

Extending the PyCBC search for gravitational waves from compact binary mergers to a global network

Gareth S. Davies¹, Thomas Dent¹, Márton Tápai², Ian Harry³, Connor McIsaac³, and Alexander H. Nitz^{4,5}

¹ *Instituto Galego de Física de Altas Enerxías, Universidade de Santiago de Compostela, 15782 Santiago de Compostela, Galicia, Spain*

² *Department of Experimental Physics, University of Szeged, Szeged, 6720 Dóm tér 9., Hungary*

³ *University of Portsmouth, Portsmouth, PO1 3FX, United Kingdom*

⁴ *Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, D-30167 Hannover, Germany and*

⁵ *Leibniz Universität Hannover, D-30167 Hannover, Germany*

(Dated: March 25, 2020)

The worldwide advanced gravitational-wave (GW) detector network has so far primarily consisted of the two Advanced LIGO observatories at Hanford and Livingston, with Advanced Virgo joining the 2016-7 O2 observation run at a relatively late stage. However Virgo has been observing alongside the LIGO detectors since the start of the O3 run; in the near future, the KAGRA detector will join the global network and a further LIGO detector in India is under construction. Gravitational-wave search methods would therefore benefit from the ability to analyse data from an arbitrary network of detectors. In this paper we extend the PyCBC offline compact binary coalescence (CBC) search analysis to three or more detectors, and describe resulting updates to the coincident search and event ranking statistic. For a three-detector network, our improved multi-detector search finds 20% more simulated signals at fixed false alarm rate in idealized colored Gaussian noise, and up to 40% more in real data, compared to the two-detector analysis previously used during O2.

I. INTRODUCTION

The recent advent of gravitational-wave astronomy with observations by the Advanced LIGO-Virgo network [1, 2] of the coalescence of binary black hole and binary neutron star systems [3–5] has been made possible by the effective operation of search algorithms to identify scarce signals among months of data dominated by noise. Although sufficiently high amplitude signals may be detectable by generic (unmodeled) searches for excess power [6, 7], the majority of binary merger signals are only identifiable by employing accurate waveform models [8, 9], as in recent results finding additional low-amplitude events in public LIGO-Virgo data [10–12].

The most sensitive searches targeting CBC sources [13–21] detect signals by correlating a bank of waveform templates with the data from each detector in a network, recording peaks in the matched filter SNR time series as *triggers*. These triggers are then compared to those from other detectors to check for *coincidences* formed with consistent times of arrival, component masses and spins. The resulting coincident events are weighted (ranked) according to apparent signal properties and compared to an estimate of the noise background in order to measure the significance of candidate coincident events.

A common challenge in the implementation of templated searches is the high computational cost of correlating months of data at kHz sample rates against 10^5 – 10^6 templates of duration up to minutes. The PyCBC search algorithm [20, 22], built on a highly modular and configurable set of libraries [23], achieves high efficiency by using Fast Fourier Transform implementations optimized for different computing platforms, and is thus able to identify candidate events with latencies

of tens of seconds at moderate computational cost [24]. PyCBC workflows have been designed to take advantage of diverse computational resources including local clusters, XSEDE, and the Open Science Grid [25] using the Pegasus workflow management system [26].

Until now, the offline (archival) PyCBC coincident search has analyzed data from the two Advanced LIGO interferometers only. In general, with data from more than two detectors available within the framework of the triggered coincident search, sensitivity will be optimized by generating triggers from all detectors and allowing all possible types of coincidence between detectors. The presence of several different types of coincidence at various times depending on which detectors are operating raises issues of complexity and of correctly ranking different types of candidate event to preserve search sensitivity.

During part of the initial GW detector era, data from the Virgo detector was analyzed in addition to the two LIGO detectors with 4 km arm length and during some of the science runs a co-located 2 km detector at LIGO Hanford, using the *ihope* [13] analysis, a predecessor of the PyCBC search pipeline architecture. The resulting four-detector search [27] was, though, complex and severely limited by computational cost in practice. Since then the extension of coincident searches to three or more detectors has been addressed for the GstLAL pipeline [28] using a ranking statistic evaluated via nearest-neighbor approximation over the parameter space of multi-detector events and by the MBTA pipeline [18].

The PyCBC Live (online) search deployed during the O2 Advanced LIGO-Virgo observing run extends the existing two-detector search, following up selected significant two-detector coincidences by calculating the corresponding matched filter time series in any additional de-

tectors and incorporating this information in candidate significance. However, this procedure is not yet optimized for sensitivity at high thresholds of significance. (Note that all available detectors are also used for source localisation and parameter estimation [29, 30].)

In this paper, we present changes to the PyCBC offline analysis to search data from three or more detectors, to allow all available detectors to generate coincident events and to compare and calculate significance for the resulting different types of coincidence. By using three or more detectors within the coincident search we increase the search duty cycle, as we can form coincidences in any time where at least two detectors are observing. Figure 1 shows when the detectors were observing during the last few weeks of the joint O2 LIGO-Virgo analysis, including the first period of three-detector observation in the Advanced detector era.

We also improve sensitivity in times where more than two detectors are operating, as it may be possible to obtain coincidences even in cases where for one detector the line of sight to the source lies in a blind spot or the data is temporarily of poor quality. Sensitivity for such times of multi-detector operation is also increased by the ranking of events where triggers are present in three or more detectors, taking advantage of the very low rates of noise for such types of coincidence.

The scheme for finding coincidences between triggers is discussed in section II. The calculation of significance requires a *ranking statistic*, a function of the matched filter SNRs and χ^2 signal-based veto values in different detectors, and of the intrinsic (mass, spin) and extrinsic (time of arrival, amplitude, phase) properties of the apparent signal, in comparison with the estimated noise background of similar templates. We will discuss the ranking statistic and its development for the case of more than two detectors in section III. In section IV we discuss how different types of coincidence can be compared for obtaining overall significance. Section V then shows how these changes to the analysis combine to improve the sensitivity of the network to compact binary coalescence gravitational-wave signals.

II. MULTI-DETECTOR COINCIDENCE

Triggers produced in multiple detectors from a common astrophysical source will occur within a short time window of each other, given by time-of-flight considerations and timing measurement uncertainty. This fact allows us to exclude the vast majority of noise triggers, which are uncorrelated in time between detectors, thus forming the basis of our test for coincidence. A signal would also produce triggers in the same waveform template in all detectors, so we only search for coincidence between detectors in a given template [20].

We check for these coincidences in all combinations of detectors in the first stages of the analysis. For notation, we will refer to coincidences by the initials of the detec-

tors involved, for example ‘HL’ coincidences are formed by the LIGO Hanford (H) and LIGO Livingston (L) detectors, and an ‘HLV’ coincidence will also incorporate a trigger from Virgo (V). Thus for the LIGO-Virgo network we form coincidences in the combinations HL, HV, LV and HLV.

A coincidence is formed if two or more triggers from different detectors are within a certain time window. This window is taken as the time of flight for gravitational waves between the sites plus a small, fixed amount to allow for timing errors. Two-detector coincidences are found by comparing the times of triggers in the two detectors; if the difference between the arrival times is less than the allowed time window, then they are considered coincident.

We form three-detector coincidences by applying the same two-detector coincidence test to trigger time differences for each pair of interferometers. The consequences of this multi-detector coincidence test are discussed in more detail in section III A

A. Multi-detector background estimation

In order to measure whether a candidate coincidence is significant we assign it a ranking statistic based on $\vec{\kappa}$, its intrinsic and extrinsic parameters, which we compare to the ranking statistics of a manufactured set of noise coincidences which form the background. By counting how many background events are ranked higher than each candidate event and dividing by the effective length of time for which background events were generated, we may calculate false alarm rates (FARs) down to one per tens of millennia for week-long stretches of data.

This manufactured set of noise coincidences is made up of combinations of single detector triggers which have been time-shifted such that the difference between arrival times at different detectors is not physically allowed. These *time-shift coincidences* are treated in the same way as candidate coincidences in order to estimate the distribution of the ranking statistic under a no-signal hypothesis [13, 20].

In the two-detector configuration, applying time shifts is straightforward as shifting one detector is entirely equivalent to shifting the other detector in the other direction. Typically we perform many background analyses with a regularly spaced set of time-shifts at intervals of 0.1 s for HL coincidences, which is a few times larger than the maximum physical time difference. However in the many-detector configuration, there is the possibility of allowing data from every detector to shift relative to the others. For just a few days of coincident data, given a 0.1 s time shift interval, applying all possible relative shifts between three detectors would result in a background analysis producing an amount of coincidences equivalent to the detector network running for the age of the universe, and the resulting data storage would become unreasonable.

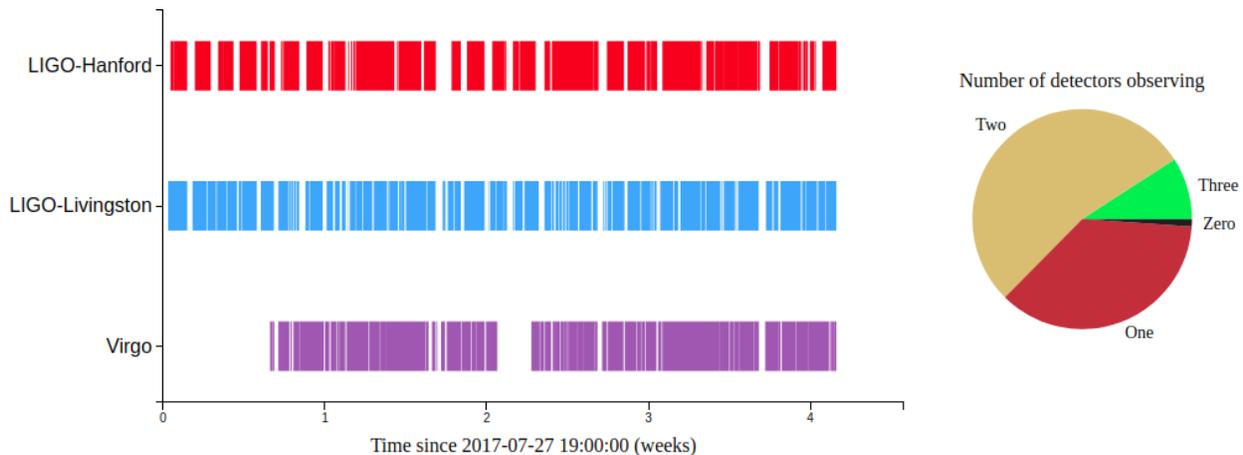


FIG. 1. Times for which the detector data was marked as ‘science ready’ during late O2. The pie chart shows the proportion of time over all of O2 where a given number of detectors are observing. The addition of Virgo in the last few weeks means that there are 15 days of data which we can analyse using three-detector coincidences.

Thus, we choose to fix all the detectors relative to one another, except for one, whose data is time-shifted relative to all the others. As a consequence, any triple-detector background coincidences will require two-detector coincidences within this fixed set. Figure 2 demonstrates how the time-shift coincidences are formed.

Other possible configurations would be to perform time shifts of randomized size on all detectors, up to a fixed total number of background coincidences; or to use different time shift increments for different detectors. These choices are expected to yield similar estimates, and our chosen configuration is the most straightforward. However, with this configuration (and for the time-shift method in general) care is required in order to minimize background contamination by loud signals, as we discuss in the next section.

B. Removal of signals from background coincidences

As discussed above, a candidate’s significance is measured by its FAR, the expected rate of noise events with a higher ranking statistic. Accurate calculation of significance requires, among other issues, separating loud signal triggers from noise triggers. Background coincidences which contain a trigger from a known signal do not accurately represent the noise distribution and therefore may bias our significance estimation.

To counter contamination of signals within our FAR calculation we remove apparent signals from the background as much as possible. We do this by successively taking the highest ranked candidate coincidence and, if its estimated FAR is below a threshold at which we consider it a confident detection, removing all triggers associated with the candidate (within ± 1 s) from the background estimate for lower ranked candidates [31]. This removal procedure is repeated until the FAR for the high-

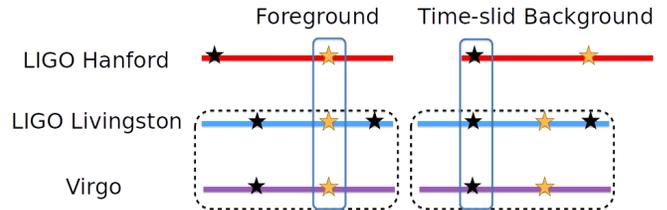


FIG. 2. Diagram showing how background coincidences are formed by time shifts for comparison with three-detector coincidences. The shifted detector here is the LIGO Hanford detector, and LIGO Livingston and Virgo are fixed to one another as shown by the dotted rectangle. The stars show triggers in each detector, with gold stars showing triggers which form candidate coincidences and black stars showing triggers which do not. On the left, a candidate coincidence is highlighted by a blue rectangle; on the right, the background time-shift procedure is illustrated, with a time-shifted coincidence highlighted. In this procedure the timestamps of Hanford detector triggers are shifted by a fixed offset, allowing a three-detector coincidence to form. This is only possible where a two-detector coincidence is formed within the fixed subset.

est ranked remaining candidate event is no longer below the confident detection threshold. The FAR calculation for a given candidate is thus inclusive of its own triggers and therefore accounts for the hypothesis that the candidate as well as all lower ranked events are noise. A more comprehensive discussion of signal removal from background is given in [32].

For our chosen method of applying time shifts to triple coincident events, we require the triggers from the fixed (non-shifted) detectors to be coincident without applying any time shift. Thus, they also form two-detector candidate events. This means that a small number of such triggers could contain signals, and therefore contaminate the background estimation.

In order to remove signals which we would not be able to see in all detectors, in addition to the steps applied to two-detector coincidences described above, we enforce that only triggers during time coincident between all detectors in that combination can form background coincidences. This means that a trigger from a signal which could not be seen in this detector combination would not be able to form background coincidences. For example, we exclude all triggers which occur in LV-only time from the HLV background, or triggers in single-detector time from the two-detector coincident background.

The removal of triggers which form significant candidate coincidences from the data, as described above, is done across all active detector combinations. For example, if a highly-significant candidate is seen in HL, and Virgo is available but did not produce a trigger, then the Hanford and Livingston triggers will be removed from the HLV, HV and LV background estimates.

1. Choice of shifted detector for background coincidences

Within our scheme, the choice of which detector to time-shift can strongly influence the amount of background contamination by signal triggers. We minimize this contamination by ensuring that the least sensitive detector is within the fixed sub-network: then, it is unlikely that a signal will generate a candidate event in the fixed sub-network without also producing a coincident trigger in the time-shifted detector

In order to show the effect of the choice of detector being shifted we compare background coincidences for each choice of shifted detector given the same set of single detector triggers. For this demonstration, we use the same data as is used in section VB. Figure 3 shows that when choosing the least sensitive detector (Virgo) to be time-shifted many background coincidences are produced at large values of the ranking statistic. This is due to candidate HL coincidences in the fixed sub-network matching to noise triggers in Virgo when time-shifted, contaminating the background. For the data chosen, this candidate HL coincidence is GW170809 [5].

The same effect is not seen in the other two cases as the triggers from candidate HL coincidences are being separated by the time-shift and candidate HV and LV coincidences generally have lower value of the ranking statistic due to the lower sensitivity of Virgo. This shows that our choice of including the least sensitive detector in the fixed sub-network is well founded.

III. MULTI-DETECTOR RANKING STATISTIC

In order to compare different types of coincidence to one another, we develop here a new ranking statistic which reflects our degree of belief that a signal is astrophysical in origin and which is consistent between detector combinations. This ranking statistic is based on the

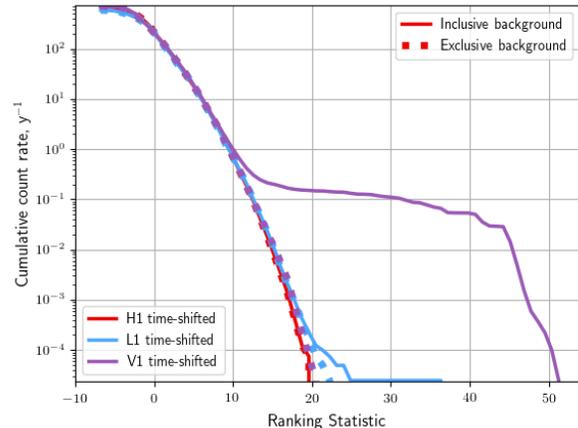


FIG. 3. Histograms of ranking statistic for three-detector noise coincidences, colored according to the detector that is time-shifted when forming the background coincidences. The solid lines show the ranking statistic including all triggers; the dotted line shows background coincidences excluding triggers that form candidate coincidences with IFAR > 0.003 years in any detector combination. We see that when Virgo is time-shifted (magenta) the inclusive background (solid line) has a significant number of loud background coincidences compared to the exclusive background (dotted line), this is not observed when shifting Hanford or Livingston.

rates of noise events in each detector, in addition to measures of detector network sensitivity and multi-detector signal consistency describing the signal event rate.

The Neyman-Pearson optimal detection statistic for triggered searches is given by the ratio of the signal and noise event rate densities [33]. Accurate description of these rates is therefore required.

A coincident event of unknown origin can be described by parameters $\vec{\kappa}$, including the trigger SNR ρ , signal-glitch discriminator(s) χ^2 and template sensitivity σ for each participating detector labelled by a ; the template intrinsic parameters, $\vec{\theta}$, comprising binary component masses and spins; and the amplitude ratio \mathfrak{A}_{ab} , time difference δt_{ab} and phase difference $\delta \phi_{ab}$ between each detector pair labelled by $a \neq b$,¹

$$\vec{\kappa} = \left\{ [\rho_a, \chi_a^2, \sigma_a], \vec{\theta}, [\mathfrak{A}_{ab}, \delta t_{ab}, \delta \phi_{ab}] \right\}. \quad (1)$$

A ranking statistic Λ is a function of this set of parameters. The optimal statistic maximises the expected number of coincident events due to signals that are recovered above a statistic threshold Λ^* corresponding to a given FAR.

The false alarm rate as a function of the threshold Λ^*

¹ For more than one pair of detectors, the amplitude ratios and time and phase differences are not all independent variables.

is

$$\text{FAR}(\Lambda^*) = \int r_N(\vec{\kappa}) \Theta(\Lambda(\vec{\kappa}) - \Lambda^*) d^n \vec{\kappa}, \quad (2)$$

where r_N is the noise rate density over $\vec{\kappa}$; the true alarm (signal) rate is given by the equivalent calculation replacing r_N with $r_S \equiv \mu_S \hat{r}_S(\vec{\kappa})$, where μ_S is an astrophysical coalescence rate per volume per time and $\hat{r}_S(\vec{\kappa})$ is a transfer function describing the recovered event distribution, which will depend on detector orientation and sensitivities.

The optimal detection statistic is given by the ratio of event densities

$$\Lambda_{\text{opt}}(\vec{\kappa}) = \mu_S \frac{\hat{r}_S(\vec{\kappa})}{r_N(\vec{\kappa})}, \quad (3)$$

and the problem reduces to finding the form of the signal and noise distributions. The signal rate μ_S is assumed constant, thus does not need to be known.

We describe in sections III A and III B how the noise and signal event rates respectively are calculated. By using the ratio of signal to noise rates as our ranking statistic, we ensure that it is comparable across different detector combinations. Considering the dynamic range of expected rates spans many orders of magnitude, it is convenient to use the logarithm of the ratio of signal and noise rates,

$$R = \log(r_{s,i}) - \log(r_{n,i}). \quad (4)$$

A. Noise Model

Here we describe the methods used in estimating the coincident noise rate, which involves calculating the distribution of triggers with a given re-weighted SNR statistic in the individual detectors, and combining the single-detector trigger rates to find the rates of noise coincidences of all possible types.

1. Single-detector trigger distributions

We estimate the rate of triggers with a given re-weighted SNR for each detector by fitting the overall distribution of triggers to a decreasing exponential function.

In Gaussian noise we would be able to analytically fit the noise distributions, and if the noise in all detectors were Gaussian, we would be able to combine the noise distribution of SNR for each detector d , $p(\rho_d|N)$ in the form $\rho_c \equiv \sum_d \rho_d^2$ for coincident triggers. However due to the presence of glitches and non-Gaussian behaviour in the data, we are unable to do this.

We use the chi-squared discriminant of [34] to ensure that candidate events have frequency evolution consistent with a binary merger signal. Although we cannot analytically predict the χ^2 distribution for glitch triggers

in each detector, we can describe the density of noise triggers via a combination of ρ and χ^2 ,

$$\hat{\rho} = \begin{cases} \frac{\rho}{[(1+(\chi^2)^{p/2})/2]^{1/p}} & \text{for } \chi_r^2 > 1 \\ \rho & \text{for } \chi_r^2 \leq 1 \end{cases}, \quad (5)$$

where the ‘index’ p is usually set to 6. This re-weighted SNR, which is approximately equal to ρ for reduced χ^2 values close to 1, can be used to describe the distribution of single-detector triggers as a simplification relative to directly modelling the density in the ρ - χ^2 plane [35]. We also implicitly include the sine-Gaussian veto as described in [36, 37] in our re-weighted SNR in order to down-weight ‘blip’ glitches in the data.

In each template, θ , and each detector d , we model the distribution of the rate of triggers r_d with respect to $\hat{\rho}$ as a falling exponential, above a threshold $\hat{\rho}_{th}$:

$$r_d(\hat{\rho}; \vec{\theta}; N) = \mu(\vec{\theta}) p(\hat{\rho} | \vec{\theta}, N), \quad (6)$$

where

$$p(\hat{\rho} | \vec{\theta}, N) = \begin{cases} \alpha(\vec{\theta}) \exp[-\alpha(\vec{\theta})(\hat{\rho} - \hat{\rho}_{th})] & \hat{\rho} > \hat{\rho}_{th} \\ 0 & \hat{\rho} \leq \hat{\rho}_{th}, \end{cases} \quad (7)$$

given model parameters of $\vec{\theta}$, the number of triggers in the template above threshold, $\mu(\vec{\theta})$, and $\alpha(\vec{\theta})$, the exponential decay rate. Before performing the fit, we remove a fixed, small number of high- $\hat{\rho}$ triggers from each detector in order to mitigate possible bias due to loud signals.

To calculate the model parameter α , we use a maximum likelihood fitting procedure. The log likelihood for obtaining a set of samples $\{\hat{\rho}\}$ is

$$\ln p(\{\hat{\rho}\}_d | \alpha, n_{tr}) = n_{tr} \ln \alpha - \alpha \sum_j^{n_{tr}} (\hat{\rho}_j - \hat{\rho}_{d,th}), \quad (8)$$

where j is the index for each trigger, and $\hat{\rho}_{d,th}$ indicates that the fit threshold value of $\hat{\rho}$ could be different for each detector; in practice we use the same fit threshold for all detectors. This likelihood is maximized in each detector by

$$\alpha_{\text{ML}} = (\bar{\hat{\rho}} - \hat{\rho}_{th})^{-1}, \quad (9)$$

where $\bar{\hat{\rho}}$ indicates the mean value of $\hat{\rho}$, and the fractional variance of the fit parameter is approximately $1/\sqrt{n_{tr}}$.

The α_{ML} and n_{tr} values are calculated for each template. There are often relatively few triggers in each individual template, so the variance can be large. We reduce variance by then taking a moving average of the fit parameters over templates which have similar parameters, under the assumption that ‘nearby’ templates will have similar noise distributions. Since α_{ML}^{-1} is a linear function of the mean $\bar{\hat{\rho}}$ for a single template, we may take a mean of α_{ML}^{-1} values (weighted by n_{tr}) over several templates and obtain an identical result to performing the ML fit directly over all triggers in those templates. We

also smooth the count of triggers above threshold by taking the mean over nearby templates.

This smoothing over template parameter was initially performed over templates with similar duration [5], however even at constant template duration the variation of fit parameters over the effective spin χ_{eff} and mass ratio η parameters is not insignificant: therefore a multi-dimensional smoothing of α and μ is performed as in [11].

2. Coincident noise event rate estimation

The optimal ranking statistic includes the expected rate of coincident noise events, which we calculate from single-detector noise trigger rates. The rate of noise coincidences in a template i can be estimated by multiplying together single-detector noise trigger rates² $r_{di}(\hat{\rho}_d)$ for each detector d , and the size of the window of allowed coincidences $A_{N\{d\}}$, where $\{d\}$ is the set of detectors involved in the coincidence.

The rate of noise coincidences for a set of detectors $\{d\}$ is

$$r_{\{d\}i} = A_{N\{d\}} \prod_d r_{di}(\hat{\rho}_d), \quad (10)$$

where r_{di} are the rates in the individual detectors for each template as a function of the reweighted SNR in that detector $\hat{\rho}_d$ and $A_{N\{d\}}$ is an allowed time window for forming coincidences. The time window is formed by the limits on trigger time differences δt_{ab} between detectors a and b . We denote these limits as τ_{ab} and specify them as the GW travel time between detectors plus a small allowance for timing error (typically $\approx \pm 2$ ms).

The allowed time window for a two-detector coincidence is simply $A_{N\{12\}} = 2\tau_{12}$. The allowed window for forming a three-detector event $A_{N\{123\}}$ via coincidence tests applied to each pair of detectors is a product of the times τ_{12} and τ_{13} , subtracting terms corresponding to disallowed regions with $|\delta t_{23}| > \tau_{23}/2$:

$$A_{N123} = 2\tau_{12}\tau_{13} + 2\tau_{12}\tau_{23} + 2\tau_{13}\tau_{23} - \tau_{12}^2 - \tau_{13}^2 - \tau_{23}^2. \quad (11)$$

The time difference window populated by signals will be different to this because of limitations on combinations of time differences in addition to those on individual two-detector differences. For example, no two-detector time differences can take maximal values from the same signal unless all detectors are in a straight line (which is not the case for detectors located on Earth's surface). This has the effect of restricting the bounding window to be an ellipse with area [38]

$$A_{S123} = \pi\tau_{12}\tau_{13} \sin \psi_{23}, \quad (12)$$

² In this section we do not explicitly include the dependence on the fitting parameters α and μ , instead returning to the general notation $r_{di}(\hat{\rho}_d)$.

where ψ_{23} is the angle between the lines-of-sight to detectors 2 and 3 as measured at detector 1. Comparison of this result with Eq. (11) given the HLV detector network shows us that around a third of noise coincidences will fall outside of this allowed signal window.

Figure 4 shows the allowed differences in arrival times of noise and signal coincidences in three detectors. Coincidences in a region disallowed for signals but allowed for noise will be heavily suppressed in the ranking statistic through the time difference consistency part of the phase-time-amplitude consistency checks of section III B 1.

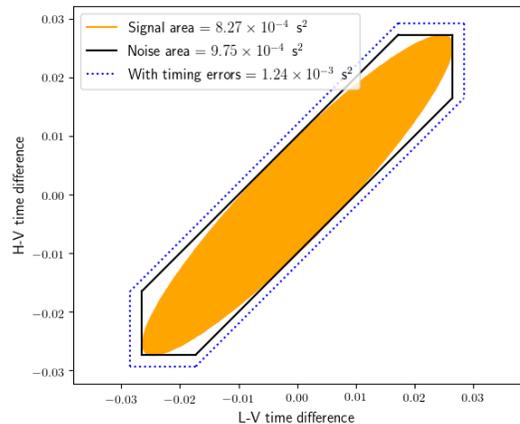


FIG. 4. Allowed time differences between LIGO-Livingston and Virgo given the time difference between LIGO-Hanford and Virgo. The black lines show the region where three-detector coincidences are allowed to form, neglecting timing error; dotted blue lines show the allowed area for noise coincidences allowing for 2 ms timing error. The orange shaded area shows where signals are expected to lie, neglecting timing error. Although noise events are allowed over a larger time window than is physically possible for signals, events far outside the signal area are suppressed by the phase-time-amplitude consistency checks of section III B 1.

Since the effective time windows for signal and noise coincidences are thus of similar size, we neglect differences between them and incorporate the rate of noise triggers in the ranking statistic as

$$\log(r_{n,i}) = \log A_{N\{d\}} + \sum_d \log r_{di}(\hat{\rho}_d). \quad (13)$$

B. Signal Model

The true rate of signals depends on the overall rate of astrophysical mergers, the location and orientation of sources, and the distribution of intrinsic signal parameters (masses, spins, etc.) [39]. Although we do not consider the distribution of intrinsic parameters over sources in this work, it has been considered in [21, 40] and employed in [11]. We instead look at the parameters we can use to predict the rate of recovered signals as a function of this (assumed constant) astrophysical merger rate.

1. Source Location

Astrophysical populations of sources are expected to be isotropically distributed over sky location and binary orientation. Similarly, for sources in the nearby universe such as LIGO and Virgo are currently detecting, the population is expected to be nearly uniform in volume. We can estimate how this population prior impacts the distribution of sources which are detectable by the gravitational-wave network, and furthermore how it affects the observed distribution of signal amplitudes, phases, and times of arrival in different detectors. For example, the probability of finding a signal is not uniformly distributed over the time and phase difference between detectors, and the two are correlated [21]. Figure 4 shows a hard boundary to the allowed signal time differences for illustration purposes, but in reality there would be timing error which would blur the edges, and the time differences within the bounding ellipse would not be uniform. For more discussion of the shape of the prior histograms, see [21]. In contrast the times of ‘arrival’ of noise triggers will be random and uniformly distributed; similarly differences in gravitational wave signal phase between the detectors will also be uniformly distributed for noise coincidences.

The ratio of signal amplitudes between detectors would also have different distributions for noise vs. signal coincidences: the relative amplitudes of signals in different detectors will be dependent upon the source’s sky position and polarization angle.

For a given set of extrinsic signal parameters $\vec{\Omega}$ comprising the relative amplitudes $\mathfrak{A}_{\{d\}}$, time delays $\delta t_{\{d\}}$ and phase differences between each site $\delta\phi_{\{d\}}$, the probability for that set of parameters to be generated by a signal $p(\vec{\Omega}|S)$ forms part of our ranking statistic for coincident events. To find the probability distribution, we perform a Monte Carlo calculation given the detector locations to produce histograms which act as look-up tables for $p(\vec{\Omega}|S)$. As an improvement over [21], we have updated these signal prior histograms to support a three-detector network for use in triple coincidences.

2. Detector sensitivity

The sensitive distance of a detector, defined as the luminosity distance at which a standard compact binary source has a given expected SNR, affects the expected rate of signals that produce triggers in a given detector. This distance varies substantially between interferometers and over time: we include this information in our ranking statistic via a term accounting for network sensitivity for a given coincidence type.

The instantaneous sensitive distance in a given detector, for sources matching a template labeled by i , is pro-

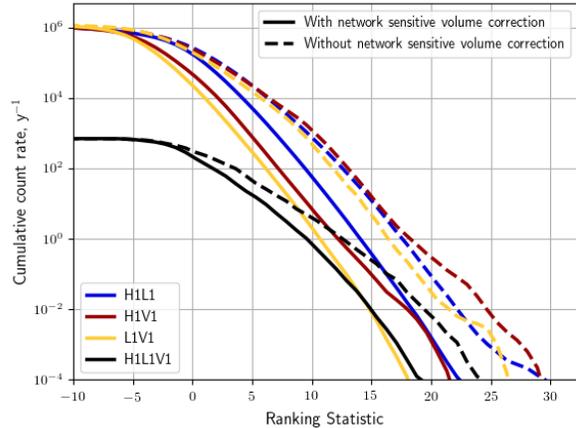


FIG. 5. Histograms of ranking statistic for different types of noise coincidences, colored by different detector combinations. The solid lines show the ranking statistic including the network sensitive volume term R_σ ; dashed lines show the ranking statistic without this term. We see that detector combinations containing Virgo (HV [crimson], LV [gold], HLV [black]) are penalized due to lower network sensitivity compared to HL-only [navy]. Note also the much lower overall rate of triple (HLV) coincidences.

portional to the quantity σ defined in [9].³ Then, network sensitivity for a given coincidence type is determined by the least sensitive detector via $\sigma_{\min,i}$. Under the assumption of a homogeneous distribution of sources in volume, the expected rate of signals for a given coincidence type is therefore proportional to $\sigma_{\min,i}^3$.

To normalize this measure of instantaneous sensitivity we compare to the rate corresponding to representative values of σ_i over the analysis time. For the LIGO-Virgo network, we choose as a representative σ value the median network sensitivity for HL coincidences $\bar{\sigma}_{HL,i}$.⁴ Thus, the time-dependent rate of signals in a given coincidence type described by $\sigma_{\min,i}$ is proportional to

$$r_{s,i} \propto p(\vec{\Omega}|S) \frac{\sigma_{\min,i}^3}{\bar{\sigma}_{HL,i}^3}, \quad (14)$$

leading to a term in the (logarithm of) the relative rate of signal vs. noise triggers

$$R_{\sigma,i} \equiv 3 (\log \sigma_{\min,i} - \log \bar{\sigma}_{HL,i}). \quad (15)$$

³ Technically, σ is the expected SNR for a face-on binary coalescence with a waveform perfectly matching a given template, located directly overhead from the detector at a luminosity distance of 1 Mpc.

⁴ I.e. the minimum over H and L of the median detector sensitivity over observing time. In general, we will use the most sensitive coincidence type as representative for this normalization.

Our statistic suppresses events in coincidence types where the least sensitive detector is significantly less sensitive than the others, as these are less likely to contain signals. During O2 the Virgo detector was much less sensitive than the LIGO detectors, thus as seen in figure 5 the distribution of background ranking statistics from HL coincidence is only mildly affected by the R_σ term, whereas HV, LV and HLV statistic values are all more heavily reduced. Note that coincidences in times of relatively poor sensitivity for H or L will also be penalized.

C. Final ranking statistic

Combining equation (13) for the noise rate and (14) for the signal rate into (4), we then obtain our final ranking statistic

$$R = -\log A_{N\{d\}} - \sum_d \log r_{di}(\hat{\rho}_d) + \log p(\vec{\Omega}|S) + R_{\sigma,i}, \quad (16)$$

where $A_{N\{d\}}$ is the allowed time window for coincidence of equation (11), $r_{di}(\hat{\rho}_d)$ is the expected rate density of triggers in template i and detector d at re-weighted SNR $\hat{\rho}_d$, $p(\vec{\Omega}|S)$ is the probability of a signal having the extrinsic parameters $\vec{\Omega}$ given by the prior histograms, and $R_{\sigma,i}$ is (the log of) network sensitive volume for a given template and coincidence type, which is proportional to the expected rate of signals.

IV. SIGNIFICANCE OF MULTI-DETECTOR CANDIDATE EVENTS

The ranking statistic described above is designed to represent the relative probability of signal vs. noise origin for a coincident event regardless of the detectors involved [33]. Thus we can compare statistic values across different coincidence types, considerably simplifying the task of producing a final list of candidates and calculating their significance.

To account for correlated events in different detector combinations produced by the same signal or noise transient, coincidences are clustered within a sliding window of ten seconds: the event with highest ranking statistic is then kept and others are discarded. In order for this clustering operation not to damage search sensitivity, it is necessary for our statistic to correspond to the relative probability of signal vs. noise origin.

The clustering over event types then determines how false alarm rates are calculated in times where more than one available detector combination is active. To do this, we add together the estimated false alarm rates from all available types of coincidence at the ranking statistic threshold of a candidate event. If our ranking statistic was not comparable across detector types, then we would

not be able to calculate an equivalent false alarm rate in other detector combinations.

As an example, in the O2 data used in section VB, if a coincident event is found from LIGO Livingston and Virgo with a ranking statistic of 10, this would have a FAR of approximately 1 per year. If LIGO Hanford data is not available at this time, then this would be the given FAR. But if LIGO Hanford is available, and did not participate in a more significant coincidence, then the false alarm rate for this ranking statistic in LV would be added to the FAR for HL coincidences (50 per year), HLV (0.8 per year) and HV (2 per year) at the same ranking statistic threshold to give an overall FAR of around 54 per year. This combination method down-ranks triggers seen in less-sensitive detector combinations when more sensitive combinations are available.

This method affects the minimum FAR we can measure for coincidences in triple-detector time, which given the minimum measurable FAR in each coincidence type is

$$\text{FAR}_{\min} = \frac{1}{t_{\text{bg,HLV}}} + \frac{1}{t_{\text{bg,HL}}} + \frac{1}{t_{\text{bg,HV}}} + \frac{1}{t_{\text{bg,LV}}}, \quad (17)$$

where t_{bg} is the total time analyzed by time shifts in a given combination. Thus, for a three- (four-, five-) detector analysis, the FAR estimate in times where all detectors are observing is limited by approximately a factor four (11, 26) relative to the minimal FAR in a comparable two-detector analysis. However as remarked above in Section IIA, the background time for combinations of three or more detectors can be extended at will by using more general multi-detector time shifts, thus only the double coincidence FAR estimates are truly limiting.

V. COMPARISON OF SENSITIVITY

In order to compare the sensitivity of this analysis with the previous search, we estimate the total number of signals out of a notional (simulated) merger population that the two would detect at a given false alarm rate threshold. We express the sensitivity of a search of a given data set as the product of volume of space and observing time, VT , under the assumption of mergers uniformly distributed in space and time: for a hypothetical signal population with local merger rate μ , the expected number of signals the search would detect over the data set is μVT . Our figure of merit is therefore the ratio between these values for different analyses.

Since we do not know the true distributions of masses and spins of merging binary objects in the local Universe, we choose instead to calculate sensitivity for simple analytic distributions lying entirely within the parameter space of the search in order to simplify the interpretation of comparisons. We also neglect the effect of redshift on simulated signals, again to simplify the interpretation of sensitivity comparisons. Injections are performed for coalescences of binary black holes (BBH) and binary neutron stars (BNS); BBH injections use the waveform

model SEOBNRv4 [41], and the BNS injections use the SpinTaylorT5 model [42].

An injection is considered to be detected if the highest ranked search event within one second of the injection merger time has an IFAR (inverse false alarm rate) value above a given threshold. We then use importance sampling over many thousands of injections to calculate the sensitive volume for the target uniform-in-volume distribution [20].

We perform two comparisons to measure sensitivity improvements due to updates described here. The first is a comparison using colored Gaussian data at design sensitivity in section V A, and then an analysis of real O2 data in section V B.

Our tests on Gaussian noise in section V A are designed to measure the improvement in sensitivity for future uses of the search given the presented improvement in the analysis, we analyse a fake three-detector network at design sensitivity and compare it to a network containing just the two LIGO detectors.

The tests on O2 data in section V B are designed to show improvements in how we deal with non-Gaussian transient noise in the data, commonly referred to as ‘glitches’. These glitches cause spikes in SNR and can therefore increase the rate of coincident noise triggers. Real data is also non-stationary with significantly time-dependent sensitivity in each detector, which our statistic is designed to account for.

A. Colored Gaussian Data

By comparing injections in Gaussian data, we see the improvement in sensitivity by using three detectors rather than two in the PyCBC analysis. The sensitivity will be increased by both allowing coincidences where either of the LIGO detectors are not operating or does not produce a trigger, and by updates to ranking of signals seen within all three detectors during triple-detector coincident time.

We use Gaussian noise which has been colored according to design sensitivity curves for advanced LIGO and Virgo [1, 2], and generate 107.3 total hours of coincident data over a five day period. We apply observational data segment definitions and data quality vetoes derived from a section of late O2 data in order to obtain realistic duty cycles of different two- and three-detector combinations and segment lengths for triggers generation. The coincident time available during this test is shown in table I.

The injections performed are separated into two bins: BNS injections with total mass between 2 – 5 M_{\odot} and BBH injections with total mass between 5 – 100 M_{\odot} . We note that since all templates have Gaussian noise event distributions in this data, we expect identical sensitivity to all signals that match the templates except for the fact that higher-mass systems produce higher amplitude GW signals at a given distance. We choose to perform

Detector Combination	Coincident time (hours)	
	Fake Data	Real Data
HL	86.5	129.6
HV	89.7	108.1
LV	95.3	120.2
HLV	82.2	99.4
Total coincident time	107.3	159.2

TABLE I. Table of times in the data used in the two tests shown in this paper for times which are coincident between LIGO Hanford (H), LIGO Livingston (L) and Virgo (V) respectively. Each of the HL, HV, LV times is inclusive of HLV time.

the importance sampling volume integral in a way that scales out this amplitude factor: thus our sensitivity estimate effectively weights every simulated signal equally at a given chirp distance[9], regardless of binary mass. We locate the simulated mergers uniformly on the sky and with uniform distribution in chirp distance between limits of 5 Mpc and 600 Mpc.

For the BBH injections the logarithms of component masses are distributed uniformly between mass limits 2.5 M_{\odot} and 50 M_{\odot} , while for BNS injections the component masses are uniformly distributed between 1 and 2.5 M_{\odot} . The component spins for BNS are distributed between 0 and 0.4, and in the BBH case, between 0 and 0.998. The BBH spins are strictly aligned with the orbital angular momentum, but the BNS spins are not.

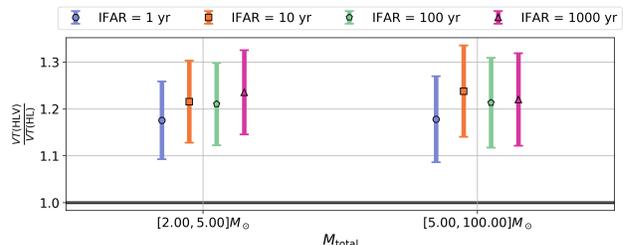


FIG. 6. Ratio of search sensitivities, VT , between the three-detector HLV analysis and two-detector HL analysis in colored Gaussian noise, plotted for four choices of IFAR threshold. Simulated signals used to estimate sensitivity have been split by total mass into BNS and BBH bins.

Figure 6 shows an increase in VT of 1.2 ± 0.1 for the three-detector analysis over the HL analysis for both source types, constant over all IFAR values considered.

A portion of this increase is due to HV or LV (two-detector) coincident time being available to the HLV analysis but not the HL analysis. We can approximate this factor via the total coincident analysis time divided by the HL time, weighted for the maximal volume sensitivity over available combinations at a given time. Using the sensitive range of an equal-mass binary black hole coalescence with component masses of 20 M_{\odot} , we find this factor to be approximately 1.12.

The remaining increase in VT is then attributed to

an increase of search sensitivity in three-detector time. This is partly due to a subset of signals that generate H1-V1 or L1-V1 events in the three-detector search, which would not be seen in the two-detector HL search; also, to the much reduced rate of noise events for three-detector coincidence, implying that even relatively low SNR signals which generate a three-detector event will be more likely to be highly ranked and significant than in the two-detector case.

B. Data from LIGO-Virgo Observing Run 2

In the next test we use real GW detector data from the second observing run of Advanced LIGO and Virgo, O2, from between 2017-08-05 and 2017-08-13. This shows the response of the analysis to signals in the presence of noise artifacts including non-Gaussian transients and time-varying detector sensitivity. Again, we characterise the difference in sensitivity through the VT ratio between two analyses: in this case we compare the HLV analysis with the HL analysis performed for the GWTC-1 catalog [5]. We see the amount of data used for this analysis in table I.

The injections used for this test are located uniformly on the sky and with uniform distribution in chirp distance between limits of 5 Mpc and 300 Mpc. This maximum distance is much less than that used in the fake data case, due to the difference in sensitive distance between the two example data sets. The BNS component masses are distributed uniformly between $1 M_{\odot}$ and $3 M_{\odot}$, and BBH are distributed uniformly in total mass, and then in primary mass for constant total mass, with component masses between $2 M_{\odot}$ and $98 M_{\odot}$ up to a maximum total mass of $100 M_{\odot}$. The spins for injections in this analysis are distributed in the same way as the fake data injections described above.

We separate the injections used here into four bins between 2, 5, 16, 50 and $100 M_{\odot}$ total mass. The motivation of this split is that different parts of the template bank are affected differently by non-Gaussian and non-stationary noise, thus we might expect injections recovered in various mass ranges to be differently affected by changes to the analysis.

Figure 7 shows an increase in sensitivity of between a factor of 1.09 ± 0.08 and 1.4 ± 0.2 , depending on the mass of the system. The strong dependence on the masses of the binary may be correlated to the presence of glitches which mimic the gravitational wave signature of very heavy binary black hole mergers [43]. By down-ranking coincidences which do not fit the time, phase and amplitude differences of section III B 1, we greatly reduce the ranking statistic of the glitch background in heavy BBH templates, this means that our injected signals will be seen with higher significance, and therefore to a greater distance.

As before, some of the increase in sensitivity is due to increased observation time, however in this case the

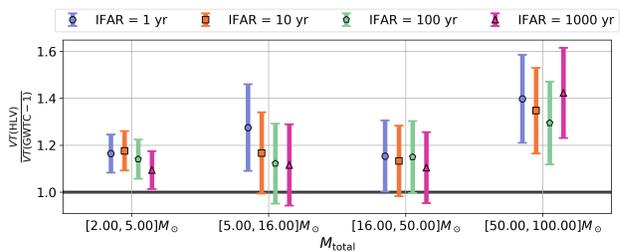


FIG. 7. Volume \times time (VT) ratio comparing analyses of real data with the updated analysis against the PyCBC analysis as performed in [5]. There is a significant increase in VT, particularly for signals from heavy binary black holes, which may be due to changes to methods to differentiate noise artefacts from signals.

expected increase in VT is only a factor ~ 1.01 , rather than the factor 1.12 estimated in section V A, despite a similar relative increase in the analysis time. This is because, in contrast to the fake data case, here the Virgo detector has a sensitive range less than half that of the LIGO detectors, thus the expected sensitivity in times when only HV or LV are observing is much smaller than for HL times.

VI. CONCLUSIONS

We have presented changes to the PyCBC offline coincidence search related to analysing data from more than two detectors within the same analysis. These changes mean that we can now search over many interferometer combinations, and take certain characteristics of detector behaviour into account within the offline search. These improvements also mean that we can suppress noise coincidences more than before, and improve our prospects of finding signals within the data.

Tests on Gaussian data have shown that by using more detectors in the analysis for a LIGO-Virgo network at design sensitivity we increase the sensitive VT by a factor of 1.2 ± 0.1 . This is largely due to better ranking of signals seen in all detectors during triple-detector time due to suppression of noise, but also partly due to an increase in the duty factor of the network, even if the network has a slightly lower sensitivity. This second factor is not as significant if the additional detectors do not have relatively equal sensitivities, but with improvements to the Virgo detector ongoing and making its sensitivity more comparable to the LIGO detectors, we expect that this will become more significant in the near future.

We have also shown that the changes made to the analysis deal better with non-Gaussian noise realisations in the data due to better signal consistency checking and noise suppression. This means that in O2 data, where the Virgo detector was available but not significantly sensitive, we obtain an increase in VT sensitivity by a factor of between 1.09 ± 0.08 and 1.4 ± 0.2 , depending

on source properties. The greatest increase in sensitivity in this test was for signals from heavier black hole binaries, which had previously been largely affected by specific types of detector glitches.

A full catalogue of the gravitational wave events identified in the O1 and O2 runs using the search described here is available in [11].

VII. FURTHER WORK

As future work, the coincidence tests and ranking method developed here could be applicable to increase the sensitivity of the the PyCBC low latency search, or be of interest to other CBC analysis pipelines looking to extend their searches to three- or more-detector analysis.

The new method for calculation of prior histograms can support an arbitrary number of detectors, however, memory requirements may become impractical for larger numbers of detectors. Alternate methods such as taking incoherent combinations of two/three detector prior histograms may provide accurate enough modelling to handle these cases and will be tested in the future.

To ensure that our event ranking is close to optimal, the modelling of noise trigger distributions as in section III A 1 may require updated fitting models. As significant multi-ifo events may involve triggers with lower SNR, the distribution of these triggers is affected by SNR thresholding, and is not easily fit by a simple function for all values of $\hat{\rho}$. Future work will investigate more accurate models for such distributions.

VIII. ACKNOWLEDGEMENTS

GSD and TD acknowledge support from the Maria de Maeztu Unit of Excellence MDM-2016-0692 and Xunta de Galicia. The PyCBC offline search software is built upon LALSuite [44], numpy [45], SciPy [46] and Astropy [47].

The analysis in this paper used PyCBC version 1.15.2 [23]. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459.

-
- [1] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, *et al.*, *Classical and Quantum Gravity* **32**, 074001 (2015).
- [2] 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 (2014).
- [3] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [4] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **119**, 161101 (2017), [arXiv:1710.05832 \[gr-qc\]](#).
- [5] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, *et al.*, *Physical Review X* **9** (2019), [10.1103/physrevx.9.031040](#).
- [6] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. D* **93**, 122004 (2016), [Addendum: *Phys. Rev. D* **94**, no.6, 069903 (2016)], [arXiv:1602.03843 \[gr-qc\]](#).
- [7] S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. Da Silva, *et al.*, *Physical Review D* **93** (2016), [10.1103/physrevd.93.042004](#).
- [8] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993), [arXiv:gr-qc/9301003 \[gr-qc\]](#).
- [9] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [10] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
- [11] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, (2019), [arXiv:1910.05331 \[astro-ph.HE\]](#).
- [12] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, (2019), [arXiv:1904.07214 \[astro-ph.HE\]](#).
- [13] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T. Cokelaer, J. D. E. Creighton, T. Dent, A. Dietz, S. Fairhurst, N. Fotopoulos, G. González, C. Hanna, I. W. Harry, G. Jones, D. Keppel, D. J. A. McKechnan, L. Pekowsky, S. Privitera, C. Robinson, A. C. Rodriguez, B. S. Sathyaprakash, A. S. Sengupta, M. Vallisneri, R. Vaulin, and A. J. Weinstein, *Phys. Rev. D* **87**, 024033 (2013).
- [14] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Keppel, L. Liao, S. Privitera, A. Searle, L. Singer, and A. Weinstein, *The Astrophysical Journal* **748**, 136 (2012).
- [15] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. E. Creighton, R. Everett, C. Hanna, D. Keppel, R. N. Lang, T. G. F. Li, D. Meacher, A. Nielsen, C. Pankow, S. Privitera, H. Qi, S. Sachdev, L. Sadeghian, L. Singer, E. G. Thomas, L. Wade, M. Wade, A. Weinstein, and K. Wiesner, *Phys. Rev. D* **95**, 042001 (2017).
- [16] S. Sachdev, S. Caudill, H. Fong, R. K. L. Lo, C. Messick, D. Mukherjee, R. Magee, L. Tsukada, K. Blackburn, P. Brady, P. Brockill, K. Cannon, S. J. Chamberlin, D. Chatterjee, J. D. E. Creighton, P. Godwin, A. Gupta, C. Hanna, S. Kapadia, R. N. Lang, T. G. F. Li, D. Meacher, A. Pace, S. Privitera, L. Sadeghian, L. Wade, M. Wade, A. Weinstein, and S. L. Xiao, “The gstlal search analysis methods for compact binary mergers in advanced ligo’s second and advanced virgo’s first observing runs,” (2019), [arXiv:1901.08580 \[gr-qc\]](#).
- [17] D. Buskulic *et al.*, **27**, 194013 (2010).
- [18] T. Adams, “Low latency search for compact binary coalescences using mbta,” (2015), [arXiv:1507.01787 \[gr-qc\]](#).

- [19] S. Hooper, S. K. Chung, J. Luan, D. Blair, Y. Chen, and L. Wen, *Physical Review D* **86** (2012), [10.1103/physrevd.86.024012](https://arxiv.org/abs/10.1103/physrevd.86.024012).
- [20] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. D. Canton, T. Dent, S. Fairhurst, *et al.*, *Classical and Quantum Gravity* **33**, 215004 (2016).
- [21] A. H. Nitz, T. Dent, T. D. Canton, S. Fairhurst, and D. A. Brown, *The Astrophysical Journal* **849**, 118 (2017).
- [22] T. Dal Canton *et al.*, *Phys. Rev. D* **90**, 082004 (2014), [arXiv:1405.6731 \[gr-qc\]](https://arxiv.org/abs/1405.6731).
- [23] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, C. Capano, L. Pekowsky, T. Dent, A. R. Williamson, S. De, G. Davies, M. Cabero, D. Macleod, B. Machenschalk, S. Reyes, P. Kumar, T. Massinger, F. Pannarale, dfinstad, M. Tápai, S. Fairhurst, S. Khan, L. Singer, A. Nielsen, shasvath, S. Kumar, idorrington92, H. Gabbard, and B. U. V. Gadre, “[gwastro/pycbc: Pycbc release v1.15.2](https://arxiv.org/abs/1905.02227),” (2019).
- [24] A. H. Nitz, T. Dal Canton, D. Davis, and S. Reyes, *Physical Review D* **98** (2018), [10.1103/physrevd.98.024050](https://arxiv.org/abs/10.1103/physrevd.98.024050).
- [25] D. Weitzel, B. Bockelman, D. A. Brown, P. Couvares, F. Würthwein, and E. Fajardo Hernandez, in *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact*, PEARC17 (ACM, New York, NY, USA, 2017) pp. 24:1–24:6, [arXiv:1705.06202 \[cs.DC\]](https://arxiv.org/abs/1705.06202).
- [26] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger, *Future Generation Computer Systems* **46**, 17 (2015).
- [27] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, *et al.*, *Physical Review D* **82** (2010), [10.1103/physrevd.82.102001](https://arxiv.org/abs/10.1103/physrevd.82.102001).
- [28] C. Hanna, S. Caudill, C. Messick, A. Reza, S. Sachdev, L. Tsukada, K. Cannon, K. Blackburn, J. D. E. Creighton, H. Fong, P. Godwin, S. Kapadia, T. G. F. Li, R. Magee, D. Meacher, D. Mukherjee, A. Pace, S. Privitera, R. K. L. Lo, and L. Wade, “Fast evaluation of multi-detector consistency for real-time gravitational wave searches,” (2019), [arXiv:1901.02227 \[gr-qc\]](https://arxiv.org/abs/1901.02227).
- [29] L. P. Singer and L. R. Price, *Physical Review D* **93** (2016), [10.1103/physrevd.93.024013](https://arxiv.org/abs/10.1103/physrevd.93.024013).
- [30] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, *Publications of the Astronomical Society of the Pacific* **131**, 024503 (2019).
- [31] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **6**, 041015 (2016), [erratum: *Phys. Rev. X*, no.3,039903(2018)], [arXiv:1606.04856 \[gr-qc\]](https://arxiv.org/abs/1606.04856).
- [32] C. Capano, T. Dent, C. Hanna, M. Hendry, Y.-M. Hu, C. Messenger, and J. Veitch, *Physical Review D* **96** (2017), [10.1103/PhysRevD.96.082002](https://arxiv.org/abs/10.1103/PhysRevD.96.082002).
- [33] R. Biswas *et al.*, *Phys. Rev. D* **85**, 122008 (2012), [arXiv:1201.2959 \[gr-qc\]](https://arxiv.org/abs/1201.2959).
- [34] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
- [35] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **88**, 024025 (2013), [arXiv:1209.0718 \[gr-qc\]](https://arxiv.org/abs/1209.0718).
- [36] A. H. Nitz, *Classical and Quantum Gravity* **35**, 035016 (2018).
- [37] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, *The Astrophysical Journal* **872**, 195 (2019).
- [38] B. Bhawal and S. V. Dhurandhar, (1995), [arXiv:gr-qc/9509042 \[gr-qc\]](https://arxiv.org/abs/gr-qc/9509042).
- [39] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Astrophys. J.* **882**, L24 (2019), [arXiv:1811.12940 \[astro-ph.HE\]](https://arxiv.org/abs/1811.12940).
- [40] T. Dent and J. Veitch, *Phys. Rev. D* **89**, 062002 (2014).
- [41] 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).
- [42] P. Ajith, *Phys. Rev. D* **84**, 084037 (2011).
- [43] M. Cabero, A. Lundgren, A. H. Nitz, T. Dent, D. Barker, E. Goetz, J. S. Kissel, L. K. Nuttall, P. Schale, R. Schofield, *et al.*, *Classical and Quantum Gravity* **36**, 155010 (2019).
- [44] LIGO Scientific Collaboration, “[LIGO Algorithm Library - LALSuite](https://arxiv.org/abs/1801.02634),” free software (GPL) (2018).
- [45] S. van der Walt, S. C. Colbert, and G. Varoquaux, *Computing in Science & Engineering* **13**, 22 (2011).
- [46] P. Virtanen *et al.*, *Nature Meth.* (2020), [10.1038/s41592-019-0686-2](https://arxiv.org/abs/10.1038/s41592-019-0686-2), [arXiv:1907.10121 \[cs.MS\]](https://arxiv.org/abs/1907.10121).
- [47] A. M. Price-Whelan *et al.*, *Astron. J.* **156**, 123 (2018), [arXiv:1801.02634](https://arxiv.org/abs/1801.02634).