Improving the efficiency of the detection of gravitational wave signals from inspiraling compact binaries: Chebyshev interpolation

S. Mitra* and S. V. Dhurandhar†

Inter-University Centre for Astronomy and Astrophysics, Ganeshkhind, Pune-411 007, India

L. S. Finn‡

Center for Gravitational Wave Physics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA

(Received 22 June 2005; published 14 November 2005)

Inspiraling compact-object binary systems are promising gravitational wave sources for ground and space-based detectors. The time-dependent signature of these sources is a well-characterized function of a relatively small number of parameters; thus, the favored analysis technique makes use of matched filtering and maximum likelihood methods. As the parameters that characterize the source model vary, so do the templates against which the detector data are compared in the matched filter. For small variations in the parameters, the filter responses are closely correlated. Current analysis methodology samples a bank of filters whose parameter values are chosen so that the correlation between successive samples from successive filters in the bank is 97%. Correspondingly, the additional information available with each successive template evaluation is, in a real sense, only 3% of that already provided by the nearby templates. The reason for such a dense coverage of parameter space is to minimize the chance that a real signal, near the detection threshold, will be missed by the parameter space sampling. Here we investigate the use of Chebyshev interpolation for reducing the number of templates that must be evaluated to obtain the same analysis sensitivity. Additionally, rather than focus on the ''loss'' of signal-to-noise associated with the finite number of filters in the template bank, we evaluate the receiver operating characteristic (ROC) as a measure of the effectiveness of an analysis technique. The ROC relates the false alarm probability to the false dismissal probability of an analysis, which are the quantities that bear most directly on the effectiveness of an analysis scheme. As a demonstration, we compare the present ''dense sampling'' analysis methodology with the ''interpolation'' methodology using Chebyshev polynomials, restricted to one dimension of the multidimensional analysis problem by plotting the ROC curves. We find that the interpolated search can be arranged to have the same false alarm and false dismissal probabilities as the dense sampling strategy using 25% fewer templates. Generalized to the two-dimensional space used in the computationally limited current analyses, this suggests a factor of 2 increase in computational efficiency; generalized to the full seven-dimensional parameter space that characterizes the signal associated with an eccentric binary system of spinning neutron stars or black holes, it suggests an order of magnitude increase in computational efficiency. Since the computational cost of the analysis is driven almost exclusively by the matched filter evaluations, a reduction in the number of template evaluations translates directly into an increase in computational efficiency; additionally, since the computational cost of the analysis is large, the increased efficiency translates also into an increase in the size of the parameter space that can be analyzed and, thus, the science that can be accomplished with the data.

DOI: 10.1103/PhysRevD.72.102001 PACS numbers: 04.80.Nn, 07.05.Kf, 95.55.Ym

I. INTRODUCTION

Inspiraling compact binaries of stellar mass neutron stars or black holes are among the most important gravitational wave sources accessible to the current generation of ground-based interferometric gravitational wave detectors [1–4]. They are also very ''clean'' systems, in the sense that the gravitational wave signal arising from the inspiral depends only on general relativity (i.e., the structure of the binary components is unimportant) and can be calculated to great accuracy by the well-understood techniques of post-Newtonian perturbation theory [5–7]. For these reasons, matched filtering and maximum likelihood techniques are well-suited for the detection and characterization of the signal from these systems [8,9] and an implementation based on these methods is currently used in the analysis of data from the LIGO and GEO detectors [10]. Current implementations of matched filtering for this problem involve a dense search strategy, where the binary parameter space is sampled by templates spaced to ensure that real signals, with signal-to-noise greater than a given threshold (typically 8) have a certain minimum correlation (typically 97%) with at least one template. Here we investigate an alternative algorithm for detection, involving interpolation between templates spaced less densely, and

*Electronic address: sanjit@iucaa.ernet.in
†Electronic address: sanjeev@iucaa.ernet.in
‡Also at the Institute for Gravitational Physics and Geometry, Department of Physics, and Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA.
Electronic address: LSFinn@PSU.Edu

compare the efficiency of this interpolated search algorithm with the dense search algorithm using the receiver operating characteristic (ROC) of each search, which takes into account both the false dismissal and the false alarm fractions of each.

The gravitational wave signature of inspiraling binary systems depends on a set of 15 parameters that characterize the system (i.e., component masses, orbital energy and angular momentum at a given epoch, component spins, orientation relative to detector line of sight). To identify an incident signal using a matched filter requires the application of a fair sampling of filter ''templates,'' each defined by a unique choice of the parameters associated with the physical system. Current implementations of matched filtering used in the analysis of gravitational wave detector data involve a very dense sampling of the two-dimensional parameter subspace corresponding to the binary component masses (intrinsic parameter space) and assuming zero eccentricity orbits and no body spins:[1] the templates are spaced so closely that the correlation between templates at neighboring points in the subspace is 97% [13,14].

We refer to this as the ''dense'' search strategy. The rationale underlying the dense search strategy is to reduce the probability that a weak signal, characterized by parameters that fall between those sampled, will be missed by the sampling. Here we describe a straightforward and practical way of using interpolation to take advantage of the correlation between the matched filter output associated with nearby points in the parameter space to significantly reduce the number of matched filter evaluations without sacrificing the efficiency with which real signals are recognized.

We are not the first to observe the significance of the high correlation between neighboring templates nor to consider the opportunity for and advantages of interpolation as part of the implementation of matched filtering for the analysis of binary inspiral signals. The significance of the high correlation as an indication that fewer templates should be able to recover signals with the same efficiency was first made in [14]. Croce *et al.* [15,16] explored the use of Cardinal interpolation with a truncated series of sinc functions to estimate the value of the matched filter output when the filter used corresponds to the actual parameters that describe the signal. They found a sampling of parameter space that would ensure the interpolated estimate would be no less than 97% of the maximum over a two-dimensional intrinsic parameter space. Their sampling and interpolation reduced by a factor of 4, compared to the dense search, the number of templates required to search over a *two-dimensional* intrinsic parameter space. Here we find that we can achieve an increase in efficiency by a factor of 3.5 for *one-dimensional* parameter space, with a simpler template spacing and a simpler and quicker to evaluate interpolation function.

Cardinal interpolation with sinc functions provides perfect reconstruction of a bandlimited function from equispaced samples. In the present case, however, the function being interpreted is not bandlimited and—in any event—we do not have the infinite number of samples that would be required for perfect reconstruction. As an alternative to cardinal interpolation with a truncated series of sinc functions, we consider interpolation using Chebyshev polynomials, which have two important properties: first, they have (very nearly) the minimum maximum error of all polynomial interpolating expressions of fixed degree; and second, they have the practical advantage of being quick and easy to calculate.

To understand and demonstrate the performance of the interpolated search, we consider a one-dimensional parameter space in the chirp mass and evaluate the minimum number of templates required to obtain a given detection efficiency using both the (Chebyshev) interpolation and dense strategies. We go further, however, and calculate also the false alarm probability of both search strategies. The relationship between the false alarm and false dismissal probabilities is referred to as the receiver operating characteristic, or ROC. Clearly, given two analysis strategies with the same efficiency, the strategy with the lower false alarm fraction has superior discriminating power. We find that interpolation strategy using Chebyshev interpolation is superior to the dense analysis strategy or interpolation using the sinc function, from the perspective of either computational efficiency or discriminating ability.

The paper is organized as follows: In Sec. II we describe the motivation behind our choice of interpolating function and the difference between our choice and the choice made in [15,16]. In Sec. III we describe in detail the dense and interpolated search strategies, the two-dimensional template space used in current gravitational wave data analyses for inspiraling binary neutron stars, the one-dimensional restriction that we use here to compare the effectiveness of the interpolating search strategy, the use of ROCs for comparing different analysis strategies, and (finally) compare the performance of the interpolated and dense search strategies by evaluating the sensitivity of each at fixed computational cost and the computational cost required by each to achieve the same sensitivity.

II. INTERPOLATING IN PARAMETER SPACE

The Wiener matched filter W, corresponding to an expected signal characterized by τ, is a scalar-valued function of the (vector-valued) instrument data $\mathbf{d}$, noise power

[1]The rationale for choosing a subspace is that the computational cost of a full parameter space search is high and that many systems are believed to be adequately represented by this subspace. Even for this two-dimensional subspace the minimum computational cost for a matched filter search over component masses in the range $0.2M_\odot < m_1 \leq m_2 < 30M_\odot$ in the LIGO detector band is several hundred GFlops/s [11]. When significant body spin is allowed, the computational cost grows by several orders of magnitude [12].

spectral density (PSD) S_n:

$$W(\mathbf{d}|\tau) = W(\tau|S_n, \mathbf{d}). \tag{2.1}$$

In our particular problem $W(\mathbf{d}|\tau)$ is a continuous function of τ and τ corresponds to the parameters that characterize our binary system model: e.g., binary system component masses, orbital energy and angular momentum, component spins, etc. Given a data set $\mathbf{d}$, we wish to find an interpolating function $\tilde{W}(\tau)$ and a set of points τ_k in the space of possible signals such that

$$W_k = \tilde{W}(\tau_k) = W(\mathbf{d}|S_n, \tau_k). \tag{2.2}$$

There are, of course, an infinite number of continuous functions $\tilde{W}(\tau)$ that take on the values W_k at the τ_k: the question is, how do we choose among them?

Focus attention first on the case where τ is a scalar x. One particular choice of interpolant $\tilde{W}(\mathbf{d}|S_n, x)$, which is especially important in the context of communication theory, is based on the Whittaker Cardinal function sinc:

$$C(x) = \sum_{k=-\infty}^{\infty} W_k \mathrm{sinc}\frac{x - x_k}{\Delta}, \tag{2.3}$$

where

$$\mathrm{sinc}\,(x) = \frac{\sin \pi x}{\pi x}, \tag{2.4}$$

$$x_k = x_0 + k\Delta. \tag{2.5}$$

Shannon [17] showed that the Cardinal interpolation $C(x)$ of $W(\mathbf{d}|S_n, x)$ is the unique interpolant $\tilde{W}$ that (i) takes on the values W_k at the x_k, (ii) has no singularities, and (iii) and whose spectrum is limited to a bandwidth $(2\Delta)^{-1}$. Correspondingly, if $W(\mathbf{d}|S_n, x)$ is bandlimited in x and has the values W_k at the equidistant sampled points x_k, then $W(\mathbf{d}|S_n, x)$ is equal to $C(x)$. In the case where τ is multidimensional the interpolation can be performed separately on each index: e.g., in the case of two dimensions [i.e., τ equal to (τ_1, τ_2)]

$$C(\tau) = \sum_{j,k=-\infty}^{\infty} W_{jk} \mathrm{sinc}\frac{\pi}{\Delta_1}(\tau_1 - \tau_{1,j}) \mathrm{sinc}\frac{\pi}{\Delta_2}(\tau_2 - \tau_{2,k}), \tag{2.6}$$

where

$$\tau_{1,j} = \tau_{1,0} + j\Delta_1, \tag{2.7}$$

$$\tau_{2,k} = \tau_{2,0} + k\Delta_2 \tag{2.8}$$

and $\tau_{1,0}$, $\tau_{2,0}$ are constants.

Cardinal interpolation using the Cardinal function sinc forms the basis of the interpolation formula used in [15,16]. If $W(\mathbf{d}|S_n, \tau)$ is bandlimited and we choose our samples of W appropriately, then we can do no better than using the Cardinal function to interpolate values of W between the samples. In our problem, however, $W(\mathbf{d}|S_n, \tau)$ is not bandlimited and we do not have an infinite number of sample points W_k; correspondingly, the Cardinal function $C(\tau)$ is at best an approximation to $W(\mathbf{d}|S_n, \tau)$. With that understanding the Cardinal interpolation $C(\tau)$ is not preferred and we are led to seek other approximations to $W(\mathbf{d}|S_n, \tau)$ that have favorable properties.[2]

One possibility, chosen from approximation (as opposed to interpolation) theory, is the use of a Chebyshev polynomial expansion to approximate $W(\mathbf{d}|S_n, \tau)$. Without loss of generality consider a continuous function $f(x)$ on $[-1, 1]$. The Weierstrass Approximation Theorem states that for any $\epsilon > 0$ we can find a polynomial P_n of order n such that

$$\max_{x\in[-1,1]} |f(x) - P_n(x)| \leq \epsilon. \tag{2.9}$$

The minimax polynomial approximation to $W(\mathbf{d}|S_n, x)$ is a natural candidate for the interpolation $\tilde{W}(x)$. Unfortunately, finding the minimax polynomial is a very difficult process; nevertheless an excellent *approximation* to the minimax polynomial does exist. Define the error $E(x|f, P_n)$ associated with the polynomial approximation $P_n(x)$ by

$$E(x|f, P_n) \equiv f(x) - P_n(x). \tag{2.10}$$

The Chebyshev Equioscillation Theorem [18] states P_n^* is the minimax polynomial if and only if there exist $n + 2$ points $-1 \leq x_0 < x_1 < \cdots < x_{n+1} \leq 1$ for which

$$E(x_k|f, P_n^*) = (-1)^k E, \tag{2.11}$$

where

$$|E| \equiv \max_{x\in[-1,1]} |E(x|f, P_n)|. \tag{2.12}$$

As a corollary, $E(x|f, P_n^*)$ vanishes for $x \in [-1, 1]$ at $n + 1$ points x_k', with $x_k < x_k' < x_{k+1}$. This result, together with the Mean Value Theorem, allows us to write the error term associated with the minimax polynomial P_n^* as

$$E(x|f, P_n^*) = \frac{f^{(n+1)}(\xi)}{(n+1)!} \prod_{k=0}^{n} (x - x_k'), \tag{2.13}$$

where $\xi \in [-1, 1]$. Correspondingly,

$$|E| \leq \max_{x\in[-1,1]} \left| \prod_{k=0}^{n} (x - x_k') \right| \max_{\xi\in[-1,1]} \frac{|f^{(n+1)}(\xi)|}{(n+1)!}. \tag{2.14}$$

[2] In fact, as noted in [15,16], the Γ is *quasibandlimited*: i.e., there exists a B_c such that the error one makes by undersampling at frequency $B > B_c$ is proportional to $\exp[-(B - B_c)]$. Nevertheless, interpolation with the Cardinal function is still an approximation and, as we are about to see, other interpolating functions can achieve equivalent accuracy at smaller computational costs.

Focus attention on the order $n+1$ polynomial

$$Q^*_{n+1}(x) = \prod_{k=0}^{n}(x - x'_k). \tag{2.15}$$

This polynomial has leading coefficient unity. A unique property of the Chebyshev polynomial T_{n+1} is that, of all order $n+1$ polynomials Q_{n+1} with leading coefficient unity,

$$\max_{x\in[-1,1]} |\frac{T_n(x)}{2^{n-1}}| \leq \max_{x\in[-1,1]} |Q_{n+1}(x)|. \tag{2.16}$$

Additionally, $T_{n+1}(x)$ has exactly $(n+2)$ extrema on $[-1, 1]$, the value of $|T_{n+1}(x)|$ at these extrema is 1, and the extrema alternate in sign. Correspondingly, if the error term $E(x|f, P^*_n)$ associated with the minimax polynomial P^*_n were polynomial—i.e., $f^{(n+1)}(\xi)$ were constant in Eq. (2.13) so that $E(x|f, P^*_n)$ was equal to Q^*_n—then by the Equioscillation Theorem Q^*_{n+1} would be equal to T_{n+1} and the x'_k—where the error vanishes—would be the $n+1$ roots of T_{n+1}. This suggests that we find the order n polynomial p^*_n such that

$$p^*_n(x'_k) = f(x'_k) \quad \forall\ k = 0 \ldots n, \tag{2.17}$$

where, again, the x'_k are the roots of T_{n+1}. The polynomial p^*_n is a *near minimax* polynomial approximation to $f(x)$. For this polynomial approximation Powell [19] showed that, as long as $f(x)$ is continuous on $[-1, 1]$,

$$1 \leq \frac{\epsilon_{\rm cheb}}{\epsilon_0} \leq \nu_n \equiv 1 + \frac{1}{n+1}\sum_{k=0}^{n} \tan\left[\frac{(k+1/2)\pi}{2(n+1)}\right], \tag{2.18}$$

where

$$\epsilon_0 = \max_{x\in[-1,1]} |E(x|f, P^*_n)|, \tag{2.19}$$

$$\epsilon_{\rm cheb} = \max_{x\in[-1,1]} |E(x|f, p^*_n)|. \tag{2.20}$$

Powell also showed that ν_n grows slowly with n: in particular,

$$\nu_n \sim \frac{2}{\pi}\log n. \tag{2.21}$$

Somewhat tighter bounds on ν_n can be placed when f is also differentiable [20].

As defined above, the near minimax polynomial p^*_n is the interpolating polynomial that agrees with f at the $n+1$ roots of T_{n+1}. Alternatively, using several properties of Chebyshev polynomials, the Chebyshev interpolating polynomial can be expressed as a linear combination of Chebyshev polynomials:

$$p^*_n(x) = \sum_{k=0}^{n} a_k T_k(x) - \frac{1}{2}a_0, \tag{2.22}$$

where

$$a_j = \frac{2}{n+1}\sum_{k=1}^{n+1} f(x'_k)T_j(x'_k), \tag{2.23}$$

where, again, the x'_k are the $n+1$ roots of T_{n+1}.

III. COMPARISON: DENSE AND INTERPOLATED SEARCH

In this section we describe the dense and interpolating search strategy and compare their efficiency when applied to the problem of identifying the gravitational wave signature of coalescing neutron star systems in the LIGO detectors.

A. Two search strategies

The conventional search strategy used in the current analyses of LIGO, GEO, and TAMA data (cf. [10,11,13,14]) begins with the placement of templates at discrete points τ_k on the parameter space τ. To choose the template locations we define the inner product of two signals $g(t)$ and $h(t)$,

$$\langle g, h\rangle = 4\int_0^\infty df \Re\left[\frac{\tilde{g}(f)\tilde{h}^*(f)}{S_n(f)}\right], \tag{3.1}$$

where $\tilde{g}(f)$ is the Fourier transform of g and S_n is the detector noise power spectral density. Denoting by $h(t|\tau)$ the signal characterized by τ, the match $\Gamma(\tau_j, \tau_k)$ is

$$\Gamma(\tau_j, \tau_k) = \frac{\langle h(t|\tau_j), h(t, \tau_k)\rangle}{\sqrt{\langle h(t|\tau_j), h(t, \tau_j)\rangle\langle h(t|\tau_k), h(t, \tau_k)\rangle}}. \tag{3.2}$$

By construction $|\Gamma| \leq 1$. The template locations are chosen so that consecutive templates in any of the directions τ_j have an overlap Γ_0, referred to as the "minimum match" and typically chosen to be 97%.

With the templates placed, the dense search strategy proceeds:

(1) Evaluate the Wiener filter $W(\mathbf{d}|S_n, \tau_k)$ at each of the template locations τ_k.
(2) Determine the template τ_j whose Wiener filter output is greatest.
(3) If the filter output at τ_j exceeds the given threshold, report an event with the parameters τ_j.

We refer to this as the dense search strategy.

Following the discussion in Sec. II we are in a position to describe an alternative strategy, which we refer to as the interpolated search strategy. First, fix the order of the interpolating polynomial. This determines the template locations τ_k on the parameter space τ. Then

(1) Evaluate the Wiener filter $W(\mathbf{d}|S_n, \tau_k)$ at each of the template locations τ_k.
(2) Form the interpolating polynomial from the $W(\mathbf{d}|S_n, \tau_k)$.

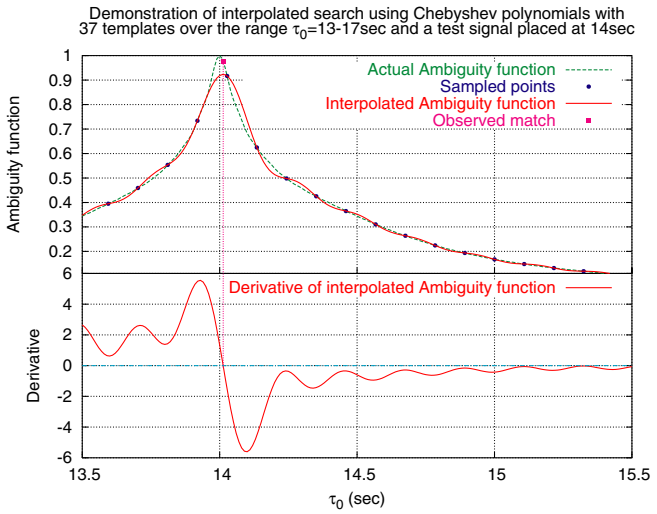

FIG. 1 (color online). This figure demonstrates the interpolated search. The ambiguity function is sampled and reconstructed over the chosen parameter space of $\tau_0 = 13$–17 sec (only a part of the parameter space has been shown in the figure) with the help of the Chebyshev interpolating polynomial. The approximate location of the peak of the interpolating function is first located and the zero of the derivative is obtained by applying successive approximations around the peak. Finally a template is placed at this estimated signal location. Note that, by placing a template at the maximum of the interpolating polynomial, the match has improved over the one obtained by simply evaluating the maximum of the interpolating polynomial.

(3) Determine the location τ' where the interpolating polynomial is maximized.
(4) Perform a final Wiener filter evaluation at τ'.
(5) If the final evaluation exceeds the given threshold, report an event with the parameters τ'.

We illustrate the interpolated search strategy using Fig. 1. In Fig. 1 we use 37 interpolating search templates, that is, we sample the ambiguity function at 37 points in the τ_0 space (the marked points on the dotted curve). We construct the interpolating function (the solid curve) and find its maximum by setting its derivative to zero. In order to avoid local extrema, we first find the approximate location of the peak of the interpolating function and then find the zero of its derivative by successive approximation near the region of the peak. One can clearly see that the proper value of the ambiguity function at the maximum of the interpolating function is more than the maximum value of the interpolating function and this is what we gain by placing a template at the maximum of the interpolating function.

B. A one-dimensional parameter space for comparative studies

We are interested in understanding the performance of the interpolated search strategy relative to the dense search strategy, which is currently used in the analysis of data from the LIGO, GEO, and TAMA detectors [10]. The current analyses focus on templates corresponding to binaries with circular orbits and no component spins. The corresponding two-dimensional parameter space is spanned by the masses of the individual components. The templates vary most rapidly, however, along the axis spanned by the so-called *chirp mass*

$$\mathcal{M} := \mu^{3/5} M^{2/5}, \tag{3.3}$$

where M is the system's total mass and μ its reduced mass. The linear density of templates needed by the dense search in the direction $\partial_{\mathcal{M}}$ is approximately 100 times the linear density needed in the orthogonal direction. For the comparison we perform here, we focus attention on the number of template evaluations needed for binaries with equal mass components that vary only in $\mathcal{M}$. We expect that the ratio of performance, measured as the number of templates required by the two search strategies to achieve the same search results, will be the same in the complementary dimension and in the other dimensions that will be introduced in future searches that accommodate component spins and orbital eccentricity.

C. Templates

The strain response of an interferometric gravitational wave detector to quadrupole formula approximation gravitational waves incident from an inspiraling binary neutron star system can be written

$$h(t|t_a, \tau_0) = h_0[\pi f(t - t_a - \tau_0)\mathcal{M}]^{2/3} \cos\Phi(t - t_a - \tau_0), \tag{3.4a}$$

where

$$f(t|t_a, \tau_0) := \frac{1}{\pi\mathcal{M}}\left(\frac{5}{256}\frac{M}{\tau_0 + t_a - t}\right)^{3/8}, \tag{3.4b}$$

$$\Phi(t|t_a, \tau_0) := \Phi_a + 2\pi\int_t^{t_a+\tau_0} dt f(t|t_a, \tau_0), \tag{3.4c}$$

for $t < t_a + \tau_0$. Here t_a is the moment when the instantaneous wave frequency is equal to f_a and τ_0 is the elapsed time from that moment until (in this approximation) the system coalesces, which is directly related to the system's chirp mass $\mathcal{M}$:

$$\tau_0 = \frac{5}{256\pi f_a}\frac{1}{(\pi\mathcal{M}f_a)^{5/3}}. \tag{3.5}$$

The elapsed time to coalescence τ_0 is a useful surrogate for the chirp mass $\mathcal{M}$: templates equispaced in τ_0 have constant cross correlation, independent of τ_0. Choosing f_a equal to 40 Hz, which is commonly taken as the lower edge of the LIGO detector bandwidth at design sensitivity [21], τ_0 ranges from approximately 43 sec for a binary system consisting of two $1M_\odot$ compact objects to 0.15 sec for a binary consisting of two $30M_\odot$ black holes.

It is convenient to work with the Fourier transform of the strain response of the detector. For neutron star binaries in the LIGO or Virgo band, the Fourier transform can be evaluated to an excellent approximation using the stationary phase approximation [9]:

$$\tilde{h}(f) = \mathcal{N} f^{-7/6} \exp\{i[-\Phi_a - \pi/4 + \Psi(f|t_a, \tau_0)]\}, \tag{3.6a}$$

where

$$\Phi_a = \Phi(t_a|t_a, \tau_0), \tag{3.6b}$$

$$\Psi(f|t_a, \tau_0) = 2\pi f t_a + f_a \tau_0 \frac{6\pi}{5}\left(\frac{f}{f_a}\right)^{-5/3}. \tag{3.6c}$$

The factor $\mathcal{N}$ is a constant amplitude.

It is important to distinguish between the nature of the parameters that characterize the template. Changes in the parameter τ_0 change the waveform *shape*: we term such parameters *dynamical* parameters. On the other hand, parameters such as t_a or Φ_a translate the waveform, but do not alter its shape: we term these *kinematical* parameters.[3] In our problem only the subspace of dynamical parameters needs to be spanned by discrete templates: the values of the kinematical parameters for the Wiener filter with the maximum output can be determined by other means. Correspondingly, at the level of approximation associated with the quadrupole formula the family of templates that must be evaluated is one dimensional.

D. Dense search template placement

There are many different ways of parametrizing the template space. Choosing τ_0 as a dynamical variable has the advantage that $\Gamma(\tau_0, \tau_0')$ depends only on the difference $\tau_0 - \tau_0'$; consequently, in the dense search templates are spaced uniformly in τ_0 [11,13,14]. To determine that spacing we evaluate

$$\mathcal{H}(\Delta\tau_0) = \Gamma(\tau_0, \tau_0 + \Delta\tau_0), \tag{3.7}$$

where now Γ has been maximized over the kinematical parameters t_a and ϕ_a. This maximization can be performed in a computationally efficient manner as shown in the literature [13]. We call $\mathcal{H}$ the dynamical ambiguity function or simply the ambiguity function. It quantifies the fractional match between the template at τ_0 and the signal at $\tau_0 + \Delta\tau_0$. Figure 2 shows $\mathcal{H}$ for power spectral density specified in the initial LIGO science requirements [21]. The requirement that $\mathcal{H}(\Delta\tau_0)$ is equal to a constant for any two consecutive templates determines the spacing $\Delta\tau_0$ between templates that differ only in τ_0. For our example problem, which has just one dynamical parameter, the requirement that $\mathcal{H}(\Delta\tau_0)$ is 97% (the conventional choice) for neighboring templates leads to a template spacing $\Delta\tau_0$ equal to 30 ms.

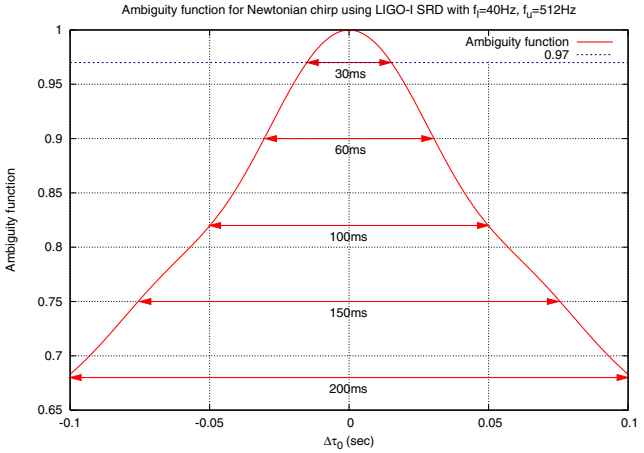

FIG. 2 (color online). Ambiguity function $\mathcal{H}$ plotted as a function of $\Delta\tau_0$. Horizontal lines are drawn for various matches; the dashed horizontal line is for a 97% match, which corresponds to an intertemplate separation of ~30 ms.

E. Interpolated search template placement

In the dense search templates are equispaced in τ_0, with the spacing between adjacent templates—and thus the number of templates—chosen such that the dynamical ambiguity function takes on a specified value. When presented with data an event is signaled when the amplitude at one of these templates exceeds a threshold.

In the interpolated search, on the other hand, the domain $[\tau_0^{\min}, \tau_0^{\max}]$ is mapped onto $[-1, 1]$ and the placement and number of templates is chosen to simplify the construction of the Chebyshev interpolating polynomial of the template output over this domain. When presented with data, the maximum value of the Chebyshev interpolating polynomial is found and an event is signaled when the amplitude at that location exceeds a threshold.

In the interpolated search our goal is to minimize the order of the interpolating polynomial (and, thus, the number of template evaluations) required for a given accuracy of interpolation. We have some control over this through the choice of mapping from $[\tau_0^{\min}, \tau_0^{\max}]$ to $[-1, 1]$. The linear map

$$\tau' = 2\frac{\tau_0 - \tau_0^{\min}}{\tau_0^{\max} - \tau_0^{\min}} - 1 \tag{3.8}$$

is the most obvious of such mapping. While we have not made an exhaustive search of all possible mappings, however, we have observed that better fits are possible with a lower-order polynomial when we use the mappings

$$\delta = \cos(\tau') = -\cos\left[\pi\frac{\tau_0 - \tau_0^{\min}}{\tau_0^{\max} - \tau_0^{\min}}\right]. \tag{3.9}$$

[3]In the literature, dynamical and kinematical parameters are also known as intrinsic and extrinsic parameters, respectively. Unlike the dynamical parameters, the kinematical parameters can be handled quickly and easily in the filtering algorithms.

Moreover, with this mapping, the roots of the Chebyshev polynomial are equispaced over the parameter range in τ_0. Note that the Chebyshev polynomials can also be written as $T_n(x) = \cos[n\theta(x)]$, where $x = \cos\theta$, so in this case $T_n(\delta) = \cos(n\pi\tau')$. Since the τ' space is sampled uniformly due to the above mapping, Chebyshev interpolation is now equivalent to a type II discrete cosine transform in $\pi\tau'$. Equispaced sampling of the parameter space could be beneficial if the signals are distributed uniformly over the τ_0 space. Nevertheless, for a nonuniform distribution of signals, different transformation could be preferred to give optimum results. Once we have fixed the order n of the interpolating polynomial, templates are placed at values of δ that are roots of the $T_{n+1}(\delta)$. This fixes the templates. The coefficients of the interpolating polynomial are found using Eq. (2.23) and then the interpolating polynomial is constructed using Eq. (2.22).

In Fig. 3 we have plotted the match by placing normalized test signals (without noise) at regular intervals of τ_0. For each injected signal τ_0, we plot the maximum of the interpolating polynomial (dashed curve) and the match obtained by placing a template at the maximum of the interpolating polynomial (solid curve). We see that the match is a (nearly) periodic function of τ_0, with the period equal to the template separation. This suggests that the detection probability is also periodic and this fact has been used in carrying out the simulations—the signals are injected within one such "period" in the parameter space. Moreover, one can see that with just 37 + 1 interpolated search templates one gets a minimal match of 0.97, whereas the dense search requires about 133 templates to achieve the same level of minimal match. This amounts to a factor of 3.5 over the dense search and this is so for just one dimension. Note that the gain factor obtained by Croce *et al.* in [15] is ~4 in *two* dimensions, which scales to ~2 per dimension. Because the metric (Fisher information matrix) determines the lattice spacing as well as the correlation, they are interdependent. This suggests that the Chebyshev interpolation method can be extended to a parameter space of higher dimensions with about a similar gain factor per dimension.

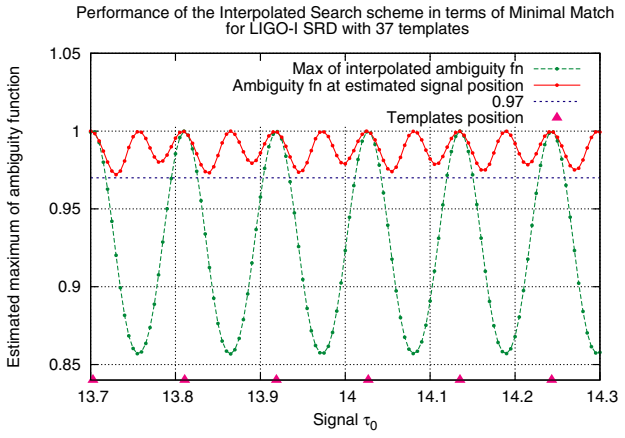

FIG. 3 (color online). Normalized test signals (without noise) were injected densely at regular intervals along the τ_0 parameter space. For each injected signal τ_0, we plot the maximum of the interpolating polynomial (dashed curve) and the match obtained by placing a template at the maximum of the interpolating polynomial (solid curve) according to the interpolated search strategy. This figure illustrates that the match is a (nearly) periodic function of τ_0 with the period equal to the template separation. Moreover, with just 37 + 1 interpolated search templates the minimal match is 0.97. To maintain the same minimal match 133 templates are needed for the usual dense search.

F. Comparison

We are interested in two, related, comparisons: first, the relative "sensitivity" of a search carried out with a fixed number of template evaluations using the dense search strategy and the interpolated search strategy; second, the number of template evaluations required by the interpolated search in order to achieve the same sensitivity as the dense search. To give meaning to the sensitivity of these two strategies we use the ROC.

The ROC is a plot of true positives as a function of the fraction of false positives for a binary classifier system as its discrimination threshold is varied. Both the dense search and the interpolated search are binary classifiers: i.e., they classify an interval of data **d** as including a signal or not including a signal. A true positive is a classification of **d** as including a signal when in fact it does; a false positive is a classification of **d** as including a signal when it does not. In both of the search strategies described here the discrimination threshold is matched filter output that must be exceeded for a data interval to be classified as including a signal. The false positive fraction is also known as the type II, or false alarm, error fraction and is denoted α. The fraction of true positives is also known as the detection efficiency ϵ, which is one minus the type I, or false positive, error fraction (which is denoted β). At fixed α a more sensitive search method has a greater ϵ. The ROC associated with a search method no better than a toss of a (possibly loaded) coin is given by the diagonal $\alpha = \epsilon$.

Using numerical simulations we have evaluated α and ϵ as a function of the detection threshold for both the interpolated search and the dense search, for different numbers of templates (dense search) and different interpolating polynomial order (interpolated search).

To evaluate the false positive fraction α we generate a large number of data segments, each 2^{15} samples long, and each consisting of Gaussian noise whose power spectrum (assuming a 1024 Hz sample rate) is that specified as the initial LIGO science requirement [21]. (The Gaussian random numbers are themselves generated using the Mersenne Twister Pseudo Random Number Generator [22] and then filtered in the Fourier domain by scaling the Fourier components by the square root of the PSD.) For

the purpose of this comparison, we look for signals in the interval $\tau_0 \in [13\text{s}, 17\text{s}]$. Both the dense and interpolated search methods are applied to this data. The ratio of the number of events signaled to the number of data segments examined as a function of the threshold η is α for that threshold. Approximately 50 000 realizations of detector noise are used to evaluate α, which gives reliable results for α greater than approximately 10^{-3}.

To compute ϵ, the true positive fraction, we proceed in a similar fashion. Now, however, with each noise instantiation we add a signal, with τ_0 drawn uniformly and randomly from the interval covered by the search: i.e., $\tau_0 \in [13\text{s}, 17\text{s}]$. We inject signals of SNR 8. In almost all cases 50 000 realizations of detector noise plus signal are used to evaluate the efficiency, which gives reliable results for efficiencies greater than approximately 10^{-3}. However, for the flat search with 40 templates and the interpolated search with 30 templates, we have used 400 000 realizations. The larger number of realizations in these cases results in smoother curves.

The top panel of Fig. 4 shows the variation of α for both methods using 40 templates: i.e., a 100 ms template spacing for the dense search and an order 39 interpolating polynomial in δ [cf. Eq. (3.9)]. For any threshold α is always greater for the dense search than for the interpolated search; similarly, as shown in the center panel of Fig. 4, for any given threshold the efficiency ϵ is always greater for the interpolated search than for the dense search. Finally, the bottom panel of Fig. 4 shows the

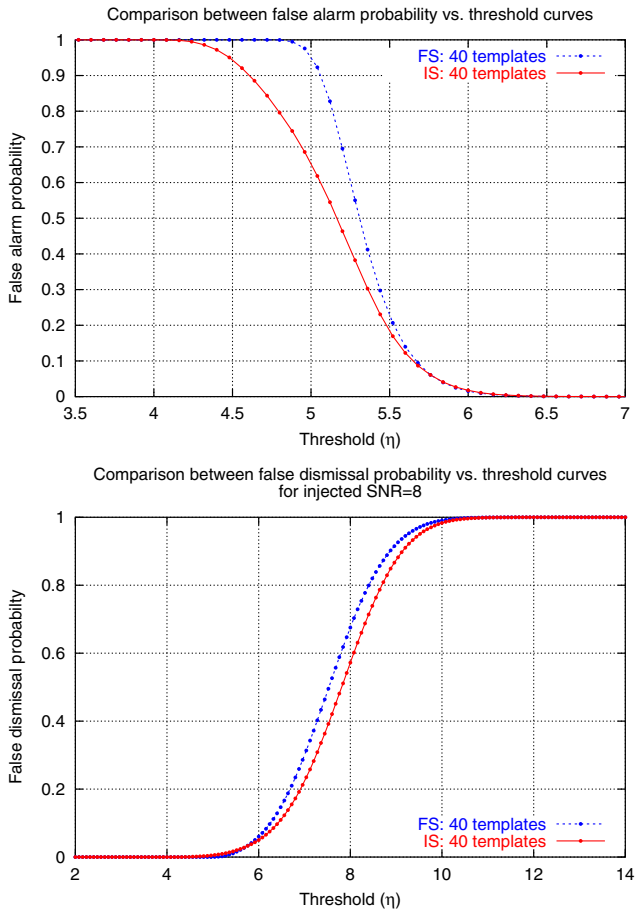

FIG. 4 (color online). The variation of the false and true positive fractions, α and ϵ with threshold η for the dense and interpolated search methods, each making use of 40 template evaluations. The top panel shows the false positive fraction. Note how the false positive falls much sooner for the interpolated search than for the dense search. The bottom panel shows ϵ when a signal of amplitude signal-to-noise 8 is present in the range $\tau_0 \in [13\text{s}, 17\text{s}]$. Note how the ϵ is always greater for the interpolated search than for the dense search. For the same computational cost (determined by the number of template evaluations), the interpolated search will always perform better than the dense search.

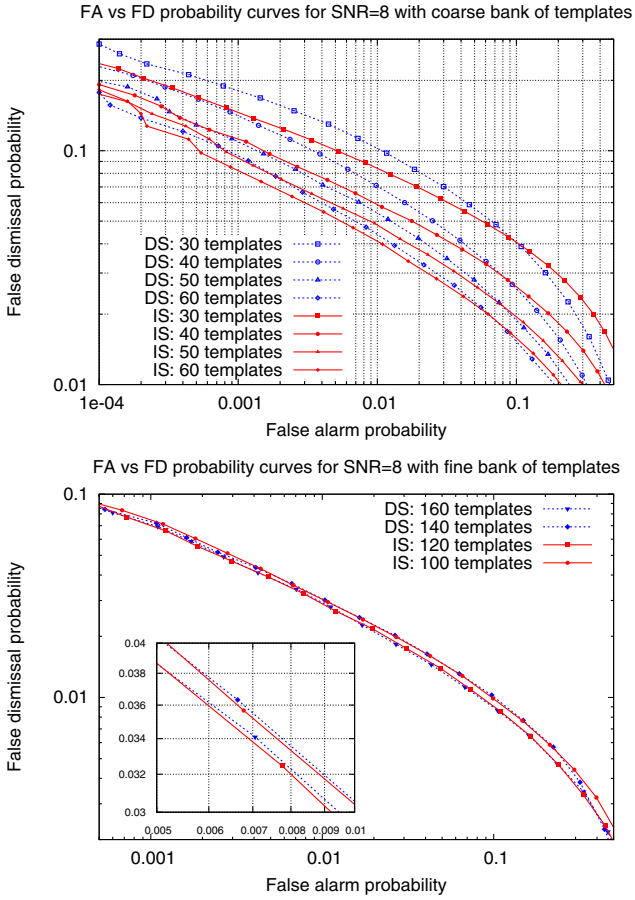

FIG. 5 (color online). ROC curves for dense searches (solid curves) and interpolating searches (dashed curves). For a given number of templates, the solid curves are "lower"—less false dismissal probability for the same false alarm—than the dotted curves in the regime of low false alarm showing that the interpolated search performs better than the dense search for low false alarm. The bottom panel shows an analogous plot for high minimal match (fine bank) ~0.98. Here the performance of the dense search with 160 and 140 templates is comparable to that of the interpolated search with 120 and 100 templates, respectively.

TABLE I. Number of template evaluations required to obtain the same efficiency at a false alarm fraction of 10^{-3} in a dense search and an interpolated search. Note how the interpolated search is computationally more efficient for the same sensitivity.

No. of templates		
Dense	Interpolated	ϵ at $\alpha = 10^{-3}$
40	31	0.859
50	41	0.890
60	49	0.905
80	64	0.919
100	89	0.924
140	105	0.927
160	115	0.929

ROC for a 40 template dense search and an order 39 interpolated search, both of which involve 40 template evaluations to decide if a signal has been detected. Comparing both ROCs it is clear that the interpolated search is more sensitive at any given α than the dense search. This is always true: i.e., for a fixed number of template evaluations the interpolated search will always have a better efficiency at a given α than the dense search, though as the number of templates grows large the fractional difference in sensitivity will decrease.

Figure 5 and Table I address the second of our two questions: the number of template evaluations required of an interpolated search to have the same sensitivity as a dense search. Figure 5 shows the ROCs for dense searches using 140 and 160 templates, together with the ROCs for interpolated searches using 120 and 100 templates. The interpolated search with an order 120 interpolating polynomial is clearly as sensitive as a dense search with 160 templates, and an interpolating search with an order 100 polynomial is as sensitive as a dense search with 140 templates. Table I shows similar pairings of the number of templates in a dense search and the number of templates in an interpolating search necessary to achieve the same sensitivity.

IV. CONCLUSION

We have shown that the use of near minimax interpolating polynomials to fit the output of matched filters to the filter parameter values can greatly improve the sensitivity of a matched filter based search for gravitational waves from compact binary inspiral. Since the lattice for dense search and the correlations are dependent on the metric (Fisher information matrix) and *any interpolation* exploits these correlations, we believe that the Chebyshev interpolation method can be extended to a parameter space of higher dimensions with about a similar gain factor per dimension. Using such a polynomial to find the parameters of the signal template leading to the best match, we can reduce the computational cost of a search over a two-dimensional parameter space by a factor of 2 compared to the methods currently in use, without any loss of sensitivity or discriminating power. This factor of 2 becomes a factor of 10 when the search is over the seven-dimensional parameter space that includes not only the masses but also the spins of the binary components [23]. This savings in computational cost is estimated under the assumption, which we believe well-founded, that we will obtain the same savings when the interpolation is extended to additional dimensions.

Other suggestions have been made for reducing the number of filter evaluations without sacrificing detection efficiency. One promising proposal involves a hierarchical search strategy, wherein a low-threshold trigger generated by the evaluation of the matched filters associated with a much coarser sampling of parameter space followed by (if necessary) a higher threshold evaluation matched filters over a much finer sampling of parameter space [24–28]. The interpolation strategy we describe here can be implemented together with the hierarchical strategies that have already been proposed to further improve the computational efficiency of binary inspiral analysis. While the gain in efficiency of the interpolated search over the dense search is approximately constant in the desired false alarm probability, the balance between the coarseness of the grids in the hierarchical steps, the number of hierarchical steps, and the gain in computational efficiency associated with the interpolation is not obvious and requires further study. Nevertheless, since the major contribution to the computational cost of a multigrid search is thought to arise in the initial stage of the search, the gain in computational efficiency—and, correspondingly, the size of the parameter space that can be studied with fixed computational resources—could be substantial.

ACKNOWLEDGMENTS

It is a pleasure to thank Innocenzo Pinto for his extremely useful comments. We are grateful to Albert Lazzarini and Anand S. Sengupta for valuable discussions. We are glad to acknowledge the use of Pleiades cluster facility at Penn State, which is supported by Penn State University, the Center for Gravitational Wave Physics, the International Virtual Data Grid Laboratory, and the National Science Foundation (PHY-00-99559). We are also glad to thank IUCAA for the use of its high performance computing facility. S. M. acknowledges the support of CSIR. S. V. D. acknowledges the support of the DST. L. S. F. acknowledges the support of the Center for Gravitational Wave Physics and the NSF under Awards No. PHY-00-99559 and No. INT-01-38459. The International Virtual Data Grid Laboratory is supported by the NSF under Cooperative Agreement No. PHY-01-22557; the Center for Gravitational Wave Physics is supported by the NSF under Cooperative Agreement No. PHY-01-14375.

[1] G. H. Sanders, in *Gravitational-Wave Detection*, edited by Mike Cruise and Peter Saulson, Proceedings of the SPIE Vol. 4856 (SPIE-International Society for Optical Engineering, Bellingham, WA, 2003), pp. 247–257.
[2] B. Abbott *et al.*, Phys. Rev. D **69**, 122004 (2004).
[3] M. Ando and K. Tsubono, in *Gravitational Waves*, edited by S. Meshkov, Proceedings of the Third Edoardo Amaldi Conference, AIP Conf. Proc. No. 523 (American Institute of Physics, Melville, New York, 2000).
[4] F. Acernese *et al.*, Classical Quantum Gravity **20**, S609 (2003).
[5] L. Blanchet, T. Damour, and B. R. Iyer, Phys. Rev. D **51**, 5360 (1995).
[6] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **63**, 044023 (2001).
[7] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **66**, 027502 (2002).
[8] L. S. Finn, Phys. Rev. D **46**, 5236 (1992).
[9] L. S. Finn and D. F. Chernoff, Phys. Rev. D **47**, 2198 (1993).
[10] B. Abbott *et al.*, Phys. Rev. D **69**, 122001 (2004).
[11] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[12] A. Buonanno, Y. Chen, and M. Vallisneri, Phys. Rev. D **67**, 104025 (2003).
[13] B. S. Sathyaprakash and S. V. Dhurandhar, Phys. Rev. D **44**, 3819 (1991).
[14] S. V. Dhurandhar and B. F. Schutz, Phys. Rev. D **50**, 2390 (1994).
[15] R. P. Croce, T. Demma, V. Pierro, I. M. Pinto, D. Churches, and B. S. Sathyaprakash, Phys. Rev. D **62**, 121101 (2000).
[16] R. P. Croce, T. Demma, V. Pierro, I. M. Pinto, and F. Postiglione, Phys. Rev. D **62**, 124020 (2000).
[17] C. E. Shannon, Proc. Inst. Radio Eng. **37**, 10 (1949).
[18] J. C. Mason and D. C. Handscomb, *Chebyshev Polynomials* (Chapman and Hall/CRC, Boca Raton, 2003).
[19] M. J. D. Powell, Computer J. **9**, 404 (1967).
[20] R.-C. Li, IEEE Trans. Comput. **53**, 678 (2004).
[21] A. Lazzarini and R. Weiss, Tech. Rep. LIGO-E950018-02-E, Laser Interferometer Gravitational Wave Observatory (1995), internal working note.
[22] M. Matsumoto and T. Nishimura, ACM Trans. Modeling Comput. Simul. **8**, 3 (1998).
[23] A. Buonanno, Y. Chen, Y. Pan, and M. Vallisneri, Phys. Rev. D **70**, 104003 (2004).
[24] S. D. Mohanty and S. V. Dhurandhar, Phys. Rev. D **54**, 7108 (1996).
[25] S. D. Mohanty, Phys. Rev. D **57**, 630 (1998).
[26] R. P. Croce, T. Demma, M. Longo, S. Marano, V. Matta, V. Pierro, and I. M. Pinto, Phys. Rev. D **70**, 122001 (2004).
[27] R. P. Croce, T. Demma, M. Longo, S. Marano, V. Matta, V. Pierro, and I. M. Pinto, Classical Quantum Gravity **21**, 4955 (2004).
[28] A. S. Sengupta, S. Dhurandhar, and A. Lazzarini, Phys. Rev. D **67**, 082004 (2003).