

Nonlinear detection algorithm for periodic signals in gravitational wave detectors

Julien Sylvestre*

LIGO Project, Massachusetts Institute of Technology, NW17-161, 175 Albany St., Cambridge, Massachusetts 02139

(Received 20 December 2000; published 23 March 2001)

We present an algorithm for the detection of periodic sources of gravitational waves with interferometric detectors that is based on a special symmetry of the problem: the contributions to the phase modulation of the signal from Earth's rotation are exactly equal and opposite at any two instants of time separated by half a sidereal day; the corresponding is true for the contributions from Earth's orbital motion for half a sidereal year, assuming a circular orbit. The addition of phases through multiplications of the shifted time series gives a demodulated signal; specific attention is given to the reduction of noise mixing resulting from these multiplications. We discuss the statistics of this algorithm for all-sky searches (which include a parametrization of the source spin-down), in particular its optimal sensitivity as a function of required computational power. Two specific examples of all-sky searches (broadband and narrowband) are explored numerically, and their performances are compared with the stack-slide technique [P. R. Brady and T. Creighton, *Phys. Rev. D* **61**, 082001 (2000)].

DOI: 10.1103/PhysRevD.63.082004

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

INTRODUCTION

Kilometer-scale gravitational wave interferometric detectors will become operational by the end of 2002, allowing observations with unprecedented sensitivity at frequencies ranging from a few tens of hertz to approximately one kilohertz [1]. Periodic sources form an interesting subclass of potentially detectable astrophysical objects, due to the possibility to improve the “visibility” of a weak periodic signal in noisy data by increasing the observation time. Good candidates for such sources are spinning galactic neutron stars with a deformation misaligned with their rotation axis; optimistic estimates suggest that the Laser Interferometric Gravitational Wave Observatory (LIGO) could detect such objects with an observation time of the order of one year [2].

Frequency modulations due to the Doppler shift induced by the detector motions, and intrinsic frequency variations due to the loss of angular momentum to gravitational waves, are two easily parametrizable examples of properties of a realistic signal that will dramatically increase the difficulty of the data analysis. In fact, for all but the simplest types of searches, the data analysis will be suboptimal because of computational limitations. More complications are likely to arise if the source is in a binary, due to the additional frequency modulation [3] and possibly larger or random intrinsic frequency variations expected if the source is accreting.

These considerations have driven serious efforts to define data analysis schemes that would maximize the sensitivity of a search (to be defined in Sec. I B), given a maximal available computational power. Matched filtering analysis, together with an in-depth discussion of signal properties, is presented in a series of papers [4]; for all-sky searches, this technique would be computationally prohibitive. The so-called coherent search, discussed in [5], uses a resampling technique to demodulate the signal and to correct for intrinsic frequency variations: the detection is performed by taking

the Fourier transform with respect to the resampling time, which amounts to a change of frame of reference from the detector to the solar system barycenter and to a “stretching” of time for frequency variations. This technique is generally not practical for all-sky searches over long time periods (longer than a few days), but a refinement of it called the stack-slide technique [6] is considered as one of the best of all known techniques: the initial data are segmented into shorter intervals, and their power spectra are built using the coherent method. For these shorter intervals, a smaller number of points in parameter space have to be used. The spectra are added (incoherently) after having been shifted according to a finer gridding of parameter space. A related technique uses Hough transforms [7] to track the frequency evolution of peaks in the spectra from the shorter intervals, which is known in advance for some choice of the source parameters. In general, it is possible to improve the sensitivity of these techniques by integrating them into more general hierarchical searches [6,7]: a first search in parameter space is done at reduced sensitivity but low confidence level (so that spurious events are likely); the vicinity of candidate events is then scrutinized at full sensitivity to get higher confidence detections.

We have explored yet another approach that exploits a particular symmetry of the problem in an attempt to reduce its complexity. First, we note that the Doppler shift in the frequency of a given source due to the rotation of Earth, independently of the source position, is exactly equal and opposite at any two instants of time separated by half a sidereal day (we will only consider linear terms in the Doppler shift, an excellent approximation). Second, in the approximation that Earth's orbital motion around the Sun is circular, the Doppler shift in frequency has the same property as above, but for a time separation of half a sidereal year. We shall use this approximation for illustrative purposes; corrections for the small eccentricity of Earth's orbit and for the influence of the moon and of other planets would have to be included in practice. Signal phases at different times can be added by properly multiplying the (analytic extensions of

*Email address: julien@ligo.mit.edu

the) signal time series, and therefore a signal free of frequency modulations may in principle be built. We will give a more precise formulation of this idea in Sec. I.

The multiplications of different stretches of the data will strongly affect the noise characteristics, principally by mixing noise at different frequencies; a useful implementation will have to bandpass filter the data to minimize this effect, and therefore a search over source parameters (frequency, frequency derivatives, source position) will be required. An example of a possible implementation will be presented in Sec. I A. The achievable sensitivity (to be defined in Sec. I B) will be studied numerically as a function of computational power requirements in Sec. II, and comparisons with the stack-slide method mentioned above will be presented.

I. ALGORITHM

We will use the following model for the signal $s(t)$ at the detector:

$$s(t) = h \cos[\phi(t)], \quad (1)$$

$$\phi(t) = 2\pi f_0[t + M_1(t) + M_2(t)], \quad (2)$$

$$M_1(t) = \frac{A_1 d}{2\pi} \cos\left(\frac{2\pi t}{d} + \alpha\right), \quad (3)$$

$$M_2(t) = \frac{A_2 \text{yr}}{2\pi} \cos\left(\frac{2\pi t}{\text{yr}} + \psi\right), \quad (4)$$

$$A_1 = \frac{2\pi R_\oplus}{d \cdot c} \cos(\lambda) \cos(\delta), \quad (5)$$

$$A_2 = \frac{2\pi A.U.}{\text{yr} \cdot c} \cos(\delta + \epsilon), \quad (6)$$

where ϵ is Earth's obliquity, λ is the detector latitude, c is the speed of light, R_\oplus is Earth's radius, A.U. is an astronomical

unit, d is a sidereal day length, yr is a sidereal year length, (α, δ) is the source angular position, f_0 is the signal frequency in the source rest frame (it may be a function of time), and ψ is a fixed number that defines the origin of time. For simplicity, since our focus is on the phase modulation of the signal, the explicit time dependence of h on the detector response pattern will not be considered, although its averaged effect will be included in our definition of sensitivity (Sec. I B). If f_0 is independent of time, a good approximation to the signal half-bandwidth Δf over a time interval of length t (i.e., half the intrinsic frequency change over t) is given by

$$\Delta f(t) = f_0[A_1 \min(1, t/d) + A_2 \min(1, t/\text{yr})]. \quad (7)$$

The numerical values of the terms multiplying the cosine terms in Eq. (5) and Eq. (6) are approximately 1.55×10^{-6} and 9.94×10^{-5} , respectively. Hence, the half-bandwidth of a 1 kHz source would be roughly 2 mHz over one day, and 0.1 Hz over a year.

The phase function $\phi(t)$ is sufficiently well-behaved that we can use the Hilbert transform of $s(t)$ to build its quadrature [8], $\hat{s}(t) = h \sin[\phi(t)]$, and therefore construct the analytic signal $S(t)$:

$$S(t) = h e^{i\phi(t)}. \quad (8)$$

Assuming that f_0 is independent of time, we have that

$$S_4(t) = S(t)S(t+d/2)S(t+\text{yr}/2)S(t+d/2+\text{yr}/2) \\ = h^4 e^{8\pi i f_0(t+d/4+\text{yr}/4)}, \quad (9)$$

that is, S_4 is monochromatic with frequency $4f_0$. If f_0 is a function of time, however, S_4 as defined above will not be demodulated because of the coupling of f_0 with the modulations; its phase will be

$$\frac{\phi_4(t)}{2\pi} = f_0(t)t + f_0(t_1)t_1 + f_0(t_2)t_2 + f_0(t_3)t_3 + M_1(t)[f_0(t) - f_0(t_1)] + M_1(t_2)[f_0(t_2) - f_0(t_3)] + M_2(t)[f_0(t) - f_0(t_2)] \\ + M_2(t_1)[f_0(t_1) - f_0(t_3)], \quad (10)$$

where $t_1 = t + d/2$, $t_2 = t + \text{yr}/2$, and $t_3 = t + d/2 + \text{yr}/2$. The signal so obtained still presents some residual phase modulation from the detector motions, but it results in a much smaller contribution Δf_{mod} to the signal bandwidth; the terms multiplying the modulation functions M_1 and M_2 are now of order of the intrinsic frequency *change* over half a day and half a year, respectively, rather than the source frequency itself. We will show below that our algorithm best performances are obtained by dividing the total data from an observation of length t_{obs} into smaller segments of length T shorter than one day. Consequently, we impose the condition

that the signal bandwidth produced by the residual modulations in Eq. (10) be smaller than the frequency resolution for a segment of length T , $\Delta f_{\text{mod}} < 1/T$, assuming a linear frequency model: $f_0(t) = f_0 \cdot [1 - t/\tau]$, where τ is the source spin-down characteristic time. We find that this condition is

$$\tau > f_0 T^2 (A_1 + A_2), \quad (11)$$

i.e.,

$$\tau > \left(\frac{f_0}{1 \text{ kHz}} \right) \left(\frac{T}{1 \text{ d}} \right)^2 \cdot 25 \text{ yr.} \quad (12)$$

This estimate shows clearly that for most realistic sources, the demodulation by multiplication is performed correctly, so that we can assume that spin down enters the phase function only through the combination $f_0(t)t + f_0(t_1)t_1 + f_0(t_2)t_2 + f_0(t_3)t_3$.

Noise will be seriously affected by the multiplications of data segments at different times, through nonlinear mixing of signal with noise and of noise with noise at different frequencies. The latter of these effects is so severe that we will have to minimize it as much as possible by bandpass filtering individually the data segments around the signal, or the technique will be of no utility. An efficient filtering requires a knowledge of the frequency content of the signal as a function of time, and will therefore involve a search over source frequencies, sky positions and spin-down parameters; this will greatly increase the computational costs. This search will be implemented by meshing this parameter space with a certain resolution, i.e. with a certain number of points N_p . Obviously, if a given signal has parameters that are not exactly coincident with any of these points, the sensitivity of our search to this signal will be reduced. We will explore below how the density of the mesh couples to the achievable mean sensitivity, and therefore give an estimate of the usefulness of our technique.

A. Implementation

From the general guidelines established in the previous section, we are able to discuss an implementation of our algorithm. This implementation has parameters that allow an optimization of the method, subject to computational power constraints.

The first step will be to prepare for the analysis data obtained from an observation of length t_{obs} . The dataset is compressed to contain only information for frequencies of interest, i.e., frequencies between f_{min} and f_{max} . The analytic signal is then constructed from this reduced dataset, using Hilbert transforms to build the quadrature, as described in Sec. I. The resulting data are divided into equal segments of length T . The computational cost of these operations is negligible, especially because they are only performed once at the beginning of the search.

Next, a bank of bandpass filtered segments is built. We choose a filter half-bandwidth $B > \Delta f$, where Δf is the maximum of the signal half-bandwidth over time T , and for every segment, we construct N_f smaller filtered segments by using a heterodyne technique [10]. These filtered segments are arranged to cover completely the frequency interval from f_{min} to f_{max} . If we impose the requirement that no signal power should be lost due to the filtering, the number N_f is a function of $B - \Delta f$: the whole frequency interval is divided into subintervals of length $2B$, displaced from each other by $2(B - \Delta f)$, so that any signal with half-bandwidth Δf is guaranteed to be fully contained within one subinterval, and $N_f = (f_{\text{max}} - f_{\text{min}}) / (2(B - \Delta f))$. In the heterodyne technique, the signal is multiplied by a complex exponential

$\exp 2\pi i(B - f_0)t$, low-pass filtered (for reasons of antialiasing), and resampled at frequency $8B$. The resulting segments have $8BT$ (complex) data points, and their Nyquist frequency is $4B$. The cost of the filtering, if implemented by an infinite impulse response digital filter, is of order $2(2n - 1)$ floating point operation per input point, where n is the number of poles in the filter, which is a function of the filter performance; we shall denote this cost by C_f . Resampling will have negligible computational cost compared to the multiplication by the complex exponential; assuming that the latter is not computed at run time, the cost for the construction of the filtered segments bank is

$$(C_f + 6)t_{\text{obs}} \frac{(f_{\text{max}} - f_{\text{min}})^2}{2(B - \Delta f)} \text{ flop.} \quad (13)$$

This cost can be adjusted by varying B and T , which enters the formula through the dependence of Δf on it.

Once the filtered segments bank is constructed, for each of the N_p points in parameter space, and for every segment at time $t < t_{\text{obs}} - T - d/2 - \text{yr}/2$, we select the appropriate filtered segments from the bank, i.e. filtered segments where the signal is expected for times t , $t_1 = t + d/2$, $t_2 = t + \text{yr}/2$, and $t_3 = t + d/2 + \text{yr}/2$, and we multiply them together. At this point, the modulations induced from the detector motion have been removed; what remain to be applied are corrections for the spin-down of the source. A refined gridding of the spin-down portion of parameter space, for the already quite narrow region used for the choice of filtered segments from the bank, is used to resample the product of the four segments according to the new time

$$t' = \frac{tf_0(t) + t_1f_0(t_1) + t_2f_0(t_2) + t_3f_0(t_3)}{f_0(0)}, \quad (14)$$

where the origin of t was arbitrarily set to zero. The resampling will only affect a few points in the time series, so its cost is negligible in itself. The power spectrum is constructed by taking the norm of the Fourier transform. Finally, all the spectra corresponding to a given point in parameter space (from the partitioning of time in segments of length T) are added together to form the final spectrum, which is searched for significant peaks up to the Nyquist frequency $4B$. All these operations are done on small segments of length $8BT$, and the total cost is the sum of the cost for building the products for all N_p points and of the cost for the spin-down corrections, power spectra construction and sum for the N'_p points corresponding to the finer grid:

$$Bt_{\text{obs}}N_p[72 + 20N'_p(1 + \log_2 8BT)] \text{ flop.} \quad (15)$$

Again, this can be adjusted by varying B and T , since N_p and N'_p depend on both B and T . The total cost of the search is the sum of Eqs. (13) and (15).

The number N_p will be the number of non-intersecting (hyper-) volume elements in parameter space required to cover completely the region of interest, such that all points within a given volume element correspond to the same choice of filtered segments from the bank, at each of the four

times of interest in this problem. Manifestly, N_p will depend on B through the way it partitions the frequency interval $[f_{\min}, f_{\max}]$, and on T , which influences the size of a volume element over which the source frequency varies by some fixed amount. We solve the problem of computing N_p numerically, by building a mean volume element: we select at random a point in parameter space, and compute its associated volume element. Repeating this procedure, we construct the average of a volume element, and divide the parameter space volume by it; the result is an estimate of N_p .

The refined gridding should only be necessary for searches with large spin-down rates (see below), and its effect will be to force to repeat N'_p times the operations following and including the resampling above, where N'_p is the number of points on the refined grid. The computation of N'_p is not of the same nature as the one for N_p : the problem now only involves the spin-down parameters, and the goal is to minimize the losses in signal power from residual frequency drifts, or misalignments of the spectra that are summed together. In that sense, it is similar to the problem of computing the number of points in parameter space for the stack-slide technique [6], and should be solved using the same geometrical approach. However, we shall rather take $N'_p = 1$, and justify this approximation by the relatively large spin-down time ($\tau = 5000$ yr) to be used in the numerical example of Sec. II. It should be kept in mind that this may not be sufficient for faster spin-down, in the sense that under this approximation N'_p may be under-estimated. With the ability we now have to compute the required computational cost as a function of B and T , the only thing that remains to be described is how we characterize the sensitivity of the search, as a function of the same choice of parameters.

B. Sensitivity

From a statistical point of view, the results in [2,5,6] are presented in a somewhat non-standard way, so to ease comparison of our results with theirs, we first briefly reformulate their definition of sensitivity. By setting the average noise power to be equal to one, we can express the signal amplitude h in units of $\sqrt{S_n(f)/t_{\text{obs}}}$, where $S_n(f)$ is the noise spectral density at frequency f . In order to be consistent with [2,5,6], we use the definition that the averaged signal power (also the square of the signal to noise ratio in amplitude) is $P_s = \langle F_+^2 \rangle h^2$, where $\langle F_+^2 \rangle = 1/5$ is the average of the detector response over all angles and polarizations [2]. This definition is useful to approximately account for variations in h over time that are produced by the motion of the source in the detector response pattern. For a true signal power P_s , computing the power spectrum of a time series as a sum of n of its subseries spectra (this describes the statistics of both the coherent and stack-slide techniques) will give a distribution $p_n(P|P_s)$ of the power P [9]. Given a false alarm probability $1 - \alpha$, one defines a threshold power P_α such that the integral of $p_n(P|0)$ from 0 to P_α is α . Observations with power exceeding P_α will be called “detections with confidence

level α ”.¹ One further defines a false dismissal probability β as a function of the signal power, as the integral of $p_n(P|P_s)$ with respect to P , from 0 to P_α . It has been customary in the field [2,5,6] to use the relation $\langle P \rangle = P_s + n$ for the average observed power $\langle P \rangle$ to interpret P_α as a lower bound on the detectable signal amplitude, by setting $P_\alpha = \langle P \rangle$. For instance, performances of sub-optimal techniques are frequently compared to the optimal case defined by $n = 1$; for $\alpha = 0.99$, the minimal detectable amplitude as defined above (the *optimal* amplitude) is $h = 4.2$, and the corresponding false dismissal rate is $\beta \approx 0.57$. In fact, it is easily checked that as n gets larger, this definition gives a value of β that approaches 1/2 from above. Therefore, for the ease of comparison with other works, we will define our minimal detectable amplitude as the smallest value of h that corresponds to a false alarm probability $1 - \alpha$ and a false dismissal rate $\beta = 1/2$. The sensitivity will be defined as $\Theta = 4.2/h$, and this definition will be approximately compatible with those of [2,5,6].

We use numerical analysis to determine h as a function of B and T . We first set the signal amplitude to zero, and generate a large number of filtered segments of pure white noise, that are multiplied together and then added by groups of $(t_{\text{obs}} - T - d/2 - \text{yr}/2)/T$ (i.e., the number of segments in the portion of the observation period that can be used for multiplications). By constructing the cumulative probability function for the resulting segments, we then evaluate the threshold P_α . We repeat the procedure for non-zero signal amplitudes: picking at random in a uniform distribution the parameters of the source, we construct the final spectrum, and evaluate the power in the expected frequency channel. For each choice of signal amplitude, we repeat this a large number of times for different source parameters, and build a distribution of the power in the expected channels; the values of β as a function of signal power are then deduced from these distributions and the previously calculated value of P_α . Therefore, the sensitivity Θ is effectively averaged over all possible source parameters. To reduce computational burden, spin-down effects are not directly included, in the sense that we only consider them in determining the bandwidth of the signal for a given T , and do not include them in the initial simulated segments that are multiplied together. By doing so, we might reduce the non-linear mixing of signal with noise, and therefore deduce a sensitivity that is too high; however, the magnitude of this mixing is such that the error so induced is expected to be small.

II. NUMERICAL SIMULATIONS

We opted for a relatively “easy” *broadband* all-sky search for our explorations. We choose $f_{\min} = 40$ Hz, $f_{\max} = 1$ kHz, $\tau = 5000$ yr, $t_{\text{obs}} = 1$ yr, and confidence level $\alpha = 99\%$. Note that $t_{\text{obs}} = 1$ yr is the minimal possible observation time, and that the data at times $t > d/2 + \text{yr}/2$ cannot be

¹The authors of [5,6] define the false alarm probability as $1 - \alpha/N_p$, where N_p is the number of points in parameter space.




FIG. 1. A contour plot of $\log_{10} N_p$ as a function of B [in units of Eq. (7)] and T . Labels on the contour lines are values of $\log_{10} N_p$.

demodulated. This obviously limits the sensitivity of the search, although it does so in a diminishing proportion as t_{obs} grows above one year. A different *narrowband* all-sky search was also considered, with parameters as for the broadband search, except for $f_{\text{min}} = 450$ Hz and $f_{\text{max}} = 500$ Hz, perhaps corresponding to the case of an interferometer made narrowband using signal recycling. For definitiveness, we choose the latitude of the detector to correspond to the LIGO Hanford detector. As described above, we computed numerically the value of N_p for different values of T and B [from now on, B will be expressed in units of $\Delta f(T)$ as defined in Eq. (7), with $f_0 = f_{\text{max}}$, and A_1, A_2 averaged over all angles δ]; see Fig. 1. Using the method described in Sec. I B, we computed the dependence of the relative sensitivity, Θ , on the same variables; the results are shown in Fig. 2.

To get a meaningful notion of the computational power P from the cost of the search [the sum of Eqs. (13) and (15)], we imposed the requirement that the analysis should be done




FIG. 2. A contour plot of Θ as a function of B [in units of Eq. (7)] and T . Labels on the contour lines are values of Θ .

in a time t_{obs} . Our map of N_p vs B and T was then reexpressed as a map of P versus B and T . We observed that for plausible filter implementations, the contribution to P from Eq. (15) was at least two orders of magnitude larger than the contribution from Eq. (13) in the region of interest of the T, B -plane; our results are highly insensitive to the real cost of constructing the filtered segments bank. Finally, we optimized our algorithm for a given computational power by choosing the values of B and T that maximize Θ along a line of constant P . The results are shown in Fig. 3, together with the corresponding sensitivity for the stack-slide technique.

Discussion

Figure 3 shows that the sensitivity of the algorithm described in this paper is roughly 30% of the sensitivity of the stack-slide technique in a broadband search, for computational power ranging from 0.4 Tflops to 8 Tflops. For the same range of computational power, the minimum amplitude detectable with 50% efficiency ranges approximately from 25 to 60 times the optimal amplitude.

Also shown in Fig. 3 is the sensitivity of an easier narrowband all-sky search. When the computational power reaches 4 Tflops, the sensitivity eventually gets better than that of the stack-slide technique, as computed for the harder, broadband search. However, it is not clear how the performances of the stack-slide technique scale with the bandwidth of the search, although they do scale strongly with the upper frequency. It is not possible, for instance, to apply the stack-slide technique on a heterodyned dataset without having to search over the additional source frequency parameter, because of the coupling of the source frequency to both the detector rest frame time and to the phases introduced by the detector motion [M_1 and M_2 in Eq. (2)]. The comparison of our algorithm to the stack-slide technique is therefore not trivial, and in particular there might be some narrowband searches where the stack-slide performances are approached or outperformed.

Our results suggest that the distortions in the noise produced by the nonlinearities inherent to our algorithm cannot be alleviated by narrowbanding the signal such that the algorithm performs in a manner that is satisfactory for a realistic broadband all-sky search. The need for narrowbanding adds one dimension (the source frequency) to the parameter space, and this greatly increases the computational burden. Moreover, the dependence of the sensitivity on the bandwidth of the bandpass filter, which is the most important parameter in determining the number of points in parameter space, is strong enough that efficient detection is computationally intensive. Consequently, our principal conclusion is that, although the symmetry described in Sec. I simplifies the problem from a formal point of view, it does so by strongly increasing the level of noise, principally by mixing components of noise of different frequencies. We have presented an algorithm that minimizes this effect, but have found that for a realistic broadband search, it did not perform better than the stack-slide technique at the same computational power.

Most detection algorithms for periodic sources, and in particular the stack-slide technique, probably do not have a




FIG. 3. Top: the relative sensitivity Θ for the optimized algorithm, as a function of computational power P (solid line), and the corresponding sensitivity for the stack-slide technique (dash-dotted line), both for the broadband search. The dotted line corresponds to the different narrowband search. Middle: the optimal segment length T in the broadband search. Bottom: the optimal filter bandwidth B [in units of Eq. (7)], also in the broadband search.

strong scaling of their performances with the bandwidth of the search, essentially because making use of the reduced bandwidth adds complexity by forcing one to consider the frequency as an explicit parameter of the search. It is therefore possible that our algorithm compares more advantageously to these other algorithms in narrowband searches, which can be considered as more natural problems for it, because it necessarily involves a search over the source frequency.

ACKNOWLEDGMENTS

The author wishes to thank Rainer Weiss for helpful discussions and comments on the work presented here. This work was supported by the National Science Foundation under cooperative agreement PHY-9210038, and by the *Fonds pour la Formation de Chercheurs et l'Aide à la Recherche* of the Province of Québec, Canada.

-
- [1] A. Abramovici *et al.*, Science **256**, 325 (1992); C. Bradachia *et al.*, Nucl. Instrum. Methods Phys. Res. A **289**, 518 (1990).
 - [2] K. S. Thorne, in *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987), pp. 330–458.
 - [3] S. V. Dhurandhar and A. Vecchio, gr-qc/9911102; gr-qc/0011085.
 - [4] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998); P. Jaranowski and A. Królak, *ibid.* **59**, 063003 (1999); P. Jaranowski and A. Królak, *ibid.* **61**, 062001 (2000).
 - [5] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
 - [6] P. R. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
 - [7] B. F. Schutz and M. A. Papa, gr-qc/9905018.
 - [8] S. L. Hahn, *Hilbert Transforms in Signal Processing* (Artech House Ed., Boston, 1996).
 - [9] E. J. Groth, Astrophys. J., Suppl. **29**, 285 (1975).
 - [10] For examples of such a heterodyne technique, see T. M. Niebauer *et al.*, Phys. Rev. D **47**, 3106 (1993).