

Bayesian detection of unmodeled bursts of gravitational waves.

Antony C. Searle

*The Australian National University, Canberra, ACT 0200, Australia and
LIGO - California Institute of Technology, Pasadena, CA 91125*

Patrick J. Sutton

School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom

Massimo Tinto

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

(Dated: September 12, 2008)

The data analysis problem of coherently searching for unmodeled gravitational-wave bursts in the data generated by a global network of gravitational-wave observatories has been at the center of research for almost two decades. As data from these detectors is starting to be analyzed, a renewed interest in this problem has been sparked. A Bayesian approach to the problem of coherently searching for gravitational wave bursts with a network of ground-based interferometers is here presented. We demonstrate how to systematically incorporate prior information on the burst signal and its source into the analysis. This information may range from the very minimal, such as best-guess durations, bandwidths, or polarization content, to complete prior knowledge of the signal waveforms and the distribution of sources through spacetime. We show that this comprehensive Bayesian formulation contains several previously proposed detection statistics as special (unphysical) limiting cases, and demonstrate that it outperforms them.

PACS numbers: 04.80.Nn, 95.55.Ym, 07.05.Kf

Keywords: Gravitational Waves, Laser Interferometry

I. INTRODUCTION

Large-scale, broad-band interferometric gravitational-wave observatories [1–4] are operating at their anticipated sensitivities, and scientists around the globe have begun to analyze the data generated by these instruments in search of gravitational wave signals [5]. Gravitational wave bursts (GWBs) are among the most exciting signals scientists expect to observe, as our present knowledge and modeling ability of GWB-emitting systems is rather limited. These signals often depend on complicated (and interesting) physics, such as dynamical gravity and the equation of state of matter at nuclear densities. While this makes GWBs an especially attractive target to study, our lack of knowledge also limits the sensitivity of searches for GWBs. Potential sources of GWBs include merging compact objects [6–14], core-collapse supernovae [15–19], and gamma-ray burst engines [20]; see [21] for an overview.

Although a gravitational wave signal is characterized by two independent polarizations, an interferometric gravitational wave observatory is sensitive to a single linear combination of them. Simultaneous observations of a gravitational wave burst by three or more observatories over-determines the waveform, permitting the source position on the sky to be determined, and a least-squares estimate of the two polarizations. This solution of the *inverse problem* for gravitational waves was first derived by Gürsel and Tinto [22] for three interferometers, subsequently generalized to four by Tinto [23], and then an arbitrary number by various authors (see e.g.,

[7, 24–27]). Subsequent developments included the introduction of regularization in the waveform reconstruction, such as by Rakhmanov [27] and Klimenko *et al.* [26, 28]. For example, Summerscales *et al.* have developed a maximum-entropy based regularization technique [29]. Another development followed the observation that the over-determination of the waveform with multiple detectors also permits consistency tests that can reject non-Gaussian noise artifacts, as noted by Wen and Schutz [30]; a robust consistency test was first demonstrated by Chatterji *et al.* [24].

This paper is concerned primarily with the problem of *signal detection*. Flanagan and Hughes [7] and later Anderson, Brady, Creighton, and Flanagan [25] generalized the Gürsel-Tinto coherent analysis using a maximum-likelihood formulation, and showed how coherent techniques could be used for signal detection as well as for solving the inverse problem of extracting signal parameters. The corresponding signal detection statistic is commonly referred to as the *standard likelihood* statistic. Several modifications of the Gürsel-Tinto standard likelihood approach have recently been proposed, such as the *constraint likelihood* approach of Klimenko *et al.* [26, 31], the SNR variability approach of Mohanty *et al.* [28], and the *Tikhonov regularization* approach demonstrated by Rakhmanov [27]. Other coherent detection algorithms have been proposed by Sylvestre [32] and by Arnaud *et al.* [33]; these methods form a coherent sum of the data from different detectors weighted so as to maximize signal energy.

These approaches to signal detection have generally

been derived by following either *ad hoc* reasoning or a maximum-likelihood criterion. In this paper we present a systematic and comprehensive Bayesian formulation [34, 35] of the problem of coherent detection of gravitational wave bursts. We demonstrate how to incorporate partial or incomplete knowledge of the signal in the analysis, to improve the probability of detection of these weak signals. This information may include time-frequency properties of the signal, polarization content, model waveform families or templates, as well as information on the distribution of the source through space-time. We also explicitly identify the prior assumptions that must be made about the signal to cause a Bayesian analysis to behave like several of the previously proposed detection statistics. This reveals that the implicit assumptions of these methods are physically unreasonable, proposing infinite or infinitesimal signals distributed unevenly across the sky. A Bayesian analysis whose prior expectations about the signal better correspond to reality will necessarily perform better, and we implement a Monte-Carlo simulation to quantify this improvement.

The paper is organized as follows. In §II we derive the Bayesian posterior detection probability of an idealized delta-like burst signal by a toy-model network of observatories. Greatly expanding on [36], we then generalize the derivation of the Bayesian odds ratio into a usable statistic for an arbitrary number of interferometers with realistic, differently colored and potentially correlated noises. We also consider a wide range of signal models corresponding to different states of knowledge about the burst, from total ignorance to complete *a priori* knowledge of the waveforms. In §III we characterize the relative performance of previously proposed statistics (the Gürsel-Tinto/standard and constraint likelihoods) and the Bayesian statistic by performing a Monte-Carlo simulation in which we add a simple binary black-hole merger waveform to simulated detector data and construct Receiver-Operating Characteristic (ROC) curves. We find that the Bayesian method increases the probability of detection for a given false alarm rate by approximately 50%, over those associated with previously proposed statistics.

II. ANALYSIS

A. Single-sample observation

We begin by investigating perhaps the simplest Bayesian coherent data analysis: detecting a signal from a known sky position in a single strain sample from each of N gravitational wave observatories. This example will show many of the basic features of the Bayesian analysis, and highlight some of the differences between the Bayesian approach and previous statistics. In the following section we will generalize to a multi-sample search for a signal arriving at an unknown time from an unknown sky position.

Consider a single strain sample from each of N detectors, each measurement taken at the moment corresponding to the passage of a postulated plane gravitational wave from some specific location on the sky, (θ, ϕ) . The measurements are then equal to

$$\mathbf{x} = \mathbf{F} \mathbf{h} + \mathbf{e}, \quad (1)$$

where \mathbf{x} is the vector of measurements $[x_1, \dots, x_N]^T$, the matrix $\mathbf{F} = [[F_1^+, F_1^\times], \dots, [F_N^+, F_N^\times]]$ contains the antenna responses of the observatories to the postulated gravitational wave strain vector $\mathbf{h} = [h_+, h_\times]^T$, and \mathbf{e} is the noise in each sample. \mathbf{F} is a known function of the direction (θ, ϕ) .

We wish to distinguish between two hypotheses: H_{noise} , that the data contains only noise, and H_{signal} , that the data contains a gravitational wave signal. The Bayesian odds ratio allows us to compare the plausibility of the hypotheses:

$$\frac{p(H_{\text{signal}}|\mathbf{x}, I)}{p(H_{\text{noise}}|\mathbf{x}, I)} = \frac{p(H_{\text{signal}}|I) p(\mathbf{x}|H_{\text{signal}}, I)}{p(H_{\text{noise}}|I) p(\mathbf{x}|H_{\text{noise}}, I)}, \quad (2)$$

where I is a set of unstated but shared assumptions (such as the detector locations, orientations and noise power spectra).

The $p(H|I)$ terms (“plausibility of H assuming I ”) are the *prior* plausibilities we assign to each hypothesis H on the basis of our knowledge I prior to considering the measurement; for example, our expectation that detectable gravitational waves are rare requires that $p(H_{\text{signal}}|I) \ll p(H_{\text{noise}}|I)$.

The $p(\mathbf{x}|H, I)$ terms (“plausibility of \mathbf{x} assuming H and I ”) are the probabilities assigned by a hypothesis to the occurrence of a particular observation \mathbf{x} . These are sometimes called likelihood functions; they represent the likelihood of a certain measurement being made.

The $p(H|\mathbf{x}, I)$ terms are the *posterior* plausibilities we assign to the hypotheses in light of the observation. The difference between the prior and posterior plausibility ratios caused by the observation is the ratio of the plausibilities those hypotheses assigned to that observation being made; the hypothesis that made the better prediction becomes more plausible.

If we make the idealized assumption that the noise in each detector is independent and normally distributed with zero mean and unit standard deviation, we can then write the following expression for the likelihood $p(\mathbf{x}|H_{\text{noise}}, I)$

$$\begin{aligned} p(\mathbf{x}|H_{\text{noise}}, I) &= \prod_{i=1}^N p(x_i|H_{\text{noise}}, I) \\ &= \prod_{i=1}^N \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}x_i^2\right) \\ &= (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{x}\right), \end{aligned} \quad (3)$$

where T denotes matrix transposition. For real detectors, the measurements can be *whitened* to achieve this

distribution. The changes introduced by whitening alter the effective response \mathbf{F} .

If we assume that there is a gravitational wave \mathbf{h}

present, then after subtracting away the response $\mathbf{F}\mathbf{h}$ the data will be distributed as noise and the likelihood $p(\mathbf{x}|\mathbf{h}, H_{\text{signal}}, I)$ becomes

$$p(\mathbf{x}|\mathbf{h}, H_{\text{signal}}, I) = (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{F}\mathbf{h})^T(\mathbf{x} - \mathbf{F}\mathbf{h})\right). \quad (4)$$

Unfortunately, we do not know the signal strain vector \mathbf{h} *a priori*. To compute the plausibility of the more general hypothesis $p(\mathbf{x}|H_{\text{signal}}, I)$ we need to marginalize away these *nuisance parameters*

$$p(\mathbf{x}|H_{\text{signal}}, I) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} p(\mathbf{h}|H_{\text{signal}}, I) p(\mathbf{x}|\mathbf{h}, H_{\text{signal}}, I) dh_+ dh_\times. \quad (5)$$

The hypothesis resulting from the marginalization integral is an average of the hypotheses for particular signals \mathbf{h} , weighted by the priors $p(\mathbf{h}|H_{\text{signal}}, I)$ we assign to those signals occurring. Choosing this prior distribution for \mathbf{h} is not simple, as it should reflect all our knowledge about sources. A convenient first guess is to use a normal distribution for each polarization, with a standard deviation σ indicative of the amplitude scale of gravitational waves we hope to detect. Under these assumptions the prior is

$$p(\mathbf{h}|H_{\text{signal}}, I) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{1}{2\sigma^2}\mathbf{h}^T\mathbf{h}\right). \quad (6)$$

This allows us to perform the marginalization integral analytically

$$\begin{aligned} p(\mathbf{x}|H_{\text{signal}}, I) &= (2\pi)^{-\frac{N}{2}-1}\sigma^{-2} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \exp\left(-\frac{1}{2}((\mathbf{x} - \mathbf{F}\mathbf{h})^T(\mathbf{x} - \mathbf{F}\mathbf{h}) + \sigma^{-2}\mathbf{h}^T\mathbf{h})\right) dh_+ dh_\times \\ &= (2\pi)^{-\frac{N}{2}} |\mathbf{I} - \mathbf{K}_{\text{ss}}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\mathbf{x}^T(\mathbf{I} - \mathbf{K}_{\text{ss}})\mathbf{x}\right), \end{aligned} \quad (7)$$

where

$$\mathbf{K}_{\text{ss}} = \mathbf{F}(\mathbf{F}^T\mathbf{F} + \sigma^{-2}\mathbf{I})^{-1}\mathbf{F}^T. \quad (8)$$

The result is a multivariate normal distribution with covariance matrix $(\mathbf{I} - \mathbf{K}_{\text{ss}})^{-1}$, which quantifies the correlations among the detectors due to the presence of a gravitational wave signal.

With both hypotheses defined, we can form the *likelihood ratio*

$$\begin{aligned} \Lambda &= \frac{p(\mathbf{x}|H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} \\ &= |\mathbf{I} - \mathbf{K}_{\text{ss}}|^{-\frac{1}{2}} \exp\left(\frac{1}{2}\mathbf{x}^T\mathbf{K}_{\text{ss}}\mathbf{x}\right) \\ &= |\mathbf{I} - \mathbf{K}_{\text{ss}}|^{-\frac{1}{2}} \exp\left(\frac{1}{2}\mathbf{x}^T\mathbf{F}(\mathbf{F}^T\mathbf{F} + \sigma^{-2}\mathbf{I})^{-1}\mathbf{F}^T\mathbf{x}\right). \end{aligned} \quad (9)$$

Multiplying the likelihood ratio by the prior plausibility ratio $p(H_{\text{signal}}|I)/p(H_{\text{noise}}|I)$ completes the Bayesian odds ratio of Equation 2.

The part of the likelihood ratio in the exponential can be directly compared to existing frequentist statistics. In particular, in the limit $\sigma \rightarrow \infty$ we find that the odds ratio contains the least-squares estimate of the strain

$$\hat{\mathbf{h}} = (\mathbf{F}^T\mathbf{F})^{-1}\mathbf{F}^T\mathbf{x}. \quad (10)$$

The odds ratio may then be rewritten in terms of a matched filter for the response to the estimated strain, $\mathbf{x}^T\mathbf{F}\hat{\mathbf{h}}$. For finite values of σ , the odds ratio contains the *Tikhonov regularized estimate* of the strain

$$\hat{\mathbf{h}} = (\mathbf{F}^T\mathbf{F} + \sigma^{-2}\mathbf{I})^{-1}\mathbf{F}^T\mathbf{x}, \quad (11)$$

and can still be rewritten as a matched filter for this estimate. We discuss the relationship of the Bayesian to Tikhonov and other statistics in more detail in Section IIF.

It is also worth noting the presence of the determinant $|\mathbf{I} - \mathbf{K}_{\text{ss}}|$ factor. It is independent of the data and depends only on the detector sensitivities and the signal model. In particular, it tells us how strongly to weight likelihoods computed for different possible sky positions of the signal. This *Occam factor* penalizes sky positions of high sensitivity relative to sky positions of lower sensitivity which give similar exponential part of the likelihood. This factor is ignored in frequentist approaches. The effect is typically small compared to the exponential in most cases if the data has good evidence for a signal, but can be important for weak signals and for parameter estimation, as we shall see in Section IIF.

B. General Bayesian model

We now generalize the analysis of the previous section to the case of burst signals of extended duration and unknown direction and arrival time.

A global network of N gravitational wave detectors each produce a time-series of M observations with sampling frequency f_s , which we pack into a single vector

$$\mathbf{x} = [x_{1,1}, x_{1,2}, \dots, x_{1,M}, x_{2,1}, \dots, x_{N,M}]^T. \quad (12)$$

We want to classify the observation as a gravitational wave detection or not. Bayesian inference does not allow us to *reject* a hypothesis in isolation, so we must propose (at least) two hypotheses and compute which is more plausible. We will consider a signal hypothesis H_{signal} and a noise hypothesis H_{noise} . The ability of the observation to distinguish between these two hypotheses is contingent upon the observation being differently distributed for each hypothesis

$$p(\mathbf{x}|H_{\text{signal}}, I) \neq p(\mathbf{x}|H_{\text{noise}}, I). \quad (13)$$

To compute these plausibility distributions, we must explicitly form a model of the experiment. For H_{signal} , we will use the model

$$\mathbf{x} = \mathbf{F}(\tau, \theta, \phi) \cdot \mathbf{h} + \mathbf{e} \quad (14)$$

where

$$\mathbf{h} = [h_{+,1}, h_{+,2}, \dots, h_{+,L}, h_{\times,1}, \dots, h_{\times,L}]^T \quad (15)$$

is a time-series of L samples describing the band-limited strain waveform (with the two polarizations packed into a single vector), (τ, θ, ϕ) are respectively the time of arrival and direction of the gravitational wave, \mathbf{e} is a random variable with probability distribution $p(\mathbf{e}|I)$ representing the instrumental noise, and $\mathbf{F}(\tau, \theta, \phi)$ is a $NM \times 2M$ response matrix describing the response of each observatory to an incoming gravitational wave.

$$\mathbf{F}(\tau, \theta, \phi) = \begin{bmatrix} F_1^+(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_1(\theta, \phi)) & F_1^\times(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_1(\theta, \phi)) \\ F_2^+(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_2(\theta, \phi)) & F_2^\times(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_2(\theta, \phi)) \\ \vdots & \vdots \\ F_N^+(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_N(\theta, \phi)) & F_N^\times(\theta, \phi)\mathbf{T}(\tau + \Delta\tau_N(\theta, \phi)) \end{bmatrix}. \quad (16)$$

Each $M \times L$ block of the response matrix is responsible for scaling and time shifting one of the waveform polarizations for one detector, so each block is the product of the scalar directional sensitivity of each detector to a polarization, $F_i^+(\theta, \phi)$ or $F_i^\times(\theta, \phi)$, and a time delay matrix, $T_{j,k}(t) = \text{sinc}(\pi(j-k - f_s t))$, for the direction dependent arrival times $\tau + \Delta\tau_i(\theta, \phi)$ at each detector.

We can restate the equality in Equation 14 as a Dirac delta-function plausibility distribution

$$p(\mathbf{x}|\mathbf{e}, \mathbf{h}, \tau, \theta, \phi, H_{\text{signal}}, I) = \delta(\mathbf{x} - \mathbf{F}(\tau, \theta, \phi) \cdot \mathbf{h} - \mathbf{e}). \quad (17)$$

We can then use the marginalization theorem to compute the likelihood of the data given the hypothesis that a burst is present [34]

$$p(\mathbf{x}|H_{\text{signal}}, I) = \int_{V_{\mathbf{e}, \mathbf{h}, \tau, \theta, \phi}} p(\mathbf{x}|\mathbf{e}, \mathbf{h}, \tau, \theta, \phi, H_{\text{signal}}, I) p(\mathbf{e}, \mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) d\mathbf{e} \dots d\phi, \quad (18)$$

where $V_{\mathbf{e}, \mathbf{h}, \tau, \theta, \phi}$ is the space of all our parameter values.

Similarly we can state the noise model hypothesis as a Dirac delta-function plausibility distribution, implying the expression

$$p(\mathbf{x}|\mathbf{e}, H_{\text{noise}}, I) = \delta(\mathbf{x} - \mathbf{e}) \quad (19)$$

$$p(\mathbf{x}|H_{\text{noise}}, I) = \int_{V_{\mathbf{e}}} p(\mathbf{x}|\mathbf{e}, H_{\text{noise}}, I) p(\mathbf{e}|H_{\text{noise}}, I) d\mathbf{e}, \quad (20)$$

where $V_{\mathbf{e}}$ is the space of the noise.

By using the above expressions we can now construct the *Bayes factor*

$$\frac{p(\mathbf{x}|H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \frac{\int_{V_{\mathbf{e}, \mathbf{h}, \tau, \theta, \phi}} p(\mathbf{x}|\mathbf{e}, \mathbf{h}, \tau, \theta, \phi, H_{\text{signal}}, I) p(\mathbf{e}, \mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) d\mathbf{e} \dots d\phi}{\int_{V_{\mathbf{e}}} p(\mathbf{x}|\mathbf{e}, H_{\text{noise}}, I) p(\mathbf{e}|H_{\text{noise}}, I) d\mathbf{e}}, \quad (21)$$

and the *posterior plausibility ratio* is equal to

$$\frac{p(H_{\text{signal}}|\mathbf{x}, I)}{p(H_{\text{noise}}|\mathbf{x}, I)} = \frac{p(\mathbf{x}|H_{\text{signal}}, I) p(H_{\text{signal}}|I)}{p(\mathbf{x}|H_{\text{noise}}, I) p(H_{\text{noise}}|I)}. \quad (22)$$

If the posterior plausibility ratio is greater than one, H_{signal} is more plausible than H_{noise} and we classify the observation as a detection. If the posterior plausibility ratio is less than one, H_{signal} is less plausible than H_{noise} and we classify the observation as a non-detection.

C. Noise model

The noise distribution is unaffected by the signal parameters,

$$p(\mathbf{e}|\mathbf{h}, \tau, \theta, \phi, H_{\text{signal}}, I) = p(\mathbf{e}|H_{\text{noise}}, I) = p(\mathbf{e}|I). \quad (23)$$

It then follows that

$$p(\mathbf{e}, \mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) = p(\mathbf{e}|I) p(\mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I). \quad (24)$$

The noise that affects gravitational wave detectors is typically modeled as stationary and colored. This can be represented with a *multivariate normal distribution*, which can be compactly represented by the notation

$$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{x}) = \frac{1}{(2\pi)^{N/2} \sqrt{|\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right). \quad (25)$$

The vector $\boldsymbol{\mu}$ is the mean of the distribution, and the positive-definite *covariance matrix* $\boldsymbol{\Sigma}$ describes the ellipsoidal shape of the distribution in terms of the pairwise covariances of the samples,

$$\boldsymbol{\Sigma}_{i,j} = \langle (e_i - \mu_i), (e_j - \mu_j) \rangle. \quad (26)$$

Using this notation, we can define the noise distribution to be

$$p(\mathbf{e}|I) = \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{e}) \quad (27)$$

for some $MN \times MN$ positive definite matrix $\boldsymbol{\Sigma}$. Under the additional assumption of stationarity over some timescale, these covariances can be estimated from previous observations.

In the case of colored noise, each detector is individually represented by a Toeplitz covariance matrix $\boldsymbol{\Sigma}^{(i)}$, and the covariance matrix for the whole network is $\boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\Sigma}^{(1)}, \boldsymbol{\Sigma}^{(2)}, \dots, \boldsymbol{\Sigma}^{(N)})$. In the simple case in which all the noises are white, we have $\boldsymbol{\Sigma} = \text{diag}(\mathbf{I}, \mathbf{I}, \dots, \mathbf{I}) = \mathbf{I}$.

We can now derive the expression for the noise likelihood

$$p(\mathbf{x}|H_{\text{noise}}, I) = \int_{V_{\mathbf{e}}} p(\mathbf{x}|\mathbf{e}, H_{\text{noise}}, I) p(\mathbf{e}|H_{\text{noise}}, I) d\mathbf{e} \quad (28)$$

$$= \int_{V_{\mathbf{e}}} \delta(\mathbf{x} - \mathbf{e}) \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{e}) d\mathbf{e} \quad (29)$$

$$= \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{x}), \quad (30)$$

and substitute the noise model expression into the signal likelihood

$$p(\mathbf{x}|H_{\text{signal}}, I) = \int_{V_{\mathbf{e}, \mathbf{h}, \tau, \theta, \phi}} p(\mathbf{x}|\mathbf{e}, \mathbf{h}, \tau, \theta, \phi, H_{\text{signal}}, I) p(\mathbf{e}, \mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) d\mathbf{e} \dots d\phi \quad (31)$$

$$= \int_{V_{\mathbf{e}, \mathbf{h}, \tau, \theta, \phi}} \delta(\mathbf{x} - \mathbf{F}(\tau, \theta, \phi) \cdot \mathbf{h} - \mathbf{e}) \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{e}) p(\mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) d\mathbf{e} \dots d\phi \quad (32)$$

$$= \int_{V_{\mathbf{h}, \tau, \theta, \phi}} \mathcal{N}(\mathbf{F}(\tau, \theta, \phi) \cdot \mathbf{h}, \boldsymbol{\Sigma}, \mathbf{x}) p(\mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) d\mathbf{h} \dots d\phi, \quad (33)$$

where $V_{\mathbf{h}, \tau, \theta, \phi}$ is the space of all signal parameters.

D. Fast signal model

The remaining integral defining $p(\mathbf{x}|H_{\text{signal}}, I)$ is over many dimensions, and in general will be unsolvable analytically. Without loss of generality we may rewrite the full prior in terms of a prior on direction and arrival time and the prior on the waveform conditional on the direction and the arrival time themselves, i.e.

$$p(\mathbf{h}, \tau, \theta, \phi|H_{\text{signal}}, I) = p(\tau, \theta, \phi|H_{\text{signal}}, I) p(\mathbf{h}|\tau, \theta, \phi, H_{\text{signal}}, I). \quad (34)$$

If we can find a way to analytically marginalize over strain

$$\frac{p(\mathbf{x}|\tau, \theta, \phi, H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \frac{\int_{\mathbb{R}^{2L}} \mathcal{N}(\mathbf{F}(\tau, \theta, \phi) \cdot \mathbf{h}, \boldsymbol{\Sigma}, \mathbf{x}) p(\mathbf{h}|\tau, \theta, \phi, H_{\text{signal}}, I) d\mathbf{h}}{\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{x})} \quad (35)$$

then the remaining integration is over only three dimensions and it can be performed numerically. This in turn allow us to evaluate the Bayes factor,

$$\frac{p(\mathbf{x}|H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \int \int \int p(\tau, \theta, \phi|H_{\text{signal}}, I) \frac{p(\mathbf{x}|\tau, \theta, \phi, H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} d\tau d\theta d\phi. \quad (36)$$

If we choose a multivariate normal distribution for the strain waveform samples, the marginalization integral over \mathbf{h} can be analytically performed. In order to do so, we first define a vector ρ containing a small (or zero) number of parameters of the strain model that might not be marginalizable analytically. These parameters may have to be numerically integrated, like τ , θ , and ϕ , and are potentially computationally very expensive. Their prior distribution will be denoted by $p(\rho|\tau, \theta, \phi, H_{\text{signal}}, I)$. Next, we consider a set of amplitude parameters \mathbf{a} mapped into strain \mathbf{h} via an $2L \times G$ matrix $\mathbf{W}(\rho, \tau, \theta, \phi)$ whose columns $\mathbf{w}_i(\rho, \tau, \theta, \phi)$ are basis waveforms, so that

$$\mathbf{h} = \mathbf{W}(\rho, \tau, \theta, \phi) \cdot \mathbf{a}. \quad (37)$$

The amplitude parameters \mathbf{a} are multivariate normal distributed with a covariance matrix $\mathbf{A}(\rho, \tau, \theta, \phi)$, so that

$$p(\mathbf{a}|\rho, \tau, \theta, \phi, H_{\text{signal}}, I) = \mathcal{N}(\mathbf{0}, \mathbf{A}(\rho, \tau, \theta), \mathbf{a}). \quad (38)$$

The resulting distribution for the waveform strain is

$$p(\mathbf{h}|\tau, \theta, \phi, H_{\text{signal}}, I) = \int_{V_\rho} \int_{\mathbb{R}^G} p(\mathbf{h}|\mathbf{a}, \rho, \tau, \theta, \phi, H_{\text{signal}}, I) p(\mathbf{a}, \rho|\tau, \theta, \phi, H_{\text{signal}}, I) d\mathbf{a} d\rho \quad (39)$$

$$= \int_{V_\rho} \int_{\mathbb{R}^G} \delta(\mathbf{h} - \mathbf{W} \cdot \mathbf{a}) \mathcal{N}(\mathbf{0}, \mathbf{A}, \mathbf{a}) p(\rho|\tau, \theta, \phi, H_{\text{signal}}, I) d\mathbf{a} d\rho, \quad (40)$$

where we have begun to omit the dependence of matrices on their parameters. If $G < 2L$ (i.e., we have fewer basis waveforms than samples in the signal timeseries) the integral over \mathbf{a} cannot be directly represented as a multivariate normal distribution.

This signal model proposes that gravitational wave signals have waveforms that are the sum of G basis waveforms with amplitudes that are normally distributed (and potentially correlated). The basis waveforms and their amplitude distributions may vary with direction, arrival time, and any other parameters we care to include in ρ . The model is capable of representing a variety of sources including the important special cases of known ‘template’ waveforms, and band-limited bursts. We will consider some concrete examples in §II E.

We can substitute the expression back into part of Equation 35 to form a multivariate normal distribution partial integral whose solution is given in [34]:

$$p(\mathbf{x}|\tau, \theta, \phi, H_{\text{signal}}, I) = \int_{V_\rho} \int_{\mathbb{R}^{G+2L}} \mathcal{N}(\mathbf{F} \cdot \mathbf{h}, \boldsymbol{\Sigma}, \mathbf{x}) \delta(\mathbf{h} - \mathbf{W} \cdot \mathbf{a}) \mathcal{N}(\mathbf{0}, \mathbf{A}, \mathbf{a}) p(\rho|\tau, \theta, \phi, H_{\text{signal}}, I) d\mathbf{a} d\rho d\mathbf{h} \quad (41)$$

$$= \int_{V_\rho} \mathcal{N}(\mathbf{0}, (\boldsymbol{\Sigma}^{-1} - \mathbf{K})^{-1}, \mathbf{x}) p(\rho|\tau, \theta, \phi, H_{\text{signal}}, I) d\rho, \quad (42)$$

where the matrix

$$\mathbf{K}(\rho, \tau, \theta, \phi) = (\boldsymbol{\Sigma}^{-1} \mathbf{F} \mathbf{W}) ((\mathbf{F} \mathbf{W})^T \boldsymbol{\Sigma}^{-1} \mathbf{F} \mathbf{W} + \mathbf{A}^{-1})^{-1} (\boldsymbol{\Sigma}^{-1} \mathbf{F} \mathbf{W})^T \quad (43)$$

will be the kernel of our numerical implementation. It is a generalization of equation (8) obtained in the single-sample case. Noting that

$$\frac{p(\mathbf{x}|\rho, \tau, \theta, \phi, H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \frac{\mathcal{N}(\mathbf{0}, (\boldsymbol{\Sigma}^{-1} - \mathbf{K})^{-1}, \mathbf{x})}{\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{x})} = \sqrt{|\mathbf{I} - \boldsymbol{\Sigma}\mathbf{K}|} \exp\left(\frac{1}{2}\mathbf{x}^T \mathbf{K}\mathbf{x}\right), \quad (44)$$

we have

$$\frac{p(\mathbf{x}|\tau, \theta, \phi, H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \int_{V_\rho} p(\rho|\tau, \theta, \phi, H_{\text{signal}}, I) \sqrt{|\mathbf{I} - \boldsymbol{\Sigma}\mathbf{K}|} \exp\left(\frac{1}{2}\mathbf{x}^T \mathbf{K}\mathbf{x}\right) d\rho. \quad (45)$$

The Bayes factor becomes

$$\frac{p(\mathbf{x}|H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} = \int_{V_{\rho, \tau, \theta, \phi}} p(\rho, \tau, \theta, \phi|H_{\text{signal}}, I) \sqrt{|\mathbf{I} - \boldsymbol{\Sigma}\mathbf{K}|} \exp\left(\frac{1}{2}\mathbf{x}^T \mathbf{K}\mathbf{x}\right) d\rho d\tau d\theta d\phi. \quad (46)$$

In other words we have reduced the task of computing the Bayes factor to an integral over arrival time, direction, and any additional signal model parameters ρ .

E. Example signal models

A simple signal model is a burst whose spectrum is white, has characteristic strain amplitude σ (at the Earth) and duration $f_s^{-1}L$

$$G = 2L \quad (47)$$

$$\mathbf{A} = \sigma^2 \mathbf{I} \quad (48)$$

$$\mathbf{W} = \mathbf{I}. \quad (49)$$

If we assert that such bursts are equally likely to come from any direction and arrive at any time in the observation window of $f_s^{-1}M$ seconds, then the priors are

$$p(\theta|H_{\text{signal}}, I) = \sin(\theta) \quad (50)$$

$$p(\phi|H_{\text{signal}}, I) = (2\pi)^{-1} \quad (51)$$

$$p(\tau|H_{\text{signal}}, I) = f_s M^{-1}. \quad (52)$$

The characteristic amplitude σ is an example of a signal model parameter we must numerically marginalize over ($\rho = [\sigma]$). If the source population is distributed uniformly in space at low redshifts, the characteristic strain of a burst, σ , will be power-law distributed,

$$p(\sigma|H_{\text{signal}}, I) = \frac{3\sigma_{\text{min}}^3}{\sigma^4}, \quad (53)$$

where σ_{min} sets a lower bound on the size of gravitational wave we are interested in. This bound is obviously somewhat arbitrary, but is a consequence of the way we distinguish between detection and non-detection. For a uniformly spatially distributed population of bursts there are of course many weak signals within the data, and the noise hypothesis is never true. In reality we are interested only in gravitational waves of at least a certain size. If σ_{min} is much smaller than the noise floor in all detectors, the expression for the noise hypothesis is an excellent approximation to the expressions of the likelihood we adopted. The classification of observations is

insensitive to different choices of σ_{min} below the noise floor.

This is an example of a relatively *uninformative* signal model. It is capable of detecting signals of any waveform (of appropriate duration). However, it incurs a large *Oc-cam penalty* for its generality, and cannot be as sensitive as a more *informed* search.

The other extreme situation is where a source's waveform is completely known, but its other parameters are not. Consider a source that produces a linearly polarized strain \mathbf{w} . If the source's orientation, inclination and amplitude are unknown, we can parameterize the system with two amplitudes \mathbf{a} mapping the strain into the observatory network's polarization basis

$$\mathbf{W} = \begin{bmatrix} \mathbf{w} & \mathbf{0} \\ \mathbf{0} & \mathbf{w} \end{bmatrix}. \quad (54)$$

This is the Bayesian equivalent of the matched filter. The template \mathbf{w} appears twice because any specific signal typically will not be aligned with the polarization basis used to describe h_+ and h_\times in the detectors, but rather will be rotated by some *polarization angle* ψ with respect to that basis. More generally, any signal model that is independent of the observatory network's polarization basis must have \mathbf{A} and \mathbf{W} composed of two identical sub-matrices on the diagonal like this, so that \mathbf{h}_+ and \mathbf{h}_\times are identically distributed. For example, if the source is not linearly polarized, but has strain described by \mathbf{w}_+ and \mathbf{w}_\times , then

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_+ & \mathbf{w}_\times & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{w}_+ & \mathbf{w}_\times \end{bmatrix}. \quad (55)$$

A more general case might be where we have a number of different predictions for a waveform, \mathbf{w}_i , numerically derived. The resulting search looks for a linear combination of these different waveforms,

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1 & \mathbf{w}_2 & \cdots & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{w}_1 & \mathbf{w}_2 & \cdots \end{bmatrix}. \quad (56)$$

Multivariate normal distributions have another useful property. The covariance matrix can be computed for any arbitrary signal model, even though it may not well describe that signal model. The multivariate normal distribution corresponding to that covariance matrix is the *least informative* distribution in that it makes the weakest assumption about the signal of any model with that covariance. Thus, for every signal model, there is a corresponding multivariate normal distribution that will not overstate what we know. It is therefore a conservative choice to replace an arbitrary signal model with the corresponding multivariate normal distribution. This is precisely the approach we follow in the Monte Carlo demonstration in §III, where we use a multivariate normal distribution model to detect binary black-hole merger waveforms.

F. Comparison with previously proposed methods

In this section we will expand on the arguments sketched in a previous paper [36].

Several previously proposed detection statistics, such as the Gürsel-Tinto (i.e. standard likelihood), the constraint likelihoods, and the Tikhonov-regularized likelihood, can be written in the form

$$\mathbf{x}^T \mathbf{J}(\rho, \tau, \theta, \phi) \mathbf{x} > \lambda, \quad (57)$$

where \mathbf{J} is an $MN \times MN$ matrix and λ is a *threshold*. The test is for an explicit direction, arrival time, and any other parameters.

The corresponding Bayesian test is not the marginalization integral we computed, but is the closely related Bayesian test for an explicit direction, arrival time and any other parameters

$$\begin{aligned} & \frac{p(H_{\text{signal}}|\mathbf{x}, \rho, \tau, \theta, \phi, I)}{p(H_{\text{noise}}|\mathbf{x}, I)} \\ &= \frac{p(\mathbf{x}|\rho, \tau, \theta, \phi, H_{\text{signal}}, I) p(H_{\text{signal}}|\rho, \tau, \theta, \phi, I)}{p(\mathbf{x}|H_{\text{noise}}, I) p(H_{\text{noise}}|I)} \quad (58) \\ &= \exp\left(\frac{1}{2} \mathbf{x}^T \mathbf{K} \mathbf{x}\right) \sqrt{|\mathbf{I} - \Sigma \mathbf{K}|} \frac{p(H_{\text{signal}}|\rho, \tau, \theta, \phi, I)}{p(H_{\text{noise}}|I)} \quad (59) \\ &> 1 \quad (60) \end{aligned}$$

which may be rearranged as

$$\begin{aligned} & \mathbf{x}^T \mathbf{K}(\rho, \tau, \theta, \phi) \mathbf{x} \\ & > -2 \ln \left(\sqrt{|\mathbf{I} - \Sigma \mathbf{K}|} \frac{p(H_{\text{signal}}|\rho, \tau, \theta, \phi, I)}{p(H_{\text{noise}}|I)} \right), \quad (61) \end{aligned}$$

suggesting that the previously proposed methods may be equivalent to the Bayesian test for certain choices of prior. When this is possible, it follows that

$$\mathbf{K} = \mathbf{J}, \quad (62)$$

$$\frac{p(H_{\text{signal}}|\rho, \tau, \theta, \phi, I)}{p(H_{\text{noise}}|I)} = \frac{\exp(-\frac{1}{2}\lambda)}{\sqrt{|\mathbf{I} - \Sigma \mathbf{K}|}}. \quad (63)$$

As λ is a constant, but \mathbf{K} , via \mathbf{F} , is a function of direction, (θ, ϕ) , this equation implies that the prior plausibility of the signal hypothesis will vary with direction. The startling implication is that any statistic of the form in Equation 57 proposes an (unphysical) anisotropic distribution of burst sources, as a consequence of neglecting the normalization term outside the exponential.

In order to compare previously proposed statistics to the Bayesian method, we will place some restrictions on the configurations considered. We will assume collocated (but differently oriented) detectors to eliminate the need to time-shift data, and we will use stationary signals and observation times that coincide with the time the signal is present. These restrictions eliminate the differences in the way previously proposed statistics and the Bayesian method handle arrival time and signal duration. For simplicity, we will further assume that the detectors are affected by white noise.

1. Tikhonov regularized statistic

The Tikhonov regularized statistic proposed in [27] for white noise interferometers can be written in the form

$$\mathbf{x}^T \mathbf{F}(\mathbf{F}^T \mathbf{F} + \alpha^2 \mathbf{I})^{-1} \mathbf{F}^T \mathbf{x} > \lambda. \quad (64)$$

The Bayesian kernel reduces to this for

$$\Sigma = \mathbf{I} \quad (65)$$

$$\mathbf{W} = \mathbf{I} \quad (66)$$

$$\mathbf{A} = \alpha^{-2} \mathbf{I}. \quad (67)$$

This is a white noise network observing a broad-band signal of characteristic amplitude $\sigma = \alpha^{-1}$. The Tikhonov regularizer α therefore places a delta function as the prior for the size of the signal,

$$p(\sigma|I) = \delta(\sigma - \alpha^{-1}). \quad (68)$$

This physical interpretation of the regularizer was not made in [27]. The prior plausibility of the signal hypothesis varies with direction

$$\begin{aligned} \frac{p(H_{\text{signal}}|\theta, \phi, I)}{p(H_{\text{noise}}|I)} &= \frac{\exp(-\frac{1}{2}\lambda)}{\sqrt{|\mathbf{I} - \mathbf{F}(\mathbf{F}^T \mathbf{F} + \alpha^2 \mathbf{I})^{-1} \mathbf{F}^T|}} \\ &= \sqrt{e^{-\lambda} |\mathbf{I} + \sigma^2 \mathbf{F}^T \mathbf{F}|}, \quad (69) \end{aligned}$$

implying that gravitational wave events to come from some directions more frequently and from others less frequently; in particular, this is in proportion to the network's sensitivity to that direction.

The Tikhonov statistic behaves like a Bayesian statistic that postulates all bursts have the same characteristic strain amplitude at the Earth and are unequally distributed across the sky. Since these priors do not reflect our knowledge of the universe, we should expect a better performance from an analysis accounting for our prior knowledge.

2. Gürsel-Tinto statistic

The Gürsel-Tinto or standard likelihood statistic [7, 22, 25] for white noise interferometers can be written in the form

$$\mathbf{x}^T \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{x} > \lambda. \quad (70)$$

For large σ , the Tikhonov statistic goes to

$$\mathbf{K} \approx \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T, \quad (71)$$

$$\frac{p(H_{\text{signal}}|\theta, \phi, I)}{p(H_{\text{noise}}|I)} \approx \sigma^{2M} \sqrt{e^{-\lambda} |\mathbf{F}^T \mathbf{F}|}. \quad (72)$$

This implies that there is no Bayesian test equivalent to the Gürsel-Tinto statistic, but that the Gürsel-Tinto statistic is rather the limit of a series of Bayesian tests for increasingly large, increasingly frequent, and increasingly directionally biased populations of gravitational wave signals.

Alternatively, we could say that the Gürsel-Tinto statistic follows from our Bayesian formulation if we adopt an *improper* (unnormalizable) strain prior $p(\mathbf{h}|H_{\text{signal}}, I) = 1$, which assigns equal plausibility to every waveform. This is what is meant by Gürsel-Tinto's independence of waveform. In practice we expect that smaller waveforms occur more frequently than larger waveforms, and this has real consequences in the analysis.

Consider a common failure mode of Gürsel-Tinto: misidentifying the direction of a gravitational wave. A moderately sized signal will come from a direction of typical sensitivity and produce a moderate response. Gürsel-Tinto will correctly declare the true direction of the injection to be plausible. However, there are directions on the sky where the global network becomes insensitive to one polarization, and near those directions $\mathbf{F}^T \mathbf{F}$ is a near-singular matrix whose inverse varies rapidly, causing the Gürsel-Tinto statistic itself to vary over a wide range. Often one of these near-pathological directions will be deemed more plausible than the true direction. These pathological directions always correspond to very low sensitivity to at least one polarization, so to explain the moderately-sized response at least one polarization of the postulated signal has to be very large. The improper Gürsel-Tinto prior says we believe very large gravitational waves to be just as plausible as moderately sized ones, so the statistic has no grounds to discount the pathological direction, and returns the wrong direction and an obviously wrong unphysically large reconstructed waveform. By contrast, the Bayesian method's prior can tell it that very large signals are very unlikely, so it does not make the same error.

This shortcoming also degrades the sensitivity of the Gürsel-Tinto statistic. Noisy observations can be consistent with a very large gravitational wave from a near-pathological direction and these false alarms force us to use a high threshold λ that limits the efficiency of the method for more physically reasonable signals. Again,

the Bayesian method does not suffer from this problem because it knows large signals are rare and it will not postulate them on the basis of weak evidence.

It is important to note that Gürsel-Tinto can detect realistically-sized signals; they form only an infinitesimal fraction of the postulated signal population, but this is canceled by the infinite prior plausibility that a signal is present, to give them a net finite plausibility. Like the Tikhonov statistic, Gürsel-Tinto works as a detection statistic; its problem is only one of efficiency.

3. Soft constraint likelihood

The soft constraint statistic [26, 31] for white noise interferometers can be written in the form

$$k^2(\theta, \phi) \mathbf{x}^T \mathbf{F} \mathbf{F}^T \mathbf{x} > \lambda, \quad (73)$$

for some function $k(\theta, \phi)$. Specifically, (73) gives the soft constraint likelihood for the choice $k^2 = (\mathbf{F}^{+T} \mathbf{F}^+)^{-1}$, where the antenna response is computed in the dominant polarization frame [26].

Consider the signal model defined by

$$\mathbf{\Sigma} = \mathbf{I} \quad (74)$$

$$\mathbf{W} = \mathbf{I} \quad (75)$$

$$\mathbf{A} = \sigma^2 k^2(\theta, \phi) \mathbf{I}. \quad (76)$$

This is a system of uncorrelated white noise detectors and a population of broad-band signals whose characteristic amplitude $\sigma k(\theta, \phi)$ varies as some known function of direction, slightly generalizing the Tikhonov statistic. For small σ ,

$$\mathbf{K} \approx \sigma^2 k^2(\theta, \phi) \mathbf{F} \mathbf{F}^T, \quad (77)$$

and the Bayesian test becomes

$$k^2(\theta, \phi) \mathbf{x}^T \mathbf{F} \mathbf{F}^T \mathbf{x} > -\frac{2}{\sigma^2} \ln \left(\sqrt{|\mathbf{I} - \sigma^2 k^2 \mathbf{F} \mathbf{F}^T|} \frac{p(H_{\text{signal}}|\theta, \phi, I)}{p(H_{\text{noise}}|I)} \right). \quad (78)$$

This implies that to mimic the soft constraint we must choose the prior

$$\frac{p(H_{\text{signal}}|\theta, \phi, I)}{p(H_{\text{noise}}|I)} \approx 1 + \frac{\sigma^2}{2} (k^2(\theta, \phi) \text{tr}(\mathbf{F} \mathbf{F}^T) - \lambda) + O(\sigma^4). \quad (79)$$

The prior is approximately unity, varying only infinitesimally with direction and the threshold λ . Even though the prior's dependence on direction is weak, the statistic's dependence on the data is equally weak. Within these assumptions equation (44) becomes

$$\frac{p(\mathbf{x}|\theta, \phi, H_{\text{signal}}, I)}{p(\mathbf{x}|H_{\text{noise}}, I)} \approx 1 + \frac{1}{2} \sigma^2 k^2(\theta, \phi) (\mathbf{x}^T \mathbf{F} \mathbf{F}^T \mathbf{x} - \text{tr}(\mathbf{F} \mathbf{F}^T)) + O(\sigma^4). \quad (80)$$

Since the expected signals are weak, and the evidence for them will also be weak, the weak prior still strongly affects the result.

The soft constraint is therefore the limit of a series of Bayesian tests for gravitational wave bursts that are common and whose amplitudes and rates are a function of direction on the sky. If we choose $k(\theta, \phi) = 1$ we can eliminate the directional amplitude bias; if we choose $k^{-2}(\theta, \phi) = \text{tr}(\mathbf{F}\mathbf{F}^T)$ we can eliminate the directional rate bias. Other choices, such as that made in [26], remove neither bias.

4. Hard constraint likelihood

Let us restrict the soft-constraint signal model to a population of *linearly polarized* signals with a known polarization angle $\psi(\theta, \phi)$ for each direction

$$\mathbf{\Sigma} = \mathbf{I} \quad (81)$$

$$\mathbf{W} = \begin{bmatrix} \cos 2\psi(\theta, \phi)\mathbf{I} \\ \sin 2\psi(\theta, \phi)\mathbf{I} \end{bmatrix} \quad (82)$$

$$\mathbf{A} = \sigma^2 k^2(\theta, \phi)\mathbf{I}. \quad (83)$$

Then with the prior

$$\frac{p(H_{\text{signal}}|\theta, \phi, I)}{p(H_{\text{noise}}|I)} = e^{-\frac{\lambda\sigma^2}{2}} \sqrt{|\mathbf{I} + \sigma^2 k^2(\mathbf{F}\mathbf{W})^T \mathbf{F}\mathbf{W}|} \quad (84)$$

for $\sigma \rightarrow 0$ the Bayesian statistic limits to

$$k^2(\theta, \phi) \mathbf{x}^T \mathbf{F}\mathbf{W}(\mathbf{F}\mathbf{W})^T \mathbf{x} > \lambda. \quad (85)$$

For the particular choice of $\psi(\theta, \phi)$ being the rotation angle between the detector polarization basis and the dominant polarization frame, and $k^2 = (\mathbf{F}\mathbf{W})^T \mathbf{F}\mathbf{W}$ (which is equal to $(\mathbf{F}^{+T} \mathbf{F}^+)^{-1}$ in the dominant polarization frame [26]), this yields the hard constraint statistic of [26].

In addition to the explicit assumptions that all signals are optimally oriented and linearly polarized, the hard constraint suffers from the same implicit problems as the soft constraint. The normalization chosen in [26] eliminates the directional rate bias but introduces a directional amplitude bias.

G. Implications

The Tikhonov, Gürsel-Tinto, and constraint methods are limits of Bayesian statistics distinguished only by different choices of prior. Yet, as Bayesian priors are *subjective*, on what basis can we critique them?

Though subjective, priors make definite statements about our *expectations* of the world; one popular paradigm relates priors to bets we would be willing to make about the outcome of an experiment. Few scientists would be willing to bet that the first detected gravitational wave burst would have strain many orders




FIG. 1: Four statistics as a function of (θ, ϕ) for identical white noise interferometers with the locations and orientations of LHO, LLO and Virgo sampled at 1024 Hz and a 1/16s white noise signal with amplitude SNR of 5. White is most plausible; black is least plausible; a circle indicates the true direction; a square indicates the most plausible direction. From top to bottom: ln Bayesian odds ratio for $\sigma = 5$; Tikhonov for $\alpha = 0.2$; Gürsel-Tinto; soft constraint with $k(\theta, \phi) = 1$.

of magnitude above (Gürsel-Tinto) or below (constraint methods) the instrumental noise. Insofar as we find these prior plausibility distributions incredible, we should expect a Bayesian method with a more credible prior to be a more *efficient* detection statistic. Quite literally, the Bayesian method does not have to waste precious

(signal) energy overcoming strong prejudices. The same logic applies to attempts to reconstruct the parameters of a detected signal, such as the sky position of the source; Fig. 1 shows an example. The above conclusions apply equally to different versions of these statistics for colored noise or different bases other than the time-domain (such as the frequency or wavelet domains).

Once may also consider what the Bayesian analysis says about how the frequentist statistics can be improved. For example, in our comparisons we have assumed that the threshold λ used in the frequentist analysis is the same for all sky positions. The Bayesian analysis shows that this implicitly imposes physically unreasonable priors on the gravitational-wave signal. However, in a frequentist analysis we are perfectly free to make λ vary with sky position. The Bayesian analysis shows how the threshold can be varied across the sky to correspond to physical priors; for example, by selecting $\lambda(\theta, \phi)$ in equation (63) to make the right-hand side constant. Analogous reasoning from the frequentist point of view would be to examine how $\lambda(\theta, \phi)$ should be chosen to balance (in some sense) the false alarm rate versus detection probability across the sky. The Bayesian formulation gives one concrete suggestion.

In summary, we have demonstrated that previously proposed methods are optimal for unreasonable choices of prior, and consequently suboptimal for reasonable choices of prior. We have not yet quantified how much worse their performance is. One way to answer this question is to perform a Monte-Carlo simulation, testing the ability of each statistic to detect thousands of simulated gravitational wave signals.

III. SIMULATIONS

We used the X-PIPELINE software package [37] to characterize the relative performance of the Gürsel-Tinto (i.e. standard likelihood), soft constraint, hard constraint, and Bayesian methods.

Our tests used a set of 4 identical detectors at the positions and orientations of the LIGO-Hanford, LIGO-Livingston, GEO 600, and Virgo detectors. The data was simulated as Gaussian noise with spectrum following the design sensitivity curve of the 4-km LIGO detectors; it was taken from a standard archive of simulated data [38] used for testing detection algorithms. Approximately 12 hours of data in total was analysed for these tests.

For the population of gravitational-wave signals to be detected we chose, somewhat arbitrarily, the ‘‘Lazarus’’ waveforms of Baker *et al.* [39]. These are fairly simple waveforms generated from numerical simulations of the merger and ringdown of a binary black-hole system. We chose to simulate a pair of 20 solar-mass black holes, which puts the peak of signal power near the frequencies of best sensitivity for LIGO. The time-series waveforms are shown in Fig. 2, while the spectra and detector noise curve are given in Fig. 3. The sources were

placed at the discrete distances $240/S$ Mpc [43], where $S = 1, 2, 2.5, 3, 10$, and with randomly chosen sky position and orientation. Approximately 5000 injections were performed for each distance.




FIG. 2: Time-series Lazarus waveforms [39] used for our simulations, from a nominal distance of 240 Mpc.




FIG. 3: Strain-equivalent noise amplitude spectral density of the simulated data used in our tests (black) with spectra of the Lazarus waveforms [39] at 240 Mpc. The Lazarus spectra have been rescaled by $T^{-1/2} = (128/\text{sec})^{1/2}$ to render them into the same units as the noise spectrum.

For the Bayesian statistic we adapt the broad-band signal prior (47)–(49) for each polarization. Specifically, we assume the simple model of a burst whose spectrum is white, with characteristic strain amplitude σ at the Earth and duration equal to our chosen FFT length:

$$G = 2M \quad (86)$$

$$\mathbf{A} = \sigma^2 \mathbf{I} \quad (87)$$

$$\mathbf{W} = \mathbf{I}. \quad (88)$$

We use a uniform prior on the signal arrival time τ , and an isotropic prior on the sky position (θ, ϕ) . The Bayesian statistic was computed for approximately logarithmically spaced discrete values of characteristic strain

$\sigma = 10^{-23}, 3 \times 10^{-23}, 10^{-22}, 3 \times 10^{-22}, 10^{-21}$, and averaged together in post-processing. This averaging approximates a single Bayesian statistic with a Jeffreys (scale invariant) prior $p(\sigma) \propto 1/\sigma$ between 10^{-23} and 10^{-21} . (Performing the combination in post-processing allowed us to maintain compatibility with the existing architecture of X-PIPELINE.)

Each likelihood statistic was computed over a fixed frequency band of [64,1088] Hz, with an FFT length of 1/128 sec. Consecutive time bins were overlapped by 75%. The detection probability as a function of false alarm probability is shown in Fig. 4. The distance used for the Lazarus simulations for this figure was $240/2.5 = 96$ Mpc; injections at other distances yielded similar results. At this distance, the total SNR deposited in the network ($\sum_{\alpha} \rho_{\alpha}^2$) was in the range $\sim 1 - 8$ with a mean value of 5, where

$$\sum_{\alpha} \rho_{\alpha}^2 = \sum_{\alpha} 4 \int_0^{\infty} df \frac{|F_{\alpha}^+ \tilde{h}_+(f) + F_{\alpha}^{\times} \tilde{h}_{\times}(f)|^2}{S(f)} \quad (89)$$

and $S(f)$ is the one-sided noise power spectral density of each interferometer.

As can be seen from Fig. 4, the best performance is achieved by the Bayesian method with the σ value most closely matching the injected signals, with the marginalized curve performing almost identically. The detection probability of the marginalized Bayesian method is significantly better than that of any of the non-Bayesian methods (standard likelihood, soft constraint, and hard constraint likelihoods) over the full range of false-alarm probabilities tested.

For a given false-alarm probability, we may compute the distance at which each likelihood achieves 50% efficiency by fitting a sigmoid curve to the simulations. The observed volume, and therefore the expected rate of detections for a uniformly distributed source population, scales as the cube of the distance. We computed the distance and volume for each statistic for two false alarm probabilities, 10^{-5} in Table I and $1/256 \approx 3.9 \times 10^{-3}$ in Table II. As we compute 512 statistics per second, these correspond to false alarm rates of 1/200 Hz and 2 Hz respectively (as the statistics are computed on 75% overlapped data, these estimates are conservative). The first rate is comparable to that used in triggered searches for gravitational waves from gamma-ray bursts [5, 40, 41], where the upper limit is set based on the largest event in a 180 s window around the trigger time. The second rate is typical for event generation at the first stage of an untriggered (all-sky, all-time) burst search [42]. At both of these false alarm probabilities, the Bayesian method can detect sources approximately 15% more distant, and consequently has an observed volume and expected detection rate approximately 50% greater, than the non-Bayesian statistics.

It is important to note that we do *not* use detailed knowledge of the signal waveform for the Bayesian analysis. Our prior is that the signal spectrum is flat over

TABLE I: Distances for false-alarm probability 10^{-5}

Statistic	Distance (Mpc)	Distance (rel.)	Volume (rel.)
Standard	72.1	1.01	1.02
Soft	71.6	1.00	1.00
Hard	72.5	1.01	1.04
Bayesian	82.0	1.15	1.50

the analysis band ([64,1088] Hz), and by imposing no phase structure or sample-to-sample correlations we are assuming that the timeseries is white noise with fixed rms amplitude over the integration time (1/128 sec). This is an example of a “least informative” signal prior, as discussed in §II E. Considering Figures 2 and 3, it is clear that these priors are not particularly accurate models for the actual gravitational-wave signal. Nevertheless, our Monte Carlo results demonstrate that even this partial and *approximate* prior information is better than the implicit priors of the Gürsel-Tinto/standard and constraint likelihoods, and improves the sensitivity of the search.




FIG. 4: Receiver-operating characteristic (ROC) curves for the Bayesian, standard likelihood / Gürsel-Tinto, soft constraint, and hard constraint likelihoods for sources at 96 Mpc. The σ^{-1} prior curve is obtained by marginalizing over probabilities associated with the discrete σ values tested. The best performance is achieved by the σ value most closely matching the injected signals, with the marginalized curve performing almost identically. The detection probability of the marginalized Bayesian method is significantly better than that of any of the non-Bayesian methods (standard/Gürsel-Tinto, soft constraint, and hard constraint likelihoods) over the full range of false-alarm probabilities tested.

IV. CONCLUSIONS AND FUTURE DIRECTIONS

We have presented a comprehensive Bayesian formulation of the problem of detecting gravitational wave bursts with a network of ground-based interferometers. We

TABLE II: Distances for false-alarm probability 1/256

Statistic	Distance (Mpc)	Distance (rel.)	Volume (rel.)
Standard	87.1	1.03	1.08
Soft	84.9	1.00	1.00
Hard	86.9	1.02	1.07
Bayesian	97.9	1.15	1.53

demonstrated how to systematically incorporate partial or approximate prior information into the analysis, such as time-frequency or polarization content, and source distributions. We have also seen that this Bayesian formulation contains several previously proposed detection statistics as special cases.

The Bayesian methodology we have derived to address the problem of detecting poorly-understood gravitational wave bursts yields a novel statistic. On theoretical grounds we expect this statistic to outperform previously proposed statistics. A Monte-Carlo analysis confirms this expectation: over a range of false alarm rates, the Bayesian statistic can detect sources at 15% greater distances and therefore observe 50% more events.

The Bayesian search requires explicitly adopting a model for the poorly understood signal. This is not a shortcoming. The model may be agnostic with respect to many features of the waveform. As we have demonstrated, existing methods are not free from their own signal models, but implicitly assume priors on the waveform properties and source distributions. We have seen that these implicit priors typically correspond to physically implausible assumptions about the gravitational wave burst population, which must necessarily impair the search sensitivity.

Coherent analyses of any kind are relatively costly, and efficient implementations must be sought. By specifying the problem as an integral, the Bayesian approach lets us leverage the extensive literature on numerical integra-

tion for techniques to accelerate the computation; one promising contender is *importance sampling*.

As second practical issue that must be dealt with is that the background noise of real gravitational wave detectors contains transient non-Gaussian features (“glitches”). As in the case of detection statistics, several *ad hoc* frequentist statistics have been proposed to distinguish glitches from gravitational wave signals (see for example [24, 30]). Again, Bayesian methodology provides us with a direction in which to proceed: augment the noise model to better reflect “glitchy” reality, and to the extent we are successful, robustness will automatically follow.

Acknowledgments

We would like to thank Shourov Chatterji, Albert Lazzarini, Soumya Mohanty, Andrew Moylan, Malik Rakhmanov, and Graham Woan for useful discussions and valuable comments on the manuscript. This work was performed under partial funding from the following NSF Grants: PHY-0107417, 0140369, 0239735, 0244902, 0300609, and INT-0138459. A. Searle was supported by the Australian Research Council and the LIGO Visitors Program. For M. Tinto, the research was also performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. M. Tinto was supported under research task 05-BEFS05-0014. P. Sutton was supported in part by STFC grant PP/F001096/1. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0107417. This document has been assigned LIGO Laboratory document number LIGO-P070114-01-Z.

-
- [1] <http://www.geo600.uni-hannover.de/>.
 - [2] <http://www.ligo.caltech.edu/>.
 - [3] <http://tamago.mtk.nao.ac.jp/>.
 - [4] <http://www.virgo.infn.it/>.
 - [5] B. Abbott et al., *Astrophys. J.* **681**, 1419 (2008), [astro-ph/0711.1163](https://arxiv.org/abs/astro-ph/0711.1163).
 - [6] E. E. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4535 (1998).
 - [7] E. E. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4566 (1998).
 - [8] F. Pretorius, *Phys. Rev. Lett.* **95**, 121101 (2005).
 - [9] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, *Phys. Rev. Lett.* **96**, 111102 (2006).
 - [10] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, *Phys. Rev. Lett.* **96**, 111101 (2006).
 - [11] P. Diener, F. Herrmann, D. Pollney, E. Schnetter, E. Seidel, R. Takahashi, J. Thornburg, and J. Ventrella, *Phys. Rev. Lett.* **96**, 121101 (2006).
 - [12] F. Herrmann, D. Shoemaker, and P. Laguna (2006), [gr-qc/0601026](https://arxiv.org/abs/gr-qc/0601026).
 - [13] F. Löffler, L. Rezzolla, and M. Ansorg (2006), [gr-qc/0606104](https://arxiv.org/abs/gr-qc/0606104).
 - [14] M. Shibata and K. Taniguchi, *Phys. Rev. D* **73**, 064027 (2006).
 - [15] T. Zwerger and E. Müller, *Astron. Astrophys.* **320**, 209 (1997).
 - [16] H. Dimmelmeier, J. Font, and E. Müller, *Astron. Astrophys.* **393**, 523 (2002).
 - [17] C. Ott, A. Burrows, E. Livne, and R. Walder, *Astrophys. J.* **600**, 834 (2004).
 - [18] M. Shibata and Y. I. Sekiguchi, *Phys. Rev. D* **69**, 084024 (2004).
 - [19] C. D. Ott, A. Burrows, L. Dessart, and E. Livne, *Phys. Rev. Lett.* **96**, 201102 (2006).

- [20] P. Mészáros, *Ann. Rev. Astron. Astrophys.* **40**, 137 (2002).
- [21] C. Cutler and K. S. Thorne (2002), gr-qc/0204090.
- [22] Y. Gursel and M. Tinto, *Phys. Rev. D* **40**, 3884 (1989).
- [23] M. Tinto, in *Proceedings of the International Conference on Gravitational Waves: Source and Detectors*. (World Scientific (Singapore), 1996).
- [24] S. Chatterji, A. Lazzarini, L. Stein, P. Sutton, A. Searle, and M. Tinto, *Phys. Rev. D* **74**, 082005 (2006).
- [25] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and E. E. Flanagan, *Phys. Rev. D* **63**, 042003 (2001).
- [26] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, *Phys. Rev. D* **72**, 122002 (2005).
- [27] M. Rakhmanov, *Class. Quant. Grav.* **23**, S673 (2006), gr-qc/0604005.
- [28] S. Mohanty, M. Rakhmanov, S. Klimenko, and G. Mitselmakher, *Class. Quant. Grav.* **23**, 4799 (2006), gr-qc/0601076.
- [29] T. Z. Summerscales, A. Burrows, L. S. Finn, and C. D. Ott, *The Astrophysical Journal* **678**, 1142 (2008), arXiv:0704.2157.
- [30] L. Wen and B. Schutz, *Class. Quantum Grav.* **22**, S1321 (2005).
- [31] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, *J. Phys. Conf. Ser.* **32**, 12 (2006).
- [32] J. Sylvestre, *Phys. Rev. D* **68**, 102005 (2003).
- [33] N. Arnaud, M. Barsuglia, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, S. Kreckelbergh, and E. K. Porter, *Phys. Rev. D* **68**, 102001 (2003).
- [34] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- [35] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences* (Cambridge University Press, 2005).
- [36] A. C. Searle, P. J. Sutton, M. Tinto, and G. Woan, *Class. Quant. Grav.* **25**, 114038 (2008).
- [37] <https://geco.phys.columbia.edu/xpipeline/wiki>.
- [38] F. Beauville et al., *J.Phys.Conf.Ser.* **32**, 212 (2006), gr-qc/0509041.
- [39] J. Baker, M. Campanelli, C. O. Lousto, and R. Takahashi, *Phys. Rev. D* **65**, 124012 (2002).
- [40] B. Abbott et al., *Phys. Rev. D* **72**, 042002 (2005), gr-qc/0501068.
- [41] B. Abbott et al., *Phys. Rev. D* **77**, 062004 (2008), arXiv:0709.0766.
- [42] B. Abbott et al., *Class. Quantum Grav.* **24**, 5343 (2007), arXiv:0704.0943.
- [43] The fiducial distance 240 Mpc is chosen for numerical convenience; it is the distance at which the sum-squared matched-filter SNR for each polarization is 1/2, assuming optimal antenna response (F^+ , $F^\times = 1$).