Implementation of barycentric resampling for continuous wave searches in gravitational wave data

Pinkesh Patel,^{1,*} Xavier Siemens,^{2,†} Rejean Dupuis,^{1,‡} and Joseph Betzwieser^{1,§}

¹*Department of Physics, California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125, USA*

²*Department of Physics, University of Wisconsin—Milwaukee, P.O. Box 413, Milwaukee, Wisconsin 53201, USA*

(Received 27 January 2010; revised manuscript received 22 March 2010; published 19 April 2010)

We describe an efficient implementation of a coherent statistic for searches of continuous gravitational wave from neutron stars. The algorithm works by transforming the data taken by a gravitational wave detector from a moving Earth bound frame to one that sits at the Solar System barycenter. Many practical difficulties arise in the implementation of this algorithm, some of which have not been discussed previously. These difficulties include constraints of small computer memory, discreteness of the data, losses due to interpolation, and gaps in real data. This implementation is considerably more efficient than previous implementations of these kinds of searches on Laser Interferometer Gravitational Wave (LIGO) detector data. The speed-up factors range from 10, when applied to EINSTEIN@HOME, to about 2000 for targeted searches which integrate over months of data.

DOI: [10.1103/PhysRevD.81.084032](https://doi.org/10.1103/PhysRevD.81.084032)

PACS numbers: 04.30.Db

I. INTRODUCTION

Rapidly rotating neutron stars are among the most promising sources of continuous gravitational waves. They can emit gravitational waves through a variety of mechanisms, including unstable oscillation modes [1,2] and deformations of the crust [1,3–6]. Neutron stars can radiate powerful beams of radio waves from their magnetic poles. If a neutron star's magnetic poles are not aligned with its rotational axis, the beams sweep through space, and if the Earth lies within the sweep of the beams, the star is observed as a point source in space emitting bursts of radio waves. Such a neutron star is called a pulsar [7,8]. Since the first discovery [9], around 2000 pulsars have been detected [10,11].

Because of magnetic dipole radiation and gravitational radiation, the rotational frequencies of neutron stars slowly decrease in time. Other than this effect, gravitational waves from isolated rotating neutron stars are essentially monochromatic in the rest frame of the star. The waves are continuous and their frequency is determined by the rotational frequency of the star. The motion of the detector as the earth rotates about its axis and around the sun, however, modulates the phase as well as the amplitude of the received signal. In order to recover the signal from interferometric data optimally, both of these effects must be taken into account. Detecting gravitational waves from neutron stars could reveal information about the strength of neutron star crusts and the equation of state of the nuclear matter that makes up the star [6]. Continuous gravitational waves may also be produced by other sources, such as cosmic strings [12,13].

There are a number of techniques available for continuous wave searches. These techniques can be loosely divided into two categories: (1) coherent methods [14,15], which keep track of the phase of the gravitational wave signals over long periods of time, and (2) semicoherent methods [16], which combine shorter periods of data without tracking the phase (for example, taking Fourier transforms of short segments of data and then summing the power).

When the sky location and phase evolution of a neutron star are known, a coherent search for continuous gravitational waves is relatively straightforward [15]. Assuming that the noise in a gravitational wave detector follows Gaussian statistics, in the presence of a signal, the signal to noise recovered in a search increases with the square root of the amount of data used in the search. This is because the signal amplitude accumulates linearly while the noise follows a random walk. Thus, with enough data, it is possible to recover any continuous signal out of noisy data.

If certain parameters of the signal (sky location, frequency, spin-downs, and binary parameters) are not known the search becomes much more involved. The reason is that the number of points needed to cover the search parameter space (and ensure no signals are lost) grows like a large power of the amount of data used [17]. This makes the sensitivity of gravitational wave searches computationally bound: One cannot simply integrate arbitrary amounts of data to gain sensitivity because there is not enough computational power available to perform the search. Thus, more efficient code and greater computing power are highly desirable, since they translate into more data being analyzed and therefore an increase in the sensitivity of gravitational wave searches.

A promising method for blind searches involves exploiting large-scale correlations in the coherent detection sta-

*ppatel@ligo.caltech.edu

†siemens@gravity.phys.uwm.edu

‡rejean@gmail.com

§josephb@ligo.caltech.edu

tistic [18]. Another method that has been successful in these kinds of searches is the hierarchical scheme of incoherently combining coherent sets of data. Some of the methods currently under use include the Hough transform and stack-slide [16].

In this paper we focus on an efficient implementation of coherent techniques. The method we present here is similar to several previous implementations [19–21], in that it uses fast Fourier transforms (FFTs) to calculate the so-called $\mathcal{F}$ -statistic [14], the logarithm of the likelihood function maximized over some of the unknown parameters of the gravitational wave produced by a neutron star. These unknown parameters include the strain amplitude h_0 , the polarization angle ψ , the initial phase Φ_0 , and the inclination angle ι . An important difference between these methods and ours, is that we resample the time domain data to the Solar System barycenter before taking a FFT. This allows us to use a single FFT to calculate the detection statistic for arbitrarily many frequencies and an arbitrary amount of observation time, while previous implementations have a maximum frequency band and observation time that can be calculated with a single FFT, which are determined by losses due to phase mismatch. Another set of techniques, described in [22–24], implement stroboscopic resampling methods described in [25]. The stroboscopic method requires data at full bandwidth, and is therefore not suitable for distributed computing applications such as EINSTEIN@HOME.

In Sec. II we review the signal properties and the nearly optimal coherent statistic that can be used to extract continuous signals from interferometric gravitational wave data. In Sec. III we discuss how to implement the calculation of the nearly optimal coherent statistic in a computationally efficient way in both the time and frequency domains. In Sec. IV we describe the results of a computer code using this algorithm on software injections of gravitational waves into Gaussian noise. Lastly, in Sec. V we address important technical issues to do with practical implementations of barycentric resampling.

II. PRELIMINARIES

In this section we closely follow the method of Jaranowski, Krolak, and Schutz [14] to provide the background on the signal and the detection statistic. Power-recycled Fabri-Perot Michelson interferometers such as those used by the Laser Interferometer Gravitational Wave Observatory (LIGO) are sensitive to the strain caused by gravitational waves passing through it. The strain measured at a detector can be written as [14]

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t), \quad (1)$$

where t is the time in the detector frame, and h_+ and $h_\times$ are the “plus” and “cross” polarizations of the gravitational wave. $F_+(t)$ and $F_\times(t)$ are the beam-pattern functions of the interferometer and are given by

$$F_+(t) = \sin\zeta[a(t)\cos 2\psi + b(t)\sin 2\psi], \quad (2)$$

and

$$F_\times(t) = \sin\zeta[b(t)\cos 2\psi - a(t)\sin 2\psi], \quad (3)$$

where ψ is the polarization angle of the wave and ζ is the angle between detector arms (which in the case of LIGO is 90°). The functions $a(t)$ and $b(t)$ both depend on time and location of source and detector, but are independent of the polarization angle ψ .

In the detector frame the phase of a gravitational wave produced by an isolated neutron star can be written as [14]

$$\begin{aligned} \Psi(t) = & \Phi_0 + 2\pi \sum_{k=0}^s f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \mathbf{n}_0 \cdot \mathbf{r}_d(t) \\ & \times \sum_{k=0}^s f_0^{(k)} \frac{t^k}{k!}, \end{aligned} \quad (4)$$

where Φ_0 is the phase at the start time of the observation, $f_0^{(k)}$ is the k th derivative of the frequency, c is the speed of light, α and δ are the right ascension and declination of the source, $\mathbf{n}_0 = \mathbf{n}_0(\alpha, \delta)$ is the unit vector of the source in the Solar System barycenter (SSB) reference frame, $\mathbf{r}_d$ is the position vector of the detector in the same frame, and s is the order of the expansion. Neglecting changes in the proper motion of the star, the third term in Eq. (4) is a correction to the phase due to the detector motion relative to the neutron star.

We can define $\Phi(t) = \Psi(t) - \Phi_0(t)$, as well as defining

$$\Phi_s(t) = 2\pi \sum_{k=1}^s f_0^{(k)} \frac{t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \mathbf{n}_0 \cdot \mathbf{r}_d(t) \sum_{k=1}^s f_0^{(k)} \frac{t^k}{k!} \quad (5)$$

and

$$t_m = \frac{\mathbf{n}_0 \cdot \mathbf{r}_d(t)}{c}. \quad (6)$$

Equations (5) and (6) let us write

$$\Phi(t) = 2\pi f[t + t_m(t; \alpha, \delta)] + \Phi_s(t; f^{(k)}, \alpha, \delta), \quad (7)$$

which has the modulation due to the detector’s motion around the SSB clearly separated from the modulation due to the gravitational wave’s intrinsic frequency, although not the derivatives of the frequency.

An almost optimal statistic for the detection of continuous gravitational wave signals is called the $\mathcal{F}$ -statistic [14,26]. It is the logarithm of the likelihood function maximized over the extrinsic and unknown signal parameters. The $\mathcal{F}$ -statistic is given by

$$\mathcal{F} = \frac{4}{S_h(f)T_0} \frac{B|F_a|^2 + A|F_b|^2 - 2CR(F_a F_b^*)}{D}, \quad (8)$$

where $S_h(f)$ is the one-sided spectral density of the detector’s noise at frequency f and T_0 is the observation time. A ,

B , C , and D are given by

$$\begin{aligned} A &= (a \parallel a); & B &= (b \parallel b); \\ C &= (a \parallel b); & D &= A \cdot B - C^2 \end{aligned} \quad (9)$$

with

$$(p \parallel q) = \frac{2}{T_0} \int_{-T_0/2}^{T_0/2} p(t)q(t)dt. \quad (10)$$

F_a and F_b are integrals defined as

$$F_a(f) = \int_{-T_0/2}^{T_0/2} a(t)x(t)e^{-i\Phi_s(t)}dt \quad (11)$$

and

$$F_b(f) = \int_{-T_0/2}^{T_0/2} b(t)x(t)e^{-i\Phi_s(t)}dt, \quad (12)$$

where $x(t)$ is the time series data output by an interferometer. We define a new time variable called t_b as follows:

$$t_b = t + t_m. \quad (13)$$

Taking a derivative with respect to t on both sides of Eq. (13), we get

$$\frac{dt_b}{dt} = 1 + \frac{dt_m}{dt}. \quad (14)$$

From Eqs. (6) and (14), we get

$$\frac{dt_m}{dt} = \frac{\mathbf{n}_0 \cdot \mathbf{v}_d(t)}{c}, \quad (15)$$

where $\mathbf{v}_d(t)$ is the velocity of the detector in the SSB frame and thus $\frac{\mathbf{n}_0 \cdot \mathbf{v}_d(t)}{c}$ is the Doppler shift of the source with respect to the detector. For a detector located on Earth, the maximum Doppler shift experienced is of the order of 10^{-4} . Using this fact and Eq. (14) we get $\delta t_b \approx \delta t$.

We can thus rewrite the Eqs. for F_a and F_b as

$$F_a(f) = \int_{-T_0/2}^{T_0/2} a(t_b)x(t_b)e^{-2\pi i f t_b}e^{-i\Phi_s(t_b)}dt_b, \quad (16)$$

and

$$F_b(f) = \int_{-T_0/2}^{T_0/2} b(t_b)x(t_b)e^{-2\pi i f t_b}e^{-i\Phi_s(t_b)}dt_b \quad (17)$$

which are just the Fourier transforms of the resampled data and the detector response, multiplied by a phase $e^{-i\Phi_s(t_b)}$ [14]. Equations (16) and (17) can be efficiently evaluated using FFTs. Details of the resampling procedure can be found in Sec. III A 2.

III. IMPLEMENTATION OF BARYCENTRIC RESAMPLING

Gravitational wave detectors collect data at the rate of about 16-20 kHz for spans of time on the order of a year. This means that typical searches for gravitational waves

will involve on the order of a terabyte (TB) of data. Computers currently have memories of a few gigabytes (GB), making it necessary to break up the data into pieces that can fit in the memory of a single computer. To analyze the full data set hundreds to thousands of these computers can then be used together in the form of a Beowulf cluster, or tens to hundreds of thousands with distributed computing systems such as EINSTEIN@HOME [27].

A. Time domain analysis

The $\mathcal{F}$ -statistic can be calculated from a time series directly by following the steps outlined in Sec. II. However, due to the large amounts of data involved, it is impractical to do this for the entire data set. One way to address this problem is to divide the data into band-limited time series, making it possible to analyze one small sub-band at a time. Time series spanning different frequency bands are then analyzed in parallel on a Beowulf cluster or a distributed computing system. In this section we provide details on how this is accomplished in the time domain, and address some of the difficulties that arise.

1. Heterodyning, low-pass filtering, and downsampling

Let the output of the instrument be the time series $x(t)$, and its Fourier transform be

$$\tilde{x}(f) = \int_{-\infty}^{\infty} x(t)e^{-2\pi i f t}dt. \quad (18)$$

If we consider the Fourier transform of the complex time series $x_h(t) = x(t)e^{-2\pi i f_h t}$,

$$\begin{aligned} \tilde{x}_h(f) &= \int_{-\infty}^{\infty} x(t)e^{-2\pi i f_h t}e^{-2\pi i f t}dt \\ &= \int_{-\infty}^{\infty} x(t) \cdot e^{-2\pi i (f+f_h)t}dt = \tilde{x}(f+f_h), \end{aligned} \quad (19)$$

it is obvious that multiplying the time series $x(t)$ by $e^{-2\pi i f_h t}$ has shifted all the frequencies in the time series $x(t)$ by f_h . This procedure is referred to as complex heterodyning.

If just a small frequency band B of data around f_h is of interest, low-pass filtering followed by downsampling can be used to reduce the bandwidth of the data appropriately. Specifically, if we wish to downsample by a factor D , the new Nyquist frequency of our time series will be given by

$$f_{\text{Nyq,new}} = \frac{f_{\text{Nyq,old}}}{D} = \frac{B}{2}. \quad (20)$$

A simple but effective downsampling technique involves picking every D th point in the time series. To avoid aliasing effects however, prior to downsampling a low-pass filter must be applied to the data with a sharp falloff around the new Nyquist frequency. The heterodyned, band-limited, downsampled complex time series will have a sampling time $\Delta t = 1/B$. For example, suppose we are only interested in analyzing data between 990 Hz and 1 kHz. By multiplying the data with the phase factor

$e^{-2\pi(995)it}$, data at 995 Hz moves to 0 Hz (DC), 990 Hz moves to -5 Hz, and 1 kHz to $+5$ Hz (we have taken t to be measured in seconds). To avoid aliasing problems when we downsample, we low-pass filter the data at 5 Hz, the new Nyquist frequency. We can then downsample by picking one point out of every 100. The resulting complex time series will be sampled at 10 Hz and contain all the information in the original time series between 990 Hz and 1 kHz.

2. Barycentric resampling and heterodyne correction

In this section we explain how to use the low bandwidth heterodyned complex time series to compute the $\mathcal{F}$ -statistic given by Eq. (8).

In the following we will work only with F_a . The procedure for F_b is completely analogous. It is easiest to begin with the integral definition for F_a in Eq. (11) with the phase explicitly written out, namely,

$$F_a(f) = \int_{-T_0/2}^{T_0/2} a(t)x(t)e^{-2\pi if(t+t_m)}e^{-i\Phi_s(t)}dt, \quad (21)$$

and a similar expression holds for F_b . The heterodyned version of F_a is

$$F_a(f + f_h) = \int_{-T_0/2}^{T_0/2} a(t)x(t)e^{-2\pi i(f+f_h)(t+t_m)}e^{-i\Phi_s(t)}dt. \quad (22)$$

If we already have a complex heterodyned time series $x_h(t)$ (heterodyned in the detector frame), we can use it to absorb some (but not all) of the heterodyne exponent in Eq. (22) as follows:

$$x(t)e^{-2\pi i(f+f_h)(t+t_m)} = x_h(t)e^{-2\pi if_h t_m}e^{-2\pi if(t+t_m)}. \quad (23)$$

This means that rather than Eq. (22), we should evaluate

$$F_a(f + f_h) = \int_{-T_0/2}^{T_0/2} a(t)z(t)e^{-2\pi if(t+t_m)}e^{-i\Phi_s(t)}dt, \quad (24)$$

where

$$z(t) = x_h(t)e^{-2\pi if_h t_m}. \quad (25)$$

At this point we have an expression which looks like Eqs. (11) and (12), and we can write the integral over t instead as an integral over t_b :

$$F_a(f + f_h) = \int_{-T_0/2}^{T_0/2} a(t_b)z(t_b)e^{-2\pi if t_b}e^{-i\Phi_s(t_b)}dt_b, \quad (26)$$

with a similar expression for F_b .

The discrete version of Eq. (26) for a time series with N points reads

$$F_a(f + f_h) = \sum_{k=1}^N a(t_b^k)z(t_b^k)e^{-2\pi if t_b^k}e^{-i\Phi_s(t_b^k)}dt_b, \quad (27)$$

and a similar expression holds for F_b :

$$F_b(f + f_h) = \sum_{k=1}^N b(t_b^k)z(t_b^k)e^{-2\pi if t_b^k}e^{-i\Phi_s(t_b^k)}dt_b, \quad (28)$$

where t_b^k is the k th datum in the time series as measured in the barycentric frame and $dt_b = t_b^{k+1} - t_b^k$. The relationship between t_b and t can be written as

$$t_b^k = t^k + t_m(t^k; \alpha, \delta). \quad (29)$$

This relationship between t^k and t_b^k can be used to calculate $z(t_b^k)$ from the time series $z(t^k)$. In practice, one starts out with $z(t^k)$, i.e. data sampled regularly in the detector frame. Then we calculate $T^k(t_b^k)$, which are detector times corresponding to regularly spaced samples in the barycentric frame. These $T^k(t_b^k)$ are irregularly sampled in the detector frame, but since we have $z(t^k)$, we can calculate $z(T^k(t_b^k))$ by using interpolation. The interpolated time series $z(T^k(t_b^k))$ is the $z(t_b^k)$ of Eqs. (27) and (28). A similar procedure may be used to calculate the $a(t_b^k)$ from $a(t^k)$, and the $b(t_b^k)$ from $b(t^k)$. The factor of $e^{i\Phi_s(t_b^k)}$ in Eqs. (27) and (28) is calculated using Eq. (5). In this case, instead of calculating $\Phi_s(t^k)$, we calculate $\Phi_s(T^k(t_b^k))$, which is equivalent to calculating $\Phi_s(t_b^k)$. While in theory one has to calculate the quantity $\mathbf{n}_0 \cdot \mathbf{r}_d(t)$ in Eq. (5), in practice this information is already encoded in $T^k(t_b^k)$ as

$$\mathbf{n}_0 \cdot \mathbf{r}_d(t) = t_m \cdot c = (t_b^k - T^k(t_b^k)) \cdot c. \quad (30)$$

With all the parts of Eqs. (27) and (28) in hand, we can compute $F_a(f + f_h)$ and $F_b(f + f_h)$.

In summary, as shown in Fig. 1, the procedure is the following:

- (1) Start with a heterodyned, band-limited, down-sampled $x_h(t^k)$ with t^k regularly spaced in time, in the frame of reference of the detector.
- (2) Correct the $x_h(t^k)$ for the heterodyning done in the detector frame by multiplying with $e^{-2\pi if_h t_m}$ to produce the $z(t^k)$.

FIG. 1 (color online). Graphical description of the resampling procedure.

- (3) The $z(t^k)$ correspond to data irregularly spaced in the barycentric frame. Calculate $T^k(t_b^k)$, which are times in the detector frame corresponding to regularly sampled solar system barycenter times.
- (4) Using interpolation, calculate $z(T^k(t_b^k))$ from $z(t^k)$, which is the $z(t_b^k)$ used in Eqs. (27) and (28).
- (5) Similarly, from $a(t^k)$ and $b(t^k)$ calculate $a(t_b^k)$ and $b(t_b^k)$, respectively.
- (6) Using FFTs, evaluate Eqs. (27) and (28) to calculate $F_a(f + f_h)$ and $F_b(f + f_h)$.
- (7) Use Eq. (8) to calculate the $\mathcal{F}$ -statistic.

B. Frequency domain analysis

In the previous section we describe a practical way of calculating the $\mathcal{F}$ -statistic from time series data. However, in practice the calculation is done in the frequency domain for a couple of reasons. One is that much of the code written in the LIGO Scientific Collaboration's (LSC) Continuous Waves working group is tailored to an analysis performed in the frequency domain and hence there exist many data processing and validation tools to process the data that are useful to this code. Another reason is that gravitational wave detectors are subject to many sources of noise, some of which change daily or even hourly, such as wind, microseism, earthquakes, anthropogenic noise, etc. These change the noise floor of any analysis as a function of time. Working in the frequency domain is a natural way to deal with this problem.

We begin a frequency domain analysis by taking short time-baseline Fourier transforms of the time domain data, called short Fourier transforms (SFTs). When we calculate the $\mathcal{F}$ -statistic, we divide by the noise in the instrument at that frequency, as shown in Eq. (8). However, Eq. (8) assumes the noise is stationary. To account for the non-stationarity of the noise we need to weight by the noise over time, which is done on a per SFT basis. This normalization process is described in the next section.

The computational cost of estimating the noise per SFT scales with the number of SFTs and thus for a fixed observation time scales inversely with the time baseline. A compromise is needed between the demands of computational time and relative stationarity of the detector for a given time baseline. In LIGO, SFTs are usually 1800 seconds long, since the detector is reasonably stationary for that time.

1. Dealing with nonstationary and colored data

To deal with nonstationarities, variations in the noise floor from SFT to SFT, and colored data, we can normalize our SFT data to absorb the $1/S_h(f)$ term in the definition of the $\mathcal{F}$ -statistic in Eq. (8). If $X_{\alpha,k}$ is the k th frequency bin of the α th SFT, then we can redefine a normalized data point $\hat{X}_{\alpha,k}$ as

$$X_{\alpha,k} \rightarrow \hat{X}_{\alpha,k} = \frac{X_{\alpha,k}}{\sqrt{S_{\alpha,k}}}, \quad (31)$$

where $S_{\alpha,k}$ is an estimate of the one-sided power spectral density for the k th frequency bin of the α th SFT. Estimators used for this purpose should be robust in the presence of spectral features in the data, such as a running median.

2. Merging SFTs into long time-baseline Fourier transforms

There are many practical difficulties that arise when dealing with SFTs. Often contiguous chunks of data have to be divided up into multiple SFTs and it is necessary to coherently combine them into one long time-baseline SFT. This is done using the Dirichlet kernel, which is the equivalent of a sinc interpolation (ideal interpolation) done in the time domain. In order to keep the computational cost down, the Dirichlet kernel is truncated at a finite number of points (usually around 16). This introduces a slight interpolation error, which cannot be avoided without sacrificing a large amount of computational power.

Suppose we divide the data $x(t)$ of length T_0 into M short chunks of length T_{SFT} each with N points, so that $T_0 = MT_{\text{SFT}}$. The discrete Fourier transform (DFT) of the data is

$$X_b = \sum_{l=0}^{NM-1} x_l e^{-2\pi i l b / NM}, \quad (32)$$

where $x_l = x(l\Delta t)$, Δt is the sampling time, and b is a long time-baseline frequency index. We can write the Fourier transform in terms of two sums:

$$X_b = \sum_{\alpha=0}^{M-1} \sum_{j=0}^{N-1} x_{\alpha,j} e^{-2\pi i b(j+N\alpha)/NM}, \quad (33)$$

where $x_{\alpha,j} = x((j + N\alpha)\Delta t)$. We can express the $x_{\alpha,j}$ in terms of an inverse DFT of a short chunk of data,

$$x_{\alpha,j} = \frac{1}{N} \sum_{k=0}^{N-1} X_{\alpha,k} e^{2\pi i j k / N}, \quad (34)$$

where the $X_{\alpha,k}$ are the starting SFT data,

$$X_{\alpha,k} = \sum_{j=0}^{N-1} x_{\alpha,j} e^{-2\pi i j k / N}. \quad (35)$$

Replacing $x_{\alpha,j}$ with Eq. (34) in Eq. (33) gives

$$\begin{aligned} X_b &= \sum_{\alpha=0}^{M-1} \sum_{j=0}^{N-1} \left(\frac{1}{N} \sum_{k=0}^{N-1} X_{\alpha,k} e^{2\pi i j k / N} \right) e^{-2\pi i b(j+N\alpha)/NM} \\ &= \frac{1}{N} \sum_{\alpha=0}^{M-1} e^{-2\pi i b \alpha / M} \sum_{k=0}^{N-1} X_{\alpha,k} \sum_{j=0}^{N-1} e^{-2\pi i j(b/M-k)/N}. \end{aligned} \quad (36)$$

The last sum in this expression can be evaluated analytically. In particular,

$$\sum_{j=0}^{N-1} z^{cj} = \frac{1 - z^{Nc}}{1 - z^c}. \quad (37)$$

We take $z = e$, $c = -iy/N$, with $y = 2\pi(b/M - k)$, so that the sum is given by

$$\sum_{j=0}^{N-1} e^{-iyj/N} = \frac{1 - e^{-iy}}{1 - e^{-iy/N}}. \quad (38)$$

In the large N limit the exponent of the denominator will be small so that

$$\begin{aligned} \frac{1 - e^{-iy}}{1 - e^{-iy/N}} &\approx \frac{1 - e^{-iy}}{1 - (1 - iy/N)} = \frac{iN}{y} (e^{-iy} - 1) \\ &= N \left(\frac{\sin y}{y} - i \frac{1 - \cos y}{y} \right). \end{aligned} \quad (39)$$

This means we can write Eq. (36) as

$$X_b = \sum_{\alpha=0}^{M-1} e^{-2\pi i b \alpha / M} \sum_{k=0}^{N-1} X_{\alpha,k} P_{b,k}, \quad (40)$$

with the Dirichlet kernel

$$P_{b,k} = \frac{\sin y}{y} - i \frac{1 - \cos y}{y}, \quad (41)$$

and $y = 2\pi(b/M - k)$. The function $P_{b,k}$ is very strongly peaked around $y = 0$, which is near a value of the frequency index $k^* = \text{floor}(b/M)$. This means one only needs to evaluate the sum over k for a few terms Δk around k^* . With this in mind we write

$$X_b \approx \sum_{\alpha=0}^{M-1} e^{-2\pi i b \alpha / M} \sum_{k=k^*-\Delta k}^{k^*+\Delta k} X_{\alpha,k} P_{\alpha,k}. \quad (42)$$

To produce a heterodyned time series a sub-band of the X_b may be selected and inverse Fourier transformed.

3. Normalized long time-baseline Fourier transforms

With the normalized SFT data $\hat{X}_{\alpha,k}$ from Eq. (31) we can construct a normalized version of the long time-baseline Fourier transform

$$\hat{X}_b \approx \sum_{\alpha=0}^{M-1} e^{-2\pi i b \alpha / M} \sum_{k=k^*-\Delta k}^{k^*+\Delta k} \hat{X}_{\alpha,k} P_{\alpha,k}, \quad (43)$$

and take a sub-band of $\hat{X}_b$, inverse Fourier transform it, and produce the heterodyned time series, and correct it to produce $\hat{z}(t_b^k)$. In terms of this time series, we can write

$$\hat{F}_a(f + f_h) = \sum_{k=1}^N \hat{z}(t_b^k) a(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \quad (44)$$

and

$$\hat{F}_b(f + f_h) = \sum_{k=1}^N \hat{z}(t_b^k) b(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)}, \quad (45)$$

and thus

$$\mathcal{F} = \frac{4}{T_0} \frac{B|\hat{F}_a|^2 + A|\hat{F}_b|^2 - 2C\Re(\hat{F}_a\hat{F}_b^*)}{D}. \quad (46)$$

4. Heterodyning

As shown before in Eqs. (18) and (19), heterodyning is a procedure by which the frequency of interest can be shifted arbitrarily. When one applies the kind of correction in Eq. (18), we effectively move all the frequencies by a set amount. By doing so, we convert the time series from a real time series to a complex time series, with the same amount of information content.

Heterodyning in the frequency domain can be done in two ways, one in which the time series produced after inverse Fourier transforming is real and another in which it is complex. A cosine transform used to heterodyne would produce a real time series, but this method is not used in an implementation of the technique (see Sec. IV). A complex heterodyned time series is produced by inverse Fourier transforming a relabelled band of the frequencies. Since in Eq. (18), all frequencies are shifted by a fixed amount, the equivalent procedure in the frequency domain is just relabelling the heterodyne frequency f_h as DC and subsequently all the other frequencies relative to this new DC.

Taking the example from Sec. III A 1, we can just internally change the labels of the 995 Hz frequency bin to DC and 1000 Hz to 5 Hz. Once this relabelling is done, the original data will have all shifted by 995 Hz, with the 10 Hz from -5 Hz to $+5$ Hz containing all the relevant information. If one were using the whole band without down-sampling or filtering, then this relabelling would have to wrap around the Nyquist frequency edge, but since the whole purpose of heterodyning is to downsample, it is never necessary to do so.

5. Downsampling and low-pass filtering

Following the time domain algorithm, after heterodyning the data, it needs to be downsampled and low-pass filtered. The downsampling and low-pass filtering is achieved by simply throwing out the data that is not in the band of interest. The heterodyning is done in such a way as to keep the center of the band of interest at DC. A Tukey window applied to the band of interest, keeping a little bit of data on both edges to facilitate the rise of the window from 0 to 1, is a good choice of a low-pass filter. Once an inverse Fourier transform is performed on this smaller subset in the frequency domain, it generates the same heterodyned, downsampled, and low-pass filtered time series as the time domain algorithm.

6. Gaps in the data

Data collected by an interferometer will have gaps due to periods of downtime. These gaps need to be dealt with in a manner that preserves the phase coherence of the segments around the gaps. The gaps increase the analysis time without contributing any power to the $\mathcal{F}$ -statistic, and thus act like a zero padding.

The data is divided up into a series of contiguous chunks and gaps. For each contiguous chunk the SFTs in that chunk are normalized, patched up and then a heterodyned, downsampled and low-pass-filtered time series is calculated from it. Heterodyning done by relabelling is equivalent to multiplying with $e^{-2\pi i f_h(t - t_c)}$, where t_c is the start time of the data chunk being heterodyned and f_h is the heterodyne frequency. If we have multiple chunks that are separately being heterodyned, then t_c is different for each chunk. In the time domain analysis, we assumed that the heterodyne reference time is the same as the start time of the analysis. In order to achieve the same kind of heterodyning, one needs to multiply each newly created time series with a correcting phase factor, namely,

$$e^{-2\pi i f_h(t_c - t_s)}, \quad (47)$$

where t_s is the start time of the overall analysis.

A Tukey window can then be applied to each of these time series to smoothly bring the data to zero at the edges, which correspond to the gaps. The gaps are then filled with zeros, as no data was collected during those times. This procedure is repeated for all the gaps and contiguous chunks. At the end, a time series is produced, which is contiguous and spans the time of the analysis. By ensuring that the time stamps of the first datum of each contiguous chunk correspond with the start time of that chunk, we ensure that the phase coherence is maintained throughout.

7. Summary

To summarize, as shown in Fig. 2, a simple algorithm to produce a time series equivalent to the one used for the time domain analysis is as follows:

- (1) Divide the data into time chunks and Fourier transform them to create SFTs.
- (2) Normalize these SFTs and assign them weights.
- (3) Identify contiguous sets of SFTs.

FIG. 2. Pictorial description of data preprocessing.

- (4) Combine each contiguous chunk of SFTs into one long time-baseline Fourier transform (FT).
- (5) Create a downsampled, heterodyned, and low-pass-filtered time series by inverse Fourier transforming the desired frequencies from the FT.
- (6) Stitch all these time domain chunks together, filling gaps with zeros.

IV. RESULTS

A. Speed

The scheme previously used to compute the $\mathcal{F}$ -statistic, involved the use of the Dirichlet kernel to combine a series of SFTs [28,29], which were calculated for 30 minutes of data taken at 16 kHz. The 30 minute window was set by the maximum Doppler shift due to the motion of the Earth. A C code called COMPUTEFSTATISTIC_V2 [30] was written in the LIGO Analysis Library (LAL) to calculate the $\mathcal{F}$ -statistic using this algorithm. The code which implements our method is also written in C and is called COMPUTEFSTATISTIC_RESAMP [31]. Henceforth we will refer to the previous implementation as the LAL implementation and our implementation as Resampling.

The $\mathcal{F}$ -statistic is calculated for a series of templates looping over various parameters such as sky location, α and δ , spin-downs f^k , and various frequencies f . We can ignore the way the two implementations deal with loops over α , δ , and f^k , since they both loop over them in the same manner. The speed of computation for a loop over frequencies f is worth comparing, however.

Assume that we have N data points (take for example 10^6 seconds of data at 100 Hz, i.e. 10^8 data points). Now assume that the number of operations per sky location and per spin-down is N_{ops} . If the number of Dirichlet kernel points used is N_{DirKer} , then the total number of operations used by the LAL implementation is

$$N_{\text{Tot}}^{\text{LAL}} = N_{\text{ops}} \cdot N_{\text{DirKer}} \cdot N_{\text{SFTs}} \cdot N, \quad (48)$$

where N_{ops} is defined as the number of operations conducted in the innermost loop and is approximately of order 10, N_{DirKer} is the number of times the Dirichlet Kernel loop is repeated, $N_{\text{SFTs}} = \frac{T_{\text{obs}}}{T_{\text{SFT}}}$ is the number of SFTs, and N is the number of data points.

Compare this to the resampling method, which consists of 4 major steps:

- (1) Calculating $t_b(t)$, given a sky location and time.
- (2) Calculating the integrands of F_a and F_b .
- (3) Interpolating and calculating the beam patterns.
- (4) Taking the Fourier transform.

Each of these steps involves order 10 operations, but all of these steps are sequential, therefore they only add, resulting in a total number of operations per data point, $N_{\text{ops}}^{\text{Resamp}}$, of approximately 30 operations. The last step is the Fourier transform, which is of order $N \log N$, therefore the total number of steps is

$$N_{\text{Tot}}^{\text{Resamp}} = (N_{\text{ops}}^{\text{Resamp}} + \log N) \cdot N. \quad (49)$$

Therefore the ratio of operations between the two methods is

$$\frac{N_{\text{Tot}}^{\text{LAL}}}{N_{\text{Tot}}^{\text{Resamp}}} = \frac{N_{\text{ops}} \cdot N_{\text{DirKer}} \cdot N_{\text{SFTs}}}{N_{\text{ops}}^{\text{Resamp}} + \log N}. \quad (50)$$

To first order, we have

$$\frac{N_{\text{Tot}}^{\text{LAL}}}{N_{\text{Tot}}^{\text{Resamp}}} \approx \frac{N_{\text{SFTs}}}{\log N}. \quad (51)$$

Therefore for large observation times, this method of calculating the $\mathcal{F}$ -Statistic is faster and, in the case of a targeted search, it allows for a large parameter space in $F^{(k)}_s$.

The speed-up in practice is reduced by a few practical issues as seen in Sec. V. However, Resampling is still considerably more efficient than the LAL implementation. For EINSTEIN@HOME, because of the relatively small coherent integration time, the speed-up is around 10. But for targeted searches that span multiple months or years, the improvement can be as high as a factor of 2000. Thus, while some targeted searches which integrate over a couple of years were impossible to do previously, they are now possible.

B. Validations

The probability density distribution of the $\mathcal{F}$ -statistic for Gaussian noise of zero mean and unity standard deviation is a χ^2 distribution with 4 degrees of freedom. In the presence of a signal, the distribution is a χ^2 of 4 degrees of freedom with a noncentrality parameter given by the $\mathcal{F}$ -statistic in the absence of noise for the particular signal.

Resampling uses various approximate methods in the calculation of the $\mathcal{F}$ -statistic, and this can lead to disagreements between the theoretical $\mathcal{F}$ -statistic probability density function and the output of the code. These changes are of the order of a few percent and are within acceptable limits. The validity of the code can be tested by using a Monte Carlo simulation of about a million different injections of the same signal in different instances of noise. The noise is generated as a Gaussian noise of zero mean and unity standard deviation, and the signal is added into this noise. For each individual injection the signal is chosen with a given set of amplitude parameters and a fixed sky location and spin-downs, and the search is conducted over these exact chosen parameters in order to avoid any mismatches. These Monte Carlos are then repeated with another set of parameters, which are themselves chosen randomly. While it is not an exhaustive test, randomly chosen parameters ensure that we are not biased in the validation test. The plot in Fig. 3 is produced by performing one such Monte Carlo simulation. In this case, both the

FIG. 3 (color online). Histogram of results of Monte Carlo simulation with signals injected in different instances of noise.

LAL implementation and Resampling were run on the same set of data. The $\mathcal{F}$ -statistic was picked out at the appropriate frequency and this was repeated about a million times. A histogram of these $\mathcal{F}$ -statistic values was then plotted. As one can see, there is very good agreement in between the expected distribution of the $\mathcal{F}$ -statistic and the two implementations.

V. PRACTICAL CONSIDERATIONS

A. Discreteness

In the implementation of the algorithm explained above, one major obstacle is the fact that the data collected by any physical instrument is discrete and thus must be handled appropriately. Take, for example, the heterodyne frequency used in the calculation. This frequency cannot be chosen arbitrarily, as only certain frequencies are sampled and thus there are only certain permitted choices.

Most major FFT computation algorithms output the frequency series in a specific format, which split the data into two parts. The first bin output by these algorithms is the DC followed by the first positive frequency bin up to positive Nyquist and then follows this up with the negative frequencies starting at the negative Nyquist frequency. This order of placing frequency bins speeds up computation and is necessary for the internal workings of these algorithms. Thus when an inverse FFT is performed on the frequency domain data in the form of SFTs, a simple reshuffling needs to be done. The frequency selected to be the first bin will become the new DC and thus the data will have been heterodyned by that said frequency. In order to ensure that the same frequency bin is chosen as DC, one needs an odd number of bins per SFT. If the number of bins are even, then upon increasing the amount of data it can shift this number to an odd number as the increase is always done by changing the number of SFTs. But if the number of bins per SFT is odd, then it will remain odd for any number of SFTs. This ensures that there is no mismatch in choosing the appropriate bin as the heterodyne frequency.

B. Interpolation issue

When the resampling algorithm is used on discrete data, one needs to interpolate between data points to go from the detector frame to the SSB frame. This interpolation acts like a nonlinear low-pass filter and destroys the power at higher frequencies in the band of analysis. Since the filter is nonlinear, the frequency response is not well defined, and thus there is no way to compensate for the power loss at high frequencies. The power loss can be significant (of order 30%) and is unacceptable in most analyses. The exact nature of the filter depends on the type of the interpolation routine used and the sky location that one resamples to. The only work around is to perform the computation over a larger band than the one desired. In practice it is sufficient to double the band and to discard the higher frequencies.

VI. SUMMARY AND CONCLUSIONS

In this paper, we describe an efficient implementation of the barycentric resampling technique, which deals with the nonstationarity of the detector and calculates the $\mathcal{F}$ -statistic. Although the calculation of the $\mathcal{F}$ -statistic has been targeted, this technique can be used for many other kinds of searches. The major contribution of this technique is to remove the Doppler shift of the Earth's motion in a gravitational wave signal. Thus, once this Doppler shift is removed, both frequentist and Bayesian techniques can be applied to the data. In the process of implementing this algorithm, a series of practical issues are dealt with, including constraints of modern computer memory, discreteness of the data taken, losses due to interpolation, and gaps in real data.

The computational savings due to this technique can be used in various ways. One such use is to increase the coherent integration time for all-sky searches like the EINSTEIN@HOME searches. Currently EINSTEIN@HOME [27] uses 40 hour long coherent integration time. The resampling code will be about 10 times faster for such integration times, and for the same computational power and keeping the same scaling for the search, we can coherently integrate 64 hours instead, which corresponds to a sensitivity increase of about 25%.

The resampling technique is most effective for long integration times, which are feasible for targeted searches like the search for gravitational waves from the Crab pulsar [32]. The computational savings can be used to search over wider parameter spaces like more spin-down parameters or to search over binary systems.

ACKNOWLEDGMENTS

We would like to thank the membership of the LIGO Virgo Collaboration's continuous waves group, especially Stefano Braccini, Vladimir Dergachev, Greg Mendell, Chris Messenger, Maria Alessandra Papa, and Reinhard Prix for helpful discussions. We are also grateful to Patrick Brady and Alan Weinstein for useful conversations. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement No. PHY-0757058. This paper has LIGO Document No. LIGO-P0900301-v1. Xavier Siemens is supported in part by NSF Grant No. PHY-0758155 and the Research Growth Initiative at the University of Wisconsin—Milwaukee.

-
- [1] L. Bildsten, *Astrophys. J.* **501**, L89 (1998).
 - [2] N. Andersson, K.D. Kokkotas, and N. Stergioulas, *Astrophys. J.* **516**, 307 (1999).
 - [3] G. Ushomirsky, C. Cutler, and L. Bildsten, *Mon. Not. R. Astron. Soc.* **319**, 902 (2000).
 - [4] C. Cutler, *Phys. Rev. D* **66**, 084025 (2002).
 - [5] A. Melatos and D.J.B. Payne, *Astrophys. J.* **623**, 1044 (2005).
 - [6] B.J. Owen, *Phys. Rev. Lett.* **95**, 211101 (2005).
 - [7] T. Gold, *Nature (London)* **218**, 731 (1968).
 - [8] F. Pacini, *Nature (London)* **219**, 145 (1968).
 - [9] A. Hewish, S.J. Bell, J.D.H. Pilkington, P.F. Scott, and R.A. Collins, *Nature (London)* **217**, 709 (1968).
 - [10] R.N. Manchester, G.B. Hobbs, A. Teoh, and M. Hobbs, The ATNF Pulsar Catalogue, <http://www.atnf.csiro.au/research/pulsar/psrcat/>.
 - [11] D. Lorimer, *Living Rev. Relativity* **8**, 7 (2005), <http://relativity.livingreviews.org/Articles/lrr-2005-7/>.
 - [12] F. Dubath and J.V. Rocha, *Phys. Rev. D* **76**, 024001 (2007).
 - [13] M.R. DePies and C.J. Hogan, [arXiv:0904.1052](https://arxiv.org/abs/0904.1052).
 - [14] P. Jaranowski, A. Krolak, and B.F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
 - [15] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 082004 (2004).
 - [16] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Phys. Rev. D* **77**, 022001 (2008).
 - [17] P.R. Brady, T. Creighton, C. Cutler, and B.F. Schutz, *Phys. Rev. D* **57**, 2101 (1998).
 - [18] H.J. Pletsch and B. Allen, *Phys. Rev. Lett.* **103**, 181102 (2009).
 - [19] P. Astone, K.M. Borkowski, P. Jaranowski, and A. Królak, *Phys. Rev. D* **65**, 042003 (2002).
 - [20] P. Astone *et al.*, *Classical Quantum Gravity* **20**, S665 (2003).
 - [21] A. Królak, M. Tinto, and M. Vallisneri, *Phys. Rev. D* **70**, 022003 (2004).
 - [22] P. Astone *et al.*, *Classical Quantum Gravity* **25**, 184012 (2008).

- (2008).
- [23] O. Torre, Ph.D. thesis, University of Pisa [LSC-Virgo Technical Report No. VIR-0013A-10, 2009].
 - [24] S. Braccini *et al.*, “Geometrical signal resampling in continuous gravitational wave search to compensate for Doppler and Spin Down effect,” *Classical Quantum Gravity* (to be published).
 - [25] B.F. Schutz, in *The Detection of Gravitational Waves*, edited by D.G. Blair (Cambridge University Press, Cambridge, 1991), Chap. 16, p. 406.
 - [26] R. Prix and B. Krishnan, *Classical Quantum Gravity* **26**, 204013 (2009).
 - [27] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **79**, 022001 (2009).
 - [28] P.R. Williams and B.F. Schutz, [arXiv:9912029v1](https://arxiv.org/abs/9912029v1).
 - [29] R. Prix, LIGO Document T0900149-v1, <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=1665>.
 - [30] ComputeFStatistic_v2, LIGO Analysis Library (LAL), <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.
 - [31] ComputeFStatistic_resamp, LAL, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lal.html>.
 - [32] B. Abbott *et al.* (The LIGO Scientific Collaboration), *Astrophys. J.* **683**, L45 (2008).