

Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement

Benjamin J. Owen

*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125
and Max Planck Institut für Gravitationsphysik (Albert Einstein Institut), Am Mühlenberg 5, 14476 Golm, Germany*

B. S. Sathyaprakash

*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125
and Department of Physics and Astronomy, Cardiff University, Cardiff, CF2 3YB, United Kingdom
(Received 27 August 1998; published 22 June 1999)*

We estimate the number of templates, computational power, and storage required for a one-step matched filtering search for gravitational waves from inspiraling compact binaries. Our estimates for the one-step search strategy should serve as benchmarks for the evaluation of more sophisticated strategies such as hierarchical searches. We use a discrete family of two-parameter wave form templates based on the second post-Newtonian approximation for binaries composed of nonspinning compact bodies in circular orbits. We present estimates for all of the large- and mid-scale interferometers now under construction: LIGO (three configurations), VIRGO, GEO600, and TAMA. To search for binaries with components more massive than $m_{\min}=0.2M_{\odot}$ while losing no more than 10% of events due to coarseness of template spacing, the initial LIGO interferometers will require about 1.0×10^{11} flops (floating point operations per second) for data analysis to keep up with data acquisition. This is several times higher than estimated in previous work by Owen, in part because of the improved family of templates and in part because we use more realistic (higher) sampling rates. Enhanced LIGO, GEO600, and TAMA will require computational power similar to initial LIGO. Advanced LIGO will require 7.8×10^{11} flops, and VIRGO will require 4.8×10^{12} flops to take full advantage of its broad target noise spectrum. If the templates are stored rather than generated as needed, storage requirements range from 1.5×10^{11} real numbers for TAMA to 6.2×10^{14} for VIRGO. The computational power required scales roughly as $m_{\min}^{-8/3}$ and the storage as $m_{\min}^{-13/3}$. Since these scalings are perturbed by the curvature of the parameter space at second post-Newtonian order, we also provide estimates for a search with $m_{\min}=1M_{\odot}$. Finally, we sketch and discuss an algorithm for placing the templates in the parameter space. [S0556-2821(99)05214-5]

PACS number(s): 04.80.Nn, 07.05.Kf, 97.80.-d

I. INTRODUCTION

Close binary systems composed of compact objects (such as black holes and neutron stars) are expected to be an important source of gravitational waves for broadband laser interferometers such as the Laser Interferometric Gravitational Wave Observatory (LIGO), VIRGO, GEO600, and TAMA [1–4]. The orbit of a compact binary decays under the influence of gravitational radiation reaction, emitting a gravitational wave signal that increases in amplitude and “chirps” upward in frequency as the objects spiral in toward each other just before their final coalescence. According to current astronomical lore [5–7], the rate of coalescences should be about three per year within 200 to 300 Mpc [5] of the Earth for neutron-star–neutron-star binaries, and within 400 Mpc to 1 Gpc for black-hole–black-hole binaries. Signals from inspiraling compact binaries at these distances are strong enough to be detected by “enhanced” LIGO [8] and VIRGO interferometers, but only if aided by a nearly optimal signal-processing technique. Fortunately, the distinctive frequency chirp has been calculated to a remarkable degree of precision using a variety of approximations to the general relativistic two-body problem (e.g. [9,10]). Because the functional form of the chirp is quite well known, a search for inspiral signals in noisy data is ideally suited to matched filtering.

Matched filtering [11] has long been known to be the

optimal linear signal-processing technique and is well discussed in the literature (e.g. [12]); therefore we will only briefly summarize it here. In the frequency domain, a matched filter is a best-guess template of the expected signal waveform divided by the interferometer’s spectral noise density. The interferometer output is cross-correlated with the matched filter at different time delays to produce a filtered output. The signal-to-noise ratio, defined as the ratio of the actual value of the filtered output to its rms value in the presence of pure noise, is compared to a predetermined threshold to decide if a signal is present in the noise. If the signal from which the matched filter was constructed is present, it contributes coherently to the cross-correlation, while the noise contributes incoherently and thus is reduced relative to the signal. Also, the weighting of the cross-correlation by the inverse of the spectral noise density emphasizes those frequencies to which the interferometer is most sensitive. Consequently, signals thousands of cycles long whose unfiltered amplitude is only a few percent of the rms noise can be detected.

A matched filtering search for inspiraling compact binaries can be computationally intensive due to the variety of possible waveforms. Although the inspiral signals are all expected to have the same functional form, this form depends on several parameters—the masses of the two objects, their spins, the eccentricity of their orbit, etc.—some of them

quite strongly. A filter constructed from a waveform template with any given parameter vector may do a very poor job of detecting a signal with another parameter vector. That is, the difference in parameter vectors can lead to a greatly reduced cross-correlation between the two wave forms; and in general, the greater the difference, the more the cross-correlation is reduced. Because the parameter vector of a signal is not known in advance, it is necessary to filter the data with a family of templates located at various points in parameter space—e.g., placed on a lattice—such that any signal will lie close enough to at least one of the templates to have a good cross-correlation with that template [13].

There are several questions that must be answered in order to determine the feasibility of a matched filtering search strategy and, if feasible, to implement it. Which parameters significantly affect the wave form? How should the spacing of the template parameters (lattice points) be chosen? Is there a parametrization that is in some sense “preferred” by the template wave forms? How many templates are needed to cover a given region of interest in the parameter space, and how much computing power and memory will it cost to process the data through them? In the case of inspiraling compact binaries, the full general-relativistic wave forms are not exactly known, but are instead approximated (e.g., using the post-Newtonian scheme); and we must also ask, what approximation to the true wave forms is good enough?

All of these questions have been addressed in recent years, at least at some level. The current state of affairs is summarized by the following brief review of the literature:

The standard measure for deciding what class of wave forms is good enough is the *fitting factor* (FF) introduced by Apostolatos [14]. The fitting factor is effectively the fraction of optimal signal-to-noise-ratio obtained when filtering the data with an approximate family of templates. Because binaries are (on large scales) uniformly distributed in space and because the signal strength scales inversely with distance, the fraction of event rate retained is approximately FF^3 . Therefore it has become conventional to regard $FF=97\%$ —i.e., 10% loss of event rate—as a reasonable goal. Using the standard post-Newtonian expansion in the test-mass case (i.e., when one body is much less massive than the other so that the wave forms can be computed with arbitrarily high precision using the Teukolsky formalism), Droz and Poisson [15] found that second post-Newtonian (2PN) signals had fitting factors of 90% or higher. Damour, Iyer, and Sathyaprakash [16] have devised a new way of rearranging the usual post-Newtonian expansion (similar to the way Padé approximants rearrange the coefficients of a Taylor expansion) to take advantage of physical intuition in constructing templates. They find fitting factors of 95% or higher for the 2PN templates. Research underway by the authors of [9] will lead to 3PN templates that should easily achieve $FF>97\%$.

Several people [17,14,18] have shown that it is insufficient to use templates that depend on just one shape parameter (the “chirp mass,” which governs the rate of frequency sweep at Newtonian-quadrupole order). To achieve $FF>90\%$ one must include the masses of both objects as template parameters, as was done in the above 2PN analyses

[15,16], and as is being done in the forthcoming 3PN templates.

Apostolatos [14,19] showed that, for binaries whose components spin rapidly about their own axes which are orthogonal to the orbital plane so that there is no precession, neglecting the spin parameters (i.e., using two-mass-parameter wave forms based on the theory of spinless binaries) degraded the fitting factors by less than 2%. With precession the situation is much more complicated, and data analysis algorithms are as yet poorly developed: It is clear that there are interesting corners of parameter space (most especially a neutron star in a substantially nonequatorial, precessing orbit around a much more massive, rapidly spinning black hole) in which the two-mass-parameter spinless wave forms give $FF\leq 90\%$; to search for such binaries will require wave forms with three or more parameters [19]. However, the 2PN (or 3PN) two-mass-parameter wave forms do appear to cover adequately a significant portion of the parameter space for precessing binaries [14].

Sathyaprakash [20] showed that in computations with the two-mass-parameter wave forms, the best coordinates to use on the parameter space are not the two masses, but rather the inspiral times from some fiducial frequency to final merger, as computed at Newtonian and first post-Newtonian order. Working with the *restricted* first post-Newtonian wave forms (see below) he found that the effective dimension of the parameter space is nearly one.

Sathyaprakash and Dhurandhar [21–23] developed a criterion for putting templates at discrete points on a grid in parameter space and numerically implemented their criterion for a one-parameter (Newtonian) family of templates and for simple noise models. They introduced the concept of what Owen [24] later called the *minimal match* (analogous to the fitting factor) as a measure of how well a discrete set of templates covers the parameter space and estimated the computational costs for an on-line search.

Owen [24], building on the work of Sathyaprakash and Dhurandhar, defined a metric on the parameter space from which one can semi-analytically calculate (i) the template spacing needed to achieve a desired minimal match, (ii) the total number of templates needed, and (iii) the computational requirements to keep up with the data—for any family of wave forms and any interferometer noise spectrum. Owen combined this metric-based formalism with computational counting procedures from Schutz [25] to estimate the computational requirements for LIGO searches based on two-parameter 1PN templates. These estimates were confirmed by Apostolatos [19] using a numerical method in the vein of (but more sophisticated than) the previous work of Sathyaprakash and Durandhar [21–23]. Apostolatos also showed that a search for precessing binaries that fully covers all the nooks and crannies of the precessional parameter space, using currently available templates and techniques, is prohibitively costly.

Mohanty and Dhurandhar [26,27] have studied hierarchical search strategies. Such strategies reduce computational costs by making a first pass of the data through a coarsely-spaced template grid and a low signal-to-noise threshold to identify candidate signals. Each candidate flagged by the first

pass is examined more closely with a second, finely-spaced grid of templates and a higher threshold to weed out false alarms. Such strategies can reduce the total computational requirements by roughly a factor 25.

The purpose of this paper is to refine and update the analyses by Owen [24] for the two-parameter, spinless templates that are likely to be used for binary-inspiral searches in ground-based interferometers. This refinement is needed because the kilometer-scale interferometers will begin taking data in about 2 years (preliminary, engineering run); people are even now designing software to implement the simplest matched filtering search algorithm; and in the context of these implementations, the factor of 3 accuracy attempted in Ref. [24] is no longer adequate. The numbers that are derived in this paper should establish a reliable baseline cost to which more sophisticated search strategies (e.g., hierarchical searches) can be compared.

The substantial differences between this paper and Ref. [24] are that we now (i) approximate the phase evolution of the inspiral wave form to 2PN rather than 1PN order; (ii) give results for the noise spectra of several more interferometers; and (iii) use a better estimate of the sampling frequency needed for each interferometer. We assume the following fiducial search: a minimal match of 97% (corresponding to 10% loss of event rate due to coarse parameter space coverage), second post-Newtonian wave forms, and templates made for objects of minimum mass $m_{\min}=0.2M_{\odot}$ and up.

Our results for the computational requirements are given in Tables II–IV. These tables show that the initial LIGO interferometers need about twice as many templates and triple the computational power estimated in Ref. [24]. These increases result mainly from using 2PN wave forms rather than the (clearly inadequate) 1PN, and from using a higher sampling rate (as, it turns out, is required to keep time-step discretization error from compromising the 97% minimal match). GEO600 requires slightly more templates and power than LIGO because of its flatter noise spectrum, while TAMA requires slightly less because its sensitivity is limited to higher frequencies where there are fewer cycles. Initial VIRGO, with its extremely broad and flat spectrum, requires about the same as advanced LIGO.

The rest of this paper is organized as follows. In Sec. II we analyze the application of matched filtering to a search for inspiraling binaries and summarize the method of Ref. [24] which uses differential geometry to answer important questions about such a search. We use this method in Sec. III to estimate the computational costs and other requirements of a matched-filtering binary search for LIGO, VIRGO, GEO600, and TAMA. In Sec. IV we illustrate a detailed example of a template placement algorithm, and in Sec. V we discuss our results.

II. FORMALISM

This section summarizes material previously presented in [24] with several incremental improvements. We begin by introducing some notation.

The Fourier transform of a function $h(t)$ is denoted by $\tilde{h}(f)$, where

$$\tilde{h}(f) \equiv \int_{-\infty}^{\infty} dt e^{i2\pi ft} h(t). \quad (2.1)$$

We write the interferometer output $h(t)$ as the sum of noise $n(t)$ and a signal $\mathcal{A}s(t)$, where we have separated the signal into a dimensionless, time-independent amplitude \mathcal{A} and a ‘‘shape’’ function $s(t)$ which is defined to have unit norm [see Eq. (2.4) below].

The strain power spectral noise density of an interferometer is denoted by $S_h(f)$. We use the one-sided spectral density, defined by

$$E[\tilde{n}(f_1)\tilde{n}^*(f_2)] = \frac{1}{2} \delta(f_1 - f_2) S_h(|f_1|), \quad (2.2)$$

where $E[\]$ denotes the expectation value over an ensemble of realizations of the noise and an asterisk denotes complex conjugation.

We use geometrized units, i.e., Newton’s gravitational constant G and the speed of light c have values of unity.

A. Matched filtering

First we flesh out the Introduction’s brief description of matched filtering. In the simplest idealization of matched filtering, the filtered signal-to-noise ratio is defined by [12]

$$\frac{S}{N} \equiv \frac{\langle h, u \rangle}{\text{rms} \langle n, u \rangle}. \quad (2.3)$$

Here u is the template wave form used to filter the data stream h , and the inner product

$$\langle a, b \rangle \equiv 4 \text{Re} \left[\int_0^{\infty} df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_h(f)} \right] \quad (2.4)$$

is the noise-weighted cross-correlation between a and b (cf. [28]). The denominator of Eq. (2.3) is equal to $\sqrt{\langle u, u \rangle}$, the norm of u (see Sec. II B of Ref. [28] for a proof). Because the norm of u cancels out of Eq. (2.3), we can simplify our calculations without loss of generality by considering all templates to have unit norm.

When searching for a parametrized family of signals the situation is somewhat more complicated. The parameter values of the signals are not known in advance; therefore one must filter the data through many templates constructed at different points in the parameter space. To develop a strategy for searching the parameter space, one must know how much the S/N is reduced by using a template whose parameter values differ from those of the signal. Neglecting fluctuations due to the noise, the fraction of the optimal S/N obtained by using the wrong parameter values is given by the *ambiguity function*

$$A(\boldsymbol{\lambda}, \boldsymbol{\Lambda}) \equiv \langle u(\boldsymbol{\lambda}), u(\boldsymbol{\Lambda}) \rangle \quad (2.5)$$

(see, e.g., Chaps. XIII and X of Ref. [12]). Here λ and Λ are the parameter vectors of the signal and template (it does not matter which is which). The ambiguity function, as its name implies, is a measure of how distinguishable two wave forms are with respect to the matched filtering process. It can be regarded as an inner product on the wave form parameter space and is fundamental to the theory of parameter estimation [12,29].

For the purposes of a search for inspiraling compact binaries, the ambiguity function isn't quite what is needed. This is because the test statistic (for a given set of parameter values θ) is not given by Eq. (2.3), but rather by

$$\max_{\phi_c, t_c} \frac{\langle h, u(\theta) e^{i(2\pi f t_c - \phi_c)} \rangle}{\text{rms} \langle n, u(\theta) \rangle}. \quad (2.6)$$

Here ϕ_c and t_c are respectively the coalescence time and coalescence phase. We separate these parameters out from the rest: $\lambda = (\phi_c, t_c, \theta)$, where θ is the vector of *intrinsic parameters* that determine the shape of the wave form and ϕ_c and t_c are *extrinsic parameters* [24] (also referred to as *kinematical* and *dynamical* parameters [20], respectively). The practical difference is that maximization over the extrinsic parameters is performed automatically by Fourier transforming, taking the absolute value, and looking for peaks. The use of Eq. (2.6) as a detection statistic suggests the definition of a modified ambiguity function known as the *match* [24]

$$M(\theta_1, \theta_2) \equiv \max_{\phi_c, t_c} \langle u(\theta_1), u(\theta_2) e^{i(2\pi f t_c - \phi_c)} \rangle, \quad (2.7)$$

where the templates u are assumed to have unit norm. The use of this match function rather than the ambiguity function takes into account the fact that a search can benefit from systematic errors in the extrinsic parameters.

B. Applications of differential geometry

The match (2.7) can be regarded as an inner product on the space of template shapes and intrinsic template parameters, and correspondingly one can define a metric on this space [24]:

$$g_{ij}(\theta) \equiv - \left. \frac{1}{2} \frac{\partial^2 M(\theta, \Theta)}{\partial \Theta^i \partial \Theta^j} \right|_{\Theta^k = \theta^k}. \quad (2.8)$$

The metric (2.8) is derived from the match (2.7) in the same way the information matrix Γ_{ij} is derived from the ambiguity function [29], and plays a role in signal detection similar to that played by the information matrix in parameter estimation [30]. The g_{ij} can be derived by expanding $M(\theta, \Theta)$ about $\Theta = \theta$, or equivalently by projecting Γ_{ij} on the subspace orthogonal to ϕ_c and t_c .

The g_{ij} can be used to approximate the match in the regime $1 - M \ll 1$ by

$$M(\theta, \theta + \Delta\theta) \approx 1 - g_{ij} \Delta\theta^i \Delta\theta^j, \quad (2.9)$$




FIG. 1. Comparison of the full match to the quadratic approximation in the case of (a) first post-Newtonian and (b) second post-Newtonian wave forms. In both cases the noise spectrum is LIGO I (see Table I). The elliptical solid line is the 97% contour of the match with a reference wave form (in the center of the ellipse) where the match is given analytically by the metric in the quadratic approximation [Eq. (2.9)]. The dots are locations of the same contour given by constructing stationary phase wave forms and numerically computing the full match from them. The reference wave form is from two $1.4M_{\odot}$ objects [the mass parameters τ_0 and τ_1 are defined in Eq. (3.2)]. The quadratic approximation is safe, in the sense that its 97% contour always lies inside the numerical contour. The quadratic approximation also works well for high mass binaries, provided the numerical and analytical contours use the same coalescence frequency.

which is simply another way of writing the Taylor expansion of $M(\theta, \theta + \Delta\theta)$ about $\Delta\theta = 0$. (The first derivative term vanishes because M takes its maximum value of unity at $\Delta\theta = 0$.) We find that the quadratic approximation (2.9) is good typically for $M \approx 0.95$ or greater, though this depends on the wave form and noise spectrum used. Experience suggests that the quadratic approximation generally underestimates the true match; and thus the spacings and numbers of templates we calculate using Eq. (2.9) err on the safe side. See Fig. 1 for an example.

TABLE I. Analytical fits to noise power spectral densities $S_h(f)$ of the interferometers treated in this paper. Here S_0 is the minimum value of $S_h(f)$, and f_0 is the frequency at which the minimum value occurs. For our purposes $S_h(f)$ can be treated as infinite below the seismic frequency f_s . The high-frequency cutoff f_u is chosen so that the loss of signal-to-noise ratio due to finite sampling rate $2f_u$ is 0.75% (see text).

Detector	Fit to noise power spectral density	S_0 (Hz ⁻¹)	f_0 (Hz)	f_s (Hz)	f_u (Hz)
LIGO I	$S_0/3 [(f_0/f)^4 + 2(f/f_0)^2]$	4.4×10^{-46}	175	40	1300
LIGO II	$S_0/11 \{2(f_0/f)^{9/2} + 9/2[1 + (f/f_0)^2]\}$	7.9×10^{-48}	110	25	900
LIGO III	$S_0/5 \{(f_0/f)^4 + 2[1 + (f/f_0)^2]\}$	2.3×10^{-48}	75	12	625
VIRGO	$S_0/4 [290(f_s/f)^5 + 2(f_0/f) + 1 + (f/f_0)^2]$	1.1×10^{-45}	475	16	2750
GEO600	$S_0/5 [4(f_0/f)^{3/2} - 2 + 3(f/f_0)^2]$	6.6×10^{-45}	210	40	1450
TAMA	$S_0/32 \{(f_0/f)^5 + 13(f_0/f) + 9[1 + (f/f_0)^2]\}$	2.4×10^{-44}	400	75	3400

In the limit of close template spacing, Eq. (2.9) leads to a simple, analytical way of placing templates on a lattice. We discuss this in some detail in Sec. IV, but for now turn our attention to the use of the quadratic approximation in calculating the number of templates needed for a lattice.

C. Computational costs

If the number \mathcal{N} of templates needed to cover a region of interest is large, it is well approximated by the ratio of the proper volume of the region of interest to the proper volume per template ΔV ,

$$\mathcal{N} = (\Delta V)^{-1} \int d^D \theta \sqrt{\det \|g_{ij}\|}, \quad (2.10)$$

where D is the dimension of the parameter space [24]. Equation (2.10) underestimates \mathcal{N} when not in the limit $\Delta V \rightarrow 0$ ($\mathcal{N} \rightarrow \infty$). The reason is template *spill over*, i.e., the fact that in any real algorithm for laying out templates, those on the boundaries of the region of interest will to some extent cover regions just outside. This effect is small in the limit of many templates because it goes as the surface-to-volume ratio of the region of interest.

The proper volume per template, ΔV , depends on the packing algorithm used, which in turn depends on the number D of dimensions (see Sec. IV). For $D=2$, the optimal packing is a hexagonal lattice, and thus

$$\Delta V = \frac{3\sqrt{3}}{2} (1 - MM), \quad (2.11)$$

where MM is the *minimal match* parameter defined in Ref. [24] as the match between signal and template in the case when the signal lies equidistant between all the nearest templates (i.e., the worst-case scenario). There is no packing scheme which is optimal for all D , but it is always possible (though inefficient) to use a hypercubic lattice, for which

$$\Delta V = (2\sqrt{(1 - MM)/D})^D. \quad (2.12)$$

For inspiraling compact binaries, Ref. [24] has spelled out a detailed prescription for obtaining the g_{ij} needed to evaluate the proper volume integral in Eq. (2.10). In summary, first one obtains a metric including the t_c parameter,

$$\gamma_{\alpha\beta} = \frac{1}{2} (\mathcal{J}[\psi_\alpha \psi_\beta] - \mathcal{J}[\psi_\alpha] \mathcal{J}[\psi_\beta]), \quad (2.13)$$

where ψ_α is the gradient of the wave form phase Ψ in the parameter space of intrinsic parameters plus t_c and the moment functionals

$$\mathcal{J}[a] \equiv \frac{\langle f^{-7/3}, a(f) \rangle}{\langle f^{-7/3}, 1 \rangle} \quad (2.14)$$

can be expanded (for binary chirp wave forms) in terms of the noise moments [31]

$$J(p) \equiv \frac{\langle (f/f_0)^{-p/3}, 1 \rangle}{\langle (f/f_0)^{-7/3}, 1 \rangle} \quad (2.15)$$

where f_0 is the frequency of the minimum of $S_h(f)$ [32]. Then one projects out the coalescence time t_c to obtain

$$g_{ij} = \gamma_{ij} - \gamma_{0i} \gamma_{0j} / \gamma_{00}. \quad (2.16)$$

Once \mathcal{N} has been found it is a relatively straightforward matter to calculate the CPU power and storage required to process all the templates in an on-line search. The interferometer data stream is broken up for processing into segments of D samples (real numbers), such that $D \gg F$ where F is the length (in real numbers) of the longest filter. (See Schutz [25] for a discussion of the optimization of D/F , taking into account the fact that successive data segments must overlap by at least F to avoid circular correlations in the Fourier transform.) Using the operations count for a real Fourier transform [25], filtering the data segment through \mathcal{N} templates of length F requires

$$\mathcal{N}D(16 + 3\log_2 F) \quad (2.17)$$

floating point operations. If we take the sampling frequency to be $2f_u$ (see Sec. III and Table I), the computational power required to keep pace with data acquisition is

$$\mathcal{P} \simeq \mathcal{N}f_u(32 + 6\log_2 F) \quad (2.18)$$

flops (floating point operations per second).

III. COMPUTATIONAL COST USING RESTRICTED 2PN TEMPLATES

In this section, using the geometric formalism summarized in Sec. II, we calculate the number \mathcal{N} of templates required to cover a region of interest as a function of the minimal match. We then use this number to calculate the computational cost of filtering a single interferometer's output through all these templates in an on-line search.

A. Functional form of the templates

We construct our wave form templates using two intrinsic parameters based on the masses of the binary's components. Inspiral wave forms in principle can be strongly affected by several other parameters: spins of the two components, orbital eccentricity, and several angles describing the orientation of the binary with respect to the interferometer. However, it is believed that two-parameter templates will be adequate to search for most binaries for the following reasons.

(i) Based on models of the evolution of currently known binary pulsars, it is expected [9] that typical NS-NS binaries will have spins of negligible magnitude ($\text{spin}/\text{mass}^2 \ll 1$). Apostolatos [19] has shown that, even if the magnitudes of the spins are large, their effect on S/N is small (reduces FF by less than 2% if the orbit and the spin vectors do not precess). He has also shown that precession will not reduce FF below 90% except in the relatively small region of parameter space containing binaries with a neutron star orbiting a more massive, rapidly rotating ($\text{spin}/\text{mass}^2 \sim 1$) black hole with orbital angular momentum inclined by more than about 30 degrees to the black hole's spin.

(ii) It has long been known [33] that gravitational radiation reaction circularizes all but the most eccentric orbits on a time scale much smaller than the lifetime of the binary if the progenitor system was the same binary. This may not be true, however, in the case of close binaries formed by capture in densely populated environments, e.g., galactic nuclei-globular clusters.

(iii) The angles make no difference in our analysis because we use the *restricted post-Newtonian* approximation [28], in which the phase evolution of the inspiral wave form is followed to a high post-Newtonian order but the amplitude is only followed to lowest order. In this approximation, the combined effect of the angles is to multiply the wave form by a constant amplitude and phase factor, which does not affect the choice of search templates [21]. Presently it is believed [28] that the restricted post-Newtonian approximation will be good enough for data analysis of ground-based interferometers.

The standard post-Newtonian expansion of the wave form phase is given as a function of mass parameters based on the standard astronomical choices M (total mass) and μ (reduced mass). In order to clearly isolate test-mass terms (i.e., those that remain when one body is much less massive than the other), the symmetric mass ratio $\eta = \mu/M$ is typically used instead of μ . In terms of M and η , the second post-Newtonian wave form phase can be calculated from the energy loss formula of Blanchet *et al.* [9] as

$$\begin{aligned} \Psi(f; M, \eta) = & \frac{3}{128} (\pi M f)^{-5/3} \eta^{-1} \left[1 + \frac{20}{9} \left(\frac{743}{336} + \frac{11}{4} \eta \right) \right. \\ & \times (\pi M f)^{2/3} - 16\pi (\pi M f) + 10 \left(\frac{3\,058\,673}{1\,016\,064} \right. \\ & \left. \left. + \frac{5\,429}{1\,008} \eta + \frac{617}{144} \eta^2 \right) (\pi M f)^{4/3} \right] \end{aligned} \quad (3.1)$$

[cf. Eq. (3.6) of Poisson and Will [31]]. However, M and η are inconvenient parameters for our purposes because, when they are used as the parameter-space coordinates, the values of the metric components vary strongly over parameter space, making calculations unnecessarily difficult and prone to numerical error. Equation (3.1) uses the stationary phase approximation to the Fourier transform of $h(t)$, which is known to simplify the wave form considerably while causing negligible loss of signal-to-noise ratio [34].

In earlier analyses [20,29,24] it was found more convenient to use as parameters the Newtonian and 1PN *chirp times*

$$\tau_0 = \frac{5}{256} M^{-5/3} (\pi f_0)^{-8/3} \eta^{-1}, \quad (3.2a)$$

$$\tau_1 = \frac{5}{192} M^{-1} (\pi f_0)^{-2} \left(\frac{743}{336\eta} + \frac{11}{4} \right), \quad (3.2b)$$

which are respectively the Newtonian and 1PN contributions to the time it takes the instantaneous gravitational-wave frequency to (formally) evolve from f_0 to infinity. The chirp times are more convenient than the usual mass parameters because, when they are chosen as parameter-space coordinates, at 1PN order the metric components are constant. Assuming the post-Newtonian expansion has reasonable convergence properties, one would expect the metric components in these coordinates to remain nearly constant at higher post-Newtonian orders (and indeed we find this is so). However, at higher than 1PN order one cannot write the wave form phase analytically in terms of τ_0 and τ_1 . To remedy this, following Mohanty [27], we base our second parameter on the 1.5PN chirp time $\tau_{1.5}$ (see [29] for a definition). More specifically, we introduce new dimensionless coordinates in parameter space. We define

$$\theta^1 = 2\pi f_0 \tau_0 = \frac{5}{128} (\pi M f_0)^{-5/3} \eta^{-1}, \quad (3.3a)$$

$$\theta^2 = 2\pi f_0 \tau_{1.5} = \frac{\pi}{4} (\pi M f_0)^{-2/3} \eta^{-1}, \quad (3.3b)$$

which can be inverted to obtain

$$M = \frac{5}{32\pi^2 f_0} \frac{\theta^2}{\theta^1}, \quad (3.4a)$$

$$\eta = \left[\frac{16\pi^5 (\theta^1)^2}{25 (\theta^2)^5} \right]^{1/3}. \quad (3.4b)$$

This choice of (θ^1, θ^2) lets us write the wave form phase analytically while keeping the metric components from varying too strongly; therefore it is convenient for calculating numbers of templates (see below). However, for other purposes (τ_0, τ_1) are just as convenient, and to be consistent with the literature we will use them.

B. Noise spectra

In this paper we consider the noise spectra of the four large- and intermediate-scale interferometer projects, LIGO, VIRGO, GEO600, and TAMA. For LIGO we use three noise spectra corresponding to three interferometer configurations, the ‘‘first interferometers’’ [1] (which are planned to perform a gravitational-wave search in 2002–2003), the ‘‘enhanced interferometers’’ [8] (which are likely to be carrying out searches in the mid 2000’s), and the ‘‘advanced interferometers’’ [1] (which are thought representative of the type of detector that might operate a few years after the enhanced ones). For convenience we abbreviate these, respectively, as LIGO I, LIGO II, and LIGO III. We use the VIRGO and GEO600 noise spectra given in Refs. [2,3]. For TAMA we use the noise spectrum given by Fujimoto in Ref. [35].

It is desirable to have simple analytical fits to the noise power spectral densities used. Kip Thorne and Scott Hughes (private communications) have provided us with fits to LIGO I and LIGO III, respectively, and a fit to LIGO II was derived in Ref. [36]. We have constructed our own fits to the remaining noise spectra listed in the previous paragraph. All of these analytical fits to the noise spectra are tabulated in Table I.

The shape of each noise spectrum determines natural low- and high-frequency cutoffs for the matched filtering integrals. The low-frequency cutoff f_s is defined as the frequency above which 99% of S^2/N^2 is obtained; we call this the *seismic frequency* and denote it with a subscript s because it is typically near the frequency at which seismic noise causes the noise power spectral density $S_h(f)$ to begin

rising sharply. Because integration below this frequency contributes little to the detectability of signals and costs much in terms of total number of templates and computational resources, we assume the templates to begin with gravitational-wave frequency f_s .

The high-frequency cutoff f_u , as pointed out by Éanna Flanagan (private communication), needs to be high enough that the S/N degradation due to discrete time steps $\Delta t = 1/(2f_u)$ in the data analysis is less than that due to discrete choices of the templates’ intrinsic parameters. Quantitatively, this requires that

$$\gamma_{00}/(4f_u)^2 < 1 - MM, \quad (3.5)$$

where γ_{00} is the t_c - t_c component of the metric before t_c is projected out [see Ref. [24], Eq. (2.23)]. At any post-Newtonian order we have $\gamma_{00} = (2\pi f_0)^2 [J(1) - J(4)^2]/2$ [see Ref. [24], Eq. (2.27)] and thus

$$f_u > \pi f_0 \sqrt{\frac{J(1) - J(4)^2}{8(1 - MM)}}. \quad (3.6)$$

[Strictly speaking, the noise moments $J(1)$ and $J(4)$ in this expression should be integrated up to f_u rather than infinity and f_u should be chosen by iteration, but this is a small correction.] We have chosen f_u to be twice the right-hand side of this expression so that the loss due to sampling is 1/4 the loss due to discrete values of the intrinsic parameters. There is probably a more clever way of optimizing f_u , but this is a first cut.

C. Number of templates and computational cost

We now proceed with the calculation of the number of templates needed to perform a single-pass, on-line search for gravitational-wave signals of the form in Eq. (3.1), for the noise spectra in Table I.

The first task is to obtain the intrinsic-parameter metric. We rewrite the wave form phase at 2PN order as

$$\begin{aligned} \Psi = & \frac{3}{5} \theta^1 \left(\frac{f}{f_0} \right)^{-5/3} + \left[\frac{11\pi}{12} \frac{\theta^1}{\theta^2} + \frac{743}{2016} \left(\frac{25}{2\pi^2} \right)^{1/3} (\theta^1)^{1/3} (\theta^2)^{2/3} \right] \left(\frac{f}{f_0} \right)^{-1} - \frac{3}{2} \theta^2 \left(\frac{f}{f_0} \right)^{-2/3} \\ & + \left[\frac{617\pi^2}{384} \frac{\theta^1}{(\theta^2)^2} + \frac{5429}{5376} \left(\frac{25\pi}{2} \right)^{1/3} \left(\frac{\theta^1}{\theta^2} \right)^{1/3} + \frac{15\,293\,365}{10\,838\,016} \left(\frac{5}{4\pi^4} \right)^{1/3} \frac{(\theta^2)^{4/3}}{(\theta^1)^{1/3}} \right] \left(\frac{f}{f_0} \right)^{-1/3}. \end{aligned} \quad (3.7)$$

From Eqs. (3.7) and Eqs. (2.13)–(2.16) it is straightforward to derive the metric components g_{ij} with symbolic manipulation software. However, the general expressions for the g_{ij} (and even the γ_{ij}) are too lengthy and complicated to display here.

Next, we obtain the proper volume

$$V = \int d^2\theta \sqrt{\det \|g_{ij}\|} \quad (3.8)$$

of the region of interest. The boundaries of this region are set by the range of masses $[m_{\min}, m_{\max}]$ of the individual objects

in binaries. It turns out that only m_{\min} has a strong influence on V [21,23,24]. For 1PN wave forms the intrinsic parameter space is flat (i.e., $\det \|g_{ij}\|$ is constant) and thus V can be obtained analytically using the coordinates (τ_0, τ_1) . Beyond 1PN order in the wave form, the intrinsic parameter space is not flat, and thus it is convenient to calculate the proper volume numerically, e.g., by a Monte Carlo method. In order

TABLE II. Numbers of templates required to cover parameter space at a minimal match of 97% with a rectangular lattice. The region of interest is that inhabited by binaries with component masses greater than m_{\min} . We use restricted post-Newtonian (PN) templates whose phase evolution is accurate to the indicated order.

Detector	$m_{\min}=0.2M_{\odot}$		$m_{\min}=1M_{\odot}$	
	1PN	1.5PN	2PN	2PN
LIGO I	2.3×10^5	5.6×10^5	4.8×10^5	1.1×10^4
LIGO II	1.0×10^6	1.9×10^6	1.7×10^6	4.0×10^4
LIGO III	5.7×10^6	7.9×10^6	6.9×10^6	1.7×10^5
VIRGO	5.8×10^6	1.1×10^7	9.3×10^6	2.2×10^5
GEO600	4.2×10^5	9.7×10^5	8.3×10^5	1.9×10^4
TAMA	5.1×10^4	1.7×10^5	1.4×10^5	3.1×10^3

to make the integrand as nearly constant as possible we use the metric tensor transformation law to switch from (θ^1, θ^2) coordinates to (τ_0, τ_1) [because in these coordinates the metric components vary more slowly over the region of interest than in (θ^1, θ^2)]. Finally, to obtain the number of templates we divide the proper volume V by the proper volume per template ΔV . For a rectangular lattice in two dimensions, $\Delta V = 2(1 - MM)$ [cf. Eq. (2.12)]. In Table II we give values of \mathcal{N} calculated for the noise spectra given in Table I assuming a rectangular lattice. For a hexagonal lattice, ΔV is given by Eq. (2.11), resulting in a reduction of about 20% in \mathcal{N} , \mathcal{P} , and \mathcal{S} . However, in practice much of this reduction may be offset by the details of actually constructing a lattice (see Sec. IV) and therefore we make our estimates for the more conservative rectangular case.

These numbers are easily translated into the computational costs shown in Tables III and IV: To a good approximation, the length of the longest filter (with $m_1 = m_2 = m_{\min}$, i.e., equal mass binaries so that $\eta = 1/4$) is given by

$$F \simeq 2f_u \tau_0 (f_0/f_s)^{8/3} = \frac{5}{32} f_u (\pi f_s)^{-8/3} (2m_{\min})^{-5/3}. \quad (3.9)$$

The storage required for all of the templates is then roughly given by

$$\mathcal{S} \simeq \mathcal{N}F, \quad (3.10)$$

TABLE III. Computational costs obtained from the numbers of templates given in Table II using 2PN templates and $m_{\min} = 0.2M_{\odot}$. The symbol F denotes the length (in real numbers) of the longest template.

Detector	$\log_2 F$	CPU power \mathcal{P} (flops)	Storage \mathcal{S} (reals)
LIGO I	21	9.9×10^{10}	1.0×10^{12}
LIGO II	22	2.5×10^{11}	7.1×10^{12}
LIGO III	25	7.8×10^{11}	2.3×10^{14}
VIRGO	26	4.8×10^{12}	6.2×10^{14}
GEO600	21	1.9×10^{11}	1.8×10^{12}
TAMA	20	7.3×10^{10}	1.5×10^{11}

TABLE IV. Computational costs as in Table III, except here we assume $m_{\min} = 1M_{\odot}$.

Detector	$\log_2 F$	CPU power \mathcal{P} (flops)	Storage \mathcal{S} (reals)
LIGO I	17	1.9×10^9	1.4×10^9
LIGO II	19	5.2×10^9	2.1×10^{10}
LIGO III	21	1.7×10^{10}	3.6×10^{11}
VIRGO	22	9.8×10^{10}	9.1×10^{11}
GEO600	17	3.7×10^9	2.5×10^9
TAMA	16	1.3×10^9	2.0×10^8

and the computational power \mathcal{P} is given by Eq. (2.18). (We use the word ‘‘storage’’ although sometimes it may be desirable to instead generate templates as needed; whether stored or generated \mathcal{S} represents a number of bytes which must be obtained at a cost of CPU cycles if not storage, and can be an obstacle if the number is too high.)

D. Scaling laws

Approximate scaling laws can be obtained as follows. At 1PN order the intrinsic parameter space is flat, and thus the proper volume of the region of interest scales as the (dimensionless) coordinate volume, i.e.,

$$V \sim f_0 \tau_0 \times f_0 \tau_1 \sim m_{\min}^{-8/3} f_0^{-8/3}. \quad (3.11)$$

At higher post-Newtonian orders the metric components with respect to the (τ_0, τ_1) coordinates are nearly constant, and thus this scaling still roughly holds.

From the scaling of the proper volume we can obtain the scalings of the other quantities of interest. Inserting the dependence (2.12) of the proper volume per template on the minimal match, we obtain

$$\mathcal{N} \sim (1 - MM)^{-1} m_{\min}^{-8/3} f_0^{-8/3}. \quad (3.12)$$

Taking Eq. (2.18), neglecting the weak logarithmic dependence, and noting that f_u is proportional to f_0 , we find that

$$\mathcal{P} \sim (1 - MM)^{-1} m_{\min}^{-8/3} f_0^{-5/3}. \quad (3.13)$$

Multiplying the number of templates \mathcal{N} by the length of each template F , the storage requirement scales as

$$\mathcal{S} \sim (1 - MM)^{-1} m_{\min}^{-13/3} f_0^{-5/3} f_s^{-8/3}. \quad (3.14)$$

This last scaling with m_{\min} presents a significant obstacle to efficiently searching for binaries composed of low-mass objects, such as massive compact halo objects (MACHOs) if they are low-mass ($\sim 0.5M_{\odot}$) black holes [37]. Searches for low-mass objects will likely need to generate templates as needed rather than store them, incurring additional CPU costs not addressed in this paper.

Note that the above scaling laws implicitly assume that f_0 or f_s is varied while holding the overall shape of the noise spectrum fixed. The noise spectrum of a real interferometer, being composed of many independent noise sources, is unlikely to change in such a manner except for fairly small

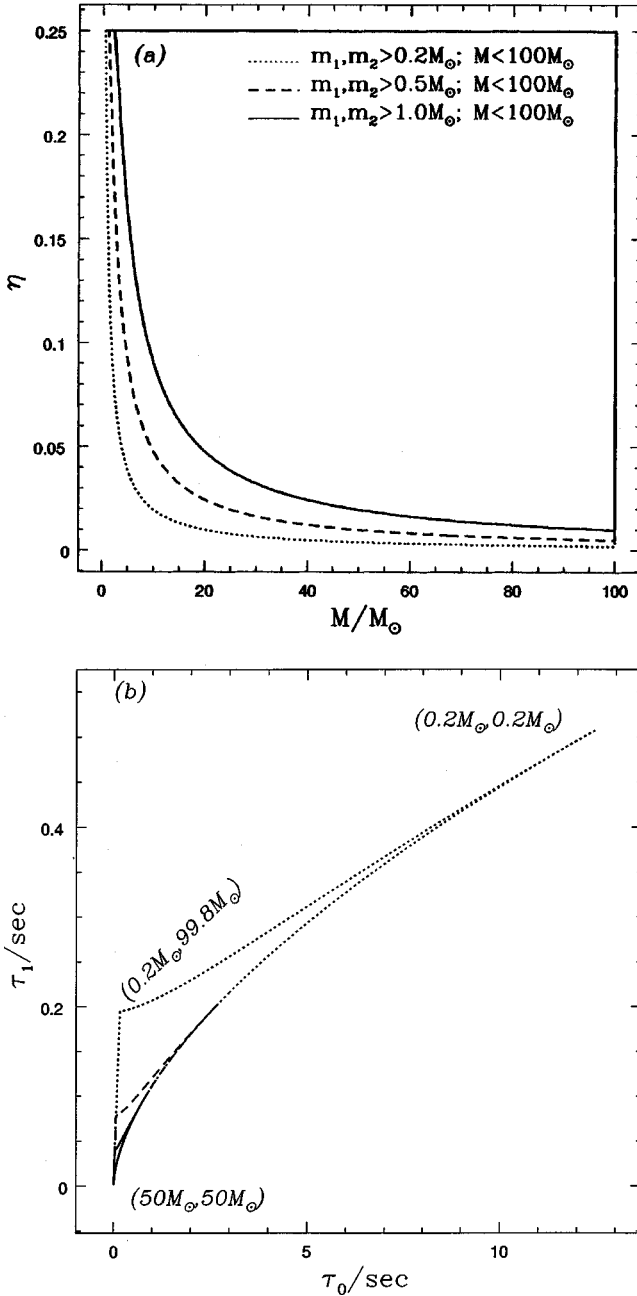


FIG. 2. The parameter space of search in terms of (a) standard binary mass parameters and (b) chirp times, for different values of the lower mass limits.

changes in f_0 or f_s . However, these scaling laws give one a rough feel for (i) how changes in a single interferometer can affect data analysis requirements, and (ii) why different interferometers have drastically different requirements.

IV. A TEMPLATE PLACEMENT ALGORITHM

In this Section we deal with the actual placement of the templates in the parameter space. The parameter space is shown in Fig. 2a in terms of the total mass M and symmetric mass ratio η and in Fig. 2b in terms of chirp times τ_0 and τ_1 , for searches in which the maximum total mass M_{\max} of

TABLE V. Dimensions of elliptical 97% match contours in the quadratic approximation [Eq. (2.9)] using the (τ_0, τ_1) coordinates. Here a and b are the semimajor and semiminor axes, and α is the angle between the τ_0 axis and the major axis of the ellipse. Here we use the LIGO I noise spectrum and second post-Newtonian templates designated in the table by the masses of the binary components.

$m_1(M_\odot)$	$m_2(M_\odot)$	a (ms)	b (ms)	α (rad)
1.4	1.4	1.19	0.177	0.802
1.4	10	1.25	0.167	0.834
10	10	1.11	0.089	0.629

the system is $100 M_\odot$ and the lower limit m_{\min} on the mass of each component star is 0.2, 0.5 or $1.0M_\odot$. The bottom line in Fig. 2b corresponds to binaries of equal mass ($\eta = 1/4$) with the rightmost point corresponding to lowest mass binaries and the leftmost to greatest mass binaries of our search. There are no binaries in the region below this line as the parameter η exceeds $1/4$ there, which is unphysical. The top and the left lines are determined by m_{\min} and M_{\max} respectively. Given the minimum mass of the component stars and the maximum total mass, the parameter space of binaries is completely fixed. The volume of the parameter space (and the corresponding number of templates required) is a sharp function of the lower cutoff in the masses of the component stars and increases, as we have seen, as $m_{\min}^{-8/3}$.

The shape of the parameter space is rather complicated and attention needs to be paid in the placement of templates so that the inevitable *spill over* (see below) is minimal. Our implementation of the filter placement is motivated by the following astrophysical consideration: The observed neutron stars are all of roughly equal mass [38]. It is therefore to be expected that many inspiral signals will come from equal mass binaries. Consequently, we optimize the filter placement for equal mass binaries. This is achieved by beginning our template placement along the $\eta = 1/4$ line. The span of each template is taken to be the largest rectangle (in a coordinate system in which the metric is locally diagonal) that can be inscribed inside the minimal match ellipse (we take $MM = 97\%$; see Table V for sample dimensions). Note that in Fig. 3 the spans do not appear rectangular because they are sheared by transforming from coordinates in which the metric is locally diagonal (see below) to the (τ_0, τ_1) coordinates. We begin with the leftmost point on the bottom edge of the parameter space of Fig. 2b, corresponding to the most massive binary of our search with the shortest chirp time. Such a system will, of course, consist of equal mass bodies and therefore our starting point is on the $\eta = 1/4$ curve. The next template is placed on the $\eta = 1/4$ curve at that location where the left edge of its rectangle touches the right edge of the previous template's rectangle. In a sense this is a straightforward generalization of placement of filters discussed in Ref. [24] along grid lines that are not necessarily straight.

This optimal translation of templates is most easily done in a coordinate system in which the metric is locally diagonal. Let (x_0, x_1) denote such a coordinate system (found by diagonalizing the two-dimensional matrix g_{ij}) and let




FIG. 3. Illustration of optimal translation of a template along an arbitrary curve.

$f(x_0, x_1) = 0$ denote a curve in the two-dimensional x_0 - x_1 plane along which templates are to be placed. For instance, $\eta = 1/4$ in Fig. 2 is one such curve. A convenient point to begin is the point $(x_0^{(1)}, x_1^{(1)})$ at one end of the curve. In Fig. 3 we have sketched an arbitrary curve together with the first template and its span. The span of a template, for minimal matches close to 1, is an ellipse. However, its effective (non-overlapping) span, when setting up a lattice of templates, is only an inscribed polygon such as a rectangle or an irregular hexagon. In the following discussion for simplicity we consider the span to be a rectangle and hence we will be setting up a rectangular lattice. By choosing a hexagonal lattice the number of templates can be reduced by about 20%, but the reduction is less when the curve along which templates need to be placed is parallel to neither x_0 nor x_1 axis.

Given the “local” distance $(dx_0^{(1)}, dx_1^{(1)})$ between templates, we can get two points on the curve $f(x_0, x_1) = 0$ which are simultaneous solutions of

$$\{f(x_0, x_1) = 0, \quad x_n = x_n^{(1)} + dx_n^{(1)}\} \quad (4.1)$$

for $n=0$ and $n=1$ respectively. In order to cover the parameter space without leaving any “holes” it is obvious that the next template should be placed at the point that is nearer to the first template. This is how one obtains the nearest neighbor of a template.

Returning to our problem of placing templates in the parameter space of chirp times, we move along the $\eta = 1/4$ line till the binary of longest chirp time is reached. The set of templates chosen on the $\eta = 1/4$ curve form the base, on top of which we construct layers to fill the parameter space. The next row of templates is set on top of the first so that the region of interest is completely covered by the spans—which means that the first few templates will be located outside that




FIG. 4. Choice of (1PN) templates in the space of chirp times for the LIGO I noise spectrum and for a search of binaries with masses of component stars larger than $1M_\odot$ and total mass no greater than $100M_\odot$.

region, and that the total number of templates must be modestly greater than naively estimated by Eq. (2.10) (see Fig. 4).

V. DISCUSSION

In this paper we have discussed the problem of optimally placing templates in a binary inspiral signal search. For our templates we have used the restricted second post-Newtonian wave form. We have made estimates of the computational costs of an on-line search of the inspiral wave form for all the large- and mid-scale interferometric detectors now under construction. We have addressed several important issues: (i) the density of templates in the parameter space, (ii) the set of parameters most suitable for easy placement of templates, and (iii) the number of templates and computational resources needed to analyze the data on-line.

These estimates should serve as a baseline for explorations of other data analysis strategies, and some of the techniques here could be incorporated into other strategies. For example, it is now recognized that a substantial reduction in computational cost can be achieved by carrying out a hierarchical search [26,27]. For a two-step hierarchical search strategy, in the first step a sparsely filled family of templates is used, with a threshold lower than what is acceptable based on the expected number of false alarms. Those events which cross this trigger threshold are further examined with a finer grid of templates chosen around the template that triggered the event. In such a hierarchical search, templates chosen in the first step will essentially be the same for each data segment. However, templates in the second step need to be changed from one data segment to the next, depending on which templates from the coarse grid family produces an

“event.” It is in the case of a coarse grid that our analytical algorithm for template placement and analytical estimates of computational requirements fail and must be replaced by numerical methods that are computationally expensive. But generating filters corresponding to the first step of the hierarchical search is more or less a one-time job. A finer grid is to be chosen quite frequently (essentially each time a possible event is selected in the first step) and in this case, fortunately, the analytical techniques discussed in this paper are quite accurate and one does not have to follow the time consuming numerical placement of templates either for estimating computational costs or for actually performing the search.

There are several important problems we have not addressed in this paper which could be the topics of future work. We have not treated the problem of searching the corners of parameter space where the precessing binaries live. Although requiring some effort to seek out, these systems could prove quite informative and astrophysically interesting. The problem of searching for precessing binaries has been addressed only in a very exploratory way [19] but could

benefit from further analysis using the techniques of this paper. Also, we have used a crude relation between the minimal match and fraction of event rate lost. This could be improved by a statistical analysis such as begun by Mohanty [27]. Now that the “P-approximants” [16] have proven a promising way of building templates, it is important to examine the computational costs of using them to conduct a search.

ACKNOWLEDGMENTS

Both authors were partially supported by NSF Grant PHY-9424337. B.J.O. was also supported by the NSF graduate program. One of us (B.S.S.) would like to thank Bernard Schutz for hospitality at the Albert Einstein Institute where some of this work was carried out. We are indebted to Bruce Allen, Kent Blackburn, Tom Prince, and Kip Thorne for fruitful conversations, and we thank many members of the LIGO, VIRGO, GEO600, and TAMA projects for helping us with their respective noise curves.

-
- [1] A. Abramovici *et al.*, *Science* **256**, 325 (1992).
 [2] B. Caron *et al.*, *Class. Quantum Grav.* **14**, 1461 (1997).
 [3] H. Lück *et al.*, *Class. Quantum Grav.* **14**, 1471 (1997).
 [4] K. Tsubono, *Gravitational Wave Experiments* (World Scientific, Singapore, 1995), p. 112.
 [5] R. Narayan, T. Piran, and A. Shemi, *Astrophys. J.* **379**, L17 (1991); E. S. Phinney, *ibid.* **380**, L17 (1991). In the intervening years astronomers have estimated event rates both lower (based on statistics of the few known binary pulsars) and higher (based on theoretical evolutions of progenitor populations), but their estimates have begun to converge again on the 1991 numbers. See [6] and [7] for recent examples of the high-rate and low-rate camps, respectively.
 [6] S. F. P. Zwart and H. N. Spreew, *Astron. Astrophys.* **312**, 670 (1996).
 [7] I. H. Stairs *et al.*, *Astrophys. J.* **505**, 352 (1998).
 [8] B. Barish *et al.*, LIGO Advanced Research and Development Proposal, Caltech/MIT, 1996 (unpublished).
 [9] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman, *Phys. Rev. Lett.* **74**, 3515 (1995).
 [10] H. Tagoshi, M. Shibata, T. Tanaka, and M. Sasaki, *Phys. Rev. D* **54**, 1439 (1996).
 [11] Some physicists are in the habit of calling this optimal filtering method simply “Wiener filtering.” However, in the larger signal processing community, “Wiener filtering” refers to an altogether different technique. In order to avoid confusion we call this method “matched filtering.”
 [12] C. W. Helstrom, *Statistical Theory of Signal Detection*, 2nd edition (Pergamon Press, London, 1968).
 [13] One might try to avoid using a large family of templates by finding a set of basis functions on which to expand the signal. However, this technique does not work very well for inspiral chirps (and many other types of signals) because the number of basis functions required to achieve nearly optimal signal-to-noise ratio is comparable to the number of templates.
 [14] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
 [15] S. Droz and E. Poisson, *Phys. Rev. D* **56**, 4449 (1997).
 [16] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **57**, 885 (1998).
 [17] R. Balasubramanian and S. V. Dhurandhar, *Phys. Rev. D* **50**, 6080 (1994).
 [18] A. Królak, K. D. Kokkotas, and G. Schäfer, *Phys. Rev. D* **52**, 2089 (1995).
 [19] T. A. Apostolatos, *Phys. Rev. D* **54**, 2421 (1996).
 [20] B. S. Sathyaprakash, *Phys. Rev. D* **50**, R7111 (1994).
 [21] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
 [22] B. S. Sathyaprakash and S. V. Dhurandhar, *J. Comput. Phys.* **109**, 215 (1993).
 [23] S. V. Dhurandhar and B. S. Sathyaprakash, *Phys. Rev. D* **49**, 1707 (1994).
 [24] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 [25] B. F. Schutz, in *The Detection of Gravitational Radiation*, edited by D. Blair (Cambridge University Press, Cambridge, England, 1991), p. 406.
 [26] S. D. Mohanty and S. V. Dhurandhar, *Phys. Rev. D* **54**, 7108 (1996).
 [27] S. D. Mohanty, *Phys. Rev. D* **57**, 630 (1998).
 [28] C. Cutler and E. É. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
 [29] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996); **54**, 1869(E) (1996).
 [30] The idea of using the Fisher information matrix as a Riemannian metric dates back to the early days of matched filtering and parameter estimation: see N. Wiener, *The Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engi-*

neering Applications (Wiley, New York, 1949); and C. R. Rao, *Bull. Calcutta Math. Soc.* **37**, 81 (1945).

- [31] E. Poisson and C. M. Will, *Phys. Rev. D* **52**, 848 (1995).
- [32] Strictly speaking, the upper limit of integration in the noise moments should be the lesser of the sampling frequency and the “coalescence frequency” f_c which represents (roughly) the point where adiabatic inspiral of the bodies ends. Both of these frequencies are much greater than f_0 and thus our approximation is quite good, except on the high-mass end of template parameter space where f_c becomes comparable to f_0 . Since most templates populate the low-mass end of parameter space, the effect on the total number of templates is small. However, detailed algorithms for placing high-mass templates will have to take f_c into account—which may be problematic

since f_c is poorly known when the binary components have comparable masses. Using an infinite upper frequency is safe in the sense that the errors cause us to space templates slightly more closely than needed, i.e., we slightly overestimate the loss of signal-to-noise ratio.

- [33] P. C. Peters, *Phys. Rev.* **136**, B1224 (1964).
- [34] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen, *Phys. Rev. D* **59**, 124016 (1999).
- [35] B. Allen *et al.*, GRASP software package, available from <http://www.lsc-group.phys.uwm.edu/>.
- [36] B. J. Owen *et al.*, *Phys. Rev. D* **58**, 084020 (1998).
- [37] T. Nakamura, M. Sasaki, T. Tanaka, and K. S. Thorne, *Astrophys. J.* **487**, L139 (1997).
- [38] J. H. Taylor, *Rev. Mod. Phys.* **66**, 711 (1994).