

Bayesian coherent analysis of in-spiral gravitational wave signals with a detector networkJ. Veitch^{1,2} and A. Vecchio¹¹*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom*²*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom*

(Received 19 November 2009; published 18 March 2010)

The present operation of the ground-based network of gravitational-wave laser interferometers in enhanced configuration and the beginning of the construction of second-generation (or advanced) interferometers with planned observation runs beginning by 2015 bring the search for gravitational waves into a regime where detection is highly plausible. The development of techniques that allow us to discriminate a signal of astrophysical origin from instrumental artefacts in the interferometer data and to extract the full range of information are therefore some of the primary goals of the current work. Here we report the details of a Bayesian approach to the problem of inference for gravitational wave observations using a network (containing an arbitrary number) of instruments, for the computation of the Bayes factor between two hypotheses and the evaluation of the marginalized posterior density functions of the unknown model parameters. The numerical algorithm to tackle the notoriously difficult problem of the evaluation of large multidimensional integrals is based on a technique known as nested sampling, which provides an attractive (and possibly superior) alternative to more traditional Markov-chain Monte Carlo methods. We discuss the details of the implementation of this algorithm and its performance against a Gaussian model of the background noise, considering the specific case of the signal produced by the in-spiral of binary systems of black holes and/or neutron stars, although the method is completely general and can be applied to other classes of sources. We also demonstrate the utility of this approach by introducing a new coherence test to distinguish between the presence of a coherent signal of astrophysical origin in the data of multiple instruments and the presence of incoherent accidental artefacts, and the effects on the estimation of the source parameters as a function of the number of instruments in the network.

DOI: [10.1103/PhysRevD.81.062003](https://doi.org/10.1103/PhysRevD.81.062003)

PACS numbers: 04.80.Nn, 02.70.Rr, 02.70.Uu

I. INTRODUCTION

Searches for gravitational waves are entering a crucial stage with the network of ground-based laser interferometers—LIGO [1,2], Virgo [3], and GEO 600 [4]—now fully operational and engaged in a new year-long data taking period [5,6] at enhanced sensitivity, which may allow the first direct detection of gravitational radiation. Construction has already begun for the upgrade of the instruments to advanced configuration (second-generation interferometers) with installation at the sites that will start at the end of 2011 [6–8]. When science observations resume at much improved sensitivity by 2015, several gravitational wave events are expected to be observed, opening a new means to explore a variety of astrophysical phenomena (see e.g. Refs. [9,10] and references therein).

Coalescing binary systems of compact objects—black holes and neutron stars—will be the workhorse source for gravitational wave observations. Ground-based laser interferometers will monitor the last seconds to minutes of the coalescence of these systems. The theoretical modelling of the (in-spiral) waveform is well in hand (see e.g. [11] and references therein), and the search algorithms are well understood [12–19]. The detection rate for ongoing searches and observations with second-generation instruments is estimated to lie in the range $9 \times$

10^{-5} yr^{-1} – 0.7 yr^{-1} and 0.2 yr^{-1} – 1000 yr^{-1} , respectively, see [20] for a review. It is likely that in a few years time ground-based laser interferometers will allow us to extract a wealth of new information ranging from the formation and evolution of binary stars, the nature of precursors of (short) gamma-ray bursts, dynamical processes in star clusters, and could yield a new set of standard candles for precise cosmography.

As instruments are beginning to operate at a meaningful sensitivity from an astrophysical, cosmological and fundamental physics point of view, much emphasis is now being placed on the development of methods that offer the maximum discriminating power to separate disturbances of instrumental origin from a true astrophysical signal, and to extract the full range of information from the detected signals. Bayesian inference provides a powerful approach to both model selection (or hypothesis testing) and parameter estimation. Despite the conceptual simplicity of the Bayesian framework, there has been only limited use of these methods for ground-based gravitational-wave data analysis due to their computational burden, in this case related to the need to compute large multidimensional integrals. Additionally, the Gaussian likelihood functions considered so far do not address the instrumental glitches which are present in data from the current generation of gravitational wave detectors.

Here we present an efficient method to compute concurrently the full set of quantities at the heart of Bayesian inference: the Bayes factor between competing hypotheses and the posterior density functions (PDFs) on the relevant model parameters. The method is based on the *nested sampling* algorithm [21–23] to perform multidimensional integrals, that present the practical and computationally intensive challenge for the implementation of Bayesian methods. We demonstrate the algorithm by considering multidetector observations of gravitational waves generated during the in-spiral phase of the coalescence of a binary system, modeled using the restricted $\text{post}^{2.0}$ -Newtonian stationary phase approximation, which is the waveform used so far for searches of nonspinning binary objects [24]. Initial results based on this method and applied to simplified gravitational waveforms were reported in [25,26]. An application to the study of different waveform approximants to detect and estimate the parameters of signals generated through the numerical integration of the Einstein’s equations for the two body problem in the context of the Numerical INjection Analysis (NINJA) Project was reported in [27–29]. In this paper we:

- (i) Provide for the first time full details about the theoretical and technical issues on which the computation of the evidence is based;
- (ii) Discuss the errors associated with the computation of the integrals and the associated computational costs;
- (iii) Show how from the nested samples one can construct at negligible computational cost the marginalized posterior PDFs on the source parameters;
- (iv) Demonstrate the performance of this technique in detecting a binary in-spiral signal against a Gaussian model of background noise in coherent observations using a network of detectors, and by introducing a new *coherence test* to distinguish between the presence of a coherent signal of astrophysical origin in the data of multiple instruments and the presence of incoherent accidental artefacts, and
- (v) Show the effects of the number of instruments in the network on the estimation of the source parameters.

The method in its present implementation can be extended in a straightforward way to the full coalescence waveform of binary systems—in fact the first example applied on the NINJA data set is described in [28,29]—and the software is available as part of the LSC Analysis Library Applications (LALapps) [30,31]. Work is already ongoing in this direction. Furthermore, the approach can also be extended to other gravitational wave signals.

The paper is organized as follows: in Sec. II we review the key concepts of Bayesian inference and the signal generated by in-spiralling gravitational wave signals; in Sec. III we describe our implementation of the numerical computation of multidimensional integrals for the computation of the marginal likelihood and marginalized poste-

rior density functions based on nested-sampling the limited likelihood function; Sec. IV contains the implementation details of the algorithm, including a quantification of the errors that affect the results as a function of the choice of the main tuning parameters of the algorithm and the effects on the scaling of the computational costs. Section V contains the results from the applications of this method to several test cases, including a Bayesian coherence test that we introduce here for the first time.

Throughout the paper we use geometric units, in which $G = c = 1$.

II. BAYESIAN INFERENCE FOR BINARY IN-SPIRALS

Statistical inference can be roughly divided into two problems: (i) Model selection, or hypothesis testing, between competing hypotheses through the computation of the *evidences* (or marginal likelihoods) of models, and (ii) parameter estimation (of the unknown parameters on which a model depends). In the context of Bayesian inference, both aspects are simply tackled through an application of Bayes’ theorem and the standard rules of probability theory.

While this approach is mathematically straightforward, its implementation is hampered by the need to explore large parameters spaces and perform what are in general computationally costly, high-dimensional integrals. For gravitational waves generated by binaries with negligible spins and eccentricity—the case considered in this paper—the number of dimensions of parameter space is 9, and for a generic binary system (described by general relativity) the total number of parameters increases up to 17. This technical aspect is one of the main factors that has limited the application of a Bayesian approach in a number of problems.

Model selection has been tackled through a number of techniques, including reversible jump Markov-chain Monte Carlo (MCMC) [32] and thermodynamic integration [33]. Parameter estimation is usually dealt with using MCMC methods [34,35], which may include advanced techniques such as parallel tempering [36,37] and delayed rejection [38–41] to enhance the exploration of the parameter space. For applications of these and other Bayesian methods in ground-based gravitational wave observations, see e.g. [26,29,42–53].

Nested sampling [21–23] is a powerful numerical technique to deal with multidimensional integrals. It differs from other Monte Carlo techniques such as MCMC methods [34,35] that are popular in applications of Bayesian inference, in that it is specifically designed to estimate the evidence integral itself, see Eq. (4), with the marginalized posterior PDFs being optional by-products.

In this section we outline the key concepts of Bayesian inference and review the signal model—in-spiral signals generated by nonspinning binary systems in circular or-

bits—on which we concentrate the development and application of the algorithm. Many of the technical implementation aspects are associated to sampling effectively the likelihood function in the sky position parameters, and are therefore completely general for any application to short-lived bursts characterized by an arbitrary waveform.

A. Model selection

In the formalism of Bayesian inference, the probability of a model or hypothesis \mathcal{H}_i , given a set of observational data \vec{d} and prior information I is given by Bayes' theorem,

$$P(\mathcal{H}_i|\vec{d}, I) = \frac{P(\mathcal{H}_i|I)P(\vec{d}|\mathcal{H}_i, I)}{P(\vec{d}|I)}. \quad (1)$$

In this expression, $P(\mathcal{H}_i|I)$ is the prior probability of \mathcal{H}_i , $P(\vec{d}|\mathcal{H}_i, I)$ is the *likelihood function* of the data, given that \mathcal{H}_i is true, and

$$P(\vec{d}|I) = \sum_i P(\vec{d}|\mathcal{H}_i, I)$$

is the marginal probability of the data set \vec{d} , which can only be calculated if there exists a complete set of independent hypotheses such that $\sum_j P(\mathcal{H}_j|\vec{d}, I) = 1$. Here we do not enumerate such a set of models, but we can still make comparisons between the models we do have by calculating the relative probabilities in the form of the posterior *odds ratio* O_{ij} between two of them,

$$O_{i,j} = \frac{P(\mathcal{H}_i|I) P(\vec{d}|\mathcal{H}_i, I)}{P(\mathcal{H}_j|I) P(\vec{d}|\mathcal{H}_j, I)} = \frac{P(\mathcal{H}_i|I)}{P(\mathcal{H}_j|I)} B_{ij}; \quad (2)$$

in the previous equation the normalization factor $P(\vec{d}|I)$ cancels out, and

$$B_{i,j} \equiv \frac{P(\vec{d}|\mathcal{H}_i, I)}{P(\vec{d}|\mathcal{H}_j, I)} \quad (3)$$

is known as the *Bayes factor* or ratio of likelihoods.

The Bayes factors can be directly found for hypotheses which have no free parameters, but the gravitational wave signal we are modelling depends on a set of parameters, $\vec{\theta} \in \Theta$, described in Section II C, where Θ is the parameter space. In this case, the likelihood of the model \mathcal{H} must be marginalized over all the parameters weighted by their prior probability distribution, giving the *marginal likelihood* or *evidence*,

$$Z = P(\vec{d}|\mathcal{H}, I) = \int_{\Theta} p(\vec{\theta}|\mathcal{H}, I) p(\vec{d}|\mathcal{H}, \vec{\theta}, I) d\vec{\theta}, \quad (4)$$

where $p(\vec{\theta}|\mathcal{H}, I)$ is the prior probability distribution over the parameter space.

The integral (4) cannot be computed analytically in all but the most trivial cases, and standard grid-based numerical approaches can take a prohibitively long time to com-

plete when the model is high-dimensional and/or very large with respect to the posterior as is the case for gravitational wave observations. By using the nested sampling algorithm, developed by Skilling [21], we have been able to solve the problem of calculating this integral—and as a consequence the desired Bayes factors *and* the marginalized posterior density functions of the unknown model parameters $\vec{\theta}$ —in a time that makes Bayesian techniques applicable in actual gravitational wave search pipelines. Section III provides an overview of the algorithm, and the implementation strategy that we have adopted. Further implementation details, as well as the characterization of its accuracy in the evaluation of the evidence integral as a function of CPU time are discussed in Sec. IV.

B. Parameter estimation

In general the hypotheses depend on a set of unknown parameters $\vec{\theta} \in \Theta$. As part of the inference process, one wants also to compute the PDF

$$p(\vec{\theta}|\vec{d}, \mathcal{H}, I) = \frac{p(\vec{\theta}|\mathcal{H}, I) p(\vec{d}|\vec{\theta}, \mathcal{H}, I)}{p(\vec{d}|\mathcal{H}, I)} \quad (5)$$

of the parameters, in this specific case of binary systems quantities such as the masses, position in the sky and distance.

The marginalized PDF on a subset $\vec{\theta}_A$ of the parameters—our notation is $\vec{\theta} \equiv \{\vec{\theta}_A, \vec{\theta}_B\}$, $\vec{\theta}_{A,B} \in \Theta_{A,B}$ —is defined as

$$p(\vec{\theta}_A|\vec{d}, \mathcal{H}, I) = \int_{\Theta_B} p(\vec{\theta}|\vec{d}, \mathcal{H}, I) d\vec{\theta}_B. \quad (6)$$

From $p(\vec{\theta}_A|\vec{d}, \mathcal{H}, I)$ it is then straightforward to compute e.g. the posterior mean

$$\langle \vec{\theta}_A \rangle = \int_{\Theta_A} \vec{\theta}_A p(\vec{\theta}_A|\vec{d}, \mathcal{H}, I) d\vec{\theta}_A. \quad (7)$$

We will show in Sec. III B that the nested sampling algorithm provides a way of computing the marginalized posterior PDFs, Eq. (5), with totally negligible additional computational costs from the results of the numerical evaluation of the evidence, Eq. (4). In this respect, nested sampling may provide advantages with respect to more traditional MCMC algorithms.

C. Target waveform

In this paper we consider observations of gravitational waves from a network (of an arbitrary number) of interferometers. The datum (in the frequency domain) at frequency f from each detector that we label with D is:

$$\tilde{d}^{(D)}(f) = \tilde{h}^{(D)}(f) + \tilde{n}^{(D)}(f), \quad (8)$$

where $\tilde{h}^{(D)}(f)$ and $\tilde{n}^{(D)}(f)$ are the gravitational wave signal and noise contribution, respectively. In Sec. V we will

consider specific choices of the interferometer network, and we will label with $D = H, L, V$, the LIGO Hanford 4-km arm instrument, the LIGO Livingston interferometer and Virgo, respectively. In practice, one works with discrete data, and we will refer to $d_k^{(D)}$ (and analogously $\tilde{h}_k^{(D)}$ and $\tilde{n}_k^{(D)}$ for the signal and noise, respectively) as the data point at discrete frequency f_k of the instrument D . As shorthand notation, we will use \vec{d} to identify the whole data set from the relevant network of instruments, and to $d^{(D)}$ to the data set from a single detector D , so that

$$\vec{d} = \{d^{(H)}, d^{(L)}, \dots\}. \quad (9)$$

Let us consider a (geocentric) reference frame, and a gravitational wave source described by the two polarization amplitudes $\tilde{h}_+(f)$ and $\tilde{h}_\times(f)$ located in the sky at (α, δ) , where α is the right ascension and δ the declination of the source. The signal as measured at the output of the detector D is therefore

$$\tilde{h}^{(D)}(f) = [F_+^{(D)}\tilde{h}_+(f) + F_\times^{(D)}\tilde{h}_\times(f)]e^{-2\pi if\Delta t^{(D)}}, \quad (10)$$

where $F_+^{(D)}(\psi, \alpha, \delta; t_0)$ and $F_\times^{(D)}(\psi, \alpha, \delta; t_0)$ are the detector response functions to each polarization, dependent on the polarization angle ψ (see e.g. Appendix B of [54] for the definition conventions), and the time of observation t_0 . These are computed using functions available in the LSC Algorithm Library [30]. Given a source at location (α, δ) , $\Delta t^{(D)}(\alpha, \delta; t_0)$ is the difference in gravitational wave arrival time between the geocenter and the detector D , computed with respect to a reference time t_0 that identifies the observation, see the text after Eq. (13) below for our specific choice of t_0 . $\Delta t^{(D)}$ depends on the time of observation, as for a fixed position in the sky, the signal impinges on the instruments with different relative time delays due to the Earth's rotation. By using the transformation (10), the waveform phase which is the most expensive part of the model, needs only to be calculated once, and is then transformed to the observed signal in each detector.

In this paper, we concentrate on the in-spiral signal generated during the coalescence of a binary system of compact objects (black holes or neutron stars) of masses m_1 and m_2 . Other mass parameters that we will use are the total mass $M = m_1 + m_2$, the symmetric mass ratio $\eta = m_1 m_2 / (m_1 + m_2)^2$ and the chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$. We assume circular orbits and we further restrict to compact objects that are nonspinning. We note however, that (an earlier implementation of) the approach discussed here was already successfully applied to the case of the full coalescence waveform generated by compact binaries [28,29]. Moreover, most of the results presented in this paper are totally general and independent of the specific waveform model, and can be applied and/or extended to any class of signals.

The model for the gravitational wave signal that we consider is the frequency domain, stationary phase,

post^{2.0}-Newtonian approximation to the waveform, and more precisely the so-called ‘‘TaylorF2’’ approximant—for an up-to-date summary of the different TaylorF/T approximants we refer the reader to [55] and references therein. The waveform is therefore described by two intrinsic parameters, the two masses or any two independent combinations of them, such as \mathcal{M} and η . We note that the specific choice of the post-Newtonian order is irrelevant for the issues discussed in this paper, as long as the waveform model used to construct the likelihood function matches the one adopted in the ‘‘injections’’ [56] to generate synthetic data sets to explore the algorithm. As a consequence the results presented here would be essentially identical if we had adopted the post^{3.5}-Newtonian order which is currently used in the analysis of the LIGO/Virgo data for the current science run (S6/VSR2). We generate the waveform directly in the frequency domain using functions of the LSC Analysis Library (LAL) [30]. The frequency domain gravitational wave polarization amplitudes are given by

$$\tilde{h}_+(f) = \mathcal{A}(1 + \cos^2\iota)f^{-7/6}e^{i\Psi(f)}, \quad (11)$$

$$\tilde{h}_\times(f) = 2\mathcal{A}\cos\iota f^{-7/6}e^{i\Psi(f)-i\pi/2}. \quad (12)$$

Here, the symbol ι denotes the inclination angle, defined as the angle between the line of sight to the source from the detector and the constant direction (as the objects are assumed to be nonspinning) of the orbital angular momentum. The gravitational wave phase $\Psi(f)$ at the post²-Newtonian order is given by

$$\Psi(\mathcal{M}, \eta, t_0, \phi_0; f) = 2\pi f t_0 - \phi_0 + \psi_N(\eta) \times \sum_{k=0}^4 \psi_k(\eta)(\pi M f)^{(k-5)/3}, \quad (13)$$

where ψ_N and ψ_k are the standard Newtonian and post-Newtonian coefficients, whose expressions can be found in e.g. Ref. [57]. In our implementation, t_0 is taken as the GPS time at the geocenter at which the frequency of the gravitational wave passes that of the nominal innermost stable circular orbit, $f_{\text{ISCO}} = (6^{3/2}\pi M)^{-1}$, and consequently ϕ_0 is the phase of the signal at this time. The amplitude of the gravitational wave $\mathcal{A} \propto \mathcal{M}^{5/6}/D_L$ and is computed by the LAL stationary phase approximation template [30]. In summary, the observed signal is therefore dependent on nine quantities, which for convenience we will write as the parameter vector

$$\vec{\theta} = \{\mathcal{M}, \eta, t_0, \phi_0, D_L, \alpha, \delta, \psi, \iota\}. \quad (14)$$

Finally, we discuss the assumptions on the noise $n^{(D)}$. We will make the standard assumption that the noise is a Gaussian and stationary process with zero mean and variance described through the one-sided noise spectral density $S_n^{(D)}(f)$:

$$\langle \tilde{n}(f)^{(D)} \rangle = 0, \quad (15)$$

$$\langle \tilde{n}^{(D)}(f) \tilde{n}^{(D)*}(f') \rangle = \frac{1}{2} \delta(f - f') S(f), \quad (16)$$

where $\langle \cdot \rangle$ stands for the ensemble average. Under these assumptions, the likelihood of a given noise realisation $n^{(D)} = n_0$ is simply given by the multivariate Gaussian distribution

$$p(n^{(D)} = n_0) \propto e^{-(n_0|n_0)/2}, \quad (17)$$

where $(\cdot | \cdot)$ stands for the usual inner product [58], see Eq. (A17) of the appendix.

We will further assume that the noise in different detectors is uncorrelated, so that we generalize Eq. (16) to

$$\langle \tilde{n}^{(D)}(f) \tilde{n}^{(D')*}(f') \rangle = \frac{1}{2} \delta(f - f') \delta_{DD'} S^{(D)}(f). \quad (18)$$

The latter assumption is appropriate for sites that are well isolated from each other, but this may not be true for the two instruments colocated at the Hanford site. Here we only consider a simulated network with no more than one instrument at any location. In terms of the elements \tilde{n}_k of the discrete Fourier series of the discretely-sampled time domain data, with sampling interval Δt and *segment length* T , this is given as $\Delta t^2 \langle |\tilde{n}_k|^2 \rangle = \frac{T}{2} S(f_k)$. Full details of the conventions used for discretely sampled data are given in the appendix.

D. Models

The problem of assessing the confidence of detection of a signal in interferometer data is the primary motivation for the nested sampling technique and implementation we present here. Translated into the Bayesian framework, assessing the confidence of detection means computing the Bayes factor between two hypotheses, and therefore we must specify exactly which models we are comparing. These models are the mathematical descriptions of the data \vec{d} , Eqs. (8) and (9), which either contain a gravitational wave signal $\vec{h}(\vec{\theta})$ parametrized by a certain vector $\vec{\theta}$ (described in section II C), or it does not. In addition, the data also contain a contribution from instrumental noise, described by Eqs. (15)–(18). The two models we will use can be written as:

- (i) \mathcal{H}_N : the noise-only model, that corresponds to the hypothesis that there is only noise (with statistical properties described in Sec. II C) present in the data set:

$$\vec{d} = \vec{n}; \quad (19)$$

note that in our application we assume that the noise spectral density is known [59] and this model has therefore no free parameters;

- (ii) \mathcal{H}_S : the signal model, that corresponds to the hypothesis that the data contain noise (as before) and a gravitational wave described by the waveform fam-

ily $h(\vec{\theta})$:

$$\vec{d} = \vec{n} + \vec{h}(\vec{\theta}). \quad (20)$$

Although in reality, there is also a wide range of instrumental glitches and artefacts which alters the evidence of each model in a variety of ways, initially we focus on characterizing the algorithm with simulated data. Strategies for distinguishing between a coherent signal and other artefacts are discussed in Sec. VC 1.

The computation of the marginal likelihood Eq. (4) and Bayes factor, Eq. (3) requires the integration of the likelihood function, $p(\vec{d}|\vec{\theta}, \mathcal{H}, I)$, where \mathcal{H} is either \mathcal{H}_N or \mathcal{H}_S , multiplied by the prior density function of the unknown parameters for the given hypothesis. We discuss the choice of prior in Sec. VA. Here we concentrate on the expression of the likelihood function. In the case of the hypothesis \mathcal{H}_N , the likelihood function is simply

$$p(d^{(D)}|\vec{\theta}, \mathcal{H}_N, I) \propto e^{-(d^{(D)}|d^{(D)})/2}, \quad (21)$$

see Eq. (17).

For the hypothesis \mathcal{H}_S , the likelihood of observing a data set $d^{(D)}$ at the output of the instrument D given the presence of a gravitational wave $h^{(D)}(\vec{\theta})$ characterized by the parameter vector $\vec{\theta}$ is

$$p(d^{(D)}|\vec{\theta}, \mathcal{H}_S, I) \propto e^{-(d^{(D)}-h^{(D)}|d^{(D)}-h^{(D)})/2}. \quad (22)$$

The constant of proportionality is equal in Eqs. (22) and (21), and cancels when the ratio of these quantities is taken. If we have a data set comprising observations from multiple interferometers, say $\vec{d} = \{\vec{d}^H, \vec{d}^L, \dots\}$, the Bayesian framework allows straightforward coherent analysis. To do this, we simply write the joint likelihood of the independent data sets in all the detectors

$$p(\vec{d}|\vec{\theta}, \mathcal{H}, I) = \prod_{(D)} p(\vec{d}^{(D)}|\vec{\theta}, \mathcal{H}, I), \quad (23)$$

where $p(\vec{d}^{(D)}|\vec{\theta}, \mathcal{H}, I)$ is either given by Eq. (17) or (22). In the appendix we provide explicit expressions for the likelihood function (23) in the case of discrete data used for the implementation in the software code.

III. THE NESTED SAMPLING ALGORITHM

The nested sampling algorithm is described by Skilling [21] as a reversal of the usual approach to Bayesian inference, in that it directly targets the computation of the evidence integral (4), producing samples from the posterior PDF, Eq. (5) of the model parameters $\vec{\theta}$ as a by-product. Although the original formulation was designed as a tool for Bayesian inference, it is actually a general method of numerical integration which could be applied to other continuous integrals. The basic algorithm, described in

[21], is therefore applicable to a wide range of problems, but in its generality it leaves considerable decisions to be made on the implementation, configuration and tuning to each specific application. In this section, we will review the core algorithm, Sec. III A, and the processing of the output of the algorithm to extract samples from the posterior PDF which can then be used for parameter estimation, Sec. III B. In Sec. IV we provide detailed information on our solution of the problem of sampling the limited prior distribution, Sec. IV A, and we will also examine the theoretical accuracy achievable with the algorithm. This result will then be compared to the practical accuracy achieved and its trade-off with computational cost in Sec. IV B.

More recently, MultiNest [60,61], based on the technique of nested sampling, has been applied to data sets primarily in the context of cosmology [62–64] and particle physics [65–70]. More recently it has been used on selected mock data sets of the Laser Interferometer Space Antenna to search for and estimate the parameters of massive black hole binary systems characterized by high ~ 1000 signal-to-noise ratio [71]. Our implementation is fundamentally different from MultiNest in the way in which the structure of the likelihood function is explored, and, in particular, it replaces the clustering algorithm or ellipsoidal rejection schemes with (nontrivial) MCMC explorations of the prior range, that are specifically tailored to the observation of gravitational waves with a network of ground-based instruments at moderate-to-low ($\lesssim 20$) signal-to-noise ratio. Moreover, given the large amount of data and the need to maximize computational efficiency while retaining the accuracy of the evaluation of the relevant integrals, our study concentrates on quantifying the errors in the evaluation of the key quantities and relating them to the computational costs.

A. Computing the evidence integral

The evidence $Z = P(\vec{d}|\mathcal{H}, I)$, given in Eq. (4) is found by integrating the product of the prior distribution with the likelihood function, in other words, Z is the expectation of the likelihood with respect to the prior. Using the product rule we can easily see the relationship between the prior, likelihood, posterior PDFs and the evidence,

$$p(\vec{\theta}|\mathcal{H}, I) \times p(\vec{d}|\vec{\theta}, \mathcal{H}, I) = Z \times p(\vec{\theta}|\vec{d}, \mathcal{H}, I) \quad (24)$$

$$\text{Prior} \times \text{Likelihood} = \text{Evidence} \times \text{Posterior.}$$

As the prior and posterior are by definition normalized, the magnitude of the evidence is governed by the likelihood function, which provides a measure of how well the data fits the hypothesis \mathcal{H} . In order to evaluate Z one must sum the product on the left side of Eq. (24) at each point $\vec{\theta} \in \Theta$ in parameter space. In our case the parameter space is a continuous manifold, and the likelihood function is a smooth function on the manifold Θ which we integrate.

When this integral is not solvable using analytic methods, we must approximate it by using a subset of points on Θ , for instance by placing a lattice on parameter space. Once such an approximation to a finite number of points is made, the result becomes subject to an integration error which is dependent on the precise means of integration.

Instead of a regular lattice of points, consider a stochastic sampling of the prior distribution to generate a basket of N samples—in the nested sampling jargon called *live points*—which we will denote $\vec{\theta}_i$, with $i = 1 \dots N$. The evidence integral [Eq. (4)] could then be expressed as

$$\begin{aligned} Z &= \int_{\Theta} p(\vec{\theta}|\mathcal{H}, I) p(\vec{d}|\mathcal{H}, \vec{\theta}, I) d\vec{\theta}, \\ &\approx \sum_{i=1}^N p(\vec{d}|\vec{\theta}_i, \mathcal{H}, I) w_i, \quad \approx \sum_{i=1}^N L_i w_i, \end{aligned} \quad (25)$$

where the “weight”

$$w_i = p(\vec{\theta}_i|\mathcal{H}, I) d\vec{\theta} \quad (26)$$

is the fraction of the prior distribution represented by the i th sample, and $L_i \equiv p(\vec{d}|\mathcal{H}, \vec{\theta}_i, I)$ is its likelihood. In the presence of a signal, the evidence integral is typically dominated by a small region of the prior where the likelihood is high, concentrated in a fraction e^{-H} of parameter space. H is called the *information* in the data, subject to the particular model and parametrization used, and is measured in nats (using base 2 instead of base e would give information measured in bits, where $1 \text{ nat} = \log_2 e \text{ bits} \approx 1.44 \text{ bits}$). H is defined as

$$H = \int p(\vec{\theta}|\vec{d}, \mathcal{H}, I) \log\left(\frac{p(\vec{\theta}|\vec{d}, \mathcal{H}, I) d\vec{\theta}}{dX}\right) d\vec{\theta}, \quad (27)$$

and will be used in Sec. IV B to quantify the accuracy [21]; X is the prior mass, and is defined in Eq. (28) below. If it is not known in advance where in parameter space the posterior is concentrated, approximately e^H points would be needed to avoid the possibility of missing the maximum of the likelihood function using a regular grid. In the case of a compact binary in-spiral signal as observed in a network of ground-based interferometers and using the parametrization given in Sec. II C, $H \gtrsim 20$ nats. If a regular grid of points was used (assuming a uniform prior), finding the weights associated with each point $w_i = 1/N$ would be simple, but the number of samples needed, N becomes prohibitively large. The key concept on which the nested sampling algorithm rests is the means to calculate the w_i of stochastically sampled points. By evolving the collection of N points to higher likelihood areas of parameter space, the algorithm simultaneously searches for the peaks of the distribution and accumulates the evidence integral as it progresses.

In order to find the weights associated with each point $\vec{\theta}_i$, it is useful to think of each point as lying on a (not necessarily closed) contour surface of equal likelihood in

the parameter space. The *prior mass*—that is the fraction of the total prior volume—enclosed by the i th contour surface is denoted X_i , with the lowest likelihood contour line enclosing the largest volume and the maximum likelihood point enclosing the smallest. With this definition, $X_0 = 1$. We can then think of a mapping between the contour lines in physical parameter space and the fractions of the prior X_i , where the likelihood $L(X)$ increases toward smaller values of X , as shown in Fig. 1, and $\Delta X_i = X_{i+1} - X_i$. The evidence, Eq. (4) or (25) can then be expressed as the one-dimensional integral

$$Z = \int L(X) dX \approx \sum_i L(X_i) \Delta X_i. \quad (28)$$

As the inverse mapping $\vec{\theta}(X)$ is not known, the analytical integral cannot be performed. However, as we know that the prior distribution is normalized to unity, the unknown prior mass enclosed by the outermost contour through X_1 has a probability distribution $P(X_1)$ which is equal to the distribution of a new variable $t_1 \in [0, 1]$, the maximum of N random numbers drawn from the uniform distribution $U(0, 1)$. If we then replace the first point with a new point sampled from the prior distribution limited to the volume lying at higher likelihood than L_1 , $X(L > L_1)$, we can repeat the process so that $X_2 = t_2 X_1$ and $X_i = t_i X_{i-1}$, where by definition $t_i \equiv X_i / X_{i-1}$ is the shrinkage ratio. The probability of t_i is $P(t_i) = N t_i^{N-1}$, where t_i is the largest of N random numbers drawn from $U(0, 1)$. The volume enclosed at each iteration therefore shrinks geometrically, ensuring the speedy convergence of the integral. The mean decrease in the volume at each iteration is

$$E[\log t] = \int_0^1 \log(t) p(t) dt = -N^{-1}, \quad (29)$$

and an estimate of the statistical variance introduced by this process is

$$\int_0^1 (\log t - E[\log t])^2 p(t) dt = N^{-2}. \quad (30)$$

The distribution of t can also be sampled by generating the

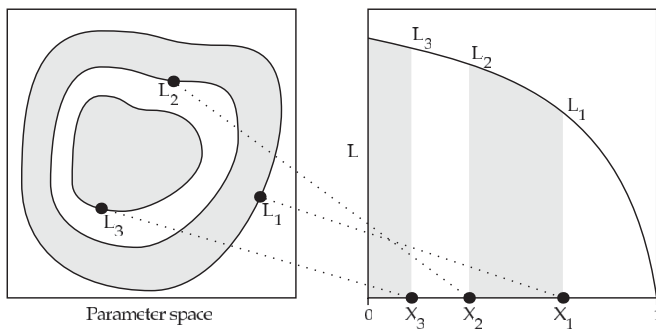


FIG. 1. Each sample in the basket of live points can be thought of as lying on a contour line of equal likelihood value. Figure reproduced from [21].

N uniform random numbers and creating many realizations of the t s for each iteration in the algorithm. In our testing with this procedure, it was found that for a reasonable number of realizations, the estimated mean and variance were very close to the expected figures. Using the approximation of the mean, we can therefore write the fractional prior volumes

$$\log X_i \approx -(i \pm \sqrt{i})/N, \quad (31)$$

and we use this approximation in the implementation, where we work with logarithmic quantities to overcome the huge range of the variables.

As we now have an approximation to the proportion of the prior mass remaining after the i th iteration, we can assign a weight to each sample as $w_i = X_i - X_{i-1}$.

As all structure of the likelihood function in the parameter space Θ is eliminated by the mapping to X , the nested sampling algorithm is in principle robust against multimodal distributions, degeneracies and problems arising from the high dimensionality of the parameter space. However, this relies on being able to sample the likelihood-limited prior distribution effectively. This difficult problem can be solved in a variety of ways, but the approach which we have found effective is described in Sec. IV A.

Note that if the bulk of the posterior is concentrated in a region of size e^{-H} of the prior, it will take approximately $NH \pm \sqrt{NH}$ iterations of the geometric shrinkage to reach the zone of high likelihood. This tells us that the algorithm will take longer to compute the integral if there is a larger amount of information in the data. This means the run time of the implementation is essentially dependent on the signal-to-noise ratio of any found signals, as well as on the number of live points N used.

Finally, we need to specify a termination condition, upon which we decide that the integral is finished. We could set a hard number for this, or a certain fraction of the prior, but the total number of points needed varies a great deal, particularly with the signal-to-noise ratio of the signal, if any. In our implementation, we keep track of the maximum likelihood point so far discovered. The algorithm will keep running until the total evidence that would be left if all the remaining points lay at the maximum likelihood so far discovered becomes less than a certain fraction of the total evidence so far accumulated. Based upon experience, we have found that continuing while $L_{\max} w_i > Z_i e^{-5}$ gives consistent results.

To summarize, the algorithm can be described in pseudocode (where $\vec{\theta} \sim p(\vec{\theta} | \mathcal{H}, I)$ means $\vec{\theta}$ is drawn from the distribution $p(\vec{\theta} | \mathcal{H}, I)$) as:

- (1) Draw N points $\vec{\theta}_a$, $a \in 1 \dots N$ from prior $p(\vec{\theta})$, and calculate their L_a 's.
- (2) Set $Z_0 = 0$, $i = 0$, $\log w_0 = 0$
- (3) While $L_{\max} w_i > Z_i e^{-5}$
 - (a) $i = i + 1$

- (b) $L_{\min} = \min(\{L_a\})$
- (c) $\log w_i = \log w_{i-1} - N^{-1}$
- (d) $Z_i = Z_{i-1} + L_{\min} w_i$
- (e) Replace $\vec{\theta}_{\min}$ with $\vec{\theta} \sim p(\vec{\theta}|\mathcal{H}, I): L(\vec{\theta}) > L_{\min}$
- (f) *Add the remaining points:* For all $a \in 1 \dots N$, $Z_i = Z_i + L(\vec{\theta}_a) w_i$

With the algorithm as it is outlined above, the crucial idea is the sorting of the likelihoods so that progressively smaller contour-lines can be assigned to each of them, and the integral (25) builds up. This leads us to a natural means of *parallelizing the algorithm*, so that we may take advantage of multiple processors or compute nodes on a cluster. If the algorithm is run in parallel with identical data and parameters, but a different random seed (this requires no internode communication), the sets of samples generated will differ. If we save each sample and its likelihood value, we can then collate the results of multiple runs, and sort the resulting samples by their likelihood values. So long as the number of parallel runs N_{runs} remains constant as the integration progresses, each subsequent sample from the limited prior distribution can be treated as being part of a collection of $N_T = \sum_{k=1}^{N_{\text{runs}}} N_k$ samples—where each parallel run has N_k live points—as we no longer know which sample belongs to which run. This then allows us to reapply the nested sampling algorithm as described above, but with a lower weight for each sample, substituting for N the number N_T .

After applying this procedure, a more accurate estimate of the evidence integral can be obtained, using our greater number of total live points. This procedure also increases the accuracy of the evaluation of the posterior PDFs that we discuss in the next section. It was found that this procedure allows the accuracy to scale with the total number of parallel live points as shown in Fig. 4, provided that each run has a sufficiently large number of live points to avoid under-sampling of the parameter space. The issue of increasing accuracy at the expense of additional run time is discussed further in Sec. IV B.

B. Extracting the posterior PDF

As the nested sampling algorithm proceeds, the list of points used in approximating the integral is stored, along with the likelihood values of each sample, the corresponding value of the parameter vector, and $\log X_i \approx i/N$. These samples are drawn from the prior distribution, limited by a likelihood contour to a fraction X_i of the full prior, meaning that the density of the samples is boosted within the contour by the probability that is excluded, as it is zero outside the contour. We can therefore write in shorthand the probability density of the i th point from the nested sampling output as

$$p(\vec{\theta}_i|\text{NS}) = \frac{p(\vec{\theta}_i|\mathcal{H}, I)}{X_i}, \quad (32)$$

whereas samples from the posterior PDF, Eq. (6) have probability density

$$p(\vec{\theta}_i|\vec{d}, \mathcal{H}, I) \propto p(\vec{\theta}_i|\mathcal{H}, I)p(\vec{d}|\vec{\theta}_i, \mathcal{H}, I). \quad (33)$$

Since the nested sampling points are independent samples, they can be reused to generate samples from the posterior PDF by resampling them. Substituting Eq. (32) into Eq. (33), it is easy to see that the probabilities are related by,

$$p(\vec{\theta}_i|\vec{d}, \mathcal{H}, I) \propto p(\vec{\theta}_i|\text{NS})p(\vec{d}|\vec{\theta}_i, \mathcal{H}, I)X_i, \quad (34)$$

and so the resampling weight of each one is $\propto p(\vec{d}|\vec{\theta}_i, \mathcal{H}, I)X_i$. As a consequence, the joint posterior PDF can be easily calculated by post-processing the output of the nested sampling algorithm (at negligible computational cost). Marginalized posterior PDFs, Eq. (6) can then be obtained as in the case of MCMC methods, by histogramming the samples. In this way we can easily perform both evidence integrals and estimation of the posterior PDF, making both model selection and parameter estimation possible. It is important to note that the method of extracting posterior samples using nested sampling is different to that in standard MCMC algorithms, as the algorithm is designed to move the ensemble of points uphill from a sampling of the entire prior toward the highest likelihood point, whereas MCMC can also move downhill (with probability <1) and requires sufficient burn-in time to fully explore the full range of the prior. If the location of the true maximum is not known, nested sampling offers the ability to home in on the true location with its geometric shrinkage of the sampling volume, making it an excellent tool for searching the parameter space for maxima.

On the other hand, the number of posterior samples generated by the nested sampling algorithm is limited by the number of live points used, with more live points corresponding to more samples within the uppermost contour lines. In order to get the posterior sampling desired, it might be necessary to increase N , or run parallel computations, which has the effect of causing the algorithm to converge more slowly (but also more accurately, see Sec. IV B).

IV. IMPLEMENTATION DETAILS

In the previous section we have described the conceptual approach to the computation of the evidence integral using a nested sampling technique; furthermore, at the end of Section III A, we have provided a pseudocode with the key steps of the algorithm. One of the key challenges in the efficient implementation of the algorithm is to replace at each iteration the active point characterized by the minimum value of the likelihood function L_{\min} , drawing a sample from the prior distribution limited to the volume that satisfy the condition $L > L_{\min}$. We do so by means of a Metropolis-Hastings Markov chain Monte Carlo with M

steps that proceeds as follows. We randomly select one of the N live points, corresponding to say $\vec{\theta}$, that we assume as the starting point of the Markov chain, the “current state.” We then propose a new state $\vec{\theta}'$ drawn from a proposal distribution (also called transition kernel) $q(\vec{\theta}, \vec{\theta}')$, i.e. the probability of $\vec{\theta}'$ given $\vec{\theta}$. The new state is accepted with probability

$$\alpha_H(\vec{\theta}, \vec{\theta}') = \begin{cases} \min[1, \frac{p(\vec{\theta}')q(\vec{\theta}, \vec{\theta}')}{p(\vec{\theta})q(\vec{\theta}', \vec{\theta})}] & L(\vec{\theta}') > L_{\min} \\ 0 & L(\vec{\theta}') \leq L_{\min} \end{cases} \quad (35)$$

and equivalently the chain remains at $\vec{\theta}$ with probability $1 - \alpha_H(\vec{\theta}, \vec{\theta}')$. We continue to evolve the chain to accumulate M states, and the last one is set to be the new live point that replace that characterized by $L = L_{\min}$. If no points have been accepted during the M proposals, then the chain will have remained at the preexisting point, and so we must rerun the chain using a different live point as an initial state.

In Sec. IVA we describe the details of the exploration of the limited prior, that is the choice of $q(\vec{\theta}, \vec{\theta}')$; in Sec. IV B we quantify how the errors in the evidence evaluation scale with the number of live points and MCMC elements, N and M , respectively, and how they are related to the CPU processing time.

A. Sampling the limited prior

In order to produce a new live point for each iteration of the nested sampling algorithm, it is necessary to draw a sample from the prior distribution, limited to volumes with likelihood greater than L_{\min} . This distribution changes from the entire prior distribution at the zeroth iteration, to a tiny fraction, typically $< 10^{-10}$, of the parameter space when the posterior mode has been located, see e.g. Fig. 2. In between these extremes, as L_{\min} increases, it will cause “islands” of probability to separate from each other and disappear, as if being submerged by a rising tide. There may also be multiple maxima of similar likelihood values, and these modes are generally curved or “banana-shaped” in the multidimensional volume, an example of which is shown in Fig. 3.

Maintaining an accurate and efficient sampling of all these islands is the biggest challenge in implementing nested sampling, as it is in other Monte Carlo methods such as MCMC’s. The character of these islands and their shapes will vary from problem to problem, but here we have attempted to proceed in a general way, through the use of a semiadaptive Markov-chain Monte Carlo algorithm to sample the prior, where the number of iterations of the chain M can be specified. This approach was augmented with custom proposals, based on some simple intuition on the structure of the likelihood function, which were found to improve the efficiency of sampling, or

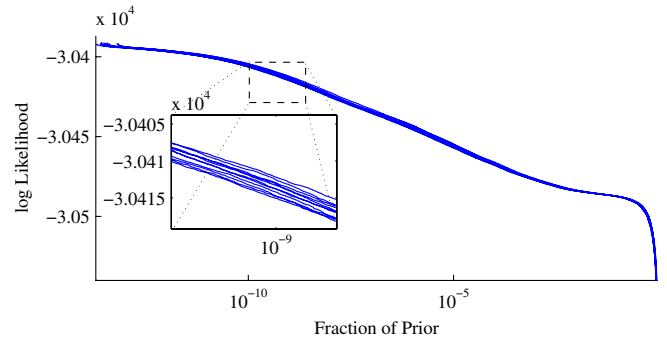


FIG. 2 (color online). An illustrative plot of the log likelihood against the fraction of the prior X_i , generated as the algorithm progresses, to be compared with Fig. 1. Reading this from right to left, as the fraction of the prior enclosed by the contour line decreases, the likelihood increases as the algorithm proceeds. The individual details of the distribution are smoothed out by the projection onto the X parameter. The bulk of the probability from the signal occurs at around $10^{-11} = e^{-25.3}$, in good agreement with the estimated information content of ~ 25.7 nats. The inset shows the slightly different results gathered by running the algorithm 10 times with a different random seed, where the precise samples used to integrate are different.

alternatively the speed of chain mixing, and will be described below in section IVA 1.

By using an MCMC sampling of the prior, we have to choose a proposal distribution $q(\vec{\theta}, \vec{\theta}')$ which will give a decent acceptance ratio at all stages of the nested sampling. As the scale of the problem varies by ten or more orders of magnitude, this is impossible to achieve with a static




FIG. 3 (color online). The structure of the likelihood function can clearly be seen in the samples from the posterior PDF. These figures show the distributions obtained when using the Livingston-Virgo (green \times), Hanford-Livingston (magenta \times), Hanford-Virgo (blue \times) and Hanford-Livingston-Virgo (red \times) networks of detectors to analyze the same stretch of simulated data. The true position of the injection is marked \star . With two detectors, the distribution lies on the circle produced by keeping the time of arrival constant in both instruments. With a network of three detectors, the timing can be kept constant by reflecting the position of the source across the plane of the three detectors. This causes two local maxima at the two places where the three circles intersect, with one of these being the true location.

choice of proposals. However, we have additional information available, in the form of the location of the N live points, which can help us select proposal distributions dynamically.

As the collection of live points shrinks at each iteration, we can obtain an estimate of the size and orientation of the area we need to sample by computing the covariance matrix C of the collection of live points,

$$C_{ij} = \langle (\theta_i - \langle \theta_i \rangle)(\theta_j - \langle \theta_j \rangle) \rangle, \quad (36)$$

where the indices i, j denote the dimensions (9 in this specific case) of the parameter space, and $\langle \cdot \rangle$ should be interpreted as the sample mean over the active live points. In the case of the cyclical angular parameters ϕ_0 and α , we need to take into account the wrapping of the boundary, and set the covariance between these and the other parameters to zero. The variances $C_{\phi_0\phi_0}$ and $C_{\alpha\alpha}$ are then computed using the circular mean and circular difference in place of the mean and difference in the above expression [72]. The matrix C_{ij} is recomputed at regular intervals (in practice, once every $N/10$ iterations) as the shrinking proceeds, scaled by a factor of 0.1, and used in the sampling of a multivariate student distribution, with mean centered on the previous point, and 2 degrees of freedom, when evolving the prior samples. In order to do this, a vector \vec{v} is drawn from the multivariate normal distribution with mean zero and covariance matrix given by the expression (36), then multiplied by a factor $\sqrt{2/x}$, where x is drawn from a χ^2 distribution with 2 degrees of freedom to yield a multivariate student deviate. The vector \vec{v} is then added to $\vec{\theta}$ to yield the new proposed point $\vec{\theta}' = \vec{\theta} + \vec{v}$.

The advantages of this somewhat *ad hoc* method are that it is fast to compute, and is applicable whenever the number of live points is greater than the dimensionality of the problem to avoid a singular matrix (which in practice is always the case). As a consequence, it can be used without modification if one wants to examine a model with additional parameters, or with some of the parameters constrained. It also adapts to the shrinking volume of the limited prior distribution as the algorithm progresses. This type of proposal is used for the majority of the jumps used in the sampler. However, there are certain types of proposals that we use specifically for short-lived bursts (such as gravitational waves generated by coalescing binaries) specifically discussed in this paper that are designed to move along the degeneracies of the distribution.

1. Custom proposals for sky position

The *coherent* analysis of data from interferometers at different locations on the Earth, as it is the case for the LIGO-Virgo network, offers a number of advantages that we will discuss in more detail in the next section. Particularly relevant for the implementation details discussed here is the fact that multiple instruments allow us to reconstruct partial or full information (depending on the

number of instruments in the network and their location/orientation) about the source position in the sky. Such information is encoded in the structure of the likelihood function, and its functional dependency on (t_0, α, δ) . However, distinctive features in the distribution also provide a challenge in exploring them in a accurate, yet efficient way. The relative timing offset between the arrival time of a gravitational wave at multiple sites encodes the majority of the information about location in the sky (of course additional information is contained in the antenna beam patterns $F_{+, \times}$). More specifically, given two detectors there exists a locus of points in the sky that traces a circle about the baseline between the two sites that yield the same time-delay at the two detectors. If one adds a third site, the three circles intersect in two points: one corresponds to the actual source position in the sky, whereas the other represents a “mirror image” that is located on the opposite side across the plane that passes through the three sites. The likelihood function therefore exhibits certain degeneracies or near-degeneracies in the (t_0, α, δ) subspace of parameter space. This results in distributions which trace out arcs of a circle on the sky, with modulation in t_0 , as shown in Fig. 3.

In the exploration of the parameter space one can therefore take advantage of the known geometrical symmetries of the problem, by suitably choosing the proposals that control the geometrical parameters. It is therefore much more efficient to move “in circles” in the sky—subject to the constraints that we have discussed above—rather than to make proposals based on the distribution discussed in Sec. IV A. We have indeed observed that custom-made proposals dramatically improve the performance of the algorithm, both in term of efficiency and accuracy. As the relative timing offsets between detectors provide the majority of the information about location on the sky, we therefore propose a fraction of new states of the MCMC chain by keeping the time of arrival of the signal constant in each detector. In the case of a network of two detectors, this constrains the jump to a ring, centered on the vector between the two detectors, which we sample by applying a rotation matrix with uniform random angle between 0 and 2π to the position vector of the source (and accounting for the relative rotation between the earth-fixed detector coordinates and the sky-fixed source coordinates). In the case of three detectors, in order to keep the same time of arrival in all detectors, the source must be reflected in their plane. As a consequence, if \hat{x} is the current and \hat{x}' the proposed Cartesian unit vector to the source, and the detectors are located at the points $\vec{x}_H, \vec{x}_L, \vec{x}_V$ in the same coordinate system, with a normal to their plane $\hat{n} = [(\vec{x}_L - \vec{x}_H) \times (\vec{x}_V - \vec{x}_H)] / |[(\vec{x}_L - \vec{x}_H) \times (\vec{x}_V - \vec{x}_H)]|$, the jump is therefore

$$\hat{x}' = \hat{x} - 2\hat{n}|\hat{n} \cdot (\hat{x} - \hat{x}_H)|. \quad (37)$$

In both cases, as the detectors are offset from the geocenter,

and it is the time of arrival at the geocenter t_0 which is used as a parameter, there is also an adjustment to be made to this parameter when moving to a new sky location, $t'_0 = t_0 + \vec{x}_H \cdot (\hat{x} - \hat{x}')$. Here, \vec{x}_H , etc. are measured in seconds.

The custom proposals discussed here apply to any likelihood exploration that involves short-lived (with respect to the time-scale of the Earth's rotation and orbit) bursts of gravitational waves, as it is simply connected to the relative time delays observed between different detectors.

2. Differential evolution

When the limited prior distribution splits into multiple isolated islands of probability, see e.g. Fig. 3, the method of using the covariance matrix of the live points as a proposal distribution leads to less efficient sampling. In order to combat this effect, it is necessary to propose jumps which have a length scale characteristic of moving between, or within the multiple modes. One possible technique is to analyze the current live points as belonging to a number of clusters, then propose new states from within these clusters, which is the approach adopted by the MultiNest algorithm [60,61].

In our implementation, we have introduced a new type of MCMC proposal which attempts to capture some of the structure of multimodal distributions, based on a simple iteration of proposals inspired by differential evolution MCMC algorithms, see e.g. Ref. [73]. From the whole set of live points $\vec{\theta}_1, \dots, \vec{\theta}_N$, we select a random point, say $\vec{\theta}_a$ that we want to evolve to point $\vec{\theta}'_a$; we then select two random existing points, say $\vec{\theta}_b$ and $\vec{\theta}_c$, such that $a \neq b \neq c$. The proposed new state is then given by

$$\vec{\theta}'_a = \vec{\theta}_a + (\vec{\theta}_c - \vec{\theta}_b), \quad (38)$$

which is accepted with the usual Hastings ratio, Eq. (35). As the probability of drawing (b, c) for the random move is equal to that of drawing (c, b) , the move is reversible, and therefore upholds the principle of detailed balance.

When used for a fraction of the jumps (10% in our case, chosen through trial and error to explore adjacent modes while still maintaining good diffusion of points through the standard jumps), this type of move allows proposals to be made at all the characteristic scales between the different modes in a multimodal distribution (as well as at the scale of the width of each mode), and so increases the efficiency when such a distribution is encountered. Note that this type of proposal needs no scale to be set by the user, and is independent of the parameters used.

B. Accuracy: Quantifying errors

Because of the probabilistic nature of the algorithm, when computing the Bayes factor there is an associated uncertainty with the result obtained by applying the nested sampling algorithm. As the evidence integral is written

$$Z = \sum_{i=1}^{N_{\text{tot}}} L_i w_i, \quad (39)$$

there is a Poissonian uncertainty arising from the variable number of iterations needed to find the region of high posterior probability, $N_{\text{tot}} = NH \pm \sqrt{NH}$, which gives rise to an uncertainty in $\log Z$ of $\pm \sqrt{H/N}$. In [21], Skilling suggests that this error will dominate other sources of uncertainty, but makes the assumption that the sampling of the limited prior distribution is done perfectly. If the sampling is not done perfectly, for example, if a small isolated mode is not properly sampled as there are no live points in its neighborhood, an additional error will be introduced into the quantities w_i . Rather than attempting to derive this quantity, we have performed Monte Carlo simulations on many sets of identical data and signals, while changing the parameters of the algorithm N and M , and examined how the distribution of estimated Z values changes. By doing this, we can also explore to which value N and M should be set to attain a given level of accuracy in the evidence computation.

Figure 4 shows the theoretical $\sqrt{H/N}$ level of uncertainty, along with the actual standard deviation of the estimates of Z , over 50 trials, for a range of N and M . The actual distribution of recovered Z scales as $\approx 1/\sqrt{N}$ as theoretically predicted. It is however noticeably larger than that predicted by the $\sqrt{H/N}$ error alone, and is dependent on both the number of live points and the number of MCMC samples used. We can see that when N and M

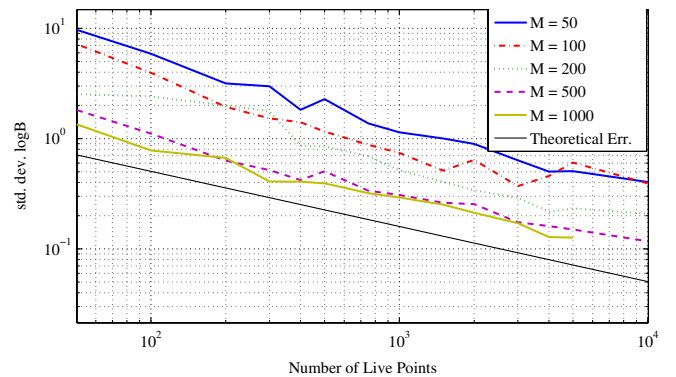


FIG. 4 (color online). The statistical error (standard deviation) computed over 50 trials with different random seeds (identical signal and noise) plotted against the number of live points used, using $M = 50, 50, 100, 200, 500, 1000$, and the theoretical prediction of the statistical error based on Skilling's estimate. These results show that the empirical error is greater than the theoretical one, indicating an additional source of uncertainty. This is greatly reduced with the number of MCMC samples used, with chains of 1000 samples (solid yellow line) approaching the theoretical limit (solid black line). This suggests that the extra uncertainty is produced by correlation between samples, caused by less than perfect mixing in the MCMC sampling of the limited prior.

increase, the variance decreases, suggesting that there is an additional source of error related to the sampling of the limited prior distribution. This is not entirely surprising, as the limited prior distribution consists of a number of isolated islands, between which it is difficult to move with an MCMC sampler. If there is residual correlation in the MCMC chain used to sample the parameter space, this will introduce an overweighting of the area of over-density of samples, which, depending on whether this is a region with higher or lower likelihood will cause an increase or decrease in the Z integral. As the correlation between start and end decreases with the length of the Markov chain, the added error decreases with increasing M . It is in fact clear that, at a given N , by increasing M the result approaches the theoretical error.

By tuning N and M appropriately, we can therefore attain any desired level of accuracy, in principle, but at the expense of increasing the computational burden, as the number of likelihood evaluations is approximately proportional to the product NM . This can be seen clearly in Fig. 5, where we show the number of likelihood evaluations required to achieve the accuracies presented in Fig. 4. The actual processing time then depends on the time taken to evaluate a single value of the likelihood function, which varies with the mass of the signal, and the length of data and sampling rate used. As an example, to compute the result with $N = 500$, $M = 200$, which gives an accuracy on $\log B$ of ± 0.8 , $\sim 3 \times 10^6$ likelihoods were evaluated. In this case the algorithm took approximately 1 h 20 mins to complete on a 2.4 GHz Intel Xeon processor. Assuming that the overhead beyond the likelihood calculation is minimal, this gives an approximate time of 1.5 ms per likelihood. The system used in these tests was a $30M_\odot - 1.4M_\odot$ binary system injected into 3 data streams, but the full parameter space (see Sec. VA) was explored so that templates across the whole low-mass range were generated. It should be noted that this number is mostly dependent on the time taken to generate the waveform, which is

lowest for the stationary phase approximation templates used here, but is higher for time domain templates and those requiring the numerical solution of differential equations to produce the waveform.

V. RESULTS

In this section we present a range of tests to demonstrate the effectiveness of a Bayesian approach in identifying gravitational wave signals in the data from a network of interferometers and estimating the associated parameters. We first investigate the detection efficiency of our algorithm, using $\log B$ as a “detection statistics”; we then introduce and characterize a new test to discriminate a coherent gravitational wave at the output of multiple instruments from the presence of incoherent instrumental artefacts. We conclude by showing the impact of the number of instruments in the recovery of the source parameters.

For the tests presented in this section, we have chosen four systems of different combination of masses for the values used in the injections, shown in Table I. The first three systems lie near the corners of the irregular prior, shown in Fig. 6 and discussed in Sec. VA, while the fourth lies somewhat towards the middle. Using these systems, the performance of the algorithm across different signals lying in different parts of parameter space is assessed.

Unless stated otherwise, we run all tests using three simulated interferometers operating as a network. These are located at the sites of LIGO Hanford, LIGO Livingston and Virgo (Cascina), and have simulated noise power

TABLE I. Mass parameters of the injections used in testing.

Name	$m_1 (M_\odot)$	$m_2 (M_\odot)$	\mathcal{M}	η
System 1	1.4	1.4	1.219	0.25
System 2	30	1.4	4.727	0.043
System 3	15	15	13.06	0.25
System 4	25	5	9.18	0.139




FIG. 5 (color online). The mean number of likelihood evaluations performed in each of the trials shown in Fig. 4. The number of likelihoods scales linearly with both the number of live points and the number of MCMC iterations used.




FIG. 6 (color online). $p(\mathcal{M}, \eta | H_s, I)$, the prior probability distribution for the parameters \mathcal{M} and η . The white area is excluded by the boundary conditions imposed on the component masses (see Sec. VA), and the distribution given by Eq. (43).

spectral densities chosen to correspond to the design sensitivities of each of these real instruments, as modeled by the appropriate LAL functions. The Gaussian and stationary colored noise is then generated in the frequency domain, with the appropriate noise spectrum. A low-frequency cut-off of 50 Hz and Nyquist frequency of 1024 Hz are used throughout the analysis.

A. Priors

The choice of prior distribution $p(\vec{\theta}|\mathcal{H}, I)$ is an important factor in Bayesian inference, and will affect the Bayes factor as it is included in the evidence integral, Eq. (4). The prior effectively determines exactly which model is being used, by incorporating the ranges of the model parameters, and probability distribution on those parameters before the data are analyzed. In the case of our implementation, the prior must be sampled using an MCMC technique, which will be more efficient when there is minimal structure in the chosen parametrization of the signal, in the sense that fewer iterations will be required to adequately sample it.

For many of the model parameters, the choice of prior distribution is obvious: we use an isotropic distribution for the source sky location (α, δ) and the direction of the orbital angular momentum (ι, ψ) over the full range on the angular parameters of the model, reflecting total ignorance and no reasons to prefer a particular geometry of a binary. We also choose a flat distribution on ϕ_0 and t_0 over the range $0 \leq \phi_0 \leq 2\pi$ and a time interval of 100 msec, respectively. For distance and masses, we use a uniform prior on $\log D_L$ in the range $D_L \in [1, 100]$ Mpc, a uniform prior in η and a prior on chirp mass of the form $p(\mathcal{M}|I) \propto \mathcal{M}^{-11/6}$. As we desire to test our approach on the mass region covered by the LIGO-Virgo low-mass searches for in-spiral signals [12,13], limits were imposed directly on the component masses, such that $m_1, m_2 \leq 35M_\odot$, where by our convention $m_1 \geq m_2$. We also place a lower limit on the mass ratio, such that of $\eta > 0.01$, and on the chirp mass, $\mathcal{M} > 0.87$, to ensure that the waveforms generated are nonzero and of a length suitable for our analysis. These constraints result in a convoluted shape for the allowed regions of parameter space in the (\mathcal{M}, η) plane, shown in Fig. 6. The specific choice of the distance and mass priors is determined by the need to ensure the accuracy of the integration, which we now discuss.

Within the core sampler, we have changed the variable used to $\log \mathcal{M}$, in order to reduce the range of the prior density, leading to better sampling of the space than using \mathcal{M} itself. The prior probability density function we use on $\log \mathcal{M}$ is based on an approximation to the Jeffreys prior, $p(\vec{\theta}|\mathcal{H}, I) \propto \sqrt{\det \Gamma(\vec{\theta})}$, where $\Gamma(\vec{\theta})$ is the Fisher information matrix, defined as

$$\Gamma_{ij}(\vec{\theta}) = \left(\frac{\partial_i \tilde{h}(\vec{\theta})}{\partial \theta_i} \left| \frac{\partial_j \tilde{h}(\vec{\theta})}{\partial \theta_j} \right. \right). \quad (40)$$

This type of prior is used when there is no information about the parameters of the signal at all, and should therefore be invariant under a change of coordinates [74,75]. For simplicity, in our initial implementation we ignore the correlations between the chirp mass and the other parameters and we just take the leading order Newtonian quadrupole approximation of the inspiral waveform $\tilde{h}(f)$ to compute the scaling of the prior. Under these assumptions we obtain

$$\frac{\partial \tilde{h}(f)}{\partial \log \mathcal{M}} = -\frac{5i}{4} (8\pi \mathcal{M} f)^{-5/3} \tilde{h}(f) \quad (41)$$

and the prior is therefore

$$p(\log \mathcal{M}|I) \propto \sqrt{\left(\frac{\partial \tilde{h}(\vec{\theta})}{\partial \log \mathcal{M}} \left| \frac{\partial \tilde{h}(\vec{\theta})}{\partial \log \mathcal{M}} \right. \right)}, \quad \propto \mathcal{M}^{-5/6}, \quad (42)$$

$$p(\mathcal{M}|I) \propto \mathcal{M}^{-11/6}. \quad (43)$$

The Jeffreys prior is based on the notion that one should assign equal probabilities to equal volume elements in the parameter space of the signal. Here the Fisher matrix is used as a metric, allowing us to calculate the volume element at each point. The Jeffreys prior encodes the fact that the volume element of a curved parameter space may vary with respect to the parametrization, and so the density of templates is greater at lower chirp masses. Failure to account for this will result in an under-sampling of certain regions of parameter space, which will increase the chances of failing to detect a signal there (or increase the number of live points needed for a given probability of detection). This effect is most significant in the \mathcal{M} parameter, which is why we have found it necessary to include it here, however a full calculation of the Fisher metric in the \mathcal{M}, η space would further improve the sampling, and is a goal for future development of this work.

Note that the use of this prior, as with the one for distance, ignores any available information about the mass distribution of neutron stars and black holes, and focusses simply on the detection of the signal with unknown parameters.

B. Detection efficiency

In order to test the detection efficiency of our implementation, we have chosen to treat the Bayes factor of the signal vs noise hypotheses

$$B_{S,N} = \frac{P(\vec{d}|\mathcal{H}_S, I)}{P(\vec{d}|\mathcal{H}_N, I)}, \quad (44)$$

see Eq. (3) and Sec. IID, as a detection statistic. By performing a large number of runs on a signal-free data set, we can find the distribution of $\log B_{S,N}$ in the absence of

a signal, and therefore choose a threshold value which will give a certain false alarm rate. This is the same approach that we have adopted in Refs. [25], however in this paper we consider coherent observations with multiple interferometers. Using this threshold, we can therefore decide whether or not the analysis of a data set which contains an actual injection yields a detection or not. To achieve this, we analyzed 5000 different realizations of Gaussian noise, and obtained the distribution of $\log B_{S,N}$. This is shown as a histogram in Fig. 7, where the vertical line represents the threshold $\log B_{S,N} = 2.786$ corresponding to a false alarm rate of 1% for any single trial.

Taking this threshold, we then classify each result as detected if $\log B > 2.786$, and otherwise not. The detection efficiency is then assessed by performing 50 injections of each test signal at varying signal-to-noise ratios (by changing the distance to the source) and determining the fraction which are detected vs those which are not. This allows us to build up a detection efficiency curve for each of the test systems, given the desired false alarm rate, which is shown in Fig. 8. From these results, we can see that the $30 - 1.4M_{\odot}$, $15 - 15M_{\odot}$, and $25 - 5M_{\odot}$ systems show consistency in their chance of being detected as a function of signal-to-noise ratio, following a typical sigmoid curve with a transition zone between signal-to-noise ratios 4 and 8 where there is an intermediate chance of detection, explained by the different noise realizations. Each of these curves crosses the 50% detection efficiency at approximately signal-to-noise ratio of 6.5, and approaches 100% detection efficiency with a signal-to-noise ratio above 8. In contrast, the algorithm performs slightly more poorly in the detection of the binary neutron star system with $1.4 - 1.4M_{\odot}$ component masses, with 50% detection efficiency at signal-to-noise ratio of 7.5 and a wider zone of transition. An examination of the raw Bayes factors output by the algorithm for this system indicated that the detected binary neutron star signals are allocated Bayes factors consistent with other signals of the same signal-to-noise




FIG. 7 (color online). The distribution of the log Bayes factor for 5000 runs on synthetic Gaussian noise. The vertical solid line represents the 1% false alarm rate threshold of $\log B_{S,N} = 2.786$.




FIG. 8 (color online). Detection efficiency curve for the nested sampling algorithm for the four test systems as a function of the coherent network signal-to-noise ratio, with a detection threshold of $\log B_{S,N} > 2.786$, corresponding to a 1% false alarm rate, see Fig. 7. Error bars indicate the 67% probability interval assuming a binomial distribution for the results of the 50 trials.

ratio, but that there is a larger fraction of sources which are not detected, producing a Bayes factor consistent with noise. This may be due to the sampler failing to identify the correct region of parameter space, as the parameter volume of signals at low mass is considerably smaller, and therefore has a lower probability of being found by a probabilistic algorithm.

It is suggested that improved performance could be obtained for these systems by incorporating the full metric into the calculation of density required in the \mathcal{M}, η subspace, which would then distribute the samples more appropriately. However, the results broadly show that the algorithm is capable of correctly analyzing and detecting signals at an SNR comparable to existing methods. Other recent work has shown that prior distributions may be found which balance the need for efficient detection with astrophysical prior information [76].

Using the distribution of Bayes factors produced by the noise-only runs, we can also examine the relationship between false alarm rate and detection efficiency. For a choice of signal-to-noise ratio, the threshold of detection is varied, causing a change in the number of detected signals but also the number of false alarms, which can be plotted in a receiver operations curve (ROC), as shown in Fig. 9. In this figure, we have used system 1, see Table I with component masses $30M_{\odot}$ and $1.4M_{\odot}$ to produce the ROC curves for SNRs below, in and above the transition region. These results were produced with 50 independent trials at each optimal SNR, and so the error bars shown represent the Poissonian error of $50^{-1/2}$.

C. Coherent analysis

Every gravitational wave search critically relies on multiple (at least two) instruments in order to make confident detections of astrophysical signals. Multiple interferometers are beneficial in two main ways. First, the signal in




FIG. 9 (color online). Receiver Operation Characteristics curve for the nested sampling algorithm, obtained with 50 trials at each signal-to-noise ratio for the $30M_{\odot} - 1M_{\odot}$ test binary.

each detector can be cross checked against those observed in the others, giving us an additional way of isolating a real gravitational wave which must yield consistent observations in all the instruments; in fact, the searches for short-lived gravitational waves employ one or more “detection confidence tests,” see e.g. [77] and references therein. Second, if the interferometers are not colocated and aligned (as is the case of the ground-based network currently in operation) then they will see a different projection of the strain tensor of the passing gravitational wave; in turn, this allows a better estimation of the parameters of the incoming signal, and can break degeneracies between parameters present with only one data set. In particular, simultaneous observations with three or more instruments allows us to determine the geometry of the source, such as its distance and location in the sky.

Using the mathematical and computational framework developed and described above, we can address both these points in a natural way. In this section we propose and demonstrate a new *coherence test* which can discriminate between coherent and incoherent signals, making optimal use of the data and we demonstrate the improvement in parameter estimation when using a detector network.

1. Coherence test

Having a network of detectors is essential to discriminate a real astrophysical signal from the background of spurious noise events which resemble gravitational wave signals, i.e. those with $B_{S,N} \gg 1$ in the language of model selection. We know that a real gravitational wave signal must be observed in all the detectors with compatible estimates for the physical parameters, and relative time delays that are consistent with the location of the instruments on Earth and a point on the celestial sphere. On the other hand, instrumental glitches will appear independently in each detector. To be more specific, the observed data in each detector must be consistent with the physical gravitational wave: the different characteristics of each detector, including their instantaneous noise levels and orientations mean that the signal-to-noise ratios will vary between them.

There will naturally be times when glitches occur simultaneously in multiple detectors. When this happens, we can use the network in a coherent manner to test whether the event is consistent with a coherent gravitational wave, or more likely to be a gravitational-wave-like glitch occurring independently in each detector. Translated into our framework of inference, we want to compare the two following hypotheses:

- (i) Coherent model, \mathcal{H}_{coh} : The data sets $\vec{d} = \{d^{(1)}, d^{(2)}, \dots, d^{(N_D)}\}$ from each of the N_D detectors contain a coherent gravitational wave signal described by the same polarization amplitudes $\tilde{h}_+(f; \vec{\theta})$ and $\tilde{h}_\times(f; \vec{\theta})$ with the same parameters $\vec{\theta}$. The posterior probability of the coherent hypothesis is as before

$$P(\mathcal{H}_{\text{coh}}|\vec{d}) = \frac{P(\mathcal{H}_{\text{coh}})}{P(\vec{d})} Z_{\text{coh}}, \quad (45)$$

where the evidence is

$$Z_{\text{coh}} = \int_{\Theta} p(\vec{\theta}|\mathcal{H}_{\text{coh}}) p(\vec{d}|\vec{\theta}, \mathcal{H}_{\text{coh}}) d\vec{\theta} \quad (46)$$

and the joint likelihood of the observation \vec{d} is $p(\vec{d}|\vec{\theta}, \mathcal{H}_{\text{coh}}) = \prod_i^{N_D} p(d^{(i)}|\vec{\theta}, \mathcal{H}_{\text{coh}})$, see Eq. (23).

- (ii) Incoherent model, \mathcal{H}_{inc} : The data set at each instrument, $d^{(1)}, d^{(2)}, \dots, d^{(N_D)}$ contains independent gravitational-wave-like glitches, characterized by the parameters $\vec{\theta}^{(1)}, \vec{\theta}^{(2)}, \dots, \vec{\theta}^{(N_D)}$, in general different for each detector. In this case, assuming that the data and signals at each instrument are independent, the marginal likelihood of the model factorizes into the marginal likelihoods of each signal in the relevant detector, and the posterior probability is

$$P(\mathcal{H}_{\text{inc}}|\vec{d}) = \frac{P(\mathcal{H}_{\text{inc}})}{P(\vec{d})} Z_{\text{inc}} = \frac{P(\mathcal{H}_{\text{inc}})}{P(\vec{d})} \prod_{i=1}^{N_D} Z^{(i)} \quad (47)$$

where the evidence for the signal in each detector is

$$Z^{(i)} = \int_{\Theta^{(i)}} p(\vec{\theta}^{(i)}|\mathcal{H}_{\text{inc}}) p(d^{(i)}|\vec{\theta}^{(i)}, \mathcal{H}_{\text{inc}}) d\vec{\theta}^{(i)}; \quad (48)$$

in the equation above $p(d^{(i)}|\vec{\theta}^{(i)}, \mathcal{H}_{\text{inc}})$ is the likelihood of the data set $d^{(i)}$ at the i th instrument output, characterized by the parameter vector $\vec{\theta}^{(i)}$ defined over the space $\Theta^{(i)}$.

In order to distinguish between these possibilities, we need to compute the odds ratio, Eq. (2), between the coherent and incoherent model

$$O_{\text{coh,inc}} = \frac{P(\mathcal{H}_{\text{coh}})}{P(\mathcal{H}_{\text{inc}})} B_{\text{coh,inc}} \quad (49)$$

where the Bayes factor is

$$B_{\text{coh,inc}} = \frac{Z_{\text{coh}}}{Z_{\text{inc}}} = \frac{Z_{\text{coh}}}{\prod_i^N Z^{(i)}}, \quad (50)$$

see Eqs. (45) and (47). Essentially the test computes the difference between the integral of the product and the product of the integrals for each data set. The incoherent model used here might be regarded as the worst possible type of glitch, in that it models a disturbance which appears exactly as a real gravitational wave would in a single detector. As usual the evaluation of the odds ratio requires to specify the prior odds, which is a subjective matter, and here we simply concentrate on the Bayes factor.

Why does this work? If we think of the parameter space of the coherent model as being embedded in the larger parameter space $\{\tilde{\theta}^{(1)}, \tilde{\theta}^{(2)}, \dots, \tilde{\theta}^{(N_D)}\}$ of the incoherent model, we see that it lies in the nine dimensional subspace where $\tilde{\theta}^{(1)} = \tilde{\theta}^{(2)} = \dots = \tilde{\theta}^{(N_D)}$. If a coherent signal is present, the distribution in the incoherent space will be peaked on or near this subspace, which will intersect a relatively large total probability mass. As the coherent subspace has fewer dimensions, in this case 9^{N_D-1} , its prior is smaller and the model is more predictive. Loosely stated, this means that it will gain whenever its prediction is correct over the more general incoherent model. This is the so-called Occam factor which arises from comparing models with different predictive power.

The larger space of the incoherent model will also capture a greater amount of evidence from the data when a large glitch is present which causes a high likelihood in many parts of the parameter space, as any signal which is not completely parametrized by the chosen parametrization will produce a broader peak on that manifold (such a glitch may provide a high evidence value when compared with the noise model only, but contains less information about the parameters). In this case (unless there is by chance a coherent glitch in the other detectors) the same argument will apply and cause the test to discriminate against the coherent signal model.

Performing model selection with the Bayes factor, Eq. (50), will then give us the optimal means of distinguishing a coherent gravitational wave from incoherent glitches which have a significant component which looks like an inspiral signal. This can be thought of as an enhanced coincidence check, which uses all the parameters to check for consistency in the observed signals, and also incorporates a check for a better explanation as simultaneous glitches.

Here we provide two examples to test this technique. First we consider how the test performs in the analysis of data sets that contain gravitational wave signals that have identical parameters, but are characterized by an unphysical time-offset; in the second case we analyze the case in which a gravitational wave is present only in the data of one of the instruments of the network.

For the first test, we inject a signal characterized by identical physical parameters—we choose $M_1 = 4M_\odot$, $M_2 = 3M_\odot$ and random position, orientation and distance such that the network optimal SNR is 17.8 in the case of a coherent injection—into the simulated data streams of LIGO-Hanford, LIGO-Livingston and Virgo; we first perform the injection coherently in the data of the three interferometers, and then we repeat it by introducing a nonphysical time shift $\Delta T \leq 10$ ms in the time of coalescence at the different sites (the coherent injection case corresponds therefore to $\Delta T = 0$). The noise realization is identical in each case; however, due to the slightly offset time of arrival of the signals and therefore the slightly different sum of data and signal, there is some spread in the recovered evidences. We compute $B_{\text{coh,inc}}$, Eq. (50) as a function of ΔT , considering the case in which the analysis is carried out using the two LIGO instruments (HL-network) and the three-instrument (HLV-network). The results are shown in Fig. 10, where we plot the Bayes factor against the time shift ΔT . As expected, when the signal is injected with $\Delta T = 0$ the evidence favors the coherent model strongly, by a factor of ~ 36461 in the case of the HLV-network and ~ 328 in the HL-network. As the time shift ΔT increases (the light-travel-time between Hanford and Livingston is ≈ 10 msec), the evidence switches to favoring the incoherent model at around 4 msec in both HLV and HL networks, and rapidly decreases to strongly exclude the coherent model with a Bayes factor of $< 10^{-10}$. With a large separation between the signals, the incoherent model becomes very strongly preferred, by a




FIG. 10 (color online). Bayes factor $B_{\text{coh,inc}}$ between the coherent and incoherent model, Eq. (50) with varying relative time delay between the arrival time of a gravitational wave at different interferometers. In solid blue, the value of $B_{\text{coh,inc}}$ when the network of the LIGO-Hanford, LIGO-Livingston and Virgo instruments (HLV) was used, and in dashed red the network of the two LIGO interferometers (HL). The time delay in milliseconds is applied equally between Hanford and Livingston, and Livingston and Virgo. The signal used had component masses $3M_\odot$ and $4M_\odot$, and a coherent network signal-to-noise ratio of 17.8. The curve shows a strong fall in coherence probability above a time-delay of ≈ 3.5 msec, allowing us to rule out a single coherent signal.

factor up to $\sim 2 \times 10^{22}$ in the case of the HLV and $\sim 1 \times 10^{24}$ in the case of the HL network. One would naively assume that HLV would yield more stringent rejection than HL, which is not the case here. We note that for $\Delta T \leq 3.5$ msec the coherent model is still favored, as the probability distribution still contains a sufficient probability of intersecting the coherent signal manifold. Although this plot is merely representative of a single test of coherence and particular details will vary with the signals and data sets, we have specifically used identical signals so that the PDF should peak at the same value in all but the time parameter. This should correspond to a situation where it is extremely difficult to determine whether or not the multiple signals are incoherent or coherent and provide a challenging test of the method.

We consider now a second test of this method. A situation which commonly arises is the presence of a ‘‘glitch’’ or instrumental artefact in a single interferometer which is not present in the others in the network. This situation is handled in the case of an incoherent search by checking that corresponding triggers, with consistent physical parameters, exist in all interferometers. As a useful sanity check for the coherence test, we examined the case in which the glitch has exactly the functional form of a gravitational wave from an inspiral signal, but is present only in one instrument. In this case, the Bayes factor of the coherent model against the Gaussian noise model is elevated, however our coherence test should allow us to exclude an event such as this due to the lack of a consistent signal in the other detectors.

Table II displays the detailed results of performing the coherent and incoherent analyses on a signal with an SNR of 9.8 injected only into LIGO-Hanford (H) simulated data. It is notable that even using the coherent model, the signal causes an elevated Bayes factor to be found, as there is some set of parameters which give a compromise signal in the network of detectors. This gives us the undesirable situation where $B_{S,N}$, the Bayes factor of signal against Gaussian noise, Eq. (44), would be triggered by an event in a single detector. However, by performing the coherence

test, we can see that the incoherent model is favored by a factor $\approx 10^3$ to the coherent model, indicating that this situation is unlikely to be a true coherent gravitational wave, and so it can be safely ruled out.

This also works in the case of a two-detector network, although in this case there is a much stronger possibility that a signal may be observed in only one of the two detectors, and so the corresponding analysis infers that there is only a 2.9 times greater chance of the incoherent model than the coherent one.

This test of coherence gives us a powerful means of distinguishing coherent from incoherent events, which can be used to quantify the additional confidence that we achieve through the use of a network. This result follows naturally from the precise statement of the hypotheses using the conceptual and computational framework that we have developed and demonstrates the power that Bayesian methodology can bring.

2. Parameter resolution

We have shown in Section III B that from the output of the nested sampling algorithm for the evidence/Bayes factor computation one can construct at no additional computational costs the marginalized posterior PDFs on the unknown source parameters. Here we show an example of the evaluation of such PDFs with nested sampling as a function of the number of instruments in the network, applied to the coherent observations of an in-spiral binary signal. We consider a system with an optimal signal-to-noise ratio of 9.3, 12.8, and 14.4, respectively, in the simulated network configurations of Hanford only, Hanford-Livingston and Hanford-Livingston-Virgo. The actual values of the parameters used for the injection are shown by the black vertical lines in Fig. 11.

Each instrument measures essentially two independent quantities—an amplitude and a phase—as a function of time. As the duration of an inspiral is negligible in comparison to the period of rotation of the Earth, there is no observable evolution of the antennae response functions during the period of observation (from which one would otherwise reconstruct the source location in the sky). From the signal strain and the time of arrival of the gravitational wave burst, one must infer the parameters \mathcal{M} , η , t_0 , D_L , α , δ , ψ , and ι . The chirp mass \mathcal{M} (and to lesser extent the symmetric mass ratio η) determines the phase evolution of the signal, which provides a large amount of information to constrain this parameter independent of the others listed here. However, in the remaining parameters a large degree of degeneracy is present, producing correlated joint posterior PDFs. For the case of observations with one interferometer, this manifests as a broad distribution in the one-dimensional marginal PDFs shown in Fig. 11. When there is insufficient information available to determine these parameters, the posterior PDFs can be influenced more

TABLE II. Results of performing the coherence test on a signal injected only into the LIGO-Hanford simulated data set. The test successfully rules out a signal which does not appear in more than one detector, despite the coherent signal vs noise comparison ($B_{S,N}$) still favoring the signal model for HLV and HL observations.

Instruments	$\log_e Z$		$\log_e B_{S,N}$		$\log_e B_{\text{coh,inc}}$
	\mathcal{H}_{inc}	\mathcal{H}_{coh}	\mathcal{H}_{inc}	\mathcal{H}_{coh}	
H	...	-5950.45	...	45.78	...
L	...	-6106.47	...	0.39	...
V	...	-6059.44	...	-0.80	...
HL	-12 056.93	-12 058.00	46.16	45.09	-1.07
HLV	-18 116.37	-18 123.29	45.36	7.52	-6.92

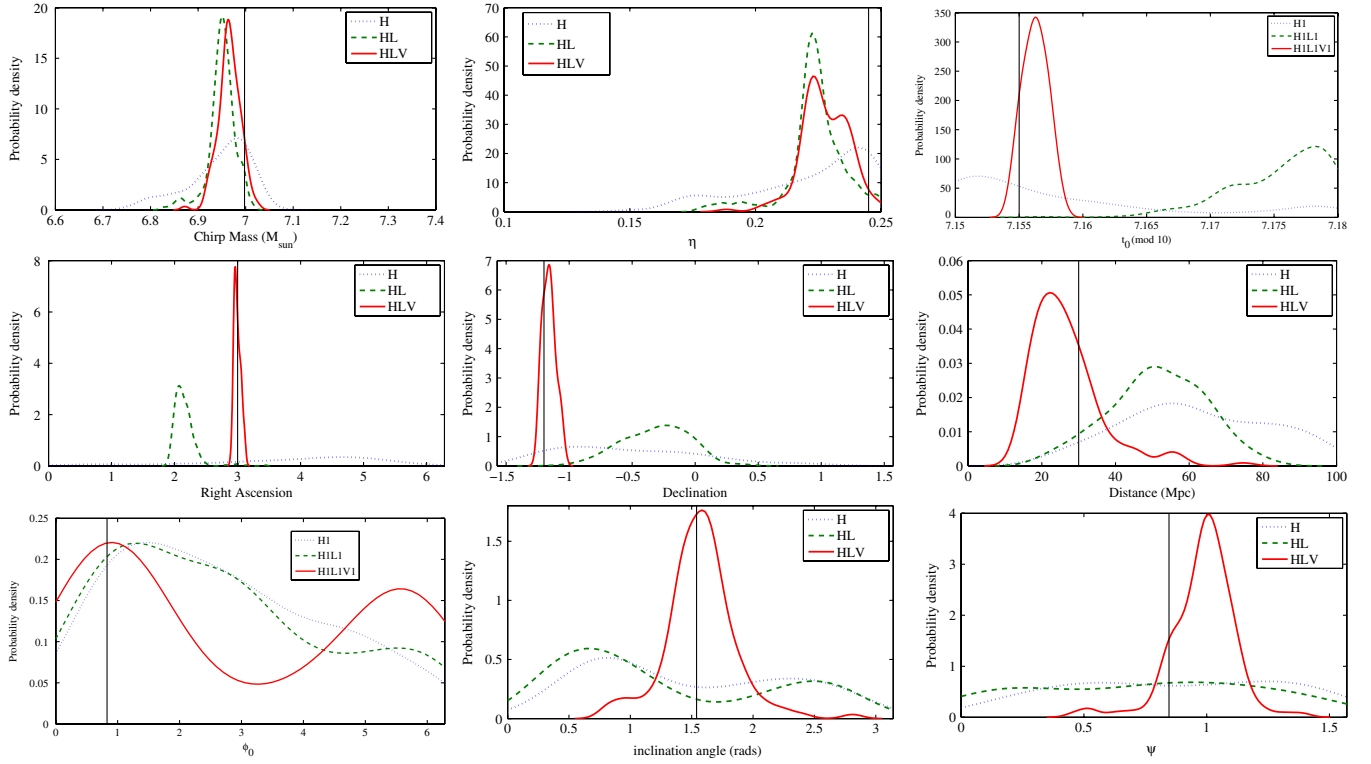


FIG. 11 (color online). The one-dimensional marginalized posterior probability density functions of the nine parameters that describe a gravitational wave inspiral signal from a circular binary of nonspinning compact objects. The plots show the effect of the coherent network analysis on the estimation of the parameters. The true values of each parameter is indicated by a vertical, black line. The signal was injected in simulated Gaussian and stationary noise representing the LIGO-Hanford, LIGO-Livingston, and Virgo instruments, with an optimal signal-to-noise ratio of 9.3 (H), 12.8 (HL), and 14.4 (HLV), respectively. As more detectors are added to the network, the parameters of the signal become better constrained, with three detectors being necessary to fully resolve all the signal parameters. Lines represent kernel density estimates of the parameters, based on the samples from the PDF generated as in Sec. III B. Edge effects from the smoothing function are responsible for lowered density estimates near the edges in the distributions of ϕ_0 and η . The density estimation was performed using the Matlab function `ksdensity`.

strongly by the prior distribution, as it is the case, in particular, for D_L , α , δ , ψ , and ι .

With the addition of a second, independent detector at a geographically different location, the possible sky locations and times of arrival are strongly constrained to lie on the surface described in Sec. IVA 1, but there is still substantial uncertainty in the marginal distributions for the chosen parameters. This is finally broken when a third detector is added to the network (red line in Fig. 11), which allows the sky location to be determined uniquely, along with the remaining parameters. In this figure we can see the evolution of the posterior PDFs with additional observations, as described above. For this injection, an inclination was used which placed the orbital plane of the binary almost edge-on to the line of sight to the Earth, meaning that only one polarization was detectable. This is reflected in the underestimation of the distance to the system for a network of less than 3 detectors, as there are many positions which are not edge-on which lead to a higher overall observed signal amplitude, and therefore must be located farther away.

This provides an example of the necessity of multiple detectors, operating as a network, if we are to make full use of the astrophysical information carried by gravitational waves. By underdetermining the parameters of the signal, biases or at least additional uncertainties may be introduced into our conclusions about the nature of observed sources.

VI. CONCLUSIONS

By taking a Bayesian approach to the analysis of data for the detection and characterization of in-spiral signals, we have been able to implement a conceptually simple yet flexible framework for drawing inference from observations. Although the Bayesian formalism calls for the evaluation and integration of high-dimensional likelihood functions, we have shown that the nested sampling technique provides us with a means to both search for and estimate the parameters of a signal. Further work remains to be done in improving the efficiency and reliability of detection at low masses, but the particular implementation

that we describe here provides a solid basis for these future improvements.

We have used our implementation to demonstrate the power of Bayesian model selection in classifying putative gravitational wave signals, through the use of the coherence test described in Sec. VC. The coherence test goes some way to implementing a robust and Bayesian defence against glitches present in gravitational wave data which do not resemble coherent gravitational waves by providing an internal consistency check within the detector network. Tests of this kind may provide a useful new additional discriminator when analyzing candidate gravitational waves, and are only achievable through the treatment of the detector network in a coherent way. Notably, this test is only possible through the use of the Occam factor and the comparison of probability distributions which differ in their dimensionality, and could not be possible with point estimates of maximum likelihood. Indeed, the maximum likelihood of the coherent and incoherent models can be trivially shown to be identical. In addition to such tests, it would also be desirable to model the detector data in a way which is nonstationary or Gaussian, but this is beyond the scope of this work. Furthermore, the use of the Bayesian parameter estimation framework is invaluable in inferring the signal parameters, and produces the full posterior probability distribution, not just maximum likelihood points and estimates of variance. The covariance and interdependence of the parameters does not prevent us from calculating consistent joint PDFs on the full parameter space even when point estimates have little meaning, and this allows us to see the benefit of using a network of detectors in a coherent fashion.

One of the main benefits in the use of the nested sampling and Bayesian evidence approach is the ease with which it can be extended or adapted to different signal models with minimal changes needed to the implementation. In other work past and ongoing, we have shown how this method may be applied to the discrimination between candidate waveforms when comparing to a numerical relativity simulation; how one may place bounds on the Compton wavelength of the graviton given a single or multiple observations of an inspiral signal, and testing the effects of including spins in the detection and estimation of inspiral signals [28,78].

Future work on the implementation of the core algorithm will focus on achieving a better reliability and efficiency in the sampling of the parameter space, which should lead to further improvements in the performance in a real world situation, and on testing on the full set of waveform approximants used in present searches for coalescing binary systems. The work presented here forms the basis of the LALAPPS_INSPNEST program, which is can be found in the LALApps software distribution, released under the terms of the GNU general public license [31].

ACKNOWLEDGMENTS

This work has benefited from many discussions with members of the LIGO Scientific Collaboration, in particular, we would like to thank Martin Hendry, Ilya Mandel, and Christian Röver for their comments and suggestions. The simulations presented in this paper were carried out on the Tsunami Beowulf clusters of the University of Birmingham and the Atlas cluster of the AEI Hannover. This work has been supported by the UK Science and Technology Facilities Council. This paper has been assigned the LIGO Document Number LIGO-P0900117.

APPENDIX: DEFINITIONS AND CONVENTIONS

We provide the definitions and conventions that we have adopted to link the continuous representation of a time series which is used in the main body of the paper with the discrete representation. The latter is what is actually used for applications, and corresponds to the expressions implemented in the LALAPPS_INSPNEST software. This appendix provides also a mapping of the notation used in our previous papers [25,26] and the one used here.

Consider a time series $a(t)$ sampled at time intervals $\Delta t = 1/(2f_{Ny})$, where f_{Ny} is the *Nyquist frequency*, for a total observation time T . The total number of data samples is therefore $N_p = T/\Delta t = 2Tf_{Ny}$. Given a generic real function $a(t)$, our conventions for the Fourier transform are [79]:

$$\tilde{a}(f) = \int_{-\infty}^{+\infty} a(t)e^{-2\pi ift} dt, \quad (\text{A1})$$

and the inverse transform is

$$a(t) = \int_{-\infty}^{+\infty} \tilde{a}(f)e^{+2\pi ift} df. \quad (\text{A2})$$

The data points at time t_j and frequency f_k are therefore:

$$a(t_j) = a(j\Delta t) = a_j \quad (\text{A3})$$

$$\tilde{a}(f_k) = \tilde{a}(k/T) = \Delta t \times \tilde{a}_k, \quad (\text{A4})$$

where we have defined the Fourier series as

$$\tilde{a}_k = \sum_j a_j e^{-2\pi ijk/N_p}, \quad (\text{A5})$$

$$a_j = \frac{1}{N_p} \sum_k \tilde{a}_k e^{+2\pi ijk/N_p}. \quad (\text{A6})$$

In the following we will indicate with $\tilde{a}(f_k)$ the (dimensionfull) approximation to the Fourier Transform of $\tilde{a}(f)$ at frequency $f = f_k$.

We consider now the statistical properties of the noise. The *one-sided* noise spectral density $S(f)$ is defined in the continuous case as:

$$S(f) = 2 \int_{-\infty}^{+\infty} \langle n(t + \tau)n(\tau) \rangle e^{-i2\pi f t} dt, \quad (\text{A7})$$

which yields

$$\langle \tilde{n}(f)\tilde{n}^*(f') \rangle = \frac{1}{2}S(f)\delta(f - f') = S^{(2)}(f)\delta(f - f'), \quad (\text{A8})$$

where the factor 1/2 comes from the fact that $S(f)$ is the one-sided noise spectral density; this is related to the *two-sided* noise spectral density $S^{(2)}(f)$ by

$$S^{(2)}(f) = \frac{1}{2}S(f). \quad (\text{A9})$$

If we explicitly write the real and imaginary part of the noise contribution in the Fourier domain, which is the notation adopted in [25,26] as

$$\tilde{n}(f_k) = \tilde{x}(f_k) + i\tilde{y}(f_k), \quad (\text{A10})$$

and we substitute into Eq. (A8), we obtain:

$$\langle |\tilde{n}(f_k)|^2 \rangle = \langle |\tilde{x}(f_k)|^2 \rangle + \langle |\tilde{y}(f_k)|^2 \rangle = \frac{T}{2}S(f_k), \quad (\text{A11})$$

where the first equality comes from the fact that $\tilde{x}(f_k)$ and $\tilde{y}(f_k)$ are independent. In terms of the elements of the Fourier series one has

$$\Delta t^2 \langle |\tilde{n}_k|^2 \rangle = \Delta t^2 [\langle |\tilde{x}_k|^2 \rangle + \langle |\tilde{y}_k|^2 \rangle]. \quad = \frac{T}{2}S(f_k) \quad (\text{A12})$$

If one defines

$$\sigma_k^2 = \langle |\tilde{n}_k|^2 \rangle \quad \zeta_k^2 = \langle |\tilde{x}_k|^2 \rangle = \langle |\tilde{y}_k|^2 \rangle \quad \sigma_k^2 = 2\zeta_k^2 \quad (\text{A13})$$

the variance of the (complex) noise and the real and imaginary part of the Fourier series elements, then

$$S(f_k) = 2 \frac{\Delta t^2}{T} \sigma_k^2 = 2 \frac{\Delta t}{N} \sigma_k^2, \quad = 4 \frac{\Delta t}{N} \zeta_k^2. \quad (\text{A14})$$

Let us now considered the usual inner product $(\cdot|\cdot)$ [58] between two (real) functions a and b and its approximation in the finite case:

$$(a|b) = 2 \int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f) + \tilde{a}^*(f)\tilde{b}(f)}{S(f)} df \quad (\text{A15})$$

$$\approx \frac{2}{T} \sum_{k>0} \frac{\tilde{a}(f_k)\tilde{b}^*(f_k) + \tilde{a}^*(f_k)\tilde{b}(f_k)}{S(f_k)} \quad (\text{A16})$$

$$\approx \sum_{k>0} \frac{\tilde{a}_k\tilde{b}_k^* + \tilde{a}_k^*\tilde{b}_k}{\sigma_k^2} \quad (\text{A17})$$

The optimal signal-to-noise ratio is just the square root of the norm of h , and using Eq. (A17) it yields

$$(h|h) = 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{S(f)} df, \quad \approx \frac{4}{T} \sum_{k>0} \frac{|\tilde{h}(f_k)|^2}{S(f_k)}, \quad \approx 2 \sum_{k>0} \frac{|\tilde{a}_k|^2}{\sigma_k^2}. \quad (\text{A18})$$

The likelihood function for the data set d given the model h , Eq. (22) therefore becomes

$$p(d|h, \mathcal{H}_S) \propto \exp\left[-\frac{1}{2}(d - h|d - h)\right] \quad (\text{A19})$$

$$\propto \exp\left[-\frac{2}{T} \sum_{k>0} \frac{|\tilde{d}(f_k) - \tilde{h}(f_k)|^2}{S(f_k)}\right], \quad (\text{A20})$$

and this is the expression used in the software implementation through Eqs. (A4), (A5), and (A14).

-
- [1] B. C. Barish and R. Weiss, *Phys. Today* **52**, 44 (1999).
[2] LIGO Scientific Collaboration, *Rep. Prog. Phys.* **72**, 076901 (2009).
[3] F. Acernese *et al.*, *Classical Quantum Gravity* **21**, S385 (2004).
[4] B. Willke *et al.*, *Classical Quantum Gravity* **21**, S417 (2004).
[5] R. Adhikari, P. Fritschel, and S. Waldman, LIGO Technical Document Report No. T060156-01, 2006.
[6] J. R. Smith (LIGO Scientific Collaboration), *Classical Quantum Gravity* **26**, 114013 (2009).
[7] <http://www.ligo.caltech.edu/advLIGO/>.
[8] <http://www.cascina.virgo.infn.it/advirgo/>.
[9] C. Cutler and K. S. Thorne, in *General Relativity and Gravitation*, edited by N. T. Bishop and D. M. Sunil (World Scientific, Singapore, 2002), p. 72.
[10] K. Kokkotas, *Reviews in modern astronomy* **20**, 140 (2008).
[11] L. Blanchet, *Living Rev. Relativity* **9**, 4 (2006).
[12] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **80**, 047101 (2009).
[13] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **79**, 122001 (2009).
[14] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **78**, 042002 (2008).
[15] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **77**, 062002 (2008).
[16] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **73**, 062001 (2006).
[17] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 082002 (2005).
[18] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 082001 (2005).
[19] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 122001 (2004).
[20] V. Kalogera, K. Belczynski, C. Kim, R. W.

- O'Shaughnessy, and B. Willems, Phys. Rep. **442**, 75 (2007).
- [21] J. Skilling, in *24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, AIP Conf. Proc. No. 735 (AIP, New York, 2004), pp. 395–405.
- [22] <http://www.inference.phy.cam.ac.uk/bayesys/>.
- [23] D. S. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial* (Oxford University Press, New York, 2006), 2nd ed..
- [24] We note that searches for low-mass (that is with total mass $\leq 35M_{\odot}$) binary systems in the data now collected by LIGO and Virgo are based on templates computed at the post^{3.5}-Newtonian order. All the results presented here apply directly to that case at no additional computational costs.
- [25] J. Veitch and A. Vecchio, Phys. Rev. D **78**, 022001 (2008).
- [26] J. Veitch and A. Vecchio, Classical Quantum Gravity **25**, 184010 (2008).
- [27] L. Cadonati *et al.*, Classical Quantum Gravity **26**, 114008 (2009).
- [28] B. Aylott *et al.*, Classical Quantum Gravity **26**, 165008 (2009).
- [29] B. Aylott, J. Veitch, and A. Vecchio, Classical Quantum Gravity **26**, 114011 (2009).
- [30] <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.
- [31] <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalapps.html>.
- [32] P. J. Green, Biometrika **82**, 711 (1995).
- [33] A. Gelman and X.-L. Meng, Stat. Sci. **13**, 163 (1998).
- [34] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Chapman and Hall/CRC, Boca Raton, 1996).
- [35] A. Gelman, J. B. Carlin, H. Stern, and D. B. Rubin, *Bayesian Data Analysis* (Chapman and Hall, Boca Raton, 1997).
- [36] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).
- [37] C. J. Geyer and E. A. Thompson, J. Am. Stat. Assoc. **90**, 909 (1995).
- [38] L. Tierney and A. Mira, Stat. Med. **18**, 2507 (1999).
- [39] A. Mira, Metron **LIX**, No. 3–4 231 (2001).
- [40] P. J. Green and A. Mira, Biometrika **88**, 1035 (2001).
- [41] M. Trias, A. Vecchio, and J. Veitch, arXiv:0904.2207.
- [42] R. Umstätter and M. Tinto, Phys. Rev. D **77**, 082002 (2008).
- [43] J. Clark, I. S. Heng, M. Pitkin, and G. Woan, Phys. Rev. D **76**, 043003 (2007).
- [44] A. C. Searle, P. J. Sutton, and M. Tinto, Classical Quantum Gravity **26**, 155017 (2009).
- [45] R. J. Dupuis and G. Woan, Phys. Rev. D **72**, 102002 (2005).
- [46] Abbott *et al.*, Astrophys. J. **683**, L45 (2008).
- [47] B. P. Abbot *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), arXiv:0909.3583.
- [48] M. V. van der Sluys, C. Roever, A. Stroeer, N. Christensen, V. Kalogera, R. Meyer, and A. Vecchio, Astrophys. J. **688**, L61 (2008).
- [49] C. Rover, R. Meyer, and N. Christensen, Classical Quantum Gravity **23**, 4895 (2006).
- [50] C. Röver, R. Meyer, and N. Christensen, Classical Quantum Gravity **23**, 4895 (2006).
- [51] C. Röver, R. Meyer, and N. Christensen, Phys. Rev. D **75**, 062004 (2007).
- [52] C. Röver, R. Meyer, G. M. Guidi, A. Viceré, and N. Christensen, Classical Quantum Gravity **24**, S607 (2007).
- [53] M. van der Sluys *et al.*, Classical Quantum Gravity **25**, 184011 (2008).
- [54] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and É. É. Flanagan, Phys. Rev. D **63**, 042003 (2001).
- [55] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Phys. Rev. D **80**, 084043 (2009).
- [56] We will indulge in the LSC jargon, and use the term *injection* to denote a gravitational wave signal added to noise.
- [57] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **57**, 885 (1998).
- [58] C. Cutler and E. E. Flanagan, Phys. Rev. D **49**, 2658 (1994).
- [59] or could be estimated using the Welch method for a real analysis.
- [60] F. Feroz and M. P. Hobson, arXiv:0704.3704.
- [61] F. Feroz, M. P. Hobson, and M. Bridges, arXiv:0809.3437.
- [62] F. Feroz, P. J. Marshall, and M. P. Hobson, arXiv:0810.0781.
- [63] F. Feroz, M. P. Hobson, J. T. L. Zwart, R. D. E. Saunders, and K. J. B. Grainge, arXiv:0811.1199.
- [64] M. Bridges, F. Feroz, M. P. Hobson, and A. N. Lasenby, arXiv:0812.3541.
- [65] F. Feroz, B. C. Allanach, M. Hobson, S. S. AbdusSalam, R. Trotta, and A. M. Weber, J. High Energy Phys. **10** (2008) 064.
- [66] R. Trotta, F. Feroz, M. P. Hobson, L. Roszkowski, and R. Ruiz de Austri, J. High Energy Phys. **12** (2008) 024.
- [67] F. Feroz, M. P. Hobson, L. Roszkowski, R. Ruiz de Austri, and R. Trotta, arXiv:0903.2487.
- [68] S. S. AbdusSalam, B. C. Allanach, F. Quevedo, F. Feroz, and M. Hobson, arXiv:0904.2548.
- [69] S. S. AbdusSalam, B. C. Allanach, M. J. Dolan, F. Feroz, and M. P. Hobson, Phys. Rev. D **80**, 035017 (2009).
- [70] D. E. Lopez-Fogliani, L. Roszkowski, R. R. de Austri, and T. A. Varley, Phys. Rev. D **80**, 095013 (2009).
- [71] F. Feroz, J. R. Gair, M. P. Hobson, and E. K. Porter, Classical Quantum Gravity **26**, 215003 (2009).
- [72] The *circular* mean of a set of angles θ_i is defined by $\bar{\theta} = \tan^{-1} \frac{\overline{\sin\theta_i}}{\overline{\cos\theta_i}}$.
- [73] C. J. F. Ter Braak, Stat. Comput. **16**, 239 (2006).
- [74] H. Jeffreys, Proc. R. Soc. A **186**, 453 (1946).
- [75] S. Amari and H. Nagaoka, Trans. Math. Monogr. **191** (1993).
- [76] C. Röver, LIGO Report No. P0900228.
- [77] R. Gouaty (LIGO Scientific Collaboration), Classical Quantum Gravity **25**, 184006 (2008).
- [78] W. del Pozzo, J. Veitch, and A. Vecchio (unpublished).
- [79] We are adopting here the convention used by the GW data analysis work and the LIGO Scientific Collaboration, which is also the engineering convention.