

**School of Physics
and Astronomy**



Gravitational Waves with Gamma-ray Bursts

Andrew Robert Williamson

Submitted for the degree of Doctor of Philosophy
School of Physics and Astronomy
Cardiff University

12 October 2016



Andrew R. Williamson, 2016

Except where otherwise noted, this work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Fig. 2.1 reprinted from *Astroparticle Physics*, Vol. 43, Meszaros, “Gamma Ray Bursts” [126], copyright (2013), with permission from Elsevier.

Fig. 2.3 reprinted by permission from Macmillan Publishers Ltd: *Nature*, Vol. 500, Tanvir et al. [165], copyright (2013).

xkcd webcomic ‘Gravitational Waves’ <https://xkcd.com/1642/> (p. xiii), copyright Randall Munroe. Licensed under the Creative Commons Attribution-NonCommercial 2.5 Generic License, <https://creativecommons.org/licenses/by-nc/2.5/>.

This thesis was typeset using the $\text{\LaTeX} 2_{\epsilon}$ release of the \LaTeX typesetting system originally developed by Leslie Lamport, based on \TeX created by Donald Knuth.

The style template used was adapted from an original written by by Duncan M. Macleod.

SUMMARY OF THESIS

Gravitational waves have now twice been detected emanating from the merging of binary black hole systems. In this thesis we detail the methods used to search for binary merger gravitational wave signals associated with short gamma-ray bursts, focusing on systems that include at least one neutron star.

We first cover the background theory behind gravitational wave emission, the means of detection via interferometry, and the types of astrophysical sources that could be detected now or in the near future. We follow this with a review of gamma-ray burst theory and observations, focusing in particular those bursts with short durations. These are likely to be caused by the mergers of binaries that include a neutron star and a black hole, or two neutron stars — events of great interest to gravitational wave astronomy.

We then discuss the methods used to search gravitational wave data in a targeted way, using the prior observation of a short gamma-ray bursts to focus the analysis and improve the chances of making a detection. We also summarise early searches of this kind and present the results of a search carried out on LIGO and Virgo data spanning 2005–2010, targeting short gamma-ray bursts detected by the InterPlanetary Network.

We then turn our attention to the current, second generation of gravitational wave detectors. We present a detailed calculation of the prospects of success for the targeted short gamma-ray burst search technique, and find that we might reasonably expect to make up to a few detections per year around the turn of the decade.

We then outline a new search structure for use during the second generation of detectors, and an astrophysical event alert system for the control rooms of gravitational wave observatories.

We end with a presentation of the results of the new and improved search carried out during the first observing run of Advanced LIGO.

CONTENTS

Contents	iv
List of Figures	vii
List of Tables	ix
Acknowledgements	x
Co-authored Works	xi
Conventions	xii
Introduction	1
I Gravitational Waves & Gamma-ray Bursts	5
1 Gravitational Waves	6
1.1 Gravitational Waves in Linearised Theory	6
1.1.1 Wave Solutions for Metric Perturbations	7
1.1.2 Choosing the Transverse–Traceless Gauge	9
1.1.3 Gravitational Waves Far from Their Sources	10
1.2 Gravitational Wave Detectors	11
1.2.1 The Tidal Effect of Gravitational Waves	11
1.2.2 Gravitational Wave Interferometry	11
1.2.3 Gravitational Wave Observatories	14
1.3 Astrophysical Sources of Gravitational Waves	17
1.3.1 Compact Binary Coalescences	18
1.3.2 Gravitational Wave Bursts	20
2 Gamma-ray Bursts	22
2.1 The Current Gamma-ray Burst Paradigm	22
2.1.1 Long Gamma-ray Bursts	25
2.1.2 Short Gamma-ray Bursts	25
2.2 Gamma-ray Burst Satellites and Observations	27
2.3 Short Gamma-ray Bursts and Gravitational Waves	29

II	Searches with First Generation Detectors	30
3	First Searches with LIGO and Virgo	31
3.1	Pipeline Summary	31
3.1.1	Multi-Detector Matched Filter	31
3.1.2	Signal Consistency	34
3.1.3	Event Significance	37
3.2	Early Results	38
3.2.1	Full Science Run Analyses	38
3.2.2	GRB 051103	38
3.2.3	GRB 070201	39
3.3	Search for InterPlanetary Network Gamma-ray Bursts	39
3.3.1	GRB Sample	39
3.3.2	Results	41
3.4	Cumulative LIGO and Virgo Results	43
4	Improved Methods	46
4.1	Background Estimation	47
4.2	Restrictions on Source Inclination	50
4.3	Searching Large Areas of Sky	53
4.3.1	Two-site Time Delay Degeneracy	59
4.4	Benefits Provided by Targeted Search	61
III	Searching with Second Generation Detectors	63
5	Prospects for Joint Observations	64
5.1	The Short Gamma-ray Burst Rate	64
5.2	Expected Rate of Joint Observations	67
5.3	Benefits of Joint Observations	73
5.3.1	Detecting a GRB with measured redshift	73
5.3.2	Constraining the jet opening angle	75
6	PyGRB: The Short Gamma-ray Burst Search for Advanced Detectors	78
6.1	PyCBC	79
6.2	The PyGRB Workflow	80
6.2.1	Increasing the Number of Analysed GRBs with Single Detector Analyses	81
6.2.2	Selecting the Analysis Period Dynamically	82
6.2.3	Reducing the Background by Refining the Source Parameter Space	84
6.2.4	Increase Sensitivity with Arbitrary Source Spins and Precession	85
6.2.5	Targeting Host Galaxy Distances with Injections	87
6.2.6	Extension to Low Frequencies	89
6.3	Responding Online	92
6.3.1	The External Trigger Alert System	95
6.3.2	The Online Workflow	96

7	PyGRB in the First Observing Run of Advanced LIGO	100
7.1	GRB Sample	100
7.1.1	GRB 150906B	102
7.2	Gravitational Wave Detectors	103
7.3	Search Methodology	105
7.4	Result for GRB 150906B	108
7.5	Cumulative Results	110
7.6	Conclusion	110
	Conclusion	114
	Bibliography	117

LIST OF FIGURES

0.1	Thirty years of orbital decay in the PSR B1913+16 system	2
0.2	GW150914: The first direct detection of gravitational waves	3
1.1	Effect of a gravitational wave on an interferometer	12
1.2	Advanced LIGO optical layout	16
1.3	Advanced LIGO quadruple pendulum suspension system	17
1.4	Comparison between Enhanced LIGO and Advanced LIGO sensitivities	18
2.1	Schematic gamma-ray burst spectrum for the leptonic photospheric-internal shock model	24
2.2	Gamma-ray bursts arranged by duration and spectral hardness	24
2.3	Observation of a kilonova associated with the short gamma-ray burst 130603B	26
3.1	Cumulative distribution of p -values from the analysis of 27 short gamma-ray bursts observed by the InterPlanetary Network	41
3.2	90 % distance exclusions for 27 short gamma-ray bursts observed by the InterPlanetary Network	42
3.3	Cumulative distribution of p -values for all 69 short gamma-ray bursts from LIGO and Virgo runs between 2005–2010	43
3.4	90 % distance exclusions for all 69 short gamma-ray bursts from LIGO and Virgo runs between 2005–2010	45
4.1	Measuring to lower false alarm probabilities through the use of time slides	49
4.2	Variation of gravitational wave polarisation amplitudes with inclination	50
4.3	Suppression of background due to restricting source inclination	53
4.4	Example patch of search sky points projected onto the celestial sphere	56
4.5	Effect of sky error on search background	57
4.6	Effect of sky error on detection efficiency for a three detector search	58
4.7	Effect of sky error on detection efficiency for a two detector search	60
5.1	Rate of binary mergers in the local universe	66
5.2	Expected rate of jointly observed gravitational wave and electromagnetic signals from short gamma-ray bursts with the second generation gravitational wave detectors	70
5.3	Expected binary neutron star merger detection rate with second generation gravitational wave detectors	71
5.4	Effect of known distance on detectability of a short gamma-ray burst	73
5.5	Effect of known distance on the odds ratio for signals of varying strength	74

5.6	Expected upper limit on the binary neutron star merger rate obtained from planned observing runs, assuming no gravitational wave detections	76
6.1	Schematic of the standard, offline PyGRB workflow	81
6.2	Dynamic calculation of analysis segments in PyGRB	83
6.3	Effect of black hole mass and spin on the ability of a neutron star–black hole merger to power a short gamma-ray burst	85
6.4	EM-bright template bank	86
6.5	Modifying injection distances according to calibration uncertainty	90
6.6	Extending search down to 30Hz	91
6.7	Trigger rate dependence on the inverse spectrum length	93
6.8	Effect of the inverse spectrum length on spectral resolution	94
6.9	Control room screen for displaying external trigger information	97
6.10	Schematic of the online PyGRB workflow	99
7.1	Durations of the gamma-ray bursts observed during the first observing run of Advanced LIGO	101
7.2	Localisation of GRB 150906B by InterPlanetary Network satellites	102
7.3	Overlay of the error box for GRB 150906B on the sky	103
7.4	Location of GRB 150906B on the $T_{50} - \log \text{HR}_{32}$ diagram for <i>Konus-Wind</i> gamma-ray bursts.	104
7.5	Vetoing a loud glitch in Advanced LIGO data with PyGRB	106
7.6	PyGRB search background around the time of GRB 150906B	109
7.7	PyGRB injection recovery around the time of GRB 150906B as a function of distance	110
7.8	PyGRB injection recovery around the time of GRB 150906B as a function of source inclination	111
7.9	Cumulative distribution of p -values for the analysis of 16 short gamma-ray bursts observed during Advanced LIGO Observing Run 1	112
7.10	90 % distance exclusions for 16 short gamma-ray bursts observed during Advanced LIGO Observing Run 1	112

LIST OF TABLES

3.1	Short gamma-ray burst sample for the InterPlanetary Network search with LIGO and Virgo	40
5.1	Expected rate of joint gravitational wave–electromagnetic short gamma-ray burst observations, assuming binary neutron star progenitors	69
5.2	Expected rate of joint gravitational wave–electromagnetic short gamma-ray burst observations, assuming neutron star–black hole binary progenitors	71
5.3	Expected bounds on short gamma-ray burst opening angles during early second generation gravitational wave detector runs	76
7.1	Exclusion distances placed on binary merger progenitors for GRB 150906B	109
7.2	Short gamma-ray burst sample for the PyGRB search during Advanced LIGO Observing Run 1	113

ACKNOWLEDGEMENTS

I owe an enormous debt of gratitude to many, many people. First and foremost, I must thank my seemingly omniscient supervisor, Steve Fairhurst. I've learnt an awful lot from you about our field and what it means to be a good scientist. You've been supportive, patient, and an invaluable guide along the way. I hope we will continue to have the chance to work together for many years to come. Thanks, Boss.

Thank you to my viva examiners, John Veitch and Mark Hannam, for taking the time to read this thesis and provide comments and feedback that have significantly improved it. Your interrogation was very fair, even when it felt like my brain was beginning to leak from my ears. Thanks also to Ant Whitworth for chairing the viva, in particular for suggesting a stop for Welsh cakes just when the brain-leaking was becoming critical.

I would never have embarked on this adventure were it not for the encouragement and incredible supervision of Martin Hendry during my undergraduate days at Glasgow. Thank you for helping me to realise this was what I wanted to do.

There are so many other collaborators who have helped me along the way, particularly in the CBC, GRB, and DetChar working groups of the LIGO and Virgo collaborations. I especially want to thank Ian, Francesco, Erin, Laura, Duncan, Patrick, and James for passing on their wisdom.

My thanks also to the rest of the Cardiff gravy wavers I've had the pleasure of working with over these past four years; Sathya, Frank, Valeriu, Chris M., Paul, Michael, Patricia, Alex, Lionel, Gernot, Max, Vaibhav, Chris N., Ioannis, Sam, Ed, Iain, Chinmay, and my room buddy for so many conferences, Seb. It was a great pleasure to work with you all, to share in the excitement of the first detection, and to sample the many culinary delights of Ernest Willows. Also, of course, special thanks to my erstwhile housemate Tom Adams — yeah, buddy!

Thank you to Brian O'Reilly and others at the LIGO Lab for agreeing to let me visit the Livingston site. It was a fantastic experience and I hope to get the chance to visit again. To Adam, Arnaud, Marie (and DD!) — thanks for putting up with me in the Pool House. I'll always remember living with you guys with great fondness. Also to Duncan (again?!), Lucy, Jon, Marissa, Shivaraj, Thomas, Terra, Anamaria, Carl, Nutsinee — thanks all for making those four months so great.

Thank you to all the other friends who have made my time in Cardiff so enjoyable. There are too many of you to name, but I've had a fantastic time here. Diolch yn fawr.

Thank you to my family and friends for their love and support. Mum, Dad, and Lisa; I love you.

Finally, Olivia. You have been my closest friend and my biggest supporter. I have loved every moment of our life together in our little Cardiff home. Thank you for all your love, for sharing the last three years with me, for basically stealing Hathor for us, for always being there. Thank you for everything. I love you.

CO-AUTHORED WORKS

Some sections of this thesis include collaborative work, some of which has been published elsewhere.

- The search method summary in Section 3.1 and the studies in Chapter 4 were previously published in a paper, of which I was lead author: [181].
- The results presented in Sections 3.3 and 3.4 were previously published in a LIGO-Virgo collaborations paper [5]. I contributed to this work by testing the analysis, and by running the final analyses for a number of gamma-ray bursts.
- The work in Chapter 5 was previously published in a short author list paper, of which I am a co-author: [57].
- The software described in Section 6.3.1 was previously described in a LIGO technical document: [180].
- The pipeline described in Chapter 6 is a project that I have led, but which has had input from many collaborators, particularly within the GRB working group of the LIGO-Virgo collaborations.
- I have also led the effort towards obtaining the results presented in Chapter 7, but have again had important contributions input from the GRB working group of the LIGO-Virgo collaborations. This will be published as part of a LIGO-Virgo collaborations paper.

CONVENTIONS

We will follow a standard convention in this thesis with spatial vectors denoted by Latin indices, and spacetime vectors denoted by Greek indices. We will also follow the Einstein summation convention, where pairs of superscript and subscript indices imply summation,

$$A^\alpha B_\alpha \equiv \sum_{\alpha=0}^3 A^\alpha B_\alpha .$$

Partial derivatives of vectors are denoted by commas,

$$A^\alpha_{,\beta} \equiv \frac{\partial}{\partial x^\beta} A^\alpha ,$$

and covariant derivatives are denoted by semicolons, i.e.

$$A^\alpha_{;\beta} \equiv A^\alpha_{,\beta} + \Gamma^\alpha_{\beta\gamma} A^\gamma ,$$

where $\Gamma^\alpha_{\beta\gamma}$ are the Christoffel symbols,


$$\Gamma^\alpha_{\beta\gamma} = \frac{1}{2} g^{\alpha\delta} (g_{\delta\gamma,\beta} + g_{\delta\beta,\gamma} + g_{\beta\gamma,\delta}) ,$$

and where $g_{\alpha\beta}$ is the spacetime metric tensor. In general, contravariant vectors are given raised indices and transform according to

$$A'^\alpha = \frac{\partial x'^\alpha}{\partial x^\beta} A^\beta ,$$

and covariant vectors, or one-forms, are given lowered indices and transform according to

$$B'_\alpha = \frac{\partial x^\beta}{\partial x'^\alpha} B_\beta .$$



To Mum, Dad, and Lisa

INTRODUCTION

In 1915 Albert Einstein published his general theory of relativity [75], the culmination of a decade of effort to incorporate gravitation into his relativistic framework introduced in the special theory [73, 74]. In the general theory gravity is the effect of the geometry of spacetime, and the geometry of spacetime is the result of the presence of matter.

“ Spacetime tells matter how to move; matter tells spacetime how to curve

– J. A. Wheeler [179] ”

Mathematically this relationship is encompassed in the Einstein field equations,

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (0.1)$$

The left hand side describes the geometry of spacetime via the Ricci tensor $R_{\mu\nu}$, the Ricci scalar R , the spacetime metric $g_{\mu\nu}$, and the cosmological constant Λ . On the right hand side the matter content of spacetime is described by the stress-energy tensor $T_{\mu\nu}$. c is the speed of light in a vacuum and G is the gravitational constant.

In the century since its publication the general theory has been a remarkably successful predictive tool. From the observation of the lensing of starlight [71], through to the measurement of the geodetic effect by Gravity Probe B [78], every empirical test of the theory has so far found it to be in agreement with nature – at least to within experimental precision.

One of the last remaining predictions of the general theory to be experimentally verified were gravitational waves, first predicted by Einstein in 1916 [76, 77]. These perturbations in the gravitational field – often termed ‘ripples in spacetime’ – propagate at the speed of light. Their existence was first supported by observations of the pulsar PSR B1913+16 discovered by Hulse and Taylor [102]. The pulsar is one of a pair of neutron stars in a tight binary, the orbital decay of which is consistent with the emission of gravitational waves [167, 178] (Fig. 0.1).

Since the 1960s there has been an effort to detect these waves. The first serious attempts by Weber made use of a number of aluminium bars [175, 176], but his claims of detections [177] were later rejected as false positives [120].




Figure 0.1: Thirty years of orbital decay in the PSR B1913+16 system. The observed decay (red data points) agrees with that predicted by general relativity if the system is emitting gravitational waves (blue solid curve). Data from [178]. Figure from [187].

Since the 1980s a new approach has dominated the effort to detect gravitational waves: laser interferometry. A number of ever larger, ever more sensitive interferometers have been built, including the 4 kilometre scale Laser Interferometer Gravitational-Wave Observatory detectors, or LIGO for short [31]. Although without any detections, these instruments operated well and drove technological advancement that is now being deployed in a new, second generation of interferometers.

On 14 September 2015, at 09:50:45 UTC, just days after becoming fully operational following a seven year upgrade project, the two LIGO detectors measured a gravitational wave signal as it passed by the Earth (Fig. 0.2), marking the first ever direct detection of gravitational waves [22]. The waves were produced by the coalescence of a binary black hole system. The black holes that merged had masses of approximately 36 and 29 solar masses, and formed a spinning black hole of approximately 62 solar masses [23]. The signal included the last few orbits of their inspiral, the merger, and final black hole ringdown, and was completely consistent with Einstein’s general theory [24]. This detection was the first direct observation of black holes and their mergers, and is the first observation supporting the existence of such large stellar mass black holes, with ramifications for star formation and stellar evolution models [25]. With this detection, gravitational wave astronomy was born.

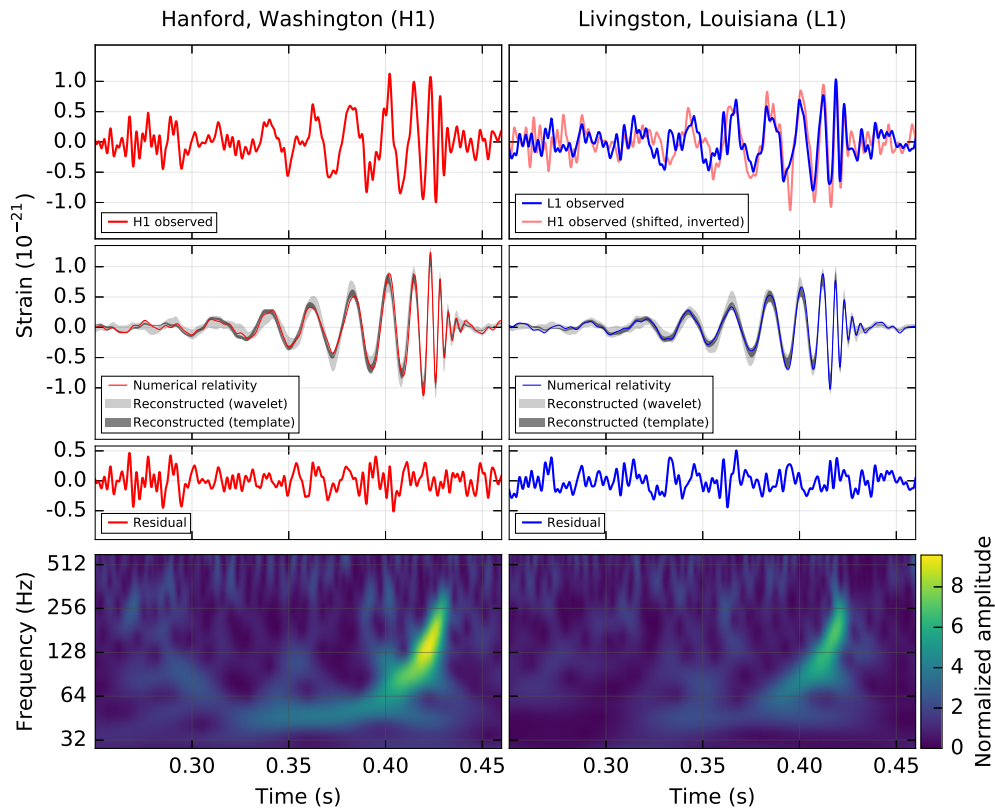


Figure 0.2: GW150914: The first direct detection of gravitational waves. The signal was detected by the two observatories of Advanced LIGO, and came from the merger of a pair of black holes around 1.3 billion light years away. Top row: the signal in the Hanford detector (left) and in the Livingston detector (right, blue), with the Hanford signal superimposed (red). Second row: the best fit numerical waveforms and reconstructed signals. Third row: Residuals from subtraction of best fit waveforms from the signals. Bottom row: Time-frequency plots of the signal in each detector, showing the significance of the signal above the background, and the upward chirp in frequency. Figure taken from [22].

During the rest of their first observing run, the two LIGO detectors saw another strong signal from a lower mass binary black hole merger [21], with another, weaker detection candidate, again seemingly from a binary black hole [17]. It is now apparent that binary black hole systems are relatively common, and will dominate the first era of gravitational wave detections.

Gravitational wave astronomy is an entirely new way for astronomers to investigate the cosmos, and one that compliments more established branches of astronomy. One such branch is high-energy transient astronomy. Many of the most energetic events in the universe not only produce copious amounts of electromagnetic emission, but are also likely to emit powerful gravitational waves. Of particular interest are short gamma-ray bursts, which may be caused by the mergers of two neutron stars or a black hole and a neutron

star. Both of these types of event would be expected to produce strong gravitational waves in the frequency band of current ground-based gravitational wave detectors, so may be among the first types of events observed via gravitational waves.

In this thesis we investigate short gamma-ray bursts as promising astrophysical events for gravitational wave follow-up, in the hope of making the first joint electromagnetic and gravitational wave astronomical observation.



**GRAVITATIONAL WAVES
& GAMMA-RAY BURSTS**

“The rumour of much toil and scheming and triumph
may never reach the stars, and what we value not at all,
are not conscious of, may break the surface of eternity
with endless ripples of good.”

Edward Thomas

1

GRAVITATIONAL WAVES

In this chapter we will summarise the theoretical and expected astrophysical origins of gravitational waves (GWs), and the motivation behind the construction of ground-based interferometers like the Laser Interferometer Gravitational Wave Observatory (LIGO).

1.1 Gravitational Waves in Linearised Theory

GWs are perturbations in the curvature of spacetime that propagate as waves at the speed of light. It may be useful to draw some analogies between electromagnetism and gravitation, but also identify where the two phenomena differ. For example, the electrostatic force acting between charges q_1 and q_2 that sit a distance d apart is proportional to $\frac{q_1 q_2}{d^2}$ (Coulomb's Law), and if accelerated the charges produce electromagnetic waves. By comparison, the gravitational force acting between masses m_1 and m_2 that sit a distance d apart is proportional to $\frac{m_1 m_2}{d^2}$ (Newton's Law). The analogy also suggests that accelerated masses might produce GWs. There are, however, important differences.

In electromagnetism there may be positive and negative charges whereas in gravitation there exists no 'negative' mass; no *anti-gravity*. This places additional constraints on the generation of GWs when compared to electromagnetic waves. Time variation of the electric dipole moment dominates the emission of electromagnetic waves, however a time-varying mass dipole moment cannot emit GWs or else the conservation of momentum would be violated. The analogue of a time-varying magnetic dipole moment, which is next dominant, cannot produce GWs either as this would violate the conservation of angular momentum.¹ We will see that the leading order of multipole that can produce

¹Monopole gravitational radiation would violate the conservation of energy, just as monopole electromagnetic radiation would violate charge conservation.

GWs is the time-varying gravitational quadrupole moment. See Chapters 35 & 36 of [127] for further discussion on this. Much of the derivations in this chapter follow those found in [124], [127], and [155].

1.1.1 Wave Solutions for Metric Perturbations

We begin by taking the linearised approximation of the Einstein field equations (Eq. (0.1)) in a weak gravitational field. We can re-write Eq. (0.1) by defining the Einstein tensor,

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}, \quad (1.1)$$

and noting that for non-cosmological considerations we can safely consider the cosmological constant Λ to be zero. Therefore, we obtain

$$G_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (1.2)$$

In a weak gravitational field the spacetime metric is almost flat, so that there exists a coordinate system where it may be described by a combination of the Minkowski metric of flat spacetime,

$$\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1), \quad (1.3)$$

plus a small perturbation, i.e.

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.4)$$

where $|h_{\mu\nu}| \ll 1$ for all μ and ν . In this linearised formulation, the Riemann-Christoffel tensor takes the form

$$R_{\mu\nu\alpha\beta} = \frac{1}{2}(h_{\mu\beta,\nu\alpha} + h_{\nu\alpha,\mu\beta} - h_{\mu\alpha,\nu\beta} - h_{\nu\beta,\mu\alpha}), \quad (1.5)$$

which we may contract to get the Ricci tensor,

$$R_{\mu\nu} = \frac{1}{2}(h_{\mu,\nu\alpha}^{\alpha} + h_{\nu,\mu\alpha}^{\alpha} - h_{\mu\nu,\alpha}^{\alpha} - h_{,\mu\nu}), \quad (1.6)$$

and the Ricci scalar, since

$$R = \eta^{\alpha\beta}R_{\alpha\beta}. \quad (1.7)$$

The linearised form of the Einstein tensor may therefore be expressed as

$$G_{\mu\nu} = \frac{1}{2} \left[h_{\mu\alpha,\nu}^{\alpha} + h_{\nu\alpha,\mu}^{\alpha} - h_{\mu\nu,\alpha}^{\alpha} - h_{,\mu\nu} - \eta_{\mu\nu} (h_{\alpha\beta}^{\alpha\beta} - h_{,\beta}^{\beta}) \right]. \quad (1.8)$$

Using the substitution

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h, \quad (1.9)$$

we may simplify Eq. (1.8) to obtain

$$G_{\mu\nu} = -\frac{1}{2} \left[\bar{h}_{\mu\nu,\alpha}{}^{,\alpha} + \eta_{\mu\nu} \bar{h}_{\alpha\beta}{}^{,\alpha\beta} - \bar{h}_{\mu\alpha,\nu}{}^{,\alpha} - \bar{h}_{\nu\alpha,\mu}{}^{,\alpha} \right]. \quad (1.10)$$

Entering this result into Eq. (1.2) we have derived the linearised form of the Einstein field equations,

$$-\bar{h}_{\mu\nu,\alpha}{}^{,\alpha} - \eta_{\mu\nu} \bar{h}_{\alpha\beta}{}^{,\alpha\beta} + \bar{h}_{\mu\alpha,\nu}{}^{,\alpha} + \bar{h}_{\nu\alpha,\mu}{}^{,\alpha} = \frac{16\pi G}{c^4} T_{\mu\nu}. \quad (1.11)$$

We can adopt the *Lorentz gauge* conditions, taking a coordinate system in which the divergence of the metric perturbations is zero, i.e.

$$\bar{h}^{\mu\alpha}{}_{,\alpha} = 0. \quad (1.12)$$

This sets the three rightmost terms in the left hand side of Eq. (1.11) to zero, so we obtain

$$-\bar{h}_{\mu\nu,\alpha}{}^{,\alpha} = \frac{16\pi G}{c^4} T_{\mu\nu}. \quad (1.13)$$

If we take the free space solutions of Eq. (1.13), i.e. where $T_{\mu\nu} = 0$, we obtain a wave equation,

$$\bar{h}_{\mu\nu,\alpha}{}^{,\alpha} = 0. \quad (1.14)$$

This can also be written as

$$\left(c^2 \nabla^2 - \frac{\partial^2}{\partial t^2} \right) \bar{h}_{\mu\nu} = 0. \quad (1.15)$$

This tells us that, to first order in nearly flat spacetime, perturbations to the spacetime metric propagate at the speed of light c as waves in free space; gravitational waves.

The simplest solutions to Eq. (1.15) are plane waves,

$$\bar{h}_{\mu\nu} = \text{Re} \left[A_{\mu\nu} \exp \{ i k_\alpha x^\alpha \} \right], \quad (1.16)$$

where $A_{\mu\nu}$ are the wave amplitude components and k_α is the wave vector. Since $\bar{h}_{\mu\nu}$ is symmetric, so is $A_{\mu\nu}$. This immediately reduces the number of independent components of $A_{\mu\nu}$ from 16 to 10. Additionally, from Eq. (1.14) we know that \mathbf{k} is a null vector,

$$k_\alpha k^\alpha = 0, \quad (1.17)$$

and from the Lorentz condition, Eq. (1.12), it follows that the components of $A_{\mu\nu}$ must be orthogonal to \mathbf{k} , i.e.

$$A_{\mu\alpha} k^\alpha = 0. \quad (1.18)$$

Eq. (1.18) constrains four of the components of $A_{\mu\nu}$, so we are left with six independent components.

1.1.2 Choosing the Transverse–Traceless Gauge

We may ask whether there are any more gauge choices that simplify this. Let us consider a gauge transformation where

$$x'^{\mu} = x^{\mu} + \xi^{\mu}. \quad (1.19)$$

The perturbation is now given by

$$h'_{\mu\nu} = h_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu}. \quad (1.20)$$

If this transformation is to maintain the Lorentz gauge conditions, the components ξ^{μ} must satisfy

$$\left(c^2 \nabla^2 - \frac{\partial^2}{\partial t^2}\right) \xi^{\mu} = \bar{h}'^{\mu\nu}{}_{,\nu}. \quad (1.21)$$

As long as this requirement is met, we can use the four degrees of freedom in the choice of our gauge (choices of ξ^{μ}) to constrain four of our remaining six degrees of freedom in $A_{\mu\nu}$. In this transformation Eq. (1.16) goes to

$$\bar{h}'_{\mu\nu} = \mathbb{R}e \left[A'_{\mu\nu} \exp \{ i k_{\alpha} x^{\alpha} \} \right], \quad (1.22)$$

where

$$A'_{\mu\nu} = A_{\mu\nu} + k_{\mu} \zeta_{\nu} + k_{\nu} \zeta_{\mu}. \quad (1.23)$$

We have freedom to adjust our coordinates in such a way as to simplify the form of $A_{\mu\nu}$. We will choose them so that we adopt the Transverse–Traceless (TT) gauge. The traceless condition can be imposed via

$$A_{\mu}^{\mu} = \eta^{\mu\nu} A_{\mu\nu} = 0, \quad (1.24)$$

and the transverse by

$$A_{\mu\nu} u^{\beta} = 0, \quad (1.25)$$

where u^{β} are the components of a constant unit four-vector.

If we set $\mathbf{u} = (1, 0, 0, 0)$ and choose a frame in which the wave is travelling in the +z-direction, i.e.

$$k^t = \omega, \quad k^x = k^y = 0, \quad k^z = \omega, \quad (1.26a)$$

$$k_t = -\omega, \quad k_x = k_y = 0, \quad k_z = \omega, \quad (1.26b)$$

then Eq. (1.22) simplifies to

$$\bar{h}'_{\mu\nu}{}^{TT} = A_{\mu\nu}{}^{TT} \cos [\omega(t - z)], \quad (1.27)$$

and

$$h'_{\mu\nu}{}^{TT} = B_{\mu\nu}{}^{TT} \cos [\omega(t - z)], \quad (1.28)$$

where the components $B_{\mu\nu}^{TT}$ are constant. Note that the traceless condition (Eq. (1.24)) ensures that $B_{11}^{TT} = -B_{22}^{TT}$. Thus the GW in this frame may be expressed by

$$h_{\mu\nu}^{TT} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}_{\mu\nu} \exp\{ik_\alpha x^\alpha\}. \quad (1.29)$$

Here we have exhausted the gauge freedom, and the two remaining degrees of freedom correspond to two polarisation states of the GW, which are quadrupolar and offset from each other by 45° . The amplitudes of these components together give the gravitational wave strain, and are denoted by h_+ and h_\times .

1.1.3 Gravitational Waves Far from Their Sources

The far-field approximation outlined below applies when considering GWs at a distance r that is much larger than the spatial scale of their source. To leading order in this approximation, waves generated by matter are described by

$$h_{\alpha\beta}^{TT}(t, \mathbf{x}) = \frac{1}{r} \frac{4G}{c^4} \Lambda_{\alpha\beta,\gamma\delta}(\hat{\mathbf{n}}) \int d^3x' T_{\gamma\delta} \left(t - \frac{r}{c} + \frac{\mathbf{x}' - \hat{\mathbf{n}}}{c}, \mathbf{x}' \right). \quad (1.30)$$

Here Λ is a tensor that transforms to the TT gauge, and $T_{\gamma\delta}$ is the stress-energy tensor.

The leading term in the production of gravitational radiation is the mass quadrupole M , as there is no gravitational dipole moment,

$$\left[h_{\alpha\beta}^{TT}(t, \mathbf{x}) \right]_{\text{quad}} = \frac{1}{r} \frac{2G}{c^4} \Lambda_{\alpha\beta,\gamma\delta}(\hat{\mathbf{n}}) \ddot{M}^{\gamma\delta} \left(t - \frac{r}{c} \right). \quad (1.31)$$

We can simply translate this from the source frame into an arbitrary ‘observer’s’ frame via a rotation operation

$$M_{\alpha\beta} = R_\alpha^\gamma M'_{\gamma\delta} R_\beta^\delta, \quad (1.32)$$

so that the two polarisations can be expressed generically as,

$$h_+ = \frac{G}{rc^4} \left[\ddot{M}'_{11} (\cos^2 \phi - \sin^2 \phi \sin^2 \iota) + \ddot{M}'_{22} (\sin^2 \phi - \cos^2 \phi \sin^2 \iota) - \ddot{M}'_{33} \cos^2 \iota \right. \\ \left. - \ddot{M}'_{12} \sin 2\phi (1 + \sin^2 \iota) + \ddot{M}'_{13} \sin \phi \sin 2\iota + \ddot{M}'_{23} \cos \phi \sin 2\iota \right], \quad (1.33a)$$

$$h_\times = \frac{G}{rc^4} \left[(\ddot{M}'_{11} - \ddot{M}'_{22}) \sin 2\phi \sin \iota + 2\ddot{M}'_{12} \cos 2\phi \sin \iota - 2\ddot{M}'_{13} \cos \phi \cos \iota \right. \\ \left. + 2\ddot{M}'_{23} \sin \phi \cos \iota \right], \quad (1.33b)$$

where ι is the angle between the observer’s z -axis and the source z' -axis, and ϕ is the angle

between the observer's x -axis and the source x' -axis.

1.2 Gravitational Wave Detectors

Having considered the theoretical basis for GWs, we will now take a brief look at the concept behind detecting them using laser interferometry.

1.2.1 The Tidal Effect of Gravitational Waves

We have seen that GWs can be understood as perturbations of the spacetime metric. Recalling Eq. (1.4), this can be expressed as

$$\begin{aligned} ds^2 &= (\eta_{\mu\nu} + h_{\mu\nu}) dx^\mu dx^\nu \\ &= -c^2 dt^2 + dx^2 (1 + h_+) + dy^2 (1 - h_+) + 2 dx dy h_\times . \end{aligned} \quad (1.34)$$

Consider a linearly polarised GW ($h_\times = 0$) moving in the z -direction. This will, to first order, affect a light ray travelling in the x -direction by

$$dt = \frac{dx}{c} \left(1 + \frac{h_+(t)}{2} \right), \quad (1.35)$$

resulting in a change to the time taken for the light to travel between points a and b , separated by a distance l ,

$$t_b - t_a = \frac{l}{c} + \frac{1}{2} \int_{t_a}^{t_b} h(t') dt' . \quad (1.36)$$

It is clear from Eq. (1.34) that the strain, acting orthogonally to the wave vector, increases the travel time in one direction and reduces it in the perpendicular direction. It is this tidal effect that L-shaped interferometers can exploit to detect a passing GW.

1.2.2 Gravitational Wave Interferometry

A simple Michelson interferometer is comprised of a laser source incident upon a half-silvered mirror called the *beam splitter*, which separates the laser beam in half. These beams then travel along perpendicular paths from the beam splitter to a pair of mirrors, where they are reflected and travel back along to the beam splitter before being combined. A photodiode is used to monitor the output for interference between the combined beams (see bottom half of Fig. 1.1). This apparatus may be set in such a way that if the *arms* (distance from beam splitter to end mirrors) are exactly the same length, there should be no output measured by the photodiode.

As it passes this apparatus, a GW may increase the proper distance of one arm, and decrease it for the other. This is easy to see if we consider the case in Eq. (1.36) where the interferometer arms lie along the x - and y -axes. A cartoon of this process is shown in Fig. 1.1. For the x -arm, length l_x , the laser will travel along from a to b and back to c in




Figure 1.1: A linearly polarised GW of strain h alternately stretches and squeezes spacetime so that the proper lengths of the arms change, inducing a phase shift in the output measured by the photodiode. Figure taken from [15].

time

$$\begin{aligned}\Delta t = t_c - t_a &= \frac{2l_x}{c} + \frac{1}{2} \int_{t_a}^{t_c} h(t') dt' \\ &= \frac{2l_x}{c} + \frac{1}{2} \int_{t_a}^{t_a + 2l_x/c} h(t') dt' .\end{aligned}\tag{1.37}$$

We can express a linearly polarised GW with frequency ω_{gw} as

$$h(t) = h_0 \cos(\omega_{gw} t) .\tag{1.38}$$

We can put Eq. (1.38) into Eq. (1.37) and integrate to obtain, to first order in h_0 ,

$$\Delta t = \frac{2l_x}{c} + \frac{h_0}{2\omega_{gw}} \left[\sin\left(\omega_{gw} \left\{t_a + \frac{2l_x}{c}\right\}\right) - \sin(\omega_{gw} t_a) \right] .\tag{1.39}$$

We can simplify this by using the trigonometric identity

$$\sin(a + 2b) - \sin a = 2 \sin b \cos(a + b) .\tag{1.40}$$

This gives us

$$\begin{aligned}\Delta t &= \frac{2l_x}{c} + \frac{\sin(\omega_{gw} l_x/c)}{\omega_{gw}} h_0 \cos\left[\omega_{gw} \left(t_a + \frac{l_x}{c}\right)\right] \\ &= \frac{2l_x}{c} + \frac{\sin(\omega_{gw} l_x/c)}{\omega_{gw}} h(t = t_a + l_x/c) .\end{aligned}\tag{1.41}$$

Since the interferometer measures the interference effect in light returning from the arms, we will be interested in the effect when the laser has completed its trip up and down the arm, i.e. at time $t_c = t_a + 2l_x/c$. This means the light that travelled along the x -arm left the

beam splitter at time $t_a^{(x)}$ given by

$$t_a^{(x)} = t_c - \frac{2l_x}{c} - \frac{\sin(\omega_{gw} l_x/c)}{\omega_{gw}} h(t = t_c - l_x/c), \quad (1.42)$$

and the light that travelled along the y -arm and back will have left the beam splitter at $t_a^{(y)}$

$$t_a^{(y)} = t_c - \frac{2l_y}{c} - \frac{\sin(\omega_{gw} l_y/c)}{\omega_{gw}} h(t = t_c - l_y/c). \quad (1.43)$$

For an interferometer using a laser of frequency ω_λ the electric fields returning from the arms at time t_c are given by

$$\begin{aligned} E^{(x)}(t_c) &= -\frac{1}{2} E_0 \exp\{-i\omega_\lambda t_a^{(x)}\} \\ &= -\frac{1}{2} E_0 \exp\{-i\omega_\lambda (t_c - 2l_x/c) + i\Delta\phi_x(t_c)\}, \end{aligned} \quad (1.44a)$$

$$\begin{aligned} E^{(y)}(t_c) &= -\frac{1}{2} E_0 \exp\{-i\omega_\lambda t_a^{(y)}\} \\ &= -\frac{1}{2} E_0 \exp\{-i\omega_\lambda (t_c - 2l_y/c) + i\Delta\phi_y(t_c)\}, \end{aligned} \quad (1.44b)$$

where the $\Delta\phi$ terms are phase shifts, given by

$$\Delta\phi_x(t_c) = \frac{\omega_\lambda \sin(\omega_{gw} l_x/c)}{\omega_{gw}} h_0 \cos\left[\omega_{gw} \left(t_c - \frac{l_x}{c}\right)\right], \quad (1.45a)$$

$$\Delta\phi_y(t_c) = -\frac{\omega_\lambda \sin(\omega_{gw} l_y/c)}{\omega_{gw}} h_0 \cos\left[\omega_{gw} \left(t_c - \frac{l_y}{c}\right)\right]. \quad (1.45b)$$

These effects are effectively equal in magnitude but of opposite sign. The total phase shift induced in the Michelson interferometer due to a GW is therefore

$$\Delta\phi_{Mich} \equiv \Delta\phi_x - \Delta\phi_y \approx 2\Delta\phi_x. \quad (1.46)$$

In order to maximise this phase shift for a GW of frequency $\omega_{gw} = 2\pi f_{gw}$ the interferometer arm length scale L will ideally be

$$\begin{aligned} L &= \frac{\pi}{2} \frac{c}{\omega_{gw}} \\ &\approx 750 \text{ km} \left(\frac{100 \text{ Hz}}{f_{gw}} \right). \end{aligned} \quad (1.47)$$

Clearly this is an impractical scale for a detector, at least on Earth. Fortunately it is possible to design an interferometer that improves on the limitations of this simple Michelson design.

1.2.3 Gravitational Wave Observatories

The ground-based gravitational wave interferometers that have been constructed and operated, or are currently under construction, have required a great many innovations in design and technology. They may be grouped into two ‘generations’.

First Generation Ground-Based Detectors

The first generation of GW observatories operated between 1999 and 2011 and included:

- **LIGO** [15, 31] (2002-2010)
 Located at two sites in the USA; LIGO Hanford Observatory (LHO) at Hanford, Washington, and LIGO Livingston Observatory (LLO) at Livingston, Louisiana. Each site housed a 4 km scale interferometer, while LHO also featured a 2 km scale interferometer. Before their final, sixth science run the 4 km detectors were given upgrades, sometimes referred to as Enhanced LIGO (eLIGO) [160].
- **Virgo** [32] (2007-2011)
 A 3 km scale detector located at Cascina, Italy, operated by the European Gravitational Observatory.
- **GEO600** [182] (2001-)
 A 600 m scale detector located near Starstedt, Germany, designed and built by a German-British collaboration. The ongoing GEO-HF upgrade project has been running since 2009, and has seen the development and implementation of new detector technologies, primarily to improve sensitivity at high frequencies [69].
- **TAMA300** [38, 161] (1999-2003)
 A 300 m instrument located in Tokyo, Japan. This was the first operational interferometer above 100 m

These L-shaped laser interferometers were of scales $L \sim O(10^2-10^3 \text{ m})$. However, the effective path length of the lasers in LIGO and Virgo, for example, were increased by the use of Fabry-Perot cavities, which consist of mirrors near the beam splitter that reflect the laser back along the arms many times before eventual recombination. For Advanced LIGO ($L = 4 \text{ km}$) this technique gives $L_{\text{eff}} \sim 1120 \text{ km}$, which compares favourably with the result from Eq. (1.47).

This alone, however, is not sufficient to make an interferometer sensitive enough to detect GWs of astrophysical origin. We will see in Section 1.3 that the most powerful sources of GWs in the frequency band of ground-based detectors are only likely to have strains of $h_0 \sim 10^{-21}$ at Earth, which was roughly the peak strain of the first detected

signal GW150914 [22]. This gives

$$\Delta L = \frac{1}{2} h_0 L \sim 2 \times 10^{-18} \text{ m}, \quad (1.48a)$$

$$\Delta L_{\text{eff}} = \frac{1}{2} h_0 L_{\text{eff}} \sim 8 \times 10^{-16} \text{ m}, \quad (1.48b)$$

which is a very small effect that would be easily swamped by instrumental and environmental noise in a standard interferometer. GW observatories have overcome these sources of noise in a number of ways.

At low frequencies ($\lesssim 40$ Hz), ground motion is the dominant noise source for GW observatories. Seismic isolation was key to minimising this effect to levels that could allow detections, with Initial LIGO (iLIGO) and Virgo using test masses suspended from pendula and actuators to dampen any motion in the mirrors. The chambers housing the mirrors were also passively damped by mechanical springs.

At middling frequencies, a number of noise sources including thermal motion and noise from electronic control systems become dominant. Thermal noise, for example, was minimised through the use of very high mechanical quality materials in the optical and suspension systems, and novel chemical coatings on the mirrors.

At the high frequency end ($\gtrsim 100$ Hz) quantum shot noise in the laser dominated. This was suppressed by increasing the number of photons in the arms of the interferometer, which equates to an increase in laser power. This was achieved by using the Fabry-Perot cavities coupled with a power recycling cavity. This reflected the majority of photons travelling back towards the laser, directing them back into the arms.

The cumulative effect of all the techniques used in the first generation detectors was to achieve, with eLIGO, sensitivities capable of detecting strains at the level of 10^{-21} . By the end of the first generation era there were no detections of GWs. LIGO and TAMA300 had carried out combined observations (e.g. [28, 27]), before LIGO and Virgo conducted joint observations in LIGO Science Run 5 and Virgo Science Run 1 (S5/VSR1) (e.g. [10]) and LIGO Science Run 6 and Virgo Science Runs 2 & 3 (S6/VSR2/VSR3) (e.g. [13, 14]). Non-detection was not unexpected [9], and these observations allowed for the calculation of upper rate limits for various predicted sources (e.g. [3, 4]).

A second generation of upgraded detectors was needed before the first direct GW detections could be made.

Second Generation Ground-Based Detectors

The second generation of gravitational wave detectors currently feature upgraded LIGO and Virgo interferometers, known as Advanced LIGO [7] and Advanced Virgo [33] respectively. A new 3 km, cryogenically cooled detector, the Kamioka Gravitational Wave Detector (KAGRA) [40] is being constructed under a mountain at the Kamioka Observatory near Hida, Japan. Additionally, there plans to locate a LIGO detector in India using much of the hardware from the retired LHO 2 km interferometer.

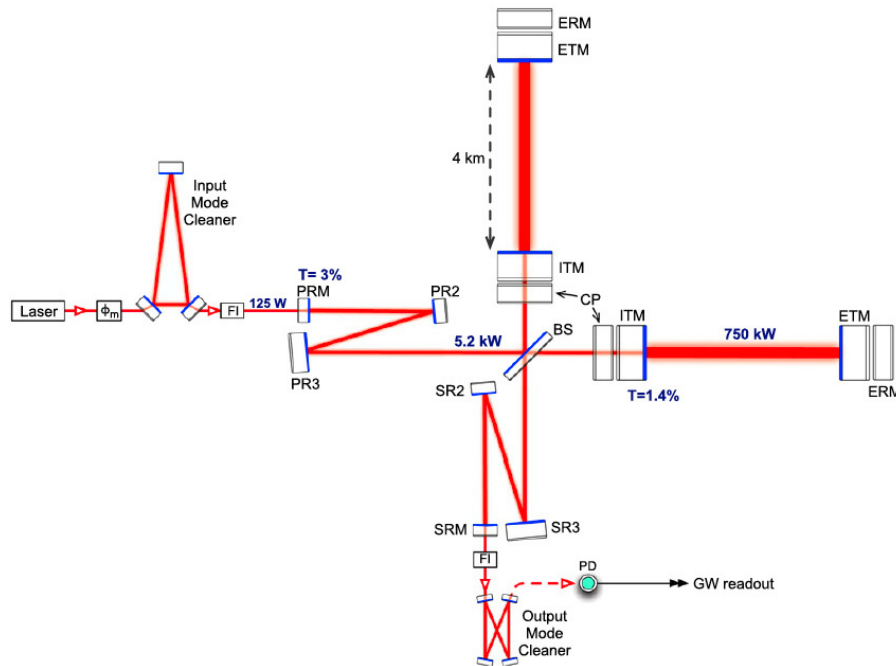


Figure 1.2: The aLIGO optical layout. The laser is passed through an Input Mode Cleaner to smooth the beam, which can contain up to 125 W of power. This travels to the beam splitter (BS) where it is split and sent down the two arms. Each half travels up one of the 4 km long arms to one of the 40 kg fused silica end test masses (ETMs), which are the final stages in quadruple pendula (see Fig. 1.3). The light resonates in each arm between the ETMs and input test masses (ITMs), which form Fabry-Perot cavities. 97 % of the light that returns towards the laser is reflected back into the arms by the power recycling cavity (PRM, PR2/3). In full power mode, up to 750 kW of laser power can build up in the Fabry-Perot cavities. A signal recycling cavity (SRM, SR2/3) effectively increases the finesse of the detector, widening the sensitivity band. The output signal is cleaned up by the output mode cleaner and measured by the photodiode (PD). Figure taken from [7].

Advanced LIGO (aLIGO) represents a complete overhaul upgrade of the original LIGO design, with new components and designs at almost every point in the instrument.

Seismic noise suppression in aLIGO is greatly improved. The main test mass mirrors are the final stages in quadruple pendula (Fig. 1.3). Each stage of the pendulum provides a frequency dependent damping factor of $1/f^2$. Additionally, each stage is actively damped by actuators. The entire pendulum system for each mirror is housed within a chamber that itself is actively isolated from seismic motion. In KAGRA (and the proposed third generation detector Einstein Telescope (ET)), the approach of building the detector underground will also help reduce the seismic noise.

Thermal noise in aLIGO is improved by the use of new mirror and suspension technologies. The aLIGO end test masses are 40 kg fused silica mirrors, suspended by fused silica wires from other, identical fused silica mirrors. The use of this material, which has an




Figure 1.3: Overview of the aLIGO detectors, highlighting the end test mass suspension systems. The quadruple pendulum design significantly reduces coupling between seismic noise and the test masses. Each of stages 1–3 can be actuated using magnetic components on the rearward pendulum (reaction) masses to dampen any motion. Stage 4 is damped via an electrostatic actuator. Figure taken from [20].

extremely high mechanical quality factor, constrains much of the thermal noise into very narrow frequency bands. In the case of KAGRA and ET, thermal noise will be suppressed by cryogenically cooling the test masses and their surroundings.

The use of a 125 W laser in aLIGO, with the potential for up to 750 kW of circulating laser power in the arms, pushes the higher frequency noise level down by an order of magnitude over iLIGO.

A particularly notable new design aspect of aLIGO is the signal recycling cavity (see Fig. 1.2). This directs much of the outgoing signal power back into the arms, which results in a greater finesse and therefore a broadening of the sensitive frequency band [7].

At their design sensitivities, the aLIGO detectors will be 10 times more sensitive than LIGO was in 2010 (Fig. 1.4). This will increase their sensitive volume by a factor of 1000, and it is now expected that they may be able to observe tens to hundreds of binary black hole (BBH) mergers per year [17].

1.3 Astrophysical Sources of Gravitational Waves

We have derived expressions for the amplitudes of both GW polarisation states from a generic source at a great distance in Eq. (1.33). While these expressions neglect the back reaction on the source due to the loss of energy as GWs, and are not appropriate in the limit where the velocity of the matter approaches the speed of light, they can nevertheless allow us to estimate the amplitudes and frequencies of gravitational waves from hypothetical source types. In this section we will consider sources that may be detectable with




Figure 1.4: The sensitivity of eLIGO during its final science run (green) compared to aLIGO during its first observing run (red). Across the full range of frequencies aLIGO is already significantly more sensitive than its predecessor. At frequencies above 100 Hz quantum noise dominates, while at lower frequencies thermal noise in optical coatings and seismic noise dominate. The spikes are predominantly a mixture of mirror suspension modes, calibration lines, and interference from power lines. The aLIGO design is represented by the dark blue curve, and a future upgrade possibility is shown in light blue. Figure taken from [20].

second generation ground-based GW detectors, and categorise them according to signal morphology.

1.3.1 Compact Binary Coalescences

The first directly detected gravitational waves were from the coalescence of a BBH system [22]. We therefore know that black holes are capable of combining in binary pairs that spiral inwards and merge due to emission of GWs, and that they can do so within timescales smaller than the current age of the universe. It is expected that the same is also true for binary neutron star (BNS), and for neutron star–black hole (NSBH) systems.

In order to estimate the strength of such signals, let us consider a system of two masses orbiting about their centre-of-mass, each in a circular orbit. Again ignoring the back-reaction of emission, which causes the orbit to decay over time, the orbit is given by

$$x_0(t) = R \cos\left(\omega_s t + \frac{\pi}{2}\right), \quad (1.49a)$$

$$y_0(t) = R \sin\left(\omega_s t + \frac{\pi}{2}\right), \quad (1.49b)$$

$$z_0(t) = 0, \quad (1.49c)$$

and the reduced mass of the system is defined as

$$\mu \equiv \frac{m_1 m_2}{m_1 + m_2}. \quad (1.50)$$

The second mass moment in the centre-of-mass frame is [124]

$$M^{ij} = \mu x_0^i(t) x_0^j(t), \quad (1.51)$$

so we are able to find the non-vanishing terms

$$M_{11} = \mu R^2 \frac{1 - \cos(2\omega_s t)}{2}, \quad (1.52a)$$

$$M_{22} = \mu R^2 \frac{1 + \cos(2\omega_s t)}{2}, \quad (1.52b)$$

$$M_{12} = -\frac{1}{2} \mu R^2 \sin(2\omega_s t). \quad (1.52c)$$

Putting Eq. (1.52) into Eq. (1.33) we then get the two polarisations of the quadrupole contribution to the strain amplitude,

$$h_+(t) = \frac{1}{r} \frac{4G\mu \omega_s^2 R^2}{c^4} \left(\frac{1 + \cos^2 \iota}{2} \right) \cos(2\omega_s t), \quad (1.53a)$$

$$h_\times(t) = \frac{1}{r} \frac{4G\mu \omega_s^2 R^2}{c^4} \cos \iota \sin(2\omega_s t), \quad (1.53b)$$

where r is the distance to the source, and ι is the inclination of the binary; the angle between the orbital axis and the line of sight.

We may express this in a different way, considering the strain at a time τ before coalescence of the system. We define the GW phase as

$$\begin{aligned} \Phi(t) &\equiv 2 \int_{t_0}^t dt' \omega_s(t') \\ &= \int_{t_0}^t dt' \omega_{gw}(t'), \end{aligned} \quad (1.54)$$

and introduce \mathcal{M} , the *chirp mass* of the system,

$$\mathcal{M} \equiv \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (1.55)$$

Now we can evaluate the strain at a time τ before the merger,

$$h_+(r, \tau, \iota) = \frac{1}{r} \left(\frac{G\mathcal{M}}{c^2} \right)^{5/4} \left(\frac{5}{c\tau} \right)^{1/4} \left(\frac{1 + \cos^2 \iota}{2} \right) \cos \Phi(\tau), \quad (1.56a)$$

$$h_\times(r, \tau, \iota) = \frac{1}{r} \left(\frac{G\mathcal{M}}{c^2} \right)^{5/4} \left(\frac{5}{c\tau} \right)^{1/4} \cos \iota \sin \Phi(\tau). \quad (1.56b)$$

where the phase $\Phi(\tau)$ is now

$$\Phi(\tau) = -2\tau^{5/8} \left(\frac{5GM}{c^3} \right)^{-5/8} + \Phi_0. \quad (1.57)$$

For a sense of scale, we can substitute in some numbers for a canonical BNS system ($m_1 = m_2 = 1.4M_\odot$) assuming a source inclination of 0° ,

$$h \sim 6 \times 10^{-23} \left(\frac{100 \text{ Mpc}}{r} \right) \left(\frac{\mathcal{M}}{1.22 M_\odot} \right)^{5/4} \left(\frac{5 \text{ s}}{\tau} \right)^{1/4}. \quad (1.58)$$

In the last seconds to minutes of an inspiral the emitted GWs from BNS, NSBH, and BBH systems (featuring stellar-mass black holes) fall within the sensitive frequency bands of aLIGO [95] and Advanced Virgo (AdV) [34]. This strain — $h \sim 6 \times 10^{-23}$ — is also within the design sensitivity of the aLIGO detectors. Although not a fully rigorous calculation, this indicates the promise that such sources may hold. We will see in Chapter 5 that the rate of BNS coalescences is thought to be in the range of 10^{-8} – $10^{-5} \text{ Mpc}^{-1} \text{ yr}^{-1}$, which translates to 10^{-2} – 10^1 yr^{-1} within a volume of radius 100 Mpc.

The sensitivity to compact binary coalescence (CBC) signals is helped thanks to relatively accurate modelling of GW waveforms from CBCs using post-Newtonian methods (e.g. [50, 49, 52, 131, 140]), and analytical and numerical relativity (e.g. [94, 110, 166]). We are able to use template waveforms in matched filter analyses [172], which can pick out CBC signals from background noise (e.g [37, 41, 98]). As such, and given the detection of GW150914 during Advanced LIGO Observing Run 1 (O1), GWs from CBCs are expected to be by far the most numerous detections in the advanced detector era [2, 17].

1.3.2 Gravitational Wave Bursts

Let us now consider bursts of GWs from potential astrophysical sources such as core collapse supernovae or long gamma-ray bursts (LGRBs)². In most cases, when core collapse occurs inside a massive star a black hole or a neutron star will be formed. The creation of a neutron star from the collapse will feature a ‘bounce’ where collapse is very suddenly halted by neutron degeneracy pressure. For spherically asymmetric collapse, as would be expected in nature, this bounce will have a non-zero mass quadrupole moment and so emit a burst of GWs. Additionally, rotational instabilities in the stellar core could produce a short burst of GWs.

The strength of the GWs will be dependent on the degree to which collapse is spherically asymmetric, which can depend upon, amongst other things, the angular momentum of the pre-collapse core, the interaction between released neutrinos and inwardly falling matter, and complex magnetohydrodynamic effects on the stellar matter distribution. As a result of such complicated physics, the expected form of the gravitational wave emission

²In Chapter 2 we will see that some LGRBs have been observed in conjunction with core collapse supernovae, and are therefore believed to be the same class of astrophysical event.

from a core collapse event is relatively unknown. This lack of knowledge about the waveform morphology means that searches for GWs from core collapse events look for generic bursts of power within the data from GW detectors.

We can approximate the strain from a source emitting energy E_{gw} in a time T as [152],

$$h \sim \frac{1}{\pi r f} \sqrt{\frac{E_{gw}}{T}}, \quad (1.59)$$

assuming the emission is at the frequency f . Therefore, we may estimate the strain of a core collapse GW burst at a distance of 100 kpc that emits $10^{-7} M_{\odot} c^2$ at a frequency of 1 kHz over a 1 ms period,

$$h \sim 6 \times 10^{-22} \left(\frac{E_{gw}}{10^{-7} M_{\odot}} \right)^{1/2} \left(\frac{1 \text{ ms}}{T} \right)^{1/2} \left(\frac{1 \text{ kHz}}{f} \right) \left(\frac{100 \text{ kpc}}{r} \right). \quad (1.60)$$

Core collapse supernovae occur at a rate of $\sim 10^{-4} \text{ Mpc}^{-3} \text{ yr}^{-1}$ [168], with LGRBs orders of magnitude rarer. Clearly, observation of GW bursts from these events would be fortunate since it requires the source to be within our Milky Way galaxy or, in some more extreme emission models, within a few Mpc. However, a detection could provide great astrophysical insights, therefore unmodelled searches are carried out for these GW bursts.

2

GAMMA-RAY BURSTS

In the late 1960s the *Vela* family of gamma-ray sensitive satellites were launched by the United States of America to monitor the Earth for gamma-ray flashes that might be produced by nuclear weapons tests. On 2 July 1967 the Vela 3 and Vela 4 satellites detected an extremely bright flash of gamma-rays that did not appear to resemble a nuclear weapons test [153]. This would later be identified as the first observed gamma-ray burst (GRB). Further observations between July 1969 and July 1972 featured sixteen bursts that were identified unambiguously as being of cosmological origin [113]. The energies involved were clearly extreme, and studies since then have concluded that GRBs are the most energetic electromagnetic events in the universe. Currently an average of approximately one GRB is seen per day.

In this chapter we will summarise the current state of GRB astronomy, and why they are events of interest for GW astronomy.

2.1 The Current Gamma-ray Burst Paradigm

GRBs are observed isotropically across the sky and up to cosmological distances [134]. The lowest confirmed redshift was 0.0085 for GRB 980425 [85, 169], corresponding to a luminosity distance of approximately 36 Mpc. By comparison, some GRBs have redshifts >8 and are among the most distant objects observed. They are also the most luminous electromagnetic events known, with typically a significant fraction of a solar rest mass emitted on timescales of seconds or less.

GRBs are characterised by extremely luminous prompt gamma-ray emission between keV–GeV, which in some cases appears as a simple flash with an exponentially decay-

ing light curve. In others it can feature multiple distinct peaks, apparent precursors and subsequent flares, and rapid variation above the background. Despite this sometimes complex structure, prompt emission spectra are generally well-fit by the empirical Band function [42]

$$f_{BAND}(E) = \begin{cases} A \left(\frac{E}{100 \text{ keV}} \right)^\alpha \exp\left(-\frac{E}{E_0}\right) & (\alpha - \beta)E_0 \leq E \\ A \left[\frac{(\alpha - \beta)E_0}{100 \text{ keV}} \right]^{\alpha - \beta} \exp(\beta - \alpha) \left(\frac{E}{100 \text{ keV}} \right)^\beta & (\alpha - \beta)E_0 \geq E \end{cases}. \quad (2.1)$$

This is a broken power law relation, with the break typically occurring at energies between 100 keV–1 MeV. The primary emission processes behind this prompt emission are generally thought to be synchrotron and/or inverse Compton scattering occurring in a highly relativistic, collimated outflow – the fireball. In fireball models, the inverse Compton process may be the result of collisional forward/reverse shocks inside the expanding material, which accelerate electrons that in turn boost photons into a power law energy distribution. In optically thin regions of the fireball, these photons may escape and be observed as the prompt emission. The exact nature of these shocks, and the manner in which they evolve via cooling, is subject to many proposed explanations.

When multiple peaks are present, a combination of non-thermal components, each with its own fit of the Band function, may fit the observed spectral evolution well (e.g. [30]), however some of the brightest observed GRBs exhibit a second component at higher energies (\geq MeV). Within the general class of leptonic fireball models there exist some that seek to explain this via upscattering of photons from the jet scattering photosphere (e.g [170]). A schematic GRB spectrum for this type of model is shown in Fig. 2.1 [126]. Here we see a Band component, a synchrotron component, and an upscattered photospheric component can produce emission over a very large range of energies, crucially up to $>$ GeV.

Other models look to baryonic and/or magnetic physics to drive the prompt emission and describe the variation from one GRB to the next. Better observations across the full energy range will help to constrain the physics of GRB prompt emission, as will multi-messenger observations.

Observations in X-ray, infrared, optical, and radio bands have shown broadband emission following the prompt emission, commonly referred to as the *afterglow*. This may also be explained by forward shocks, this time when the outflow collides with interstellar material [137].

GRBs vary in duration from approximately 10^{-3} s to over 10^4 s. For a time the consensus was that GRBs were probably caused by the mergers of neutron star binaries [72]. However, a bimodality observed in the duration and spectral hardness of GRBs suggested a split between ‘long-soft’ and ‘short-hard’ varieties (Fig. 2.2). This implies more than one class of progenitors [116]. Indeed, more recently evidence has emerged that suggests there may be a third, ‘ultra-long’ class of GRBs [119], although this is not conclusive [185].




Figure 2.1: A schematic gamma-ray burst spectrum for the leptonic photospheric-internal shock model [170]. The photospheric component follows a Band function [42], and some of the higher energy photons from this component are upscattered by electrons in internal shocks to produce an upscattered photospheric component (UP). Classical synchrotron and synchrotron self-Compton (SSC) processes may also contribute. In some cases it may be possible for the UP and photospheric components to overlap and appear as a single component. Figure reprinted from [126].




Figure 2.2: GRBs from the Fourth BATSE Gamma-ray Burst Catalogue, showing a partial separation into ‘long-soft’ and ‘short-hard’ populations. This was an early sign that at least two physical phenomena were responsible for GRBs. Figure taken from [134].

2.1.1 Long Gamma-ray Bursts

Long gamma-ray bursts are linked to core collapse supernovae. Gamma-ray emission from supernovae were predicted before the discovery of GRBs [58]. It was also noted that if GRBs were to lie at cosmological distances, they would require energies comparable to those of supernovae ($\sim 10^{51}$ erg) [135]. Later, it was proposed that LGRBs could be explained by accretion onto a black hole formed from a massive stellar core collapse [183], and early numerical work supported this as a viable central engine [122]. Some tentative evidence existed that connected GRBs with star forming regions [136], however the discovery of SN 1998bw in association with GRB 980425 [85] was the first strong piece of observational evidence that the two phenomena were linked [86]. Following this, the observation of SN 2003dh in association with GRB 030329 made the observational evidence for a link compelling [101].

For a more detailed review of the evidence in support of a connection between LGRB and supernovae see [184].

This connection has made LGRBs of interest to GW astronomy, since in Section 1.3.2 we saw that core collapse events may produce detectable GWs. In this thesis, however, we will concentrate primarily on CBC GW signals (Section 1.3.1), and therefore LGRBs are not the subject of our focus.

2.1.2 Short Gamma-ray Bursts

The compact binary merger progenitor model, involving two neutron stars or a black hole and a neutron star, is still favoured for short gamma-ray bursts (SGRBs). SGRB light curves vary on very short timescales, which indicates a compact source smaller than a stellar-scale object [147]. They are also seen to have larger displacements with respect to their assumed host galaxies than LGRBs [56]. This is suited to a progenitor involving old astrophysical objects that evolve over long periods of time after having been subjected to significant ‘kicks’, exactly the expectation for compact binary mergers. It is not consistent with the supernovae of massive – and therefore young – stars.

In this model the burst is again powered by accretion onto a black hole [72, 129]. Immediately after merger, energy comparable in scale to a stellar rest mass is emitted in neutrinos, photons, and GWs. A neutrino driven, highly relativistic jet expands into the dense medium around the black hole, powered by accretion, causing shock fronts to form. Prompt gamma-ray emission may be powered by synchrotron and/or inverse Compton scattering of relativistic electrons in the jet. Various shock fronts from interaction of the jet with surrounding material, which may propagate both outwards and inwards, are thought to be largely responsible for the afterglow emission in X-rays and at lower energies.

More recently, the detection of a kilonova [121] associated with the SGRB 130603B [46, 165] has provided strong support for the compact merger hypothesis. This late time emission in the infrared band is predicted to come from the decay of r-process elements produced in the merger of a compact binary.




Figure 2.3: The observation of a kilonova associated with the short gamma-ray burst 130603B. The blue data points (optical) are fit very well by a smoothly broken power law decays (dashed line). The black data points (X-ray) also exhibit this sharp decay. However, the near infrared data (red) show a similar decay at early times, but an excess approximately one week after the GRB, which then disappears quickly. This was interpreted as a kilonova [121], the emission from decaying r-process elements in the ejecta of a compact binary merger. Figure taken from [165].

However, the physics of SGRB central engines are still relatively poorly understood. This is largely due to the difficulty in observing the late-time emission, or *afterglow*. SGRBs are generally less energetic than LGRBs. They are also characterised by harder spectra, so emit proportionally less of their energy in the X-ray and optical bands where afterglows are likely to be seen. In fact, SGRB X-ray and optical afterglows are roughly 7 times fainter than those of LGRBs with the same isotropic gamma-ray energy [48].

One of the current goals of electromagnetic SGRB observations is to observe a *jet break*, a steepening of the power law decay of the afterglow flux. This is interpreted as being a sign of a highly relativistic jet. Due to relativistic beaming, an observer will see emission from a fraction of the jet. As the outflow in the jet slows, the Lorentz factor drops, the beaming lessens, and the observer sees emission from a larger and larger fraction of the jet. The material in the jet also begins to expand faster as the opposite sides of the jet become causally connected. Eventually the observer sees the entire jet and the expansion

will no longer be partially compensated by seeing more of the jet. The flux will drop more rapidly as a result, producing a break in the light curve. The time at which this occurs after the prompt GRB is linked to the angular size of the jet.

If the emission in SGRBs is collimated, as it is for LGRBs, the true energy E_γ may be inferred from the time of an observed jet break (in days), $t_{j,d}$. The angular size of the collimated jet is given by [151]

$$\theta_j = 0.13 \left(\frac{t_{j,d}}{1+z} \right)^{3/8} \left(\frac{n_0 10^{52} \text{ erg}}{E_{iso}} \right)^{1/8}, \quad (2.2)$$

where E_{iso} is the energy if isotropic emission is assumed, n_0 is the number density of the material surrounding the central engine (cm^{-3}), and z is the GRB redshift. Therefore

$$E_\gamma = (1 - \cos \theta_j) E_{iso}. \quad (2.3)$$

Unfortunately there have been very few observations of jet breaks in SGRBs to date. Analysis of the few measurements that there are suggest a median value of $\theta_j \approx 10^\circ$, and energies $E_\gamma \approx 10^{49}$ erg [83].

For further background on SGRBs see the very good review articles [128, 48] and references therein.

2.2 Gamma-ray Burst Satellites and Observations

A number of satellites are currently used to monitor the sky for GRBs. These may be dedicated GRB observatories or simply spacecraft carrying gamma-ray sensitive detectors on board.

Swift

Swift is named for its rapid slewing capabilities, which allow it to autonomously target GRBs with its on board instruments within roughly a minute of first gamma-ray detection. It features three instruments, of which the Burst Alert Telescope (BAT) is the primary tool for identifying GRBs [43]. The BAT observes between 15–150 keV and has a field of view of ~ 2 sr. It observes around 10 SGRBs per year. In total, *Swift* has detected over 1,000 GRBs between 2004 and the present time, of which approximately 90 have measured BAT $t_{90} \leq 2$ s – in other words, are likely SGRBs¹.

GRBs are generally localised on the sky with an accuracy of 1–4 arcminutes by the BAT, and better if the burst is followed up with the onboard X-Ray Telescope (XRT) (3–5 arcseconds) or Ultraviolet/Optical Telescope (UVOT) (0.3 arcseconds). This ability makes *Swift* a crucial tool for multi-wavelength studies of GRBs and their host environments.

¹http://swift.gsfc.nasa.gov/archive/grb_table/stats/

Fermi

The *Fermi* satellite² was launched in 2008 and carries on board the Gamma-ray Burst Monitor (GBM) [125] and Large Area Telescope (LAT). The GBM is essentially an all-sky telescope with a field of view of 9.5 sr, covering the frequency range between \sim keV– \sim 30 MeV. This coverage comes from a combination of 12 sodium iodide scintillators (\sim keV– \sim MeV), which identify the transient time and location, and two bismuth germanate scintillators (\sim 150 keV– \sim 30 MeV), which span the gap to, and slightly overlap with, the high energy sensitive band of the LAT (20 MeV–300 GeV). The GBM typically observes around 45 SGRBs per year, of which only a small fraction are seen in the LAT. The *Fermi* localisation is typically accurate to tens or hundreds of square degrees [111], making optical followup of these events challenging. To date, no afterglow has been observed for an SGRBs seen only by *Fermi*, and consequently the redshifts of these bursts are not known.

Recent work has demonstrated that *Swift* and *Fermi* are observing the same population of SGRBs [54] – essentially every burst observed by *Swift* that was in the field of view of *Fermi* was observed by the GBM, and vice versa. *Fermi* is currently operational, with its 10 year funding cycle ending in 2018, though it may continue operations further.

The InterPlanetary Network

The set of instruments which make up the InterPlanetary Network (IPN)³ are not dedicated GRB satellites, but instead have GRB monitors on board [104]. The majority of satellites in the network are unable to localise the bursts individually but it is possible to localise bursts observed in numerous satellites using triangulation. The sizes and shapes of these error regions vary greatly, depending upon the number of satellites and their locations (more distant satellites greatly improve localisation). The IPN provides essentially all-sky coverage for GRBs, although, given the sensitivity of the detectors, the GRBs observed tend to be closer.

SVOM

The Space-based multi-band astronomical Variable Objects Monitor (SVOM) satellite⁴ is a recently approved Chinese-French mission, scheduled for launch in 2021. SVOM will have a similar sky coverage to *Swift*, and will also carry X-ray, optical and ultraviolet telescopes that can be rapidly and automatically slewed to observe afterglows [45, 60]. Additionally, it will have dedicated resources on the ground for wide-angle, rapid follow-up. In the era of aLIGO and AdV this is likely to be an important source of GRB detections.

²<http://fermi.gsfc.nasa.gov/>

³<http://heasarc.gsfc.nasa.gov/W3Browse/all/ipngrb.html>

⁴<http://www.svom.fr/>

2.3 Short Gamma-ray Bursts and Gravitational Waves

SGRBs have long been discussed as promising targets for gravitational wave astronomy due to their probable binary merger progenitors [115]. They likely represent some of the most intense electromagnetic *and* GW events in the universe [62, 157]. There is a realistic chance of jointly observing an SGRB with aLIGO/AdV in conjunction with gamma-ray sensitive telescopes. The closest spectroscopically confirmed SGRB was 080905A, with a redshift of 0.12 [149], corresponding to a luminosity distance of ~ 550 Mpc. An SGRB at this distance is likely to be close to the sensitive threshold of aLIGO. However, the rate of SGRBs may be as high as $10^{-8} \text{ Mpc}^{-3} \text{ yr}^{-1}$, and this could equate to as many as 2–3 joint observations per year when the GW detectors are at their design sensitivities (see Chapter 5).

GW observations of SGRBs will make possible direct observation of the central engines that power these events, a feat that electromagnetic observations alone cannot achieve due to circumburst material and ejecta [44]. In addition to the astrophysical insights a joint observation could deliver, SGRBs represent an attractive trigger for GW follow-up searches. The prior observation of an SGRBs provides the time and sky position of a potential GW source. A targeted search for a binary merger GW signal, informed by the SGRB observation, need only search a small fraction of the parameter space of a generic binary merger search – one that searches the whole sky at all times and for all binary merger signals (including BBHs). Consequently, it is possible to significantly reduce the detection threshold for the targeted SGRB search (see Chapters 4 and 5), thereby increasing the sensitivity of the search.



**SEARCHES WITH
FIRST GENERATION
DETECTORS**

“Astronomy? Impossible to understand and madness to investigate.”

Sophocles

3

FIRST SEARCHES WITH LIGO AND VIRGO

In Part I we have outlined the reasons why SGRBs are interesting events for GW follow-up. In this chapter we will describe the methods used to search GW detector data for signals associated with SGRBs. We will summarise results from LIGO and Virgo science runs prior to 2010 for SGRBs observed by *Swift* and *Fermi*, then describe a later analysis for IPN SGRBs.

3.1 Pipeline Summary

In [98] the authors introduced an analysis pipeline that performed a targeted search for CBC GW signals in LIGO and Virgo data associated with SGRBs seen by *Swift* and *Fermi* in S6/VSR2/VSR3 [13]. In later chapters we will describe how this pipeline has since been developed, however here we simply review the underlying methodology.

3.1.1 Multi-Detector Matched Filter

The pipeline makes use of the well understood gravitational waveforms emitted during BNS or NSBH mergers [49] to perform a modelled search for a CBC signal in data from operational GW detectors. A bank of template waveforms [132] that densely cover the mass parameter space was used to perform a *matched-filter* analysis [172]. The analysis makes use of the known sky location of the SGRB and the relative GW detector sensitivities to appropriately time shift and weight the data streams from the individual detectors to perform a *coherent* analysis. In this framework a network signal-to-noise ratio (SNR) is

calculated directly.

This is in contrast to the majority of searches for CBC signals, which perform the matched-filter independently on individual interferometer data streams before comparing the resulting *triggers* to search for coincident events (see e.g. [41]). Such searches are termed *coincident* analyses.

The coherent approach affords several benefits. First, by performing the analysis coherently, we combine the detector data to produce two data streams which are sensitive to the two GW polarisations. Any other, orthogonal data streams will necessarily contain only noise and can either be ignored, or used to eliminate noise transients which will often contribute power to these *null streams*. Additionally, by combining the data from the detectors at the time of analysis, we will accumulate power from all detectors, not just those which produced a trigger above a pre-determined SNR threshold.

It was shown in [98] that a coherent analysis provides an improvement in sensitivity over the coincident one, but is more computationally expensive. A targeted GRB search, where both the sky location and arrival time of the signal are constrained is ideal for performing the more sensitive, coherent analysis.

The amplitude of a GW signal from a non-precessing binary may be decomposed into two polarisations h_+ and h_\times ,

$$h_+(t) = \mathcal{A}^1 h_0(t) + \mathcal{A}^3 h_{\pi/2}(t), \quad (3.1a)$$

$$h_\times(t) = \mathcal{A}^2 h_0(t) + \mathcal{A}^4 h_{\pi/2}(t). \quad (3.1b)$$

Here, h_0 and $h_{\pi/2}$ denote the two phases of the waveform, which depend upon the binary masses as well as the coalescence time of the signal. These are calculated using the post-Newtonian formalism [49]. In S6/VSR2/VSR3, the analyses were restricted to non-spinning systems. However, the search is easily extended to binaries with spins aligned with the orbital angular momentum by simply generating additional templates to cover the spin parameter space (see e.g. [51, 97]). The amplitude terms for an inspiral GW signal are

$$\mathcal{A}^1 = \frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \cos 2\phi_0 \cos 2\psi - \frac{D_0}{D} \cos \iota \sin 2\phi_0 \sin 2\psi, \quad (3.2a)$$

$$\mathcal{A}^2 = \frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \cos 2\phi_0 \sin 2\psi + \frac{D_0}{D} \cos \iota \sin 2\phi_0 \cos 2\psi, \quad (3.2b)$$

$$\mathcal{A}^3 = -\frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \sin 2\phi_0 \cos 2\psi - \frac{D_0}{D} \cos \iota \cos 2\phi_0 \sin 2\psi, \quad (3.2c)$$

$$\mathcal{A}^4 = -\frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \sin 2\phi_0 \sin 2\psi + \frac{D_0}{D} \cos \iota \cos 2\phi_0 \cos 2\psi. \quad (3.2d)$$

These terms are dependent on four variables: the source distance, D ; the coalescence phase, ϕ_0 ; the polarisation angle, ψ ; and the inclination angle, ι . D_0 is a scaling distance (usually 1 Mpc). It is worth noting that, for any set of amplitudes \mathcal{A}^u , there is

a unique set of $\{D, \iota, \phi_0, \psi\}$, up to reflection and rotation symmetry – for example, the transformation $\{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^3, \mathcal{A}^4\} \rightarrow \{-\mathcal{A}^2, \mathcal{A}^1, \mathcal{A}^4, -\mathcal{A}^3\}$ is equivalent to $\phi \rightarrow \phi + \pi/4$, and $\{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^3, \mathcal{A}^4\} \rightarrow \{-\mathcal{A}^1, -\mathcal{A}^2, -\mathcal{A}^3, \mathcal{A}^4\}$ is equivalent to $\{\phi, \psi\} \rightarrow \{\phi + \pi/4, \psi + \pi/4\}$.

The GW signal seen by a detector X is a combination of the two polarisations, each weighted by an antenna power pattern factor $F_{\{+, \times\}}^X$ [108], which describes the relative response of the detector to each polarisation,

$$h^X(t) = F_+^X h_+(t^X) + F_\times^X h_\times(t^X). \quad (3.3)$$

Here, t^X is the time of arrival of the signal at detector X , which will depend upon a fiducial arrival time (for example at the geocentre) and the relative location of the detector and source.

In matched-filtering analysis the inner products between a template gravitational waveform time series $h(t)$ and detector data stream time series $s(t)$ are calculated. In general, the inner product between two such time series, a^X and b^X , is given by

$$(a^X | b^X) = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}^X(f) \cdot \tilde{b}^X(f)^*}{S_n^X(f)} df, \quad (3.4)$$

where $S_n^X(f)$ is the noise power spectral density in detector X , and $\tilde{a}(f)$ denotes the Fourier transform of the time series $a(t)$. For binary merger signals, the two phases h_0 and $h_{\pi/2}$ are orthogonal, in the sense that

$$(h_0 | h_{\pi/2}) = 0. \quad (3.5)$$

For a network of detectors, we define the multi-detector inner product as the sum of the single detector inner products,

$$(\mathbf{a} | \mathbf{b}) \equiv \sum_{X=1}^d (a^X | b^X), \quad (3.6)$$

where d denotes the number of detectors in the network. The multi-detector log-likelihood is then defined as,

$$\begin{aligned} \ln \Lambda &= (\mathbf{s} | \mathbf{h}) - \frac{1}{2} (\mathbf{h} | \mathbf{h}) \\ &= \left[\mathcal{A}^\mu (\mathbf{s} | \mathbf{h}_\mu) - \frac{1}{2} \mathcal{A}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}^\nu \right], \end{aligned} \quad (3.7)$$

where \mathbf{s} is the time series containing the quadrature sum of individual detector data streams, $\mathbf{h} = (\mathbf{F}_+ \mathbf{h}_0, \mathbf{F}_\times \mathbf{h}_0, \mathbf{F}_+ \mathbf{h}_{\pi/2}, \mathbf{F}_\times \mathbf{h}_{\pi/2})$, and the matrix

$$\mathcal{M}_{\mu\nu} \equiv (\mathbf{h}_\mu | \mathbf{h}_\nu). \quad (3.8)$$

Maximising this likelihood ratio over the amplitude parameters \mathcal{A}_μ , we obtain the max-

imised coherent SNR,

$$\rho_{\text{coh}}^2 \equiv 2 \ln \Lambda|_{\text{max}} = \left[(\mathbf{s}|\mathbf{h}_\mu) \mathcal{M}^{\mu\nu} (\mathbf{s}|\mathbf{h}_\nu) \right], \quad (3.9)$$

where $\mathcal{M}^{\mu\nu}$ is the inverse of the matrix $\mathcal{M}_{\mu\nu}$.

The coherent SNR forms the basis of the detection statistic and has a χ^2 background distribution with four degrees of freedom. The four degrees of freedom correspond to the four components of the gravitational wave signal – the 0 and $\pi/2$ phases of the two polarisations. This becomes more transparent if we work in the dominant polarisation frame. In this frame, the network is maximally sensitive to the + polarisation and the two polarisations are orthogonal. Then, the coherent SNR can be re-expressed as

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s}|\mathbf{F}_+\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_+\mathbf{h}_0|\mathbf{F}_+\mathbf{h}_0)} + \frac{(\mathbf{s}|\mathbf{F}_\times\mathbf{h}_0)^2 + (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_\times\mathbf{h}_0|\mathbf{F}_\times\mathbf{h}_0)}. \quad (3.10)$$

In Gaussian noise, the coherent SNR would be the detection statistic. Events with a larger coherent SNR would be less likely to be due to noise fluctuations and consequently more likely to be due to a GW signal. However, in real data GW signals are not the only cause of deviations from the background distribution. Noise transients, or *glitches*, also contribute to the background. Although glitches will not typically mimic template waveforms, if they are large enough they will still produce a large SNR. Consequently, we must use a number of consistency tests to eliminate or down-weight triggers that are unlikely to be due to a GW signal incident upon the detector network.

3.1.2 Signal Consistency

Matched filtering alone leads to the identification of a large number of triggers, many of which are purely due to non-Gaussian noise transients present in the data. Such noise transients may be discarded by performing signal consistency tests across the individual detectors that make up the network. Here, we briefly describe the different tests used in the analysis.

Null Stream Consistency

Null stream consistency makes use of one or more null data streams or, in the case of this pipeline, the related null SNR statistic. This is simply the SNR observed in the detector network that is *not* consistent with the signal model;

$$\rho_{\text{null}}^2 \equiv \sum_X \rho_X^2 - \rho_{\text{coh}}^2, \quad (3.11)$$

where ρ_X is the SNR in detector X . For a signal which matches the template waveform, there will be no signal power in the null SNR, so for a population of signals null is expected to be χ^2 distributed with $2d-4$ degrees of freedom due to the presence of noise (where again

d is the number of detectors in the network). Therefore, for networks with fewer than three detectors, this consistency test is not used. An incoherent, non-Gaussian transient noise event will contribute to the null SNR and consequently a large null SNR is used to eliminate spurious events via a hard cut if

$$\begin{aligned} \rho_{\text{null}} &> 5.25, & \rho_{\text{coh}} &\leq 20 \\ \rho_{\text{null}} &> \frac{\rho_{\text{coh}}}{5} + 5.25, & \rho_{\text{coh}} &> 20 \end{aligned} \quad (3.12)$$

Single Detector Thresholds

Noise transients are, by their nature, events which occur in a single detector. Conversely, gravitational wave events will lead to signal power being distributed among all detectors in the network. We can use this difference to further reduce the background due to glitches. The most effective, and most straightforward, method is simply to require that a signal is observed with an SNR above threshold (typically four) in at least two detectors. This serves to eliminate the majority of glitches, which have power in only one detector, with very little effect on signals.

χ^2 Tests

When matched-filtering identifies a trigger with a large SNR there is necessarily some component of the data which matches the signal $h(t)$. If the trigger is caused by a noise glitch, there is likely to be an additional, orthogonal component of the data which is not well described by Gaussian noise. χ^2 tests are designed to eliminate glitch triggers by identifying power that is not consistent with either signal or Gaussian noise. To do so, we introduce a set of basis waveforms T^i which are orthonormal and also orthogonal to the signal waveform $h(t)$. Specifically, we require

$$(\mathbf{T}_\mu^i | \mathbf{T}_\nu^j) = \delta^{ij} \delta_{\mu\nu} \quad \text{and} \quad (\mathbf{T}_\mu^i | \mathbf{h}_\nu) = 0, \quad (3.13)$$

where μ, ν refer to the waveform components and i, j the waveforms that comprise the basis for the χ^2 test. We then construct a χ^2 statistic as

$$\chi^2 = \sum_{\mu=1}^4 \sum_{i=1}^N (\mathbf{T}_\mu^i | \mathbf{s})^2. \quad (3.14)$$

In the presence of a signal that matches the template waveform (or no signal), the statistic will be χ^2 distributed with $4N$ degrees of freedom. If the data contains some additional, non-Gaussian noise the χ^2 value will be elevated provided that the set of templates T^i captures at least a fraction of the power contained in the glitch. Triggers with a large χ^2 value are discarded. In practice it is far from trivial to choose the set of waveforms T^i so that they are both orthonormal and orthogonal to $h(t)$, and match a variety of non-Gaussianities. Three different χ^2 tests have been implemented in the analysis:

- i *Frequency bins*: The test waveforms T^i are generated by chopping up the template $h(t)$ into $(N+1)$ sub-templates in the frequency domain, each of which contains an equal amount of power. From these, we generate N orthonormal waveforms which are also orthogonal to $h(t)$.
- ii *Template bank*: The test waveforms T^i are taken from the template-bank of binary merger waveforms used in the search. In general, these will not be orthogonal to $h(t)$, but it is straightforward to subtract the part proportional to $h(t)$. However, it is more difficult to render the waveforms T^i orthonormal. In practice we do not attempt to do so, but instead use an empirical threshold based on an effective number of degrees of freedom. The templates are chosen so that they cover the mass parameter space of the search.
- iii *Autocorrelation*: The test waveforms T^i are simply copies of the waveform $h(t)$ offset in time from the original. As with the template bank, it is straightforward to remove the component of T^i that is proportional to $h(t)$. We do not attempt to orthonormalise the T^i and again empirically set the threshold.

Re-weighted SNR

In addition to discarding triggers which fail the signal consistency test described above, we also re-weight the SNR of triggers based on the values of the χ^2 tests and null SNR. This allows us to better differentiate signals from noise background. The re-weighting is chosen such that the SNR of signals will be unaffected while those noise triggers which do not match well with the template waveform will be down-weighted. We perform two sets of down-weighting. Firstly, with the χ^2 values,

$$\rho_{\chi^2} = \begin{cases} \rho_{\text{coh}} & \chi^2 \leq n_{\text{dof}} \\ \frac{\rho_{\text{coh}}}{\left\{ \left[1 + \left(\frac{\chi^2}{n_{\text{dof}}} \right)^3 \right] / 2 \right\}^{1/6}} & \chi^2 > n_{\text{dof}} \end{cases}, \quad (3.15)$$

then with the null SNR,

$$\rho_{\text{rw}} = \begin{cases} \rho_{\chi^2} & \rho_{\text{null}} \leq 4.25 \\ \frac{\rho_{\chi^2}}{\rho_{\text{null}} - 3.25} & \rho_{\text{null}} > 4.25 \end{cases}. \quad (3.16)$$

This re-weighted SNR value is the detection statistic used for evaluating candidate events.

We note that the χ^2 re-weighted SNR given in Eq. (3.15) is different from the one used in the original paper [98], and is in fact the same as the weighting applied in a number of past coincident CBC searches (see e.g. [14]). In particular, the exponents in the denominator have been changed. In the process of developing an all-sky, all-time coherent analysis [123], it was found that the original re-weighting left a small tail of high SNR

noise events. These had not been observed in the GRB search previously, due to the limited amount of data used in the analyses. By using a re-weighted SNR identical to the one used in the all-sky coincidence search [41], we were able to eliminate the high SNR events. The same re-weighting has now been applied in the GRB search.

3.1.3 Event Significance

In S6/VSR2/VSR3 this targeted, coherent search was carried out whenever an observed SGRB was detected during a time that at least two GW detectors were operating and had good quality data for a sufficiently long period of time either side of the SGRB. We search for a signal in a 6 s window covering 5 s before to 1 s after the Earth crossing time of the SGRB called the *on-source* window. The analysis is performed for all template waveforms in the template bank covering the mass parameter space. For each template the re-weighted SNR is calculated, and the template producing the largest re-weighted SNR during the on-source window is retained as the event candidate.

However, we require additional data around this time in order to ensure that the detectors were operating stably at the time of the SGRB, and to provide a good estimate of the detector sensitivity. Our ability to detect a GW signal associated with an SGRB depends upon both the stationary noise background and also the non-stationary noise transients in the data which might mask a signal. The data surrounding the on-source time is used to evaluate both of these. This time is designated *off-source*. This data will not contain a signal corresponding to the SGRB and is also unlikely to contain a GW signal from the same sky position which is unassociated with the SGRB, thus any events occurring in the off-source will be due to background noise.

In a typical search we use approximately an hour of data for the off-source, and split this into trials with durations equal to that of the on-source window. This gives us a means of characterising the background noise in our detector network around the time of the SGRB. The significance of the on-source event is determined by calculating the false alarm probability, or p-value. This is simply the fraction of off-source trials with an event of equal or greater significance than in the on-source.

To evaluate the sensitivity of the pipeline to finding GW signals in the data around the time of the SGRB, we inject a number of simulated signals into the off-source data. The simulated signals are drawn randomly from an astrophysically motivated distribution of distances, component masses and spins and binary inclination. The simulated signals are compact binary merger waveforms at 3.5 post-Newtonian order [50, 49], where one component of the binary is taken to be a neutron star and the second either a neutron star or black hole. The efficiency of the analysis at recovering these signals provides a measure of pipeline performance and produces an estimate of the distance to which the pipeline is sensitive.

3.2 Early Results

The earliest analyses looking for GWs in association with GRBs used search methods that focused on correlated excess power, including the search for a signal associated with the LGRB 030329 [26]. Later, however, the analysis method described in Section 3.1 was used to carry out numerous SGRB searches on data from the LIGO and Virgo detectors.

3.2.1 Full Science Run Analyses

During the era of first generation GW detectors, searches were performed for SGRBs observed by *Swift* [43] and *Fermi* [125] during S5/VSR1, and S6/VSR2/VSR3.

In S5/VSR1 a search was performed for 22 SGRBs, which found no candidate events [10]. For 21 of these, analysis was only run for the two most sensitive detectors in the network, which at the time featured four detectors (L1 at LLO, H1 and H2 at LHO, and the Virgo instrument, V1). Only GRB 070923 was analysed using data from three detectors (H1, L1, and V1). The GRB 070201 was of particular interest, and is further described below, as is GRB 051103, which just preceded the analysis period covered by [10].

In S6/VSR2/VSR3 a population of 26 SGRBs were analysed, with one apparently significant candidate associated with GRB 100328A [13]. This had a false alarm probability (FAP) of only 1 %. However, further analysis identified a noise transient in the H1 detector as being responsible for the GW trigger.

The lack of detection was not surprising given the sensitivity of the initial detectors – tens of Mpc for binary merger signals – and the typical distances to GRBs – a median redshift of 0.5 and a closest measured redshift of 0.1, implying a distance of 500 Mpc.

3.2.2 GRB 051103

The SGRB 051103 was observed at 09:25:42 UTC on 3 November 2005 [106]. It was localised by the IPN to a region of the sky that overlapped with the outer spiral arms of the galaxy M81 [88], which is some 3.6 Mpc away.

The possible proximity to Earth motivated a specific, focused search for a GW signal associated with it [12]. This involved analysis of data from the L1 and H2 detectors, with a most significant CBC candidate event in the on-source with FAP of 76 %. If a BNS progenitor was assumed, and the gamma-ray emission was beamed into a jet with 30° opening angle, it could be excluded from being in M81 with 98 % confidence. Similarly, if an NSBH progenitor was assumed with a 30° opening angle, it could be excluded with >99 %.

This confident exclusion of a CBC source in M81 is consistent with the hypothesis that GRB 051103 was actually a soft gamma repeater (SGR) [163].

3.2.3 GRB 070201

The SGRB 070201 was observed at 15:23:10 UTC on 1 February 2007 [89]. It was localised by the IPN to a region including a strip through the spiral arms of M31 (the Andromeda galaxy) [105]. This is the nearest large spiral galaxy to the Milky Way, at only approximately 780 kpc.

Analysis was performed on data from the H1 and H2 detectors and no candidate events were found [29]. Therefore, a BNS or NSBH source in M31 was excluded at >99 % confidence, and a BNS source within 3.5 Mpc was excluded at 90 % confidence (with no assumptions on beaming).

This was entirely consistent with the estimated energetics of the event, which would have been only $\sim 10^{45}$ erg if in M31 at 780 kpc, similar in scale to an SGR. If GRB 070201 was of typical SGRB energy ($\sim 10^{48}$ – 10^{52} erg) then it would have had to have been at a much greater distance than the exclusion distance of the search ($\gg 20$ Mpc) [29]. Again, as with GRB 051103, these results are consistent with the hypothesis that GRB 070201 was in fact an SGR [163].

3.3 Search for InterPlanetary Network Gamma-ray Bursts

During S5/VSR1 and S6/VSR2/VSR3, there were 27 additional, well-localised SGRBs detected by the IPN that did not feature in the analyses described in Section 3.2, but for which there was science-quality GW data. We performed a separate analysis of these as part of an analysis of 223 IPN GRBs.

3.3.1 GRB Sample

At the time of this analysis, nine spacecraft contributed data as part of the IPN: *Wind*, *Mars Odyssey*, *MESSENGER*, *INTEGRAL*, *RHESSI*, *Swift*, *Suzaku*, *AGILE*, and *Fermi*. The sky coordinates of each GRB were determined by comparing the relative arrival times of the signal at multiple spacecraft. The precision afforded by this approach is inversely proportional to the spacecraft separations, among other factors, meaning the localisation accuracy of a network with a baseline of thousands of light-seconds can be equal to or greater than that of any other technique. The process for constructing the full sky error boxes for these GRBs is described in [103] and, more specifically for the purposes of this search, in [144]. The light curves, energy spectra, and localisations of all the bursts in our sample were examined to eliminate the possibility of contamination by magnetar bursts or solar flares. None of these events have been followed up by X-ray or optical telescopes, so no information is available on afterglows or possible host galaxies and associated redshifts.

Only the SGRBs that occurred when two or more GW detectors were taking science-quality data are included in the final sample. This reduces the number in our sample from over 600 to 223, which includes both long and short GRBs.

It is not possible to rely on a single value of the t_{90} duration statistic to classify these GRBs as either short or long, since multiple satellites observed each GRB and each may measure a different value of the t_{90} due to their different sensitive energy bands. Wherever possible we have used the classification provided by [139], based on observations by *Konus-Wind*. We note that the set of SGRBs observed by *Konus-Wind* is split into two types: I – likely merger scenario, and II – collapsar. Only those classified as type I are therefore analysed as possible CBC events. For those not observed by *Konus-Wind*, the t_{90} measured by *Suzaku* was used. If this was also unavailable, an estimate was made by studying the light curves from another mission with good sensitivity, such as *Swift* or *INTEGRAL*. In these cases, and GRB with t_{90} ascertained to be less than 2 s was classified as short.

From the analysable sample of 223 GRBs, 27 are thus classified as short and analysed with our search. The full list of these and their parameters can be found in Table 3.1.

Table 3.1: The SGRB sample – 17 GRBs ‘well-localised’ ($\lesssim 200$ deg², non-H1H2); 10 H1H2-only SGRBs. These results show the 90 % exclusion distances for both possible progenitor models, either BNS or NSBH for a jet opening angle of 30°.

GRB Name	UTC Time	IPN Satellites ^a	GW Network	90 % Exclusion Distances (Mpc)	
				BNS	NSBH
051111B	07:47:51	K/MO/Sw	H1H2	3.4	5.8
051127A	22:54:30	K/MO/Sw	H1H2	8.1	14.6
060103A	08:42:17	MO/I	H1H2L1	5.2	9.2
060203B	07:28:58	K/MO/H	H1H2L1	5.7	10.3
060306C	15:22:38	K/H/I	H1H2	10.0	17.7
060415B	18:14:44	K/MO/S	H1H2L1	13.2	23.7
060522C	10:10:19	K/MO/S	H1H2L1	26.0	44.1
060601A	07:55:40	I/S	H1H2	4.3	6.4
060708B	04:30:38	K/MO/H	H1H2L1	17.1	30.3
061006B	08:43:34	MO/K	H1H2L1	26.6	47.3
061201B	08:11:29	K/Sw	H1H2	15.5	27.1
070113A	11:56:23	K/I/S	H1H2	1.5	3.1
070129B	22:09:26	K/S	H1H2	4.9	7.9
070222A	07:31:56	K/MO	H1H2	6.7	11.9
070321A	18:52:15	K/MO/I	H1H2	20.1	36.1
070413A	20:37:55	I/S	H1H2	7.1	12.5
070414A	17:19:52	S/M	H1H2L1	24.4	45.3
070516A	20:41:25	K/M	H1H2L1	17.6	30.7
070614A	05:05:09	K/H	H1H2L1V1	17.0	29.1
070910A	17:33:29	K/S	H1H2L1V1	11.0	22.6
070915A	08:34:48	K/I/M/Sw	H1H2L1V1	17.6	31.5
070927A	16:27:55	I/M/Sw	L1V1	1.7	2.8
090721A	05:59:21	K/I/Sw	H1L1	11.9	20.0
091114A	03:07:49	K/I/S	L1V1	7.7	14.1
100826B	19:06:36	K/Sw/M	H1L1V1	30.0	52.9
100827A	10:55:49	K/S/Fermi	H1L1V1	12.0	21.7
101009A	06:54:18	K/M/MO/I	H1L1V1	18.5	34.2

^a The detecting satellites: S - *Suzaku*, Sw - *Swift*, I - *INTEGRAL*, M - *MESSENGER*, MO - *Mars Odyssey*, K - *Konus-Wind*, H - *HESSI (RHESSI)*.




Figure 3.1: Cumulative distribution of p -values from the analysis of 27 IPN SGRBs for evidence of a binary merger GW signal. The expected distribution under the no-signal hypothesis is indicated by the dashed line. For those SGRBs with no event in the on-source window, we provide upper bounds on the p -value equal to 1.

3.3.2 Results

For each of the 27 SGRBs we estimate the FAP, or p -value, of the most significant on-source candidate event. The distribution of observed p -values is shown in Fig. 3.1. No significant candidates were found. For a number of GRBs, particularly those observed by the two co-located detectors at LHO, the search yields no candidate gravitational wave events after background rejection cuts. For these GRBs we cannot quote an exact p -value, and instead provide a range bounded below by the fraction of all trials with an event, and above by 1. The result of the weighted binomial population detection test yields a background probability of $\approx 98\%$, strongly favouring the no-signal hypothesis. In conclusion, no noteworthy individual events were found by this search, nor is there evidence for a collective population of weak GW signals.

Given that no events were found in the analysis, we place limits on GW emission associated with each GRB. For a given signal type (BNS or NSBH) the search is sensitive out to a certain distance, which depends on the sensitivity of the detectors at the time of the search and in the direction of the GRB. We may therefore quote 90% confidence lower limit on the distance to the SGRB progenitor, assuming it was a binary merger event of




Figure 3.2: Histograms across the sample of IPN SGRBs of the distance exclusions at the 90 % confidence level for both BNS and NSBH systems.

some stated type. This is the distance at which we recover 90 % of simulated signals with greater significance than any on-source event. The quoted values are marginalised over systematic error inherent in the analysis: mismatches between a true GW signal and the waveforms used in the simulated signals [16]; uncertainties in the calibration of the GW detectors [11].

For both BNS and NSBH signal types, we assume that the prompt gamma-ray emission is collimated along the total angular momentum axis of the binary within a jet of opening angle $\leq 30^\circ$, since SGRB jets are not thought to exceed this angular size [83, 92]. The median exclusion distance for BNSs is 12 Mpc, and for NSBHs is 22 Mpc. A histogram of their values is shown in Fig. 3.2. The neutron star masses are chosen from a Gaussian distribution centred at $1.4 M_\odot$ [112, 133] with a width of $0.2 M_\odot$ for the BNS case, and a broader spread of $0.4 M_\odot$ for the NSBH systems, to account for larger uncertainties given the lack of observations for such systems. The black hole masses are drawn from a Gaussian distribution with a mean of $10 M_\odot$ and a width of $6 M_\odot$. The black hole mass is restricted such that the total mass of the system is less than $25 M_\odot$. For masses greater than this, the neutron star would be swallowed whole by the black hole, and no massive torus would form to power a GRB [70, 81, 158]. The dimensionless neutron star spins are drawn uniformly from the interval $[0, 0.4]$, and the black hole spins are drawn uniformly from the interval $[0, 0.98)$ with tilt angle $< 60^\circ$.

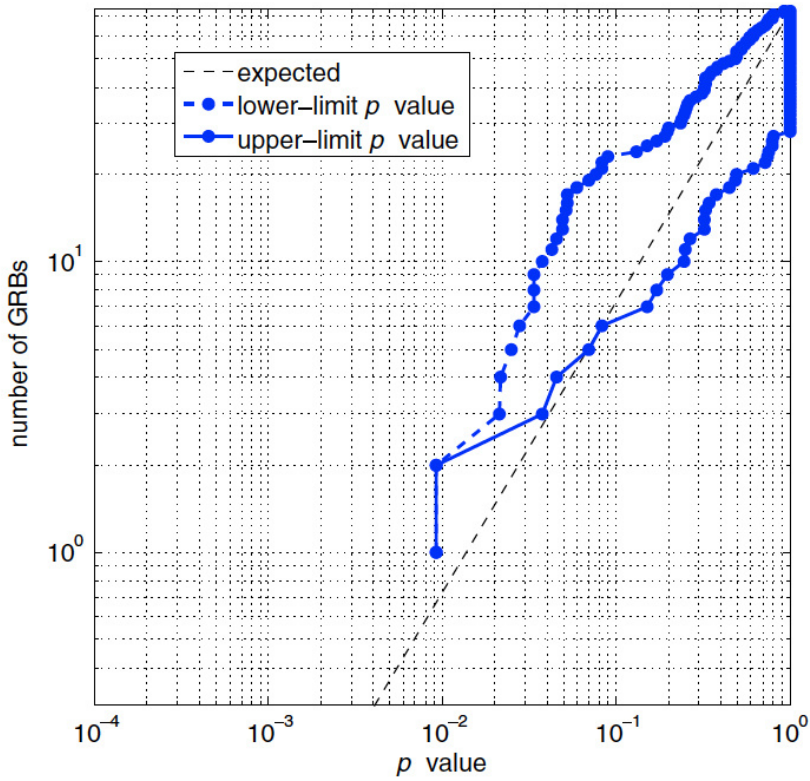


Figure 3.3: Cumulative distribution of p -values for all 69 SGRBs from S5/VSR1 and S6/VSR2/VSR3. The expected distribution under the no-signal hypothesis is indicated by the dashed line. For SGRBs with no event in the on-source window, we provide upper bounds on the p -value of 1.

3.4 Cumulative LIGO and Virgo Results

Here we present the combination of the IPN results from Section 3.3 along with those from the full S5/VSR1 and S6/VSR2/VSR3 searches for coincident SGRB and GW signals [10, 13] (Section 3.2.1). The algorithms used in the S5/VSR1 results were adjusted and reviewed to ensure compatibility with the results of the later analyses.

In total, 69 SGRBs were analysed, with no evidence for a population of weak events. Fig. 3.3 shows the distribution of p -values. The weighted binomial test applied to the full population of SGRBs confirms that the observed distributions are consistent with the null hypothesis (no observed signal).

Next, we use the full sample of SGRBs to place exclusions on the progenitor population. To do this we use a simple population model where all progenitors have the same GW emission, and perform exclusion on cumulative distance distributions. We parameterise the distance distribution with two components: a fraction F of SGRB distributed with a constant comoving density rate up to a luminosity distance R , and a fraction $1 - F$ at effectively infinite distance. This simple model yields a parameterisation of astrophysical GRB distance distribution models that predict a uniform local rate density and a more

complex dependence at redshift >0.1 , as the high-redshift part of the distribution is well beyond the sensitivity of current GW detectors. The exclusion is then performed in the (F, R) plane. For details of this method, see Appendix B of [13].

In Fig. 3.4 we show the exclusion for the BNS and NSBH sources, as well as the redshift distribution of SGRBs as observed by *Swift*. The exclusion line does not come close to the observed population redshift for either source type, indicating that we would have been unlikely to observe an event in these analyses. Indeed, an analysis of all IPN bursts shows that their average redshift is 1.7, and that it detects SGRBs with good efficiency up to a redshift of about 0.45.

We may use these results to extrapolate and predict what might be expected with the aLIGO and AdV detectors. In S5/VSR1 and S6/VSR2/VSR3 there were around 21 months of two (or more) detector duty cycle. Over that period, the detectors' reach varied by approximately a factor of 4, from a 5 Mpc sensitive distance to BNS sources for the H2 detector in early S5, to 20 Mpc for H1 and L1 by the end of S6. Similarly for the advanced detectors, the current scenario calls for around 18 months of science runs of increasing sensitivity during commissioning, before extended running at design sensitivity of approximately 10 times greater than that achieved in S6/VSR2/VSR3.

To approximate the expected advanced detector results, we scale the exclusion distances obtained here by a factor of ten and also increase by a factor of two the number of observed SGRBs to account for the increased run time of a few years. These extrapolated curves are also shown in Fig. 3.4. We see that the exclusion curves now compare favourably with the observed redshift distribution. If BNSs are the progenitors of SGRBs, we may expect perhaps one (or fewer) signals associated with an SGRB. However, if NSBHs are the progenitors, we might expect several.

These extrapolations are in broad agreement with those obtained using only the S6/VSR2/VSR3 results [13], the slight disparity coming from a more realistic estimate of the evolution of detector sensitivity. However, in Chapter 5 we will more thoroughly investigate the prospects for joint observations and obtain far more robust estimates.




Figure 3.4: 90 % distance exclusions for 69 analysed SGRBs for both BNS and NSBH models. The exclusion distance is given for this test then extrapolated by a factor of two in number and ten in sensitivity for the advanced detector era expectations. For reference, the red staircase curve shows the cumulative distribution of measured redshifts for *Swift* SGRBs.

4

IMPROVED METHODS

With the realistic prospect of a joint GW–SGRB observation with aLIGO and AdV, we have made a number of changes and improvements to the analysis pipeline since the publication of [98]. These enhancements are critical to optimising the search to take full advantage of the potential for joint observations.

In this chapter we will illustrate the pipeline improvements using example analyses based upon GRB 100928A, which was observed by the *Swift* BAT [63, 117] during S6/VSR2/VSR3. No other *Swift* instrument observed this SGRB as the spacecraft was unable to slew to the sky position of the prompt burst due to a Sun observing constraint. It was not detected by *Fermi* or any other gamma ray sensitive instrument.

We have chosen this SGRB for a number of reasons. Virgo and both LIGO detectors were operational and had ample science quality data either side of the SGRB time. Specifically, 5264 s of coherent network data between 01:34:35 and 03:02:19 UTC on 28 September 2010 was available for analysis purposes. Additionally, the BAT localised the burst to a point on the sky ($RA = 223.037^\circ$, $Dec = -28.542^\circ$) where both LIGO detectors were approximately equally sensitive, and where Virgo had good sensitivity. Furthermore, this position was known accurately, with a 90 % confidence radius of only 2.3 arcminutes.

It should be emphasised that the following results are dependent on the data at the time of GRB 100928A, which features a number of very large glitches in all three detectors. In particular, two glitches in LIGO Livingston Observatory have $SNR > 400$ and therefore dominate the coherent SNR background. The effect of the improvements will vary slightly from GRB to GRB, but the example results presented here are representative of what we can expect in general.

In performing the coherent analysis, we search the full space of BNS and NSBH sys-

tems. Specifically, we make use of a bank of template waveforms that cover the space of binaries with non-spinning components, with masses between $1\text{--}25 M_{\odot}$, and a maximum chirp mass (Eq. (1.55)) of $8 M_{\odot}$, as was done in previous searches, e.g. [5]. Binaries outside this range are unlikely to produce electromagnetic emission as they will either be comprised of two black holes or, in the case of NSBH systems, the neutron star will be swallowed whole [142].

When evaluating the sensitivity of the search, we perform simulations of BNS systems which are added to the data prior to the analysis. We make use of SpinTaylor waveforms at 3.5 post-Newtonian order [50, 49]. These waveforms approximate the inspiral phase of the binary merger by expanding the equations of motion in terms of v/c (e.g. see terms in integral of Eq. (1.30)). The simulated signals have component masses between $1\text{--}3 M_{\odot}$ drawn from the normal distribution with mean $1.4 M_{\odot}$ and standard deviation $0.2 M_{\odot}$, inclinations drawn uniformly from the intervals $[0^{\circ}, 30^{\circ}]$ and $[150^{\circ}, 180^{\circ}]$, and dimensionless spins ≤ 0.4 . The waveforms were placed uniformly in distance between $2\text{--}45$ Mpc. In total, 2,500 such injected waveforms were used per example analysis.

4.1 Background Estimation

To make a confident detection statement, we must establish that the probability of an observed event being due to noise alone is very small. This requires a detailed understanding of the search background generated by both Gaussian detector noise and non-stationary transients. We do this by looking at the data around the time of the SGRB. We make the reasonable assumption that the off-source data contains no GW signal originating from the same location on the sky and has, on average, the same statistical properties as the detector network background during the on-source period. Thus, the off-source data provides a means of characterising the background noise in the detector network at the time of an SGRB. We have improved the ability of the pipeline to estimate the significance of rare events by introducing the ability to perform time-shifted analyses, where the data from the different detectors are shifted by several seconds relative to each other and the analysis is repeated. This allows us to measure the background of the search to lower than 1 part in 10^5 , a level that would be required for an unambiguous detection claim [14].

The FAP associated to the on-source event, with re-weighted SNR ρ^* , is the probability of having a more significant event in any randomly chosen 6 s of data. This is calculated by counting the fraction of background trials which have an event with $\rho > \rho^*$,

$$\text{FAP} = \frac{N(\rho > \rho^*)}{N_{\text{BG}}}, \quad (4.1)$$

where N_{BG} denotes the total number of background trials. In the standard approach, we simply split the background into as many 6 s trials as possible, so the number of background trials is given by $N_{\text{BG}} = T_{\text{off}}/T_{\text{on}}$.

The standard analysis makes use of approximately an hour of data around the time of

the GRB, leading to a lower limit on the FAP of around 10^{-3} . For the majority of GRBs, this will be sufficient to demonstrate that there is no candidate GW event associated to a particular GRB. However, when there is an interesting candidate, a FAP of 10^{-3} is not sufficient to warrant a detection claim, and further background trials are required to more accurately evaluate the significance.

What would be an acceptable FAP to support a detection claim? In particle physics, the standard level is a 5σ observation, or 1 in 3 million. Prior to the the first GW observation, it was generally agreed that a similar significance threshold would be required before claiming a discovery. In a CBC search from S6/VSR2/VSR3 [14], a simulated signal was added to the data and recovered with a false alarm rate of 1 in 7000 years, which was deemed sufficient to claim evidence for a detection. Translating this to the SGRB search equates to a FAP of $\sim 3 \times 10^{-6}$ for one of the 50 SGRBs observed each year.

For the first detection of a GW signal, GW150914, there was a measured false alarm rate of $< 6 \times 10^{-7}$ per year, corresponding to a FAP of less than 7.5×10^{-8} [22, 17]. This was equivalent to a minimum significance of 5.3σ . The second detection, GW151226, had the same bounds on false alarm rate, FAP, and significance [21, 17].

Conversely, the candidate event known as LVT151012 had a false alarm rate of 0.37 per year, for a FAP of 0.045 and significance of 1.7σ [17]. As such, this was not deemed a detection.

Alternatively, we might consider the chance of there being an observable signal around the time of a GRB. In Chapter 5 we estimate this to be around 1 % for the second generation detector network operating at design sensitivity. Clearly, a detection candidate would require a FAP much lower than the probability of observing a signal. All arguments point to requiring a minimum of approximately 10^5 background trials to assess the significance of a detection candidate, with ideally more than 3×10^5 trials.

To reach a significance level of better than 10^{-5} , we require further background trials. The most straightforward approach would be to simply extend the off-source analysis to incorporate one week of data. While in principle this is possible, the typical duration of continuous operation for the detectors is on the order of hours. Furthermore, the data quality is known to change between different stretches of data [1, 8], so a week of off-source data may not accurately characterise the data at the time of the GRB. In addition, extending the off-source data to one week would increase the computational cost of the analysis by a factor of several hundred, rendering it impractical to estimate the background promptly. Consequently, an alternative method is required. To obtain an improved estimate of the network background, we instead artificially time shift the data from the different detectors and repeat the analysis. These time shifts are always significantly longer than the light travel time between detectors (~ 10 ms) and the signal auto-correlation time and typical glitch durations (both well under one second), so that GW signals will not appear coherently in the time-shifted analysis.

We are able to increase the number of background trials performed by an order of magnitude, with minimal impact on the computational cost, thereby allowing us to estimate




Figure 4.1: FAP as a function of the re-weighted SNR detection statistic for a search performed for GRB 100928A, using time slides to reach FAP of $<10^{-5}$. The figure shows the background estimated with off-source only (787 trials) plotted in orange Y; the short slide analysis (8,917 trials) plotted in green x; both long and short slides (267,185 trials) plotted in blue +. With short slides alone, we can estimate a significance of 1 part in 10^4 while long and short slides give a background estimate to 1 in 3.7×10^6 . The shaded regions show the 95 % Jeffreys credible interval for each case, which assumes each time slide is a statistically independent trial. For clarity of presentation we have only plotted the 20 loudest trials for each search.

FAPs to around 10^{-4} . This is achieved by time shifting the SNR time series of the individual detectors prior to performing the coherent analysis. In the analysis, the detector data is split into sections, typically of 128 s length, which are match filtered to produce a (complex) SNR time series for each detector. These are then combined according to Eq. (3.9) to calculate the coherent SNR time series. A *short slide* is performed by introducing relative time-shift between the detectors' SNR time series prior to computing the coherent SNR. For the example GRB, we leave the H1 data alone, shift the L1 data by multiples of 6 s and the V1 data by multiples of 12 s. This allows for ten time shifted analyses to be performed. Since calculating the single detector SNR time series is the most computationally costly part of the analysis, short slides have a relatively small computational cost. In Fig. 4.1, we show the improvement in background estimation afforded by the inclusion of the short slides.

We have also implemented *long slides* which involve permuting the data segments prior to analysis. Unfortunately, this does require repeating the analysis, so the computational cost increases linearly with the number of long slides. However, it is possible to perform short slides within each long slide. Thus, we only require around ten long slides in order to achieve a background estimate of 10^{-5} .

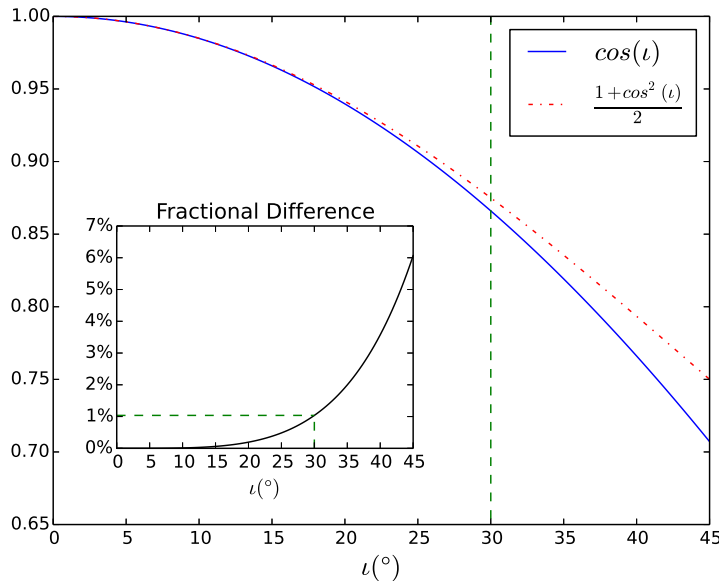


Figure 4.2: Comparison between the amplitudes of h_+ and h_\times as a function of inclination angle ι . Note that even at 30° the difference is only $\sim 1\%$.

Fig. 4.1 shows FAP as a function of re-weighted SNR for the analysis of GRB 100928A. This shows that any on-source event with $\rho_{\text{rw}} > 8.5$ would have a FAP at the 10^{-5} level. We have, however, assumed that all time slides are independent. In reality, all time slides are formed from different combinations of the same detector data streams, and so are not statistically independent at all. A more rigorous treatment of FAP uncertainty when dealing with time slides would likely show far larger 95 % credible intervals for all cases, however it is not clear how to implement such a treatment for this search [174].

4.2 Restrictions on Source Inclination

The search introduced in [98] makes use of the sky location of the source, but places no restrictions on the orientation of the binary. SGRBs are believed to be beamed phenomena [83, 141], with prompt gamma-ray emission concentrated in collimated jets along the axis of angular momentum. These jets are expected to have opening angles not exceeding 30° [83, 92]. Therefore, it may be reasonable to assume that observed SGRB progenitor systems have their orbital angular momenta nearly parallel with the line-of-sight, corresponding to system orbital inclinations $\iota \sim 0$ or $\iota \sim \pi$ with respect to the observer. We can incorporate this into the search by restricting it to binaries which have small inclination angles. This restriction reduces the parameter space of the search, providing an increase in sensitivity.

In Eq. (3.2), we see that the GW amplitudes depend linearly on $\cos \iota$ and $(1 + \cos^2 \iota)/2$. For a binary inclination close to $\iota = 0$, both of these tend towards unity. In Fig. 4.2, we plot both amplitude factors as a function of ι . This serves to highlight the fact that the

amplitudes vary almost identically with ι , up to an angle of 30° , by which time they differ by only $\sim 1\%$. Even at 45° , the two amplitudes differ by only 6% . Consequently for SGRB signals, it is reasonable to treat the amplitude factors as equal and to approximate the signal as left circularly polarised. Similarly, when $\iota \sim 180^\circ$, the two terms agree up to an overall sign and the signal is right circularly polarised.

It is therefore convenient to introduce a single amplitude and phase to describe the signal as

$$\tilde{D} = \frac{D}{\cos \iota} \quad \text{and} \quad \chi_{l,r} = \phi_0 \pm \psi. \quad (4.2)$$

Then, for $\iota \approx 0$, the amplitudes simplify to

$$\mathcal{A}^1 \approx \mathcal{A}^4 \approx -\frac{D_0}{\tilde{D}} \cos 2\chi_l \equiv \mathcal{B}_1, \quad (4.3a)$$

$$\mathcal{A}^2 \approx -\mathcal{A}^3 \approx \frac{D_0}{\tilde{D}} \sin 2\chi_l \equiv \mathcal{B}_2, \quad (4.3b)$$

and similar for $\iota \approx 180^\circ$. As expected, the circularly polarised GW signal is then dependent upon two amplitudes \mathcal{B}_1 and \mathcal{B}_2 (or, equivalently, a single overall amplitude and phase),

$$h_+(t) = \mathcal{B}_1 h_0(t) - \mathcal{B}_2 h_{\pi/2}(t), \quad (4.4a)$$

$$h_\times(t) = \mathcal{B}_2 h_0(t) + \mathcal{B}_1 h_{\pi/2}(t). \quad (4.4b)$$

rather than the original four amplitudes \mathcal{A}^u .

Substituting these expressions into Eq. (3.7), and working in the dominant polarisation, we obtain,

$$\begin{aligned} \ln \Lambda = & \mathcal{B}_1 (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0 + \mathbf{F}_\times \mathbf{h}_{\pi/2}) + \mathcal{B}_2 (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0 + \mathbf{F}_+ \mathbf{h}_{\pi/2}) \\ & - \frac{1}{2} [\mathcal{B}_1^2 + \mathcal{B}_2^2] [(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{F}_\times \mathbf{h}_0 | \mathbf{F}_\times \mathbf{h}_0)] \end{aligned} \quad (4.5)$$

It is straightforward to maximise over the amplitude parameters $\mathcal{B}_{1,2}$ to obtain

$$\rho_{\text{coh}}^2 = \frac{\alpha^2 + \beta^2}{(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{F}_\times \mathbf{h}_0 | \mathbf{F}_\times \mathbf{h}_0)}, \quad (4.6)$$

where

$$\alpha = (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_{\pi/2}), \quad (4.7a)$$

$$\beta = (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0) - (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_{\pi/2}). \quad (4.7b)$$

The calculation proceeds in an analogous manner for $\iota \sim 180^\circ$, with the signal now right, rather than left, polarised. After maximisation, the coherent SNR takes the same

form as Eq. (4.6), but with

$$\alpha = (\mathbf{s}|\mathbf{F}_+\mathbf{h}_0) - (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_{\pi/2}), \quad (4.8a)$$

$$\beta = (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_0) + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\pi/2}). \quad (4.8b)$$

The motivation for performing the search for only circularly polarised waveforms is to further reduce the noise background and thereby increase the sensitivity of the search. Additionally, restricting to circularly polarised waveforms provides us with an additional *null stream* that can be used to reject noise glitches. Prior to assessing the improvement in real data, it is useful to evaluate the expected benefit in Gaussian noise. The original search has four free amplitude parameters \mathcal{A}^μ , and the coherent SNR in the absence of a signal is χ^2 distributed with four degrees of freedom. When restricting to circular polarisation, there are two free parameters \mathcal{B}_μ and the coherent SNR in Gaussian noise will be χ^2 distributed with two degrees of freedom. In other words, this restriction places a strict relationship between h_+ and h_\times — one lags the other by 45° — whereas before there was an arbitrary relationship between the two. Thus, if we know the amplitude and phase of h_+ , there are only two possible configurations for h_\times corresponding to left and right circular polarisation. We must therefore search over both left and right circularly polarised signals, which leads to a doubling of the number of trials.¹ Comparison of these distributions, for a large number of trials, suggests restricting to circular polarisation should result in at a decrease in FAP of around one order of magnitude at fixed SNR, or an increase in sensitivity at fixed FAP of roughly 5%.

In Fig. 4.3 we plot the FAP as a function of SNR for the circularly polarised and un-restricted searches. Over a broad range of SNRs we observe a reduction in the background of a factor of three, corresponding to an increase in sensitivity of around 3% at a given FAP. This improvement is less significant than might have been expected in Gaussian data, and may either be due to the non-Gaussian features in the data or simply a statistical fluctuation observed in this analysis.

Interestingly, we have noticed that the most significant background triggers in the circular search do not correspond to outliers in the un-restricted search. This is likely due to how the pipeline selects triggers. It first applies a clustering method to choose the trigger with the largest coherent SNR in a given time window, before applying signal consistency tests to the trigger which may lead to it being discarded or the SNR re-weighted. Consequently, it is possible that loud events in the un-restricted search do not survive in the circular analysis, and vice versa.

We have demonstrated that restricting to circularly polarised signals can provide a small improvement in the search sensitivity and, furthermore, that it is a reasonable ap-

¹Left and right circular signals will appear identical to a network sensitive only to a single polarisation, but appear exactly orthogonal to a network equally sensitive to both polarisations. For other configurations there is some degree of overlap between the two cases. For most sky locations and networks, they will therefore not be orthogonal, and the two trials are *not* independent leading to a further reduction in the expected background.




Figure 4.3: The background significance against detection statistic for a search performed for GRB 100928A. In red \times , we plot the background calculated using the circular polarisation restriction and in blue $+$ we plot the background from the un-restricted search. In both cases, we perform time shifts of the data as discussed in Section 4.1. Over a broad range of SNR values, the circular polarisation restriction reduces the background by a factor of three. Equivalently, the required SNR to achieve a given FAP is reduced by about 0.25, equating to a 3% increase in the distance sensitivity of the search. For clarity of presentation we have only plotted the loudest 50 trials for each search.

proximation given our current understanding of GRB beaming. We note that a 3% improvement in distance reach corresponds to a 10% increase in the rate of observable signals.

4.3 Searching Large Areas of Sky

SGRBs are localised to sky error boxes of varying sizes by different satellites. This has implications for the targeted GW search following up on these events. For example, the BAT instrument aboard NASA’s *Swift* satellite is capable of localising to 1-4 arcminutes [43], while the typical GW localisation region is several square degrees or larger [2, 79]. Thus, we may follow up a BAT trigger by searching only a single point on the sky since the SGRB localisation is significantly better than the sky resolution of the GW search. However, the GBM aboard NASA’s *Fermi* satellite often localises SGRBs to far larger patches of the sky [125]. The 3σ confidence regions are roughly circular, with a radius of several degrees. Additionally, the IPN localises SGRBs by triangulation with a number of satellites [104]. Depending upon the number of satellites observing the event and their relative positions, the localisations can range from under a square degree to hundreds or even thousands of square degrees. For poorly localised SGRBs observed by *Fermi* or IPN, the

SGRB localisation will be comparable to, or larger than, the typical GW localisation region. Consequently, it is no longer appropriate to treat the SGRB localisation as a single point in the sky, and we must extend the GW search to cover the entire confidence region.

The improved method described in this section has already been used in the searches described in [5, 13].

The targeted, coherent GW search makes use of the sky location in two ways. Firstly, and most importantly, it is the sky location which determines the relative arrival time of a signal at the detectors in the network. These time delays are used to appropriately shift the data prior to coherently combining them in the search. Using the incorrect sky location will cause the signals from different detectors to be mis-aligned in time. Secondly, the detector sensitivities, encoded in the antenna response factors $F_{\{+, \times\}}$, depend upon the location of the source relative to the detector. The use of incorrect $F_{\{+, \times\}}$ will lead to the wrong weighting of detector data streams in the coherent SNR and signal power being present in the null stream.

We can estimate when the single sky point search will not be sufficient. To do so, let us consider only the loss in SNR arising from timing offsets. Following [79], the posterior distribution for a timing offset dt given data s for a matched-filter search is

$$P(dt|s) \propto \exp \left\{ \frac{\rho^2}{2} \left[(h_0|h_0(dt))^2 + (h_{\pi/2}|h_0(dt))^2 \right] \right\}, \quad (4.9)$$

where

$$(h_0|h_0(dt))^2 \approx 1 - 4\pi^2 (dt)^2 \bar{f}^2, \quad (4.10a)$$

$$(h_{\pi/2}|h_0(dt))^2 \approx 2\pi \bar{f} dt, \quad (4.10b)$$

and

$$\bar{f}^n \equiv 4 \int_0^\infty f^n \frac{|\tilde{h}(f)|^2}{S(f)} df. \quad (4.11)$$

We therefore obtain

$$P(dt|s) \propto \exp \left\{ -2(\rho\pi\sigma_f)^2 dt^2 \right\}, \quad (4.12)$$

where ρ is the SNR, and $\sigma_f = \bar{f}^2 - f^2$ is the signal bandwidth (typically around 100 Hz for a binary merger signal). As a result, the recovered SNR in a detector falls off as

$$\rho(dt)^2 \approx \rho_o^2 [1 - (2\pi\sigma_f)]^2 dt^2, \quad (4.13)$$

Thus, a timing offset of $\delta t = 0.5$ ms will lead to a 5% loss in SNR in a single detector.

Given a network of N detectors, $D_{\{1, \dots, N\}}$, let \mathbf{r}_i denote the location of the detector and t_i be the arrival time of the GW signal at detector i from a SGRB at the central location of the sky patch. The distance between two detectors is

$$d_{ij} = \|\mathbf{r}_j - \mathbf{r}_i\|, \quad (4.14)$$

and the light travel time between them is

$$T_{ij} = d_{ij}/c. \quad (4.15)$$

The difference in the arrival time of the signal at two detectors, τ_{ij} , is calculated as [145],

$$\tau_{ij} = t_i - t_j = \frac{1}{c} (\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{w} \equiv T_{ij} \cos \alpha, \quad (4.16)$$

where \mathbf{w} is the unit wave vector describing the direction of propagation of the source, and α is the angle between the line connecting the detectors and the direction to the source.

It is then straightforward to calculate the change in time delay with a change in the angle α as

$$\delta\tau_{ij} = \sqrt{T_{ij}^2 - \tau_{ij}^2} \delta\alpha. \quad (4.17)$$

So, for a source lying on the line connecting the two detectors, the time delay τ_{ij} between detectors is maximal and changes only quadratically with the change in the location of the source. In contrast, for a source which lies on the zero time delay plane, $\tau_{ij} = 0$, a change in location will induce the largest time offset.

Once we select the maximum time offset δt that we are willing to tolerate, it is straightforward to calculate the required angular spacing of the sky points as

$$\delta\alpha = \min_{i,j} \left[\frac{2\delta t}{\sqrt{T_{ij}^2 - \tau_{ij}^2}} \right]. \quad (4.18)$$

Here, the factor of two arises because δt is the largest single detector time offset. We typically choose $\delta t = 0.5$ ms. The two LIGO detectors are separated by a light travel time of 10 ms, while LIGO and Virgo are separated by around 25 ms, which sets the angular scale to around 2° for the LIGO detectors and 1° between LIGO and Virgo. In practice, the resolution is usually determined by the detector pair (D_i, D_j) for which the SGRB target location has smallest relative arrival time difference.

The circular grid is generated by placing rings of points spaced by $\delta\alpha$, starting at the centre, with the final ring passing the 3σ confidence radius. An example of such a grid is shown in Fig. 4.4 (full grid). Each ring will have $2\pi n/\delta\alpha$ points, where $n = 0$ labels the central point and increases as we move outwards. The method of covering the patch is based upon the one introduced in [6]. In the analysis, each point in the grid is treated independently, with the single-detector data streams time shifted appropriately for the given sky location. The coherent SNR and signal consistency tests are calculated with the appropriate detector responses, F_+ and F_\times , for that sky point. As with the background estimation, searching over points in the sky patch is performed *after* the computationally dominant step of calculating the single detector SNR time series. Consequently, SGRBs observed by *Fermi* GBM, requiring around hundred sky points, are processed in approximately double the time required for the *Swift* SGRBs with a single sky point.




Figure 4.4: An example patch of search sky points projected onto the celestial sphere. The blue filled circles show the full grid, while the empty circles are those few points that map to unique differences in signal arrival time between LIGO’s Hanford and Livingston detectors. The parsed points do not form a straight line, but this is simply due to an artefact of the parsing routine and has no effect on the grid reduction.

To demonstrate the efficacy of searching over a sky patch, we repeated the analysis of GRB 100928A, but used a typical *Fermi* GBM 3σ localisation uncertainty radius of 15° [125], with the centre of the *Fermi* patch offset by a few degrees from the *Swift* location. The sky patch for the search contained 178 search points in total (parsed to 41 points for a 2-site, HL network). When performing simulations, the location of each source was chosen randomly from a normal distribution with width 5° , i.e. $\sim 99\%$ of simulated signals were within the 15° radius 3σ localisation region. As previously, we use a search which covers the full BNS and NSBH parameter space, but use only BNS signals when performing simulations.

In Fig. 4.5, we show the search background as a function of detection statistic for both point and patch searches. The background from searching over the sky patch is seen to be about a factor of 20 higher than for a single point. We expect an increase in the background as we have increased the number of trials by searching over the sky patch, however since signals from neighbouring sky points are correlated, the factor of increase is expected to be smaller than the total number of sky points. We have not strictly measured the degree to which these search points are independent. At larger SNR there is a slight deficit of events in the point search, but this is consistent with the statistical uncertainties. While the difference appears significant, it remains statistically consistent with the measured background of the original point. As a consequence, the loudest background event for the patch search has a re-weighted SNR value of 8.33 compared to 7.44 for the original single




Figure 4.5: The background significance against detection statistic for a search performed for GRB 100928A. In red +, we plot the background measured for a single point in the direction of GRB 100928A. In blue \times , we show the background for a sky patch of radius 15° (178 points), encompassing the location of the GRB. In green Y, we show the background for a different single-point search. The point was chosen as it contributed two of the ten loudest events in the patch search. For SNRs between 6.5–7.5, the background of the patch is around a factor of 20 above the single point searches. The increase is expected as we are searching a large number of points, but they are not all independent. At low SNR the increase is smaller, due to clustering effects in the analysis. At larger SNR, the variations between the different analyses are all consistent with statistical fluctuations.

point.

In Fig. 4.6, we show the search efficiency as a function of distance for three different searches: a single point search with simulations spread over the 0.036° *Swift* BAT sky patch; a single point search with simulations spread over a typical 15° *Fermi* GBM sky patch; and a grid of points covering the GBM sky patch with simulations spread over the patch. In all cases the efficiency is calculated at the SNR of the loudest background event in the *short slide* analysis. If we perform the search using only a point at the centre of the *Fermi* localisation region, the results are poor: across the whole range of distances, the search efficiency is never greater than 40%, even for nearby signals which have large SNRs. The reason for this lies in the signal consistency tests discussed in Section 3.1. At the incorrect sky location, the signal does not match the template due to inevitable time offsets between them and the signal will be recovered with a different phase in each of the detectors. Consequently, the coherent SNR will not correctly reflect the total signal power and this will lead to increased values of the signal consistency tests. At all SNRs,




Figure 4.6: The fraction of artificially injected binary neutron star signals found louder than the loudest background event as a function of injected distance. The three curves represent three observational scenarios for a three detector network comprised of Virgo and both LIGO interferometers. In the scenario mimicking a BAT SGRB (black solid line, error radius = 0.036°) the pipeline searches a single point on the sky and finds 90% of signals within 20 Mpc. In the two scenarios mimicking a GBM SGRB we see that by searching over a patch of points covering the large error box of 15° radius (red dashed line) the pipeline performs nearly as well as for the BAT SGRB for signals below 15 Mpc. This is in stark contrast to the previous treatment for GBM-like SGRBs (blue dotted line), which searched a single point at the centre of the error box resulting in very poor rates of injection recovery. In this example, recovery does not quite reach 100 % in the search over the patch because of poor data quality. Since the average injection lies further from its nearest search point, those overlapping times of poor data quality are more likely to be rejected by the signal consistency tests. For GRBs with better data quality, and perhaps more significantly with better data quality monitoring, this effect will be all but eradicated. The increased number of trials resulting from multiple sky points leads to a tail of background events louder than any seen of the BAT single point search, reducing the overall sensitivity of the patch search.

this will lead to a down-weighting of signals due to increased χ^2 and null SNR values. Furthermore, at high SNR the power in the null stream will be sufficient to cause the trigger to be rejected outright due to the null stream cut Eq. (3.12). This explains the, somewhat counter-intuitive, result that the search efficiency actually decreases at small distances.

The sensitivity of the search over the *Fermi* error region is almost the same as the search over just the *Swift* point at small distances, but decreases more rapidly for quieter signals at larger distances. For example, the distance at which we achieve 50 % efficiency is reduced by 10 %. This loss in sensitivity can be attributed to the fact that the background of the *Fermi* search is increased due to the necessity of searching over the sky patch. The reduction in sensitivity is consistent with the 10 % increase in the SNR of the loudest background event.

This method of placing a grid of points in the sky has already been used in the analysis of *Fermi*-detected SGRBs during S6/VSR2/VSR3. An analogous method was used to perform the search over the irregular sky patches produced by the IPN [144].

4.3.1 Two-site Time Delay Degeneracy

In the case of a two-site detector network, for example the LIGO-only network, the ability to resolve independent sky locations is vastly reduced. With a single baseline between sites, multiple sky locations will map to the same difference in signal arrival time. Thus, when moving across the sky patch, there will be one direction where only the antenna response factors $F_{\{+, \times\}}$ change, and not the time delays, while in the orthogonal direction both will change. With two detectors, after maximising over the \mathcal{A}^u , the values of $F_{\{+, \times\}}$ drop out of the coherent SNR expression. This is not immediately obvious, but can be understood by noting that for a two detector search, there are four degrees of freedom in both the coincident and coherent searches. Therefore, *any* observed amplitude and phase in the two detectors is consistent with an astrophysical signal; there is no null stream. Then, the size of the sky grids can be significantly reduced, to represent only those sky locations that map to unique time-delays between observatory sites. Fig. 4.4 shows an example result of parsing the circular sky maps to remove degeneracies in time-delay. For the map shown, only 20% of the points are required to uniquely span the allowed time-delays between the LIGO sites, allowing a reduction in cost in the analysis for two-site SGRB analyses.

Unfortunately, once we restrict to circularly polarised signals, as described in Section 4.2, the restriction to a single time-delay line is no longer appropriate. Now, there are only two free signal amplitudes, which cannot match arbitrary amplitude and phase measurements in the two detectors. Thus the detector response functions again enter into the construction of the coherent SNR, and there is again a null stream.

In Fig. 4.7, we show the sensitivity of the search performed using only the two LIGO detectors in Hanford and Livingston and incorporating an inclination restriction. As before, we plot the *Swift* search results – where both the simulated signals and search are

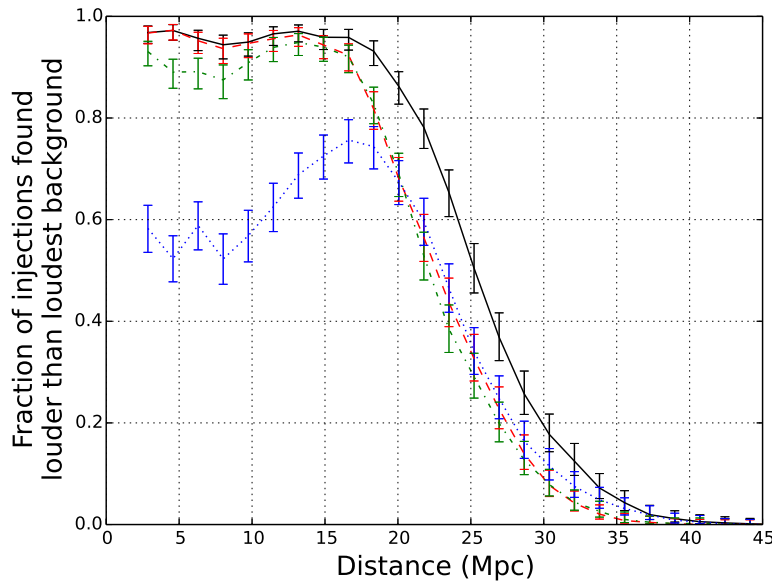


Figure 4.7: The fraction of artificially injected binary neutron star signals found louder than the loudest background event using only the LIGO Hanford and Livingston detectors, plotted as a function of injected distance. As in Fig. 4.6, we plot a scenario mimicking a BAT SGRB (black solid line, error radius = 0.036°) where the pipeline searches a single point on the sky. In this case, the pipeline finds 90 % of signals within 18 Mpc. In the scenario where a GBM SGRB with error box of 15° radius is searched at a single point (blue dotted line), we see poor signal recovery performance at small distances due to signal consistency effects, similar to the three detector case. The difference between the full patch of search points (red dashed line) and a set of points covering unique time delays between sites (green dot-dashed line) is noticeable at small distances, with the use of incorrect antenna response factors causing a drop in performance for the parsed patch. Again, the increased number of trials resulting from multiple sky points leads to a tail of background events louder than any seen in the BAT single point search, reducing the overall sensitivity of multiple point searches.

restricted to a single sky point – as a reference. Next we consider the SGRB localised to a typical *Fermi* GBM error region. When searching over the full *Fermi* sky patch, there is again a degradation of the sensitivity due to a tail of loud background events (a maximum SNR of 8.12 compared to 7.25 for the single point search). However, searching a single sky point leads to a dramatic loss of sensitivity, with only 60 % of nearby signals being recovered. By searching over only the one dimensional time-delay space, we recover the majority of this sensitivity, but do observe a small drop in efficiency at low distances due to the use of incorrect antenna response factors.

4.4 Benefits Provided by Targeted Search

It is interesting to compare the background for the SGRB search with the all-sky coincidence search [14]. This will allow us to estimate the sensitivity improvement offered by the targeted, coherent search. For the all-sky search, the background is one event per year at an SNR of 10, decreasing by two orders of magnitude per unit increase in SNR². Interestingly, the background for the targeted, coherent search, as shown in Fig. 4.1, falls off at the same rate. In both cases, this is significantly slower than expected in Gaussian noise, suggesting that both pipelines are affected in a similar way by the non-Gaussian transients in the data. The background for the all-sky coincidence search translates to a FAP of 10^{-3} in six seconds of data at an SNR of 8.2. In comparison, the targeted, coherent search achieves this background at an SNR of 7.3, as seen in Section 4.1. While both of these are re-weighted SNR measurements, and the details of the pipelines differ, the analysis methods have much in common, so it is reasonable to compare the results. Thus, the coherent analysis provides approximately a 13 % reduction in the SNR at a given FAP.

We can use this to estimate the benefit of performing the SGRB search. To do so, we compare against a simple analysis that just examines the results of the all-sky search for triggers within the 6 s on-source window. The comparison of FAPs above shows that the targeted, coherent search would identify a candidate event with a 13 % lower SNR, or equivalently at a 13 % greater distance. In addition, the targeted, coherent search applies lower single detector SNR thresholds of 4, rather than 5.5, and it includes the SNR contribution from all detectors, even if they did not produce a trigger above threshold. For the case of GRB 100928A, a signal near the detection threshold would be unlikely to register as a trigger in the Virgo detector, and the coherent analysis would register about 10 % greater SNR by incorporating the power from Virgo³. With an additional 3 % increase in sensitive distance afforded by the inclination angle restriction in Section 4.2, the targeted, coherent search provides approximately a 25 % increase in distance sensitivity for a well-localised SGRB over a search that simply looks for a coincident GW trigger from the all-sky search. This equates to a doubling of the event rate.

For an SGRB localised to an extended region on the sky, we observe an increase of a factor of 20 in the background that comes from repeating the search over the sky patch. This translates to a reduction in distance sensitivity of around 10 %. Thus, for SGRBs observed by the *Fermi* GBM, the improvement over the all-sky search is around 15 % in distance sensitivity, corresponding to a 50 % increase in the number of observable sources.

Our example analyses show that the improvement is variable on a case-by-case basis, depending on the data analysed and, in practice, may be reduced by large glitches

²This is taken from Figure 3 in [14], which shows a background of around 0.2 events per year at SNR of 10. However, we must also apply a trials factor of six, as described in the paper, to give a background of 1 event per year at this SNR.

³This is consistent with what is seen when we perform the coherent search using only the two LIGO detectors (Fig. 4.7). The distance at which the search achieves 50 % efficiency (for both the single point and sky patch analysis) is ~ 10 % lower than what is achieved with the full LIGO-Virgo network, (Fig. 4.6).

contributing to a loud tail in the background.

III

**SEARCHING WITH
SECOND GENERATION
DETECTORS**

“We are all in the gutter, but some of us are looking at the stars.”

Oscar Wilde

5

PROSPECTS FOR JOINT OBSERVATIONS

In Part I we have discussed why SGRBs are of interest to GW astronomy, and in Part II we have discussed the methods and results of searches for SGRB-related signals with LIGO and Virgo between 2005–2010. In the coming years we hope to make the first joint GW and SGRB observations, but just how many might we expect to make? In this chapter we will carry out a detailed investigation into the expected rate of these observations in the era of the aLIGO and AdV detectors.

5.1 The Short Gamma-ray Burst Rate

In Section 2.2 we have summarised the current and near-future state of SGRB observations. If we are to accurately evaluate the prospects for joint SGRB–GW observations, we must first work out the rate of SGRBs in the local universe, and the proportion of these we might reasonably expect to detect with gamma-ray sensitive detectors.

There have been numerous recent studies that have attempted to estimate the rate of SGRBs, based primarily on redshift measurements by *Swift* [61, 159, 173].¹ Here, we follow [173], who use the observed SGRB populations, and measured redshifts in *Swift*, in order to derive a luminosity function and local rate density for SGRBs.

The energy spectra of SGRBs is modelled, following [42], as a power law decay with exponential cutoff at low energy and a steeper power law at higher frequencies. The parameters used in the Band function are $\alpha_{\text{BAND}} = -0.5$, $\beta_{\text{BAND}} = -2.25$ and $E_{\text{peak}} = 800$ keV.

¹A nice summary of recent rate estimates is provided in Table 4 of [173].

For an SGRB at a given distance/redshift, the peak photon count in a detector can be related to the peak luminosity in a straightforward way [148, 173]. The detection threshold is taken to be 2.5 photons per second in the 15–150 keV band for *Swift* and 2.37 photons per second in the 50–300 keV band for *Fermi*.

The SGRB luminosity function is taken to be a broken power law, with a logarithmic distribution

$$\phi_o(L) = \begin{cases} \left(\frac{L}{L_\star}\right)^{-\alpha_L} & L < L_\star \\ \left(\frac{L}{L_\star}\right)^{-\beta_L} & L > L_\star \end{cases} \quad (5.1)$$

where L is the peak luminosity in the source frame between 1 keV and 10 MeV, and α_L and β_L give the power law decay below and above the break at L_\star .² The other important parameter is the minimum SGRB luminosity, which determines the lower cutoff of the luminosity distribution. This is poorly constrained as only nearby, low luminosity SGRBs would be observable. The minimum luminosity is taken to be $L_{min} = 5 \times 10^{49}$ erg s⁻¹.

The parameters α_L , β_L , L_\star are fitted jointly with the SGRB rate. Best fit values are $\alpha_L = 1$, $\beta_L = 2$ and $L_\star = 2 \times 10^{52}$ erg s⁻¹, with a local SGRB rate of 4.1 Gpc⁻³ yr⁻¹. The SGRB rate evolves with redshift, peaking at $z \approx 1$.

Other works take a similar approach to estimating the rate of SGRBs [61, 159], although the assumptions they make vary. Consequently there is some variation in the rate estimates. Typically they lie in the range 10⁻⁹–10⁻⁸ Mpc⁻³ yr⁻¹, with a median rate around 3 × 10⁻⁹ Mpc⁻³ yr⁻¹. These rates are somewhat lower than earlier estimates based on a smaller sample of SGRBs [93, 128]. For the remainder of this work, we make use of the Band function and luminosity distribution parameters of [173], but allow for a constant rate per comoving volume in the range 1–10 Gpc⁻³ yr⁻¹. We do not include any variation of SGRB rate with redshift as we found it had little impact on the overall results, due to the limited range of the GW detectors.

Given the evidence for a binary merger progenitor for SGRBs, it is interesting to compare the observed and predicted rates of SGRBs and binary mergers. To do so, we must take into account the beaming of the SGRB jet. The evidence for beaming in SGRBs comes primarily from the observation of jet breaks, at which time the material in the jet starts to spread out, leading to a break in the light curve (see Chapter 2 for more discussion on this). The observation of such a break can be used to infer the jet's opening angle [151]. The observation of a jet break in a number of SGRB afterglows (e.g. [83, 92, 141]) has been used to infer opening angles between 3° and 8°. In others, the lack of an observed break has been used to set a lower limit on the beaming angle. In many cases this leads to a limit of only a few degrees. However, GRB 050724 had no observed break after 22 days, leading to an inferred opening angle of at least 20°. See [48] for a recent summary of observations.

The rate of observed SGRBs can be related to the all sky rate of binary mergers via

$$R_{\text{GRB}} = f_\gamma(1 - \cos \theta_j) R_{\text{merger}}, \quad (5.2)$$

²Other papers use a smaller energy band when defining the luminosity, and this has an impact on the value of L_\star , although not on the slopes of the power law components.




Figure 5.1: The rate of binary mergers in the local universe. The figure shows the predicted rates of BNS (upper, grey band) and NSBH (lower, blue band) mergers, taken from [9]. The shaded regions mark the range of reasonable values, while the dashed lines show the best estimate of the rate. We can also infer the rate of SGRB progenitors, given an opening angle, as plotted in red.

where θ_j is the average jet opening angle of the gamma-ray emission, and the factor f_γ encodes the fraction of binary mergers which produce an SGRB. The rate of BNS mergers, inferred from binary pulsar observations and population synthesis modelling, is taken to lie in the range 10^{-5} – 10^{-8} $\text{Mpc}^{-3} \text{ yr}^{-1}$ (see [9] and references therein). To date, no NSBH systems have been observed as binary pulsars. However, the rate can still be predicted through population synthesis modelling constrained by the observations of BNS. This gives 10^{-6} – 6×10^{-10} $\text{Mpc}^{-3} \text{ yr}^{-1}$.

In Fig. 5.1, we compare the observed and predicted rates for SGRBs to those for BNS and NSBH mergers. As has been observed elsewhere, there is a remarkable concordance between the SGRB and BNS rates [93]. Observed beaming angles are compatible with the best guess BNS rate, with a lower rate of BNS mergers requiring larger SGRB opening angles.³ For NSBH, the rates are not in such good agreement. A 5° jet angle requires an NSBH rate at the highest end of the predicted range. A lower NSBH rate would only be consistent with a wider opening angle than has been inferred from observations. Furthermore, it is likely that a reasonable fraction of NSBH mergers will not produce any electromagnetic emission as the neutron star will not be tidally disrupted, instead being swallowed whole. This would leave no material with which to form an accretion disk to power an SGRB jet [84, 142]. Indeed, in [162], the fraction f_γ of NSBH mergers that produce SGRBs is argued to be in the range 0.1–0.3, depending upon black hole mass and spin distributions.

³For this discussion, we have implicitly been assuming that *all* BNS mergers, produce SGRBs, i.e. $f_\gamma = 1$ in Eq. (5.2). There are, however, arguments that only a subset of BNS mergers will produce SGRBs (e.g. [87]).

Thus, based on rate estimates, it seems unlikely that NSBH mergers can account for *all* observed SGRBs.

5.2 Expected Rate of Joint Observations

The sensitivities of aLIGO and AdV are expected to evolve according to the timeline set out in [2]. Using the predicted rates of BNS mergers we may predict expected rates of joint SGRB–GW observations in the near future.

In Chapter 4 we have discussed the benefits afforded by targeting a search at individual GRBs, thereby reducing the search parameter space. In what follows, we will deliberately avoid the question of what will be required for a first GW detection of its type – where a “ 5σ ” observation may well be required [14, 20]. Instead, we will consider a later observation for which we might require a specific false positive rate: i.e. a limit on the fraction of GW observations that are spurious. In that case, the threshold for announcing a detection is tied to the true signal rate. Since neither the SGRB or BNS rates are known with great accuracy, for this discussion we will adopt the “realistic” rates of $10^{-6} \text{ Mpc}^{-3} \text{ yr}^{-1}$ for BNS mergers, and $3 \times 10^{-9} \text{ Mpc}^{-3} \text{ yr}^{-1}$ for SGRBs.

A detailed evaluation of the expected rate of BNS observations is provided in [2]. There, a false rate of one event per century is chosen, corresponding to an SNR of 12 in the advanced detectors. When the aLIGO and AdV are operating at design sensitivity, the expected rate of observed BNS mergers is 20 per year. Thus the threshold corresponds to a false positive rate of 1 in 2000.

To obtain a comparable SNR threshold for our search, we need to evaluate both the expected foreground and background around the time of an SGRB. Using the results of Chapter 4 we estimate a background rate of 1 in 1000 for events with an SNR above 8 in the SGRB search, with the background decreasing by a factor of 100 for a unit increase in SNR:⁴

$$P_{\text{BG}}(\rho > \rho_{\star}) = \begin{cases} 10^{-(5+2[\rho_{\star}-9])} & \rho_{\star} > 6.5 \\ 1 & \rho_{\star} \leq 6.5. \end{cases} \quad (5.3)$$

Next, we must determine the probability of any given SGRB occurring at a low enough redshift that the GW signal will be observable by aLIGO and AdV. The sky and binary orientation averaged sensitivity of the network is 200 Mpc. However, it is natural to assume that the SGRB jet is beamed perpendicular to the plane of the binary’s orbit (see e.g. [143]). The GW signal is also weakly beamed in this direction – the amplitude for a face-on signal is a factor of 1.5 greater than the orientation averaged amplitude.⁵ The GW beaming is

⁴The analysis in Chapter 4 was performed for the initial LIGO and Virgo detectors and, assuming that SGRB emission is beamed and the jet is perpendicular to the plane of the binary, we obtain a background of 1 in 10^3 above an SNR of 8. However, we must include a trials factor since we will require a larger template bank for the advanced detectors [132]. Consequently we (somewhat conservatively) increase the background by a factor of 100 as was done in [2]

⁵The sensitivity of a detector to binary mergers is typically quoted in two different ways: either the range – the sky and orientation averaged sensitivity; or the horizon – the maximal sensitivity to binaries which

rather weak and the amplitude falls off slowly with opening angle. Even with opening angles up to 30° the mean amplitude is only reduced by 5 % from the face-on case [68]. Thus, the nominal sensitivity for SGRB signals in the advanced detector network is 300 Mpc, rather than 200 Mpc for signals of arbitrary orientation. The sensitive distance scales inversely with the SNR threshold, i.e.

$$D_\star = \left(\frac{12}{\rho_\star}\right) 300 \text{ Mpc} . \quad (5.4)$$

There are around 50 SGRBs observed annually (approximately 10 by *Swift* BAT and 45 by *Fermi* GBM, of which several are observed by both instruments). Assuming a local SGRB rate of $3 \times 10^{-9} \text{ Mpc}^{-3} \text{ yr}^{-1}$, we would expect around one event per year to be detected at a distance of 500 Mpc or less, taking into account detector sensitivities, sky coverage, and live times. Thus, the chance of any SGRB occurring within a distance D_\star can be approximated as

$$P_{\text{GRB}}(D < D_\star) \approx \frac{1}{50} \left(\frac{D_\star}{500 \text{ Mpc}}\right)^3 \quad D_\star \lesssim 500 \text{ Mpc} . \quad (5.5)$$

We have ignored the impact of detector sensitivity since, assuming the GRB model from the previous section, the majority of SGRBs within this range would be observed by *Swift* or *Fermi* if they were in the field of view. This is broadly consistent with the observed redshifts from *Swift*, where the smallest of 30 measurements is $z = 0.12$, corresponding to a distance of 550 Mpc. Obviously, this relationship will break down at larger distances where cosmological effects, variation of the intrinsic SGRB rate, and detection efficiencies all become significant.

In the SGRB search, the chance of a noise event giving an SNR above 9.1 is 5×10^{-6} . At this SNR, the sky averaged sensitivity to face-on BNS mergers is 400 Mpc so, from Eq. (5.5), there is a 1 % chance of the GW signal from an SGRB being observable. This gives a false positive rate of 1 in 2000 as desired. Therefore, the observation of an SGRB allows us to lower the threshold in a GW search by 25 % while maintaining a fixed false positive rate. We note that neither the astrophysical rate of BNS or SGRBs nor the noise background of the advanced detectors are known at this time. Nonetheless, the predicted increase in sensitivity of the SGRB search is relatively robust. The observed background for the BNS and SGRB searches is very similar in nature and, in particular, both show the same, rapid rate of falloff at large acSNR. Thus, changes in the required detection confidence will affect both searches in the same way.

Reducing the detection threshold by 25 % will more than double the number of detectable signals. In other words, less than half of the GW signals associated with SGRBs

are directly overhead the detector and face-on. The horizon distance is a factor of 2.26 greater than the range [82]. Here, we are assuming all sources are face-on, but still averaging over sky positions. It turns out that the averaging over orientation and sky give the same factor, so performing just one average increases the sensitivity by $\sqrt{2.26} = 1.51$.

will be detected based on the GW signal alone – it is only with a joint search that makes use of the SGRB observation that these additional signals will be seen.

It is instructive to ask why the detection threshold can be lowered by 25 % for the SGRB search. Firstly, the expected rate of signals is significantly higher in the data around the time of an SGRB. In Eq. (5.5) we gave the probability of there being an observable signal in the 6 s of data around the time of an SGRB, as a function of the sensitive distance. Within the nominal range of 300 Mpc (at $\text{SNR} = 12$), there is a 1 in 250 chance of observing a signal associated with the SGRB. Meanwhile, for an arbitrary 6 s of data, assuming a BNS rate of $10^{-6} \text{ Mpc}^{-3} \text{ yr}^{-1}$, there is a 1 in 150,000 chance of observing a signal associated with a BNS merger. Thus, assuming that BNSs are SGRB progenitors, it is around a thousand times more likely that we observe a signal within the 6 s around an SGRB than in an arbitrary 6 s of data. Secondly, the background is further reduced because searching a small time window makes a fully coherent search feasible [98], and this increases the sensitivity relative to the all-sky search [41]. These factors combine to give the 25 % reduction in threshold that can be achieved by the search.

We will consider three GRB observing scenarios, each corresponding to a different proportion of the sky covered and different sensitivities to SGRBs. These are with only *Swift* observations, only *Fermi* observation, and an all-sky, full sensitivity coverage. While the latter is, of course, somewhat optimistic, it serves to provide an upper bound on the joint observation rate. For *Swift* and *Fermi*, we use the sky coverage and detection thresholds outlined in Section 2.2 and assume an 80 % detector duty cycle for both detectors due to passage through the South Atlantic Anomaly.

Table 5.1: The expected rate of joint GW–SGRB observations in the second generation GW detector observing runs, assuming that the progenitor of every SGRB is an BNS merger.^a

Epoch	Run Duration	BNS Range (Mpc)		Number of BNS detections	Number of GW–SGRB detections		
		LIGO	Virgo		All Sky	Fermi GBM	Swift BAT
2015	3 months	40 – 80	–	0.0004 – 3	$2 \times 10^{-4} - 0.02$	$2 \times 10^{-4} - 0.02$	$3 \times 10^{-5} - 0.003$
2016/17	6 months	80 – 120	20 – 60	0.006 – 20	0.004 – 0.2	0.003 – 0.1	$3 \times 10^{-4} - 0.03$
2017/18	9 months	120 – 170	60 – 85	0.04 – 100	0.02 – 0.8	0.01 – 0.5	$7 \times 10^{-4} - 0.1$
2019+	(per year)	200	65 – 130	0.2 – 200	0.1 – 2	0.07 – 1	0.01 – 0.2
2022+	(per year)	200	130	0.4 – 400	0.2 – 3	0.1 – 2	0.02 – 0.3

^a Sensitivities, run durations, and BNS rates taken from [2]. We assume a fiducial BNS with a neutron star masses of $1.4 M_{\odot}$.

The expected rates of SGRB observations, assuming a BNS progenitor, are given in Table 5.1. For each observing run we quote a range of possible detector sensitivities to take into account the uncertain nature of commissioning and operating the second generation GW detectors [2]. The rate of observed BNS mergers is calculated for merger rates in the range $10^{-5} - 10^{-8} \text{ Mpc}^{-3} \text{ yr}^{-1}$. The range of predicted rates reflects the uncertainty in both the detector sensitivities and the rate of sources. For joint GW–SGRB observations, we take the SGRB rate to lie in the range $10^{-8} - 10^{-9} \text{ Mpc}^{-3} \text{ yr}^{-1}$. As discussed in Section 4.4, we

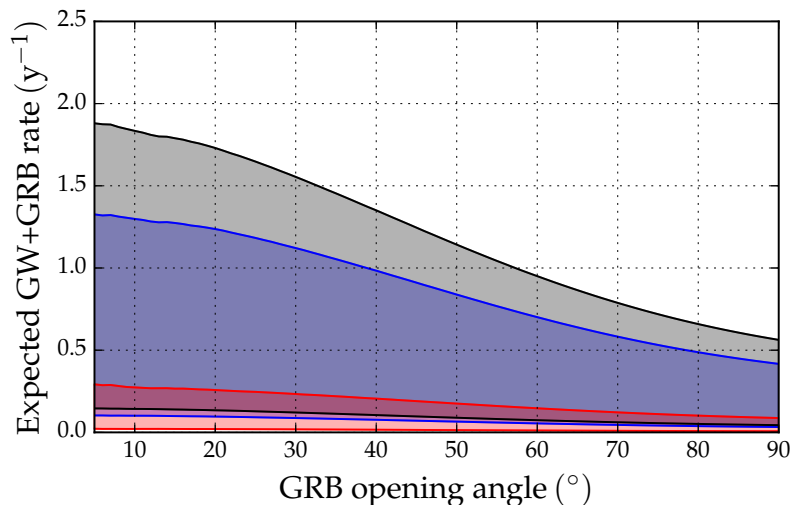


Figure 5.2: The expected rate of observed GW–SGRB signals with aLIGO and AdV operating at design sensitivities. We take the intrinsic SGRB rate to be in the range 10^{-9} – 10^{-8} $\text{Mpc}^{-3} \text{yr}^{-1}$ and assume that BNS mergers are the progenitors of all SGRBs. The grey region shows the range of expected rates with all-sky SGRB coverage. The observed rate increases with a small opening angle as the systems will necessarily have small inclination angles, and thus have the maximum GW emission. The blue region shows the expected rate for joint observations with *Fermi* GBM and the red region for *Swift* BAT. For preferred opening angles ($<30^\circ$) we expect to see at least one SGRB per year in coincidence with *Fermi* GBM.

allow for a 25 % decrease in detection threshold associated with a dedicated SGRB search when compared to an all-sky all-time GW search. When calculating the *Swift* and *Fermi* rates, we use the SGRB luminosity distribution and energy spectra described in Section 5.1. These thresholds, however, have little effect on the rate since the majority of SGRBs within the sensitive range of aLIGO and AdV will have a peak luminosity sufficient to be observed by the BAT and GBM.

The expected number of joint observations in early second generation detector observing runs is much less than one. However, by the 2017/18 observing run there is a real chance of a joint observation. With the network operating at design sensitivity there is an excellent chance of joint GW–SGRB observations during an extended observing run. It is critical, however, to continue monitoring the sky for SGRBs. Only with the sky coverage provided by *Fermi* (and the IPN) might we expect to make joint observations.

Figure 5.2 shows the expected annual rate of joint observations as a function of SGRB opening angle for the 2019+ configuration in Table 5.1. The dependence of the rate on the SGRB opening angle is due to the beaming of the gravitational wave signal; the amplitude for a ‘face-on’ signal is a factor of 1.5 greater than the orientation averaged signal, giving a factor of 3.4 between small opening angles and no beaming. Figure 5.3 shows the expected all-sky BNS merger rate as a function of SGRB opening angle, under the assumption that *all*




Figure 5.3: The expected rate of observed BNS signals when aLIGO and AdV are operating at their design sensitivities. We take the intrinsic SGRB rate to be in the range 10^{-9} – 10^{-8} $\text{Mpc}^{-3} \text{yr}^{-1}$. The rate increases with smaller opening angles as this implies a greater fraction of sources which are not observed as SGRBs. The horizontal lines bound the predicted number of observations based upon estimates of the BNS merger rate. At the largest opening angles, only the higher SGRB rates are consistent with the BNS predictions.

BNS mergers produce gamma-ray emission. As discussed in [55], there is a crossover point where we see more GWs associated with SGRBs than in an all-sky, all-time search. This will clearly depend upon the sky coverage and sensitivity of operational GRB satellites but, assuming full sky coverage, this occurs around 40° . If the beaming angle is larger than this, the SGRB search will detect more signals than the all-sky, all-time search due to the ability to lower thresholds around the time of observed SGRBs. Of course, based on astrophysical measurements of SGRB opening angles, this is unlikely to be the case.

Table 5.2: The expected rate of joint GW–SGRB observations in the second generation GW detector observing runs, assuming that the progenitor of every SGRB is an NSBH merger.^a

Epoch	Run Duration	BNS Range (Mpc)		Number of NSBH detections	Number of GW–SGRB detections		
		LIGO	Virgo		All Sky	Fermi GBM	Swift BAT
2015	3 months	70 – 130	–	0.0001 – 1	3×10^{-4} – 0.06	2×10^{-4} – 0.03	4×10^{-5} – 0.007
2016/17	6 months	130 – 200	30 – 100	0.002 – 10	0.005 – 0.5	0.003 – 0.3	7×10^{-4} – 0.07
2017/18	9 months	200 – 280	100 – 140	0.01 – 40	0.03 – 2	0.02 – 1	0.004 – 0.3
2019+	(per year)	330	110 – 220	0.05 – 100	0.2 – 6	0.1 – 2	0.02 – 0.5
2022+	(per year)	330	220	0.1 – 200	0.4 – 10	0.2 – 3	0.03 – 0.7

^a Sensitivities and run durations taken from [2]. We assume a fiducial NSBH with a neutron star mass of $1.4 M_\odot$ and a black hole mass of $5.0 M_\odot$.

The expected rates of SGRB observations, assuming an NSBH progenitor, are given in

Table 5.2. For NSBH mergers the masses and spins of the system have a stronger effect upon the expected rates of observation. Higher masses and large, aligned spins result in greater GW emission increasing the distance to which the sources can be observed. For simplicity we take the system to be a neutron star of mass $1.4 M_{\odot}$ and a non-spinning black hole of mass $5.0 M_{\odot}$. Following the same procedure as before, we assume that *all* SGRB progenitors are NSBH binaries and use the SGRB model discussed in Section 2.2 to determine the fraction of SGRB signals that are observed by *Swift* and *Fermi*. This has a significant impact on the rate of observable signals, particularly in the epochs after 2019.

As we have discussed previously, there is already a tension between the observed SGRB rate and predicted NSBH rate. Specifically, as is clear from Fig. 5.1, for *all* SGRBs to have an NSBH origin requires a merger rate at the high end of the predicted range, a relatively large SGRB opening angle, or both. Additionally, numerical simulations indicate that for a large fraction of NSBH mergers, there will not be sufficient matter in the accretion disk to power an SGRB, making the rates even less compatible [84]. Thus, the assumption that *all* SGRBs are due to NSBH mergers seems difficult to accommodate, meaning that the highest rates in Table 5.2 are not realistic. Nonetheless, even if 15 % of SGRBs have NSBH progenitors, this would double the expected rate of joint observations. Alternatively, the absence of a joint GW–SGRB observation could be used to limit the fraction of SGRBs that have an NSBH progenitor.

We can compare our results with other recently published works. In [173] the authors calculate the rate of joint GW–SGRB detections by simply assuming a 300 Mpc range for the network featuring aLIGO and AdV. They obtain a rate of joint *Fermi* (*Swift*) observations of 0.4 ± 0.2 (0.06 ± 0.03) assuming a minimum peak luminosity of 5×10^{49} erg s $^{-1}$. This is entirely consistent with the rates for BNS in the 2019+ epoch given in Table 5.1. The fact that they have neglected the directional sensitivity of the GW detector network has little impact because essentially all SGRBs within the aLIGO/AdV range will be observable by *Swift* and *Fermi*. By varying the luminosity threshold they obtain rates that span the same range as ours. For NSBH systems they assume a 1 Gpc range for the second generation GW detectors, as opposed to our range of 660 Mpc. Consequently they obtain a significantly higher rate (5 ± 2 for *Fermi* and 0.7 ± 0.3 for *Swift*). In [148] the authors have also calculated joint detection rates of GW–SGRB signals. They predict rates of joint observations with *Swift* of 0.01–0.5 yr $^{-1}$ for BNS and 0.004–0.16 yr $^{-1}$ for NSBH. The rates are broadly comparable to those presented here, although the range goes somewhat higher for BNS and lower for NSBH. These differences arise due to different choices of parameters in the Band function, SGRB luminosity distribution, and detector thresholds. Additionally, the authors choose a fixed BNS rate of 6×10^{-8} Mpc $^{-3}$ yr $^{-1}$, and a fixed NSBH rate of 3×10^{-9} Mpc $^{-3}$ yr $^{-1}$, as well as a range of opening angles between 5° and 30°. With these rates, NSBH signals could only account for a fraction of SGRBs. This explains why their numbers are lower than the ones in Table 5.2, where we have assumed that *all* SGRBs have NSBH progenitors.




Figure 5.4: The probability of obtaining an event of a given SNR for: noise only (red, dot-dashed line); an SGRB progenitor at an unknown distance (black, dashed line); and a known distance (blue, solid line). In this example, we have used the parameters from GRB 080905A, with a distance of 550 Mpc which gives a signal SNR of 7.7. The top plot shows the probability distribution function, while the bottom plot gives the cumulative probability of observing an event as loud or louder.

5.3 Benefits of Joint Observations

We have already discussed some of the benefits of joint GW–SGRB observations. To summarise again, they will enable us to confirm or rule out the binary merger progenitor model for SGRBs [72]. Assuming this model to be true, measurements of time-delay between the prompt gamma-ray emission and merger time could help to understand the physics of jet breakout. We should gain the ability to probe SGRB jet opening angles [55, 67], and may even be able to provide an independent measurement of distance and redshift for use as a probe of cosmology [130, 156]. Here, we do not discuss all of these in detail. Instead, we focus on two issues. First, we discuss how the measurement of an SGRB redshift may actually assist in the detection of a GW counterpart. Then, we discuss prospects for measuring or constraining opening angles.

5.3.1 Detecting a GRB with measured redshift

The advanced detector network will, on average, be sensitive to a BNS merger associated with an SGRB within a distance of 400 Mpc, or $z \lesssim 0.1$. The closest observed SGRB is GRBs 080905A, which had a measured redshift of $z = 0.12$ [149]. It is interesting to ask whether this SGRB could have been observed by the aLIGO/AdV network. While at a distance of




Figure 5.5: The odds ratio between the signal and noise models. We consider two signal models: an SGRB at an unknown distance (black, dashed line) and an SGRB at a known distance (blue, solid line). In this example, we have used the parameters from GRB 080905A, with a distance of 550 Mpc, which gives an expected SNR of 7.7. The blue curve gives the odds ratio for a BNS signal at that distance, as a function of SNR. The black curve gives the odds ratio for a BNS signal at an unknown distance. At low SNR, knowledge of the distance increases the odds ratio significantly; at higher SNRs it decreases the odds because the observed SNR is no longer consistent with the distance.

550 Mpc, it was at a favourable sky position for the GW detector network. A BNS merger associated with GRB 080905A may have been marginally detectable, but only once the known redshift is incorporated in our calculations.

Let us consider the expected distribution of the observed SNR in the GW search under three distinct scenarios: no observed GW signal; a BNS merger signal associated with an SGRB at an unknown distance; a BNS merger signal at 550 Mpc. To obtain the distribution in the absence of a signal, we simply use the empirical estimate provided in Eq. (5.3). For a signal at 550 Mpc in the direction of GRB 080905A, a BNS merger will generate an expected network SNR of 7.7. The expected, maximum SNR observed in the gravitational wave search then follows a non-central χ^2 with four degrees of freedom [98] overlaid on the noise background given in Eq. (5.3). Finally, for an SGRB with unmeasured redshift, we use the distance distribution as given in Eq. (5.5), i.e. signals distributed uniformly in D^3 at low redshift, with only a small probability of the SGRB occurring within the aLIGO/AdV sensitive range.

In Fig. 5.4 we show the probability distribution for the SNR of the GW event under these three scenarios. The figure shows the probability distribution, as well as the cumulative

probability of observing an event above a given SNR. In this example the knowledge of the distance greatly increases the chance of observing a signal with a moderate SNR. For example, the chance of observing an event with $\text{SNR} > 7.5$ due to noise alone is around 1 %. If there is a BNS merger at unknown distance this rises to 3 %. However, when the distance is known to be 550 Mpc it increases to 60 %.

One way to visualise the benefit of a redshift measurement is through the odds ratio – the ratio of the signal probability to the noise probability. This is plotted in Fig. 5.5. For an observed SNR above 7.5, the signal model is favoured over the noise by a factor of 10, increasing to 100 at SNR of 8. Even at these low SNRs, this would be an interesting event. However, if the distance is not known, a larger SNR (8.5–9) is required before the signal model is strongly favoured over the noise. Thus, if this SGRB had occurred during the second generation detector era, there is a real chance that measuring the redshift would make the difference between identifying a GW candidate or not.

5.3.2 Constraining the jet opening angle

A joint GW–SGRB observation would provide a measurement of the binary’s inclination angle and, consequently, would provide a constraint on the jet opening angle of SGRBs. However, the majority of observed GW signals are likely to be weak, with an $\text{SNR} \leq 10$, and this will make accurate parameter recovery difficult. Accurate measurement of the binary inclination angle is further complicated by the fact that it is highly degenerate with the distance, particularly when the signal is close to face-on. Specifically, the overall amplitude of the two polarisations scale as $(1+\cos^2\iota)/2D$ and $|\cos\iota|/D$ and, at an SNR of 10, we would expect to measure these amplitudes with an accuracy of roughly 10 % [18].

For a face-on signal (with $\iota \approx 0$ or π), the two amplitudes are equal. They differ by 1 % for an inclination angle of 30° and by 10 % for an inclination of 50° . Thus, while the GW observation will constrain opening angles, it is most likely to limit the angle to be $\lesssim 45^\circ$. In the case where the redshift, and hence distance D , is known there will still be a $\sim 10\%$ uncertainty in $\cos\iota$ corresponding to a constraint on the opening angle of $\lesssim 25^\circ$. Even for the loudest signals, we are faced with an uncertainty in the Hubble constant of 1 % and a likely instrumental calibration error of at least a few percent [18], making it difficult to constrain the opening angle to less than 10° .

It is more likely that the observed populations of SGRBs and binary mergers will allow us to place an upper limit on the opening angle of SGRB jets, which is clear if we consider again Fig. 5.1. If the permitted range of BNS merger rates can be reduced from three orders of magnitude to a factor of two, then the SGRB opening angle will be highly constrained.

First, we consider the case where the early observing runs do not yield a GW–SGRB detection. We assume that, in the absence of a detection, the loudest event is consistent with background and estimate the expected upper limit on the rate as

$$R^{\text{ul}} = \frac{2.3}{VT}, \quad (5.6)$$

Table 5.3: The expected bounds on SGRB opening angle during the early aLIGO/AdV observing runs. These results assume that no GW signal is observed and use the observed SGRB rate to infer the minimum jet opening angle consistent with the lack of GW detection.^a

Epoch	Run Duration	BNS Range (Mpc)		Limit on SGRB opening angle (°)	
		aLIGO	AdV	BNS	NSBH
2015	3 months	40 - 80	-	0 - 3	0 - 6
2016/17	6 months	80 - 120	20 - 60	1 - 8	3 - 15
2017/18	9 months	120 - 170	60 - 85	3 - 15	7 - 35

^a We assume that all SGRBs are either BNS (1st column) or NSBH (2nd column). In both cases, the range quoted takes into account both the uncertainty in the detector performance in these runs as well as the uncertainty in the local rate of SGRBs.




Figure 5.6: The expected upper limit on the BNS rate obtained from planned observing runs, assuming no GW detections: 2015 in black; 2016–17 in blue; and 2017–18 in green. The ranges correspond to the uncertainties in detector sensitivity as detailed in Table 5.1. The expected rate of SGRB progenitors as a function of the SGRB opening angle is overlaid in red. Assuming that all SGRBs correspond to BNS, we can read off the lower limit on opening angle that would be obtained at the end of each run.

where V is the volume searched, and T the analysis time [80]. From this we can read off the smallest opening angle consistent with the upper limit on the rate by re-arranging Eq. (5.2) to obtain

$$1 - \cos \theta_j \geq \frac{R_{\text{GRB}}}{f_\gamma R_{\text{BNS}}^{\text{ul}}}. \quad (5.7)$$

Thus the tightest limit on θ_j is given by assuming the maximum BNS rate, i.e. right at the upper limit, and an f_γ of unity, i.e. that all BNS mergers produce SGRBs.

In Fig. 5.6 we plot the expected upper limits in the absence of a GW detection during

the early observing runs. The bands here correspond to the uncertainties in detector sensitivities as given in Table 5.1. For example, at the end of the 2016–17 run, the lack of a detection can place a constraint on the SGRB opening angle between 2° and 8° , depending upon the detector sensitivity achieved and assumed SGRB rate. Thus, even in the absence of an observation, we are starting to impact measurements with GW–SGRB observations.

In Table 5.3 we summarise these results for both BNS and NSBH sources. In both cases, we are assuming that all SGRBs are produced by one particular type of merger. This, of course, is unrealistic. Although we cannot know the fraction of SGRBs that have a BNS or NSBH progenitor, we might reasonably assume that all SGRB progenitors are mergers. Since the sensitivity to BNS mergers is less than for NSBH, the conservative limit comes from assuming that all are due to BNS mergers. Alternatively, it is possible to make reasonable assumptions in our priors for the various parameters, and then marginalise over them to obtain a distribution for the opening angle.

Of course, we hope to observe GWs from CBCs. Even a handful of observations will provide a measurement of the rate within a factor of two, which will correspond to a much tighter horizontal band on Fig. 5.1. If, for example, the rate is $10^{-6} \text{Mpc}^{-3} \text{yr}^{-1}$ then this will restrict the SGRB opening angle to be between about 3° and 8° .

6

PYGRB: THE SHORT GAMMA-RAY BURST SEARCH FOR ADVANCED DETECTORS

We wish to develop a search for SGRBs that can be run shortly after an SGRB alert is received. The search should be flexible, both in terms of the number of detectors used and the amount of time analysed. This will allow us to search for GWs around the times of as many SGRBs as possible.

We want to run the search within a few hours so that we may inform observing partners of any interesting events, but also run it once the final data quality and calibration information is available. For some GRBs, such as those observed by the IPN that require manual processing, we may not be able to run on short timescales. The search must be able to calculate the background down to detection level as well as performing simulations to evaluate the sensitivity to GWs from BNSs and NSBHs, using improved techniques akin to those described in Chapter 4.

This motivates the creation of a flexible search pipeline structure. Flexibility will not only meet these needs, but allow us to meet future needs that are likely to arise, such as rising computational costs as the matched filter template banks grow in size, or a narrowing search focus as emission models are excluded or detections are made. In this chapter we introduce this new pipeline for use during O1 and beyond, which we call PyGRB.¹

¹<http://ligo-cbc.github.io/pycbc/latest/html/workflow/pygrb.html>

6.1 PyCBC

First we must introduce the PyCBC project [65].² The aim of the PyCBC project is to provide a framework, written primarily in the PYTHON programming language,³ for the development of software tools that can be used to search data from GW detectors for CBC signals.

The PyCBC package is designed to make it possible to create custom analysis tools that may be considerably more flexible and optimised than those used in analyses prior to O1 [171].

Past CBC analyses (e.g. [41]) have largely been written in the programming language C [109], and are catalogued in the LIGO Scientific Collaboration Algorithm Library (LAL).⁴ However, the higher level PYTHON language is arguably more user friendly for both creating and using software tools. A large part of its strength is its object-oriented nature. For example, data arrays representing GW data time series may have arithmetic operations performed on them within a single, simple line of code. The same may not be true for equivalent C code. These data arrays may even have pre-defined built-in methods specific to their object class, or inherited from other classes in a hierarchical manner, which can make code development easier.

Furthermore, PYTHON is free, open-source, and cross-platform. Indeed, the collaborative ethos of the PYTHON user community has seen the production of many widely-used and well-maintained free software packages (e.g. NUMPY⁵). These resources provide pre-compiled functions and new classes, some very complex, that developers of new projects such as PyCBC may call upon, a simple and time-saving process. In turn, PyCBC itself can be considered as just such a resource for those looking to write a new CBC analysis tool, since it contains methods and classes that may be used or modified depending on the needs of individual developers.

Operations that require lower level, pre-compiled C binaries to run in a computationally efficient manner (e.g. fast Fourier transforms (FFTs)) can still be called from within PYTHON code via ‘wrapping’ with an interface such as SWIG.⁶ This means that CBC search specific operations that have already been written and tested in LAL, and which would not benefit from being re-written in PYTHON, may be used directly inside PyCBC tools.

Advances in parallel computing architectures have seen a rise in the number of GPUs in supercomputing clusters that are used for non-graphical purposes. For operations like FFTs, GPUs are likely to vastly outperform CPUs. This can be handled from within PYTHON code via packages such as PyCUDA [114], which delegate these operations to GPU cores on computing clusters while the remainder of the analysis is processed on CPU cores. This is a promising avenue for future exploitation in PyCBC analysis pipelines.

²<http://ligo-cbc.github.io/pycbc/latest/html/>

³<http://www.python.org/>

⁴<https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>

⁵<http://www.numpy.org/>

⁶<http://www.swig.org/>

PyCBC, therefore, ought to make it not only possible but relatively straightforward to create a faster, more flexible, automatable SGRB-triggered search than the previous incarnation, one that is ultimately more sensitive, computationally efficient, and easier to develop further to meet future needs. This new version of the search will not only increase the likelihood of making a joint detection, but will also maximise the benefits of multi-messenger SGRB observations.

6.2 The PyGRB Workflow

PyGRB and other PyCBC pipelines are built by *workflow* generation PYTHON scripts. A workflow is an abstract representation of a full analysis. PyCBC workflows are comprised of many interrelated *nodes*, each node representing an analysis job. Associated with each job is an executable or process, and in most cases a set of arguments or options, perhaps including some data inputs and/or outputs. Any two nodes may be independent of one another, or they may have a relationship defined by, for example, one node’s output being another’s input. In PyCBC, this often complex interrelationship is implicitly handled via tracking of data inputs and outputs as PYTHON objects during the workflow generation.

A schematic of the overall structure of the workflow is shown in Fig. 6.1. Data from all available detectors is retrieved and a decision is made on what data are to be analysed. A template bank of waveforms is then used to perform matched filtering on the data, both with and without simulated signals added, and the results are collated.

This overall design is set by the form of the workflow generation script, and is further controlled by workflow options, which are set either in configuration files or given individually on the command line. These can include all manner of details such as which GW detectors to include, the range of time to analyse, where to store results, etc. PyCBC allows configuration files to be stored remotely under version control and downloads them at run time, which makes it easier to impose uniformity between analyses. Likewise, pre-compiled executables may also be downloaded at run time.

Some of the processes that comprise the analysis may also happen at the time of workflow generation. In the case of PyGRB, the standard behaviour is to calculate the availability of science data at this time. It would also be possible to generate other primary analysis inputs at this stage, including injection sets and a template bank of waveforms. However, at time of writing the PyGRB standard is to import a pre-generated template bank from a given URL, and to generate injection sets by submitting these nodes to computing cluster resources as analysis jobs.

Once generated, the workflow must be executed, a process overseen by PEGASUS [66]. PEGASUS transforms the abstract workflow — which at this stage is little more than a list of nodes, their attributes, and dependencies — into a batch of analysis jobs. It intelligently manages the running and data handling tasks for these jobs, and will deal with job failures by retrying parts of the analysis. Progress can be monitored through a number of PEGASUS tools and a web-based dashboard.

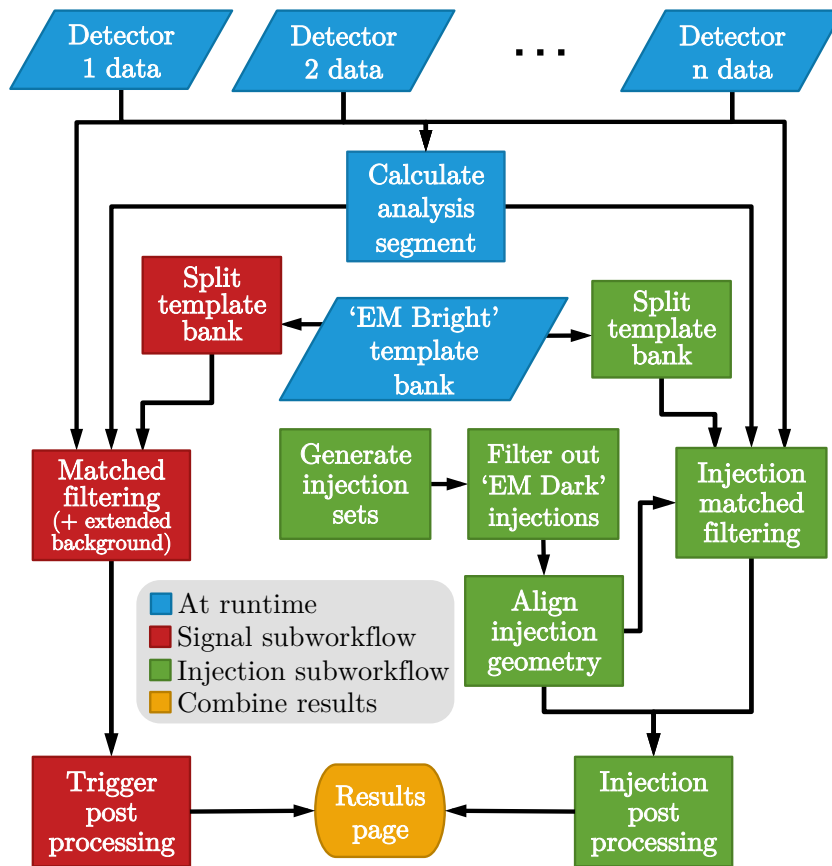


Figure 6.1: A schematic of the standard, offline PyGRB workflow used during O1. The data inputs that exist prior to workflow generation, the EM-bright template bank and GW detector data, are shown as parallelograms. The retrieval of these data and the calculation of the analysis segment are done at time of workflow generation (blue). After this stage the workflow may be thought of as two essentially independent workflows; a standard matched filtering workflow (red) and an injection workflow (green). Once all processing and post processing jobs are complete, the results are collated and a summary web page is generated (gold).

As well as having a new workflow structure, generation procedure, and submission/monitoring tools, the PyGRB analysis differs from the previous version of the SGRB analysis in a number of key aspects.

6.2.1 Increasing the Number of Analysed GRBs with Single Detector Analyses

During LIGO and Virgo runs prior to O1, SGRBs were only analysed if they occurred when at least two detectors were in full science mode. With only one data stream the analysis cannot use time slides to generate extended background statistics as described in Section 4.1. We can increase the number of trials by allowing the search to incorporate longer stretches of data than in a multi-detector search, however even with 24 hours of

unbroken science-quality data split into 6 s trials we would only be able to reach FAP estimates of 6.94×10^{-5} at best. Furthermore, unbroken stretches of this length or more are quite rare, and the noise power spectrum in the aLIGO detectors can change significantly over such a long period of time even if the detector remains in science mode. We also cannot include the inclination restriction from Section 4.2, since there are too few degrees of freedom. The signal based vetoes (Section 3.1.2) will also be less effective because we do not have any coherence conditions that we can place on search triggers. These limitations will inevitably result in a less sensitive search compared with a multi-detector equivalent.

However, despite these limitations, we may still want the analysis to run when only a single detector has enough data around the GRB time. For example, an SGRB candidate may be observed close to or overlapping the disk of a nearby galaxy — as happened for GRBs 051103 and 070201 (see Sections 3.2.2 and 3.2.3) — at a time when only one GW detector is taking science quality data. A lack of candidate signal in that detector might at least allow us to reject the hypothesis of a binary merger within that particular galaxy. This is an astrophysically interesting statement that could be made if we allow for single detector analyses. More generally, we can place an exclusion distance using a single detector for any SGRB. One could even argue that if electromagnetic observations strongly imply a nearby SGRB, and a single GW *does* see a signal candidate, this information would be very scientifically compelling. Indeed, given the ease with which a single detector search can be carried out, there seems little reason not to run this configuration as standard practice.

Being able to run a single detector search is therefore a potentially astrophysically important development that has been implemented in PyGRB.

6.2.2 Selecting the Analysis Period Dynamically

Ground based GW detectors typically operate in a state of stable lock for periods of $O(\text{hours})$. Lock loss can be caused by a great many reasons, from internal instabilities in the instrumentation itself to external influences like earthquakes. This means that for a significant amount of time during an observing run, one or more of the detectors will not be taking science quality data. For example, during the early part of O1, H1 had a duty cycle of 70 %, and L1 55 %, with coincident data 48 % of the time [20]. Typically, a few percent of the total time is further lost to data quality vetoing [19]. As a result, there will not necessarily be any science data at the time of a GRB, and if there is there may still not be enough to run a PyGRB search.

The workflow generator queries a data server for all available data within a window of $\sim 12,000$ s centred on the GRB time. The software must then decide which of these data to analyse. This is calculated based upon the availability of science-quality data, and upon criteria set in the configuration files. These criteria include a minimum required amount of contiguous data, and minimum amounts of data before and after the SGRB time, required for data conditioning. If one or more of these criteria are not met by coincident data, PyGRB can be told to consider running a single detector search via a simple flag.

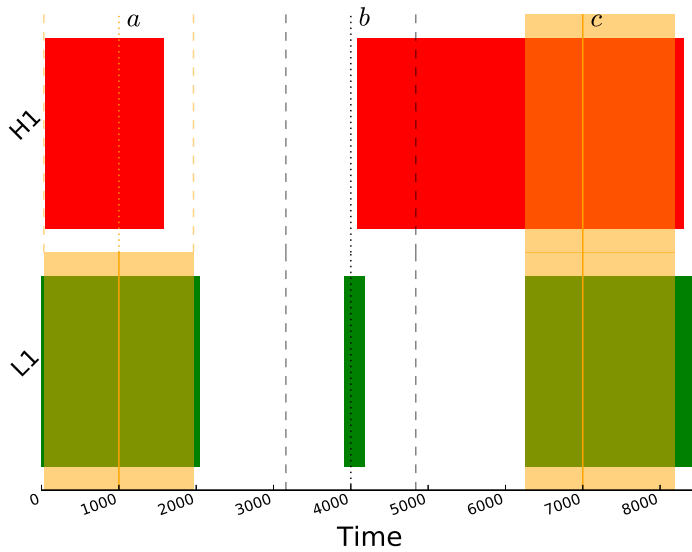


Figure 6.2: Examples of PyGRB calculating which data to analyse. *a*) For the GRB trigger at $t = 1000$, it finds that not enough coincident data is available due to a lack of H1 data (1664 s required), so it generates an analysis for L1 (analysis segment shown in orange). *b*) For the trigger at $t = 4000$, H1 has no data and L1 does not have enough, so no analysis is performed (minimum duration requirement shown in black). *c*) For the trigger at $t = 7000$, a coherent search may be performed. The software will attempt to use as much data as possible within the given constraints. A relative lack of L1 data before the SGRB means the analysis segment (orange) is asymmetric.

The minimum required stretch of data is generally set to 1664 s (roughly half an hour) to ensure there is enough data to accurately estimate the power spectral density (PSD). Furthermore, if an interferometer is only in science mode for a short period of time, it can suggest the instrument is not entirely stable.

Fig. 6.2 shows how the algorithm deals with three example scenarios, each one featuring a restricting condition: *a*) where only the L1 detector has enough data; *b*) where H1 has no data at the SGRB time and L1 does not have enough data; *c*) where a coherent segment meets the minimum requirements but is still limited by L1 data availability, becoming asymmetric about the SGRB time. In all cases the PyGRB workflow generation script reports the reasons for its decision to the user.

This dynamic approach to the analysis segment is a new PyGRB feature, and facilitates automation of the entire analysis by removing the need for a human to check whether there is adequate data before launching a search.

6.2.3 Reducing the Background by Refining the Source Parameter Space

Not all NSBH mergers are expected to emit electromagnetically [84, 138, 162]. In most GRB emission models it is necessary for a central black hole to form with an accreting torus. In an NSBH merger this requires that the neutron star be disrupted before merger, rather than falling into the black hole intact. The tidal force acting across the bulk of the neutron star must overcome the internal forces of the neutron star itself. Many factors influence whether this will happen.

A greater black hole mass decreases the tidal force at the innermost stable circular orbit (ISCO), making it more likely that the neutron star will not be disrupted. A higher black hole spin shrinks the radius of the ISCO, potentially beyond the Roche limit of the neutron star. A ‘softer’ neutron star equation of state leads to a larger neutron star radius, which will increase the tidal force acting across the diameter of the star. A ‘stiffer’ neutron star equation of state therefore makes it more difficult to power an SGRB with an NSBH merger.

The combinations of black hole masses and spins that might give rise to a torus was investigated in [142]. Around half of the NSBH parameter space used in past searches will not lead to the formation of a torus of mass $0.03 M_{\odot}$ or more. This is true for any reasonable model of the neutron star equation of state, resulting in a very conservative cut. As such, these sources can safely be considered *EM-dark* and ignored by the SGRB search (Fig. 6.3).

We do this simply by removing all corresponding template waveforms from the analysis. This reduces both the computational cost and the search background. This has been implemented in PyCBC in the form of an *EM-bright* template bank (Fig. 6.4), which the PyGRB pipeline uses to match filter the data.

The collisions or tidal disruptions associated with the coalescences of BNS and NSBH systems will produce extremely complex GW emission, which we cannot accurately model in our template waveforms. However, the majority of the signal power from these systems is emitted during the inspiral. We therefore populate the template bank with SPINTAYLORT4 waveforms, which are post-Newtonian descriptions of the inspiral phase only.

The template bank is designed to have a maximal loss in SNR of 3% due to discretisation effects. This is ensured for binaries with spins aligned to the orbital angular momentum over the entire valid parameter space. In this valid space, at least one constituent of the binary is assumed to be a neutron star, because no viable mechanisms are currently known that allow for stellar mass BBH mergers to produce significant gamma-ray emission, required to appear as an SGRB. The neutron star can have mass $1 M_{\odot} \leq m_{\text{NS}} \leq 2.8 M_{\odot}$ and dimensionless spin magnitude of ≤ 0.05 , which corresponds to the dimensionless spin of the fastest known pulsar in a double neutron star system (J0737–3039A [53]). For the companion object, we test masses in the range $1 M_{\odot} \leq m_{\text{comp}} \leq 25 M_{\odot}$ and dimensionless spins up to 0.999. In cases where $m_{\text{comp}} > 2.8 M_{\odot}$, we treat the system as an NSBH and therefore apply the EM-bright condition [142]. In total the bank contains approximately




Figure 6.3: The black hole mass–spin parameter space for an NSBH merger demarcated into regions where an SGRB is expected (white), possible (grey), and not expected (black). These calculations assume the softest neutron star equation of state. For sources where the black hole spin is aligned with the orbital axis, the boundaries are shown by the dashed lines. Figure taken from [142].

135,000 templates.

PyGRB also passes NSBH injection sets through the same EM-bright condition so that only potentially EM-bright source waveforms are used to tune the pipeline and test its efficiency at making detections.

6.2.4 Increase Sensitivity with Arbitrary Source Spins and Precession

Past GW–SGRB searches made use of template waveforms appropriate for binaries with non-spinning components. For neutron stars, this is a reasonable approximation as they are expected to have low spins which will not greatly affect the waveform [51]. However, in an NSBH system the black hole spin can have a significant effect on the emitted waveform. The component of the spin aligned with the orbital angular momentum will affect the rate at which the binary inspirals [154], while the orthogonal spin components will lead to precession of the system [39].

It has been shown that using waveforms which incorporate the effects of aligned spins can greatly enhance the sensitivity of a search to NSBH systems [35, 97]. This improvement Furthermore, when the spin is aligned with the orbital angular momentum, the wave-




Figure 6.4: The EM-bright template bank viewed in the black hole mass-spin plane. A significant portion of the parameter space (lower-right) is devoid of any templates, since greater black hole masses and larger anti-aligned spins will not result in disruption of the neutron star before merger. This template bank was generated using the softest reasonable neutron star equation of state, making the EM-bright cut used here conservative.

forms simplify to the form given in Eq. (3.1). It is therefore straightforward to extend the template bank to include these waveforms and incorporate the effects of aligned spins, which could improve the sensitivity of the search by up to an order of magnitude [64].

It is not as straightforward to incorporate precession effects. Precession typically has a less significant effect on the waveform when the binary is observed at small inclinations ($\iota \sim 0$ or $\iota \sim \pi$) [94]. This reduces the importance of precession for the SGRB search. Nonetheless, in PyGRB we wish to cover the full range of possible source parameters. This means that these effects will be somewhat present.

Currently, no complete treatment of precession is available for PyGRB, and so only a partial incorporation of precession has been completed. However, in [96] the authors investigated a method of extending the search to waveforms with precession. This could further increase search sensitivity by a factor between tens of percent and a factor of 2 [64]. In the future, we will identify the regions of parameter space where the spin-aligned waveforms do not provide good sensitivity to precessing signals and develop PyGRB further to provide a sensitive search over these parts of the parameter space. Some work has already been done towards this for CBC searches (see e.g. [99]).

The populations of injected signals used to test search sensitivity now include arbitrary spins (within the bounds of our EM-bright condition). We note that, for NSBHs, the sensitive distance to systems with arbitrary spins tends to be $\sim 80\text{--}85\%$ that for systems with

aligned spins, highlighting the weakness that comes from having inadequate precession incorporated into the template bank.

6.2.5 Targeting Host Galaxy Distances with Injections

If we wish to estimate the detection efficiencies of GW detectors, and therefore calculate exclusion distances, we use families of injected signals placed at a large range of distances. When calculating efficiencies we marginalise over a number of sources of uncertainty, including detector calibration uncertainty. However, if we want to test a specific signal model that hypothesises a source at a certain distance, this approach is not optimal. A much better method would be to inject signals at only the hypothesised distance. In order to do this in a way that correctly incorporates calibration uncertainty, it is best to include the effects of this uncertainty *before* injecting the signals, not afterwards in post processing.

GW interferometers are calibrated using complex models of their frequency response to GWs. These models are validated using many measurements, and the output data calibrated accordingly. The uncertainties on the accuracy of the calibration are quoted as upper bounds on the uncertainties on the amplitude and phase of the interferometer output. These uncertainties are in fact frequency dependent, and so the quoted figures are always the upper bounds across the whole sensitive frequency band, although the uncertainty is roughly constant over this band [18], being slightly greater at the most sensitive frequencies.

In order to understand how to account for these in the placing of injections, we must understand how these uncertainties affect our ability to detect real signals in the data.

Consider a GW signal h in miscalibrated data. The miscalibration will bias the signal that is observed, h_c , which is related to the true signal by

$$h_c = (1 + \delta\alpha)e^{i\delta\phi}h. \quad (6.1)$$

where $\delta\alpha$ and $\delta\phi$ are the amplitude and phase uncertainties respectively. The miscalibration therefore affects the result of a matched filter,

$$\begin{aligned} z_c &= 4 \int_0^\infty \frac{(h_c|h)}{S(f)|h|} df, \\ &= 4 \int_0^\infty (1 + \delta\alpha)e^{i\delta\phi} \frac{|h|^2}{S(f)|h|} df, \\ &\simeq 4 \int_0^\infty (1 + \delta\alpha) \left(1 + i\delta\phi - \frac{\delta\phi^2}{2!}\right) \frac{|h|^2}{S(f)|h|} df, \\ &\simeq 4 \int_0^\infty \left(1 + \delta\alpha + i\delta\phi - \frac{\delta\phi^2}{2!}\right) \frac{|h|^2}{S(f)|h|} df. \end{aligned} \quad (6.2)$$

The biased SNR is therefore

$$\rho_c^2 = |z_c|^2 \simeq \rho_0^2 \left[1 + 2\overline{\delta\alpha} - \sigma_\phi^2 \right], \quad (6.3)$$

where ρ_0 is the unbiased SNR. Here we have defined

$$\overline{\delta X^n} \equiv \frac{4 \int_0^\infty \delta X^n \frac{|h|^2}{S(f)} df}{4 \int_0^\infty \frac{|h|^2}{S(f)} df}, \quad (6.4)$$

for $X = \alpha$ or ϕ , and

$$\sigma_\phi^2 \equiv \left(\overline{\delta\phi^2} \right) - \left(\overline{\delta\phi} \right)^2. \quad (6.5)$$

Miscalibration of amplitude can make the recovered SNR larger or smaller than we might expect. Miscalibration of phase, however, will always lead to a poorer match than we might expect between the templates and the signal, and so will always reduce the recovered SNR. This will also detrimentally affect the effectiveness of our signal consistency tests, but we do not consider that here.

For more discussion on the effect of calibration on matched-filter searches see [36].

Amplitude Uncertainty

Let us minimise the match m due to a given $\delta\alpha_{max}$.

$$\begin{aligned} m &= \frac{(h_c|h)}{|h_c||h|}, \\ &= \frac{\int \frac{(1+\delta\alpha)|h|^2}{S(f)} df}{\left[\int \frac{(1+\delta\alpha)^2|h|^2}{S(f)} df \right]^{\frac{1}{2}} \left[\int \frac{|h|^2}{S(f)} df \right]}, \\ &\simeq \frac{1 + \overline{\delta\alpha}}{\left(1 + 2\overline{\delta\alpha} + \overline{\delta\alpha^2} \right)^{\frac{1}{2}}}, \\ &\simeq (1 + \overline{\delta\alpha}) \left(1 - \overline{\delta\alpha} - \frac{1}{2}\overline{\delta\alpha^2} + \frac{3}{2}\overline{\delta\alpha^2} \right), \\ &\simeq 1 - \frac{1}{2} \left(\overline{\delta\alpha^2} - \overline{\delta\alpha}^2 \right), \\ &\simeq 1 - \frac{1}{2} \sigma_\alpha^2. \end{aligned} \quad (6.6)$$

Therefore, if we have 20 % calibration uncertainty ($\sigma_\alpha = 0.2$), we would get a worst-case match of 0.98 due to miscalibration.

Phase Uncertainty

For phase, $\delta\phi$ leading terms don't have an effect. However, σ_ϕ does play a role in biasing the SNR,

$$\rho_c^2 = \rho_0^2(1 - \sigma_\phi^2), \quad (6.7)$$

and also the match,

$$m = (1 - \sigma_\phi^2). \quad (6.8)$$

A 20° phase error corresponds to $\sigma_\phi^2 \simeq 0.1$, therefore at most a 10 % loss in SNR and match.

Combined Effects on SNR and Match

We have derived the expressions for the biased SNR and match due to phase and amplitude calibration uncertainties. The combined effects are:

$$\rho_c = \rho_0 \left(1 + \delta\alpha - \frac{1}{2}\sigma_\phi^2 \right) \quad (6.9)$$

$$m = 1 - \frac{1}{2}\sigma_\alpha^2 - \frac{1}{2}\sigma_\phi^2 \quad (6.10)$$

Implementation

The effect of $\delta\alpha$ is as a noise that can either increase or decrease the SNR/match between template and signal. To counteract this we must spread our signals out in injected distance accordingly.

The effect of $\delta\phi$ is primarily to artificially reduce SNR. To counteract this we must systematically inject signals at inflated distances.

In general, these calibration uncertainties affect coherent searches such as PyGRB by both altering the SNR and affecting the amplitude and phase consistency between detectors. As such, in future we will wish to inject our signals independently in each detector data stream, each time with an appropriately different alteration to better simulate calibration uncertainties.

6.2.6 Extension to Low Frequencies

The second generation detectors will be sensitive to signals from 10 Hz upwards when at their design sensitivities [95], compared with 40 Hz for their predecessors. This represents a challenge for matched filtering analyses because of the increased time that CBC signals will spend in the sensitive band, which will require much longer template waveforms. The lowest mass system we consider in our search, and therefore the longest waveform, is a $m_1 = m_2 = 1.0 M_\odot$ BNS merger. This system takes about 44 s to evolve from 40 Hz to merger, but 1777 s to do so from 10 Hz.

The SGRB-triggered search used in the past has generally split the data into 256 s long segments when matched filtering, however this is clearly not long enough to filter down to

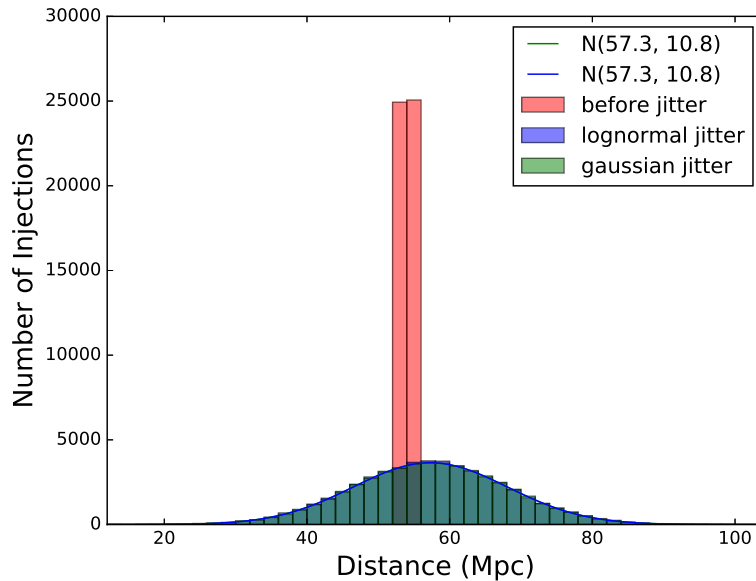


Figure 6.5: Modifying injection distances according to calibration uncertainty. Here a population of 50,000 signals at distances drawn uniformly from between 53–55 Mpc have been adjusted assuming $\delta\alpha = 20\%$ and $\delta\phi = 20^\circ$, using both a ‘Gaussian’ and a ‘lognormal’ method. After adjustment, the signals have distances that follow a normal distribution centred at ~ 57.3 Mpc, with a standard deviation of ~ 10.8 Mpc. The result is the same for both methods, although the lognormal method protects against injections potentially being given an unphysical, negative distance. We can therefore mimic the effect of miscalibration by injecting signals at a range of distances. We interpret our ability to recover these injections as our efficiency at detecting signals between 53–55 Mpc.

10 Hz. Consequently, the search will need to be extended to use longer segments. The longer these segments are, however, the more memory is required to perform the analysis. We will require a better set of methods for reading in data and processing it with the matched filter, as the memory footprint will increase considerably. By being incorporated into the PyCBC project, PyGRB will be able to share a common solution to this problem with the other PyCBC searches in time for runs with increased low frequency sensitivity.

However, the early observing runs are not expected to obtain the full low-frequency sensitivity [2]. For these runs a search beginning at 30 Hz will be sufficient to recover the available signal power. A $m_1 = m_2 = 1.0 M_\odot$ system takes only about 96 s to merge from 30 Hz, which is short enough that we will still be able to use the standard 256 s segments. All that is required is that we change the part of the segment used for trigger generation.

Each 256 s segment overlaps each of its neighbours by 128 s. The data in each segment are then filtered using templates that terminate within a 128 s window, ensuring that there are no gaps. In Fig. 6.6 we show the difference between the previous approach and the new approach. Previously the middle half of a given segment was used for trigger generation.




Figure 6.6: Changing the portion of each data segment used to produce triggers (hatched) allows for the search to be extended down to 30 Hz. The waveforms plotted are for $m_1 = m_2 = 1.0 M_\odot$ systems. The upper, green waveforms show the length of the inspiral from 40 Hz, and the lower, red waveforms show this from 30 Hz. By extending the search down to 30 Hz the old method (top) of analysing the middle half of the segment would see the loss of much of the early inspiral power when filtering the early part of the segment. The new method (bottom) shifts the focus to the latter part of the analysable data segment, and avoids this issue almost entirely.

Now we make use of most of the latter half of a segment to generate triggers.

The previous, ‘segment-middle’ approach was sufficient for a search beginning at 40 Hz since a $1.0 M_\odot$ – $1.0 M_\odot$ template does not stray into the portion of data that is corrupted due to data conditioning. Extending the template down to 30 Hz would mean that roughly half of the inspiral template is lost, and a considerable amount of the total potential SNR is not recovered.

The new, ‘segment-end’ approach provides at least 94 s of available data, so for the longest templates only a negligible amount of signal power will ever be lost.

We may also ask whether we might reduce the amount of data lost to conditioning. A total of 17 s is lost at both the beginning and end of the segments to conditioning. If we could reduce the amount of data used in this process without detrimental effects, this might be desirable for the search down to 30 Hz.

One second of this allows for the sliding of data due to light travel time delay between sites. The remaining 16 s are lost due to the calculation of the truncated inverse power spectrum.

Recall Eq. (3.4):

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

We must calculate the inverse power spectrum, $1/S_n(f)$, in order to do the matched filter calculation. However, if we use the whole analysis segment to calculate $S_n(f)$, sharp spectral features can have impulse responses as long as the analysis segment when they are Fourier transformed. This can corrupt the entire segment. Instead, we take the inverse of the noise power spectral density in the temporal domain, $1/S_n(t)$, truncate its duration (the inverse spectrum length (ISL)), then apply an FFT to obtain $1/S_n(f)$. The presence of sharp features will now only corrupt data of length equal to the ISL. This is a two-sided calculation, so that a 32 s ISL corrupts 16 s of data on either side of the template when matched filtering. This inverse spectrum is then upsampled to the data sample rate for matched filtering. See [37] for more technical discussion on this procedure.

The ISL is given to the matched filter code as an argument, so it can be reduced. But what, if any, are the effects?

To investigate this we take an example analysis segment and filter with a small template bank containing just five waveforms spanning a range of chirp masses. We do this a number of times, changing the ISL each time.

In Fig. 6.7 we show the number of triggers as a function of SNR for three values of the ISL: 8 s, 16 s, and 32 s. As the value used is reduced, the number of triggers increases dramatically. For a value of 16 s the number of triggers is larger than for 32 s. When using 8 s the number of triggers is a factor of ~ 100 larger than for 16 s. This becomes computationally prohibitive, since for each trigger we must calculate signal consistency tests and read/write from/to file. The increase is not strongly correlated to any specific templates or times, but is not seen when searching down to only 40 Hz. The increase in trigger rate is also more pronounced at lower SNR.

The reason for this is that, by reducing the ISL, narrow features in $S_n(f)$ become less well resolved. In this case, line features present in the O1 aLIGO data between 30–40 Hz are barely resolved at all when using an ISL of 8 s (Fig. 6.8). When calculating the matched filter SNR, any underestimation of the noise at a given frequency will make it appear like there is excess power there. Given the degradation in the spectral features and increase in trigger rate, even for an ISL of 16 s, and how little of the low mass templates we currently lose at the start of each segment, we will continue to use a 32 s ISL as default for as long as we search down only to 30 Hz.

6.3 Responding Online

In order to maximise the astrophysical potential of the SGRB analysis, we must minimise the time between an SGRB observation and the PyGRB search result. With an observed GRB rate of ~ 1 per day, it is not practical for all relevant telescopes to follow up all GRBs.

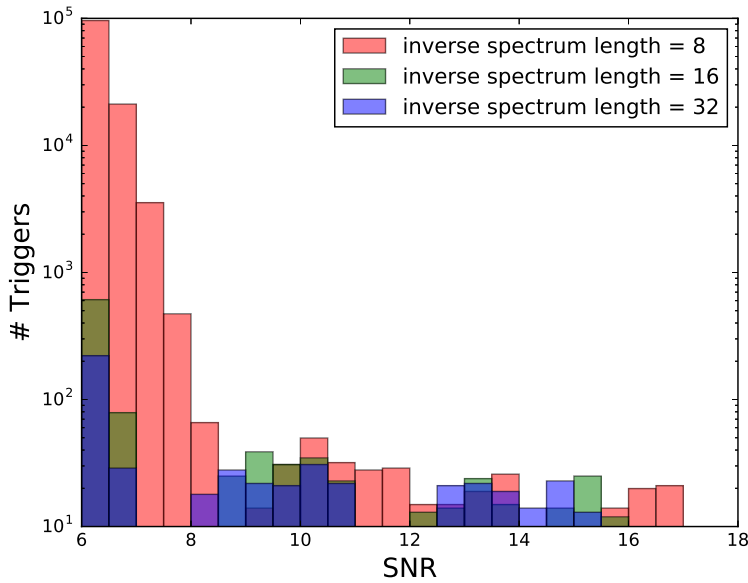


Figure 6.7: Trigger rate dependence on the inverse spectrum length. A reduction in the inverse spectrum length from 32 s to 16 s sees a modest increase in the trigger rate. However, reducing it further to 8 s (red) leads to an explosion in triggers ($\sim \times 100$), which places a prohibitively heavy memory load on the analysis.

However, it is highly likely that the report of a GW detection in association with a GRB will motivate astronomers, who might otherwise have ignored the GRB, to target it. The sooner such a report is received, the better the suite of observations are likely to be.

Unfortunately, as described in Section 4.1, the ability to make a definitive claim on detection comes from running many time shifted analyses, which is time consuming. Similarly, accurate measures of sensitivity are obtained by injecting $\sim 10^4$ simulated signals in the data and searching for them, again a time consuming process.

Therefore, a desirable set up is to run PyGRB in two modes: the standard, or *offline* mode; and a specialised *online* mode. The online configuration minimises computational cost and is run soon after the SGRB, sacrificing some statistical accuracy for speed, with the result used to inform observational follow-up strategies. By comparison, the offline configuration is more computationally intensive and is run some time later, making use of improved knowledge of data quality and calibration to obtain a final result, available days to weeks after the SGRB.

Naturally, the online mode is most effective when automated. Removing the human from the loop not only reduces the time between SGRB and result, but also helps to ensure a consistency across all online analyses by removing the possibility of human error in launching an analysis.

Options such as the time and sky position of an SGRB will clearly differ for each instance of the online search but can be parsed automatically from an event alert, such as a




Figure 6.8: The effect of the inverse spectrum length on spectral resolution. Features in the spectrum are poorly resolved when smaller lengths are used. This is particularly pronounced for sharper features, such as those between 30–40 Hz (inset). Poor resolution of these features in particular are predominantly responsible for the large increase in trigger rate when using an 8 s inverse spectrum length (Fig. 6.7), since underestimating the noise at a given frequency will make it appear that there is excess power there when matched filtering.

Gamma-ray Coordination Network (GCN) alert, and provided as command line arguments upon execution of the PyGRB online workflow generator. The other workflow-specific options may be given in configuration files tailored to the online workflow, which are kept under version control in a repository and are copied to the local machine at run time. Automation of the workflow generation is therefore relatively simple.

Finally, the generation of a result web page and copying of the final result data products to a web-facing location can be done by the final job in the workflow, tasks that used to require human intervention.

Through these features it is possible to have a pipeline that will run and be monitored completely without human intervention. The first time a human need even be aware of the analysis is when the software sends out a notification that the result is ready to be reviewed.

Only the most crucial information for astronomical follow-up purposes should be needed before issuing a first alert to observing partners. This can include whether there is a significant on-source candidate after short time slides, achieving a FAP measured down to to $\lesssim 10^{-4}$, and a small population of injected signals for an early estimate of how detection efficiency varies with distance. With current hardware and software technology these can be achieved within an hour or two of the SGRB being detected, but a number of

developments were needed in the infrastructure to make this possible.

6.3.1 The External Trigger Alert System

Since many transient astrophysical events, including supernovae and GRBs, are of great interest to GW astronomy, we want to do whatever possible to provide periods of stable data around the times of these astrophysical triggers. Specifically, planned downtime can be cancelled or postponed for a period of time after an astrophysical trigger. We will also cancel planned hardware injections, when the mirrors are actuated to simulate a passing GW.

We have written a software system, based on a PYTHON script, that runs in GW detector control rooms, which alerts on-duty controllers of the trigger, and interfaces with hardware injection software to suspend planned injections. At time of writing this is used at both LIGO observatories and at GEO-600.

Once an astrophysical trigger is detected by an observatory a GCN alert is issued. This alert is digested by a listener, which creates a Gravitational Wave Candidate Event Database (GRACEDB)⁷ event associated with the trigger. This typically occurs within a few minutes of detection.

We query this database at regular intervals from a continuous PYTHON process by utilising the `ligo.gracedb` PYTHON module, written specifically for interfacing with GRACEDB. For our purposes we request the information on all ‘External’ group events between the present time and T seconds before present. This information is compiled into a list. These currently include events reported by *Fermi*, *Swift*, and the SuperNova Early Warning System (SNEWS).⁸ This database query may sometimes fail to complete, so we repeat the query every z seconds, where $z < \frac{T}{N_f}$. This allows for N_f failed/skipped queries before a trigger risks going unnoticed by the alert system. If the query is successful and no event has occurred in the time period T , the code sleeps before repeating the query.

If there have been one or more events within T we take the most recent event as our event of interest. If this is a new event — i.e. the GRACEDB entry was created after the last successful query — the event’s unique GRACEDB ID, GPS time, type, and source are parsed by the code. If the type and source identify this as a relevant astrophysical trigger, the software will alert the controller and pause hardware injections automatically. The current default time period for the disabling of hardware injections for all GRB and supernova alerts is 3 hours.

The latency for this whole process is only a few seconds for reasonable lookback times T (of order tens of minutes). To mitigate against untimely code or query failures, a value of T of order one hour may be used without increasing the latency. A standard value for z is 5 s.

The control room systems run EPICS software⁹ with a graphical interface pro-

⁷<https://gracedb.ligo.org/>

⁸<http://snews.bnl.gov/>

⁹<http://www.aps.anl.gov/epics/>

vided via MEDM screens.¹⁰ These systems allow people to interact with and control the interferometers, and monitor the states of all the detector subsystems. We modified an MEDM screen that was dedicated to external triggers and hardware injections, so that it now displays information on the latest event of interest (Fig. 6.9). In order to hold this information we created four EPICS channels per interferometer, where <IFO> represents either H1 or L1 depending on the interferometer:

- <IFO> : CAL-INJ_EXTTRIG_ALERT_ID
The Unique GRACEDB ID
- <IFO> : CAL-INJ_EXTTRIG_ALERT_TIME
The GPS time of the trigger
- <IFO> : CAL-INJ_EXTTRIG_ALERT_TYPE
The trigger type: {GRB, Supernova, ...}
- <IFO> : CAL-INJ_EXTTRIG_ALERT_SOURCE
Origin of the alert: {*Fermi*, *Swift*, SNEWS, ...}

The above channels may be easily read and written to from a running PYTHON process by making use of the PYTHON package PYEPICS.¹¹ All changes are visible in real time on the MEDM screen and software may easily monitor these channels and react accordingly when they are updated.

In order to monitor the running of this PYTHON code, a fifth new EPICS channel was added, <IFO> : CAL-INJ_EXTTRIG_ALERT_QUERY_TIME. This channel is populated with the GPS time whenever a successful query is performed. The MEDM screen monitors this channel and will display a warning message if it has been more than one minute since the last successful query.

This system was used to ensure the integrity of the data around external triggers during O1, and will continue to be used at the sites for the foreseeable future.

6.3.2 The Online Workflow

A GCN alert will contain a preliminary calculation of the sky position and duration of the burst. This is parsed automatically by a script that listens for alerts. If early indications are that this is an SGRB, the script generates an online PyGRB workflow for the event using this information.

The offline search is run on GRBs with $t_{90} - \Delta t_{90} \leq 4$ s, where t_{90} is the time interval during which the middle 90 % of the total integrated flux is received. However, the t_{90} statistic often takes some time to be calculated since it is based upon the full gamma-ray light curve. For the online search we must base the decision to analyse on another factor, one which is available within minutes and that will rarely discard a short-hard burst but still discard the majority of the long-soft ones.

¹⁰<http://www.aps.anl.gov/epics/extensions/medm/>

¹¹<http://cars9.uchicago.edu/software/python/pyepics3/>

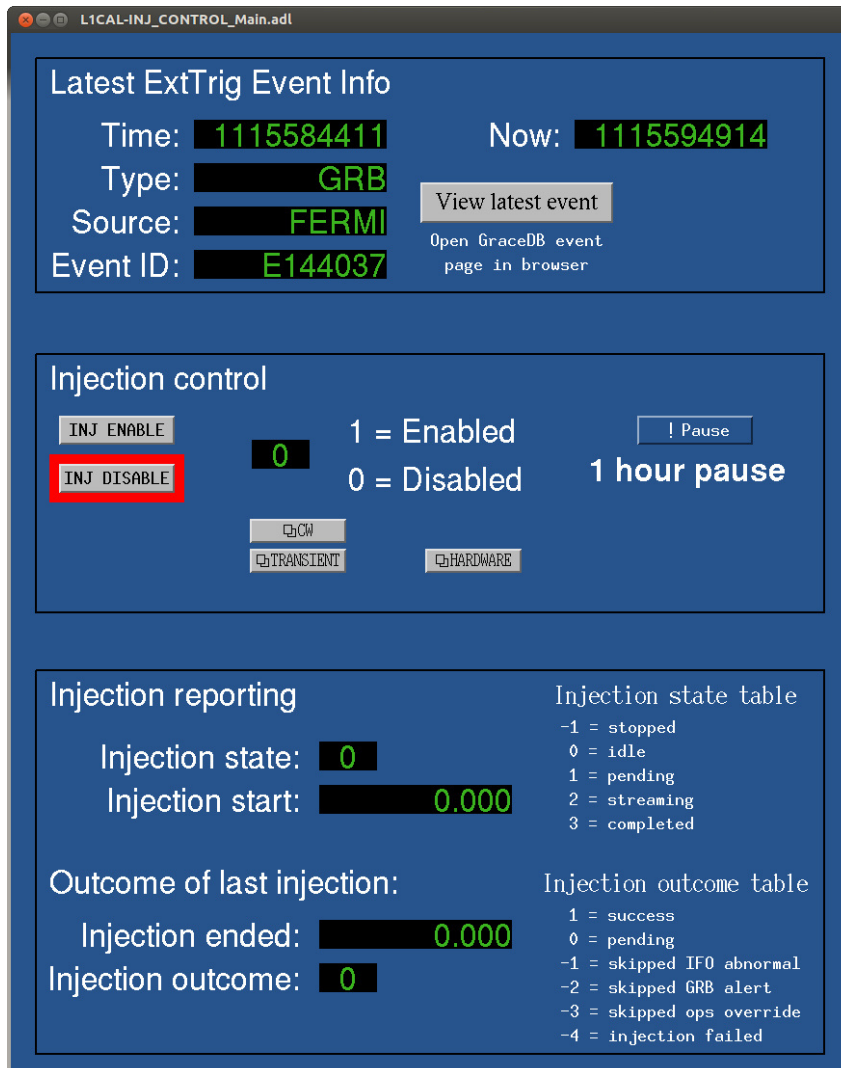


Figure 6.9: Screenshot of the injection control MEDM screen, featuring information on the latest external trigger event as processed by the PYTHON script. If the event occurred within a pre-defined period of time before the present, the top rectangular box will flash red to alert the controller. The ‘View latest event’ button opens the relevant GRACEDB event page in the default web browser.

GCN alerts generally report the trigger duration; the amount of time that the spacecrafts’ onboard analysis spends integrating the light curve before the event is identified as having high significance. *Swift* and *Fermi* report this information in their ground notices as ‘Integ_Time’ and ‘Data_Integ’ respectively. These notices are generally issued on timescales of a minute after the burst. If this value is more than 1.024 s the GRB is very unlikely to be short-hard [59], so we use this as the threshold for launching the online analysis. As a result, the online search should run on almost all SGRBs and some LGRBs (based upon t_{90}). When updated information becomes available, we may cancel ongoing analyses of LGRBs or launch offline analyses for misclassified SGRBs.

In order to run online the pipeline requires the data to have been transferred from the

observatories to the analysis machines soon after data collection. This online data pathway is different from that used by the offline search, where data are packaged in longer segments. The workflow generator will attempt to build a viable workflow as soon as possible, however there may not yet have been enough elapsed time for all the data requirements to have been met (e.g. too little data after GRB time or too little contiguous analysable data). In such cases the workflow generation will exit gracefully, and the online infrastructure will retry the generation a short time later.

It is also possible that the transfer process itself will miss out segments of data. PyCBC handles this by checking for consistency between the data marked as being of science quality and the data that are stored on the machine. This again leads to a failed attempt that can be retried later.

If a workflow is successfully generated it will include a matched filter subworkflow, without extended background, which is executed first. Once this is finished a decision is made depending on the FAP of the loudest on-source event. If the FAP is above a threshold, a minimal number of simulated signals are analysed to test the search response to signals. If the FAP is below the threshold, simulated signals and extended background analyses are carried out to further assess the significance of the event, and the sensitivity of the search. A diagram of this workflow structure is shown in Fig. 6.10.

During O1, this online workflow generally achieved latencies of between a few and around ten hours. This is likely to be reduced significantly in future runs, both due to improved workflow development, tuning, and optimisation, and by the use of dedicated, high performance computing nodes.

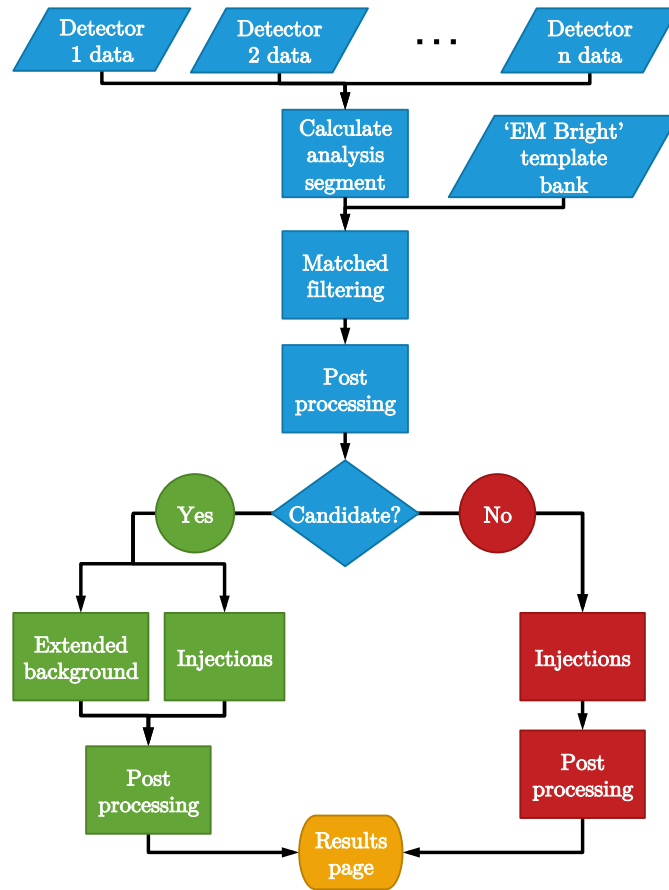


Figure 6.10: A schematic of the online PyGRB workflow. The first part of the analysis features matched filter jobs, with the aim of obtaining a relatively quick estimate of the significance of the loudest on-source event. This does not feature extended background estimation with timeslides. The FAP of the loudest event is used to decide what path to take for the remainder of the analysis. If there is no significant event, the analysis will finish with a minimal set of simulated signals, used to estimate the search sensitivity (red). If however there is a (potentially) significant event, sets of simulated signals and further background time slides are done (green). Note that the same EM-bright template bank is used for all GRBs, and is generated based upon the data in the engineering run immediately preceding each observing run.

7

PYGRB IN THE FIRST OBSERVING RUN OF ADVANCED LIGO

Advanced LIGO Observing Run 1 began on 12 September 2015, and continued through to 12 January 2016. During the run, the two LIGO detectors were operating at a sensitivity of 3–5 times better than in LIGO Science Run 6 (S6) in the most sensitive band (100–300 Hz), and at lower frequencies were at least 10 times better than in S6 [20].

In this chapter we present the results of the PyGRB search for GWs associated with 16 SGRBs that were observed during O1. These include GRB 150906B, which – although it occurred 6 days before the official beginning of O1 – was a potentially nearby burst observed by the IPN, and thus of particular interest.

7.1 GRB Sample

Our sample contains GRB triggers from the GCN, supplemented by the *Swift* and *Fermi* trigger pages. These GRBs were ingested into GRACEDB within seconds of detection through a GCN listener. Their time and localisation parameters were later updated and cross-checked against the satellite databases and published catalogues through a dedicated vetting process based upon Vetting Automation and Literature Informed Database (VALID). The information collected in this database is automatically retrieved and updated by querying satellite databases and by performing automated literature searches via the arXiv system.

The classification of GRBs into short and long is somewhat ambiguous. Since binary mergers are particularly strong sources of gravitational radiation, we make use of a more




Figure 7.1: Durations of the GRBs observed during the first observing run of aLIGO. These classifications are based on conservative constraints on the T_{90} statistic. A GRB is classified as unambiguously short if its $T_{90} + \Delta T_{90} < 2$ s. A GRB is classified unambiguously long if its $T_{90} - \Delta T_{90} > 4$ s. The remaining GRBs are classified as ambiguous. These thresholds can be changed at any time, depending on the needs of the search. Additionally, there are two GRBs (151107A and 151212A) that do not have T_{90} information. This histogram does not include IPN-only GRBs.

lenient classification to identify GRBs which may originate from a binary merger. Our selection is based on the T_{90} statistic, which is the time interval over which the middle 90 % of the total background-subtracted photon counts are observed. Specifically, we consider all GRBs that have $T_{90} - \Delta T_{90} > 4$ s as being short or potentially short. This choice, as opposed to the more standard 2 s cutoff for SGRBs, is made so that we do not neglect SGRBs in the tail of the duration distribution, even if it means we do analyse some LGRBs.

In total there were 20 SGRB candidates observed by these satellites during the period of interest, as shown in Fig. 7.1. Of these, only those that occurred when at least one of the aLIGO detectors was operating in a stable configuration are analysed. GW data segments that are flagged as being of poor quality are excluded from the analysis. 14 of the 20 met these data requirements and were analysed. Two further SGRB were included in our analysis, despite both occurring on 6 September 2015, 6 days before the beginning of O1. At this time the aLIGO detectors were in a stable engineering run just prior to O1, and so had data of sufficient quality to include these SGRBs in our sample. GRB 150906944 was detected by the *Fermi* GBM, and GRB 150906B was detected by the IPN as a potentially nearby burst. In total, we analysed 16 SGRBs with PyGRB, and these are listed in Table 7.2.




Figure 7.2: The localisation of GRB 150906B by IPN satellites, indicating the localisation constraints due to triangulation between different pairs of satellites. The black box in the centre is the overall 3σ error box.

7.1.1 GRB 150906B

GRB 150906B was a short-duration, hard-spectrum burst that occurred at 08:42:25 UTC on 6 September 2015, and was detected by the IPN [90, 107]. It was observed by the *Konus-Wind*, *INTEGRAL*, *Mars-Odyssey*, and *Swift* satellites. It was outside the coded field of view of the *Swift* BAT, so localisation was achieved by triangulation of the signals observed in the IPN satellites, as shown in Fig. 7.2. It was localised to within a 3σ error box with an area of 210 square arcminutes.

The localisation region of the GRB lies close to the local spiral galaxy NGC 3313, which lies close to face-on at a redshift of $z = 0.0124$, corresponding to a luminosity distance of ~ 54 Mpc [118]. At this distance, the galaxy lies 130 kpc in projection from the GRB error box — a distance that is consistent with observed offsets of SGRBs from galaxies, and consistent with the expected supernova kicks imparted on BNS systems [47]. NGC 3313 is the most luminous of a group of galaxies. Other, fainter members of the group also lie close to the GRB error region as shown in Fig. 7.3 and are also possible hosts of the progenitor. There are also a number of galaxies at around 500 Mpc within the error region of the GRB. Unfortunately, follow-up electromagnetic observations of 150906B were not possible due to its proximity to the Sun.

The *Konus-Wind* observation of GRB 150906B was further used to classify the GRB [164]. This was observed to have a duration of $T_{50} = (0.952 \pm 0.036)\text{s}$ ¹ and $T_{90} = (1.642 \pm 0.076)\text{s}$, which places it at the longer end of the SGRB distribution. The

¹Similarly to T_{90} , T_{50} is the time interval over which the middle 50% of the total background-subtracted photon counts are observed.




Figure 7.3: Overlay of the error box for GRB 150906B on the sky. The circled galaxies are at around 54 Mpc, while the others are at 500 Mpc. The largest galaxy in the field is NGC 3313.

spectral hardness in the *Konus-Wind* satellite, $\log HR_{32}$, gives the logarithm of the ratio of counts in the 200–760 keV and 50–200 keV bands. The spectral hardness of GRB 150906B also lies between the peaks of the short and long GRB distributions, as shown in Fig. 7.4. Thus, it is problematic to firmly classify the GRB as either long or short.

If we assume GRB 150906B occurred in NGC 3313 it would have had an isotropic-equivalent gamma-ray energy $E_{\text{iso}} \sim 10^{49}$ erg, which is consistent with the lower end of inferred luminosities of SGRBs with measured redshifts [47]. Some theoretical arguments proposed that the energetics would have fit better with a more distant system, possibly originating from one of the galaxies at 500 Mpc [150, 186]. Therefore, the energetics alone cannot rule out either of these hypotheses.

7.2 Gravitational Wave Detectors

During O1, the aLIGO detectors were operating in a configuration that provided sensitivity to astrophysical sources in the frequency range ~ 30 –2000 Hz, with maximum sensitivity around 150 Hz. With a strain sensitivity of $10^{-23} / \sqrt{\text{Hz}}$ at 100 Hz, the typical strain sensitivity of this data represents about a factor of 3 improvement over initial LIGO through most of the sensitive frequency band. This yields approximately a factor of 27 improvement in the sensitive volume.

Detection and astrophysical source parameter estimation require a calibrated estimate of the GW strain sensed by the detectors. This estimate is made possible by a frequency




Figure 7.4: The location of GRB 150906B on the $T_{50} - \log HR_{32}$ diagram for *Konus-Wind* GRBs. It lies between the peaks of the long and short GRB distributions, so it cannot be distinctly classified as either short or long.

domain model of the detector’s GW response, which is validated by making measurements of the response of the detector to deliberate actuation and various cross-checks. During O1 proper, aLIGO calibration uncertainty was less than 10% in amplitude and less than 10° in phase across the 20 Hz–1 kHz band [18]. We analyse data after removing time segments during which an identified instrumental or environmental noise source coupled to the GW strain signal.

Before O1 the aLIGO detectors were undergoing final preparations in an engineering test phase – Advanced LIGO Engineering Run 8 (ER8). The beginning of O1 was a smooth transition between ER8 and the observing phase on 12 September 2015. The GRBs 150906B and GRB 150906944 occurred 6 days before this, firmly placing them within ER8. Both detectors were operational at the time of 150906944, and H1 was operational for GRB 150906B. However, at this stage in ER8 the calibration uncertainty was greater than in O1, with upper bounds of 20% in amplitude and 20° in phase across the 20 Hz–1 kHz band.

Data Quality Issues

During O1, a number of instrumental glitch types contributed significantly to the background in PyGRB analyses. The most significant were saturations of digital-to-analogue converters (DACs) associated with the y-arm end test mass (ETMYs) at both sites, though more frequently in the H1 detector. These glitches are typically short, usually less than 1 s

duration, but result in triggers with significant SNR in CBC searches, and therefore large drops in the sensitive range of the detectors to CBC signals.

Since these glitches have a known instrumental cause, we can mark the data at these times as being compromised, a process known as vetoing. The all-sky coincident PyCBC search [171] takes the approach of ‘gating’ these vetoed times. This involves smoothly setting the strain data to 0 within the vetoed window, thus erasing all noise. This is an ideal approach for such short glitches, as the matched filter will no longer pick up any contribution to the SNR from these times, and in future runs PyGRB will also gate this type of glitches.

However, during O1 the coherent matched filter engine used by PyGRB was the same as that used by the SGRB-triggered search in S6/VSR2/VSR3 [98]. Unfortunately, this engine does not have the capability to gate the data around glitch times, and instead applies the vetoing for these glitches in post processing. This means the code match filters the data, produces an SNR time series of triggers, and then rejects triggers that coincide with the template waveform ending within the vetoed time.

In Fig. 7.5 we show the burst of triggers in the H1 detector due to an ETMY saturation glitch that occurred in the background time of GRB 151024179, plotting the template end times against SNRs. The triggers are coloured by the chirp mass of the template that produced them. An orange band shows the two second window identified by automatic veto scripts as the time the DAC was saturated. However, it is clear that the glitch contributes SNR when the data are filtered by templates that end slightly before and, more significantly, after these two seconds. This means that, to adequately remove an ETMY saturation glitch in PyGRB post processing, we must veto considerably more time than this.

The colouring shows that at earlier times, the glitch rings off against relatively high mass templates. We have significant triggers for approximately 16 s before the onset of the glitch, corresponding to the glitch falling within the 16 s used to calculate the truncated inverse spectrum (see Section 6.2.6). Then at later times the glitch overlaps longer duration, lower mass templates. With a longest template of duration 96 s, and again 16 s used for the inverse spectrum, we have significant triggers for approximately 112 s after the glitch. Taking into account an additional second at either side to accommodate for time shifting for light travel time between detectors, we must therefore veto from 17 s before to 113 s after the original veto window. This is shown in grey.

7.3 Search Methodology

For the O1 search, we ran PyGRB in both online and offline modes. The online search launched within about twenty minutes following the receipt of an SGRB detection notice, and results became available between a few hours to a day later. These results were then automatically uploaded to GRACEDB and checked by *event advocates*, volunteers who were automatically assigned this responsibility once the results became available. This was the

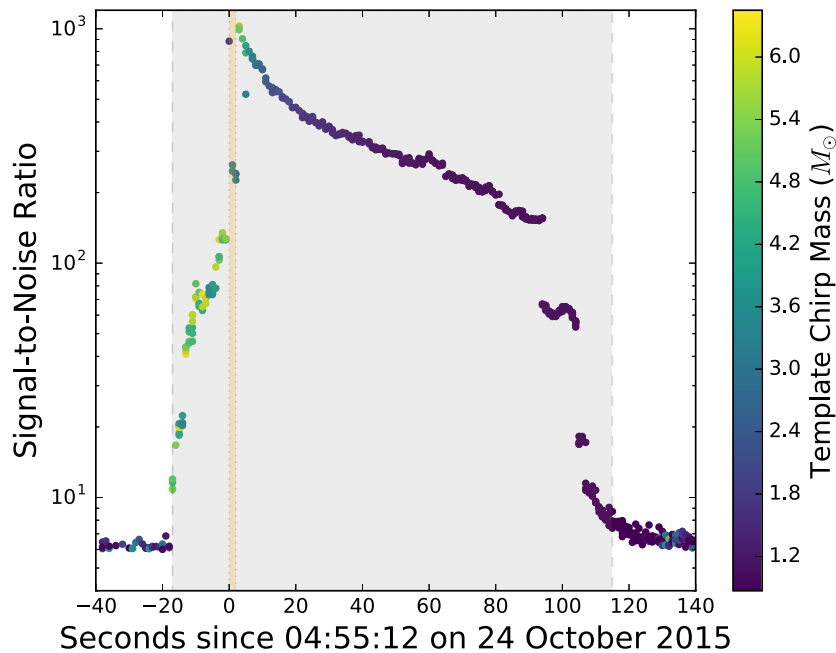


Figure 7.5: An example of vetoing a loud glitch in aLIGO data with PyGRB. We plot triggers due to an ETMY saturation glitch in the H1 detector within the background time of GRB 151024179. The points are coloured by the template chirp mass, and the reported glitch window is shown as an orange band. PyGRB produces significant triggers between about 16 s before and 112 s after this time window. Before the window we see higher mass templates returning high SNRs, due to the glitch falling within the 16 s inverse spectrum truncation time after the template end (see Section 6.2.6). After, ever lower mass templates return high SNR until 112 s after the end of the glitch. These latest triggers are due to the the lowest mass, ~ 96 s templates preceded by another 16 s inverse spectrum truncation time. In grey we show the PyGRB veto window for this glitch, spanning 17 s before to 117 s after the glitch window, which successfully removes the effect of the glitch.

first time such an approach had been taken for a GRB-triggered search for CBC signals.

The final results presented in the rest of this chapter, however, were produced by re-running the offline search with final data quality and calibration.

Configuration: The search was carried out down to a lower frequency of 30 Hz. For those SGRBs detected by *Fermi* GBM, we assume systematic 1σ error circles for the Fermi GBM sky localisations, with radii of 3.7° with 90 % probability, and 14° with 10 % probability. This systematic is added in quadrature to the reported statistical error.

Template Bank: For searching O1 data, we made use of the EM-bright template bank described in Section 6.2.3. Briefly, this comprises SPINTAYLORT4 waveforms [52, 140] for BNS and NSBH systems where $1 M_\odot \leq m_{\text{NS}} \leq 2.8 M_\odot$, $1 M_\odot \leq m_{\text{comp}} \leq 25 M_\odot$, $\chi_{\text{NS}} \leq 0.05$,

$\chi_{\text{BH}} \leq 0.999$, and the EM-bright condition is met [142]. We note that this is the first time a GRB follow-up search has used a spinning template bank.

Simulated Signals: The efficiency of the search at recovering relevant GW signals is evaluated by the addition of simulated signals to the data in software. These data are then filtered with the same EM-bright bank of templates to assess performance. This also provides a means of placing constraints on the SGRB progenitor in the event of no detection in the on-source window. All simulated signals are modelled using the SPINTAYLORT2 waveform approximant (see e.g. [49, 131]). We note that this approximant differs from the one used to build the template bank. This choice is designed to reflect the disagreement between existing NSBH inspiral waveform models in our efficiency assessment (see [131] for more on the challenge this poses for the construction of optimal searches, and on the biases it may cause in measuring parameters from detected signals).

We inject three sets of simulated inspiral signals corresponding to:

1. Generic spin BNS
2. Generic spin NSBH
3. Aligned spin NSBH

We build both generic and aligned spin injection sets for NSBH systems in order to assess the impact of precession on the search sensitivity for rapidly-spinning and highly-precessing systems (BNS systems will not precess significantly within the aLIGO band). The properties of the injected binary populations are determined as follows:

- Black hole masses are drawn from a Gaussian distribution with mean $10 M_{\odot}$ and width $6 M_{\odot}$.
- Black hole dimensionless spins are drawn uniformly from the interval $[0.00, 0.98]$.
- Neutron star masses are drawn from a Gaussian distributions with mean $1.4 M_{\odot}$, and widths of $0.2 M_{\odot}$ for BNS systems and $0.4 M_{\odot}$ for NSBH systems [112, 133]. The larger width for NSBH binaries is chosen in order to reflect the greater uncertainty arising from a lack of observed NSBH systems.
- Neutron star dimensionless spins are drawn uniformly from the interval $[0.0, 0.4]$, which is compatible with the spin of the fastest observed millisecond pulsar [100]. We note that, if we were to design the EM-bright template bank to cover the whole $[0.0, 0.4]$ neutron star spin range, we would require about twice as many templates as we have.
- For sets 1 and 2 (generic spins), both component spins are drawn uniformly from all possible orientations.

- In all three sets, we assume that the SGRB is emitted in the direction of the binary total angular momentum. Relativistic beaming and collimation due to the ambient medium confines the GRB jet to a half-opening angle θ_j . The observation of prompt gamma-ray emission is, therefore, indicative that the inclination of the total angular momentum with respect to the line of sight to the detectors lies within the jet cone. Studies of observed jet breaks in *Swift* GRB X-ray afterglows find a mean (median) value of $\theta_j = 6.5^\circ(5.4^\circ)$, with a tail extending to almost 25° [146]. In at least one case where no jet break was observed, the inferred lower limit was 25° , and the upper limit was 79° [91]. In all three injection sets we assume $\theta_j \leq 30^\circ$; the angle between the total angular momentum and the line-of-sight is therefore uniformly drawn from the interval $[0^\circ, 30^\circ]$.
- The total mass of NSBH systems is restricted to be less than $18 M_\odot$, because the disagreement between the SPINTAYLORT2 and SPINTAYLORT4 approximants is known to dominate beyond this region of the parameter space [131].
- A further astrophysically motivated cut is imposed on the NSBH populations. We do not inject systems that will not produce an accretion torus, even under the most optimistic assumptions regarding the neutron star equation of state and the amount of tidally disrupted neutron star material required to power the GRB emission [142].
- Injections are distributed uniformly in distance within the reach of the detectors, given the time and sky position of each trigger.

7.4 Result for GRB 150906B

If NGC 3313 were indeed the host of a binary merger progenitor of GRB 150906B, we would reasonably expect aLIGO to have detected an associated GW signal since this galaxy lies at a luminosity distance of only about 54 Mpc. A similar hypothesis was tested with the LIGO detectors for the SGRBs 051103 and 070201 [12, 29] (see Sections 3.2.2 and 3.2.3). In both cases, a binary merger scenario was excluded with greater than 90 % confidence, and the preferred scenario is that these events were extra-galactic SGR flares.

PyGRB found no evidence for a GW signal produced by an EM-bright CBC at the time and sky position of 150906B. The most significant candidate event in the on-source window had a re-weighted SNR of 6.19, corresponding to a FAP of 0.526. In other words, in 52.6 % of background trials there was a noise fluctuation as loud or louder than the candidate, so we estimate that to be the probability of seeing a candidate at least as significant in a randomly chosen 6 s of data. The search background is shown in Fig. 7.6.

This null-detection result allows us to compute the frequentist confidence with which we may exclude binary coalescence in NGC 3313 as the progenitor event.

In Fig. 7.7 we show the detection efficiency of the search for all three primary injection sets. For each set we recover all injections within a distance of 54 Mpc, the distance of

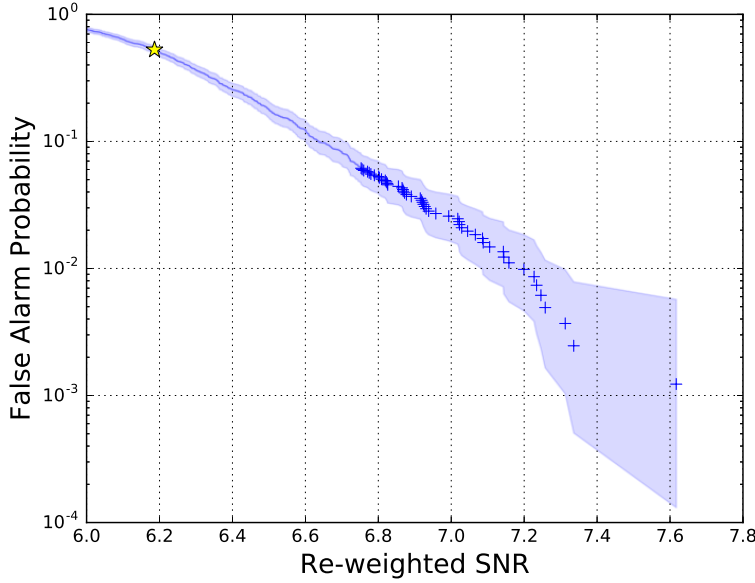


Figure 7.6: PyGRB search background around the time of GRB 150906B. The loudest on-source event ($\rho_{\text{rw}} = 6.19$, FAP = 0.526) is superimposed with a yellow star. This result shows that 52.6% of the 814 background trials contained at least one trigger with re-weighted SNR of at least 6.19.

NGC 3313. We therefore have a very high confidence that, had there been a binary merger in NGC 3313 – or one of its companion galaxies – in the on-source window around the time of GRB 150906B, we would have had a very strong probability of detecting the signal.

In Fig. 7.8 we show the detection efficiency as a function of source inclination angle, this time for generic spin BNSs and NSBHs. For values within the expected range of SGRB beaming angles we find $\geq 99\%$ of simulated signals of both types. Additionally, we recover at least 90% of BNS systems with inclinations up to 60° , and NSBH systems with inclinations up to 80° .

We may also take the loudest on-source event and ask how many injected signals we find louder than this event as a function of distance. In this way we may place exclusion distances on a binary merger event in the on-source window and in the direction of 150906B. These distances are listed in Table 7.1.

Table 7.1: The exclusion distances placed on binary merger progenitors for GRB 150906B

Injection Set	90% Exclusion (Mpc)	50% Exclusion (Mpc)
Generic Spin BNS	101.5	159.3
Generic Spin NSBH	170.2	256.2
Aligned Spin NSBH	186.2	298.5




Figure 7.7: PyGRB injection recovery around the time of GRB 150906B as a function of distance, showing the fraction found louder than the loudest background event. For all three injection sets we see that 100 % of signals are recovered within 54 Mpc. Note that all these injected signals had inclinations of $\leq 30^\circ$, and neutron star dimensionless spins are restricted to below 0.4.

7.5 Cumulative Results

We find no significant events associated with any of the 16 O1 SGRBs analysed by PyGRB. These include GRB 150906B and a further 15 seen only by *Fermi* and *Swift*.

The most significant on-source event occurred for GRB 151228857 and had a FAP of 0.054. We see no strong evidence for a population of weak GW signals.

For each SGRB we obtain 90 % exclusion distances for each of the three populations of simulated signals. In Fig. 7.10 we plot these distances for the generic spin NSBH and BNS populations. The median values are 92.8 Mpc for generic spin BNSs, 147.1 Mpc for generic spin NSBHs, and 157.1 Mpc for aligned spin NSBHs. These are about 7 times larger than the equivalent results from the search discussed in Section 3.3. This compares very favourably with the projections discussed in Section 3.4, where a factor of 10 increase was predicted with aLIGO and AdV operating at design sensitivities.

The results are summarised in Table 7.2.

7.6 Conclusion

We have analysed data from the first aLIGO observing run with PyGRB, searching for GWs associated with 15 SGRBs seen by *Fermi* and *Swift*. No such signals were detected, and there is no strong evidence for a population of weak signals. We set 90 % confidence




Figure 7.8: PyGRB injection recovery around the time of GRB 150906B as a function of source inclination, showing the fraction found louder than the loudest background event. Here we use two populations, generic spin BNS and NSBH, following the same mass and spin constraints outlined in Section 7.3, but remove the constraint on the angle between the line of sight and total angular momentum axes. The injections were placed between 53–55 Mpc following the method described in Section 6.2.5. We have greater than 90 % detection efficiencies NSBH systems with inclinations up to 80°, and for BNS systems with inclinations up to 60°. For values in the most likely range of SGRB beaming angles ($\lesssim 30^\circ$ – e.g. [146]) detection efficiency for both source families is $\gtrsim 99\%$.

lower limits on the distance for each SGRB assuming various progenitor models.

We also analysed data from the H1 detector to look for a GW signal associated with the short-hard GRB 150906B, observed by the IPN. No evidence was found for a GW signal associated with this GRB. The sensitivity of the PyGRB search allows us to confidently exclude the hypothesis that the progenitor system was a BNS or NSBH merger in NGC 3313. Specifically, assuming an outflow jet opening angle $\theta_j = 30^\circ$, we exclude a BNS or NSBH merger in NGC 3313 with $>99\%$ confidence. More generally, we can exclude a BNS progenitor within 100 Mpc with 90 % confidence, and a NSBH progenitor within 170 Mpc, also with 90 % confidence.

We conclude that it is highly unlikely that the progenitor for GRB 150906B was a binary merger in NGC 3313, or one of its smaller companion galaxies. If the event was in fact associated with NGC 3313, then this may have been a core collapse LGRB. Alternatively, it could have originated from a compact binary merger in one of the more distant galaxies at 500 Mpc.




Figure 7.9: Cumulative distribution of p -values for the analysis of 16 SGRBs observed during O1. There is no evidence for a GW signal associated with any of the 16 bursts, with the lowest p -value of 0.054 occurring in the on-source window of GRB 151228857. Additionally, we see no strong evidence for a population of weak sub threshold GW signals based upon a weighted binomial test, as described in Appendix A of [13]. The shaded areas denote the Jeffreys credible intervals corresponding to a $\geq 95\%$ deviation from the null hypothesis.

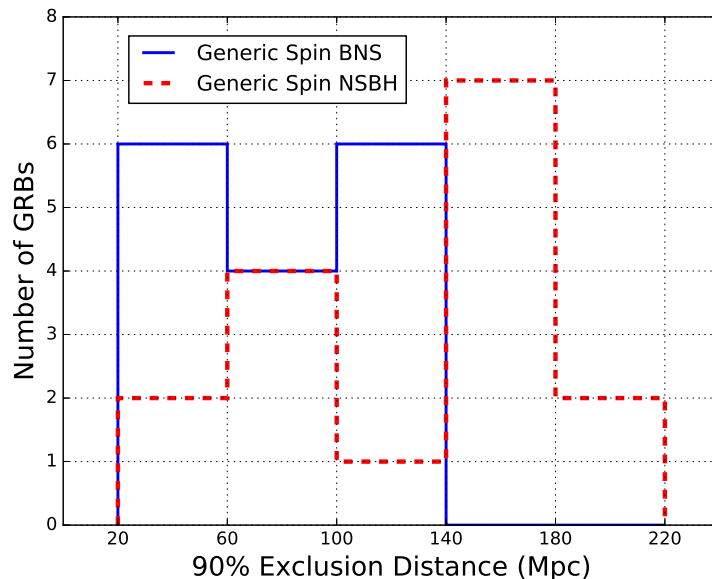


Figure 7.10: 90% distance exclusions for 16 SGRBs observed during O1. The median exclusion distances are 92.8 Mpc for generic spin BNSs, 147.1 Mpc for generic spin NSBHs, and 157.1 Mpc for aligned spin NSBHs.

Table 7.2: The 16 SGRBs that comprise our O1 sample. For each we show the p -values of the loudest on-source event and quote the 90 % exclusion distances for our three injection sets, all with jet opening angles restricted to within 30° . The horizontal line marks the official boundary between ER8 and O1, after which better calibration was available.

GRB Name	UTC Time	Obs. Sat.	GW Dets.	p -value	90 % Exclusion Distances (Mpc)		
					BNS	P. NSBH	A. NSBH
150906B	08:42:25	IPN ^a	H1	0.526	101.5	170.2	186.2
150906944	22:38:47	<i>Fermi</i>	H1L1	0.740	116.9	172.2	199.9
150912600	14:24:31	<i>Fermi</i>	H1L1	0.594	88.4	150.2	149.7
150922883	21:11:32	<i>Fermi</i>	H1L1	0.530	70.7	121.7	122.0
150923297	07:07:36	<i>Fermi</i>	H1L1	0.056	97.8	144.1	186.5
150923429	10:18:17	<i>Fermi</i>	H1L1	0.166	136.2	212.5	241.1
151022577	13:51:02	<i>Fermi</i>	H1L1	-	114.7	178.9	205.7
151024179	04:17:53	<i>Fermi</i>	H1	0.202	24.7	30.2	48.1
151114A	09:59:50	<i>Swift</i>	L1	0.090	41.7	61.4	74.9
151127A	09:08:49	<i>Swift</i>	H1L1	0.823	97.3	151.7	164.5
151202565	13:33:49	<i>Fermi</i>	H1	-	120.6	198.0	226.1
151218857	20:33:31	<i>Fermi</i>	H1L1	0.054	20.7	37.9	35.1
151227072	01:44:07	<i>Fermi</i>	H1L1	-	57.4	97.4	107.5
151228A	03:05:12	<i>Swift</i> ^b	H1	0.492	121.5	168.8	199.6
151229486	11:40:06	<i>Fermi</i>	H1	0.148	57.1	86.4	92.7
151231568	13:38:08	<i>Fermi</i>	L1	0.479	57.6	85.0	96.3

^a IPN satellites: *Konus-Wind*, *INTEGRAL*, *Mars Odyssey*, and *Swift*.

^b 151228A was also seen by *Fermi* and is catalogued as 151228129. However, *Swift* localisation was considerably better so was used for our analysis.

CONCLUSION

The era of gravitational wave astronomy is here. With the clear detection of two binary black hole mergers during the first Advanced LIGO observing run alone [21, 22] (and an additional, weaker detection candidate), it is very likely that these systems will dominate the observations made by ground-based gravitational wave detectors, with probably tens to hundreds of detections per year by the end of the decade [17].

However, at time of writing there has yet to be a detection of a binary merger, or indeed any event, involving anything other than black holes. One of the next milestones will be the detection of a signal from a merger involving one or a pair of neutron stars. In Chapter 2 we have discussed why such events are good candidates for the progenitors of short gamma-ray bursts, amongst the most luminous events in the universe, and could well be the first joint observations made by gravitational wave and electromagnetic observatories. We have investigated the prospects of such a joint detection in the coming years in Chapter 5, and find that they are positive, with perhaps as many as a few per year by the early 2020s. Such observations would provide rich new insights into physics in extremely strong gravitational and magnetic fields, and at nuclear densities. We may, for example, constrain the angular beaming of short gamma-ray bursts, which has implications for the mechanisms that drive the emission of these extreme events.

In Chapter 3 we outlined the methods used for targeted searching of LIGO and Virgo data after the observation of a short gamma-ray burst, and in Chapter 4 we presented important improvements to these methods, which provide an increase in sensitivity of up to about 25 % when compared with a standard all-sky, all-time binary merger search. This roughly doubles the possible rate of joint detections.

We have also summarised of the results of targeted gamma-ray burst follow-up searches with the first generation LIGO and Virgo detectors in Chapter 3, in which some of these improvements were used for the first time. Specifically, in Section 3.3 we presented the results of a search for 27 short gamma-ray bursts observed by the InterPlanetary Network between 2005 and 2010. Although no detections were made, the cumulative results in Section 3.4 for all 69 short gamma-ray bursts analysed with LIGO and Virgo also suggest that joint detections could be possible with the second generation of detectors.

In Chapter 6 we introduced a new search pipeline for targeting short gamma-ray bursts, called PyGRB. This gathers together the improved methods used in the past and

builds upon them to provide a faster and more flexible search. New features allow PyGRB to analyse more short gamma-ray bursts than was previously possible, to do automatically very soon after the gamma-ray burst occurs, and with greater sensitivity and shorter analysis times. Crucially, its flexible, modular nature will make it relatively easy to improve and diversify this search in the future.

We also introduced an astrophysical event alert system in Section 6.3.1, which notifies gravitational wave detector control rooms when a gamma-ray burst, supernova, or other transient event of interest occurs. This helps to ensure the integrity of the gravitational wave data around the time of such an event, and has been operational at the Advanced LIGO and GEO-600 sites since the summer of 2015.

Finally, in Chapter 7, we presented the results of a PyGRB search that targeted 16 short gamma-ray bursts during the first observing run of Advanced LIGO. No associated detections were made, and there is no strong evidence that there is a weak population of associated signals. However, the astrophysical reach of PyGRB during this time was a vast improvement over past searches, and allowed us to make a particularly interesting astrophysical statement regarding GRB 150906B. This short gamma-ray burst occurred near on the sky to a group of galaxies that lie within the Advanced LIGO sensitive range, and the non-detection by PyGRB allowed us to reject the hypothesis that it was associated with the galaxies in the group.

In the upcoming observing runs, gravitational wave astronomy will play an ever more important role in how we view the universe. With Virgo, KAGRA, and a third LIGO detector in India joining the hunt, there will surely be many more discoveries made, and a good chance that a gravitational wave signal associated with a short gamma-ray burst will be one of them. We will be able to obtain highly sensitive results from PyGRB within hours and, hopefully, provide a unique view that will revolutionise our understanding of some of the most mysterious events in the universe.

BIBLIOGRAPHY

- [1] J. Aasi et al. “The characterization of Virgo data and its impact on gravitational-wave searches”. *Class. Quant. Grav.* 29 (2012), p. 155002. DOI: [10.1088/0264-9381/29/15/155002](https://doi.org/10.1088/0264-9381/29/15/155002). arXiv: [1203.5613](https://arxiv.org/abs/1203.5613) [[gr-qc](#)] (cited on p. 48).
- [2] J. Aasi et al. “Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories” (2013). arXiv: [1304.0670](https://arxiv.org/abs/1304.0670) [[gr-qc](#)] (cited on pp. 20, 53, 67, 69, 71, 90).
- [3] J. Aasi et al. “Search for gravitational waves from binary black hole inspiral, merger, and ringdown in LIGO-Virgo data from 2009–2010”. *Phys. Rev. D* 87.2 (2013), p. 022002. DOI: [10.1103/PhysRevD.87.022002](https://doi.org/10.1103/PhysRevD.87.022002). arXiv: [1209.6533](https://arxiv.org/abs/1209.6533) [[gr-qc](#)] (cited on p. 15).
- [4] J. Aasi et al. “Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009–2010 LIGO and Virgo Data”. *Phys. Rev. Lett.* 113.23 (2014), p. 231101. DOI: [10.1103/PhysRevLett.113.231101](https://doi.org/10.1103/PhysRevLett.113.231101). arXiv: [1406.4556](https://arxiv.org/abs/1406.4556) [[gr-qc](#)] (cited on p. 15).
- [5] J. Aasi et al. “Search for gravitational waves associated with gamma-ray bursts detected by the InterPlanetary Network”. *Phys. Rev. Lett.* 113 (2014), p. 011102. DOI: [10.1103/PhysRevLett.113.011102](https://doi.org/10.1103/PhysRevLett.113.011102). arXiv: [1403.6639](https://arxiv.org/abs/1403.6639) [[astro-ph.HE](#)] (cited on pp. xi, 47, 54).
- [6] J. Aasi et al. “Methods and results of a search for gravitational waves associated with gamma-ray bursts using the GEO600, LIGO, and Virgo detectors”. *Phys. Rev. D* 89 (2014), p. 122004. DOI: [10.1103/PhysRevD.89.122004](https://doi.org/10.1103/PhysRevD.89.122004). arXiv: [1405.1053](https://arxiv.org/abs/1405.1053) [[astro-ph.HE](#)] (cited on p. 55).
- [7] J. Aasi et al. “Advanced LIGO”. *Class. Quant. Grav.* 32 (2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). arXiv: [1411.4547](https://arxiv.org/abs/1411.4547) [[gr-qc](#)] (cited on pp. 15–17).
- [8] J. Aasi et al. “Characterization of the LIGO detectors during their sixth science run”. *Class. Quant. Grav.* 32.11 (2015), p. 115012. DOI: [10.1088/0264-9381/32/11/115012](https://doi.org/10.1088/0264-9381/32/11/115012). arXiv: [1410.7764](https://arxiv.org/abs/1410.7764) [[gr-qc](#)] (cited on p. 48).

-
- [9] J. Abadie et al. “Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors”. *Class. Quant. Grav.* 27 (2010), p. 173001. DOI: [10.1088/0264-9381/27/17/173001](https://doi.org/10.1088/0264-9381/27/17/173001). arXiv: [1003.2480](https://arxiv.org/abs/1003.2480) [[astro-ph.HE](#)] (cited on pp. [15](#), [66](#)).
- [10] J. Abadie et al. “Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO’s fifth and Virgo’s first science run”. *Astrophys. J.* 715 (2010), pp. 1453–1461. DOI: [10.1088/0004-637X/715/2/1453](https://doi.org/10.1088/0004-637X/715/2/1453). arXiv: [1001.0165](https://arxiv.org/abs/1001.0165) [[astro-ph.HE](#)] (cited on pp. [15](#), [38](#), [43](#)).
- [11] J. Abadie et al. “Calibration of the LIGO Gravitational Wave Detectors in the Fifth Science Run”. *Nucl. Instrum. Meth.* A624 (2010), pp. 223–240. DOI: [10.1016/j.nima.2010.07.089](https://doi.org/10.1016/j.nima.2010.07.089). arXiv: [1007.3973](https://arxiv.org/abs/1007.3973) [[gr-qc](#)] (cited on p. [42](#)).
- [12] J. Abadie et al. “Implications For The Origin Of GRB 051103 From LIGO Observations”. *Astrophys. J.* 755 (2012), p. 2. DOI: [10.1088/0004-637X/755/1/2](https://doi.org/10.1088/0004-637X/755/1/2). arXiv: [1201.4413](https://arxiv.org/abs/1201.4413) [[astro-ph.HE](#)] (cited on pp. [38](#), [108](#)).
- [13] J. Abadie et al. “Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3”. *Astrophys. J.* 760 (2012), p. 12. DOI: [10.1088/0004-637X/760/1/12](https://doi.org/10.1088/0004-637X/760/1/12). arXiv: [1205.2216](https://arxiv.org/abs/1205.2216) [[astro-ph.HE](#)] (cited on pp. [15](#), [31](#), [38](#), [43](#), [44](#), [54](#), [112](#)).
- [14] J. Abadie et al. “Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO’s Sixth Science Run and Virgo’s Science Runs 2 and 3”. *Phys. Rev. D* 85 (2012), p. 082002. DOI: [10.1103/PhysRevD.85.082002](https://doi.org/10.1103/PhysRevD.85.082002). arXiv: [1111.7314](https://arxiv.org/abs/1111.7314) [[gr-qc](#)] (cited on pp. [15](#), [36](#), [47](#), [48](#), [61](#), [67](#)).
- [15] B. P. Abbott et al. “LIGO: The Laser interferometer gravitational-wave observatory”. *Rept. Prog. Phys.* 72 (2009), p. 076901. DOI: [10.1088/0034-4885/72/7/076901](https://doi.org/10.1088/0034-4885/72/7/076901). arXiv: [0711.3041](https://arxiv.org/abs/0711.3041) [[gr-qc](#)] (cited on pp. [12](#), [14](#)).
- [16] B. P. Abbott et al. “Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO’s S5 Data”. *Phys. Rev. D* 79 (2009), p. 122001. DOI: [10.1103/PhysRevD.79.122001](https://doi.org/10.1103/PhysRevD.79.122001). arXiv: [0901.0302](https://arxiv.org/abs/0901.0302) [[gr-qc](#)] (cited on p. [42](#)).
- [17] B. P. Abbott et al. “Binary Black Hole Mergers in the first Advanced LIGO Observing Run” (2016). arXiv: [1606.04856](https://arxiv.org/abs/1606.04856) [[gr-qc](#)] (cited on pp. [3](#), [17](#), [20](#), [48](#), [114](#)).
- [18] B. P. Abbott et al. “Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914” (2016). arXiv: [1602.03845](https://arxiv.org/abs/1602.03845) [[gr-qc](#)] (cited on pp. [75](#), [87](#), [104](#)).
- [19] B. P. Abbott et al. “Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914”. *Class. Quant. Grav.* 33.13 (2016), p. 134001. DOI: [10.1088/0264-9381/33/13/134001](https://doi.org/10.1088/0264-9381/33/13/134001). arXiv: [1602.03844](https://arxiv.org/abs/1602.03844) [[gr-qc](#)] (cited on p. [82](#)).
-

- [20] B. P. Abbott et al. “GW150914: The Advanced LIGO Detectors in the Era of First Discoveries”. *Phys. Rev. Lett.* 116.13 (2016), p. 131103. DOI: [10.1103/PhysRevLett.116.131103](https://doi.org/10.1103/PhysRevLett.116.131103). arXiv: [1602.03838](https://arxiv.org/abs/1602.03838) [[gr-qc](#)] (cited on pp. [17](#), [18](#), [67](#), [82](#), [100](#)).
- [21] B. P. Abbott et al. “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence”. *Phys. Rev. Lett.* 116.24 (2016), p. 241103. DOI: [10.1103/PhysRevLett.116.241103](https://doi.org/10.1103/PhysRevLett.116.241103). arXiv: [1606.04855](https://arxiv.org/abs/1606.04855) [[gr-qc](#)] (cited on pp. [3](#), [48](#), [114](#)).
- [22] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. *Phys. Rev. Lett.* 116.6 (2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837](https://arxiv.org/abs/1602.03837) [[gr-qc](#)] (cited on pp. [2](#), [3](#), [15](#), [18](#), [48](#), [114](#)).
- [23] B. P. Abbott et al. “Properties of the Binary Black Hole Merger GW150914”. *Phys. Rev. Lett.* 116.24 (2016), p. 241102. DOI: [10.1103/PhysRevLett.116.241102](https://doi.org/10.1103/PhysRevLett.116.241102). arXiv: [1602.03840](https://arxiv.org/abs/1602.03840) [[gr-qc](#)] (cited on p. [2](#)).
- [24] B. P. Abbott et al. “Tests of general relativity with GW150914”. *Phys. Rev. Lett.* 116.22 (2016), p. 221101. DOI: [10.1103/PhysRevLett.116.221101](https://doi.org/10.1103/PhysRevLett.116.221101). arXiv: [1602.03841](https://arxiv.org/abs/1602.03841) [[gr-qc](#)] (cited on p. [2](#)).
- [25] B. P. Abbott et al. “Astrophysical Implications of the Binary Black-Hole Merger GW150914”. *Astrophys. J.* 818.2 (2016), p. L22. DOI: [10.3847/2041-8205/818/2/L22](https://doi.org/10.3847/2041-8205/818/2/L22). arXiv: [1602.03846](https://arxiv.org/abs/1602.03846) [[astro-ph.HE](#)] (cited on p. [2](#)).
- [26] B. Abbott et al. “A Search for gravitational waves associated with the gamma ray burst GRB030329 using the LIGO detectors”. *Phys. Rev. D* 72 (2005), p. 042002. DOI: [10.1103/PhysRevD.72.042002](https://doi.org/10.1103/PhysRevD.72.042002). arXiv: [gr-qc/0501068](https://arxiv.org/abs/gr-qc/0501068) [[gr-qc](#)] (cited on p. [38](#)).
- [27] B. Abbott et al. “Upper limits from the LIGO and TAMA detectors on the rate of gravitational-wave bursts”. *Phys. Rev. D* 72 (2005), p. 122004. DOI: [10.1103/PhysRevD.72.122004](https://doi.org/10.1103/PhysRevD.72.122004). arXiv: [gr-qc/0507081](https://arxiv.org/abs/gr-qc/0507081) [[gr-qc](#)] (cited on p. [15](#)).
- [28] B. Abbott et al. “Joint LIGO and TAMA300 search for gravitational waves from inspiralling neutron star binaries”. *Phys. Rev. D* 73 (2006), p. 102002. DOI: [10.1103/PhysRevD.73.102002](https://doi.org/10.1103/PhysRevD.73.102002). arXiv: [gr-qc/0512078](https://arxiv.org/abs/gr-qc/0512078) [[gr-qc](#)] (cited on p. [15](#)).
- [29] B. Abbott et al. “Implications for the Origin of GRB 070201 from LIGO Observations”. *Astrophys. J.* 681 (2008), pp. 1419–1428. DOI: [10.1086/587954](https://doi.org/10.1086/587954). arXiv: [0711.1163](https://arxiv.org/abs/0711.1163) [[astro-ph](#)] (cited on pp. [39](#), [108](#)).
- [30] Aous A. Abdo et al. “Fermi Observations of High-Energy Gamma-Ray Emission from GRB 080916C”. *Science* 323 (2009), pp. 1688–1693. DOI: [10.1126/science.1169101](https://doi.org/10.1126/science.1169101) (cited on p. [23](#)).
- [31] Alex Abramovici et al. “LIGO: The Laser interferometer gravitational wave observatory”. *Science* 256 (1992), pp. 325–333. DOI: [10.1126/science.256.5055.325](https://doi.org/10.1126/science.256.5055.325) (cited on pp. [2](#), [14](#)).

-
- [32] T. Accadia et al. “Virgo: a laser interferometer to detect gravitational waves”. *JINST* 7 (2012), P03012. DOI: [10.1088/1748-0221/7/03/P03012](https://doi.org/10.1088/1748-0221/7/03/P03012) (cited on p. 14).
- [33] F. Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. *Class. Quant. Grav.* 32.2 (2015), p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). arXiv: [1408.3978](https://arxiv.org/abs/1408.3978) [gr-qc] (cited on p. 15).
- [34] F. Acernese et al. “The Advanced Virgo detector”. *J. Phys. Conf. Ser.* 610.1 (2015), p. 012014. DOI: [10.1088/1742-6596/610/1/012014](https://doi.org/10.1088/1742-6596/610/1/012014) (cited on p. 20).
- [35] P. Ajith. “Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins”. *Phys. Rev. D* 84 (2011), p. 084037. DOI: [10.1103/PhysRevD.84.084037](https://doi.org/10.1103/PhysRevD.84.084037). arXiv: [1107.1267](https://arxiv.org/abs/1107.1267) [gr-qc] (cited on p. 85).
- [36] Bruce Allen. “LIGO Calibration Accuracy” (1996). URL: <https://dcc.ligo.org/LIGO-T960189/public> (cited on p. 88).
- [37] Bruce Allen et al. “FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries”. *Phys. Rev. D* 85 (2012), p. 122006. DOI: [10.1103/PhysRevD.85.122006](https://doi.org/10.1103/PhysRevD.85.122006). arXiv: [gr-qc/0509116](https://arxiv.org/abs/gr-qc/0509116) [gr-qc] (cited on pp. 20, 92).
- [38] Masaki Ando et al. “Stable operation of a 300-m laser interferometer with sufficient sensitivity to detect gravitational wave events within our galaxy”. *Phys. Rev. Lett.* 86 (2001), p. 3950. DOI: [10.1103/PhysRevLett.86.3950](https://doi.org/10.1103/PhysRevLett.86.3950). arXiv: [astro-ph/0105473](https://arxiv.org/abs/astro-ph/0105473) [astro-ph] (cited on p. 14).
- [39] Theodoros A. Apostolatos et al. “Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries”. *Phys. Rev. D* 49 (1994), pp. 6274–6297. DOI: [10.1103/PhysRevD.49.6274](https://doi.org/10.1103/PhysRevD.49.6274) (cited on p. 85).
- [40] Yoichi Aso et al. “Interferometer design of the KAGRA gravitational wave detector”. *Phys. Rev. D* 88.4 (2013), p. 043007. DOI: [10.1103/PhysRevD.88.043007](https://doi.org/10.1103/PhysRevD.88.043007). arXiv: [1306.6747](https://arxiv.org/abs/1306.6747) [gr-qc] (cited on p. 15).
- [41] S. Babak et al. “Searching for gravitational waves from binary coalescence”. *Phys. Rev. D* 87 (2013), p. 024033. DOI: [10.1103/PhysRevD.87.024033](https://doi.org/10.1103/PhysRevD.87.024033). arXiv: [1208.3491](https://arxiv.org/abs/1208.3491) [gr-qc] (cited on pp. 20, 32, 37, 69, 79).
- [42] D. Band et al. “BATSE observations of gamma-ray burst spectra. 1. Spectral diversity.” *Astrophys. J.* 413 (1993), pp. 281–292. DOI: [10.1086/172995](https://doi.org/10.1086/172995) (cited on pp. 23, 24, 64).
- [43] S. D. Barthelmy et al. “The Burst Alert Telescope (BAT) on the Swift MIDEX mission”. *Space Sci. Rev.* 120 (2005), p. 143. DOI: [10.1007/s11214-005-5096-3](https://doi.org/10.1007/s11214-005-5096-3). arXiv: [astro-ph/0507410](https://arxiv.org/abs/astro-ph/0507410) [astro-ph] (cited on pp. 27, 38, 53).
-

- [44] Imre Bartos, Patrick Brady, and Szabolcs Marka. “How Gravitational-wave Observations Can Shape the Gamma-ray Burst Paradigm”. *Class. Quant. Grav.* 30 (2013), p. 123001. DOI: [10.1088/0264-9381/30/12/123001](https://doi.org/10.1088/0264-9381/30/12/123001). arXiv: [1212.2289](https://arxiv.org/abs/1212.2289) [[astro-ph.CO](#)] (cited on p. 29).
- [45] S. Basa et al. “The Chinese-French SVOM mission for Gamma-Ray Burst studies” (2008). arXiv: [0811.1154](https://arxiv.org/abs/0811.1154) [[astro-ph](#)] (cited on p. 28).
- [46] E. Berger, W. Fong, and R. Chornock. “An r-process Kilonova Associated with the Short-hard GRB 130603B”. *Astrophys. J.* 774 (2013), p. L23. DOI: [10.1088/2041-8205/774/2/L23](https://doi.org/10.1088/2041-8205/774/2/L23). arXiv: [1306.3960](https://arxiv.org/abs/1306.3960) [[astro-ph.HE](#)] (cited on p. 25).
- [47] Edo Berger. “The Environments of Short-Duration Gamma-Ray Bursts and Implications for their Progenitors”. *New Astron. Rev.* 55 (2011), pp. 1–22. DOI: [10.1016/j.newar.2010.10.001](https://doi.org/10.1016/j.newar.2010.10.001). arXiv: [1005.1068](https://arxiv.org/abs/1005.1068) [[astro-ph.HE](#)] (cited on pp. 102, 103).
- [48] Edo Berger. “Short-Duration Gamma-Ray Bursts”. *Ann. Rev. Astron. Astrophys.* 52 (2014), pp. 43–105. DOI: [10.1146/annurev-astro-081913-035926](https://doi.org/10.1146/annurev-astro-081913-035926). arXiv: [1311.2603](https://arxiv.org/abs/1311.2603) [[astro-ph.HE](#)] (cited on pp. 26, 27, 65).
- [49] Luc Blanchet. “Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries”. *Living Rev. Rel.* 17 (2014), p. 2. DOI: [10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2). arXiv: [1310.1528](https://arxiv.org/abs/1310.1528) [[gr-qc](#)] (cited on pp. 20, 31, 32, 37, 47, 107).
- [50] Luc Blanchet et al. “Gravitational radiation from inspiralling compact binaries completed at the third post-Newtonian order”. *Phys. Rev. Lett.* 93 (2004), p. 091101. DOI: [10.1103/PhysRevLett.93.091101](https://doi.org/10.1103/PhysRevLett.93.091101). arXiv: [gr-qc/0406012](https://arxiv.org/abs/gr-qc/0406012) [[gr-qc](#)] (cited on pp. 20, 37, 47).
- [51] Duncan A. Brown et al. “Detecting binary neutron star systems with spin in advanced gravitational-wave detectors”. *Phys. Rev. D* 86 (2012), p. 084017. DOI: [10.1103/PhysRevD.86.084017](https://doi.org/10.1103/PhysRevD.86.084017). arXiv: [1207.6406](https://arxiv.org/abs/1207.6406) [[gr-qc](#)] (cited on pp. 32, 85).
- [52] Alessandra Buonanno et al. “Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors”. *Phys. Rev. D* 80 (2009), p. 084043. DOI: [10.1103/PhysRevD.80.084043](https://doi.org/10.1103/PhysRevD.80.084043). arXiv: [0907.0700](https://arxiv.org/abs/0907.0700) [[gr-qc](#)] (cited on pp. 20, 106).
- [53] Marta Burgay et al. “An Increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system”. *Nature* 426 (2003), pp. 531–533. DOI: [10.1038/nature02124](https://doi.org/10.1038/nature02124). arXiv: [astro-ph/0312071](https://arxiv.org/abs/astro-ph/0312071) [[astro-ph](#)] (cited on p. 84).
- [54] Eric Burns et al. “Do the Fermi Gamma-Ray Burst Monitor and Swift Burst Alert Telescope see the Same Short Gamma-Ray Bursts?” (2015). arXiv: [1512.00923](https://arxiv.org/abs/1512.00923) [[astro-ph.HE](#)] (cited on p. 28).

-
- [55] Hsin-Yu Chen and Daniel E. Holz. “Gamma-Ray-Burst Beaming and Gravitational-Wave Observations”. *Phys. Rev. Lett.* 111.18 (2013), p. 181101. DOI: [10.1103/PhysRevLett.111.181101](https://doi.org/10.1103/PhysRevLett.111.181101). arXiv: [1206.0703](https://arxiv.org/abs/1206.0703) [[astro-ph.CO](#)] (cited on pp. [71](#), [73](#)).
- [56] Ross P. Church et al. “Implications for the origin of short gamma-ray bursts from their observed positions around their host galaxies”. *Mon. Not. R. Astron. Soc.* 413 (2011), pp. 2004–2014. DOI: [10.1111/j.1365-2966.2011.18277.x](https://doi.org/10.1111/j.1365-2966.2011.18277.x). arXiv: [1101.1088](https://arxiv.org/abs/1101.1088) [[astro-ph.HE](#)] (cited on p. [25](#)).
- [57] J. Clark et al. “Prospects for joint gravitational wave and short gamma-ray burst observations”. *Astrophys. J.* 809.1 (2015), p. 53. DOI: [10.1088/0004-637X/809/1/53](https://doi.org/10.1088/0004-637X/809/1/53). arXiv: [1409.8149](https://arxiv.org/abs/1409.8149) [[astro-ph.HE](#)] (cited on p. [xi](#)).
- [58] S. A. Colgate. “Prompt gamma rays and X-rays from supernovae”. *Can. J. Phys.* 46 (1968), p. 476 (cited on p. [25](#)).
- [59] Valerie Connaughton. Personal communication. Discussion at LIGO-Fermi workshop. Mar. 15, 2015. URL: <https://labcit.ligo.caltech.edu/~jkanner/ligo-fermi/> (cited on p. [97](#)).
- [60] Bertrand Cordier et al. “The SVOM gamma-ray burst mission”. *PoS SWIFT10* (2015), p. 005. arXiv: [1512.03323](https://arxiv.org/abs/1512.03323) [[astro-ph.IM](#)] (cited on p. [28](#)).
- [61] David Coward et al. “The Swift short gamma-ray burst rate density: implications for binary neutron star merger rates”. *Mon. Not. Roy. Astron. Soc.* 425 (2012), p. 1365. DOI: [10.1111/j.1365-2966.2012.21604.x](https://doi.org/10.1111/j.1365-2966.2012.21604.x). arXiv: [1202.2179](https://arxiv.org/abs/1202.2179) [[astro-ph.CO](#)] (cited on pp. [64](#), [65](#)).
- [62] Curt Cutler et al. “The Last three minutes: issues in gravitational wave measurements of coalescing compact binaries”. *Phys. Rev. Lett.* 70 (1993), pp. 2984–2987. DOI: [10.1103/PhysRevLett.70.2984](https://doi.org/10.1103/PhysRevLett.70.2984). arXiv: [astro-ph/9208005](https://arxiv.org/abs/astro-ph/9208005) [[astro-ph](#)] (cited on p. [29](#)).
- [63] V. D’Elia et al. “GRB 100928A: Swift detection of a burst.” *GRB Coordinates Network* 11310 (2010), p. 1 (cited on p. [46](#)).
- [64] Tito Dal Canton, Andrew P. Lundgren, and Alex B. Nielsen. “Impact of precession on aligned-spin searches for neutron-star–black-hole binaries”. *Phys. Rev.* D91.6 (2015), p. 062010. DOI: [10.1103/PhysRevD.91.062010](https://doi.org/10.1103/PhysRevD.91.062010). arXiv: [1411.6815](https://arxiv.org/abs/1411.6815) [[gr-qc](#)] (cited on p. [86](#)).
- [65] Tito Dal Canton et al. “Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors”. *Phys. Rev.* D90.8 (2014), p. 082004. DOI: [10.1103/PhysRevD.90.082004](https://doi.org/10.1103/PhysRevD.90.082004). arXiv: [1405.6731](https://arxiv.org/abs/1405.6731) [[gr-qc](#)] (cited on p. [79](#)).
-

- [66] Ewa Deelman et al. “Pegasus: a Workflow Management System for Science Automation”. *Future Generation Computer Systems* 46 (2015). Funding Acknowledgements: NSF ACI SDCI 0722019, NSF ACI SI2-SSI 1148515 and NSF OCI-1053575, pp. 17–35. DOI: [10.1016/j.future.2014.10.008](https://doi.org/10.1016/j.future.2014.10.008). URL: <http://pegasus.isi.edu/publications/2014/2014-fgcs-deelman.pdf> (cited on p. 80).
- [67] Alexander Dietz. “Estimation of Compact Binary Coalescence Rates from Short Gamma-Ray Burst Redshift Measurements”. *Astron. Astrophys.* 529 (2011), A97. DOI: [10.1051/0004-6361/201016166](https://doi.org/10.1051/0004-6361/201016166). arXiv: 1011.2059 [astro-ph.HE] (cited on p. 73).
- [68] Alexander Dietz et al. “Outlook for detection of GW inspirals by GRB-triggered searches in the advanced detector era”. *Phys. Rev. D* 87.6 (2013), p. 064033. DOI: [10.1103/PhysRevD.87.064033](https://doi.org/10.1103/PhysRevD.87.064033). arXiv: 1210.3095 [gr-qc] (cited on p. 68).
- [69] K. L. Dooley et al. “GEO 600 and the GEO-HF upgrade program: successes and challenges”. *Class. Quant. Grav.* 33 (2016), p. 075009. DOI: [10.1088/0264-9381/33/7/075009](https://doi.org/10.1088/0264-9381/33/7/075009). arXiv: 1510.00317 [physics.ins-det] (cited on p. 14).
- [70] Matthew D. Duez. “Numerical relativity confronts compact neutron star binaries: a review and status report”. *Class. Quant. Grav.* 27 (2010), p. 114002. DOI: [10.1088/0264-9381/27/11/114002](https://doi.org/10.1088/0264-9381/27/11/114002). arXiv: 0912.3529 [astro-ph.HE] (cited on p. 42).
- [71] F. W. Dyson, A. S. Eddington, and C. Davidson. “A Determination of the Deflection of Light by the Sun’s Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919”. *Phil. Trans. Roy. Soc. Lond.* A220 (1920), pp. 291–333. DOI: [10.1098/rsta.1920.0009](https://doi.org/10.1098/rsta.1920.0009) (cited on p. 1).
- [72] David Eichler et al. “Nucleosynthesis, Neutrino Bursts and Gamma-Rays from Coalescing Neutron Stars”. *Nature* 340 (1989), pp. 126–128. DOI: [10.1038/340126a0](https://doi.org/10.1038/340126a0) (cited on pp. 23, 25, 73).
- [73] Albert Einstein. “Zur Elektrodynamik bewegter Körper”. *Annalen der Physik* 17.4 (1905), pp. 891–921. URL: http://www.physik.uni-augsburg.de/annalen/history/einstein-papers/1905_17_891-921.pdf (cited on p. 1).
- [74] Albert Einstein. “Ist die Trägheit eines Körpers von seinem Energieinhalt abhängig?” *Annalen der Physik* 18.4 (1905), pp. 639–641. URL: http://www.physik.uni-augsburg.de/annalen/history/einstein-papers/1905_18_639-641.pdf (cited on p. 1).
- [75] Albert Einstein. “Die Feldgleichungen der Gravitation”. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)* 2 (1915), pp. 844–847. URL: <http://echo.mpiwg-berlin.mpg.de/MPIWG:ZZB2HK6W> (cited on p. 1).
- [76] Albert Einstein. “Näherungsweise Integration der Feldgleichungen der Gravitation”. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)* 1 (1916), 688–696. URL: <http://echo.mpiwg-berlin.mpg.de/MPIWG:RA6W5W65> (cited on p. 1).

-
- [77] Albert Einstein. “Über Gravitationswellen”. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)* 1 (1918), 154–167. URL: <http://echo.mpiwg-berlin.mpg.de/MPIWG:8HSP60BU> (cited on p. 1).
- [78] C. W. F. Everitt et al. “Gravity Probe B: Final Results of a Space Experiment to Test General Relativity”. *Phys. Rev. Lett.* 106 (2011), p. 221101. DOI: [10.1103/PhysRevLett.106.221101](https://doi.org/10.1103/PhysRevLett.106.221101). arXiv: [1105.3456](https://arxiv.org/abs/1105.3456) [gr-qc] (cited on p. 1).
- [79] Stephen Fairhurst. “Triangulation of gravitational wave sources with a network of detectors”. *New J. Phys.* 11 (2009). [Erratum: *New J. Phys.*13,069602(2011)], p. 123006. DOI: [10.1088/1367-2630/11/12/123006](https://doi.org/10.1088/1367-2630/11/12/123006), [10.1088/1367-2630/13/6/069602](https://doi.org/10.1088/1367-2630/13/6/069602). arXiv: [0908.2356](https://arxiv.org/abs/0908.2356) [gr-qc] (cited on pp. 53, 54).
- [80] Stephen Fairhurst and Patrick Brady. “Interpreting the results of searches for gravitational waves from coalescing binaries”. *Class. Quant. Grav.* 25 (2008), p. 105002. DOI: [10.1088/0264-9381/25/10/105002](https://doi.org/10.1088/0264-9381/25/10/105002). arXiv: [0707.2410](https://arxiv.org/abs/0707.2410) [gr-qc] (cited on p. 76).
- [81] V. Ferrari, L. Gualtieri, and F. Pannarale. “Neutron star tidal disruption in mixed binaries: the imprint of the equation of state”. *Phys. Rev. D* 81 (2010), p. 064026. DOI: [10.1103/PhysRevD.81.064026](https://doi.org/10.1103/PhysRevD.81.064026). arXiv: [0912.3692](https://arxiv.org/abs/0912.3692) [gr-qc] (cited on p. 42).
- [82] Lee Samuel Finn and David F. Chernoff. “Observing binary inspiral in gravitational radiation: One interferometer”. *Phys. Rev. D* 47 (1993), pp. 2198–2219. DOI: [10.1103/PhysRevD.47.2198](https://doi.org/10.1103/PhysRevD.47.2198). arXiv: [gr-qc/9301003](https://arxiv.org/abs/gr-qc/9301003) [gr-qc] (cited on p. 68).
- [83] W. Fong et al. “Short GRB 130603B: Discovery of a Jet Break in the Optical and Radio Afterglows, and a Mysterious Late-time X-Ray Excess”. *Astrophys. J.* 780 (2014), p. 118. DOI: [10.1088/0004-637X/780/2/118](https://doi.org/10.1088/0004-637X/780/2/118). arXiv: [1309.7479](https://arxiv.org/abs/1309.7479) [astro-ph.HE] (cited on pp. 27, 42, 50, 65).
- [84] Francois Foucart. “Black Hole-Neutron Star Mergers: Disk Mass Predictions”. *Phys. Rev. D* 86 (2012), p. 124007. DOI: [10.1103/PhysRevD.86.124007](https://doi.org/10.1103/PhysRevD.86.124007). arXiv: [1207.6304](https://arxiv.org/abs/1207.6304) [astro-ph.HE] (cited on pp. 66, 72, 84).
- [85] T. J. Galama et al. “An unusual supernova in the error box of the gamma-ray burst of 25 April 1998”. *Nature* 395 (1998), p. 670. DOI: [10.1038/27150](https://doi.org/10.1038/27150). arXiv: [astro-ph/9806175](https://arxiv.org/abs/astro-ph/9806175) [astro-ph] (cited on pp. 22, 25).
- [86] T. J. Galama et al. “On the possible association of SN 1998bw and GRB 980425”. *Astron. Astrophys.* 138 (1999), pp. 465–466. DOI: [10.1051/aas:1999311](https://doi.org/10.1051/aas:1999311) (cited on p. 25).
- [87] Bruno Giacomazzo et al. “Compact Binary Progenitors of Short Gamma-Ray Bursts”. *Astrophys. J.* 762 (2013), p. L18. DOI: [10.1088/2041-8205/762/2/L18](https://doi.org/10.1088/2041-8205/762/2/L18). arXiv: [1210.8152](https://arxiv.org/abs/1210.8152) [astro-ph.HE] (cited on p. 66).
- [88] S. Golenetskii et al. “IPN triangulation and konus spectrum of the bright short/hard GRB051103.” *GRB Coordinates Network* 4197 (2005) (cited on p. 38).
-

- [89] S. Golenetskii et al. “IPN localization of very intense short GRB 070201.” *GRB Coordinates Network* 6088 (2007) (cited on p. 39).
- [90] S. Golenetskii et al. “Konus-Wind observation of GRB 150906B.” *GRB Coordinates Network* 18259 (2015) (cited on p. 102).
- [91] Dirk Grupe et al. “Jet breaks in short gamma-ray bursts. 1. the uncollimated afterglow of grb 050724”. *Astrophys. J.* 653 (2006), pp. 462–467. DOI: [10.1086/508739](https://doi.org/10.1086/508739). arXiv: [astro-ph/0603773](https://arxiv.org/abs/astro-ph/0603773) [[astro-ph](#)] (cited on p. 108).
- [92] A. Nicuesa Guelbenzu et al. “Multi-color observations of short GRB afterglows: 20 events observed between 2007 and 2010”. *Astron. Astrophys.* 548 (2012), A101. DOI: [10.1051/0004-6361/201219551](https://doi.org/10.1051/0004-6361/201219551). arXiv: [1206.1806](https://arxiv.org/abs/1206.1806) [[astro-ph.HE](#)] (cited on pp. 42, 50, 65).
- [93] Dafne Guetta and Tsvi Piran. “The batse-swift luminosity and redshift distributions of short-duration grbs”. *Astron. Astrophys.* 453 (2006), p. 823. DOI: [10.1051/0004-6361:20054498](https://doi.org/10.1051/0004-6361:20054498). arXiv: [astro-ph/0511239](https://arxiv.org/abs/astro-ph/0511239) [[astro-ph](#)] (cited on pp. 65, 66).
- [94] Mark Hannam et al. “A simple model of complete precessing black-hole-binary gravitational waveforms”. *Phys. Rev. Lett.* 113 (2014), p. 151101. DOI: [10.1103/PhysRevLett.113.151101](https://doi.org/10.1103/PhysRevLett.113.151101). arXiv: [1308.3271](https://arxiv.org/abs/1308.3271) [[gr-qc](#)] (cited on pp. 20, 86).
- [95] Gregory M. Harry. “Advanced LIGO: The next generation of gravitational wave detectors”. *Class. Quant. Grav.* 27 (2010), p. 084006. DOI: [10.1088/0264-9381/27/8/084006](https://doi.org/10.1088/0264-9381/27/8/084006) (cited on pp. 20, 89).
- [96] I. W. Harry and S. Fairhurst. “A coherent triggered search for single spin compact binary coalescences in gravitational wave data”. *Class. Quant. Grav.* 28 (2011), p. 134008. DOI: [10.1088/0264-9381/28/13/134008](https://doi.org/10.1088/0264-9381/28/13/134008). arXiv: [1101.1459](https://arxiv.org/abs/1101.1459) [[gr-qc](#)] (cited on p. 86).
- [97] I. W. Harry et al. “Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO”. *Phys. Rev. D* 89 (2014), p. 024010. DOI: [10.1103/PhysRevD.89.024010](https://doi.org/10.1103/PhysRevD.89.024010). arXiv: [1307.3562](https://arxiv.org/abs/1307.3562) [[gr-qc](#)] (cited on pp. 32, 85).
- [98] Ian W. Harry and Stephen Fairhurst. “A targeted coherent search for gravitational waves from compact binary coalescences”. *Phys. Rev. D* 83 (2011), p. 084002. DOI: [10.1103/PhysRevD.83.084002](https://doi.org/10.1103/PhysRevD.83.084002). arXiv: [1012.4939](https://arxiv.org/abs/1012.4939) [[gr-qc](#)] (cited on pp. 20, 31, 32, 36, 46, 50, 69, 74, 105).
- [99] Ian Harry et al. “Searching for Gravitational Waves from Compact Binaries with Precessing Spins”. *Phys. Rev. D* 94.2 (2016), p. 024012. DOI: [10.1103/PhysRevD.94.024012](https://doi.org/10.1103/PhysRevD.94.024012). arXiv: [1603.02444](https://arxiv.org/abs/1603.02444) [[gr-qc](#)] (cited on p. 86).
- [100] Jason W. T. Hessels et al. “A radio pulsar spinning at 716-hz”. *Science* 311 (2006), pp. 1901–1904. DOI: [10.1126/science.1123430](https://doi.org/10.1126/science.1123430). arXiv: [astro-ph/0601337](https://arxiv.org/abs/astro-ph/0601337) [[astro-ph](#)] (cited on p. 107).

-
- [101] Jens Hjorth et al. “A Very energetic supernova associated with the gamma-ray burst of 29 March 2003”. *Nature* 423 (2003), pp. 847–850. DOI: [10.1038/nature01750](https://doi.org/10.1038/nature01750). arXiv: [astro-ph/0306347](https://arxiv.org/abs/astro-ph/0306347) [[astro-ph](#)] (cited on p. 25).
- [102] R. A. Hulse and J. H. Taylor. “Discovery of a pulsar in a binary system”. *Astrophys. J.* 195 (1975), pp. L51–L53. DOI: [10.1086/181708](https://doi.org/10.1086/181708) (cited on p. 1).
- [103] K. Hurley et al. “Precise interplanetary network localization of the bursting pulsar gro j1744-28”. *Astrophys. J.* 537 (2000), p. 953. DOI: [10.1086/309083](https://doi.org/10.1086/309083). arXiv: [astro-ph/9912506](https://arxiv.org/abs/astro-ph/9912506) [[astro-ph](#)] (cited on p. 39).
- [104] K. Hurley et al. “The current performance of the third interplanetary network”. *AIP Conf. Proc.* 662 (2003), pp. 473–476. DOI: [10.1063/1.1579405](https://doi.org/10.1063/1.1579405). arXiv: [astro-ph/0201374](https://arxiv.org/abs/astro-ph/0201374) [[astro-ph](#)] (cited on pp. 28, 53).
- [105] K. Hurley et al. “IPN triangulation of GRB070201 (refined error box using MESSENGER data).” *GRB Coordinates Network* 6103 (2007) (cited on p. 39).
- [106] K. Hurley et al. “A new analysis of the short-duration, hard-spectrum GRB 051103, a possible extragalactic SGR giant flare”. *Mon. Not. Roy. Astron. Soc.* 403 (2010), p. 342. DOI: [10.1111/j.1365-2966.2009.16118.x](https://doi.org/10.1111/j.1365-2966.2009.16118.x). arXiv: [0907.2462](https://arxiv.org/abs/0907.2462) [[astro-ph.HE](#)] (cited on p. 38).
- [107] K. Hurley et al. “IPN Triangulation of GRB 150906B (short/intense).” *GRB Coordinates Network* 18258 (2015) (cited on p. 102).
- [108] Piotr Jaranowski, Andrzej Krolak, and Bernard F. Schutz. “Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection”. *Phys. Rev. D* 58 (1998), p. 063001. DOI: [10.1103/PhysRevD.58.063001](https://doi.org/10.1103/PhysRevD.58.063001). arXiv: [gr-qc/9804014](https://arxiv.org/abs/gr-qc/9804014) [[gr-qc](#)] (cited on p. 33).
- [109] B. W. Kernighan and D. M. Ritchie. *The C Programming Language*. Prentice-Hall software series. Prentice-Hall, 1978. ISBN: 9780131101630. URL: <https://books.google.co.uk/books?id=va1QAAAAMAAJ> (cited on p. 79).
- [110] Sebastian Khan et al. “Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era”. *Phys. Rev. D* 93.4 (2016), p. 044007. DOI: [10.1103/PhysRevD.93.044007](https://doi.org/10.1103/PhysRevD.93.044007). arXiv: [1508.07253](https://arxiv.org/abs/1508.07253) [[gr-qc](#)] (cited on p. 20).
- [111] Andreas von Kienlin et al. “The Second Fermi GBM Gamma-Ray Burst Catalog: The First Four Years”. *Astrophys. J. Suppl.* 211 (2014), p. 13. DOI: [10.1088/0067-0049/211/1/13](https://doi.org/10.1088/0067-0049/211/1/13). arXiv: [1401.5080](https://arxiv.org/abs/1401.5080) [[astro-ph.HE](#)] (cited on p. 28).
- [112] Bülent Kiziltan et al. “The Neutron Star Mass Distribution”. *Astrophys. J.* 778 (2013), p. 66. DOI: [10.1088/0004-637X/778/1/66](https://doi.org/10.1088/0004-637X/778/1/66). arXiv: [1309.6635](https://arxiv.org/abs/1309.6635) [[astro-ph.SR](#)] (cited on pp. 42, 107).
-

- [113] Ray W. Klebesadel, Ian B. Strong, and Roy A. Olson. “Observations of Gamma-Ray Bursts of Cosmic Origin”. *Astrophys. J.* 182 (1973), pp. L85–L88. DOI: [10.1086/181225](https://doi.org/10.1086/181225) (cited on p. 22).
- [114] Andreas Klöckner et al. “PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation”. *Parallel Computing* 38.3 (2012), pp. 157–174. ISSN: 0167-8191. DOI: [10.1016/j.parco.2011.09.001](https://doi.org/10.1016/j.parco.2011.09.001) (cited on p. 79).
- [115] Christopher S. Kochanek and Tsvi Piran. “Gravitational waves and gamma-ray bursts”. *Astrophys. J.* 417 (1993), pp. L17–L20. DOI: [10.1086/187083](https://doi.org/10.1086/187083). arXiv: [astro-ph/9305015](https://arxiv.org/abs/astro-ph/9305015) [[astro-ph](https://arxiv.org/abs/astro-ph)] (cited on p. 29).
- [116] Chryssa Kouveliotou et al. “Identification of two classes of gamma-ray bursts”. *Astrophys. J.* 413 (1993), pp. L101–104. DOI: [10.1086/186969](https://doi.org/10.1086/186969) (cited on p. 23).
- [117] H. A. Krimm et al. “GRB 100928A, Swift-BAT refined analysis.” *GRB Coordinates Network* 11312 (2010), p. 1 (cited on p. 46).
- [118] A. J. Levan, N. R. Tanvir, and J. Hjorth. “Short GRB 150906B: proximity to NGC 3313 galaxy group.” *GRB Coordinates Network* 18263 (2015) (cited on p. 102).
- [119] A. J. Levan et al. “A new population of ultra-long duration gamma-ray bursts”. *Astrophys. J.* 781 (2013), p. 13. DOI: [10.1088/0004-637X/781/1/13](https://doi.org/10.1088/0004-637X/781/1/13). arXiv: [1302.2352](https://arxiv.org/abs/1302.2352) [[astro-ph](https://arxiv.org/abs/astro-ph). HE] (cited on p. 23).
- [120] J. L. Levine and R. L. Garwin. “New negative result for gravitational wave detection, and comparison with reported detection”. *Phys. Rev. Lett.* 33 (1974), pp. 794–797. DOI: [10.1103/PhysRevLett.33.794](https://doi.org/10.1103/PhysRevLett.33.794) (cited on p. 1).
- [121] Li-Xin Li and Bohdan Paczynski. “Transient events from neutron star mergers”. *Astrophys. J.* 507 (1998), p. L59. DOI: [10.1086/311680](https://doi.org/10.1086/311680). arXiv: [astro-ph/9807272](https://arxiv.org/abs/astro-ph/9807272) [[astro-ph](https://arxiv.org/abs/astro-ph)] (cited on pp. 25, 26).
- [122] A. MacFadyen and S. E. Woosley. “Collapsars: Gamma-ray bursts and explosions in ‘failed supernovae’”. *Astrophys. J.* 524 (1999), p. 262. DOI: [10.1086/307790](https://doi.org/10.1086/307790). arXiv: [astro-ph/9810274](https://arxiv.org/abs/astro-ph/9810274) [[astro-ph](https://arxiv.org/abs/astro-ph)] (cited on p. 25).
- [123] D. M. Macleod, I. W. Harry, and S. Fairhurst. “A fully-coherent all-sky search for gravitational-waves from compact binary coalescences” (2015). arXiv: [1509.03426](https://arxiv.org/abs/1509.03426) [[gr-qc](https://arxiv.org/abs/gr-qc)] (cited on p. 36).
- [124] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. OUP Oxford, 2007. ISBN: 9780198570745. URL: <https://books.google.it/books?id=AqVpQgAACAAJ> (cited on pp. 7, 19).
- [125] Charles Meegan et al. “The Fermi Gamma-Ray Burst Monitor”. *Astrophys. J.* 702 (2009), pp. 791–804. DOI: [10.1088/0004-637X/702/1/791](https://doi.org/10.1088/0004-637X/702/1/791). arXiv: [0908.0450](https://arxiv.org/abs/0908.0450) [[astro-ph](https://arxiv.org/abs/astro-ph). IM] (cited on pp. 28, 38, 53, 56).

-
- [126] Peter Meszaros. “Gamma Ray Bursts”. *Astropart. Phys.* 43 (2013), pp. 134–141. DOI: [10.1016/j.astropartphys.2012.03.009](https://doi.org/10.1016/j.astropartphys.2012.03.009). arXiv: [1204.1897](https://arxiv.org/abs/1204.1897) [[astro-ph.HE](#)] (cited on pp. [i](#), [23](#), [24](#)).
- [127] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. Gravitation 3. W. H. Freeman, 1973. ISBN: 9780716703440. URL: <https://books.google.co.uk/books?id=w4Gigq3tY1kC> (cited on p. [7](#)).
- [128] Ehud Nakar. “Short-Hard Gamma-Ray Bursts”. *Phys. Rept.* 442 (2007), pp. 166–236. DOI: [10.1016/j.physrep.2007.02.005](https://doi.org/10.1016/j.physrep.2007.02.005). arXiv: [astro-ph/0701748](https://arxiv.org/abs/astro-ph/0701748) [[astro-ph](#)] (cited on pp. [27](#), [65](#)).
- [129] Ramesh Narayan, Bohdan Paczynski, and Tsvi Piran. “Gamma-ray bursts as the death throes of massive binary stars”. *Astrophys. J.* 395 (1992), pp. L83–L86. DOI: [10.1086/186493](https://doi.org/10.1086/186493). arXiv: [astro-ph/9204001](https://arxiv.org/abs/astro-ph/9204001) [[astro-ph](#)] (cited on p. [25](#)).
- [130] Samaya Nissanke et al. “Exploring short gamma-ray bursts as gravitational-wave standard sirens”. *Astrophys. J.* 725 (2010), pp. 496–514. DOI: [10.1088/0004-637X/725/1/496](https://doi.org/10.1088/0004-637X/725/1/496). arXiv: [0904.1017](https://arxiv.org/abs/0904.1017) [[astro-ph.CO](#)] (cited on p. [73](#)).
- [131] Alexander H. Nitz et al. “Accuracy of gravitational waveform models for observing neutron-star–black-hole binaries in Advanced LIGO”. *Phys. Rev. D* 88.12 (2013), p. 124039. DOI: [10.1103/PhysRevD.88.124039](https://doi.org/10.1103/PhysRevD.88.124039). arXiv: [1307.1757](https://arxiv.org/abs/1307.1757) [[gr-qc](#)] (cited on pp. [20](#), [107](#), [108](#)).
- [132] Benjamin J. Owen and B. S. Sathyaprakash. “Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement”. *Phys. Rev. D* 60 (1999), p. 022002. DOI: [10.1103/PhysRevD.60.022002](https://doi.org/10.1103/PhysRevD.60.022002). arXiv: [gr-qc/9808076](https://arxiv.org/abs/gr-qc/9808076) [[gr-qc](#)] (cited on pp. [31](#), [67](#)).
- [133] Feryal Ozel et al. “On the Mass Distribution and Birth Masses of Neutron Stars”. *Astrophys. J.* 757 (2012), p. 55. DOI: [10.1088/0004-637X/757/1/55](https://doi.org/10.1088/0004-637X/757/1/55). arXiv: [1201.1006](https://arxiv.org/abs/1201.1006) [[astro-ph.HE](#)] (cited on pp. [42](#), [107](#)).
- [134] W. S. Paciesas et al. “The Fourth batse gamma-ray burst catalog (revised)”. *Astrophys. J. Suppl.* 122 (1999), pp. 465–495. DOI: [10.1086/313224](https://doi.org/10.1086/313224). arXiv: [astro-ph/9903205](https://arxiv.org/abs/astro-ph/9903205) [[astro-ph](#)] (cited on pp. [22](#), [24](#)).
- [135] Bohdan Paczyński. “Gamma-ray bursters at cosmological distances”. *Astrophys. J.* 308 (1986), pp. L43–L46 (cited on p. [25](#)).
- [136] Bohdan Paczyński. “Are gamma-ray bursts in star forming regions?” *Astrophys. J.* 494 (1998), p. L45. DOI: [10.1086/311148](https://doi.org/10.1086/311148). arXiv: [astro-ph/9710086](https://arxiv.org/abs/astro-ph/9710086) [[astro-ph](#)] (cited on p. [25](#)).
- [137] Bohdan Paczynski and James E. Rhoads. “Radio transients from gamma-ray bursters”. *Astrophys. J.* 418 (1993), p. L5. DOI: [10.1086/187102](https://doi.org/10.1086/187102). arXiv: [astro-ph/9307024](https://arxiv.org/abs/astro-ph/9307024) [[astro-ph](#)] (cited on p. [23](#)).
-

- [138] Carlos Palenzuela et al. “Electromagnetic and Gravitational Outputs from Binary-Neutron-Star Coalescence”. *Phys. Rev. Lett.* 111.6 (2013), p. 061105. DOI: [10.1103/PhysRevLett.111.061105](https://doi.org/10.1103/PhysRevLett.111.061105). arXiv: [1301.7074](https://arxiv.org/abs/1301.7074) [[gr-qc](#)] (cited on p. 84).
- [139] V. D. Pal’shin et al. “IPN localizations of Konus short gamma-ray bursts”. *Astrophys. J. S* 207 (2013), p. 38. DOI: [10.1088/0067-0049/207/2/38](https://doi.org/10.1088/0067-0049/207/2/38). arXiv: [1301.3740](https://arxiv.org/abs/1301.3740) [[astro-ph.HE](#)] (cited on p. 40).
- [140] Yi Pan et al. “A Physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single spin binaries”. *Phys. Rev. D* 69 (2004). [Erratum: *Phys. Rev. D* 74,029905(2006)], p. 104017. DOI: [10.1103/PhysRevD.69.104017](https://doi.org/10.1103/PhysRevD.69.104017), [10.1103/PhysRevD.74.029905](https://doi.org/10.1103/PhysRevD.74.029905). arXiv: [gr-qc/0310034](https://arxiv.org/abs/gr-qc/0310034) [[gr-qc](#)] (cited on pp. 20, 106).
- [141] Alin Panaitescu. “The energetics and environment of the short-grb afterglows 050709 and 050724”. *Mon. Not. R. Astron. Soc. Lett.* 367 (2006), pp. L42–L46. DOI: [10.1111/j.1745-3933.2005.00134.x](https://doi.org/10.1111/j.1745-3933.2005.00134.x). arXiv: [astro-ph/0511588](https://arxiv.org/abs/astro-ph/0511588) [[astro-ph](#)] (cited on pp. 50, 65).
- [142] Francesco Pannarale and Frank Ohme. “Prospects for joint gravitational-wave and electromagnetic observations of neutron-star–black-hole coalescing binaries”. *Astrophys. J.* 791 (2014), p. L7. DOI: [10.1088/2041-8205/791/1/L7](https://doi.org/10.1088/2041-8205/791/1/L7). arXiv: [1406.6057](https://arxiv.org/abs/1406.6057) [[gr-qc](#)] (cited on pp. 47, 66, 84, 85, 107, 108).
- [143] Tsvi Piran, Ehud Nakar, and Stephan Rosswog. “The Electromagnetic Signals of Compact Binary Mergers”. *Mon. Not. Roy. Astron. Soc.* 430.3 (2013), pp. 2121–2136. DOI: [10.1093/mnras/stt037](https://doi.org/10.1093/mnras/stt037). arXiv: [1204.6242](https://arxiv.org/abs/1204.6242) [[astro-ph.HE](#)] (cited on p. 67).
- [144] V. Predoi and Kevin Hurley. “Search for gravitational waves associated with the InterPlanetary Network short gamma ray bursts”. *J. Phys. Conf. Ser.* 363 (2012), p. 012034. DOI: [10.1088/1742-6596/363/1/012034](https://doi.org/10.1088/1742-6596/363/1/012034). arXiv: [1112.1637](https://arxiv.org/abs/1112.1637) [[gr-qc](#)] (cited on pp. 39, 59).
- [145] Olivier Rabaste, Eric Chassande-Mottin, and Archana Pai. “Sparse sky grid for the coherent detection of gravitational wave bursts” (2009). arXiv: [0905.4832](https://arxiv.org/abs/0905.4832) [[gr-qc](#)] (cited on p. 55).
- [146] J. L. Racusin et al. “Jet breaks and Energetics of Swift GRB X-ray Afterglows”. *Astrophys. J.* 698 (2009), pp. 43–74. DOI: [10.1088/0004-637X/698/1/43](https://doi.org/10.1088/0004-637X/698/1/43). arXiv: [0812.4780](https://arxiv.org/abs/0812.4780) [[astro-ph](#)] (cited on pp. 108, 111).
- [147] M. J. Rees and P. Meszaros. “Unsteady outflow models for cosmological gamma-ray bursts”. *Astrophys. J.* 430 (1994), pp. L93–L96. DOI: [10.1086/187446](https://doi.org/10.1086/187446). arXiv: [astro-ph/9404038](https://arxiv.org/abs/astro-ph/9404038) [[astro-ph](#)] (cited on p. 25).
- [148] T. Regimbau et al. “Revisiting coincidence rate between Gravitational Wave detection and short Gamma-Ray Burst for the Advanced and third generation”. *Astrophys. J.* 799.1 (2015), p. 69. DOI: [10.1088/0004-637X/799/1/69](https://doi.org/10.1088/0004-637X/799/1/69). arXiv: [1410.2739](https://arxiv.org/abs/1410.2739) [[astro-ph.HE](#)] (cited on pp. 65, 72).

-
- [149] A. Rowlinson et al. “Discovery of the afterglow and host galaxy of the low redshift short GRB 080905A”. *Mon. Not. Roy. Astron. Soc.* 408 (2010), pp. 383–391. arXiv: 1006.0487 [astro-ph.HE] (cited on pp. 29, 73).
- [150] R. Ruffini et al. “GRB 150906B: theoretical estimation of redshift and isotropic energy.” *GRB Coordinates Network* 18296 (2015) (cited on p. 103).
- [151] Re’em Sari, Tsvi Piran, and Jules Halpern. “Jets in GRBs”. *Astrophys. J.* 519 (1999), p. L17. DOI: 10.1086/312109. arXiv: astro-ph/9903339 [astro-ph] (cited on pp. 27, 65).
- [152] B. S. Sathyaprakash and B. F. Schutz. “Physics, Astrophysics and Cosmology with Gravitational Waves”. *Living Rev. Rel.* 12 (2009), p. 2. DOI: 10.12942/lrr-2009-2. arXiv: 0903.0338 [gr-qc] (cited on p. 21).
- [153] G. Schilling. *Flash!: The Hunt for the Biggest Explosions in the Universe*. Cambridge University Press, 2002. ISBN: 9780521800532. URL: <https://books.google.co.uk/books?id=6mGGX6AGBTAC> (cited on p. 22).
- [154] Patricia Schmidt, Mark Hannam, and Sascha Husa. “Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals”. *Phys. Rev. D* 86 (2012), p. 104063. DOI: 10.1103/PhysRevD.86.104063. arXiv: 1207.3088 [gr-qc] (cited on p. 85).
- [155] B. F. Schutz. *A First Course in General Relativity*. Cambridge University Press, 2009. ISBN: 9781139479004. URL: <https://books.google.co.uk/books?id=GgRRt7AbdwQC> (cited on p. 7).
- [156] Bernard F. Schutz. “Determining the Hubble Constant from Gravitational Wave Observations”. *Nature* 323 (1986), pp. 310–311. DOI: 10.1038/323310a0 (cited on p. 73).
- [157] Masaru Shibata and Keisuke Taniguchi. “Merger of black hole and neutron star in general relativity: Tidal disruption, torus mass, and gravitational waves”. *Phys. Rev. D* 77 (2008), p. 084015. DOI: 10.1103/PhysRevD.77.084015. arXiv: 0711.1410 [gr-qc] (cited on p. 29).
- [158] Masaru Shibata and Keisuke Taniguchi. “Coalescence of Black Hole-Neutron Star Binaries”. *Living Rev. Rel.* 14 (2011), p. 6. DOI: 10.12942/lrr-2011-6 (cited on p. 42).
- [159] K. Siellez, M. Boër, and B. Gendre. “Simultaneous event detection rates by electromagnetic and gravitational wave detectors in the Advanced Era of LIGO and Virgo”. *Mon. Not. Roy. Astron. Soc.* 437.1 (2014), pp. 649–655. DOI: 10.1093/mnras/stt1915. arXiv: 1310.2106 [astro-ph.HE] (cited on pp. 64, 65).
- [160] Joshua R Smith. “The Path to the enhanced and advanced LIGO gravitational-wave detectors”. *Class. Quant. Grav.* 26 (2009), p. 114013. DOI: 10.1088/0264-9381/26/11/114013. arXiv: 0902.0381 [gr-qc] (cited on p. 14).
-

- [161] “Status of TAMA project”. *Class. Quant. Grav.* 14 (1997), pp. 1477–1480. DOI: [10.1088/0264-9381/14/6/013](https://doi.org/10.1088/0264-9381/14/6/013) (cited on p. 14).
- [162] Nicholas Stone, Abraham Loeb, and Edo Berger. “Pulsations in Short GRBs from Black Hole-Neutron Star Mergers”. *Phys. Rev. D* 87 (2013), p. 084053. DOI: [10.1103/PhysRevD.87.084053](https://doi.org/10.1103/PhysRevD.87.084053). arXiv: [1209.4097](https://arxiv.org/abs/1209.4097) [[astro-ph.HE](#)] (cited on pp. 66, 84).
- [163] D. S. Svinkin et al. “A search for giant flares from soft gamma-repeaters in nearby galaxies in the Konus-Wind short burst sample”. *Mon. Not. Roy. Astron. Soc.* 447 (2015), p. 1032. DOI: [10.1093/mnras/stu2436](https://doi.org/10.1093/mnras/stu2436). arXiv: [1411.5589](https://arxiv.org/abs/1411.5589) [[astro-ph.HE](#)] (cited on pp. 38, 39).
- [164] D. Svinkin, D. Frederiks, and V. Pal’Shin. “GRB 150906B: Further analysis of the Konus-Wind data.” *GRB Coordinates Network* 18321 (2015) (cited on p. 102).
- [165] N. R. Tanvir et al. “A “kilonova” associated with short-duration gamma-ray burst 130603B”. *Nature* 500 (2013), pp. 547–549. DOI: [10.1038/nature12505](https://doi.org/10.1038/nature12505). arXiv: [1306.4971](https://arxiv.org/abs/1306.4971) [[astro-ph.HE](#)] (cited on pp. i, 25, 26).
- [166] Andrea Taracchini et al. “Effective-one-body model for black-hole binaries with generic mass ratios and spins”. *Phys. Rev. D* 89.6 (2014), p. 061502. DOI: [10.1103/PhysRevD.89.061502](https://doi.org/10.1103/PhysRevD.89.061502). arXiv: [1311.2544](https://arxiv.org/abs/1311.2544) [[gr-qc](#)] (cited on p. 20).
- [167] J. H. Taylor and J. M. Weisberg. “A new test of general relativity: Gravitational radiation and the binary pulsar PS R 1913+16”. *Astrophys. J.* 253 (1982), pp. 908–920. DOI: [10.1086/159690](https://doi.org/10.1086/159690) (cited on p. 1).
- [168] Matt Taylor et al. “The Core Collapse Supernova Rate from the SDSS-II Supernova Survey”. *Astrophys. J.* 792 (2014), p. 135. DOI: [10.1088/0004-637X/792/2/135](https://doi.org/10.1088/0004-637X/792/2/135). arXiv: [1407.0999](https://arxiv.org/abs/1407.0999) [[astro-ph.SR](#)] (cited on p. 21).
- [169] C. Tinney et al. “GRB 980425”. *IAU Circ.* 6896 (1998) (cited on p. 22).
- [170] Kenji Toma, Xue-Feng Wu, and Peter Meszaros. “A Photosphere-Internal Shock Model of Gamma-Ray Bursts: Case Studies of Fermi/LAT Bursts”. *Mon. Not. Roy. Astron. Soc.* 415 (2011), pp. 1663–1680. DOI: [10.1111/j.1365-2966.2011.18807.x](https://doi.org/10.1111/j.1365-2966.2011.18807.x). arXiv: [1002.2634](https://arxiv.org/abs/1002.2634) [[astro-ph.HE](#)] (cited on pp. 23, 24).
- [171] Samantha A. Usman et al. “An improved pipeline to search for gravitational waves from compact binary coalescence” (2015). arXiv: [1508.02357](https://arxiv.org/abs/1508.02357) [[gr-qc](#)] (cited on pp. 79, 105).
- [172] L. A. Wainstein and V. D. Zubakov. *Extraction of Signals from Noise*. Prentice-Hall, Englewood Cliffs, NJ, 1962. URL: <https://books.google.co.uk/books?id=krU8AAAAIAAJ> (cited on pp. 20, 31).
- [173] David Wanderman and Tsvi Piran. “The rate, luminosity function and time delay of non-Collapsar short GRBs”. *Mon. Not. Roy. Astron. Soc.* 448.4 (2015), pp. 3026–3037. DOI: [10.1093/mnras/stv123](https://doi.org/10.1093/mnras/stv123). arXiv: [1405.5878](https://arxiv.org/abs/1405.5878) [[astro-ph.HE](#)] (cited on pp. 64, 65, 72).

-
- [174] Michal Was et al. “On the background estimation by time slides in a network of gravitational wave detectors”. *Class. Quant. Grav.* 27 (2010), p. 015005. DOI: [10.1088/0264-9381/27/1/015005](https://doi.org/10.1088/0264-9381/27/1/015005). arXiv: [0906.2120](https://arxiv.org/abs/0906.2120) [[gr-qc](#)] (cited on p. 50).
- [175] J. Weber. “Detection and Generation of Gravitational Waves”. *Phys. Rev.* 117 (1960), pp. 306–313. DOI: [10.1103/PhysRev.117.306](https://doi.org/10.1103/PhysRev.117.306) (cited on p. 1).
- [176] J. Weber. “Observation of the Thermal Fluctuations of a Gravitational-Wave Detector”. *Phys. Rev. Lett.* 17 (1966), pp. 1228–1230. DOI: [10.1103/PhysRevLett.17.1228](https://doi.org/10.1103/PhysRevLett.17.1228) (cited on p. 1).
- [177] J. Weber. “Evidence for discovery of gravitational radiation”. *Phys. Rev. Lett.* 22 (1969), pp. 1320–1324. DOI: [10.1103/PhysRevLett.22.1320](https://doi.org/10.1103/PhysRevLett.22.1320) (cited on p. 1).
- [178] Joel M. Weisberg and Joseph H. Taylor. “Relativistic binary pulsar B1913+16: Thirty years of observations and analysis”. *ASP Conf. Ser.* 328 (2005), p. 25. arXiv: [astro-ph/0407149](https://arxiv.org/abs/astro-ph/0407149) [[astro-ph](#)] (cited on pp. 1, 2).
- [179] J. A. Wheeler. *A Journey Into Gravity and Spacetime*. Library series. Scientific American Library, 1990. ISBN: 9780716750161. URL: https://books.google.co.uk/books?id=Wn_FQgAACAAJ (cited on p. 1).
- [180] A. R. Williamson. “An External Trigger Alert System for the LIGO Control Rooms” (2015). URL: <https://dcc.ligo.org/LIGO-T1500197/public> (cited on p. xi).
- [181] A. R. Williamson et al. “Improved methods for detecting gravitational waves associated with short gamma-ray bursts”. *Phys. Rev. D* 90.12 (2014), p. 122004. DOI: [10.1103/PhysRevD.90.122004](https://doi.org/10.1103/PhysRevD.90.122004). arXiv: [1410.6042](https://arxiv.org/abs/1410.6042) [[gr-qc](#)] (cited on p. xi).
- [182] B. Willke et al. “The GEO 600 gravitational wave detector”. *Class. Quant. Grav.* 19 (2002), pp. 1377–1387. DOI: [10.1088/0264-9381/19/7/321](https://doi.org/10.1088/0264-9381/19/7/321) (cited on p. 14).
- [183] S. E. Woosley. “Gamma-ray bursts from stellar mass accretion disks around black holes”. *Astrophys. J.* 405 (1993), p. 273. DOI: [10.1086/172359](https://doi.org/10.1086/172359) (cited on p. 25).
- [184] S. E. Woosley and J. S. Bloom. “The Supernova Gamma-Ray Burst Connection”. *Ann. Rev. Astron. Astrophys.* 44 (2006), pp. 507–556. DOI: [10.1146/annurev.astro.43.072103.150558](https://doi.org/10.1146/annurev.astro.43.072103.150558). arXiv: [astro-ph/0609142](https://arxiv.org/abs/astro-ph/0609142) [[astro-ph](#)] (cited on p. 25).
- [185] Bin-Bin Zhang et al. “How Long does a Burst Burst?” *Astrophys. J.* 787 (2014), p. 66. DOI: [10.1088/0004-637X/787/1/66](https://doi.org/10.1088/0004-637X/787/1/66). arXiv: [1310.2540](https://arxiv.org/abs/1310.2540) [[astro-ph.HE](#)] (cited on p. 23).
- [186] F. W. Zhang, B. Zhang, and B. Zhang. “GRB 150906B - NGC 3313 association: Constraints from Ep-Eiso and Ep-Liso Relations.” *GRB Coordinates Network* 18298 (2015) (cited on p. 103).
- [187] user:Inductiveload. Wikimedia commons. Public domain. 2010. URL: <https://commons.wikimedia.org/w/index.php?curid=9538634> (cited on p. 2).
-