



Detecting Binary Compact-object Mergers with Gravitational Waves: Understanding and Improving the Sensitivity of the PyCBC Search

Alexander H. Nitz^{1,2} , Thomas Dent¹ , Tito Dal Canton^{1,3,6} , Stephen Fairhurst⁴ , and Duncan A. Brown^{2,5}

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

²Department of Physics, Syracuse University, Syracuse NY 13244, USA

³Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁴School of Physics and Astronomy, Cardiff University, Cardiff, UK

⁵Kavli Institute for Theoretical Physics, University of California at Santa Barbara, Santa Barbara, CA 93106, USA

Received 2017 June 20; revised 2017 September 21; accepted 2017 September 24; published 2017 November 7

Abstract

We present an improved search for binary compact-object mergers using a network of ground-based gravitational-wave detectors. We model a volumetric, isotropic source population and incorporate the resulting distribution over signal amplitude, time delay, and coalescence phase into the ranking of candidate events. We describe an improved modeling of the background distribution, and demonstrate incorporating a prior model of the binary mass distribution in the ranking of candidate events. We find an $\sim 10\%$ and $\sim 20\%$ increase in detection volume for simulated binary neutron star and neutron star black hole systems, respectively, corresponding to a reduction of the false alarm rates assigned to signals by between one and two orders of magnitude.

Key words: black hole physics – gravitational waves – stars: neutron

1. Introduction

The observation of binary black hole mergers (Abbott et al. 2017b, 2016f, 2016h) by the Advanced Laser Interferometer Gravitational-wave Observatory (LIGO; Abbott et al. 2016e) has provided a new way of exploring the origin and evolution of compact objects. The measured masses, spins, and orientations of merging binaries can tell us about the nature and formation of compact-object remnants (Abbott et al. 2016a). Beyond the astrophysics that gravitational waves alone can reveal, a campaign is underway to search for electromagnetic counterparts to binary mergers (Abbott et al. 2016g, 2016i, 2016j). Joint electromagnetic and gravitational-wave observations will yield a tremendous amount of information about mergers, their host environments, and their ejecta (Metzger & Berger 2012).

The purpose of LIGO’s binary-merger search is to identify astrophysical events in the detectors’ noise-dominated data. Several different methods were used to identify gravitational-wave signals (Klimenko et al. 2016; Usman et al. 2016; Messick et al. 2017) in Advanced LIGO’s first observing run (Abbott et al. 2016b, 2017c). In this paper, we discuss the PyCBC binary-merger search (Dal Canton et al. 2014; Usman et al. 2016; Nitz et al. 2017), which uses matched filtering (Wiener 1949; Cutler et al. 1993; Allen et al. 2012) with accurate models of gravitational waveforms (Buonanno & Damour 1999; Blanchet 2002; Faye et al. 2012; Taracchini et al. 2014; Bohé et al. 2016) to identify signals. The PyCBC search has been used to detect binary black hole mergers in the first Advanced LIGO observing run (Abbott et al. 2016d) and to place upper limits on the rate of binary neutron star and neutron star–black hole mergers (Abbott et al. 2016k).

Since LIGO detector noise contains both non-stationary and non-Gaussian components, matched filtering will report high signal-to-noise ratios for events that are not astrophysical (Nuttall et al. 2015; Abbott et al. 2016c). To suppress noise, the PyCBC search requires that a source is seen consistently in both LIGO detectors (Usman et al. 2016) and performs an additional consistency test between the data and the target waveform (Allen 2005). Using the signal-to-noise ratio and the signal-consistency tests, each coincident event is assigned a number determined by the search’s *detection statistic*. By construction, the greater the detection statistic’s value, the more likely the event is to be an astrophysical signal. The detection statistic alone is typically not sufficient to determine if an event is astrophysical, since we do not know the distribution of this statistic due to detector noise; this noise background cannot be accurately modeled and must be empirically measured for the search. To determine an event’s significance, we calculate the *false alarm rate* of the search (Abbott et al. 2016d). The search’s false alarm rate measures how often it would report a non-astrophysical event with a detection statistic value as high as a given candidate event. The smaller this false alarm rate is, the more likely the candidate event is to be astrophysical. The PyCBC search consists of two components: a low-latency search that identifies triggers with an estimate of the false alarm rate for rapid follow-up (A. H. Nitz et al. 2017, in preparation), and an offline search that reprocesses the data with additional detector data-quality information and provides the final statement of candidate significance (Usman et al. 2016).

There are two ways to improve the sensitivity of the search for binary-merger signals: we can increase the measured signal-to-noise ratio for a signal, or we can reduce the false alarm rate of the search at a given signal-to-noise ratio. The signal-to-noise ratio of a merger could be increased by reducing instrumental noise in the detectors (Martynov et al. 2016) or improving the fit between the model waveforms used in the search and the true gravitational-wave signal. Work to reduce noise toward the Advanced LIGO detectors’ design sensitivity (and beyond) is ongoing (Abbott et al. 2016l;

⁶ NASA Postdoctoral Program Fellow.



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Abbott et al. 2017a). For binary neutron stars, the fit of the waveform models is already very high for detection searches (Brown et al. 2012), although improvements will help parameter measurement (Lackey & Wade 2015). For binary black holes and neutron star–black hole binaries, improvements to the theoretical waveforms, and extension of searches to cover a larger parameter space (e.g., the effects of binary spin–orbit precession and higher-order modes of gravitational-wave emission) are active areas of research. This paper considers the second of the two approaches to improve search sensitivity: improving the separate ability to separate astrophysical signals from detector noise. Here we restrict our study to the waveforms and search parameter space used in Advanced LIGO’s first observing run (Abbott et al. 2016d).

We present a new detection statistic that reduces the false alarm rate of the search across the entire target signal space: binary neutron stars, binary black holes, and neutron star–black hole binaries. This detection statistic uses information from both the physical parameters of the gravitational-wave signal and the expected distribution of the noise background to rescale the “combined effective signal-to-noise ratio” used previously by the PyCBC search. This is similar to detector-sensitivity weighting of the signal-to-noise ratio proposed by Biswas et al. (2012a, 2012b) and used in the Initial LIGO and Virgo binary-merger searches (Abadie et al. 2010b). We test this detection statistic using data from Advanced LIGO’s first observing run. We explain the meaning of the PyCBC search’s false alarm rate and detection statistics, and connect these to the traditional matched filter signal-to-noise ratio and to the source’s luminosity distance. Using simulated signals we show that the improvements described here yield an $\sim 10\%$ increase in the detection rate of binary neutron stars and an $\sim 20\%$ increase in the detection rate of neutron star–black hole binaries compared to the search used in the first Advanced LIGO observing run (Abbott et al. 2016b, 2016d). These techniques are currently in use in binary-merger searches in Advanced LIGO’s second observing run and was used to evaluate events such as GW170104 (Abbott et al. 2017b). Finally, we propose an additional improvement—not yet implemented—that can increase the binary black hole detection rate by $\sim 30\%$ beyond the search used in Advanced LIGO’s first observing run, without negatively impacting the sensitivity for lower-mass mergers.

2. Measuring Event Significance

If the LIGO detector noise was stationary and Gaussian, a signal time-of-arrival test would suffice to identify candidate events in the detector network, and the quadrature sum of the matched filter signal-to-noise ratio ρ of an event in each Advanced LIGO detector would be an effective detection statistic (Wainstein & Zubakov 1962; Cutler et al. 1993; Pai et al. 2001; Allen et al. 2012). In this simple case, we could construct an analytic model of how often the detector background noise produces an event that is as loud as a given event candidate (Finn 1992; Finn & Chernoff 1993); this is the search’s false alarm rate. In reality, noise transients and other non-stationarity in the detectors produce excursions in the signal-to-noise ratio that can mimic astrophysical events and prevent us from analytically modeling the search’s false alarm rate.

To suppress transient noise, the PyCBC search uses a χ^2 test that checks that the time-frequency evolution of an event

matches that of the target waveform (Allen 2005). This test is constructed so that $\chi^2 \sim 1$ for a real signal and is larger for transient noise. If $\chi^2 \geq 1$ for a trigger, we down-weight the signal-to-noise ratio according to $\hat{\rho} = \rho/[1 + (\chi^2)^3]^{1/6}$ (Babak et al. 2013). The detection statistic used in Advanced LIGO’s first observing run was the quadrature sum of the re-weighted signal-to-noise ratio $\hat{\rho}$ observed in each detector, denoted $\hat{\rho}_c$ (Abbott et al. 2016d).

The non-stationarity of the detectors’ noise means that it is impossible to construct an analytic model of the search’s false alarm rate. We also cannot shield the detectors from gravitational waves and directly measure the signal-free background noise. To “turn off the signals” we artificially shift the timestamps of one detector’s data by an offset that is large compared to the gravitational-wave propagation time between the observatories and produce a new set of coincident events from this time-shifted data set. Real signals will no longer be time-coincident, and the remaining events will be produced by the detector noise alone. This process is repeated with many different time shifts to measure the search’s false alarm rate. Under the assumption that signals are relatively rare and that instrumental noise is uncorrelated in time between detectors, this is an effective way to measure the search’s false alarm rate.

In practice, the search background varies across the target signal space. Higher-mass compact-object mergers produce shorter waveforms, which look more like detector glitches; thus, the χ^2 test is less effective at rejecting noise transients in this region of the signal space. This means that the search’s false alarm rate at a given $\hat{\rho}_c$ value is not the same for binary neutron stars as it is for binary black holes. In Advanced LIGO’s first observing run, candidate and background events were divided into three search classes based on template length. The significance of any given event was determined by comparing its detection statistic value to the false alarm rate in its own class; for a candidate seen in two or more classes, the significance reported is the maximum of the three possible values from the three classes. To account for searching over multiple classes, the false alarm rate is increased by a trials factor equal to the number of classes (Lyons 2008).

Figure 1 (left) shows the false alarm rate of the PyCBC search as a function of the detection statistics considered in this paper, as measured using data from Advanced LIGO’s first observing run. The dotted lines in Figure 1 show the false alarm rate as a function of $\hat{\rho}_c$, i.e., the detection statistic used in the PyCBC offline search of Abbott et al. (2016b, 2016f, 2016h, 2016k). The lower (orange) dotted line shows the false alarm rate for binary neutron star sources (defined as mergers with total mass $M = m_1 + m_2 \leq 4 M_\odot$), and the upper (green) line shows the false alarm rate for binary mergers with total mass $4 \leq M = m_1 + m_2 \leq 100$ and a gravitational-wave frequency $f_{\text{peak}} \geq 100$ Hz at the peak amplitude of the template (Abbott et al. 2016b); this class includes both binary black hole and neutron star–black hole sources. The search’s false alarm rate can be understood as follows. For a given astrophysical source, the signal-to-noise ratio scales linearly as the inverse of the distance $\rho \propto 1/d_L$. Since, by construction, $\hat{\rho} \sim \rho$ for real signals, it follows that $\hat{\rho} \propto 1/d_L$ for real signals. As d_L decreases, the detection statistic increases and the false alarm rate is reduced according to the curves shown in Figure 1; the event is therefore more significant.




Figure 1. Left: the PyCBC search’s false alarm rate measured using data from Advanced LIGO’s first observing run. The false alarm rate is shown as a function of the detection statistic used by the search. The dotted lines show the search’s false alarm rate as a function of the combined re-weighted signal-to-noise ratio $\hat{\rho}$ used to detect the binary mergers in Advanced LIGO’s first observing run. The orange line shows this false alarm rate for binary neutron star sources ($M \leq 4 M_{\odot}$) and the green line for binary black holes and neutron star–black hole sources ($M \geq 4 M_{\odot}$). The solid blue line shows the false alarm rate of the offline PyCBC search using the new signal- and noise-weighted detection statistic ϱ proposed here. The solid gray line shows the false alarm rate of the low-latency PyCBC search, designed to rapidly identify triggers for electromagnetic follow-up, at a typical point in time during the first observing run. Note that this curve will vary over time due to changing noise characteristics. The low-latency search currently uses the $\hat{\rho}$ detection statistic of Equation (2), which suppresses triggers with suboptimal values of $p^5(\theta)$, and so will have a reduced detection statistic value at a fixed false alarm rate. Right: using simulated merger signals in data from Advanced LIGO’s first observing run, we construct a map between the new detection statistic ϱ and the combined re-weighted signal-to-noise ratio $\hat{\rho}$. The inset shows this map for binary neutron star sources. For these signals, the $\hat{\rho}$, and hence ϱ , scale approximately linearly with the inverse luminosity distance to the source. Together, these figures show that our new detection statistic gives around one order of magnitude improvement in the significance of a binary neutron star merger at a given luminosity distance. The improvement for binary black holes depends on the masses of the compact objects with higher-mass systems showing less improvement, as indicated by the scatter below the diagonal line seen in the right-hand figure.

The solid lines of Figure 1 (left) show the search false alarm rate for the new detection statistics proposed here measured using the same data from Advanced LIGO’s first observing run. The dark (blue) line shows the full statistic used by the PyCBC offline search and the light (gray) line shows the false alarm rate of the PyCBC low-latency search. These results differ as the offline search uses a large amount of data to measure the noise background, whereas the low-latency search only uses the previous five hours of data. Consequently, the offline search can measure false alarm rates to a precision of better than 1 in 10,000 years, whereas the low-latency search measures the false alarm rate to 1 in 100 years. This is sufficient to identify events for electromagnetic follow-up, with the final event significance measured by the offline search. The low-latency search also uses a simplification of the full detection statistic proposed here, as described in Section 3 below.

A key result of Figure 1 is that for all of these searches the false alarm rate is an extremely steep function of the detection statistic value. To illustrate the effect of the steep change in false alarm rate, consider a hypothetical binary neutron star source in Advanced LIGO’s first observing run. The signal-to-noise ratio of the source depends on the orientation and sky location of the source. Suppose a source was oriented so that it produced a signal-to-noise ratio of 7 in each LIGO detector. In the first observing run, the average luminosity distance to such a source would be $d_L \sim 85$ Mpc, but it could be as far away as $d_L \sim 194$ Mpc, if it were favorably oriented. The combined detection statistic for this source would be $\hat{\rho}_c = 9.9$. For a binary neutron star signal this loud, the search’s false alarm rate is $\sim 10^{-2} \text{ yr}^{-1}$, meaning that noise is expected to produce an event with $\hat{\rho}_c \geq 9.9$ in the $M \leq 4 M_{\odot}$ class once every ~ 300 years, which becomes once every 100 years after accounting for the trials factor. If the source’s luminosity distance was increased by 10%, then $\hat{\rho}_c$ decreases to 9: at this statistic value,

the search’s false alarm rate is $\sim 10 \text{ yr}^{-1}$. Since the noise background drops very quickly as a function of signal amplitude, the edge of the observable signal space is very sharp in signal amplitude.

Figure 1 also shows that small changes in the detection statistic value due to noise can have a large effect on the measured false alarm rate of a signal. This can cause two sources with similar strain amplitudes to have significant differences between their false alarm rates due to fluctuations in the detector noise at the time of each event. A change of half a unit in signal-to-noise ratio can result in an order-of-magnitude difference in the measured false alarm rate of a signal. Similarly, slight differences in implementation between low-latency and offline analyses (e.g., in the computation of the noise power spectral density, or in the number of χ^2 -veto bins) can cause substantial differences in false alarm rate. Highly significant signals, for example GW150914 and GW151226, will be robust to such changes, but the false alarm rate of marginal signals may change significantly.

3. Improving Search Sensitivity

Previously, the PyCBC search used four quantities to construct its detection statistic: the matched filter signal-to-noise ratio and the value of the χ^2 test in each of the two detectors. However, the search also records the difference in time of arrival between the LIGO Hanford and Livingston observatories $\delta t = t_H - t_L$, and the difference in the phase of the gravitational waveform $\delta\phi = \phi_H - \phi_L$. These time and phase differences, as well as the signal-to-noise ratio in each detector $\rho_{H,L}$, depend on the location of the source relative to the detectors. Therefore, an astrophysical population of signals will have a nontrivial distribution of events over the parameters $(\rho_H, \rho_L, \delta\phi, \delta t)$. However, under the assumption that noise is

not correlated between the LIGO Hanford and Livingston detectors, events due to background noise will be distributed uniformly in $(\delta\phi, \delta t)$; this information can be used to improve the detection statistic.

We can also use the measured noise background as a function of the template parameters to account for the variation of the search background over the target signal space. In the search in LIGO's first observing run, the search space divided the signal space into three classes: this provides a crude accounting of the mass-dependent variation in the search background. Here, we show that it is possible to construct a better model of the noise background over the search space, as a function of the masses and spins (m_1, m_2, s_1, s_2) of the merging compact objects.

We describe a new detection statistic that uses the full information $\theta = (\rho_H, \rho_L, \chi_H^2, \chi_L^2, \delta\phi, \delta t, m_1, m_2, s_1, s_2)$ for a candidate event that significantly improves the sensitivity of the PyCBC search. We construct this statistic by approximating the probability densities of both signal events and noise events over these parameters and forming the *ratio of densities*. This is the equivalent of the full likelihood ratio for a reduced data set consisting of only the signal-to-noise ratio maxima. A general detection statistic resulting from approximations to the density of signal $p^S(\theta)$ and noise $p^N(\theta)$ over the parameter space θ of coincident events can be written as⁷

$$\rho^2 \propto 2[\log p^S(\theta) - \log p^N(\theta)] + \text{constant}. \quad (1)$$

We choose this form for the detection statistic since in the case that the detectors' noise was stationary and Gaussian, we have $p^N(\theta) \sim \exp[-(\rho_H^2 + \rho_L^2)/2]$ and we recover the standard quadrature sum signal-to-noise ratio statistic.

We first consider the dependence on $\delta\phi$ and δt . For a single detector, the expected distribution of signals over ϕ, t is uniform, and thus does not aid in separating signals from noise. In a multi-detector network, the phase and time *differences* between detectors are determined by a source's sky location and orientation in relation to the detector locations. However, the distribution of noise events will be uniform in δt and $\delta\phi$. To use this information, we must first construct the probability distribution in this parameter space for coincident signals from an astrophysical population. We perform a Monte-Carlo simulation with a population of binary mergers that is uniform in spatial volume and isotropic over the sky and over source orientation. We use the antenna patterns of the detectors located at the LIGO Hanford and Livingston Observatories to calculate the expected signal-to-noise ratio of a source in each detector. If the single detector signal-to-noise ratio falls below a threshold imposed in the search (a value of $\rho = 5.5$ has been used in recent searches), the source is considered not detectable. For each of the remaining sources, we record δt , $\delta\phi$, ρ_H , and ρ_L , building up a multi-dimensional distribution. We account for a possible difference in sensitivity between the detectors by repeating the simulation over a range of relative detector sensitivities. The resulting multi-dimensional histogram is then smoothed with a Gaussian kernel whose width is determined by the expected measurement errors in each parameter, to obtain our approximation of the signal distribution $p^S(\theta)$. This Monte-Carlo simulation is performed once for




Figure 2. Unnormalized distribution of phase differences and time delays between the LIGO Hanford and Livingston Observatories for a simulated population of sources isotropically distributed over the sky and uniformly distributed in space. Noise events have a uniform distribution of time delays and phase differences. The events from LIGO's first observing run are overlaid; they are consistent with the signal distribution.

each detector network and stored for future use as an efficient look-up table.

In Figure 2, we show the resulting (unnormalized) signal distribution over the two most important parameters, the Livingston–Hanford time delay and phase difference, having marginalized over the parameters not shown. While the expected background distribution is uniform over δt and $\delta\phi$, the signal distribution clearly is not. This information about $p^S(\theta)$ can improve the separation of signal from noise events and thus increase search sensitivity. The peak of the phase difference around π is due to the relative positions and orientations of the two LIGO detectors, which are close to coplanar and have a $\sim 90^\circ$ relative rotation angle of the interferometer arms. For most sources they observe approximately the same polarization with opposite sign of the gravitational-wave strain. However, for arrival directions close to the line joining the two detectors, depending on the polarization, there may be partial cancellation or reversal of sign in the detector responses allowing a full range of relative phases between 0 and 2π . The phase difference distribution will, in general, be different for other combinations of observatories. We can construct an improved detection statistic using this information, defined as

$$\tilde{\rho}^2 = \hat{\rho}_c^2 + 2 \log \left(\frac{p^S(\theta)}{p_{\max}^S} \right), \quad (2)$$

where $p_{\max}^S$ is the most likely (highest) value in the multi-dimensional histogram. This is the detection statistic used by the PyCBC low-latency search (A. H. Nitz et al. 2017, in preparation).

So far, we are still considering an idealized noise distribution, $p^N \sim \exp[-(\hat{\rho}_L^2 + \hat{\rho}_H^2)/2]$, which neglects how the background may vary with the target waveform parameters and/or between detectors. However, the noise distributions in real data are quite different across different templates and between different detectors (Abbott et al. 2016d). The previous method of dividing the search into separate background classes has the drawbacks of requiring a somewhat arbitrary choice of

⁷ We may rescale the detection statistic by a constant (positive) factor and add a constant without affecting the relative ranking of events.

class boundaries; failing to model background variation within a class, thus potentially hurting sensitivity to some types of signal; and, if a large number of classes are used, a high trials factor diluting the significance of an event in the search as a whole compared to an estimate from a single class.

We obtain a more accurate and flexible description over the whole search space via an explicit model of the rate of occurrence of single-detector signal-to-noise ratio maxima due to noise:

$$\lambda_{Aw}^N(\hat{\rho}_A) = \mu_{Aw} \alpha_{Aw} \exp[-\alpha_{Aw}(\hat{\rho}_A - \hat{\rho}_{th})], \quad (3)$$

where μ_{Aw}^N is the total noise event rate in a given template w and detector A having $\hat{\rho}$ above a threshold value $\hat{\rho}_{th}$, and α_{Aw} describes how steeply the noise distribution falls off. The values of μ and α in each template and detector are empirically determined by a two-step procedure: first, a maximum-likelihood fit is performed to the signal-to-noise ratio maxima in each template and detector, which are assumed to be dominated by noise. The distribution at large $\hat{\rho}$ values is likely to be dominated by signals; hence, when performing the fit we remove a predetermined number of the highest- $\hat{\rho}$ maxima. Then to reduce variance due to small number statistics, we perform kernel smoothing of the λ and μ values, in each detector, over the duration of the template in the LIGO sensitive band τ_w , since templates with similar τ have similar responses to transient noise artifacts.

We expect noise events to be uncorrelated in time between different detectors: thus, the expected distribution of coincident noise over $\hat{\rho}_{H,L}$ in a given template is the product of single-detector distributions. For a template of duration τ , we have

$$p^N(\theta) \propto \lambda_H^N(\hat{\rho}_H, \tau) \lambda_L^N(\hat{\rho}_L, \tau), \quad (4)$$

where each factor on the right-hand side is given by the exponential model of Equation (3) using the smoothed fit parameters $\mu_A(\tau)$, $\alpha_A(\tau)$ ($A = H, L$).

Using this approximation for the noise distribution in the general statistic Equation (1) has the effect of smoothly rescaling the statistic values of events depending on the steepness of the noise distributions in the template where the events appear. Figure 1 (right) shows the new statistic value for a set of simulated signal events added to O1 data. We use the freedom to rescale the statistic value by a constant factor to set the values for low-mass candidate events, recovered in templates with steep background distributions, to be similar to $\hat{\rho}_c$. Then given the elevated rate of high- $\hat{\rho}$ noise events in higher-mass templates, the statistic values are suppressed for more massive candidates. In Figure 1 (left), we see that the background classes used in the first observing run are combined into a single noise distribution, with a noise level significantly suppressed relative to the search used in O1, even for binary neutron star templates where the χ^2 test is most effective at suppressing noise artifacts. This statistic is used in the PyCBC offline search in Advanced LIGO's second observing run.

We have not yet included the parameters of the waveform template (binary component masses and spins) in the signal part of our detection statistic. Not considering these parameters is only an optimal choice if the distribution of signal events over them, at a given ρ , is approximately equal to the distribution of waveform templates (Dent & Veitch 2014). In binary-merger searches, we use a discrete set of templates laid

out to bound the maximum possible loss of detection rate for any of our target sources (Brown et al. 2012; Ajith et al. 2014; Capano et al. 2016). The density of noise events over the parameter space of masses and spins varies as the density of templates: due to the way the binary parameters affect the waveform, this density varies by many orders of magnitude between low- and high-mass regions. The statistics described above then have an implicit prior that the rate of detectable signals is much higher in the low-mass region—although they certainly allow the detection of high-mass signals.

We can introduce an explicit model of the mass (and/or spin) distribution by including a term in the general statistic Equation (1) that depends on the binary parameters of the template waveform w . This term is the logarithmic ratio of the expected density of signals over the parameter space to the density of templates over the same space. Here we neglect statistical errors in the template parameters compared to the true parameters of the merging system; this approximation is likely to be valid for signal distributions that do not vary very abruptly over the binary system's parameters (Dent & Veitch 2014). As an example of a possible prior, we take a distribution of sources that is uniform over the logarithms of binary component masses; the ratio of template density to signal density is found via a simple kernel density estimate applied to the set of template $(\log m_1, \log m_2)$ values.

The relative rates of signals arriving at Earth from binary mergers of different masses are still highly uncertain (Abadie et al. 2010a), so applying a complex or fully realistic model prior is not appropriate. Note that a prior might be chosen to maximize either the total expected number of detections, or some other figure of merit for the results of a search.

4. Improvement in Detection Rate and Significance

In this section, we evaluate how the use of the improved detection statistics described in the previous sections affects the sensitivity of searches for binary neutron star, neutron star black hole, and binary black hole mergers. We simulate a population of sources and insert these into data from Advanced LIGO's first observing run. Using the PyCBC analysis to recover these sources, we determine which would be significant candidates, corresponding to a false alarm rate of less than 1 per 100 years. The volume of space that the search is sensitive to, which is proportional to the expected number of detections, is then obtained by Monte Carlo integration (Usman et al. 2016).

We simulate an isotropic distribution of sources in sky location, orientation, and spin angle, i.e., orientation of the binary component angular momenta relative to the orbital angular momentum; thus, the binaries have generically precessing orbital dynamics. The magnitude of the spin vector is restricted to <0.4 for sources with component masses less than $3 M_\odot$, and <0.989 otherwise. We consider a population of sources that is uniformly distributed in component mass. Compared to the detection statistic used in the first observing run (Abbott et al. 2016b), we show in Figure 3 how sensitivity is improved using our new statistics that take into account (left) the location and orientation of sources, (center) a more accurate background model, and (right) a prior over the source mass distribution uniform in $(\log m_1, \log m_2)$. We achieve an $\sim 10\%$ and $\sim 20\%$ increase in the number of detections of binary neutron star mergers and neutron star–black hole mergers, respectively.




Figure 3. Comparison of the expected rate of detections with the PyCBC pipeline using different detection statistics, estimated via a set of simulated signals added to Advanced LIGO data. Comparisons are relative to the statistic used in the first observing run. Left: with a statistic using information on source sky location and orientation sensitivity increases for all sources. Center: including also an improved model of the background distribution, the excess noise in higher-mass templates is taken into account and we see further sensitivity increases for BNS and NSBH sources. Right: using a mass prior uniform in the logarithm of component masses sensitivity is significantly increased for higher-mass sources.

An explicit prior on the mass distribution of detectable sources is not currently being employed in the analysis of Advanced LIGO data; however, if introduced, an $\sim 30\%$ improvement in the rate of detections for binary black hole mergers with masses similar to those identified in Abbott et al. (2016b) could be achieved while maintaining or improving upon the sensitivity to binary neutron star and neutron star–black hole mergers.

5. Conclusions

Although we have described relatively straightforward improvements to our search detection statistic, we may still not have achieved an optimal event ranking: various issues remain to be addressed by future work. We have approximated the signal and noise event rates as constant (stationary) over time for the duration of each experiment; however, despite the application of data-quality vetoes (Nuttall et al. 2015; Abbott et al. 2016c) to remove times of known environmental and instrumental disturbances from the search, we occasionally observe large upward fluctuations in the noise event rate on the scale of seconds to minutes due to non-stationary or non-Gaussian data. Search efficiency could be further improved by down-ranking these times via a model incorporating a time-dependent noise rate. Also, it may be possible to develop other diagnostics besides the χ^2 test (Allen 2005) to distinguish signals from noise transients in single-detector data.

The reader may ask what the effect these detection statistics have on the significance of the candidate events seen in O1. GW150914 and GW151226 have high enough SNR that their statistic values would lie above all available background samples for any of the statistics we present, so we can only place a bound on their significance in all cases. For the marginal event LVT151012, its phase difference, time difference, and amplitude parameters lie well within the expected signal distribution. An evaluation of its significance using the signal–noise model ranking statistic ρ would require a reanalysis of the O1 data, which is beyond the scope of this paper. To estimate the possible reduction in false alarm rate, we may use the expected signal distribution of Section 2 to find the relative likelihood of a signal event with a specific $\delta\phi$ and δt value as compared to a noise event. As in Figure 2, we marginalize over the other dimensions of the signal distribution, and note that noise events are uniformly distributed over

$\delta\phi$, δt . At the values measured for LVT151012, we find that the density (normalized pdf) of signal events is ~ 10 times higher than that of noise events, suggesting that if this information were included in the ranking statistic of the search the event would be assigned a false alarm rate lower by approximately an order of magnitude.

The significance of marginal events is also very sensitive to the relative weighting assigned to different regions of parameter space. As an example, LVT151012 would have been assigned a false alarm rate of less than 1 per 1000 years if the proposed logarithmic-in-component-mass prior was incorporated into the ranking statistic. We warn, though, that a nontrivial mass prior such as that proposed at the end of Section 3 is not an unbiased choice with respect to previously identified events such as LVT151012.

In this study, we have considered template waveforms for dominant-mode gravitational-wave emission from non-precessing (and quasi-circular) binaries. Such templates are in general not expected to have a uniformly optimal detection efficiency for more complex binary-merger signals: in particular, those including non-dominant modes (Capano et al. 2014), relevant for high-mass mergers with significantly unequal component masses, and those showing strong amplitude modulation as a result of spin-orbital precession (Harry et al. 2016, 2014). To best detect such signals, templates with more degrees of freedom may be required (Capano et al. 2014; Harry et al. 2014), implying that each candidate event may be described by numerous independent matched filter outputs. Optimizing the efficiency of searches over a larger event parameter space will likely require modeling the expected distributions of signal and noise over that space: techniques similar to those presented here will be applicable to this task.

The authors thank the LIGO Scientific Collaboration for access to the Advanced LIGO data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. We thank Christopher Biwer, Soumi De, Ian Harry, Greg Mendell, Frank Ohme, Peter Saulson, Bernard Schutz, and Alan Weinstein for

helpful discussions. T.D.C. is supported by an appointment to the NASA Postdoctoral Program at the Goddard Space Flight Center, administered by Universities Space Research Association under contract with NASA. D.A.B. acknowledges support from National Science Foundation awards PHY-1404395, ACI-1443047. D.A.B. performed part of this work at the Kavli Institute for Theoretical Physics, which is supported by NSF award PHY-1125915. S.F. acknowledges support from STFC awards ST/L000962/1. We thank the Atlas cluster computing team at AEI Hannover where this analysis was carried out.

ORCID iDs

Alexander H. Nitz  <https://orcid.org/0000-0002-1850-4587>
 Thomas Dent  <https://orcid.org/0000-0003-1354-7809>
 Tito Dal Canton  <https://orcid.org/0000-0001-5078-9044>
 Stephen Fairhurst  <https://orcid.org/0000-0001-8480-1961>
 Duncan A. Brown  <https://orcid.org/0000-0002-9180-5765>

References

- Abadie, J., Abbott, B. P., Abbott, R., et al. 2010a, *CQGra*, 27, 173001
 Abadie, J., Abbott, B. P., Abbott, R., et al. 2010b, *PhRvD*, 82, 102001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *ApJL*, 818, L22
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvX*, 6, 041015
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, *CQGra*, 33, 134001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016d, *PhRvD*, 93, 122003
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016e, *PhRvL*, 116, 131103
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016f, *PhRvL*, 116, 241103
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016g, *ApJL*, 826, L13
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016h, *PhRvL*, 116, 061102
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016i, *ApJS*, 225, 8
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016j, *ApJL*, 833, L1
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016k, *ApJL*, 832, L21
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016l, *LRR*, 19, 1
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *CQGra*, 34, 044001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *PhRvL*, 118, 221101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, *PhRvD*, 96, 022001
 Ajith, P., Fotopoulos, N., Privitera, S., Neunzert, A., & Weinstein, A. J. 2014, *PhRvD*, 89, 084041
 Allen, B. 2005, *PhRvD*, 71, 062001
 Allen, B., Anderson, W. G., Brady, P. R., Brown, D. A., & Creighton, J. D. E. 2012, *PhRvD*, 85, 122006
 Babak, S., Biswas, R., Brady, P., et al. 2013, *PhRvD*, 87, 024033
 Biswas, R., et al. 2012a, *PhRvD*, 85, 122009
 Biswas, R., et al. 2012b, *PhRvD*, 85, 122008
 Blanchet, L. 2002, *LRR*, 5, 3
 Bohé, A., et al. 2016, arXiv:1611.03703
 Brown, D. A., Harry, I., Lundgren, A., & Nitz, A. H. 2012, *PhRvD*, 86, 084017
 Buonanno, A., & Damour, T. 1999, *PhRvD*, 59, 084006
 Capano, C., Harry, I., Privitera, S., & Buonanno, A. 2016, *PhRvD*, 93, 124007
 Capano, C., Pan, Y., & Buonanno, A. 2014, *PhRvD*, 89, 102003
 Cutler, C., Apostolatos, T. A., Bildsten, L., et al. 1993, *PhRvL*, 70, 2984
 Dal Canton, T., Nitz, A. H., Lundgren, A. P., et al. 2014, *PhRvD*, 90, 082004
 Dent, T., & Veitch, J. 2014, *PhRvD*, 89, 062002
 Faye, G., Marsat, S., Blanchet, L., & Iyer, B. R. 2012, *CQGra*, 29, 175004
 Finn, L. S. 1992, *PhRvD*, 46, 5236
 Finn, L. S., & Chernoff, D. F. 1993, *PhRvD*, 47, 2198
 Harry, I., Privitera, S., Bohé, A., & Buonanno, A. 2016, *PhRvD*, 94, 024012
 Harry, I. W., Nitz, A. H., Brown, D. A., et al. 2014, *PhRvD*, 89, 024010
 Klimentenko, S., et al. 2016, *PhRvD*, 93, 042004
 Lackey, B. D., & Wade, L. 2015, *PhRvD*, 91, 043002
 Lyons, L. 2008, *AnApS*, 2, 887
 Martynov, D. V., et al. 2016, *PhRvD*, 93, 112004
 Messick, C., Blackburn, K., Brady, P., et al. 2017, *PhRvD*, 95, 042001
 Metzger, B. D., & Berger, E. 2012, *ApJ*, 746, 48
 Nitz, A. H., Harry, I., Biwer, C. M., et al. 2017, PyCBC Software, Zenodo, doi:10.5281/zenodo.545845
 Nuttall, L. K., Massinger, T. J., Areeda, J., et al. 2015, *CQGra*, 32, 245005
 Pai, A., Dhurandhar, S., & Bose, S. 2001, *PhRvD*, 64, 042004
 Taracchini, A., Buonanno, A., Pan, Y., et al. 2014, *PhRvD*, 89, 061502
 Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, *CQGra*, 33, 215004
 Wainstein, L. A., & Zubakov, V. D. 1962, *Extraction of Signals from Noise* (Englewood Cliffs, NJ: Prentice-Hall)
 Wiener, N. 1949, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (New York: Wiley)