

HIGH-ENERGY ELECTROMAGNETIC OFFLINE FOLLOW-UP OF LIGO-VIRGO GRAVITATIONAL-WAVE BINARY COALESCENCE CANDIDATE EVENTS

L. BLACKBURN^{1,2}, M. S. BRIGGS³, J. CAMP¹, N. CHRISTENSEN⁴, V. CONNAUGHTON³, P. JENKE³, R. A. REMILLARD⁶, AND J. VEITCH⁷

¹NASA/Goddard Space Flight Center, Greenbelt, MD, USA

²University of Maryland/CRESST, College Park, MD, USA

³University of Alabama in Huntsville, Huntsville, AL, USA

⁴Carleton College, Northfield, MN, USA

⁵NASA/Marshall Space Flight Center, Huntsville, AL, USA

⁶Massachusetts Institute of Technology, Cambridge, MA, USA

⁷University of Birmingham, Birmingham, UK

Received 2014 September 16; accepted 2015 January 24; published 2015 March 5

ABSTRACT

We present two different search methods for electromagnetic counterparts to gravitational-wave (GW) events from ground-based detectors using archival NASA high-energy data from the Fermi Gamma-ray Burst Monitor (GBM) and *RXTE* All-sky Monitor (ASM) instruments. To demonstrate the methods, we use a limited number of representative GW background noise events produced by a search for binary neutron star coalescence over the last two months of the LIGO-Virgo S6/VSR3 joint science run. Time and sky location provided by the GW data trigger a targeted search in the high-energy photon data. We use two custom pipelines: one to search for prompt gamma-ray counterparts in GBM, and the other to search for a variety of X-ray afterglow model signals in ASM. We measure the efficiency of the joint pipelines to weak gamma-ray burst counterparts, and a family of model X-ray afterglows. By requiring a detectable signal in either electromagnetic instrument coincident with a GW event, we are able to reject a large majority of GW candidates. This reduces the signal-to-noise ratio of the loudest surviving GW background event by around 15–20%.

Key words: gravitational waves – gamma-ray burst: general

1. INTRODUCTION

1.1. Initial and Advanced LIGO and Virgo Detectors

The initial LIGO (Abbott et al. 2009) and Virgo (Accadia et al. 2012) gravitational-wave (GW) detectors took their last science data between 2009 July and 2010 October before going offline for several years for major upgrades to advanced detector configurations (Acernese 2009; Harry 2010). For its sixth science run (S6), LIGO consisted of two 4 km baseline interferometric detectors—one instrument H1 at Hanford, WA, and another instrument L1 at Livingston, LA. The Virgo 3 km interferometer V1 in Cascina, Italy took concurrent data during their second (VSR2 covering 2009 July–2010 January) and third (VSR3 covering 2010 August–October) science runs. Together, they represent the most sensitive worldwide network of GW detectors to date. Figure 1 shows typical strain sensitivity of each instrument as a function of frequency, as well as the distance at which an optimally oriented merger of two compact objects (black holes or neutron stars) would produce a nominal signal-to-noise ratio (S/N) of eight in each detector.

Compact binary coalescence (CBC) is the most anticipated GW source for first and second-generation ground-based detectors. These systems are very strong emitters of GWs, and we are confident of their existence through the discovery of a handful of galactic NS/NS binary systems where one object is a radio pulsar which is modulated by the orbit—the most famous of these systems is the Hulse–Taylor pulsar PSR1316+16 (Weisberg & Taylor 2005). The number and lifetime of these systems can be used to obtain an estimate of about one NS/NS merger event per Mpc^3 per million years (Abadie et al. 2010), with up to two orders of magnitude uncertainty due largely to our limited knowledge of the

pulsar luminosity function and limited statistics. The merger rate translates into an estimate of ~ 0.02 detectable NS/NS merger events per year for the initial detector network, and ~ 40 per year for the advanced detectors once they reach design sensitivity. NS/BH systems are also of interest for electromagnetic (EM) follow-up and are stronger GW emitters. However we have not yet observed any NS/BH binary systems, and have generally poor knowledge of the black hole mass distribution.

While no NS/NS or NS/BH binary coalescence events have been detected during S6/VSR2+3 using GW data alone (Abadie et al. 2012), it is conceivable that an EM counterpart might allow us to resolve a rare event otherwise too weak to distinguish from background in the first generation detectors. In the rapidly approaching advanced detector era, searching for EM counterparts can play an important role in increasing our confidence of detection for otherwise marginal events and will provide astrophysical context for GW detections.

Advanced LIGO is likely to begin taking its first science data in 2015, with Virgo following a year later (Aasi et al. 2013). As they reach design sensitivity, the advanced detectors are expected to begin an era of regular detection of GW events, making the search for EM counterparts triggered by GW detections an enticing possibility. In this study, we demonstrate strategies for searching high-energy archival EM data for counterparts using representative GW background events from real data taken during the most recent LIGO-Virgo science runs and high-energy EM survey data recorded at the same time. We then measure EM detection efficiencies under various plausible emission models and the probability of accidental coincidence given realistic GW background and sky-localization accuracy.

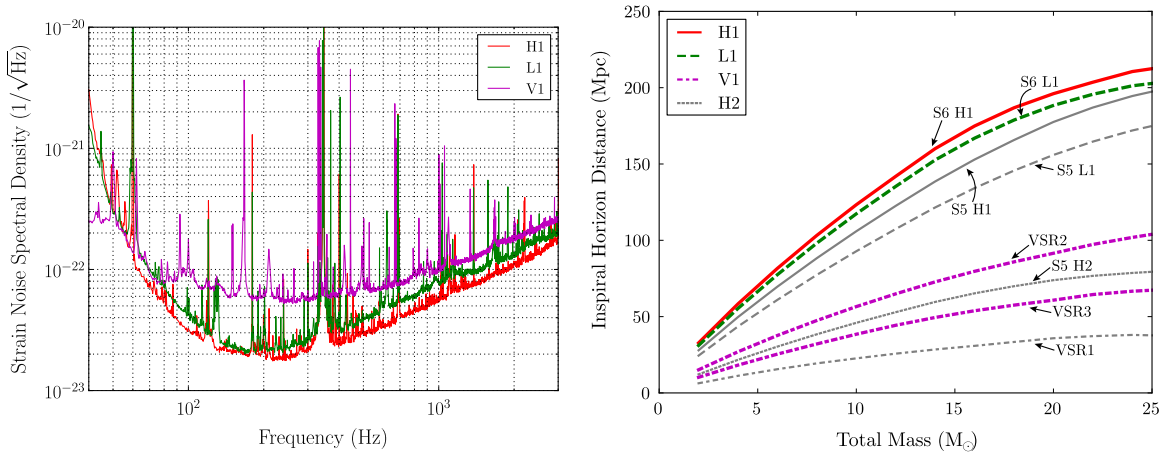


Figure 1. Typical detector strain noise spectral density for the LIGO S6 and Virgo VSR2+3 runs, as well as the best equal-mass binary coalescence horizon distance achieved for each run as a function of total system mass. The horizon distance is the distance at which an optimally oriented binary merger would produce an expected signal-to-noise ratio of 8. Figures reproduced from (Abadie et al. 2012).

1.2. High-energy Photon Survey Instruments

The two instruments chosen for this EM follow-up study include the Gamma-ray Burst Monitor (GBM; Meegan et al. 2009) on board the *Fermi* spacecraft, and the soft X-ray All-sky Monitor (ASM; Levine et al. 1996) on board *RXTE*. Both are particularly well suited to an offline follow-up of GW events because of their large, regular coverage of the entire sky (Table 1), and because they save a large amount of archival survey data which allows for sensitive offline searches.

The high-energy sky itself offers the advantage of being relatively clean compared to optical wavelengths, which is important for maintaining a low probability of false-coincidence given the large GW sky-location uncertainty. The use of offline survey data for follow-up has many practical differences from other *triggered* searches in the EM spectrum for GW counterparts. (Evans et al. 2012; Kanner et al. 2012; Aasi et al. 2014). The offline search does not rely on rapid GW data analysis and continuous coordination with EM observational facilities. An important consequence is that the offline coincidence can generally be applied to many more events, allowing the follow-up of weaker, more marginal candidates.

1.3. Short Gamma-ray Bursts (sGRBs) and Afterglows as Counterparts to GWs

Gamma-ray bursts (GRBs) are flashes of gamma-rays observed approximately once per day. Their isotropic distribution in the sky was the first evidence of an extra-galactic origin, and indicated that they were extremely energetic events. The duration of prompt gamma-ray emission shows a bi-modal distribution which naturally groups GRBs into two categories (Kouveliotou et al. 1993). Long GRBs emit 90% of their prompt radiation on timescales longer than 2 s. They have been associated with the collapse of rapidly rotating massive stars (Hjorth & Bloom 2012). sGRBs with prompt emission < 2 s and a generally harder spectrum are thought to arise from the merger of two neutron stars, or a neutron star and black hole (Nakar 2007). It is this favored progenitor model which makes sGRBs and associated afterglow emission a promising counterpart to GW observations.

The Swift satellite has revolutionized our understanding of short GRBs over the last several years by the rapid observation of X-ray afterglows, providing the first localization, host

identification, and redshift information. The beaming angle for short GRBs is highly uncertain, although limited observations of jet breaks in some afterglows imply half-opening angles of $\theta_j \sim 3^\circ\text{--}14^\circ$ (Liang et al. 2008; Fong et al. 2012). The absence of an observable jet break sets a lower limit on the opening angle which is generally weak (due to limits in sensitivity), though in the case of GRB 050724A, late-time *Chandra* observations were able to constrain $\theta_j \gtrsim 25^\circ$ (Grupe et al. 2006). The observed spatial density of short GRBs and limits on beaming angle result in an NS/NS merger event rate roughly consistent with that derived from galactic binary pulsar measurements.

Although the beaming factor of $\sim \theta_j^2/2$ means most merger events seen by the advanced GW detectors will not be accompanied by a GRB, this is somewhat compensated by the fact that the ones that are beamed toward us have stronger GW emission. Current estimates for coincident GW-sGRB observation for advanced LIGO-Virgo are a few per year assuming an NS/NS progenitor model (Coward et al. 2012; Metzger & Berger 2012). The rate increases by a factor of eight if all observed short GRBs are instead due to NS/BH ($10 M_\odot$) mergers which are detectable in GWs to twice the distance. Coincident rates for the initial detectors go down by a factor a thousand due to their factor of ten worse sensitivity.

About 80% of short GRBs seen by Swift are accompanied by some kind of X-ray afterglow (Gehrels & Razzaque 2013), which are often observable for about a day. The observed X-ray afterglows are thought to occur from synchrotron emission at the shock front where the outgoing jet meets the local interstellar medium. Simulations show that such afterglow emission becomes very weak off-axis, with little possibility of detection at twice the opening angle (van Eerten & MacFadyen 2011). Searches for orphan afterglow signals, without the presence of detected prompt GRB emission, are quite difficult to confirm due to various sources of transient background. However with incomplete coverage of the gamma-ray sky, it is quite possible that the first GW-EM association will be with an afterglow signal.

In addition to the jet-driven burst and afterglow, other EM emission associated with a compact merger can be a promising channel for GW-EM coincidence, particularly if the EM radiation is less-beamed or even isotropic. A few short GRBs

(~10%) have shown clear evidence of high-energy flares which precede the primary burst by 1–10 s, and possibly up to 100 s (Troja et al. 2010). The precursors can be interpreted as evidence of some activity during or before merger, such as the resonant shattering of NS crusts (Tsang et al. 2012), which depends on the NS equation-of-state and could radiate isotropically. A class of short GRBs (~25%) also contain a period of extended soft X-ray emission on a timescale of ~10–100 s after the initial spike, which is difficult to explain with the standard jet scenario. It has been proposed that the extended emission could arise from a relativistic wind powered by a short-lived rapidly rotating protomagnetar star surviving post-merger (Metzger et al. 2008; Gao et al. 2013; Nakamura et al. 2013; Zhang 2013), which could also be considerably more isotropic than the jet afterglow. Since short GRBs themselves may be subject to a large beaming factor, the more numerous nearby NS/NS mergers seen in GWs will provide a unique opportunity to search for these potentially weak exotic EM phenomena, and isolate them from jet behavior.

Detection of the characteristic signature of a CBC in GWs prior to an associated sGRB will provide unambiguous support for the compact merger progenitor model. Along with this critical piece of the sGRB puzzle, GWs provide a largely complimentary set of information: component masses (whether black holes are involved) and spins, system inclination, luminosity distance; while the EM counterpart provides information about EM energetics, a precise location, local and host environment, and redshift. This makes joint GW-EM detections particularly valuable and a key goal in GW astronomy.

1.4. Methodology of GW-EM Coincidence Search

GW-EM coincidence can be approached from a variety of angles. In this analysis, we use a collection of all-sky, all-time GW events to conduct a targeted search in high-energy EM archival survey data. Another strategy demands real-time analysis and localization of the GW data (Abadie et al. 2012; Cannon et al. 2012; Singer et al. 2014), suitable for targeted observation with narrow-field instruments with the goal of catching a short-lived afterglow signal. This rapid EM-follow-up approach was successfully tested in initial LIGO/Virgo for a handful of online GW events using both a collection of ground-based optical telescopes (Aasi et al. 2014) and Swift-XRT (Evans et al. 2012). The online strategy places heavy demands on both the rapid processing and interpretation of GW data, as well as the capability of pointed telescopes to effectively cover the large GW-determined sky location (Abbott et al. 2012; Kanner et al. 2012; Singer et al. 2013). In return they provide the deepest observations, with sufficient sensitivity to observe, for example, the faint but promising optical/IR kilonova merger counterpart (Metzger & Berger 2012; Kasliwal & Nissanke 2014).

Known EM transients can also trigger specialized searches in the GW data itself, where the precise time, sky-location, and potentially model information from the EM event may greatly constrain the space of GW signals to look for. The reduced search can be performed at a lower threshold in the GW data due to both computational and statistical (false-alarm probability) considerations (Was et al. 2012). In the case of known GRBs, such searches have been used to place minimum distance limits dependent on GW emission model (Abadie et al. 2012a; Aasi et al. 2014), and for two short GRBs have

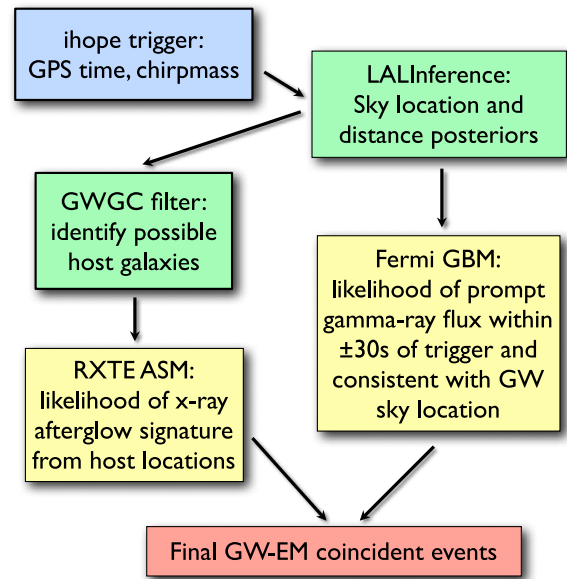


Figure 2. Flowchart for GW-EM coincidence analysis. We begin with GW compact binary coalescence events detected with a standard matched-filter analysis (ihope) (Babak et al. 2013). A coherent Bayesian follow-up using the LALInference package (Veitch et al. 2014) estimates distance and sky location parameters using data from multiple detectors. From this, we obtain a list of possible nearby host galaxies from the galaxy catalog (GWGC) and their relative probabilities. We search nearby GBM data for prompt gamma-ray bursts consistent with the GW skymap, as well as ASM light curve data at the positions of potential host galaxies for a family of parameterized afterglow signals beginning at the coalescence time. From this we gather a filtered list of GW-EM candidate events.

been able to uniquely exclude the possibility of a binary merger progenitor from a plausible nearby host (Abbott et al. 2008; Abadie et al. 2012b).

For this analysis, the offline coincidence between GW and EM observation is done in time and sky location, providing an ability to exclude non-coincident background from either search. Timing for a GW signal is very good, and model spiral waveforms can be placed to fractions of a millisecond (Fairhurst 2009). The coincidence search window in time is determined by the expected delay between coalescence and the emission of EM radiation, which ranges from seconds for prompt gamma-rays to hours for afterglow emission. Sky localization using GWs is based primarily on triangulation across a light-travel baseline of 10–30 ms (Fairhurst 2011), but can be improved by including the signal amplitude and phase (Sidery et al. 2014; Singer et al. 2014). For the initial detector network, sky position for a merger event can be determined to several tens of square degrees (Fairhurst 2009).

Figure 2 shows an outline of the GW-EM coincidence pipeline. We begin with GW events found by a standard matched-filter analysis for coincident CBC signals across two or all three of the GW detectors (Abadie et al. 2012). The original results of this analysis identified a loud signal that had been injected into the data as a blind test of the pipeline (as seen in Figure 3), but did not find any other outliers that stood out above the coincident background of the GW instruments. The matched filter analysis gives some information about sky location (from timing) and distance (from amplitude), but better parameters are obtained through a coherent Bayesian follow-up (Veitch & Vecchio 2010; Veitch et al. 2014) which calculates the posterior probability of all possible physical




Figure 3. The S6/VSR2-3 science run contained a simulated NS/BH merger signal that was injected into the data without knowledge of the analysis teams. This plot from (Abadie et al. 2012) shows the cumulative rate of events coincident in the H1 and L1 detectors with chirp mass $3.48 \leq M_c/M_\odot < 7.40$ as seen by the matched filter pipeline in four months of data around the simulation. Chirp mass is a function of the individual masses: $M_c = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, and is the dominant mass parameter driving binary evolution during inspiral. The injection, shown by the blue triangle with ranking statistic $\rho_c = 12.5$, is clearly resolved against the expected background distributions derived from time-shift analysis, both with (black) and without (gray) single-detector contributions from the simulated event in one of the detectors (coincident with time-shifted noise in another). Additional coincidence with an EM counterpart can help resolve any astrophysical events present in the data in regions with otherwise large background ($\rho_c \sim 9.5$).

signals (with varying component masses, location, distance, inclination, etc) using their coherent overlap with the data from all detectors. This is particularly useful for the relatively common case of a signal detected in H1 and L1, but too weak to be detected in Virgo, thus missing V1 timing information. The Bayesian analysis is still able to use the Virgo data to exclude regions of the sky where, for example, Virgo is particularly sensitive relative to the LIGO instruments.

The sky location and distance estimates can be further refined by assuming that compact mergers arise from a known galaxy, for which we have reasonably complete catalogs out to the LIGO horizon (White et al. 2011). Each galaxy in the catalog can be assigned a host probability which is proportional to its blue-light luminosity (as a proxy for its mass, which we assume is proportional to the rate of compact binary mergers), as well as its overlap with the GW-derived posterior distributions in distance and sky-location. This reduces the search from several hundred square degrees to a handful of essentially point-source locations (given the angular resolution of ASM), at the risk of the true host being absent from the catalog.

The offline search for an EM counterpart from GBM or ASM is then triggered by the time of coalescence, and the GW-derived skyamp or the list of possible host galaxies with corresponding probabilities. We search over the parameter space of expected EM signals, and set nominal thresholds above which we consider GW-EM candidate coincident events.

2. LIGO-VIRGO EVENT GENERATION

2.1. Matched Filter CBC Search

As two compact objects orbit, they lose orbital energy and angular momentum to GWs. The orbit responds by shrinking

and the orbital period decreases, causing even more rapid loss of orbital energy to gravitational radiation. This process ultimately leads in a runaway process to merger. At any moment during the inspiral, the GW emission is well characterized by high-order analytic post-Newtonian approximations to general-relativity. By following the orbit adiabatically through the detector bandwidth, a characteristic chirp waveform of increasing amplitude is produced that sweeps through the sensitive band (~ 40 – 2000 Hz) of ground-based GW detectors.

LIGO-Virgo GW data from S6/VSR2+3 has been searched for CBC events with total system mass $< 25 M_\odot$, and the results have been reported in Abadie et al. (2012). The search made use of matched filtering, correlating the GW data in each detector against a bank of theoretical templates composed of model inspiral chirp signals for non-spinning systems with various component masses and orientations. Matched filtering describes the optimal linear filter for maximizing S/N in the presence of stationary Gaussian noise. For the filter corresponding to a frequency domain signal $\tilde{h}(f)$, the matched filter produces an expected single-detector S/N of,

$$\rho^2 = \int_0^\infty \frac{4|\tilde{h}(f)|^2}{S_n(f)} df \quad (1)$$

where $S_n(f)$ is the one-sided power spectral density of the stationary noise.

Due to the presence of non-Gaussian noise in the detector data, a χ^2 consistency test (Allen 2005) is applied to events which are initially identified as local maxima in matched-filter S/N. A re-weighted S/N $\hat{\rho}$ is obtained by downgrading events using an empirically determined relationship that takes the level of inconsistency with the model signal into account (Babak et al. 2013),

$$\hat{\rho} = \begin{cases} \frac{\rho}{\left[\frac{1 + (\chi_r^2)^3}{2} \right]^{1/6}} & \text{for } \chi_r^2 > 1, \\ \rho & \text{for } \chi_r^2 \leq 1. \end{cases} \quad (2)$$

Events from two or more detectors are searched for coincidence in time and mass parameters. Coincident events are ranked by the quadrature sum, ρ_c , of their re-weighted S/N. This *combined* S/N for each event is compared against an estimated background distribution derived from running the same coincidence analysis many times while applying unphysical relative time-shifts between the detector data (as in Figure 3). When the event has a ρ_c value above a predetermined threshold corresponding to some low false-alarm rate (e.g., one false alarm from background per several 1000 yr), the event is considered a GW candidate. Background estimation is performed separately for various two-week epochs, detector combinations, and mass bins in order to isolate configurations with varying background levels.

The S6/VSR2+3 analysis identified one outlier event with an estimated false-alarm-rate of $\sim 1/7000$ yr $^{-1}$. The event was ultimately revealed as a blind simulated signal injected in the data, which served as an end-to-end test of the detection strategy.⁸ Otherwise, no events stood out clearly from the expected background.

⁸ Details at <http://www.ligo.org/science/GW100916/>.

2.2. Bayesian Parameter Estimation

Bayesian integration is a natural way to determine whether or not a coherent GW signal is present in the data from a network of instruments, and if so, what are the likely parameters. While the coincident matched-filter analysis described in the previous section imposes some consistency on time and mass, a coherent analysis requires from the start a common physical gravitational waveform projected onto each instrument, making it particularly useful for a network of detectors.

If we let \mathcal{H}_1 represent the hypothesis that a CBC signal is present in the data, and let \mathcal{H}_0 represent only noise, the likelihood ratio provides the optimal statistic to distinguish the two,

$$\Lambda = \frac{P(d|\mathcal{H}_1, I)}{P(d|\mathcal{H}_0, I)} \quad (3)$$

where I represents any prior information we have about the system. Given that our signal hypothesis \mathcal{H}_1 represents a large population of possible signals with different waveforms determined by their set of binary parameters $\vec{\theta}$, proper evaluation of the likelihood involves marginalizing over the parameter space,

$$P(d|\mathcal{H}_1, I) = \int P(\vec{\theta}|\mathcal{H}_1, I)P(d|\mathcal{H}_1, \vec{\theta}, I)d\vec{\theta} \quad (4)$$

where $P(\vec{\theta})$ represents the prior probability distribution of parameters $\vec{\theta}$ in our signal population—for example it may reflect that we expect systems to be uniformly distributed in volume with random orientations with respect to the detectors.

In addition to the question of whether or not a signal is present, we are also interested in determining the physical parameters of the binary system that are implied by a set of data. This is represented by the posterior probability distribution over the parameters $\vec{\theta}$,

$$P(\vec{\theta}|d, \mathcal{H}_1, I) = \frac{P(\vec{\theta}|\mathcal{H}_1, I)P(d|\vec{\theta}, \mathcal{H}_1, I)}{P(d|\mathcal{H}_1, I)}. \quad (5)$$

If we are interested in the probability distribution only over a subset $\vec{\theta}_A$ of the parameters, where $\theta \equiv \{\theta_A, \theta_B\}$, we marginalize over those that remain,

$$P(\vec{\theta}_A|d, \mathcal{H}_1, I) = \int P(\vec{\theta}|d, \mathcal{H}_1, I)d\vec{\theta}_B. \quad (6)$$

Nested sampling (Skilling 2004) is a computationally efficient alternative to MCMC techniques (Christensen & Meyer 2001), both of which are designed to intelligently sample and integrate the probability distribution of a high-dimensional parameter space of possible signals. For the case of CBC, we assume the spin of each neutron star is small and there are nine relevant physical parameters: two for the component masses, sky location, and orientation, as well as distance, time, and phase. We apply a Bayesian coherent follow-up based on nested sampling (Veitch & Vecchio 2010) to a large number of coincident events identified from the matched filter analysis. The output of the nested sampling routine is a set of sample vectors which trace the estimated 9-dimensional posterior probability distribution of system parameters (Equation (5)), as well as the integrated likelihood obtained by summing over them all (Equation (4)).

2.3. Galaxy Targeting

Extra-galactic events detectable by the initial LIGO-Virgo network are close enough so that a galaxy catalog can be used to identify probable hosts. While this is only marginally useful for a search using GW data alone due to the poor angular resolution of current GW detector networks, it becomes quite important for EM follow-ups where the EM sky resolution is significantly better (Nissanke et al. 2013; Hanna et al. 2014). Inclusion of a galaxy catalog for this offline EM search greatly reduces the computational cost and false-positive rate of the ASM analysis by reducing an area of ~ 150 square degrees to just tens of individual points.

The Gravitational-wave Galaxy Catalog (GWGC; White et al. 2011) was designed specifically to aid in GW searches with the initial detector network. It contains distance, type, location, geometry, and blue-light luminosity information for ~ 50 k galaxies within 100 Mpc, taken from the HyperLEDA database and other sources. To choose probable host galaxies from the catalog, we compare their distance and sky location against the probability distribution derived from the Bayesian analysis on GW data. Galaxies which occur in regions of high probability are considered for EM follow-up. At advanced LIGO/Virgo distances, galaxy catalogs suffer from incompleteness, with only the brighter galaxies making the flux limits of surveys, and deep surveys often covering only fractions of the sky. Hanna et al. (2014) outline some of the considerations necessary to effectively use incomplete catalogs in that regime. For this initial LIGO/Virgo study, we assume the GWGC is 100% complete out to the distance of detectable NS/NS mergers.

The GW posterior probability distribution must be estimated from the discrete sampling provided by the nested sampling procedure. We estimate the density of posterior samples at each galaxy location and distance using Gaussian kernel density estimation (KDE). We treat distance and sky location independently due to limited sampling of the posterior, typically thousands of points.

For a galaxy with blue-light luminosity $L_{10}(10^{10} L_{\odot})$ at distance $r < 100$ Mpc and location Ω , the estimated distribution from the N posterior samples is

$$P_{\text{GW}}(r) = \frac{1}{N} \sum_i \frac{1}{r^2} w(r_i - r, \sigma_r = 10 - 20\%) \quad (7)$$

$$P_{\text{GW}}(\Omega) = \frac{1}{N} \sum_i w(\vec{x}_i - \vec{x}, \sigma_{\Omega} = 3^{\circ}) \quad (8)$$

$$P(L_{10}, r, \Omega) = L_{10} P_{\text{GW}}(r) P_{\text{GW}}(\Omega) \quad (9)$$

where $w(\mu, \sigma)$ represents the 1D and 2D Gaussian kernels, normalized to give $P = 1$ in the case of a uniform posterior. Figure 4 shows an example of the kernel density estimation applied to the distance parameter, as well as a sample of matching galaxies, weighted by all three factors in Equation (9), against the GW skymap. The distance KDE bandwidth is chosen according to the estimated individual galaxy distance errors, which are around 10–20% depending on the original source catalog. The 3° sky bandwidth is chosen to balance statistical errors (resulting Poisson errors are at the level of a few %) with resolution of the GW network. Finally, we have undone the r^2 distance prior from the assumption of homogeneously distributed sources during the Bayesian




Figure 4. Distance (top) and sky position (bottom, with spherical coordinates θ , ϕ in radians) posterior distributions derived from the Bayesian follow-up of GW data are used to filter a galaxy catalog to check for likely nearby hosts (bottom). The area of the weighted galaxies corresponds to their relative probability of being the host. It is directly proportional to the blue-light luminosity (as a proxy for mass, or merger rate), and the estimated posterior probability distribution at the corresponding distance and sky location. For this simulation of a misaligned NS/NS merger in the Virgo cluster, the maximum probability galaxy corresponded to the true host.

integration as the galaxies themselves already include this volume effect.

3. GAMMA-RAY COUNTERPARTS WITH FERMI GBM

3.1. GBM Instrument and Data Products

The GBM (Meegan et al. 2009) on board the *Fermi* spacecraft measures photon rates from 8 keV–40 MeV. The instrument consists of 12 semi-directional NaI scintillation detectors and 2 BGO scintillation detectors which cover the entire sky not occluded by the Earth (about 65%). The lower-energy NaI detectors have an approximately $\cos \theta$ response relative to angle of incidence, and relative rates across detectors are used to reconstruct the source location to a few degrees. The BGO detectors are much less directional, and can be used to detect and resolve the higher energy spectrum above ~ 200 keV.

GBM produces on-board triggers for GRB events by looking for multi-detector rate excess over background across various energy bands and timescales. In the case of a trigger, individual photon information is sent to the ground and the event is publicly reported. Those events which have been confirmed as

GRBs have already been studied in coincidence with LIGO-Virgo data (Abadie et al. 2012). In addition to the triggered events, survey data is available which records binned photon counts over all time. For this offline analysis of short transients, we use the CTIME (time-resolved) daily data, which is binned at 0.256 s over 8 energy channels for each detector. A new GBM data product for continuous time-tagged events was introduced in 2012 that provides complete individual photon information with $2 \mu\text{s}$ and 128 energy channel resolution. This data product was not available for un-triggered times during the initial LIGO-Virgo science runs, but will enhance offline sensitivity to short bursts when the advanced GW detectors begin operation.

3.2. Coherent Analysis of GBM Data

In this section, we develop a procedure to coherently search GBM detector data for modeled events (Blackburn et al. 2013). The idea is that by processing multiple detector data coherently, we can obtain a greater sensitivity than when considering one detector at a time. Greater computational resources available offline (versus on-board) also allow for more careful background estimation to be done. For this analysis, we can relax to some extent the strict 2-detector coincidence requirement used to veto spurious events on board as the GW trigger means much less time and sky area is considered. These advantages help to balance out the coarse time resolution of CTIME data, which reduces offline sensitivity to very short bursts prior to 2013.

Each detector is subject to a substantial time-varying background from bright high-energy sources that come in and out of the wide field of view (FOV), as well as location-dependent particle and Earth atmospheric effects. This background must be estimated and subtracted out to look for any prompt excess. In this analysis where we are interested in the background estimate for a short foreground interval $[-T/2, +T/2]$ where $T \sim 1$ s, we estimate the background using a polynomial fit to local data from $[-10T, +10T]$ (minimum ± 5 s), excluding time $[-3T/2, +5T/2]$ around the foreground interval to avoid bias from an on-source excess. The polynomial degree is determined by the interval length to account for more complicated background variability over longer intervals. It ranges from 2 (minimum) to $1 + 0.5 \log_2 T$. A separate fit is done for each detector/channel combination, and determines the background rate estimation over the foreground interval, as well as its systematic uncertainty from fitting error. Data from channels with poor fits (large χ^2) are excluded from the analysis.

High-energy cosmic rays striking a NaI crystal can result in long-lived phosphorescent light emission. The detector may interpret this as a rapid series of events, creating a short-lived jump in rates for one or multiple channels, and severely distorting the background fit if not accounted for. They are identified with a simple procedure that scans for rapid 1-bin spikes in the background interval. The affected bins are removed from the background fitting. Cosmic rays affecting the foreground interval are handled differently, as described in Section 5.1.

A likelihood ratio combines information about sources and noise into a single variable. It is defined as the probability of measuring the observed data, d , in the presence of a particular true signal H_1 (with some source amplitude $s > 0$) divided by the probability of measuring the observed data in noise alone

$H_0(s = 0)$,

$$\Lambda(d) = \frac{P(d|H_1)}{P(d|H_0)}. \quad (10)$$

When signal parameters such as light curve, spectrum, amplitude s , and sky-location α, δ are unknown, one can either marginalize over the unknown parameters, or take the maximum likelihood over the range to obtain best-fit values.

For binned, uncorrelated Gaussian noise,

$$P(d_i|H_1) = \prod_i \frac{1}{\sqrt{2\pi}\sigma_{d_i}} \exp\left[-\frac{(\tilde{d}_i - r_i s)^2}{2\sigma_{d_i}^2}\right] \quad (11)$$

$$P(d_i|H_0) = \prod_i \frac{1}{\sqrt{2\pi}\sigma_{n_i}} \exp\left[-\frac{\tilde{d}_i^2}{2\sigma_{n_i}^2}\right] \quad (12)$$

where we have used $\tilde{d}_i = d_i - \langle n_i \rangle$ to represent the background-subtracted measurements in each detector-time-energy bin, σ_{n_i} and σ_{d_i} for the standard deviation of the background and expected data (background+signal), r_i for the location/spectrum-dependent instrumental response, and s for the intrinsic source amplitude at the Earth. Maximizing the likelihood ratio is the same as maximizing the log-likelihood ratio $\mathcal{L} = \ln \Lambda$,

$$\mathcal{L} = \sum_i \ln \left[\frac{\sigma_{n_i}}{\sigma_{d_i}} + \frac{\tilde{d}_i^2}{2\sigma_{n_i}^2} - \frac{(\tilde{d}_i - r_i s)^2}{2\sigma_{d_i}^2} \right]. \quad (13)$$

The dependence of the response factors r_i on sky location is complicated, so the likelihood ratio is calculated over a sample grid of all possible locations. Assuming a single location, the remaining free parameter is the source amplitude s . The variance in the background-subtracted detector data includes both background and source contributions,

$$\sigma_{d_i}^2 = \sigma_{n_i}^2 + r_i s + \sigma_{r_i}^2 s^2 \quad (s \geq 0) \quad (14)$$

with $\sigma_{r_i}^2$ representing Gaussian-modeled systematic uncertainty in the instrumental response. Source terms are only included for physical $s \geq 0$, else their contribution is zero. The background contributes Poisson error, as well as any systematic variance $\sigma_{b_i}^2$ from poor background fitting which is also assumed to be Gaussian,

$$\sigma_{n_i}^2 = \langle n_i \rangle + \sigma_{b_i}^2. \quad (15)$$

We find s_{best} which maximizes \mathcal{L} by setting the derivative $d\mathcal{L}/ds$ to zero using Newton's method.

To consider all possible source amplitudes we need to integrate the likelihood $P(d|s)$ (Equation (11)) over a prior on s ,

$$P(d) = \int P(d|s)P(s)ds. \quad (16)$$

For a given set of detector data d , the likelihood $P(d|s)$ over s is almost the product of individual Gaussian distributions (not quite Gaussian because σ_d depends on s). The product of Gaussian distributions with mean values μ_i and standard

deviations σ_i is itself Gaussian with mean and variance,

$$\mu_{\text{prod}} = \frac{\sum \mu_i / \sigma_i^2}{\sum 1 / \sigma_i^2}, \quad \sigma_{\text{prod}}^2 = \frac{1}{\sum 1 / \sigma_i^2}. \quad (17)$$

In this case, $\mu_i = \tilde{d}_i / r_i$ and $\sigma_i = \sigma_{d_i} / r_i$. We estimate the variance of \mathcal{L} over s as,

$$\sigma_{\mathcal{L}}^2 = \frac{1}{\sum r_i^2 / \sigma_{d_i}^2}, \quad \sigma_{d_i}^2 \text{ evaluated at } s_{\text{best}} \quad (18)$$

and choose a scale-free prior with fixed $\beta = 1$, so that our choice of form for an amplitude prior at the Earth does not translate back into a luminosity distribution that varies with distance.

One difficulty with any power-law prior is that it diverges for $s \rightarrow 0$. We enforce a finite and well-behaved prior by multiplying by a prefactor,

$$P(s) = \left[1 - e^{-(s/\sigma_{\mathcal{L}})^\beta} \right] s^{-\beta} \quad (19)$$

so that $P(s)$ reaches a maximum constant value of $(1/\gamma\sigma_{\mathcal{L}})^\beta$ for small s . The tunable parameter γ sets the number of standard deviations at which the prior begins to plateau, and we use $\gamma = 2.5$. We then approximate $P(s)$ as constant over a range of $\sigma_{\mathcal{L}}$ for any $s > 0$, and include a correction to account for clipping of the Gaussian for non-physical $s < 0$, which can be represented by the error function. The final approximation for the amplitude-marginalized log-likelihood becomes,

$$\begin{aligned} \mathcal{L}(d) = \ln \sigma_{\mathcal{L}} + \ln \left[1 + \text{Erf} \left(\frac{s_{\text{best}}}{\sqrt{2} \sigma_{\mathcal{L}}} \right) \right] + \mathcal{L}(d|s_{\text{best}}) \\ + \begin{cases} \ln \left[1 - e^{-(s_{\text{best}}/\sigma_{\mathcal{L}})^\beta} \right] - \beta \ln s_{\text{best}} & s_{\text{best}} > 0 \\ -\beta \ln(\gamma\sigma_{\mathcal{L}}) & s_{\text{best}} \leq 0 \end{cases} \end{aligned} \quad (20)$$

which contains factors from the Gaussian width, fractional overlap with $s > 0$, maximum likelihood at s_{best} , and scaling from $P(s)$, respectively. Finally we are free to calibrate the log-likelihood by subtracting the expected $\mathcal{L}(d)$ calculated for no signal at a reference sensitivity: $\mathcal{L}_{\text{ref}} = -\beta \ln \gamma + (1 - \beta) \ln \sigma_{\text{ref}}$. $\sigma_{\mathcal{L}}$ represents the source amplitude required for a 1σ excess in the combined data, and is around $0.05 \text{ photons/s/cm}^2 \times (T/1\text{s})^{-1/2}$ [50–300 keV] for a typical source spectrum and background level. Figure 5 shows the coherent S/N expected from all detectors for a 0.512 s-long event with normal GRB spectrum and constant amplitude of $1.0 \text{ photons/s/cm}^2$, and compares it to the S/N expected from individual detectors alone in the 50–300 keV band.

3.3. GBM Prompt Coincidence with GW Events

A GW trigger provides an accurate time of coalescence t_c as well as an approximate sky location to tens or hundreds of square degrees. This matches well the time and sky location constraints provided by a prompt observation by GBM, so that coincidence between the two instruments is particularly effective. To make use of the coherent analysis described in the previous section, we must first obtain an expected all-sky GBM response for a given counterpart. We make use of three representative source spectra using a set of precomputed response tables. The representative spectra are Band functions with the parameters listed in Table 2. For the soft spectrum,

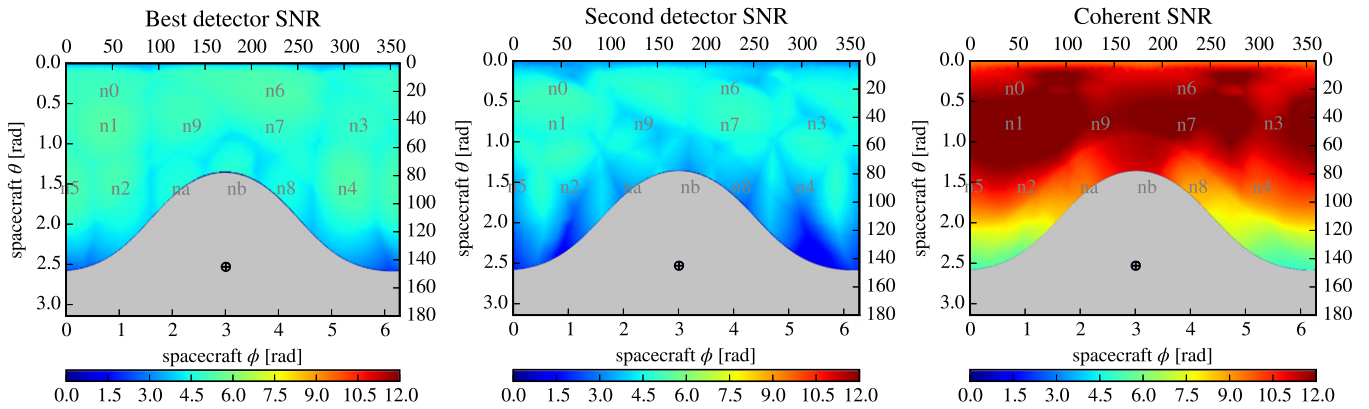


Figure 5. Signal-to-noise ratio in GBM data expected from a 0.512 s long signal with a normal GRB spectrum. The signal is normalized to 1.0 photons/s/cm² in the 50–300 keV band, and the background rates and Earth position are selected from 10 s prior to GRB 090305 A. The model response includes contributions from atmospheric scattering. The first two maps show the signal-to-noise ratio expected from the best and second-best detectors assuming data collected across 50–300 keV. The interpretation is as follows: given a hypothetical source at any sky location (ϕ , θ), what is the maximum (left) or second-best (middle) signal-to-noise ratio seen across *all* NaI detectors. Generally one would assume that the maximum S/N would come from the best-aligned NaI detector (n0, n1, ... as shown), and the second-best S/N is from one that is nearby. The final plot shows the S/N expected from a coherent analysis of the data using all detectors, also as a function of source location. The maximum coherent S/N is achieved where multiple detectors are favorably oriented toward the source. The coherent analysis samples all sky positions, which means it is subject to a trials factor, but this is difficult to evaluate because of the large degree of correlation between nearby sky locations. The “best-detector” S/N similarly has a trials factor equal to the number of detectors (12). Use of the “second-best” detector, which corresponds to the on-board GBM triggering strategy, implies two detectors above a given S/N threshold, which means a much lower false-alarm probability at equivalent S/N, as well as improved rejection of non-Gaussian noise.

E_{peak} and β are undefined. The tables provide channel-dependent expected counts for the direct and spacecraft-scattered response as a function of source location to $\sim 1^\circ$ resolution. Contributions from atmospheric scattering are available as counts summed over the two CTIME channels covering 50–300 keV as a function of source and Earth relative position for the two most common spacecraft rocking angles.

For each GW candidate event, we search a window $[-30, 30]$ s relative to t_c for prompt excess between 0.256 and 8 s long. Emission outside of the standard accretion timescale covering the first seconds following merger is speculative. However, any events detectable in GWs will be much closer than known sGRBs to date, and therefore they provide an opportunity to search below threshold for weak and possibly less-beamed precursor or extended emission. To appropriately tile the search in time and duration T , we use rectangular windows with T spaced by powers of two (0.256, 0.512, 1 s, etc). Their central times are sampled along the search interval in units of $T/4$ to provide an even mismatch in S/N across search windows. For each emission model tested, the likelihood ratio is then marginalized over all windows and spectra.

The GW data also provides a rough sky-location, which can be represented as a prior probability distribution over the sky $P_{\text{GW}}(\Omega)$. This GW prior is multiplied by the GBM likelihood ratio, which is also a function of sky position, before marginalization over sky location is done,

$$\Lambda_{\text{GW-GBM}} = \int d\Omega P_{\text{GW}}(\Omega) \Lambda_{\text{GBM}}(\Omega) \quad (21)$$

where the GBM likelihood ratio depends on source spectrum, on-source time window, and location. While the sky prior can in principle be improved by incorporating knowledge of the anisotropic local mass distribution (using a galaxy catalog), we note that as neither the GW network nor GBM have precise (sub-degree) sky localization, the gains from sharpening the prior in this way are small, and catalog incompleteness and luminosity-rate relationships introduce unnecessary

complications. The situation is different when considering better-localized X-ray (and optical/ratio) counterparts.

3.4. GW-GBM Coincident Background Estimation

The background for GW-GBM analysis is characterized by the probability of observing a given $\Lambda_{\text{GW-EM}}$ from random coincidence between GW and GBM data. We calculate the expected distribution of $\Lambda_{\text{GW-EM}}$ by running the coincidence search defined by Equation (21) on GW background events derived from time-shift analysis. The time-shift sample represents 100 artificial time-shifts (\sim seconds) applied to data between the LIGO Hanford and Livingston sites, as well as a further 10 time-shifts (\sim day) applied between the GW and GBM data. Thus we obtain a background sample representing characteristic noise from $\times 1000$ the original foreground live-time.

3.5. Simulation of Prompt Gamma-ray Counterparts

GRB simulated signals are injected into the data by using the instrument response model to predict the signal contribution from a source with a given spectrum and fluence. In our analysis, we simulate a standard candle signal following the normal-type Band spectrum with a fluence at the Earth of 1 photon/cm²/s in 50–300 keV assuming a source at 30 Mpc. The signal assumes normal GRB spectral parameters according to Table 2 and lasts for 1 s, corresponding to an isotropically emitted energy of about 6.6×10^{46} erg. When simulating a signal at distance D , the fluence is reduced by a factor of $(30 \text{ Mpc}/D)^2$.

4. X-RAY COUNTERPARTS WITH RXTE ASM

4.1. ASM Instrument and Data Products

The ASM (Levine et al. 1996) on the RXTE surveyed the X-ray sky between 1.5–12 keV from 1996–2011. The instrument consisted of three long-and-narrow shadow-mask X-ray cameras rigidly attached to a common motorized axis of

rotation. Each camera saw a 6 by 90 degree FOV, and together covered about 3% of the sky at any one time. As the cameras scanned across the sky while tiling 90 s dwells, they were able to localize a steady source to about 0'.1. The duty cycle of the ASM along with constraints from the observing schedule of other *RXTE* instruments resulted in a randomly chosen location being scanned by an ASM camera a few times per day.

An ASM camera observation consists of a superposition of the shadow pattern of all X-ray sources within the FOV. The data can be modeled as the contributions from active sources with known locations along with a diffuse background. Amplitudes for sources are treated as free parameters, and estimated using a linear least-squares fit to the data (Levine et al. 1996), with errors dominated by photon counting statistics. The amplitude from any additional position of interest within the FOV can also be estimated by adding the location as a source during the fit. For a weak source, the error on estimated amplitude is dominated by the diffuse background which typically contributes a 3- σ error of ~ 20 mCrab (4.8×10^{-10} erg/cm²/s between 2–10 keV).

4.2. Modeled Search for Afterglow Light Curves

The ASM measurements consist of irregularly spaced flux observations with varying uncertainties from specific regions in the sky, which makes a search strategy for an arbitrary flux excess complicated. We use limited knowledge of the sGRB X-ray afterglow signal to narrow the search to signals that decay like broken power laws in flux beginning from the initial time-of-merger t_0 as determined by the GW data.

Following Zhang et al. (2006), we model a canonical sGRB X-ray afterglow by a double-broken power law with a short region of rapid decline (power law index $\alpha \approx 3$) lasting 1e2–1e3 seconds, followed by a standard afterglow decay ($\alpha \approx -1.2$), and finally a break and rapid decay ($\alpha \approx -2$) after 10^4 to 10^5 seconds. Additionally there may be a plateau or extended emission somewhere after the rapid decay phase around 10^2 to 10^3 seconds. We attempt to create a generic light-curve with minimal free parameters which covers possible observed light curve scenarios, in context of the limited sampling of ASM data, and also allows for the possibility of delayed or non-standard X-ray emission due to off-axis observation or some other unobserved phenomenon.

For this we sample double-broken model power-law light curves with break times sampled on a grid (Figure 6) and freely varying power-law indices. We label N local measurements of flux excess d_i by ASM at times t_i with standard deviations σ_i . A given light curve parameterization defines expected relative counts at the measurement times $L_i = L(t_i)$. We can then define the weighted sum of measurements c_i which maximizes the S/N of the summed data,

$$D = \sum_i \frac{L_i d_i}{\sigma_i^2}. \quad (22)$$

The measurements d_i are already assumed to be zero-mean, giving a measured S/N for the coherent sum D of,

$$\rho = \frac{D}{\sqrt{\sum_i L_i^2 / \sigma_i^2}}. \quad (23)$$

The expected S/N of each measurement is added in quadrature

for the total expected S/N of the coherent sum,

$$\langle \rho \rangle = \sqrt{\sum_i \frac{L_i^2}{\sigma_i^2}}. \quad (24)$$

If we are fortunate to have an observation soon after the beginning of the afterglow, the first measurement is likely to dominate the sum due to the rapid decay. If only late-time measurements are available, the coherent sum naturally bins the data in order to gain in sensitivity. A positive power-law index for the first segment of the light curve can handle situations where the afterglow does not begin immediately after the burst, for example due to beaming effects.

For each pair of power-law breakpoints (t_1, t_2), we fit the three power-law indices ($\alpha_1, \alpha_2, \alpha_3$) using a minimization routine under the constraints $-3 < \alpha_1 < 2$, $-2 < \alpha_2 < 1$, and $-3 < \alpha_3 < 0$. This allows for a wide variety of afterglow waveforms including the standard canonical double-broken power law, an initially rising light curve which then decays normally, or an isolated burst of extended emission that might occur minutes to days after merger. For each set of ASM measurements from a location of interest [0, 4 d] about t_0 , we record the fit parameters which give the maximum S/N, as well as a variety of auxiliary parameters useful for characterizing non-Gaussian noise.

4.3. ASM Afterglow Coincidence with GW Events

The ASM point source flux estimation is well-suited to follow up locations of likely host galaxies for GW events. The shadow-mask flux reconstruction isolates contributions to about 0'.1 on the sky. This is small enough that a large fraction of the sky can be excluded by requiring known galaxy coincidence, but also large enough so that all but the most extended nearby galaxies can be targeted by a single location. For each host location, we generate ASM flux measurements from available camera dwells by adding that single location as a test source along with other known active X-ray sources during the fit of each dwell's shadow-mask data. Derived flux measurements from that location are then searched numerically for the five parameter light curve (two breakpoints and three decay factors) which maximize S/N of the coherent sum. Because the data are sparse, it is difficult to maximize S/N numerically for freely varying breakpoints, so they are sampled on a fixed grid and only the three power-law indices are searched with standard numerical minimization routines as described above.

We take the maximum S/N ρ observed across all model light curves to weigh each candidate host galaxy according to likelihood. The conversion from ρ to likelihood is determined by a global S/N distribution of X-ray background derived from the same methods. The distribution is fit to a function $f(\rho)$ parameterized by a Gaussian (maximum value 1) with power-law tail (Figure 7), and this empirical fit is used to calculate the likelihood= $1/f(\rho)$ at each S/N measurement. For a given GW event with a set of possible hosts, we obtain a final rank for the coincident observation,

$$\Lambda_{\text{GW-ASM}} = \sum_{\text{host}} \left[P_{\text{GW}}(\text{host}) \times f(\rho_{\text{host}})^{-1} \right] \quad (25)$$

where the likelihood $\Lambda_{\text{GW-ASM}}$ represents the probability of an afterglow signal in the ASM data, marginalized over all possible hosts. The hosts must maintain a $P_{\text{GW}} > 1$ (a threshold




Figure 6. (Top) double-broken power-law break-point parameters for model afterglow light curves used when fitting ASM data. (Bottom) typical 2–10 keV sGRB afterglow light curve with decay indexes of -2 , -1.2 , -2 between breakpoints at 3×10^2 and 5×10^4 s. The amplitude is based on Swift XRT observations of actual sGRB afterglows, which can vary in intrinsic luminosity by over two orders of magnitude. Also shown is a family of ad hoc X-ray bursts, which have a total 2–10 keV energy release of 10^{49} erg over their durations. This amplitude is chosen so that the bursts at peak amplitude rise just above the afterglow signal, as in the case for the extended emission seen in some sGRB afterglows. The bursts are also parameterized by a double-broken power law, this time with decay indexes of $+2$, -1 , -2 between breakpoints at T and $2T$ where T is the duration. The yellow band represents an intrinsic X-ray luminosity required to be detectable (>20 mCrab) in a single 90 s dwell by ASM for sources between 20 and 80 Mpc.

Table 1

High-energy Photon Survey Instruments Used to Search for EM Counterparts

Instrument	Energy	FOV	Resolution	Cadence
Fermi GBM	8 keV–40 MeV	65%	$>5^\circ$	1.5 h
RXTE ASM	1 keV–10 keV	3%	$0:1$	1.5 h

Note. Field of View (FOV) is represented as a percentage of the entire sky.

representing the case for a galaxy of typical \sim MW mass and no location or distance prior information), and at maximum 200 hosts are scanned to limit computational expense. $P_{\text{GW}}(\text{host})$ is derived, as in Equation (9), from the sky and distance overlap of each host with the prior distributions determined by the GW data and host mass (luminosity). To convert the likelihood into a more S/N-like quantity, we define rank $r = \sqrt{2\Lambda_{\text{GW-ASM}}}$.

Table 2

Model Band Spectral Parameters Used to generate GBM Detector Response Tables Used in the Coherent Analysis

Spectrum	E_{peak}	α	β
normal	230 keV	1	2.3
hard	1 MeV	0	1.5
soft	...	2	...

Note. The soft and hard spectra were originally designed to produce a representative response over the primary GBM-GRB localization band 50–300 keV, and a number of more realistic response models are currently being developed which will be more appropriate for full-spectral all-sky analysis. The construction of these all-sky, all-time models is difficult computationally due to the many degrees of freedom in the detector response.

4.4. Simulation of X-ray Afterglow Counterparts

We simulate a family of X-ray afterglow signals by adding model light curves directly to the derived ASM flux measurements for a given location. We assume that the error estimate for each flux measurement remains dominated by contributions from the diffuse background, so that Poisson or other systematic error from the source can be neglected. The model light curve is then injected into the data from a reconstructed location at a time of interest (Figure 8). For a given light curve shape, we recover the distribution of maximum S/N measurements from the template search as a function of injected amplitude, or equivalently injected distance for a standard candle source.

For the case of a GW-ASM coincidence, we begin with a set of simulated NS/NS coalescence events detected with the GW analysis procedure. For S6/VSR2+3, a large number of simulations were done at random sky locations and distances in order to evaluate the overall detection efficiency of the pipeline (Abadie et al. 2012). To characterize the sensitivity of the joint GW-ASM coincidence analysis which assumes signals originating from a set of host galaxies, we first create an artificial galaxy at the location and distance of the GW simulation, with a luminosity taken from a (luminosity-weighted) random draw of galaxies from the catalog within some compatible volume. We then synthesize ASM data from this fake location and distance by appropriating the ASM flux measurements from the randomly chosen galaxy. Reconstructed ASM flux measurements from the remaining true galaxies within the GW-derived sky region are included in the coincidence analysis as before, as well as the fake galaxy with the simulated lightcurve added to the ASM data. We then obtain a distribution of the joint GW-ASM detection statistic for a given CBC system as a function of distance and intrinsic EM amplitude.

We test performance of the method against a standard X-ray afterglow typical of Swift XRT observations of those from short GRBs (Figure 6). This light curve has power-law decay indices of -2 , -1.2 , and -2 between breakpoints at 3×10^2 and 5×10^4 s, and a 1–10 keV luminosity of 10^{46} erg s^{-1} at the first break-point. We also assume a Crab-like spectrum, for convenient conversion into ASM band counts which are Crab calibrated.

In addition, we inject a representative sample of X-ray burst signals which occur minutes to days after the merger, motivated by theoretical magnetar wind-driven scenarios (Gao et al. 2013; Zhang 2013). We test a short, medium, and late-time burst beginning with a rise $\propto t^2$ until $t_1 = 10^2, 10^3,$




Figure 7. All-sky X-ray background from the generic light curve fitting procedure. The background distribution is largely Gaussian, but shows a significant tail which can be captured by a power law. The background fit is used to estimate the significance of coincident signal-to-noise ratio measurements ρ in the analysis.

and 10^4 s, decaying as t^{-1} until $t_2 = 2t_1$, and then falling as t^{-2} afterward. The simple wind-driven emission scenarios consider a total amount of energy released as the magnetar spins down, and then either dissipating via internal or external shocks, releasing some fraction of the dissipated energy in X-rays. We set the total energy released in X-rays over the entire period of the burst to be 10^{49} erg in 2–10 keV, resulting in bursts with instantaneous luminosities slightly greater than that for the standard afterglow at peak. In this way, they can be thought of representing the extended emission/flares seen in some sGRB afterglows. As in the standard afterglow model, we assume a Crab-like spectrum for convenient translation into ASM counts.

4.5. GW-ASM Coincident Background Estimation

The GW-ASM background distribution is found by running the coincidence search using GW background events derived from time-shift analysis, similar to that for GBM. The same 10 additional time-shifts (spaced in multiples of 4 days) are applied between the GW and ASM data to increase statistics. We then obtain the distribution of the GW-ASM rank statistic expected from accidental coincidence between GW and ASM background. This distribution, as well as that for the family of simulated afterglow signals, are shown in Figure 9.

5. RESULTS AND CONCLUSIONS

5.1. Selection Cuts and Background Rejection

We have tested the end-to-end GW-EM pipelines for both GBM and ASM on a selection of real GW background noise events from the LIGO-Virgo S6/VSR3 joint science run, as well as simulated GW-EM events injected in coincidence in both the GW and EM data streams. By comparing the properties of GW background and simulations, we identify a number of selection cuts that help reduce the non-Gaussian background present in the EM data. The ASM selection cut is based on the reduced χ^2 of the ASM data after the best-fit double-broken power-law light curve has been subtracted (Figure 10). However, due to the sporadic coverage of the ASM measurements, it is difficult to select on any morphological properties of the light curves themselves, beyond the




Figure 8. ASM data with injected afterglow light curve at 14 Mpc. The best-fit light curve provides the optimal coherent sum signal-to-noise ratio measurement.

weak requirement that the measurements be consistent with one or more of the possible templates.

We place a nominal threshold of 18 on $\Lambda_{\text{GW-ASM}}$, the GW-ASM coincident likelihood ratio. Because the likelihood is largely determined by the Gaussian-like shape of the ASM background distribution, we represent this as a threshold of 6 on the alternative rank parameter $r = \sqrt{2\Lambda}$. The distribution of r for the GW background events, as well as a variety of simulations is shown in Figure 9. With the ASM coincidence requirement applied, the loudest GW background event drops from a combined S/N of 11.0 to 9.13 (Figure 14), a reduction of $\sim 17\%$. ASM is generally sensitive to afterglow signals within the LIGO horizon, however detectability varies widely since it is highly dependent on exactly when, or if ever, the ASM observation(s) occur relative to the peak in the light curves (Figure 11). Similarly, detections may fail if the GW localization is not sufficient to pick out the correct galaxy host among the top 200 candidates, which happens about 7% of the time at 10 Mpc, and 32% of the time at 30 Mpc.

In the GBM detector, most short background events are due to phosphorescent events in the NaI detectors, which can arise from cosmic rays. While we implemented a rough cut to exclude such events from contaminating the background fits during the analysis, we implement a further selection to remove remaining particle events from the foreground interval. The particle events are unique in that they show up in a single detector, and are generally soft in reconstructed energy. Thus we implement two cuts based on the ratio of detected channel 0 S/N in the loudest and second-loudest NaI detectors, as well as between channel 0 and channel 1 in the loudest detector. With these cuts we are able to remove most of the remaining soft particle background.

We also implement a sky-coincidence cut between the GW sky location and GBM by comparing the GBM likelihood ratio with and without the GW sky prior folded in. Since the sky prior typically covers a fractionally small region of the sky for a well-localized event (100 s of square degrees), we expect an appropriately normalized coincident observation to have a correspondingly higher likelihood ratio than one that assumes an isotropic all-sky prior. We expect a factor that can be roughly in line with the fractional reduction in sky area (~ 400), but since this can be moderated by many effects (averaging




Figure 9. Cumulative distribution of ASM rank (monotonic with GW-ASM likelihood ratio) for background events, as well as a variety of GW-ASM afterglow simulations, volume-weighted to resemble a spatially homogeneous distribution in distance up to 40 Mpc. The standard afterglow follows the parameterization in Figure 6 with an intrinsic luminosity of 10^{46} erg s^{-1} at the first breakpoint at 300 s. The other three simulated afterglows represent X-ray bursts at various timescales, with characteristic onset and duration of 100, 1000, and 10,000 s, and a total integrated energy release of 10^{49} erg.




Figure 10. Reduced χ^2 vs. max-S/N in ASM data of background and simulations (all waveforms) for GW-ASM events. The scatter points represent full end-to-end implementation of the coincident GW-ASM pipeline, including joint simulation of NS/NS and afterglow signals from nearby host galaxies. We cut at $\chi^2 > 3$ (gray shade) to remove loud background events inconsistent with one of the template waveforms. For example, they may have excess flux before the merger time, or sporadic flux afterward inconsistent with a power-law decay. The cut is relaxed at high S/N due to template mismatch.

windows, systematic errors in reconstructed location), we choose an empirical factor of e^2 as a coincidence requirement. This cut is still able to reject a large fraction of the background where we do not expect any systematic increase in likelihood ratio due to incorporation of a GW sky prior (Figure 12).

Finally we place a nominal threshold of 30 on the coincident GBM likelihood ratio, which corresponds roughly to that of the weakest un-triggered GRBs detectable by this method (Figure 13). After this threshold and the previously mentioned selection cuts, there is only one coincident GW-GBM event remaining, which has a combined S/N ρ_c of 8.1 (Figure 14), $\sim 26\%$ lower than the previous loudest event.

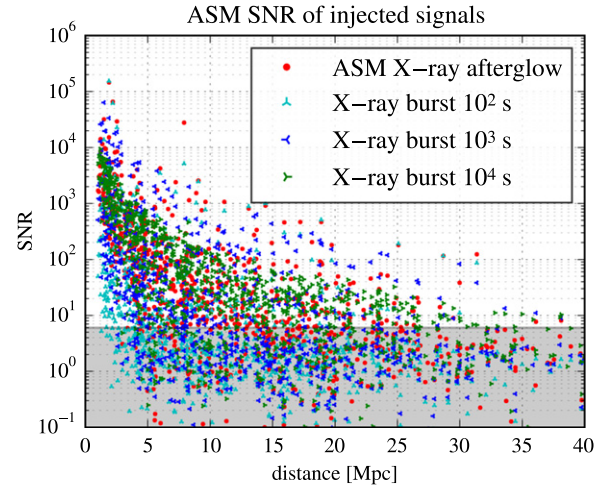


Figure 11. Best-fit S/N of ASM injections vs. distance for joint GW-ASM NS/NS-afterglow injections split by model waveform. The ASM S/N is highly variable mainly due to the uncertain time-to-first observation. Missed detections in ASM are due to several factors including lack of ASM data or poor sky localization due to antenna pattern effects. Rather than cut directly on S/N, ASM events are selected using the derivative *rank* parameter (Figure 9) which folds in S/N information from all possible host locations and their probabilities, as well as information about the ASM background distribution. For reference, we show a shaded region at S/N < 6 which encompasses the majority of ASM background (Figure 10).

The cumulative histograms of S/N of the time-shift GW background events before and after EM coincidence are shown in Figure 14. This can be qualitatively compared to the published background distribution shown in Figure 3.

The original low-threshold GW events used in this analysis were subject to a single detector S/N threshold of 5.5. This means that for a double-coincident H1L1 event, the absolute minimum network S/N is $\sqrt{2} \times 5.5^2 = 7.78$, before any reduction by the χ^2 consistency check (Equation (2)). Near this threshold, one must be careful when assuming simple range scaling relationships between combined S/N thresholds and effective range of the analysis because of the upstream cuts.

Past the knee in the original GW background distribution $\rho_c \gtrsim 9$, the inverse relationship between distance and combined S/N should be robust, so we can estimate a factor of $\gtrsim 15\%$ increase in range for this sample at equivalent false-alarm rate given events where EM coincidence can provide at least a factor of 1/1000 background rejection. While the background sample used in this study reliably probes about three orders of magnitude in rate, the extended background studies from Abadie et al. (2012) and represented in Figure 3 show that for H1L1 events near the end of the LIGO S6 with chirp mass $3.48 \leq \mathcal{M}_c < 7.40$, the background distribution scales roughly as $\sim 100^{-\rho_c}$ in combined S/N over eight orders of magnitude in rate. For the special case in which both a prompt gamma-ray and X-ray afterglow counterpart are observable, the accidental coincidence factors measured here give a large joint rate of background rejection around $\sim 10^7$, which implies a factor of ~ 1.5 increased range to such events if maintaining a fixed false-alarm-rate given an original threshold of $\rho_c > 10$ (Camp et al. 2013) before EM coincidence.




Figure 12. Sky coincidence cut for GW-EM search. We require the likelihood ratio of GBM signal vs. no-signal to be a factor of e^2 greater when the GW skymap prior is used, vs. an all-sky prior. The sky-coincidence is effective at removing loud GBM background, although some signal injections are also rejected due to poor GW or GBM sky localization.

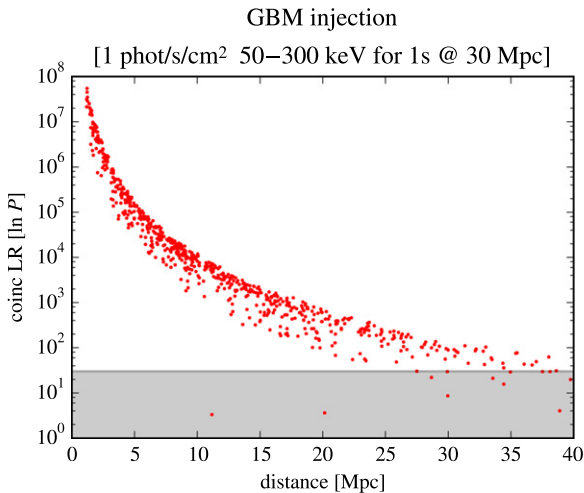


Figure 13. GBM log-likelihood ratio after folding in GW sky prior information, as a function of injection distance for a standard-candle weak GRB simulation calibrated to a flux of 1 photon/s/cm² in 50–300 keV at 30 Mpc (about 6.6×10^{46} erg). The shaded region shows a nominal likelihood-ratio threshold of 30.

5.2. Efficiency for Simulated Signals

The efficiency of the EM pipelines to our standard-candle signal injections is shown in Figure 15. The efficiency fractions are calculated from the end-to-end simulation of a joint NS/NS inspiral signal with corresponding prompt and afterglow EM counterpart, and the process is triggered by a low-threshold ($\rho_c \gtrsim 8.5$) GW candidate. The presence of a GW trigger is required, and the represented fraction does not include efficiency factors from the analysis of GW data itself. However, the EM follow-up efficiency is still generally influenced by the quality of the GW sky localization.

GBM views the entire unocultured sky (65%) when not in the South Atlantic Anomaly ($\sim 15\%$), and this duty-cycle dominates the efficiency factor out to 40 Mpc. This is not surprising as our injection amplitude was chosen to be moderately detectable at 30 Mpc (and still several hundred times weaker than a typical sGRB).




Figure 14. Cumulative distribution of combined S/N (ρ_c) for time-shifted background events observed in GW data before and after selection cuts are made on the requirement of an EM coincidence from either GBM or ASM. An additional 10 time-shifts are applied between GW and EM data, resolving the expected distribution after coincidence to 0.1 events. The coincidence rejection factors are $\sim 10^{-3}$ and 10^{-4} for ASM and GBM, respectively. The corresponding loudest events are at a combined S/N of 11.0, 9.1, and 8.1. However below a combined S/N of ~ 9 , the effects of other analysis cuts take effect. Additional factors of background rejection from tighter time and sky coincidence could further dig into the noise distribution as suggested in Camp et al. (2013), but demonstrating that robustly would require a much larger, or lower-threshold GW background set than was used this study.

The chance of detecting an X-ray afterglow signal with ASM is much more variable (Figure 11) due primarily to the large variability in delay between onset of the afterglow and the first available measurement. The ASM follow-up is also more sensitive to the sky localization accuracy from the GW trigger due to the choice to follow-up only the most probable 200 individual galaxy host locations. Increased distance both increases the sky area uncertainty (due to decreased GW S/N), and increases the area-density of galaxies on the sky. We do not observe ASM counterparts above threshold beyond ~ 30 Mpc. For the X-ray burst model waveforms, longer duration bursts (at equivalent total fluence) were relatively easier to detect as they better matched the ASM cadence.

5.3. Discussion

In this paper, we introduced a strategy to search for high-energy EM counterparts to GW binary coalescence events in archival satellite survey data. We designed an end-to-end GW candidate follow-up pipeline and tested it on a large number of background binary coalescence GW events from time-shifted initial LIGO/Virgo data during their most recent science runs (S6/VSR2+3). The representative noise events were subject to a fully coherent Bayesian parameter estimation analysis in order to obtain sky and distance posterior probability distributions. These distributions were used to obtain a set of probable hosts from a catalog of nearby galaxies.

Two custom follow-up methods were designed to search for both a prompt gamma-ray counterpart in offline data from the Fermi GBM within ± 30 s of the GW coalescence time and consistent with the GW sky location, as well as a generic family of X-ray afterglow lightcurves in data from the RXTE ASM at the locations of possible host galaxies, and parametrized by a generic broken power law with characteristic timescales of minutes to days. The requirement of a GBM or




Figure 15. Fraction of EM counterpart injections to simulated GW triggers which pass thresholds in the EM analysis, as a function of source distance. The EM efficiency is limited by coverage, EM sensitivity, and the success of GW sky localization. The gamma-ray simulation shown is a weak prompt signal lasting 1 s with a standard-candle amplitude of 1 photon/cm²/s in 50–300 keV at 30 Mpc, and following a normal band spectrum according to Table 2. This corresponds to a total energy release of 6.6×10^{46} erg. The ASM X-ray afterglow injection represents a typical X-ray short GRB afterglow shown in Figure 6, and the X-ray burst injections with varying durations are as shown in Figure 6.

ASM coincident counterpart reduced the number of background events by factors of 10^{-4} and 10^{-3} , respectively, reducing the GW amplitude of the loudest surviving background event by ~ 15 – 20% .

Both EM pipelines were tested on a set of joint GW-EM simulated signals, where GW simulations corresponded to NS/NS mergers at random times and sky locations/orientations. The GBM follow-up pipeline maintained sensitivity to weak coincident prompt bursts with a normal GRB spectrum and a fluence corresponding to an isotropic-equivalent energy release of $E_{\text{iso}} \sim 6.6 \times 10^{46}$ erg over 1 s, detecting about 55% of the brightest counterparts (due to Earth occultation) and about 20% of counterparts placed at 40 Mpc (the most distant simulations due to GW sensitivity). The simulated signal was several hundred times weaker than a typical sGRB, implying a wide range of possible GRB luminosities that could be probed using GW triggers up to advanced LIGO distances. The ASM follow-up pipeline was able to detect a typical sGRB afterglow signal with about 40% efficiency at 10 Mpc.

The strategy of targeted offline follow-up of GW candidates in EM data presented here is one of a couple possibilities for joint GW-EM observation—others being rapid online GW analysis to alert pointed telescopes, and using known EM events like GRB alerts to do a targeted specialized search in GW data. All of these strategies will likely be in use in the advanced detector era. The main benefits of having archival EM data at hand is that survey data is continuously available, and not rate or schedule limited. Thousands or even millions of GW candidates can be followed-up offline, provided there is a low-enough probability of accidental coincidence. In the case of a confident GW detection, automated archival EM searches will also be able to look for subthreshold/untriggered (Gruber 2012; Blackburn et al. 2013), potentially exotic EM emission scenarios (e.g., soft flares which are difficult to target on-board due to large background variation), and in the case of

no counterpart, place upper limits on the EM emission through software simulations.

While Fermi will continue operating in the foreseeable future, the *RXTE* mission ended in early 2012 after 16 yr of operation. In the advanced LIGO era, the MAXI mission (Matsuoka et al. 2009) on board the ISS has a similar wide-field soft X-ray camera and continuously surveys the X-ray sky at a sensitivity several times better than *RXTE*-ASM. Another promising candidate for X-ray follow-up is ISS-Lobster (Camp et al. 2013)—a NASA-proposed wide field ($30^\circ \times 30^\circ$) soft X-ray imaging telescope which would also be able to survey the sky (as well as do directed observations). The wide field imager on ISS-Lobster would be ~ 100 times as sensitive as *RXTE*-ASM, matching well the expected factor of 10 increase in GW horizon distance for advanced LIGO. Moreover the ability of ISS-Lobster to repoint for an online GW trigger would allow it to begin observing a rapidly fading afterglow signal as soon as it appears on the right side of the Earth. Both Fermi-GBM and ISS-Lobster would be able to search for counterparts with or without the help of a galaxy catalog due to their extremely wide FOV.

We thank Colleen Wilson-Hodge for assistance with the GBM direct response model. We would also like to thank Collin Capano, Kipp Cannon, Thomas Dent, Leo Singer, Peter Shawhan, Ruslan Vaulin, Xilong Fan and the LIGO Burst and CBC working groups for helpful discussion and ideas. LB did much of this work under the support of a the NASA Postdoctoral Program fellowship at GSFC, administered by Oak Ridge Associated Universities through a contract with NASA. J.V. was supported by the research programme of the Foundation for Fundamental Research on Matter (FOM), which is partially supported by the Netherlands Organization for Scientific Research (NWO), and by Science and Technology Facilities Council (STFC) grant ST/K005014/1. The authors would also like to acknowledge the support of the NSF through grant PHY-1204371. Finally we thank the Albert Einstein Institute in Hannover, supported by the Max-Planck-Gesellschaft, for use of the Atlas high-performance computing cluster.

REFERENCES

- Aasi, J., et al. (LIGO Scientific Collaboration, Virgo Collaboration) 2013, arXiv:1304.0670
- Aasi, J., Abadie, J., Abbott, B. P., et al. 2014, *ApJS*, **211**, 7
- Aasi, J., et al. (The LIGO Scientific Collaboration, the Virgo Collaboration) 2014, *PhRvL*, **113**, 011102
- Abadie, J., Abbott, B. P., Abbott, R., et al. 2012, *ApJ*, **760**, 12
- Abadie, J., et al. (LIGO Scientific Collaboration, Virgo Collaboration) 2010, *CQGra*, **27**, 173001
- Abadie, J., et al. (LIGO Scientific Collaboration, Virgo Collaboration) 2012, arXiv:1205.2216
- Abadie, J., et al. (LIGO Scientific Collaboration) 2012, *ApJ*, **755**, 2
- Abbott, B., et al. (LIGO Scientific Collaboration) 2008, *ApJ*, **681**, 1419
- Abadie, J., et al. (LIGO Scientific Collaboration, Virgo Collaboration) 2012, *PhRvD*, **85**, 082002
- Abadie, J., Abbott, B. P., Abbott, R., et al. 2012, *A&A*, **541**, A155
- Abbott, B., et al. (LIGO Scientific Collaboration) 2009, *RPPh*, **72**, 076901
- Abbott, B., et al. (LIGO Scientific Collaboration, Virgo Collaboration) 2012, *A&A*, **539**, A124
- Accadia, T., et al. (VIRGO Collaboration) 2012, *JINST*, **7**, P03012
- Acernese, F. e. a. (Virgo Collaboration) Advanced Virgo Baseline Design, Technical Report VIR-027A-09 (Cascina, 2009)
- Allen, B. 2005, *PhRvD*, **71**, 062001
- Babak, S., Biswas, R., Brady, P., et al. 2013, *PhRvD*, **87**, 024033
- Blackburn, L., Briggs, M. S., Camp, J., et al. 2013, arXiv:1303.2174

- Cannon, K., Cariou, R., Chapman, A., et al. 2012, *ApJ*, 748, 136
- Camp, J., Barthelmy, S. D., Blackburn, L., et al. 2013, *ExA*, 36, 505
- Christensen, N., & Meyer, R. 2001, *PhRvD*, 64, 022001
- Coward, D., Howell, E., Piran, T., et al. 2012, *MNRAS*, 425, 1365
- Evans, P. A., Fridriksson, J. K., Gehrels, N., et al. 2012, *ApJS*, 203, 28
- Fairhurst, S. 2009, *NJPH*, 11, 123006
- Fairhurst, S. 2011, *CQGra*, 28, 105021
- Fong, W.-f., Berger, E., Margutti, R., et al. 2012, *ApJ*, 756, 189
- Gao, H., Ding, X., Wu, X.-F., Zhang, B., & Dai, Z.-G. 2013, *ApJ*, 771, 86
- Gehrels, N., & Razzaque, S. 2013, *FrPhC*, 8, 661
- Gruber, D. (Fermi/GBM Collaboration) 2012, *PoS, GRB2012*, 036
- Grupe, D., Burrows, D. N., Patel, S. K., et al. 2006, *ApJ*, 653, 462
- Hanna, C., Mandel, I., & Vousden, W. 2014, *ApJ*, 784, 8
- Harry, G. M. (LIGO Scientific Collaboration) 2010, *CQGra*, 27, 084006
- Hjorth, J., & Bloom, J. S. 2012, *CAPS*, 51, 169
- Kanner, J., Camp, J., Racusin, J., Gehrels, N., & White, D. 2012, *ApJ*, 759, 22
- Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, *ApJL*, 413, L101
- Kasliwal, M. M., & Nissanke, S. 2014, *ApJL*, 789, L5
- Levine, A. M., Bradt, H., Cui, W., et al. 1996, *ApJL*, 469, L33
- Liang, E.-W., Racusin, J. L., Zhang, B., Zhang, B.-B., & Burrows, D. N. 2008, *ApJ*, 675, 528
- Matsuoka, M., Kawasaki, K., & Ueno, S. 2009, *PASJ*, 61, 999
- Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, 702, 791
- Metzger, B. D., Quataert, E., & Thompson, T. A. 2008, *MNRAS*, 385, 1455
- Metzger, B., & Berger, E. 2012, *ApJ*, 746, 48
- Nakamura, T., Kashiyama, K., Nakauchi, D., et al. 2013, arXiv:1312.0297
- Nakar, E. 2007, *PhR*, 442, 166
- Nissanke, S., Kasliwal, M., & Georgieva, A. 2013, *ApJ*, 767, 124
- Sidery, T., Aylott, B., & Christensen, N. 2014, *PhRvD*, 89, 084060
- Singer, L. P., Cenko, S. B., Kasliwal, M. M., et al. 2013, *ApJL*, 776, L34
- Skilling, J. 2004, in *AIP Conf. Ser. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 395
- Singer, L. P., Price, L. R., Farr, B., et al. 2014, arXiv: 1404.5623
- Troja, E., Rosswog, S., & Gehrels, N. 2010, *ApJ*, 723, 1711
- Tsang, D., Read, J. S., Hinderer, T., Piro, A. L., & Bondarescu, R. 2012, *PhRvL*, 108, 011102
- van Eerten, H., & MacFadyen, A. 2011, *ApJL*, 733, L37
- Veitch, J., & Vecchio, A. 2010, *PhRvD*, 81, 062003
- Veitch, J., Raymond, V., Farr, B., et al. 2014, arXiv: 1409.7215
- Was, M., Sutton, P. J., Jones, G., & Leonor, I. 2012, *PhRvD*, 86, 022003
- Weisberg, J. M., & Taylor, J. H. 2005, in *ASP Conf. Ser. 328, Binary Radio Pulsars*, ed. F. A. Rasio, & I. H. Stairs (San Francisco, CA: ASP), 25
- White, D. J., Daw, E., & Dhillon, V. 2011, *CQGra*, 28, 085016
- Zhang, B., Fan, Y. Z., Dyks, J., et al. 2006, *ApJ*, 642, 354
- Zhang, B. 2013, *ApJL*, 763, L22