

The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs

Surabhi Sachdev,^{1,2,3} Sarah Caudill,^{4,5} Heather Fong,^{6,7} Rico K. L. Lo,^{8,3} Cody Messick,^{1,2} Debnandini Mukherjee,⁴ Ryan Magee,^{1,2} Leo Tsukada,^{7,9} Kent Blackburn,³ Patrick Brady,⁴ Patrick Brockill,⁴ Kipp Cannon,^{6,7} Sydney J. Chamberlin,^{1,2} Deep Chatterjee,⁴ Jolien D. E. Creighton,⁴ Patrick Godwin,^{1,2} Anuradha Gupta,^{1,2} Chad Hanna,^{1,10,2,11} Shasvath Kapadia,⁴ Ryan N. Lang,⁴ Tjonnie G. F. Li,⁸ Duncan Meacher,^{1,2,4} Alexander Pace,^{1,2} Stephen Privitera,¹² Laleh Sadeghian,⁴ Leslie Wade,¹³ Madeline Wade,¹³ Alan Weinstein,³ and Sophia Liting Xiao³

¹*Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA*

²*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA*

³*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

⁴*Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics,
University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA*

⁵*Nikhef, Science Park, 1098 XG Amsterdam, Netherlands*

⁶*Canadian Institute for Theoretical Astrophysics, 60 St. George Street,
University of Toronto, Toronto, Ontario, M5S 3H8, Canada*

⁷*RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan*

⁸*Department of Physics, The Chinese University of Hong Kong, Shatin, New Territories, Hong Kong*

⁹*Department of Physics, Graduate School of Science,
The University of Tokyo, Tokyo, 113-0033, Japan*

¹⁰*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA*

¹¹*Institute for CyberScience, The Pennsylvania State University, University Park, PA 16802, USA*

¹²*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany*

¹³*Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA*

(Dated: January 25, 2019)

After their successful first observing run (September 12, 2015 - January 12, 2016), the Advanced LIGO detectors were upgraded to increase their sensitivity for the second observing run (November 30, 2016 - August 26, 2017). The Advanced Virgo detector joined the second observing run on August 1, 2017. We discuss the updates that happened during this period in the GstLAL-based inspiral pipeline, which is used to detect gravitational waves from the coalescence of compact binaries both in low latency and an offline configuration. These updates include deployment of a zero-latency whitening filter to reduce the over-all latency of the pipeline by up to 32 seconds, incorporation of the Virgo data stream in the analysis, introduction of a single-detector search to analyze data from the periods when only one of the detectors is running, addition of new parameters to the likelihood ratio ranking statistic, increase in the parameter space of the search, and introduction of a template mass-dependent glitch-excision thresholding method.

I. INTRODUCTION

The observations of the binary neutron star merger GW170817 [1] and the associated short gamma ray burst GRB 170817A by the LIGO-Virgo Scientific Collaboration and the Fermi-GBM monitor and INTEGRAL satellite [2] led to an electromagnetic follow-up on an unprecedented scale [3] and marked the dawn of a new era of multi-messenger astronomy.

The gravitational-wave event was identified in low-latency by the GstLAL-based inspiral pipeline [4–6]. The low-latency detection of the gravitational-wave event made it possible for alerts to be sent out to the electromagnetic facilities in the timely fashion required for the identification of the optical transient. These gravitational-wave and electromagnetic observations support the hypothesis that neutron star mergers are progenitors of short Gamma-Ray Bursts and are followed by transient electromagnetic events known as kilonovae [3]. The joint observations allow us to set strong constraints on the fundamental physics of gravity [2], and also pro-

vide a new method of probing the cosmological parameters [7]. Besides being vital for multi-messenger astronomy, low-latency detection is also useful to freeze the detector state in case of a candidate, so as to collect sufficient data for background estimation before any configuration changes are made in the detectors.

The GstLAL-based inspiral pipeline (henceforth referred to as the GstLAL pipeline) is a matched-filtering analysis pipeline that can detect gravitational waves from compact binary mergers in near real time, and provide point estimates for binary parameters. Matched-filtering is performed by cross-correlating data against a bank of waveform templates formed using general relativity. The GstLAL pipeline is built on the GstLAL library, a collection of GStreamer [8] libraries and plug-ins that make use of the LIGO Algorithm Library, LALSuite [9]. It uses the GStreamer library to stream the gravitational strain data in real time, performs matched-filtering in the time-domain [4, 5], as opposed to the more traditional frequency-domain [10] method, and uses a time-domain rather than a frequency-domain signal consis-

tency test [10]. It also employs multi-banding and singular value decomposition on signal templates [4, 11] to reduce the number of filters and samples used in matched-filtering. A multidimensional likelihood-ratio statistic is used to rank the gravitational-wave candidates according to the properties of noise and signal [4, 12]. Instead of performing time-slides [10] for background estimation, a technique that is based on tracking the noise distributions and allows for rapid significance estimation [13] is used. Several other low-latency gravitational-wave pipelines are used by the LIGO-Virgo Collaboration to analyze their data, MBTA Online [14], GstLAL-spiir [15], PyCBC Live [16], and coherentWaveBurst [17].

In this work, we describe the GstLAL-based inspiral pipeline, which was used to detect the binary neutron star merger event GW170817 [1] and several binary black hole mergers [18–21] during the second observing run of the Advanced LIGO detectors and the first observing run of the Advanced Virgo detector (henceforth referred to as the second observing run or simply, O2). We focus on the updates made in the GstLAL pipeline since the first observing run of the LIGO detectors (O1) [4]. These updates were aimed at decreasing the latency, increasing the sensitivity, and expanding the parameter space of the pipeline.

II. OUTLINE OF PIPELINE METHODS

The GstLAL pipeline makes use of matched filtering [22–30], which is a method to extract signals from noisy data by cross-correlating the detector output with a predicted waveform signal. In case of gravitational waves from compact binary mergers, the predicted waveforms come from models that use the Post-Newtonian approach [31–35], and the effective-one-body (EOB) formalism [36, 37] to model the inspiral phase of the waveform. These models are then hybridized with the results from black hole perturbation theory and numerical relativity to obtain the full inspiral-merger-ringdown waveforms [38–44].

The compact binary coalescence waveforms depend on a number of parameters, some that are intrinsic to the source, such as the masses and spins of the binary components, and some that are extrinsic, such as the distance, inclination angle, etc. which are related to the position of the source with respect to the observer. When performing matched filtering, our goal is to maximize the matched-filter output over all these parameters. The extrinsic parameters enter in the waveform only as an overall amplitude and an overall phase.¹ The maximization

over the overall coalescence phase in the detector frame is done analytically, and that over the time of coalescence and all the intrinsic parameters is done by brute force. For maximizing over the intrinsic parameters, we create a template bank [45–48] containing a discrete set of waveforms spanning the intrinsic parameter space of our search. These discrete points in the search parameter space are chosen such that the mismatch between any signal and the best matching template from the template bank (arising from the discrete nature of our template bank) is less than a predecided tolerance set as 3% in O2. In the template bank for O2, we consider the spins of the objects to be aligned to the total angular momentum of the system [6], although efforts are ongoing to also include systems which have component spins that are not aligned to the orbital angular momentum [49–52]. The template bank that was used in O2 is described in more detail in III A. Validation of the template bank is discussed in [53].

The output of the matched-filter is the signal-to-noise ratio (SNR), the inner product of the whitened data with the whitened template. In the GstLAL pipeline, it is calculated in the time-domain [4]:

$$x_i(t) = \int_{-\infty}^{\infty} d\tau \hat{h}_i(\tau) \hat{d}(\tau + t), \quad (1)$$

where

$$\hat{d}(\tau) = \int_{-\infty}^{\infty} df \frac{\tilde{d}(f)}{\sqrt{S_n(|f|)/2}} e^{2\pi i f \tau} \quad (2)$$

is the whitened data and the whitened template $\hat{h}_i(\tau)$ is defined similarly. The subscript, i , runs over each set of template parameters in our template bank. $S_n(f)$ is the single-sided noise power spectral density (PSD). The whitener used in the pipeline is described in III B. The whitened data undergoes data quality checks and conditioning as described in III C. For each set of parameter values in the template bank, we have two real waveforms - one corresponding to the + polarization, and the other corresponding to a ‘quadrature phase-shifted +’ waveform, which is equal to the \times polarised waveform barring an overall amplitude factor.

We construct a complex signal-to-noise ratio (SNR) time series, real part of which is the SNR time series from the ‘+ polarized template ($x_i(t)$), and the complex part is the SNR time series from the ‘quadrature-phase shifted +’ polarized template ($y_i(t)$). We maximize over the unknown time and phase by maximizing over the absolute value of the complex SNR time series over time,

$$\rho = \max_t |x_i(t) + iy_i(t)|. \quad (3)$$

The GstLAL pipeline makes use of the LLOID [4, 5] algorithm, which combines singular valued decomposition

¹ This is true for systems in which spins of the component objects are aligned or anti-aligned to their orbital angular momenta. However, when the component spins are misaligned with the orbital angular momenta, the inclination and polarization angles become time dependent and lead to a non-trivial amplitude and phase modulation of the observed signals.

with near-critical sampling to construct a reduced set of orthonormal filters with far fewer samples.

Detector data often contain glitches [54] (see also III C), which can produce high peaks in the SNR time series. SNR is not sufficient to distinguish noise from transient signals in presence of non-Gaussian data. Therefore, in addition to recording the peaks in the SNR time series, the pipeline performs a signal consistency check whenever it records an SNR above a certain threshold. This is done by determining how similar the SNR time series of the data around the peak value is to the SNR time series expected from a real signal. The SNR time series is predicted by calculating the auto-correlation between the complex template waveform and itself, and scaling it by the peak complex SNR. This predicted SNR time series is equal to the SNR time series under the assumption that the signal matches the template waveform exactly in absence of any detector noise. This signal consistency test value, ξ^2 , is computed by integrating the amplitude squared of the difference between the complex SNR time series and the predicted SNR time series over a δt time window around the peak, and normalizing it appropriately [4],

$$\xi^2 = \frac{\int_{-\delta t}^{\delta t} dt |z(t) - z(0)R(t)|^2}{\int_{-\delta t}^{\delta t} dt (2 - 2|R(t)|^2)}. \quad (4)$$

Here $z(t)$ is the complex SNR time series, $z(0)$ its peak, and $R(t)$ is the auto-correlation series. Whenever the GstLAL pipeline records a peak in the SNR time-series that is greater than a preset threshold², it records the SNR and ξ^2 , the masses and the spins of the template that returned those values upon matched filtering, and the phase and the time of coalescence. Together these quantities form a ‘trigger’. These triggers are divided into smaller bins based on the spins and masses of the templates. Triggers from different detectors corresponding to the same template that are coincident in time within a time window which takes into account the maximum light travel time between detectors and statistical fluctuations in the measured event time due to detector noise are called coincident triggers. In O2, a time window of 5 ms plus light travel time between the detectors was used.

The set of triggers that did not participate in a coincidence when more than one detector was operating are used to characterize noise for their respective (mass-spin) bin. The GstLAL pipeline builds a signal and a noise model as a function of the SNR, ξ^2 , the detector(s) that participated in making the set of trigger(s), and the horizon distances of all the detectors at the time of the

event. The signal model is constructed by assuming that the signals follow their expected distribution in Gaussian noise [13]. The expectation value of SNR can be obtained by assuming uniform in volume distribution of sources, and the distribution of ξ^2 is obtained by assuming a maximum of 10% loss in SNR due to waveform mismatch [13, 55]. The triggers are then assigned a log likelihood-ratio which is the pipeline’s detection statistic based on the noise and the signal model for that particular bin. A more detailed description of the likelihood-ratio statistic can be found in III D. A Monte Carlo sampler is used to draw from these background probability density functions formed from the single detector triggers to construct a mapping from the log likelihood-ratio (\mathcal{L}) to the false-alarm-probability, $FAP(\mathcal{L})$, which is the probability of having at least one event with a log likelihood-ratio greater than or equal to \mathcal{L} under the noise hypothesis [12, 13].

For a detailed description of these methods we refer the reader to [4, 5, 11, 12]. In the following subsections we describe the developments in these methods used by the GstLAL pipeline since the configuration used during Advanced LIGO’s first observing run [4].

III. LATEST DEVELOPMENTS

A. Template bank decomposition

As described in II, template banks are discrete sets of waveforms that ensure that the SNR loss due to signal and template waveform mismatch will not be greater than a pre-specified threshold.

For Advanced LIGO’s second observing run, the parameter space of the template bank was increased from a total mass of $2M_\odot - 100M_\odot$ to $2M_\odot - 400M_\odot$. Neutron stars were assumed to have masses less than $2M_\odot$ with dimensionless spin parameters in the range (-0.05, 0.05). Black holes were assumed to have masses greater than $2M_\odot$ with dimensionless spins in the range (-0.999, 0.999). Individual masses of the systems lie in the range $1 - 399M_\odot$ with mass ratios in the range $1 - 97.989$. The templates were placed in two stages, below a total mass of $4M_\odot$ inspiral-only templates of Post Newtonian approximation [56] were laid down first by using a geometric technique [57] and then using these templates as a seed for the stochastic algorithm, while above a total mass of $4M_\odot$ full inspiral-merger-ringdown templates of effective-one-body approximation hybridized with numerical relativity and black hole perturbation theory [44] were placed using stochastic methods [48]. The original templates that were placed are shown in green, red, and blue in Fig. 1. The colors denote the range of spins covered by the template bank for those masses.

For the purpose of background estimation in the GstLAL pipeline, we first divide the entire template bank into several sub banks, known as $\hat{\theta}$ bins, that contain ‘‘similar’’ waveforms. This division is done based on their

² This SNR threshold was set to 4 for the Hanford and Livingston detectors in O2. For the online analysis, the minimum trigger SNR in Virgo was not determined by an explicit threshold, but instead by a restriction to record at most 1 trigger per second in a given template. For the offline analysis, the minimum SNR threshold was set to 3.5 for Virgo.




FIG. 1. A visual representation of the bank used in the second observing run in the component mass space. Each point here represents a template in the bank. All regions shown in the figure are discrete points, they continuous regions being the highly dense regions. The green points are the binary neutron star templates, blue are the binary black hole templates, and red are the neutron star-black hole templates. Magenta and black points represent additional templates that were added to aid in the background estimation for the scarcely populated, high-mass region of the template bank.

intrinsic parameters. We assume that the templates that belong in the same sub bank respond similarly to noise and noise statistics are collected for all templates in a bin as a whole. The sub banks are typically formed by combining a certain number of split banks”, which are formed for performing the SVD [4]. In O2, we originally grouped together 2 split banks containing 500 templates each to form the $\bar{\theta}$ bins. We want the binning method to, a) group together similar templates to use the LLOID method [4, 5] for computationally-efficient time-domain searches, and b) group together templates with similar noise backgrounds for appropriate FAR estimates. Prior to O2, the pipeline used two composite parameters, which are a combination of the four intrinsic parameters, in order to group the waveforms into split banks - the chirp mass and the effective spin. The chirp mass and the effective spin parameter are the leading order terms that describe the phase evolution of the inspiral part of the waveform according to the Post Newtonian expansion.

The chirp mass \mathcal{M} is defined as:

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

The effective spin parameter is given by:

$$\chi_{\text{eff}} \equiv \frac{m_1 \chi_1 + m_2 \chi_2}{m_1 + m_2}, \quad (6)$$

where m_1, m_2 are the masses and $\chi_i = \vec{S}_i \cdot \hat{L} / m_i^2$ are the dimensionless spin parameters of each component in the binary. \vec{S}_i are the component-spin vectors and \hat{L} is the orbital angular momentum unit vector.

While analyzing the data from O2, it was found that for some analyses, there was an excess in the closed box result compared to the background predicted by the GstLAL pipeline. The closed box result is prepared by looking at time-shifted coincidences between the two detectors, and is expected to closely follow the background model when the pipeline is only treating coincident events as candidates. After investigations, it was found that the template bank, together with the grouping scheme of templates that was used in the first observing run of the Advanced LIGO detectors [4] for collecting noise statistics was inadequate to construct an accurate background model for the high mass region that was added to the parameter space for O2. As a result, a different grouping scheme was introduced in O2, and the template bank was also modified. In Fig. 1 we see that the density of the templates decrease with the increase in the masses of the binaries. This is a typical feature seen in all template banks. The waveform of a system with smaller masses is longer in the frequency band Advanced LIGO and Virgo detectors are most sensitive in. Therefore even a small change in masses can lead to a big mismatch between two waveforms in this region since there are more cycles in band for which the match has to be performed. This means that we need more waveforms to cover the lower mass region of the parameter space.

The extremely low density of the templates in the high mass region of the bank, which was added for O2 caused templates vastly different in \mathcal{M} , but similar in χ_{eff} to be grouped together in the low template density region, since the number of templates in each bin was required to be the same. This in turn led to sub-optimal estimation of noise properties for these bins, which led to inaccurate significance estimation of noise events. In Fig. 2, these are the highest \mathcal{M} groups in each χ_{eff} bin.

We can also see this in the background SNR- ξ^2 PDFs for these “bad” bins that the pipeline uses to assign the likelihood-ratio statistic, see Fig. 3. The features marked in white in the PDF happen when there are template present in a bin whose properties are not well represented by the rest of the templates in the bin.

To solve this issue described above, the grouping scheme was changed for the high mass templates with total mass $> 80M_{\odot}$. Instead of using \mathcal{M} and χ_{eff} , we now use template duration in the LIGO sensitivity band for grouping these high mass templates. \mathcal{M} and χ_{eff} are the leading order parameters that describe the inspiral part of the waveforms, but not the merger and ringdown. For high mass systems, a significant amount of power in the LIGO band comes from merger and ringdown phases of

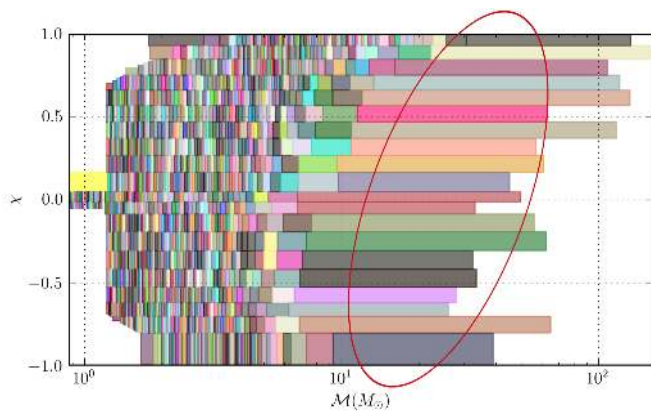


FIG. 2. A visual representation of the background bins according to the old binning scheme of the templates in $\mathcal{M} - \chi_{\text{eff}}$ space. The bins circled were identified as the ones that contained the templates whose noise properties were not being described by the noise properties of the bins.

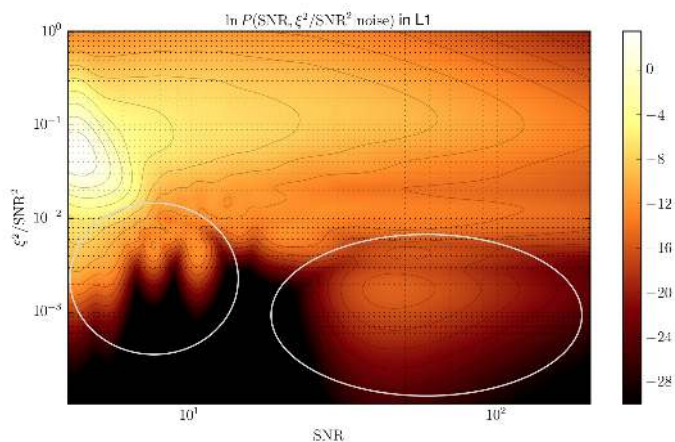


FIG. 3. SNR- ξ^2 PDF for one of the “bad” bins. We draw the reader’s attention to some of the island features inside the white circles. These occur when the pipeline doesn’t have enough data about the behavior of some of the templates in the bin, because respond to noise differently from the rest of the templates in the bin.

the coalescence, therefore \mathcal{M} and χ_{eff} are no longer the best parameters to group the templates. The number of templates in these high-mass bins was also reduced to 200, 400, or 800, instead of 1000 everywhere else, to account for the sparse density of templates in this region. Hence only those templates that have the similar noise properties are now grouped together. 15,665 extra templates were also added to the bank above a total mass of 80, by increasing the convergence threshold from 0.97 to 0.98 and by adding a grid of templates. This was done to increase the template density in that region. This brought the total number of templates to 677000. The templates that were added by increasing the convergence threshold of the stochastic placement process are shown in magenta and those that were added as a grid are shown

in black in Fig. 1. Refer [53] for a detailed description of the construction of the bank.

B. Zero Latency Whitener

During the first observing run of Advanced LIGO, the pipeline detected a binary black hole merger event, GW151226, with a latency of 70 seconds [58], which means that the signal was detected a mere 70 seconds after it arrived on Earth. This was the first time a gravitational-wave event had been identified by a matched-filtering pipeline in low-latency and was a huge success for the pipeline, which aims to send prompt alerts to electromagnetic observatories in order to facilitate studies correlating astrophysical phenomena. But from the recent joint gravitational-wave and gamma-ray detection of GW170817, we now know that the time delay between gravitational-wave emission and the onset of the following SGRB is approximately 2 seconds, motivating achieving alert latencies below 2 seconds in order to capture the earliest associated electromagnetic phenomena in other bands.

The original whitening filter employed in the GstLAL pipeline contributes to one of the bottleneck processes in the pipeline’s latency, adding up to 32 seconds to the total latency time [4]. The zero latency whitener was introduced in spirit of reducing the latency of the pipeline [59].

As described in II, a matched filter is the inner product between a template waveform and the strain data. However, since the strain data are strongly colored by the frequency-dependent noise of the detector, we ‘whiten’ the inner product by weighting both the data and the template waveform by a factor of $1/\sqrt{S_n(f)}$ each, where $S_n(f)$ is the single-sided power spectral density of the detector noise. The method of a frequency-domain whitening filter described in [4] has discrete Fourier transforms and window functions applied to 32-second blocks of input data, with the PSD being updated every 16 seconds. Since a 32-second block is processed every 16 seconds, this filter has a latency between 16 to 32 seconds, which is more than half of the pipeline’s total latency.

To reduce the latency, we need to whiten in the time-domain. To this end, a Finite Impulse Response (FIR) filter-based algorithm to the frequency-domain whitening is introduced. The square root of inverse PSD of a given strain data is used to construct the FIR of a linear phase filter. This filter still requires 16 seconds of data from the future for its evaluation [60]. It is not possible to further reduce the latency of this filter without changing the whitening transformation. Therefore an approximation to the original filter is introduced [60], which derives a minimum-phase approximation of the desired filter [61]. The matched filter output is insensitive to the phase response of the whitening filter. The minimum-phase whitening filter accurately approximates the amplitude response of the FIR-based filter, introducing most errors only in the phase response. This filter does not

use any information from future samples for its evaluation, and is therefore called a “zero-latency whitening filter”. A new windowing process has also been implemented for the new whitening method. The PSD transition is now allowed to occur continuously, and the resulting filter is a linear combination of the newest and next newest filters during their transition. This is recursively applied to the zero-latency algorithm when a new whitening filter becomes available. It is shown in [60] that the causal FIR filter approximation successfully whitens the data, producing zero mean, unit variance, white Gaussian noise. Matched-filter outputs (SNR and ξ^2 , II) produced using the zero-latency whitener are compared to those produced using the frequency-domain whitener for both noise and simulated signal triggers, and we see a good agreement for both. The time-domain zero-latency filter was used by the GstLAL low-latency analysis of O2, and was a significant contributor in bringing down the latency of the pipeline as compared to O1 4'. The details of the filter, and the consistency checks with the old filter are described in [60].

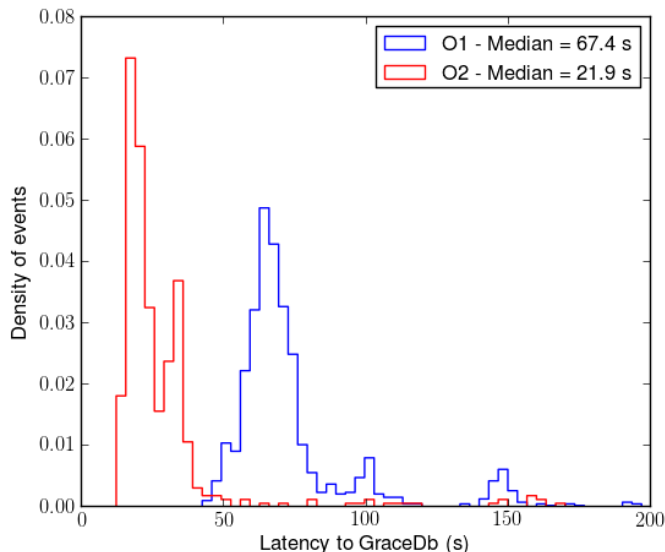


FIG. 4. Comparison of the latency of detections by the GstLAL online pipeline in O1 (blue) vs. in O2 (red). We can see that the latency of the pipeline has been reduced by ≈ 40 s. The application of the zero-latency whitener played a significant contribution in decreasing the latency of the pipeline.

C. Data Conditioning

The output of the matched filter is the SNR (II), which is the optimal detection statistic under the assumption that noise is stationary and Gaussian. For offline analyses, where the GstLAL pipeline processes archival gravitational-wave data, we use data quality vetoes to flag poor data [54]. However, such information is not available for online analyses. Gating on the

whitened strain data, whitened $h(t)$, is one of the techniques adopted by the pipeline to eliminate short transient instrumental noise fluctuations. These fluctuations, also known as glitches, can cause unreasonably high values of SNRs in the data, mimicking gravitational-wave signals, and causing false alarm triggers.

In presence of glitches, the whitened $h(t)$ may have values higher than the expected values from the coalescences of binary systems that the pipeline is aiming to detect. By construction, whitened $h(t)$ should have a unit variance. Whenever whitened $h(t)$ is momentarily greater than a threshold value, set as some multiple of the standard deviation σ of $h(t)$, we gate that piece of data by setting the samples around the peak with a window of 0.25 s on each side to zero [4].

The amplitude of a signal increases with the chirp mass \mathcal{M} (III A) of a binary system. Therefore we set the threshold based on the highest masses we are sensitive to, such that it is higher than the whitened strain amplitude we expect from such systems. We want the threshold to be such that it removes the maximum number of glitches from our data without gating out real signals. During the first observing run, the threshold value was set to 50σ . At lower threshold values, it was seen that the pipeline started to gate some of the high mass simulated BBH signals that were injected in the data.

As described in III A, in the second observing run, the parameter space of our search was increased from a maximum total mass of $100M_{\odot}$ to $400M_{\odot}$. In order to avoid gating the highest mass signals, we would have to increase our value of the gate threshold. This would cause an increase in the number of glitches that pass through without being gated, and therefore an increase in the number of false alarm triggers. Hence a linear gating scheme, in which the gate threshold value is a linear function of chirp mass instead of a constant for all masses, was introduced. The pipeline computes an appropriate threshold value according to the linear scale provided by the user and the highest chirp mass template in a sub-bank. Fig. 5 shows that linear gating scheme helps in removing more glitches as compared to the constant gating scheme, while still recovering all the simulated signals injected in the data.

D. Likelihood-Ratio Statistic

Once all the candidate events are identified by the pipeline, each of them is assigned a likelihood-ratio statistic in order to determine its significance. The likelihood ratio as defined in [4, 12] was used in the analysis of data from the first observing run of Advanced LIGO during which only the Hanford and Livingston detectors were operating. It is the ratio of probability of certain observables given the signal hypothesis versus the noise hypothesis. Included observables were terms to account for the set of detectors involved in the coincidence $\{H1, L1\}$, the detector horizon distances (for a $1.4M_{\odot}$ - $1.4M_{\odot}$ sys-




FIG. 5. Comparison results between a run with constant gating scheme (left) and linear gating scheme (right). The yellow stars highlight the additional background glitches present in the constant gating run that have been removed in the run with the linear gating scheme applied.

tem) at the time of coincidence $\{D_{H1}, D_{L1}\}$, the SNRs for each trigger $\{\rho_{H1}, \rho_{L1}\}$, and the ξ^2 -signal-based-veto values for each trigger $\{\xi_{H1}^2, \xi_{L1}^2\}$. For Advanced LIGO's second observing run, the likelihood-ratio statistic was modified to allow the pipeline to rank single detector events when only one detector is operational and to include additional parameters when both advanced LIGO facilities are operational. The inclusion of additional parameters was done by two methods during O2. The first method which we will describe shortly in III D 2 was used to analyze the data in the beginning of O2 [1, 18–20], and the second method described in [62] was used to analyze final calibrated version of O1 and O2 [21]. The binning for SNR and ξ^2 histograms, used to compute the (ρ, ξ^2) PDFs, was also changed O2. The $\tan^{-1} \ln$ binning described in [12] is still used, to collect SNR – ξ^2 but now $x_{lo} = 2.6$, $x_{hi} = 26$, and $n = 300$; the values used in O1 were $x_{lo} = 3.6$, $x_{hi} = 70$, and $n = 600$. The binning was changed so that the shape of the density kernel estimation matches the natural shape of the SNR – ξ^2 PDFs it is being used to estimate.

1. Single Detector Events

In Advanced LIGO's first observing run, the pipeline could only identify gravitational-wave events when both

the detectors were operating. Gravitational-wave candidates were formed by demanding coincidence between the two detectors. This meant that we were blind to signals occurring during single-detector time (defined as time when only one detector is operational). For O2, the online analysis also looked at the non-coincident candidates, and these were also assigned a log likelihood ratio statistic.

Non-coincident triggers found during single-detector time (defined as time when only one detector is operational) are now excluded from informing the background model, since these could potentially be loud signals that were found as non-coincident triggers in absence of data from multiple detectors. Since we use an SNR threshold of 4 for triggers, there are too many non-coincident triggers to write to disk. Therefore the non-coincident triggers are first assigned a preliminary log likelihood-ratio, and only those that have $\log \mathcal{L}_{prelim} > 2$ are considered as gravitational-wave candidates. The likelihood-ratio defined in [4, 12] is still valid for single detector events,

$$\begin{aligned} \mathcal{L}(\{D_{\text{IFO}_{\text{net}}}\}, \{\text{IFO}\}, \rho_{\text{IFO}}, \xi_{\text{IFO}}^2, \bar{\theta}) &= \mathcal{L}(\{D_{\text{IFO}_{\text{net}}}\}, \{\text{IFO}\}, \rho_{\text{IFO}}, \xi_{\text{IFO}}^2 | \bar{\theta}) \mathcal{L}(\bar{\theta}) \\ &= \frac{P(\{D_{\text{IFO}_{\text{net}}}\}, \{\text{IFO}\}, \rho_{\text{IFO}}, \xi_{\text{IFO}}^2 | \bar{\theta}, \text{signal})}{P(\{D_{\text{IFO}_{\text{net}}}\}, \{\text{IFO}\}, \rho_{\text{IFO}}, \xi_{\text{IFO}}^2 | \bar{\theta}, \text{noise})} \mathcal{L}(\bar{\theta}). \end{aligned} \quad (7)$$

Here, $\{D_{\text{IFO}_{\text{net}}}\}$ is the set of horizon distances for all

instruments in the network at the time the event is ob-

served, $\{\text{IFO}\}$ is the detector that produced the non-coincident trigger, ρ_{IFO} , and ξ_{IFO}^2 are the SNR and ξ^2 values of the trigger. The numerator and denominator

$$\mathcal{L}(\{D_{\text{IFOnet}}\}, \{\text{IFO}\}, \rho_{\text{IFO}}, \xi_{\text{IFO}}^2, \bar{\theta}) = \dots \times \frac{P(\{\text{IFO}\} | \{D_{\text{IFOnet}}\}, \bar{\theta}, \text{signal})}{P(\{\text{IFO}\} | \bar{\theta}, \text{noise})} \times \dots \quad (8)$$

The probability that a signal yields a trigger above-threshold in only one of the detectors depends on the horizon distances of all the detectors operating at the time, and the duty cycles of all the detectors in the network. The probability that noise yields an above-threshold trigger in one of the detectors is computed from the trigger rates, coincidence window size, and the duty cycles of all the detectors in the network. If more than one detector is operating, and the signal is only seen above-threshold in one of the detectors then that becomes a constraint on the SNR distribution. The single-variable SNR distribution is computed using a Monte Carlo generation of samples described in [12]. In the case where only one detector is operating, the SNR distribution reduces to $P(\rho) \propto \rho^{-4}$.

For assigning the false-alarm rate (FAR), in case of coincident triggers, the pipeline calculates the probability of accidental coincidence by drawing events from each of the single-detector background PDFs. This allows us to measure the likelihood-ratio distribution under the noise hypothesis to values of likelihood-ratio higher than we have actually observed in the experiment.

of the fraction in Eq. 7 are factored in the same way as described in [4, 12]. In particular the factorization leads to a form,

But the non-coincident triggers cannot benefit from the boost due to low probability of an accidental coincidence and therefore the FAR cannot be measured less than $1/(\text{time of experiment})$. GW170817 is an example of the success of including single detector candidate events in the pipeline. It was first identified as a single-detector Hanford event, because of the presence of a glitch in Livingston at the time of the event.

2. Inclusion of phase and time delay terms in the likelihood statistic

Two additional parameters were added to the likelihood-ratio statistic for the case where both the Hanford (H1) and Livingston (L1) advanced LIGO detectors are operational: Δt , the difference in end times between the H1 and L1 triggers, and $\Delta\phi$, the difference in coalescence phase between the H1 and L1 triggers. We require $\Delta\phi \in [-\pi, \pi]$ and compute the modulus $\Delta\phi \pmod{2\pi}$ to enforce a cyclic distribution.

This new likelihood-ratio statistic is defined as follows

$$\mathcal{L} = \frac{P(\{D_{\text{H1}}, D_{\text{L1}}\}, \{\text{H1}, \text{L1}\}, \rho_{\text{H1}}, \rho_{\text{L1}}, \xi_{\text{H1}}^2, \xi_{\text{L1}}^2, \Delta\phi, \Delta t | \text{signal})}{P(\{D_{\text{H1}}, D_{\text{L1}}\}, \{\text{H1}, \text{L1}\}, \rho_{\text{H1}}, \rho_{\text{L1}}, \xi_{\text{H1}}^2, \xi_{\text{L1}}^2, \Delta\phi, \Delta t | \text{noise})} \quad (9)$$

This statistic only supports ranking coincidences found with the H1L1 network. Future work for the third observing run will add support for the H1L1V1 network with a goal towards a generalized N-detector network statistic [62]. Additionally, we make several assumptions when factoring the dependencies of the probability density functions in Eq. 9. We assume that the noise distributions of Δt and $\Delta\phi$ are independent of each other. We assume that the ξ^2 statistic is dominated by instrumental noise, thus the ξ^2 term reduces to its previous form given in [4]. We expect that the signal (but not the noise) distributions for Δt depend on trigger SNRs as well as on detector sensitivities. The SNR ratio from the two detectors for a signal depends on the position of the source with respect to the detectors and the inherent sensitivities of the detectors. Δt for a signal depends only on the position of the source and the location of

the observatories. Thus we model the Δt distributions as a function of a ratio of SNRs normalized by horizon distances, to factor out the inherent sensitivities of the detectors, so this term only depends on the position of the source with respect to the detectors. Furthermore, we define it such that it is always smaller than 1,

$$\rho_{\text{ratio}} = \min\left(\frac{\rho_{\text{H1}}/D_{\text{H1}}}{\rho_{\text{L1}}/D_{\text{L1}}}, \frac{\rho_{\text{L1}}/D_{\text{L1}}}{\rho_{\text{H1}}/D_{\text{H1}}}\right) \quad (10)$$

On the other hand, we do not consider dependence on the detector sensitivities when modeling the $\Delta\phi$ distributions. We only consider dependence of $\Delta\phi$ on Δt and network SNR, defined as

$$\rho_{\text{network}} = \sqrt{\rho_1^2 + \rho_2^2} \quad (11)$$

With these assumptions, the factor describing the de-

pendence of the likelihood-ratio in terms of ρ , Δt , and $\Delta\Phi$ can be written as,

$$\begin{aligned} \mathcal{L} &\propto \frac{P(\Delta t|\{\rho\}, \{\text{Horizon}\}, \text{signal}) P(\Delta\phi|\Delta t, \{\rho\}, \{\text{Horizon}\}, \text{signal})}{P(\Delta t|\text{noise}) P(\Delta\phi|\text{noise})} \\ &\approx \frac{P(\Delta t|\{\rho_{\text{ratio}}\}, \text{signal}) P(\Delta\phi|\Delta t, \{\rho_{\text{network}}\}, \text{signal})}{P(\Delta t|\text{noise}) P(\Delta\phi|\text{noise})} \end{aligned} \quad (12)$$




FIG. 6. Distribution of $P(\Delta t|\{\rho_{\text{ratio}}\}, \text{signal})$

In order to construct distributions for $P(\Delta t|\{\rho_{\text{ratio}}\}, \text{signal})$ and $P(\Delta\phi|\Delta t, \{\rho_{\text{network}}\}, \text{signal})$, we performed an injection (simulated signal injected in the data to test the performance of the pipeline) campaign in the data from the first observing run covering the parameter space of binary neutron star and binary black hole systems. We discarded any injection recovered with $|\Delta t| > 0.01$ s, assuming that this regime is dominated by accidental coincidences. The Δt histograms were smoothed with a gaussian kernel such that they taper to zero above $|\Delta t| = 0.01$ s. Smoothed, normalized Δt histograms were modeled as a function of ρ_{ratio} using Chebyshev polynomials. Fig. 6 shows the logarithm of the PDF for $P(\Delta t|\{\rho_{\text{ratio}}\}, \text{signal})$ as defined over a range of Δt and ρ_{ratio} . Normalized histograms of $\Delta\phi$ were modeled using von Mises distributions which is a continuous probability distribution on a circle since $\Delta\phi$ is cyclic, plus a uniform noise background as a function of Δt and ρ_{network} . In Fig. 7, we see the logarithm of the PDF for $P(\Delta\phi|\Delta t, \{\rho_{\text{network}} = 14\})$ as defined over a range of Δt and $\Delta\phi$.

In order to construct the distributions for $P(\Delta t|\text{noise})$ and $P(\Delta\phi|\text{noise})$, we simply assume Δt and $\Delta\phi$ are uniformly distributed for noise triggers that form false coincidences.




FIG. 7. Distribution of $P(\Delta\phi|\Delta t, \{\rho_{\text{network}}\})$ where we have set $\rho_{\text{H1}} = 10$ and $\rho_{\text{L1}} = 10$.

3. Computing joint SNR PDF for different horizon distance ratios

The horizon distance D_h is the effective distance at which a binary system is observed with a nominal SNR of 8 [10]. This means that horizon distance is a measure of a detector's sensitivity to a particular system. The horizon distances included in the ranking statistic of the pipeline are computed for a $1.4M_{\odot} - 1.4M_{\odot}$ binary neutron star system, and their fluctuations reflect fluctuations in the noise spectrum.

One of the factors in the numerator of the likelihood-ratio ranking statistic is the joint SNR PDF given a set of horizon distances of all the detectors at the time of the event, and the set of detectors that observed the event with an SNR above the threshold. [12]. For candidates that arise from genuine signals, the joint SNR PDF depends only on the ratios of the horizon distances. In the first observing run, the pipeline was limited to two detectors and assumed a fixed joint SNR PDF, corresponding to equal horizon distances. This assumption implies that the fractional change in horizon distance is approximately independent of the mass of the system being observed, which is not valid because changes to the

noise spectrum do not rescale equally over the entire frequency range. In the second observing run, this assumption was relaxed by pre-computing joint SNR PDFs for a collection of discrete horizon distance ratios.

E. Software injections

The sensitivity of the pipeline is measured by injecting simulated gravitational-wave signals into the data and measuring the fraction that is recovered, as described in [4]. To reduce the computational cost of searching for injections, the option to search only a subset of the template bank was added to the pipeline. The subset is specified by providing a lower and upper bound on chirp mass (III A). Low-mass injections, where the gravitational wave is dominated by the inspiral stage, are typically recovered with more accurate chirp masses than high mass injections. For this reason, the bounds are computed as a function of injection chirp mass, \mathcal{M}_{inj} . The function was written based on visual inspection of chirp mass recovery plots from the searches in the first observing run, and subsequent tests confirmed a loss of less than 1% of the total found injections.

$$\mathcal{M}_{\text{lower bound}} = \begin{cases} 0.65\mathcal{M}_{\text{inj}}, & \mathcal{M}_{\text{inj}} < 10 \\ 0.5\mathcal{M}_{\text{inj}}, & 10 \leq \mathcal{M}_{\text{inj}} < 20 \\ 0.5\mathcal{M}_{\text{inj}}, & \mathcal{M}_{\text{inj}} \geq 20 \end{cases} \quad (13a)$$

$$\mathcal{M}_{\text{upper bound}} = \begin{cases} 1.35\mathcal{M}_{\text{inj}}, & \mathcal{M}_{\text{inj}} < 10 \\ 1.5\mathcal{M}_{\text{inj}}, & 10 \leq \mathcal{M}_{\text{inj}} < 20 \\ 2\mathcal{M}_{\text{inj}}, & \mathcal{M}_{\text{inj}} \geq 20 \end{cases} \quad (13b)$$

F. Introducing Virgo

The Advanced Virgo [63] detector joined the second observing run of the Advanced LIGO detectors on August 1st, 2017. It operated at a lower sensitivity relative to the Advanced LIGO detectors for the observing run reported here. Initially only the online search filtered over the Virgo data stream. Due to this, there was no minimum SNR threshold for a trigger in Virgo. Instead, it was restricted to record at most one trigger per second per template [64].

In the first observing run, gravitational-wave candidates were formed by demanding coincidence (both in time and template) between the LIGO Hanford and LIGO Livingston triggers [4], [II]. In the second observing run, this was generalized so that a candidate can be formed by an arbitrary number of detectors. For a network of detectors, we can define different types of coincidences, based on the number of detectors participating in the coincidence. However, due to the lower sensitivity of Virgo in the second observing run and the incompatibility of the likelihood-ratio statistic in the beginning of

O2, if Virgo participated in a coincidence with either one of the two detectors, the set of triggers was still considered as a single-detector, non-coincident event. In other words *HV* and *LV* doubles were treated as *H* and *L* singles respectively. And if Virgo participated in a triple event (*HLV* triple), it was still treated as an *HL* double. Nonetheless a network of three detectors improves the sky localization of the source. GW170814 was the first gravitational-wave event that had a significant SNR in Virgo. Including Virgo in the analysis helps in reducing the area of the 90% credible region from 1160 deg² when using only the two LIGO detectors to 60 deg² when using all three detectors for the binary black hole event GW170814 [20].

For the final offline reanalysis of the O1 and O2 data calibrated with the final versions [21], the Advanced Virgo data stream was treated in the same manner the same as the Advanced LIGO detectors and also used to inform the likelihood-ratio statistic. To this end, a new method for calculating the likelihood-ratios was introduced which also uses the information from Virgo. This method is fast, computationally efficient and can also be adapted to addition of new detectors in the network. For the detailed description of this method, we refer the readers to [62]. In the offline reanalysis, GW170818 was detected as a triple coincident event with an SNR of 4.2 in Virgo, 4.1 in Hanford and 9.7 in Livingston [21].

IV. CONCLUSION

The GstLAL pipeline is a stream-based matched-filtering pipeline that has detected gravitational waves from several compact binary mergers in near real time since Advanced LIGO's first observing run. In this work, we have described the advancements made in the techniques of the pipeline to reduce the latency, increase the parameter space, analyze single-detector time, incorporate the Advanced Virgo detector's data stream in the analysis, add new parameters to the likelihood-ratio statistic, and introduce a template mass-dependent glitch-excision thresholding method. These methods were successfully deployed in Advanced LIGO's second and Advanced Virgo's first observing run.

V. FUTURE WORK

Active development tasks aim to maximize the science from future observations, including the nature of prompt electromagnetic emissions associated with binary neutron star and neutron star-black hole mergers. To this end, we are working on further reducing the latency of the pipeline and providing early warning alerts for binary neutron stars and neutron star-black hole coalescences 10 s to a minute before merger, with an approximate sky location so that the electromagnetic facilities can start the process of setting up their observations in advance of

the merger making it possible to capture the earliest possible light with narrow field instruments [5]. We are also working to add a factor to the likelihood-ratio that takes in as an input a source population mass model [65], which should make the pipeline more robust if we believe that these population models are correct. With an increase in the number of compact binary detections made by the ground-based interferometers, our population models are expected to improve [66], and including these as a factor in the likelihood-ratio will in turn aid in an increase in the number of detections.

VI. ACKNOWLEDGEMENTS

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation (NSF) and operates under cooperative agreement PHY-0757058. We would like to thank Tito Dal Canton for

providing helpful comments and suggestions. We would also like to thank Graham Woan for his help in the review of the O2 pipeline. SS was supported in part by the LIGO Laboratory and in part by the Eberly Research Funds of Penn State, The Pennsylvania State University, University Park, PA 16802, USA. DM, JC, PB, and DC were supported by the the NSF grant PHY-1607585. SC was supported by the research programme of the Netherlands Organisation for Scientific Research (NWO). HF was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). CH was supported in part by the NSF through PHY-1454389. Funding for this project was provided by the Charles E. Kaufman Foundation of The Pittsburgh Foundation. TGFL was partially supported by a grant from the Research Grants Council of the Hong Kong (Project No. CUHK 14310816 and CUHK 24304317) and the Direct Grant for Research from the Research Committee of the Chinese University of Hong Kong. AG acknowledges support from NSF Grants AST-1716394 and AST-1708146. MW was supported by NSF grant PHY-1607178. This paper carries LIGO Document Number ligo-p1700411.

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017).
- [2] B. P. Abbott *et al.* (Virgo, Fermi-GBM, INTEGRAL, LIGO Scientific), *Astrophys. J.* **848**, L13 (2017), arXiv:1710.05834 [astro-ph.HE].
- [3] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.*, *The Astrophysical Journal Letters* **848**, L12 (2017).
- [4] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. Creighton, R. Everett, *et al.*, *Physical Review D* **95**, 042001 (2017).
- [5] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Keppel, L. Liao, *et al.*, *The Astrophysical Journal* **748**, 136 (2012).
- [6] S. Privitera, S. R. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, *Physical Review D* **89**, 024003 (2014).
- [7] B. P. Abbott *et al.* (The LIGO Scientific and The Virgo, The Las Cumbres Observatory, The 1M2H, The DLT40, The MASTER, The VINROUGE, The Dark Energy Camera GW-EM and the DES), *Nature* (2017), 10.1038/nature24471, arXiv:1710.05835 [astro-ph.CO].
- [8] “Gstreamer,” <https://gstreamer.freedesktop.org>.
- [9] “Lalsuite,” <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>, accessed: 2015-07-01.
- [10] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, *Physical Review D* **85**, 122006 (2012).
- [11] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, *Physical Review D* **82**, 044025 (2010).
- [12] K. Cannon, C. Hanna, and J. Peoples, arXiv preprint arXiv:1504.04632 (2015).
- [13] K. Cannon, C. Hanna, and D. Keppel, *Physical Review D* **88**, 024025 (2013).
- [14] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, *Classical and Quantum Gravity* **33**, 175012 (2016).
- [15] X. Guo, Q. Chu, S. K. Chung, Z. Du, and L. Wen, arXiv e-prints, arXiv:1702.02256 (2017), arXiv:1702.02256 [astro-ph.IM].
- [16] A. H. Nitz, T. Dal Canton, D. Davis, and S. Reyes, *Phys. Rev. D* **98**, 024050 (2018), arXiv:1805.11174 [gr-qc].
- [17] S. Klimentenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. F. Da Silva, and G. Mitselmakher, *Phys. Rev. D* **93**, 042004 (2016).
- [18] B. P. Abbott *et al.* (VIRGO, LIGO Scientific), *Phys. Rev. Lett.* **118**, 221101 (2017), [Erratum: *Phys. Rev. Lett.* 121, no.12, 129901 (2018)], arXiv:1706.01812 [gr-qc].
- [19] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.*, *The Astrophysical Journal Letters* **851**, L35 (2017).
- [20] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 141101 (2017).
- [21] (2018), arXiv:1811.12907 [astro-ph.HE].
- [22] L. Waistein and V. Zubakov, *European Science Notes* **1** (1970).
- [23] K. S. Thorne, *Three hundred years of gravitation* **1**, 330 (1987).

- [24] B. S. Sathyaprakash and S. V. Dhurandhar, Phys. Rev. D **44**, 3819 (1991).
- [25] C. Cutler *et al.*, Phys. Rev. Lett. **70**, 2984 (1993), arXiv:9208005 [astro-ph].
- [26] L. S. Finn, Physical Review D **46**, 5236 (1992).
- [27] L. S. Finn and D. F. Chernoff, Physical Review D **47**, 2198 (1993).
- [28] S. V. Dhurandhar and B. S. Sathyaprakash, Phys. Rev. D **49**, 1707 (1994).
- [29] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996), arXiv:gr-qc/9508011.
- [30] E. E. Flanagan and S. A. Hughes, Phys. Rev. D **57**, 4535 (1998), arXiv:9701039 [gr-qc].
- [31] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman, Physical Review Letters **74**, 3515 (1995).
- [32] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, Classical and Quantum Gravity **13**, 575 (1996).
- [33] L. Blanchet, G. Faye, B. R. Iyer, and B. Joguet, Physical Review D **65**, 061501 (2002).
- [34] L. Blanchet, T. Damour, G. Esposito-Farese, and B. R. Iyer, Physical review letters **93**, 091101 (2004).
- [35] L. Blanchet, Living Reviews in Relativity **9**, 4 (2006).
- [36] A. Buonanno and T. Damour, Physical Review D **59**, 084006 (1999).
- [37] A. Buonanno and T. Damour, Phys. Rev. D **62**, 064015 (2000).
- [38] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, Phys. Rev. Lett. **106**, 241101 (2011).
- [39] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, Phys. Rev. D **82**, 064016 (2010).
- [40] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044006 (2016).
- [41] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044007 (2016).
- [42] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, Phys. Rev. D **86**, 024011 (2012).
- [43] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu, Phys. Rev. D **89**, 061502 (2014).
- [44] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond, T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, M. A. Scheel, and B. Szilágyi, Phys. Rev. D **95**, 044028 (2017).
- [45] B. J. Owen, Physical Review D **53**, 6749 (1996).
- [46] B. J. Owen and B. Sathyaprakash, Physical Review D **60**, 022002 (1999).
- [47] I. W. Harry, B. Allen, and B. Sathyaprakash, Physical Review D **80**, 104014 (2009).
- [48] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. Weinstein, Physical Review D **89**, 084041 (2014).
- [49] P. Schmidt, M. Hannam, and S. Husa, Physical Review D **86**, 104063 (2012).
- [50] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, *et al.*, arXiv preprint arXiv:1308.3271 (2013).
- [51] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Phys. Rev. D **89**, 084006 (2014).
- [52] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, Phys. Rev. D **94**, 024012 (2016), arXiv:1603.02444 [gr-qc].
- [53] D. Mukherjee, S. Caudill, R. Magee, C. Messick, S. Privitera, S. Sachdev, K. Blackburn, P. Brady, P. Brockill, K. Cannon, S. J. Chamberlin, D. Chatterjee, J. D. E. Creighton, H. Fong, P. Godwin, C. Hanna, S. Kapadia, R. N. Lang, T. G. F. Li, R. K. L. Lo, D. Meacher, A. Pace, L. Sadeghian, L. Tsukada, L. Wade, M. Wade, A. Weinstein, and L. Xiao, arXiv e-prints, arXiv:1812.05121 (2018), arXiv:1812.05121 [astro-ph.IM].
- [54] B. P. Abbott *et al.* (LIGO Scientific, Virgo), (2016), arXiv:1602.03844 [gr-qc].
- [55] B. Allen, Physical Review D **71**, 062001 (2005).
- [56] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Physical Review D **80**, 084043 (2009).
- [57] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Physical Review D **86**, 084017 (2012).
- [58] B. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, Physical Review Letters **116**, 241103 (2016).
- [59] L. Tsukada, K. Cannon, C. Hanna, D. Keppel, D. Meacher, and C. Messick, ArXiv e-prints (2017), arXiv:1708.04125 [astro-ph.IM].
- [60] L. Tsukada, K. Cannon, C. Hanna, D. Keppel, D. Meacher, and C. Messick, Phys. Rev. D **97**, 103009 (2018).
- [61] N. Damera-Venkata, B. L. Evans, and S. R. McCaslin, IEEE Transactions on Signal Processing **48**, 1491 (2000).
- [62] C. Hanna *et al.*, (2019), arXiv:1901.02227 [gr-qc].
- [63] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, *et al.*, Classical and Quantum Gravity **32**, 024001 (2015).
- [64] B. P. Abbott *et al.* (LIGO Scientific, Virgo), (2019), arXiv:1901.03310 [astro-ph.HE].
- [65] H. K. Y. Fong, *From simulation to signals: Analyzing gravitational waves from compact binary coalescences*, Ph.D. thesis, University of Toronto (2018).
- [66] B. P. Abbott *et al.* (LIGO Scientific, Virgo), (2018), arXiv:1811.12940 [astro-ph.HE].