Early Warnings of Binary Neutron Star Coalescence Using the SPIIR Search

Manoj Kovalam[1,2] , Md Anwarul Kaium Patwary[1,3] , Anala K. Sreekumar[1,2], Linqing Wen[1,2] , Fiona H. Panther[1,2] , and Qi Chu[1,2]

[1] Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Australia; manoj.kovalam@research.uwa.edu.au, linqing.wen@uwa.edu.au
[2] Department of Physics, University of Western Australia, Crawley WA 6009, Australia
[3] Department of Computer Science and Software Engineering, University of Western Australia, Crawley WA 6009, Australia

Abstract

Gravitational waves from binary neutron star mergers can be used as alerts to enable prompt follow-up observations. In particular, capturing prompt electromagnetic and astroparticle emissions from the moment of a binary merger presents unique constraints on the timescale and sky localization for online gravitational-wave detection. Here we present the expected performance of the SPIIR online detection pipeline that is designed for this purpose in the upcoming international LIGO–Virgo's 4th Science Run (O4). Using simulated Gaussian data for the two LIGO observatories with expected O4 sensitivity, we demonstrate that there is a nonnegligible opportunity to deliver premerger warnings at least 10 s before the final plunge. These alerts are expected to be issued at a nominal rate of one binary neutron star coalescence per year and localized within a median searched area of 300 deg^2. We envision such detection to be extremely useful for follow-up observatories with a large field of view such as the Murchison Widefield Array radio facility in Western Australia.

1. Introduction

The era of gravitational-wave (GW) astronomy began with its first-ever detection from a compact binary coalescence (CBC) in 2015 (Abbott et al. 2016), during the first advanced LIGO observing run. This was achieved using the two LIGO observatories at Hanford (H1) and Livingston (L1) (Aasi et al. 2015). In 2017, Virgo (V1; Acernese et al. 2015) joined the duo in the second observing run. The iconic detection of a binary neutron star (BNS) merger (Abbott et al. 2017a) and associated electromagnetic (EM) emission signaled the beginning of a new era of GW multimessenger astronomy (MMA). This was a joint detection of GWs and a short gamma-ray burst (sGRB) detected ∼2 s after the binary merger by the Fermi-GBM and INTEGRAL space telescopes (Abbott et al. 2017b). This event was followed by numerous other EM observations spanning the entire EM spectrum. This detection has had a tremendous impact in astronomy with several studies emerging from it, including estimation of the Hubble constant (Abbott et al. 2017c), constraints on the neutron star equation of state (Abbott et al. 2019), and also connecting sGRBs and kilonovae to BNS mergers (Abbott et al. 2017b).

GW detections were reported publicly in real-time for the first time during the third LIGO–Virgo observing run O3 (Acernese et al. 2019; Buikema et al. 2020), leading to the search for EM counterparts of GW observations. However, the typical delay between a GW detection and the associated GCN alert was on the order of at least minutes (Magee et al. 2021), and hence the discovery of prompt EM emissions, which are expected during a binary merger (Rezzolla et al. 2011), continued to rely on serendipitous discovery. For example, in the case of GW170817, the sGRB was observed only 2 s after the GW detection. Thus, an advance warning of such events is crucial to alert conventional EM telescopes for prompt follow-up observations.

Five GW detection pipelines have processed GW data in real time for past LIGO–Virgo–Kagra collaboration (LVK) science runs. Four modeled search pipelines—SPIIR (Summed Parallel Infinite Impulse Response) (Hooper et al. 2012; Luan et al. 2012; Chu et al. 2022), GstLAL (Messick et al. 2017; Sachdev et al. 2019), PyCBC (Nitz et al. 2018; Dal Canton et al. 2021), and MBTA (Adams et al. 2016; Aubin et al. 2021)—use known CBC waveforms to identify a signal within the detector data, while burst signals are recovered by the cWB (Klimenko et al. 2016) pipeline via a coherent analysis. SPIIR uses a time-domain filtering method equivalent to matched filtering to detect GWs. The SPIIR pipeline uses GPU acceleration for parallel processing to reduce latency and improve computational efficiency (Liu et al. 2012; Guo et al. 2018). Out of the 38 detections reported by SPIIR through public alerts in O3 (Chu et al. 2022), 90% (35) of them were confirmed by offline searches and later added to the GW catalogs (Abbott et al. 2021a, 2021b, 2021c). The overall latency for the SPIIR pipeline, defined as the time delay between a GW event merger and its detection, is ∼9 s,[4] which corresponds to the H1L1 two-detector analysis. This paper describes a new feature for the SPIIR search, targeting "early warning" to detect GWs and generate alerts before the merger or at negative latencies. Note that a binary coalescence involving at least one neutron star tends to spend a few minutes of its time during the inspiral phase within the LIGO–Virgo sensitive frequency band (15–2000 Hz).

Early-warning (EW) alert systems associated with other matched filtering pipelines have been tested for their localization

---

[4] Not encompassing latencies associated with data transfer from detectors and sky-map/alert generation.

accuracy and detection efficiency; see footnote 4 (Nitz et al. 2020; Sachdev et al. 2020). In 2020, a mock data challenge (Magee et al. 2021) was conducted within the LVK to test the capability of the low-latency infrastructure to send EW alerts. SPIIR participated in this test by sending out one of the five EW mock GCNs and successfully demonstrated the feasibility for it to send premerger alerts.[5] This study also estimated the expected rate of BNS mergers and their localization areas using Bayestar rapid localization (Singer & Price 2016) by simulating a four-detector network—H1, L1, V1, and Kagra (K1; Akutsu et al. 2021)—in O4.

In this work, we analyze the performance of the SPIIR EW pipeline in a simulated O4 environment. We, in particular, study the performance of the two-LIGO-detector network (H1L1), which has the minimum expected overall latency caused by data transfer among all available detectors. The main motivation behind the choice of this network is to make detection as early as possible, thereby assisting EM telescopes in observing possible short transients ($\sim$1 s) right at the merger (Rezzolla et al. 2011). Because most follow-up telescopes need time to orient themselves to the source direction, saving an additional 4 s of delay from Virgo is extremely useful. We wish to take advantage of this faster network even though it results in a possibly worse localization, which can still be beneficial to observatories with large fields of view (FOVs). It should be noted that Virgo will be added to the full-bandwidth search, which will produce an enhanced localization within tens of seconds. We then demonstrate the accuracy of parameters including signal-to-noise ratio (S/N) and the chirp mass, which are estimated internally and used to classify the source type via p_astro (Kapadia et al. 2020) and infer properties via hasNS (Foucart et al. 2018). We also ensure the reliability of localization areas at various latencies, which are published via SPIIR EW alerts.

This paper is organized as follows. In Section 2, we show details about the SPIIR EW method and provide information about the simulated O4 data and GW signal injections. Section 3 discusses the results of this simulation run. Finally, in Section 4, we discuss the results in the context of EM follow-up by ground- and space-based facilities.

2. Method

EW alerts are expected to be issued publicly for the first time by LVK in O4. Through this study, we demonstrate the expected performance of the SPIIR online pipeline to deliver EW alerts. We only use the LIGO detectors in this work because this two-detector network has a very-high-coincident duty-cycle time (62% of total) reported in O3 (Davis et al. 2021) and high sensitivity with nearly aligned antenna beam patterns and, most importantly, saves an additional 4 s data-transfer latency caused by Virgo (Magee et al. 2021), making it the most viable combination to promptly report a premerger alert in O4, at the cost of poorer localization for some of the bright events at early detections compared to other combinations.

2.1. SPIIR Early-warning Pipeline

Matched filtering is the optimal method used by CBC pipelines to detect the presence of a GW signal amidst noise. This method involves cross-correlating known waveforms, also known as templates, with detector data to output S/N (Finn 1992; Cutler & Flanagan 1994). SPIIR filtering is a time-domain equivalent to matched filtering (Luan et al. 2012; Hooper et al. 2012), which uses first-order IIR filters to approximate GW templates to a high accuracy, which are then used in a time-domain convolution with the detector data, constructing the SPIIR S/N. The best-matched template would maximize the S/N in the presence of a signal, assuming that the noise is a stationary Gaussian.

The SPIIR pipeline has the capability to process GW data with extremely low latency. Because it is a time-domain convolution, it takes SPIIR theoretically close to 0 s to produce S/Ns (Luan et al. 2012; Hooper et al. 2012). However, in our actual implementation, the design costs SPIIR $\sim$1 s to compute the S/N and $\sim$4 s to identify candidates (Chu et al. 2022). A coherent network S/N across all active detectors is calculated for any single-detector event above a particular S/N threshold (set to $\rho > 4$ for real-time searches). Trigger candidates are then ranked based on two quantities: (a) the reduced chi-square (ξ^2), indicating a goodness of fit of the S/N time series against expectation, and (b) the coherent network S/N (ρ_c). A false-alarm rate (FAR) is calculated based on the probability of an event associated with the background.

In this work, we use a template bank consisting of 100,000 binary mass and projected spin pairs. These waveform templates have component masses ranging from $1.1\,M_\odot < m_1$, $m_2 < 3.0\,M_\odot$, focusing on binaries with two neutron stars with promising opportunities to observe EM emissions (Rezzolla et al. 2011). Waveforms for these masses (based on the SpinTaylorT4 (Buonanno et al. 2009) time-domain approximation) are truncated at specific time intervals before merger and used to construct the SPIIR EW filters. The SPIIR filters used in the work are set to have a high ($>$97%) overlap, which is the inner product of the original waveform and the approximate waveform.

We conduct seven parallel EW searches on the simulated data using the SPIIR pipeline, with the searches having their templates truncated between 10 s and 70 s, at 10 s intervals, respectively. For simplicity, we label these configurations with their premerger truncation time. For example, the EW search with its templates truncated 10 s premerger is simply addressed as the -10 s run. It should be noted that this number only represents the premerger template truncation time and does not correspond to the overall detection latency, which will be addressed in Section 3.3. The EW searches are processed at a sampling rate of 256 Hz in conjunction with a full-bandwidth simulation (nontruncated templates), which is processed at 2048 Hz for comparison. This is because EW detections are recovered at a frequency of $<$128 Hz at 10 s before the merger and earlier, in the inspiral phase.[6] We label the full-bandwidth simulation as the 0 s run to stay consistent across. For each of these searches, the FARs associated with candidate triggers are computed independently from one another.

2.2. Simulated Data

We inject signals into stationary Gaussian noise, simulating the strain data produced by the LIGO detectors in O4. The estimated power spectral density (PSD) of LIGO in O4 is given in Abbott et al. (2018), with an expected BNS range of 190 Mpc. We used

---

6   The downsampling also reduces the computational complexity of EW simulations, thereby decreasing the amount of resources used.
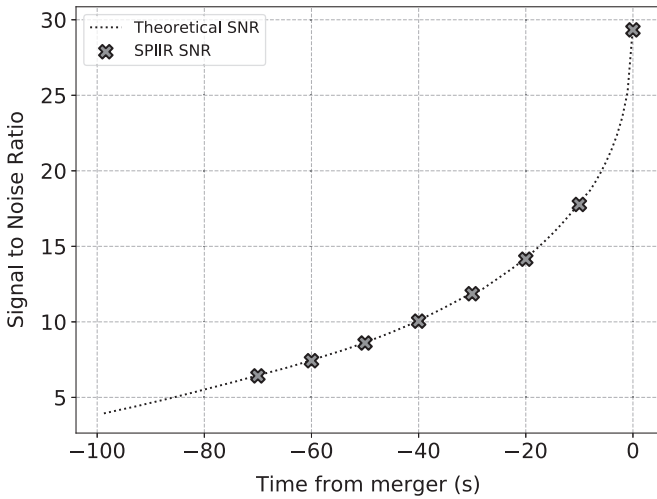
Figure 1. Evolution of S/Ns at different premerger latencies in simulated LIGO O4 data for a $1.4\,M_\odot$–$1.4\,M_\odot$ BNS system at a distance of 100 Mpc. The dotted line represents the theoretical S/Ns calculated using matched filtering while the symbols show S/Ns from the SPIIR method. A strong event with a full-bandwidth S/N of 30 can be detected at subminute premerger latencies.

the `gstlal_fake_frames` package (Messick et al. 2017) to generate these strain data using the given PSD. Both LIGO-Hanford and LIGO-Livingston are expected to reach a similar sensitivity in O4 and hence the output strain is expected to be similar in both. The whole data segment spans three weeks and is generated in the `gwf`[7] format.

The injected signals used in this analysis are generated using the `lalsuite` package (LIGO Scientific Collaboration 2018). The component masses of the injections are sampled between $1.0\,M_\odot < m_1$, $m_2 < 2.3\,M_\odot$, with a uniform distribution. The spin is restricted to be below 0.4 for both components and sampled with an isotropic distribution. The injected signals are distributed uniformly within a comoving volume of redshift $z = 0.2$. The final population set, using the constraints mentioned above, has $\mathcal{O}(10^4)$ number of injections.

Expected S/Ns for GW signals can be theoretically calculated using their mass, spin, and distance parameters. By truncating the waveforms, we can also estimate the S/Ns at different latencies before the merger. Figure 1 shows the evolution of S/Ns at different premerger latencies for a fiducial BNS reference $1.4\,M_\odot + 1.4\,M_\odot$ source at 100 Mpc in O4 sensitivity. We also show that the SPIIR S/Ns at the truncated intervals indicate consistency with the expected S/Ns. Similarly, expected S/Ns at different premerger latencies are theoretically calculated for all the injected signals in our simulation.

3. Results

Recovered signals from the searches are considered to be a detection only if the FAR reported by the associated search is less than one per month (or 3.85×10^{-7} Hz), which is also the threshold used for reporting EW candidates in the MDC (Magee et al. 2021). We find that 61% of all injected signals were detected by the 0 s run. Using 320^{+490}_{-240} Gpc^{-3}yr^{-1} as the local BNS merger rate, estimated in Abbott et al. (2021a), we expect a median detection rate of ~18 BNSs per year in O4 using the full-bandwidth configuration. It should be noted that our choice of the FAR threshold also includes subthreshold events when compared

---
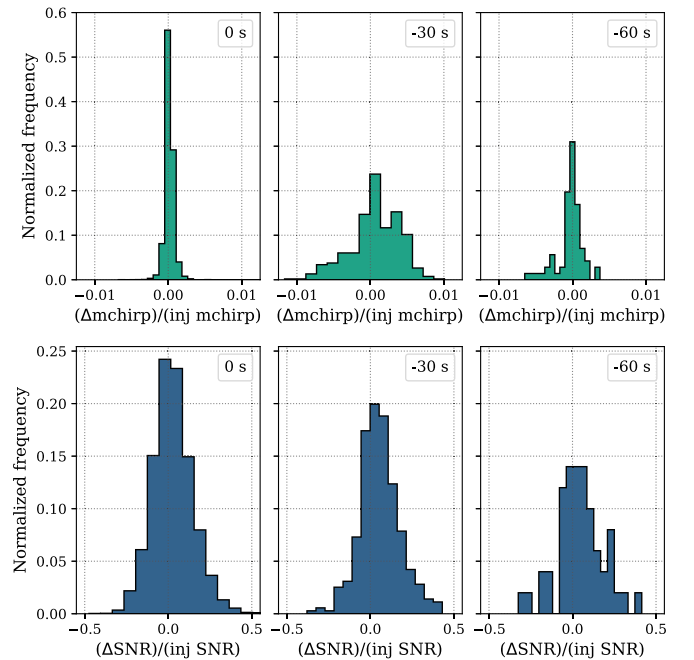
Figure 2. Accuracy of chirp mass and S/N recovery for the 0 s, −30 s, and −60 s configurations. Top panel: distributions for fractional differences in the recovered and injected chirp masses. The error margin is within 2% for all the runs. Bottom panel: distributions for fractional differences in recovered and expected S/Ns. About half of all events fall within the 10% error margin (in the bottom panel), of which 3% can be attributed to the SPIIR GW waveform approximation.

to the threshold of one every two months used to select the open public alerts during O3. In comparison, Magee et al. (2021) use a threshold of S/N > 12 and anticipate the detection rate to be ~9 BNSs per year for an H1L1V1K1 four-detector network. If we apply the same criterion, we also see a similar rate but this is for the H1L1 two-detector network. This is probably due to the fact that Magee et al. (2021) uses a Gaussian distribution for their mass model, while we use a uniform distribution, which leads to a larger detection range for our search and hence a comparatively larger rate.

3.1. Accuracy

Online estimations of parameters like chirp mass and S/Ns are calculated internally by the pipeline and not published directly to the public. However, these parameters are essential to compute the published values of FARs, p_astro (Kapadia et al. 2020) and hasNS (Foucart et al. 2018), which are used to infer source properties and classify detections. Thus, it is very important for pipelines to determine the accuracy of these parameters to ensure the reliability of the published alerts. Figure 2 compares the S/Ns and chirp masses recovered from the searches to the injection set. We demonstrate the performance of the SPIIR EW pipelines by first comparing the chirp mass of the recovered signals to the expected values (Figure 2, upper panel). We find the fractional difference between the two to be less than 2% for all runs, indicating an accurate recovery. We see that the chirp mass recovered from the 0 s run is within an error margin of 0.5%, with subsequently wider margins for other EW configurations. Next, we find that, on average, the recovered S/Ns are within an error margin of ~10% with the expected S/Ns (Figure 2, lower panel). It should be noted that a 3% deviation is expected in S/Ns

because of the SPIIR waveform approximation. The rest of the discrepancy could be attributed to factors like the discreteness in template banks, leading to a mismatch with injections, noise influences in low-S/N events, and relatively minor PSD mismatches while calculating the S/Ns using moving PSD estimations. The accuracy of chirp mass and S/N estimates for the full-bandwidth run has been well studied in Chu et al. (2022) with similar results reported. Figure 2 shows a histogram of the fractional differences in the injected and recovered parameters for the 0 s, 30 s, and 60 s runs.[8]

3.2. Localization

The accurate and rapid localization of GW signals is crucial to enable prompt follow-up observations. Providing source direction to EM observatories via premerger alerts would help enable a prompt follow-up observation. For this study, we use LVK's rapid localization software, Bayestar (Singer & Price 2016), to construct sky maps. As we go down in latency, EW simulations trigger at lower frequencies, so only partial information of the detected signal is recovered. This lowers the estimated S/Ns and also the accuracy with which the temporal and phase information is recovered, thus resulting in larger localization areas (Fairhurst 2018). Figure 3 shows the cumulative histogram of the searched areas (defined as the smallest area needed to be searched to encompass an event) normalized w.r.t. the total number of detection in the full-bandwidth run (18).
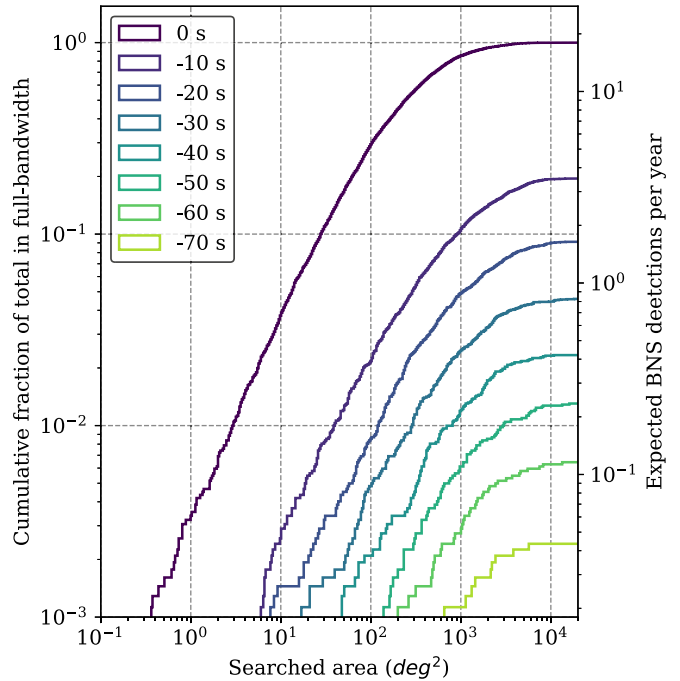
Based on our Gaussian noise simulation for O4, we expect to deliver premerger BNS alerts at a rate of 1 (2) per year in O4 with a searched area below 300 (1000) deg^2 using the LIGO network alone. We also expect to detect at least one premerger detection in two years localized within a searched area of 100 deg^2 in O4. However, the chance to have a subminute premerger detection is seen to be less than 0.2%. We also compare the searched areas of the candidates with their 90% credible region areas in Figure 4. Table 1 shows the median values for the searched and 90% credible areas and also the expected detection rate at all EW latencies in O4. We find that the median values of the searched areas are several times smaller than the 90% credible areas, and the difference becomes more prominent as we go down in latency for the EW runs.
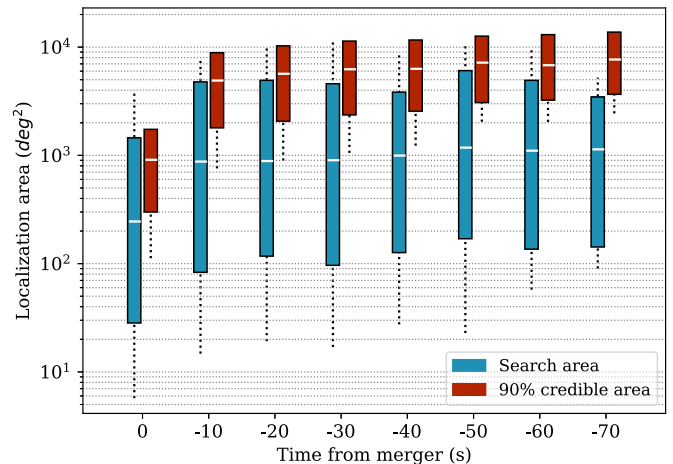
3.3. Latency

The overall pipeline latency for each configuration is measured on an online data streaming platform, and the median values are recorded in Table 1. These values correspond to the internal latency of the pipeline, which includes delays in the SPIIR filtering algorithm and calculating the FARs. We find that the full-bandwidth configuration has a median intrinsic latency of 9–10 s,[9] consistent with results from Chu et al. (2022). Additional latencies, including data transfer, localization, and alert generation, depend on the online infrastructure and are believed to be significantly lower in O4 as compared to O3 (Magee et al. 2021). The EW runs have a negative median latency indicating premerger detection.

4. Discussion

In this work, we investigated the prospects of using the SPIIR EW search, on the two-detector LIGO network, to detect

---

[8] The −60 s simulation has a lower number of statistics compared to others.
[9] With the potential to go below 5 s.

Figure 3. Searched areas vs. the number of detected events in each EW configuration as the cumulative fraction of the total number of full-bandwidth detection (left y-axis). The secondary y-axis shows the same as a function of the expected number of BNS detections per year in O4. Each color represents one EW configuration. Using a simulated Gaussian environment for O4, we expect to recover at least 1 (6%) premerger event per year with a localization area of <300 deg^2 with the H1L1 configuration in O4.

Figure 4. Box and whiskers plot comparing the searched and 90% credible areas for all the runs. The y-axis represents the localization areas, and the x-axis represents the truncation latency for each EW run. The boxes encompass 95% of the events and the whiskers extend up to the rest. The white lines within the boxes represent the median values of the respective data sets. Using simulated Gaussian data, we demonstrate that the median searched areas are several times lower than the 90% credible areas, implying that a lot less area (as compared to the 90% credible region) can be searched to localize the GW event in O4.

GWs from BNSs before their merger in O4. This configuration saves an additional 4 s of data-transfer latency from Virgo, making it the most probable combination to issue negative latency alerts in O4. However, the major drawback of using this network is that the source localization areas for bright events are poorer compared to three- or four-detector networks. This is evident from the full-bandwidth detection of GW170817 during O2, where the localization area improved

Table 1
Expected BNS Detection Rate (N_{events}), Median Searched Area (Area$_{\text{se}}$),
Median 90% Credible Area (Area$_{90}$), and the Median Latency (with 90% Error
Margins) Associated with Different EW Simulations for O4

EW	$N_{\text{events}}(\text{yr}^{-1})$	Area$_{\text{se}}(\text{deg}^2)$	Area$_{90}(\text{deg}^2)$	Latency (s)
0 s	18	245	911	$9.46^{+1.42}_{-0.76}$
−10 s	3.51	876	4912	$-2.30^{+1.15}_{-1.20}$
−20 s	1.64	886	5673	$-12.36^{+1.31}_{-1.15}$
−30 s	0.83	902	6243	$-22.35^{+1.33}_{-1.18}$
−40 s	0.42	995	6299	$-32.32^{+1.29}_{-1.17}$
−50 s	0.23	1178	7187	$-42.27^{+1.19}_{-1.18}$
−60 s	0.12	1104	6803	$-52.25^{+1.13}_{-1.14}$
−70 s	0.04	1136	7683	$-61.86^{+0.81}_{-0.82}$

by a factor of ∼6 with the inclusion of an additional detector (Abbott et al. 2017a). The main motivation behind our choice of a two-detector network in O4 is to facilitate the rapid follow-up detection of short transient signals (only ∼1 s long) within seconds of a BNS merger (Rezzolla et al. 2011) that was not possible before. Given that (a) most major follow-up telescopes require a response time of ≳10 s to reorient themselves to a particular sky direction, (b) the anticipated event rate is small for early detection of BNS events in O4 (even just 10 s before their final merger), and (c) the latencies of the existing pipelines are larger than 10 s, it is extremely challenging to observe short transients right at the merger. Thus, saving a few seconds of latency might prove to be extremely useful to capture such events using instruments with larger viewing areas and fast response times. Configurations including other detectors can also be run in parallel for better localizations but with a slight delay.

Using a Gaussian noise simulation, we show that EW pipelines have an exceptional recovery of chirp mass, which is helpful in classifying the source properties such as p_astro and embright. We also compare the 90% credible areas and the searched areas associated with the EW triggers and demonstrate that the actual searched area for a detection is about one order of magnitude better. These areas are still a few times larger than the three- and four-detector network localizations reported in Magee et al. (2021), Sachdev et al. (2020), and Nitz et al. (2020), but this trade-off of several to tens of seconds in latency is something we hope to take in O4. It is important to note that a more accurate three-detector and four-detector localization will be provided by the full-bandwidth SPIIR search.

We demonstrate the possibility of issuing alerts for at least one BNS merger per year ∼12 s before the merger in Table 1. Adding an additional 8 s for data transfer and localization, this would still be a prospect for a negative latency alert. The best localized EW alert is likely to be from the −10 s search, and we expect to deliver alerts localized within $300\,\text{deg}^2$ at a nominal rate of one detection per year from this search.

These alerts are especially useful for follow-up observatories that have large FOVs. Radio telescopes, like the Murchison Widefield Array (MWA; Tingay et al. 2013), which have the capability of viewing a quarter of the sky, are one such recipient. James et al. (2019) have investigated the response time of this MWA observational mode to be about 10 s. In the future, EW alerts produced by SPIIR can also be ingested by X-ray and gamma-ray missions with large FoVs. For example, the Neil Gehrels Swift Observatory (FoV ∼ $4600\,\text{deg}^2$) has

developed a fully autonomous, extremely low-latency onboard commanding pipeline (GUANO; Tohuvavohu et al. 2020) capable of recovering subthreshold Burst Alert Telescope (BAT) triggers. Similarly, the large FoVs of the FERMI gamma-ray telescope's gamma-ray burst monitor (FoV ∼ 3.2e4 deg^2; Meegan et al. 2009) and INTEGRAL's gamma-ray burst detection subsystems (FoV all-sky for SPI-ACS, except for regions occulted by Earth; Rau et al. 2005) are also well suited to following up SPIIR triggers. Thus, the SPIIR EW pipeline will contribute significantly to the possibility of targeted follow-up of sGRB signals by the global astronomy community.

ORCID iDs

Manoj Kovalam ⓘ https://orcid.org/0000-0001-8143-9696
Md Anwarul Kaium Patwary ⓘ https://orcid.org/0000-0003-0760-3835
Linqing Wen ⓘ https://orcid.org/0000-0001-7987-295X
Fiona H. Panther ⓘ https://orcid.org/0000-0002-2618-5627

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, CQGra, 32, 074001
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, PhRvL, 116, 061102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, PhRvL, 119, 161101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, ApJL, 848, L12
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, Natur, 551, 85
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, LRR, 21, 3
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, PhRvX, 9, 011001
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, PhRvX, 11, 021053
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021b, arXiv:2108.01045
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021c, arXiv:2111.03606
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001
Acernese, F., Agathos, M., Aiello, L., et al. 2019, PhRvL, 123, 231108
Adams, T., Buskulic, D., Germain, V., et al. 2016, CQGra, 33, 175012
Akutsu, T., Ando, M., Arai, K., et al. 2021, PTEP, 2021, 05A101
Aubin, F., Brighenti, F., Chierici, R., et al. 2021, CQGra, 38, 095004
Buikema, A., Cahillane, C., Mansell, G. L., et al. 2020, PhRvD, 102, 062003
Buonanno, A., Iyer, B., Ochsner, E., Pan, Y., & Sathyaprakash, B. S. 2009, PhRvD, 80, 084043
Chu, Q., Kovalam, M., Wen, L., et al. 2022, PhRvD, 105, 024023
Cutler, C., & Flanagan, E. E. 1994, PhRvD, 49, 2658
Dal Canton, T., Nitz, A. H., Gadre, B., et al. 2021, ApJ, 923, 254
Davis, D., Areeda, J. S., Berger, B. K., et al. 2021, CQGra, 38, 135014
Fairhurst, S. 2018, CQGra, 35, 105002
Finn, L. S. 1992, PhRvD, 46, 5236
Foucart, F., Hinderer, T., & Nissanke, S. 2018, PhRvD, 98, 081501
Guo, X., Chu, Q., Chung, S. K., et al. 2018, CoPhC, 231, 62
Hooper, S., Chung, S. K., Luan, J., et al. 2012, PhRvD, 86, 024012
James, C. W., Anderson, G. E., Wen, L., et al. 2019, MNRAS, 489, L75

Kapadia, S. J., Caudill, S., Creighton, J. D. E., et al. 2020, CQGra, 37, 045007
Klimenko, S., Vedovato, G., Drago, M., et al. 2016, PhRvD, 93, 042004
LIGO Scientific Collaboration 2018, LIGO Algorithm Library - LALSuite, free software (GPL), doi:10.7935/GT1W-FZ16
Liu, Y., Du, Z., Chung, S. K., et al. 2012, CQGra, 29, 235018
Luan, J., Hooper, S., Wen, L., & Chen, Y. 2012, PhRvD, 85, 102002
Magee, R., Chatterjee, D., Singer, L. P., et al. 2021, ApJL, 910, L21
Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, ApJ, 702, 791
Messick, C., Blackburn, K., Brady, P., et al. 2017, PhRvD, 95, 042001

Nitz, A. H., Dal Canton, T., Davis, D., & Reyes, S. 2018, PhRvD, 98, 024050
Nitz, A. H., Schäfer, M., & Dal Canton, T. 2020, ApJL, 902, L29
Rau, A., Kienlin, A. V., Hurley, K., & Lichti, G. G. 2005, A&A, 438, 1175
Rezzolla, L., Giacomazzo, B., Baiotti, L., et al. 2011, ApJL, 732, L6
Sachdev, S., Caudill, S., Fong, H., et al. 2019, arXiv:1901.08580
Sachdev, S., Magee, R., Hanna, C., et al. 2020, ApJL, 905, L25
Singer, L. P., & Price, L. R. 2016, PhRvD, 93, 024013
Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, PASA, 30, 7
Tohuvavohu, A., Kennea, J. A., DeLaunay, J., et al. 2020, ApJ, 900, 35