

Searching for periodic sources with LIGO. II. Hierarchical searches

Patrick R. Brady

*Institute for Theoretical Physics, University of California, Santa Barbara, California 93106
and Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, Wisconsin 53201*

Teviet Creighton

Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125

(Received 4 December 1998; published 29 February 2000)

The detection of quasi-periodic sources of gravitational waves requires the accumulation of signal to noise over long observation times. This represents the most difficult data analysis problem facing experimenters with detectors such as those at LIGO. If not removed, Earth-motion induced Doppler modulations and intrinsic variations of the gravitational-wave frequency make the signals impossible to detect. These effects can be corrected (removed) using a parametrized model for the frequency evolution. In a previous paper, we introduced such a model and computed the number of independent parameter space points for which corrections must be applied to the data stream in a coherent search. Since this number increases with the observation time, the sensitivity of a search for continuous gravitational-wave signals is computationally bound when data analysis proceeds at a similar rate to data acquisition. In this paper, we extend the formalism developed by Brady *et al.* [Phys. Rev. D **57**, 2101 (1998)], and we compute the number of independent corrections $N_p(\Delta T, N)$ required for incoherent search strategies. These strategies rely on the method of *stacked* power spectra—a demodulated time series is divided into N segments of length ΔT , each segment is Fourier transformed, a power spectrum is computed, and the N spectra are summed up. This method is incoherent; phase information is lost from segment to segment. Nevertheless, power from a signal with fixed frequency (in the corrected time series) is accumulated in a single frequency bin, and amplitude signal to noise accumulates as $\sim N^{1/4}$ (assuming the segment length ΔT is held fixed). For fixed available computing power, there are optimal values for N and ΔT which maximize the sensitivity of a search in which data analysis takes a total time $N\Delta T$. We estimate that the optimal sensitivity of an all-sky search that uses incoherent stacks is a factor of 2–4 better than achieved using coherent Fourier transforms, assuming the same available computing power; incoherent methods are computationally efficient at exploring large parameter spaces. We also consider a two-stage hierarchical search in which candidate events from a search using short data segments are followed up in a search using longer data segments. This hierarchical strategy yields a further 20–60% improvement in sensitivity in all-sky (or directed) searches for old (≥ 1000 yr) slow (≤ 200 Hz) pulsars, and for young (≥ 40 yr) fast (≤ 1000 Hz) pulsars. Assuming enhanced LIGO detectors (LIGO-II) and 10^{12} flops of effective computing power, we examine the sensitivity to sources in three specialized classes. A limited area search for pulsars in the Galactic core would detect objects with gravitational ellipticities of $\epsilon \geq 5 \times 10^{-6}$ at 200 Hz; such limits provide information about the strength of the crust in neutron stars. Gravitational waves emitted by unstable r -modes of newborn neutron stars would be detected out to distances of ~ 8 Mpc, if the r -modes saturate at a dimensionless amplitude of order unity and an optical supernova provides the position of the source on the sky. In searches targeting low-mass x-ray binary systems (in which accretion-driven spin up is balanced by gravitational-wave spin down), it is important to use information from electromagnetic observations to determine the orbital parameters as accurately as possible. An estimate of the difficulty of these searches suggests that objects with x-ray fluxes exceeding 2×10^{-8} erg cm $^{-2}$ s $^{-1}$ would be detected using the enhanced interferometers in their broadband configuration. This puts Sco X-1 on the verge of detectability in a broadband search; the amplitude signal to noise would be increased by a factor of order ~ 5 – 10 by operating the interferometer in a signal-recycled, narrow-band configuration. Further work is needed to determine the optimal search strategy when limited information is available about the frequency evolution of a source in a targeted search.

PACS number(s): 04.80.Nn, 95.55.Ym, 95.75.Pq, 97.60.Gb

I. INTRODUCTION

The detection of gravitational waves from periodic sources is seemingly the most straightforward data analysis problem facing gravitational-wave astronomers. It is also the most computationally intensive. The long observation times required to detect these waves mean that Earth-motion induced Doppler effects and intrinsic frequency drifts degrade the signal to noise if not removed. Since these effects depend

sensitively on the location and intrinsic properties of the source, searches for periodic (or quasi-periodic) sources will be limited primarily by the computational resources available for data analysis, rather than the duration of the signals or the lifetime of the instrument. For this reason, it is of paramount importance to explore different search strategies and to determine the optimal approach before the detectors go on line at the end of the century.

In a previous paper [1], hereafter referred to as paper I, we presented a detailed discussion of issues that arise when one searches for these sources in the detector output. Using a parametrized model for the expected gravitational wave signal, we also presented a method to determine the number of independent parameter values that must be sampled in a search using coherent Fourier transforms (which accumulate the signal to noise in an optimal fashion). The results were presented in the context of single-sky-position directed searches and all-sky searches, although the method outlined in paper I is applicable to *any* search over a specified region of parameter space. Livas [2], Jones [3] and Niebauer *et al.* [4] have implemented variants of the coherent search technique without the benefit of the optimization advocated in paper I.

In this paper, we discuss alternative search algorithms that can better detect quasi-periodic gravitational waves using broadband detectors. These algorithms achieve better sensitivities than a coherent search with equivalent available computational resources. This improvement is accomplished by combining coherent Fourier transforms with incoherent addition of power spectra and by using hierarchical searches that follow up the candidate detections from a first pass search.

The most likely sources of quasi-periodic gravitational waves in the frequency bands of terrestrial interferometric detectors are rapidly rotating neutron stars. We use these objects as guides when choosing the scope of the example searches considered below. Nevertheless, the search algorithms are sufficient to detect any source of continuous gravitational waves with slowly changing frequency.

A rotating neutron star will radiate gravitational waves if its mass distribution (or mass-current distribution) is not symmetric about its rotation axis. Several mechanisms that may produce non-axisymmetric deformations of a neutron star, and hence lead to gravitational wave generation, have been discussed in the literature [5–10]. A neutron star with non-zero quadrupole moment produces gravitational waves at a frequency equal to twice its rotation frequency if it rotates about a principle axis. Equally strong gravitational waves can be emitted at other frequencies when the rotation axis is not aligned with a principal axis of the source [7,11]. If the star also precesses, the gravitational waves will be produced at three frequencies: the rotation frequency and the rotation frequency plus and minus the precession frequency [9].

For concreteness, we consider a model gravitational-wave signal with one spectral component. This is not a limitation of our analysis since the search strategy presented below is inherently broadband; it can be used to detect sources that emit gravitational waves at any frequency in the detector passband. Additional knowledge of the spectral characteristics of a signal might allow us to improve our sensitivity in the case when multiple spectral components have similar signal-to-noise ratio. In such a circumstance, a modified search algorithm would sum the power at all of the appropriate frequencies. In a background of Gaussian noise, the sensitivity would improve as (number of spectral lines)^{1/4} for only a moderate increase in computational cost.

Finally, we mention several other works that consider searching for quasi-periodic signals in the output of gravitational wave detectors. Data from the resonant bar detectors around the world have been used in searches for periodic sources. New *et al.* [12] have discussed issues in searching for gravitational waves from millisecond pulsars. Krolak [13] and Jaranowski *et al.* [14,15] have considered using matched filtering to extract information about the continuous wave sources from the data stream. Finally, work is ongoing in the Albert Einstein Institute, Gölm, to investigate line-tracking algorithms based on the Hough transform [16]; this technique looks promising, although we must await results on the computational cost and statistical behavior before we can make a detailed comparison to the techniques described in this paper.

A. Gravitational waveform

The long observation times required to detect continuous sources of gravitational waves make it necessary to account for changes in the wave frequency. Physical processes responsible for these changes, and the associated time scales, were discussed in paper I. In addition, the detector moves with respect to the solar system barycenter (which we take to be approximately an inertial frame), introducing Doppler modulations of the gravitational-wave frequency. To account for these two effects, we introduce a parametrized model for the gravitational-wave frequency $f(t;\boldsymbol{\lambda})$ and phase $\phi(t;\boldsymbol{\lambda}) = 2\pi \int f(t;\boldsymbol{\lambda}) dt$ measured at the detector:

$$f(t;\boldsymbol{\lambda}) = f_0 \left(1 + \frac{\vec{v}}{c} \cdot \hat{n} \right) \left(1 + \sum_{k=1} f_k \left[t + \frac{\vec{x}}{c} \cdot \hat{n} \right]^k \right), \quad (1.1)$$

$$\phi(t;\boldsymbol{\lambda}) = 2\pi f_0 \left(t + \frac{\vec{x}}{c} \cdot \hat{n} + \sum_{k=1} \frac{f_k}{k+1} \left[t + \frac{\vec{x}}{c} \cdot \hat{n} \right]^{k+1} \right). \quad (1.2)$$

Here f_0 is the initial, intrinsic gravitational-wave frequency, $\vec{x}(t)$ is the detector position, $\vec{v}(t)$ is the detector velocity, \hat{n} is a unit vector in the direction of the source, and f_k are arbitrary coefficients that we call *spindown parameters*. (We refer the reader to paper I for a detailed discussion of this model and its physical origin.) The vector $\boldsymbol{\lambda}$ denotes the *search parameters* — the parameters of the frequency model that are (generally) unknown in advance. In the most general case that we consider below, the search parameters include frequency f_0 , the polar angles (θ, φ) used to specify \hat{n} , and the spin-down parameters f_k :

$$\boldsymbol{\lambda} = (\lambda^0, \lambda^1, \lambda^2, \lambda^3, \lambda^4, \dots) = (f_0, \theta, \varphi, f_1, f_2, \dots). \quad (1.3)$$

We note that the parameter $\lambda^0 = f_0$ defines an overall frequency scale, whereas the remaining parameters define the shape of the phase evolution. It is convenient to introduce the projected vector $\vec{\lambda} = (\lambda^1, \lambda^2, \lambda^3, \lambda^4, \dots)$ of shape parameters alone.

The strain measured at the interferometer is a linear combination of the $+$ and \times polarizations of the gravitational waves, and is given by the real part of

$$h(t; \boldsymbol{\lambda}) = \mathcal{A} e^{-i[\phi(t; \boldsymbol{\lambda}) + \Psi]}. \quad (1.4)$$

The time-dependent amplitude \mathcal{A} and phase Ψ depend on the detector response functions and the orientation of the source; they vary gradually over the course of a day (see Refs. [7,14]). In what follows, we treat \mathcal{A} and Ψ as constants. Our analysis may be generalized to include the additional phase modulation. This effectively increases the dimension of the parameter space, by 1, and the number of points that must be sampled, by ~ 4 , which translates into a reduction in relative sensitivity of $\sim 6\%$.

B. Parameter ranges

The computational difficulty of a search for quasi-periodic signals depends on the range of parameter values that are considered in the search. The intrinsic gravitational wave frequency f_0 ranges from (near) zero to some cutoff frequency f_{\max} . If gravitational waves are emitted at twice the rotation frequency, theoretical estimates [17,18] suggest that

$$f_{\max} \lesssim 1.2 - 4 \text{ kHz} \quad (1.5)$$

depending on the equation of state adopted in the neutron star model. Observational evidence—the coincidence of the periods of PSR 1937+21 and PSR 1957+20—favors the lower bound on gravitational wave frequency $f_{\max} \approx 1.2 \text{ kHz}$ [19]. The spin-down parameters f_j are allowed to take any value in the range $|f_j| \leq (1/\tau_{\min})^j$ where $\tau_{\min} \sim f/\dot{f}$ is the characteristic time scale over which the frequency might be expected to change by a factor of order unity. Observations of radio pulsars provide rough guidance about the time scales τ_{\min} . In paper I we considered two fiducial classes of sources: (i) young, fast pulsars, with $f_{\max} = 1000 \text{ Hz}$ and $\tau_{\min} = 40 \text{ yr}$, and (ii) old, slow pulsars, with $f_{\max} = 200 \text{ Hz}$ and $\tau_{\min} = 1000 \text{ yr}$. To facilitate direct comparison with the achievable sensitivities quoted in paper I, we again use these two classes to present our results.

The two extremes of sky area to be searched are (i) zero steradians for a *directed* search in which we know the source location in advance, e.g., a supernova remnant, and (ii) 4π steradians (Sr) for an *all-sky* search. We consider both of these cases, as well as the intermediate case of a 0.004 Sr search about the galactic center.

It has been suggested recently that the gravitational wave frequency of newborn, rapidly spinning neutron stars may evolve on a time scale of months rather than decades [20–23]; an active r -mode instability can radiate away most of a neutron star’s angular momentum in the form of gravitational waves within a year. Thus, newborn neutron stars may be loud enough to be detected in other galaxies, in which case optical detection of a supernova can serve as a trigger for a targeted search. Therefore we consider the case of a directed search for sources with frequencies of $f_{\max} = 200 \text{ Hz}$ and evolution time scales of $\tau_{\min} = 1 \text{ yr}$.

A final class of sources that we consider is accreting neutron stars in binary systems. Several such binary systems have been identified via x-ray observations; the rotation frequencies of the accreting neutron stars are inferred to be $\sim 250 - 350 \text{ Hz}$ ($f_{\max} = 700 \text{ Hz}$). Bildsten [24] has argued that these accreting objects in low mass x-ray binaries (LMXBs) may emit detectable amounts of gravitational radiation. Since the positions of these sources are well localized in the sky by their x-ray emissions, the Earth-motion induced Doppler modulations of the gravitational waves can be precisely determined. The difficulty with these sources is the unknown, or poorly known, orbits of the neutron stars about their stellar companions and the stochastic accretion-induced variations in their spin. We estimate the size of these effects, and outline a search algorithm in Sec. VII C. These issues deserve further study in an effort to improve the search strategy.

C. Search technique

In searches for continuous gravitational waves, our sensitivity will be limited by the computational resources available, rather than the duration of the signal or the total amount of data. Therefore the computational efficiency of a search technique is extremely important. For example, matched filtering (convolution of noise-whitened detector output with a noise-whitened template) may detect a signal with the greatest signal-to-noise ratio for any given stretch of data; however, it becomes computationally prohibitive to search over large parameter spaces with long data stretches. A sub-optimal, but more efficient, algorithm might achieve the best overall sensitivity for fixed computational resources.

We present two possible search strategies to accumulate signal to noise from the data stream. Central to both of these methods is the technique we adopt to demodulate the signal. We can remove the effects of Doppler and spin-down modulations by defining a canonical time coordinate

$$t_b[t; \vec{\lambda}] = t + \frac{\vec{x}(t)}{c} \cdot \hat{n} + \sum_{k=1} \frac{f_k}{k+1} \left[t + \frac{\vec{x}(t)}{c} \cdot \hat{n} \right]^{k+1}, \quad (1.6)$$

with respect to which the signal, defined in Eq. (1.4), is perfectly sinusoidal:

$$h(t_b[t; \vec{\lambda}]) = \mathcal{A} e^{-2\pi i f_0 t_b[t; \vec{\lambda}] - i\Psi}. \quad (1.7)$$

(Remember, we treat \mathcal{A} and Ψ as constant in time.) The introduction of the new time coordinate can be achieved by re-sampling the data stream at equal intervals in t_b . The power spectrum, computed from the Fourier transform of the re-sampled data, will consist of a single monochromatic spike, whose amplitude (relative to broadband noise) increases in proportion to the length of the data stretch. In practice the data will be sampled in the detector frame, so that a sample may not occur at the desired value of t_b . Consequently, we advocate the use of nearest-neighbor (stroboscopic) resampling [25]. This method will not substantially

reduce the signal to noise in a search provided the detector output is sampled at a sufficiently high frequency. (See Appendix B.)

When the waveform shape parameters $\vec{\lambda}$ are not known in advance, one must search over a mesh of points in parameter space. The result of a phase corrected Fourier transform will be sufficiently monochromatic only if the true signal parameters lie close enough to one of the mesh points. In Sec. II we rigorously define what is meant by “close enough,” and show how to determine the number of points for which corrections should be applied. This number depends on the adopted search strategy, and increases as a large power of the total observation time T_{data} . Note that the method of resampling followed by a Fourier transform has the benefit that a single Fourier transform automatically searches over all frequencies f_0 , leaving only the shape parameters $\vec{\lambda}$ to be searched explicitly. Other demodulation techniques, such as matched filtering, must apply separate corrections for each value of f_0 in addition to the $\vec{\lambda}$. This increases the computational cost dramatically.

A signal can also be accumulated *incoherently* from successive stretches of data by adding their power spectra [26]. However, even if each data stretch is demodulated to sufficient precision that the power from a signal is focused in a single Fourier frequency bin, residual errors in $\vec{\lambda}$ may cause the power to be at different frequencies between successive spectra. A more precise knowledge of the phase evolution is required to correct for this drift, i.e. a finer mesh in parameter space. Once a set of parameter corrections $\Delta\vec{\lambda}$ is assumed, it is relatively easy to correct for the frequency drift: successive power spectra are shifted in frequency by a correction factor Δf , where Δf is computed by differencing $f(t; \vec{\lambda})$, in Eq. (1.1), between the initial and corrected guesses for $\vec{\lambda}$, as a function of the start time of each data stretch. Once the spectra have been corrected by Δf , they can be added together. This accumulates signal to noise less efficiently than coherent phase corrections and FFT’s, but is computationally cheaper.

1. Stack-slide search

The search techniques considered in this paper are variants on the following scheme. First, the data stream is divided into shorter lengths, called *stacks*. Each stack is phase corrected using a mesh of correction points sufficient to confine a putative signal to ~ 1 frequency bin in each stack. This procedure is complicated somewhat by the evolution of the spin-down parameters with time. The following recipe determines the parameters used to phase correct each stack. Choose $\vec{\lambda}$ from the fine mesh, compute the *barycentered time* $t + \vec{x}(t) \cdot \hat{n}/c$ for the start of the entire search and the start of a given stack. Call the difference Δt_{start} . The appropriate spindown parameters to use for the given stack are those on the coarse grid which are closest to

$$f_k = \frac{\sum_{n=k} \lambda^{(2+n)} (n/k) (\Delta t_{\text{start}})^{n-k}}{1 + \sum_{n=1} \lambda^{(2+n)} (\Delta t_{\text{start}})^n}. \quad (1.8)$$

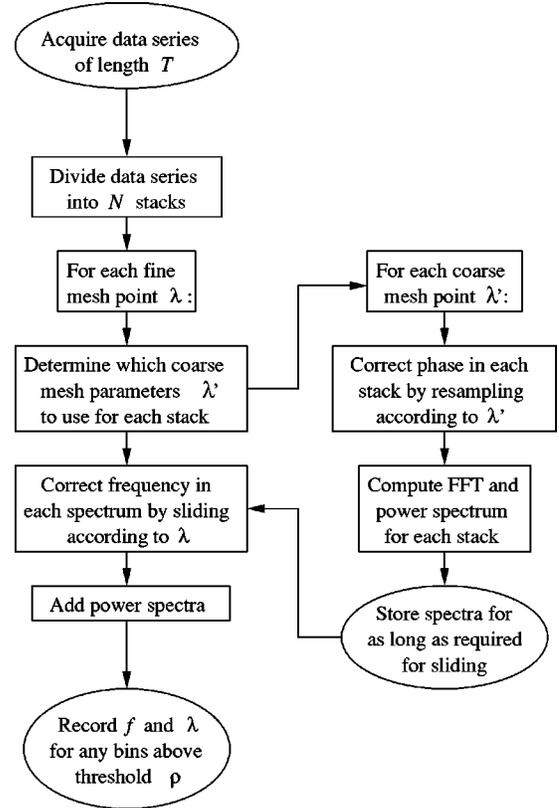


FIG. 1. A flowchart representation of the *stacked slide* algorithm to search for sources of continuous gravitational waves. Notice that the computational cost of sampling the fine grid is reduced by sliding the power spectra, rather than re-computing a fast Fourier transform (FFT) for each point on the fine grid.

The phase corrected stacks are fast Fourier transformed and the power spectra stored for as long as they are required in the next step. The individual power spectra are then shifted, relative to each other, to correct for residual frequency drift. The corrected power spectra are summed, and searched for spikes that exceed some specified significance threshold [26]. The procedure is summarized in the flowchart in Fig. 1.

2. Hierarchical search

We also consider a two-pass hierarchical search strategy. In this case, one performs an initial search of the data using a low threshold that allows for many false alarms. This is followed by a second pass, using longer stretches of data, but searching the parameter space only in the vicinity of the candidate detections of the first pass. This procedure is summarized in Fig. 2. The advantage of a hierarchical search is twofold: (i) the low threshold on the first pass allows detection of low-amplitude signals that would otherwise be rejected, and (ii) the second pass can search longer data stretches on a limited computing budget, because of the reduced parameter space being searched, thus excluding false positives from the first pass. If the thresholds and mesh points are optimally chosen between the first and second passes, this technique achieves the best sensitivity of the strategies considered here and in paper I for given computational resources.

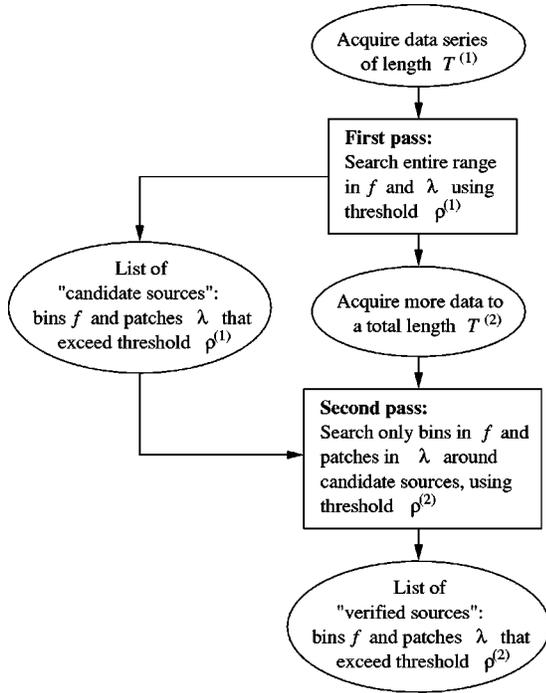


FIG. 2. A flowchart representation of the *hierarchical* algorithm to search for sources of continuous gravitational waves. It should be noted that while this approach will almost certainly be incorporated into the eventual search algorithm for gravitational waves, the real benefit of such an approach will be to increase the confidence in a detection made using some other technique.

D. Results

The sensitivity $\Theta = 1/h_{\text{th}}$ of a search is defined in Eqs. (3.3) and (3.4). The threshold strain amplitude h_{th} is defined such that there is a 1% *a priori* probability that detector noise alone will produce an event during the analysis, and therefore is the minimum characteristic strain detectable in the search. We compare our results for the sensitivity Θ to a canonical sensitivity determined by the search threshold $h_{3/\text{yr}} = 4.2\sqrt{S_n(f)} \times 10^{-7}$ Hz, where $S_n(f)$ is the one-sided power spectral density of the noise in the detector. This threshold is the characteristic amplitude of the weakest source detectable with 99% confidence in a coherent search of 10^7 s of data, if the frequency and phase evolution of the signal are known. The relative sensitivity Θ_{rel} is given by $\Theta_{\text{rel}} \equiv h_{3/\text{yr}}/h_{\text{th}}$; a relative sensitivity $\Theta_{\text{rel}} = 0.1$ for a search means that a signal must have a characteristic amplitude $h_c \geq 10 \times h_{3/\text{yr}}$ to be detected in that search. Figure 3 shows $h_{3/\text{yr}}$ based on noise spectral estimates for three detector systems in the Laser Interferometer Gravitational-wave Observatory (LIGO): the initial detectors are expected to go on line in the year 2000, with the first science run from 2002 to 2004; the upgrade to the enhanced detectors should begin in ~ 2004 , with subsequent upgrades leading to, and perhaps past, the advanced detector sensitivity. The expected amplitudes h_c of several putative sources are also shown; we use the definition of h_c given in Eq. (50) of Ref. [27] and Eq. (3.5) of paper I. The strengths of gravitational waves from the Crab and Vela radio pulsars are upper limits assuming all the rotational en-

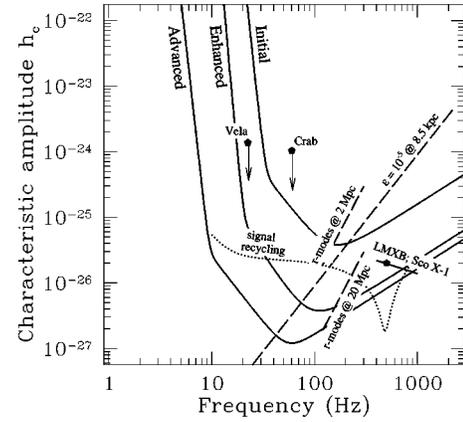


FIG. 3. Characteristic amplitudes h_c [see Eq. (3.5) in [1]] for several postulated periodic sources, compared with sensitivities $h_{3/\text{yr}}$ of the initial, enhanced (broadband and narrow band) and advanced detectors in LIGO. ($h_{3/\text{yr}}$ corresponds to the amplitude h_c of the weakest source detectable with 99% confidence in $\frac{1}{3}\text{yr} = 10^7$ s integration time, if the frequency and phase of the signal, as measured at the detector, are known in advance.) Long-dashed lines show the expected signal strength as a function of frequency for pulsars at a distance of 8.5 kpc assuming a gravitational ellipticity $\epsilon = 10^{-5}$ of the source (see Ref. [1]). Upper limits are plotted for the Crab and Vela pulsars, assuming their entire measured spin down is due to gravitational wave emission. The characteristic amplitude of waves from r -modes is also shown. These signals are not precisely periodic; rather, they chirp downward through a frequency band of ~ 200 Hz in 2×10^7 s. Finally, the strength of the gravitational waves from LMXB's, normalized to the observed x-ray flux from Sco X-1, is plotted under the assumption that gravitational waves are entirely responsible for their angular momentum loss.

ergy is lost via gravitational waves. The estimates of waves from the r -mode instability are based on Owen *et al.* [23], and those from Sco X-1 are based on the recent analysis by Bildsten [24].

The results presented in this paper are based on three reasonable assumptions about any search strategy: (i) data analysis should proceed at roughly the same rate as data acquisition, (ii) finite computational resources will be available to perform the data analysis, and (iii) in a search, the overall false alarm probability should be small (usually taken to be 1%). Given these assumptions, there is an optimal choice for the duration T_{data} of data to be analyzed and how many stacks it should be divided into. The optimal choice is that which *maximizes* the final sensitivity of a search subject to the constraints (i)–(iii) [28]. To demonstrate the existence of this optimum point, we have plotted the relative sensitivity of a search for young, fast pulsars as a function of the number of stacks and the available computing power in Fig. 4. The optimal number of stacks can be read off the plot for fixed computing power. Note that the maximum sensitivity in this plot is quite flat, especially in the regime where one is most computationally bound. This may be extremely relevant when implementing these search techniques; data management issues may impose more severe constraints on the size and number of stacks than computational requirements do. This remains to be explored when the data analysis platforms have been chosen.

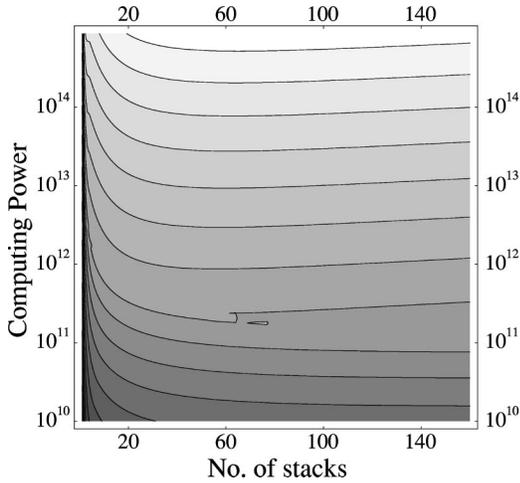


FIG. 4. A contour plot of the relative sensitivity defined in Sec. ID as a function of available computing power and the number of stacks in the search. The plot indicates the sensitivity that can be achieved using a stack-slide search for sources with $f \leq 1000$ Hz and $\tau \geq 40$ yr. The darkest shading represents the worst sensitivity. For a fixed number of stacks the sensitivity improves with increasing computing power as expected. Notice that for fixed computing resources there is generally a point of optimal sensitivity; indeed the number of stacks at this optimal operating point should be compared with those given in Table I. It is important to notice that the maximum falls off very slowly as the number of stacks increases.

Figure 5(a) shows the optimal sensitivities that can be achieved, as a function of available computing power, using a stack-slide search. The results are presented for both fiducial classes of pulsars: old ($\tau \geq 1000$ yr) slow ($f \leq 200$ Hz) pulsars and young ($\tau \geq 40$ yr) fast ($f \leq 1000$ Hz) pulsars. In each case, we have considered both directed and all-sky searches for the sources. The results should be compared with those of paper I, in which we considered coherent searches without stacking: the use of stacked searches gains a factor of $\sim 2-4$ in sensitivity.

The use of a two-pass hierarchical search can further improve sensitivities by balancing the computational requirements between the two passes. Figure 5(b) shows the sensitivities achievable when each pass uses a stack-slide strategy. The sensitivities achieved exceed those of one-pass stack-slide searches by $\sim 20-60\%$.

The computational requirements for all-sky, all-frequency surveys are sufficiently daunting that we explore three restricted searches in Sec. VII: (i) a directed search for a newborn neutron star in the young (≤ 1 yr old) remnant of an extra-galactic supernova, (ii) an area search of the galactic core for pulsars with $\tau \geq 100$ yr and $f \leq 500$ Hz, and (iii) a directed search for an accreting neutron star in a binary system (such as Sco X-1). Figure 6 shows the relative sensitivities attainable in such searches. With computational resources capable of 1 Tflop, we expect to see galactic core pulsars with enhanced LIGO if they have non-axisymmetric strains of $\epsilon \geq 5 \times 10^{-6}$ at frequencies of ~ 200 Hz. Estimates of the characteristic strain of gravitational waves from an active r -mode instability in a newborn neutron star suggest that these sources will be detectable by the enhanced inter-

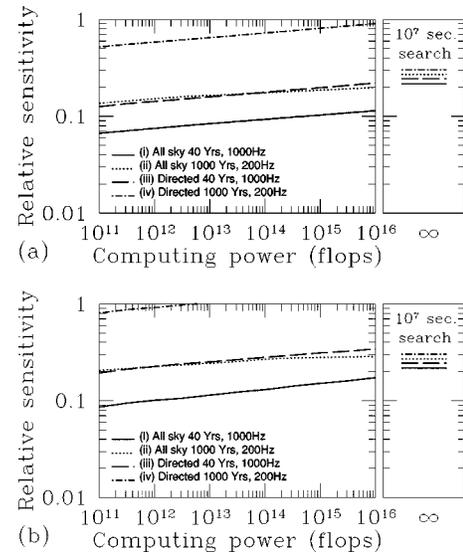


FIG. 5. Relative amplitude sensitivities $\Theta_{\text{rel}} = h_{3\text{yr}}/h_{\text{th}}$ achievable with given computational resources for (a) one pass stack-slide search strategies and (b) two-pass hierarchical strategies (using the stack-slide algorithm in each pass). The results are presented for our fiducial classes of sources: (i) all-sky search for young ($\tau \geq 40$ yr), fast ($f \leq 1000$ Hz) pulsars, (ii) all-sky search for old ($\tau \geq 1000$ yr), slow ($f \leq 200$ Hz) pulsars, (iii) directed search for young, fast pulsars, and (iv) directed search for old, slow pulsars. For a given computational power, we have determined the optimum observation time and number of stacks as described in Secs. IV and VI. Thus h_{th} is the expected sensitivity of the detector for an optimal stack-slide search, with 99% confidence. For comparison, the relative sensitivity achievable in a 10^7 s search with infinite computing power is shown in the bar on the right of each figure. When the relative sensitivity of a search exceeds this reference value, using finite computational resources, it indicates that the search uses more than 10^7 s of data.

ferometers in LIGO out to distances ~ 8 Mpc; the rate of supernovae is ~ 0.6 per year within this distance. Finally, gravitational waves from accreting neutron stars in LMXB's may be detectable by enhanced interferometers in LIGO if we can obtain sufficient information about the binary orbit from electromagnetic observations. Sco X-1 is on the margins of detectability using the enhanced LIGO interferometers operating in broadband configuration. We estimate that the amplitude signal to noise from these sources could be improved by a factor of $\sim 5-10$ by operating the interferometer in a signal-recycled, narrow-band configuration.

E. Organization of the paper

In Sec. II we extend the metric formalism that was developed in paper I to determine the number of parameter space points that must be sampled in a search that accumulates signal to noise by summing up power spectra. This method can then be used to compute the number of correction points needed in a stack-slide search. Approximate formulas, useful for estimating the computational cost of a search, are presented for the number of corrections needed in an all-sky search and also in directed searches of a single sky position.

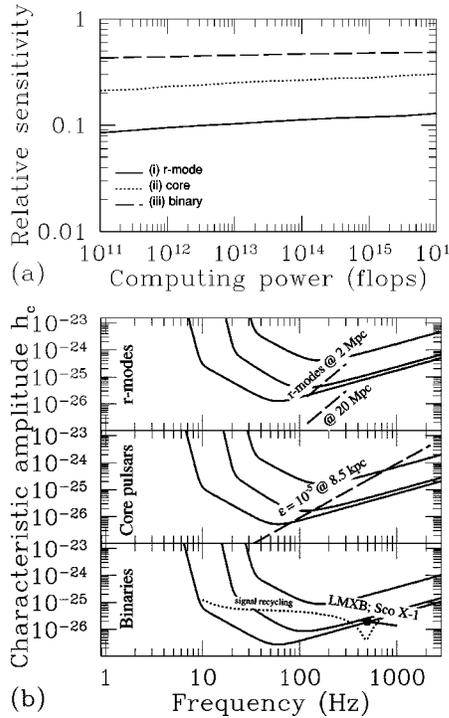


FIG. 6. Panel (a) represents the relative amplitude sensitivities $\Theta_{\text{rel}} = h_{3\text{yr}}/h_{\text{th}}$ achievable with given computational resources, in three specialized searches: (i) A search for a newborn neutron star (whose direction is determined by observing an optical supernova) that is spinning down by gravitational wave emission via an active r -mode instability. We took $\tau \geq 1$ yr and $f \leq 200$ Hz. (ii) A search for pulsars in a region extending 0.004 sr about the galactic core, with $\tau \geq 100$ yr and $f \leq 500$ Hz. (iii) A source in a binary system, e.g. Sco X-1. We assume the orbit is characterized by two orthogonal velocity parameters, known to within a total error of 17 (km/s) 2 ; we further assume that the frequency $f \leq 500$ Hz experiences a random walk typical of Eddington-rate accretion. For each of these sources, panel (b) shows h_{th} for initial (upper lines), enhanced (middle lines), and advanced (lower lines) interferometers in LIGO, assuming 1 Tflops of computing power. Thus this is the characteristic amplitude of the weakest source that can be detected with 99% confidence using a two pass-hierarchical search strategy. We have also indicated the threshold sensitivity h_{th} for enhanced LIGO in narrow-band searches for LMXB's; the center frequency coincides with the expected wave frequency for Sco X-1.

We discuss the the issue of thresholding in Sec. III. Then we present the computational cost estimates, and determine the optimal parameters for single-pass, stack-slide searches in Sec. IV.

Section V contains a general discussion of hierarchical searches for periodic sources using a single interferometer. Schutz [29] has emphasized the potential of hierarchical strategies in searches for periodic sources. The relationship between the threshold in the second stage of the search and the threshold required in the first stage is discussed in detail. We also present the computational cost of each stage of the search. These results are used in Sec. VI to determine the optimal search parameters in hierarchical searches for our fiducial classes of sources.

We discuss three specialized searches in Sec. VII. In particular, we discuss issues that arise when the source of gravitational waves is in a binary system (e.g. an LMXB). Introducing a parametrized model of the binary orbit, we estimate the number of parameter space points that must be sampled in a search for the gravitational waves from one of the objects in the binary. When the gravitational-wave emitter is accreting material from its companion, we also allow for stochastic changes in frequency due to fluctuations in the accretion rate.

Detailed formulas for the number of points in parameter space when dealing with a stacked search are presented in Appendix A. In Appendix B we discuss the loss in signal to noise that can occur when using nearest neighbor resampling to apply corrections to the detector output. If the data is sampled at 16384 Hz, we demonstrate that this method will lose less than 1% of amplitude signal to noise for a signal with gravitational wave frequency ≤ 1000 Hz.

II. MISMATCH

In a detection strategy that searches over a discrete mesh of points in parameter space, the search parameters and signal parameters will never be precisely matched. This mismatch will reduce the signal to noise since the signal will not be precisely monochromatic. It is desirable to quantify this loss and to choose the grid spacing so that the loss is within acceptable limits. This can be achieved by defining a distance measure on the parameter space based on the fractional losses in detected signal power due to parameter mismatch. In paper I we derived such a measure in the case where the search was performed using coherent Fourier transforms; this method was modeled after Owen's computation of a metric on the parameter space of coalescing binary waveforms [30]. We extend this approach to the case of incoherent searches, in which several power spectra are added incoherently, or *stacked*, and then searched for spikes.

Let $h(t; \boldsymbol{\lambda})$ be a hypothetical signal given by Eq. (1.4) with true signal parameters $\boldsymbol{\lambda} = (f_0, \vec{\lambda})$; we use the complex form of $h(t; \boldsymbol{\lambda})$ without loss of generality. If the data containing this signal are corrected for some nearby set of shape parameters $\vec{\lambda} + \Delta\vec{\lambda}$, the signal will take the form

$$h_b(t; \boldsymbol{\lambda}, \Delta\boldsymbol{\lambda}) = A e^{-i\{2\pi f_0 t_b[\vec{\lambda}; \vec{\lambda}] + \phi[t; \boldsymbol{\lambda}] - \phi[t; (f_0, \vec{\lambda} + \Delta\vec{\lambda})]\}}, \quad (2.1)$$

where the subscript b is used to indicate the corrected waveform, and $\phi[t; \boldsymbol{\lambda}]$ is defined in Eq. (1.2). In a stacked search, the data are divided into N segments of equal length ΔT seconds, each of these segments is Fourier transformed, and then a power spectrum $P_h(f; \boldsymbol{\lambda}, \Delta\vec{\lambda})$ is computed according to the formula

$$H(f; \boldsymbol{\lambda}, \Delta\vec{\lambda}) = 2 \sum_{k=1}^N |\tilde{h}_k(f; \boldsymbol{\lambda}, \Delta\vec{\lambda})|^2. \quad (2.2)$$

The Fourier transform of each individual segment is defined to be

$$\tilde{h}_k(f; \boldsymbol{\lambda}, \Delta \vec{\lambda}) = \frac{A}{\sqrt{\Delta T}} \int_{(k-1)\Delta T}^{k\Delta T} e^{i\Delta\phi[t; \boldsymbol{\lambda}, \Delta \vec{\lambda}]} dt_b \quad (2.3)$$

where $\Delta\phi$ is given by

$$\Delta\phi[t; \boldsymbol{\lambda}, \Delta \vec{\lambda}] = 2\pi(f - f_0)t_b + \phi[t; (f_0, \vec{\lambda} + \Delta \vec{\lambda})] - \phi[t; \boldsymbol{\lambda}]. \quad (2.4)$$

Here $\Delta\boldsymbol{\lambda} = (f - f_0, \Delta \vec{\lambda})$ denotes the error in matching the modulation shape parameters *and* the error in sampling the resulting power spectrum at the wrong frequency. Both of these errors lead to a reduction in the detected power relative to the optimum case where the carrier frequency and the phase modulation are precisely matched.

The *mismatch* $m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda})$, which is the fractional reduction in power due to imperfect phase correction *and* sampling at the wrong Fourier carrier frequency, is defined to be

$$m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda}) = 1 - \frac{H(f; \boldsymbol{\lambda}, \Delta \vec{\lambda})}{H(f_0; \boldsymbol{\lambda}, \vec{0})}. \quad (2.5)$$

Remember $\boldsymbol{\lambda} = (\lambda^0, \vec{\lambda}) = (f_0, \lambda^1, \lambda^2, \dots)$. Substituting the expressions for H from Eq. (2.2) into Eq. (2.5), we find

$$m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda}) = 1 - \frac{1}{N A^2} \sum_{k=1}^N |\tilde{h}_k(f; \boldsymbol{\lambda}, \Delta \vec{\lambda})|^2. \quad (2.6)$$

It is easily shown that $m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda})$ has a local minimum of zero when $\Delta \boldsymbol{\lambda} = 0$. We therefore expand the mismatch in powers of $\Delta \boldsymbol{\lambda}$ to find

$$m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda}) = \sum_{\alpha, \beta} g_{\alpha\beta}(\boldsymbol{\lambda}) \Delta \lambda^\alpha \Delta \lambda^\beta + \mathcal{O}(\Delta \boldsymbol{\lambda}^3), \quad (2.7)$$

where (α, β) are summed over $0, 1, \dots, j, \dots$. The quantity $g_{\alpha\beta}$ is a local distance metric on the parameter space. This metric is explicitly given by

$$g_{\alpha\beta}(\boldsymbol{\lambda}) \equiv \frac{1}{2} \partial_{\Delta \lambda^\alpha} \partial_{\Delta \lambda^\beta} m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda}) \Big|_{\Delta \boldsymbol{\lambda} = 0}, \quad (2.8)$$

where $\partial_{\Delta \lambda^\alpha}$ denotes a partial derivative with respect to $\Delta \lambda^\alpha$. It is convenient to express $g_{\alpha\beta}$ as a sum of metrics computed for the individual stacks, that is

$$g_{\alpha\beta}(\boldsymbol{\lambda}) = \frac{1}{N} \sum_{k=1}^N g_{\alpha\beta}^{(k)}(\boldsymbol{\lambda}) \quad (2.9)$$

where the individual stack metrics $g_{\alpha\beta}^{(k)}(\boldsymbol{\lambda})$ are explicitly given by

$$g_{\alpha\beta}^{(k)}(\boldsymbol{\lambda}) = \langle \partial_{\Delta \lambda^\alpha} \Delta \phi \partial_{\Delta \lambda^\beta} \Delta \phi \rangle_k - \langle \partial_{\Delta \lambda^\alpha} \Delta \phi \rangle_k \langle \partial_{\Delta \lambda^\beta} \Delta \phi \rangle_k. \quad (2.10)$$

The phase error $\Delta\phi$ is given in Eq. (2.4), and we use the notation

$$\langle \dots \rangle_k = \frac{1}{\Delta T} \int_{(k-1)\Delta T}^{k\Delta T} (\dots) dt_b \Big|_{\Delta \boldsymbol{\lambda} = 0}. \quad (2.11)$$

In a search, we will look for spikes in the power spectrum $H(f; \boldsymbol{\lambda}, \Delta \vec{\lambda})$ computed from the detector output; that is, we will look for local maxima in the frequency parameter f . The relevant measure of distance in the space of shape parameters $\vec{\lambda}$ is the fractional loss in power due to mismatched parameters $\Delta \vec{\lambda}$, but after maximizing over frequency. We therefore define the *projected mismatch* $\mu(\vec{\lambda}, \Delta \vec{\lambda})$ to be

$$\mu(\vec{\lambda}, \Delta \vec{\lambda}) = \min_f m(\boldsymbol{\lambda}, \Delta \boldsymbol{\lambda}) = \sum_{i,j} \gamma_{ij}(\vec{\lambda}) \Delta \lambda^i \Delta \lambda^j, \quad (2.12)$$

where

$$\gamma_{ij} = \left(g_{ij} - \frac{g_{i0} g_{j0}}{g_{00}} \right)_{\lambda^0 = f_{\max}} \quad (2.13)$$

is the mismatch metric projected onto the subspace of shape parameters, and f_{\max} is the maximum frequency that we include in the search. The meaning of the minimization \min_f is clear from the definition of the mismatch in Eq. (2.6).

Technically, γ_{ij} should be computed from $g_{\alpha\beta}$ evaluated at the specific value of f at which the minimum projected mismatch occurred. Since this number is unknown in advance of detection, we evaluate γ_{ij} for the maximum frequency that we include in the search. In this way we never underestimate the projected mismatch.

The distance function, and in particular the metric in Eq. (2.13), can be used to determine the number of discrete mesh points that must be sampled in a search. Let \mathcal{P} be the space of all parameter values $\vec{\lambda}$ to be searched over, and define the maximal mismatch μ_{\max} to be the largest fractional loss of power that we are willing to tolerate from a putative source with parameters in \mathcal{P} . For the model waveform in Eqs. (1.1), (1.2), and (1.4) this parameter space is coordinatized by $\vec{\lambda} = (\theta, \phi, f_1, f_2, \dots)$ where θ, ϕ denote location of the source on the sky, and f_j are related to the time derivative of the intrinsic frequency of the source. Each correction point of the mesh is considered to be at the center of a cube with side $2\sqrt{\mu_{\max}/n}$, where n is the dimension of \mathcal{P} ; this ensures that all points in \mathcal{P} are within a proper distance μ_{\max} of a discrete mesh point as measured with the metric γ_{ij} . The number of patches required to fill the parameter space is

$$N_p(\Delta T, \mu_{\max}, N) = \frac{\int_{\mathcal{P}} \sqrt{\det|\gamma_{ij}|} d^n \boldsymbol{\lambda}}{(2\sqrt{\mu_{\max}/n})^n}. \quad (2.14)$$

Since μ_{\max} is the maximum loss in detected power *after* the power spectra have been added, $N_p(\Delta T, \mu_{\max}, N)$ is the number of patches required to construct the *fine* mesh in the stacked search strategy described in Sec. I C. The coarse mesh in the stacked search strategy requires that the spikes in

the individual power spectra be reduced by no more than μ_{\max} ; consequently, the number of points in such a mesh is simply $N_p(\Delta T, \mu_{\max}, 1)$.

A. Directed search

In most cases, the forms of Eqs. (1.2) and (1.6) are sufficiently complicated to defy analytical solution, especially since \vec{v} in Eq. (1.2) should properly be taken from the true ephemeris of the Earth during the period of observation. However, for the case of a directed search, that is a search in just a single sky direction, the phase correction is polynomial in t , and the metric can be computed analytically. To a good approximation, the metric is flat — the spacing of points in parameter space is independent of the value of the spin-down parameters (f_1, f_2, \dots, f_s). For a given number s of spin-down parameters in a search, the right hand side of Eq. (2.14) can be evaluated analytically. The result is expressed as a product $\mathcal{N}_s G_s$, where

$$\mathcal{N}_s = \frac{f_{\max}^s (\Delta T)^{s(s+3)/2}}{(\mu_{\max}/s)^{s/2} \tau_{\min}^{s(s+1)/2}} \quad (2.15)$$

depends on the maximum frequency f_{\max} (in Hz), the length of each stack ΔT (in seconds), the maximal mismatch μ_{\max} , and the minimum spindown age τ_{\min} (in seconds) considered in the search. The dependence on the number of stacks N is contained in $G_s(N)$, which are given by

$$G_0(N) = 1, \quad (2.16)$$

$$G_1(N) \approx 0.524N, \quad (2.17)$$

$$G_2(N) \approx 0.0708N^3, \quad (2.18)$$

$$G_3(N) \approx 0.00243N^6, \quad (2.19)$$

when $N \gg 4$. The detailed expressions for $G_s(N)$ are presented in Appendix A. For up to 3 spin-down terms in the search, the number of patches is then

$$N_p(\Delta T, \mu_{\max}, N) = \max_{s \in \{0,1,2,3\}} [\mathcal{N}_s G_s(N)]. \quad (2.20)$$

The maximization accounts for the situation where increasing s , the dimension of the parameter space, decreases the value of $\mathcal{N}_s G_s$ because the parameter space extends less than one patch width in the new spin-down coordinate f_s ; one should not search over this coordinate.

B. Sky search

For signal modulations that are more complicated than simple power-law frequency drift, it is impossible to compute N_p analytically. In an actual search over sky positions as well as spin down, one should properly compute the mismatch metric numerically, using the exact ephemeris of the Earth in computing the detector position. In paper I we computed $N_p(\Delta T, \mu_{\max}, 1)$ numerically, with the simplification that both the Earth's rotation and orbital motion were taken

to be circular. However, in this paper we are concerned also with the dependence of N_p on the number of stacks N . This significantly complicates the calculation of the metric and its determinant, and makes it necessary to adopt some approximations in the calculation. Fortunately, the results of interest here are insensitive to small errors in N_p .

In paper I we mentioned that there are strong correlations between sky position and spindown parameters. This requires the use of the full $(s+2)$ -dimensional metric. However, these correlations are due primarily to the Earth's orbital motion, which has a low-order Taylor approximation for times much less than a year. Therefore we treat the number of patches as the product of the number of spin-down patches times the number of sky positions \mathcal{M}_s , computed analytically using only the Earth's rotational motion. We note that this approximation is appropriate only for computing the number of patches; *when actually demodulating the signals, the true orbital motion would have to be included*. This approximation works well as long as the orbital residuals (the remaining orbital modulations after correction on this sky mesh) are much smaller than the spin-down corrections being made at the same power in t . The residual orbital velocity at any power t^k is roughly

$$\sim \frac{\alpha_k}{k! \sqrt{\mathcal{M}_s}} \frac{r\Omega}{c} (\Omega t)^k, \quad (2.21)$$

where α_k is a number of order unity, \mathcal{M}_s is the number of sky patches, and $r=1\text{ AU}$ and $\Omega=2\pi/\text{yr}$ are the Earth's orbital radius and angular velocity. When the range in this residual is comparable to or larger than the range in the corresponding spin-down term $f_k t^k$, the ‘‘spin-down’’ parameter space must be expanded to include the orbital residuals. The range in α_k is difficult to arrive at analytically. We have found that using a maximum value of ≈ 0.3 gives good agreement with the numerical results of paper I (i.e. for $N=1$), to within factors of ~ 2 .

One other approximation was made in computing the number of sky patches. We found that the measure $\sqrt{|\det \gamma_{ij}|}$ for the sky position metric is almost constant in the azimuth φ , and has a polar angle dependence that is dominantly of the form $\sin 2\theta$. When performing the integral over sky positions, we approximated the measure by $\sqrt{|\det \gamma_{ij}|} \approx \text{const} \times \sin 2\theta$; this approximation is accurate to about 1 part in 10^4 .

Given these approximations, the number of patches for a sky search is

$$N_p = \max_{s \in \{0,1,2,3\}} \left[\mathcal{M}_s \mathcal{N}_s G_s \prod_{k=0}^s \left(1 + \frac{0.3 r \Omega^{k+1} \tau_{\min}^k}{c k! \sqrt{\mathcal{M}_s}} \right) \right]. \quad (2.22)$$

The number of sky patches, \mathcal{M}_s , in the $(s+2)$ -dimensional search, is given approximately by

$$\mathcal{M}_s \approx \frac{f_{\max}^2}{4 \mu_{\max} / (s+2)} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (2.23)$$

where

$$A = 0.014, \quad (2.24)$$

$$B = 0.046(\Delta T/1 \text{ day})^2, \quad (2.25)$$

$$C = 0.18(\Delta T/1 \text{ day})^5 N^3. \quad (2.26)$$

This is a fit to the analytic result given in Appendix A. The number of spin-down patches $\bar{N}_s G_s$ in the $(s+2)$ -dimensional search is

$$\bar{N}_s G_s = \frac{s^{s/2}}{(s+2)^{s/2}} N_s G_s \quad (2.27)$$

where N_s and G_s are given in Eqs. (2.15)–(2.19), and the prefactor on the right corrects for the sky dimensions. The remaining product terms Π in Eq. (2.22) represent the increase in the size of the spin-down space in order to include the orbital residuals.

III. THRESHOLDS AND SENSITIVITIES

The thresholds for a search are determined under the assumption that the detector noise is a stationary, Gaussian random process with zero mean and power spectral density $S_n(f)$. In the absence of a signal, the power $P_n(f) = 2|\tilde{n}(f)|^2$ at each sampled frequency is exponentially distributed with probability density function $e^{-P_n/S_n}/S_n$. The statistic for stacked spectra is $\rho = \sum_1^N P_n(f)$. The cumulative probability distribution function for ρ , in the absence of a signal, is

$$\text{CDF}[\rho/S_n, N] = \int_0^{\rho/S_n} e^{-r} \frac{r^{N-1}}{(N-1)!} dr = \frac{\gamma(N, \rho/S_n)}{(N-1)!} \quad (3.1)$$

where $\gamma(N, \rho/S_n)$ is an incomplete gamma function.

A (candidate) detection occurs whenever ρ in some frequency bin exceeds a pre-specified threshold ρ_c chosen so that the probability of a false trigger due to noise alone is small. There are $f_{\max} \Delta T$ Fourier bins in each spectrum and $N_p(\Delta T, \mu_{\max}, N)$ spectra in the entire search. Therefore we assume that a search consists of $N_p f_{\max} \Delta T$ independent trials of the statistic ρ , and compute the expected number of false events F to be

$$F = f_{\max} \Delta T N_p(\Delta T, \mu_{\max}, N) (1 - \text{CDF}[\rho_c/S_n, N]). \quad (3.2)$$

[In reality, there will be correlations between the statistic computed for different frequencies and different patches. Since this will reduce the number of independent trials, Eq. (3.2) overestimates the number of false events. This is a small effect that should not change the overall sensitivity of a search by much. It is only in the case that the number of trials is initially small that one should be concerned with this effect; unfortunately, we operate in the other extreme.] If $F \ll 1$, the number of false events F is approximately equal to the probability that an event is caused by noise in the detec-

tor. Consequently, $\alpha = 1 - F$ can be thought of as the confidence of detection. In a non-hierarchical search, the threshold ρ_c is set by specifying α and then inverting Eq. (3.2).

Finally, how does the threshold ρ_c affect the sensitivity of our search? We define a threshold amplitude h_{th} to be the minimum dimensionless signal amplitude that we expect to register as a detection in the search, that is

$$h_{\text{th}} = \sqrt{\frac{(\rho_c/N - S_n)}{\langle F_+^2 \rangle (1 - \langle \mu \rangle) \Delta T}} \quad (3.3)$$

where $\langle F_+^2 \rangle$ is the square of the detector response averaged over all possible source positions and orientations, and $\langle \mu \rangle = \mu_{\max}/3$ is the expected mismatch of a signal that is randomly located within a patch [31]. The *sensitivity* Θ of the search is then defined by

$$\Theta \equiv \frac{1}{h_{\text{th}}} \propto \sqrt{\frac{(1 - \mu_{\max}/3) \Delta T}{\rho_c/N - S_n}}. \quad (3.4)$$

For any given search strategy, the goal of optimization is to maximize the final sensitivity of the search, given limited computational power.

IV. STACK-SLIDE SEARCH

A stack-slide search is the simplest alternative to coherent searches we consider here. The main steps involved in the algorithm are shown in the flow chart of Fig. 1. In this section we estimate the computational cost of each step, and determine the ultimate sensitivity of this technique.

The first step, before the search begins, is to specify the size of the parameter space to be searched (i.e. choose f_{\max} , τ , and a region of the sky), the computational power P that will be available to do the data analysis, and an acceptable false alarm probability. From these, one can determine optimal values for the maximal mismatch μ_{\max} for a patch, the number of stacks N , and the length ΔT of each stack, using the optimization scheme discussed at the end of this section. For now, we treat these as free parameters.

Coarse and fine grids are laid down on the parameter space with $N_{pc} = N_p(\Delta T, \mu_{\max}, 1)$ and $N_{pf} = N_p(\Delta T, \mu_{\max}, N)$ points, respectively. The data stream is low-pass filtered to the upper cutoff frequency f_{\max} , and broken into N stretches of length ΔT .

Each of the steps above have negligible computational cost since they are done only once for the entire search. The subsequent steps, on the other hand, must be executed for each of the N_{pc} correction points.

Each stretch of data is re-sampled (at the Nyquist frequency $2f_{\max}$) and simultaneously demodulated by stroboscopic sampling for a set of demodulation parameters selected from the coarse grid. The result is N demodulated time series, each one consisting of $n = 2f_{\max} \Delta T$ samples. Since stroboscopic demodulation only shifts one in every few thousand data points (assuming a sampling rate at the detector of 16 384 Hz), the computational cost of the demodulation itself is negligible.

Each stretch of data is then Fourier transformed using a

FFT algorithm with a computational cost of $3nN \log_2(n)$ floating point operations. Power spectra are computed for each Fourier series, costing 3 floating point operations per frequency bin, i.e. a total cost of $1.5nN$ floating point operations.

For demodulation parameters in the coarse grid, the power of a matched signal will be confined to ~ 1 Fourier bin in each power spectrum, but not necessarily the *same* bin in different spectra. To ensure that power from a signal is accumulated by summing the N spectra, we must apply the following steps for each of the N_{pf} correction points in the *fine* mesh.

For a given set of parameters from the *fine* mesh, one determines which power spectra from the coarse grid are to be summed, using Eq. (1.8). For each of the N spectra to be stacked, the frequency of a putative signal with initial frequency f_{\max} is computed using Eq. (1.1). These can be read from a look-up table generated when the meshes were laid out. Each spectrum is re-indexed so that the power from such a signal would be in the same frequency bin (we ignore the computational cost of this step), and the spectra are added [$0.5n(N-1)$ floating point operations]. We automatically account for corrections at other frequencies by applying the fine grid corrections in this way. It may be possible to reduce the computational cost of this portion of the search by noting, for example, that we over count the fine grid corrections for signals with frequency $f_{\max}/2$ by a factor of 2^n where n is the dimension of the parameter space being explored. Since it is difficult to assess the feasibility of using this in a real search, we simply mention it so that it might be explored at the time of implementation.

The resulting stacked spectrum is scanned for peaks that exceed the threshold ρ_c . Since this has negligible computational cost, the number of floating point operations required for the entire search is

$$C = 3nNN_{pc} [\log_2(n) + 0.5 + N_{pf}(N-1)/(6NN_{pc})]. \quad (4.1)$$

If data analysis proceeds at the same rate as data acquisition, the computational power P required to complete a search is $P = C/N\Delta T$ floating-point operations per second (flops). Equation (4.1) and the definition of $n = 2f_{\max}\Delta T$ imply that the computational power is

$$P = 6f_{\max}N_{pc} [\log_2(n) + 0.5 + N_{pf}(N-1)/(6NN_{pc})]. \quad (4.2)$$

The final sensitivity Θ , defined in Eq. (3.4), of the search is determined once we know the function N_p , the frequency f_{\max} , the maximal mismatch μ_{\max} , and the confidence level $\alpha = 1 - F$. An optimized algorithm will maximize Θ as a function of μ_{\max} , N , and ΔT , subject to the constraints imposed by fixing the false alarm probability F , and the computational power P .

The results of the optimization procedure are given in Tables I and II for the fiducial classes of pulsar defined in Sec. I B. In each case we have set the probability of a false alarm threshold at $F = 0.01$ (indicating a 99% confidence that detector noise will not produce an event above threshold),

TABLE I. The optimum stack length ΔT , number of stacks N , maximal projected mismatch μ_{\max} , and the relative sensitivity $\Theta_{\text{rel}} = h_{3\text{yr}}/h_{\text{th}}$ as functions of available computational power for directed, stack-slide searches. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value). The optimization was performed numerically using simulated annealing, which accounts for some of the fluctuations in the observation times.

Young ($\tau_{\min} = 40$ yr), fast ($f_{\max} = 1000$ Hz) pulsars				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.52	157	0.49	0.13
3.16×10^{11}	1.15	56	0.43	0.13
1.00×10^{12}	1.24	60	0.41	0.14
3.16×10^{12}	1.51	55	0.43	0.15
1.00×10^{13}	1.65	58	0.43	0.16
3.16×10^{13}	1.79	62	0.43	0.17
1.00×10^{14}	2.01	63	0.42	0.18
3.16×10^{14}	2.31	62	0.42	0.19
1.00×10^{15}	2.57	64	0.43	0.20
3.16×10^{15}	2.93	64	0.43	0.21
1.00×10^{16}	3.17	69	0.44	0.22
Old ($\tau_{\min} = 10^3$ yr), slow ($f_{\max} = 200$ Hz) pulsars				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^7	1.15	236	0.49	0.23
3.16×10^7	1.58	206	0.49	0.25
1.00×10^8	2.04	199	0.50	0.28
3.16×10^8	2.63	192	0.49	0.31
1.00×10^9	3.54	172	0.49	0.34
3.16×10^9	4.29	183	0.49	0.38
1.00×10^{10}	5.26	192	0.49	0.42
3.16×10^{10}	6.52	197	0.49	0.47
1.00×10^{11}	12.62	95	0.47	0.52
3.16×10^{11}	19.41	59	0.43	0.55
1.00×10^{12}	21.28	62	0.42	0.58

and have determined optimal values of μ_{\max} , N , and ΔT for a range of values of the available computational power. Table I shows the results for directed searches for young, fast ($\tau_{\min} = 40$ yr, $f_{\max} = 1000$ Hz) and old, slow ($\tau_{\min} = 1000$ yr, $f_{\max} = 200$ Hz) pulsars, respectively. Table II shows the results for an all-sky search for the same two classes of source. The optimal sensitivities achieved by these searches are summarized in Fig. 5(a) of the Introduction.

It is worth pointing out here that we have assumed memory and I/O requirements to be irrelevant in determining the computational cost. Specifically, power spectra generated on the coarse grid are stored as long as they are needed for the sliding portion of the search. In practice, it may prove necessary to recompute power spectra or to retrieve them from slow-access data storage. Such hardware- and implementation-specific details are beyond the scope of this paper.

TABLE II. Same as in Table I, but for an all-sky search.

Young ($\tau_{\min}=40$ yr), fast ($f_{\max}=1000$ Hz) pulsars				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.04	1232	0.50	0.07
3.16×10^{11}	0.06	890	0.49	0.07
1.00×10^{12}	0.08	651	0.49	0.08
3.16×10^{12}	0.11	475	0.49	0.08
1.00×10^{13}	0.16	329	0.50	0.08
3.16×10^{13}	0.21	250	0.48	0.09
1.00×10^{14}	0.30	166	0.50	0.09
3.16×10^{14}	0.36	152	0.49	0.10
1.00×10^{15}	0.37	179	0.50	0.10
3.16×10^{15}	0.41	184	0.50	0.11
1.00×10^{16}	0.51	154	0.50	0.11
Old ($\tau_{\min}=10^3$ yr), slow ($f_{\max}=200$ Hz) pulsars				
Compute power (flops)	ΔT (days)	N	μ_{\max}	Θ_{rel}
1.00×10^{11}	1.18	54	0.42	0.14
3.16×10^{11}	1.33	55	0.42	0.14
1.00×10^{12}	1.55	53	0.42	0.15
3.16×10^{12}	2.40	31	0.31	0.16
1.00×10^{13}	2.84	27	0.27	0.16
3.16×10^{13}	3.10	27	0.28	0.17
1.00×10^{14}	3.22	29	0.30	0.17
3.16×10^{14}	3.38	30	0.28	0.18
1.00×10^{15}	3.77	29	0.29	0.19
3.16×10^{15}	4.67	24	0.27	0.19
1.00×10^{16}	5.81	19	0.21	0.20

V. HIERARCHICAL SEARCH: GENERAL REMARKS

The basic hierarchical strategy involving a two pass search is represented schematically in Fig. 2. In the first pass, $N^{(1)}$ stacks of data of length $\Delta T^{(1)}$ are demodulated on a coarse and fine mesh of correction points computed for some mismatch level $\mu^{(1)}$, and then searched by stacked Fourier transforms. A threshold signal-to-noise level is chosen that will, in general, admit many false alarms. In the second stage, $N^{(2)}$ stacks of length $\Delta T^{(2)}$ are searched on a finer mesh of points computed at a mismatch level $\mu^{(2)}$, but only in the vicinity of those events that passed the first-stage threshold. The second stage will involve fewer correction points than the first, so the second-stage transforms can be made longer and more sensitive. The goal of optimization is to find some combination of $\Delta T^{(1)}$, $\Delta T^{(2)}$, $\mu^{(1)}$, $\mu^{(2)}$, $N^{(1)}$, and $N^{(2)}$ that maximizes the final sensitivity for fixed computational power P and second pass false alarm probability $F^{(2)}$.

A. Thresholds

In the first pass of a hierarchical search, each of $N_f^{(1)} = f_{\max} \Delta T^{(1)}$ frequency bins in $N_{pf}^{(1)} = N_p(\Delta T^{(1)}, \mu^{(1)}, N^{(1)})$ stacked power spectra will be scanned for threshold crossing

events. If (as we assume) all of these trials are statistically independent, the number of false events above the threshold $\rho^{(1)}$ will be

$$F^{(1)} = N_{pf}^{(1)} N_f^{(1)} (1 - \text{CDF}[\rho^{(1)}/S_n^{(1)}, N^{(1)}]). \quad (5.1)$$

We assume that the number of false events will significantly exceed the number of true signals in this pass; consequently the number of events to be analyzed in the second pass will be $F^{(1)}$.

The second stage uses a coarse grid with $N_{pc}^{(2)} = N_p(\Delta T^{(2)}, \mu^{(2)}, 1)$ points and a fine grid with $N_{pf}^{(2)} = N_p(\Delta T^{(2)}, \mu^{(2)}, N^{(2)})$ points. On average each false alarm will require $N_{pc}^{(2)}/N_{pf}^{(1)}$ coarse grid points and $N_{pf}^{(2)}/N_{pf}^{(1)}$ fine grid points in the second stage. (When a second-pass mesh is coarser than the first pass's parameter determination, the corresponding ratio should be taken as unity.) Furthermore, since the first stage will identify the candidate signal's frequency to within ~ 2 frequency bins, the second-stage search should be over the $2\Delta T^{(2)}/\Delta T^{(1)}$ second-stage frequency bins that lie in this frequency range. Once again, we assume that the noise in all frequency bins (and over all grid points) is independent, so the number of false events that exceed the threshold $\rho^{(2)}$ in the second stage is

$$\begin{aligned} 1 - \alpha &= F^{(2)} \\ &= 2F^{(1)} \frac{N_{pf}^{(2)} \Delta T^{(2)}}{N_{pf}^{(1)} \Delta T^{(1)}} (1 - \text{CDF}[\rho^{(2)}/S_n^{(2)}, N^{(2)}]) \\ &= 2f_{\max} \Delta T^{(2)} N_{pf}^{(2)} (1 - \text{CDF}[\rho^{(1)}/S_n^{(1)}, N^{(1)}]) \\ &\quad \times (1 - \text{CDF}[\rho^{(2)}/S_n^{(2)}, N^{(2)}]), \end{aligned} \quad (5.2)$$

where α is our desired confidence level for the overall search.

The thresholds $\rho^{(1)}$ and $\rho^{(2)}$ cannot be assigned independently; rather, they should be chosen so that any *true* signal buried in the noise that would exceed (in expectation value) the second-stage threshold will have passed the first-stage threshold. In other words, it serves no purpose to set $\rho^{(2)}$ any lower than the weakest signal that would have exceeded $\rho^{(1)}$. A signal that is expected to pass the second-stage threshold exactly has an amplitude $|\tilde{h}^{(2)}|^2 = \rho^{(2)} - N^{(2)} S_n^{(2)}$. We define the *false dismissal probability* D to be the probability that such a signal will be falsely rejected in the first pass. Since the spectral power of a true signal increases with $N\Delta T$, the signal seen in the first pass has amplitude $|\tilde{h}^{(1)}|^2 = |\tilde{h}^{(2)}|^2 (N^{(1)} \Delta T^{(1)}) / (N^{(2)} \Delta T^{(2)})$, and the thresholds satisfy the relation

$$\begin{aligned} D &= \text{CDF} \left[\frac{\rho^{(1)} - |\tilde{h}^{(1)}|^2}{S_n^{(1)}}, N^{(1)} \right] \\ &= \text{CDF} \left[\frac{\rho^{(1)}}{S_n^{(1)}} - \left(\frac{\rho^{(2)}}{S_n^{(2)}} - N^{(2)} \right) \frac{S_n^{(2)}}{S_n^{(1)}} \frac{N^{(1)} \Delta T^{(1)}}{N^{(2)} \Delta T^{(2)}}, N^{(1)} \right]. \end{aligned} \quad (5.3)$$

Now, for any choice of $\Delta T^{(1)}$, $\Delta T^{(2)}$, etc., the thresholds $\rho^{(1)}$ and $\rho^{(2)}$ are completely constrained by our choices of the final confidence level α and false dismissal probability D . The false dismissal probability is fixed at $D=0.01$ in our optimization; this is an acceptably low level, meaning that only one signal in a hundred is expected to be lost in this type of search.

B. Computational costs

The computational cost $C^{(1)}$ of the first stage of the search follows the same formula as for a simple non-hierarchical search, that is

$$C^{(1)} = 6f_{\max} N^{(1)} \Delta T^{(1)} N_{pc}^{(1)} [\log_2(2f_{\max} \Delta T^{(1)}) + 0.5] + N_{pf}^{(1)} (N^{(1)} - 1) / (6N^{(1)} N_{pc}^{(1)}). \quad (5.4)$$

For each of the $F^{(1)}$ first-stage triggers, the second stage requires $N_{pc}^{(2)}/N_{pf}^{(1)}$ (minimum 1) coarse grid corrections (each involving $N^{(2)}$ FFT's of length $\Delta T^{(2)}$), along with $N_{pf}^{(2)}/N_{pf}^{(1)}$ (minimum 1) frequency shifts and spectrum additions. Each of the coarse grid corrections requires the usual $2f_{\max} N^{(2)} \Delta T^{(2)} [3 \log_2(2f_{\max} \Delta T^{(2)}) + 0.5]$ floating-point operations. The incoherent frequency shifts and spectrum additions require only $2(N^{(2)} - 1) \Delta T^{(2)} / \Delta T^{(1)}$ floating point operations since the frequency correction and power summation need only be applied over a bandwidth of ~ 2 first-pass frequency bins. The total cost of the second pass is therefore

$$C^{(2)} = \frac{2F^{(1)} N^{(2)} \Delta T^{(2)} N_{pc}^{(2)}}{N_{pf}^{(1)}} \left\{ 3f_{\max} [\log_2(2f_{\max} \Delta T^{(2)}) + 0.5] + \frac{N_{pf}^{(2)} (N^{(2)} - 1)}{N^{(2)} N_{pc}^{(2)} \Delta T^{(1)}} \right\}. \quad (5.5)$$

We require that data analysis proceed at the rate of data acquisition. Since the amount of data used in the second-stage of the search will generally be greater than that used in the first, we require that the analysis be completed in $N^{(2)} \Delta T^{(2)}$ seconds. Thus the computational power is given by

$$P = (C^{(1)} + C^{(2)}) / N^{(2)} \Delta T^{(2)}. \quad (5.6)$$

Our final sensitivity Θ is given by Eq. (3.4), using the observation time, mismatch level, and threshold of the *second* stage of the search. Optimization then consists of maximizing this function over the six parameters $\Delta T^{(1)}$, $\Delta T^{(2)}$, $\mu^{(1)}$, $\mu^{(2)}$, $N^{(1)}$, and $N^{(2)}$, for specified α , D , and P .

VI. HIERARCHICAL SEARCH WITH STACKING

It turns out that the optimization described in the previous section is only weakly sensitive to the parameters $\mu^{(1)}$ and $\mu^{(2)}$; that is, even if we choose values for $\mu^{(1)}$ and $\mu^{(2)}$ quite different from the optimal ones, we can recover nearly all of the sensitivity by adjusting the other parameters for the same computational power P . In particular, if we arbitrarily fix

$\mu^{(1)} = \mu^{(2)} = 0.3$ and re-optimize, we obtain sensitivities within 20% of the optimal.

This becomes very useful when we consider the generalized two-stage hierarchical search *with* stacking. Normally this would involve optimizing over six variables ($\mu^{(1),(2)}$, $N^{(1),(2)}$, and $\Delta T^{(1),(2)}$) with one constraint on P . However, by assuming that we can continue to set $\mu^{(1)} = \mu^{(2)} = 0.3$ with minimal loss of sensitivity, we can reduce our degrees of freedom back down to four minus one constraint.

The results of this optimization for our four canonical example searches are given in Tables III and IV. We have chosen a final confidence level $\alpha = 0.99$ and a false dismissal probability of $D = 0.01$. Table III shows the results for directed searches for young, fast ($\tau_{\min} = 40$ yr, $f_{\max} = 1$ 000 Hz) and old, slow ($\tau_{\min} = 1$ 000 yr, $f_{\max} = 200$ Hz) pulsars, respectively. Table IV shows the results for all-sky searches for the same two classes of source. The optimal sensitivities achieved by these searches are summarized in Fig. 5(b) in the Introduction.

VII. SPECIALIZED SEARCHES

The strongest sources of continuous gravitational waves are likely to be the most difficult to detect since the frequency of the waves will be changing significantly as the source radiates angular momentum. As we have seen in the previous sections, an all sky search for these sources is unlikely to achieve the desired sensitivity with available computational resources. To reach better sensitivity levels, it will be useful to consider targeted searches for specific types of source. In this section, we consider three such searches: (i) neutron stars in the galactic core as an example of a limited area sky survey, (ii) newborn neutron stars triggered on optically observed extra-galactic supernovas, and (iii) low mass x-ray binary systems such as Sco X-1.

A. Galactic core pulsars

Area surveys of the sky will certainly begin with the region most likely to hold a large number of nearby sources. Based on population models of radio pulsars in our Galaxy [32], there should be many rapidly rotating neutron stars in the galactic bulge. As an example of a limited area search, we therefore consider the optimal strategy for searching an area of 0.004 sr about the galactic core, for sources with frequencies $f \leq 500$ Hz and spin-down ages $\tau \geq 100$ yr. The choice of a 0.004 sr search is arbitrary; it includes the entire molecular cloud complex at the core of the galaxy (~ 300 pc radius at a distance of ~ 8.5 kpc).

It is easy to include a correction factor, to allow for this limited area, in our calculation of the number of patches by reducing the ranges of the integral over \mathcal{P} in Eq. (2.14). Given the approximations in Sec. II B, this amounts to reducing N_p in Eq. (2.22) by

$$0.97 \times \left(\frac{0.004}{4\pi} \right), \quad (7.1)$$

where the multiplicative factor 0.97 is the correction for the difference in functional form between the mismatch metric

TABLE III. The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of directed, hierarchical searches. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max} = 0.3$. The last column gives the overall relative sensitivity $\Theta_{\text{rel}} = h_{3\text{yr}}/h_{\text{th}}$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	Young ($\tau_{\min} = 40$ yr), fast ($f_{\max} = 1000$ Hz) pulsars				μ_{\max}	Θ_{rel}
	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$		
1.00×10^{11}	2.12	16	2.40	31	0.30	0.19
3.16×10^{11}	2.18	20	2.87	35	0.30	0.21
1.00×10^{12}	4.52	9	4.59	22	0.30	0.23
3.16×10^{12}	4.18	12	5.24	22	0.30	0.24
1.00×10^{13}	8.87	5	8.93	12	0.30	0.25
3.16×10^{13}	5.19	13	5.62	29	0.30	0.27
1.00×10^{14}	6.61	11	8.11	20	0.30	0.28
3.16×10^{14}	9.41	8	11.70	13	0.30	0.29
1.00×10^{15}	9.88	9	10.57	19	0.30	0.31
3.16×10^{15}	8.15	13	8.93	31	0.30	0.32
1.00×10^{16}	12.17	9	15.78	15	0.30	0.34

Compute power (flops)	Old ($\tau_{\min} = 10^3$ yr), slow ($f_{\max} = 200$ Hz) pulsars				μ_{\max}	Θ_{rel}
	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$		
1.00×10^7	2.37	58	2.39	185	0.30	0.36
2.51×10^7	3.43	47	3.46	132	0.30	0.39
6.31×10^7	4.31	47	4.69	110	0.30	0.42
1.58×10^8	6.66	32	7.07	76	0.30	0.46
3.98×10^8	4.77	81	5.69	152	0.30	0.50
1.00×10^9	10.24	31	11.69	62	0.30	0.55
2.51×10^9	9.75	48	9.80	115	0.30	0.59
6.31×10^9	14.35	34	15.00	75	0.30	0.64
1.58×10^{10}	17.71	32	20.65	60	0.30	0.69
3.98×10^{10}	19.78	36	23.01	69	0.30	0.74
1.00×10^{11}	25.44	32	28.76	64	0.30	0.81

and the angular area metric $d\Omega^2 = \sin^2\theta d\theta d\phi$ in the direction of the galactic center (i.e. -28.9° declination).

The optimal choices of $N^{(1,2)}$ and $\Delta T^{(1,2)}$ for a hierarchical stacked search are shown in Table V as a function of available computing power; the relative sensitivity of this search is shown in Fig. 6 of the Introduction.

We note from Eq. (3.6) of paper I that gravitational waves from rapidly rotating neutron stars might be expected to have a characteristic amplitude of

$$h_c = 2.3 \times 10^{-25} \frac{\epsilon}{10^{-5}} \frac{I_{zz}}{10^{45} \text{ g cm}^2} \frac{8.5 \text{ kpc}}{r} \left(\frac{f}{500 \text{ Hz}} \right)^2, \quad (7.2)$$

where $\epsilon = (I_{xx} - I_{yy})/I_{zz}$ is the non-axisymmetric strain, I_{ij} is the moment of inertia tensor, r is the distance to the source, f is the gravitational wave frequency, and h_c has been averaged over the detector responses to various source inclinations [27]. Theoretical estimates of the strength of the crystalline neutron star crust suggest that it can support static

deformations of up to $\epsilon \sim 10^{-5}$, though most neutron stars probably support smaller deformations. From Figs. 3 and 6, we see that 1 Tflops of computing power should allow us to detect pulsars with strains as small as $\epsilon \sim 5 \times 10^{-6}$ at 8.5 kpc using enhanced LIGO detectors.

B. Newborn neutron stars

Several recent papers [20–22] have indicated that newly-formed fast-spinning neutron stars may be copious emitters of gravitational radiation. If the newborn neutron star is rotating sufficiently fast, its r -modes (axial-vector current oscillations whose restoring force is the Coriolis force) are unstable to gravitational radiation reaction. As the star cools, viscous interactions eventually damp the modes in isolated neutron stars. Numerical studies [23] indicate that neutron stars born with rotational frequencies above several hundred Hz will radiate away most of their angular momentum in the form of gravitational waves during their first year of life. Estimates of the viscous time scales and the superfluid transition temperature suggest that the r -modes are stabilized

TABLE IV. Same as in Table III, but for an all-sky search.

Compute power (flops)	Young ($\tau_{\min}=40$ yr), fast ($f_{\max}=1000$ Hz) pulsars					
	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	0.09	169	0.10	462	0.30	0.09
3.16×10^{11}	0.12	160	0.12	451	0.30	0.09
1.00×10^{12}	0.16	105	0.16	349	0.30	0.10
3.16×10^{12}	0.20	83	0.26	204	0.30	0.11
1.00×10^{13}	0.28	56	0.34	174	0.30	0.11
3.16×10^{13}	0.50	30	0.56	97	0.30	0.12
1.00×10^{14}	0.60	31	0.78	72	0.30	0.13
3.16×10^{14}	0.94	21	1.07	55	0.30	0.14
1.00×10^{15}	1.17	22	1.30	47	0.30	0.15
3.16×10^{15}	1.33	23	1.77	40	0.30	0.16
1.00×10^{16}	2.64	11	2.75	26	0.30	0.17

Compute power (flops)	Old ($\tau_{\min}=10^3$ yr), slow ($f_{\max}=200$ Hz) pulsars					
	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	3.68	9	4.10	15	0.30	0.20
3.16×10^{11}	5.44	6	7.28	8	0.30	0.21
1.00×10^{12}	8.09	4	8.55	8	0.30	0.23
3.16×10^{12}	9.78	4	10.14	7	0.30	0.24
1.00×10^{13}	14.78	3	14.95	5	0.30	0.25
3.16×10^{13}	16.75	3	19.54	3	0.30	0.26
1.00×10^{14}	14.88	4	16.18	6	0.30	0.27
3.16×10^{14}	20.80	3	23.89	3	0.30	0.28
1.00×10^{15}	22.66	3	25.47	3	0.30	0.28
3.16×10^{15}	35.03	2	36.25	3	0.30	0.29
1.00×10^{16}	20.79	4	23.89	5	0.30	0.29

when the star cools below $\sim 10^9$ K and are rotating at ~ 100 – 200 Hz. During the evolutionary phase when most of the angular momentum is lost, the amplitude and spin-down time scale are expected to be

$$h_c = 1.2 \times 10^{-24} \sqrt{\kappa} \left(\frac{f}{1 \text{ kHz}} \right)^3 \left(\frac{20 \text{ Mpc}}{r} \right) \quad (7.3)$$

$$\tau \approx \frac{580 \text{ s}}{\kappa} \left(\frac{1 \text{ kHz}}{f} \right)^6 \approx 6t. \quad (7.4)$$

These estimates are based on Eqs. (4.9) and (5.13) in Ref. [23]. (We note that the ‘‘characteristic amplitude’’ used in Ref. [23] is appropriate to estimate the strength of burst sources, and is different from our h_c .) Here κ is a dimensionless constant of order unity; it parametrizes our ignorance of the non-linear evolution of the r -mode instability. The distance to the neutron star is r , and t is the actual age of the star. Figure 3 shows h_c as a function of frequency with $\kappa=1$ at distances $r=2$ Mpc and 20 Mpc.

Sources outside our Galaxy are potentially detectable due to the high gravitational luminosity of a newborn neutron star with an active r -mode instability. Nevertheless, it is a significant challenge to develop a feasible search strategy for

TABLE V. The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of an hierarchical search for pulsars located in a sky region of 0.004 sr about the Galactic center, with $\tau \geq 100$ yr and $f \leq 500$ Hz. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max}=0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	4.72	7	5.50	13	0.30	0.22
3.16×10^{11}	5.56	7	5.65	14	0.30	0.23
1.00×10^{12}	7.76	5	8.64	10	0.30	0.24
3.16×10^{12}	9.05	5	10.97	7	0.30	0.24
1.00×10^{13}	15.81	3	16.84	5	0.30	0.26
3.16×10^{13}	18.33	3	18.37	5	0.30	0.27
1.00×10^{14}	19.90	3	23.86	3	0.30	0.27
3.16×10^{14}	30.07	2	31.67	3	0.30	0.28
1.00×10^{15}	33.21	2	35.33	3	0.30	0.29
3.16×10^{15}	35.26	2	39.99	3	0.30	0.30
1.00×10^{16}	38.13	2	43.57	3	0.30	0.31

TABLE VI. The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of a hierarchical for newborn neutron stars spinning down due to an active r -mode instability. We assume that a supernova has been identified and accurately located in the sky, so this is a directed search for an object with $\tau_{\min}=1$ yr and $f \leq 200$ Hz. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max}=0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	1.08	4	1.16	6	0.30	0.08
3.16×10^{11}	1.22	4	1.25	7	0.30	0.09
1.00×10^{12}	1.22	5	1.26	8	0.30	0.09
3.16×10^{12}	1.67	4	2.16	5	0.30	0.10
1.00×10^{13}	2.51	3	2.79	3	0.30	0.10
3.16×10^{13}	2.63	3	3.39	3	0.30	0.11
1.00×10^{14}	2.28	4	2.32	6	0.30	0.11
3.16×10^{14}	4.58	2	4.97	3	0.30	0.11
1.00×10^{15}	3.94	3	4.04	3	0.30	0.12
3.16×10^{15}	5.94	2	6.07	3	0.30	0.12
1.00×10^{16}	4.63	3	5.00	3	0.30	0.13

these signals since the frequency evolves on such short time scales (compared to those considered above). One approach is to perform directed searches on optically observed supernova explosions. Although some supernovas may not be optically visible, and this instability may not operate in all newborn neutron stars, the computational benefits of targeting supernovas are substantial (if not essential). Based on the estimates in Ref. [23], most of the signal to noise is accumulated during the final stages of spin down. With limited computational resources, it seems best to limit the directed searches to frequencies ≤ 200 Hz, when the spin-down time scale is ~ 1 yr. Table VI shows the optimal search criteria in a hierarchical stacked search for neutron stars aged 2 months or older; the upper frequency cutoff is $f_{\max}=200$ Hz and the minimum spin-down time scale is $\tau_{\min}=1$ yr. The sensitivities achievable in a search are shown in Fig. 6 of the Introduction.

Figure 6 shows that 1 Tflops of computing power will *not* suffice to detect newborn neutron stars as far away as the Virgo cluster (≈ 20 Mpc); however, such sources will be marginally detectable within ~ 8 Mpc by enhanced LIGO detectors. The NBG catalog [33] lists 165 galaxies within this distance (assuming a Hubble expansion of 75 km/s/Mpc, retarded by the Virgo cluster). From the Hubble types and luminosities of these galaxies, and the supernova event rates in [34], we estimate a total supernova rate of ~ 0.6 per year in this volume, of which $\sim 10\%$ would be of type Ia, $\sim 20\%$ of type Ib or Ic, and $\sim 70\%$ of type II. (We note that the total rate is consistent with values given in Ref. [35].) At present, it is not known what fraction of these will produce neutron stars with unstable r -modes.

C. X-ray binaries

A low-mass x-ray binary is a neutron star orbiting around a stellar companion from which it accretes matter. The accretion process deposits both energy and angular momentum onto the neutron star. The energy is radiated away as x-rays, while the angular momentum spins the star up. Bildsten [24] has suggested that the accretion could create non-axisymmetric temperature gradients in the star, resulting in a substantial mass quadrupole and gravitational wave emission. The star spins up until the gravitational waves are strong enough to radiate away the angular momentum at the same rate as it is accreting; according to Bildsten's estimates the equilibrium occurs at a gravitational-wave frequency ~ 500 Hz. The characteristic gravitational-wave amplitudes from these sources would be

$$h_c \gtrsim 4 \times 10^{-27} \left(\frac{R}{10 \text{ km}} \right)^{3/4} \left(\frac{M}{1.4 M_{\odot}} \right)^{-1/4} \times \left(\frac{F}{10^{-8} \text{ erg cm}^{-2} \text{ s}^{-1}} \right)^{1/2} \left(\frac{f}{600 \text{ Hz}} \right)^{-1/2}, \quad (7.5)$$

where R and M are the radius and mass of the neutron star, and F is the observed x-ray flux at the Earth.

The amplitude of the gravitational waves from these sources makes them excellent candidates for targeted searches. If the source is an x-ray binary pulsar—an accreting neutron star whose rotation is observable in radio waves—then one can apply the exact phase correction deduced from the radio timing data to optimally detect the gravitational waves. (In this process, one must assume a relationship between the gravitational-wave and radio pulsation frequencies.) Unfortunately, radio pulsations have not been detected from the rapidly rotating neutron stars in all LMXB's (i.e. neutron stars that rotate hundreds of times a second). In the absence of direct radio observations, estimates of the neutron-star rotation rates are obtained from high-frequency periodic, or quasi-periodic, oscillations in the x-ray output during type I x-ray bursts. (See Ref. [36] for a summary.) But this does not provide precise timing data for a coherent phase correction. To detect gravitational waves from these sources, one must search over the parameter space of Doppler modulations due to the neutron-star orbit around its companion and fluctuations in the gravitational-wave frequency due to variable accretion rates. The Doppler effects of the gravitational-wave detector's motion can be computed exactly, because the sky position of the source is known.

In most cases, the orbital period of an x-ray binary can be deduced from periodicity in its x-ray or optical light curve. In some cases, the radial component of the orbital velocity can be computed by observing an optical emission line from the accretion disk, as was done with the bright x-ray binary Sco X-1 [37]. Such observations do not determine the phase modulation of the gravitational-wave signal with sufficient precision to make the search trivial; however, they do substantially constrain the parameter space of modulations.

In this subsection, we consider a directed search for an

x-ray binary in which the orbital parameters are known up to an uncertainty δv in the radial velocity v_r of the neutron star and an uncertainty $\delta\phi$ in the orbital phase. It is assumed that long-term photometric observations of the source can give the orbital period P to sufficient precision that we need not search over it explicitly. We therefore parametrize the phase modulations as follows:

$$\phi(t; \boldsymbol{\lambda}) = 2\pi f_0 \left(t + \frac{v_1 P}{2\pi c} \cos 2\pi t/P + \frac{v_2 P}{2\pi c} \sin 2\pi t/P \right), \quad (7.6)$$

where $\boldsymbol{\lambda} = (f_0, v_1, v_2)$ are our search parameters, the gravitational-wave frequency f_0 is constrained to be $\leq f_{\max}$, and the pair (v_1, v_2) is constrained to lie within an annular arc of radius v_r , width δv , and arc angle $\delta\phi$.

Applying the formalism developed in Sec. II to this problem gives essentially the same result as for a sky search over Earth-rotation-induced Doppler modulations if one converts time units by the ratio P/day . In the case of the Earth's rotation, a search over sky positions \hat{n} corresponds to a search over an area $\pi v_{\text{rot}}^2 \cos^2(\lambda)$ in the equatorial components of the source's velocity relative to the detector, whereas in the case of a binary orbit, the search is over a coordinate area $v_r \delta v \delta\phi$. So we can simply multiply Eq. (2.22) by the ratio of these coordinate areas to obtain the number of grid points N_p in the parameter space:

$$N_p \approx \frac{(f_{\max} P)^2}{2\mu_{\max}} \frac{v_r \delta v \delta\phi}{c^2} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (7.7)$$

$$A = 0.5, \quad (7.8)$$

$$B = 1.6(\Delta T/P)^2, \quad (7.9)$$

$$C = 6.4(\Delta T/P)^5 N^3. \quad (7.10)$$

Accounting for the intrinsic phase variations of the spinning neutron star itself is problematic since changes in the accretion rate may lead to stochastic variations in the rotation frequency of the star. Consider a typical neutron star (with radius 10 km, mass $1.4M_\odot$ and moment of inertia 10^{45} g cm^2) in an accreting system. If the torque from accretion turns off, or doubles, for some reason, we can expect the gravitational-wave frequency to drift by

$$\Delta f \approx 3.6 \times 10^{-8} \left(\frac{\dot{M}}{1.5 \times 10^{-9} M_\odot \text{ yr}^{-1}} \right) \left(\frac{t}{\text{day}} \right) \text{ Hz}, \quad (7.11)$$

over a time t , where \dot{M} is the rate of accretion onto the neutron star. For the sources of interest to us, accretion proceeds at (or near) the Eddington rate $\dot{M} \approx 1.5 \times 10^{-9} M_\odot \text{ yr}^{-1}$. If we require that the frequency drifts by less than one Fourier bin $\Delta f \leq 1/\Delta T$ during a coherent observation, the observing time $t = \Delta T$ must satisfy

$$\Delta T \leq 18 \text{ days}. \quad (7.12)$$

The frequency will make a random walk as the accretion rate fluctuates about its mean value. This type of random walk cannot be modeled as a low-order polynomial in time. Nevertheless, the stack-slide technique is well suited to search for these sources since corrections for the stochastic changes in frequency can be applied by shifting the stacks by $+1$, 0 , or -1 frequency bins when required. In a search using N stacks, each of length ΔT , this kind of correction would be applied after a time t such that

$$3.6 \times 10^{-8} \text{ Hz } (t/\text{day}) = 1/\Delta T. \quad (7.13)$$

The number of times that these corrections must be applied is then $N\Delta T/t$, and the number of distinct frequency evolutions traced out in this procedure is $3^{N\Delta T/t}$. Monte Carlo simulations of stacked FFT's of signals undergoing random walks in frequency have shown that one can increase t by up to a factor of 4, i.e. allowing drifts of up to ± 2 frequency bins, while incurring only $\sim 20\%$ losses in the final summed power; however, we have not yet studied in detail how this combines with the mismatches generated from other demodulations or how to search over all demodulations together in an optimal way. For now we assume a factor of $3^{0.003N(\Delta T/\text{day})^2}$ extra points in our search mesh for mismatches of $\mu_{\max} \approx 0.3$.

As an example, we consider a search for gravitational waves from the neutron star in Sco X-1. This system has an orbital period $P = 0.787313 \pm 0.000001$ days [38], a radial orbital velocity amplitude of $v_r = 58.2 \pm 3.0$ km/s, an orbital phase known to ± 0.10 rad [37], and an inferred gravitational-wave frequency $f_{\max} \approx 500$ Hz [24]. We note that the uncertainty in P is basically negligible over the < 18 day coherent integrations expected. The remaining uncertainties give

$$N_p = \frac{10^5}{\mu_{\max}} 3^{0.003N(\Delta T/\text{day})^2} (A^{-2} + B^{-2} + C^{-2})^{-1/2}, \quad (7.14)$$

$$A = 0.5, \quad (7.15)$$

$$B = 2.6(\Delta T/\text{day})^2, \quad (7.16)$$

$$C = 21(\Delta T/\text{day})^5 N^3, \quad (7.17)$$

where it is understood that $\Delta T \leq 18$ days, in order for the random-walk stack-slide corrections to achieve maximum sensitivity.

Table VII shows the optimal search criteria for a hierarchical stacked search for the Sco X-1 pulsar under these assumptions. The sensitivities achievable in such a search are shown in Fig. 6 of the Introduction. We see that 1 Tflops of computing power may be sufficient to detect this source using enhanced LIGO detectors if it is radiating most of the accreting angular momentum as gravitational waves. The sensitivity to these sources would be enhanced by a factor of $\sim 5-10$ if the interferometer is operated in a signal-recycled, narrow-band configuration during the search.

TABLE VII. The optimum stack length $\Delta T^{(1,2)}$ and number of stacks $N^{(1,2)}$ for the first and second stages of a hierarchical search for sources of continuous gravitational waves that are in binary systems. We assume the orbit is characterized by two orthogonal velocity parameters that are known to within a total error of 17 (km/s)^2 , and that the frequency $f \leq 500 \text{ Hz}$ is experiencing a random walk typical of Eddington-rate accretion. For numerical convenience the maximal projected mismatch was chosen in advance to be $\mu_{\max} = 0.3$. The threshold was chosen to give an overall statistical significance of 99% to a detection (although the results are insensitive to the precise value).

Compute power (flops)	$\Delta T^{(1)}$ (days)	$N^{(1)}$	$\Delta T^{(2)}$ (days)	$N^{(2)}$	μ_{\max}	Θ_{rel}
1.00×10^{11}	7.12	11	8.60	53	0.30	0.40
3.16×10^{11}	7.33	14	8.51	57	0.30	0.40
1.00×10^{12}	4.72	49	8.31	58	0.30	0.40
3.16×10^{12}	5.13	50	8.47	58	0.30	0.41
1.00×10^{13}	5.84	57	8.27	56	0.30	0.41
3.16×10^{13}	6.14	56	8.46	57	0.30	0.41
1.00×10^{14}	6.78	55	8.38	57	0.30	0.42
3.16×10^{14}	6.69	48	9.17	57	0.30	0.42
1.00×10^{15}	7.40	49	8.93	58	0.30	0.42
3.16×10^{15}	7.53	55	8.99	56	0.30	0.42
1.00×10^{16}	7.00	51	9.71	58	0.30	0.42

VIII. FUTURE DIRECTIONS

We have presented in this manuscript the rudiments of a search algorithm for sources of continuous gravitational waves. We have completely neglected overheads associated with memory and I/O in our analysis. These are potentially serious issues that depend on the platform chosen for the analysis. To investigate them requires implementation of the schemes presented in this paper. This will test the algorithms with real data, and, hopefully, highlight the shortcomings of the computational cost estimates presented here. It will also allow direct comparison with other approaches such as the line-tracking method that is being explored by Papa *et al.* [16].

Several issues remain to be explored. The two-stage hierarchical searches discussed in this paper use only 60 days of

data to achieve optimal sensitivity. At LIGO, it is planned to acquire $\sim 1 \text{ yr}$ of data in the first science run. Should we discard all but ~ 60 days of it? Clearly not. We have explored two possible search algorithms in this paper, but there are infinitely many algorithms to choose from. The goal is to find the one that uses *all* of the data to achieve optimal sensitivity in the sense that we have used it. The algorithm that achieves this goal will be more complex than the two-stage hierarchical algorithm — we suspect it will be a multi-stage hierarchical algorithm — we suspect it will be a multi-stage hierarchical algorithm. At this point, the urgency is to implement any reasonable search algorithm. Experimental advances will probably out pace gains achieved by optimizing the algorithm to use all the data.

Further theoretical work is also need to determine the parameter space that should be searched, especially in the case of active *r*-mode instabilities and radiating neutron stars in LMXB's. We have performed preliminary estimates in this paper; however, the promising nature of these sources should make them high priorities in targeted searches.

Finally, it would be worthwhile to consider what advantage, if any, can be gained by using data from multiple interferometers at the initial detection stage of a search for continuous gravitational waves.

ACKNOWLEDGMENTS

This work was supported by NSF grant PHY94-24337. P.R.B. is supported by NSF grant PHY94-07194 at the ITP, and he is grateful to the Sherman Fairchild Foundation for financial support while at the California Institute of Technology. We especially thank Kip Thorne for his help and encouragement throughout this work, Lars Bildsten for detailed discussions about LMXB's, and Stuart Anderson for many illuminating discussions. This work has also benefited from interactions with Bruce Allen, Curt Cutler, Jolien Creighton, Sam Finn, Scott Hughes, Andrzej Krolak, Ben Owen, Tom Prince, Bernard Schutz, and Alan Wiseman.

APPENDIX A: PATCH NUMBER FORMULAS

The approximate formulas given in Eqs. (2.17)–(2.19) are valid when $N \gg 4$. General expressions for the G_s can be derived by setting $\vec{x} = \vec{0}$ in Eqs. (1.2) and (1.6), and using Eqs. (2.4), (2.9), (2.10), (2.11), (2.13), and (2.14):

$$G_1(N) = \frac{\pi}{6\sqrt{5}} \sqrt{5N^2 - 4} \quad (\text{A1})$$

$$G_2(N) = \frac{\pi^2}{180\sqrt{105}} \sqrt{175N^6 - 840N^4 + 1100N^2 - 432} \quad (\text{A2})$$

$$G_3(N) = \frac{\pi^3}{75600\sqrt{105}} \sqrt{3675N^{12} - 58800N^{10} + 363160N^8 - 1053360N^6 + 1484336N^4 - 987840N^2 + 248832}. \quad (\text{A3})$$

In Eq. (2.14) we approximate the metric as having constant determinant and evaluate it at the point of zero spin down; this introduces small errors of order $([N\Delta T]/\tau)^2$.

Equations (2.23)–(2.26) for the number of sky patches ignoring spin down and orbital motions provide an empirical fit to a numerical evaluation of the metric determinant. The determinant was found to have an approximate functional dependence $\sqrt{|\gamma_{ij}|} \sim |\sin 2\theta|$ with corrections of order $v/c \approx 10^{-4}$. Assuming this dependence to be exact, Eqs. (2.13) and (2.14) give

$$\mathcal{M}_s = \frac{2/\pi}{4\mu_{\max}^{s+2}} \sqrt{\det \left| g_{ij} - \frac{g_{0i}g_{0j}}{g_{00}} \right|_{\substack{f=f_{\max} \\ \theta=\pi/2}}}. \quad (\text{A4})$$

Here $g_{\alpha\beta}$ is the mismatch metric, defined in Eqs. (2.9)–(2.11), computed using only the Earth-rotation-induced Doppler modulation. Since the Earth's rotation is a simple circular motion, and since we are evaluating the metric at a single point in parameter space, we can carry out the integrals in Eq. (2.11) analytically, to obtain

$$g_{\alpha\beta} = b_{\alpha\beta} - \frac{1}{N} \sum_{k=1}^N a_{k\alpha} a_{k\beta}, \quad (\text{A5})$$

where

$$a_{k0} = \left(k - \frac{1}{2} \right) \Delta\Phi + \frac{v/c}{\Delta\Phi} \{ \sin(k\Delta\Phi) - \sin([k-1]\Delta\Phi) \} \quad (\text{A6})$$

$$a_{k1} = \frac{f_{\max}v/c}{\Delta\Phi} \{ \sin(k\Delta\Phi) - \sin([k-1]\Delta\Phi) \} \quad (\text{A7})$$

$$a_{k2} = \frac{f_{\max}v/c}{\Delta\Phi} \{ \cos([k-1]\Delta\Phi) - \cos(k\Delta\Phi) \} \quad (\text{A8})$$

$$b_{00} = \frac{1}{12\Phi} \left\{ -24 + \frac{v}{c} + 6 \left(\frac{v}{c} \right)^2 \Phi + 4\Phi^3 + 24 \frac{v}{c} \cos \Phi + 24 \frac{v}{c} \Phi \sin \Phi + 3 \left(\frac{v}{c} \right)^2 \sin 2\Phi \right\} \quad (\text{A9})$$

$$b_{11} = \frac{f_{\max}^2(v/c)^2}{4\Phi} (2\Phi + \sin 2\Phi) \quad (\text{A10})$$

$$b_{22} = \frac{f_{\max}^2(v/c)^2}{4\Phi} (2\Phi - \sin 2\Phi) \quad (\text{A11})$$

$$b_{01} = b_{10} = \frac{f_{\max}v}{4\Phi} \left\{ -4 + 2 \frac{v}{c} \Phi + 4 \cos \Phi + 4\Phi \sin \Phi + \frac{v}{c} \sin 2\Phi \right\} \quad (\text{A12})$$

$$b_{02} = b_{20} = \frac{f_{\max}v/c}{2\Phi} \left\{ -2\Phi \cos \Phi + \sin \Phi \left(2 + \frac{v}{c} \sin \Phi \right) \right\} \quad (\text{A13})$$

$$b_{12} = b_{21} = \frac{f_{\max}^2(v/c)^2}{2\Phi} \sin^2 \Phi. \quad (\text{A14})$$

Here $\Phi = 2\pi N\Delta T / (1 \text{ day})$ is the total angle over which the Earth rotates during the observation, $\Delta\Phi = \Phi/N$ is the angle rotated during each stretch of the data, and v is the maximum radial velocity relative to the detector at latitude $\lambda = 45^\circ$ of a point at a polar angle $\theta = \pi/4$ on the sky, that is

$$v = \frac{2\pi R_{\text{Earth}} \cos \lambda \sin \theta}{1 \text{ day}}. \quad (\text{A15})$$

APPENDIX B: RESAMPLING ERROR

In this paper, we have assumed that coherent phase corrections are achieved through stroboscopic resampling: a demodulated time coordinate $t_b[t]$ is constructed, and the data stream $h(t)$ is sampled at equal intervals in t_b at the Nyquist rate for the highest frequency signal present, $f_{\text{Nyquist}} = 2f_{\max}$. However, since the data stream is initially sampled at some finite rate $f_s = Rf_{\text{Nyquist}}$ (where R is the *oversampling factor*), this can introduce errors: in general, there will not be a data point exactly at a given value of t_b , so the nearest (in time) datum must be substituted. Consequently, there will be residual phase errors $\Delta\Phi(t) \in [-\pi/2R, \pi/2R]$ caused by rounding to the nearest datum even if one chooses a phase model whose frequency and modulation parameters exactly match the signal. The phase of the resampled signal drifts until the timing error is $\geq 1/2f_s$, at which point one corrects the phase by sampling an adjacent datum, which shifts in time by $1/f_s$. These residual phase errors reduce the Fourier amplitude of the signal by a fraction

$$F = \left| \frac{1}{N} \sum_{k=1}^N e^{i\Delta\Phi(k/f_{\text{Nyquist}})} \right|, \quad (\text{B1})$$

where $N = f_{\text{Nyquist}}\Delta T$ is the total number of points in the resampled data stream.

The uncorrected signal will in general drift by many radians in phase, which is the reason why we must apply phase corrections in the first place. This means that $\Delta\Phi(t)$ will sweep through the range $[-\pi/2R, \pi/2R]$ many times over

TABLE VIII. The percentage reduction $(1-F)$ in amplitude of a signal as a function of the oversampling factor R . The LIGO interferometers will collect data at $f_s = 16384$ Hz, so that the data will be oversampled by $R \geq 4$ compared to the maximum gravitational wave frequency that we expect on physical grounds. In fact, it seems more likely that $R \approx 8$ for real signals.

$R =$	2	3	4	5	6	7	8
$1-F =$	10.0%	4.5%	2.6%	1.6%	1.1%	0.8%	0.6%

the course of the observation. So, regardless of the precise form of the phase evolution, we expect $\Delta\Phi(k/f_{\text{Nyquist}})$ to be essentially evenly distributed over this interval. Thus, replacing the sum with an expectation integral, we have

$$F \approx \left| \frac{R}{\pi} \int_{-\pi/2R}^{\pi/2R} e^{i\Phi} d\Phi \right| = \frac{\sin(\pi/2R)}{\pi/2R}. \quad (\text{B2})$$

The fractional losses in amplitude $1 - F$ for a few values of the oversampling factor R are given in Table VIII.

The highest gravitational-wave frequencies we consider are 1000 Hz, requiring a resampling rate of $f_{\text{Nyquist}} = 2000$ Hz. Since LIGO will acquire data at a rate of 16384 Hz, corresponding to an oversampling factor of $R > 8$, we have a maximum signal loss due to resampling of $1 - F = 0.6\%$. Resampling errors will increase if the number of data samples is reduced by some factor *before* phase correcting.

-
- [1] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, *Phys. Rev. D* **57**, 2101 (1998).
- [2] J. C. Livas, Ph.D. thesis, Massachusetts Institute of Technology, 1987.
- [3] G. S. Jones, Ph.D. thesis, University of Wales, 1995.
- [4] T. M. Niebauer *et al.*, *Phys. Rev. D* **47**, 3106 (1993).
- [5] S. Chandrasekhar, *Phys. Rev. Lett.* **24**, 611 (1970).
- [6] J. L. Friedman and B. F. Schutz, *Astrophys. J.* **222**, 281 (1978).
- [7] S. Bonazzola and E. Gourgoulhon, *Astron. Astrophys.* **312**, 675 (1996).
- [8] M. Zimmermann, *Phys. Rev. D* **21**, 891 (1980).
- [9] M. Zimmermann and E. Szedenits, Jr., *Phys. Rev. D* **20**, 351 (1979).
- [10] R. V. Wagoner, *Astrophys. J.* **278**, 345 (1984).
- [11] D. V. Gal'tsov, V. P. Tsvetkov, and A. N. Tsurlev, *Zh. Eksp. Teor. Fiz.* **86**, 809 (1984) [*Sov. Phys. JETP* **59**, 472 (1984)].
- [12] K. C. B. New, G. Chanmugam, W. W. Johnson, and J. E. Tohline, *Astrophys. J.* **450**, 757 (1995).
- [13] A. Krolak, "Searching data for periodic signals," gr-qc/9803055.
- [14] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- [15] P. Jaranowski and A. Krolak, *Phys. Rev. D* **59**, 063003 (1999).
- [16] M. A. Papa, B. F. Schutz, S. Frasca, and P. Astone, in *Proceedings of the LISA Symposium, 1998*; M. A. Papa, P. Astone, S. Frasca, and B. F. Schutz, in the *Proceedings of Gravitational Wave Data Analysis Workshop, 1997*; M. A. Papa (private communication).
- [17] S. L. Shapiro, S. A. Teukolsky, and I. Wasserman, *Astrophys. J.* **272**, 702 (1984).
- [18] J. L. Friedman, J. R. Ipser, and L. Parker, *Nature (London)* **312**, 255 (1984).
- [19] S. R. Kulkarni, *Philos. Trans. R. Soc. London* **A341**, 77 (1992).
- [20] N. Andersson, *Astrophys. J.* **502**, 708 (1998).
- [21] J. L. Friedman and S. M. Morsink, *Astrophys. J.* **502**, 714 (1998).
- [22] L. Lindblom, B. J. Owen, and S. M. Morsink, *Phys. Rev. Lett.* **80**, 4843 (1998).
- [23] B. J. Owen *et al.*, *Phys. Rev. D* **58**, 084020 (1998).
- [24] L. Bildsten, *Astrophys. J. Lett.* **501**, L89 (1998).
- [25] B. F. Schutz, in *The Detection of Gravitational Waves*, edited by D. G. Blair (Cambridge University Press, Cambridge, England, 1991), Chap. 16, pp. 406–451.
- [26] The method of stacking power spectra has been used by radio astronomers in deep searches for millisecond pulsars, although all corrections were applied to the data stream via resampling and not sliding the spectra. For more information on the implementation, see S. B. Anderson, Ph.D. thesis, California Institute of Technology, 1993.
- [27] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987), Chap. 9, pp. 330–458.
- [28] It may seem that the optimal duration of data to analyze is equal to the amount of data taken by the instrument. This is not necessarily true for a given algorithm (e.g. the stack-slide search or the two-stage hierarchical search). Suppose, under the stated assumptions, we determine the optimal amount of data T_{data} to analyze using a given algorithm. Now, hold the computational resources fixed, but increase the amount of data by a factor of 10, so that we have 10 times as long to analyze it. Unfortunately the computational cost increases by more than a factor of 10 because the number of parameter-space corrections increases faster than T_{data} . Thus, we cannot complete our analysis in the time it takes to acquire the data. Implications of this point are further discussed in Sec. VIII.
- [29] B. F. Schutz, "Sources of radiation from neutron stars," gr-qc/9802020; talk given at Gravitational Wave Data Analysis Workshop, MIT, 1996.
- [30] B. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [31] The average expected power loss for a source randomly placed within a cubical patch is $\langle \mu \rangle = \mu_{\text{max}}/3$. In paper I we quoted an average that was computed for ellipsoidal patches; this is not appropriate to the cubical grid that will likely be used in a real search.
- [32] S. J. Curran and D. R. Lorimer, *Mon. Not. R. Astron. Soc.* **276**, 347 (1995).
- [33] R. B. Tully, *Nearby Galaxy Catalog* (Cambridge University Press, Cambridge, England, 1988).
- [34] S. van den Bergh and R. D. McClure, *Astrophys. J.* **425**, 205 (1994).
- [35] R. Talbot, *Astrophys. J.* **205**, 535 (1976).
- [36] M. van der Klis, in *The Many Faces of Neutron Stars*, edited by A. Alpar, L. Buccheri, and J. van Paradijs (Kluwer, Dordrecht, in press).
- [37] A. P. Cowley and D. Crampton, *Astrophys. J., Lett. Ed.* **201**, L65 (1975).
- [38] E. W. Gottlieb, E. L. Wright, and W. Liller, *Astrophys. J., Lett. Ed.* **195**, L33 (1975).