

Fast evaluation of multi-detector consistency for real-time gravitational wave searches

Chad Hanna,^{1,2,3,4} Sarah Caudill,^{5,6} Cody Messick,^{1,3} Surabhi Sachdev,^{1,3,7} Kipp Cannon,^{8,9} Kent Blackburn,⁷ Jolien D. E. Creighton,⁵ Heather Fong,⁸ Patrick Godwin,^{1,3} Shasvath Kapadia,⁵ Tjonnie G. F. Li,¹⁰ Ryan Magee,^{1,3} Duncan Meacher,^{1,3} Debnandini Mukherjee,⁵ Alex Pace,^{1,3} Stephen Privitera,¹¹ Rico K. L. Lo,^{10,7} and Leslie Wade¹²

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

²*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA*

³*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA*

⁴*Institute for CyberScience, The Pennsylvania State University, University Park, PA 16802, USA*

⁵*Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA*

⁶*Nikhef, Science Park, 1098 XG Amsterdam, Netherlands*

⁷*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

⁸*Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, Ontario, M5S 3H8, Canada*

⁹*RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan*

¹⁰*Department of Physics, The Chinese University of Hong Kong, Shatin, New Territories, Hong Kong*

¹¹*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany*

¹²*Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA*

(Dated: January 9, 2019)

Gravitational waves searches for compact binary mergers with LIGO and Virgo are presently a two stage process. First, a gravitational wave signal is identified. Then, an exhaustive search over signal parameters is performed. The identification stage must be efficient in order to maximize the number of gravitational wave sources that are identified. Initial identification of gravitational wave signals with LIGO and Virgo happens in real-time, which requires that less than one second of computational time must be used for each second of gravitational wave data collected. In contrast, subsequent parameter estimation may require hundreds of hours of computational time to analyze the same one second of gravitational wave data. The real-time identification requirement necessitates efficient and often approximate methods for signal analysis. We describe one piece of a real-time gravitational-wave identification algorithm: an efficient method for ascertaining a signal's consistency between multiple gravitational wave detectors. This technique was used in analyses of Advanced LIGO's second observing run and Advanced Virgo's first observing run.

I. INTRODUCTION

Advanced LIGO [1, 2] and Virgo [3] gravitational-wave observatories have seen great success since the discovery of gravitational waves in September 2015 [4]. With the detection of gravitational waves from ten binary black holes and one binary neutron star in just 170 days of observing, gravitational wave detections are now happening at a rate of about two per month [5]. The rate is expected to increase by more than an order of magnitude as the world-wide network of gravitational wave detectors improve in sensitivity [6]. Furthermore, over the next decade, the gravitational wave detector network will grow to include KAGRA [7, 8] in Japan and LIGO-India [9]. It is anticipated that real-time gravitational wave searches will use data from the entire gravitational wave detector network in order to increase the detection rate and better localize the gravitational wave sources [6].

The current generation of interferometric gravitational wave detectors measure only one of two possible gravitational wave polarizations h_+ and h_\times . However, given the detector orientation with respect to incoming signals, they each measure a different linear combination of polarizations defined by a suitable Earth-centered coordinate system. Furthermore, the response to the gravitational

wave signal changes as a function of the incident angle. While the detector response function is broad, some incident angles produce no response in a given gravitational wave detector. Thus, each detector's response to a given signal is a function of the location and polarization state of a gravitational wave source [1, 3, 10]. Given a particular gravitational wave signal incident on Earth, it is possible to predict precisely what the projection of the signal will be on each gravitational wave detectors [11]. A corollary of this is that only certain arrival times, phases and amplitudes measured in a set of gravitational wave detectors are consistent with a real gravitational wave signal.

Typically, gravitational wave searches have relied on one of several mechanisms for imposing consistency of gravitational wave detection among a network of detectors. Each method has a varying degree of fidelity traded against complexity or computational cost. The ideal method requires explicitly evaluating the likelihood of measuring the projected waveform in each detector while sampling over physically reasonable prior distributions, e.g., isotropically distributed on the sky, using Markov Chain Monte Carlo techniques [12]. This is the method used in parameter estimation as the second stage of gravitational wave astronomy. While regarded as the most accurate assessment of consistency among gravita-

tional wave detectors, the full likelihood evaluation has the practical limitation of being a computationally costly procedure. However, approximate methods have been developed [13], that are used to e.g., quickly compute the sky location of gravitational wave events. The approximate methods are orders of magnitude faster and still extremely accurate. Nevertheless, the approximate methods used to compute sky maps are not suitably fast to use for the initial identification stage of gravitational wave events.

Another widely adopted mechanism to ensure consistency between gravitational wave observatories is to form a coherent combination of the detector data [14, 15] in order to create a sky location dependent reconstruction of the two gravitational wave polarizations. This technique typically uses a predefined grid of sky positions. Coherence is used presently in searches for gravitational wave bursts [16, 17] and triggered searches for gamma ray bursts [18, 19] which only focus on specific time intervals and sky regions that overlap with known gamma ray bursts.

The final mechanism involves searching each detector data stream independently, triggering on potential signal times by finding peaks in filter output over some time interval, and then comparing the triggers between each detector to see if they are consistent with a signal [20–22]. The simplest implementations of this procedure impose that triggers occur within the gravitational wave propagation time between detectors, but don't otherwise place constraints on the amplitude or phase measured in each detector. However, not all gravitational-wave arrival times are probable or even possible when considering a network of gravitational wave detectors. Furthermore, arrival times are strongly correlated with the measured amplitude and phase of the gravitational wave signals. Ignoring these effects means that one increases the background (false-positive rate) of a given gravitational-wave search by considering unphysical gravitational wave signals.

Work to check consistency of the distribution of measured amplitude, time and phase of gravitational wave triggers from a network of detectors has been explored in several contexts. The early work used only one-dimensional marginalized distributions for these parameters [23], which did not accurately describe the joint probability for the network but was nevertheless demonstrably more effective than simply using coarse windows in arrival time to define coincidences. A joint probability distribution of signal amplitudes for the LIGO and Virgo network was developed for compact binary searches in [24] and used during advanced LIGO's first observing run, however time of arrival and phase parameters were ignored. Explicitly constructing probabilities for time, phase and amplitude for the two LIGO detectors and using this information in ranking of candidate events was explored in the second Advanced LIGO observing run [25].

In this work, we present a method for accounting for

signal consistency in compact binary searches across a network of gravitational wave detectors that is computationally efficient and will scale to the eventual 5 detector gravitational wave network. This is accomplished through a factorization of the likelihood ratio (LR) ranking statistic first described in [24] and used in the analysis of candidates from the compact binary coalescence search in the first observing run O1 [4, 22, 26–29] that generalizes to multiple gravitational wave detectors, and which takes into account the correlated trigger amplitudes, times and phases between multiple gravitational wave observatories. While the main objective of this work is to document a novel gravitational wave detection algorithm, it also serves as a reference for methods used in LIGO and Virgo searches in Advanced LIGO's second observing run and advanced Virgo's first observing run [5].

II. THE LIKELIHOOD RATIO

In this section we describe the likelihood ratio (LR) ranking statistic as implemented in the GstLAL inspiral pipeline [22, 26, 30] with an emphasis on terms relevant to this particular discussion. In brief, the GstLAL inspiral pipeline analyzes data from multiple gravitational wave detectors with time-domain matched filtering [31] over a collection of gravitational wave templates. First, the peaks in each template filter output are identified for each detector. Then, these peaks, which are called triggers, are combined to look for coincident triggers in the gravitational wave network. If triggers from the same template are within the gravitational wave propagation time between observatories, the collection of triggers is called a coincident event [32].

Coincident events are common and the vast majority are not gravitational waves. This is due to the fact that the threshold in each detector for identifying a trigger is very low leading to potentially billions of triggers per analysis and tens of millions found in coincidence. Therefore, we have to rank the collection of coincident events from least to most probable of being a gravitational wave. To do this, we explicitly evaluate the LR over the trigger parameters. We consider a LR ranking statistic for gravitational-wave candidates defined as ¹

$$\mathcal{L} = \frac{P(\vec{D}_H, \vec{O}, \vec{\rho}, \xi^2, \vec{\phi}, \vec{t} | \mathbf{s})}{P(\vec{D}_H, \vec{O}, \vec{\rho}, \xi^2, \vec{\phi}, \vec{t} | \mathbf{n})}. \quad (1)$$

Each vector of parameters is used to denote detector specific information. To be concrete we will assume that LIGO Hanford, H , LIGO Livingston, L , and Virgo, V are all in observing mode and being analyzed. However, we note that the method described can be generalized to

¹ This only includes terms relevant to this paper.

add more detectors. \vec{D}_H is a vector of horizon distances for each observatory $D_H = \{D_{H1}, D_{L1}, D_{V1}\}$, which accounts for how sensitive the detectors are at the time of the event. \vec{O} is the set of detectors that observed the event in coincidence. Since it is possible for a gravitational wave to not register in one or more detectors above threshold, \vec{O} will take on one of six possible values in this case: $\vec{O} \in \{\{H1, L1, V1\}, \{H1, L1\}, \{H1, V1\}, \{L1, V1\}, \{H1\}, \{L1\}, \{V1\}\}$. $\vec{\rho}$ is the vector individual signal-to-noise ratios (SNRs) in each detector, e.g., $\vec{\rho} = \{\rho_{H1}, \rho_{L1}, \rho_{V1}\}$ and $\vec{\xi}^2$ is the vector of ξ^2 -signal-based-veto values for each detector (described in [22]), e.g., $\vec{\xi}^2 = \{\xi_{H1}^2, \xi_{L1}^2, \xi_{V1}^2\}$. Likewise, \vec{t} and $\vec{\phi}$ are the time and phases measured at each gravitational wave detector [33], $\vec{t} = \{t_H, t_L, t_V\}$ and $\vec{\phi} = \{\phi_H, \phi_L, \phi_V\}$.

The independence between detectors for noise events implies that the denominator of (1), $P(\dots | n)$, factors to the product of one-dimensional and two-dimensional distributions and is therefore generally easy to estimate and evaluate numerically. The numerator, $P(\dots | s)$ is not easily factorable. Reference [24] outlines the factorization of Eq. 1 without the $\vec{\Delta}\phi$ and $\vec{\Delta}t$ terms, though we point out that it does include the joint distribution of SNR. This form was used in the ranking of gravitational-wave candidates from the first observing run, O1, during which only the H1 and L1 advanced LIGO detectors were operating [4, 28]. Here we extend the technique used in [24] and propose the following factorization of the numerator,

$$\begin{aligned} P(\dots | s) &= P(\vec{D}_H | s) \\ &\times P(\vec{O} | \vec{D}_H, s) \\ &\times P(\vec{\xi}^2 | \vec{\rho}, s) \\ &\times P(\vec{\rho}, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s), \end{aligned} \quad (2)$$

where we assume that the distribution of the ξ^2 -signal-based-veto values are independent of $\vec{\phi}$ and \vec{t} but not $\vec{\rho}$.

In this work we are concerned with an approximation to the final term: $P(\vec{\rho}, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s)$, which accounts for the amplitude, time, and phase measured in each gravitational wave detector. We begin by isolating an overall term that scales with the detector network SNR, $|\vec{\rho}|$, to the negative fourth power [34] so that

$$\begin{aligned} P(\vec{\rho}, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s) &\approx \\ P(\vec{\rho}/|\vec{\rho}|, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s) &\times |\vec{\rho}|^{-4} \end{aligned} \quad (3)$$

This approximation ignores the fact that the accuracy of the measured $\vec{\phi}$ and \vec{t} depend on the SNR, $\vec{\rho}$. In other words, at low SNR the distribution of time and phase is broad and at high SNR it is narrow. However, since it is most critical for the detection process to catch signals that are near threshold, we assume that the uncertainty in \vec{t} and $\vec{\phi}$ take on values consistent with a network SNR ~ 10 signal. Our goal is to adequately describe

the $P(\vec{\rho}/|\vec{\rho}|, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s)$ when the network SNR is about 10.

Next, we note that the absolute time and phase of the signal is arbitrary. Assuming isotropically distributed gravitational waves, the distribution of time and phase at a given detector is uniform. However, the relative times and phases between detectors make up a nontrivial correlated probability density function. This suggests that we can reduce the dimensionality by computing parameters relative to a fiducial instrument either by difference or ratio. We pick the first instrument in alphabetical order (denoted with 0 subscript). Furthermore, since $P(\vec{\rho}/|\vec{\rho}|, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s)$ is conditional on the detector horizon distances, we can evaluate the effective distances, \vec{D}_{eff} [33] instead of $\vec{\rho}$. The advantage to switching to effective distances, is that the distribution of effective distance ratios is constant. Thus we have,

$$\begin{aligned} P(\vec{\rho}, \vec{\phi}, \vec{t} | \vec{O}, \vec{D}_H, s) &\approx \\ P(\Delta\vec{D}_{\text{eff}}, \Delta\vec{\phi}, \Delta\vec{t} | \vec{O}, s) &\times |\vec{\rho}|^{-4}, \end{aligned} \quad (4)$$

where concretely for the H, L, V network

$$\begin{aligned} \Delta\vec{D}_{\text{eff}} &\equiv \{d_L/d_H, d_V/d_H\}, \\ \Delta\vec{t} &\equiv \{t_L - t_H, t_V - t_H\}, \\ \Delta\vec{\phi} &\equiv \{\phi_L - \phi_H, \phi_V - \phi_H\}. \end{aligned}$$

In interest of simplifying the subsequent notation, we define a single vector of parameters that contains $\Delta\vec{D}_{\text{eff}}$, $\Delta\vec{t}$ and $\Delta\vec{\phi}$,

$$\vec{\lambda} \equiv \{\Delta\vec{D}_{\text{eff}}, \Delta\vec{t}, \Delta\vec{\phi}\}. \quad (5)$$

In order to construct the distribution of these parameters for a signal, we assert that gravitational waves have a uniform distribution in Earth-based coordinates: right ascension α , declination $\cos(\delta)$, inclination angle $\cos(\iota)$, and polarization angle ψ . We lay down a uniform, densely sampled grid in $\{\alpha, \cos(\delta), \cos(\iota), \psi\}$ and further assert that any signal should “exactly” land on one of the grid points. We transform that regularly-spaced grid into a grid of irregularly spaced points in $\vec{\lambda}$, which we denote as $\vec{\lambda}_{\text{mi}}$ for the i^{th} model vector. We consider that the only mechanism to push a signal away from one of these exact i^{th} grid points is Gaussian noise. Furthermore we assume the probability distribution is of the form:

$$\begin{aligned} P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{\text{mi}}) &= \\ \frac{1}{\sqrt{(2\pi)^k |\mathbf{\Sigma}|}} \exp \left[-\frac{1}{2} \Delta\vec{\lambda}_i^T \mathbf{\Sigma} \Delta\vec{\lambda}_i \right], \end{aligned} \quad (6)$$

where $\Delta\vec{\lambda}_i \equiv \vec{\lambda} - \vec{\lambda}_{\text{mi}}$. We make a further simplifying assumption that the the covariance matrix $\mathbf{\Sigma}$ is constant and diagonal. This is not a very good approximation for highly significant triggers, however, the goal of this work is to improve efficiency of detecting near threshold

triggers, since significant triggers are likely to be identified easily anyway. Thus, we choose a Σ tuned for near threshold triggers. Relaxing the assumption of a diagonal covariance matrix is future work and will likely make the approximation better. For the present work we will replace Σ with σ^2 to denote its diagonal status.

Next we will define a rescaled set of coordinates $\vec{x} \equiv \vec{\lambda}/\sigma$, and similarly $\vec{x}_{mi} \equiv \vec{\lambda}_{mi}/\sigma$ such that $\Delta\vec{x}_i \equiv (\vec{\lambda} - \vec{\lambda}_{mi})/\sigma$. Furthermore, for brevity we will drop any overall normalizations giving,

$$P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{mi}) \propto \exp\left[-\frac{1}{2}\Delta\vec{x}_i^2\right]. \quad (7)$$

Then we assert that

$$\begin{aligned} P(\vec{\lambda} | \vec{O}, s) &\propto \sum_i P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{mi}) P(\vec{\lambda}_{mi}) \\ &\propto \sum_i P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{mi}) \\ &= \sum_i \exp\left[-\frac{1}{2}\Delta\vec{x}_i^2\right], \end{aligned} \quad (8)$$

where the second line holds by construction since $p(\vec{\lambda}_{mi})$ does not depend on i as they were chosen uniform in prior signal probability. Computing probability (8) in real-time for each gravitational wave candidate is not feasible since the grid might have millions of points. Therefore, we make another simplifying assumption: we assume that the noise only adds a contribution which is orthogonal to the hyper-surface defined by the signal. In other words, we assume that noise cannot push a signal from one grid point towards another along the signal hyper-surface which implies,

$$\Delta\vec{x}_i^2 \approx \Delta\vec{x}_0^2 + g_{0i}^2, \quad (9)$$

where $\Delta\vec{x}_0$ refers to the distance between the candidate parameter and the *nearest-neighbor* grid point, and g_{0i} is the distance between the i th grid point and the nearest neighbor grid point. In this case we can simplify the marginalization step with a precomputation since

$$\begin{aligned} P(\vec{\lambda} | \vec{O}, s) &\approx \exp\left[-\frac{1}{2}\Delta\vec{x}_0^2\right] \\ &\times \sum_i \exp\left[-\frac{1}{2}g_{0i}^2\right]. \end{aligned} \quad (10)$$

The entire sum over i can be precomputed and stored. In order for this endeavor to be successful, we still need a fast way to find the nearest neighbor. We use SciPy KDTree to accomplish this [35]. Fig. 1 shows an example probability density function for time delays and phase differences between LIGO Hanford and Virgo computed from the above method with a comparison shown from a direct Monte Carlo. We note that each evaluation of (10) required less than a millisecond to compute on a modern CPU. Thus, this method is suitable for evaluating hundreds of candidates per second per CPU core in real-time.

III. RESULTS

To assess the effect of the procedure described in II, we conducted a simulation of synthetic signals and non-stationary noise in order to produce a Receiver Operating Characteristics (ROC) curve. We constructed synthetic signal and noise models. For the signal model, we assumed sources uniform in the volume of space and isotropically oriented. We constructed 100000 such signals and calculated their corresponding SNRs, arrival times, and phases in the LIGO Hanford and LIGO Livingston detectors using the LAL Simulation package [36]. For noise, we constructed 1×10^5 simulated glitches with SNRs that were independent in each of the LIGO Hanford and LIGO Livingston detectors given by an exponential distribution

$$P(\rho_H, \rho_L) = \frac{1}{6} \exp(-\frac{\rho_H}{6}) \frac{1}{6} \exp(-\frac{\rho_L}{6}), \quad (11)$$

where we placed an additional constraint that we only kept samples which had both $\rho_H \geq 4$ and $\rho_L \geq 4$. We chose arrival time differences that were uniform within the GW travel time between the two detectors and phase differences that were uniform between 0 and 2π .

Figure 2 shows the ROC curves for three cases: the SNR-only terms in the LR, which was used in initial advanced LIGO searches [4, 27, 28] (green); a previous implementation of time and phase consistency used for subsequent gravitational wave searches [26, 37–40] terms (orange); and finally the current implementation used in the production of the GWTC-1 gravitational wave catalog [5] (blue). Here we see that even in the two-detector situation, the implementation described in this work achieves an improvement.

IV. CONCLUSION

We have demonstrated a computationally efficient likelihood ratio statistic for use in gravitational wave searches for compact binary systems with LIGO and Virgo. The work here specifically presents a way to assess the probability that a given set of measurements independently made at each observatory in a worldwide network of observatories is consistent with a signal. This is useful for real-time identification of candidates where the latency of the computation is critical and the overall computational scale cannot be too large. Our method allows for computing the likelihood ratio for hundreds of candidates per second on a single modern CPU core. This method was used to produce results for advanced LIGO's second observing run and Advanced Virgo's first observing run with the GstLAL pipeline.

We note that this work could be improved by relaxing the condition that the covariance matrix is diagonal for $\{\Delta\vec{D}_{\text{eff}}, \Delta\vec{t}, \Delta\vec{\phi}\}$. A better approximation will exploit the correlations between time and phase uncertainty and should make the method more efficient.

V. ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation through PHY-1454389, OAC-1841480, PHY-1700765, and PHY-1607585. Funding for this project was provided by the Charles E. Kaufman Foundation of The Pittsburgh Foundation. We thank the LIGO Scientific Collaboration for input on this work. Specifically, CH would like to thank Patrick Brady for several illuminating discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science, and Economic Development, and by the Province

of Ontario through the Ministry of Research and Innovation. Computations for this research were performed on the Pennsylvania State Universitys Institute for Cyber-Science Advanced CyberInfrastructure (ICS-ACI). We are grateful for computational resources provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at the University of Wisconsin-Milwaukee and supported by National Science Foundation Grants PHY-1626190 and PHY-1700765. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. This paper has LIGO document number: P1800362.

-
- [1] G. M. Harry, L. S. Collaboration, *et al.*, *Classical and Quantum Gravity* **27**, 084006 (2010).
- [2] J. Aasi, B. Abbott, R. Abbott, T. Abbott, M. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Classical and quantum gravity* **32**, 074001 (2015).
- [3] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardín, *et al.*, *Classical and Quantum Gravity* **32**, 024001 (2014).
- [4] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Physical review letters* **116**, 061102 (2016).
- [5] B. P. Abbott *et al.*, arXiv preprint arXiv:1811.12907 (2018).
- [6] B. Abbott, R. Abbott, T. Abbott, *et al.*, *Living Reviews in Relativity* **21** (2018).
- [7] K. Somiya, *Classical and Quantum Gravity* **29**, 124007 (2012).
- [8] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, H. Yamamoto, K. Collaboration, *et al.*, *Physical Review D* **88**, 043007 (2013).
- [9] C. Unnikrishnan, *International Journal of Modern Physics D* **22**, 1341010 (2013).
- [10] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, S. Anderson, W. Anderson, M. Arain, *et al.*, *Reports on Progress in Physics* **72**, 076901 (2009).
- [11] W. G. Anderson, J. T. Whelan, P. R. Brady, J. D. Creighton, D. Chin, and K. Riles, .
- [12] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, *et al.*, *Physical Review D* **91**, 042003 (2015).
- [13] L. P. Singer and L. R. Price, *Physical Review D* **93**, 024013 (2016).
- [14] A. Pai, S. Bose, and S. Dhurandhar, *Classical and Quantum Gravity* **19**, 1477 (2002).
- [15] D. Macleod, I. Harry, and S. Fairhurst, *Physical Review D* **93**, 064004 (2016).
- [16] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, *Classical and Quantum Gravity* **25**, 114029 (2008).
- [17] P. J. Sutton, G. Jones, S. Chatterji, P. Kalmus, I. Leonor, S. Poprocki, J. Rollins, A. Searle, L. Stein, M. Tinto, *et al.*, *New Journal of Physics* **12**, 053034 (2010).
- [18] I. W. Harry and S. Fairhurst, *Physical Review D* **83**, 084002 (2011).
- [19] A. R. Williamson, C. Biwer, S. Fairhurst, I. Harry, E. Macdonald, D. Macleod, and V. Predoi, *Physical Review D* **90**, 122004 (2014).
- [20] S. Babak, R. Biswas, P. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T. Cokelaer, J. D. Creighton, T. Dent, *et al.*, *Physical Review D* **87**, 024033 (2013).
- [21] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst, *et al.*, *Classical and Quantum Gravity* **33**, 215004 (2016).
- [22] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. Creighton, R. Everett, *et al.*, *Physical Review D* **95**, 042001 (2017).
- [23] K. C. Cannon, *Classical and Quantum Gravity* **25**, 105024 (2008).
- [24] K. Cannon, C. Hanna, and J. Peoples, ArXiv e-prints (2015), arXiv:1504.04632 [astro-ph.IM].
- [25] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, *The Astrophysical Journal* **849**, 118 (2017).
- [26] S. Sachdev *et al.*, *in prep.*.
- [27] B. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Physical Review X* **6**, 041015 (2016).
- [28] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Physical review letters* **116**, 241103 (2016).
- [29] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Physical Review D* **93**, 122003 (2016).
- [30] LIGO Scientific Collaboration and Virgo Collaboration, “Gstlal,” (2018).
- [31] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Koppel, L. Liao, *et al.*, *The Astrophysical Journal* **748**, 136 (2012).

- [32] S. Privitera, S. R. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, *Physical Review D* **89**, 024003 (2014).
- [33] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, *Physical Review D* **85**, 122006 (2012).
- [34] B. F. Schutz, *Classical and Quantum Gravity* **28**, 125023 (2011).
- [35] E. Jones, T. Oliphant, and P. Peterson, (2014).
- [36] LIGO Scientific Collaboration and Virgo Collaboration, “Lalsuite: Lsc algorithm library suite,” (2018).
- [37] L. Scientific, B. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, *Physical Review Letters* **118**, 221101 (2017).
- [38] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, V. Adya, *et al.*, *The Astrophysical Journal Letters* **851**, L35 (2017).
- [39] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, V. Adya, *et al.*, *Physical review letters* **119**, 141101 (2017).
- [40] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, V. Adya, *et al.*, *Physical Review Letters* **119**, 161101 (2017).




FIG. 1. Example probability density function for the difference in gravitational wave arrival time and phase between LIGO Hanford and Virgo. Here we fix the horizon distances between the two detectors to be similar and show a slice of the probability density function where the SNR ratios are close to 1. The full two dimensional distribution (lower left) was computed using the method described in section II. The top and bottom right bar plots use the data from the two dimensional distribution marginalized over phase and time respectively. For comparison, we computed the marginal distributions via a direct Monte Carlo method to obtain the orange traces. The agreement is excellent, however, we note that the same error assumptions were used in the Monte Carlo method and that changing the assumptions about the covariance matrix will lead to poorer agreement. Improving the covariance matrix approximation is future work.

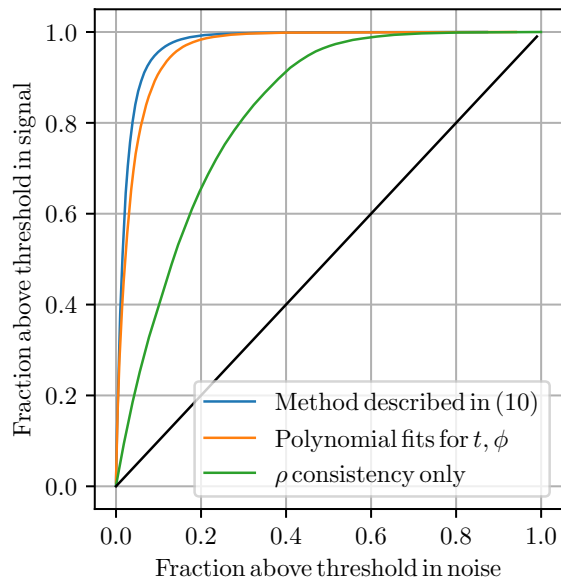


FIG. 2. ROC curve showing performance of the signal consistency check from the simulation described in III. The ROC curve plots the fraction of simulated events found above a given threshold of LR vs the fraction of noise found above that threshold. The black diagonal line would indicate that the method is not helpful at discriminating signals from noise. The green curve shows the likelihood ratio including only the amplitude consistency terms ($\hat{\rho}$), which was used in the first advanced LIGO observing run [24]. The orange curve shows the inclusion of $\hat{\Delta}t$ and $\hat{\Delta}\phi$ terms used in the beginning of advanced LIGO's second observing run [26] and the blue curve shows the method described in II used to produce results for advanced LIGO's second observing run and Advanced Virgo's first observing run [5].