



Searches for Modulated γ -Ray Precursors to Compact Binary Mergers in Fermi-GBM Data

Cosmin Stachie, Tito Dal Canton, Nelson Christensen, Marie-Anne Bizouard,
Michael Briggs, Eric Burns, Jordan Camp, Michael Coughlin

► To cite this version:

Cosmin Stachie, Tito Dal Canton, Nelson Christensen, Marie-Anne Bizouard, Michael Briggs, et al.. Searches for Modulated γ -Ray Precursors to Compact Binary Mergers in Fermi-GBM Data. *The Astrophysical Journal*, 2022, 930 (1), pp.45. [⟨10.3847/1538-4357/ac5f53⟩](#). [⟨hal-03664006⟩](#)

HAL Id: hal-03664006

<https://hal.science/hal-03664006v1>

Submitted on 25 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



Searches for Modulated γ -Ray Precursors to Compact Binary Mergers in Fermi-GBM Data

Cosmin Stachie¹ , Tito Dal Canton^{2,3} , Nelson Christensen¹ , Marie-Anne Bizouard¹ , Michael Briggs⁴, Eric Burns⁵ , Jordan Camp², and Michael Coughlin⁶

¹ Artemis, Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Nice 06300, France; scosmin@oca.eu

² NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771, USA

³ Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France

⁴ Department of Space Science, University of Alabama in Huntsville, Huntsville, AL 35899, USA

⁵ Department of Physics & Astronomy, Louisiana State University, Baton Rouge, LA 70803, USA

⁶ School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

Received 2021 December 8; revised 2022 March 13; accepted 2022 March 18; published 2022 May 3

Abstract

GW170817 is the only gravitational-wave event for which a confirmed γ -ray counterpart, GRB 170817A, has been detected. Here, we present a method to search for another type of γ -ray signal, a γ -ray burst precursor, associated with a compact binary merger. If emitted shortly before the coalescence, a high-energy electromagnetic (EM) flash travels through a highly dynamical and relativistic environment, created by the two compact objects orbiting each other. Thus, the EM signal arriving at an Earth observer could present a somewhat predictable time-dependent modulation. We describe a targeted search method for light curves exhibiting such a modulation, parameterized by the observer-frame component masses and binary merger time, using Fermi-GBM data. The sensitivity of the method is assessed based on simulated signals added to GBM data. The method is then applied to a selection of potentially interesting compact binary mergers detected during the second (O2) and third (O3) observing runs of Advanced LIGO and Advanced Virgo. We find no significant modulated γ -ray precursor signal associated with any of the considered events.

Unified Astronomy Thesaurus concepts: Neutron stars (1108); Black holes (162); Compact binary stars (283)

1. Introduction

Multimessenger astronomy started with the detection of the core-collapse supernova 1987A in both the electromagnetic (EM) and neutrino channels (Arnett et al. 1989). Since then, only one other unambiguous example of astrophysical event heralded by different messengers occurred: the simultaneous detection of the gravitational waves (GWs) from the binary neutron star (BNS) coalescence GW170817 (Abbott et al. 2017), and of several EM counterpart signals: the high-energy photons of GRB 170817A (Goldstein et al. 2017; Savchenko et al. 2017), the ultraviolet, optical and infrared radiation of the kilonova AT 2017gfo (Coulter et al. 2017; Smartt et al. 2017) and the X-ray, optical, and radio afterglow of the γ -ray burst (Lamb et al. 2019; D'Avanzo et al. 2018). Additionally, there is a convincing claim for the coincident detection of the high-energy neutrino IceCube-170922A and the multiwavelength EM radiation coming from the γ -ray blazar TXS 0506+056 (IceCube Collaboration et al. 2018).

In the observable universe, compact objects such as neutron stars and black holes are often found in pairs, forming binaries. During the inspiral, they lose angular momentum and binding energy, by emission of GWs (Taylor & Weisberg 1982). This implies a narrowing of the distance separating the binary components, leading in some cases to a merger in less than a Hubble time (Phinney 1991). Both the frequency and the amplitude of the GWs increase as the merger approaches; during the seconds prior to the merger, the frequency of the GWs sweeps from tens of hertz to above a kilohertz. If the binary is close enough to us, the GW

strain is above the sensitivity threshold of and thereby detectable by GW detectors such as Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acemese et al. 2015).

One of the EM counterparts to compact binary mergers is the flash of γ - and X-rays from the short γ -ray burst (Kochanek & Piran 1993; Berger 2014), lasting less than 2 s (Kouveliotou et al. 1993) and possessing an isotropic-equivalent energy up to 10^{51} erg (D'Avanzo 2015). This short, high-energy EM signal is followed by a longer-lasting, less energetic radiation, the γ -ray burst afterglow (Fox et al. 2005) covering a broad range of the EM spectra, from X-ray (Vietri 1997), through optical (Meszaros & Rees 1997) to radio (Paczynski & Rhoads 1993). The kilonova comprises the ultraviolet, optical, and infrared radiation associated with the radioactive decay of heavy elements (Lattimer & Schramm 1974; Li & Paczynski 1998; Metzger et al. 2010; Kasen et al. 2017); it has a quasi-thermal (Kasen & Barnes 2019) and quasi-isotropic emission (Darbha & Kasen 2020), very different from the γ -ray burst afterglows, which display large anisotropies (Beloborodov et al. 2011), and consequently require nearly aligned observer positions in order to be detectable.

While the association between the merger of compact objects and some γ -ray burst related messengers, such as the γ -ray prompt emission, the afterglow, and the kilonova, has been unambiguously highlighted by GW170817, the presence of other kinds of EM emission is still debated. One such example is the precursor activity to short γ -ray bursts. Troja et al. (2010) claim that up to 10% of short γ -ray bursts, detected by Swift (Barthelmy et al. 2005), possess EM precursors, lasting less than 1 s and whose starting emission might go back about 100 s before the main γ -ray burst. The existence of such precursor signals is also supported by Zhu (2015). The investigation by Li et al. (2021) concludes that the precursors have shorter duration than the prompt γ -ray emission, but seem



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

to be produced by similar central engine activity. 16 precursors to short γ -ray bursts in Fermi-GBM data have been found by Wang et al. (2020), who infer comparable duration for the main and precursor emissions, and possible fits with blackbody, nonthermal cutoff power-law models. Precursor emission is also motivated on a theoretical basis. Two popular models are the resonant shattering of the crust of a neutron star during the inspiral (Tsang et al. 2012; Neill et al. 2021) and the *black hole battery* model (McWilliams & Levin 2011). Theoretical work has also been done to identify possible features of an EM signal emitted during the inspiral, namely, a modulation induced by the orbital motion (Schnittman et al. 2018).

The prompt γ emission, afterglows, and kilonovae described above are currently expected to follow from the disruption of a neutron star and the formation of an accretion disk during the merger of the binary. However, whether a neutron star will actually disrupt before merging with the companion strongly depends on the properties of the two objects, in particular on their masses, spins, and structure (Foucart et al. 2018). Although BNS mergers are generally always expected to radiate the whole plethora of EM signals observed with GW170817, this is far from being guaranteed for neutron star–black hole (NSBH) mergers. In addition, the prompt γ emission is expected to be detectable only for very specific orientations of the binary with respect to the observer. Hence, the idea of a premerger, precursor emission that does not require disruption and may not be strongly anisotropic becomes particularly interesting: it may be the only EM signal systematically emitted by NSBH mergers, even those involving non-spinning black holes more massive than $\approx 10 M_{\odot}$.

Motivated by the above consideration, and by the present ambiguity regarding the possibility of precursor emission, we propose a method to analyze archival γ -ray data in temporal proximity to GW events associated with BNS mergers. Following the idea of Schnittman et al. (2018), the analysis aims at detecting pulsations in the γ -ray data having the same orbital phase evolution like the GW inspiral signal, during the last several seconds before merger. The method is an extension of a more generic search for γ -ray transients based on a likelihood approach, introduced by Blackburn et al. (2015) and further optimized in Goldstein et al. (2016, 2019). We demonstrate how the existing and proposed methods recover simulated signals added to archival data, and we apply the proposed method to data around a selection of compact binary merger events detected by Advanced LIGO and Advanced Virgo.

We employ data from the γ -ray Burst Monitor (GBM) instrument on board the Fermi spacecraft (Meegan et al. 2009). GBM is an ideal instrument to study rapidly evolving γ -ray counterparts to compact binary mergers. It is able to perform time-resolved spectroscopy of high-energy EM radiation, thanks to its 12 sodium iodide (NaI) and two bismuth germanate (BGO) scintillation detectors, covering an energy range from 8 keV to 40 MeV (Atwood 1994), and distributed around the spacecraft. It has a temporal resolution of $2 \mu\text{s}$, suitable to study the rapid modulations we are interested in, and can discriminate between 128 energy ranges/channels. Its field of view is only limited by Earth, and its observing time is mostly limited by passages through the South Atlantic Anomaly (SAA), implying it can witness an astrophysical transient event 75% of the time.

The paper is structured as follows: in Section 2, we discuss the possible physical mechanisms responsible for the γ -ray precursor activity, we introduce a modulated EM waveform model and present how it can be used to simulate data from a γ -ray detector.

The statistical search method is presented in Section 3. The sensitivity of the search is presented in Section 4, and the results of the search around some of the LIGO-Virgo events appear in Section 5. Finally, Section 6 offers the conclusion.

2. Modulated γ -Ray Precursors

2.1. Physical Mechanisms and Signal Models

While the physical mechanism responsible for such a γ -ray precursor is still in question, several possibilities have been suggested in the literature. Palenzuela et al. (2013) and Most & Philippov (2020) show that magnetosphere interaction in a BNS can power EM radiation, prior to the main emission, with luminosities reaching $10^{45} \text{ erg s}^{-1}$. Copious EM radiation, prior to the merger of a magnetized neutron star-spinning black hole binary, can also be emitted through the unipolar inductor mechanism (McWilliams & Levin 2011; Palenzuela et al. 2011; D’Orazio & Levin 2013; D’Orazio et al. 2016): indeed an electric circuit is formed, where the roles of battery, resistor, and electrical wires are played by the black hole, the neutron star and its magnetosphere, and the magnetic field lines; in this way, black hole rotational energy is extracted by the magnetic field and sent far away by means of powerful, collimated Poynting flux.

During the last orbits of a quasicircular inspiral, or a periastron passage in an eccentric or hyperbolic close encounter, a neutron star can experience tidal forces that excite some of its many oscillation modes, e.g., core, shear, and crustal discontinuity modes (Lai 1994; Shibata 1994). This process, especially when resonant, can cause quakes and/or shattering of the neutron star crust, followed by the release of a huge amount of energy in the form of EM radiation (Reisenegger & Goldreich 1994; Tsang et al. 2012; Tsang 2013; Suvorov & Kokkotas 2020).

While traditionally γ -ray burst emissions are detected by searching for an excess of photons with respect to the background (Meegan et al. 2009; Kocevski et al. 2018; Burgess et al. 2016), in the present work we attempt to increase the sensitivity of the pipeline by hunting for well-characterized EM waveforms. Schnittman et al. (2018) propose a method to calculate the EM precursor light curve from an NSBH, where the surface of the neutron star is uniformly shining. A similar method is proposed in Haiman (2017) for the case of a supermassive black hole binary, which are future LISA (Amaro-Seoane et al. 2017) signal progenitors. In the preceding papers, the light curve is modulated by physical processes such as relativistic beaming and gravitational lensing. There is a parallelism to the compact binary coalescence GW waveforms as the EM light curve could be locked in phase with the GW signal. As the merger time approaches, the amplitude and the frequency of the presumed EM signal increase. Radio precursor light curves to compact binaries containing at least one magnetized neutron star are also suggested in Sridhar et al. (2021).

The association of GWs and γ -ray burst precursors may represent an unique class of multimessenger events, in the coming years. Indeed, the theoretically motivated signals presented above are expected to be visible only for nearby events. As GW detections will be nearby too, this is a huge gain for this kind of work. In addition, NSBH systems, for which the mass ratio is too high, are not expected to generate neither short γ -ray bursts nor kilonovae because the neutron star component is swallowed by the black hole companion before being disrupted by the tidal field. On the other hand, these systems might power precursors of γ -ray burst, during the inspiral phase. Moreover, light-curve models like the one

proposed in Schnittman et al. (2018) are interesting because they give information about the potential γ -ray signal we want to detect, based on the GW detection.

2.2. Orbital-modulation Model

An important feature of an EM light curve is the evolution of the brightness with time. In this work, given a binary inspiral, we are interested in the orbital phase dependence of the luminosity. For a compact binary, the first post-Newtonian expansion (Blanchet 2014) term of the orbital angular frequency evolution can be written as

$$\Omega(t) = \left(\frac{5}{256} \frac{1}{t_c - t} \right)^{3/8} \left(\frac{GM}{c^3} \right)^{-5/8}, \quad (1)$$

where $\mathcal{M} = (m_1 m_2)^{3/5} (m_1 + m_2)^{-1.5}$, G and c are the observer-frame chirp mass, the gravitational constant, and the speed of light in vacuum, respectively. t_c and t are the merger and variable time measured in the observer frame. m_1 and m_2 are the binary component masses. We define the orbital phase, $\Phi_{\text{orbit}}(t) \in [0, 2\pi]$, with the origin at $t = -30$ s, by $\Phi_{\text{orbit}}(t) = \int_{-30}^t \Omega(x) dx \bmod 2\pi$.

We also introduce the *simplified light curve*, a light curve which, besides the chirp mass, depends on two other parameters, namely, θ_{peak} and θ_{width} . The normalized luminosity has the following expression:

$$\frac{I(\Phi_{\text{orbit}})}{I_{\text{max}}} = \begin{cases} 1 & \text{if } |\Phi_{\text{orbit}} - \theta_{\text{peak}}| \leq \theta_{\text{width}}/2 \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

This expression corresponds to a compact binary emitting only during the orbital phase window centered at θ_{peak} , with width θ_{width} . In Figure 1, there is an illustration of a simplified light curve. Additionally, for comparison purposes, we show a more realistic (albeit still approximate) light curve that attempts to explicitly model the relativistic beaming and gravitational lensing effects described in Schnittman et al. (2018). Although we will not use this model for the analyses described later, more details of its construction are presented in the Appendix.

Figure 1 shows that the blue simplified light curve and the more realistic orange light curve are similar. The choice of the simplified light curve for the present paper is motivated by the pronounced modulation feature, limiting the spread of incoming photons in regions where the light flux is negligible. While the exacerbated modulation of the simplified light curve is most likely unrealistic, the preference for this model is mainly motivated by our attempt to check the validity of the search method in the case of signals with optimistically large modulation.




Figure 1. Normalized light intensity vs. time. In blue, the *simplified light curve* with $\theta_{\text{peak}} = 270^\circ$ and $\theta_{\text{width}} = 40^\circ$. In orange, a light curve obtained by the combination of relativistic beaming and gravitational lensing effects, for which we chose an inclination angle of 45° and a spectral index equal to 0. For both light curves, the compact objects are assumed to have masses $m_1 = 10 M_\odot$ and $m_2 = 1.4 M_\odot$.

2.3. Simulating a Fermi-GBM Observation

We use Fermi-GBM's Time-Tagged Event (TTE) data,⁷ which consists of a list of photons characterized by their arrival time (to microsecond precision) and energy channel. In our analysis, for simplicity and to increase the photon statistic, we coalesce the 128 possible energy channels into eight main channels, larger in energy.

In order to simulate a simplified light-curve-like signal, one needs to convert the spectra into detector counts. To this end, one needs the detector response to radiation, characterized by the energy of its photons, the arrival direction, and the amplitude of the light curve. This function includes two components: the response due to the direct radiation (Kippen et al. 2007) as well as the response due to the scattering from both Earth's atmosphere and the spacecraft (Pendleton et al. 1999). The spectral and directional dependence was validated experimentally by ground-based calibration (Bissaldi et al. 2009). As explained in Connaughton et al. (2015), based on the numerical values of the detector response taken at 41,168 grid points (accounting for 272 sky directions), the response to any arrival direction is constructed by interpolation among the three closest grid points, obtained by Delaunay triangulation. Regarding the radiation energetics, we make use of three spectral templates, which we refer to as *soft* (lowest energy), *normal*, and *hard* (highest energy). For the soft and normal templates, we use the Band parameterized functions (Band et al. 1993), introduced in Connaughton et al. (2015). Thus, the flux of photons having an energy in $[E, E + dE]$ is proportional to

$$\begin{cases} \left(\frac{E}{100 \text{ keV}} \right)^{\alpha_{\text{spec}}} e^{-\frac{(\alpha_{\text{spec}}+2)E}{E_{\text{peak}}}} & \text{if } E < \frac{(\alpha_{\text{spec}} - \beta_{\text{spec}})E_{\text{peak}}}{\alpha_{\text{spec}} + 2} \\ \left(\frac{E}{100 \text{ keV}} \right)^{\beta_{\text{spec}}} e^{(\beta_{\text{spec}} - \alpha_{\text{spec}}) \left[\frac{(\alpha_{\text{spec}} - \beta_{\text{spec}})E_{\text{peak}}}{(\alpha_{\text{spec}} + 2)100 \text{ keV}} \right]} & \text{otherwise} \end{cases}, \quad (3)$$

⁷ <https://fermi.gsfc.nasa.gov/ssc/data/access/gbm/>

with $(\alpha_{\text{spec}}, \beta_{\text{spec}}, E_{\text{peak}})$ equal to $(-1.9, -3.7, 70 \text{ keV})$ and $(-1, -2.3, 230 \text{ keV})$ for the soft and normal spectrum. With respect to the hard spectrum, we employ the Comptonized template proposed in Goldstein et al. (2016), such that the photon flux is proportional to

$$\left(\frac{E}{E_{\text{piv}}}\right)^{\alpha_{\text{spec}}} \exp\left[-\frac{(\alpha_{\text{spec}} + 2)E}{E_{\text{peak}}}\right], \quad (4)$$

where E_{piv} is a constant and $(\alpha_{\text{peak}}, E_{\text{peak}})$ equals $(-0.5, 1.5 \text{ MeV})$. However, the γ -ray bursts detected to date reveal a richer set of spectra than the three templates presented here. This fact motivates our choice of considering injected light curves parameterized by the variable $\kappa \in [0, 2]$, such that $\kappa = 0, 1$, and 2 correspond to the hard, normal, and soft spectrum. For all the other values, we consider linear combinations of the three templates with the following weights:

$$(w_0, w_1, w_2) = \begin{cases} (1 - \kappa, \kappa, 0) & \text{if } \kappa \in (0, 1) \\ (0, 2 - \kappa, \kappa - 1) & \text{if } \kappa \in (1, 2) \end{cases} \quad (5)$$

where w_0, w_1 , and w_2 are the weights attributed to the hard, normal and soft spectrum templates.

The detector output to an input flux of high-energy photons is captured by the response matrix. More precisely, the response matrix is a 14×8 array, where each row designates one of the 14 detectors and each column stands for an energy channel. Finally, each element of the response matrix represents the photon rate as a function of time. As the TTE data counts are assigned arrival times, we convert the photon rate function into a time histogram using a Poisson distribution. Moreover, for technical reasons, the data is binned. More precisely, the time is divided in intervals of size equal to 0.5 ms , and the photons found in the same interval/bin and belonging to the same energy channel are summed. According to the Nyquist–Shannon sampling theorem, such a binning allows the preservation of signal frequency components up to 1000 Hz . According to Equation (1), the orbital frequency is always lower than this upper limit, for all binaries where the heavier compact object weighs more than $1.4 M_{\odot}$ and for a binary evolution up to the last millisecond before the merger. Finally a simplified light-curve injection is completely parameterized by the tuple $(t_c, m_1, m_2, f_{\text{amp}}, t_{\text{start}}, \Delta t_{\text{dur}}, \theta_{\text{peak}}, \theta_{\text{width}}, \theta_{\text{ra}}, \theta_{\text{dec}}, \kappa)$, where the variables $f_{\text{amp}}, t_{\text{start}}, \Delta t_{\text{dur}}, \theta_{\text{ra}}$, and θ_{dec} designate the light-curve amplitude factor, the EM signal start time, the signal duration, the right ascension, and the declination associated with the sky location. It is worth mentioning that the photon flux is obtained by the multiplication of the amplitude factor f_{amp} with the expression of the normalized simplified light curve proposed in Equation (2).

3. Statistical Search Method

We propose a statistical framework that adapts the search methods presented in Goldstein et al. (2016) and Blackburn et al. (2015). Hereafter, we refer to the original search as the *generic targeted search*, since it aims at detecting a generic transient excess of high-energy photons above the detector background associated with a particular target time, regardless of its temporal morphology. We refer to our modified method as the *chirp targeted search* instead, as it aims at detecting excesses of photons that exhibit repetitions locked in phase with the orbital evolution of a compact binary that is about to merge. The main

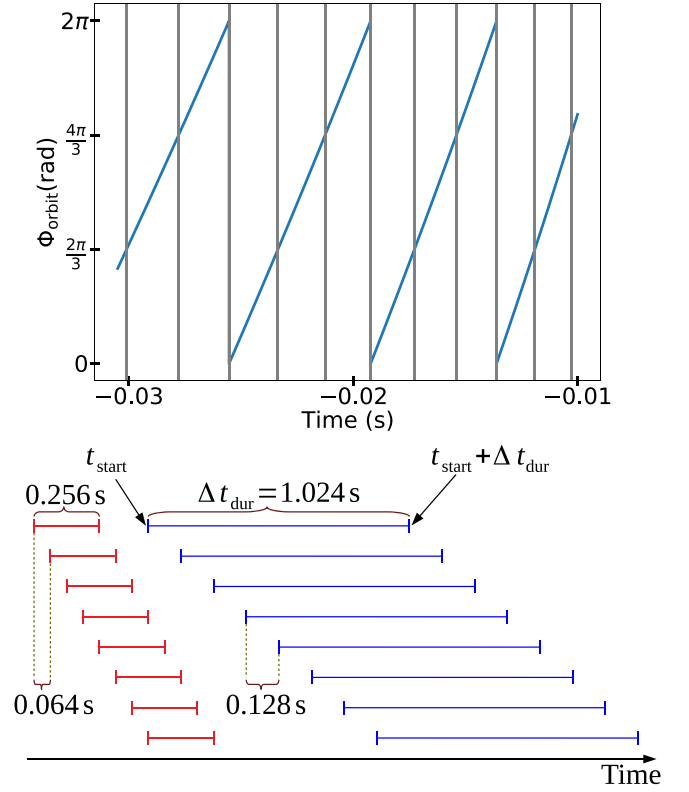


Figure 2. Top panel: time dependence of the orbital phase and positions of I_k intervals in the case $N_{\text{bins}} = 3$; the binary has $(m_1, m_2) = (10 M_{\odot}, 1.4 M_{\odot})$ and the origin of the x-axis coincides with the merger time. Bottom panel: positions of the generic targeted search time windows $[t_{\text{start}}, t_{\text{start}} + \Delta t_{\text{dur}}]$.

difference with respect to the generic targeted search is that we apply its statistical formalism in the orbital phase space, instead of the time space. The transformation from time space to orbital phase space is provided, approximately, by the coalescence time t_c and the component masses m_1 and m_2 inferred from the GW signal. For a fixed number of bins N_{bins} , the orbital phase is split into N_{bins} equal intervals, i.e., $I_0 = \left[0, \frac{2\pi}{N_{\text{bins}}}\right]$, $I_1 = \left[\frac{2\pi}{N_{\text{bins}}}, 2\frac{2\pi}{N_{\text{bins}}}\right]$, ..., $I_{N_{\text{bins}}-1} = \left[(N_{\text{bins}} - 1)\frac{2\pi}{N_{\text{bins}}}, 2\pi\right]$. The top panel of Figure 2 illustrates the conversion from the time space to the orbital phase space. Once this transformation is defined, we can rebin the photons registered by GBM (the TTE data) into the orbital phase intervals: using each photon's arrival time t , we identify its orbital phase interval I_k such that $\Phi_{\text{orbit}}(t) \in I_k$.

Next, we need an estimate of the background photon rate in each phase interval, i.e., the rate of photons registered by GBM in the absence of a modulated transient. To this end, we first estimate the background rate over time using the unbinned Poisson maximum likelihood technique introduced in Goldstein et al. (2016): at a given time t_0 , the background photon rate $\lambda_{\text{max}}(t_0)$ is defined as the ratio between the number of photons N_{photons} contained in a large enough time window of duration T (in our case $T = 100 \text{ s}$) and the width of the window, i.e., $\lambda_{\text{max}}(t_0) = \frac{N_{\text{photons}}}{T}$. Assuming the background can be described as a stationary Poisson process over the search interval, the uncertainty in its rate can be written as $\sigma_{\lambda_{\text{max}}}^2(t_0) = \lambda_{\text{max}}(t_0)/T$. The time window is slid over the time range of interest; thus, a background photon rate (assigned

with standard deviation) is calculated for an array of times, and finally the photon rates (and their uncertainties) are interpolated over time. A chi-squared statistic χ^2 is computed in order to evaluate the quality of the fit. If the fit is poor, i.e., χ^2 is too large, for one of the GBM energy channels, that channel is excluded from the search. The background rate is then transformed to the orbital phase space to predict the rate in each phase interval.

Once the foreground photon histogram (background photon fitting) is calculated (estimated) for the $[0, 2\pi]$ orbital phase range, we aim to search for a subset of adjacent intervals I_k presenting an excess of photons. Such a behavior is equivalent to saying that the binary, during an orbit, emits the majority of the radiation in a specific orbital phase window. This feature is characteristic to the simplified light curves.

A quantity combining information about both source and noise is the likelihood ratio, defined as

$$\Lambda(d) = \frac{P(d|H_1)}{P(d|H_0)}, \quad (6)$$

where d , H_1 , and H_0 are the observed data, the signal presence hypothesis, and the hypothesis of noise alone. As in Blackburn et al. (2015), the assumption of uncorrelated Gaussian noise allows us to write the preceding probabilities in the following way:

$$P(d|H_1, s) = \prod_i \frac{1}{\sqrt{2\pi} \sigma_{d_i}} \exp\left(-\frac{(\tilde{d}_i - r_i s)^2}{2\sigma_{d_i}^2}\right), \quad (7)$$

$$P(d|H_0) = \prod_i \frac{1}{\sqrt{2\pi} \sigma_{n_i}} \exp\left(-\frac{\tilde{d}_i^2}{2\sigma_{n_i}^2}\right). \quad (8)$$

In those previous formulas, $\tilde{d}_i = d_i - \langle n_i \rangle$, where d_i and $\langle n_i \rangle$ are the foreground and the estimated background photons. $\sigma_{n_i}^2$ and $\sigma_{d_i}^2$ represent the variances of the background and the expected data (background plus signal). The variances appearing in Equations (7) and (8) are computed in the orbital phase space, and are obtained by the summation of time space variances. Lastly the time space standard deviations are calculated as in Blackburn et al. (2015). r_i stands for the detector-energy response, which depends on both the EM source sky location and spectrum. Finally s is the amplitude (measured on Earth) of the signal. Moreover, for all these quantities, the index i designates a pair (detector, energy channel). Given that σ_{d_i} , σ_{n_i} , and \tilde{d}_i are measured quantities, while r_i is calculated for a sample grid, accounting for all possible locations, as explained in Kocevski et al. (2018), then the amplitude parameter s is the only variable over which the marginalization needs to be done. Thus, the expression of the likelihood ratio becomes $\Lambda(d) = \int \frac{P(d|H_1, s)}{P(d|H_0)} P(s) ds$, where $P(s)$ is the amplitude signal prior. Maximizing the likelihood ratio is the same as maximizing the logarithm of it, $\mathcal{L}(d) = \ln \Lambda(d)$. As explained in Blackburn et al. 2013, $\ln \Lambda(d|s)$ is almost a Gaussian function with variance $\sigma_{\ln \Lambda(d|s)}^2 = \frac{1}{\sum_i r_i^2 / \sigma_{d_i}^2}$. Therefore the maximum of $\ln \Lambda(d|s)$ is reached for s_{best} , obtained by means of the iterative Newton's method. The $(k+1)$ th step in the Newton's method consists in refining the k th estimate by using the analytic second derivative, i.e., $s_{k+1} \approx s_k - \frac{\partial \mathcal{L} / \partial s}{\partial^2 \mathcal{L} / \partial s^2}$,

while the initial guess is $s_0 = \frac{\sum_i r_i \tilde{d}_i / \sigma_{d_i}^2}{\sum_i r_i^2 / \sigma_{d_i}^2}$. We assume the same well-behaved prior as in Kocevski et al. (2018), i.e.,

$$P(s) = \left[1 - \exp\left(-\left(\frac{s}{\gamma_{\text{prior}} \sigma_{\ln \Lambda(d|s)}}\right)^{\beta_{\text{prior}}}\right) \right] s^{-\beta_{\text{prior}}}, \quad (9)$$

where $\gamma_{\text{prior}} = 2.5$ and $\beta_{\text{prior}} = 1$. The value of γ_{prior} ensures a prior almost constant over a range of $\sigma_{\mathcal{L}}$, while the value of β_{prior} translates in a luminosity distribution independent of distance. The log-likelihood ratio becomes

$$\begin{aligned} \mathcal{L}(d) = & \ln \sigma_{\ln \Lambda(d|s)} + \ln \left[1 + \operatorname{erf}\left(\frac{s_{\text{best}}}{\sqrt{2} \sigma_{\ln \Lambda(d|s)}}\right) \right] + \ln \Lambda(d|s_{\text{best}}) \\ & + \begin{cases} \ln \left[1 - \exp\left(-\frac{s_{\text{best}}}{\gamma_{\text{prior}} \sigma_{\ln \Lambda(d|s)}}\right) \right] - \beta_{\text{prior}} \ln s_{\text{best}} & \text{if } s_{\text{best}} \geq 0 \\ -\beta_{\text{prior}} \ln (\gamma_{\text{prior}} \sigma_{\ln \Lambda(d|s)}) & \text{if } s_{\text{best}} \leq 0 \end{cases}. \end{aligned} \quad (10)$$

In Equation (10), $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$ is the error function. And finally, as explained in Blackburn et al. (2015), we calibrate the log-likelihood ratio (hereafter LLR) by subtracting the quantity $\mathcal{L}_{\text{ref}} = \beta_{\text{prior}} \ln \gamma_{\text{prior}} + (1 - \beta_{\text{prior}}) \ln \sigma_{\text{ref}}$. Here, we do not care about the value of σ_{ref} because $\beta_{\text{prior}} = 1$, and so the term $(1 - \beta_{\text{prior}}) \ln \sigma_{\text{ref}}$ cancels.

Various kinds of transients commonly appear in Fermi-GBM data and produce large LLR values, despite being certainly unrelated to GW events. One such class is represented by high-energy cosmic-rays hitting the NaI detectors responsible for long-lived phosphorescent light emission. By means of the technique introduced in Blackburn et al. (2015), we remove most undesirable high-LLR triggers from this class. A second class is represented by sharp photon rate changes due to Fermi approaching the SAA. These are discarded instead as explained in Goldstein et al. (2016).

Both the generic targeted search and our chirp targeted search take as input a GPS time t_c , which is the compact binary merger time, measured at Fermi. In this work, the generic targeted search is performed over the following exact timescales: 0.064, 0.128, 0.256, 0.512, 1.024, 2.048, 4.096, and 8.192 s. It uses a time displacement of 64 ms for the four shortest timescales and a time displacement of factor 8 for the four longest timescales (e.g., the 8.192 s search windows are separated by 1.024 s). Thus, in total there are 2270 such windows. The positions of some search windows on the time axis are shown in the bottom panel of Figure 2. For the chirp targeted search, we shorten/extend the same search windows by a small amount in such a way that each window contains an integer number of orbits. This procedure is important in order to avoid artificially unequal number of photons in different I_k intervals, which would produce artificially large LLR values. Additionally, in the case of the chirp targeted search, for a fixed $(t_{\text{start}}, \Delta t_{\text{dur}})$, searches are done over 47 subsets of adjacent intervals I_k . More precisely, we consider subsets of any length in $\{1, 2, \dots, N_{\text{bins}}\}$ and we use a phase factor of 2 (e.g., subsets of length 2, 3, and 6 are separated by 1, 2, and 3 intervals I_k). Finally the most significant trigger, i.e., with the highest LLR, is reported.

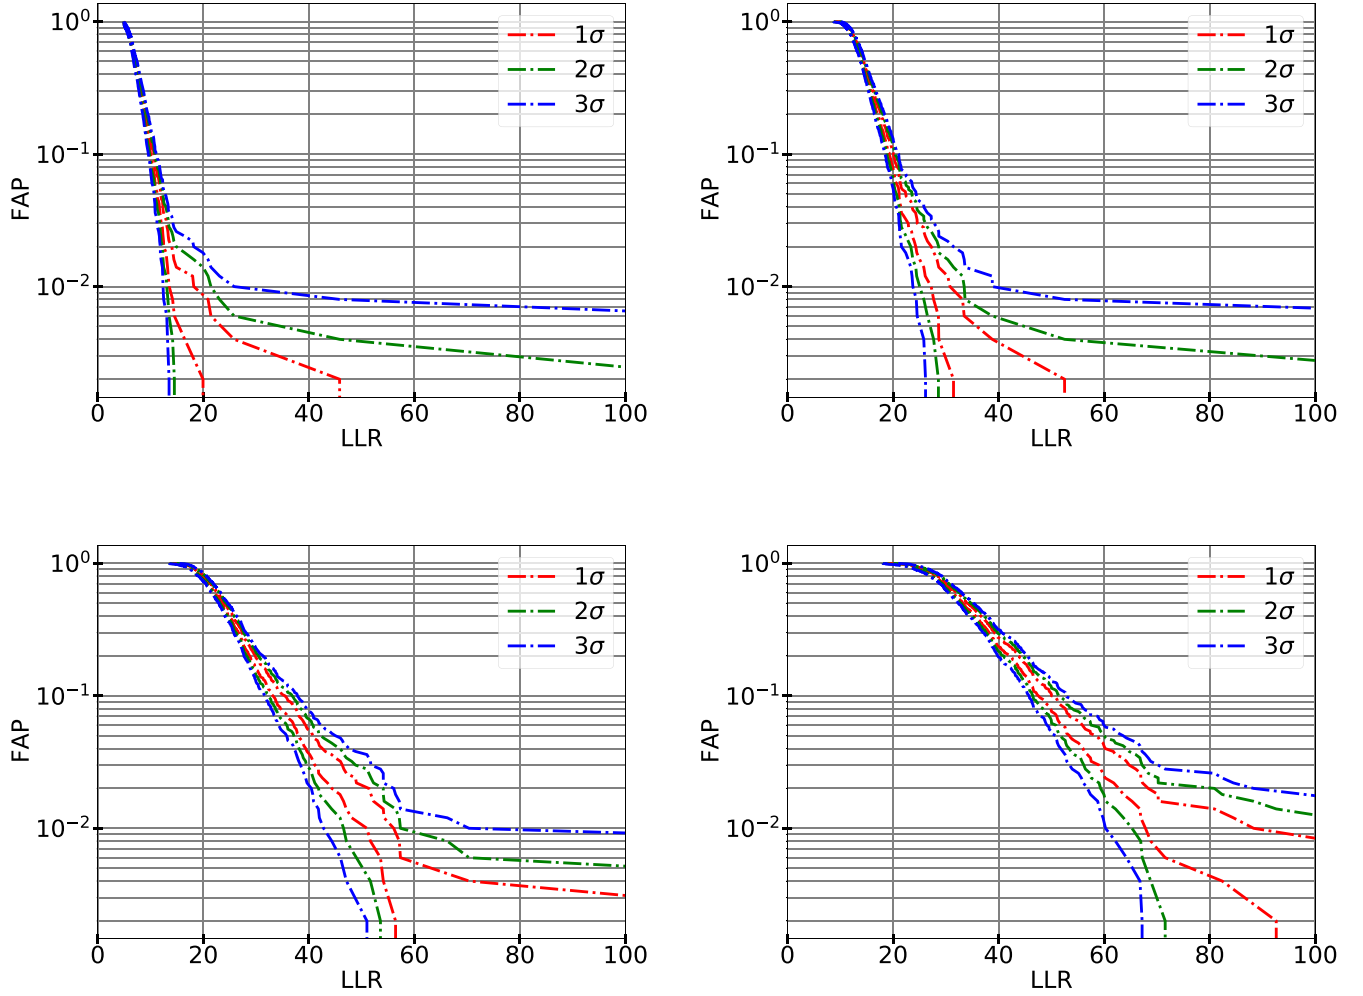


Figure 3. FAP vs. LLR of triggers by trying the two searches on random times. Plotted are the background distributions assigned with $\pm 1\sigma$, $\pm 2\sigma$, and $\pm 3\sigma$ uncertainties. The top left panel corresponds to the distribution of the generic targeted search output triggers, while the top right, bottom left, and bottom right panels represent the chirp targeted search distributions with settings $N_{\text{bins}} = 5, 10$, and 15 . For all panels of the chirp targeted search, the setting $(m_1, m_2) = (1.6 M_\odot, 1.4 M_\odot)$ is used.

3.1. Null Distribution of LLR

It is worth mentioning that a statistically significant trigger is not necessarily a GW-related signal, and sometimes not even the effect of an EM radiation intercepted by the NaI and/or BGO detectors. Despite the filtering strategy discussed in Section 3, large LLR spurious signals survive. For this reason, an empirical measure of the false alarm probability (FAP) distribution is extremely useful. In this paper, for a given LLR_0 , the value of $\text{FAP}(\text{LLR}_0)$ represents the fraction of noise events in which at least one set member has a statistical significance higher than LLR_0 . In Figure 3, we show the FAP distribution of LLR, obtained by running the search on 1000 random times, spread over the period of 2019 April 1–2020 April 1. The random times have been chosen such that they are at least 30 s away from the SAA entrance/exit. This choice is motivated by the exclusion of those situations where the 30 s time windows would otherwise analyze non-science times. Based on Figure 3, at least two important remarks should be made: (i) even though similar, the distributions corresponding to the chirp targeted search have higher LLRs with respect to the generic targeted search distribution; and (ii) the noise output triggers have higher LLR as the number of bins N_{bins} increases. Remark (i) is an expected behavior, because the generic targeted search is

included in the chirp targeted search. In fact, the evaluation of the statistical significance of all I_k intervals together is equivalent to performing the generic targeted search. Regarding (ii), there are at least two reasons favoring this behavior: (a) for two bins numbers $N_{\text{bins},1}$ and $N_{\text{bins},2}$, with $N_{\text{bins},2} > N_{\text{bins},1}$, such that $N_{\text{bins},2}$ is a multiple of $N_{\text{bins},1}$, the chirp targeted search with setting $N_{\text{bins},1}$ is included in the chirp targeted search with setting $N_{\text{bins},2}$; (b) the higher the number of bins N_{bins} , the less true is the approximation of Gaussian background noise in the high-energy detectors, on the scale of a I_k interval.

4. Search Sensitivity

In this section, we test the sensitivity of our chirp targeted search and we compare it to the sensitivity of the generic targeted search. In the case of modulated γ -ray signals, in phase with GWs, we expect better performance for the chirp targeted search.

4.1. Properties of the Simulated Signals

In this subsection, we describe the simplified light-curve injections in the GBM data. We consider signals with durations Δt_{dur} log-uniformly distributed in $[0.064, 8.192 \text{ s}]$. For a fixed Δt_{dur} , the beginning of the radiation, t_{start} , is sampled uniformly




Figure 4. Fraction of detected injections vs. FAP for the simplified light-curve model. The left (right) panels correspond to injections with $(m_1, m_2) = (1.6 M_\odot, 1.4 M_\odot)$ (respectively $(m_1, m_2) = (10 M_\odot, 1.4 M_\odot)$). The injected signals are recovered with the generic targeted search and the chirp targeted search having the correct (m_1, m_2) setting, while N_{bins} is varied.

in $[t_c - 30 \text{ s}, t_c - \Delta t_{\text{dur}}]$. The injections are uniformly distributed in the sky with $\theta_{\text{ra}} \in [-90^\circ, 90^\circ]$ and $\theta_{\text{dec}} \in [-180^\circ, 180^\circ]$, while the spectral index κ is uniform in $[0, 2]$. For all the injections, we fix $\theta_{\text{width}} = 10^\circ$, while θ_{peak} is random in $[0^\circ, 360^\circ]$. The light-curve amplitude factor f_{amp} is uniformly distributed in $[20 \sqrt{\frac{0.064 \text{ ms}}{\Delta t_{\text{dur}}}}, 50 \sqrt{\frac{0.064 \text{ ms}}{\Delta t_{\text{dur}}}}]$. This choice was found empirically in order to respect the following requirements: the lower and upper limits impose for the majority of the injected signals to have statistical significance right above the background LLR distribution; the dependence on Δt_{dur} causes signals with different durations to have similar LLRs. Finally, all the modulated signals considered in this study are injected at random times spread over a 1 yr period starting on 2019 April 1 in such a way that the merger time is always at least $\pm 30 \text{ s}$ away from the closest SAA episode.

4.2. Comparison with the Generic Targeted Search

In this subsection, we simulate signals with (m_1, m_2) equal to $(10 M_\odot, 1.4 M_\odot)$ and $(1.6 M_\odot, 1.4 M_\odot)$, and recover them with the chirp targeted search, where (m_1, m_2) is fixed at the same value. We also apply the generic targeted search to these injection sets. For each injection, the most significant trigger is selected and a FAP is derived according to the results summarized in Figure 3. The cumulative distribution functions (CDF) of the FAP for these injections sets are illustrated in Figure 4. The first remark one should draw from Figure 4 is that the chirp targeted search is more sensitive than the generic targeted search in the case of simplified light curves, as long as $N_{\text{bins}} > 1$. The case of the chirp targeted search with setting $N_{\text{bins}} = 1$ should be equivalent to the case of the generic targeted search. This relation is verified here, the small discrepancy between the two distributions being due to the different technical implementations.

The N_{bins} dependence of the chirp targeted search sensitivity warrants some discussion. First, Figure 3 suggests that the statistical significance of the noise triggers increases with N_{bins} , which has as an effect the degradation of the sensitivity with the augmentation of N_{bins} . Second, the performance of the pipeline is expected to depend on the relation between the width of intervals I_k and the orbital phase length of the chirping signals we want to detect. Following this idea, the sensitivity

should increase with N_{bins} , as long as the width of I_k is higher than the orbital phase length of the recovered signal. Indeed, for a given EM chirp radiation, if the signal is included in only one interval I_k , the smaller the I_k the higher the signal-to-noise ratio because the higher the percentage of I_k where the foreground photon rate is above the background photon rate. Therefore, one should expect a compromise between the two regimes described above: an increase of the sensitivity with the number of I_k intervals at low N_{bins} , then a saturation followed by a degradation of the performance at high N_{bins} . This behavior is verified in Figure 4. For N_{bins} equal to 1, 5, and 15, the width of the interval I_k is 360° , 72° and 24° . Given that $\theta_{\text{width}} = 10^\circ$, such a signal could a priori be included in one interval I_k , unless it is situated at the border of two adjacent I_k intervals. From Figure 4, one can note an increase of the sensitivity between the cases $N_{\text{bins}} = 1$ and 5. However, the performances of the pipeline seem to be quite similar in the cases $N_{\text{bins}} = 5$ and 15.

4.3. Detectability of Signals with Different Parameters

Figure 5 shows the statistical significance dependence of the chirp targeted search with the parameters of the simplified light-curve injections. The top left panel of Figure 5 proves that our choice for the Δt_{dur} dependence of the amplitude factor f_{amp} puts on an equal footing the signals with different durations. One should note that an injection with $f_{\text{amp}} = 1$ corresponds to an energy flux of $1 \text{ erg s}^{-1} \text{ cm}^{-2}$ in the 50–300 keV band. Thus, the same panel indicates that a signal spread over around 1 s and emitting during 10° orbital phase window, with a flux reaching $10 \text{ erg s}^{-1} \text{ cm}^{-2}$ in the Fermi-GBM band, might be recovered by our pipeline with an important statistical significance ($\text{LLR} > 100$). The top right panel shows that the pipeline sensitivity is independent of the signal position in the orbital phase space. From the bottom left panel, one can note better performance of the chirp targeted search when the signal has a spectrum either close to $\kappa = 0$ (hard) or close to $\kappa = 2$ (soft). This might be explained by the fact that although for the injections we have considered κ anywhere in the interval $[0, 2]$, the search recovery is realized with only the three spectrum templates corresponding to $\kappa \in \{0, 1, 2\}$. Therefore, an injection with κ not near to 0 or 2,

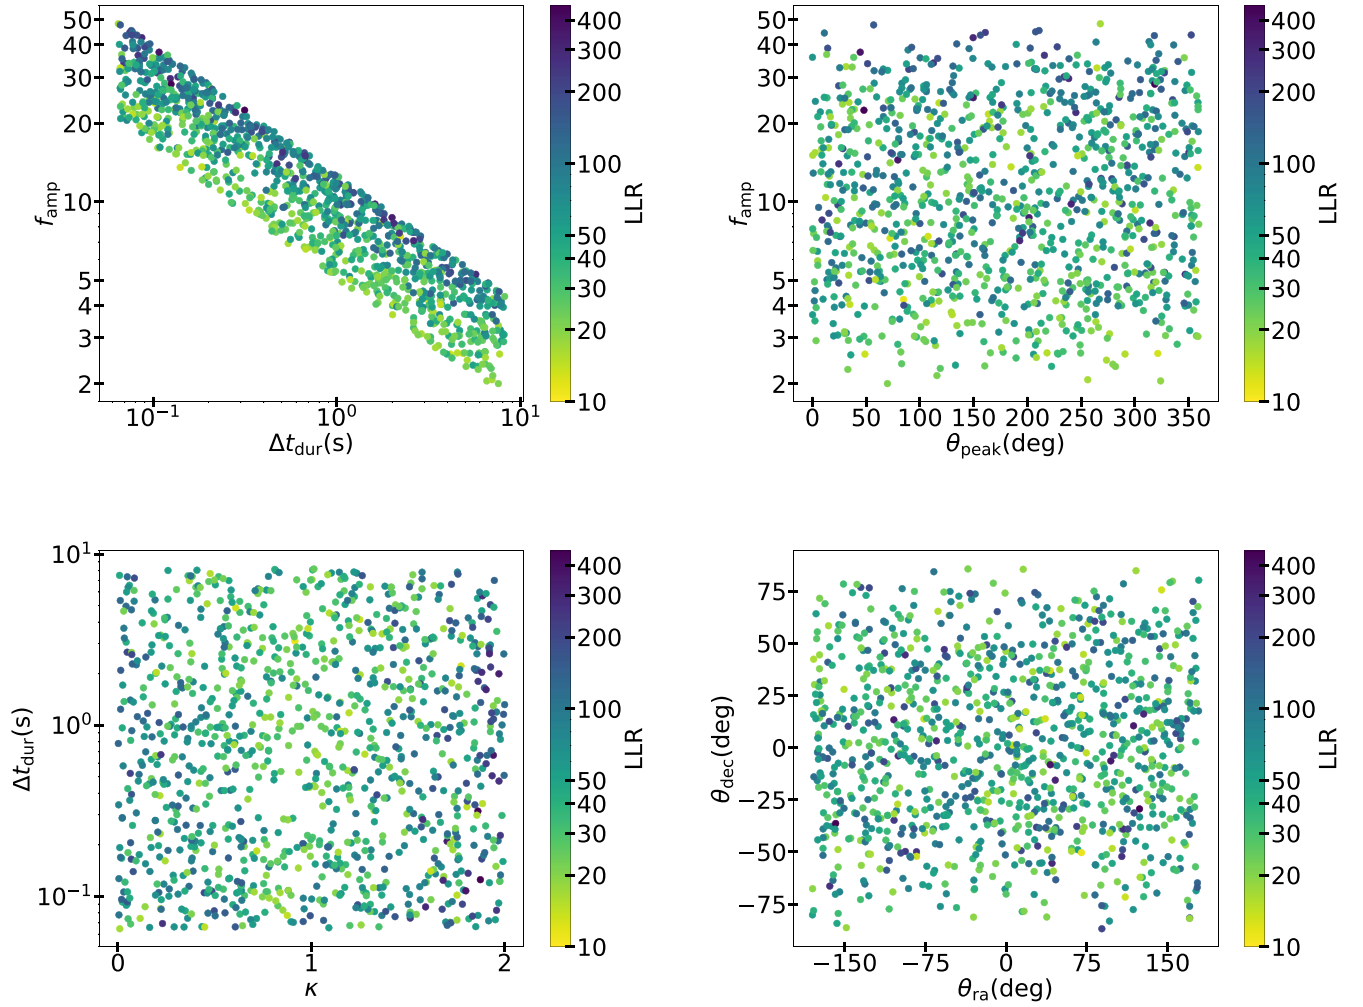


Figure 5. Variation of the LLR statistic produced by the search with the parameters of the simulated simplified light curves. All the panels correspond to the injection set with $(m_1, m_2) = (10 M_\odot, 1.4 M_\odot)$, while the merger times are uniformly distributed over the 1 yr period starting on 2019 April 1. The chirp targeted search is used with setting $N_{\text{bins}} = 10$.

represents an heterogeneous weighted addition of quite different spectral templates, and which is more difficult to be recovered with only one template in between *{soft, normal, hard}*. Finally, the bottom right panel of Figure 5 illustrates the independence of the sensitivity with the position in the sky of the injected signal.

4.4. Impact of Orbital Phase Uncertainty

The properties of a compact binary merger inferred from the GW data always carry some uncertainty. In particular, there are uncertainties on the chirp mass, the merger time and the sky location. The chirp mass uncertainty will impact our knowledge about the orbital phase evolution. The uncertainties on the geocentric merger time and the sky location together reflect on the uncertainty of the merger time at the position of the Fermi satellite, which in turn affects our knowledge of the positions of the I_k intervals. Concerning the chirp mass detection errors, the O1, O2, and O3 LIGO-Virgo observing runs showed that the uncertainties are below $0.1 M_\odot$ for BNS or NSBH like objects, and of the order of a few M_\odot for binary black hole mergers. Concerning the geocenter merger time detection, the uncertainty is of the order of 1 ms. However the sky locations of GW mergers are often poorly constrained, especially in the case of single-interferometer detections, which means that the

uncertainty of the merger time measured at Fermi is of the order of the photon flight time between Fermi and the center of Earth, i.e., ≈ 23 ms. Before applying our method to LIGO-Virgo events, we have to evaluate the impact of imprecise GW measurements on the sensitivity of our search.

The variation of the sensitivity with the uncertainty on the Fermi location merger time and on the chirp mass is illustrated in Figure 6. In order to investigate the impact of the simultaneous inaccuracies of the merger time and chirp mass measurements on the recovery efficiency of the search, we use a set of 1000 injections. Given the total mass $M_{\text{total}} = m_1 + m_2$ and the chirp mass \mathcal{M} , we consider injections with $(t_c, M_{\text{total}}, \mathcal{M}) = (t_c^i + \delta t_c^i, M_{\text{total}}^0, \mathcal{M}_0 + \delta \mathcal{M}^i)$, where M_{total}^0 and \mathcal{M}_0 correspond to a binary formed of a $10 M_\odot$ black hole and a $1.4 M_\odot$ neutron star, while δt_c^i and $\delta \mathcal{M}^i$ are uniformly sampled in $[-\sigma_{\Delta t}, \sigma_{\Delta t}]$ and $[-\sigma_{\Delta \mathcal{M}}, \sigma_{\Delta \mathcal{M}}]$. Then we recover the injected signal by means of the chirp targeted search, with the setting $(t_c^i, M_{\text{total}}^0, \mathcal{M}_0)$. This is equivalent to a GW trigger having a merger time uncertainty at Fermi of $\sigma_{\Delta t}$ and a chirp mass uncertainty equal to $\sigma_{\Delta \mathcal{M}}$. In the previous expressions the exponent i stands for the i th trigger. According to the results presented in Figure 6, when used with uncertainties $\sigma_{\Delta t} \leq 0.01$ s and $\sigma_{\Delta \mathcal{M}} \leq 0.01 M_\odot$, the CDF chirp targeted search has a relative error with respect to the case of perfect




Figure 6. CDF vs. LLR for simplified light-curve injections assigned with merger time and chirp mass uncertainties $\sigma_{\Delta t}$ and $\sigma_{\Delta \mathcal{M}}$. The injections are uniformly spread over the period between 2019 April 1 and 2020 April 1, and the setting $N_{\text{bins}} = 10$ is used.

measurements, i.e., $(\sigma_{\Delta t}, \sigma_{\Delta \mathcal{M}}) = (0 \text{ s}, 0 M_{\odot})$, lower than 10% over the range $\text{LLR} \in [0, 200]$. This means that if we have a measurement with such low uncertainties, and if the output of the chirp targeted search, when used with setting $N_{\text{bins}} \leq 10$, indicates a trigger with $\text{LLR} \geq 100$, then the EM candidate is promising because it is above the LLR background range corresponding to the -1σ lower limit.

5. Results on LIGO-Virgo Detections

We finally apply the search method described previously to a few LIGO-Virgo detections, which could plausibly be associated with a modulated γ -ray counterpart: the BNS mergers GW170817 (Abbott et al. 2017) and GW190425 (Abbott et al. 2020a); GW190814 (Abbott et al. 2020b), whose heavier object is a black hole, while the lighter object is either the heaviest neutron star or the lightest black hole observed to date; and the neutron star–black hole mergers GW200105 and GW200115 (Abbott et al. 2021a). Binary black hole mergers, although much more plentiful, are not considered as likely sources of detectable γ -ray counterparts in this work, and are not investigated. Moreover, the search for EM counterparts to such binary systems would impose a large trials factor due to the high rate of these mergers and the large uncertainties on both the chirp mass and the sky localizations (Abbott et al. 2021b). Hence, brighter EM counterparts may be required for a detection.

Observer-frame chirp masses and mission elapsed times (MET) since 2001.0 UTC (decimal), as well as the corresponding uncertainties, for the five GW events, are given in Table 1. Concerning the MET of GW170817 (respectively GW190814), 3 ms (respectively 20 ms) have been subtracted from the geocenter merger time posterior, in order to account for the angle between the direction to NGC 4993 (respectively the directions representing 90% credible regions of the GW190814 skymap) and the Fermi–Earth center baseline. For the other events, whose sky localizations are poorly constrained, we make the conservative simplifying assumption that the sky location is completely unknown. We then simply increase (decrease) the upper (lower) limits on the merger time

Table 1

The MET, Corresponding to the Merger Time at the Fermi Satellite, and the Observer-frame Chirp Mass, for the GW Detections Explored in this Work as Possible Sources of Modulated γ -Ray Precursors

Event	Merger MET (s)	$\mathcal{M} (M_{\odot})$
GW170817	$524666469.424^{+0.001}_{-0.002}$	1.198
GW190425	$577873090.009^{+0.056}_{-0.031}$	$1.487^{+0.001}_{-0.015}$
GW190814	$587509843.970^{+0.003}_{-0.003}$	$6.413^{+0.012}_{-0.015}$
GW200105	$599934271.048^{+0.058}_{-0.030}$	$3.620^{+0.009}_{-0.007}$
GW200115	$600754994.755^{+0.029}_{-0.055}$	$2.582^{+0.005}_{-0.005}$

Note. The values appearing here are the median (50th percentile), the upper limits (90th percentile), and the lower limits (10th percentile). A missing MET (respectively chirp mass) lower/upper limit means that the limit value is away from the median value by less than 1 ms (respectively $10^{-3} M_{\odot}$). The parameter estimates are given in Romero-Shaw et al. (2020) for GW170817, Abbott et al. (2019) for GW190425 and GW190814, and Abbott et al. (2021a) for GW200115 and GW200115.

Table 2

Results of the Chirp Search and Targeted Search, in Terms of Highest LLR and FAP with $\pm 3\sigma$ Uncertainties, for the GW Events Explored in this Work

Event	Chirp Search		Targeted Search	
	LLR	FAP $^{-3\sigma}_{+3\sigma}$	LLR	FAP $^{-3\sigma}_{+3\sigma}$
GW170817	52.7	$0.034^{+0.021}_{-0.029}$	7.0	$0.554^{+0.064}_{-0.065}$
GW190425	26.9	$0.290^{+0.061}_{-0.062}$	7.2	$0.500^{+0.069}_{-0.065}$
GW190814	21.3	$0.135^{+0.044}_{-0.049}$	8.5	$0.217^{+0.055}_{-0.057}$
GW200105	25.4	$0.127^{+0.042}_{-0.048}$	10.7	$0.053^{+0.027}_{-0.033}$
GW200115	24.3	$0.237^{+0.057}_{-0.059}$	8.9	$0.168^{+0.047}_{-0.052}$

Note. The setting $N_{\text{bins}} = 10$ has been used for the chirp search. The FAP values indicate that the LLRs are compatible with background fluctuations.

by 23 ms, which is approximately the light travel time between Earth’s center and the Fermi satellite.

To increase our chance of picking merger time and chirp mass values close enough to the true ones, we make use of a grid in the following way: if the MET (respectively the chirp mass) upper and lower limits are MET_{min} and MET_{max} (respectively \mathcal{M}_{min} and \mathcal{M}_{max}), we consider 2D grid points $(\text{MET}_i, \mathcal{M}_j)$, with MET_i (respectively \mathcal{M}_j) ranging from MET_{min} to MET_{max} (respectively from \mathcal{M}_{min} to \mathcal{M}_{max}), such that $\text{MET}_{i+1} - \text{MET}_i = 0.02 \text{ s}$ and $\mathcal{M}_{j+1} - \mathcal{M}_j = 0.02 M_{\odot}$. In this way, if the true values MET^0 and \mathcal{M}^0 are indeed in between the upper and the lower limits, there is at least a grid point $(\text{MET}_{i_0}, \mathcal{M}_{j_0})$ such that $|\text{MET}^0 - \text{MET}_{i_0}| \leq 0.01 \text{ s}$ and $|\mathcal{M}^0 - \mathcal{M}_{j_0}| \leq 0.01 M_{\odot}$. This working method is motivated by the results obtained in the previous section and summarized in Figure 6. We create the grids corresponding to the five GW events and then we run both the generic and chirp targeted searches with the setting defined by these grid points.

The maximum LLR obtained for each run is converted into a FAP in the following way. For a given GW event with merger time t_c , the chirp targeted search (and the generic targeted search) background distributions are obtained by running the search on 1000 random off-source times covering the interval $[t_c - \frac{1}{2} \text{ month}, t_c + \frac{1}{2} \text{ month}] \setminus [t_c - 60 \text{ s}, t_c + 30 \text{ s}]$. In this way, the estimated noise distribution samples the high-energy EM activity in the few weeks around the time of interest, but we avoid the 30 s of on-source Fermi–GBM data preceding and

following the merger. This prevents any possible candidate counterpart from contaminating the background. Of course, when constructing the off-source background of the chirp targeted search, mass parameters consistent with the particular on-source event of interest are used.

The maximum LLR and FAP associated with each GW event are reported in Table 2. None of the FAP values have upper limit below 0.01, indicating that the LLR values associated with all events are compatible with the off-source background fluctuations. We conclude that there are neither statistically significant excesses of photons, nor significant modulated signals prior to the mergers we considered.

6. Conclusion

In this work we present a method, the chirp targeted search, to detect modulated γ -ray precursors to compact binary mergers in Fermi-GBM data. The existence of such signals has not been confirmed so far. If they exist, there are several physical mechanisms that might be responsible for their emission. This fact makes the light-curve amplitude dependence difficult to predict. Despite these difficulties, the presented method is very general. It aims to look for an excess of photons in the orbital phase space, while the GW frequency evolution is defined by the first order term in the post-Newtonian expansion. The sensitivity of the method has been tested on simplified light curves, for which the EM emission takes place only during the same orbital phase window. The performance of the chirp targeted search has been compared to that of an existing, more generic targeted search, which aims to detect an excess of photons in the time space. It has been displayed that the chirp targeted search has higher sensitivity than the generic targeted search, when the signal is modulated by the GW frequency of the binary. Finally, both pipelines have been used to search for EM precursors associated with confident GW events having a non-negligible probability to contain a neutron star, namely GW170817, GW190425, GW190814, GW200105, and GW200115. We found no significant candidate precursor signals associated with any of those events. However, given the potential of new physics provided by the presence of precursor signals, the proposed method, here demonstrated, will be applied to future BNS and NSBH observations by LIGO-Virgo-KAGRA, especially since an increased number of events is expected in the coming years.

It is possible that this method could also be important for stellar mass–binary black hole mergers. As an example, these mergers could occur in gas-rich environments, such as AGN disks, and could produce EM counterparts with orbital modulations, similarly to what has been proposed for more massive binaries in the LISA band (e.g., Tang et al. 2018). Potential EM counterparts have been reported for both GW150914 (Connaughton et al. 2016) and GW190521 (Graham et al. 2020). As stated above, however, in order to apply this method to the large number of binary black hole events detected by current ground-based interferometers, it will be necessary to address the statistical and computational challenges.

Although a sensitivity gain has been achieved by means of the chirp targeted search, improvements can be envisaged in the future. The actual method is trying to recover signals with spectra described by the related Band function. One can argue that this is not the most optimal choice. The Band functions are very appropriate to recover γ -ray burst prompt emission, at the origin of which most likely the synchrotron radiation and the

inverse Compton scattering are at play in producing photons. In the case of γ -ray precursors to compact binary mergers, one can imagine other physical emission mechanisms, like thermal emission. A higher sensitivity might then be obtained by including additional spectral templates in the search. Additional filtering strategies may also turn out to be effective. A cleaning of the search output is synonymous with decreasing the FAP, and as a consequence, the sensitivity of the search gets higher.

T.D.C. was supported by an appointment to the NASA Postdoctoral Program at the Goddard Space Flight Center, administered by Universities Space Research Association under contract with NASA, during part of this work. M. W. Coughlin acknowledges support from the National Science Foundation with grant Nos. PHY-2010970 and OAC-2117997.

Appendix EM Realistic Light Curve

We present here the analytical formulae used for the derivation of the light curve highlighted in orange, in Figure 1. The system is assumed to be an NSBH binary and the flaring is due to the uniform and isotropic emission of the neutron star surface. The amplitude of the luminosity is modulated by relativistic beaming, gravitational lensing, and orbital separation shrinking.

Concerning the relativistic Doppler beaming, the same magnification factor as in Dubus et al. (2010) is used. This factor depends on the neutron star velocity $v_{\text{NS}}^{\text{DUB}}$, the angle ζ_1^{DUB} between $v_{\text{NS}}^{\text{DUB}}$ and the line of sight, and the thermal light spectral index α^{DUB} . The expression of the amplification factor is



$$\left[\frac{1 - \frac{v_{\text{NS}}}{c} \cos(\zeta_1^{\text{DUB}})}{\sqrt{1 - \left(\frac{v_{\text{NS}}^{\text{DUB}}}{c}\right)^2}} \right]^{\alpha^{\text{DUB}} - 3}. \quad (\text{A1})$$

Regarding the gravitational lensing factor, it is assumed to be the same as that in Narayan & Bartelmann (1996), depending on the binary separation D^{NAR} , the angle ζ_2^{NAR} between the orbital separation and the line of sight and the Einstein radius R_E^{NAR} . Moreover, for the cases where the neutron star is close or inside the Einstein ring, following Liebes (1964), we pretend it is exactly behind the BH and approximate its image as a ring with angular width equal to the neutron star angular diameter. Therefore, the gravitational lensing factor can be written as

$$\begin{cases} \frac{(u^{\text{NAR}})^2 + 2}{u^{\text{NAR}} \sqrt{(u^{\text{NAR}})^2 + 4}} & \text{for } \|u^{\text{NAR}}\| \geq 1 \\ \min \left(\frac{(u^{\text{NAR}})^2 + 2}{u^{\text{NAR}} \sqrt{(u^{\text{NAR}})^2 + 4}}, 2 \sqrt{\frac{4Gm_{\text{BH}}^{\text{NAR}} D^{\text{NAR}}}{c^2 r_{\text{NS}}^{\text{NAR}}}} \right) & \text{for } \|u^{\text{NAR}}\| \leq 1 \end{cases}, \quad (\text{A2})$$

where $u^{\text{NAR}} = D^{\text{NAR}} \sin \zeta_2^{\text{NAR}} (R_E^{\text{NAR}})^{-1}$, while $m_{\text{BH}}^{\text{NAR}}$ and $r_{\text{NS}}^{\text{NAR}}$ are the black hole mass and the neutron star radius.

ORCID iDs

Cosmin Stachie  <https://orcid.org/0000-0002-5723-2342>
Tito Dal Canton  <https://orcid.org/0000-0001-5078-9044>
Nelson Christensen  <https://orcid.org/0000-0002-6870-4202>

Marie-Anne Bizouard  <https://orcid.org/0000-0002-4618-1674>
 Eric Burns  <https://orcid.org/0000-0002-2942-3379>
 Michael Coughlin  <https://orcid.org/0000-0002-8262-2924>

References

- Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, *CQGra*, **32**, 074001
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **119**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *PhRvX*, **9**, 031040
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020a, *ApJL*, **892**, L3
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020b, *ApJL*, **896**, L44
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, *ApJL*, **915**, L5
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, *PhRvX*, **11**, 021053
- Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *CQGra*, **32**, 024001
- Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv:1702.00786
- Arnett, W. D., Bahcall, J. N., Kirshner, R. P., & Woosley, S. E. 1989, *ARA&A*, **27**, 629
- Atwood, W. 1994, *NIMPA*, **342**, 302
- Band, D., Matteson, J., Ford, L., et al. 1993, *ApJ*, **413**, 281
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *SSRv*, **120**, 143
- Beloborodov, A. M., Daigne, F., Mochkovitch, R., & Uhm, Z. L. 2011, *MNRAS*, **410**, 2422
- Berger, E. 2014, *ARA&A*, **52**, 43
- Blackburn, L., Briggs, M. S., Camp, J., et al. 2013, arXiv:1303.2174
- Bissaldi, E., von Kienlin, A., Lichti, G., et al. 2009, *ExA*, **24**, 47
- Blackburn, L., Briggs, M. S., Camp, J., et al. 2015, *ApJS*, **217**, 8
- Blanchet, L. 2014, *LRR*, **17**, 2
- Burgess, J. M., Yu, H.-F., Greiner, J., & Mortlock, D. J. 2016, arXiv:1610.07385
- Connaughton, V., Briggs, M. S., Goldstein, A., et al. 2015, *ApJS*, **216**, 32
- Connaughton, V., Burns, E., Goldstein, A., et al. 2016, *ApJL*, **826**, L6
- Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, *Sci*, **358**, 1556
- Darbha, S., & Kasen, D. 2020, *ApJ*, **897**, 150
- D'Avanzo, P. 2015, *JHEAp*, **7**, 73
- D'Avanzo, P., Campana, S., Salafia, O. S., et al. 2018, *A&A*, **613**, L1
- D'Orazio, D. J., & Levin, J. 2013, *PhRvD*, **88**, 064059
- D'Orazio, D. J., Levin, J., Murray, N. W., & Price, L. 2016, *PhRvD*, **94**, 023001
- Dubus, G., Cerutti, B., & Henri, G. 2010, *A&A*, **516**, A18
- Foucart, F., Hinderer, T., & Nissanke, S. 2018, *PhRvD*, **98**, 081501
- Fox, D. B., Frail, D. A., Price, P. A., et al. 2005, *Natur*, **437**, 845
- Goldstein, A., Burns, E., Hamburg, R., et al. 2016, arXiv:1612.02395
- Goldstein, A., Hamburg, R., Wood, J., et al. 2019, arXiv:1903.12597
- Goldstein, A., Veres, P., Burns, E., et al. 2017, *ApJL*, **848**, L14
- Graham, M. J., Ford, K. E. S., McKernan, B., et al. 2020, *PhRvL*, **124**, 251102
- Haiman, Z. 2017, *PhRvD*, **96**, 023004
- IceCube Collaboration, Aartsen, M. G., Ackermann, M., et al. 2018, *Sci*, **361**, eaat1378
- Kasen, D., & Barnes, J. 2019, *ApJ*, **876**, 128
- Kasen, D., Metzger, B., Barnes, J., Quataert, E., & Ramirez-Ruiz, E. 2017, *Natur*, **551**, 80
- Kippen, R., Hoover, A., Wallace, M., et al. 2007, AIP Conf. Proc. 921, The First Glast Symposium, ed. S. Ritz, P. Michelson, & C. A. Meegan, (Melville, NY: AIP)
- Kocevski, D., Burns, E., Goldstein, A., et al. 2018, *ApJ*, **862**, 152
- Kochanek, C. S., & Piran, T. 1993, *ApJL*, **417**, L17
- Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, *ApJL*, **413**, L101
- Lai, D. 1994, *MNRAS*, **270**, 611
- Lamb, G. P., Lyman, J. D., Levan, A. J., et al. 2019, *ApJL*, **870**, L15
- Lattimer, J. M., & Schramm, D. N. 1974, *ApJL*, **192**, L145
- Li, L.-X., & Paczynski, B. 1998, *ApJL*, **507**, L59
- Li, X. J., Zhang, Z. B., Zhang, X. L., & Zhen, H. Y. 2021, *ApJS*, **252**, 16
- Liebess, S. 1964, *PhRv*, **133**, B835
- McWilliams, S. T., & Levin, J. 2011, *ApJ*, **742**, 90
- Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, **702**, 791
- Meszáros, P., & Rees, M. J. 1997, *ApJ*, **476**, 232
- Metzger, B. D., Martínez-Pinedo, G., Darbha, S., et al. 2010, *MNRAS*, **406**, 2650
- Most, E. R., & Philippov, A. A. 2020, *ApJL*, **893**, L6
- Narayan, R., & Bartelmann, M. 1996, arXiv:astro-ph/9606001
- Neill, D., Tsang, D., van Eerten, H., Ryan, G., & Newton, W. G. 2021, arXiv:2111.03686
- Paczynski, B., & Rhoads, J. E. 1993, *ApJL*, **418**, L5
- Palenzuela, C., Bona, C., Lehner, L., & Reula, O. 2011, *CQGra*, **28**, 134007
- Palenzuela, C., Lehner, L., Ponce, M., et al. 2013, *PhRvL*, **111**, 061105
- Pendleton, G. N., Briggs, M. S., Kippen, R. M., et al. 1999, *ApJ*, **512**, 362
- Phinney, E. S. 1991, *ApJL*, **380**, L17
- Reisenegger, A., & Goldreich, P. 1994, *ApJ*, **426**, 688
- Romero-Shaw, I., Talbot, C., Biscoveanu, S., et al. 2020, *MNRAS*, **499**, 3295
- Savchenko, V., Ferrigno, C., Kuulkers, E., et al. 2017, *ApJL*, **848**, L15
- Schnittman, J. D., Dal Canton, T., Camp, J., Tsang, D., & Kelly, B. J. 2018, *ApJ*, **853**, 123
- Shibata, M. 1994, *PTHPh*, **91**, 871
- Smartt, S. J., Chen, T. W., Jerkstrand, A., et al. 2017, *Natur*, **551**, 75
- Sridhar, N., Zrake, J., Metzger, B. D., Sironi, L., & Giannios, D. 2021, *MNRAS*, **501**, 3184
- Suvorov, A. G., & Kokkotas, K. D. 2020, *PhRvD*, **101**, 083002
- Tang, Y., Haiman, Z., & Macfadyen, A. 2018, *MNRAS*, **476**, 2249
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, **253**, 908
- Troja, E., Rosswog, S., & Gehrels, N. 2010, *ApJ*, **723**, 1711
- Tsang, D. 2013, *ApJ*, **777**, 103
- Tsang, D., Read, J. S., Hinderer, T., Piro, A. L., & Bondarescu, R. 2012, *PhRvL*, **108**, 011102
- Vietri, M. 1997, *ApJL*, **478**, L9
- Wang, J.-S., Peng, Z.-K., Zou, J.-H., Zhang, B.-B., & Zhang, B. 2020, *ApJL*, **902**, L42
- Zhu, S. 2015, PhD thesis, Univ. of Maryland, <http://hdl.handle.net/1903/17258>