

Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars

Karl Wette

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Callinstraße 38, D-30167 Hannover, Germany*

This paper presents an in-depth study of how to estimate the sensitivity of searches for gravitational-wave pulsars – rapidly-rotating neutron stars which emit quasi-sinusoidal gravitational waves. It is particularly concerned with searches over a wide range of possible source parameters, such as searches over the entire sky and broad frequency bands. Traditional approaches to estimating the sensitivity of such searches use either computationally-expensive Monte Carlo simulations, or analytic methods which sacrifice accuracy by making an unphysical assumption about the population of sources being searched for. This paper develops a new, analytic method of estimating search sensitivity which does not rely upon this unphysical assumption. Unlike previous analytic methods, the new method accurately predicts the sensitivity obtained using Monte Carlo simulations, while avoiding their computational expense. The change in estimated sensitivity due to properties of the search template bank, and the geographic configuration of the gravitational wave detector network, are also investigated.

PACS numbers: 02.50.Tt, 04.80.Nn, 95.55.Ym, 97.60.Jd

I. INTRODUCTION

The last decade saw the successful construction and operation of the first generation of large-scale gravitational-wave interferometric detectors, at the observatories of LIGO [1] in the United States, and of Virgo [2] and GEO 600 [3] in Europe. Ongoing improvements to the sensitivities of these detectors (e.g. [4]), combined with the construction of new large-scale gravitational-wave interferometers in Japan [5] and potentially elsewhere, are widely anticipated to result in the first ground-based detection of gravitational waves within the next decade. From these first detections will follow new tests of the fundamental physics of gravity, as well as new tools for exploring the astrophysics of compact objects.

Gravitational-wave pulsars are one class of sources which may be detected by the next generation of gravitational-wave interferometers. They are rapidly-rotating neutron stars which emit long-lived, narrow-band, quasi-sinusoidal gravitational waves, and are often referred to as *continuous* or *periodic* sources. The dominant gravitational-wave emission is expected to be due to non-axisymmetric distortions of the neutron star; other possible emission mechanisms are unstable oscillation modes such as *r*-modes, and free precession due to misaligned rotation and symmetry axes; see [6] for a review. An isolated neutron star may have acquired a non-axisymmetric deformation during birth in a supernova, although it is uncertain for how long such a deformation might be retained. For neutron stars in binary systems, non-axisymmetry may result from e.g. differential heating from accreted matter leading to differential density gradients [7], or from the confinement of the accreted matter by the star's magnetic field in a magnetic

mountain [8, 9].

Signals from gravitational-wave pulsars are being actively searched for in data from the LIGO and Virgo detectors. Although none have been found to date, many of these searches have reached sensitivities where (albeit under optimistic conditions) such signals could be detected. The most recent search for gravitational waves from known radio- and X-ray pulsars [10] constrained the gravitational-wave power radiated by the Crab pulsar to be less than 2% of the total power available from the loss of rotational energy. A search for gravitational waves from the Vela pulsar [11] set energy constraints of 35–45%, depending on assumptions about Vela's orientation. Upper limits on the amplitude of gravitational waves from the neutron star in the supernova remnant Cassiopeia A, set in [12], are below the level expected (optimistically) from the total conversion of rotational energy into gravitational waves, assuming that the (as-yet unknown) rotation period of the neutron star is within the searched frequency band. Recent searches for undiscovered neutron stars which may be radiating gravitational waves [13, 14] have set gravitational-wave amplitude limits comparable to upper limits hypothesized for a population of such stars [15]. Searches for gravitational-wave pulsars in binary systems have so far focused on the most promising known target, the low-mass X-ray binary Scorpius X-1 [16, 17].

To best assess the prospects of future searches for gravitational-wave pulsars, it is important to be able to accurately estimate the sensitivity such searches can achieve. (What is meant here by *sensitivity* is defined in Section II C.) In particular, designing searches for continuous gravitational waves which cover a wide range of possible signal parameters (e.g. searches for undiscovered neutron stars) commonly requires constructing a hierarchical pipeline comprised of different data analysis techniques, each with different trade-offs, such as better sensitivity but increased computational cost, or vice versa.

* karl.wette@aei.mpg.de

An accurate estimate of the overall sensitivity of such a pipeline is important, therefore, for identifying the optimal combination of its elements.

Obtaining the sensitivity of a search targeting a *single* source, such as a known pulsar, is relatively straightforward (e.g. [6]). The calculation becomes more difficult, however, for searches over wide signal parameter spaces, e.g. searches for undiscovered neutron stars, or searches targeting known neutron stars with unknown rotation periods. This difficulty has resulted in two different approaches to sensitivity estimation. Wide-parameter-space searches of LIGO and Virgo data [12–14, 16, 18–22] set upper limits on gravitational-wave amplitude (which in turn characterize the sensitivity of the search) by performing Monte Carlo simulations, where the search is re-performed on computer-generated data containing simulated signals (see Section II C). While Monte Carlo simulations are appropriate for accurately computing the sensitivity of searches of real gravitational-wave detector data, which often contain e.g. non-Gaussian instrumental noise artifacts, they are usually too computationally expensive to be useful for theoretical studies of the sensitivities of different data analysis techniques.

Instead, theoretical studies, e.g. [23–31], commonly make certain assumptions about the distribution of the gravitational-wave signals being searched for, in order to simplify the sensitivity calculation (see Section III C). These assumptions, however, result in a measure of sensitivity that, as shown in Section III D, is quantitatively different from that arrived at using Monte Carlo simulations. To date, there has been little published work in the gravitational-wave literature on the discrepancy between these two approaches. Furthermore, an accurate, computationally cheap (i.e. suitable for theoretical studies) estimator of sensitivity, as obtained by Monte Carlo simulations, has yet to be proposed (although combinations of numerical and analytic sensitivity estimation methods have been developed; see [32, 33]). It is these two issues that the present work intends to address.

Section II of this paper presents an overview of gravitational-wave pulsar searches, and describes the most common method by which their sensitivities are estimated. Section III contains the main result of this paper: an analytic expression which may be used to quickly and accurately estimate the sensitivity of wide-parameter-space searches for gravitational-wave pulsars. Section IV verifies the accuracy of the analytic sensitivity estimator, and Section V assesses the validity of some assumptions that were made during its derivation. Section VI discusses the results presented in this paper, as well as possible avenues for future research.

II. GRAVITATIONAL-WAVE PULSAR SEARCHES

This section is an overview of the signal model of gravitational-wave pulsars (Section II A), the data analy-

sis techniques used to search for them (Section II B), and the method by which search sensitivities are estimated (Section II C). See also [6] for a review of gravitational-wave pulsar data analysis, and [34] for an overview of the current data-analysis activities of the LIGO and Virgo scientific collaborations.

A. Signal model

The signal from a gravitational-wave pulsar is written as a time series $h(t)$ of the dimensionless strain amplitude h , which for a ground-based interferometric detector is proportional to the differential change in the length of its arms (which are assumed to be much shorter than the gravitational wave-length). We assume that the signal contains only a single frequency component, although it is possible for it to contain multiple frequency components arising from free precession [35]. Following [23], the time series $h(t)$ may be written as the summed products of four amplitudes A_i and four time-dependent functions $h_i(t)$:

$$h(t) = \sum_{i=1}^4 A_i h_i(t) \quad (2.1)$$

The amplitudes A_i are related to the four *amplitude parameters* of the signal: its overall strain amplitude h_0 ; its initial phase ϕ_0 ; the inclination angle ι between the neutron star angular momentum and wave propagation vectors; and the polarization angle ψ , which fixes the orientation of the neutron star about the wave propagation vector. The functions $h_i(t)$ are functions of the signal's remaining *phase parameters*: its sky position, given by its right ascension α and declination δ ; and its frequency evolution, given by an initial frequency f , and frequency time-derivatives or *spindowns* \dot{f} , \ddot{f} , $f^{(3)}$, etc. The number of spindowns required generally depends on the age of the sources being targeted [36–38].

The time series $h(t)$ may also be written in a form which illustrates the two polarizations, *plus* and *cross*, of a gravitational wave:

$$h(t) = A_+ F_+(t) \cos \Phi(t) + A_\times F_\times(t) \sin \Phi(t), \quad (2.2)$$

where A_+ and A_\times are the amplitudes of their respective polarizations, and $\Phi(t)$ is the signal phase. The *antenna-pattern* functions $F_+(t)$ and $F_\times(t)$ give the response of the detector to each polarization, and are modulated by the sidereal motion of the Earth. Expressions for $F_+(t)$ and $F_\times(t)$ are given in [23, 39–42], and in Appendix A.

The signal-to-noise ratio (SNR) of a signal, ρ , is found by integrating $h(t)$ over the observation time T , which gives:

$$\rho^2 = \frac{h_0^2 T}{S_h} (a_+^2 \langle F_+^2 \rangle_t + a_\times^2 \langle F_\times^2 \rangle_t), \quad (2.3)$$

where $\langle F_+^2 \rangle_t = (1/T) \int_{-T/2}^{T/2} dt F(t)^2$ is the time average of $F_+(t)^2$ (similarly for $\langle F_\times^2 \rangle_t$), $a_+ = A_+/h_0$ (similarly

for a_\times), and S_h is the one-sided detector noise power spectral density. The SNR is independent of the signal's phase modulation; when T is longer than several days, the SNR's dependence on the signal's amplitude modulation also vanishes, and ρ^2 becomes a linear function of T . Expressions for $\langle F_+^2 \rangle_t$ and $\langle F_\times^2 \rangle_t$ are given in Appendix A; see also [23] for expressions for ρ^2 .

B. Search techniques

Gravitational-wave pulsars are very weak sources; the amplitude of their signals is likely to be several orders of magnitude smaller than the noise amplitude of current- and even next-generation interferometric detectors. Nevertheless, gravitational-wave pulsar signals can be recovered using the well-known technique of *matched filtering*, where the data are correlated against a template which models the signals' amplitude and phase evolution over time. Matched filtering was first applied to the detection of gravitational-wave pulsars in [23], and extended to multiple detectors in [43]. The detection statistic derived in these papers, known as the \mathcal{F} -statistic, maximizes the signal SNR over the four amplitude parameters, but requires values to be chosen for the phase parameters.

A search for gravitational-wave pulsars therefore consists of performing matched filtering against a bank of templates, whose phase parameters are chosen to cover the parameter space of interest, e.g. over the whole sky and a broad range of frequencies for a search for unknown neutron stars. It is almost certain, however, that any signal in the data will possess parameters which are different from any one of the searched templates; consequentially, no template will perfectly match the signal, and the signal SNR will be degraded. While some loss in SNR is unavoidable, template banks are constructed such that the fractional loss in SNR, also known as the *mismatch*, can never be greater than some prescribed maximum. To accomplish this, a *metric* is often used to determine how closely the templates must be spaced in each parameter [36, 44]. How to construct a bank which minimizes the number of templates is known in theory [45], but is often difficult to accomplish in practice.

Unfortunately, the number of templates which must be matched filtered increases rapidly with the length of data being analyzed. If T denotes the time-span of the analyzed data, the number of templates which must be placed in each parameter dimension scales as follows: $T^{\sim 2}$ for the sky position (α, δ) , T for frequency, and T^{k+1} for each spindown parameter $f^{(k)}$. For example, the number of templates for an all-sky search requiring one spindown scales as $T^{\sim 5}$. In contrast, the sensitivity achievable by matched filtering increases only as $T^{1/2}$ (assuming that the data are contiguous in time). Due to the rapid increase in the number of required templates, matched filtering quickly becomes too computationally expensive for searching long data sets and large parameter spaces [25].

The solution is to resort to a hierarchical pipeline, where typically the data are broken into short segments, each of which are matched filtered separately. The results from each segment are then combined using *semi-coherent* analysis methods, which resemble matched filtering but do not require full amplitude and phase consistency of the signal template between data segments: typically, only the derivative of the phase (i.e. the frequency) must be consistent. (To distinguish it from semi-coherent methods, matched filtering is also referred to as *coherent* matched filtering.) For fixed T , and assuming no limits on computational cost, semi-coherent methods are less sensitive than coherent matched filtering: their sensitivity scales roughly as $N_s^{1/4} T_s^{1/2}$, where N_s and T_s are the number and time-span of each segment, while a coherent search (assuming contiguous data) scales as $T^{1/2} = N_s^{1/2} T_s^{1/2}$. On the other hand, wide-parameter-space gravitational-wave pulsar searches are almost always computationally limited, due to the large parameter spaces which must be searched. The number of templates, and hence the computational cost, of semi-coherent methods scale with $T_s \ll T$, instead of T , making them computationally cheaper than a fully-coherent search. This in turn permits a semi-coherent search to use more data, improving its sensitivity, while remaining computationally tractable. Some hierarchical searches, e.g. [20, 22], use $T_s \sim 30$ -minute segments, so that the coherent matched-filtering step closely resembles the computation of a power spectrum; other searches, e.g. [13, 21] use segments of $T_s \sim 1$ day.

Examples of semi-coherent methods are the Stack-Slide [24, 27, 28], Hough [26, 29], PowerFlux [46], cross-correlation [47], and global correlation transform [48] methods. Recently, methods which blend together aspects of semi- and fully-coherent methods have been developed [49–51]. The problem of how to optimise hierarchical searches is studied in [27, 31].

C. Search sensitivity

The sensitivity of a wide-parameter-space search for gravitational-wave pulsars has traditionally been characterized by the method presented in this section; two alternative methods of estimating sensitivity are discussed in Section VI.

The method follows from the canonical framework for statistical hypothesis testing formulated in [52, 53]. It is commonly referred to as the *frequentist method* in the gravitational-wave literature. In essence, the method provides the answer to the following question: if there were a population of gravitational-wave pulsar signals present in the searched data, each with the same amplitude h_0 , how large would h_0 need to be before we would be confident of detecting a very large fraction of them, e.g. 95%?

In order to make the above question more precise, we must first define what is meant here by *detection*. This

is complicated by the unavoidable fact that the output of any real gravitational-wave detector contains noise, in addition to any signal. When noisy data is analyzed, e.g. using matched filtering, the results of the analysis may be subject to *false alarms* and *false dismissals*. A false alarm is when the results of the analysis falsely indicate the presence of a signal, which is instead simply due to spurious noise fluctuations; a false dismissal is when a real signal present in the data is sufficiently corrupted by noise that it fails to be identified as a signal in the analysis results. We define a *detection statistic* to be a number quantifying a single search result, e.g. a correlation of the data against a single signal template. Typically, the detection statistic increases with the probability that the data contains a signal which matches the template. As an example, the \mathcal{F} -statistic represents the result of a matched-filtering analysis maximized over signal amplitude parameters. We say that a signal has been *detected*¹ if the value of its detection statistic s exceeds some threshold s_{fa} ².

Suppose that we perform a wide-parameter-space search of data which is known to contain no signal, e.g. the data may comprise only computer-generated Gaussian noise. Because the data used to compute the detection statistic are combined differently for each template, the values of s returned by the search will not be identical, but will instead follow a certain probability distribution. The probability of a detection in this data, i.e. the probability of a false alarm, denoted p_{fa} , is the probability of the detection statistic s exceeding the threshold s_{fa} , under the assumption that no signal is present in the data:

$$p_{\text{fa}} = p(s > s_{\text{fa}} \mid \text{no signal in data}), \quad (2.4)$$

where $p(A|B)$ denotes the probability of the statement A being true, given that we already know that B is true.

We now perform the same search using data which is known to contain a single signal, e.g. a simulated signal added to computer-generated Gaussian noise. Assuming that the majority of the searched templates are insensitive to the signal (which would be true of a properly constructed template bank), the majority of the returned values of s will follow the same probability distribution as before. For templates whose parameters are close to those of the signal, however, the values of s will follow a different probability distribution, which instead assumes that a signal is present. Assume that the SNR of the sig-

nal is large enough that these templates, close to the signal, are distinguishable from the remaining search templates; this assumption is satisfied by requiring a low false dismissal probability. The probability of this particular signal not being detected, i.e. the probability of its false dismissal, denoted $p_{\text{fd}}(\mathbf{p})$, is the probability of s , in the neighborhood of the signal, falling below the threshold s_{fa} , under the assumption that a signal with parameters \mathbf{p} is present in the data:

$$p_{\text{fd}}(\mathbf{p}) = p(s \leq s_{\text{fa}} \mid \text{signal}(\text{param. } \mathbf{p}) \text{ in data}). \quad (2.5)$$

Note that we assume here that the search template is perfectly matched to the signal, whereas in reality there will be some mismatch between them. (The loss of sensitivity due to mismatch between template and signal is considered in Section V C.) We repeat the search for a large number of signals, each with different \mathbf{p} drawn from a chosen distribution of signals parameters, $p(\mathbf{p})$. The overall false dismissal probability, denoted p_{fd} , is the average false dismissal probability obtained from each of the sampled signals:

$$p_{\text{fd}} = \langle p_{\text{fd}}(\mathbf{p}) \rangle_{\mathbf{p}} = \int \cdots \int d\mathbf{p} p(\mathbf{p}) p_{\text{fd}}(\mathbf{p}). \quad (2.6)$$

Equation (2.6) is computed using Monte Carlo integration, i.e. by the repeated computation of Eq. (2.5), with different parameters \mathbf{p} drawn from the distribution $p(\mathbf{p})$.

We now return to the question stated at the beginning of this section: how large would the amplitude h_0 of a population of gravitational wave signals need to be in order for them to be detected, e.g. 95% of the time? First, Eq. (2.4) is solved for the threshold s_{fa} that would result in a desired false alarm probability p_{fa} . Then, given a target false dismissal probability p_{fd} , Eq. (2.6) is solved for the overall strain amplitude h_0 . Suppose $h_0^{95\%}$ is the value of h_0 which solves Eq. (2.6) for $p_{\text{fd}} = 5\%$; then, if a population of signals have constant amplitudes $h_0^{95\%}$, then a fraction $1 - p_{\text{fd}} = 95\%$ of them will be detected. The amplitude $h_0^{95\%}$ thus characterizes the amplitude of signals a particular search method is able to confidently detect, and hence gives a useful measure of the search's sensitivity. Where a search finds no credible gravitational-wave signal, this sensitivity is interpreted as an upper limit, with confidence $1 - p_{\text{fd}}$, on the amplitude of signals present in the searched data.

The above procedure is used to set upper limits on the amplitude of gravitational waves for wide-parameter-space searches of LIGO and Virgo data (see Section I for references). The search parameter space is typically partitioned into small frequency bands, and upper limits are set separately for each frequency band. One important difference to the procedure described in this section is that s_{fa} is determined by the largest value of s returned by the search (after instrumental noise artifacts have been removed); an effective false alarm probability can then be determined from s_{fa} using Eq. (2.4).

¹ It is important to note that, when performing a search of real gravitational-wave detector data, any signal thus detected would never be automatically claimed as a genuine gravitational-wave signal; extensive follow-up investigations would first be performed to e.g. consider possible contamination by instrumental artifacts.

² The hierarchical searches in [13, 21] employ a slightly different definition of *detection*: the detection statistic must exceed a threshold in a given number of data segments.

III. ANALYTIC SENSITIVITY ESTIMATION

Theoretical studies of gravitational-wave pulsar search pipelines require a method of efficiently and accurately estimating the sensitivity achievable by such searches. The procedure described in Section II C is generally unsuited to this task, due to the computational expense of repeatedly generating and searching data containing simulated signals. In this section, we construct an analytic expression which accurately reproduces sensitivity estimates computed using the frequentist method, using the procedure described in Section II C, but is simpler to implement and computationally cheaper to calculate. Expressions for the false alarm and false dismissal probabilities are presented in Sections III A and III B respectively. Section III C presents a commonly used, but inaccurate, analytic estimator of the sensitivity, and Section III D presents a new, more accurate expression.

We restrict our attention to detection statistics s which follow χ^2 distributions (e.g. [54]). This implies that s can be written as the sum of squares of some number of normally-distributed random variables. Examples of such statistics are the StackSlide power [28], and the \mathcal{F} -statistic [23, 27]. Our prototypical search is a single-stage hierarchical search of N_s data segments, each of which span a time T_s . (A fully-coherent search is then given by the special case $N_s = 1$.) A coherent analysis is performed for each segment i , returning detection statistics s_i , which are then summed using a semi-coherent method to attain the final detection statistic $s = \sum_{i=1}^{N_s} s_i$.

In the absence of a signal, s is distributed according to a central χ^2 distribution, which takes a single parameter: the number of degrees of freedom of the statistic. Since s is the sum of the N_s values s_i , its number of degrees of freedom is given by $N_s\nu$, where each of the s_i has ν degrees of freedom. A detection statistic derived from the power of a signal (e.g. the StackSlide power) has $\nu = 2$; the \mathcal{F} -statistic has $\nu = 4$. We denote that s is distributed according to a central χ^2 distribution with $N_s\nu$ degrees of freedom by $s \sim \chi^2(N_s\nu, 0)$.

In the presence of a signal, s is distributed according to a non-central χ^2 distribution, which take two parameters: the number of degrees of freedom $N_s\nu$ as before, and the non-centrality parameter

$$N_s\rho^2 = h_0^2 T (a_+^2 \langle F_+^2 \rangle_t + a_\times^2 \langle F_\times^2 \rangle_t) \sum_{i=1}^{N_s} \sum_{j=1}^{N_{\text{det.}}} (S_h)_{ij}^{-1}. \quad (3.1)$$

The non-centrality parameter is the accumulated SNR of a signal analyzed over N_s data segments from $N_{\text{det.}}$ detectors, assuming perfect match between signal and template. The noise power spectral density of the i th data segment from the j th detector is labeled $(S_h)_{ij}$. It is convenient to define $N_s S_h^{-1} = \sum_{i=1}^{N_s} \sum_{j=1}^{N_{\text{det.}}} (S_h)_{ij}^{-1}$, whereupon we recover the expression for ρ^2 given in Eq. (2.3). For simplicity, therefore, we can ignore the summation of $(S_h)_{ij}$ over segments and detectors, and take ρ^2 to be

given by Eq. (2.3) for some appropriate value of S_h . For a real detector, S_h also varies as a function of frequency; we therefore assume that we are considering the sensitivity of a search over a frequency band sufficiently narrow, so that S_h can be assumed constant.

We denote that s is distributed according to a non-central χ^2 distribution with $N_s\nu$ degrees of freedom and non-centrality parameter $N_s\rho^2$ by $s \sim \chi^2(N_s\nu, N_s\rho^2)$. Note that, as in Section II C, we assume that the search template is perfectly matched to the signal, and delay considering the sensitivity lost due to template-signal mismatch until Section V C.

A. False alarm probability

We start by determining the threshold s_{fa} on the detection statistic required to give a certain false alarm probability p_{fa} .

Assuming that no signal is present, the probability that a single value of the detection statistic s falls below s_{fa} is $p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s\nu, 0])$. Suppose that the search returns N_t values of s , i.e. N_t templates³ are searched. We assume that each value of s is statistically independent, i.e. that the joint probability of obtaining any two values s_1 and s_2 is just the product of the probabilities of obtaining s_1 and s_2 individually. (The validity of this assumption is examined in Section V A.) The probability that N_t values of s fall below s_{fa} is then given by N_t multiples of the single-value probability: $p(\dots)^{N_t}$. Its complement, $1 - p(\dots)^{N_t}$, is the probability that in a search of N_t templates, one or more values of s will be returned above the threshold, i.e. the probability of a false alarm, p_{fa} . In short, we have

$$p_{\text{fa}} = 1 - [p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s\nu, 0])]^{N_t}. \quad (3.2)$$

Simple re-arrangement gives

$$(1 - p_{\text{fa}})^{1/N_t} = p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s\nu, 0]). \quad (3.3)$$

Since typically either $N_t = 1$ (e.g. for a search for a known pulsar) or $N_t \gg 1$, the left-hand side of Eq. (3.3) can be replaced with the first-order binomial expansion $1 - p_{\text{fa}}/N_t$. We now have

$$p_{\text{fa}}/N_t = p(s > s_{\text{fa}} | s \sim \chi^2[N_s\nu, 0]). \quad (3.4)$$

Note that p_{fa}/N_t can be interpreted as the false alarm probability for a single template out of the N_t templates searched.

We must now solve Eq. (3.4) for s_{fa} , i.e. we must compute the inverse function of the central χ^2 distribution.

³ Note that we are referring to the number of templates searched in the semi-coherent stage of a hierarchical pipeline, which is distinct from the number of templates searched in each individual data segment; see e.g. [27].

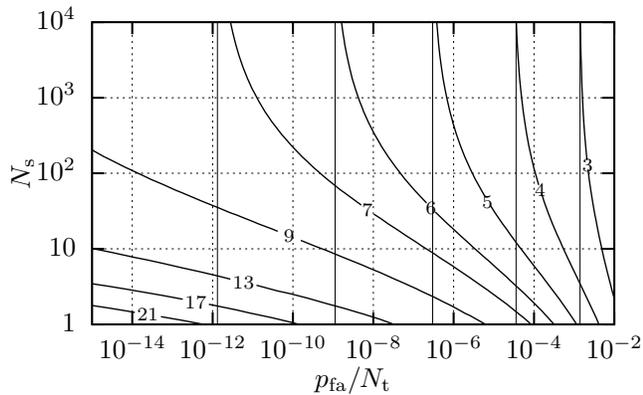


FIG. 1. Normalized false alarm threshold z_{fa} , as a function of p_{fa}/N_t and N_s , with $\nu = 4$ (thick lines). Its limiting values for large N_s , given by the normal distribution approximation, are plotted as thin vertical lines for $z_{\text{fa}} = 3$ –7.

For small values of p_{fa}/N_t (typically, p_{fa} is chosen to be 1%, and $N_t \gg 1$), an analytic expression for s_{fa} derived from [55] is given in Appendix B. We define the normalized false alarm threshold z_{fa} to be the difference between s_{fa} and the mean of a central χ^2 distribution with $N_s\nu$ degrees of freedom, in units of the distribution's standard deviation:

$$z_{\text{fa}} = \frac{s_{\text{fa}} - N_s\nu}{\sqrt{2N_s\nu}}. \quad (3.5)$$

In the limit of large N_s , $z_{\text{fa}} \approx \sqrt{2} \operatorname{erfc}^{-1}(2p_{\text{fa}}/N_t)$ and is independent of N_s (see Appendix B). Note that this limiting value for z_{fa} is identical to the expression obtained if the central χ^2 distribution is approximated by a normal distribution, as is commonly done (e.g. in [26]). Figure 1 plots z_{fa} as a function of p_{fa}/N_t and N_s , and illustrates its convergence to the normal distribution approximation for large N_s .

B. False dismissal probability

Having determined the threshold s_{fa} appropriate for a desired false alarm probability p_{fa} , we now attempt to solve Eq. (2.6) for some quantity which usefully characterizes the sensitivity of the search. The quantity most often used for this purpose has been the dimensionless expression $h_0\sqrt{T_s/S_h}$, which is sometimes referred to as the *statistical factor* [30, 38]. The statistical factor quantifies the sensitivity of a search to a population of signals of amplitude h_0 , relative to the performance of the detector given by its noise power spectral density S_h , and the length of (coherently) analyzed data T_s .

In this paper, we propose instead to use the root-mean-square SNR, $\sqrt{\langle\rho^2\rangle}$, to quantify sensitivity. Here, $\langle\rangle$ denotes averaging over α , $\sin\delta$, ψ [see Eq. (A9)], and $\xi \equiv \cos\iota$, given by $\langle\rho^2\rangle_\xi = (1/2) \int_{-1}^1 d\xi \rho^2$. This choice

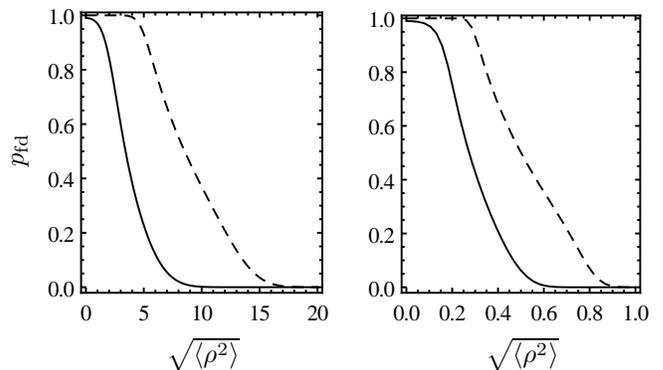


FIG. 2. False dismissal probability p_{fd} as a function of $\sqrt{\langle\rho^2\rangle}$, given by Eq. (3.9), with $\nu = 4$, and for $p_{\text{fa}}/N_t = 0.01$ (solid) and 10^{-12} (dashed), and $N_s = 1$ (left) and 10^4 (right).

of averaging implies that the population of gravitational-wave pulsars being searched for are isotropically distributed over the sky, and are isotropically oriented. Using Eq. (2.3), we have

$$\langle\rho^2\rangle = \frac{h_0^2 T_s}{S_h} (\langle a_+^2 \rangle_\xi \langle F_+^2 \rangle_{\alpha, \sin\delta, \psi, t} + \langle a_\times^2 \rangle_\xi \langle F_\times^2 \rangle_{\alpha, \sin\delta, \psi, t}). \quad (3.6)$$

The averages of F_+^2 and F_\times^2 over α , $\sin\delta$, and ψ are given by Eq. (A11) (where we assume $\zeta = \pi/2$); after this averaging F_+^2 and F_\times^2 are independent of time. We assume a signal generated by a non-axisymmetrically deformed neutron star, for which (e.g. [23, 35]):

$$a_+ = \frac{1 + \xi^2}{2}, \quad \langle a_+^2 \rangle_\xi = \frac{7}{15}; \quad (3.7a)$$

$$a_\times = \xi, \quad \langle a_\times^2 \rangle_\xi = \frac{1}{3}. \quad (3.7b)$$

Finally we have

$$\sqrt{\langle\rho^2\rangle} = \frac{2}{5} h_0 \sqrt{\frac{T_s}{S_h}}, \quad (3.8)$$

i.e. $\sqrt{\langle\rho^2\rangle}$ is directly proportional to the statistical factor. Unlike the statistical factor, however, $\sqrt{\langle\rho^2\rangle}$ relates directly to a property of the population of signals being searched for (i.e. their mean SNR), and hence is a more directly physical quantity. It also has a clearer interpretation as a measure of sensitivity: for example, to improve the sensitivity of a search, we must make the search able to detect signals with weaker SNR (at the same false alarm and dismissal probabilities), and hence we must lower the mean SNR, i.e. $\sqrt{\langle\rho^2\rangle}$, of the population of signals which the search can detect. It is convenient to write ρ in terms of $\sqrt{\langle\rho^2\rangle}$ and a factor R , defined such that $\rho = \sqrt{\langle\rho^2\rangle} R$, which implies $\langle R^2 \rangle = 1$.

Assuming that a signal with parameters \mathbf{p} is present, the probability that the detection statistic s (in the neighborhood of the signal, as discussed in Section II C) falls

below s_{fa} is $p_{\text{fd}}(\mathbf{p}) = p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s \nu, N_s \langle \rho^2 \rangle R^2])$. Since h_0 , S_h , and T_s are taken to be constants, $\langle \rho^2 \rangle$ is also a constant. The only quantity which depends on the signal parameters is therefore R , which is a function of $\mathbf{p} = (\alpha, \delta, \psi, \xi)$. By Eq. (2.6), the overall false dismissal probability p_{fd} is

$$p_{\text{fd}} = \left\langle p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s \nu, N_s \langle \rho^2 \rangle R^2]) \right\rangle_{\alpha, \sin \delta, \psi, \xi}. \quad (3.9)$$

Equation (3.9) is plotted in Fig. 2 for different choices of p_{fa}/N_t and N_s .

To proceed, we must now solve Eq. (3.9) for $\langle \rho^2 \rangle$. Unfortunately, an analytic solution is difficult to obtain, due to the complicated dependence of R on the signal parameters. Two approaches to solving Eq. (3.9) analytically are presented in Sections III C and III D.

C. Sensitivity to constant-SNR signal populations

To solve Eq. (3.9), it is common to assume that every signal, in the population of signals being searched for, has the *same* SNR, which we denote by $\bar{\rho}$. An alternative interpretation of this approximation is that the population of signals can be replaced by a *single* signal, whose SNR is $\bar{\rho} = \sqrt{\langle \rho^2 \rangle}$. While this assumption does not accurately model a physically reasonable population of gravitational-wave pulsar signals (see Section III D), it does allow Eq. (3.9) to be readily solved for $\bar{\rho}$: because the signals all have the same SNR, no averaging over signal parameters is required, and $\langle \rho^2 \rangle R^2$ is simply replaced by $\bar{\rho}^2$. The accuracy of this approximation is examined in Section III D.

Another common simplification is to approximate the non-central χ^2 distribution by a normal distribution $\mathcal{N}(\mu, \sigma)$ with the same mean and standard deviation, which are $\mu = N_s(\nu + \bar{\rho}^2)$ and $\sigma = \sqrt{2N_s(\nu + 2\bar{\rho}^2)}$ respectively. This approximation introduces an error in $\bar{\rho}$ of $\lesssim 2.5\%$ (at $p_{\text{fa}}/N_t = 0.01$, $N_s = 1$) which decreases with decreasing p_{fa}/N_t and increasing N_s . With this approximation, Eq. (3.9) reduces to

$$p_{\text{fd}} \approx p(s \leq s_{\text{fa}} | s \sim \mathcal{N}[\mu, \sigma]) \quad (3.10)$$

$$= \frac{1}{2} \operatorname{erfc} \left(\frac{N_s \bar{\rho}^2 - z_{\text{fa}} \sqrt{2N_s \nu}}{2\sqrt{N_s(\nu + 2\bar{\rho}^2)}} \right), \quad (3.11)$$

where we have substituted s_{fa} with the normalized threshold z_{fa} defined by Eq. (3.5), and erfc is the complementary error function. The solution to this equation is

$$\bar{\rho}^2 = \sqrt{\frac{2\nu}{N_s}} z_{\text{fa}} + \frac{2q^2}{N_s} \left[1 + (1 + \mathcal{Q})^{\frac{1}{2}} \right], \quad (3.12)$$

where $\mathcal{Q} = (N_s \nu + z_{\text{fa}} \sqrt{8N_s \nu}) / (2q^2)$, $q = \sqrt{2} \operatorname{erfc}^{-1} 2p_{\text{fd}}$, and erfc^{-1} is the inverse complementary error function.

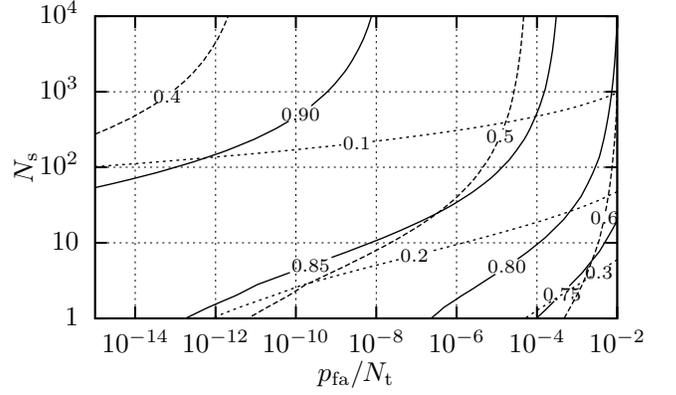


FIG. 3. Contours of contributions of the terms \mathcal{T} (see the text) to the second factor of Eq. (3.13): $\sqrt{\mathcal{T}_1}/\sum_i \mathcal{T}_i$ (solid contours), $\sqrt{\mathcal{T}_2}/\sum_i \mathcal{T}_i$ (dashed contours), and $\sqrt{\mathcal{T}_3}/\sum_i \mathcal{T}_i$ (dotted contours).

For $p_{\text{fa}}/N_t \leq 0.01$, $p_{\text{fd}} \geq 0.05$, $N_s \geq 1$, and $\nu \geq 2$, $\mathcal{Q} \gtrsim 5$, and it is reasonable to approximate $\sqrt{1 + \mathcal{Q}}$ with $\sqrt{\mathcal{Q}}$. This permits further simplification to

$$\bar{\rho} = \left[\frac{2\nu}{N_s} \right]^{\frac{1}{4}} \left[z_{\text{fa}} + q \left(1 + \frac{z_{\text{fa}} \sqrt{8}}{\sqrt{N_s \nu}} \right)^{\frac{1}{2}} + \frac{q^2 \sqrt{2}}{\sqrt{N_s \nu}} \right]^{\frac{1}{2}}. \quad (3.13)$$

The first factor of $\bar{\rho}$ scales with $N_s^{-1/4}$, which is a well-known property of hierarchical searches [26, 28]. The second factor contains, inside the square root, a constant term, $\mathcal{T}_1 = z_{\text{fa}}$, a term which scales approximately with $N_s^{-1/4}$, $\mathcal{T}_2 = q(1 + z_{\text{fa}} \sqrt{8/N_s \nu})^{1/2}$, and a term which scales with $N_s^{-1/2}$, $\mathcal{T}_3 = q^2 \sqrt{2/N_s \nu}$. These additional terms appear because we do not employ the weak-signal approximation of [26], which Taylor-expands p_{fd} [Eq. (3.11)] to first order in $\bar{\rho}$. This approximation is not valid here because we are interested in small false dismissal probabilities, which implies that a large fraction of signals will be strong, i.e. enough to cross the detection threshold. Each of the terms \mathcal{T} contribute to the value of $\bar{\rho}$ over the ranges $p_{\text{fa}}/N_t \leq 0.01$ and $N_s \geq 1$, as shown in Fig. 3; over the plotted ranges of p_{fa}/N_t and N_s , the contributions of each term to the sum $\sqrt{\sum_i \mathcal{T}_i}$ is 70–90% for $\sqrt{\mathcal{T}_1}$, 40–60% for $\sqrt{\mathcal{T}_2}$, and $\leq 40\%$ for $\sqrt{\mathcal{T}_3}$. For small p_{fa}/N_t and large N_s , $\sqrt{\mathcal{T}_1}$ dominates the sum, and $\bar{\rho}$ begins to scale purely with $N_s^{-1/4}$. A general power-law scaling of $\bar{\rho}$ with N_s is utilised in [31].

D. Sensitivity to isotropically-distributed signals

This section presents a new, more accurate sensitivity estimator than that presented in Section III C. The important difference is that we will no longer assume that all signals being searched for have the same SNR, and instead perform the correct averaging of $p_{\text{fd}}(\mathbf{p})$ as given

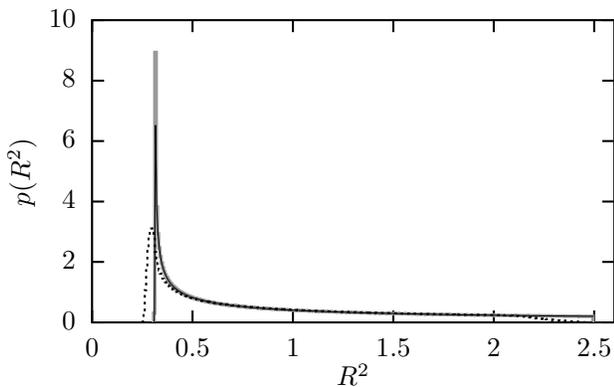


FIG. 4. The distribution of R_{iso}^2 , calculated analytically using Eq. (3.15) (thin black line), and numerically (thick gray line). The black dashed line plots the distribution of $R_{\text{iso,mism}}^2 = R_{\text{iso}}^2(1 - \mu)$, where the mismatch μ is drawn from a distribution appropriate for a lattice-based template bank (see Section V C).

in Eq. (3.9). This implies that we are searching for a population of signals isotropically distributed in sky position (α, δ) and orientation (ψ, ξ) parameters. Unlike the assumption of constant SNR, this is a more physically reasonable assumption. The observed distribution of nearby (millisecond) radio pulsars is roughly isotropic in the sky [56], and the angular momentum of neutron stars (which determines ψ and ξ) is not expected to have a preferred direction.

We also assume that we are searching data from a network comprising a large number of gravitational-wave detectors, evenly distributed over the Earth, such the network is isotropically sensitive to gravitational wave arriving from all directions. (The validity of this assumption is investigated in Section V B.) We can therefore approximate R^2 by [see Eqs. A9 and A10]

$$\begin{aligned} R_{\text{iso}}^2(\xi) &= \langle R^2 \rangle_{\Phi_s, \sin \lambda, \gamma} = \langle R^2 \rangle_{\alpha, \sin \delta, \psi} \\ &= \frac{5}{16}(\xi^4 + 6\xi^2 + 1). \end{aligned} \quad (3.14)$$

Substituting into Eq. (3.9), we see that we now need only to average $p_{\text{fd}}(\mathbf{p})$ over ξ . Figure 4 plots the probability distribution of R_{iso}^2 assuming a uniform distribution in ξ , given analytically by

$$p(R_{\text{iso}}^2)^{-2} = \frac{1}{10}(50 + 20R_{\text{iso}}^2)^{\frac{3}{2}} - \frac{3}{4}(50 + 20R_{\text{iso}}^2), \quad (3.15)$$

where R_{iso}^2 ranges from $5/16$ ($\xi = 0$) to $5/2$ ($|\xi| = 1$). Note that the most probable values of R_{iso}^2 are those from linearly-polarized signals ($\xi = 0$), with a rapid fall-off towards circularly-polarized ($|\xi| = 1$) signals. It is clear from this plot that the assumption that all signals have the same SNR (i.e. that $R_{\text{iso}}^2 = 1$), as assumed in Section III C, is not a reasonable one.

We next approximate the non-central χ^2 distribution by a normal distribution \mathcal{N} , as per Section III C. Here,

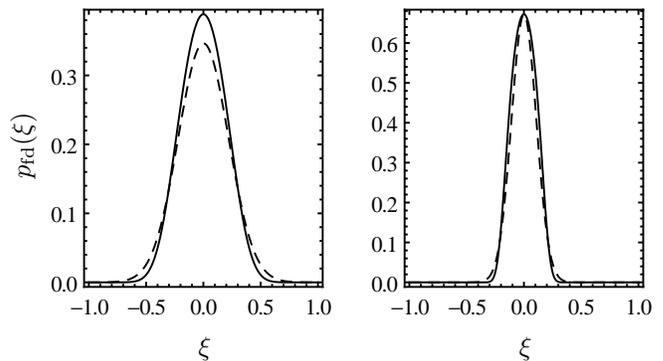


FIG. 5. False dismissal probability $p_{\text{fd}}(\xi)$ of a single signal observed in an isotropically-sensitive detector network, as a function of the cosine of the signal's inclination angle, $\xi = \cos \iota$. The exact result of Eq. (3.17) (solid line) is compared against the Gaussian function approximation of Eq. (3.18) (dashed line), for $p_{\text{fd}} = 1\%$, $\nu = 4$, and: $p_{\text{fa}}/N_t = 0.01$, $N_s = 1$ (left), and $p_{\text{fa}}/N_t = 10^{-12}$, $N_s = 10^4$ (right).

the approximation introduces a smaller error of $\lesssim 0.5\%$ (at $p_{\text{fa}}/N_t = 0.01$, $N_s = 1$) which also decreases with decreasing p_{fa}/N_t and increasing N_s . Equation (3.9) now reads

$$p_{\text{fd}} = \langle p_{\text{fd}}(\xi) \rangle_{\xi}, \quad (3.16)$$

where

$$p_{\text{fd}}(\xi) = \frac{1}{2} \operatorname{erfc} \left(\frac{N_s(\nu + \hat{\rho}^2 R_{\text{iso}}^2) - s_{\text{fa}}}{2\sqrt{N_s(\nu + 2\hat{\rho}^2 R_{\text{iso}}^2)}} \right), \quad (3.17)$$

and we define $\hat{\rho} \equiv \sqrt{\langle \rho^2 \rangle}$.

We find that $p_{\text{fd}}(\xi)$ is well-approximated by a Gaussian function:

$$p_{\text{fd}}(\xi) \approx A e^{-B\xi^2}, \quad (3.18)$$

for suitable choices of A and B . This is illustrated in Fig. 5, where we plot Eq. (3.17) against Eq. (3.18) for two different choices of p_{fa}/N_t and N_s . We note that it is important only that the integral of Eq. (3.18) accurately approximate that of Eq. (3.17); thus, the slight underestimation of $p_{\text{fd}}(\xi)$ by Eq. (3.18) at $\xi \approx 0$ is partly offset by its overestimation at $|\xi| \gtrsim 0.5$. We note that $p_{\text{fd}}(\xi) \approx 0$ at $|\xi| = 1$; numerical investigations confirm that this property holds true for $p_{\text{fa}}/N_t \lesssim 0.01$ and $N_s \gtrsim 1$. We can therefore replace the integral of Eq. (3.18) over $-1 \leq \xi \leq 1$ with one over $-\infty \leq \xi \leq \infty$, since the integral over $|\xi| > 1$ contributes little to the value of p_{fd} . The integral of Eq. (3.18) then simplifies to

$$p_{\text{fd}} \approx \frac{1}{2} \int_{-\infty}^{\infty} d\xi A e^{-B\xi^2} = \frac{A}{2} \sqrt{\frac{\pi}{B}}. \quad (3.19)$$

We choose A by setting $\xi = 0$ in Eq. (3.18), and obtain

$$A = \frac{1}{2} \operatorname{erfc} \left(\frac{N_s \hat{\rho}^2 R_0^2 - z_{\text{fa}} \sqrt{2N_s \nu}}{2\sqrt{N_s(\nu + 2\hat{\rho}^2 R_0^2)}} \right), \quad (3.20)$$

where $R_0^2 = R_{\text{iso}}^2(\xi = 0)$. Next, we choose $\xi_1 > 0$ to be the value of ξ such that $p_{\text{fd}}(\xi = \xi_1)$ [Eq. (3.17)] equals the target false dismissal probability p_{fd} . By equating Eqs. (3.11) and (3.17), we deduce that ξ_1 satisfies $\hat{\rho}^2 R_{\text{iso}}^2(\xi = \xi_1) = \hat{\rho}^2$, and is given by

$$\xi_1 = \left[2\sqrt{2 + \frac{4}{5} \left(\frac{\hat{\rho}}{\hat{\rho}}\right)^2} - 3 \right]^{\frac{1}{2}}. \quad (3.21)$$

Substituting ξ_1 into Eq. (3.18) gives the solution for

$$B = \frac{1}{\xi_1^2} \ln \left(\frac{A}{p_{\text{fd}}} \right). \quad (3.22)$$

Combining Eqs. (3.19), (3.20), and (3.22), and taking logarithms, we have

$$\ln \left(\frac{2p_{\text{fd}}}{\xi_1 \sqrt{\pi}} \right) = \ln A - \frac{1}{2} \ln \left[\ln \left(\frac{A}{p_{\text{fd}}} \right) \right]. \quad (3.23)$$

The terms on the right-hand side of Eq. (3.23) may be Taylor-expanded to second order in $x = \text{erfc}^{-1} 2A$:

$$\ln A = -\ln 2 - \frac{2}{\sqrt{\pi}} x - \frac{2}{\pi} x^2 + \mathcal{O}(x^3), \quad (3.24)$$

$$\begin{aligned} \frac{1}{2} \ln \left[\ln \left(\frac{A}{p_{\text{fd}}} \right) \right] &= \ln \sqrt{|\ln 2p_{\text{fd}}|} + \frac{1}{\sqrt{\pi} \ln 2p_{\text{fd}}} x \\ &\quad - \frac{1 - \ln 2p_{\text{fd}}}{\pi (\ln 2p_{\text{fd}})^2} x^2 + \mathcal{O}(x^3), \end{aligned} \quad (3.25)$$

where the second expansion is valid only for $p_{\text{fd}} < 1/2$. For $p_{\text{fa}}/N_t \leq 0.01$, $p_{\text{fd}} \geq 0.05$, $N_s \geq 1$, and $\nu \geq 2$, these expansions introduce errors of $\lesssim 1\%$ and $\lesssim 16\%$ respectively. Substituting Eqs. (3.24) and (3.25) into Eq. (3.23), and using Eq. (3.20), we solve for

$$\begin{aligned} x &= \frac{N_s \hat{\rho}^2 R_0^2 - z_{\text{fa}} \sqrt{2N_s \nu}}{2\sqrt{N_s(\nu + 2\hat{\rho}^2 R_0^2)}} \\ &= \frac{\sqrt{\pi}}{2} \Gamma^{-1} \left(\sqrt{1 - 2(1 - \Delta) \ln(2p_{\text{fd}} \Xi)} - 1 \right), \end{aligned} \quad (3.26)$$

where

$$\Gamma = 1 - \frac{1}{\ln 2p_{\text{fd}}} + \frac{2}{1 + 2 \ln 2p_{\text{fd}}}, \quad (3.27a)$$

$$\Delta = \frac{1}{1 + 2 \ln 2p_{\text{fd}}} + \frac{2}{(1 + 2 \ln 2p_{\text{fd}})^2}, \quad (3.27b)$$

$$\Xi = \frac{2}{\xi_1} \sqrt{\frac{|\ln 2p_{\text{fd}}|}{\pi}}. \quad (3.27c)$$

Now, suppose that $p'_{\text{fd}} = (1/2) \text{erfc } y$, and Taylor-expand $\ln p'_{\text{fd}}$ to second order in y , as in Eq. (3.24); solving for y gives an expression (valid only for $p'_{\text{fd}} < \sqrt{e}/2$) for

$$y = \text{erfc}^{-1} 2p'_{\text{fd}} = \frac{\sqrt{\pi}}{2} \left(\sqrt{1 - 2 \ln 2p'_{\text{fd}}} - 1 \right). \quad (3.28)$$

Note the similarity of the right-hand sides of Eqs. (3.26) and (3.28). We now define the *effective* false dismissal probability p'_{fd} , normalized false alarm threshold z'_{fa} , and number of segments N'_s :

$$p'_{\text{fd}} = p_{\text{fd}} \frac{\Xi}{(2p_{\text{fd}} \Xi)^\Delta}, \quad z'_{\text{fa}} = z_{\text{fa}} \Gamma, \quad N'_s = N_s \Gamma^2. \quad (3.29)$$

Using these quantities, we equate Eqs. (3.26) and (3.28), obtaining

$$\frac{N'_s \hat{\rho}^2 R_0^2 - z'_{\text{fa}} \sqrt{2N'_s \nu}}{2\sqrt{N'_s(\nu + 2\hat{\rho}^2 R_0^2)}} = \text{erfc}^{-1} 2p'_{\text{fd}}. \quad (3.30)$$

Finally, by noting that Eq. (3.30) is similar in form to Eq. (3.11), it follows that its solution for $\hat{\rho}$ is given by analogy to Eq. (3.13):

$$\hat{\rho} = \frac{1}{R_0} \left[\frac{2\nu}{N'_s} \right]^{\frac{1}{4}} \left[z'_{\text{fa}} + q' \left(1 + \frac{z'_{\text{fa}} \sqrt{8}}{\sqrt{N'_s \nu}} \right)^{\frac{1}{2}} + \frac{q'^2 \sqrt{2}}{\sqrt{N'_s \nu}} \right]^{\frac{1}{2}}. \quad (3.31)$$

where $q' = \sqrt{2} \text{erfc}^{-1} 2p'_{\text{fd}}$.

Evaluation of Eq. (3.31) is complicated by the fact that its right-hand side is itself a function of $\hat{\rho}$, through p'_{fd} [Eq. (3.29)], Ξ [Eq. (3.27c)], and ξ_1 [Eq. (3.21)]. Nevertheless, Eq. (3.31) may be iteratively solved for $\hat{\rho}$ using the following scheme. First, a reasonable initial guess, $\hat{\rho}_0$ is chosen: an appropriate choice is $\hat{\rho}_0 \approx 1.4\bar{\rho}$ (see Fig. 7). Next, $\hat{\rho}_0$ is used to compute an updated value, $\hat{\rho}_1$, by substituting into Eq. (3.31): $\hat{\rho}_1 = \hat{\rho}(\hat{\rho}_0)$. Thereafter, new values of $\hat{\rho}$ are obtained using the mean of the previous two values, i.e. the n th value of $\hat{\rho}$ is

$$\hat{\rho}_n = \hat{\rho} \left(\frac{\hat{\rho}_{n-1} + \hat{\rho}_{n-2}}{2} \right). \quad (3.32)$$

Using this scheme, the sequence of values $\{\hat{\rho}_n\}$ reliably converges to an accurate value of $\hat{\rho}$: typically, 20–80 iterations are required to achieve sufficient accuracy. The use of the mean of the previous two $\hat{\rho}$ in Eq. (3.32) suppresses divergent oscillations in the sequence $\{\hat{\rho}_n\}$.

Figure 6 plots $\hat{\rho}$ as a function of p_{fa}/N_t and N_s , with $p_{\text{fd}} = 0.1$, for a detection statistic with $\nu = 4$ degrees of freedom (e.g. the \mathcal{F} -statistic). For a single-template search at 1% false alarm and 10% false dismissal probabilities, we see that $p_{\text{fa}} = 0.01$, $N_t = N_s = 1$, and $\hat{\rho} = 6.3$, which should be interpreted as the *average* sensitivity of a collection of single-template searches which cover the parameter space uniformly in sky position (α, δ) and orientation (ψ, ξ) . This value of $\hat{\rho}$ is equivalent to a statistical factor of 15.7. This differs from the often-quoted (e.g. [18]) statistical factor for a single-template search of 11.4, because that statistical factor is calculated assuming a signal with an *average* SNR, i.e. it is calculated from $\bar{\rho}$ instead of $\hat{\rho}$. For the search for Cassiopeia A presented in [12], $p_{\text{fa}} = 0.01$, $N_t \sim 10^{13}$, $N_s = 1$, and

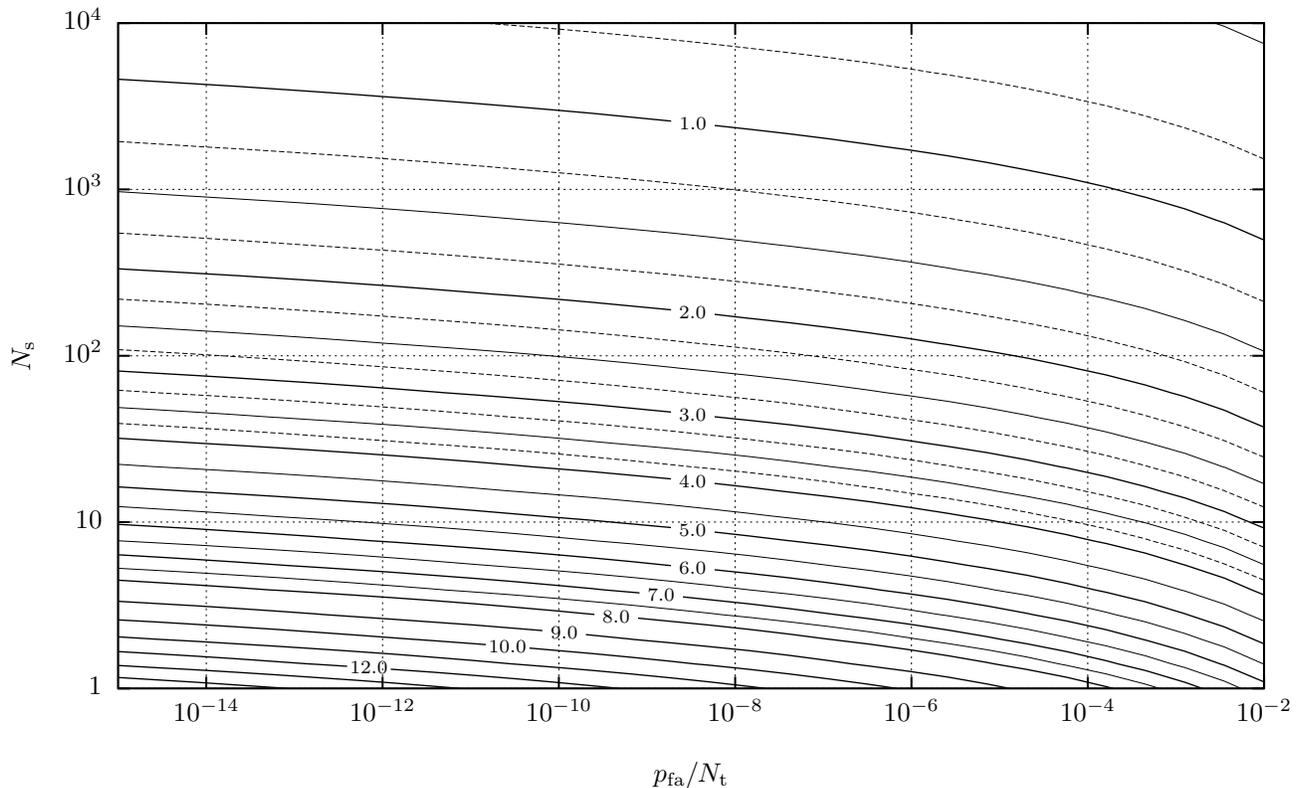


FIG. 6. Contours of $\hat{\rho}$ as a function of p_{fa}/N_t and N_s , with $p_{fd} = 0.1$ and $\nu = 4$. Thick solid contours are in units of 1.0, thin solid contours are in units of 0.5, and thin dashed contours are in units of 0.25.

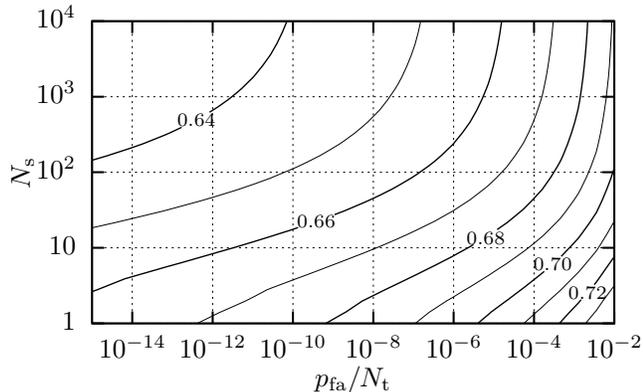


FIG. 7. Contours of the ratio of $\bar{\rho}/\hat{\rho}$ as a function of p_{fa}/N_t and N_s , with $p_{fd} = 0.1$ and $\nu = 4$. Unlabeled thin contours are in units of 0.01.

$\hat{\rho} \approx 14$, which implies a statistical factor of 35 consistent with that quoted in [38].

Figure 7 plots the ratio of $\bar{\rho}$ to $\hat{\rho}$, as a function of p_{fa}/N_t and N_s . Relative to $\hat{\rho}$, we see that $\bar{\rho}$ underestimates the mean SNR detectable by a search by $\sim 29 \pm 5\%$, i.e. it overestimates the search's sensitivity by the same amount. For comparison, the typically amplitude cali-

bration error of the LIGO detectors is $\sim 10\%$ [57], and the ratio of the best upper limits on gravitational waves from Cassiopeia A to the indirect limits from energy conservation is $\sim 60\%$ [12]. Thus, an error of $\sim 30\%$ in estimating a search's sensitivity is a significant discrepancy. Note, however, that the change in $\bar{\rho}/\hat{\rho}$ is small over the ranges of p_{fa}/N_t and N_s plotted in Fig. 7. We conclude from this that, while $\bar{\rho}$ does not predict the correct sensitivity, it does capture the correct scaling of sensitivity with the false alarm probability, template count, and number of segments. This conclusion also follows from the similarity in form between Eqs. (3.13) and (3.31).

To illustrate the relationship between $\hat{\rho}$ and $\bar{\rho}$, Fig. 8 plots z'_{fa}/z_{fa} and N'_s/N_s as functions of p_{fd} , and Fig. 9 plots p'_{fd}/p_{fd} as a function of p_{fd} and $\bar{\rho}/\hat{\rho}$. We see that both z'_{fa}/z_{fa} and N'_s/N_s are less than unity for $p_{fd} > 0$, while p'_{fd} is generally greater than p_{fd} by a factor of ~ 4 – 8 . Note too that the right-hand side of Eq. (3.31) is divided by $R_0 \approx 0.56$ [Eq. (3.14)]. This implies that $\hat{\rho}$ is dominated by contributions from linearly-polarised signals, as seen in the distribution of R_{iso}^2 plotted in Fig. 4. Therefore, we may think of $\hat{\rho}$ as estimating the sensitivity of a search to a population of constant-SNR signals (as for $\bar{\rho}$), but where the signals are linearly polarised (hence the division by R_0), and where the search is performed with reduced false alarm threshold z'_{fa} , a reduced number

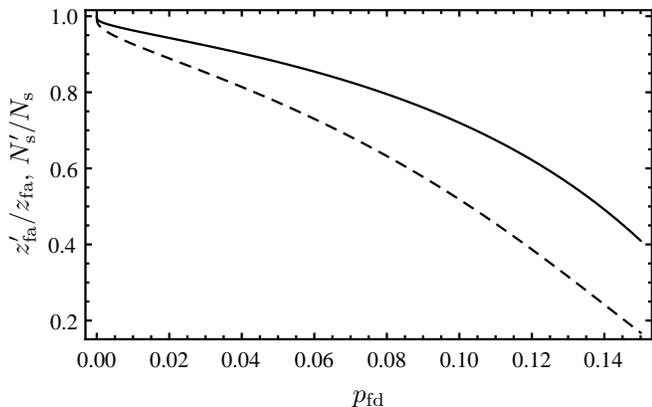


FIG. 8. Ratios of the effective false alarm threshold $z'_{\text{fa}}/z_{\text{fa}}$ (solid) and number of segments N'_s/N_s (dashed) as functions of p_{fd} .

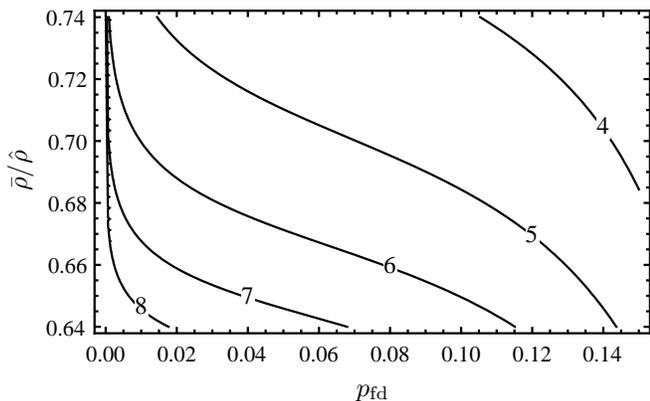


FIG. 9. Contours of the effective false dismissal probability $p'_{\text{fd}}/p_{\text{fd}}$ as a function of p_{fd} and $\tilde{\rho}/\hat{\rho}$.

of segments N'_s , and a greatly increased false dismissal probability p'_{fd} .

We note that Eqs. (3.25) and (3.28) impose restrictions on permissible values of $p_{\text{fd}} < 0.5$ and $p'_{\text{fd}} < \sqrt{e}/2 \approx 0.82$ respectively. The latter restriction and Fig. 9 implies that, depending on the value of $\tilde{\rho}/\hat{\rho}$, p_{fd} is further restricted to be less than $0.82/4$ – $0.82/8 \approx 0.1$ – 0.2 . In general, therefore, the use of Eq. (3.31) is restricted to values of $p_{\text{fd}} \lesssim 10$ – 20% . In practice this is not an onerous restriction, as we are generally only interested in small false dismissal probabilities of either 10% or 5%. Gravitational-wave pulsar searches of LIGO and Virgo data (see Section I for references) have set upper limits with corresponding upper limit confidences of 90% and 95% respectively.

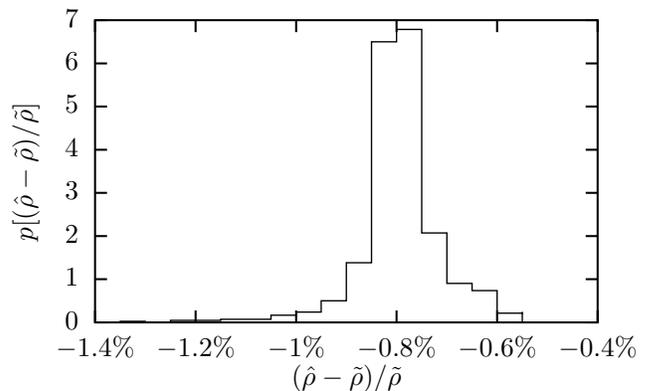


FIG. 10. Histogram of the relative difference between $\hat{\rho}$, given by Eq. (3.31), and the numerical solution to Eq. (3.9), $\tilde{\rho}$, over the ranges $10^{-15} \leq p_{\text{fa}}/N_t \leq 10^{-2}$ and $1 \leq N_s \leq 10^4$.

IV. ACCURACY OF ANALYTIC SENSITIVITY ESTIMATOR

In this section, we validate the accuracy of the analytic sensitivity estimator $\hat{\rho}$, derived in Section III D, against the sensitivity calculated using two methods: numerically solving Eq. (3.9), in Section IV A; and performing software injections, in Section IV B. The validation is performed using a detection statistic with $\nu = 4$, and for a target false dismissal probability of $p_{\text{fd}} = 0.1$.

A. Numerical solution of false dismissal equation

We first compare the sensitivity predicted by $\hat{\rho}$ against the sensitivity $\sqrt{\langle \rho^2 \rangle}$ calculated by solving Eq. (3.9) numerically. We denote by $\tilde{\rho}$ the value of $\sqrt{\langle \rho^2 \rangle}$ which solves Eq. (3.9) for a target false dismissal probability p_{fd} . We denote by $p_{\text{fd}}(\tilde{\rho}_{\text{trial}})$ the result of computing Eq. (3.9) for a given trial value of $\tilde{\rho}$, denoted $\tilde{\rho}_{\text{trial}}$. For $\tilde{\rho}_{\text{trial}} = 0$, $p_{\text{fd}}(\tilde{\rho}_{\text{trial}}) > p_{\text{fd}}$, otherwise Eq. (3.9) has no solution for the chosen p_{fd} . We start by determining a $\tilde{\rho}_{\text{max}}$ such that $p_{\text{fd}}(\tilde{\rho}_{\text{max}}) < p_{\text{fd}}$, thus bracketing $\tilde{\rho}$ to between 0 and $\tilde{\rho}_{\text{max}}$. We then use a simple bifurcation search to converge to $\tilde{\rho}$, which terminates when the relative error between $p_{\text{fd}}(\tilde{\rho})$ and the target p_{fd} is less than 10^{-3} . Equation (3.9) is solved for a grid of logarithmically-spaced values of p_{fa}/N_t and N_s : 30 values in the range $10^{-15} \leq p_{\text{fa}}/N_t \leq 10^{-2}$, and 28 values in the range $1 \leq N_s \leq 10^4$.

Figure 10 shows the relative difference between the analytic $\hat{\rho}$, and the numerically calculated $\tilde{\rho}$. The maximum relative error between $\hat{\rho}$ and $\tilde{\rho}$ is $\lesssim 1.4\%$ over the given range of p_{fa}/N_t and N_s . Note that $\hat{\rho}$ is consistently smaller than $\tilde{\rho}$, i.e. $\hat{\rho}$ slightly overestimates the sensitivity calculated using Eq. (3.9).

B. Software injections

We next determine whether the sensitivity predicted by $\hat{\rho}$, and calculated numerically from Eq. (3.9) in the previous section, correctly predict the performance of a real gravitational-wave pulsar search pipeline. To do so, we perform software injection studies similar (with some simplifications) to the Monte Carlo simulations used to set upper limits for gravitational-wave searches of LIGO and Virgo data (see Section I for references). We use software from the LALSuite⁴ repository.

First, the `Makefakedata_v4` program is used to generate gravitational wave strain data, of timespan $N_s T_s$ (see Table I), containing Gaussian noise (with a power spectral density of $S_h = 1$), and a simulated gravitational-wave pulsar signal, as it would be observed in the LIGO Livingston detector. The strain amplitude h_0 is given by [c.f. Eq. (3.8)]

$$h_0 = \frac{5}{2} \sqrt{\frac{S_h}{T_s}} \hat{\rho} = \frac{5}{2} \sqrt{\frac{S_h}{T_s}} \tilde{\rho}, \quad (4.1)$$

where $S_h = 1$, T_s is given values from Table I, and either $\hat{\rho}$ or $\tilde{\rho}$ are calculated for $p_{\text{fd}} = 0.1$ and values for p_{fa}/N_t and N_s given in Table I. Other parameters of the simulated signal are chosen uniformly from the following ranges: $-1 \leq \xi \leq 1$, $0 \leq \psi < 2\pi$, $0 \leq \phi_0 < 2\pi$, $0 \leq \alpha < 2\pi$, $-1 \leq \sin \delta \leq 1$, $f = 100$ Hz, and 10^{-8} Hz s⁻¹ $\leq \dot{f} \leq 0$ (with higher-order spindowns set to zero).

Next, the `ComputeFStatistic_v2` program, an implementation of the \mathcal{F} -statistic, is used to perform a single-template search of the data generated by `Makefakedata_v4`, at precisely the sky position (α, δ) and frequency evolution (f, \dot{f}) of the simulated signal. (We consider the loss in sensitivity due to mismatch between the search template and signal in Section V C.) A single value of the \mathcal{F} -statistic is returned by `ComputeFStatistic_v2` and stored. For $N_s > 1$, since `ComputeFStatistic_v2` cannot combine searches of multiple data segments, we instead run `ComputeFStatistic_v2` N_s times on successive segments, of timespans T_s , and add together the N_s returned \mathcal{F} -statistic values. Thus, the (summed) \mathcal{F} -statistic values will follow χ^2 distributions with $N_s \nu = 4N_s$ degrees of freedom.

Finally, we repeat the injection procedure 5000 times. We record the number of (summed) \mathcal{F} -statistic values which are below the false alarm threshold s_{fa} , as calculated in Section III A. This fraction equals the false dismissal probability as determined by the software injections, which we denote $p_{\text{fd,inj}}$. If the sensitivity of `ComputeFStatistic_v2` is accurately estimated by $\hat{\rho}$ and/or $\tilde{\rho}$, $p_{\text{fd,inj}}$ should be close to the target false dismissal probability of $p_{\text{fd}} = 0.1$. Different choices of p_{fa}/N_t are used

TABLE I. Validation of the analytic sensitivity estimator $\hat{\rho}$, and the numerically-computed sensitivity $\tilde{\rho}$, using software injections. The injections are performed for three values of p_{fa}/N_t (listed in row 1), five combination of N_s and T_s (listed in columns 1 and 2), and once each using either $\hat{\rho}$ (top panel) or $\tilde{\rho}$ (bottom panel). Values of $\Delta\hat{\rho}$ (or $\Delta\tilde{\rho}$) and $p_{\text{fd,inj}}$ are given, for each of the three values of p_{fa}/N_t , in columns 4–5, 5–6, and 7–8 respectively.

N_s	T_s	$p_{\text{fa}}/N_t = 10^{-2}$			$p_{\text{fa}}/N_t = 10^{-6}$			$p_{\text{fa}}/N_t = 10^{-10}$		
		Analytic $\hat{\rho}$								
		$\Delta\hat{\rho}$	$p_{\text{fd,inj}}$	$\Delta\hat{\rho}$	$p_{\text{fd,inj}}$	$\Delta\hat{\rho}$	$p_{\text{fd,inj}}$	$\Delta\hat{\rho}$	$p_{\text{fd,inj}}$	
1	1	-5.0%	13.2%	-2.5%	12.4%	-5.0%	14.9%			
1	5	-2.4%	11.8%	-2.8%	12.6%	-2.9%	13.0%			
1	10	-4.9%	13.2%	-4.7%	14.0%	-2.5%	12.7%			
10	1	-1.3%	11.1%	-1.7%	11.9%	-1.7%	12.2%			
25	1	-0.16%	10.5%	-0.93%	11.4%	-1.5%	12.2%			
Numerical $\tilde{\rho}$										
		$\Delta\tilde{\rho}$	$p_{\text{fd,inj}}$	$\Delta\tilde{\rho}$	$p_{\text{fd,inj}}$	$\Delta\tilde{\rho}$	$p_{\text{fd,inj}}$	$\Delta\tilde{\rho}$	$p_{\text{fd,inj}}$	
1	1	-3.0%	11.5%	-4.1%	12.9%	-3.8%	13.1%			
1	5	-4.7%	12.4%	-3.5%	12.4%	-3.4%	12.8%			
1	10	-2.9%	11.4%	-3.0%	12.0%	-2.6%	12.1%			
10	1	-0.2%	10.1%	-1.6%	11.2%	-1.8%	11.5%			
25	1	0.36%	9.8%	-1.2%	10.9%	-0.99%	10.9%			

to set different thresholds s_{fa} , and thus simulate the sensitivity of a wide-parameter-space search over N_t templates.

Ideally, we would then refine the injected value of h_0 , increasing it if $p_{\text{fd,inj}} > p_{\text{fd}}$, decreasing it if $p_{\text{fd,inj}} < p_{\text{fd}}$, and then repeat the entire injection procedure, until $p_{\text{fd,inj}} = p_{\text{fd}}$. The relative error between the value of h_0 calculated using Eq. (4.1), and the value of h_0 arrived at by repeating the injection procedure, would then be equal to the relative error in the estimation of the sensitivity of `ComputeFStatistic_v2`. Since the injection procedure is time-consuming and computationally intensive, however, we instead re-compute $\hat{\rho}$ or $\tilde{\rho}$, as appropriate, using the false dismissal probability $p_{\text{fd,inj}}$ determined by the injections. We denote by $\Delta\hat{\rho} = |\hat{\rho}(p_{\text{fd,inj}}) - \hat{\rho}(p_{\text{fd}})|/\hat{\rho}(p_{\text{fd}})$ the relative error between the value of $\hat{\rho}$ calculated using $p_{\text{fd,inj}}$, and the value of $\hat{\rho}$ calculate using p_{fd} ; similarly for $\Delta\tilde{\rho}$. These quantities serve as a reasonable estimate of the error in the sensitivities estimated by $\hat{\rho}$ and $\tilde{\rho}$ respectively.

The results of the software injections are shown in Table I. Both $\hat{\rho}$ and $\tilde{\rho}$ underestimate, by $\lesssim 5\%$ and $\lesssim 4.7\%$ respectively, the h_0 required to achieve a false dismissal probability of 10%; except in one instance (last row of column 4), $p_{\text{fd,inj}} > 10\%$. Nevertheless, these errors are still within the $\sim 10\%$ typical calibration error of gravitational wave detectors, e.g. LIGO [57], and hence can be considered small. The difference between the predicted and actual sensitivity of `ComputeFStatistic_v2` is likely because the \mathcal{F} -statistic values returned by `ComputeFStatistic_v2` do not strictly follow a χ^2 distribution; small

⁴ Available from <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.

errors, on the order of a few percent, are introduced due to implementation details of the code [58, 59]; see Section V A and Figs. 11 and 12.

V. VALIDITY OF ASSUMPTIONS

In the previous two sections, we have derived an analytic estimator of the sensitivity of wide-parameter-space gravitational-wave pulsar searches (Section III), and validated its accuracy (Section IV). In doing so, we made certain assumptions: that the values of the detection statistic returned by the search are statistically independent, that the network of gravitational wave detectors being searched is sensitive to gravitational waves from all sky locations, and that any signal present in the data is perfectly matched by at least one of the searched templates. In this section, we investigate to what extent these assumptions are valid for real gravitational-wave pulsar search pipelines.

A. Statistical independence of templates

It was assumed, in deriving an expression for the false alarm probability in Section III A, that the N_t values of the detection statistic s returned by the search are statistically independent. This is not necessarily the case in practice. Gravitational-wave pulsar searches typically use template banks with small, e.g. 20% mismatches; templates nearby in parameter space will therefore have similarly-shaped waveforms. Matched filtering of nearby templates will therefore combine the same data with nearly-identical waveforms to produce corresponding values of the detection statistic. There is the potential, therefore, for values of the detection statistic computed from nearby templates to be correlated with each other.

To assess to what extent this effect is important, we perform four fully-coherent searches, using `ComputeFStatistic.v2` (see Section IV B) of 3 days of computer-generated Gaussian noise. The search parameter space is sky and frequency, with higher spindowns set to zero. Template banks were generated using the `gridType=2` option, which places templates over the sky using an adaptive mesh. The first three searches are performed using template banks with the following mismatches: an unrealistically-small mismatch of 1%, a realistic mismatch of 20%, and an unrealistically-high mismatch of 500%. The fourth search repeats the first search using the 1% mismatch template bank, with the following modification to the `ComputeFStatistic.v2` program: before each value of the \mathcal{F} -statistic is computed, the input Gaussian noise data is regenerated, so that each \mathcal{F} -statistic is computed from independent instances of Gaussian noise. The frequency bands of the searches are chosen such that the searches return 4.2×10^6 values of the \mathcal{F} -statistic.

We partition the returned \mathcal{F} -statistic values into 2100 blocks of $N = 2000$ values, contiguous in frequency, and

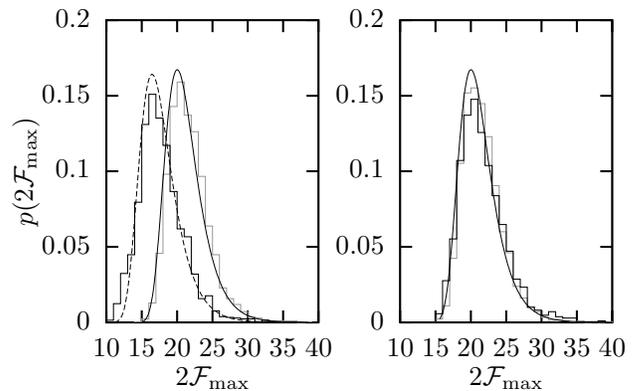


FIG. 11. Histograms of the maximum values of the \mathcal{F} -statistic obtained in blocks of 2000 templates returned by an all-sky-frequency search of 3 days of computer-generated Gaussian noise. (Left) Histogram of the 1% mismatch template bank, computed from the same Gaussian noise (black), and from regenerated independent Gaussian noise (gray). The theoretical distribution is plotted for $N = 2000$ (solid black curve), and $N = 400_t$ (dashed black curve). (Right) Histograms of the 20% mismatch template bank (black), and the 500% mismatch template bank (gray). The theoretical distribution is plotted for $N = 2000$ (solid black curve).

select the maximum value of the \mathcal{F} -statistic in each block. If all the values of the \mathcal{F} -statistic are mutually independent, the distribution of the 2100 \mathcal{F} -statistic maxima is expected to be [33]

$$p(2\mathcal{F}_{\max}) = Np(2\mathcal{F} = 2\mathcal{F}_{\max} | 2\mathcal{F} \sim \chi^2[4]) \times \left[p(2\mathcal{F} < 2\mathcal{F}_{\max} | 2\mathcal{F} \sim \chi^2[4]) \right]^{N-1}. \quad (5.1)$$

If, however, the values of the \mathcal{F} -statistic exhibit some mutual correlation, then we expect the distribution of the 2100 \mathcal{F} -statistic maxima to be well-modeled by Eq. (5.1), but with an *effective* number of statistically independent templates $N \leq 2000$. This procedure was used by the gravitational-wave search for Cassiopeia A [12] to estimate the statistical correlation of the template bank.

Figure 11 plots histograms of the \mathcal{F} -statistic maxima obtained from the four searches. In the left-hand plot of Fig. 11, we see that the \mathcal{F} -statistic values from the 1% mismatch search (without regenerating the input data) are highly correlated; they are best fitted by Eq. (5.1) with $N = 400$ effectively statistically-independent templates. When the input data is regenerated before computing each \mathcal{F} -statistic value, the distribution is much closer to the expected distribution with $N = 2000$ statistically independent templates. This demonstrates that the origin of the correlations is that the same data are being used to compute the \mathcal{F} -statistic values of nearby templates. The small differences between the computed and expected distributions is likely due to the implementation details of `ComputeFStatistic.v2` (see Section IV B). In the right-hand plot of Fig. 11, we see that the distribution of the \mathcal{F} -statistic maxima, for both the 20%

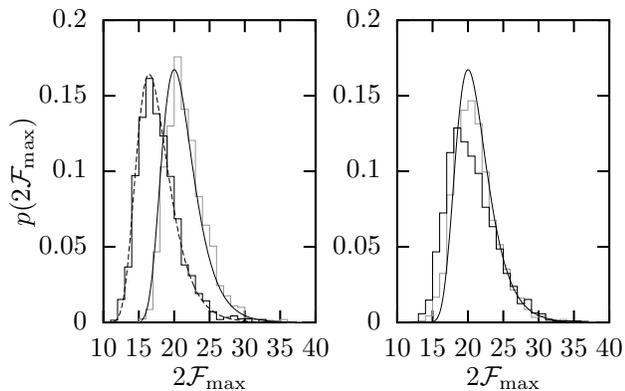


FIG. 12. Histograms of the maximum values of the \mathcal{F} -statistic obtained in blocks of 2000 templates returned by an frequency–spindown search of 7 days of computer-generated Gaussian noise. Details are the same as Fig. 11.

and 500% mismatch template banks, are both close to the expected distribution with $N = 2000$. This suggests that, while statistical correlation between \mathcal{F} -statistic values is noticeable for very closely-spaced templates, it is less significant for realistic template bank mismatches.

We repeat the above four searches, instead using 7 days of data and searching over frequency and spindown. Template banks were generated using the `gridType=8` option to `ComputeFStatistic.v2`, which places templates over frequency and spindown using a lattice. The resulting distributions of the \mathcal{F} -statistic maxima, plotted in Fig. 12, are similar to those presented in Fig. 11, except that the 20% mismatch distribution is less well-fitted by the $N = 2000$ distribution. The same template bank was used in the search for Cassiopeia A [12], which found the number of statistically-independent templates to be $\sim 88\%$ of the total number of templates.

While it is difficult to precisely quantify the effect of statistical correlations between templates on search sensitivity, for template banks with realistic mismatches the effect is likely to be small. As may be deduced from Fig. 6, a change in N_t of e.g. 10% does not significantly alter the predicted $\hat{\rho}$. Given that, as shown in Table I, both analytic and numerical sensitivity estimators accurately predict the sensitivity of a real gravitational-wave search pipeline to $\sim 5\%$, it is reasonable to neglect the smaller effect due to statistical correlations between templates.

B. Isotropic sensitivity of detector network

We now investigate whether it is reasonable to assume that a real gravitational-wave detector network, e.g. that of the LIGO and Virgo detectors, can be modeled by a network which is isotropically sensitive to gravitational waves arriving from all directions. While this assumption is reasonable for all-sky searches, since averaging

over sky position is equivalent to averaging over detector orientation (see Appendix A), it may not be reasonable for searches targeting a single sky position.

In Section IV A, we numerically solved Eq. (3.9) for $\tilde{\rho} = \sqrt{\langle \rho^2 \rangle}$, assuming an isotropically-sensitive detector network, and averaging over the sky position (α, δ) and polarization angle ψ ; we denote the values of $\hat{\rho}$ thus obtained by $\tilde{\rho}_{\text{iso}}$. We now solve Eq. (3.9), using the same algorithm detailed in Section IV A, for three detector networks: the LIGO Livingston detector, the two LIGO (Livingston and Hanford) detectors, and the three-detector network comprising the LIGO and Virgo detectors; detector locations and orientations are taken from [60]. We assume that all detectors in the network are equally sensitive, since we are concerned only with the effect of the geographic configuration of the network. We consider the following five search scenarios:

1. A search using a long data segment length of $T_s \gg \text{days}$, covering the entire sky. For large T_s , the signal SNR becomes independent of α (see Appendix A). We compute $\tilde{\rho}$ averaged over the remaining parameters, $\sin \delta$ and ψ .
2. A search using a long data segment length of $T_s \gg \text{days}$, and targeting a source at a known sky position but with an unknown polarization. We compute $\tilde{\rho}$ for a grid of 41 linearly-spaced values of $\sin \delta$ in the range $-1 \leq \sin \delta \leq 1$, and average over ψ .
3. A search using a long data segment length of $T_s \gg \text{days}$, and targeting a source with both a known sky position and polarization. We compute $\tilde{\rho}$ for the same grid of $\sin \delta$ values as in scenario 2, and a grid of 10 linearly-spaced values of ψ in the range $-\pi/4 \leq \psi < \pi/4$.
4. A search using a short data segment length of $T_s = 0.5 \text{ days}$, and targeting a source at a known sky position but with an unknown polarization. For T_s shorter than a day, the signal SNR is a function of α , $\sin \delta$, and ψ . We compute $\tilde{\rho}$ for a grid of 20 linearly-spaced values of α in the range $0 \leq \alpha < 2\pi$, 19 linearly-spaced values of $\sin \delta$ in the range $-1 < \sin \delta < 1$, and average over ψ .
5. A search using a short data segment length of $T_s = 0.5 \text{ days}$, and targeting a source with both a known sky position and polarization. We compute $\tilde{\rho}$ for the same grids of α and $\sin \delta$ values as in scenario 4, and the same grid of ψ values as in scenario 3.

For each of the above search scenarios, we record the minimum, mean, standard deviation, and maximum of the relative error $|\tilde{\rho} - \tilde{\rho}_{\text{iso}}|/\tilde{\rho}_{\text{iso}}$ and relative difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$, over the grids of α , $\sin \delta$, and ψ values given above, and over the grids of p_{fa}/N_t and N_s values given in Section IV A.

TABLE II. Relative errors and differences in the sensitivity estimated assuming an isotropically-sensitive detector network, under five search scenarios (see the text for details). For each detector network (column 1), the mean and standard deviation of the relative error $|\tilde{\rho} - \tilde{\rho}_{\text{iso}}|/\tilde{\rho}_{\text{iso}}$ (column 2), and the minimum and maximum difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ (columns 3 and 4) are listed. The abbreviations *L*, *LH*, and *LHV* indicate the LIGO Livingston detector, the LIGO detector network, and the LIGO–Virgo detector network respectively. Each block of the table corresponds to a search scenario, and a summary of each scenario is given just above each block.

Network	Error $ \tilde{\rho} - \tilde{\rho}_{\text{iso}} /\tilde{\rho}_{\text{iso}}$ mean \pm stdv.	Difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ minimum maximum	
1. $T_s \rightarrow \infty$; averaged $\sin \delta$, and ψ			
L	0.0058 \pm 0.0037%	-0.018%	+0.017%
LH	0.0062 \pm 0.0042%	-0.02%	+0.017%
LHV	0.0063 \pm 0.0044%	-0.02%	+0.014%
2. $T_s \rightarrow \infty$; known $\sin \delta$; averaged ψ			
L	3.4 \pm 1.8%	-4.5%	+6.5%
LH	2.6 \pm 1.3%	-4.9%	+3.8%
LHV	3.7 \pm 2.2%	-7.5%	+5.4%
3. $T_s \rightarrow \infty$; known $\sin \delta$ and ψ			
L	4.3 \pm 2.4%	-11.0%	+6.5%
LH	2.8 \pm 2.2%	-4.9%	+8.7%
LHV	3.8 \pm 2.4%	-7.5%	+9.3%
4. $T = 0.5$ days; known α , $\sin \delta$; averaged ψ			
L	5.1 \pm 3.6%	-6.5%	+16.0%
LH	4.7 \pm 3.0%	-10.0%	+10.0%
LHV	4.1 \pm 2.3%	-8.8%	+7.4%
5. $T = 0.5$ days; known α , $\sin \delta$, and ψ			
L	6.3 \pm 4.7%	-14.0%	+20.0%
LH	5.6 \pm 3.7%	-10.0%	+18.0%
LHV	4.4 \pm 3.1%	-9.9%	+15.0%

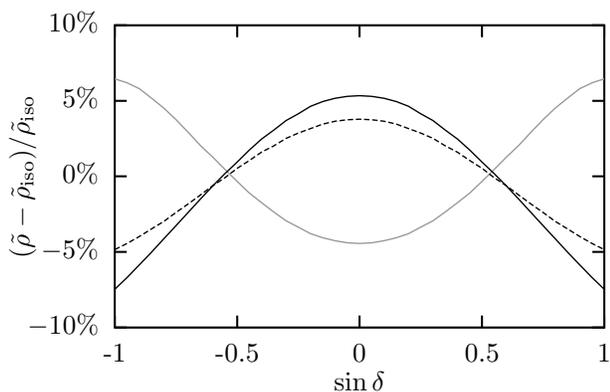


FIG. 13. Mean relative difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ in the sensitivity estimated assuming an isotropically-sensitive detector network, for search scenario 2 (see the text), as a function of $\sin \delta$, for the LIGO Livingston detector (gray), the LIGO detector network (black, dashed), and the LIGO–Virgo detector network (black).

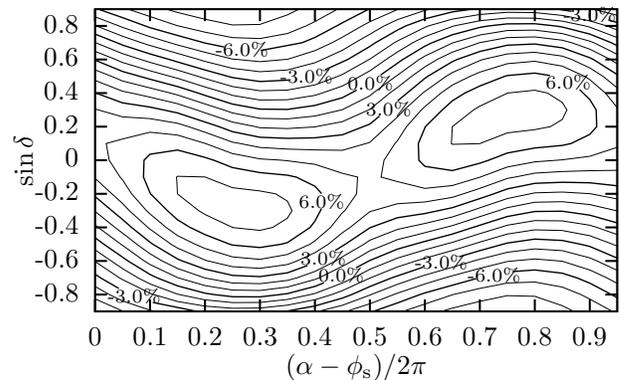


FIG. 14. Mean relative difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ in the sensitivity estimated assuming an isotropically-sensitive detector network, for scenario 4 (see the text), as a function of α and $\sin \delta$, for the LIGO–Virgo detector network. The sidereal time at Greenwich at the mid-point of the observation time is denoted by ϕ_s .

Table II shows the mean and standard deviation of the relative error, and the minimum and maximum relative difference obtained under the five search scenarios. For an all-sky search (scenario 1), the assumption of an isotropic detector network is an excellent one. For the remaining search scenarios, the error in assuming an isotropically-sensitive detector network increases as T_s is reduced, and as more parameters are set to fixed values. While the mean error is limited to $\lesssim 6.3\%$ for all four scenarios, the maximum difference can be up to 20% for a single detector, although it reduces to 15% for a three-detector network. There are only a few (potential) gravitational-wave sources for which the polarization angle may be determined with any accuracy; noted examples are the Crab and Vela pulsars, where X-ray observations of the pulsar wind nebula provide information on the pulsar’s orientation [10, 11]. Thus, scenarios 3 and 5 are less likely to arise in practice (but see Section VI for a discussion of these scenarios in relation to the Power-Flux upper limit procedure). We conclude that, while the error in assuming an isotropically-sensitive detector network may be acceptable in many cases, it can be significant for particular choices of fixed search parameters.

Figure 13 plots the mean relative difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ for scenario 2, as a function of $\sin \delta$, for the three detector networks considered. The difference between $\tilde{\rho}$ and $\tilde{\rho}_{\text{iso}}$ is smallest at $\sin \delta \sim 0.5$ – 0.55 , the approximate latitudes of the three detectors (~ 30 – 46°); a signal originating from these declinations would therefore be located at the approximate maximum sensitivities of the detectors. For the same reason, a single detector at low latitude (LIGO Livingston) is more sensitive to signals arriving at low declinations, explaining the increased sensitivity relative to that of an isotropic detector network. The addition of two detectors at higher latitudes (LIGO Hanford and Virgo) shifts the network configura-

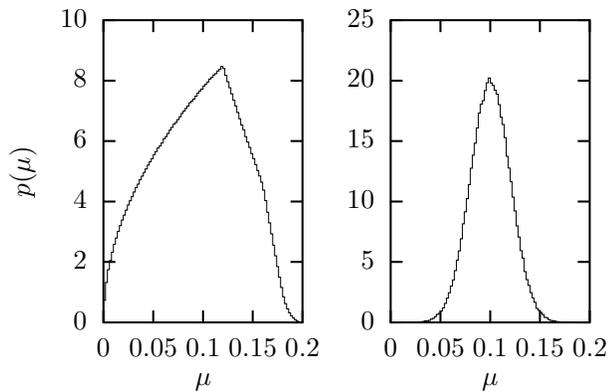


FIG. 15. Histograms of the mismatch distribution of a 3-dimensional body-centered cubic lattice template bank (left), and an *ad-hoc* Gaussian mismatch distribution (right).

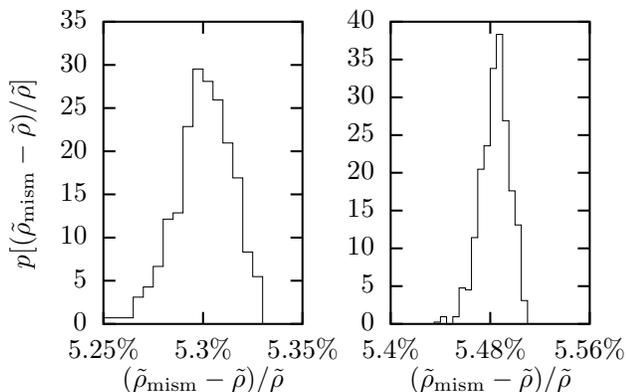


FIG. 16. Histograms of the relative difference between $\tilde{\rho}_{\text{mism}}$, calculated using the lattice template bank (left) and Gaussian mismatch distributions (right), and $\tilde{\rho}$ (zero mismatch) over the ranges $10^{-15} \leq p_{\text{fa}}/N_t \leq 10^{-2}$ and $1 \leq N_s \leq 10^4$.

tion towards improved sensitivity at higher declinations.

Figure 14 plots the mean relative difference $(\tilde{\rho} - \tilde{\rho}_{\text{iso}})/\tilde{\rho}_{\text{iso}}$ for scenario 4, as a function of α and $\sin \delta$, for the LIGO–Virgo network. Consistent with Fig. 13, the network is more sensitive to signals arriving at the poles, and less sensitive to signals arriving at the equator, than an isotropically-sensitive network. The change in sensitivity as a function of α and δ is comparable to previous studies of gravitational-wave detector network configuration; see e.g. Fig 3 in [61].

C. Template bank mismatch

Finally, we consider the reduction in sensitivity due to mismatch between the searched templates and any signal which may be present. The mismatch μ is related to the difference between the SNR $\rho^2(\mathbf{p})$ of a signal with parameters \mathbf{p} , and the SNR $\rho^2(\mathbf{p}')$ recovered by a search template with mismatched parameters \mathbf{p}' . It is given by

(e.g. [44])

$$\mu = \frac{\rho^2(\mathbf{p}) - \rho^2(\mathbf{p}')}{\rho^2(\mathbf{p})}. \quad (5.2)$$

Substituting $\sqrt{\langle \rho^2 \rangle} R$ for $\rho(\mathbf{p})$ and $\sqrt{\langle \rho'^2 \rangle} R_{\text{mism}}$ for $\rho(\mathbf{p}')$, we find that $R_{\text{mism}}^2 = R^2(1 - \mu)$; thus the effect of mismatch is to reduce R^2 by a factor $1 - \mu$. For a large template bank, μ can be considered as a random variable drawn from a distribution $p(\mu)$ characterizing the geometric arrangement of the templates in parameter space. Equation (3.9) is then modified to additionally average over mismatch:

$$p_{\text{fd}} = \left\langle p(s \leq s_{\text{fa}} | s \sim \chi^2[N_s \nu, N_s \langle \rho^2 \rangle R^2(1 - \mu)]) \right\rangle_{\alpha, \sin \delta, \psi, \xi, \mu}, \quad (5.3)$$

where $\langle \dots \rangle_{\mu} = \int_0^{\mu_{\text{max}}} d\mu p(\mu) \dots$, and μ_{max} is the maximum allowed mismatch.

We solve Eq. (5.3) numerically for $\tilde{\rho}_{\text{mism}} = \sqrt{\langle \rho^2 \rangle}$, using a modified version of the algorithm described in Section IV A. We compute $\tilde{\rho}_{\text{mism}}$ for two examples of $p(\mu)$: the mismatch distribution for a template bank constructed using a 3-dimensional body-centered cubic lattice with $\mu_{\text{max}} = 20\%$, from [33], and an *ad hoc* Gaussian mismatch distribution with a mean of 10% and standard deviation of 2%, restricted to the range $0 \leq \mu \leq 20\%$. Histograms of these distributions are plotted in Fig. 15; mismatch distributions for lattice template banks in other dimensions are plotted in [62]. The lattice template bank mismatch distribution has a mean mismatch of 10% and standard deviation of 4%.

We find that taking mismatch into account increases $\tilde{\rho}_{\text{mism}}$, relative to the equivalent (zero-mismatch) $\tilde{\rho}$ computed in Section IV A, by on average $5.3 \pm 0.10\%$ for the lattice template bank mismatch distribution, and $5.5 \pm 0.10\%$ for the Gaussian distribution. This effect is of the same magnitude as the error in the analytic/numerical estimators compared to software injections (see Table I). We note that, by substituting $R_0 \sqrt{1 - \langle \mu \rangle}$ (where $\langle \mu \rangle$ is the mean of the mismatch distribution) for R_0 in Eq. (3.31), we can predict the observed relative increase in $\tilde{\rho}_{\text{mism}}$ reasonably accurately; for $\langle \mu \rangle = 10\%$, $\tilde{\rho}$ is increased by 5.4%. This suggests that it is reasonable to model the effect of template bank mismatch as a uniform reduction in recovered SNR (by $\sqrt{1 - \langle \mu \rangle}$), and that the mean of the mismatch distribution is sufficient to quantify this reduction; this approach is also taken in [31]. Figure 16 plots histograms of the relative difference between $\tilde{\rho}_{\text{mism}}$ and $\tilde{\rho}$; the histograms are narrow, with no long tails, confirming that the most typical reduction in SNR is close to the average reduction.

Figure 4 plots the distribution of $R_{\text{iso,mism}}^2 = R_{\text{iso}}^2(1 - \mu)$, where μ are drawn from the lattice template bank mismatch distribution. Relative to the zero-mismatch distribution, $R_{\text{iso,mism}}^2$ is reduced for circularly-polarized

signals (at high $R_{\text{iso,mism}}^2$), and exhibits a broader peak at linearly-polarized signals (at low $R_{\text{iso,mism}}^2$).

VI. DISCUSSION

In this paper, we develop a new analytic method of estimating the sensitivity of wide-parameter space searches for gravitational-wave pulsars. The new method avoids the unphysical assumption of constant-SNR signals, used by previous analytic sensitivity estimation methods, and hence can accurately (to $\lesssim 5\%$) reproduce the sensitivity estimated using Monte Carlo simulations, but without their computational cost. Additional effects, not included in the analytic model but which may be important for real gravitational-wave pulsar searches, are investigated: the statistical correlation between values of the detection statistic due to closely-spaced template banks; the deviation of the directional sensitivity of real gravitational wave detector networks from an ideal, isotropically-sensitive network; and the loss in SNR due to mismatch between a signal and the searched templates. While the first of these effects can be neglected, the remaining two may need to be accounted for where accurate sensitivity estimates are required.

It is important to note that an accurate prediction of the *absolute* sensitivity of a search may not always be necessary. For studies of the optimal configuration of a hierarchical search (e.g. [27, 31]), the *relative* sensitivity of different search schemes is usually more important. For these studies, use of the constant-SNR sensitivity estimator $\bar{\rho}$ may be acceptable; Fig. 7 shows that $\bar{\rho}$ reproduces to $\lesssim 5\%$ the correct scaling (i.e. that of $\hat{\rho}$) of sensitivity with respect to p_{fa}/N_t and N_s . On the other hand, given the similarity between the expressions for $\bar{\rho}$ and $\hat{\rho}$ [compare Eqs. (3.13) and (3.31)], these studies may also be able to utilise the more accurate sensitivity scaling of $\hat{\rho}$ with little modification.

This work has restricted its attention to detection statistics which follow χ^2 distributions, which does not cover all gravitational-wave search methods. In particular, the Hough semi-coherent method [26] computes a statistic, the number count, which is the number of coherently-analyzed segments where the coherent detection statistic (which may either be power or the \mathcal{F} -statistic) exceeds a set threshold. It follows that the number count follows a binomial distribution. Since, however, the Hough method is generally used to analyze a large number of data segments (i.e. N_s is large), the binomial distribution can generally be approximated by a normal distribution (see [26]), an approximation also used in Section III. Thus, the analytic sensitivity estimation method developed here may be applicable to the Hough method with minor modifications.

Searches performed using the PowerFlux method [14, 20, 22] employ a slightly different formulation of the frequentist procedure described in Section II C. (To distinguish these two frequentist procedures, in the follow-

ing discussion we refer to the upper limits produced by the frequentist procedure described in Section II C as *population-averaged* upper limits.) An upper limit on h_0 is set, using the Feldman-Cousins method of confidence interval construction [63], for every searched parameter (typically sky position, frequency, and first-order spindown), and assuming either linearly-polarized signals ($\xi = 0$, with a range of ψ) or a circularly-polarized signal ($|\xi| = 1$). For each polarization case, the maximum value of h_0 over the searched parameters is chosen. Two upper limits are then quoted: a best-case upper limit, which assumes circular polarization; and a worst-case⁵ upper limit, where h_0 is further maximized over the linear polarizations.

The PowerFlux upper limit procedure differs from the population-averaged upper limit procedure in two respects. (In the following comparison, we assume that the data being searched is relatively free of non-Gaussian instrumental noise.) First, while the population-average procedure first selects the maximum value of the detection statistic s over the parameter space, and then computes an upper limit on h_0 using s , the PowerFlux procedure first computes an upper limit on h_0 for every value of s , assuming a fixed polarisation, and then selects the maximum upper limit. By computing upper limits assuming a fixed sky position (α, δ), corresponding to each value of s , and a fixed ψ , PowerFlux is essentially performing search scenario 5 from Table II (although with different search parameters, e.g. $T_s = 30$ minutes, $p_{\text{fd}} = 0.05$). From the maximum relative difference in sensitivity quoted in Table II for this scenario, we estimate that the effect of maximising over h_0 instead of over s raises PowerFlux upper limits by $\sim 10\text{--}20\%$ relative to population-averaged upper limits.

Second, the population-average procedure produces a single upper limit averaged over all parameters; the PowerFlux procedure instead quotes upper limits for best-case (circular) and worst-case (linear) polarizations. By setting upper limits at fixed polarizations (i.e. fixed values of ξ), the PowerFlux procedure is setting upper limits on a population of signals with a fixed value of R^2 , i.e. a fixed SNR. We can therefore predict the upper limits on h_0 set by the PowerFlux procedure by replacing $\langle \rho^2 \rangle R^2$ with $\rho_{\text{P.F.}}^2 R_{\text{iso}}^2(\xi)$ in Eq. (3.9); it follows that $\rho_{\text{P.F.}}$ is given by $\bar{\rho}/R_{\text{iso}}(\xi)$, where $\bar{\rho}$ is given by Eq. (3.13). The best- and worst-case sensitivities are then given by

$$\rho_{\text{P.F. best}} = \frac{\bar{\rho}}{R_{\text{iso}}(|\xi| = 1)} = \sqrt{\frac{2}{5}} \bar{\rho} \quad (6.1)$$

$$\rho_{\text{P.F. worst}} = \frac{\bar{\rho}}{R_{\text{iso}}(\xi = 0)} = \frac{4}{\sqrt{5}} \bar{\rho} \quad (6.2)$$

Ratios of Power-Flux best- and worst-case sensitivities to the sensitivity predicted by $\hat{\rho}$ range over $0.39 \leq$

⁵ Actually, PowerFlux worst-case upper limits, as quoted in e.g. [14], also account for the worst-case mismatch; in this discussion however we assume that no mismatch is present.

$\rho_{\text{P.F. best}}/\hat{\rho} \leq 0.46$, and $1.1 \leq \rho_{\text{P.F. worst}}/\hat{\rho} \leq 1.3$, for the range of $\bar{\rho}/\hat{\rho}$ plotted in Fig. 7. An injection study by Dergachev [64] found ratios consistent with these ranges: $\rho_{\text{P.F. best}}/\hat{\rho} \sim 0.43$, and $\rho_{\text{P.F. worst}}/\hat{\rho} \sim 1.2$. Further investigation into the differences between the PowerFlux and population-averaged upper limits methods would facilitate direct comparison of the upper limits produced by different gravitational wave searches, e.g. between [22] and [13].

Recent searches for gravitational waves from known pulsars [10, 11] use Bayesian inference to set upper limits. The frequentist method described in this paper constructs confidence intervals which are derived from the probability of the *data*, given a particular hypothesis e.g. that a signal is or is not present in the data. In contrast, Bayesian inference directly calculates the probability of the *hypothesis*, given the particular data that were observed. Despite their very different interpretations, a study in [65] found that the frequentist and Bayesian methods produce similar upper limits, in the limit of large signal amplitudes. Further research is needed to understand fully the relationship between these two methods.

ACKNOWLEDGMENTS

I thank Bruce Allen, Joseph Betzwieser, Vladimir Dergachev, Evan Goetz, David Keitel, Ben Owen, Chris Messenger, Reinhard Prix, Letizia Sammut, and members of the LIGO Scientific Collaboration's Continuous Wave Working Group for many informative discussions, and for helpful comments on the manuscript. Numerical simulations were performed on the ATLAS computer cluster of the Max-Planck-Institut für Gravitationsphysik. This paper has LIGO document number P1100151-v3.

Appendix A: Antenna-pattern functions and their averages

We present expressions for the antenna-pattern functions $F_+(t)$ and $F_\times(t)$ of an interferometric detector, and the averages of $F_+(t)^2$ and $F_\times(t)^2$ over various parameters; see also the references in Section II A.

The antenna-pattern functions can be written in terms of a time-dependent detector response matrix $\mathbf{R}(t)$ and constant polarization matrices \mathbf{H}_+ and \mathbf{H}_\times :

$$F_+(t) = \text{tr}(\mathbf{R}(t)^T \mathbf{H}_+), \quad F_\times(t) = \text{tr}(\mathbf{R}(t)^T \mathbf{H}_\times), \quad (\text{A1})$$

where tr is the matrix trace and T denotes transposition. We implicitly assume a coordinate system at rest with respect to the gravitational radiation, e.g. the Solar System barycenter.

Let $\hat{\mathbf{n}}_1(t)$ and $\hat{\mathbf{n}}_2(t)$ be unit vectors along the interferometer's arms, such that the cross product $\hat{\mathbf{n}}_1(t) \times \hat{\mathbf{n}}_2(t)$ points toward zenith. The directions along which the interferometer is insensitive to any gravitational radiation

are given by the unit vectors $\pm\hat{\mathbf{a}}(t)$ and $\pm\hat{\mathbf{b}}(t)$, where

$$\hat{\mathbf{a}}(t) = \frac{\hat{\mathbf{n}}_1(t) - \hat{\mathbf{n}}_2(t)}{2 \sin(\zeta/2)}, \quad \hat{\mathbf{b}}(t) = \frac{\hat{\mathbf{n}}_1(t) + \hat{\mathbf{n}}_2(t)}{2 \cos(\zeta/2)}, \quad (\text{A2})$$

and ζ is the angle between the interferometer arms. Note that $\hat{\mathbf{a}}(t)$ and $\hat{\mathbf{b}}(t)$ are orthogonal, i.e. $\hat{\mathbf{a}}(t) \cdot \hat{\mathbf{b}}(t) = 0$. Assuming that the gravitational wavelength is much larger than the interferometer arm-length, as is the case for ground-based detectors, the response matrix \mathbf{R} is given in terms of these vectors by (e.g. [66])

$$\mathbf{R}(t) = \frac{\sin \zeta}{2} [\hat{\mathbf{a}}(t) \otimes \hat{\mathbf{b}}(t) + \hat{\mathbf{b}}(t) \otimes \hat{\mathbf{a}}(t)], \quad (\text{A3})$$

where \otimes is the vector outer product. The vectors can be expressed in terms of time-independent components:

$$\hat{\mathbf{a}}(t) = \hat{\mathbf{a}}_1 \cos \Omega_s(t - t_0) + \hat{\mathbf{a}}_2 \sin \Omega_s(t - t_0) + \hat{\mathbf{a}}_3, \quad (\text{A4})$$

where the Earth rotates about the unit vector $\hat{\Omega}_s$ in a right-handed sense with angular frequency Ω_s , and $\hat{\mathbf{a}}_1 = \hat{\mathbf{a}}_2 \times \hat{\Omega}_s$, $\hat{\mathbf{a}}_2 = \hat{\Omega}_s \times \hat{\mathbf{a}}_0$, $\hat{\mathbf{a}}_3 = (\hat{\Omega}_s \cdot \hat{\mathbf{a}}_0) \hat{\Omega}_s$, and $\hat{\mathbf{a}}_0 = \hat{\mathbf{a}}(t = t_0)$; similarly for $\hat{\mathbf{b}}(t)$.

The polarization matrices \mathbf{H}_+ and \mathbf{H}_\times may also be written in terms of certain vectors [58, 67]. The directions along which a gravitational wave creates no space-time perturbation are given by the mutually orthogonal unit vectors $\pm\hat{\mathbf{x}}_+$ and $\pm\hat{\mathbf{y}}_+$, for a purely plus-polarized wave, and $\pm\hat{\mathbf{x}}_\times$ and $\pm\hat{\mathbf{y}}_\times$, for a purely cross-polarized wave. The two sets of vectors are related by

$$\hat{\mathbf{x}}_+ \sqrt{2} = \hat{\mathbf{x}}_\times - \hat{\mathbf{y}}_\times, \quad \hat{\mathbf{y}}_+ \sqrt{2} = \hat{\mathbf{x}}_\times + \hat{\mathbf{y}}_\times. \quad (\text{A5})$$

The cross-polarisation vectors are given by

$$\hat{\mathbf{x}}_\times = \hat{\xi} \cos \psi + \hat{\eta} \sin \psi, \quad \hat{\mathbf{y}}_\times = \hat{\eta} \cos \psi - \hat{\xi} \sin \psi, \quad (\text{A6})$$

where $\hat{\xi} = (\hat{\mathbf{n}} \times \hat{\Omega}_s) / \|\hat{\mathbf{n}} \times \hat{\Omega}_s\|$, $\hat{\eta} = \hat{\xi} \times \hat{\mathbf{n}}$, and $-\hat{\mathbf{n}} = \hat{\mathbf{x}}_+ \times \hat{\mathbf{y}}_+ = \hat{\mathbf{x}}_\times \times \hat{\mathbf{y}}_\times$ is the direction of propagation of the wave. The polarization matrices are then given by

$$\mathbf{H} = \hat{\mathbf{x}} \otimes \hat{\mathbf{y}} + \hat{\mathbf{y}} \otimes \hat{\mathbf{x}}, \quad (\text{A7})$$

where we henceforth take $F(t)$, \mathbf{H} , $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ to mean either $F_+(t)$, \mathbf{H}_+ , $\hat{\mathbf{x}}_+$, $\hat{\mathbf{y}}_+$ or $F_\times(t)$, \mathbf{H}_\times , $\hat{\mathbf{x}}_\times$, $\hat{\mathbf{y}}_\times$, as appropriate.

Combining Eqs. (A1), (A3), and (A7), we find that the antenna-pattern functions F can be written as:

$$\frac{F(t)}{\sin \zeta} = [\hat{\mathbf{a}}(t) \cdot \hat{\mathbf{x}}] [\hat{\mathbf{b}}(t) \cdot \hat{\mathbf{y}}] + [\hat{\mathbf{a}}(t) \cdot \hat{\mathbf{y}}] [\hat{\mathbf{b}}(t) \cdot \hat{\mathbf{x}}]. \quad (\text{A8})$$

Note that when $\hat{\mathbf{a}}(t)$, $\hat{\mathbf{b}}(t)$ and $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ are parallel to each other (e.g. when $\hat{\mathbf{a}}(t)$ is parallel to $\hat{\mathbf{x}}$, implying $\hat{\mathbf{b}}(t)$ is parallel to $\hat{\mathbf{y}}$), $F(t)$ achieves its maximal values of $\pm \sin \zeta$.

The averages of $F(t)^2$ over isotropic source sky position (α, δ) and orientation ψ , and over isotropic detector location (Φ_s, λ) and orientation γ , are given by:

$$\langle F^2 \rangle_{\alpha, \sin \delta, \psi} = \int_0^{2\pi} d\alpha \int_{-1}^1 d(\sin \delta) \int_0^{2\pi} d\psi \frac{F^2}{8\pi^2}, \quad (\text{A9})$$

$$\langle F^2 \rangle_{\Phi_s, \sin \lambda, \gamma} = \int_0^{2\pi} d\Phi_s \int_{-1}^1 d(\sin \lambda) \int_0^{2\pi} d\gamma \frac{F^2}{8\pi^2}, \quad (\text{A10})$$

where $\Phi_s = \Omega_s(t - t_0)$ is the local sidereal time at the detector. Note that these two equations can be transformed into each other if one makes the exchanges $\alpha \leftrightarrow \Phi_s$, $\delta \leftrightarrow \lambda$, and $\psi \leftrightarrow \gamma$. Note too that, since $\langle F^2 \rangle_{\alpha, \sin \delta, \psi}$ averages over all possible orientations of $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, it must be invariant to the orientation of $\hat{\mathbf{a}}(t)$ and $\hat{\mathbf{b}}(t)$; for the same reason, $\langle F^2 \rangle_{\Phi_s, \sin \lambda, \gamma}$ must be invariant to the orientation of $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$. Finally, note that $F(t)$ is invariant if one exchanges $\hat{\mathbf{a}}(t)$ and $\hat{\mathbf{b}}(t)$ for $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, and vice versa. We conclude that the averages $\langle F^2 \rangle_{\alpha, \sin \delta, \psi}$ and $\langle F^2 \rangle_{\Phi_s, \sin \lambda, \gamma}$ must be equal. To calculate e.g. $\langle F^2 \rangle_{\alpha, \sin \delta, \psi}$, we choose $\hat{\mathbf{a}}(t) = (1, 0, 0)$, $\hat{\mathbf{b}}(t) = (0, 1, 0)$ for convenience, and substituting $\hat{\mathbf{n}} = (\cos \alpha \cos \delta, \sin \alpha \cos \delta, \sin \delta)$ into Eqs. (A6) obtain

$$\langle F^2 \rangle_{\alpha, \sin \delta, \psi} = \langle F^2 \rangle_{\Phi_s, \sin \lambda, \gamma} = \frac{1}{5} \sin^2 \zeta . S \quad (\text{A11})$$

To compute the average of $F(t)^2$ over time, $\langle F^2 \rangle_t$, we write $\hat{\mathbf{a}}(t) = \sum_{i=1}^3 c_i(t) \hat{\mathbf{a}}_i$, where $c_1(t) = \cos \Omega_s(t - t_0)$, $c_2(t) = \sin \Omega_s(t - t_0)$, $c_3(t) = 1$, and t_0 is chosen to be the mid-point of the observation time, at which $\hat{\mathbf{a}}_0$ is calculated; similarly for $\hat{\mathbf{b}}(t)$. We can now write

$$\langle F^2 \rangle_t = \sin^2 \zeta \sum_{i,j=1}^3 \frac{1}{T} \int_{-T/2}^{T/2} dt [c_i(t) c_j(t) B_{ij}]^2, \quad (\text{A12})$$

where we define

$$B_{ij} = (\hat{\mathbf{a}}_i \cdot \hat{\mathbf{x}}) (\hat{\mathbf{b}}_j \cdot \hat{\mathbf{y}}) + (\hat{\mathbf{a}}_i \cdot \hat{\mathbf{y}}) (\hat{\mathbf{b}}_j \cdot \hat{\mathbf{x}}). \quad (\text{A13})$$

To write down the result of the integration in Eq. (A12), we first define

$$J_{\pm i} = B_{i' i''} \pm B_{i'' i'}, \quad (\text{A14a})$$

$$K_{\pm i} = B_{i' i''} \pm B_{i'' i'}, \quad (\text{A14b})$$

$$S_{\pm i} = \sqrt{B_{i' i''} B_{i'' i'} \pm B_{i' i''} B_{i'' i'}}, \quad (\text{A14c})$$

where $i = 1, 2, 3$, and \prime denotes a permutation operation: $1 = 3' = 2''$, $2 = 1' = 3''$, and $3 = 2' = 1''$. We also note the following identities:

$$J_{-1} = -B_{11}, \quad J_{-2} = B_{22}, \quad J_{+3} = B_{33}. \quad (\text{A15})$$

Employing Eqs. (A14) and (A15), we find that

$$\langle F^2 \rangle_t = \sin^2 \zeta \sum_{n=0}^4 C_n \text{sinc} \frac{n \Omega_s T}{2}, \quad (\text{A16})$$

where

$$C_0 = \frac{19J_{+3}^2 + K_{-3}^2}{8} + \frac{K_{+1}^2 + K_{+2}^2 + S_{-3}^2}{2}, \quad (\text{A17a})$$

$$C_1 = \left(\frac{5}{2} J_{+3} - J_{-1} \right) K_{+2} + \frac{1}{2} K_{+1} K_{+3}, \quad (\text{A17b})$$

$$C_2 = \frac{3}{2} J_{-3} J_{+3} + \frac{1}{2} (K_{+2}^2 - K_{+1}^2), \quad (\text{A17c})$$

$$C_3 = \frac{1}{2} (J_{-3} K_{+2} - K_{+1} K_{+3}), \quad (\text{A17d})$$

$$C_4 = \frac{1}{8} (J_{-3}^2 - K_{+3}^2). \quad (\text{A17e})$$

For $T \gtrsim 1$ sidereal day, $\langle F^2 \rangle_t \approx C_0$, which may also be written as

$$C_0 = 9a_{\parallel}^2 b_{\parallel}^2 x_{\parallel}^2 y_{\parallel}^2 + 8a_{\perp}^2 b_{\perp}^2 x_{\perp}^2 y_{\perp}^2 + 2(a_{\parallel}^2 b_{\perp}^2 + a_{\perp}^2 b_{\parallel}^2 - a_{\parallel}^2 b_{\parallel}^2) \times (x_{\parallel}^2 y_{\perp}^2 + x_{\perp}^2 y_{\parallel}^2 - x_{\parallel}^2 y_{\parallel}^2), \quad (\text{A18})$$

where

$$a_{\parallel} = \boldsymbol{\Omega}_s \cdot \hat{\mathbf{a}}_0, \quad 2a_{\perp}^2 = 1 - a_{\parallel}^2, \quad (\text{A19})$$

$$x_{\parallel} = \boldsymbol{\Omega}_s \cdot \hat{\mathbf{x}}, \quad 2x_{\perp}^2 = 1 - x_{\parallel}^2, \quad (\text{A20})$$

and similarly for b_{\parallel} , b_{\perp} , y_{\parallel} and y_{\perp} . Note that a_{\parallel} , b_{\parallel} , x_{\parallel} , and y_{\parallel} (and consequentially a_{\perp} , b_{\perp} , x_{\perp} , and y_{\perp}) are independent of the Earth's angular displacement. We deduce that C_0 cannot depend on parameters defined relative to the Earth's angular displacement, which are the detector's local sidereal time Φ_s , and the source's right ascension α .

Appendix B: Limited inverse of the central χ^2 distribution

The right-hand side of Eq. (3.4) is equal to the normalized upper incomplete gamma function:

$$p(s > s_{\text{fa}} | s \sim \chi^2[N_s \nu, 0]) = \frac{\Gamma(N_s \nu / 2, s_{\text{fa}} / 2)}{\Gamma(N_s \nu / 2)}. \quad (\text{B1})$$

We use a limiting form of the asymptotic inverse of the incomplete gamma function given in [55], which is sufficiently accurate for small values of p_{fa}/N_t . It gives the following expression for s_{fa} :

$$s_{\text{fa}} = N_s \nu \lambda(\eta), \quad (\text{B2})$$

where

$$\eta = \eta_0 + \frac{2}{N_s \nu \eta_0} \ln \left(\frac{\eta_0}{\lambda(\eta_0) - 1} \right), \quad (\text{B3})$$

$$\eta_0 = \frac{2}{\sqrt{N_s \nu}} \text{erfc}^{-1}(2p_{\text{fa}}/N_t), \quad (\text{B4})$$

$$\lambda(x) = -W_{-1} \left[-\exp \left(-1 - \frac{x^2}{2} \right) \right], \quad x \geq 0, \quad (\text{B5})$$

and W_{-1} is the -1 branch of the Lambert W -function (see e.g [68]). For the purposes of this work, $\lambda(x)$ is well approximated by the following two functions [see [69], Eqs. (4.13.6) and (4.13.11)]:

$$\lambda_1(x) = 1 + x + \frac{x^2}{3} + \frac{x^3}{36} - \frac{x^4}{270} \quad (\text{B6})$$

for $x \lesssim 2\sqrt{\pi}$, and

$$\lambda_2(x) = y + (1 + y^{-1} + y^{-2}) \ln y, \quad y = 1 + \frac{x^2}{2}, \quad (\text{B7})$$

as $x \rightarrow \infty$. A sufficiently accurate piecewise expression for $\lambda(x)$ is

$$\begin{cases} \lambda_1(x), & x < 2, \\ g(x)\lambda_1(x) + [1 - g(x)]\lambda_2(x), & 2 \leq x \leq 4, \\ \lambda_2(x), & x > 4, \end{cases} \quad (\text{B8})$$

where $g(x) = [1 - \tanh 5(x - 3)]/2$. Equation (B2) is accurate to $\lesssim 0.3\%$ for values of $p_{\text{fa}}/N_t \leq 0.01$.

As $N_s \rightarrow \infty$, $\eta_0 \rightarrow 0$ and $\lambda(\eta_0) \rightarrow \lambda_1(\eta_0) \approx 1 + x$. With this approximation, $\eta \approx \eta_0$, $s_{\text{fa}} \approx N_s \nu(1 + \eta_0)$, and the normalized false alarm threshold [Eq. (3.5)] is $z_{\text{fa}} \approx \sqrt{2} \operatorname{erfc}^{-1}(2p_{\text{fa}}/N_t)$. Thus, in the limit of large N_s , z_{fa} is independent of N_s .

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Rep. Prog. Phys. **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
- [2] T. Accadia *et al.*, Class. Quant. Grav. **28**, 114002 (2011).
- [3] H. Grote *et al.* (LIGO Scientific Collaboration), Class. Quant. Grav. **27**, 084003 (2010).
- [4] G. M. Harry *et al.* (LIGO Scientific Collaboration), Class. Quant. Grav. **27**, 084006 (2010).
- [5] K. Kuroda *et al.* (LIGO Scientific Collaboration), Class. Quant. Grav. **27**, 084004 (2010).
- [6] R. Prix, in *Neutron Stars and Pulsars*, Astrophysics and Space Science Library, Vol. 357, edited by W. Becker (2009) p. 651, preprint available from <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=P060039>.
- [7] L. Bildsten, Astrophys. J **501**, L89 (1998), arXiv:astro-ph/9804325.
- [8] A. Melatos and D. J. B. Payne, Astrophys. J **623**, 1044 (2005), arXiv:astro-ph/0503287.
- [9] M. Vigelius and A. Melatos, M. Not. R. Astron. Soc **395**, 1972 (2009), arXiv:0902.4264 [astro-ph.HE].
- [10] B. P. Abbott *et al.* (LIGO Scientific Collaboration *et al.*), Astrophys. J **713**, 671 (2010), arXiv:0909.3583 [astro-ph.HE].
- [11] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration *et al.*), Astrophys. J **737**, 93 (2011), arXiv:1104.2712 [astro-ph.HE].
- [12] J. Abadie *et al.* (LIGO Scientific Collaboration), Astrophys. J **722**, 1504 (2010), arXiv:1006.2535 [gr-qc].
- [13] B. P. Abbott *et al.* (LIGO Scientific Collaboration *et al.*), Phys. Rev. D **80**, 042003 (2009), arXiv:0905.1705 [gr-qc].
- [14] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “All-sky Search for Periodic Gravitational Waves in the Full S5 LIGO Data,” (2011), arXiv:1110.0208 [gr-qc].
- [15] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008), arXiv:0804.3075 [gr-qc].
- [16] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **76**, 082001 (2007), arXiv:gr-qc/0605028.
- [17] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007), arXiv:astro-ph/0703234.
- [18] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004), arXiv:gr-qc/0308050.
- [19] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 102004 (2005), arXiv:gr-qc/0508065.
- [20] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008), arXiv:0708.3818 [gr-qc].
- [21] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009), arXiv:0804.1747 [gr-qc].
- [22] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. Lett **102**, 111102 (2009), arXiv:0810.0283 [gr-qc].
- [23] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998), arXiv:gr-qc/9804014.
- [24] P. R. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000), arXiv:gr-qc/9812014.
- [25] P. Jaranowski and A. Królak, Phys. Rev. D **61**, 062001 (2000), arXiv:gr-qc/9901013.
- [26] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, Phys. Rev. D **70**, 082001 (2004), arXiv:gr-qc/0407001.
- [27] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005), arXiv:gr-qc/0505082.
- [28] G. Mendell and M. Landry, *StackSlide and Hough Search SNR and Statistics*, Tech. Rep. T050003-x0 (LIGO, 2005) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T050003>.
- [29] B. Krishnan and A. M. Sintes, *Hough search with improved sensitivity*, Tech. Rep. T070124-x0 (LIGO, 2007) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T070124>.
- [30] A. L. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, M. Not. R. Astron. Soc **389**, 839 (2008), arXiv:0803.4097 [astro-ph].
- [31] R. Prix and M. Shaltev, “Continuous gravitational waves searches at fixed computing cost: The single stage of hierarchical search,” (2011), in preparation.
- [32] J. Betzwieser, *Analysis of spatial mode sensitivity of a gravitational wave interferometer and a targeted search for gravitational radiation from the Crab pulsar*, Ph.D. thesis, Massachusetts Institute of Technology (2007), <http://hdl.handle.net/1721.1/45422>.
- [33] K. W. Wette, *Gravitational waves from accreting neutron stars and Cassiopeia A*, Ph.D. thesis, The Australian National University (2009), <http://hdl.handle.net/1885/7354>.
- [34] LSC and VC (LIGO Scientific Collaboration and Virgo Collaboration), *The 2011-2012 data analysis, software and computing, detector characterization white paper*, Tech. Rep. T1100322-v3 (LSC-VC, 2011) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T1100322>.
- [35] C. Van Den Broeck, Class. Quant. Grav. **22**, 1825 (2005), gr-qc/0411030.
- [36] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998), arXiv:gr-qc/9702050.
- [37] P. Jaranowski and A. Królak, Phys. Rev. D **59**, 063003 (1999), arXiv:gr-qc/9809046.
- [38] K. Wette *et al.*, Class. Quant. Grav. **25**, 235011 (2008), arXiv:0802.3332 [gr-qc].

- [39] B. F. Schutz and M. Tinto, *M. Not. R. Astron. Soc* **224**, 131 (1987), <http://adsabs.harvard.edu/abs/1987MNRAS.224..131S>.
- [40] P. Jaranowski and A. Krolak, *Phys. Rev. D* **49**, 1723 (1994).
- [41] S. Bonazzola and E.ourgoulhon, *Astron. Astrophys.* **312**, 675 (1996), arXiv:astro-ph/9602107.
- [42] D. C. Srivastava and S. K. Sahay, *M. Not. R. Astron. Soc* **337**, 305 (2002), arXiv:gr-qc/0111050.
- [43] C. Cutler and B. F. Schutz, *Phys. Rev. D* **72**, 063006 (2005), arXiv:gr-qc/0504011.
- [44] R. Prix, *Phys. Rev. D* **75**, 023004 (2007), arXiv:gr-qc/0606088.
- [45] R. Prix, *Class. Quant. Grav.* **24**, 481 (2007), arXiv:0707.0428 [gr-qc].
- [46] V. Dergachev, *Description of PowerFlux 2 algorithms and implementation*, Tech. Rep. T1000272-v5 (LIGO, 2010) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T1000272>.
- [47] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, *Phys. Rev. D* **77**, 082001 (2008), arXiv:0712.1578 [gr-qc].
- [48] H. J. Pletsch and B. Allen, *Phys. Rev. Lett* **103**, 181102 (2009), arXiv:0906.0023 [gr-qc].
- [49] V. Dergachev, *Class. Quant. Grav.* **27**, 205017 (2010).
- [50] H. J. Pletsch, *Phys. Rev. D* **83**, 122003 (2011), arXiv:1101.5396 [gr-qc].
- [51] C. Cutler, “An improved, “phase-relaxed” F-statistic for gravitational-wave data analysis,” (2011), arXiv:1104.2938 [gr-qc].
- [52] J. Neyman and E. S. Pearson, *Phil. Trans. R. Soc. Lond. A* **231**, 289 (1933).
- [53] J. Neyman, *Phil. Trans. R. Soc. Lond. A* **236**, 333 (1937).
- [54] N. L. Johnson, S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*, 2nd ed., Wiley Series in Probability and Mathematical Statistics (Wiley, N.Y., 1994).
- [55] N. M. Temme, *Math. Comp.* **58**, 755 (1992).
- [56] D. R. Lorimer, in *Neutron Stars and Pulsars*, Astrophysics and Space Science Library, Vol. 357, edited by W. Becker (2009) p. 1.
- [57] J. Abadie *et al.* (LIGO Scientific Collaboration), *Nucl. Instrum. Meth. A* **624**, 223 (2010), arXiv:1007.3973 [gr-qc].
- [58] R. Prix, *The F-statistic and its implementation in ComputeFStatistic_v2*, Tech. Rep. T0900149-v3 (LIGO, 2010) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T0900149>.
- [59] R. Prix, *F-statistic bias due to noise-estimator*, Tech. Rep. T1100551-v1 (LIGO, 2011) <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=T1100551>.
- [60] B. F. Schutz, *Class. Quant. Grav.* **28**, 125023 (2011), arXiv:1102.5421 [astro-ph.IM].
- [61] A. C. Searle, S. M. Scott, and D. E. McClelland, *Class. Quant. Grav.* **19**, 1465 (2002), arXiv:gr-qc/0110053.
- [62] C. Messenger, R. Prix, and M. A. Papa, *Phys. Rev. D* **79**, 104017 (2009), arXiv:0809.5223 [gr-qc].
- [63] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998), arXiv:physics/9711021.
- [64] V. Dergachev, “Conversion of PowerFlux-style limits to 95% population based limits,” (2010), unpublished.
- [65] C. Röver, C. Messenger, and R. Prix, “Bayesian versus frequentist upper limits,” (2011), arXiv:1103.2987 [physics.data-an].
- [66] L. S. Finn, *Phys. Rev. D* **79**, 022002 (2009), arXiv:0810.4529 [gr-qc].
- [67] R. Prix and J. T. Whelan, *Class. Quant. Grav.* **24**, 565 (2007), arXiv:0707.0128 [gr-qc].
- [68] R. Corless, G. Gonnet, D. Hare, D. Jeffrey, and D. Knuth, *Adv. Comput. Math.* **5**, 329 (1996).
- [69] DLMF, “Digital Library of Mathematical Functions,” National Institute of Standards and Technology from <http://dlmf.nist.gov>, released July 1 (2011).