

Optimal Search for an Astrophysical Gravitational-Wave Background

Rory Smith^{1,2,3,*} and Eric Thrane^{1,2,†}

¹*Monash Centre for Astrophysics, School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

²*OzGrav: The ARC Centre of Excellence for Gravitational-Wave Discovery, Monash University, VIC 3800, Australia*

³*LIGO, California Institute of Technology, Pasadena, California 91125, USA*



(Received 5 December 2017; revised manuscript received 23 February 2018; published 16 April 2018)

Roughly every 2–10 min, a pair of stellar-mass black holes merge somewhere in the Universe. A small fraction of these mergers are detected as individually resolvable gravitational-wave events by advanced detectors such as LIGO and Virgo. The rest contribute to a stochastic background. We derive the statistically optimal search strategy (producing minimum credible intervals) for a background of unresolved binaries. Our method applies Bayesian parameter estimation to all available data. Using Monte Carlo simulations, we demonstrate that the search is both “safe” and effective: it is not fooled by instrumental artifacts such as glitches and it recovers simulated stochastic signals without bias. Given realistic assumptions, we estimate that the search can detect the binary black hole background with about 1 day of design sensitivity data versus ≈ 40 months using the traditional cross-correlation search. This framework independently constrains the merger rate and black hole mass distribution, breaking a degeneracy present in the cross-correlation approach. The search provides a unified framework for population studies of compact binaries, which is cast in terms of hyperparameter estimation. We discuss a number of extensions and generalizations, including application to other sources (such as binary neutron stars and continuous-wave sources), simultaneous estimation of a continuous Gaussian background, and applications to pulsar timing.

DOI: [10.1103/PhysRevX.8.021019](https://doi.org/10.1103/PhysRevX.8.021019)

Subject Areas: Astrophysics, Cosmology, Gravitation

I. INTRODUCTION

Observations of gravitational waves from binary black hole mergers imply that stellar-mass black holes coalesce somewhere in the visible Universe every 223_{-115}^{+352} s [1]. Binary neutron stars merge every 13_{-9}^{+49} s [1]. The vast majority of these events are too distant to be individually resolved by the current generation of detectors. The most distant event yet observed, GW170104, was measured to have a redshift of $z = 0.18_{-0.07}^{+0.08}$ [2]. Nonetheless, unresolved compact binary mergers contribute to a stochastic background of gravitational waves, which may be detectable with current detectors [3]. Measuring the stochastic background from compact binaries has the potential to provide information about high-redshift binary black holes

and neutron stars, which complements observations of local mergers [4].

The stochastic background is typically characterized by the gravitational-wave energy density spectrum,

$$\Omega_{\text{gw}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\text{gw}}}{d \ln f}, \quad (1)$$

where $d\rho_{\text{gw}}$ is the energy density of gravitational waves between f and $f + df$ and ρ_c is the critical energy density for a flat universe [5]. Searches for the stochastic background seek to measure $\Omega_{\text{gw}}(f) > 0$. In the LIGO-Virgo band (10–2000 Hz), the best current limits on the gravitational-wave energy density spectrum are $\Omega_{\text{gw}}(f) < 1.7 \times 10^{-7}$ (95% confidence, measured in the band 20–86 Hz) [6,7].

To date, all LIGO and Virgo searches for the stochastic background have relied on the cross-correlation method described by Allen and Romano [5]. A similar cross-correlation technique is employed by pulsar timing arrays operating in the nanohertz band [8]. The cross-correlation method has two nice features. It is computationally cheap and it yields a statistically optimal (minimum-variance) measurement of $\Omega_{\text{gw}}(f)$ for the case of a persistent, Gaussian background.

* rory.smith@monash.edu

† eric.thrane@monash.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

A Gaussian background is characterized entirely by $\Omega_{\text{gw}}(f)$. The stochastic background from stellar-mass binary black holes is highly non-Gaussian; it is rare for different events to overlap in time [1] (within the advanced detector band). The stochastic backgrounds from binary neutron stars are quasi-Gaussian since the signals from individual events overlap in time [1]. We focus initially on the highly non-Gaussian background from stellar-mass binary black holes, but return below to consider quasi-Gaussian backgrounds from binary neutron stars.

The binary black hole background consists of (in principle), clearly distinguishable, deterministic signals [9]. Since distant events are not resolvable with current detectors, most events are not distinguishable in practice, but they would be with a more sensitive detector. This is in contrast to, for example, the stochastic background from white dwarf binaries in the millihertz band, which cannot be distinguished in principle; see, e.g., Ref. [10]. Cross-correlation searches are suboptimal for non-Gaussian backgrounds. It is possible to improve the sensitivity of the stochastic search by including a more accurate description of the signal model.

A number of studies dating back to 2002 have outlined a variety of strategies for developing a non-Gaussian pipeline. Drasco and Flanagan derived an algorithm suitable for bursting sources [11] observed by colocated detectors with white noise [12]. While the assumptions of white noise and colocated detectors are not realistic, the study was important for showing that a non-Gaussian analysis could achieve a potentially significant improvement in sensitivity compared to a cross-correlation search. Subsequently, Thrane outlined a method for bursting sources that can be applied in the more realistic case of spatially separated detectors with colored noise [13]. Work by Martellini and Regimbau has explored additional strategies [14,15]. None of these methods have yet been used in a published search. For a comprehensive review of stochastic background methodology up until this point, see Ref. [16].

In this paper, we take a step back and ask a new question: What is the *optimal* method for detecting a stochastic background of binary black holes (or any other astrophysical background)? Formally, the optimal method yields a minimum credible interval posterior. In practical terms, the optimal method can be more sensitive than suboptimal techniques. Conceptually, deriving the optimal method amounts to implementing a likelihood function, which describes the salient features of the signal model and measurement noise. Turning the crank of Bayesian statistics, the resulting posterior distributions yield minimum credible intervals by construction.

It turns out that this question of optimality is interesting for several reasons. First and most obvious, it is desirable to derive the most sensitive possible search. Depending on the magnitude of improvement over the cross-correlation search, an optimal pipeline could significantly reduce the

time to the detection of a stochastic background. Second, we show below that the optimal Bayesian search yields, as a by-product, information about the population properties of high-redshift black holes. In particular, we automatically obtain a posterior distribution for the coalescence rate of binary black holes at high redshift. This information is lost during the process of cross-correlation.

Third, our formalism provides a natural framework for unifying the stochastic search with measurements of unambiguous detections and even “silver-plated” candidates [17] like LVT151012 [18]. As the number of detected black hole mergers increases, the gravitational-wave community is likely to be increasingly interested in population statements. The framework proposed here is the natural means of combining all possible data to make statements about populations of binary black holes (and other astrophysical phenomena). The method is free of Malmquist bias [19] since there is no selection of events.

Last, we outline how the method can be generalized in order to solve a number of important problems in gravitational-wave astronomy including (1) measurement of a primordial Gaussian background in the presence of an astrophysical foreground and (2) measurement of the population properties of binary black holes and neutron stars, e.g., their mass and spin distributions.

The remainder of this paper is organized as follows. In Sec. II, we derive the optimal search for an ensemble of binary black hole mergers. In Sec. III, we present the results of a Monte Carlo study that demonstrates the method. In Sec. IV, we investigate scaling relations in order to estimate the performance of the search in various contexts. In Sec. V, we consider complications arising from non-Gaussian noise. In Sec. VI, we consider the feasibility of the search given plausible computing resources.

In Sec. VII, we discuss how this method can be used to measure the population properties of unresolved binary black holes using hyperparameters. In Sec. VIII, we discuss how the method can be generalized to similar problems including binary neutron stars, neutron star black hole binaries, continuous waves from rotating neutron stars, and supermassive black hole binaries. In Sec. IX, we discuss the simultaneous measurement of a primordial gravitational-wave background. In Sec. X, we provide an overall assessment of the feasibility of an optimal stochastic search and the prospects for detection of a stochastic background.

II. METHOD

A. Likelihood function

We break the strain data s into convenient-sized segments. The data for segment i is denoted s_i . By assumption, each segment is big enough that it might include one signal (specifically, the merger part of the signal), but small enough that it is unlikely to contain two such signals. In this paper, we satisfy these criteria by using segments with

a duration of 4 s. For binary black holes like GW150914, the signal is in band for only ≈ 0.2 s, less than the segment duration. Since the probability of observing just one event in any given segment is small $\approx 2\%$, the probability of observing two at once is negligibly small, $\approx 10^{-4}$.

Assuming Gaussian noise, the log-likelihood for a single segment s_i containing a compact binary signal with parameters θ is

$$\log[\mathcal{L}(s_i|\theta_i)] \propto -\frac{1}{2} \langle s_i - h(\theta_i), s_i - h(\theta_i) \rangle. \quad (2)$$

Here, $h(\theta_i)$ is the signal model, which depends on a vector of parameters θ_i , e.g., sky location, component masses, the time of coalescence within the segment, and so on. We have introduced the usual noise-weighted inner product:

$$\langle a, b \rangle \equiv 4\text{Re}\Delta f \sum_k \frac{a^*(f_k)b(f_k)}{S_n(f_k)}, \quad (3)$$

where Δf is the frequency-bin size and Re is the real part of the sum. The variable $S_n(f_k)$ is the detector noise power spectral density. The sum runs over k frequency bins.

When combining data from a network of M detectors, the likelihood function for segment i becomes

$$\mathcal{L}(\vec{s}_i|\theta) = \prod_{j=1}^M \mathcal{L}(s_i^{(j)}|\theta_i), \quad (4)$$

where \vec{s}_i represents the data from all M detectors and j indexes each of the M detectors. Stochastic searches typically rely on $M \geq 2$ networks in order to distinguish astrophysical signals from poorly modeled noise [5]. Henceforth, we assume $M \geq 2$ detectors because this enables a *coherent* search for subthreshold signals. The notation \vec{s}_i is used to indicate the strain data from M detectors associated with segment i .

We can generalize the likelihood function in Eq. (2) to account for the fact that we expect the data to contain a signal plus Gaussian noise with probability ξ , or pure Gaussian noise with probability $(1 - \xi)$. (We defer discussion of non-Gaussian noise until Sec. V.) Hence, we define a “generalized likelihood” for segment i , modified to include the signal probability hyperparameter ξ [12], which we refer to as the “duty cycle”:

$$\mathfrak{L}(\vec{s}_i|\theta_i, \xi) = \xi \mathcal{L}(\vec{s}_i|\theta_i) + (1 - \xi) \mathcal{L}(\vec{s}_i|0). \quad (5)$$

We note that this astrophysical duty cycle should not be confused with the detector duty cycle, which is the fraction of time during which a detector collects science-quality data. Here, $\mathcal{L}(\vec{s}_i|0)$ is the likelihood function given the hypothesis that no signal is present:

$$\log[\mathcal{L}(\vec{s}_i|0)] \propto -\frac{1}{2} \langle s_i, s_i \rangle. \quad (6)$$

We marginalize over the astrophysical parameters of the event θ_i [with a prior distribution $\pi(\theta_i)$] in order to obtain a likelihood for the data given the duty cycle ξ for segment i :

$$\mathfrak{Z}(\vec{s}_i|\xi) = \xi \mathcal{Z}_S^i + (1 - \xi) \mathcal{Z}_N^i. \quad (7)$$

We have introduced two new terms,

$$\mathcal{Z}_S^i \equiv \int d\theta \mathcal{L}(\vec{s}_i|\theta_i) \pi(\theta_i), \quad (8)$$

$$\mathcal{Z}_N^i \equiv \mathcal{L}(\vec{s}_i|0), \quad (9)$$

corresponding, respectively, to the signal evidence and the noise evidence. These evidences are the raw output of LIGO parameter estimation algorithms such as LALINFERENCE [20]. We do not marginalize over the detector noise power spectral density. The power spectral densities that enter Eq. (3) are empirically estimated for each data segment.

The next step is to combine data from a large set of n segments $\{\vec{s}\}$. The combined likelihood for the data given ξ is

$$\begin{aligned} \mathfrak{Z}^{\text{tot}}(\{\vec{s}\}|\xi) &= \prod_i^n \mathfrak{Z}(\vec{s}_i|\xi) \\ &= \prod_i^n [\xi \mathcal{Z}_S^i + (1 - \xi) \mathcal{Z}_N^i]. \end{aligned} \quad (10)$$

The posterior for the duty cycle $p(\xi|\{\vec{s}\})$ is simply

$$p(\xi|\{\vec{s}\}) \propto \mathfrak{Z}^{\text{tot}}(\{\vec{s}\}|\xi) \pi(\xi), \quad (11)$$

where $\pi(\xi)$ is the prior distribution on ξ . For simplicity we assume a flat, uniform prior for $\pi(\xi)$, so that

$$p(\xi|\{\vec{s}\}) \propto \mathfrak{Z}^{\text{tot}}(\{\vec{s}\}|\xi), \quad (12)$$

though our analysis can incorporate any suitable choice of $\pi(\xi)$, informed, for instance, by expectations about the average time between binary black hole mergers [1].

B. Detection statistic

In order to search for an astrophysical stochastic background, we calculate a “stochastic background evidence”:

$$\mathcal{Z}_{\text{stoch}} = \int d\xi \mathfrak{L}(\{\vec{s}\}|\xi) \pi(\xi). \quad (13)$$

The null hypothesis (there is no stochastic background) is described by a null evidence:

$$\mathcal{Z}_0 = \mathcal{L}(\{\vec{s}\}|\xi = 0). \quad (14)$$

We construct a Bayes factor to compare the two hypotheses:

$$\text{BF} = \mathcal{Z}_{\text{stoch}}/\mathcal{Z}_0. \quad (15)$$

This variable is an optimal detection statistic for an astrophysical background of compact binaries. The optimality follows from the fact that the likelihood function describing the data given the signal and noise models is complete. Hence, the search produces a minimum credible-interval posterior distribution on the duty cycle. In the remainder of this paper, we adopt the convention that a log Bayes factor of ≈ 8 represents a statistically significant preference for one hypothesis over the other [21].

Up to this point, we have, for the sake of convenience, written our likelihoods in terms of duty cycle ξ . In the subsequent subsections (Secs. II C–II E), we discuss how ξ is related to other quantities including rate and energy density.

C. Rate

In this section, we take the likelihood $\mathfrak{L}^{\text{tot}}(\{\vec{s}\}|\xi)$ —a function of duty cycle ξ —and recast it as a function of R : the number of mergers per segment. In particular, R is the rate of events throughout the visible Universe per segment. This allows us to more easily relate our analysis to observations of individual merger events. It is useful to contrast ξ and R . While duty cycle ξ is defined on $(0,1)$, the rate of events per unit segment is defined on $(0, \infty)$. Even perfect knowledge of ξ does not determine R , because the latter is subject to cosmic variance arising from the fact that events take place randomly following Poisson statistics.

The number of compact binary mergers N is given by the product of the duty cycle ξ and the number of segments n :

$$N = n\xi. \quad (16)$$

We perform a change of variable in order to recast the likelihood variable in terms of N :

$$\mathfrak{L}^{\text{tot}}(\{\vec{s}\}|N) = \int d\xi \mathfrak{L}^{\text{tot}}(\{\vec{s}\}|\xi) \pi(\xi|N), \quad (17)$$

where

$$\pi(\xi|N) = \left| \frac{dN}{d\xi} \right| \pi(N) = n\pi(N). \quad (18)$$

Having recast the likelihood in terms of the number of compact binary mergers N , we can marginalize over N to obtain

$$\mathcal{L}^{\text{tot}}(\{\vec{s}\}|R) = \sum_N \mathfrak{L}^{\text{tot}}(\{\vec{s}\}|N) \pi(N|R), \quad (19)$$

where $\pi(N|R)$ is a conditional prior for the number of compact binary mergers N given a rate R (with units of mergers per segment). The conditional prior is given by a Poisson distribution:

$$\pi(N|R) = e^{-R} \frac{R^N}{N!}. \quad (20)$$

The total evidence is a likelihood function for the data given the rate hyperparameter:

$$\mathfrak{L}^{\text{tot}}(\{\vec{s}\}|R) = \mathcal{Z}^{\text{tot}}(R) \quad (21)$$

$$= \sum_N \mathfrak{L}^{\text{tot}}(\{\vec{s}\}|N) e^{-R} \frac{R^N}{N!}. \quad (22)$$

The rate posterior is

$$p(R|\{\vec{s}\}) \propto \mathfrak{L}^{\text{tot}}(\{\vec{s}\}|R) \pi(R), \quad (23)$$

where $\pi(R)$ is some suitable prior on the rate R .

D. Local rate

The variable that we have been referring to as “the rate” R —by which we mean the number of compact binary mergers in the visible Universe per segment—is distinct from the local merger rate R_0 with units of $\text{Gpc}^{-3} \text{yr}^{-1}$. However, they are related. We derive the local rate R_0 from R :

$$\frac{R}{\delta t} = \int \frac{dz}{1+z} \left(\frac{dV}{dz} \right) R(z). \quad (24)$$

Here, $R(z)$ is the comoving merger rate in units of $\text{Gpc}^{-3} \text{yr}^{-1}$ as a function of redshift and δt is the segment duration. The factor dV/dz describes how an element of volume evolves in an expanding universe while the factor of $1+z$ comes about transforming the time variable in the source frame to the detector frame; see Ref. [1].

The shape of $R(z)$ is determined by a model, which takes into account, e.g., the stellar formation rate as a function of redshift and the time delay between formation and coalescence; see, e.g., Ref. [3]. However, we can treat the overall normalization as a free parameter so that

$$R(z) \equiv R_0 \mathfrak{G}(z), \quad (25)$$

where $\mathfrak{G}(z)$ is a model-dependent, dimensionless shape function [normalized so that $\mathfrak{G}(z=0) = 1$] and the local rate $R_0 \equiv R(z=0)$ is the normalization. Combining Eqs. (24) and (25), we obtain

$$R_0 = \frac{R}{\delta t} / \int \frac{dz}{1+z} \left(\frac{dV}{dz} \right) \mathfrak{G}(z) \quad (26)$$

It will be interesting to compare the local rate inferred from a stochastic detection—and assuming some shape model $\mathfrak{S}(z)$ —with the local rate measured directly by resolvable mergers. Tension in these two measurements could indicate an inadequate shape model $\mathfrak{S}(z)$ among other things.

E. Energy density

Given some model, the local rate R_0 (and therefore R) can be converted into dimensionless energy density $\Omega_{\text{gw}}(f)$ defined in Eq. (1); see, e.g., Refs. [22,23]. The fully general expression is a bit unwieldy, so we employ two simplifying assumptions in order to obtain an intuitive initial expression. First, we assume that the (source-frame) energy spectrum of the event dE_{gw}/df_s —a function of the source frame frequency f_s —is determined primarily by the chirp mass of the binary M_c . Second, we assume that the comoving merger rate does not depend on mass. Given these two assumptions,

$$\Omega_{\text{gw}}(f) = \left(\frac{f}{\rho_c}\right) \left(\int_0^\infty \frac{dz}{1+z} \frac{R_0 \mathfrak{S}(z)}{H(z)}\right) \times \left(\int dM_c \frac{dE_{\text{gw}}}{df_s}(f_s|M_c) \pi(M_c)\right), \quad (27)$$

where

$$H(z) = H_0 \sqrt{\Omega_M(1+z)^3 + \Omega_\Lambda}. \quad (28)$$

H_0 is the Hubble constant, Ω_M is the energy density of matter, and Ω_Λ is the energy density of the cosmological constant. In the first set of parentheses, ρ_c is the critical energy density for a flat universe. The variable $\pi(M_c)$ is the mass distribution.

Thus, we can relate the number of events per segment R (and/or the local rate R_0) to the energy density spectrum $\Omega_{\text{gw}}(f)$, but only by employing a model to describe the distribution of events in redshift, mass, and so on. In the derivation above, we assume that $\Omega_{\text{gw}}(f)$ depends primarily on chirp mass. However, a more general expression for $\Omega_{\text{gw}}(f)$ would include integrals over every variable that can affect the energy spectrum, for example, the mass ratio. Also, in general, $\mathfrak{S}(z)$ can depend on variables such as M_c , in which case $\mathfrak{S}(z|M_c, \dots)$ cannot be taken out of the Eq. (27) integral over M_c .

III. DEMONSTRATION USING GAUSSIAN NOISE

In this section, we carry out a Monte Carlo simulation in order to demonstrate the method described in Sec. II. Our goal is to calculate the duty cycle posterior $p(\xi|\{\vec{s}\})$ using simulated data containing a population of subthreshold black hole binaries. We seek to fulfill three criteria. First, the method should be “safe”: it should return a null result when applied to pure Gaussian noise. Second, the method should be effective: it must yield a positive detection in the

presence of a sufficiently loud stochastic signal. Third, the method should be unbiased. The duty cycle posterior ought to, on average, peak at the injected value.

We assume a two-detector network consisting of the LIGO Hanford and Livingston observatories operating at design sensitivity [24]. It is straightforward to extend the method to include additional detectors [25], but we begin with two for the sake of simplicity. For each detector, we generate two data sets. The *noise* data set consists of 1000 4-s segments of Gaussian noise. The *signal* data set consists of 300 4-s segments with a binary black hole signal added to Gaussian noise.

The signals are coherently generated (and later recovered) using the IMRPhenomPv2 approximant [26]. The parameters of each merger are drawn randomly. The orientation angles and sky position are drawn from isotropic distributions. We marginalize over the time of coalescence, which is drawn from a uniform distribution. The total mass is drawn from a uniform distribution on $(48M_\odot, 80M_\odot)$ while the mass ratio is drawn from a uniform distribution on $(1, 8)$. This mass range produces signals that fit conveniently into our 4-s segments, given a minimum frequency of 20 Hz. The mass ratio is within the domain of validity for IMRPhenomPv2. The dimensionless spin magnitudes (a_1, a_2) are drawn from a uniform distribution on $(0, 0.89)$, which is within the domain of validity for IMRPhenomPv2. The spin unit vectors are drawn from an isotropic distribution. The luminosity distances are drawn from a uniform-in-volume distribution on the interval $(0.50 \text{ Gpc}, 5 \text{ Gpc})$. The lower limit of 0.50 Gpc removes most gold-plated detections, the remainder of which are eliminated with an additional cut described below.

Reconstructed masses and distance are affected by redshift. The *lab-frame* masses $m_{i=1,2}^l$ and the luminosity distance d_L are given by

$$m_i^l = (1+z)m_i^s, \quad (29)$$

$$d_L^l = (1+z)d_M, \quad (30)$$

where m_i^s is measured in the *source frame* and d_M is the comoving distance. The mass distributions described above apply to quantities measured in the lab frame.

Luminosity distances of $(0.50 \text{ Gpc}, 5 \text{ Gpc})$ imply a redshift interval of $(0.10, 0.77)$ within the framework of the standard Λ CDM cosmology. In principle, we can (and eventually will) extend the interval to $d \gtrsim 30 \text{ Gpc}$ ($z \gtrsim 3.3$), beyond which the stochastic signal is expected to be marginal due to the low stellar formation rate and expanding Universe [4]. The uniform-in-volume distribution is then modified by the imposition of a “Malmquist prior” [19]; we exclude gold-plated detections that can be observed with a statistically significant coherent matched-filter signal-to-noise ratio $\rho_{\text{network}} \geq 12$:




FIG. 1. Histogram of the coherent network matched-filter signal-to-noise ratio [Eq. (31)] for a data set containing simulated signals plus Gaussian noise. The network signal-to-noise ratio is calculated by maximizing over all astrophysical parameters. This is why the distribution falls off sharply near $\rho_{\text{network}} = 1$; one of the many available templates tends to produce a reasonable fit to the data such that $\rho_{\text{network}} \gtrsim 1$, even for weak signals. The signal distribution is generated from a population of mergers uniform in volume between (0.50 Gpc, 5 Gpc). Gold-plated detections with $\rho_{\text{network}} \geq 12$ are excluded.

$$\rho_{\text{network}} = \max_{\theta} \sqrt{\sum_j \frac{\langle s_j, h(\theta) \rangle^2}{\langle h(\theta), h(\theta) \rangle}}. \quad (31)$$

Here, the sum over j runs over detectors. The maximum over θ determines the matched-filter template that best fits the data, and so ρ_{network} is the maximum-likelihood signal-to-noise ratio. Since most binaries merge near the edge of our 5 Gpc sphere, this cut removes $\lesssim 1\%$ of the events.

Eventually, gold-plated events should be included in the analysis in order to achieve a unified approach to compact binary population inference and stochastic backgrounds. However, we exclude them here for the sole purpose of demonstrating that we can recover the signal from a stochastic background of subthreshold events. A histogram of ρ_{network} is shown in Fig. 1.

Having generated mock data for the noise and signal populations, we create mixed populations for arbitrary duty cycles ξ by selecting at random a mixture of entries from each distribution. We construct three data sets corresponding to $\xi = (0, 0.05, 1)$ with $n = (500, 525, 300)$ segments, respectively. We compute the signal evidence \mathcal{Z}_S and noise evidence \mathcal{Z}_N for every event in each data set [Eqs. (8) and (9)]. The calculation is carried out using LALINFERENCE. We employ reduced order modeling and reduced order quadrature methods to control the computational cost of the analysis [27].

In Fig. 2, we plot $p(\xi|\{\vec{s}\})$ for the three values of duty cycle. The posterior for pure noise ($\xi_{\text{true}} = 0$) is indicated by the hatched orange area. The fact that it peaks at $\xi = 0$ shows that the method is safe. The posterior for pure signal ($\xi_{\text{true}} = 1$) is indicated by the hatched blue area. The fact that it peaks at $\xi = 1$, clearly excluding $\xi = 0$, indicates

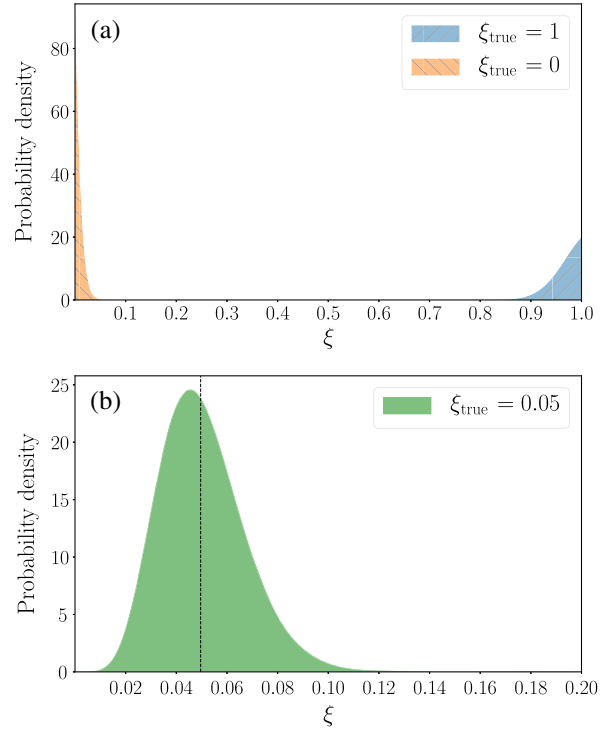


FIG. 2. Duty cycle posteriors $p(\xi|h)$ [Eq. (11)] for Monte Carlo data sets. The orange data are pure Gaussian noise, the blue data are a population of subthreshold binary black hole events added to Gaussian noise, and the green data correspond to a mixture with a true duty cycle of $\xi_{\text{true}} = 0.05$. The fact that each posterior peaks at the appropriate value shows that the method is safe, effective, and unbiased. (a) Safe and effective duty cycle posteriors. (b) Unbiased duty cycle posterior for Gaussian noise containing software injections with a duty cycle $\xi_{\text{true}} = 0.05$. The vertical line shows the true duty cycle.

that the method is effective. Finally, the posterior for a mixed distribution ($\xi_{\text{true}} = 0.05$) is indicated by green. The posterior peaks near the true value of $\xi_{\text{true}} = 0.05$, which shows that the method is unbiased.

IV. TIME TO DETECTION

In this section, we estimate the time it will take to make a confident detection of an astrophysical background ($\ln \text{BF} > 8$). In doing so, we study how the Bayes factor [Eq. (15)] scales as a function of the true duty cycle ξ_{true} and the number of data segments N_{segs} . Recent estimates place the binary black hole background at $\Omega_{\text{gw}}(f = 25 \text{ Hz}) \approx 1.1 \times 10^{-9}$ [1]. In the Appendix, we show that, given our mass and distance distributions (described above in Sec. III), $\xi_{\text{true}} = 4 \times 10^{-4}$ provides a realistic duty cycle. We assume this value of ξ_{true} for the remainder of this section. Because of computational constraints, we create 1000 noise segments for this preliminary study. This allows us to carry out a mock study using ≈ 1 h of data and to probe duty cycles $\xi \gtrsim 0.1\%$. In order to estimate how the




FIG. 3. Probability density functions for signal and noise evidences. (a) Fit for the signal data set and (b) fit for the noise data set. The horizontal axis is the log signal evidence while the vertical axis is the log noise evidence.

algorithm will perform when applied to longer data sets and/or lower duty cycles, it is necessary to extrapolate. Our extrapolation provides an initial performance estimate, which should be checked with a more computationally expensive mock data challenge.

Our extrapolation uses a Gaussian mixture model (GMM) [28] to fit the distributions of signal and noise evidences $[\pi(\mathcal{Z}_S), \pi(\mathcal{Z}_N)]$ using the (1000, 300) data segments that we have already simulated. A GMM models the data as a superposition of independent Gaussian distributions. The GMM fits are displayed in Fig. 3(a) (signal data set) and Fig. 3(b) (noise data set). In each panel, the horizontal axis is the log signal evidence while the vertical axis is the log noise evidence. The color bar indicates the probability density.

Using the GMM fits, we can generate large extrapolated data sets with arbitrarily low duty cycles. In Fig. 4, we show how the log Bayes factor [Eq. (15)] scales with the injected duty cycle ξ_{true} and the number of segments N_{segs} , the latter of which is equivalent to the observation time. We average over 1000 realizations of $(\mathcal{Z}_S, \mathcal{Z}_N)$ drawn from the GMMs




FIG. 4. Contours of average $\ln \text{BF}$ as a function of the simulated true duty cycle ξ_{true} and number of 4-s segments N_{segs} . The blue contour corresponds to $\langle \ln \text{BF} \rangle = 8$, where an astrophysical background is detectable. The color bar saturates at $\ln \text{BF} = 20$. The horizontal and vertical lines correspond to $\xi_{\text{true}} = 4 \times 10^{-4}$ and $N_{\text{segs}} = 17\,500$, respectively.

(created assuming a two-detector LIGO network operating at design sensitivity). We find that an astrophysical background with a realistic effective duty cycle $\xi_{\text{true}} \approx 4 \times 10^{-4}$ can be detected with $N_{\text{segs}} \approx 17\,000$ data segments, corresponding to 20 h. This can be compared to a detection time of $\gtrsim 1$ yr using cross-correlation [1]. The signal is created by ≈ 7 subthreshold events with $z < 0.77$.

We expect that the improvement in sensitivity results from two effects. First, the likelihood includes information about the non-Gaussian nature of the binary black hole background. Second, the likelihood includes information about the deterministic nature of compact binary waveforms. Additional work is required to determine the relative importance of these two factors.

V. NON-GAUSSIAN NOISE

Up until this point, we have chosen to model our data as either Gaussian noise or Gaussian noise plus signal. However, real gravitational-wave detectors are subject to non-Gaussian transient noise called glitches. It is therefore necessary to extend our algorithm to account for non-Gaussian noise in order to ensure that glitches are not mistaken for gravitational-wave signals. Failure to account for glitches can bias the duty cycle posterior, or worse, yield a false positive. We compute the duty cycle posterior for a set of 600 4-s segments of data from LIGO's first observing run (O1). We introduce an unphysical time shift to ensure that any real gravitational-wave signals do not produce coherent signals in the Hanford and Livingston detectors. The true duty cycle should therefore be zero, but the posterior peaks at around $\xi = 0.4$ and excludes zero. Nonstationary noise is producing a significant bias. An example of the undesirable effect of glitches is shown in Fig. 5.

In order to take into account non-Gaussian noise, we extend the likelihood expression from Eq. (7) to include contributions from glitches. For simplicity, we consider a two-detector network, but the results generalize. For our glitch model, we conservatively suppose that glitches look exactly like binary black hole waveforms except that the




FIG. 5. Biased duty cycle posterior, generated with O1 background data and computed using the Gaussian likelihood function, Eq. (5).

waveform in one detector is completely uncorrelated with the waveform in the other. Any part of a glitch that is orthogonal to the binary black hole signal manifold will not contribute any signal evidence \mathcal{Z}_S .

Some additional notation is necessary. We introduce parameters $\xi_g^{(1)}$ and $\xi_g^{(2)}$ corresponding to the glitch duty cycle in detectors one and two, respectively. The variables $\mathcal{Z}_g^{(1)}$ and $\mathcal{Z}_g^{(2)}$ are the single-detector evidences for the glitch hypothesis given by

$$\mathcal{Z}_g^{(1)} \equiv \int d\theta^{(1)} \mathcal{L}(s^{(1)}|\theta^{(1)})\pi(\theta^{(1)}), \quad (32)$$

$$\mathcal{Z}_g^{(2)} \equiv \int d\theta^{(2)} \mathcal{L}(s^{(2)}|\theta^{(2)})\pi(\theta^{(2)}). \quad (33)$$

Here, $\theta^{(1)}$ and $\theta^{(2)}$ are the signal parameters for (uncorrelated) glitches in detectors 1 and 2, respectively. The variables $s^{(1)}$ and $s^{(2)}$ are the strain data in each detector. We note that the glitch evidences are equivalent to single-detector *signal* evidences given our conservative glitch model. A key feature of this glitch model is that the parameters in each detector are independent. We introduce the variables $\mathcal{Z}_N^{(1)}$ and $\mathcal{Z}_N^{(2)}$ for the single-detector noise evidences [see Eq. (9)]. The variables $\mathcal{Z}_{S+g}^{(1)}$ and $\mathcal{Z}_{S+g}^{(2)}$ are the two-detector evidences for a coherent signal with a glitch in one detector. Finally, the variable $\mathcal{Z}_{S+g}^{(1,2)}$ is the two-detector evidence for a coherent signal with a glitch in both detectors. We do not include explicit expressions for $\mathcal{Z}_{S+g}^{(1)}$, $\mathcal{Z}_{S+g}^{(2)}$, and $\mathcal{Z}_{S+g}^{(1,2)}$ because—as we shall see momentarily—they can be approximated as zero. For the sake of compact notation, we suppress the segment number index i on every evidence term.

Given our new definitions, the likelihood for coherent merger events in glitchy data is

$$\begin{aligned} \mathfrak{L}(\vec{s}_i|\xi, \xi_g^{(1)}, \xi_g^{(2)}) &= \xi(1 - \xi_g^{(1)})(1 - \xi_g^{(2)})\mathcal{Z}_S \\ &+ (1 - \xi)(1 - \xi_g^{(1)})(1 - \xi_g^{(2)})\mathcal{Z}_N \\ &+ (1 - \xi)\xi_g^{(1)}(1 - \xi_g^{(2)})\mathcal{Z}_g^{(1)}\mathcal{Z}_N^{(2)} \\ &+ (1 - \xi)(1 - \xi_g^{(1)})\xi_g^{(2)}\mathcal{Z}_N^{(1)}\mathcal{Z}_g^{(2)} \\ &+ (1 - \xi)\xi_g^{(1)}\xi_g^{(2)}\mathcal{Z}_g^{(1)}\mathcal{Z}_g^{(2)} \\ &+ \xi\xi_g^{(1)}(1 - \xi_g^{(2)})\mathcal{Z}_{S+g}^{(1)} \\ &+ \xi(1 - \xi_g^{(1)})\xi_g^{(2)}\mathcal{Z}_{S+g}^{(2)} \\ &+ \xi\xi_g^{(1)}\xi_g^{(2)}\mathcal{Z}_{S+g}^{(1,2)}. \end{aligned} \quad (34)$$

Each line corresponds to a distinct possibility. A probability tree corresponding to the likelihood function is shown in Fig. 6. Starting from the first line and reading toward the bottom, the possibilities are (i) signal and no glitches; (ii) no signal and no glitches; (iii) no signal and a glitch in detector (1); (iv) no signal and a glitch in (2); (v) no signal and a glitch in (1) and (2); (vi) signal and a glitch in (1); (vii) signal and a glitch in (2); (viii) signal with glitches in (1) and (2).

Note the priors for $\xi, \xi_g^{(1)}, \xi_g^{(2)}$ all run from (0,1). They are all independent. Each of the above possibilities corresponds to a branching path in Fig. 6.

For the sake of simplicity, we hypothesize that the last three terms in this likelihood contain evidences that are small enough to safely ignore in practice: $\mathcal{Z}_{S+g}^{(1)}, \mathcal{Z}_{S+g}^{(2)}$, and $\mathcal{Z}_{S+g}^{(1,2)}$. We expect these three \mathcal{Z} to be small because they employ overzealous models, which tend to overfit the data. While terms like $\mathcal{Z}_S, \mathcal{Z}_g^{(1)}$, and $\mathcal{Z}_g^{(2)}$ employ 15 parameters to fit a merger event or mergerlike glitch, $\mathcal{Z}_{S+g}^{(1)}$ and $\mathcal{Z}_{S+g}^{(2)}$ employ 30 parameters to simultaneously fit a merger event *and* a glitch. The $\mathcal{Z}_{S+g}^{(1,2)}$ term employs 45 parameters to fit a merger event and two glitches. Since we expect that mergers and glitches are easily fit with 15-parameter models, the final three \mathcal{Z} incur large Occam factors, which results in small \mathcal{Z} . This reasoning may not apply to the special case of data that actually contain a merger and a glitch, but we expect such events to be rare. The recent binary neutron star event GW170817 *was* detected coincident with a significant glitch [29], but this does not change our expectations. Binary neutron star signals are in band for $\gtrsim 100$ s versus $\lesssim 0.3$ s for high-mass binary black holes, and so the chance of a coincident glitch is relatively higher. Of course, one is free to retain the $S+g$ terms, but this requires modification of LALINFERENCE. We anticipate that such modifications will be well worth pursuing for a number of applications, including work to extend this analysis to low-mass systems. The subtraction of a glitch associated with GW170817 [29] using a wavelet reconstruction is a step in this direction [30].




FIG. 6. Probability tree for Eq. (34) assuming a two-detector network. The left-hand branches correspond to “yes” and the right-hand branches to “no.” The probability of each branch is labeled accordingly. The “leaves” at the base of the tree show the evidence associated with each path.

If we assume that the final three \mathcal{Z} are small enough to ignore, the glitchy likelihood becomes

$$\begin{aligned} \mathfrak{L}(\vec{s}|\xi, \xi_g^{(1)}, \xi_g^{(2)}) &\approx \xi(1 - \xi_g^{(1)})(1 - \xi_g^{(2)})\mathcal{Z}_S \\ &+ (1 - \xi)(1 - \xi_g^{(1)})(1 - \xi_g^{(2)})\mathcal{Z}_N \\ &+ (1 - \xi)\xi_g^{(1)}(1 - \xi_g^{(2)})\mathcal{Z}_g^{(1)}\mathcal{Z}_N^{(2)} \\ &+ (1 - \xi)(1 - \xi_g^{(1)})\xi_g^{(2)}\mathcal{Z}_N^{(1)}\mathcal{Z}_g^{(2)} \\ &+ (1 - \xi)\xi_g^{(1)}\xi_g^{(2)}\mathcal{Z}_g^{(1)}\mathcal{Z}_g^{(2)}. \end{aligned} \quad (35)$$

Below, we show that the “small- \mathcal{Z} ” approximation works well when we apply this likelihood to real data. This is convenient because it enables us to obtain results using only existing 15-dimensional signal models.

We now show that our method is robust when applied to real LIGO noise. Because the LIGO detectors cannot be shielded from gravitational-wave signals, we ensure that the analyzed data cannot contain coincident real signals by performing “time slides” in which a relative time offset, longer than the light-travel time between sites, is applied between the data from the detectors. Time slides are a common bootstrap technique for generating realistic background noise.

As before, we generate two data sets: a noise-only data set consisting of 670 4-s background data segments from O1, and an injection data set consisting of 60 software injections into 4-s background data segments from O1.

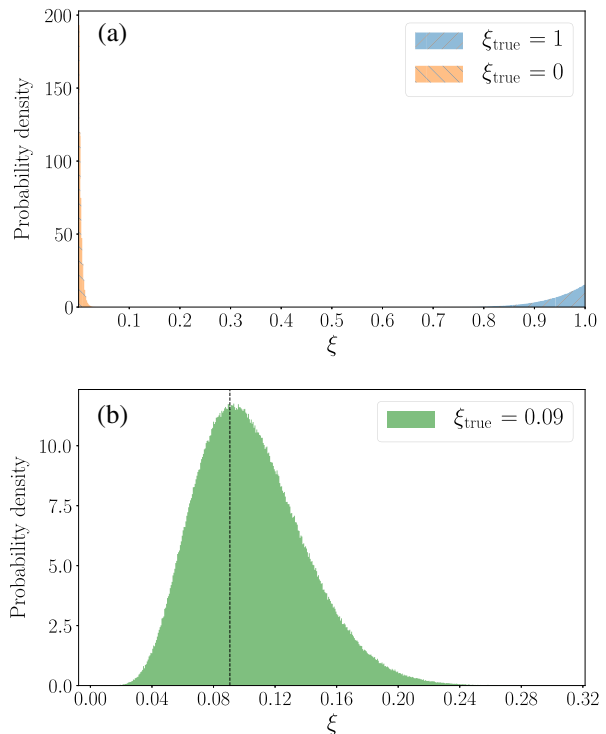


FIG. 7. Duty cycle posteriors computing using the “glitchy” likelihood function Eq. (35). (a) Safe and effective duty cycle posteriors for O1 background. (b) Unbiased duty cycle posterior for O1 background containing software injections with a duty cycle $\xi_{\text{true}} = 0.09$. The vertical line shows the true duty cycle.

The signals are generated according to the prescription described in Sec. III. Injections are drawn from a uniform-in-volume distribution. They are all subthreshold. We construct three populations corresponding to $\xi = (0, 0.09, 1)$ to test for effectiveness, safety, and bias. We compute the duty cycle posterior using the glitchy likelihood function [Eq. (35)]. In Fig. 7, we demonstrate that the glitch model employed in the likelihood function Eq. (35) is safe, effective, and unbiased, as required.

VI. COMPUTATIONAL REQUIREMENTS

Figure 4 implies that, at design sensitivity, a realistic astrophysical background with an effective $z < 0.77$ duty cycle of ($\xi_{\text{true}} = 4 \times 10^{-4}$) could be detected using $\approx 17\,000$ 4-s data segments. At the time of writing, a typical run time of LALINFERENCE on 4-s data segments is usually no more than 10 CPU hours [27]. To produce evidences in a two-detector network, we perform three separate runs: a coherent analysis, which produces $(\mathcal{Z}_S^i, \mathcal{Z}_N^i)$, and two incoherent analyses, which produce $(\mathcal{Z}_S^{i,(j)}, \mathcal{Z}_N^{i,(j)})$, where the index j labels the detector, $j = (1, 2)$. We therefore estimate that an astrophysical background is detectable in 500 000 CPU hours. This works out to one week with 3000 dedicated CPUs. Detecting an astrophysical background of binary black holes is therefore feasible with current computing resources of the LIGO Data Grid (LDG), which consists of tens of thousands of CPUs. This cost estimate assumes relatively high-mass signals split into 4-s data segments. Additional work is required to investigate the cost and performance of the search as the analysis is extended to lower-mass events. As the minimum mass is decreased, the event rate is expected to increase while the signal duration increases, resulting in a higher duty cycle. We discuss some of the associated challenges in Sec. VIII A.

While it may be possible to observe a stochastic background by analyzing a small amount of data (with a small computational cost), there is strong motivation for analyzing all available data. By analyzing an entire observing run, we can make inferences about the population properties of binary black holes (and with additional work, eventually binary neutron stars). We describe this population inference in the subsequent section. Carrying out full parameter estimation on a year of data with existing software and hardware would require around 15 000 continuously operational CPUs in order to process the data in real time. While this is technically possible, it would be a costly proposition.

In order to reduce the cost, it is worthwhile to consider the development of new computational methods and implementation on new hardware architecture in order to realize the ultimate goals of the search. These advances might come, for instance, by exploring greater parallelization of parameter estimation methods through the use

of, e.g., graphical processor unit clusters or larger super-computer clusters; see, e.g., Ref. [31]. We note in passing that it might be possible to detect some individually resolvable sources using parameter estimation if they were just below the threshold for detection with matched filter pipelines [32].

VII. POPULATION INFERENCE

In addition to inferring the astrophysical duty cycle, which can be related to the local merger rate (Sec. II C) and GW energy density (Sec. II E), we can also use this framework to infer properties of the binaries that contribute to the astrophysical background. These population properties may be encoded as hyperparameters, which affect the prior distributions for various binary parameters. For example, instead of assuming a flat prior on chirp mass, we can assume that the total mass follows a power-law distribution with some spectral index α :

$$\pi(M) \rightarrow \pi(M|\alpha) \propto M^\alpha. \quad (36)$$

Using Bayesian hierarchical modeling [33,34], we can marginalize over M to obtain a posterior on α .

Before sketching how this works, we note that there are three good reasons for eventually developing a sophisticated hyperparametrization scheme. First, there is interesting information encoded in the population properties of binary black hole coalescences, which can be used, for example, to study binary black hole formation channels; see, e.g., Refs. [35–41]. Second, we should hyperparameterize theoretically uncertain prior distributions in order to obtain an unbiased estimate of the coalescence rate. For example, a systematic error in the assumed mass spectrum will lead to a systematic error in the inferred rate posterior. Finally, by acknowledging theoretical uncertainty with hyperparametrization, we should improve the sensitivity of the search. This should generally be true when *necessary* parameters are added to the search. For example, marginalizing over the unknown mass spectrum index is likely to yield a higher Bayes factor than just assuming some value (unless we have a strong prior belief in some particular value).

The likelihood function for a vector of hyper-parameters Λ is obtained by introducing a conditional prior $\pi(\theta|\Lambda)$ and then marginalizing over θ

$$\mathcal{L}(h|\Lambda) = \int d\theta \mathcal{L}(h|\theta) \pi(\theta|\Lambda) \quad (37)$$

The distribution $\pi(\theta|\Lambda)$ is the hyperparameterized conditional prior. The above integral can be computed without performing any extra sampling by “recycling” the posterior samples already computed by LALINFERENCE. Replacing the likelihood by a sum over samples we obtain:

$$\mathcal{L}(h|\Lambda) \approx \prod_i^n \frac{1}{n_i} \sum_{k=1}^{n_i} \frac{\pi(\theta_{i,k}|\Lambda)}{\pi(\theta_{i,k}|\text{LAL})}. \quad (38)$$

The distribution $\pi(\theta|\text{LAL})$ is the prior distribution used for the generation of the posterior samples by LALINFERENCE. The product over i runs over the number of data segments from 1 to n . The sum over k runs over the number of posterior samples from 1 to n_i . (Each segment has a different number of posterior samples.)

We note that Eq. (38) can be applied as a postprocessing step. That is, one need only run LALINFERENCE once. As long as we use relatively uninformative priors for $\pi(\theta|\Lambda)$, the resulting posterior samples can be reweighted in order to obtain the likelihood for different hyperparameter values. The posterior for Λ and the (hypermarginalized) Bayes factor are calculated the usual way. The first prior distributions to hyperparametrize are (1) those which shed light on the mechanisms of binary black hole formation and (2) those which are subject to significant theoretical uncertainty. This probably means hyperparameter descriptions of the binary black hole mass spectrum and the distribution of black hole spins. The former is the subject of work in preparation [42].

VIII. EXTENSIONS OF THE SEARCH

In this section, we consider five possible extensions to this search method: searches for binary neutron stars, continuous waves, supermassive black hole binaries, bursts, and glitch classification. We also consider the simultaneous estimation of Gaussian background with a non-Gaussian foreground, but the discussion requires some depth, and so we discuss that separately in the next section. Our goal here is twofold. First, we seek to illustrate promising directions for future research. Second, we endeavor to showcase that, for every astrophysical background, there is an appropriately complete likelihood function, which can serve as the basis for an optimal search.

A. Binary neutron stars

An obvious extension of this technique is to search for backgrounds of binary neutron stars. Recent observations of a binary neutron star inspiral [29] suggest that the gravitational-wave energy density Ω_{gw} from such systems is roughly comparable to the energy density from binary black holes [3]. Binary neutron stars are less massive than typical binary black holes, but they coalesce more frequently. Thus, while the two backgrounds are expected to produce comparable Ω_{gw} , the rate of binary neutron star mergers is much higher (by a factor of ≈ 17). Moreover, binary neutron star waveforms are much longer than binary black hole waveforms. The first binary black hole event GW150914 was in band for ≈ 0.3 s [43] versus 100 s for the first binary neutron star GW170817 [29]. It is important to understand that while the binary neutron star background is

continuous (always present), it is still non-Gaussian: the mergers are, for the most part, clearly separated in time [9].

These two effects all but ensure that a signal is present in the data at any given time. At any given moment, there are likely to be 15 binary neutron stars somewhere in the Universe, producing gravitational waves with $f > 10$ Hz (versus 0.06 for binary black holes) [1]. This violates an assumption in our formulation of the binary black hole search: that the vast majority of segments contain either no signal or one signal.

All is not lost. There are probably a number of solutions. One possibility is to treat the number of events in some data segment as a free parameter using the local rate to inform the prior. The formalism of reversible-jump Markov chain Monte Carlo, where the dimensionality of the parameter space is itself a free parameter, is potentially suited for this problem; see, e.g., Ref. [30]. Using this plan of attack, the algorithm could attempt, for example, to simultaneously fit dozens of binary neutron signals simultaneously, each with 15 parameters. It is not clear if such a high-dimensional search would be computationally feasible, but further investigation is warranted.

B. Continuous waves

There are thought to be $\approx 50\,000$ isolated neutron stars in the Milky Way emitting gravitational waves in the observing band of advanced detectors [44,45]. The computational power required to search the full parameter space is so vast that semicoherent techniques are required. The data are analyzed coherently in small chunks of manageable size, typically 1800 s [46]. The results from each chunk are combined incoherently.

Since the signals overlap in time, it would not make sense to define a duty cycle as fraction of segment that include a signal. However, it may be possible to define a duty cycle equal to the fraction of *frequency bins* that contain a signal. To get a (very rough) idea of the duty cycle, we can assume that continuous-wave sources are roughly evenly distributed throughout the advanced detector observing band from 10 to 1500 Hz [46]. Assuming 1800-s segments (with frequency resolution 0.56 mHz), there are 2.7×10^6 frequency bins. This implies a duty cycle of $\xi \approx 2\%$, which is much less than one. We therefore expect that this formalism can be applied relatively straightforwardly to search for a population of Galactic neutron stars. This proposal bears similarities to Ref. [47], which proposes an ensemble search for *known* pulsars.

C. Supermassive black hole binaries

The background from supermassive black hole binaries is analogous to the background from isolated neutron stars, except the measurements are carried out by pulsar timing arrays; for a review, see Refs. [48,49]. Most supermassive black hole binaries are expected to produce nearly

monochromatic signals that evolve slowly over the decade-long observation period. At any one time, there might be 10^4 – 10^5 such binaries emitting gravitational waves in the pulsar-timing band [50]. The probability that an inspiralling binary emits gravitational waves on the interval $(f, f + df)$ is [51]

$$\pi(f) \propto f^{-11/3}. \quad (39)$$

Thus, most binaries emit near the lower limit of the observing band. Pulsar timing arrays observe in a band 1–100 nHz with a resolution of ≈ 1 nHz. We expect 90% of the binaries to be observed in a narrow band of (1–2.4 nHz). This is also the most sensitive part of the band. Thus, we expect many thousands of binaries per frequency bin in the relevant part of the band.

We are therefore unable to define a duty cycle as the fraction of frequency bins containing a signal. The frequency bins that contribute the signal do not contain a signal, they contain thousands. One can imagine defining the duty cycle in terms of sky location: the fraction of patches of sky which contain a signal. However, given current pulsar timing arrays, it seems unlikely that there are enough quasi-independent patches of sky so that the expected duty cycle is less than one; see, e.g., Ref. [53]. If so, and if we are not missing some other means of distinguishing supermassive black hole binary signals, it seems that the pulsar timing background is, for all intents and purposes, Gaussian in nature, at least as measured by foreseeable detectors.

D. Bursts

Gravitational-wave bursts are unmodeled transients, which can be contrasted with well-modeled signals from compact binaries. There are expected bursts from objects like supernovae, but gravitational-wave astronomers also search for unexpected bursts. Given the significant theoretical uncertainties about the loudness and rate of different gravitational-wave bursts, it is hard to say if there is a significant stochastic background from bursts, and if so, whether or not it is Gaussian. If, for example, there is a detectable background from supernovae [54], one might expect a Gaussian background since supernovae explode in the Universe at a rate of ≈ 30 s $^{-1}$. Alternatively, the burst background might be dominated by louder but less frequent signals creating a non-Gaussian background from, e.g., cusps of cosmic strings.

Given these theoretical uncertainties, it seems worthwhile to carry out a non-Gaussian search for bursts. The formalism described here can be extended. Instead of marginalizing over compact binary parameters, one can imagine marginalizing over arbitrary combinations of wavelets; see, e.g., Ref. [30]. Burst waveforms that can be easily parametrized, such as cosmic string bursts [55]

and gravitational-wave memory [56], are relatively straightforward to implement in this formalism.

E. Glitch classification

The method can also be repurposed to identify populations of glitches. To do this, we employ the formalism for non-Gaussian noise. We assume that the glitches are described by a hyperparametrized prior. Following the method described in Sec. VII, we can estimate the hyperparameters of the glitch population in order to classify populations of glitches. For example, one might find that there is a population of transients in the LIGO Hanford detector that are best fit by templates corresponding to $50 + 50M_{\odot}$ mergers. This idea is only a sketch, but we can envision developing a practical tool, which could be useful for commissioning and detector characterization. For a related discussion, see Ref. [57].

IX. SIMULTANEOUS ESTIMATION OF GAUSSIAN BACKGROUND WITH A NON-GAUSSIAN FOREGROUND

In the long run, we are interested in uncovering the primordial background likely to be lurking underneath the astrophysical background from compact binary mergers. Measurement of a primordial background is considered a holy grail of gravitational-wave astronomy, potentially allowing us to probe times well before the formation of the cosmic microwave background and to test energy scales that are not accessible through any other means; see, e.g., Ref. [58]. The problem of disentangling primordial backgrounds from astrophysical foregrounds is therefore important. The primordial background is likely to be Gaussian. In this section, we sketch out how the analysis could be extended to simultaneously measure a Gaussian background in the presence of a non-Gaussian foreground. For the sake of readability, we suppress frequency dependence as well as indices denoting segment number. The reader should consider both of these to be implied.

As a first step, we derive the likelihood for a purely Gaussian background. This derivation will be helpful in order to see how the result is generalized to include simultaneous Gaussian and non-Gaussian signals. The likelihood of obtaining strain data s given a persistent Gaussian signal \mathbf{h}_G is

$$\log[\mathcal{L}(s|\mathbf{h}_G)] \propto -\frac{1}{2} \langle s - \mathbf{h}_G, s - \mathbf{h}_G \rangle. \quad (40)$$

The inner product, defined in Eq. (3), includes an implicit sum over detectors. Variables in bold face are vectors with a different entry for each detector; e.g., $\mathbf{h}_G = (h_G^{(1)}, h_G^{(2)}, \dots)$. The strain induced in two detectors α and β is not in general the same, though the two strains are related via the overlap reduction function $\gamma_{\alpha\beta}$ [59]:

$$\langle (h_G^{(\alpha)})^\dagger h_G^{(\beta)} \rangle = \gamma_{\alpha\beta} S_h. \quad (41)$$

We do not have a template for \mathbf{h}_G because it is described by a stochastic process. Our prior for \mathbf{h}_G is

$$\pi(\mathbf{h}_G | S_h) = \frac{1}{\det(2\pi\Sigma)^{1/2}} \exp\left(-\frac{1}{2} \mathbf{h}_G^\dagger \Sigma^{-1} \mathbf{h}_G\right), \quad (42)$$

$$\mathbf{h} = \begin{pmatrix} h_G^{(1)} \\ h_G^{(2)} \end{pmatrix}, \quad (43)$$

$$\Sigma = S_h \begin{pmatrix} \gamma_{11} & \gamma_{12} \\ \gamma_{21} & \gamma_{22} \end{pmatrix}, \quad (44)$$

where S_h is the signal power spectral density. Note that both S_h and \mathbf{h}_G are implicit functions of frequency. (This frequency dependence may be described by additional parameters such as a spectral index.) The prior in Eq. (42) states that \mathbf{h}_G is a Gaussian field characterized by a power spectrum S_h and with covariance described by Eq. (41).

Before proceeding further, we introduce a simplifying assumption: that the detector network consists of two colocated detectors. This assumption is by no means necessary, but it will facilitate straightforward comparison with Ref. [16]. Employing this assumption, we can make the following substitutions:

$$\mathbf{h}_G \rightarrow h_G \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad (45)$$

$$\Sigma \rightarrow \gamma S_h \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}. \quad (46)$$

Here, γ with no indices is the overlap reduction function for a colocated detector pair. For comparison with Ref. [16], we work for the time being with ‘‘effective power’’ $S_h^{\text{eff}} \equiv \gamma S_h$.

Next, following Ref. [16] (see their Sec. IV.2), we marginalize over h_G in Eq. (42) to obtain a marginalized likelihood:

$$\mathcal{L}(\mathbf{s} | S_h^{\text{eff}}) = \frac{1}{\sqrt{2\pi \det(\mathbf{C})}} \exp\left(-\frac{1}{2} \mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}\right), \quad (47)$$

where

$$\mathbf{s} = \begin{pmatrix} s_1 \\ s_2 \end{pmatrix}, \quad (48)$$

$$\mathbf{C} = \begin{pmatrix} P + S_h^{\text{eff}} & S_h^{\text{eff}} \\ S_h^{\text{eff}} & P + S_h^{\text{eff}} \end{pmatrix}. \quad (49)$$

Following Ref. [16], we define $P \equiv \sigma^2$. Remember that there is an implied sum over frequency bins in the

expression $\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}$ and that $\det(\mathbf{C})$ includes an implied product over frequency bins. After combining data from multiple segments, there is, additionally, an implied sum over segments in the expression $\mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}$ and an implicit product over segments in the expression $\det(\mathbf{C})$.

There are two terms in the exponential of Eq. (47): an autopower term containing $s_1^2 + s_2^2$ and a cross-power term containing $s_1^* s_2$. The autopower terms are considered unreliable because a detection relying on autopower would require a precise noise budget. If the noise power spectral density includes a component that exceeds the noise budget, experimentalists assume it is an unmodeled noise, not a stochastic background. The typical solution is to start over with a likelihood constructed only out of cross-power [5]:

$$\mathcal{L}(\mathbf{s} | S_h^{\text{eff}}) = \frac{1}{\sqrt{2\pi P}} \exp[-(s_1^* s_2 - S_h^{\text{eff}})^2 / 2P^2]. \quad (50)$$

This prescription ensures that autopower from unknown noise does not create a false signal.

The prescription in Eq. (50) does not lend itself to our present purposes. Fortunately, there is a Bayesian approach that does. The notion that interferometer noise budgets are not trustworthy enough to detect excess autopower can be framed in terms of a prior belief. We treat P as a parameter with a prior distribution $\pi(P)$ peaked at P_0 with some width σ_P , which is wide compared to S_h :

$$\pi(P) \propto \exp[-(P - P_0)^2 / 2\sigma_P^2]. \quad (51)$$

That is to say, we do not trust our measurement of the noise at a level comparable to the size of the stochastic signal power S_h . Marginalizing over P , we obtain an evidence (marginalized likelihood) for each S_h^{eff} , which cannot be tricked by excess autopower:

$$\mathcal{Z}(S_h^{\text{eff}}) = \int dP \pi(P) \frac{1}{\sqrt{2\pi \det(\mathbf{C})}} \exp\left(-\frac{1}{2} \mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}\right). \quad (52)$$

When $\sigma_P \gg S_h^{\text{eff}}$, the stochastic signal encoded in the autopower is lost, and we recover something close to the cross-correlation likelihood; see Fig. 8.

In Fig. 8, we show likelihoods for a demonstration calculation with mock data. The cross-correlation likelihood [Eq. (50), hatched purple area] and the marginalized likelihood from Eq. (52) (hatched red area) are nearly identical. Both are consistent with the injected value of S_h . If we calculate the likelihood from Eq. (47), but do not marginalize over uncertainty in P , we obtain the green distribution. Since this distribution includes information from both cross- and autopower, it is narrower than the other distributions. However, it is also biased because an error in P leads to an overestimate of S_h . If we do not include a systematic error in P , the posterior peaks in the correct place. For the purposes of this sketch, we model

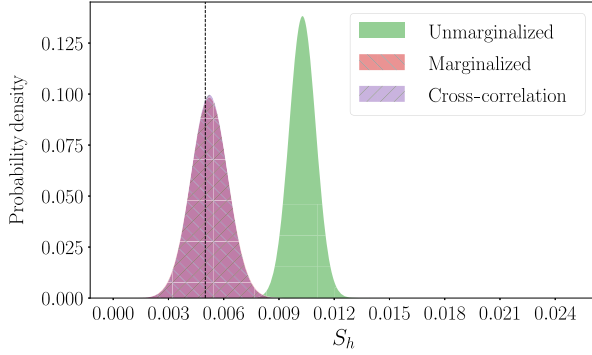


FIG. 8. A comparison of likelihood and marginalized likelihood functions for stochastic power spectral density S_h . We carry out a numerical experiment with toy-model data. We plot likelihood functions $\mathcal{L}(s|S_h)$ using cross- and autopower [Eq. (47), “Unmarginalized,” green], using cross-power only [Eq. (50), “Cross correlation,” hatched purple], and using cross- and autopower, but marginalizing over large uncertainty in the autopower [Eq. (52), “Marginalized,” hatched red]. The purple is nearly indistinguishable from the red. The green unmarginalized distribution is narrower, but it is also biased because it does not take into account uncertainty in P . (If we do not include a systematic error in P , the posterior peaks in the correct place.) The horizontal axis is strain power in arbitrary units and the vertical axis is the likelihood. The black vertical line indicates the injected value.

uncertainty in the noise power spectral density with a simple expression for $\pi(P)$. We note in passing that significantly more sophisticated models are available in order to take into account, e.g., frequency-dependent artifacts such as instrumental lines [60].

Now we are ready to sketch out a method to detect a background with both Gaussian and non-Gaussian components. The likelihood includes a Gaussian component h_G and a non-Gaussian component $h_{\text{NG}}(\theta)$:

$$\log[\mathcal{L}(s|\mathbf{h}_G, \mathbf{h}_{\text{NG}})] \propto -\frac{1}{2} \langle \mathbf{s} - \mathbf{h}_{\text{NG}} - \mathbf{h}_G, \mathbf{s} - \mathbf{h}_{\text{NG}} - \mathbf{h}_G \rangle. \quad (53)$$

The non-Gaussian signal depends implicitly on binary parameters θ . The bold face indicates a vector with entries for different detectors.

Following the same reasoning we used to calculate the Gaussian likelihood in Eq. (47), the Gaussian plus non-Gaussian likelihood is

$$\mathcal{L}(s|\mathbf{h}_{\text{NG}}, S_h) = \frac{1}{\sqrt{\det 2\pi(\mathbf{C})}} \times \exp\left(-\frac{1}{2}(\mathbf{s} - \mathbf{h}_{\text{NG}})^T \mathbf{C}^{-1}(\mathbf{s} - \mathbf{h}_{\text{NG}})\right). \quad (54)$$

By introducing a suitable prior $\pi(P)$ and marginalizing over uncertainty in the detector noise, it should be possible

to simultaneously infer the existence of a population of unresolved binaries and a continuous Gaussian background:

$$\mathcal{Z}(\mathbf{h}_{\text{NG}}, S_h) = \int dP \pi(P) \mathcal{L}(s|\mathbf{h}_{\text{NG}}, S_h). \quad (55)$$

We note that, by marginalizing over uncertainty in P , we should not hamper our ability to measure binary signals. Individual binary signals contribute to our evidence by matching with waveform templates, not by inducing excess power. Indeed, parametrized noise models are a fixture of recent transient analysis [30].

Using the Gaussian plus non-Gaussian likelihood function in Eq. (55), we can marginalize over the implicit binary parameters θ to obtain signal and “noise” evidences analogous to Eqs. (8) and (9):

$$\mathcal{Z}_S(S_h) = \int d\theta \mathcal{Z}(\mathbf{h}_{\text{NG}}, S_h), \quad (56)$$

$$\mathcal{Z}_N(S_h) = \mathcal{Z}(\mathbf{h}_{\text{NG}} = \mathbf{0}, S_h). \quad (57)$$

These evidences are functions of the hyperparameter S_h , which describes the Gaussian background. We refer to the “noise” evidence with quotation marks because it contains a Gaussian background signal. In order to calculate them in LALINFERENCE, one must implement the revised likelihood function [Eq. (55)] with a new parameter S_h and a suitable prior $\pi(P)$. The next step is to introduce the duty cycle as in Eq. (5) as in Sec. II. Retracing our steps, we obtain a posterior $p(\xi, S_h | \{\vec{s}\})$ analogous to Eq. (11). We may convert S_h into Ω_G , the Gaussian energy density; see, e.g., Ref. [61]:

$$\Omega_G(f) = \frac{2\pi^2 f^3}{3H_0^2} S_h, \quad (58)$$

where H_0 is the Hubble parameter.

While significant work is required to go beyond this sketch and demonstrate this technique, we hope this proposal provides a useful outline to construct the optimal method to simultaneously detect Gaussian and non-Gaussian backgrounds. In particular, we expect the optimal method to improve upon various schemes in which astrophysical events are fit separately and then subtracted; see, e.g., Ref. [62]. It is worth investigating as a promising tool for future efforts to measure primordial backgrounds.

X. CONCLUSIONS

Preliminary estimates suggest that advanced detectors, operating at design sensitivity, can detect a stochastic background from binary black holes in about 1 day. These estimates rely on extrapolation using Gaussian mixture modeling of our Bayesian evidence distributions.

The next step is to carry out a mock data challenge in which we demonstrate the safety and efficacy of the search using ≈ 1 day of design sensitivity Monte Carlo data. Such a demonstration would allow us to verify the extrapolations made here with a modest computational cost $\approx 500\,000$ core hours.

We have highlighted new directions worthy of deeper investigation beyond the overview we provide here. It will be interesting to more fully develop this method for other audio-band sources of gravitational waves including binary neutron stars, continuous waves, unmodeled bursts, and glitches. The method does not appear to be helpful for pulsar timing. We look forward to demonstrating the simultaneous detection of a Gaussian background in the presence of a non-Gaussian foreground.

This formulation of binary black hole detection provides a unified framework for the analysis of both resolvable signals and a stochastic background of unresolvable signals. It is also a natural framework to carry out analysis of the population properties of binary black holes. Since the resolved and unresolved binaries are analyzed as a single data set, it is possible to eliminate selection effects.

ACKNOWLEDGMENTS

We thank Joe Romano, Letizia Sammut, Tom Callister, Max Isi, Xingjiang Zhu, Pablo Rosado, and Kent Blackburn for helpful comments. R. S. and E. T. are supported by CE170100004. E. T. is supported through ARC FT150100281. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This manuscript has LIGO Document ID LIGO-P1700407.

APPENDIX: DERIVATION OF ξ_{true} FOR MONTE CARLO STUDY

In this appendix, we justify the choice of $\xi_{\text{true}} = 4 \times 10^{-4}$ as a reasonable choice of duty cycle for our Monte Carlo study in order to simulate a plausible stochastic background. We show $\xi_{\text{true}} = 4 \times 10^{-4}$ implies an energy density spectrum that is roughly consistent with the prediction of Ref. [1], and that such a background cannot be detected with cross-correlation in 1 year of design sensitivity Hanford-Livingston data.

The first step is to calculate the average gravitational-wave power spectral density $S_h(f)$ for data segments in which a signal is present:

$$S_h(f) = \int d\theta \pi(\theta) \frac{|\tilde{h}_+(f|\theta)|^2 + |\tilde{h}_\times(f|\theta)|^2}{\mathcal{N}}. \quad (\text{A1})$$

Here, $\tilde{h}_+(f|\theta)$ and $\tilde{h}_\times(f|\theta)$ are the + and \times components of the Fourier transform of the metric perturbation (embedded

in a 4-s segment) given binary parameters θ . Meanwhile, $\pi(\theta)$ is the prior distribution assumed for our analysis (described in Sec. III) and \mathcal{N} is a normalization factor to ensure that $S_h(f)$ is a power spectral density with units of Hz^{-1} . In practice, we carry out the integral numerically using N Monte Carlo draws:

$$S_h(f) = \frac{1}{N} \sum_{k=1}^N \frac{|\tilde{h}_+(f|\theta_k)|^2 + |\tilde{h}_\times(f|\theta_k)|^2}{\mathcal{N}}. \quad (\text{A2})$$

Given a duty cycle ξ , the average power spectral density in a data set with noise and signal is given by

$$S'_h(f) = \xi S_h(f). \quad (\text{A3})$$

The signal amplitude spectral density $[S_h(f)]^{1/2}$ is plotted in Fig. 9(a) (blue line) alongside the noise amplitude spectral density (orange line).

The energy density associated with S_h is

$$\Omega_{\text{gw}}(f) = \frac{2\pi^2 f^3}{3H_0^2} S'_h(f), \quad (\text{A4})$$

where H_0 is the Hubble parameter. In Fig. 9(b), we plot $\Omega_{\text{gw}}(f)$ for the Monte Carlo data set used in our analysis given $\xi = 4 \times 10^{-4}$ (solid blue line). Alongside, in the thick black line, we plot a power-law integrated sensitivity curve [61], which shows the 1σ sensitivity of the cross-correlation search for the Hanford-Livingston network operating at design sensitivity for 1 year.

Since our data set includes only binaries with luminosity distances $d_L \leq 5$ Gpc, our data set includes only a fraction of the full stochastic background. Using Fig. 2(c) from Ref. [4], we estimate that the full stochastic background has an energy density about 3 times larger than the stochastic background from events with $z < 0.77$. The red dashed curve shows $\Omega_{\text{gw}}(f)$ scaled by a factor of 3 to roughly estimate the total gravitational-wave energy density including the contribution from high redshifts. The dashed red line reaches a value of 3×10^{-10} at 25 Hz, which is just below the allowed range predicted in Ref. [1]: $\Omega_{\text{gw}}(f = 25 \text{ Hz}) = 1.1^{+1.2}_{-0.7} \times 10^{-9}$. The fact that the dashed red line does not intersect the black sensitivity curve indicates that the background is not detectable with a year of data using cross-correlation.

The signal-to-noise ratio of the cross-correlation search is given by

$$\rho_{\text{cc}} = \sqrt{2n} \left[\sum_k \frac{\Gamma_{12}^2(f_k) S_h^2(f_k)}{P_1(f_k) P_2(f_k)} \right]^{1/2}. \quad (\text{A5})$$

See Eq. (22) in Ref. [61]. The sum over k ranges over frequency bins. Here, P_1, P_2 are the noise power spectral

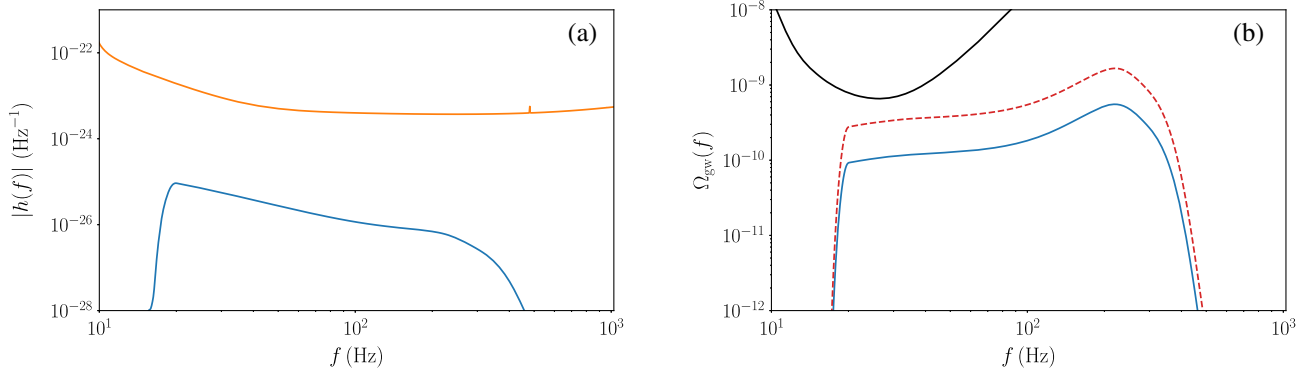


FIG. 9. (a) Amplitude spectral density $[S_h(f)]^{1/2}$. Blue is for the data set used in our mock data set ($z < 0.77$). The drop at 20 Hz is an artifact from our cutoff frequency. The orange curve is the LIGO design sensitivity noise curve. (b) Energy density $\Omega_{\text{gw}}(f)$. Solid blue is for the data set used in our mock data set ($z < 0.77$) while dashed red is the same as solid blue, but including binaries out to arbitrarily high redshifts. The dashed red background is 3 times higher than the solid blue. The black curve is the 1σ power-law sensitivity curve [61], indicating the sensitivity of the stochastic cross-correlation search assuming a LIGO Hanford-Livingston network operating at design sensitivity for 1 year. Since the red dashed curve is below the black sensitivity curve, the background is not detectable with cross-correlation after 1 year.

densities of detectors 1 and 2. The variable Γ_{12} is the overlap reduction function given in Eq. (15) of Ref. [61] and derived in Ref. [59]. Assuming 1 year of data at design sensitivity, we obtain $\rho_{\text{cc}} = 0.15$ for the data set used in our analysis, which is limited to redshifts $z < 0.77$. Extrapolating to arbitrarily high redshifts, we estimate $\rho_{\text{cc}} = 0.46$.

-
- [1] B. P. Abbott *et al.*, *GW170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences*, *Phys. Rev. Lett.* **120**, 091101 (2018).
- [2] B. P. Abbott *et al.*, *GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2*, *Phys. Rev. Lett.* **118**, 221101 (2017).
- [3] B. P. Abbott *et al.*, *GW150914: Implications for the Stochastic Gravitational-Wave Background from Binary Black Holes*, *Phys. Rev. Lett.* **116**, 131102 (2016).
- [4] T. Callister, L. Sammut, S. Qiu, I. Mandel, and E. Thrane, *Limits of Astrophysics with Gravitational-Wave Backgrounds*, *Phys. Rev. X* **6**, 031018 (2016).
- [5] B. Allen and J. D. Romano, *Detecting a Stochastic Background of Gravitational Radiation: Signal Processing Strategies and Sensitivities*, *Phys. Rev. D* **59**, 102001 (1999).
- [6] B. P. Abbott *et al.*, *Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run*, *Phys. Rev. Lett.* **118**, 121101 (2017).
- [7] B. P. Abbott *et al.*, *Directional Limits on Persistent Gravitational Waves from Advanced LIGO's First Observing Run*, *Phys. Rev. Lett.* **118**, 121102 (2017).
- [8] R. W. Hellings and G. S. Downs, *Upper Limits on the Isotropic Gravitational Radiation Background from Pulsar Timing Analysis*, *Astrophys. J. Lett.* **265**, L39 (1983).
- [9] P. A. Rosado, *Gravitational Wave Background from Binary Systems*, *Phys. Rev. D* **84**, 084004 (2011).

- [10] C. Cutler, *Angular Resolution of the LISA Gravitational Wave Detector*, *Phys. Rev. D* **57**, 7089 (1998).
- [11] A practical working definition of “bursting” sources is transient signals that are not compact binary coalescences and which are not modeled with a bank of templates.
- [12] S. Drasco and E. E. Flanagan, *Detection Methods for Non-Gaussian Gravitational Wave Stochastic Backgrounds*, *Phys. Rev. D* **67**, 082003 (2003).
- [13] E. Thrane, *Measuring the Non-Gaussian Stochastic Gravitational-Wave Background: A Method for Realistic Interferometer Data*, *Phys. Rev. D* **87**, 043009 (2013).
- [14] L. Martellini and T. Regimbau, *Semiparametric Approach to the Detection of Non-Gaussian Gravitational Wave Stochastic Backgrounds*, *Phys. Rev. D* **89**, 124009 (2014).
- [15] L. Martellini and T. Regimbau, *Efficiency of the Cross-Correlation Statistic for Gravitational Wave Stochastic Background Signals with Non-Gaussian Noise and Heterogeneous Detector Sensitivities*, *Phys. Rev. D* **92**, 104025 (2015).
- [16] J. D. Romano and N. J. Cornish, *Detection Methods for Stochastic Gravitational-Wave Backgrounds: A Unified Treatment*, *Living Rev. Relativity* **20**, 2 (2017).
- [17] K. Cannon, C. Hanna, and D. Keppel, *Method to Estimate the Significance of Coincident Gravitational-Wave Observations from Compact Binary Coalescence*, *Phys. Rev. D* **88**, 024025 (2013).
- [18] B. P. Abbott *et al.*, *Binary Black Hole Mergers in the First Advanced LIGO Observing Run*, *Phys. Rev. X* **6**, 041015 (2016).
- [19] G. Malmquist, *On Some Relations in Stellar Statistics*, *Medd. Lunds Astron. Obs. Ser. I* **100**, 1 (1922).
- [20] J. Veitch *et al.*, *Parameter Estimation for Compact Binaries with Ground-Based Gravitational-Wave Observations Using the LALInference Software Library*, *Phys. Rev. D* **91**, 042003 (2015).
- [21] H. Jeffreys, *Theory of Probability*, 3rd ed. (Oxford, Oxford, England, 1961).

- [22] X.-J. Zhu, E. Howell, T. Regimbau, D. Blaire, and Z.-H. Zhu, *Stochastic Gravitational Wave Background from Coalescing Binary Black Holes*, *Astrophys. J.* **739**, 86 (2011).
- [23] E. S. Phinney, *A Practical Theorem on Gravitational Wave Backgrounds*, arXiv:astro-ph/0108028.
- [24] J. Aasi *et al.*, *Advanced LIGO*, *Classical Quantum Gravity* **32**, 115012 (2015).
- [25] To include additional detectors, one simply uses the expression for the likelihood function in Eq. (4) setting M to the required number of detectors.
- [26] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms*, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [27] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, *Fast and Accurate Inference on Gravitational Waves from Precessing Compact Binaries*, *Phys. Rev. D* **94**, 044031 (2016).
- [28] D. Reynolds, in *Encyclopedia of Biometrics*, edited by S. Z. Li and A. Jain (Springer, Boston, 2009) pp. 659–663.
- [29] B. P. Abbott *et al.*, *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [30] N. J. Cornish and T. B. Littenberg, *Bayeswave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches*, *Classical Quantum Gravity* **32**, 135012 (2015).
- [31] E. A. Huerta, R. Haas, E. Fajardo, D. S. Katz, S. Anderson, P. Couvares, J. Willis, T. Bouvet, J. Enos, W. T. C. Kramer, H. W. Leong, and D. Wheeler, *BOSS-LDG: A Novel Computational Framework that Brings Together Blue Waters, Open Science Grid, Shifter and the LIGO Data Grid to Accelerate Gravitational Wave Discovery*, in *Proceedings of the 2017 IEEE 13th Annual Conference on e-Science* (<http://ieeexplore.ieee.org/document/8109152/>).
- [32] M. Isi, R. Smith, S. Vitale *et al.* (to be published).
- [33] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. (Chapman & Hall, Boca Raton, 2013).
- [34] M. R. Adams, N. J. Cornish, and T. B. Littenberg, *Astrophysical Model Selection in Gravitational Wave Astronomy*, *Phys. Rev. D* **86**, 124032 (2012).
- [35] C. Talbot and E. Thrane, *Determining the Population Properties of Spinning Black Holes*, *Phys. Rev. D* **96**, 023012 (2017).
- [36] S. Stevenson, F. Ohme, and S. Fairhurst, *Distinguishing Compact Binary Population Synthesis Models Using Gravitational Wave Observations of Coalescing Binary Black Holes*, *Astrophys. J.* **810**, 58 (2015).
- [37] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tio, and J. Veitch, *Model-Independent Inference on Compact-Binary Observations*, *Mon. Not. R. Astron. Soc.* **465**, 3254 (2017).
- [38] M. Fishbach, D. Holz, and B. Farr, *Are LIGO's Black Holes Made From Smaller Black Holes?*, *Astrophys. J. Lett.* **840**, L24 (2017).
- [39] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, *Measuring the Spin of Black Holes in Binary Systems Using Gravitational Waves*, *Phys. Rev. Lett.* **112**, 251101 (2014).
- [40] S. Stevenson, C. Berry, and I. Mandel, *Hierarchical Analysis of Gravitational-Wave Measurements of Binary Black Hole Spin-Orbit Misalignments*, *Mon. Not. R. Astron. Soc.* **471**, 2801 (2017).
- [41] C. Talbot and E. Thrane, *Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization*, arXiv:1801.02699.
- [42] C. Callister, M. Isi, R. Smith, and E. Thrane (to be published).
- [43] B. P. Abbott *et al.*, *Observation of Gravitational Waves from a Binary Black Hole Merger*, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [44] D. R. Lorimer, L. Nicastrò, A. G. Lyne, M. Bailes, R. N. Manchester, S. Johnston, J. F. Bell, N. D'Amico, and P. A. Harrison, *Four New Millisecond Pulsars in the Galactic Disk*, *Astrophys. J.* **439**, 933 (1995).
- [45] D. Talukder, E. Thrane, S. Bose, and T. Regimbau, *Measuring Neutron-Star Ellipticity with Measurements of the Stochastic Gravitational-Wave Background*, *Phys. Rev. D* **89**, 123008 (2014).
- [46] B. P. Abbott *et al.*, *Comprehensive All-Sky Search for Periodic Gravitational Waves in the Sixth Science Run LIGO Data*, *Phys. Rev. D* **94**, 042002 (2016).
- [47] X. Fan, Y. Chen, and C. Messenger, *Method to Detect Gravitational Waves from an Ensemble of Known Pulsars*, *Phys. Rev. D* **94**, 084029 (2016).
- [48] G. Hobbs and S. Dai, *A Review of Pulsar Timing Array Gravitational Wave Research*, arXiv:1707.01615.
- [49] P. A. Rosado, A. Sesana, and J. Gair, *Expected Properties of the First Gravitational Wave Signal Detected with Pulsar Timing Arrays*, *Mon. Not. R. Astron. Soc.* **451**, 2417 (2015).
- [50] A. Sesana, *Systematic Investigation of the Expected Gravitational Wave Signal from Supermassive Black Hole Binaries in the Pulsar Timing Band*, *Mon. Not. R. Astron. Soc. Lett.* **433**, L1 (2013).
- [51] The probability distribution $\pi(f)$ can be derived from the post-Newtonian approximation for the time to coalescence from a frequency f , sometimes known as the “chirp time.” To leading order, the chirp time τ is related to frequency by $\tau \propto f^{-8/3}$ [52]. By conservation of probability $\pi(f)df = \pi(\tau(f))d\tau$, and assuming $\pi(\tau)$ is constant, which is to say there is no preferred coalescence time, we arrive at $\pi(f) = f^{-11/3}$.
- [52] L. Blanchet, *Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries*, *Living Rev. Relativity* **17**, 2 (2014).
- [53] X.-J. Zhu, L. Wen, J. Xiong, Y. Xu, Y. Wang, S. D. Mohanty, G. Hobbs, and R. N. Manchester, *Detection and Localization of Continuous Gravitational Waves with Pulsar Timing Arrays: The Role of Pulsar Terms*, *Mon. Not. R. Astron. Soc.* **461**, 1317 (2016).
- [54] K. Crocker, T. Prestegard, V. Mandic, T. Regimbau, K. Olive, and E. Vangioni, *Phys. Rev. D* **95**, 063015 (2017).
- [55] T. Damour and A. Vilenkin, *Gravitational Wave Bursts from Cosmic Strings*, *Phys. Rev. Lett.* **85**, 3761 (2000).
- [56] M. Favata, *The Gravitational-Wave Memory Effect*, *Classical Quantum Gravity* **27**, 084036 (2010).

- [57] J. Powell, D. Trifiro, E. Cuoco, I.-S. Heng, and M. Cavaglia, *Classical Quantum Gravity* **32**, 215012 (2015).
- [58] M. Maggiore, *Gravitational Wave Experiments and Early Universe Cosmology*, *Phys. Rep.* **331**, 283 (2000).
- [59] N. Christensen, *Measuring the Stochastic Gravitational-Radiation Background with Laser-Interferometric Antennas*, *Phys. Rev. D* **46**, 5250 (1992).
- [60] C. Cahillane, J. Betzwieser, D. A. Brown, E. Goetz, E. D. Hall, K. Izumi, S. Kandhasamy, S. Karki, J. S. Kissel, G. Mendell, R. L. Savage, D. Tuyenbayev, A. Urban, A. Viets, M. Wade, and A. J. Weinstein, *Calibration Uncertainty for Advanced LIGO's First and Second Observing Runs*, *Phys. Rev. D* **96**, 102001 (2017).
- [61] E. Thrane and J. D. Romano, *Sensitivity Curves for Searches for Gravitational-Wave Backgrounds*, *Phys. Rev. D* **88**, 124032 (2013).
- [62] C. Cutler and J. Harms, *Big Bang Observer and the Neutron-Star-Binary Subtraction Problem*, *Phys. Rev. D* **73**, 042001 (2006).