



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Rapid detection of gravitational waves from compact binary mergers with PyCBC Live

Alexander H. Nitz, Tito Dal Canton, Derek Davis, and Steven Reyes

Phys. Rev. D **98**, 024050 — Published 30 July 2018

DOI: [10.1103/PhysRevD.98.024050](https://doi.org/10.1103/PhysRevD.98.024050)

PyCBC Live: Rapid Detection of Gravitational Waves from Compact Binary Mergers

Alexander H. Nitz¹, Tito Dal Canton², Derek Davis³, Steven Reyes³

¹ *Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, D-30167 Hannover, Germany*

² *NASA Postdoctoral Program Fellow, Goddard Space Flight Center, Greenbelt, MD 20771, USA and*

³ *Syracuse University, Syracuse, NY 13244, USA*

(Dated: July 9, 2018)

We introduce an efficient and straightforward technique for rapidly detecting gravitational waves from compact binary mergers. We show that this method achieves the low latencies required to alert electromagnetic partners of candidate binary mergers, aids in data monitoring, and makes use of multidetector networks for sky localization. This approach was instrumental to the analysis of gravitational-wave candidates during the second observing run of Advanced LIGO, including the period of coincident operation with Advanced Virgo, and in particular the analysis of the first observed binary neutron star merger GW170817, where it led to the first tightly localized sky map (31 deg²) used to identify AT 2017gfo. Operation of this analysis also enabled the initial discovery of GW170104 and GW170608 despite non-nominal observing of the instrument.

I. INTRODUCTION

The era of gravitational-wave multi-messenger astronomy began with the near simultaneous detection of a binary neutron star merger through gravitational waves (GW170817) [1] and gamma rays (GRB 170817A) [2], spawning numerous astronomers and telescopes to perform observations [3], and led to the first identification of an optical counterpart to gravitational-waves (AT 2017gfo) [4]. This single observation has confirmed long held suggestions that neutron star mergers can produce short gamma-ray bursts [5–8] and r-process fueled kilonova [9–11].

In this paper we present a new method to detect gravitational-wave events, employed during the second observing run of Advanced LIGO (O2) and during the coincident observation with Virgo in 2017, known as **PyCBC Live**. This analysis is designed to find gravitational waves from binary neutron star, neutron star-black hole, and binary black hole mergers in Advanced LIGO and Virgo data for the purpose of alerting astronomers to perform observational follow-up. Existing methodologies to search for gravitational-waves in low latency rely on elaborate use of multi-rate filtering [12–17], basis construction by means of singular value decomposition [13], or even decomposition into a series of infinite impulse response filters [18]. We show that reuse of the existing infrastructure for doing straightforward frequency-domain matched filtering used by the full deep PyCBC-based offline analysis [19–21] is similarly able to achieve high-performance, low-latency identification of gravitational-waves, with a substantially simpler architecture and approach.

At the time of this writing, in addition to the binary neutron star merger GW170817, three binary black hole mergers have been reported by Advanced LIGO and Virgo [22–24] during the O2 observing run and another three were previously detected in the first observing run [25–27]. Notably, GW170104 and GW170608 were identified using **PyCBC Live** during adverse observational conditions [28, 29]. GW170814 was detected dur-

ing standard autonomous analysis [30]. We note that the final significance statements for events are typically made using deep offline analysis such as those discussed in [13, 19, 20]. Unlike low latency analysis, these are able to employ more comprehensive data quality information [31–33] and updated instrumental calibration [34]. The initial identification of GW170817 was based on detection in the LIGO Hanford observatory alone using the method presented in [13, 35]. As a single detector analysis alone cannot provide a well constrained sky localization, the precise localization, which enabled the identification of the optical counterpart to GW170817, required the rapid analysis of data from the LIGO and Virgo observatories by **PyCBC Live** [36] and subsequent localization using the **BAYESTAR** algorithm [37].

In this paper, we discuss how **PyCBC Live** extracts gravitational-wave candidates under the requirement to produce results in low latency, how these candidate events are ranked, and how their statistical significance is estimated. We show that this technique can produce alerts with accurate sky localizations using the information from multiple detectors, which was demonstrated during the analysis of GW170817. Furthermore, using data from the first observing run of Advanced LIGO as a test bed, we demonstrate that the sensitivity of this analysis compares favorably to the full deep PyCBC-based offline analysis. We will discuss applications of this analysis to searching for gravitational waves from a single detector along with future enhancements. We note that during the O2 observing period, this analysis operated primarily on data from LIGO Hanford and LIGO Livingston when extracting signals and determining their significance. We will discuss the analysis in this context, however, the procedures described here are generalizable to additional observatories, which will be the configuration used in future observing runs.




FIG. 1. A diagram of how data is processed by the PyCBC Live low-latency analysis. Data is analyzed in discrete *analysis segments* with a fixed *analysis stride* between them, in this case 8 seconds. The amount of data required varies by the template being analyzed. The maximum latency of this method is 20s for the configuration used in the diagram above. The average latency from merger to identification is 16s. In practice, additional time is required to calibrate and distribute the strain data to computing resources where the low latency analysis takes place, typically ~ 10 s. Furthermore, critical information such as the sky localization, nearly takes an additional 10s to be generated after an event is initially identified and uploaded to the Gravitational-wave Candidate Database (GraceDB) [38]. This occurs as a separate follow-up process.

II. METHODOLOGY

The most sensitive searches for gravitational waves from compact binary mergers of LIGO/Virgo data have relied upon matched filtering [19, 39–41], which is an optimal technique for extracting signals from stationary colored Gaussian noise. Since the data is not Gaussian, but instead contains many classes of transient non-Gaussian noise, numerous additional signal consistency tests are also employed [41] to mitigate the large signal-to-noise ratios that can be produced [42, 43]. This method works when there are accurate models of the gravitational waveform [44–47]. A discrete set of waveform templates is chosen carefully to search over a wide range of binary component masses and spins [48–51]. The parameter space searched and the templates chosen for the first and second observing runs of Advanced LIGO are described in [52, 53].

In this section we will detail how we adapt these methods for low latency analysis to extract candidate events from gravitational-wave data (II A), rank candidates found in one or more detectors (II B), and empirically estimate the background to determine the rate of false alarms (II C). We find that we are able to achieve a high throughput low-latency analysis, using only straightforward frequency-domain matched-filtering, based on the methods described in [19] and [41]. The philosophy behind the design of this analysis has been to achieve low latencies while maintaining a straightforward filtering design and leveraging the tools built into the PyCBC gravitational-wave analysis library [54] which was used to construct the deep offline analysis that determines the significance of LIGO/Virgo detections [19, 20].

The latency goal of PyCBC Live is determined by ex-

ternal latencies such as the production, aggregation and transfer of the strain data, the sky localization of candidates, and the generation and distribution of the resulting alerts. All these processes set the appropriate latency time scale to achieve at $O(10$ s), which is achieved by PyCBC Live as described next.

A. Extracting Gravitational-wave Signals

This section focuses on the extraction of gravitational-wave signals from observatory data: the aim of this step is to generate a matched-filter signal-to-noise-ratio (SNR) time series for each template in a predefined bank of waveforms. Each template represents a specific choice of signal parameters the analysis is searching for. Low-latency design requires that the analysis of a portion of data be completed for the entire template bank within a short, fixed amount of time, and be able to keep pace with the incoming data. This portion of the analysis is highly parallel and may conceptually proceed independently for each template and each observatory’s data being analyzed. Finite impulse response (FIR) filters are used to guarantee a fixed latency with predetermined time invariant properties. As such, the overall latency of our analysis is limited by the duration of all the filtering operations, along with the *analysis stride*, which determines the pace at which data is analyzed. For the filter configuration used in the O2 LIGO/Virgo observing run, the maximum latency incurred by this analysis method is 20s, with an average latency of only 16s.

In O2 the *analysis stride* was chosen to be 8 seconds. New data is read in and preconditioned by highpassing and resampling using zero-phase FIR filters generated

by applying a Kaiser window to the ideal frequency response [55]. The highpass filter corner frequency is set at 15 Hz, primarily to reduce the dynamic range of the data. This allows the computationally-intensive part of the analysis to be performed in single-precision floating point arithmetic. We also downsample the data, which is originally recorded at 16384 Hz, to 2048 Hz. We can safely discard the higher frequency data as the bulk of the signal power is accumulated well below the Nyquist frequency of 1024 Hz, even for binary neutron star signals which reach several kHz. The portion of time-domain data corrupted due to filtering is then discarded and a continuous stream of preconditioned data, free from boundary artifacts, is produced by combining with previous data. The preconditioned stream is delayed by ~ 0.5 s due to the chosen duration of the highpass and resampling FIR filters.

The matched filter signal-to-noise ratio ρ is constructed according to [40, 41], for a template waveform h , and data s , which can be defined as,

$$\rho^2 \equiv \frac{\| \langle s|h \rangle \|^2}{\langle h|h \rangle} \quad (1)$$

where the inner product is

$$\langle a|b \rangle = 4 \int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} df \quad (2)$$

and $S_n(f)$ is the estimated one-sided power spectral density (PSD) of the noise around the time of a candidate event. The power spectral density is estimated from each observatory's data using Welch's method and median averaging [56]. The prior minute of data is divided into 4s intervals which overlap by 50%, and are then hann-windowed and converted to the frequency domain. The median of each frequency bin is then taken across all intervals, thus preventing short transient signals from biasing the $S_n(f)$ estimate. The median estimate is also robust to the presence of compact-binary-merger signals in the data, since the monotonic increase in frequency of such signals over time causes them to only affect one sample of any frequency bin.

The signal-to-noise time series is,

$$\rho^2(t) = \frac{4}{\langle h|h \rangle} \int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*(f)}{S_n(f)} e^{2\pi ift} df \quad (3)$$

which can be efficiently computed by use of the inverse fast Fourier transform.

Fig 1 provides an overview of the filtering stages and their data use requirements. The *analysis segment*, which contains the complete amount of data needed to produce a portion of uncorrupted SNR time series is shown. We can consider the action of the matched filter itself as an FIR filter with a length determined by the signal duration in the analyzed frequency band. For the analysis of O2 data, this duration is determined by the lower frequency cutoff chosen independently for each template.

Using this lower frequency cutoff accumulates at least 99.5% of signal power relative to using a much lower frequency cutoff[53]. A fixed number of segment sizes (typically multiples of 16 s) is chosen, allowing templates to be grouped together into sets of similar duration, and analyzed together using batched operations.

Overwhitened data is produced for each *analysis segment* size by constructing an FIR filter from $1/S_n(f)$ via a truncation in the time domain and multiplying the frequency-domain data by the corresponding transfer function [41]. The amount of data corrupted by this filter is discarded, introducing an additional ~ 3.5 s of latency. We can finally compute an *analysis stride* worth of valid SNR outputs for each template by way of batched inverse FFT.

B. Selecting and Ranking Gravitational-wave Candidate Events

In this section we will present how candidate events are identified and how they are ranked based on their signal and astrophysical consistency. We will restrict the discussion to the LIGO-Hanford and LIGO-Livingston observatories, as this was the configuration used in O2, but we note that these procedures can be extended to additional observatories.

Using the SNR time series for each template and observatory, we first identify the peak SNR in each analysis stride. This defines a single-detector *trigger*. Generating triggers for *each* local SNR maximum would lead to a prohibitively large number of triggers. An effective and widely-used means of controlling the number of triggers we need to analyze with minimal loss of sensitivity is to threshold on the peak SNR at ~ 5.5 . The bank of templates is divided among dozens of compute nodes for parallel analysis. Only the loudest O(10) triggers from the set of templates being analyzed by a particular compute node are kept. In all we accept O(10^3) single detector triggers from each observatory for every analysis segment. Comparison of this procedure to the one used for the deep PyCBC-based offline analysis presented in [19, 20] shows this procedure is robust, and recovers 98% of identified signals from a simulated population of sources.

Due to time constraints, low-latency analyses do not have access to the full set of data quality information used to determine when the detectors are in suboptimal states. In addition to making use of the minimal set of low latency data quality information, which include times of electronic control system overflows [57], we make additional cuts based on a continuous local estimation of detector's sensitivity, the variability of the PSD over time, and the measurement of an unreasonably large SNR. These cuts serve to limit times where the instrument is clearly misbehaving in ways that would hinder the analysis, but are not currently flagged by other automated processes.

In addition, we apply a more robust signal consistency

test that has been successfully employed in the deep of-line analyses [19, 58]. For each single detector trigger we calculate a χ^2 signal consistency test defined as

$$\chi_r^2 = \frac{1}{2p-2} \sum_{i=1}^{i=p} \|\langle s|h_i \rangle - \langle h_i|h_i \rangle\|^2, \quad (4)$$

where the template h is divided into p frequency bands which each contribute equally to the signal power. If the data, s , is adequately approximated by Gaussian noise with an embedded signal that is described by the waveform template h , this statistic will follow a reduced χ^2 distribution with $2p-2$ degrees of freedom. Many classes of non-Gaussian noise, however, tend to produce exceedingly large χ_r^2 values [58, 59], allowing the classification of triggers as due either to astrophysical signals or non-Gaussian transient noise. This test only needs to be performed for the few SNR peaks above threshold so it is not a significant fraction of the total cost of the analysis.

We use the re-weighted SNR, introduced in [59] to suppress triggers caused by non-Gaussian noise transients. This re-weighted SNR, $\tilde{\rho}$, is given as

$$\tilde{\rho} = \begin{cases} \rho & \text{for } \chi_r^2 \leq 1 \\ \rho \left[\frac{1}{2} \left(1 + (\chi_r^2)^3 \right) \right]^{-1/6} & \text{for } \chi_r^2 > 1 \end{cases}. \quad (5)$$

To combine triggers from multiple detectors, we employ the statistic first described in [20], which includes the astrophysical probability of a trigger observed with properties $p^S(\vec{\theta})$, where $\vec{\theta}$ includes relative amplitudes, phases, and time differences between observatories. This statistic can be expressed as,

$$\tilde{\rho}_c^2 = \tilde{\rho}_H^2 + \tilde{\rho}_L^2 + 2 \ln \left(p^S(\vec{\theta}) \right), \quad (6)$$

where $\tilde{\rho}_L$ and $\tilde{\rho}_H$ are the re-weighted SNRs of the triggers in the LIGO Livingston and LIGO Hanford observatories respectively. This was the statistic used by the low latency PyCBC Live analysis used in O2. In future analyses, this statistic may be extended as in [20] to account for noise variation over intrinsic parameters such as the masses and spins of a binary system which was used by the PyCBC-based offline analysis, and could take into account priors over the intrinsic parameters. We further require that only single detector triggers that arise from the same template in different observatories may form a coincident candidate event. Finally, we impose a maximum 15 ms time delay between triggers from the two observatories, as the largest possible Hanford-Livingston time-of-flight is ~ 10 ms. If there are multiple candidate events from different templates, we choose the one with the largest $\tilde{\rho}_c$ value.

C. Background Estimation

As the next step, we need to determine if a candidate is significant enough for consideration by astronomers and




FIG. 2. Variability of the relationship between the ranking statistic and the false alarm rate (FAR) for independent analysis periods during O1. A green line is placed at the threshold for astronomical alerts (1 per 2 months). Background estimates were computed over 5 hour analysis periods. The mapping between the ranking statistic and the FAR is stable with the notable exception of classes of astrophysical (blue) and instrumental (red) outliers.

intermediate follow-up processes. To determine the statistical significance of a candidate event, we calculate the false alarm rate of the analysis producing an event with as large a statistic value as the candidate event. The background of false coincidences is measured empirically by the resampling procedure of the data.

The data in one detector is shifted in time relative to the other detector to create a simulated analysis. Since we use time shifts that are greater than the maximum physical time-of-flight between observatories, which is 10 ms between the LIGO observatories, this analysis will only produce non-astrophysical coincidences. We repeatedly shift the data by increments of 100 ms to generate as much background as possible. Since the analysis of each detector is independent of the time shift, we can in practice calculate the background of false coincidences by many recombinations of the single-detector triggers, which involve computationally trivial operations. This is the same procedure used by the offline analysis [19].

For the low-latency analysis, a rolling buffer of single detector triggers is stored, which is approximately 5 hours in duration. This is used to estimate the significance of a given event down to a false alarm rate of 1 per 100 years. The limit of the inverse FAR that can be estimated, IFAR, is related to the amount of past data stored, T_{buffer} , and the time shift size, T_{shift} , by

$$IFAR = \frac{T_{buffer}^2}{T_{shift}} \quad (7)$$

Since only 5 hours of past data is kept, the background estimate will adapt to changes in the detector state on

this time scale. This is important given that the analysis will inevitably include times of highly nonstationary data, which may be later identified and vetoed from the full offline, archival analyses.

To demonstrate how different data quality affects the background, independent 5-hour data sections from the first Advanced LIGO observing run were analyzed with the PyCBC Live analysis for a total of 14.8 days. A single background estimate from each was recorded and the variability of the background estimates calculated across these sections can be seen in Fig 2. Beyond a few outliers, the background estimates for independent analysis periods is stable, with the threshold for astronomical alerts (1 false alarm per 2 months) corresponding to a ranking statistic of 8.55 to 8.75.

The two cases when the background estimate is not stable are in the presence of a loud coincident event (blue) and a broadband disturbance in the PSD (red). Triggers participating in coincident events are included in the background estimate to ensure an unbiased rate of false alarms [60]. Similar behavior is seen in estimates from the high-latency searches [19, 60].

Due to the low-latency nature of this pipeline, fewer data quality products are available to be utilized by the search. The data quality information distributed in low-latency includes only short term disturbances that are often vetoed by the consistency checks the search incorporates [32]. Data quality products that indicate broadband biases to the PSD, which generally necessitate the removal of that time from analyses, are not processed in low-latency, and hence not available to this search. These disturbances lead to an increase in the rate of loud non-coincident events that are incorporated into the background. In O2, the amount of time that the low-latency search processed, but was not searched by the PyCBC-based offline analysis due to additional data quality inputs is approximately 2.1 days, indicating that 2% of data which the low-latency search would analyze is later thrown out in the offline analysis [57].

D. Architecture and Computational Considerations

In this section we will discuss the high level architecture of the PyCBC Live analysis, along with the computational cost and scaling. Fig 3 provides a high level diagram of how the analysis is ordered. We see that data from the observatories is first distributed to computing clusters, such as the ATLAS supercomputer [61], onto cluster nodes via multicast.

Once data is distributed to the cluster nodes, the PyCBC Live analysis takes over and calculates the SNR time series for each template and observatory, which is a highly parallel process. To ensure the analysis completes in a time shorter than the *analysis stride*, we distribute the work over multiple computing nodes by use of task parallelization via Message Passing Interface (MPI [62]), represented in blue in Fig 3.

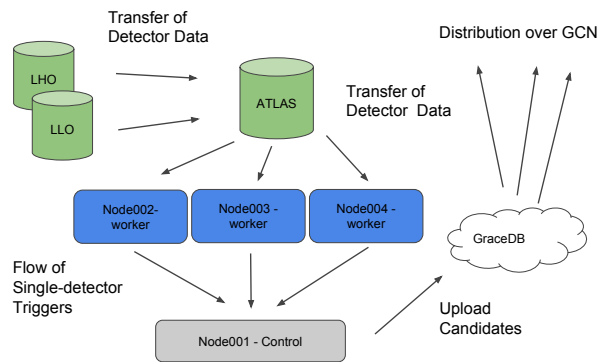


FIG. 3. The high level overview of how data is flowing from the observatories through the PyCBC Live analysis, to GraceDB[38], and finally to astronomers by way of the Gamma-ray Coordinates Network (GCN). The work nodes (blue) process different portions of a template bank, while the control node (gray) collates all these results and determines if there is a significant candidate event.

In the early advanced-detector era, template banks typically contain $O(10^5)$ templates. However, template waveforms are sufficiently short that they can be generated once at the beginning of an analysis and stored in system memory indefinitely. We see in Fig 1 how the duration of the template affects the length of an individual *analysis segment*. Templates are grouped in batches of similar duration and analyzed together. This allows for the calculation of the SNR time series to be further parallelized on each compute node using OpenMP by taking advantage of the well optimized batched FFT algorithms provided by FFTW [63] and Intel’s MKL library.

Conceptually, each compute node handles its own section of the template bank independently of the others and produces its own set of single detector triggers. This computation comprises the vast majority of the overall computing cost of the analysis. It is efficient to simply transfer the recorded triggers from all the nodes to a single control node (gray in Fig 3) where triggers from separate observatories are combined into candidate events, ranked, and assigned a statistical significance after the background is estimated. This same process is responsible for finally uploading a candidate event to the Gravitational-wave Candidate Database (GraceDB) [38] if it passes nominal false alarm rate thresholds (typically 1 / day).

To ensure that the analysis keeps up with the incoming data, it is sufficient to simply use enough worker nodes such that it completes the previous set of analysis segments before it must begin the next. In practice, the analysis is configured to only take 80% of this time. During an observing run there are disruptions in the data stream, or other delays, which cause the data distribution and analysis to fall behind. When the disruption subsides, this extra headroom allows the analysis to catch up rapidly.

A common measurement of computational perfor-

mance for a template-based gravitational-wave search is how many templates a CPU core can analyze in real time. During O2 this analysis achieved an in-situ performance of 5000 templates per core at 80% load, or 6300 templates per core at full machine load. The computational scaling of this analysis is related to the length of templates which need to be analyzed, and the latency desired. For a similar analysis to the one performed in O2, reducing the *analysis stride* from 8 to 4 seconds, decreases the average analysis-induced latency from 16 to 10 seconds, while increasing the computational cost by just under a factor of 2. This occurs because the computational cost C has roughly the following scaling relation,

$$C \propto \frac{\bar{T} \ln(\bar{T})}{T_{stride}} \quad (8)$$

where \bar{T} is the template duration averaged over the template bank, and T_{stride} is the *analysis stride*. \bar{T} is primarily determined by the shape of the PSD during a particular observing run. There is a balance between the latency of the analysis and the computational cost. Also note that a reduction in the analysis stride below a certain value no longer results in a proportionate drop in the average overall latency, due to unavoidable latency incurred from data transfers and conditioning filters.

III. SENSITIVITY OF THE ANALYSIS




FIG. 4. Binary component masses of the samples of the offline analysis injection set. BNS population is constrained between component masses $m_1, m_2 \in (1.0, 3.0) M_\odot$, the NSBH population ranges between $m_1 \in (1.0, 3.0)$, $m_2 \in (1.0, 97.)$, and the BBH population ranges between $m_1, m_2 \in (1.0, 99.0)$ and $m_2 + m_1 < 100$.

The sensitivity of a gravitational-wave search is estimated by simulating an astrophysical population of

Source Category	Relative Sensitivity
BNS	0.87 ± 0.05
NSBH	0.98 ± 0.04
BBH	1.05 ± 0.09

TABLE I. Relative sensitivity of the PyCBC Live low-latency analysis and the PyCBC-based offline analysis for different source classes evaluated using data from LIGO’s first observing run, and using a false alarm rate threshold of 1 per 2 months. We see that the PyCBC Live analysis is close to the sensitivity of the full offline analysis which benefits from improved data quality information, improved calibration, and a larger set of signal consistency tests [20, 65].

sources, adding the signals to a data set, and observing how many signals are detected. The source parameters of each simulated signal are selected from three binary populations that are expected to be the most significant sources of gravitational-waves from compact binary mergers: binary neutron stars (BNS), binary black holes (BBH), and neutron star-black hole binaries (NSBH). We model these sources using an effective-one-body, inspiral-merger-ringdown waveform model for NSBH and BBH signals and a post-Newtonian, inspiral-only model for BNS systems [44–47]. Once these signals are added into detector data from LIGO’s first observing run, the full PyCBC Live search is conducted and signals are recovered with an estimated false alarm rate. The sensitivity of the analysis can be determined by a weighted Monte Carlo integration of the distances at which sources are found above a given detection threshold [52, 64].

Here we take a subset of the software injections used to evaluate the offline PyCBC analysis [20] and compare the sensitivity of the low latency PyCBC Live and PyCBC-based offline analyses. To minimize differences where possible, we consider the same detection statistic for both, however, we note that the offline analysis used in the first observing run of Advanced LIGO treated background separately for templates corresponding to mergers with total mass less than four solar masses [52], whereas the PyCBC Live analysis treats them together. We expect this difference to increase the sensitivity of the offline analysis to lower-mass sources. We consider $O(1000)$ BNS, BBH, and NSBH-like sources. These injections sufficiently cover the parameter space investigated during the O1 offline analysis as can be seen in Figure 4. The source orientation and sky location is chosen isotropically, and sources are weighted based on their distance so as to achieve an effective uniform-in-volume distribution. We choose to evaluate the sensitivity at a false alarm rate of 1 per 2 months, to match the threshold used for alert generation by LIGO/Virgo. Table I shows that despite the differences in the PyCBC Live and PyCBC-based offline analysis the sensitivity is comparable and matches our expectation.

IV. MULTIDETECTOR SKY LOCALIZATION

The addition of a third instrument to the detector network dramatically reduces the uncertainty on the sky localization of candidate events, in some cases going from hundreds to tens of square degrees [66]. The additional detector does not need to be as sensitive as the others to provide a noticeable contribution, and in fact does not need to detect the candidate at all, as demonstrated by GW170817. The rapid advance of the Virgo commissioning during 2017, therefore, prompted the need to include its data in the rapid localization of low-latency candidates from PyCBC Live.

An optimal search of three-detector data involves matched-filtering all detectors separately and combining the result in a way that varies over a grid of many possible source sky locations [67]. Thus, its implementation is considerably more complex than the two-detector coincidence algorithm described earlier. However, if two of the detectors have a comparable sensitivity and the third is much less sensitive—a configuration expected for the LIGO/Virgo network towards the end of O2—applying the simple coincidence to the first two detectors and ignoring the third provides a reasonable approximation to the optimal search in terms of detection rate [67].

The rapid localization of candidates identified in low latency is performed by the BAYESTAR algorithm [37]. The most recent version of BAYESTAR takes as input a candidate’s complex SNR time series and a local estimate of the noise PSD from any number of detectors, regardless of which detectors actually reported a trigger for that candidate. Therefore, PyCBC Live proceeds as follows. If a candidate from the LIGO detectors is found significant enough to submit to GraceDB, we check the availability and quality status information of Virgo data for the entire duration of the template that produced the candidate. If the checks pass, we calculate the Virgo complex SNR time series around the time of the candidate, using the same template. This incurs negligible computation cost, since it involves a single template and less than 0.1s of data. The resulting set of SNR time series from LIGO and possibly Virgo is handed over to BAYESTAR to generate the low-latency localization, regardless of the SNR of the candidate in Virgo. The significance of the candidate does not use the Virgo SNR and is based on LIGO data only. In this way, we can rapidly produce a three-detector localization without having to calculate a three-detector ranking statistic and estimate its background distribution. Note that the same procedure is also applied to significant single-detector LIGO candidates (see Sec. V A), such that their localization is actually based on two detectors, if Virgo data are available.

Virgo is expected to be significantly more sensitive in O3 and will thus require a three-detector ranking statistic in order for the search to achieve its maximum sensitivity. Nevertheless, the technique described above can still be applied for new detectors—such as KAGRA—which may join the network with an initially reduced sensitivity as




FIG. 5. Reliability of the rapid sky localization of simulated signals detected by PyCBC Live and localized by BAYESTAR. The blue curve shows signals detected and localized using the LIGO Hanford (H) and Livingston (L) detectors only, and only using point estimates of the arrival time, arrival phase and SNR at the two detectors. In the orange curve, the point estimates are replaced by the full SNR time series from the two detectors. In the green curve, the SNR time series from Virgo (V) is included as well, regardless of the presence of a Virgo trigger. The gray band is the binomial 95% interval given the number of detected simulations. The fraction of signals whose true location falls within a given credible region is close to the expectation, although the HLV localizations tend to be slightly larger than expected.

their commissioning proceeds.

We test the described multidetector localization procedure by simulating 3000 BNS, NSBH and BBH signals immersed in Gaussian noise, detecting them with PyCBC Live and generating localization maps of the resulting candidates as described. Waveform models, similar to those used in Sec. III are used, and all sources are distributed uniformly in volume, sky location and spatial orientation. The detector noise is simulated from analytic PSD models, qualitatively comparable to estimates from the tenth engineering run of the LIGO detectors, and a factor of ~ 3 less sensitive for Virgo¹. Such models are chosen to represent the sensitivity of the detectors at the end of O2. The reliability of the produced localization maps can be tested by plotting the fraction of simulated signals whose true location falls within a certain probability of the posterior sky location. Figure 5 shows such a plot for three configurations: (i) the BAYESTAR algorithm employed during O1 and early O2, which only used point estimates of the candidate event time, ampli-

¹ Specifically the aLIGOMidLowSensitivityP1200087 and AdVEarlyLowSensitivityP1200087 models from LALSimulation [68].

tude and phase instead of the full SNR time series (blue), (ii) BAYESTAR in its late-O2 version, using the full SNR time series, but from LIGO detectors only (orange), and (iii) BAYESTAR using the full LIGO and Virgo SNR time series (green).

We find that using the SNR time series leads to consistent sky localizations, although the posteriors with Virgo are slightly conservative: about 60% of the 50% credible regions contain the simulated location. We attribute this small discrepancy to a tunable parameter in the BAYESTAR algorithm, which we have left at the default value for the purpose of this test and could in principle be optimized to achieve exactly consistent posteriors for the three-detector network as well. We conclude that the standard BAYESTAR settings produce consistent localizations for events observed by LIGO only, and slightly conservative localizations for events observed by Virgo as well. We find a typical reduction in uncertainty with the inclusion of Virgo to be roughly an order of magnitude, similarly to that found in [66].

V. OTHER APPLICATIONS

A. Single Detector Search

We have focused on the detection of gravitational-waves that are observed in multiple detectors. It may be difficult to claim a detection based solely on a single detector analysis, due to the background incurred from a detector’s non-Gaussian noise transients. However, as the sensitivity of the instrument increases, and the detection rate increases, it has been found that some of the most significant single detector triggers may be astrophysical in origin [69]. Maximizing the scientific output of a compact binary merger requires electromagnetic observations as close to the merger as possible [70], so rapidly identifying and distributing the most promising candidates when only a single detector is operating is warranted. Unfortunately, sky localization is a significant challenge in this case from gravitational-wave observation alone. On the other hand, observations of associated prompt gamma-ray emissions, for instance with wide-field observatories like Fermi/GBM [71], may be able to localize the event with sufficient precision for subsequent followup.

To allow for the detection of gravitational waves from compact binary mergers even when only a single detector is running, we can employ most of the tools we have already described. The ranking of gravitational-wave triggers at a single detector level, proceeds in the same manner. However, a few additional cuts are made to help cope with the decreased ability to reject noise. Candidates which are consistent with very short-duration signals are excluded, due to the difficulty in separating short duration transient noise from high mass binary black hole mergers [65]. In addition, our signal consistency test is more stringently applied and a cut directly on the re-

duced $\chi^2 < 4$ is used. The primary goal is to enable the detection of binary neutron star mergers, which are well modelled by the templates and chirp for a relatively long time, making our signal consistency tests particularly effective. In addition, binary neutron star candidates have the highest chance of being aided in sky localization by coincident wide-field gamma-ray observations. Finally, candidates are chosen and distributed based on an empirically-measured threshold which produces false alarms at a rate of less than 1 / month.

B. Data Monitoring

The PyCBC Live analysis was employed and ran continuously through the second observing run of Advanced LIGO, including the period of coincident observation with Virgo. Along with generating detection alerts, it was used as a low-latency monitoring tool of the data. Triggers, which include information such as time, signal-to-noise, template masses and spins, along with signal consistency tests, are displayed through the LIGO detector characterization summary pages [72]. These are regularly monitored by analysts for anomalous behavior, which could be a sign of adverse detector behavior. This regular monitoring has also been critical to the detection of two gravitational-wave events during O2; both GW170104 and GW170608 were initially discovered by visual inspection of PyCBC Live results from these monitoring pages [28, 29]. In particular, GW170608, was discovered during time where the LIGO-Hanford instrument was in a commissioning phase [24].

VI. CONCLUSIONS

In this paper we have presented PyCBC Live, a simple and effective analysis designed to rapidly detect gravitational-waves from compact binary mergers, built upon the existing technologies used for the well-established offline analysis [19], and distributed as free and open software in the PyCBC library [54]. PyCBC Live’s successful operation during O2, where it has been involved in either the detection or analysis of all currently published gravitational-wave events [28–30, 36], demonstrates that well-known, standard frequency-domain matched filtering is a successful technique for the low-latency detection of gravitational waves.

During the production conditions of O2, PyCBC Live reported the detection of gravitational-wave candidates with a median latency of 29s and a full range of 20-38s. This includes the O(10) seconds of latency expected from data distribution and network communication, and can be directly compared to the reported median latencies (57 and 66s) of the two analyses reported in [73].

We have primarily focused on the coincident detection of a gravitational-wave observations at the two current LIGO interferometers. In the coming years, the

Virgo detector will improve in sensitivity, while KAGRA and LIGO India will also join the worldwide network of gravitational-wave observatories [74]. The analysis can already produce accurate sky localizations from any number of additional observatories. The analysis can also be straightforwardly extended to analyze more than two detectors in a symmetric manner. A similar technique to the multidetector compact binary merger analysis performed in the fifth and sixth observing runs of initial LIGO/Virgo [59] will be investigated in the future.

ACKNOWLEDGMENTS

We thank Thomas Dent, Duncan Brown, and Collin Capano for their input, Leo Singer for his help with BAYESTAR, along with Laura Nuttall, Gregory Mendell and Alan Weinstein for their useful comments and feedback during the internal LIGO Scientific Collaboration review process.

We acknowledge the Max Planck Gesellschaft for support and the Atlas cluster computing team at AEI Hannover. Computations were also supported by Syracuse University and NSF award OAC-1541396. We also thank the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council.

DD acknowledges support from NSF award PHY-1607169, SR acknowledges support from NSF awards PHY-1707954 and PHY-1607169. TDC 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.

-
- [1] B. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017), [arXiv:1710.05832 \[gr-qc\]](#).
- [2] A. Goldstein *et al.*, *Astrophys. J.* **848**, L14 (2017), [arXiv:1710.05446 \[astro-ph.HE\]](#).
- [3] B. P. Abbott *et al.* (GROND, SALT Group, OzGrav, DFN, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAAS-TRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), Dark Energy Survey, MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, ALMA), *Astrophys. J.* **848**, L12 (2017), [arXiv:1710.05833 \[astro-ph.HE\]](#).
- [4] D. A. Coulter *et al.*, *Science* (2017), [10.1126/science.aap9811](#), [Science358,1556(2017)], [arXiv:1710.05452 \[astro-ph.HE\]](#).
- [5] B. Paczynski, *ApJL* **308**, L43 (1986).
- [6] J. Goodman, *ApJL* **308**, L47 (1986).
- [7] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, *Nature* **340**, 126 (1989).
- [8] R. Narayan, B. Paczynski, and T. Piran, *ApJL* **395**, L83 (1992), [astro-ph/9204001](#).
- [9] L.-X. Li and B. Paczyński, *ApJL* **507**, L59 (1998), [astro-ph/9807272](#).
- [10] S. Rosswog, *ApJ* **634**, 1202 (2005), [astro-ph/0508138](#).
- [11] B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner, *Mon. Notices Royal Astron. Soc.* **406**, 2650 (2010), [arXiv:1001.5029 \[astro-ph.HE\]](#).
- [12] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, *Class. Quant. Grav.* **33**, 175012 (2016), [arXiv:1512.02864 \[gr-qc\]](#).
- [13] C. Messick *et al.*, *Phys. Rev. D* **95**, 042001 (2017), [arXiv:1604.04324 \[astro-ph.IM\]](#).
- [14] F. Marion, in *Proceedings of Rencontres de Moriond on Gravitational Waves and Experimental Gravity 2003* (2004) pp. 23–29.
- [15] D. Buskulic (LIGO Scientific Collaboration, Virgo Collaboration), *Proceedings, 14th Workshop on Gravitational wave data analysis (GWDAW-14): Rome, Italy, January 26-29, 2010*, *Class. Quant. Grav.* **27**, 194013 (2010).
- [16] K. Cannon *et al.*, *The Astrophysical Journal* **748**, 136 (2012).
- [17] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, *Phys. Rev. D* **82**, 044025 (2010).
- [18] S. Hooper, S. K. Chung, J. Luan, D. Blair, Y. Chen, and L. Wen, *Phys. Rev.* **D86**, 024012 (2012), [arXiv:1108.3186 \[gr-qc\]](#).
- [19] S. A. Usman *et al.*, *Class. Quant. Grav.* **33**, 215004 (2016), [arXiv:1508.02357 \[gr-qc\]](#).
- [20] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, *Astrophys. J.* **849**, 118 (2017), [arXiv:1705.01513 \[gr-qc\]](#).
- [21] T. Dal Canton *et al.*, *Phys. Rev.* **D90**, 082004 (2014), [arXiv:1405.6731 \[gr-qc\]](#).
- [22] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **119**, 141101 (2017),

- arXiv:1709.09660 [gr-qc].
- [23] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **118**, 221101 (2017), arXiv:1706.01812 [gr-qc].
- [24] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Astrophys. J.* **851**, L35 (2017), arXiv:1711.05578 [astro-ph.HE].
- [25] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al., *Physical Review Letters* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [26] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016), arXiv:1606.04855 [gr-qc].
- [27] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev.* **X6**, 041015 (2016), arXiv:1606.04856 [gr-qc].
- [28] LIGO Scientific Collaboration, Virgo Collaboration, *GCN Circ.* **20364** (2017).
- [29] LIGO Scientific Collaboration, Virgo Collaboration, *GCN Circ.* **21221** (2017).
- [30] LIGO Scientific Collaboration, Virgo Collaboration, *GCN Circ.* **21474** (2017).
- [31] L. K. Nuttall, *Phil. Trans. Roy. Soc. Lond.* **A376**, 20170286 (2018), arXiv:1804.07592 [astro-ph.IM].
- [32] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Class. Quant. Grav.* **35**, 065010 (2018), arXiv:1710.02185 [gr-qc].
- [33] T. J. Massinger, *Detector Characterization for Advanced LIGO*, Ph.D. thesis, Syracuse U. (2016).
- [34] C. Cahillane *et al.* (LIGO Scientific), *Phys. Rev.* **D96**, 102001 (2017), arXiv:1708.03023 [astro-ph.IM].
- [35] The LIGO Scientific Collaboration, The Virgo Collaboration, *et al.*, *GCN* **21505**, 1 (2017).
- [36] The LIGO Scientific Collaboration, The Virgo Collaboration, *et al.*, *GCN* **21513**, 1 (2017).
- [37] L. P. Singer and L. R. Price, *Phys. Rev.* **D93**, 024013 (2016), arXiv:1508.03634 [gr-qc].
- [38] G. Developers, “GraceDB,” <https://www.lsc-group.phys.uwm.edu/daswg/projects/gracedb.html>.
- [39] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (Wiley, 1949).
- [40] C. Cutler *et al.*, *Phys. Rev. Lett.* **70**, 2984 (1993), arXiv:9208005 [astro-ph].
- [41] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012), arXiv:0509116 [gr-qc].
- [42] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Class. Quant. Grav.* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].
- [43] L. K. Nuttall *et al.*, *Class. Quantum Grav.* **32**, 245005 (2015), arXiv:1508.07316 [gr-qc].
- [44] L. Blanchet, *Living Rev. Rel.* **5**, 3 (2002), arXiv:gr-qc/0202016.
- [45] G. Faye, S. Marsat, L. Blanchet, and B. R. Iyer, *Class. Quant. Grav.* **29**, 175004 (2012), arXiv:1204.1043 [gr-qc].
- [46] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, *et al.*, *Phys. Rev.* **D89**, 061502 (2014), arXiv:1311.2544 [gr-qc].
- [47] A. Bohé *et al.*, (2016), arXiv:1611.03703 [gr-qc].
- [48] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [49] S. Babak, R. Balasubramanian, D. Churches, T. Coke-laer, and B. S. Sathyaprakash, *Class. Quantum Grav.* **23**, 5477 (2006).
- [50] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein, *Phys. Rev. D* **89**, 084041 (2014), arXiv:1210.6666 [gr-qc].
- [51] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev.* **D86**, 084017 (2012), arXiv:1207.6406 [gr-qc].
- [52] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2016), arXiv:1602.03839 [gr-qc].
- [53] T. Dal Canton and I. W. Harry, (2017), arXiv:1705.01845 [gr-qc].
- [54] A. H. Nitz, I. W. Harry, J. L. Willis, C. M. Biwer, D. A. Brown, L. P. Pekowsky, T. Dal Canton, A. R. Williamson, T. Dent, C. D. Capano, T. J. Massinger, A. K. Lenon, A. B. Nielsen, and M. Cabero, “PyCBC Software,” <https://github.com/gwastro/pycbc> (2018).
- [55] E. Jones, T. Oliphant, P. Peterson, *et al.*, “SciPy: Open source scientific tools for Python,” (2001–).
- [56] P. Welch, *IEEE Transactions on Audio and Electroacoustics* **15**, 70 (1967).
- [57] L. S. Collaboration, (2018), in Preparation.
- [58] B. Allen, *Phys. Rev. D* **71**, 062001 (2005), arXiv:gr-qc/0405045 [gr-qc].
- [59] S. Babak, R. Biswas, P. Brady, D. Brown, K. Cannon, *et al.*, *Phys. Rev. D* **87**, 024033 (2013), arXiv:1208.3491 [gr-qc].
- [60] C. Capano, T. Dent, Y.-M. Hu, M. Hendry, C. Messenger, and J. Veitch, (2016), arXiv:1601.00130 [astro-ph.IM].
- [61] H. Fehrmann and C. Aulbert, “Gravity wave search with one of the world’s largest supercomputers.” (2009).
- [62] L. D. Dalcin, R. R. Paz, P. A. Kler, and A. Cosimo, *Advances in Water Resources* **34**, 1124 (2011), new Computational Methods and Software Tools.
- [63] M. Frigo and S. G. Johnson, in *PROCEEDINGS OF THE IEEE* (2005) pp. 216–231.
- [64] C. Capano, I. Harry, S. Privitera, and A. Buonanno, *Phys. Rev.* **D93**, 124007 (2016), arXiv:1602.03509 [gr-qc].
- [65] A. H. Nitz, *Class. Quant. Grav.* **35**, 035016 (2018), arXiv:1709.08974 [gr-qc].
- [66] L. P. Singer, L. R. Price, B. Farr, A. L. Urban, C. Pankow, *et al.*, *Astrophys. J.* **795**, 105 (2014), arXiv:1404.5623 [astro-ph.HE].
- [67] D. Macleod, I. W. Harry, and S. Fairhurst, *Phys. Rev.* **D93**, 064004 (2016), arXiv:1509.03426 [gr-qc].
- [68] “LSC Algorithm Library Suite,” <https://git.ligo.org/lscsoft/lalsuite>.
- [69] T. A. Callister, J. B. Kanner, T. J. Massinger, S. Dhurandhar, and A. J. Weinstein, *Class. Quant. Grav.* **34**, 155007 (2017), arXiv:1704.00818 [astro-ph.IM].
- [70] I. Arcavi, *The Astrophysical Journal Letters* **855**, L23 (2018).
- [71] C. Meegan, G. Lichti, P. N. Bhat, E. Bissaldi, M. S. Briggs, V. Connaughton, R. Diehl, G. Fishman, J. Greiner, A. S. Hoover, A. J. van der Horst, A. von Kienlin, R. M. Kippen, C. Kouveliotou, S. McBreen, W. S. Paciesas, R. Preece, H. Steinle, M. S. Wallace, R. B. Wilson, and C. Wilson-Hodge, *The Astrophysical Journal* **702**, 791 (2009).
- [72] D. Macleod, “Gravitational-wave interferometer summary information system,” <https://github.com/gwpy/gwsumm> (2017).
- [73] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J.* **832**, L21 (2016), arXiv:1607.07456 [astro-ph.HE].

[74] B. P. Abbott *et al.* (LIGO Scientific Collaboration,

Virgo Collaboration), [Living Rev. Relat. **19**, 1 \(2016\)](#),
[arXiv:1304.0670 \[gr-qc\]](#).