore, ill defined. See Hjorth & Bloom (2012) and references therein for a more complete census.

⁴<http://gcn.gsfc.nasa.gov>

⁵Usually, the *Fermi* team suppresses the notices if the burst is detected and localized more accurately by *Swift* BAT.

is less than 10° , the TOO Marshal selects a set of ten P48 fields that optimally cover the error region.⁶ It converts the GBM position estimate and radius into a probability distribution by applying a well-known empirical prescription of the systematic errors of the GBM localization. [Paciesas et al. \(2012b\)](#) states the total effective error radius in the FERMI_GBM_FIN_POS localizations is well described by the quadrature sum of the statistical radius and a systematic contribution, where the systematic is 2.6° for 72% of bursts and 10.4° for 28% of bursts. We use the weighted RMS of these two values, $r_{\text{sys}} = \sqrt{0.72(2.6)^2 + 0.28(10.4)^2} \approx 6^\circ$. The total error radius is then $r_{\text{eff}} = \sqrt{r_{\text{stat}}^2 + r_{\text{sys}}^2}$. We construct a Fisher–von Mises distribution, centered on the best-estimate position, with a concentration parameter of

$$\kappa = \left[1 - \cos \left(\frac{\pi}{180^\circ} r_{\text{eff}} \right) \right]^{-1}. \quad (7.1)$$

With the FERMI_GBM_FIN_POS alert, the *Fermi* GBM team also distributes a detailed localization map that accounts for the systematic effects ([Connaughton et al., 2014](#)). The TOO Marshal retrieves from the *Fermi* data archive a file that describes the 1-, 2-, and 3- σ significance contours. If the localization has significant asymmetry, we also retrieve a 2D FITS image whose pixel values correspond to the GBM localization significance, and use this instead of the Fisher–von Mises distribution.

Giving preference to fields for which deep coadded reference images exist, the TOO Marshal selects ten P48 fields spanning an area of $\approx 72 \text{ deg}^2$ to maximize the probability of enclosing the true (but as yet unknown) location of the source, assuming the above distribution.

The Marshal then immediately contacts a team of humans (the authors) by SMS text message, telephone, and e-mail. The humans are directed to a mobile-optimized web application to trigger the P48 (see Fig. 7.1).

7.2.2 Triggering the P48

Within the above constraints, we decide whether to follow up the burst based on the following criteria. The event must be $\lesssim 12$ hours old when it first becomes observable from Palomar and we must cover enough of the error circle to have a $\gtrsim 30\%$ chance of enclosing the position of the source. We discard any bursts that are detected and accurately localized by *Swift* BAT, because

⁶We followed up but did not detect afterglows of two bursts that were beyond our nominal cutoff error radius of 10° .




Figure 7.1 Screen shot of the iPTF TOO Marshal shortly after a *Fermi* GBM detection. At this stage, the application presents the recommended P48 fields, the time window of observability, and the history of GCN notices and circulars related to the trigger. It gives the human participants the option to customize the P48 sequence by adding or removing P48 fields and tuning the airmass limit, cadence, or number of images.

these are more efficiently followed up by conventional means. We also give preference to events that are out of the Galactic plane and that are observable for at least 3 hours.

There are some exceptional circumstances that override these considerations. If the burst’s position estimate is accessible within an hour after the burst, we may select it even if the observability window is very brief. If the burst is very well localized or has the possibility of a substantially improved localization later due to a LAT or IPN detection, we may select it even if it is in the Galactic plane.

The default observing program is three epochs of P48 images at a 30-minute cadence. The human may shorten or lengthen the cadence if the burst is very young or old, change the number of epochs, or add and remove P48 fields. When the human presses the “Go” button, the TOO Marshal sends a machine-readable e-mail to the P48 robot. The robot adds the requested fields to the night’s schedule with the highest possible priority, ensuring that they are observed as soon as visible.

7.2.3 Automated candidate selection

As the night progresses, the TOO Marshal monitors the progress of the observations and the iPTF real-time image subtraction pipeline (Nugent et al., in preparation). The real-time pipeline creates difference images between the new P48 observations and coadded references composed of observations from months or years earlier. It generates candidates by performing source extraction on the difference images. A machine learning classifier assigns a *real/bogus* score (RB2; Brink et al. 2013) to each candidate that predicts how likely the candidate is to be a genuine astrophysical source (rather than a radiation hit, a ghost, an imperfect image subtraction residual, or any other kind of artifact).

Table 7.1 lists the number of candidates that remain after each stage of candidate selection. First, requiring candidates to have $\text{SNR} > 5$ gives us a median of 35 000 candidates. This number varies widely with galactic latitude and the area searched (a median of $\sim 500 \text{ deg}^{-2}$). Second, we only select candidates that have $\text{RB2} > 0.1$, reducing the number of candidates to a median of 36% of the original list.⁷ Third, we reject candidates which coincide with known stars in reference catalogs (SDSS and the PTF reference catalog), cutting the list to 17%. Fourth, we eliminate

⁷This RB2 threshold is somewhat deeper than that which is used in the iPTF survey. An improved classifier, RB4 (Bue et al., 2014), entered evaluation in August 2014 shortly before GRB 140808A.

Table 7.1. Number of optical transient candidates surviving each vetting stage

GRB	SNR > 5	RB2 > 0.1	not stellar	not in MPC ^a	detected twice	saved for follow-up
130702A	14 629	2 388	1 346	1 323	417	11
131011A	21 308	8 652	4 344	4 197	434	23
131231A	9 843	2 503	1 776	1 543	1 265	10
140508A	48 747	22 673	9 970	9 969	619	42
140606B	68 628	26 070	11 063	11 063	1 449	28
140620A	152 224	50 930	17 872	17 872	1 904	34
140623A	71 219	29 434	26 279	26 279	442	23
140808A	19 853	4 804	2 349	2 349	79	12
median reduction		36%	17%	16%	1.7%	0.068%

^aNot in Minor Planet Center database

asteroids catalogued by the Minor Planet Center, reducing the list to 16%. Fifth, we demand at least two secure P48 detections after the GBM trigger, reducing the list to a few percent, or ~ 500 candidates.

When the image subtraction pipeline has finished analyzing at least two successive epochs of any one field, the TOO Marshal contacts the humans again and the surviving candidates are presented to the humans via the Treasures portal.

7.2.4 Visual scanning in Treasures Portal

The remaining candidate vetting steps currently involve human participation, and are informed by the nature of the other transients that iPTF commonly detects: foreground SNe (slowly varying and in low- z host galaxies), active galactic nuclei (AGNs), cataclysmic variables, and M-dwarf flares.

In the Treasures portal, we visually scan through the automatically selected candidates one P48 field at a time, examining ~ 10 objects per field (see Figure 7.2 for a screen shot of the Treasures portal). We visually assess each candidate’s image subtraction residual compared to the neighboring stars of similar brightness in the new image. If the residual resembles the new image’s PSF, then the candidate is considered likely to be a genuine transient or variable source.

Next, we look at the photometric history of the candidates. Given the time t of the optical observation relative to the burst and the cadence, δt , we expect that a typical optical afterglow that decays as a power law $F_\nu \propto t^{-\alpha}$, with $\alpha = 1$, would fade by $\delta m = 2.5 \log_{10}(1 + \delta t/t)$ mag over the course of our observations. Any source that exhibits statistically significant fading ($\delta m/m \gg 1$) consistent with an afterglow decay becomes a prime target.⁸

Note that a 1σ decay in brightness requires such a source to be

$$-2.5 \log_{10} \left(\frac{\delta t}{t\sqrt{2}} \right) \quad (7.2)$$

brighter than the 1σ limiting magnitude of the exposures. For example, given the P48's typical limiting magnitude of $R = 20.6$ and the standard cadence of $\delta t = 0.5$ hour, if a burst is observed $t = 3$ hours after the trigger, its afterglow may be expected to have detectable photometric evolution only if it is brighter than $R = 18.3$. Noting that long GRBs preferentially occur at high redshifts and in intrinsically small, faint galaxies (Svensson et al., 2010), we consider faint sources that do not display evidence of fading if they are not spatially coincident with any sources in SDSS or archival iPTF observations.

Therefore, we consider faint sources that do not display evidence of fading if they have no plausibly associated host galaxy in iPTF reference images or SDSS (indicating either a faint quiescent stellar source or a distant host galaxy).

If a faint source is near a spatially resolved galaxy, then we compute its distance modulus using the galaxy's redshift or photometric redshift from SDSS. We know that long GRB optical afterglows at $t = 1$ day typically have absolute magnitudes of $-25 < M_B < -21$ (1σ range; see Figure 9 of Kann et al. 2011). Most SNe are significantly fainter: type Ia are typically $M_B \sim -19$ whereas Ibc and II are $M_B \sim -17$, with luminous varieties of both Ibc and II extending to $M_B \sim -19$ (Richardson et al., 2002; Li et al., 2011). Therefore, if the candidate's presumed host galaxy would give it an absolute magnitude $M_R < -20$, it is considered promising. This criterion is only useful for long GRBs because short GRB afterglows are typically ~ 6 mag fainter than long GRB afterglows (Kann et al., 2011).

The human saves all candidates that are considered promising by these measures to the iPTF Transient Marshal database. This step baptizes them with an iPTF transient name, which consists

⁸A source that exhibits a statistically significant rise is generally also followed up, but as part of the main iPTF transient survey, rather than as a potential optical afterglow.

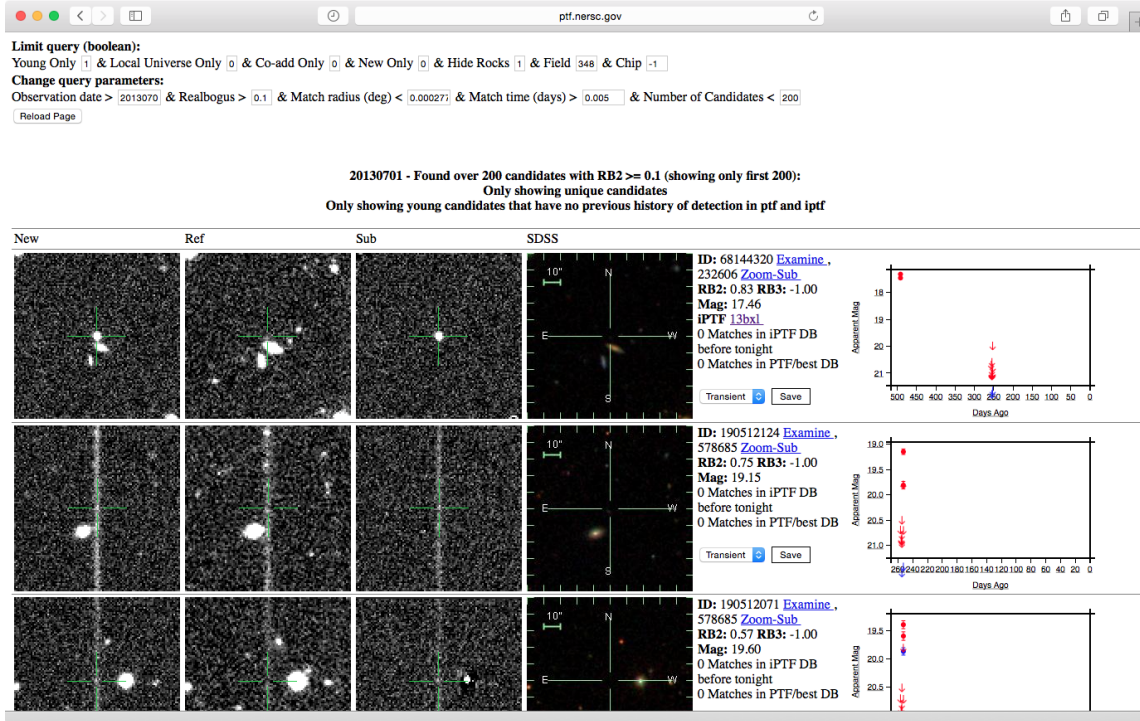


Figure 7.2 Screen shot of the Treasures portal, showing new, reference, subtraction, and archival SDSS images as well as P48 light curves. This page is for the date and field containing GRB 130702A / iPTF13bxl.

of the last two digits of the year and a sequential alphabetic designation.

7.2.5 Archival vetting in the Transient Marshal

Once named in the Transient Marshal, we perform archival vetting of each candidate using databases including VizieR (Ochsenbein et al., 2000), NED⁹, the High Energy Astrophysics Science Archive Research Center (HEASARC)¹⁰, and Catalina Real-Time Sky Survey (CRTS; Drake et al. 2009), in order to check for any past history of variability at that position (see Figure 7.3 for a screen shot of the Transient Marshal).

We check for associations with known quasars or active galactic nuclei in Véron-Cetty & Véron (2010) or with AGN candidates in Flesch (2010).

⁹<http://ned.ipac.caltech.edu>

¹⁰<http://heasarc.gsfc.nasa.gov>

M-dwarfs can produce bright, blue, rapidly fading optical flares than can mimic optical afterglows. To filter our M-dwarfs, we check for quiescent infrared counterparts in WISE (Cutri & et al., 2014). Stars of spectral type L9–M0 peak slightly blue-ward of the WISE bandpass, with typical colors (Wright et al., 2010)

$$\begin{aligned} 3 &\lesssim [R - W1] \lesssim 12 \\ 0.1 &\lesssim [W1 - W2] \lesssim 0.6 \\ 0.2 &\lesssim [W2 - W3] \lesssim 1 \\ 0 &\lesssim [W3 - W4] \lesssim 0.2. \end{aligned}$$

Therefore, a source that is detectable in WISE but that is either absent from or very faint in the iPTF reference images suggests a quiescent dwarf star.

7.2.6 Photometric, spectroscopic, and broad-band follow-up

The above stages usually result in ~ 10 promising optical transient candidates that merit further follow-up. If, by this point, data from *Fermi* LAT or from IPN satellites is available, we can use the improved localization to select an even smaller number of follow-up targets.

For sources whose photometric evolution is not clear, we perform photometric follow-up. We may schedule additional observations of some of the P48 fields if a significant number of candidates are in the same field. We may also use the P48 to gather more photometry for sources that are superimposed on a quiescent source or galaxy, in order to make use of the image subtraction pipeline to automatically obtain host-subtracted magnitudes. For isolated sources, we schedule one or more epochs of *r*-band photometry with the P60. If, by this point, any candidates show strong evidence of fading, we begin multicolor photometric monitoring with the P60.

Next, we acquire spectra for 1–3 candidates per burst using the P200, Gemini, Keck, Magellan, or Himalayan Chandra Telescope (HCT). A spectrum that has a relatively featureless continuum and high redshift absorption lines secures the classification of the candidate as an optical afterglow.

Once any single candidate becomes strongly favored over the others based on photometry or spectroscopy, we trigger X-ray and UV observations with *Swift* and radio observations with VLA, CARMA, and AMI. Detection of a radio or X-ray afterglow typically confirms the nature of the

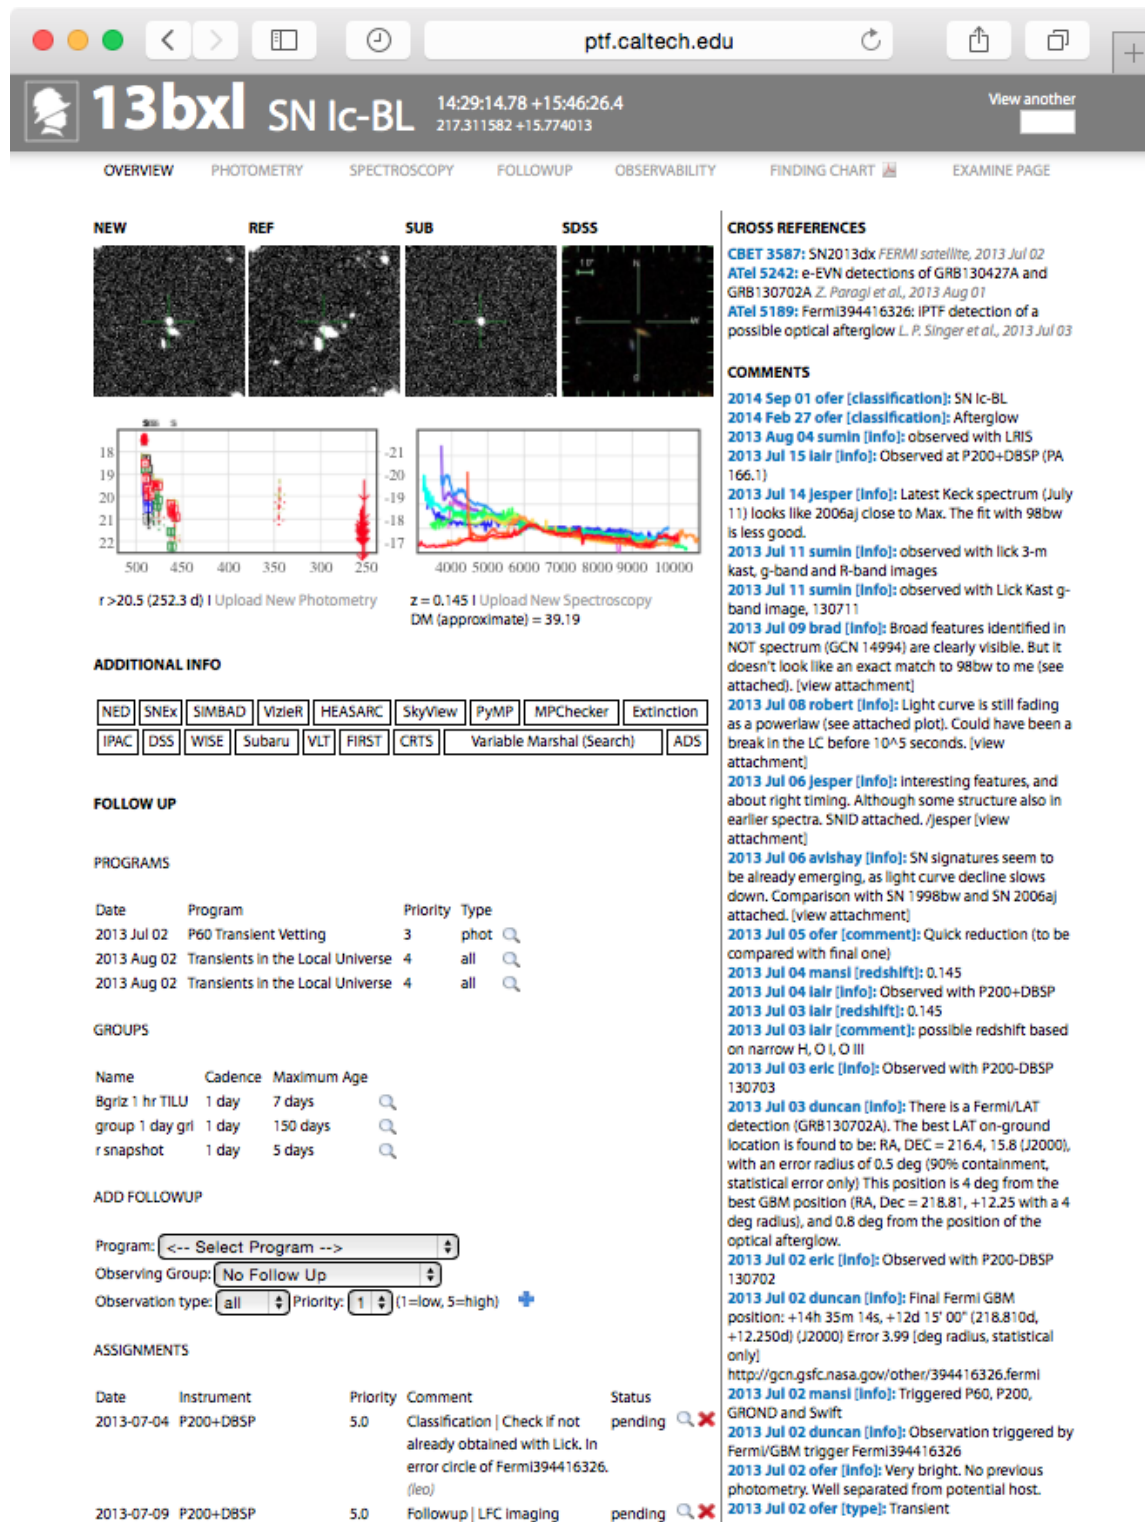


Figure 7.3 Screen shot of the iPTF Transient Marshal, showing GRB 130702A / iPTF13bxi.

optical transient, even without spectroscopy.

Finally, we promptly release our candidates, upper limits, and/or confirmed afterglow discovery in GCN Circulars.

7.2.7 Long-term monitoring and data reduction

To monitor the optical evolution of afterglows identified by our program, we typically request nightly observations in *ri* (and occasionally *gz*) filters for as long as the afterglow remained detectable. Bias subtraction, flat-fielding, and other basic reductions are performed automatically at Palomar by the P60 automated pipeline using standard techniques. Images are then download and stacked as necessary to improve the SNR. Photometry of the optical afterglow is then performed in IDL using a custom aperture-photometry routine, calibrated relative to SDSS secondary standards in the field (when available) or using our own solution for secondary field standards constructed during a photometric night (for fields outside the SDSS footprint).

For some bursts (GRB 140606B), we also obtain photometry with the Large Monolithic Imager (LMI) mounted on the 4.3 m Discovery Channel Telescope (DCT) in Happy Jack, Arizona. Standard CCD reduction techniques (e.g., bias subtraction, flat-fielding, etc.) are applied using a custom IRAF pipeline. Individual exposures are aligned with respect to astrometry from the 2 Micron All-Sky Survey (2MASS; [Skrutskie et al. 2006](#)) using SCAMP ([Bertin, 2006](#)) and stacked with SWarp ([Bertin et al., 2002](#)).

We also usually monitor GBM-iPTF afterglows with CARMA, a millimeter-wave interferometer located at Cedar Flat near Big Pine, California. All observations are conducted at 93 GHz in single-polarization mode in the array’s C, D, or E configuration. Targets are typically observed once for 1–3 hours within a few days after the GRB, establishing the phase calibration using periodic observations of a nearby phase calibrator and the bandpass and the flux calibration by observations of a standard source at the start of the track. If detected, we acquire additional observations in approximately logarithmically-spaced time intervals until the afterglow flux falls below detection limits. All observations are reduced using MIRIAD using standard interferometric flagging and cleaning procedures.

VLA observations are reduced using the Common Astronomy Software Applications (CASA) package. The calibration is performed using the VLA calibration pipeline. After running the pipeline, we inspect the data (calibrators and target source) and apply further flagging when

Table 7.2. GBM–iPTF detections.

GRB	OT	RA (J2000)	Dec (J2000)	Gal. lat. ^a	<i>z</i>	E_{peak} (keV, rest)	$E_{\gamma,\text{iso}}$ (10^{52} erg, rest) ^{b,c}	T_{90} (s)	$m_R(t_{\text{P48}})$ ^d
GRB 130702A	iPTF13bxl	14 ^h 29 ^m 15 ^s	+15°46′26″	65°	0.145	18±3	<0.065±0.001	58.9±6.2	17.38
GRB 131011A	iPTF13dsw	02 ^h 10 ^m 06 ^s	-4°24′40″	-61°	1.874	625±92	14.606±1.256	77.1±3	19.83
GRB 131231A	iPTF13ekl	00 ^h 42 ^m 22 ^s	-1°39′11″	-64°	0.6419	291±6	23.015±0.278	31.2±0.6	15.85
GRB 140508A	iPTF14aue	17 ^h 01 ^m 52 ^s	+46°46′50″	38°	1.03	534±28	24.529±0.86	44.3±0.2	17.89
GRB 140606B	iPTF14bfu	21 ^h 52 ^m 30 ^s	+32°00′51″	-17°	0.384	801±182	0.468±0.04	22.8±2.1	19.89
GRB 140620A	iPTF14cva	18 ^h 47 ^m 29 ^s	+49°43′52″	21°	2.04	387±34	7.28 ±0.372	45.8±12.1	17.60
GRB 140623A	iPTF14cyb	15 ^h 01 ^m 53 ^s	+81°11′29″	34°	1.92	834±317	3.58 ±0.398	114.7±9.2	18.04
GRB 140808A	iPTF14eag	14 ^h 44 ^m 53 ^s	+49°12′51″	59°	3.29	503±35	8.714±0.596	4.5±0.4	19.01

^aGalactic latitude of optical afterglow. This is one of the main factors that influences the number of optical transient candidates in Table 7.1.

^b $E_{\gamma,\text{iso}}$ is given for a 1 keV–10 MeV rest frame bandpass.

^cThe rest-frame spectral properties, E_{peak} and $E_{\gamma,\text{iso}}$, for GRB 130702A are reproduced from Amati et al. (2013). For all other bursts, we calculated these quantities from the spectral fits (the scat files) in the *Fermi* GBM catalog (Goldstein et al., 2012) using the *k*-correction procedure described by Bloom et al. (2001).

^dR-band apparent magnitude in initial P48 detection.

needed. The VLA measurement errors are a combination of the RMS map error, which measures the contribution of small unresolved fluctuations in the background emission and random map fluctuations due to receiver noise, and a basic fractional error (here estimated to be $\approx 5\%$) which accounts for inaccuracies of the flux density calibration. These errors are added in quadrature and total errors are reported in Table E.2.

AMI is composed of eight 12.8 m dishes operating in the 13.9–17.5 GHz range (central frequency of 15.7 GHz) when using frequency channels 3–7 (channels 1, 2, and 8 are disregarded due to their currently susceptibility to radio interference). For further details on the reduction and analysis performed on the AMI observations please see (Anderson et al., 2014b).

7.3 The GBM–iPTF bursts

To date, we have successfully followed up 35 *Fermi* GBM bursts and detected eight optical afterglows. The detections are listed in Table 7.2, and all of the P48 tilings are listed in Table 7.3. In Figure 7.4, the light curves are shown in the context of a comprehensive sample of long and short GRB afterglows compiled by D. A. Kann (private communication).

The outcome of an individual afterglow search is largely determined by two factors: how much




Figure 7.4 Optical light curves of *Fermi*-iPTF afterglows to date. The light curves of the eight iPTF/GBM bursts are shown in red. For comparison, the gray lines show a comprehensive sample of long GRB optical light curves from [Cenko et al. \(2009\)](#), [Kann et al. \(2010\)](#), [Perley et al. \(2014c\)](#), and D. A. Kann (private communication). The white area outside of the light gray shading illustrates the range of GRB afterglows that are accessible given a half-hour cadence and the P48's 60 s limiting magnitude of $R = 20.6$. The two light curves shown in blue are other related iPTF transients. The first is PTF11agg, an afterglow-like transient with no detected gamma-ray emission ([Cenko et al., 2013b](#)). The second is GRB 140226A / iPTF14yb, reported initially by iPTF from its optical afterglow ([Cenko et al. 2014](#); Cenko et al., in preparation), and later by IPN from its gamma-ray emission ([Hurley et al., 2014a](#)).

Table 7.3. Log of P48 tilings for *Fermi* GBM bursts.

GRB time ^a	GBM fluence ^b	t_{P48} $-t_{\text{burst}}^{\text{c}}$	P48 area ^d	Prob. ^e
2013-06-28 20:37:57	10 \pm 0.1	10.02	73	32%
→ 2013-07-02 00:05:20	57 \pm 1.2	4.20	74	38%
2013-08-28 07:19:56	372 \pm 0.6	20.28	74	64%
2013-09-24 06:06:45	37 \pm 0.6	23.24	74	28%
2013-10-06 20:09:48	18 \pm 0.6	15.26	74	18%
→ 2013-10-11 17:47:30	89 \pm 0.6	11.56	73	54%
2013-11-08 00:34:39	28 \pm 0.5	4.69	73	37%
2013-11-10 08:56:58	33 \pm 0.3	17.47	73	44%
2013-11-25 16:32:47	5.5 \pm 0.3	11.72	95	26%
2013-11-26 03:54:06	17 \pm 0.3	6.94	109	59%
2013-11-27 14:12:14	385 \pm 1.4	13.46	60	50%
2013-12-30 19:24:06	41 \pm 0.4	7.22	80	38%
→ 2013-12-31 04:45:12	1519 \pm 1.2	1.37	30	32%
2014-01-04 17:32:00	333 \pm 0.6	18.57	15	11%
2014-01-05 01:32:57	6.4 \pm 0.1	7.63	74	22%
2014-01-22 14:19:44	9.1 \pm 0.5	11.97	75	34%
2014-02-11 02:10:41	7.4 \pm 0.3	1.77	44	19%
2014-02-19 19:46:32	28 \pm 0.5	7.01	71	14%
2014-02-24 18:55:20	24 \pm 0.6	7.90	72	30%
2014-03-11 14:49:13	40 \pm 1.2	12.18	73	54%
2014-03-19 23:08:30	71 \pm 0.3	3.88	74	48%
2014-04-04 04:06:48	82 \pm 0.2	0.11	109	69%
2014-04-29 23:24:42	6.2 \pm 0.2	10.99	74	15%
→ 2014-05-08 03:03:55	614 \pm 1.2	6.68	73	67%
2014-05-17 19:31:18	45 \pm 0.4	8.60	95	69%
2014-05-19 01:01:45	39 \pm 0.5	4.42	73	41%
→ 2014-06-06 03:11:52	76 \pm 0.4	4.08	74	56%
2014-06-08 17:07:11	19 \pm 0.6	11.20	73	49%
→ 2014-06-20 05:15:28	61 \pm 0.6	0.17	147	59%
→ 2014-06-23 05:22:07	61 \pm 0.6	0.18	74	4%
2014-06-28 16:53:19	18 \pm 1.0	16.16	76	20%
2014-07-16 07:20:13	2.4 \pm 0.3	0.17	74	28%
2014-07-29 00:36:54	81 \pm 0.7	3.43	73	65%
2014-08-07 11:59:33	13 \pm 0.1	15.88	73	54%
→ 2014-08-08 00:54:01	32 \pm 0.3	3.25	95	69%

^aTime of *Fermi* GBM trigger. →**Afterglow detections** are marked with an arrow and set in bold face. The corresponding entries in Table 7.2 can be found by matching the date to the GRB name (GRB YYMMDDA).

^bObserved *Fermi* GBM fluence in the 10–1000 keV band, in units of 10^{-7} erg cm $^{-2}$. This quantity is taken from the bcat files from the *Fermi* GRB catalog at HEASARC.

^cAge in hours of the burst at the beginning of the P48 observations.

^dArea in deg 2 spanned by the P48 fields.

^eProbability, given the *Fermi* GBM localization, that the source is contained within the P48 fields.




Figure 7.5 Prior probability of containing the burst’s location within the P48 fields versus age of the burst at the beginning of P48 observations. Afterglow detections are shown in red and non-detections are shown in gray.

probability is contained within the P48 footprints, and how bright the afterglow is at the time of the observations (see Figure 7.5). We calculate the expected success rate as follows. For each burst, we find the prior probability that the position is contained within the P48 fields that we observed. We then compute the fraction of afterglows from Kann’s sample (which has a mean and standard deviation of 22 ± 2 mag at $t = 1$ day) that are brighter than $R = 20.6$ mag at the same age as when the P48 observations started. The product of these two numbers is the prior probability of detection for that burst. By summing over all of the iPTF/GBM bursts, we obtain the expected number of detections. Within 95% confidence bootstrap error bars, we find an expected 5.5–8.5 detections or a success rate of 16%–24%. This is consistent with the actual success rate of 23%.

This suggests that the success rate is currently limited by the survey area and the response time (dictated by sky position and weather). We could increase the success rate by decreasing the maximum time since trigger at which we begin follow-up. We could increase the success rate without adversely affecting the number of detections by simply searching a greater area for coarsely localized events.

Over the next few sections, we summarize the observations and general physical interpretation of all of the GBM–iPTF afterglows detected to date. Figures 7.19–7.26 show light curves and SEDs spanning X-ray, UV, optical, IR, and radio frequencies. In Appendix E, Table E.1 lists a

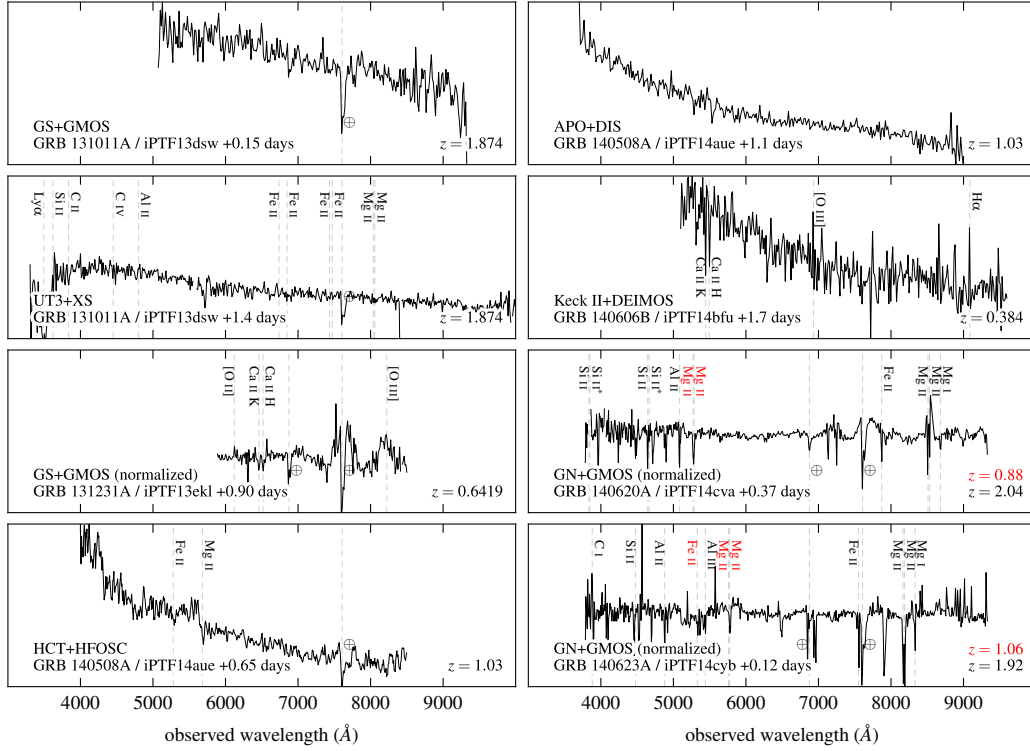


Figure 7.6 Afterglow spectra. The horizontal axis shows wavelength in vacuum in the observer frame, and the vertical axis shows scaled flux. Lines at the redshift of the putative host are labeled in black; lines corresponding to any intervening absorbing systems are labeled in red. Note that in cases where one or fewer lines are discernible in our spectra, the redshifts have been reported in GCNs by other groups.

selection of UV/O/IR observations including all of our P48 and P60 observations. Table E.2 lists all of our radio detections.

The reported P48 magnitudes are all in the Mould R band and in the AB system (Oke & Gunn, 1983), calibrated with respect to either r' point sources from SDSS or for non-SDSS fields using the methods described in Ofek et al. (2012).

7.3.1 GRB 130702A / iPTF13bxl

This is the first GBM burst whose afterglow we discovered with iPTF (Singer et al., 2013), indeed the first afterglow ever to be pinpointed based solely on a *Fermi* GBM localization. It is also the

Table 7.4. Log of spectroscopic observations

Date	Telescope	Instrument	Wavelengths (Å)	Lines	References
GRB 131011A / iPTF13dsw					
2013-10-12 08:56	Gemini South	GMOS	5100–9300	none	Kasliwal et al. (2013)
2013-10-13 03:59	ESO/VLT UT3	X-shooter	3100–5560	Ly α , Si II, C II, C IV, Al II	Rau et al. (2013)
...	5550–10050	Fe II, Mg II	...
GRB 131231A / iPTF13ekl					
2014-01-01 02:15	Gemini South	GMOS	6000–10000	[O II], [O III], Ca II H+K	Cucchiara (2014)
GRB 140508A / iPTF14aue					
2014-05-08 18:55	HCT	HFOSC	3800–8400	Fe II, Mg II	Bhalerao & Sahu (2014)
2014-05-09 06:33	APO	DIS	3200–9800	none	none
GRB 140606B / iPTF14bfu					
2014-06-07 19:16	Keck II	DEIMOS	4500–9600	[O II], [O III], H α , Ca II H+K	Perley et al. (2014a)
GRB 140620A / iPTF14cva					
2014-06-20 14:00	Gemini North	GMOS	5090–9300	Mg I, Mg II, Fe II, Al II, Si II, Si II*	Kasliwal et al. (2014)
...	4000–6600
GRB 140623A / iPTF14cyb					
2014-06-23 08:10	Gemini North	GMOS	4000–6600	Mg II, Fe II, Al II, Si II, Al III, C I, C IV	Bhalerao et al. (2014)

lowest redshift GRB in our sample, so it has the richest and most densely sampled broadband afterglow data. It has two other major distinctions: its associated SN (SN 2013dx, [Schulze et al. 2013](#); [Pozanenko et al. 2013](#); [Cenko et al. 2013a](#); [D’Elia et al. 2013](#)) was detected spectroscopically, and its prompt energetics are intermediate between II GRBs and standard cosmic bursts (see below).

Based on the *Fermi* GBM ground localization with an error radius of 4° , we imaged ten fields twice with the P48 at $\Delta t = t - t_{\text{GBM}} = 4.2$ hr after the burst.¹¹ We scheduled P60 imaging and P200 spectroscopy for three significantly varying sources. Of the three, iPTF13bxi showed the clearest evidence of fading in the P48 images. Its spectrum at $\Delta t = 1.2$ days consisted of a featureless blue continuum. We triggered *Swift*, which found a bright X-ray source at the position of iPTF13bxi ([Singer et al., 2013](#); [D’Avanzo et al., 2013](#)). Shortly after we issued our GCN circular ([Singer et al., 2013](#)), [Cheung et al. \(2013\)](#) announced that the burst had entered the FOV of LAT at $\Delta t = 250$ s. The LAT error circle had a radius of 0.5° , and its center was 0.8° from iPTF13bxi. An IPN triangulation with *MESSENGER* (GRNS), *INTEGRAL* (SPI-ACS), *Fermi*–GBM, and Konus–WIND ([Hurley et al., 2013](#)) yielded a 0.46° -wide annulus that was also consistent with the OT.

The afterglow’s position is $0''.6$ from an $R = 23.01$ mag source that is just barely discernible in the P48 reference images. A NOT+ALFOSC spectrum ([Leloudas et al., 2013](#)) determined a redshift of $z = 0.145$ for a galaxy $7''.6$ to the south of iPTF13bxi. At $\Delta t = 2.0$ days, we obtained a Magellan+IMACS spectrum ([Mulchaey et al., 2013](#)) and found weak emission lines at the location of the afterglow that we interpreted as $H\alpha$ and [O III] at the same redshift. [Kelly et al. \(2013\)](#) characterized the burst’s host environment in detail, and concluded that it exploded in a dwarf satellite galaxy.

Joining the two P48 observations at $\Delta t < 1$ day to the late-time P60 light curve requires a break at $\Delta t = 1.17 \pm 0.09$ days, with slopes $\alpha_{\text{O},1} = 0.57 \pm 0.03$ and $\alpha_{\text{O},2} = 1.05 \pm 0.03$ before and after the break respectively. The XRT light curve begins just prior to this apparent break and seems to follow the late-time optical decay (until the SN begins to dominate at $\Delta t = 5$ days), although

¹¹Our tiling algorithm at the time selected fields based on an empirical calibration of *Fermi* GBM’s systematic errors. We had selected bursts that were detected by both *Swift* and *Fermi*, and constructed a fit to a cumulative histogram of the number of bursts whose BAT or XRT positions were within a given number of nominal 1σ statistical radii of the center of the *Fermi* error circle. Our tiling algorithm scaled this fit by the 1σ radius of the burst in question, and then constructed a 2D angular probability distribution from it. For sufficiently large error radii, this prescription produced probability distributions that had a hole in the middle. For this reason, the tiling algorithm picked out P48 fields that formed an annulus around the GBM $1\text{-}\sigma$ error circle (not, as we stated in [Singer et al. 2013](#), because of a lack of reference images).

the automated *Swift* light curve analysis (Evans et al., 2009) also suggests a possible X-ray break with about the same time and slopes. This hints at an achromatic break, normally a signature of a jet. However, the late slope and the change in slope are both unusually shallow for a jet break. The change in slope is also a little too large for cooling frequency crossing the band (for which one would expect $\Delta\alpha = 1/4$). An energy injection or a structured jet model may provide a better fit (Panaitescu, 2005).

Late-time $\Delta t > 1$ day observations include several P60 *gri* observations, three RATIR *r'i'ZYJH* epochs, an extensive *Swift* XRT and UVOT light curve, as well as radio observations with VLA and CARMA (although of the VLA data, we only have access to the first observation). The optical and X-ray spectral slopes are similar, $\beta_O = 0.7 \pm 0.1$ and $\beta_X = 0.8 \pm 0.1$. An SED at $2 < \Delta t < 2.3$ days is well explained by the standard external shock model (Sari et al., 1998) in the slow cooling regime, with ν_m lying between the VLA and CARMA frequencies and ν_c in the optical. This fit requires a relatively flat electron spectrum, $dn_e/d\gamma_e \propto \gamma_e^{-p}$ with $p \approx 1.6$, cut off at high energies. Applying the relevant closure relations (for the case of $1 < p < 2$, see Dai & Cheng 2001) to α_X and β_X permits either an ISM or wind environment. However, the fact that the X-ray and optical continue to track each other in time suggests that the ν_c is decreasing with time, arguing for an ISM density profile.

This model has difficulty explaining the constant or slowly varying 93 GHz CARMA light curve.

Our late-time spectroscopy and analysis of the SN will be published separately (Cenko et al., in preparation).

7.3.2 GRB 131011A / iPTF13dsw

We started P48 observations of *Fermi* trigger 403206457 (Jenke, 2013) about 11.6 hours after the burst. The optical transient iPTF13dsw (Kasliwal et al., 2013) faded from $R = 19.7$ mag to $R = 20.2$ mag from 11.6 to 14.3 hours. The latest pre-trigger image on 25 September 2013 had no source at this location to a limit of $R > 20.6$ mag. The optical transient continued to fade as it was monitored by several facilities (Xu et al., 2013,a; Perley et al., 2013; Sudilovsky et al., 2013; Volnova et al., 2013).

At 15.1 hours after the burst, we obtained a spectrum of iPTF13dsw with the Gemini Multi-Object Spectrograph (GMOS) on the Gemini-South telescope. GMOS was configured with the




Figure 7.7 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), 3σ IPN triangulation (blue), LAT 1σ error circle (green), and discovery images for GRB 130702A / iPTF13bxl.

R400 grating with a central wavelength of 7200 \AA and the $1''$ slit, providing coverage over the wavelength range from $5100\text{--}9300 \text{ \AA}$ with a resolution of $\approx 3 \text{ \AA}$. No prominent features were detected over this bandpass, while the spectrum had a typical SNR of $\approx 3\text{--}4$ per 1.4 \AA pixel.

Rau et al. (2013) observed the optical transient with the X-Shooter instrument on the ESO 8.2-m Very Large Telescope (VLT). In their spectrum extending from ~ 3000 to $\sim 24000 \text{ \AA}$, they identified several weak absorption lines from which they derived a redshift of $z = 1.874$.

The source was detected by *Swift* XRT (Page, 2013), but with insufficient photons for spectral analysis. There are no radio observations. Largely because in our sample this is the oldest afterglow at the time discovery, there are not enough broadband data to constrain the blast wave physics.

7.3.3 GRB 131231A / iPTF13ekl

GRB 131231A was detected by *Fermi* LAT (Sonbas et al., 2013) and GBM (Jenke & Xiong, 2014), with photons detected up to 9.7 GeV . Xu et al. (2013) observed the LAT error circle with the 1-m telescope at Mt. Nanshan, Xinjiang, China. At 7.9 hours after the burst, they detected a single $R \approx 17.6$ mag source that was not present in SDSS images. At 17.3 hours after the burst, Malesani et al. (2013) observed the afterglow candidate with the MOSaic CAmera (MOSCA) on the 2.56-m Nordic Optical Telescope (NOT). The source had faded to $R = 18.6$.




Figure 7.8 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), and discovery images for GRB 131011A / iPTF13dsw.

Although we had imaged ten P48 fields shortly after the *Fermi* trigger (Singer et al., 2013), due to the short visibility window at Palomar we were only able to obtain one epoch. At 1.45 hours after the burst, we detected an $R = 15.7$ mag optical transient iPTF13ekl at the position of the Nanshan candidate. Though our single detection of iPTF13ekl could not by itself rule out that the source was a moving solar system object, the Nanshan detection at 6.46 hours, fitting a decay with a power law index of $\alpha = 1.03$, was strong evidence that the transient was the optical afterglow of GRB 131231A.

On January 1.09 UT (21.5 hours after the trigger), we observed the afterglow with Gemini-South using the GMOS camera (Hook et al. 2004) in Nod&Shuffle mode: we obtained 32 dithered observations of 30 seconds each at an average airmass of 2. We analyzed this dataset using the dedicated GEMINI package under the IRAF environment and extracted the 1-dimensional spectrum using the APALL task. We determined the redshift of the GRB, based on the simultaneous identification of forbidden nebular emission lines ([O II], [O III]) and absorption features (CaH&K) at the same redshift of $z = 0.6419$. In Figure 7.6, we show the normalized spectrum. In black we show a smoothed continuum. Emission lines are indicated as well as CaH&K absorption features. Also atmospheric bands are marked.

The source was also detected by *Swift* XRT (Mangano et al., 2014b) and UVOT (Holland & Mangano, 2014), as well as CARMA (Perley, 2014).

With only the CARMA observations, the SED is highly degenerate. Contributing to the




Figure 7.9 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), LAT 1σ error circle (green), and discovery images for GRB 131231A / iPTF13ekl.

degeneracy, the X-ray and optical observations appear to fall on the same power-law segment. It is consistent with either fast or slow cooling if the greater of ν_c or ν_m is near the optical, assuming a flat electron distribution with $p \sim 1.5$. It is also consistent with slow cooling if ν_c is above the X-ray band and $p \sim 2.6$.

7.3.4 GRB 140508A / iPTF14aue

This burst was detected by *Fermi* GBM and *INTEGRAL* SPI-ACS (Yu & Goldstein, 2014), as well as by *Konus-WIND*, Mars Odyssey (not included in the GCN circular), *Swift* BAT (outside the coded field of view), and *MESSENGER*, yielding a $1^\circ 5' \times 12'$ IPN error box (Hurley et al., 2014b).

Due to poor weather early in the night, P48 observations started 6.7 hours after the trigger (Singer et al., 2014a). We found one optical transient candidate within the IPN triangulation, iPTF14aue, which faded from $r = 17.89 \pm 0.01$ mag with a power law fit of $\alpha = 1.12 \pm 0.1$.

We triggered a *Swift* TOO. From 0.8 to 8.1 days after the trigger, *Swift* XRT detected a coincident X-ray source that faded with a power law $\alpha = 1.48 (+0.15, -0.14)$ (Amaral-Rogers, 2014a,b). The source was also detected by *Swift* UVOT (Marshall & Amarel-Rogers, 2014).

Moskvitin et al. (2014) obtained a 20 min, 3800–7200 Å spectrum of iPTF14aue with the 6-m BTA telescope in Zelenchukskaia. Exhibiting no absorption features, this established an upper limit of $z < 2.1$. Malesani et al. (2014) used the Andalucia Faint Object Spectrograph and Camera




Figure 7.10 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), 3σ IPN triangulation (blue), and discovery images for GRB 140508A / iPTF14aue.

(ALFOSC) on NOT to get an 1800 s spectrum spanning 3200–9100 Å, and found several absorption features at redshift $z = 1.03$. Consistent redshifts were reported by [Wiersema et al. \(2014\)](#) with the ACAM instrument on the 4.2-m William Herschel Telescope and by [Bhalerao & Sahu \(2014\)](#) with Himalaya Faint Object Spectrograph and Camera (HFOSC) on the 2-m HCT.

Due to the brightness of the optical transient, optical photometry was available from several facilities up to 4.5 days after the burst ([Gorosabel et al., 2014b](#); [Moskvitin et al., 2014](#); [Malesani et al., 2014](#); [Masi, 2014](#); [Butler et al., 2014b,a](#); [Fujiwara et al., 2014](#); [Volnova et al., 2014a](#)).

[Horesh et al. \(2014\)](#) detected the source with VLA 5.2 days after the Fermi trigger, at 6.1 GHz (C-band) and at 22 GHz (K-band). A broadband SED constructed from P60 and XRT data from around this time is consistent with $p \approx 2$. Because p is not distinguishable from 2, we cannot discriminate between fast and slow cooling based on this one time slice. However, given the late time of this observation, the slow cooling interpretation is more likely, putting ν_m between the radio and optical bands and ν_c between the optical and X-ray. Because the VLA light curve is decreasing with time, an ISM circumburst density profile is favored.

7.3.5 GRB 140606B / iPTF14bfu

Fermi trigger 423717114 ([Burns, 2014](#)) was observable from Palomar for several hours, starting about 4.3 hours after the time of the burst. Based on the final GBM localization, we searched ten

P48 fields and found several plausible optical transient candidates (Singer et al., 2014c).

iPTF14bfu had no previous detections in iPTF between 23 May and 13 October 2013. Its position was outside the SDSS survey footprint, but it had no plausible host associations in VizieR (Ochsenbein et al., 2000). From 4.3 to 5.5 hours after the burst, it faded from $R = 19.89 \pm 0.10$ to 20.32 ± 0.14 mag, fitting a power law of $\alpha = -1.6 \pm 0.7$ relative to the time of the GBM trigger. iPTF14bfw ($R = 19.96 \pm 0.06$ mag) was coincident with an $r = 21.27$ galaxy in SDSS DR10, and displayed no statistically significant photometric variation over the course of our P48 observations. iPTF14bgc ($R = 18.44 \pm 0.02$ mag) was coincident with an $R = 21.07 \pm 0.08$ mag point source in our coadded reference image composed of exposures from July 31 through 24 September 2013. iPTF14bga ($R = 19.75 \pm 0.06$ mag) was likewise coincident with a $R = 20.42 \pm 0.17$ mag point source in our reference image composed of exposures from 29 July through 20 October 2011.

On the following night, we observed all four candidates again with P48 and P60 (Perley & Singer, 2014). iPTF14bfw and iPTF14bga had not faded relative to the previous night. iPTF14bgc had faded to $R = 20.68 \pm 0.21$ mag, consistent with the counterpart in our reference images but significantly fainter than the previous night. A power law fit to the decay gave a temporal index of $\alpha = -1.1 \pm 0.1$, entirely consistent with typical GRB afterglows. iPTF14bfu was not detected in our P48 images to a limiting magnitude of $R < 21.1$, but it was detected in stacked P60 images ($r = 21.1 \pm 0.2$), consistent with a power law of $\alpha \sim -0.5$.

An IPN triangulation from *Fermi*, Konus-WIND, and MESSENGER yielded a long, slender $14^\circ 18' \times 0^\circ 41'$ error box that contained iPTF14bfu and iPTF14bfw (Hurley et al., 2014c).

We obtained two 900 s spectra with the DEIMOS spectrograph on the Keck II 10-m telescope (Perley et al., 2014a). On a blue continuum, we found O II, O III, and $H\alpha$ emission features, and Ca II absorption features, at a common redshift of $z = 0.384$. A galaxy offset by $\sim 2''$ along the slit showed the same emission lines at the same redshift.

Swift XRT observed the location of iPTF14bfu for a total of 9 ks from 2.1 to 9.3 days after the GBM trigger, and found a source that faded with a power-law fit of $\alpha = -1.0 (+0.7, -0.6)$ (Mangano et al., 2014a; Mangano & Burrows, 2014; Mangano, 2014).

At 18.4 days after the trigger, we obtained a 1200 s spectrum of iPTF14bfu with the Low Resolution Imaging Spectrometer (LRIS) on the Keck I 10-meter telescope (Perley et al., 2014b). The spectrum had developed broad emission features. A comparison using Superfit (Howell et al., 2005) showed a good match to SN 1998bw near maximum light, indicating that the source had




Figure 7.11 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), 3σ IPN triangulation (blue), and discovery images for GRB 140606B / iPTF14bfu.

evolved into an SN Ic-BL. Our late-time photometry and spectroscopy will be published separately (Cano et al., in preparation).

Although there were three radio detections of GRB 140606B, only during the first CARMA detection does the optical emission appear to be dominated by the afterglow. We can construct an SED around this time using nearly coeval DCT and XRT data. Because of the faintness of the X-ray afterglow, the spectral slopes β_X and β_{OX} are only weakly determined. As a result, there is a degeneracy between two plausible fits. The first has ν_m anywhere below the CARMA band, ν_c just below the X-rays, and $p \approx 2$. The second has ν_m just above the radio and ν_c in the middle of the XRT band, with $p \approx 2.2$.

The early P48 observations do not connect smoothly with the P60 and DCT observations from $\Delta t = 1$ to 4 days. This may indicate a steep–shallow–steep time evolution requiring late time energy injection, or may just indicate that the afterglow is contaminated by light from the host galaxy or the SN at relatively early times.

7.3.6 GRB 140620A / iPTF14cva

This burst is distinctive in our sample for two reasons. First, it is the earliest afterglow detection in the iPTF sample at $\Delta t = 0.25$ hours. Second, its broadband SED is not clearly explainable by the standard forward shock model.

Fermi trigger 424934131 (Fitzpatrick & Connaughton, 2014) was observable from Palomar for about 6 hours from the time of the burst. Based on the ground localization, we started observing ten P48 fields about ten minutes after the trigger. Based on the final localization, we added ten more fields, for a total of twenty, about an hour after the trigger.

The candidate iPTF14cva (Kasliwal et al., 2014) was contained within one of the early ten fields. From 14.9 to 87.2 minutes after the trigger, the candidate faded from $R = 17.60 \pm 0.01$ to 18.80 ± 0.02 mag, consistent with a somewhat slow power law of $\alpha = 0.62 \pm 0.01$.

We observed the candidate with GMOS on the 8-m Gemini North telescope. Starting 8.8 hours after the trigger, we obtained two 900 s spectra extending from 4000 to 9300 Å. We detected Mg II and Fe II absorption lines at $z = 0.88$ and many absorption features at a common redshift of $z = 2.04$. The lack of Ly- α absorption implied an upper limit of $z \sim 2.3$, and suggested that $z = 2.04$ was the redshift of the source.

We triggered *Swift* and VLA follow-up. In a 3 ks exposure starting 10.4 hours after the *Fermi* trigger, *Swift* XRT detected an X-ray source with a count rate of 1.2×10^{-1} cts s $^{-1}$ (De Pasquale, 2014b). Over the next four days of *Swift* observations, the X-ray source faded with a slope $\alpha = 1.32 \pm 0.16$ (De Pasquale, 2014a). A fading source was also detected by *Swift* UVOT (Siegel & De Pasquale, 2014).

The source was detected by VLA on June 23 at 6.1 GHz (C band) at 108 ± 15 μ Jy and at 22 GHz (K band) at 62 ± 15 μ Jy. On June 30, there was a marginal detection in C band with 48 ± 12 μ Jy and no detection in K-band with a noise level of 15 μ Jy RMS.

The optical transient was also observed in *R* band by the Konkoly Observatory (Kelemen, 2014) and the 1-m telescope at the Tien Shan Astronomical Observatory (Volnova et al., 2014b).

The SED of this afterglow cannot be explained by a standard forward shock model. If we place the peak frequency near the radio band, the optical and X-ray fluxes are drastically under-predicted, whereas if we place the peak frequency between the optical and X-ray bands, we miss the radio observations by orders of magnitude. This seems to require an additional component. One possibility is that there is a forward shock peak in the UV and a reverse shock peak at low frequencies (similar to GRB 130427A; see Laskar et al. 2013; Perley et al. 2014c). Another possibility is that there is an inverse Compton peak in the UV (similar to GRB 120326A; Urata et al. 2014).




Figure 7.12 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), and discovery images for GRB 140620A / iPTF14cva.

7.3.7 GRB 140623A / iPTF14cyb

Fermi trigger 425193729 (von Kienlin, 2014) was observable from Palomar for about 6 hours from the time of the burst. Based on the ground localization, we started imaging ten fields 11 minutes after the trigger. The final *Fermi* localization, which was available 2.6 hours later, shifted by $13^\circ.4$. Due to the large change in the localization, we calculated only a 4% chance that the source was contained within the P48 fields.

Candidate iPTF14cyb (Kasliwal et al., 2014), situated at an extreme edge of the P48 tiling, was within the $1\text{-}\sigma$ confidence region for both the ground and final localizations. From 16 to 83 minutes after the trigger, the source faded from $R = 18.04 \pm 0.01$ to 19.69 ± 0.06 mag, consistent with a power-law decay with an index $\alpha = 0.94 \pm 0.03$.

Starting 2.8 hours after the trigger, we obtained two 900 s GMOS spectra extending from 4000 to 9300 Å. We detected Mg II and Fe II absorption lines at $z = 1.06$ and many absorption features at $z = 1.92$. The lack of Ly- α absorption implied that this was the redshift of the burst.

We triggered *Swift*, VLA, and CARMA follow-up. In a 3 ks exposure starting 10.7 hours after the burst, *Swift* XRT detected an uncatalogued X-ray source with a count rate of $(2.2 \pm 0.6) \times 10^{-3}$ cts s $^{-1}$ (D’Elia et al., 2014). By 79 hours after the trigger, the source was no longer detected in a 5 ks exposure (D’Elia & Izzo, 2014). No radio source was detected with VLA in C band (6.1 GHz) to an RMS level of 17 μ Jy, or in K band (22 GHz) to an RMS level of 18 μ Jy.




Figure 7.13 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), and discovery images for GRB 140623A / iPTF14cyb.

Because of the lack of radio detections and the extreme faintness of the X-ray afterglow, the broadband behavior of the afterglow does not constrain the shock physics.

7.3.8 GRB 140808A / iPTF14eag

Fermi trigger 429152043 (Zhang, 2014) was observable from Palomar about 3 hours after the burst. We imaged 13 fields with P48 and found one promising optical transient. iPTF14eag was situated on the extreme edge of one of the P48 tiles that was just outside the GBM 1-sigma contour. It faded from $R = 18.91 \pm 0.06$ to 19.29 ± 0.10 mag from 3.35 to 4.91 hours after the trigger and had no archival counterparts in SDSS or in our own reference images.

We were unable to use our TOO programs on Keck or Gemini North because Hawaii was being struck by Hurricane Iselle. We requested photometric confirmation of the fading from HCT (Sahu et al., 2014), submitted a *Swift* TOO, and sent our GCN circular (Singer et al., 2014b) to encourage others to obtain a spectrum.

Swift observed the position of iPTF14eag from 11.6 to 14.4 hours after the burst (Page et al., 2014). An X-ray source was detected with a count rate of 1.5×10^{-2} counts s^{-1} . In a second observation starting 62.2 hours after the trigger (Page & Cenko, 2014), the source had faded to below 2.46×10^{-3} counts s^{-1} . No source was detected by UVOT (Oates & Cenko, 2014).

Meanwhile, from 20.8 to 21.9 hours after the burst, Gorosabel et al. (2014a) obtained a spectrum




Figure 7.14 *Fermi* GBM localization (black contours), P48 tiling (gray rectangles), and discovery images for GRB 140808A / iPTF14eag.

from 3630 to 7500 Å with the OSIRIS instrument on the Gran Telescopio Canarias (GTC) and determined a redshift of $z = 3.29$.

The source was detected in radio with VLA (Corsi & Horesh, 2014) and AMI (Anderson et al., 2014a). The broadband SED around the time of the VLA detection broadly fits a forward shock model, but is poorly constrained due to the lack of a contemporaneous X-ray detection. The spectral slope between the two VLA bands is somewhat steeper than the standard low-frequency value of $\beta = -1/3$, possibly indicating that the radio emission is self-absorbed. We obtained 14 AMI observations every 2 or 3 days from 2014 August 8 until 2014 Sept 12. Observations were 2–4 hours in duration. AMI first detected the afterglow 4.6 days post-burst. The AMI light curve peaked ~ 10.6 days post-burst at 15.7 GHz, which is characteristic of forward shock emission at radio wavelengths (Chandra & Frail, 2012).

A peculiar feature of the optical light curve is that the P60 r and i band observations at $\Delta t \approx 2$ days appears to be inverted, with a rising rather than falling spectral shape, compared to the earlier P60 photometry at $\Delta t \approx 1$ day. However, this feature is within the error bars and may be merely a statistical fluctuation.

This is the highest redshift burst in our sample, and also had the weakest prompt emission in terms of the fluence observed by GBM.

7.4 The population in context

7.4.1 Selection effects

First, we investigate the properties of the subset of GBM bursts followed up by iPTF compared to the GBM bursts as a whole. It is known that, on average, GRBs with larger prompt fluences have brighter optical afterglows, though the correlation is very weak (Nysewander et al., 2009). In Figure 7.15, we plot the fluence in the 10–1000 keV band and $1\text{-}\sigma$ localization radius of all GBM bursts from the beginning of our experiment, retrieved from the *Fermi* GBM Burst Catalog at HEASARC¹². As expected, there is a weak but clearly discernible correlation between fluence and radius, $F \propto r^{-1.3}$, with a Pearson correlation coefficient of $R = 0.64$.¹³ The subset of bursts that we followed up spans a wide range in fluence, and error radii up to $\sim 10^\circ$. The bursts for which we detected optical afterglows are preferentially brighter, with the faintest burst having a fluence as low as $3 \times 10^{-6} \text{ erg cm}^{-2}$. There are some bright ($> 3 \times 10^{-5} \text{ erg cm}^{-2}$) and well confined ($< 1.8^\circ$) events for which we did not find afterglows: those at 2013-08-28 07:19:56, 2013-11-27 14:12:14, and 2014-01-04 17:32:00 (see Table 7.3). However, these non-detections are not constraining given their ages of 20.28, 13.46, and 18.57 hours respectively. Conversely, there were two especially young bursts (followed up at $\Delta t = 0.11$ and 0.17 hours) for which we did not detect afterglows. The non-detection of the burst at 2014-07-16 07:20:13 makes sense because we searched only 28% of the GBM localization. The non-detection on 2014-04-04 04:06:48, for which we observed 69% of the localization, is a little more surprising, especially given its relatively high fluence of $8 \times 10^{-6} \text{ erg cm}^{-2}$; this is a possible candidate for a “dark GRB”. On the whole, however, we can see that (1) we have followed up bursts with a large range of error radii and fluences, (2), there is a weak preference toward detecting bursts with small error radii, and (3) the detections tend toward bursts with high fluences. Naively one might expect higher fluences to translate into lower redshifts, but the interplay between the GRB luminosity function and detector threshold greatly complicate such inferences (Butler et al., 2010).

Second, the rich sample of all of the GRB afterglows that we have today is undeniably the result of the success of the *Swift* mission. It is therefore interesting to consider how the GBM–iPTF

¹²<http://heasarc.gsfc.nasa.gov/W3Browse/fermi/fermigbrst.html>

¹³In a separate sample of GBM GRBs compiled by Connaughton et al. (2014), the correlation between error radius and photon fluence is slightly stronger than the correlation between error radius and fluence. However, we use fluence rather than photon fluence here because the latter is not available for all bursts in the online *Fermi* GBM archive.




Figure 7.15 Fluence and statistical error radius of GBM bursts. Red dots mark bursts that were followed up with iPTF; black circles around red dots mark bursts whose afterglows were detected by iPTF. The black line is a power-law fit to the fluence–error radius relation.




Figure 7.16 Cumulative distribution of redshifts of long GRBs observed by *Swift* BAT (gray) and the GBM–iPTF experiment (red).

sample is similar to or different from the *Swift* sample, given the differences in bandpasses and our increased reliance on the optical afterglow. In Figure 7.16, we plot the cumulative redshift distribution of our sample, alongside the distribution of redshifts of long GRBs detected by *Swift*.¹⁴ Indeed, we find that our sample is at lower redshifts; the former distribution lies almost entirely to the left of the latter, and the ratio of the median redshifts ($z = 1.5$ versus $z = 1.9$) of the two populations is about 0.75. However, with the small sample size, the difference between the two redshift distributions is not significant: a two-sample Kolmogorov–Smirnov test yields a p -value of 0.26, meaning that there is a 26% chance of obtaining these two empirical samples from the same underlying distribution. More GBM–iPTF events are needed to determine whether the redshift distribution is significantly different.

7.4.2 GRBs as standard candles?

Amati et al. (2002) pointed out a striking empirical correlation in the rest-frame prompt emission spectra of BeppoSAX GRBs, with the peak energy (in the νF_ν sense) E_{peak} related to the bolometric, isotropic-equivalent energy release E_{iso} by $E_{\text{peak}} \propto E_{\text{iso}}^{0.52 \pm 0.06}$. It was quickly realized that such a relation, if intrinsic to the bursts, could be used to measure the redshifts of GRBs

¹⁴This sample was extracted from the *Swift* GRB Table, <http://swift.gsfc.nasa.gov/archive/grb.table/>.




Figure 7.17 Rest-frame energetics of GBM-iPTF bursts (red) in comparison to an illustrative sample of previous GRB-SNe¹⁵ (black). A general long GRB sample from Amati (2006); Amati et al. (2008, 2009) is shown in gray. The dashed lines denote 1σ confidence bands around the Amati relation (Amati et al., 2002).

non-spectroscopically (Atteia, 2003). As with the Phillips relation for Type Ia SNe (Phillips, 1993), with such a relation GRBs could serve as standardizable candles in order to measure cosmological parameters (Dai et al. 2004; Friedman & Bloom 2005; Ghirlanda et al. 2006; etc.).

However, there has been a vigorous debate about whether the Amati relation and related correlations are innate to GRBs or reflect a detector-dependent selection bias (Band & Preece, 2005; Nakar & Piran, 2005; Ghirlanda et al., 2005; Sakamoto et al., 2006; Butler et al., 2007; Cabrera et al., 2007; Schaefer & Collazzi, 2007; Butler et al., 2009; Firmani et al., 2009; Krimm et al., 2009; Butler et al., 2010; Shahmoradi & Nemiroff, 2011; Collazzi et al., 2012; Kocevski, 2012). One alternative interpretation is that bursts to the upper-left boundary of the Amati relation are selected against by photon-counting instruments because, being relatively hard, there are fewer photons. The lack of bursts to the lower-right of the Amati line may be due to a genuine lack of relativistic explosions that are much softer than, but as energetic as, standard GRBs.

It has been difficult to directly test the Amati relation in the context of *Fermi* bursts because most lack known redshifts, aside from bursts that were coincidentally also observed and localized by the *Swift* BAT, which do not directly sample the selection bias of *Fermi* GBM. However, Heussaff et al. (2013) showed that many *Fermi* bursts without known redshifts would be inconsistent with the Amati relation at any distance. (See also Urata et al. 2012 for outlier events detected by *Fermi* LAT and *Suzaku* WAM.) Here, we have a small sample of *Fermi* bursts with *known* redshifts. One of them, GRB 140606B / iPTF14bfu at $z = 0.384$, is a clear outlier, over 2σ away from the mean Amati relation. This burst is not alone: in Figure 7.17, we have marked a selection previous long GRBs with spectroscopically identified SNe. Three among them are also outliers. (A possible caveat is that the prompt emission mechanism for GRB 140606B could be different from typical cosmological bursts; we explore this in the next section.) To be sure, most of the bursts in our GBM–iPTF sample fall within a 1σ band of the Amati relation. This includes the nearest event to date, GRB 130702A / iPTF13bxi at $z = 0.145$. However, the one outlier in our admittedly small sample strengthens the case that the boundary of the Amati relation is somewhat influenced by the detector thresholds and bandpasses.

7.4.3 Shock breakout

Two GRBs in our sample, GRB 130702A / iPTF13bxi and GRB 140606B / iPTF14bfu, have $E_{\text{iso}} \sim 10^{51}$ erg (rest frame), energetically intermediate between “standard” luminous, cosmically

distant bursts and nearby llGRBs. Prototypes of the latter class include GRB 980425 / SN 1998bw (Galama et al., 1998; Kulkarni et al., 1998), which was also the first SN discovered in association with a GRB. They offer an interesting test case for competing theories to explain the wide range of prompt gamma-ray energy releases observed from GRBs (e.g., Schulze et al. 2014b).

It has been suggested that the two luminosity regimes correspond to different prompt emission mechanisms (Bromberg et al., 2011). The llGRBs could be explained by the breakout of a mildly relativistic shock from the progenitor envelope (Nakar & Sari, 2012). High-luminosity bursts, on the other hand, are thought to be produced by internal shocks within an ultra-relativistic jet (Rees & Meszaros, 1994) that has successfully punched through the star. A central engine that sometimes fails to launch an ultra-relativistic jet is one way to unify the luminosity functions of standard GRBs and llGRBs (Pescalli et al., 2014).

The smoking gun for the relativistic shock breakout model is a cooling, thermal component to the prompt X-ray emission, as in the case of GRB 060218 (Campana et al., 2006). Unfortunately, this diagnostic is not possible for GRBs 130702A and 140606B because we lack early-time *Swift* observations.

However, Nakar & Sari (2012) propose a closure relation (their Equation 18) between the prompt energy, temperature, and timescale that is valid for shock breakout-powered GRBs. We reproduce it here:

$$t_{\text{bo}}^{\text{obs}} \sim 20 \text{ s} \left(\frac{E_{\text{bo}}}{10^{46} \text{ erg}} \right)^{\frac{1}{2}} \left(\frac{T_{\text{bo}}}{50 \text{ keV}} \right)^{-\frac{9+\sqrt{3}}{4}}. \quad (7.3)$$

If we very crudely assume that all of the prompt emission is from a shock escaping from the progenitor envelope, then we can use E_{iso} , E_{peak} , and T_{90} as proxies for these observables. This gives us a simple discriminator of which bursts are plausible shock breakout candidates, the ratio

$$\xi = (1+z)t_{\text{bo}}^{\text{obs}}/T_{90}, \quad (7.4)$$

which should be close to 1. As expected, most of the energetic ($E_{\text{iso}} > 10^{52} \text{ erg}$), cosmic ($z > 0.5$) GRBs in our sample are inconsistent with the closure relation. They are all much *shorter* in duration, given their γ -ray spectra, than would be expected for a shock breakout. The exception is GRB 140623A / iPTF14cyb, which yields $\xi = 0.5 \pm 0.5$.

Surprisingly, of the two low-luminosity, low-redshift bursts in our sample, GRB 130702A / iPTF13bxl's

prompt emission was also much too brief to be consistent with this shock breakout model, with $\xi = (1.6 \pm 0.7) \times 10^3$. Most likely, this means that the prompt emission of GRB 130702A is simply a very soft, very sub-luminous version of an otherwise ‘ordinary’ long GRB. Any early-time shock breakout signature, if present, was unobserved either because it occurred at energies below GBM’s bandpass, or because it was much weaker than the emission from the standard GRB mechanism. However, GRB 140606B / iPTF14bfu’s prompt emission is consistent with the closure relation, with $\xi = 0.5 \pm 0.3$. Though we must interpret this with caution because we cannot disentangle a thermal component from the GBM data, if we naively apply linear least squares to (the logarithm of) Equations (14, 16, 17) of [Nakar & Sari](#),

$$E_{\text{bo}} \approx 2 \times 10^{45} R_5^2 \gamma_{f,0}^{\frac{1+\sqrt{3}}{2}} \text{ erg}, \quad (7.5)$$

$$T_{\text{bo}} \sim 50 \gamma_{f,0} \text{ keV}, \quad (7.6)$$

$$t_{\text{bo}}^{\text{obs}} \approx 10 \frac{R_5}{\gamma_{f,0}^2} \text{ s}, \quad (7.7)$$

then we find the breakout radius and Lorentz factor to be:

$$\begin{aligned} R_{\text{bo}} &= (1.3 \pm 0.2) \times 10^3 R_{\odot}, \\ \gamma_{f,0} &= 14 \pm 2. \end{aligned}$$

The breakout radius is comparable to that which [Nakar & Sari \(2012\)](#) find for GRBs 060218 and 100316D, suggestive of breakout from a dense wind environment, rather than the star itself. However, the derived Lorentz factor of GRB 140606B is a bit higher than those of the other two examples.

Another way to constrain the nature of the explosion is to look at the kinetic energy $E_{k,\text{iso}}$ of the blast compared to the promptly radiated energy $E_{\gamma,\text{iso}} \equiv E_{\text{iso}}$, and the radiative efficiency $\eta = E_{\gamma,\text{iso}} / (E_{k,\text{iso}} + E_{\text{iso}})$. After the end of any plateau phase, the X-ray flux is a fairly clean diagnostic of $E_{k,\text{iso}}$ assuming the X-rays are above the cooling frequency ([Freedman & Waxman, 2001](#)). During the slow-cooling phase and under the typical conditions where $p \approx 2$ and $\nu_c < \nu_X$, the X-ray flux is only weakly sensitive to global parameters such as the fraction of the internal energy partitioned to electrons and to the magnetic (ϵ_e, ϵ_B). Even the radiative losses, necessary for extrapolating from the late-time afterglow to the end of the prompt phase, are minor, amounting

to order unity at $\Delta t = 1$ day (Lloyd-Ronning & Zhang, 2004). We calculate the isotropic-equivalent rest frame X-ray luminosity from the flux at $\Delta t = 1$ day using Equation (1) of Racusin et al. (2011), reproduced below:

$$L_X(t) = 4\pi D_L^2 F_X(t) (1+z)^{-\alpha_X + \beta_X - 1}. \quad (7.8)$$

Then we estimate the kinetic energy at the end of the prompt emission phase using Equation (7) of Lloyd-Ronning & Zhang (2004):

$$E_{k,\text{iso}} = \left(10^{52} \text{ ergs}\right) \times R \times \left(\frac{L_X(1 \text{ day})}{10^{46} \text{ ergs s}^{-1}}\right)^{-4/(p+2)} \left(\frac{1+z}{2}\right)^{-1} \\ \times \epsilon_{e,-1}^{4(1-p)/(2+p)} \epsilon_{B,-2}^{(2-p)/(2+p)} t_{10 \text{ hr}}^{(3p-2)/(p+2)} \nu_{18}^{2(p-2)/(p+2)}. \quad (7.9)$$

The correction factor R for radiative losses is given by Equation (8) of Lloyd-Ronning & Zhang (2004), adopted here:

$$R = \left(\frac{t}{T_{90}}\right)^{(17/16)\epsilon_e}. \quad (7.10)$$

The numeric subscripts follow the usual convention for representing quantities in powers of 10 times the cgs unit, i.e., $\epsilon_{e,-1} = \epsilon_e/10^{-1}$, $\epsilon_{B,-2} = \epsilon_B/10^{-2}$, and $\nu_{18} \equiv \nu/(10^{18} \text{ Hz})$. We assume $\epsilon_e = 0.1$ and $\epsilon_B = 0.01$. For bursts that have XRT detections around $\Delta t = 1$ day (GRBs 130702A, 131231A, 140508A, 140606B, and 140620A), we calculate L_X by interpolating a least squares power law fit to the X-ray light curve. Some of our bursts (GRBs 131011A, 140623A, and 140808A) were only weakly detected by XRT; for these we extrapolate from the mean time of the XRT detection assuming a typical temporal slope of $\alpha_X = 1.43 \pm 0.35$ (Racusin et al., 2011). The kinetic and radiative energies of our eight bursts are shown in Figure 7.18. Half of our bursts are reasonably well constrained in E_k - E_γ space; these are shown as red points. The other half (GRBs 131011A, 131231A, 140620A, and 140623A) have highly degenerate SEDs, so their position in this plot is highly sensitive to model assumptions; these are shown as gray points. Dotted lines are lines of constant radiative efficiency.

Within our sample, there are at least three orders of magnitude of variation in both $E_{k,\text{iso}}$ and $E_{\gamma,\text{iso}}$. The two GRB-SNe have radiative and kinetic energies of $\sim 10^{51}$ ergs, both two to three orders of magnitude lower than the other extreme in our sample or the average values for *Swift* bursts. In our sample, they have two of the lowest inferred radiative efficiencies of $\eta \sim 0.1$ – 0.5 , but




Figure 7.18 Fireball kinetic energy $E_{k,\text{iso}}$ at $t = T_{90}$ as estimated from X-ray flux versus rest-frame isotropic-equivalent gamma-ray energy $E_{\gamma,\text{iso}}$. Red points denote bursts for which $E_{k,\text{iso}}$ can be reliably estimated from the *Swift* XRT data; gray points denote bursts for which the calculation of $E_{k,\text{iso}}$ may have extreme model dependence. Dashed lines are lines of constant radiative efficiency $\eta = E_{\gamma,\text{iso}} / (E_{k,\text{iso}} + E_{\gamma,\text{iso}})$. The gray, blue, and red rectangles show the 1σ parameter ranges of *Swift* BAT, BAT+GBM, and BAT+LAT long GRBs from [Racusin et al. \(2011\)](#).

these values are not atypical of BATSE bursts (e.g., [Lloyd-Ronning & Zhang 2004](#)) and are close to the median value for *Swift* bursts. These are, therefore, truly less energetic than cosmological bursts, not merely less efficient at producing gamma-rays.

7.5 Looking forward

In this experiment, we have followed up 35 *Fermi* GBM bursts, scanning areas from 30 to 147 deg². To date, we have detected eight afterglows with apparent optical magnitudes as bright as $R \approx 16$ and as faint as $R \approx 20$. We have found redshifts as nearby as $z = 0.145$ and as distant as $z = 3.29$. A continuation of the project should reveal more low-redshift events, more GRB-SNe, and more relatively hard GRBs.

We aim to uncover the much fainter afterglows of short, hard bursts by using stacked P48 exposures and integrating a co-addition stage into the real-time pipeline, and by honing our follow-up to sift through the increased number of candidates. The greatest factor limiting discoveries is, of course, that *Fermi* detects bursts all over the sky, only a fraction of which are visible from Palomar. Given our success so far, we enthusiastically suggest that other wide-field surveys implement a similar program. Furthermore, automatically sharing lists of candidates between longitudinally separated instruments would facilitate rapid identification and follow-up of the fastest fading events.

It is uncertain what directions future gamma-ray space missions will take. Some may be like *Swift*, able to rapidly train multiple on-board follow-up instruments on new targets. Even if they lack these capabilities, we should be able to routinely locate GRB afterglows and find their redshifts using targeted, ground-based optical transient searches similar to the one that we have described.

Looking beyond GRBs, the current iPTF TOO pipeline serves as a prototype for searching for optical counterparts of GW transients. Near the end of 2015, Advanced LIGO will begin taking data, with Advanced Virgo soon following suit. The first detections are anticipated by 2016 or later (Aasi et al., 2013c). On a similar timescale, iPTF will become the Zwicky Transient Facility, featuring a new 47 deg^2 survey camera that can reach $R = 20.4 \text{ mag}$ in 30 s. The prime GW sources, BNS mergers, may also produce a variety of optical transients: on- or off-axis afterglows (van Eerten & MacFadyen, 2011), kilonovae (Li & Paczyński, 1998; Barnes & Kasen, 2013a), and neutron-powered precursors (Metzger et al., 2015); see Figure 1.4 for some examples.

There will be two key challenges. First, GW localizations can be even coarser than *Fermi* GBM error circles. Starting around $\sim 600 \text{ deg}^2$ in the initial (2015) two-detector configuration (Kasliwal & Nissanke, 2014; Singer et al., 2014), the areas will shrink to $\sim 200 \text{ deg}^2$ with the addition of Virgo in 2016. They should reach $\sim 10 \text{ deg}^2$ toward the end of the decade as the three detectors approach final design sensitivity and can approach $\sim 1 \text{ deg}^2$ as additional planned GW facilities come online (LIGO–India and KAGRA; see Schutz 2011; Veitch et al. 2012; Fairhurst 2014; Nissanke et al. 2013; Aasi et al. 2013c). Since the detection efficiency of our GBM–iPTF afterglow search is consistent with the areas that we searched, we expect that even the earliest Advanced LIGO localizations will present no undue difficulties for ZTF when we consider its 15-fold increase in areal survey rate as compared to iPTF.

However, there is a second challenge that these optical signatures are predicted to be fainter than perhaps 22 mag (with the exception of on-axis afterglows, which should be rare but bright due to beaming). To confront this, we must perform deep (10 min–1 hour) integrations with the P48. To work through the larger number of contaminating sources, we propose the following strategies:

1. We may adopt a mix of two or more exposure depths in order to cover the area both rapidly and deeply, to be sensitive to both bright on-axis events and the fainter isotropic signatures. For instance, we may cover the entire GW localization once at a single-exposure depth, then repeat with deeper 10 min–1 hour exposures as the visibility allows. Note that because the LIGO antenna pattern is preferentially sensitive above and directly opposite of North America, we are optimistic that most early Advanced LIGO events should be promptly accessible from Palomar with long observability windows ([Kasliwal & Nissanke, 2014](#)).
2. We will increasingly automate the selection of targets for photometric follow-up. With more advanced machine learning algorithms and better leveraging of light curve history across multiple surveys, we plan to robotically trigger follow-up by multiple telescopes.
3. With longer TOO time blocks, we will have to begin follow-up of the most promising candidates while the P48 observations are ongoing, to capitalize on accessibility from Palomar.
4. To further prioritize candidates for follow-up and severely reduce false positives, we can use spatial proximity to nearby galaxies ([Nissanke et al., 2013](#)).
5. Our first experiences with detections and non-detections will guide decisions about the optimal filter. At the moment, kilonova models prefer redder filters (suggesting *i*-band), and precursor models prefer bluer (suggesting *g*-band).

The combination of LIGO, Virgo, *Fermi*, *Swift*, and iPTF/ZTF is poised to make major discoveries over the next few years, of which we have provided a small taste in this work. We offer both lessons learned and a way forward in this multimessenger effort. The ultimate reward will be joint observations of a compact binary merger in gamma, X-rays, optical, and GWs, giving us an exceptionally complete record of a complex astrophysical process: it will be almost as good as being there.

Acknowledgements

L.P.S. thanks generous support from the NSF in the form of a Graduate Research Fellowship.

The National Radio Astronomy Observatory is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc.

This paper is based on observations obtained with the Palomar 48-inch Oschin telescope and the robotic Palomar 60-inch telescope at the Palomar Observatory as part of the Intermediate Palomar Transient Factory 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. The present work is partly funded by *Swift* Guest Investigator Program Cycle 9 award 10522 (NASA grant NNX14AC24G) and Cycle 10 award 10553 (NASA grant NNX14AI99G).

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.

We thank Thomas Krühler for reducing the X-shooter spectrum of GRB 131011A / iPTF13dsw.

We thank the staff of the Mullard Radio Astronomy Observatory for their invaluable assistance in the operation of AMI. G.E.A., R.P.F., and T.D.S. acknowledge the support of the European Research Council Advanced Grant 267697, “4 Pi Sky: Extreme Astrophysics with Revolutionary Radio Telescopes.”

Support for CARMA construction was derived from the Gordon and Betty Moore Foundation, the Kenneth T. and Eileen L. Norris Foundation, the James S. McDonnell Foundation, the Associates of the California Institute of Technology, the University of Chicago, the states of California, Illinois, and Maryland, and the NSF. Ongoing CARMA development and operations are supported by the NSF under a cooperative agreement, and by the CARMA partner universities.

These results made use of Lowell Observatory’s DCT. Lowell operates the DCT in partnership with Boston University, Northern Arizona University, the University of Maryland, and the University of Toledo. Partial support of the DCT was provided by Discovery Communications. LMI was built by Lowell Observatory using funds from the NSF (AST-1005313).

A portion of this work was carried out at the Jet Propulsion Laboratory under a Research and Technology Development Grant, under contract with NASA. Copyright 2014 California Institute of Technology. All Rights Reserved. US Government Support Acknowledged.

K.H. acknowledges support for the IPN under the following NASA grants: NNX07AR71G, NNX13AP09G, NNX11AP96G, and NNX13AI54G.

IRAF is distributed by the National Optical Astronomy Observatory, which is operated by the Association of Universities for Research in Astronomy (AURA) under cooperative agreement with the NSF.

This research has made use of data, software and/or web tools obtained from HEASARC, a service of the Astrophysics Science Division at NASA/GSFC and of the Smithsonian Astrophysical Observatory's High Energy Astrophysics Division.

This research has made use of NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA.

This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester including the *Swift* XRT GRB catalog and light curve repository (Goad et al., 2007; Evans et al., 2007, 2009).

This research made use of Astropy¹⁶ (Robitaille et al., 2013), a community-developed core Python package for Astronomy. Some of the results in this paper have been derived using HEALPix (Górski et al., 2005).

¹⁶<http://www.astropy.org>

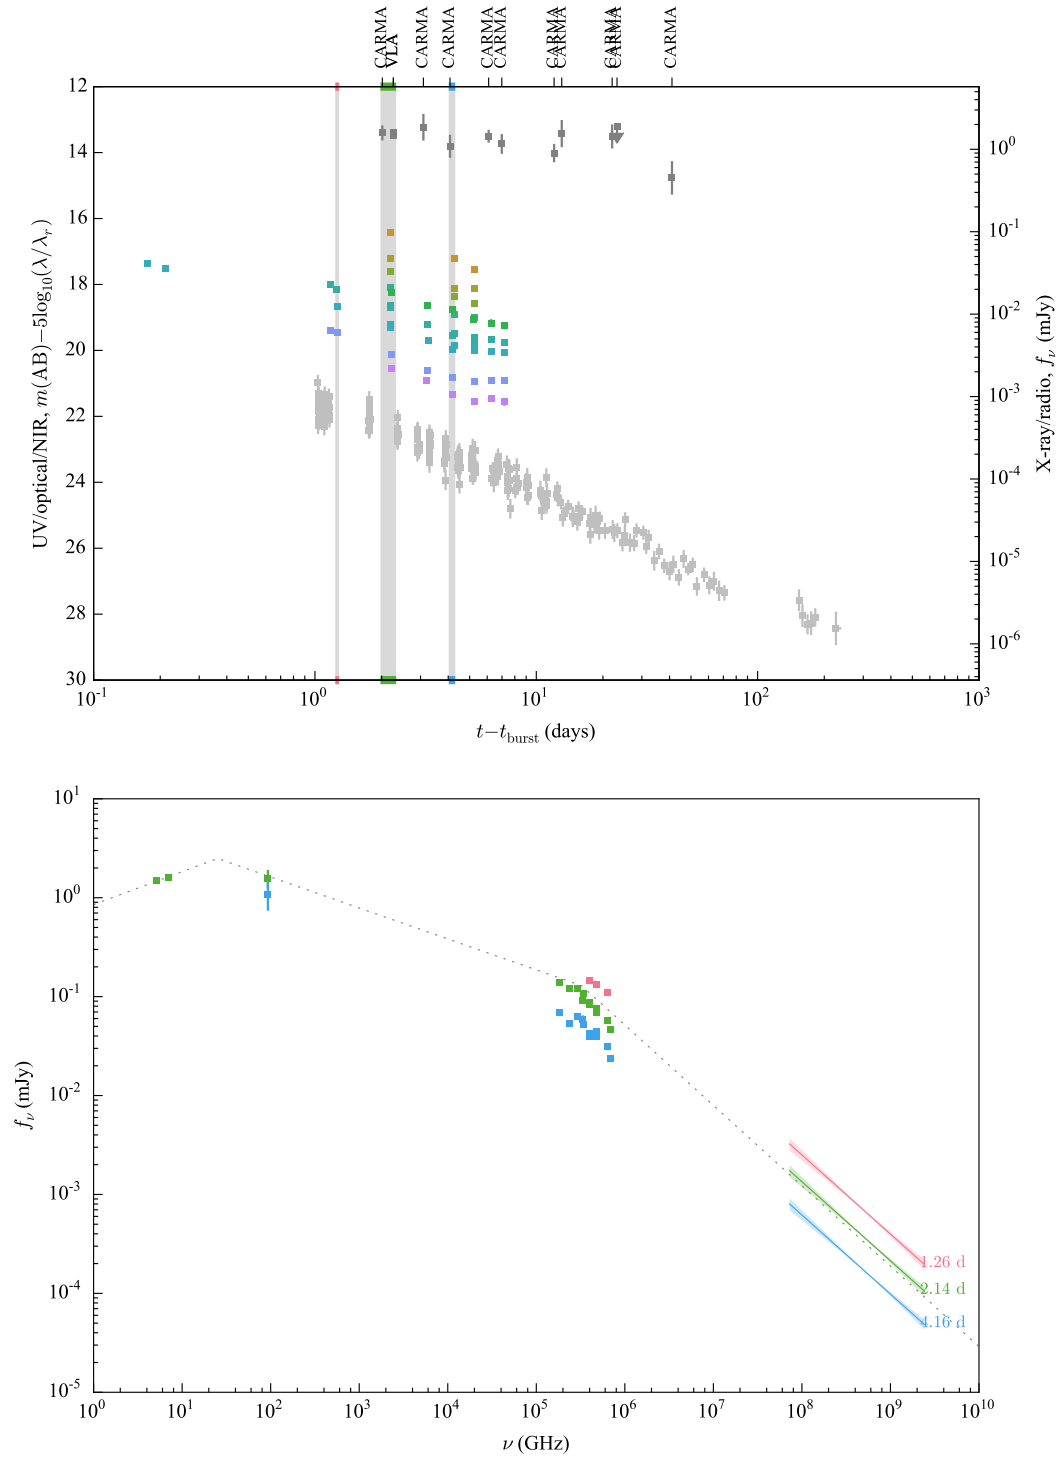


Figure 7.19 Light curve and SED of GRB 130702A / iPTF13bxl.




Figure 7.20 Light curve of GRB 131011A / iPTF13dsw.

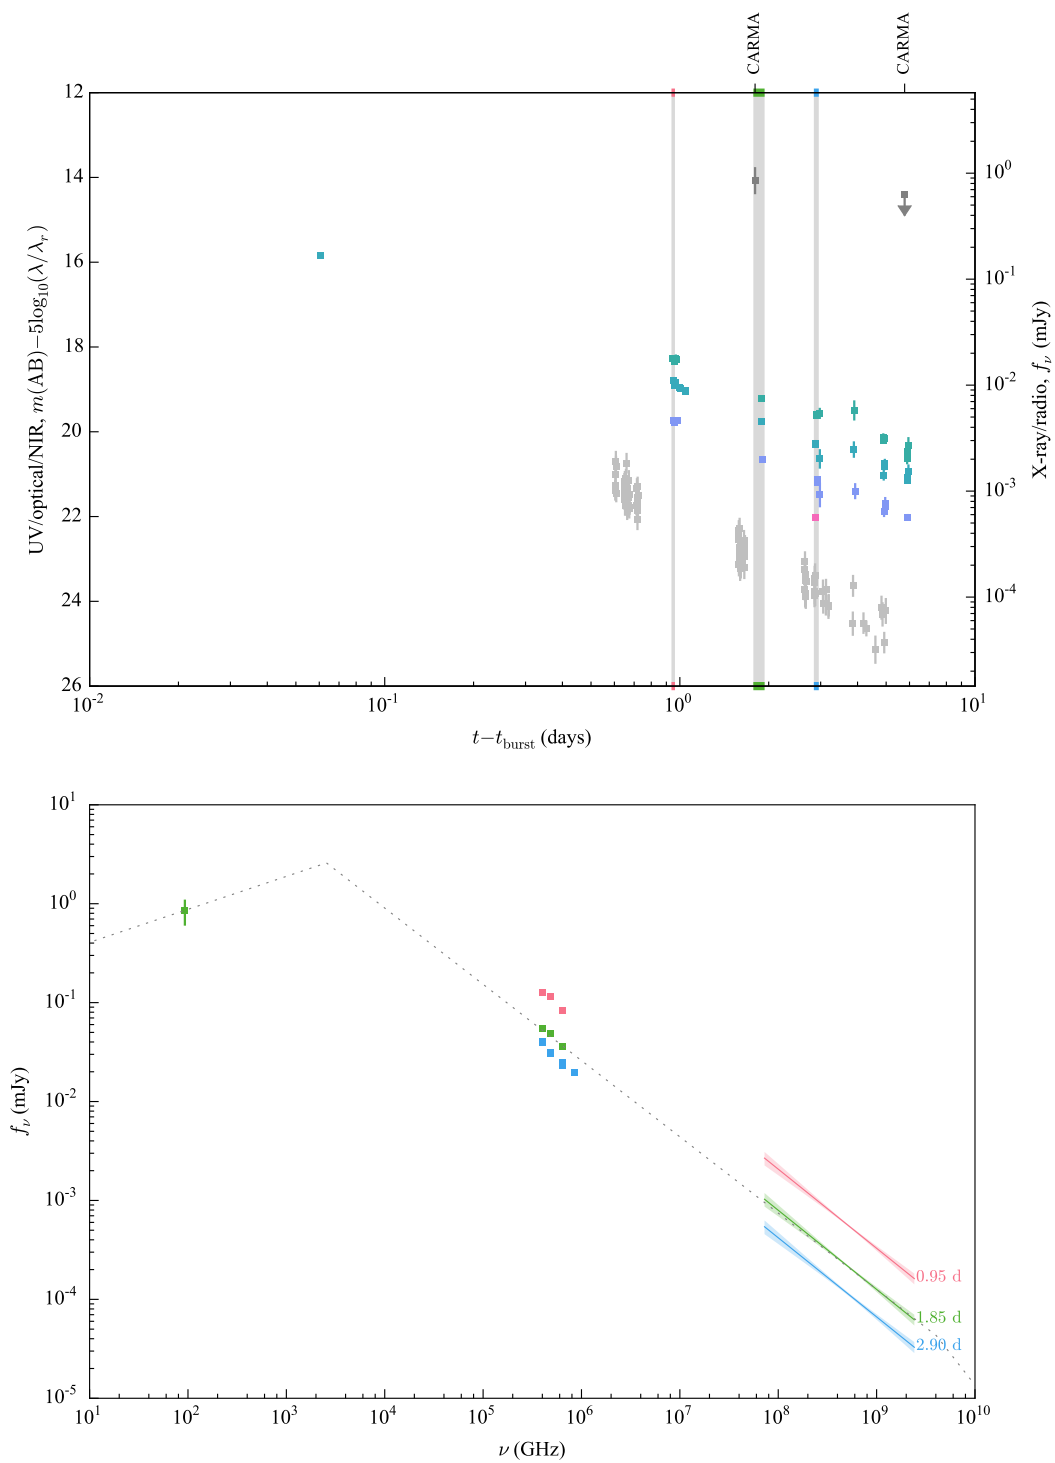


Figure 7.21 Light curve and SED of GRB 131231A / iPTF13ekl.

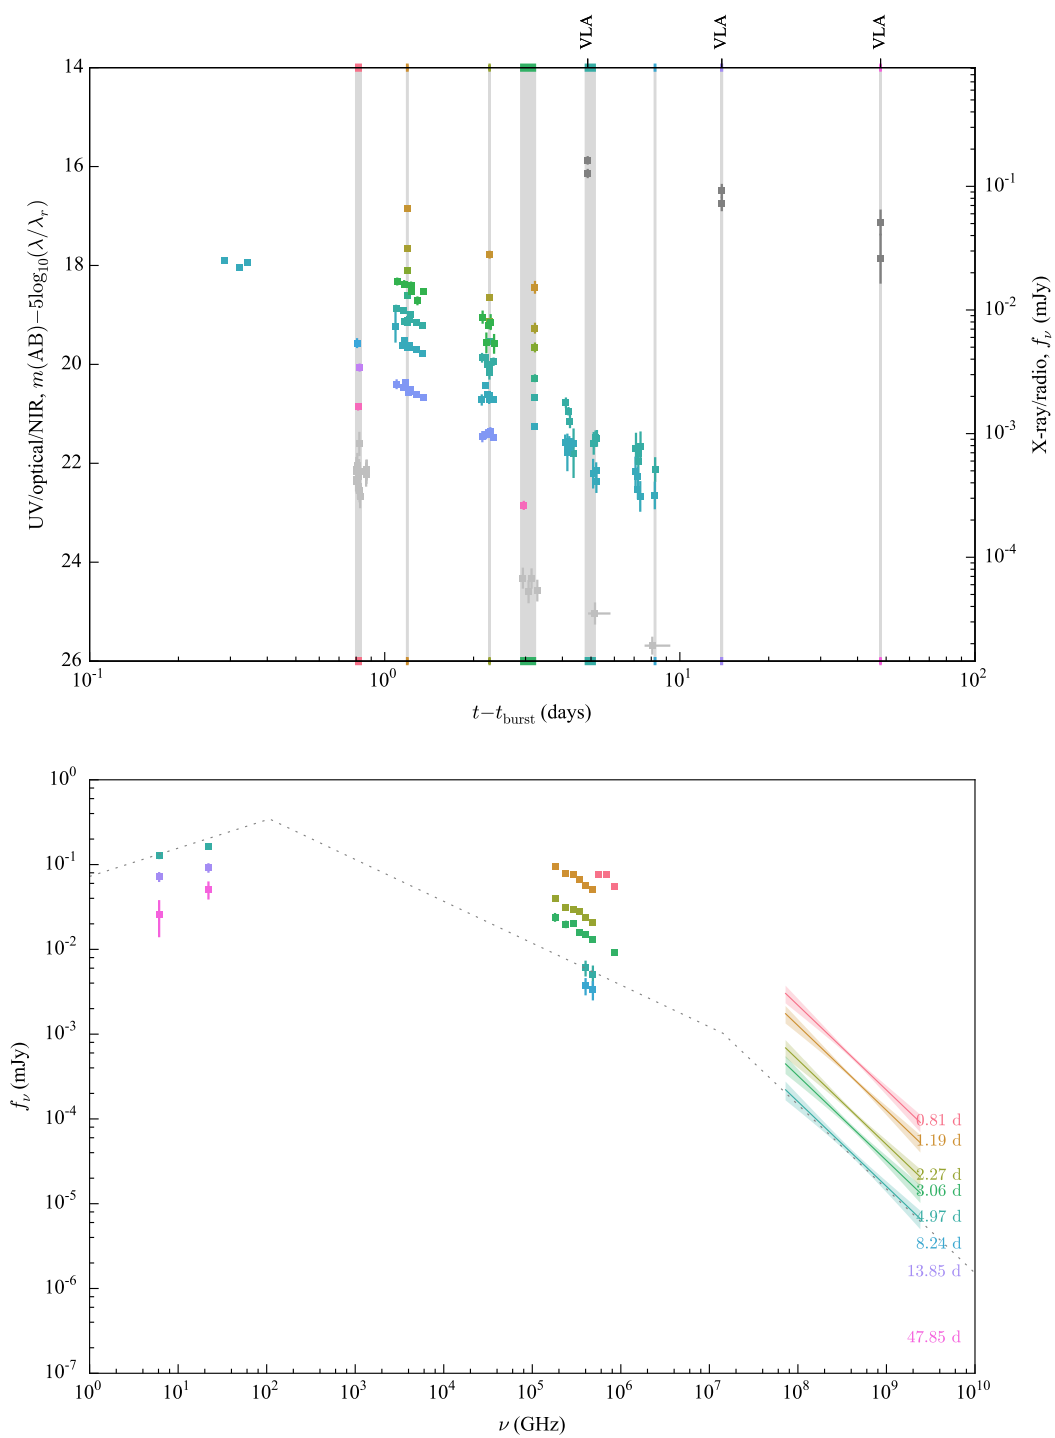


Figure 7.22 Light curve and SED of GRB 140508A / iPTF14aue.

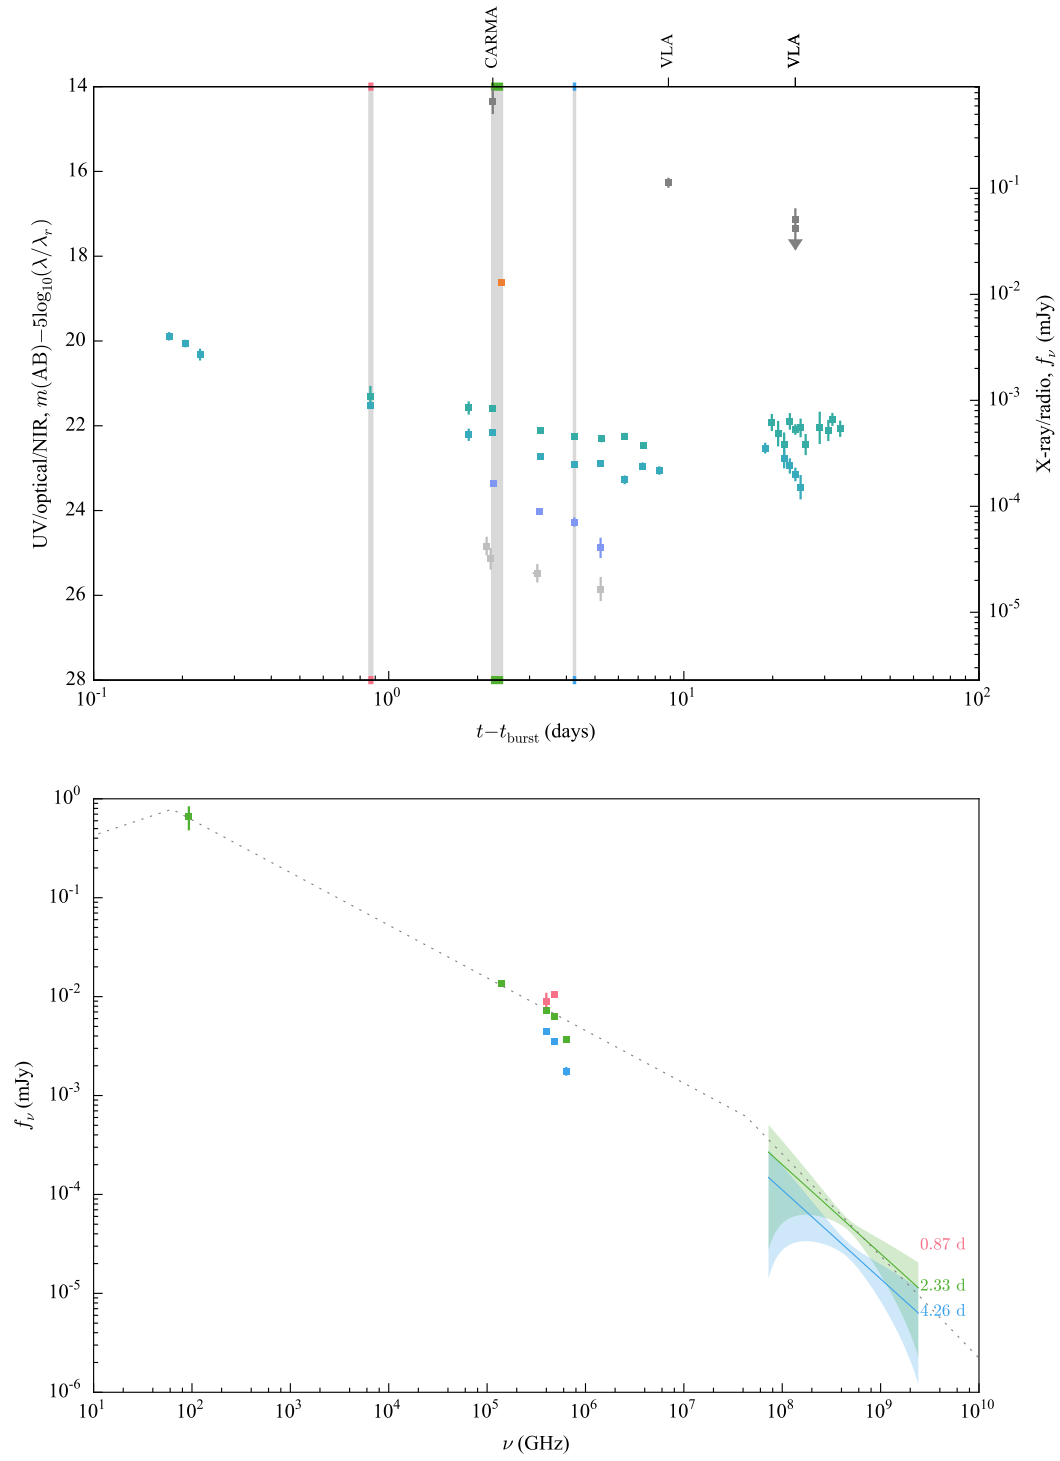


Figure 7.23 Light curve and SED of GRB 140606B / iPTF14bfu.

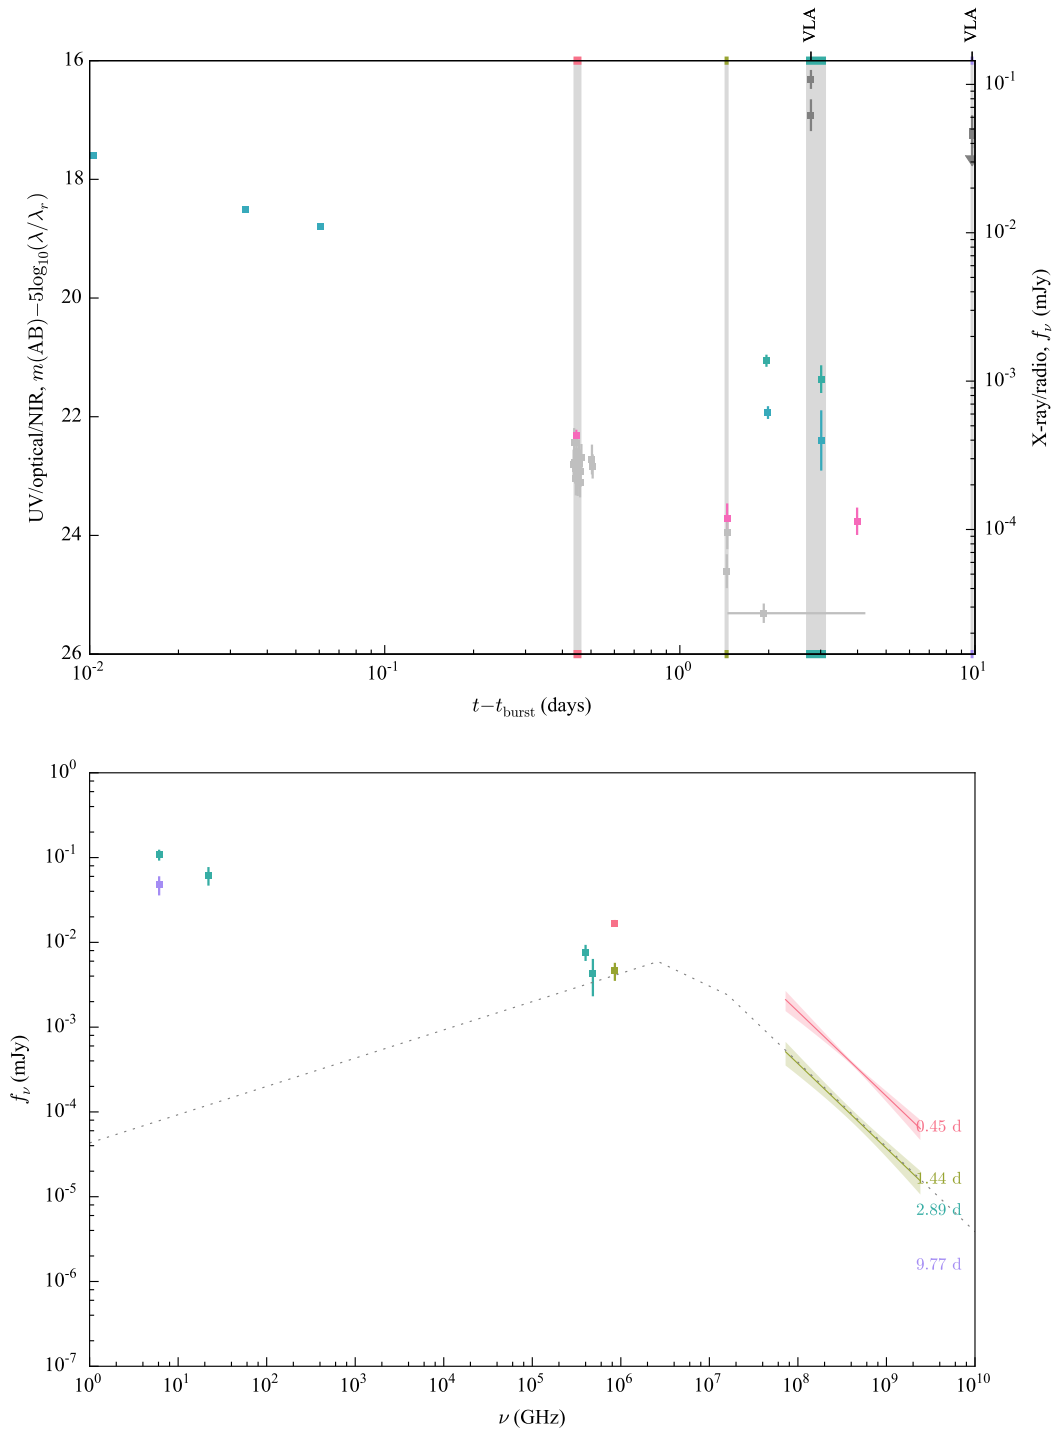


Figure 7.24 Light curve and SED of GRB 140620A / iPTF14cva.

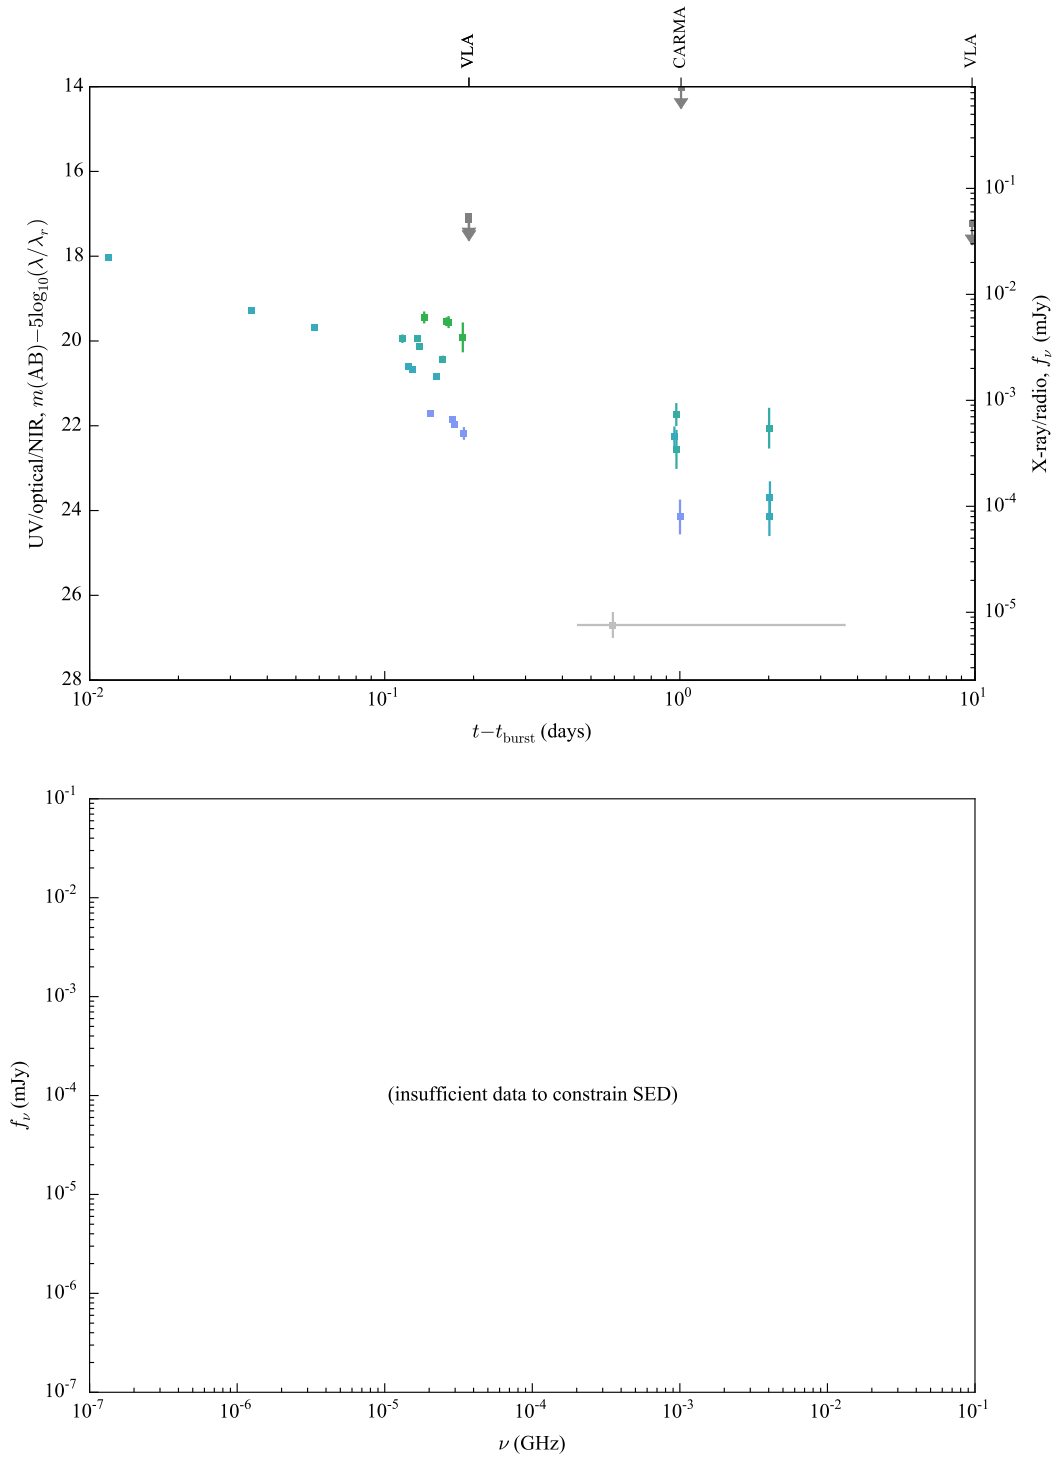


Figure 7.25 Light curve of GRB 140623A / iPTF14cyb.




Figure 7.26 Light curve and SED of GRB 140808A / iPTF14eag.

Chapter 8

Conclusion

In this thesis, I have described two of the main components of the Advanced LIGO real-time search for GWs from BNS mergers. First, I described a novel filter design and detection pipeline that can handle long-duration inspiral waveforms while remaining computationally tractable and without sacrificing latency. Curiously, the unique multiband structure lends itself to producing detection candidates a few seconds before merger for most events, and tens to a hundred seconds before merger for exceptional events.

Second, I developed a rapid Bayesian sky localization algorithm that is a thousand times faster than the full MCMC parameter estimation, while producing sky maps that are just as accurate. As a result of my contributions, the total response time of the whole online Advanced LIGO analysis is now low enough that it will be possible to search for all of the predicted optical counterparts of BNS mergers, including kilonovae, neutron-powered precursors, and afterglows.

I described in detail the detection and sky localization capabilities of the early two- and three-detector configurations of the global GW detector network. With the computational advances described above, we were able to simulate several months of Advanced LIGO observations and detect and localize thousands of BNS mergers. We show how the interplay between priors and polarization break parameter degeneracies. This overturns an old misconception that sky localization requires three or more detectors. In the case of two detectors, we elucidate one surviving degeneracy between the island of probability containing the true location of the source, and its polar opposite location. We quantify the improvement in position accuracy that is gained when Virgo comes online with reduced sensitivity. The large sample size, the astrophysically motivated source population, the detector configurations based on the commissioning and observ-

ing schedule, and the use of Bayesian parameter estimation instead of the Fisher matrix or other approximations, make this work a realistic description of the position reconstruction capabilities of advanced GW detectors. It also constitutes a demonstration of a complete pipeline for rapid detection and localization of BNS mergers in Advanced LIGO.

Switching from GWs to photons, within iPTF we have developed a unique wide-field TOO capability with which we can now routinely pinpoint out of areas of $\sim 100 \text{ deg}^2$ the afterglows of GRBs detected by *Fermi* GBM. So far, this has been a relatively rich source of low redshift, low-luminosity GRBs, with spectroscopically detectable SNe. Two of these bursts fit more cleanly in the picture of mildly relativistic shock breakout, rather than the standard picture of an internal shock within an ultra-relativistic jet. On a technical level, this serves as a demonstration and a prototype for finding optical counterparts of GW events with ZTF.

8.1 Next steps

Activities over the next few years are clearly critical to maximizing the early scientific returns from Advanced LIGO and ZTF. I will list a few clear next steps for the Advanced LIGO real-time BNS search and the iPTF TOO program.

8.1.1 Spin and NSBH mergers

For the reasons discussed in Section 5.1, I have ignored the effect of spin on BNS merger waveforms. NSs in field binaries formed through isolated binary stellar evolution should have small spins (the most rapidly spinning pulsar known in a BNS system having $\chi = |\chi| = |\mathbf{S}|/m^2 \lesssim 0.05$; Burgay et al. 2003) that can be neglected for the purposes of detection and sky localization with Advanced LIGO (Berry et al., 2014). However, in the dense environment of a globular cluster, a BNS system with appreciable component spins could be formed through dynamical capture of a young or accretion-recycled NS (up to the NS breakup limit of $\chi \sim 0.7$; see references in Brown et al. 2012). Furthermore, in NSBH binaries, the spin of the BH is expected to be an important effect because stellar-mass BHs are known to have large or nearly extremal spins ($\chi \sim 0.98$; see Miller & Miller 2014 and references therein). Effort is underway to incorporate spin effects into offline Advanced LIGO CBC searches using either non-precessing templates for spins that are aligned with the orbital angular momentum (Ajith et al., 2014; Brown et al., 2012) or precessing

templates for systems where one spin dominates (Pan et al., 2004; Harry & Fairhurst, 2011a). It is known that failing to account for spin can result in a loss of SNR (Brown et al., 2012) or can introduce biases in parameter estimation (Raymond et al., 2009). Therefore, incorporating spin at some level into the online Advanced LIGO detection pipeline may be essential to providing accurate rapid sky localizations for NSBH merger events.

8.1.2 Sub-threshold signals in rapid localization

As we have noted, the condition in which the rapid localization is *not* as accurate as the full parameter estimation is when one or more of the detectors is not represented by a trigger. This is especially an issue in the 2016 configuration, when Virgo is present but with a third of the sensitivity of the LIGO detectors. In this situation, the signal will often be below threshold in Virgo and represented by triggers in only the Hanford and Livingston detectors. In these cases, using the extra information from the weak signal in Virgo can reduce the area by a factor of one third. One remedy would be to reduce or eliminate the single-detector threshold, such that times, amplitudes, and phases on arrival are available for every detector and every event. This would probably require redesigning the coincidence stage of the detection pipeline. A simpler and more useful approach might be to have the detection pipeline send to BAYESTAR an excerpt from the SNR time series of the best matching template, extending a small fraction of a second before and after the time of the event. Within BAYESTAR, these small portions of the SNR time series would be used in the likelihood instead of the template autocorrelation function. Mathematically, this trivial change would make the BAYESTAR likelihood mathematically equivalent to what is used by the MCMC analysis. This simple change ought to be implemented before 2016 when Virgo comes online.

8.1.3 Distance-resolved rapid localizations

Initial LIGO EM counterparts relied upon imaging nearby galaxies that were contained within the GW localizations. Within the expanded detection volume of Advanced LIGO, there are enough galaxies that this approach will be of little help. However, Nissanke et al. (2013) have proposed to reduce the number of optical false positives by combining a galaxy catalog with the *joint* posterior probability distribution of sky location *and distance*. This could be an especially

important technique in the early years of Advanced LIGO, due to the initially coarse localizations of hundreds of deg^2 . Currently, BAYESTAR supplies flat, two-dimensional posterior distributions on the sky. This is simply because distance is treated as a marginal variable, and integrated away. There is no difficulty in calculating the distance posterior. For example, even without modifying the integration scheme one could simply run several instances of BAYESTAR in parallel with different distance limits spanning a sequence of shells. The challenge in supplying the distance information to astronomers is simply one of data management: the sky maps already take up several hundred kilobytes per event, and encoding the distance information naively would make the data products cumbersome. Some small effort must go into picking an appropriate representation of the full three-dimensional probability distributions, and determining whether some two-plus-one dimensional distribution (for example, sky posterior plus the mode of the marginal distance posterior at every pixel) would be sufficient.

8.1.4 GRBs beyond the *Fermi* bandpass

Still within the context of GRBs, there are several other possible applications of the iPTF TOO capability. At even higher energy scales than are probed by gamma-ray satellites, the High-Altitude Water Čerenkov Gamma-Ray Observatory (HAWC) recently began normal observations, at any instant monitoring about 15% of the sky for air showers induced by gamma-rays and cosmic rays with energies of $\sim 30\text{ GeV}$ – 100 TeV [Abeysekara et al. \(2014\)](#). HAWC is likely to detect 50–500 GeV emission for one or two GRBs per year ([Abeysekara et al., 2012](#); [Taboada & Gilmore, 2014](#)), with a typical localization uncertainty radius of $\lesssim 1^\circ$ (Ukwatta, private communication). Although this is too large for conventional optical follow-up, it is well within the reach of iPTF, the Rapid Telescopes for Optical Response (RAPTOR) network, or even tiled *Swift* XRT observations. We have submitted a *Swift* Guest Investigator proposal for X-ray follow-up of HAWC GRBs (P.I.: Ukwatta, LANL).

8.1.5 *Fermi* GBM and iPTF as a short GRB factory

As we noted in Chapter 7, *Fermi* GBM detects a large number of short bursts, about with perhaps 40 out of a total of 240 GRBs year^{-1} likely to be associated with CBC events ([Connaughton et al., 2014](#)). The sample of short GRB optical afterglows is still so small that even a few of these

Fermi bursts per year would be a significant contribution. Over the past year, with iPTF we have followed up a handful of *Fermi* short GRBs, but did not detect their afterglows. In these few cases, the optical limits were not constraining because the GBM localizations were only observable from Palomar only after a significant fraction of a day from the trigger. This is just due to chance, but we do expect it to be more difficult to find afterglows of short bursts than long bursts for two reasons. First, short GRBs are typically less well localized, because they have lower photon fluences. Second, their afterglows are about 6 mag fainter in an absolute sense, and the brightest short GRBs afterglows have apparent magnitudes that are comparable with the faintest long GRBs (Kann et al., 2011).

To address the first issue, we have increased the standard P48 TOO tiling from 10 to 20 fields. The larger FOV of the ZTF camera (see Figure 8.2) will also be a help here. For the second issue, we need deeper P48 observations. Since the P48 does not have a guide camera, the longest useful exposure time is about 60 s; we need to stack images. The real-time transient pipeline now has an automatic co-add stage prior to the image subtraction, but we need to modify the standard TOO program to take advantage of this. Due to the relatively long ~ 46 s overhead between exposures, for short GRBs we should probably tile the GBM localization several times without pauses between epochs. ZTF, on the other hand, will be able to reach deeper limits faster; current plans call for the ZTF transient survey to adopt a standard exposure duration of 30 s (Bellm, 2014), but variable exposure times of 10 min or more have been considered in the context of following up LIGO events (Kasliwal & Nissanke, 2014). Furthermore, with faster readout electronics and a reduced overhead of ~ 15 s between exposures, ZTF will be able to tile the localizations faster or reach deeper limits with stacked images.

In the context of iPTF co-adding consecutive two or three consecutive P48 images will modestly increase the number of candidates that need to be sorted. We will need to automate some of the vetting procedures that are currently done by humans. We will also need to speed up the human vetting by streamlining how candidates are presented. These modest improvements will be an important step in the evolution of the TOO program toward a depth that is relevant for Advanced LIGO and achievable with ZTF.

8.2 Future directions

In this section, I list some future directions for related research. I will first discuss two applications related to early-warning detection and sky localization. Then, I will end with some thoughts on TOOs with ZTF.

8.2.1 Early warning and dynamically tuned squeezing

The LIGO detectors are designed to be sensitive across a broad range of frequencies, but there exist several techniques that can improve the sensitivity across a narrow range of frequencies while sacrificing sensitivity at other frequencies. Since CBC signals are chirps, if we coupled a time-domain filtering scheme like LLOID to the control of the detector, then we could dynamically tune the narrowband sensitivity to track the frequency evolution of the signal. [Meers et al. \(1993\)](#) proposed dynamically detuning the signal recycling cavity, and [Simakov \(2014\)](#) showed with a detailed time-domain analysis that this scheme could result in SNR gains of a factor of 17. A disadvantage of this proposal is that the detector has a nontrivial transient response to changes in length of the signal recycling cavity.

An approach without this difficulty could use squeezed vacuum states of light. Squeezed states of light have less amplitude variation but more phase variation, or vice-versa, than a minimum-uncertainty coherent state. By injecting squeezed vacuum states into the dark port of the interferometer, the noise can be reduced below the standard quantum limit in a narrow frequency band that is determined by the squeezing angle. This has recently been demonstrated on the GEO 600 ([Ligo Scientific Collaboration et al., 2011](#)) and LIGO Hanford detectors ([Aasi et al., 2013a](#)). With an optimally frequency-dependent squeezing angle, it would be possible to beat the standard quantum limit across a broad range of frequencies. Unfortunately, frequency-dependent squeezing has an added cost in terms of infrastructure because current proposals involve external 100–1000 m filter cavities ([Kimble et al., 2002](#)). An alternative is to use the already-demonstrated frequency-independent squeezing, but dynamically tune the squeezing angle so that the narrowband sensitivity tracks the frequency evolution of a GW chirp. This would be simpler in some ways than tuning the signal recycling cavity because the detector would have no transient response to the change in squeezing angle.

Noise PSDs for Advanced LIGO with no squeezing, frequency-independent squeezing, and




Figure 8.1 Sensitivity curves for various squeezing configurations. The blue line represents the baseline Advanced LIGO design with no squeezing. The black and gray lines show configurations with 11 dB of squeezing injected with 15% losses, with squeezing angles varying from 90° to 0° in 15° increments. The orange line shows a frequency-dependent squeezing configuration obtained by sending the squeezed light through a filter cavity (length: 100 m, input mirror transmission: 1.85×10^{-4} , end mirror transmission: 3×10^{-6} , round trip losses: 1.5×10^{-5}).

frequency-dependent squeezing are shown in Figure 8.1. The filter cavity design with the fiducial parameters in the caption would increase the range for BNS mergers by 24%, almost doubling the detection rate. A dynamically tuned squeezing scheme would increase the BNS range by 32%, slightly better because the filter cavity itself has optical losses. The improvement in sensitivity would also enhance the estimation of all of the GW parameters (Lynch et al., 2014). In both scenarios, much of the improved sensitivity relative to no squeezing is at high frequencies, with the amplitude spectral density reduced by a factor of ≈ 0.56 at $f \gtrsim 500$ Hz. The GW signal from a BNS merger starts to be affected by tidal effects, deviating from the inspiral of point particles, at $f \sim 500$ – 1000 Hz. Therefore, the injection of squeezed vacuum states could make it possible to detect tidal effects and constrain the NS equation of state with Advanced LIGO.

For dynamically tuned squeezing, it remains to be shown is that the binary parameters can be

estimated with sufficient accuracy from the early inspiral to track the frequency evolution through the late inspiral. One must consider in what configuration to hold the detector while waiting for the start of a chirp, and at what threshold to activate squeezing or to start varying the squeezing angle. A first step would be to study the costs and benefits of dynamically tuned squeezing using Fisher matrix considerations similar to Chapter 2 and Section 3.1. For a practical implementation, one would have to think about whether the squeezing is activated based on an a coincidence of early-warning triggers from all of the sites, and whether feedback on the squeezing angle arises from the data from a single detector, from the network, or some intermediate combination. One would also need take into account that there will be some delay, due to technical sources of latency listed in Chapter 3, between the detection pipeline and actuation on the squeezer.

8.2.2 LIGO as a short GRB early-warning system

In Section 3.1, we briefly described using an early-warning GW detection to position *Swift* BAT in advance of the merger. This would require a radically new, fully autonomous TOO mode, and the rate of suitable early-warning detections would be constrained to about 1% of BNS events due to *Swift*'s ≈ 80 s slew time¹. Extremely low latency triggers might be more suitable for an instrument such as ISS-Lobster (Camp et al., 2013), a proposed X-ray imager on the International Space Station. ISS-Lobster would have a sensitivity of $1.3 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ over a band similar to *Swift* XRT, 0.3–5 keV. This would be sufficient to detect a short GRB X-ray afterglow at Advanced LIGO distances. It is envisioned as having an instantaneous FOV of 820 deg^2 and a slew time of ≈ 25 s (for a 60° slew, assuming a slew rate of 4° s^{-1} and acceleration and settling times of 5 s; Camp, personal communication). Although still only a few percent of BNS events will be detectable 25 s before merger, for about 20% of events it could be possible to have Lobster on settled on the GW localization within a 15 seconds after merger², especially if the initial pre-merger position estimate is updated with a refined localization every few seconds. Even using ordinary post-merger detections, with modest improvements to GW data handling one could begin soft X-ray observations within 25 s after merger, earlier than is usually possible with *Swift* XRT.

¹http://swift.gsfc.nasa.gov/proposals/tech.appd/swiffta_v10/node15.html

²Subject, of course, to the position of the ISS and any structural constraints on pointing.




Figure 8.2 The ZTF camera. Rendering reproduced from E. Bellm (private communication).

8.2.3 Optical counterpart search with ZTF

The prime sources for Advanced LIGO and Virgo are BNS mergers, which are also thought to be the progenitors of short GRBs (Paczynski, 1986; Eichler et al., 1989; Narayan et al., 1992; Rezzolla et al., 2011). As we have noted, short bursts and their afterglows are much fainter than long bursts. However, the redshifts of short GRBs that we might find in connection with GW detections would be limited to Advanced LIGO's range. In Figure 1.4, we show Kann's complete sample of short GRB afterglows with known redshifts and van Eerten & MacFadyen (2011)'s synthetic afterglow models, both shifted to Advanced LIGO's range at final design sensitivity of 200 Mpc or $z = 0.045$. By comparison with Figure 7.4, we see that short GRB afterglows at this distance will be just as bright as a typical long GRB afterglow at cosmological distances.

An added challenge is that GW localizations of BNS mergers can be more uncertain than *Fermi* GBM bursts. In the early (2015) two-detector network configuration, we expect probability maps that consist of multimodal arcs spanning up to 600 deg^2 (Kasliwal & Nissanke, 2014; Singer et al., 2014), though with the addition of Virgo in 2016 these shrink to 200 deg^2 (Aasi et al., 2013c). The typical areas should reach $\lesssim 10 \text{ deg}^2$ toward the end of the decade as the detectors approach

final design sensitivity and as additional planned GW facilities come online (LIGO–India and KAGRA; see [Schutz 2011](#); [Veitch et al. 2012](#); [Fairhurst 2014](#); [Nissanke et al. 2013](#); [Aasi et al. 2013c](#)). Since we have shown that our iPTF afterglow search has exactly the detection efficiency that is predicted by the areas that we search and the afterglow luminosity distribution, we expect that even the earliest Advanced LIGO localizations would present no undue difficulties for ZTF when we consider its 15-fold increase in areal survey rate as compared to iPTF.

Though *Swift* may detect a GRB in coincidence with a LIGO trigger, a *Fermi* GBM or IPN detection is much more likely on the basis of sky coverage. In these cases, efforts such as ours may still be necessary to pinpoint the event and initiate deeper follow-up.

Although afterglows are the brightest and most distinctive expected optical counterparts of BNS merger, they will probably also be quite rare considering that, like the GRBs themselves, they are tightly collimated within a jet with half-opening angle θ_j . If short GRBs have typical values of $\theta_j = 3\text{--}8^\circ$ as is suggested by recent jet break observations ([Fong et al., 2012](#)), then only a few percent of BNS events could be accompanied by GRBs.

[Metzger & Berger \(2012\)](#) point out that “kilonovae”, radioactively powered transients sourced in the hot *r*-process ejecta, may be a more promising counterpart because they are not expected to be beamed. When one considers realistic opacities, these signatures can be faint and very red, rising to only $R = 21\text{--}25$ mag but peaking in the infrared ([Barnes & Kasen, 2013a](#)). These would be challenging but sometimes possible to capture with ZTF using stacked exposures of 600 s or longer ([Kasliwal & Nissanke, 2014](#)). The “kilonova precursor” powered by free neutrons in the fast-moving outer ejecta ([Metzger et al., 2015](#)) may be a better prospect for ZTF: this signature could be much bluer and might rise as high as $R \approx 22$ mag in a matter of hours after the merger. A fast, bright, blue peak could also arise if the remnant persists as a hypermassive NS for $\gtrsim 100$ ms, supplying neutrinos to the ejecta and keeping the electron fraction too high to form the high-opacity Lanthanide elements ([Metzger & Fernández, 2014](#)). We have plotted some representative kilonova precursor models in Figure 1.4.

Contaminants for a kilonova search will be numerous due to the required survey depth, and may be different in nature from those which we discussed in Section 7.2. To deal with them, we will need to further automate the candidate vetting process and develop new target selection criteria. On the other hand, a search designed for kilonova precursors may be similar in many ways to our present afterglow effort, aside from using much longer and deeper exposures.




Figure 8.3 Range of kilonova absolute magnitudes detectable with various optical instruments within two-detector, HL type, GW localization areas of hundreds of deg^2 . The instruments shown include LSST, ZTF, Pan-STARRS (PS1), BlackGEM (BG4), Dark Energy Camera (DECAM), Hyper Suprime-Cam (HSC), and the Asteroid Terrestrial-impact Last Alert System (ATLAS). Reproduced with permission from [Kasliwal & Nissanke \(2014\)](#).

As scheduled, the first two years of Advanced LIGO will roughly coincide with the last two years of iPTF. The first detections could occur as early as 2016, but will have (median 90% confidence) positions that are uncertain by 200–600 deg², depending on when Advanced Virgo comes online and with what sensitivity. These localizations could take an hour or more to tile with iPTF, leaving little time for deep co-added observations. Given this constraint, except in the case of an exceptionally nearby and well-localized BNS merger, iPTF will be sensitive to on-axis afterglows but probably not kilonovae.

However, ZTF’s fast readout and much expanded FOV will enable much more rapid tiling; its guide camera will enable deep exposures. ZTF is expected to see first light in early 2017 and be fully commissioned by mid 2017 to early 2018. Considering also the expected improvements in the GW sensitivity, ZTF should be able to tile the GW localizations in a handful of pointings. Given present theoretical predictions, our ZTF follow-up strategy should be optimized for searching for kilonovae, without sacrificing early-time observations that could capture an optical afterglow. A mixed depth approach involving rapidly tiling with one or two epochs of 30 s exposures and then revisiting with one-hour exposures should be sufficient to capture almost any coincident on-axis afterglow, the fiducial kilonova precursor model, and a significant fraction of the predicted kilonova phase space (see Figure 8.3).

8.3 Conclusion

The next few years should be an exciting time, seeing both the dawn of GW transient astronomy and the maturing of synoptic optical transient surveys. Advanced LIGO should begin observing by late 2015, with Advanced Virgo following soon thereafter. The first direct detections of GW transients seem likely over the next few years. As a result of this thesis, we will be able to deliver GW detection candidates and accurate sky localizations promptly enough to look for all of the predicted optical counterparts of BNS mergers, including afterglows and kilonovae. Given current predictions for the EM signatures, some LIGO sources should be detectable by present and planned synoptic optical surveys, including PTF and ZTF, operating in TOO mode. The GBM–iPTF afterglow search serves as a prototype that must now rapidly evolve to fulfill its role in the search for optical signatures of BNS mergers.

Appendix A

Computer codes

A.1 Astropy

All chapters of this thesis made use of Astropy ([Robitaille et al. 2013](#); <http://www.astropy.org>), a community-developed core Python package for Astronomy. The author's contributions to Astropy included a numerical code for inverting World Coordinate System (WCS) transformations ([Greisen & Calabretta, 2002](#); [Calabretta & Greisen, 2002](#); [Greisen et al., 2006](#)), and an implementation of the image distortion corrections used in PTF images.

A.2 GSTLAL

Source code for GSTLAL is at <http://www.lsc-group.phys.uwm.edu/cgit/gstlal/tree>.

A.3 GStreamer

GStreamer is used in Chapter 3. Its project page is at <http://gstreamer.freedesktop.org>. It is widely installed on nearly all Linux and Unix desktop configurations. The author's contributions to GStreamer include enhancements to the upsampling/downsampling element `audioresample` to support rapidly skipping over gaps.

A.4 HEALPix

Some of the results in this paper have been derived using HEALPix ([Górski et al. 2005](#); <http://healpix.sourceforge.net>). The author's contributions to HEALPix include build infrastructure enhancements to support deployment on LIGO computing clusters and improve the availability of OpenMP-accelerated HEALPix routines in Python on the Macintosh platform.

A.5 LALSuite

BAYESTAR is part of the LALINFERENCE parameter estimation library, which is in turn part of the LIGO Algorithm Library Suite (<http://www.lsc-group.phys.uwm.edu/cgit/lalsuite/tree>). BAYESTAR's C language source code is in `lalinference/src/bayestar_sky_map.c`. Python bindings and high-level driver codes are in the directories `lalinference/python` and `lalinference/python/lalinference/bayestar`.

A.6 Code listing: sky resolution from Fisher matrices

The following Python listing calculates the coherent Fisher matrix described in Section 2.6.3.

```
"""
Fisher matrix calculation of detector resolution
Copyright (C) 2014 Leo Singer <leo.singer@ligo.org>
"""

from __future__ import division
import collections
import numpy as np
import scipy.optimize
import scipy.stats
import scipy.linalg

SkyLocalizationAccuracyBase = collections.namedtuple(
    'SkyLocalizationAccuracyBase',
    ('coherent_fisher_matrix', 'timing_fisher_matrix', 'snrs'))
```

```

def exp_i(phi):
    return np.cos(phi) + np.sin(phi) * 1j

class SkyLocalizationAccuracy(SkyLocalizationAccuracyBase):

    @staticmethod
    def _h(psi, u, r):
        """Polarization for phic=0 (arbitrary?)"""
        phic = 0
        cos2psi = np.cos(2*psi)
        sin2psi = np.sin(2*psi)
        f = 0.5 * (1 + np.square(u))
        pre = exp_i(2*phic) / r
        hp = pre * (f * cos2psi + 1j * u * sin2psi)
        hc = pre * (f * sin2psi - 1j * u * cos2psi)
        return hp, hc

    @staticmethod
    def _rot(phi, theta):
        """Active transformation from radiation to Earth frame."""
        c1 = np.cos(phi)
        c2 = np.cos(theta)
        s1 = np.sin(phi)
        s2 = np.sin(theta)
        return np.asarray(
            [[c1*c2, -s1, c1*s2],
             [c2*s1, c1, s1*s2],
             [-s2, 0, c2]])

    @staticmethod
    def _rescale(r, locations, r1s, w1s, w2s):
        """Rescale distance and time units."""
        wbar = np.sqrt(np.mean(w2s))
        rbar = scipy.stats.gmean(np.concatenate(([r], r1s)))
        return (
            r / rbar,
            locations * wbar,
            r1s / rbar,

```

```

        w1s / wbar,
        w2s / np.square(wbar))

    @staticmethod
    def _term(hp, hc, R, response, location, r1, w1, w2):
        # Transformed detector response tensor and position
        D = r1 * np.dot(np.dot(R.T, response), R)
        d = np.dot(R.T, location)

        # Amplitude
        z = hp*(D[0, 0] - D[1, 1]) + 2 * hc*D[0, 1]
        a = np.real(z)
        b = np.imag(z)
        snr2 = np.square(a) + np.square(b)

        # Jacobian matrix
        dz_dth = -2 * hp * D[0, 2] - 2 * hc * D[1, 2]
        dz_dph = -2 * hc * D[0, 2] + 2 * hp * D[1, 2]
        dz_drehp = D[0, 0] - D[1, 1]
        dz_drehc = 2 * D[0, 1]
        J = np.asarray([
            [np.real(dz_dth), np.real(dz_dph), dz_drehp, 0, dz_drehc, 0, 0],
            [np.imag(dz_dth), np.imag(dz_dph), 0, dz_drehp, 0, dz_drehc, 0],
            [-d[0], -d[1], 0, 0, 0, 0, 1]])

        I = np.asarray([
            [1, 0, w1*b],
            [0, 1, -w1*a],
            [w1*b, -w1*a, snr2*w2]])
        coherent_term = np.dot(np.dot(J.T, I), J)

        J = np.asarray([[-d[0], -d[1], 1]])
        I = snr2 * (w2 - np.square(w1))
        timing_term = np.dot(J.T, J) * I

        return coherent_term, timing_term, np.sqrt(snr2)

    @staticmethod
    def marginal_information(I, n):

```

```

"""Given an information matrix I, find the information matrix for the
marginal distribution of just the first n parameters by using
partitioned matrix inversion (the Schur complement)."""
A = I[:n, :n]
B = I[:n, n:]
C = I[n:, n:]
return A - np.dot(B, scipy.linalg.solve(C, B.T, sym_pos=True))

def __new__(
    cls, phi, theta, psi, u, r, responses, locations, r1s, w1s, w2s):

    r, locations, r1s, w1s, w2s = cls._rescale(r, locations, r1s, w1s, w2s)
    hp, hc = cls._h(psi, u, r)
    R = cls._rot(phi, theta)

    # Loop over detectors
    coherent_terms, timing_terms, snrs = zip(*[
        cls._term(hp, hc, R, response, location, r1, w1, w2)
        for response, location, r1, w1, w2
        in zip(responses, locations, r1s, w1s, w2s)])

    I_coherent = cls.marginal_information(
        np.sum(coherent_terms, axis=0), 2)
    I_timing = cls.marginal_information(
        np.sum(timing_terms, axis=0), 2)
    snrs = tuple(snrs)

    # Done!
    return super(SkyLocalizationAccuracy, cls).__new__(
        cls, I_coherent, I_timing, snrs)

    @staticmethod
    def _area(I, quantile):
        return -2*np.log(1-quantile)*180*180/np.pi / np.sqrt(np.linalg.det(I))

    @staticmethod
    def _width(I, quantile):
        ppf = scipy.stats.halfnorm.ppf(quantile)
        return 2 * ppf * 180 / np.pi / np.sqrt(np.max(np.linalg.eigvalsh(I)))

```

```
def coherent_area(self, quantile=0.9):
    if len(self.snrs) <= 2:
        raise ValueError(
            'Networks of <= 2 detectors are always degenerate. ' +
            'Use coherent_width instead to get the annulus width.')
    return self._area(self.coherent_fisher_matrix, quantile)

def timing_area(self, quantile=0.9):
    if len(self.snrs) <= 2:
        raise ValueError(
            'Networks of <= 2 detectors are always degenerate. ' +
            'Use timing_width instead to get the annulus width.')
    return self._area(self.timing_fisher_matrix, quantile)

def coherent_width(self, quantile=0.9):
    return self._width(self.coherent_fisher_matrix, quantile)

def timing_width(self, quantile=0.9):
    return self._width(self.timing_fisher_matrix, quantile)
```

Appendix B

Low frequency cutoff for inspiral searches

This Appendix is reproduced in part from [Cannon et al. \(2012\)](#), © 2012 The American Astronomical Society.

Ground-based GW detectors are unavoidably affected at low frequencies by seismic and anthropogenic ground motion. The LIGO test masses are suspended from multiple-stage pendula, which attenuate ground motion down to the pole frequency. In the detector configuration in place during S6, seismic noise dominated the instrumental background below about 40 Hz. Considerable effort is being invested in improving seismic attenuation in Advanced LIGO using active and passive isolation ([Harry & the LIGO Scientific Collaboration, 2010](#)), so that suspension thermal noise will dominate around 10–15 Hz. Inspiral waveforms are chirps of infinite duration, but since an interferometric detector’s noise diverges at this so-called “seismic wall,” templates for matched filter searches are truncated at a low-frequency cutoff f_{low} in order to save computational overhead with negligible loss of SNR.

The expected matched-filter SNR, integrated from f_{low} to f_{high} , is given by Equation (3.2). The high-frequency cutoff for the inspiral is frequently taken to be the GW frequency at the LSO; for non-spinning systems, $f_{\text{LSO}} = 4400(M_{\odot}/M)$ Hz, where M is the total mass of the binary (Section 3.4.1 of [Sathyaprakash & Schutz, 2009](#)). The choice of f_{low} is based on the fraction of the total SNR that is accumulated at frequencies above f_{low} . To illustrate the relative contributions to the SNR at different frequencies for a $(1.4, 1.4) M_{\odot}$ binary, we normalized and plotted the

integrand of Equation (3.2), the noise-weighted power spectral density of the inspiral waveform, in Figure B.1(b). This is the quantity

$$\frac{1}{\rho^2} \frac{d\rho^2}{df} = \frac{f^{-7/3}}{S(f)} \left(\int_0^{f_{\text{LSO}}} \frac{f'^{-7/3}}{S(f')} df' \right)^{-1},$$

which is normalized by the total SNR squared in order to put detectors with different absolute sensitivities on the same footing. We used several different noise power spectra: all of the envisioned Advanced LIGO configurations from Shoemaker (2010); the best-achieved sensitivity at LIGO Hanford Observatory (LHO) in LIGO’s fifth science run (S5), measured by Abadie et al. (2010a); and the best-achieved sensitivity at LHO during S6, measured by Abadie et al. (2012d). (The noise spectra themselves are shown in Figure B.1(a).) It is plain that during S5 and S6 the greatest contribution to the SNR was between 100 and 150 Hz, but for all of the proposed LIGO configurations the bulk of the SNR is accumulated below 60 Hz. This information is presented in a complementary way in Figure B.1(c), as the square root of the cumulative integral from f_{low} to f_{LSO} , interpreted as a fraction of the total “available” SNR,

$$\rho_{\text{frac}}(f_{\text{low}}) = \sqrt{\left(\int_{f_{\text{low}}}^{f_{\text{LSO}}} \frac{f^{-7/3}}{S(f)} df \right) \left(\int_0^{f_{\text{LSO}}} \frac{f^{-7/3}}{S(f)} df \right)^{-1}}.$$

Table B.1 shows the fractional accumulated SNR for four selected low-frequency cutoffs, 40 Hz, 30 Hz, 20 Hz, and 10 Hz. In S5 and S6, all of the SNR is accumulated above 40 Hz. For the ‘high frequency’ Advanced LIGO configuration, scarcely *half* of the SNR is accumulated above 40 Hz. For the preferred final configuration, ‘zero detuning, high power,’ 86.1% of the SNR is above 40 Hz, 93.2% is above 30 Hz, and 98.1% is above 20 Hz. (Since SNR accumulates in quadrature, this means, on the other hand, that under the ‘high frequency’ configuration a template encompassing *just the early inspiral* from 10 to 40 Hz would accumulate $\sqrt{1 - 0.533^2} \approx 84.6\%$ of the total SNR! In the ‘zero detuning, high power,’ configuration, integration from 10 to 40 Hz alone would yield 50.9% of the total SNR, from 10 to 30 Hz, 36.2%, and from 10 to 20 Hz, 19.4%.)

Since the GW amplitude is inversely proportional to the luminosity distance of the source, and the sensitive volume is proportional to distance cubed, the rate of detectable coalescences depends on the choice of low-frequency cutoff. An inspiral search that is designed with a low-frequency cutoff at the seismic wall would gain an increase in detection rate of $\rho_{\text{frac}}^{-3}(f_{\text{low}})$ relative to a search

Table B.1. Fractional accumulated SNR for selected low frequency cutoffs

Noise model	40 Hz	30 Hz	20 Hz	10 Hz
LHO (best S5)	100.0	100.0	100.0	100.0
LHO (best S6)	100.0	100.0	100.0	100.0
High frequency	53.3	80.1	97.6	100.0
No SRM	87.8	95.1	98.7	100.0
BHBH 20°	71.1	84.2	96.2	100.0
NSNS optimized	91.5	96.3	99.0	100.0
Zero detuning, low power	67.9	80.0	93.5	100.0
Zero detuning, high power	86.1	93.2	98.1	100.0

with a low-frequency cutoff of f_{low} . This would represent almost a twofold increase in the rate of detection over a search with a fractional accumulated SNR of 80%, and still a 37% increase over a search with $\rho_{\text{frac}} = 90\%$. Existing coalescing binary detection pipelines strive to sacrifice no more than 3% of the available SNR; this forfeits less than a 10% gain in detection rate. In order to satisfy this constraint, the low-frequency cutoff would have to be placed below 30 Hz for all of the conceived Advanced LIGO configurations.

The instantaneous GW frequency, given by Equation (3.1), is a power law function of time, so the amount of time for the GW frequency to evolve from f_{low} to f_{LSO} depends strongly on f_{low} . The duration of a $(1.4, 1.4) M_{\odot}$ inspiral is shown in Figure B.1(d). The inspiral takes only 25 s to evolve from 40 Hz to f_{LSO} , but takes 54 s to evolve from 30 Hz to f_{LSO} , 158 s from 20 Hz, and 1002 s from 10 Hz.

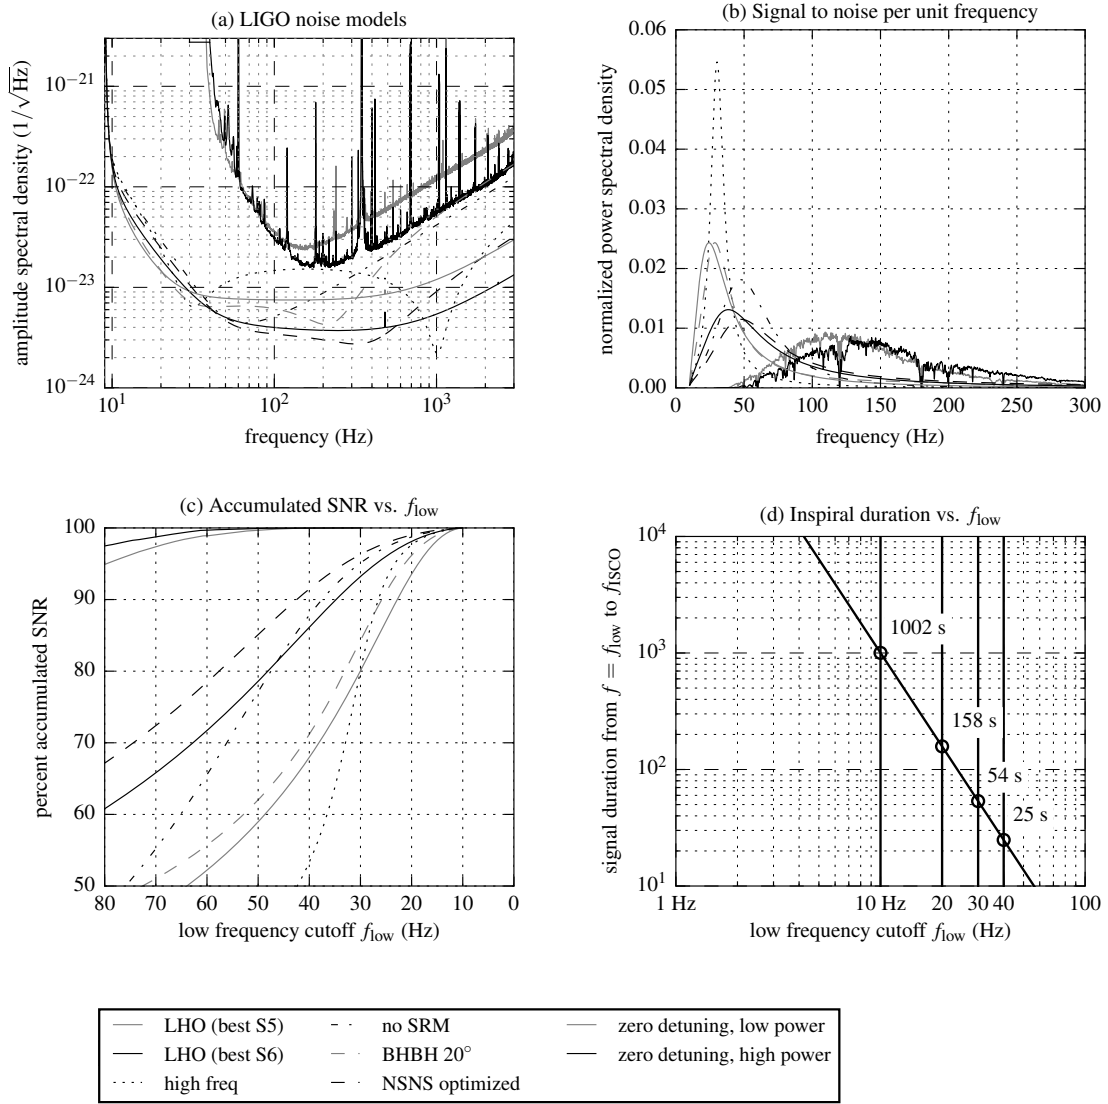


Figure B.1 From top left: (a) noise amplitude spectral density for a variety of Advanced LIGO noise models, S5, and S6. (b) Normalized signal-to-noise per unit frequency, $(d\rho^2/df)/\rho^2$, for a $(1.4, 1.4) M_{\odot}$ inspiral. (c) Percentage of SNR that is accumulated from f_{low} to f_{LSO} , relative to SNR accumulated from $f_{\text{low}} = 0$ Hz to f_{LSO} . (d) Amount of time for a NS-NS inspiral signal to evolve from frequency f_{low} to f_{LSO} , as a function of f_{low} . For (a)–(c), the line style indicates which noise model was used.

Appendix C

Sky map file format

This Appendix is reproduced from the technical document LIGO-T1300512-v7.

We are proposing to adopt HEALPix in FITS files as the official format for LIGO/Virgo GW sky maps, in order to facilitate exchange with astronomers. We suggest that we adopt the following conventions.

Note: See http://archive.stsci.edu/fits/fits_standard/ for a reference to the FITS standard and http://heasarc.gsfc.nasa.gov/docs/fcg/standard_dict.html for a dictionary of standard FITS header keywords.

C.1 Mandatory

1. Sky may use either the RING or NESTED indexing scheme.
2. Sky maps must use IMPLICIT indexing (full HEALPix array; not sparse indices).
3. Sky maps must use equatorial coordinates (COORDSYS = C in the FITS header).
4. Pixel data should be placed in a binary table with dimensions reshaped to rows of 1024 elements for backwards compatibility with IDL (pass `use_idl=True` when calling `healpy.write_map`).
5. Pixel values should denote the probability per square degree.
6. The name of the column containing the HEALPix pixel data should be PROB.

7. The units of the column containing the HEALPix pixel data should be given in the header as `pix-1`, to indicate that the sky map is given in units of probability per pixel. This goes in the `TUNITn` keyword in the header.
8. Skymaps should be viewable and reasonable in the Aladin¹ sky browser.

C.2 Optional

1. Sky maps may use any HEALPix resolution, at the creator's discretion.
2. Pixels that contribute negligibly to the total probability may be set to 0.
 - (a) If all pixels are nonzero, then the sum of all of the pixel values should be equal to 1.
 - (b) If some pixels are zero, then the sum of all pixel values should be less than or equal to 1.
3. The following optional entries may be added to the FITS header:
 - (a) `DATE`, UTC file creation date in ISO 8601.
 - (b) `DATE-OBS`, UTC trigger time in ISO 8601.
 - (c) `INSTRUME`, a comma-separated string of the style `H1,L1,V1` denoting which detectors contributed data to the event.
 - (d) `MJD-OBJ`, trigger time in modified Julian days.
 - (e) `OBJECT`, the GraceDB ID of the event, or any other unique identifier.
 - (f) `CREATOR`, the name of the program that created the FITS file.
 - (g) `REFERENC`, the URL of the GraceDB entry, or any other URL that provides further information about the event.
 - (h) `ORIGIN`, the organization responsible for the data (in most cases, probably LIGO/Virgo).
 - (i) `RUNTIME` (nonstandard), the elapsed time in seconds that the program took to create the sky map.
4. Sky maps may be gzip-compressed, signaled by the file extension `.fits.gz`.

¹<http://aladin.u-strasbg.fr/>

C.3 Example code in Python

Example code to read and write HEALPix/FITS sky maps with all optional metadata is in the Python module `lalinference.bayestar.fits`².

C.4 Specimen

Below is an example of what the FITS headers should look like.

```
$ funhead -a skymap.fits.gz
```

```
SIMPLE =                T / conforms to FITS standard
BITPIX =                8 / array data type
NAXIS  =                0 / number of array dimensions
EXTEND =                T
END
```

```
      Extension: xtension
```

```
XTENSION= 'BINTABLE'      / binary table extension
BITPIX  =                8 / array data type
NAXIS   =                2 / number of array dimensions
NAXIS1  =               4096 / length of dimension 1
NAXIS2  =               192 / length of dimension 2
PCOUNT  =                0 / number of group parameters
GCOUNT  =                1 / number of groups
TFIELDS =                1 / number of table fields
TTYPE1  = 'PROB          '
TFORM1  = '1024E         '
TUNIT1  = 'pix-1         '
PIXTYPE = 'HEALPIX      ' / HEALPIX pixelisation
ORDERING= 'RING          ' / Pixel ordering scheme, either RING or NESTED
COORDSYS= 'C             ' / Ecliptic, Galactic or Celestial (equatorial)
```

²<http://www.lsc-group.phys.uwm.edu/cgit/lalsuite/tree/lalinference/python/bayestar/fits.py>

```
EXTNAME = 'xtension'          / name of this binary table extension
NSIDE   =                    128 / Resolution parameter of HEALPIX
FIRSTPIX=                    0 / First pixel # (0 based)
LASTPIX =                   196607 / Last pixel # (0 based)
INDXSCHM= 'IMPLICIT'          / Indexing: IMPLICIT or EXPLICIT
OBJECT  = 'T73435'           / Unique identifier for this event
REFERENC= 'https://gracedb.ligo.org/events/T73435' / URL of this event
INSTRUME= 'H1,L1,V1'         / Instruments that triggered this event
DATE-OBS= '2013-02-16T03:19:17.438259' / UTC date of the observation
MJD-OBS =   56339.13839627615 / modified Julian date of the observation
DATE    = '2013-06-28T18:32:57' / UTC date of file creation
CREATOR  = 'bayestar_localize_lvalert' / Program that created this file
ORIGIN   = 'LIGO/Virgo'       / Organization responsible for this FITS file
RUNTIME  =  0.6783878803253174 / Runtime in seconds of the CREATOR program
END
```

Appendix D

“First Two Years” data release

This Appendix is reproduced from [Singer et al. \(2014\)](#), copyright © 2014 The American Astronomical Society.

We describe a catalog of all simulated events, detections, and sky maps that were generated for Chapter 5.

For the 2015 scenario, parameters of simulated signals are given in Table D.1. In the same order, parameters of the detection including the operating detector network, false alarm rate, ρ_{net} , SNR in each detector, recovered masses, and sky localization areas are given in Table D.2. For the 2016 scenario, the simulated signals are recorded in Table D.3 and the detections in Table D.4. In the print journal, parameters are given for just the four sample events that appear earlier in the text (see Figures 5.6, 5.7, 5.10, and 5.11). In the machine readable tables in the online journal, parameters are given for all detected signals.

The tables give two integer IDs. The “event ID” column corresponds to a field in that scenario’s full `gstlal` output that identifies one coincident detection candidate. The “simulation ID” likewise identifies one simulated signal. In the full `gstlal` output, there may be zero or many event candidates that match any given simulated signal. However, in our catalog there is one-to-one correspondence between simulation and event IDs because we have retained only simulated signals that are detected above threshold, and only the highest SNR detection candidate for each signal.

Note that the arbitrary dates of the simulated signals range from August 21 through 2010 October 19. This reflects the two-month duration of the simulated data stream, not the dates or

durations of the anticipated Advanced LIGO/Virgo observing runs.

For convenience, we also provide a browsable sky map catalog¹. This Web page provides a searchable version of Tables D.1–D.4, with posterior sky map images from both the rapid parameter estimation and the stochastic samplers.

The Web page also provides, for each localization, a FITS file representing the posterior in the HEALPix projection (Górski et al., 2005) using the NESTED indexing scheme. For reading these files, the authors recommend the Python package Healpy² or the HEALPix C/C++/IDL/Java/Fortran library³. They can also be displayed by many standard imaging programs such as DS9⁴ and Aladin⁵.

Synthetic GW time series data and posterior sample chains are available upon request.

¹<http://www.ligo.org/scientists/first2years>

²<http://healpy.readthedocs.org>

³<http://healpix.sourceforge.net>

⁴<http://ds9.si.edu>

⁵<http://aladin.u-strasbg.fr>

Table D.1. Simulated BNS signals for 2015 scenario

Event ID ^a	sim ID ^b	MJD ^c	Orientation ^d					d (Mpc)	Masses (M_{\odot})		Spin 1			Spin 2		
			α	δ	i	ψ	ϕ_c		m_1	m_2	S_1^x	S_1^y	S_1^z	S_2^x	S_2^y	S_2^z
18951	10807	55442.25864	137.8	-39.9	139	43	42	75	1.40	1.51	-0.01	-0.01	-0.04	-0.05	-0.01	+0.01
20342	21002	55454.76654	19.8	-23.7	145	197	145	75	1.34	1.48	-0.03	+0.01	-0.03	-0.01	+0.02	-0.01
...

^aIdentifier for detection candidate. This is the same value as the `coinc_event_id` column in the `gstlal` output database and the OBJECT cards in sky map FITS headers, with the `coinc_event:coinc_event_id:` prefix stripped.

^bIdentifier for simulated signal. This is the same value as the `simulation_id` column in the `gstlal` output database, with the `sim_inspiral:simulation_id:` prefix stripped.

^cTime of arrival at geocenter of GWs from last stable orbit.

^d α : RA, δ : Dec (J2000), i : binary orbital inclination angle, ψ : polarization angle ([Anderson et al., 2001](#), Appendix B), ϕ_c : orbital phase at coalescence.

Table D.2. Detections and sky localization areas for 2015 scenario

Event ID	sim ID	Network	Net ^a	SNR			Masses ^b		BAYESTAR			LALINFERENCE_NEST		
				H	L		m_1	m_2	50%	90%	Searched	50%	90%	Searched
18951	10807	HL	15.0	10.3	10.9		1.67	1.27	159	630	127	158	683	81.2
20342	21002	HL	12.7	7.3	10.3		1.59	1.25	126	526	16.9	168	618	12.3
...

^aNetwork SNR, or root-sum-squared SNR over all detectors.

^bMaximum likelihood estimate of masses as reported by `gstlal`.

Table D.3. Simulated BNS signals for 2016 scenario

Event ID	sim ID	MJD	Orientation					d (Mpc)	Masses (M_{\odot})		Spin 1			Spin 2		
			α	δ	ι	ψ	ϕ_c		m_1	m_2	S_1^x	S_1^y	S_1^z	S_2^x	S_2^y	S_2^z
655803	45345	55484.63177	79.2	+5.0	121	321	69	66	1.60	1.29	+0.00	+0.00	-0.00	+0.00	+0.00	-0.00
821759	8914	55439.93634	18.3	-15.1	158	257	230	187	1.60	1.45	-0.00	+0.02	-0.01	+0.04	+0.03	-0.02
...

Table D.4. Detections and sky localization areas for 2016 scenario

Event ID	sim ID	Network	Net	SNR			Masses		BAYESTAR			LALINFERENCE.MCMC		
				H ^a	L ^a	V ^a	m_1	m_2	50%	90%	Searched	50%	90%	Searched
655803	45345	HV	12.2	11.5		4.2	1.52	1.35	478	4570	65.5	304	3960	20.6
821759	8914	HLV	13.4	8.5	10.4		1.57	1.47	336	1070	473	91.0	515	93.8
...

^aBlank if SNR < 4 or detector is not online.

Appendix E

Observations of iPTF/GBM afterglows

*This Appendix contains the optical and radio observations from Chapter 7, which is in preparation as an article for the
The Astrophysical Journal.*

Table E.1. Optical observations of GBM-iPTF afterglows

Date (mid)	Inst. ^a	Δt^b	mag ^c
GRB 130702A / iPTF13bxi			
2013-07-02 04:18	P48	0.18	$R = 17.38 \pm 0.04$
2013-07-02 05:10	P48	0.21	$R = 17.52 \pm 0.04$
2013-07-03 04:13	P60	1.17	$g = 18.80 \pm 0.04$
2013-07-03 04:15	P60	1.17	$i = 18.42 \pm 0.04$
2013-07-03 06:16	P60	1.26	$i = 18.56 \pm 0.06$
2013-07-03 06:17	P60	1.26	$r = 18.66 \pm 0.05$
2013-07-03 06:20	P60	1.26	$g = 18.86 \pm 0.04$
2013-07-04 04:53	RATIR	2.20	$H = 18.56 \pm 0.03$
2013-07-04 04:53	RATIR	2.20	$J = 18.72 \pm 0.03$
2013-07-04 04:53	RATIR	2.20	$Y = 18.70 \pm 0.03$
2013-07-04 04:53	RATIR	2.20	$Z = 18.85 \pm 0.03$
2013-07-04 04:53	RATIR	2.20	$i = 19.06 \pm 0.02$
2013-07-04 04:53	RATIR	2.20	$r = 19.22 \pm 0.01$
2013-07-04 05:00	P60	2.20	$i = 19.12 \pm 0.04$
2013-07-04 05:01	P60	2.21	$r = 19.32 \pm 0.04$
2013-07-04 05:08	P60	2.21	$g = 19.52 \pm 0.04$
2013-07-04 05:10	P60	2.21	$B = 19.74 \pm 0.06$
2013-07-04 05:11	P60	2.21	$z = 19.02 \pm 0.08$
2013-07-05 05:02	P60	3.21	$B = 20.12 \pm 0.08$
2013-07-05 05:40	P60	3.23	$g = 20.02 \pm 0.05$
2013-07-05 05:47	P60	3.24	$i = 19.64 \pm 0.05$
2013-07-05 05:48	P60	3.24	$z = 19.43 \pm 0.09$
2013-07-05 06:04	P60	3.25	$r = 19.70 \pm 0.08$
2013-07-06 04:13	P60	4.17	$i = 19.97 \pm 0.06$
2013-07-06 04:24	P60	4.18	$r = 19.97 \pm 0.04$
2013-07-06 04:34	P60	4.19	$z = 19.53 \pm 0.10$
2013-07-06 04:44	P60	4.19	$g = 20.22 \pm 0.04$
2013-07-06 04:54	P60	4.20	$B = 20.53 \pm 0.07$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2013-07-06 06:05	RATIR	4.25	$H = 19.36 \pm 0.08$
2013-07-06 06:05	RATIR	4.25	$J = 19.64 \pm 0.07$
2013-07-06 06:05	RATIR	4.25	$Y = 19.46 \pm 0.05$
2013-07-06 06:05	RATIR	4.25	$Z = 19.68 \pm 0.05$
2013-07-06 06:05	RATIR	4.25	$i = 19.89 \pm 0.03$
2013-07-06 06:05	RATIR	4.25	$r = 19.86 \pm 0.02$
2013-07-07 04:55	P60	5.20	$z = 19.84 \pm 0.12$
2013-07-07 05:59	P60	5.25	$B = 20.75 \pm 0.08$
2013-07-07 06:20	RATIR	5.26	$H = 19.69 \pm 0.08$
2013-07-07 06:20	RATIR	5.26	$J = 19.64 \pm 0.06$
2013-07-07 06:20	RATIR	5.26	$Y = 19.69 \pm 0.05$
2013-07-07 06:20	RATIR	5.26	$Z = 19.76 \pm 0.04$
2013-07-07 06:20	RATIR	5.26	$i = 20.02 \pm 0.02$
2013-07-07 06:20	RATIR	5.26	$r = 19.94 \pm 0.02$
2013-07-07 06:20	P60	5.26	$i = 20.13 \pm 0.05$
2013-07-07 06:30	P60	5.27	$r = 20.00 \pm 0.04$
2013-07-07 06:40	P60	5.27	$g = 20.35 \pm 0.04$
2013-07-08 06:03	P60	6.25	$B = 20.68 \pm 0.06$
2013-07-08 06:19	P60	6.26	$z = 19.95 \pm 0.12$
2013-07-08 06:29	P60	6.27	$i = 20.09 \pm 0.05$
2013-07-08 06:37	P60	6.27	$r = 20.04 \pm 0.04$
2013-07-08 06:46	P60	6.28	$g = 20.31 \pm 0.04$
2013-07-09 04:09	P60	7.17	$B = 20.77 \pm 0.12$
2013-07-09 04:19	P60	7.18	$z = 20.04 \pm 0.11$
2013-07-09 04:29	P60	7.18	$i = 20.17 \pm 0.05$
2013-07-09 04:39	P60	7.19	$r = 20.06 \pm 0.04$
2013-07-09 04:49	P60	7.20	$g = 20.32 \pm 0.04$
GRB 131011A / iPTF13dsw			

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2013-10-12 05:26	P48	0.49	$R = 19.83 \pm 0.10$
2013-10-12 06:08	P48	0.51	$R = 19.86 \pm 0.07$
2013-10-12 06:43	P48	0.54	$R = 20.14 \pm 0.09$
2013-10-12 08:08	P60	0.60	$i = 19.80 \pm 0.06$
2013-10-12 08:11	P60	0.60	$r = 20.09 \pm 0.04$
2013-10-12 08:13	P60	0.60	$g = 20.40 \pm 0.03$
2013-10-12 10:20	P60	0.69	$B = 20.78 \pm 0.16$
2013-10-12 10:23	P60	0.69	$z = 20.24 \pm 0.30$
2013-10-12 10:56	P60	0.71	$B = 21.03 \pm 0.14$
2013-10-12 10:58	P60	0.72	$z = 20.01 \pm 0.20$
2013-10-13 04:15	GROND	1.44	$H = 20.50 \pm 0.20$
2013-10-13 04:15	GROND	1.44	$J = 21.10 \pm 0.10$
2013-10-13 04:15	GROND	1.44	$K = 20.20 \pm 0.40$
2013-10-13 04:15	GROND	1.44	$g' = 21.90 \pm 0.10$
2013-10-13 04:15	GROND	1.44	$i' = 21.50 \pm 0.10$
2013-10-13 04:15	GROND	1.44	$r' = 21.70 \pm 0.10$
2013-10-13 04:15	GROND	1.44	$z' = 21.30 \pm 0.10$
GRB 131231A / iPTF13ekl			
2013-12-31 06:12	P48	0.06	$R = 15.85 \pm 0.01$
2014-01-01 03:30	P60	0.95	$i = 18.68 \pm 0.03$
2014-01-01 03:32	P60	0.95	$r = 18.80 \pm 0.04$
2014-01-01 03:34	P60	0.95	$g = 19.14 \pm 0.03$
2014-01-01 03:43	P60	0.96	$i = 18.75 \pm 0.02$
2014-01-01 03:45	P60	0.96	$g = 19.17 \pm 0.04$
2014-01-01 03:46	P48	0.96	$R = 18.92 \pm 0.04$
2014-01-01 03:48	P48	0.96	$R = 18.89 \pm 0.03$
2014-01-01 03:53	P60	0.96	$r = 18.83 \pm 0.03$
2014-01-01 03:54	P60	0.96	$r = 18.84 \pm 0.01$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-01-01 04:00	P60	0.97	$i = 18.69 \pm 0.10$
2014-01-01 04:01	P60	0.97	$g = 19.14 \pm 0.05$
2014-01-01 04:03	P60	0.97	$i = 18.71 \pm 0.04$
2014-01-01 04:19	P60	0.98	$g = 19.14 \pm 0.02$
2014-01-01 04:46	P48	1.00	$R = 18.96 \pm 0.04$
2014-01-01 04:48	P48	1.00	$R = 18.99 \pm 0.04$
2014-01-01 05:47	P48	1.04	$R = 19.06 \pm 0.07$
2014-01-01 05:48	P48	1.04	$R = 19.03 \pm 0.06$
2014-01-02 02:01	P60	1.89	$r = 19.75 \pm 0.03$
2014-01-02 02:12	P60	1.89	$i = 19.62 \pm 0.03$
2014-01-02 02:32	P60	1.91	$g = 20.06 \pm 0.02$
2014-01-03 01:56	UVOT	2.88	$U = 20.78 \pm 0.08$
2014-01-03 02:02	P60	2.89	$r = 20.29 \pm 0.07$
2014-01-03 02:04	P60	2.89	$r = 20.27 \pm 0.08$
2014-01-03 02:16	P60	2.90	$i = 20.01 \pm 0.04$
2014-01-03 02:18	P60	2.90	$i = 20.02 \pm 0.06$
2014-01-03 02:50	P60	2.92	$g = 20.53 \pm 0.04$
2014-01-03 02:52	P60	2.92	$g = 20.60 \pm 0.07$
2014-01-03 04:15	P60	2.98	$i = 19.98 \pm 0.13$
2014-01-03 04:18	P60	2.98	$r = 20.64 \pm 0.23$
2014-01-03 04:20	P60	2.98	$g = 20.89 \pm 0.29$
2014-01-04 02:03	P60	3.89	$r = 20.42 \pm 0.19$
2014-01-04 02:20	P60	3.90	$i = 19.91 \pm 0.24$
2014-01-04 02:59	P60	3.93	$g = 20.80 \pm 0.19$
2014-01-05 02:08	P60	4.89	$i = 20.54 \pm 0.09$
2014-01-05 02:31	P60	4.91	$r = 21.02 \pm 0.13$
2014-01-05 02:38	P60	4.91	$i = 20.63 \pm 0.06$
2014-01-05 02:49	P60	4.92	$g = 21.29 \pm 0.12$
2014-01-05 03:01	P60	4.93	$r = 20.83 \pm 0.08$
2014-01-05 03:07	P60	4.93	$i = 20.58 \pm 0.07$
2014-01-05 03:31	P60	4.95	$r = 20.74 \pm 0.10$
2014-01-05 03:47	P60	4.96	$g = 21.17 \pm 0.11$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-01-05 03:54	P60	4.96	$g = 21.08 \pm 0.14$
2014-01-06 02:06	P60	5.89	$i = 21.05 \pm 0.13$
2014-01-06 02:08	P60	5.89	$i = 20.89 \pm 0.12$
2014-01-06 02:31	P60	5.91	$r = 21.07 \pm 0.07$
2014-01-06 02:33	P60	5.91	$r = 21.15 \pm 0.08$
2014-01-06 02:53	P60	5.92	$g = 21.42 \pm 0.08$
2014-01-06 03:31	P60	5.95	$i = 20.73 \pm 0.19$
2014-01-06 03:34	P60	5.95	$r = 20.93 \pm 0.16$
GRB 140508A / iPTF14aue			
2014-05-08 09:57	P48	0.29	$R = 17.89 \pm 0.01$
2014-05-08 10:47	P48	0.32	$R = 18.04 \pm 0.01$
2014-05-08 11:19	P48	0.34	$R = 17.95 \pm 0.02$
2014-05-08 22:24	UVOT	0.81	$V = 19.27 \pm 0.10$
2014-05-08 22:34	UVOT	0.81	$U = 19.62 \pm 0.07$
2014-05-08 22:48	UVOT	0.82	$B = 19.27 \pm 0.08$
2014-05-09 05:10	P60	1.09	$r = 19.24 \pm 0.32$
2014-05-09 05:17	P60	1.09	$i = 19.28 \pm 0.07$
2014-05-09 05:24	P60	1.10	$g = 19.80 \pm 0.10$
2014-05-09 05:31	P60	1.10	$z = 19.11 \pm 0.08$
2014-05-09 06:40	P60	1.15	$r = 19.62 \pm 0.04$
2014-05-09 06:46	P60	1.15	$i = 19.33 \pm 0.06$
2014-05-09 06:53	P60	1.16	$g = 19.86 \pm 0.05$
2014-05-09 07:00	P60	1.16	$z = 19.16 \pm 0.08$
2014-05-09 07:07	P60	1.17	$i = 19.55 \pm 0.07$
2014-05-09 07:09	P60	1.17	$r = 19.52 \pm 0.05$
2014-05-09 07:12	P60	1.17	$g = 19.77 \pm 0.06$
2014-05-09 07:43	RATIR	1.19	$H = 18.98 \pm 0.03$
2014-05-09 07:43	RATIR	1.19	$J = 19.18 \pm 0.03$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-05-09 07:43	RATIR	1.19	$Y = 19.20 \pm 0.03$
2014-05-09 07:43	RATIR	1.19	$Z = 19.36 \pm 0.03$
2014-05-09 07:43	RATIR	1.19	$i = 19.52 \pm 0.02$
2014-05-09 07:43	RATIR	1.19	$r = 19.65 \pm 0.02$
2014-05-09 07:50	P60	1.20	$i = 19.56 \pm 0.07$
2014-05-09 07:52	P60	1.20	$r = 19.62 \pm 0.05$
2014-05-09 07:55	P60	1.20	$g = 19.96 \pm 0.08$
2014-05-09 08:09	P60	1.21	$r = 19.66 \pm 0.02$
2014-05-09 08:11	P60	1.21	$r = 19.62 \pm 0.07$
2014-05-09 08:16	P60	1.22	$i = 19.50 \pm 0.04$
2014-05-09 08:18	P60	1.22	$i = 19.40 \pm 0.03$
2014-05-09 08:23	P60	1.22	$g = 19.94 \pm 0.08$
2014-05-09 08:25	P60	1.22	$g = 19.91 \pm 0.06$
2014-05-09 08:32	P60	1.23	$z = 19.32 \pm 0.04$
2014-05-09 08:32	P60	1.23	$z = 19.19 \pm 0.08$
2014-05-09 09:42	P60	1.28	$r = 19.70 \pm 0.02$
2014-05-09 09:49	P60	1.28	$i = 19.57 \pm 0.05$
2014-05-09 09:56	P60	1.29	$g = 20.01 \pm 0.06$
2014-05-09 10:03	P60	1.29	$z = 19.50 \pm 0.09$
2014-05-09 11:13	P60	1.34	$r = 19.77 \pm 0.02$
2014-05-09 11:20	P60	1.34	$i = 19.63 \pm 0.03$
2014-05-09 11:27	P60	1.35	$g = 20.07 \pm 0.06$
2014-05-09 11:34	P60	1.35	$z = 19.32 \pm 0.06$
2014-05-10 06:14	P60	2.13	$r = 20.72 \pm 0.11$
2014-05-10 06:21	P60	2.14	$i = 20.28 \pm 0.10$
2014-05-10 06:28	P60	2.14	$g = 20.86 \pm 0.11$
2014-05-10 06:35	P60	2.15	$z = 19.83 \pm 0.13$
2014-05-10 07:46	P60	2.20	$r = 20.42 \pm 0.06$
2014-05-10 07:53	P60	2.20	$i = 20.27 \pm 0.07$
2014-05-10 08:00	P60	2.21	$g = 20.81 \pm 0.10$
2014-05-10 08:07	P60	2.21	$z = 20.35 \pm 0.21$
2014-05-10 08:31	P60	2.23	$r = 20.60 \pm 0.04$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-05-10 08:38	P60	2.23	$i = 20.42 \pm 0.06$
2014-05-10 08:53	P60	2.24	$g = 20.79 \pm 0.08$
2014-05-10 09:00	P60	2.25	$z = 19.99 \pm 0.08$
2014-05-10 09:16	P60	2.26	$r = 20.72 \pm 0.08$
2014-05-10 09:23	P60	2.26	$i = 20.58 \pm 0.14$
2014-05-10 09:28	RATIR	2.27	$H = 19.92 \pm 0.08$
2014-05-10 09:28	RATIR	2.27	$J = 20.17 \pm 0.06$
2014-05-10 09:28	RATIR	2.27	$Y = 20.24 \pm 0.07$
2014-05-10 09:28	RATIR	2.27	$Z = 20.29 \pm 0.05$
2014-05-10 09:28	RATIR	2.27	$i = 20.47 \pm 0.03$
2014-05-10 09:28	RATIR	2.27	$r = 20.63 \pm 0.03$
2014-05-10 09:47	P60	2.28	$g = 20.75 \pm 0.08$
2014-05-10 09:54	P60	2.28	$z = 19.93 \pm 0.16$
2014-05-10 11:03	P60	2.33	$r = 20.72 \pm 0.05$
2014-05-10 11:10	P60	2.34	$i = 20.36 \pm 0.08$
2014-05-10 11:17	P60	2.34	$g = 20.88 \pm 0.06$
2014-05-10 11:24	P60	2.35	$z = 20.37 \pm 0.20$
2014-05-11 02:04	UVOT	2.96	$U = 21.61 \pm 0.09$
2014-05-11 08:34	RATIR	3.23	$H = 20.58 \pm 0.13$
2014-05-11 08:34	RATIR	3.23	$J = 20.79 \pm 0.11$
2014-05-11 08:34	RATIR	3.23	$Y = 20.76 \pm 0.10$
2014-05-11 08:34	RATIR	3.23	$Z = 21.04 \pm 0.08$
2014-05-11 08:34	RATIR	3.23	$i = 21.08 \pm 0.05$
2014-05-11 08:34	RATIR	3.23	$r = 21.25 \pm 0.05$
2014-05-12 05:09	P60	4.09	$r = 21.58 \pm 0.16$
2014-05-12 05:44	P60	4.11	$i = 21.19 \pm 0.11$
2014-05-12 06:51	P60	4.16	$r = 21.78 \pm 0.38$
2014-05-12 06:54	P60	4.16	$g > 21.68$
2014-05-12 07:44	P60	4.19	$i = 21.36 \pm 0.09$
2014-05-12 08:20	P60	4.22	$r = 21.68 \pm 0.07$
2014-05-12 08:22	P60	4.22	$r = 21.55 \pm 0.10$
2014-05-12 08:55	P60	4.24	$i = 21.57 \pm 0.13$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-05-12 11:29	P60	4.35	$r = 21.60 \pm 0.08$
2014-05-12 11:47	P60	4.36	$i = 22.21 \pm 0.50$
2014-05-13 00:36	UVOT	4.90	$U > 22.28$
2014-05-13 05:05	P60	5.08	$r = 22.21 \pm 0.30$
2014-05-13 05:38	P60	5.11	$i = 22.01 \pm 0.23$
2014-05-13 07:23	P60	5.18	$i = 21.89 \pm 0.14$
2014-05-13 07:25	P60	5.18	$r > 21.42$
2014-05-13 07:27	P60	5.18	$g > 21.13$
2014-05-13 08:00	P60	5.21	$r = 22.14 \pm 0.16$
2014-05-13 08:12	P60	5.21	$r = 22.36 \pm 0.24$
2014-05-13 08:45	P60	5.24	$i = 21.91 \pm 0.17$
2014-05-15 04:52	P60	7.07	$r = 22.17 \pm 0.30$
2014-05-15 05:30	P60	7.10	$i = 22.11 \pm 0.31$
2014-05-15 07:38	P60	7.19	$r > 21.25$
2014-05-15 07:41	P60	7.19	$g > 21.77$
2014-05-15 08:05	P60	7.21	$r = 22.53 \pm 0.20$
2014-05-15 08:06	P60	7.21	$r = 22.27 \pm 0.20$
2014-05-15 08:28	P60	7.23	$i = 22.22 \pm 0.19$
2014-05-15 08:40	P60	7.23	$i = 22.37 \pm 0.26$
2014-05-15 11:16	P60	7.34	$r = 22.67 \pm 0.31$
2014-05-15 11:39	P60	7.36	$i = 22.07 \pm 0.30$
2014-05-15 17:29	UVOT	7.60	$U > 21.84$
2014-05-15 17:47	UVOT	7.61	$V > 20.23$
2014-05-16 08:29	P60	8.23	$r = 22.65 \pm 0.28$
2014-05-16 08:59	P60	8.25	$i = 22.54 \pm 0.25$
2014-05-17 07:53	UVOT	9.20	$U > 21.75$
2014-05-19 01:27	UVOT	10.93	$U > 21.70$
GRB 140606B / iPTF14bfu			

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-06-06 07:31	P48	0.18	$R = 19.89 \pm 0.10$
2014-06-06 08:07	P48	0.20	$R = 20.06 \pm 0.09$
2014-06-06 08:42	P48	0.23	$R = 20.32 \pm 0.14$
2014-06-07 00:00	P60	0.87	$i = 21.72 \pm 0.25$
2014-06-07 00:00	P60	0.87	$r = 21.53 \pm 0.07$
2014-06-08 00:00	P60	1.87	$i = 21.99 \pm 0.16$
2014-06-08 00:00	P60	1.87	$r = 22.21 \pm 0.15$
2014-06-08 09:08	DCT	2.25	$r = 22.16 \pm 0.04$
2014-06-08 09:16	DCT	2.25	$i = 22.02 \pm 0.05$
2014-06-08 09:33	DCT	2.26	$g = 22.76 \pm 0.04$
2014-06-08 13:04	Keck I	2.41	$Ks = 21.34 \pm 0.06$
2014-06-09 09:04	DCT	3.24	$g = 23.41 \pm 0.06$
2014-06-09 09:27	DCT	3.26	$r = 22.72 \pm 0.04$
2014-06-09 09:28	DCT	3.26	$i = 22.53 \pm 0.07$
2014-06-10 09:16	DCT	4.25	$g = 23.67 \pm 0.11$
2014-06-10 09:34	DCT	4.27	$r = 22.91 \pm 0.06$
2014-06-10 09:44	DCT	4.27	$i = 22.66 \pm 0.05$
2014-06-11 08:33	DCT	5.22	$g = 24.28 \pm 0.24$
2014-06-11 08:39	DCT	5.23	$r = 22.89 \pm 0.07$
2014-06-11 09:37	DCT	5.27	$i = 22.72 \pm 0.08$
2014-06-12 10:24	DCT	6.30	$g > 23.41$
2014-06-12 10:30	DCT	6.30	$r = 23.28 \pm 0.10$
2014-06-12 10:35	DCT	6.31	$i = 22.66 \pm 0.08$
2014-06-13 09:08	DCT	7.25	$r = 22.95 \pm 0.09$
2014-06-13 10:38	DCT	7.31	$g > 23.88$
2014-06-13 10:50	DCT	7.32	$i = 22.87 \pm 0.07$
2014-06-14 09:23	DCT	8.26	$r = 23.06 \pm 0.10$
2014-06-25 00:00	P60	18.87	$r = 22.53 \pm 0.13$
2014-06-26 00:00	P60	19.87	$g > 22.71$
2014-06-26 00:00	P60	19.87	$i = 22.34 \pm 0.20$
2014-06-27 00:00	P60	20.87	$i = 22.60 \pm 0.30$
2014-06-27 00:00	P60	20.87	$r > 22.53$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-06-28 00:00	P60	21.87	$i = 22.85 \pm 0.28$
2014-06-28 00:00	P60	21.87	$r = 22.77 \pm 0.23$
2014-06-29 00:00	P60	22.87	$g > 23.74$
2014-06-29 00:00	P60	22.87	$i = 22.31 \pm 0.20$
2014-06-29 00:00	P60	22.87	$r = 22.95 \pm 0.18$
2014-06-30 00:00	P60	23.87	$g > 23.88$
2014-06-30 00:00	P60	23.87	$i = 22.50 \pm 0.13$
2014-06-30 00:00	P60	23.87	$r = 23.15 \pm 0.16$
2014-07-01 00:00	P60	24.87	$g > 23.41$
2014-07-01 00:00	P60	24.87	$i = 22.46 \pm 0.22$
2014-07-01 00:00	P60	24.87	$r = 23.45 \pm 0.29$
2014-07-02 00:00	P60	25.87	$i = 22.86 \pm 0.25$
2014-07-05 00:00	P60	28.87	$i = 22.46 \pm 0.38$
2014-07-07 00:00	P60	30.87	$i = 22.52 \pm 0.25$
2014-07-08 00:00	P60	31.87	$i = 22.26 \pm 0.15$
2014-07-10 00:00	P60	33.87	$i = 22.48 \pm 0.19$
GRB 140620A / iPTF14cva			
2014-06-20 05:30	P48	0.01	$R = 17.60 \pm 0.01$
2014-06-20 06:04	P48	0.03	$R = 18.51 \pm 0.02$
2014-06-20 06:43	P48	0.06	$R = 18.80 \pm 0.02$
2014-06-20 15:57	UVOT	0.45	$U = 21.08 \pm 0.10$
2014-06-21 15:58	UVOT	1.45	$U = 22.48 \pm 0.26$
2014-06-22 04:25	P60	1.97	$i = 21.47 \pm 0.10$
2014-06-22 05:00	P60	1.99	$r = 21.93 \pm 0.11$
2014-06-23 05:32	P60	3.01	$i = 21.78 \pm 0.23$
2014-06-23 05:35	P60	3.01	$r = 22.40 \pm 0.51$
2014-06-24 04:56	UVOT	3.99	$U = 22.52 \pm 0.23$
2014-06-24 06:30	P60	4.05	$i > 21.63$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-06-24 06:33	P60	4.05	$r > 22.65$
2014-06-28 01:50	UVOT	7.86	$U > 22.60$
GRB 140623A / iPTF14cyb			
2014-06-23 05:39	P48	0.01	$R = 18.04 \pm 0.01$
2014-06-23 06:13	P48	0.04	$R = 19.28 \pm 0.04$
2014-06-23 06:46	P48	0.06	$R = 19.69 \pm 0.06$
2014-06-23 08:08	P60	0.11	$i = 20.36 \pm 0.10$
2014-06-23 08:16	P60	0.12	$r = 20.60 \pm 0.06$
2014-06-23 08:21	P60	0.12	$r = 20.67 \pm 0.04$
2014-06-23 08:28	P60	0.13	$i = 20.35 \pm 0.05$
2014-06-23 08:31	P60	0.13	$i = 20.55 \pm 0.04$
2014-06-23 08:38	P60	0.14	$z = 20.23 \pm 0.14$
2014-06-23 08:48	P60	0.14	$g = 21.11 \pm 0.05$
2014-06-23 08:59	P60	0.15	$r = 20.83 \pm 0.06$
2014-06-23 09:09	P60	0.16	$i = 20.84 \pm 0.08$
2014-06-23 09:16	P60	0.16	$z = 20.33 \pm 0.10$
2014-06-23 09:19	P60	0.16	$z = 20.34 \pm 0.14$
2014-06-23 09:26	P60	0.17	$g = 21.25 \pm 0.04$
2014-06-23 09:29	P60	0.17	$g = 21.36 \pm 0.06$
2014-06-23 09:47	P60	0.18	$z = 20.70 \pm 0.35$
2014-06-23 09:49	P60	0.19	$g = 21.58 \pm 0.15$
2014-06-24 04:22	P60	0.96	$r = 22.25 \pm 0.23$
2014-06-24 04:43	P60	0.97	$i = 22.15 \pm 0.27$
2014-06-24 04:44	P60	0.97	$i = 22.97 \pm 0.46$
2014-06-24 05:03	P60	0.99	$z > 22.00$
2014-06-24 05:05	P60	0.99	$z > 22.06$
2014-06-24 05:24	P60	1.00	$g = 23.55 \pm 0.41$
2014-06-24 06:23	P60	1.04	$i > 21.60$

Table E.1 (cont'd)

Date (mid)	Inst. ^a	Δt^b	mag ^c
2014-06-24 06:27	P60	1.05	$z > 20.69$
2014-06-25 05:18	P60	2.00	$r > 23.50$
2014-06-25 05:31	P60	2.01	$i = 22.47 \pm 0.48$
2014-06-25 05:38	P60	2.01	$r = 24.14 \pm 0.46$
2014-06-25 05:50	P60	2.02	$r = 23.70 \pm 0.39$
2014-06-25 05:57	P60	2.02	$i > 21.94$
2014-06-25 06:23	P60	2.04	$i > 21.91$
GRB 140808A / iPTF14eag			
2014-08-08 04:15	P48	0.14	$R = 19.01 \pm 0.04$
2014-08-08 05:01	P48	0.17	$R = 19.27 \pm 0.04$
2014-08-08 05:49	P48	0.20	$R = 19.46 \pm 0.06$
2014-08-09 04:38	P60	1.16	$i = 21.85 \pm 0.19$
2014-08-09 05:01	P60	1.17	$r = 22.20 \pm 0.23$
2014-08-10 04:30	P60	2.15	$r = 22.84 \pm 0.35$
2014-08-10 04:53	P60	2.17	$i = 23.02 \pm 0.40$
2014-08-11 04:59	P60	3.17	$i = 23.01 \pm 0.24$

^aRATIR data is from [Butler et al. \(2013b,a, 2014b\)](#).
GROND data is from [Sudilovsky et al. \(2013\)](#).

^bTime in days relative to GBM trigger.

^cMagnitudes are in the AB system ([Oke & Gunn, 1983](#)).

Table E.2. Radio observations of GBM–iPTF afterglows

Date (start)	Inst. ^a	Δt^b	flux density ^c
GRB 130702A / iPTF13bxl			
2013-07-04	CARMA	2	$f_\nu(93) = 1580 \pm 330$
2013-07-04	VLA	2.3	$f_\nu(5.1) = 1490 \pm 75$
2013-07-04	VLA	2.3	$f_\nu(7.1) = 1600 \pm 81$
2013-07-05	CARMA	3.1	$f_\nu(93) = 1850 \pm 690$
2013-07-06	CARMA	4.1	$f_\nu(93) = 1090 \pm 350$
2013-07-08	CARMA	6.1	$f_\nu(93) = 1440 \pm 260$
2013-07-08	CARMA	7	$f_\nu(93) = 1160 \pm 320$
2013-07-14	CARMA	12	$f_\nu(93) = 900 \pm 230$
2013-07-15	CARMA	13	$f_\nu(93) = 1550 \pm 590$
2013-07-24	CARMA	22	$f_\nu(93) = 1430 \pm 480$
2013-07-25	CARMA	23	$f_\nu(93) < 1890$
2013-08-12	CARMA	41	$f_\nu(93) = 450 \pm 210$
GRB 131011A / iPTF13dsw			
2013-10-17	ATCA	5.7	$f_\nu(34) < 120$
GRB 131231A / iPTF13ekl			
2014-01-01	CARMA	1.8	$f_\nu(93) = 850 \pm 250$
2014-01-05	CARMA	5.8	$f_\nu(93) < 630$
GRB 140508A / iPTF14aue			

Table E.2 (cont'd)

Date (start)	Inst. ^a	Δt^b	flux density ^c
2014-05-13	VLA	4.9	$f_\nu(6.1) = 127 \pm 11$
2014-05-13	VLA	4.9	$f_\nu(22) = 162 \pm 13$
2014-05-22	VLA	14	$f_\nu(6.1) = 72 \pm 10$
2014-05-22	VLA	14	$f_\nu(22) = 92 \pm 12$
2014-06-25	VLA	48	$f_\nu(6.1) = 26 \pm 12$
2014-06-25	VLA	48	$f_\nu(22) = 51 \pm 12$
GRB 140606B / iPTF14bfu			
2014-06-08	CARMA	2.3	$f_\nu(93) = 660 \pm 180$
2014-06-15	VLA	8.9	$f_\nu(6.1) = 113 \pm 12$
2014-06-30	VLA	24	$f_\nu(6.1) = 51 \pm 12$
2014-06-30	VLA	24	$f_\nu(22) < 42$
GRB 140620A / iPTF14cva			
2014-06-23	VLA	2.8	$f_\nu(6.1) = 108 \pm 16$
2014-06-23	VLA	2.8	$f_\nu(22) = 62 \pm 15$
2014-06-30	VLA	9.8	$f_\nu(6.1) = 48 \pm 12$
2014-06-30	VLA	9.8	$f_\nu(22) < 45$
GRB 140623A / iPTF14cyb			
2014-06-23	VLA	0.19	$f_\nu(6.1) < 51$

Table E.2 (cont'd)

Date (start)	Inst. ^a	Δt^b	flux density ^c
2014-06-23	VLA	0.19	$f_\nu(22) < 54$
2014-06-24	CARMA	1	$f_\nu(93) < 900$
2014-07-03	VLA	9.8	$f_\nu(6.1) < 47$
GRB 140808A / iPTF14eag			
2014-08-08	AMI	0.48	$f_\nu(15) < 271$
2014-08-09	AMI	1.5	$f_\nu(15) < 127$
2014-08-10	VLA	2.1	$f_\nu(6.1) = 38 \pm 14$
2014-08-10	VLA	2.1	$f_\nu(22) = 130 \pm 20$
2014-08-12	AMI	4.6	$f_\nu(15) = 151 \pm 42$
2014-08-14	AMI	6.7	$f_\nu(15) = 185 \pm 80$
2014-08-16	AMI	8.6	$f_\nu(15) < 105$
2014-08-18	VLA	10	$f_\nu(6.2) = 81 \pm 17$
2014-08-18	VLA	10	$f_\nu(22) = 209 \pm 24$
2014-08-18	AMI	11	$f_\nu(15) = 241 \pm 58$
2014-08-20	AMI	13	$f_\nu(15) = 184 \pm 39$
2014-08-22	AMI	15	$f_\nu(15) = 134 \pm 35$
2014-08-27	AMI	20	$f_\nu(15) = 170 \pm 40$
2014-08-30	AMI	23	$f_\nu(15) = 147 \pm 43$
2014-09-03	VLA	26	$f_\nu(6.1) = 104 \pm 12$
2014-09-03	VLA	26	$f_\nu(22) = 112 \pm 14$
2014-09-04	AMI	28	$f_\nu(15) < 126$
2014-09-06	AMI	29	$f_\nu(15) < 206$
2014-09-09	AMI	32	$f_\nu(15) < 144$
2014-09-12	AMI	35	$f_\nu(15) < 202$

^aThe ATCA observation is from [Hancock et al. \(2013\)](#).

^bTime in days relative to GBM trigger.

^cFlux density in μJy as a function of frequency in GHz. For detections, the confidence intervals are 1σ statistical uncertainties added in quadrature with an estimated 5% systematic error. For non-detections, we show 3σ upper limits.

Acronyms

2MASS 2 Micron All-Sky Survey

AMI Arcminute Microkelvin Imager

AGN active galactic nucleus

ATCA Australia Telescope Compact Array

ATLAS Asteroid Terrestrial-impact Last Alert System

BAT Burst Alert Telescope (instrument on *Swift*)

BAYESTAR BAYESian TriAngulation and Rapid localization

BBH binary black hole

BH black hole

BNS binary neutron star

CARMA Combined Array for Research in Millimeter-wave Astronomy

CASA Common Astronomy Software Applications

CFH12k Canada–France–Hawaii 12 288 × 8 192 pixel CCD mosaic (instrument formerly on the Canada–France–Hawaii Telescope, now on the P48)

CRTS Catalina Real-Time Sky Survey

CTIO Cerro Tololo Inter-American Observatory

CBC compact binary coalescence

CCD charge coupled device

CRLB Cramér–Rao lower bound

cWB Coherent WaveBurst

DBSP Double Spectrograph (instrument on P200)

DCT Discovery Channel Telescope

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

EM electromagnetic

FD frequency domain

FAR false alarm rate

FFT fast Fourier transform

FIR finite impulse response

FITS Flexible Image Transport System

FLOPS floating point operations per second

FOV field of view

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

GCN Gamma-ray Coordinates Network

GMOS Gemini Multi-Object Spectrograph (instrument on the Gemini telescopes)

GRB gamma-ray burst

GSL GNU Scientific Library

GW gravitational wave

HAWC High-Altitude Water Čerenkov Gamma-Ray Observatory

HCT Himalayan Chandra Telescope

HEALPix Hierarchical Equal Area isoLatitude Pixelization

HEASARC High Energy Astrophysics Science Archive Research Center

HFOSC Himalaya Faint Object Spectrograph and Camera (instrument on HCT)

HMXB high-mass X-ray binary

HSC Hyper Suprime-Cam (instrument on the 8.2-m Subaru telescope)

IIR infinite impulse response

IMACS Inamori-Magellan Areal Camera & Spectrograph (instrument on the Magellan Baade telescope)

IPAC Infrared Processing and Analysis Center

IPN InterPlanetary Network

iPTF Intermediate Palomar Transient Factory

ISM interstellar medium

LAT Large Area Telescope

LHO LIGO Hanford Observatory

LIGO Laser Interferometer GW Observatory

llGRB low-luminosity GRB

LLOID Low Latency Online Inspiral Detection

LLO LIGO Livingston Observatory

LMI Large Monolithic Imager (instrument on DCT)

LSC LIGO Scientific Collaboration

LSO last stable orbit

LSST Large Synoptic Survey Telescope

MAP maximum a posteriori

MBTA Multi-Band Template Analysis

MCMC Markov chain Monte Carlo

MLE ML estimator

ML maximum likelihood

NED NASA/IPAC Extragalactic Database

NSBH neutron star–black hole

NSBH neutron star–black hole

NSF National Science Foundation

NS neutron star

P48 Palomar 48-inch Oschin telescope

P60 robotic Palomar 60-inch telescope

P200 Palomar 200-inch Hale telescope

PSD power spectral density

PTF Palomar Transient Factory

RAPTOR Rapid Telescopes for Optical Response

RMS root mean square

ROTSE Robotic Optical Transient Search

S5 LIGO’s fifth science run

S6 LIGO’s sixth science run

SNR signal-to-noise ratio

SDSS Sloan Digital Sky Survey

SED spectral energy distribution

SN supernova

SN Ic-BL broad-line type Ic SN

SVD singular value decomposition

TD time domain

TOO target of opportunity

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

VLA Karl G. Jansky Very Large Array

WAM Wide-band All-sky Monitor (instrument on *Suzaku*)

WCS World Coordinate System

XRF X-ray flash

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

ZTF Zwicky Transient Facility

Bibliography

- Aasi, J., et al. 2013a, *Nature Photonics*, 7, 613
- . 2013b, *Phys. Rev. D*, 88, 062001
- . 2013c, *arXiv:1304.0670*
- . 2014, *ApJS*, 211, 7
- Abadie, J., et al. 2010a, LIGO-T0900499-v19
- . 2010b, *Class. Quantum Grav.*, 27, 173001
- . 2012a, *A&A*, 541, A155
- . 2012b, *A&A*, 539, A124
- . 2012c, *Phys. Rev. D*, 85, 082002
- . 2012d, *arXiv:1203.2674*
- Abeysekara, A. U., et al. 2012, *Astroparticle Physics*, 35, 641
- . 2014, *ApJ*, 796, 108
- Acernese, F., et al. 2013, *Tech. Rep. VIR-0027A-09*
- Adhikari, R. X. 2014, *Rev. Mod. Phys.*, 86, 121
- Ajith, P., & Bose, S. 2009, *Phys. Rev. D*, 79, 084032
- Ajith, P., et al. 2014, *Phys. Rev. D*, 89, 084041
- Allen, B. 2005, *Phys. Rev. D*, 71, 062001

- Allen, B., et al. 2012, *Phys. Rev. D*, 85, 122006
- Amaral-Rogers, A. 2014a, *GCN*, 16232, 1
- . 2014b, *GCN*, 16254, 1
- Amati, L. 2006, *MNRAS*, 372, 233
- Amati, L., et al. 2013, *GCN*, 15025, 1
- Amati, L., Frontera, F., & Guidorzi, C. 2009, *A&A*, 508, 173
- Amati, L., et al. 2008, *MNRAS*, 391, 577
- . 2002, *A&A*, 390, 81
- Anderson, G. E., et al. 2014a, *GCN*, 16725, 1
- . 2014b, *MNRAS*, 440, 2059
- Anderson, W. G., et al. 2001, *Phys. Rev. D*, 63, 042003
- Atteia, J.-L. 2003, *A&A*, 407, L1
- Atwood, W. B., et al. 2009, *ApJ*, 697, 1071
- Babak, S., et al. 2013, *Phys. Rev. D*, 87, 024033
- Balasubramanian, R., Sathyaprakash, B. S., & Dhurandhar, S. V. 1996, *Phys. Rev. D*, 53, 3033
- Band, D. L., & Preece, R. D. 2005, *ApJ*, 627, 319
- Barankin, E. 1949, *Annals Math. Stat.*, 477
- Barnes, J., & Kasen, D. 2013a, *ApJ*, 775, 18
- . 2013b, *ApJ*, 775, 18
- Barsotti, L., & Fritschel, P. 2012, *LIGO-T1200307-v4*
- Barthelmy, S. D., et al. 2005, *Space Sci. Rev.*, 120, 143
- Belczynski, K., Bulik, T., & Bailyn, C. 2011, *ApJ*, 742, L2

- Belczynski, K., et al. 2013, *ApJ*, 764, 96
- Belczynski, K., Kalogera, V., & Bulik, T. 2002, *ApJ*, 572, 407
- Bellm, E. C. 2014, arXiv:1410.8185
- Berry, C. P. L., et al. 2014, arXiv:1411.6934
- Bertin, E. 2006, in ASP Conf. Series, Vol. 351, *Astronomical Data Analysis Software and Systems XV*, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 112
- Bertin, E., et al. 2002, in ASP Conf. Series, Vol. 281, *Astronomical Data Analysis Software and Systems XI*, ed. D. A. Bohlender, D. Durand, & T. H. Handley, 228
- Bhalerao, V., & Sahu, D. K. 2014, *GCN*, 16244, 1
- Bhalerao, V., et al. 2014, *GCN*, 16442, 1
- Blackman, J., et al. 2014, *Phys. Rev. Lett.*, 113, 021101
- Bloom, J. S., Frail, D. A., & Sari, R. 2001, *AJ*, 121, 2879
- Bloom, J. S., et al. 2012, *PASP*, 124, 1175
- Bock, D. C.-J., et al. 2006, in SPIE Conf. Series, Vol. 6267, *SPIE Conf. Series*, 13
- Bork, R., et al. 2001, in Proc. 8th Int. Conf. on Accelerator and Large Experimental Physics Control Systems, ed. H. Shoaee, Vol. C011127 (Menlo Park, CA: Stanford Linear Accelerator Center Technical Publications), 19–23
- Brink, H., et al. 2013, *MNRAS*, 435, 1047
- Bromberg, O., Nakar, E., & Piran, T. 2011, *ApJ*, 739, L55
- Brown, D. A., et al. 2012, *Phys. Rev. D*, 86, 084017
- Bue, B. D., et al. 2014, in Proceedings of the 2014 conference on Big Data from Space (BiDS'14)
- Bufano, F., et al. 2011, *Astronomische Nachrichten*, 332, 262
- Buonanno, A., & Chen, Y. 2001, *Phys. Rev. D*, 64, 042006

- Buonanno, A., Chen, Y., & Vallisneri, M. 2003, *Phys. Rev. D*, 67, 104025
- . 2006, *Phys. Rev. D*, 74, 029904
- Buonanno, A., et al. 2009, *Phys. Rev. D*, 80, 084043
- Burgay, M., et al. 2003, *Nature*, 426, 531
- Burns, E. 2014, *GCN*, 16363, 1
- Burrows, D. N., et al. 2005, *Space Sci. Rev.*, 120, 165
- Buskulic, D., the LIGO Scientific Collaboration, & the Virgo Collaboration. 2010, *Class. Quantum Grav.*, 27, 194013
- Butler, N., et al. 2013a, *GCN*, 14993, 1
- . 2013b, *GCN*, 14980, 1
- . 2014a, *GCN*, 16246, 1
- . 2014b, *GCN*, 16236, 1
- Butler, N. R., Bloom, J. S., & Poznanski, D. 2010, *ApJ*, 711, 495
- Butler, N. R., Kocevski, D., & Bloom, J. S. 2009, *ApJ*, 694, 76
- Butler, N. R., et al. 2007, *ApJ*, 671, 656
- Cabrera, J. I., et al. 2007, *MNRAS*, 382, 342
- Calabretta, M. R., & Greisen, E. W. 2002, *A&A*, 395, 1077
- Cameron, E. 2011, *PASA*, 28, 128
- Camp, J., et al. 2013, *Experimental Astronomy*, 36, 505
- Campana, S., et al. 2006, *Nature*, 442, 1008
- Canizares, P., et al. 2013, *Phys. Rev. D*, 87, 124005
- Cannon, K., et al. 2012, *ApJ*, 748, 136

- . 2010, *Phys. Rev. D*, 82, 44025
- Cannon, K., Hanna, C., & Keppel, D. 2013, *Phys. Rev. D*, 88, 024025
- Cannon, K., et al. 2011, *Phys. Rev. D*, 83, 084053
- Cano, Z., et al. 2014, *A&A*, 568, A19
- Cenko, S. B., et al. 2013a, *GCN*, 14998, 1
- Cenko, S. B., et al. 2014, *GCN*, 15883, 1
- Cenko, S. B., Ofek, E. O., & Nugent, P. E. 2012, *GCN*, 13489, 1
- Cenko, S. B., et al. 2006, *PASP*, 118, 1396
- . 2009, *ApJ*, 693, 1484
- . 2013b, *ApJ*, 769, 130
- . 2014, *GCN*, 15883, 1
- Chandra, P., & Frail, D. A. 2012, *ApJ*, 746, 156
- Cheung, T., et al. 2013, *GCN*, 14971, 1
- Chevalier, R. A., & Li, Z.-Y. 1999, *ApJ*, 520, L29
- Chornock, R., et al. 2010, *arXiv:1004.2262*
- Cokelaer, T. 2007, *Phys. Rev. D*, 76, 102004
- Collazzi, A. C., & Connaughton, V. 2013, *GCN*, 14972, 1
- Collazzi, A. C., et al. 2012, *ApJ*, 747, 39
- Connaughton, V., et al. 2014, *arXiv:1411.2685*
- Corder, S. A., Wright, M. C. H., & Carpenter, J. M. 2010, in *SPIE Conf. Series*, Vol. 7733, *SPIE Conf. Series*, 3
- Corsi, A., & Horesh, A. 2014, *GCN*, 16694, 1

- Corsi, A., et al. 2012, *ApJ*, 747, L5
- Costa, E., et al. 1997, *Nature*, 387, 783
- Cucchiara, A. 2014, *GCN*, 15652, 1
- Cutri, R. M., & et al. 2014, *VizieR Online Data Catalog*, 2328, 0
- Dai, Z. G., & Cheng, K. S. 2001, *ApJ*, 558, L109
- Dai, Z. G., Liang, E. W., & Xu, D. 2004, *ApJ*, 612, L101
- Dalal, N., et al. 2006, *Phys. Rev. D*, 74, 063006
- D’Avanzo, P., et al. 2013, *GCN*, 14973, 1
- De Pasquale, M. 2014a, *GCN*, 16455, 1
- . 2014b, *GCN*, 16428, 1
- Del Pozzo, W., et al. 2013, *Phys. Rev. Lett.*, 111, 071101
- D’Elia, V., & Izzo, L. 2014, *GCN*, 16464, 1
- D’Elia, V., Marshall, F. E., & Malesani, D. 2014, *GCN*, 16451, 1
- D’Elia, V., et al. 2013, *GCN*, 15000, 1
- DLMF. 2014, NIST Digital Library of Mathematical Functions, <http://dlmf.nist.gov/>, Release 1.0.8 of 2014-04-25, online companion to [Olver et al. \(2010\)](#)
- Drake, A. J., et al. 2009, *ApJ*, 696, 870
- Dressler, A., et al. 2011, *PASP*, 123, 288
- East, W. E., et al. 2013, *Phys. Rev. D*, 87, 043004
- Eichler, D., et al. 1989, *Nature*, 340, 126
- Evans, P. A., et al. 2007, *A&A*, 469, 379
- . 2009, *MNRAS*, 397, 1177

- . 2012, *ApJS*, 203, 28
- Fairhurst, S. 2009, *New J. Phys.*, 11, 123006
- . 2011, *Class. Quantum Grav.*, 28, 105021
- . 2014, *J. Phys. Conf. Series*, 484, 012007
- Farr, B., & Kalogera, V. 2013, in *APS Meeting Abstracts*, 10003
- Farr, B., Kalogera, V., & Luijten, E. 2014, *Phys. Rev. D*, 90, 024014
- Field, S. E., et al. 2011, *Phys. Rev. Lett.*, 106, 221102
- Firmani, C., et al. 2009, *MNRAS*, 393, 1209
- Fitzpatrick, G., & Connaughton, V. 2014, *GCN*, 16426, 1
- Flesch, E. 2010, *PASA*, 27, 283
- Fong, W., et al. 2012, *ApJ*, 756, 189
- Frail, D. A., et al. 1997, *Nature*, 389, 261
- Frater, R. H., Brooks, J. W., & Whiteoak, J. B. 1992, *Journal of Electrical and Electronics Engineering Australia*, 12, 103
- Freedman, D. L., & Waxman, E. 2001, *ApJ*, 547, 922
- Friedman, A. S., & Bloom, J. S. 2005, *ApJ*, 627, 1
- Fujiwara, T., et al. 2014, *GCN*, 16259, 1
- Gabriel, K. R., & Zamir, S. 1979, *Technometrics*, 21, 489
- Gal-Yam, A., et al. 2011, *ApJ*, 736, 159
- . 2014, *Nature*, 509, 471
- Galama, T. J., et al. 1998, *Nature*, 395, 670
- Gardner, W. G. 1995, *J. Audio Eng. Soc*, 43, 127

- Gehrels, N., et al. 2004, *ApJ*, 611, 1005
- Ghirlanda, G., Ghisellini, G., & Firmani, C. 2005, *MNRAS*, 361, L10
- . 2006, *New J. Phys.*, 8, 123
- Goad, M. R., et al. 2007, *A&A*, 476, 1401
- Goldstein, A., et al. 2012, *ApJS*, 199, 19
- Gorosabel, J., et al. 2014a, *GCN*, 16671, 1
- . 2014b, *GCN*, 16227, 1
- Górski, K. M., et al. 2005, *ApJ*, 622, 759
- Graff, P., et al. 2012, *MNRAS*, 421, 169
- . 2014, *MNRAS*, 441, 1741
- Greisen, E. W., & Calabretta, M. R. 2002, *A&A*, 395, 1061
- Greisen, E. W., et al. 2006, *A&A*, 446, 747
- Grover, K., et al. 2014, *Phys. Rev. D*, 89, 042004
- Hancock, P., et al. 2013, *GCN*, 15395, 1
- Hanna, C. 2008, PhD thesis, Louisiana State University
- Hannam, M., et al. 2014, *Phys. Rev. Lett.*, 113, 151101
- Harry, G. M. 2010, *Class. Quantum Grav.*, 27, 084006
- Harry, G. M., & the LIGO Scientific Collaboration. 2010, *Class. Quantum Grav.*, 27, 084006
- Harry, I. W., Allen, B., & Sathyaprakash, B. S. 2009, *Phys. Rev. D*, 80, 104014
- Harry, I. W., & Fairhurst, S. 2011a, *Class. Quantum Grav.*, 28, 134008
- . 2011b, *Phys. Rev. D*, 83, 084002
- Harry, I. W., et al. 2014, *Phys. Rev. D*, 89, 024010

- Hessels, J. W. T., et al. 2006, *Science*, 311, 1901
- Heussaff, V., Atteia, J.-L., & Zolnierowski, Y. 2013, *A&A*, 557, A100
- Hinderer, T., et al. 2010, *Phys. Rev. D*, 81, 123016
- Hinshaw, G., et al. 2013, *ApJS*, 208, 19
- Hjorth, J., & Bloom, J. S. 2012, in *Cambridge Astrophysics Series*, Vol. 51, *Gamma-Ray Bursts*, ed. C. Kouveliotou, R. A. M. J. Wijers, & S. Woosley (Cambridge: Cambridge University Press), 169–190
- Holland, S. T., & Mangano, V. 2014, *GCN*, 15673, 1
- Holz, D. E., & Hughes, S. A. 2005, *ApJ*, 629, 15
- Hooper, S., et al. 2010, in *AIP Conf. Proc.*, Vol. 1246, *Frontiers of Fundamental and Computational Physics: 10th Intl. Symp.*, ed. J. G. Hartnett & P. C. Abbott (Melville, NY: AIP), 211–214
- Horesh, A., et al. 2014, *GCN*, 16266, 1
- Howell, D. A., et al. 2005, *ApJ*, 634, 1190
- Hughey, B. 2011, in *Gravitational Wave Physics and Astronomy Workshop (GWPAW)*
- Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, 195, L51
- Hurley, K., et al. 2013, *GCN*, 14974, 1
- . 2014a, *GCN*, 15888, 1
- . 2014b, *GCN*, 16225, 1
- . 2014c, *GCN*, 16369, 1
- Ivezic, Z., et al. 2008, *arXiv:0805.2366*
- Jackson, J. 2003, *A User's Guide to Principal Components*, *Wiley Series in Probability and Mathematical Statistics* (New York: Wiley-Interscience)
- Jenke, P. 2013, *GCN*, 15331, 1

- Jenke, P., & Xiong, S. 2014, GCN, 15644, 1
- Johnson, S., & Frigo, M. 2007, IEEE Trans. Signal Process., 55, 111
- Jovanovic-Dolecek, G. 2002, Multirate Systems: Design and Applications (Hershey, PA: Idea Group Publ.)
- Kaiser, N., et al. 2010, in SPIE Conf. Series, Vol. 7733, SPIE Conf. Series, 0
- Kalogera, V., et al. 2004, ApJ, 601, L179, erratum in 614, L137
- Kann, D. A., et al. 2010, ApJ, 720, 1513
- . 2011, ApJ, 734, 96
- Kanner, J., et al. 2008, Class. Quantum Grav., 25, 184034
- Kasliwal, M. M. 2011, PhD thesis, California Institute of Technology
- Kasliwal, M. M., Cenko, S. B., & Singer, L. P. 2014, GCN, 16425, 1
- Kasliwal, M. M., & Nissanke, S. 2014, ApJ, 789, L5
- Kasliwal, M. M., Singer, L. P., & Cenko, S. B. 2013, GCN, 15324, 1
- Kasliwal, M. M., et al. 2012, ApJ, 755, 161
- Kelemen, J. 2014, GCN, 16440, 1
- Kelly, P. L., et al. 2013, ApJ, 775, L5
- Keppel, D. 2012, Phys. Rev. D, 86, 123010
- . 2013, arXiv:1307.4158
- Kimble, H. J., et al. 2002, Phys. Rev. D, 65, 022002
- Klimenko, S., et al. 2011, Phys. Rev. D, 83, 102001
- Klose, S., et al. 2013, GCN, 15320, 1
- Kocevski, D. 2012, ApJ, 747, 146

- Kopparapu, R. K., et al. 2008, *ApJ*, 675, 1459
- Krimm, H. A., et al. 2009, *ApJ*, 704, 1405
- Kulkarni, S. R. 2012, in *IAU Symposium*, Vol. 285, *IAU Symposium*, ed. E. Griffin, R. Hanisch, & R. Seaman, 55–61
- Kulkarni, S. R. 2013, *ATEL*, 4807, 1
- Kulkarni, S. R., et al. 1998, *Nature*, 395, 663
- Laskar, T., et al. 2013, *ApJ*, 776, 119
- Law, N. M., et al. 2009, *PASP*, 121, 1395
- Leloudas, G., et al. 2013, *GCN*, 14983, 1
- Li, L.-X., & Paczyński, B. 1998, *ApJ*, 507, L59
- Li, W., et al. 2011, *MNRAS*, 412, 1441
- Li, X., & Hjorth, J. 2014, *arXiv:1407.3506*
- Ligo Scientific Collaboration et al. 2011, *Nature Physics*, 7, 962
- Lloyd-Ronning, N. M., & Zhang, B. 2004, *ApJ*, 613, 477
- Luan, J., et al. 2012, *Phys. Rev. D*, 85, 102002
- Lynch, R., et al. 2014, *arXiv:1410.8503*
- Maggiore, M. 2008, *Gravitational Waves: Volume 1: Theory and Experiments*, *Gravitational Waves* (OUP Oxford)
- Malesani, D., et al. 2014, *GCN*, 16229, 1
- . 2013, *GCN*, 15642, 1
- Manca, G. M., & Vallisneri, M. 2010, *Phys. Rev. D*, 81, 024004
- Mangano, V. 2014, *GCN*, 16412, 1
- Mangano, V., & Burrows, D. N. 2014, *GCN*, 16373, 1

- Mangano, V., Evans, P. A., & Goad, M. R. 2014a, GCN, 16366, 1
- Mangano, V., Page, K., & Malesani, D. 2014b, GCN, 15648, 1
- Marion, F., & the Virgo Collaboration. 2003, in Proc. Rencontres de Moriond 2003, ed. J. Dumarchez, Les Arcs
- Marshall, F. E., & Amarel-Rogers, A. 2014, GCN, 16243, 1
- Masi, G. 2014, GCN, 16235, 1
- McAulay, R., & Hofstetter, E. 1971, IEEE Trans. Inf. Theory, 17, 669
- McAulay, R., & Seidman, L. 1969, IEEE Trans. Inf. Theory, 15, 273
- Meegan, C., et al. 2009a, ApJ, 702, 791
- . 2009b, ApJ, 702, 791
- Meegan, C. A., et al. 1992, Nature, 355, 143
- Meers, B. J., Krolak, A., & Lobo, J. A. 1993, Phys. Rev. D, 47, 2184
- Melandri, A., et al. 2012, A&A, 547, A82
- Messenger, C., & Read, J. 2012, Physical Review Letters, 108, 091101
- Metzger, B. D., et al. 2015, MNRAS, 446, 1115
- Metzger, B. D., & Berger, E. 2012, ApJ, 746, 48
- Metzger, B. D., & Fernández, R. 2014, MNRAS, 441, 3444
- Miller, M. C., & Miller, J. M. 2014, arXiv:1408.4145
- Mirabal, N., Halpern, J. P., & O'Brien, P. T. 2007, ApJ, 661, L127
- Modjaz, M., et al. 2006, ApJ, 645, L21
- Moskvitin, A. S., et al. 2014, GCN, 16228, 1
- Mulchaey, J., et al. 2013, GCN, 14985, 1

- Nakar, E., & Piran, T. 2005, *MNRAS*, 360, L73
- . 2011, *Nature*, 478, 82
- Nakar, E., & Sari, R. 2012, *ApJ*, 747, 88
- Narayan, R., Paczynski, B., & Piran, T. 1992, *ApJ*, 395, L83
- Nissanke, S., et al. 2010, *ApJ*, 725, 496
- Nissanke, S., Kasliwal, M., & Georgieva, A. 2013, *ApJ*, 767, 124
- Nissanke, S., et al. 2011, *ApJ*, 739, 99
- Nitz, A. H., et al. 2013, *Phys. Rev. D*, 88, 124039
- Nugent, P. E., et al. 2011, *Nature*, 480, 344
- Nuttall, L. K., & Sutton, P. J. 2010, *Phys. Rev. D*, 82, 102002
- Nysewander, M., Fruchter, A. S., & Pe'er, A. 2009, *ApJ*, 701, 824
- Oates, S. R., & Cenko, S. B. 2014, *GCN*, 16672, 1
- Ochsenbein, F., Bauer, P., & Marcout, J. 2000, *A&AS*, 143, 23
- Ofek, E. O., et al. 2012, *PASP*, 124, 62
- . 2013, *Nature*, 494, 65
- Oke, J. B., & Gunn, J. E. 1983, *ApJ*, 266, 713
- Olver, F. W. J., et al., eds. 2010, *NIST Handbook of Mathematical Functions* (New York, NY: Cambridge University Press), print companion to [DLMF \(2014\)](#)
- Oppenheim, A., Willsky, A., & Nawab, S. 1997, *Signals and systems*, Prentice Hall Signal Processing Series (Upper Saddle River, NJ: Prentice Hall)
- Owen, B. J., & Sathyaprakash, B. S. 1999, *Phys. Rev. D*, 60, 022002
- Özel, F., et al. 2012, *ApJ*, 757, 55
- Paciesas, W. S., et al. 2012a, *ApJS*, 199, 18

- . 2012b, *ApJS*, 199, 18
- Paczynski, B. 1986, *ApJ*, 308, L43
- Page, K. L. 2013, *GCN*, 15329, 1
- Page, K. L., & Cenko, S. B. 2014, *GCN*, 16682, 1
- Page, K. L., Evans, P. A., & Cenko, S. B. 2014, *GCN*, 16670, 1
- Pan, Y., et al. 2004, *Phys. Rev. D*, 69, 104017
- Panaitescu, A. 2005, *MNRAS*, 362, 921
- Pejcha, O., Thompson, T. A., & Kochanek, C. S. 2012, *MNRAS*, 424, 1570
- Penn, S., et al. 2007, LIGO-G070553-00-Z
- Perley, D. A. 2014, *GCN*, 15680, 1
- Perley, D. A., et al. 2014a, *GCN*, 16365, 1
- Perley, D. A., Cenko, S. B., & Kasliwal, M. M. 2013, *GCN*, 15327, 1
- Perley, D. A., et al. 2014b, *GCN*, 16454, 1
- Perley, D. A., & Singer, L. 2014, *GCN*, 16362, 1
- Perley, D. A., et al. 2014c, *ApJ*, 781, 37
- Pescalli, A., et al. 2014, *ArXiv e-prints*
- Peters, P. C. 1964, *Phys. Rev.*, 136, B1224
- Phillips, M. M. 1993, *ApJ*, 413, L105
- Pian, E., et al. 2000, *ApJ*, 536, 778
- . 2006, *Nature*, 442, 1011
- Piran, T., Nakar, E., & Rosswog, S. 2013, *MNRAS*, 430, 2121
- Pozanenko, A., et al. 2013, *GCN*, 14996, 1

- Press, W. H., et al. 2007a, *Numerical Recipes*, 3rd edn. (Cambridge Univ. Press), 81–82
- . 2007b, *Numerical Recipes*, 3rd edn. (Cambridge: Cambridge Univ. Press)
- Privitera, S., et al. 2014, *Phys. Rev. D*, 89, 024003
- Quimby, R. M., et al. 2011, *Nature*, 474, 487
- Racusin, J. L., et al. 2011, *ApJ*, 738, 138
- Rahmer, G., et al. 2008, in *SPIE Conf. Series*, Vol. 7014, *SPIE Conf. Series*, 4
- Rau, A., Kruehler, T., & Greiner, J. 2013, *GCN*, 15330, 1
- Rau, A., et al. 2009, *PASP*, 121, 1334
- Raymond, V., et al. 2009, *Class. Quantum Grav.*, 26, 114007
- Read, J. S., et al. 2009, *Phys. Rev. D*, 79, 124033
- Rees, M. J., & Meszaros, P. 1994, *ApJ*, 430, L93
- Rezzolla, L., et al. 2011, *ApJ*, 732, L6
- Rhoads, J. E. 1997, *ApJ*, 487, L1
- Richardson, D., et al. 2002, *AJ*, 123, 745
- Robitaille, T. P., et al. 2013, *A&A*, 558, A33
- Rodriguez, C. L., et al. 2014, *ApJ*, 784, 119
- Roming, P. W. A., et al. 2005, *Space Sci. Rev.*, 120, 95
- Sahu, D. K., Bhalerao, V., & Anupama, G. C. 2014, *GCN*, 16684, 1
- Sakamoto, T., et al. 2006, *ApJ*, 636, L73
- Sari, R., Piran, T., & Narayan, R. 1998, *ApJ*, 497, L17
- Sathyaprakash, B., & Schutz, B. 2009, *Living Rev. Relativity* 12, arXiv:0903.0338v1
- Savaglio, S., Glazebrook, K., & Le Borgne, D. 2009, *ApJ*, 691, 182

- Schaefer, B. E., & Collazzi, A. C. 2007, *ApJ*, 656, L53
- Schlaflly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, 737, 103
- Schulze, S., et al. 2013, *GCN*, 14994, 1
- . 2014a, *A&A*, 566, A102
- . 2014b, *A&A*, 566, A102
- Schutz, B. F. 1986, *Nature*, 323, 310
- . 2011, *Class. Quantum Grav.*, 28, 125023
- Shahmoradi, A., & Nemiroff, R. J. 2011, *MNRAS*, 411, 1843
- Shoemaker, D. 2010, *LIGO-T0900288-v3*
- Sidery, T., et al. 2014, *Phys. Rev. D*, 89, 084060
- Siegel, M. H., & De Pasquale, M. 2014, *GCN*, 16432, 1
- Simakov, D. A. 2014, *Phys. Rev. D*, 90, 102003
- Singer, L. P., Cenko, S. B., & Kasliwal, M. M. 2013, *GCN*, 15643, 1
- Singer, L. P., et al. 2014a, *GCN*, 16226, 1
- . 2014b, *GCN*, 16668, 1
- Singer, L. P., Kasliwal, M. M., & Cenko, S. B. 2014c, *GCN*, 16360, 1
- Singer, L. P., et al. 2013, *ApJ*, 776, L34
- Singer, L. P., et al. 2013, *GCN*, 14967, 1
- . 2014, *ApJ*, 795, 105
- Skrutskie, M. F., et al. 2006, *AJ*, 131, 1163
- Smith, R. J. E., et al. 2014a, *Phys. Rev. D*, 90, 044074
- Smith, R. M., et al. 2014b, in *SPIE Conf. Series*, Vol. 9147, *SPIE Conf. Series*, 79

- Soderberg, A. M., et al. 2004, *Nature*, 430, 648
- . 2010, *Nature*, 463, 513
- Sollerman, J., et al. 2006, *A&A*, 454, 503
- Sonbas, E., et al. 2013, *GCN*, 15640, 1
- Strain, K., & Cagnoli, G. 2006, LIGO-T050267-01-K
- Sudilovsky, V., Tanga, M., & Greiner, J. 2013, *GCN*, 15328, 1
- Svensson, K. M., et al. 2010, *MNRAS*, 405, 57
- Taboada, I., & Gilmore, R. C. 2014, *Nuclear Instruments and Methods in Physics Research A*, 742, 276
- Taracchini, A., et al. 2014, *Phys. Rev. D*, 89, 061502
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, 253, 908
- Turin, G. 1960, *IRE Trans. Inf. Theory*, 6, 311
- Urata, Y., et al. 2012, *ApJ*, 748, L4
- . 2014, *ApJ*, 789, 146
- van den Broeck, C., et al. 2009, *Phys. Rev. D*, 80, 024009
- van der Sluys, M., et al. 2008a, *Class. Quantum Grav.*, 25, 184011
- van der Sluys, M. V., et al. 2008b, *ApJ*, 688, L61
- van Eerten, H. J., & MacFadyen, A. I. 2011, *ApJ*, 733, L37
- van Paradijs, J., et al. 1997, *Nature*, 386, 686
- Veitch, J., & Vecchio, A. 2010, *Phys. Rev. D*, 81, 062003
- Veitch, J., et al. 2012, *Phys. Rev. D*, 85, 104045
- Véron-Cetty, M.-P., & Véron, P. 2010, *A&A*, 518, A10

- Vitale, S., & Zanolin, M. 2011, *Phys. Rev. D*, 84, 104020
- Vogt, R. E. 1989, A Laser Interferometer Gravitational-Wave Observatory, Tech. Rep. PHY-8803557, Proposal to the National Science Foundation
- Volnova, A., et al. 2013, *GCN*, 15341, 1
- . 2014a, *GCN*, 16260, 1
- . 2014b, *GCN*, 16453, 1
- von Kienlin, A. 2014, *GCN*, 16450, 1
- Wen, L., & Chen, Y. 2010, *Phys. Rev. D*, 81, 082001
- White, D. J., Daw, E. J., & Dhillon, V. S. 2011, *Class. Quantum Grav.*, 28, 085016
- Wiersema, K., et al. 2014, *GCN*, 16231, 1
- Will, C. M. 2006, *Living Rev. Rel.*, 9
- Wright, E. L., et al. 2010, *AJ*, 140, 1868
- Xu, D., et al. 2013, *GCN*, 15325, 1
- Xu, D., et al. 2013, *GCN*, 15641, 1
- Xu, D., Zhang, C.-M., & Cao, C. 2013a, *GCN*, 15326, 1
- Xu, D., et al. 2013b, *ApJ*, 776, 98
- Yaron, O., & Gal-Yam, A. 2012, *PASP*, 124, 668
- Yu, H.-F., & Goldstein, A. 2014, *GCN*, 16224, 1
- Zhang, B. 2014, *GCN*, 16669, 1
- Zhang, B.-B., et al. 2012, *ApJ*, 756, 190
- Zwart, J. T. L., et al. 2008, *MNRAS*, 391, 1545