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

Pinkesh Patel,[1] Xavier Siemens,[2] and Rejean Dupuis

[1] *California Institute of Technology*[*]
[2] *University of Wisconsin Milwaukee*[†]

(Dated: January 13, 2009)

We describe an efficient implementation of a coherent statistic for continuous gravitational wave searches from neutron stars.

PACS numbers:

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 called a pulsar [7, 8]. Since the first discovery [9], around 1700 pulsars have been detected [10–12].

Due to magnetic dipole 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].

There are a number of techniques available for pulsar searches. These techniques can be loosely divided into two categories: (1) coherent methods [13, 14], which keep track of the phase of the gravitational wave signals over long periods of time, and (2) semi-coherent methods [15] 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 gravitational waves is relatively straightforward [14]. If a signal is present 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 grows linearly while the noise follows a random walk. Thus, with enough data, it is possible to recover any signal out of the noise.

If certain parameters of the signal (sky location, frequency, etc.) are not known the search becomes much more involved. The reason is that the number of points needed cover the search parameter space (and ensure no signals are lost) grows like a large power of the amount of data used. This makes the sensitivity of our 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 increased sensitivity of our searches.

The most promising method for blind searches involves a hierarchical combination of analysis techniques. In particular a coherent step followed by an semi-coherent step such as the Hough transform or stack-slide [15] appears to be the most promising.

In this paper we focus on the efficient implementation of coherent techniques. The sensitivity of searches for continuous waves is computationally bound, and increasing the speed of the search is an economical way to improve search sensitivities. In particular we show how to compute the so-called $\mathcal{F}$-statistic [13], the logarithm of the likelihood function maximized over the intrinsic (and unknown) parameters of the gravitational wave produced by a neutron star, using barycentric resampling and address several important technical issues essential to the computation.

II. THE SIGNAL

The strain produced by a gravitational-wave in Power-recycled Fabri-Perot Michelson interferometer can be represented as

$$h(t_d) = F_+(t_d)h_+(t_d) + F_\times(t_d)h_\times(t_d), \qquad (1)$$

[*] ppatel@ligo.caltech.edu
[†] siemens@phys.uwm.edu

where, h_+ and $h_\times$ are the 'plus' and 'cross' polarizations of a gravitational wave, and $F_+(t_d)$ and $F_\times(t_d)$ are the beam-pattern functions of the detector and are specified by the detector's location, the polarization angle, source location and the angle between the arms of the detector. The detector response may also be expressed through functions $a(t_d)$ and $b(t_d$ as follows

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

and

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

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

The phase of the gravitational wave is given by

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

where Φ_0 is the phase at the start time, $f_0^{(k)}$ is the k^{th} derivative of the frequency and $\mathbf{n}_0$ 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.

We can write $\Psi(t_d) = \Phi_0(t_d) + \Phi(t_d)$ and we can also rewrite

$$\Phi(t_d) = 2\pi f[t_d + \Phi_m(t_d;\alpha,\delta) + \Phi_s(t_d; f^{(k)},\alpha,\delta)] \quad (5)$$

The maximum likelihood statistic for the detection of pulsar signals is called the $\mathcal{F}$-statistic and as referenced in the JKS paper, the $\mathcal{F}$-statistic is given by

$$\mathcal{F} = \frac{4}{S_h(f)T_0}\frac{B|F_a|^2 + A|F_b|^2 - 2CR(F_aF_b^*)}{D}. \quad (6)$$

The A, B, C, and D are given by

$$A = (a\|a); B = (b\|b); C = (a\|b); D = A\cdot B - C^2 \quad (7)$$

with

$$(x\|y) = \frac{2}{T_0}\int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} x(t_d)y(t_d)dt_d \quad (8)$$

and

$$F_a(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} x(t_d)a(t_d)e^{-i\Phi(t_d)}dt_d, \quad (9)$$

and

$$F_b(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} x(t_d)b(t_d)e^{-i\Phi(t_d)}dt_d, \quad (10)$$

If we define a new time variable called t_b as follows

$$t_b = t_d + \Phi_m \quad (11)$$

since $\delta t_b \approx \delta t_d$, to very good accuracy, we can rewrite the equations for F_a and F_b as

$$F_a(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} x(t_b)a(t_b)e^{-2\pi i f t_b}e^{i\Phi_s(t_b)}dt_b, \quad (12)$$

and

$$F_b(f) = \int_{\frac{-T_0}{2}}^{\frac{T_0}{2}} x(t_b)b(t_b)e^{-2\pi i f t_b}e^{i\Phi_s(t_b)}dt_b \quad (13)$$

which are just the Fourier transform of the resampled data and the detector response, multiplied by a phase $e^{i\Phi_s(t_b)}$ [13].

III. IMPLEMENTATION OF BARYCENTRIC RESAMPLING

Modern scientific computers have memories which range from 1 Giga-Bytes to a few Giga-Bytes. A practical way of using such computers for scientific computing is to combine them in the form a Beowulf cluster. These clusters typically are made up of hundreds to thousands of such computers. Gravitational-Wave detectors collect data at the rate of about 16-20 KHz and typically have data-sets spanning years. This roughly converts to a few Tera-Bytes of data for any given analysis. It is important to break the data down into parts the fit the in the memory of a single computer.

The method we use to break up the data into small and processable pieces is by processing sub-bands one at a time. We heterodyne and downsample the data such that the frequencies of interest are brought to near DC, and analyse only a small sub-band. Parceling out all the data this way allows it to be ananlysed in paralell by a Beowulf cluster.

Gravitational-wave detectors are subject to many sources of noise; some of which change daily or even hourly, such as wind, microseism, earthquakes, antropogenic noise etc. These change the noise floor of any analysis as a function of time. Continuous wave searches are sensitive to such changes in the noise floor. A practical way to deal with non-stationarity is to divide the data into short segments within which the noise is stationary and estimate the noise for each segment. The length of the segments can be decided based on the particular interferometer and noise sources which are prominent. For the LIGO detectors 1800 seconds is the usual choice, which is sufficient to track non-stationarities in the data.

A. Time Domain Analysis

The $\mathcal{F}$-statistic can be calculated from a time series directly by following the steps outlined in te background

section. But this is impractical due to the size of the data-sets involved in any real analysis. One way to solve this problem is by dividing the data into band limited, heterodyned time series.

1. Heterodyning

Consider a Fourier transform

$$\int_{-\infty}^{\infty} X(t) \cdot e^{-2\pi f t} = \tilde{X}(f), \tag{14}$$

where $X(t)$ is an arbitrary function. If we multiply $X(t)$ by $e^{2\pi i f_0 t}$, then we get the following expression

$$\begin{aligned}\int_{-\infty}^{\infty} X(t) \cdot e^{-2\pi i f t} \cdot e^{2\pi i f_0 t} dt &= \int_{-\infty}^{\infty} X(t) \cdot e^{-2\pi i (f-f_0) t} dt \\ &= \tilde{X}(f - f_0). \end{aligned} \tag{15}$$

Thus by multiplying by a phase, we can shift the frequency of interest arbitrarily. In practice one shifts the middle of the band of interest to D.C and then down-samples the data. This ensures that down-sampling is possible even for very large frequencies. This procedure is called Heterodyning.

If one is interested in analyzing data from say 990 Hz to 1000 Hz, a minimum sampling rate of 2 KHz is required. But this high sampling rate is only needed if one needs to analyze the whole band from D.C to 1 KHz. If instead the frequencies of interest lie in a 10 Hz band starting at 990 Hz, we can heterodyne the data such that 995 Hz moves to D.C and thus moving 990 Hz to -5 Hz and 1000 Hz to +5 Hz. This heterodyning can be achieved in the time domain by multiplying by the phase factor $e^{-2\pi 995 t}$, where t is measured in units of seconds.

In the frequency domain, heterodyning can be achieved by just relabelling the various frequency bins. In the example above, we can just internally change the labels of the 995 Hz frequency bin to be 0 Hz and 100 Hz to be 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 +5Hz containing all the relevant information. All the other frequencies can then be ignored, since they are of no consequence in our analysis.

2. Low Pass Filtering and Downsampling

After heterodyning the middle of the band to D.C, it is imperative to downsample the data, since it will reduce the size of the data set. But before downsampling, a low pass filter must be applied with a sharp fall off near and around the new nyquist frequency. The new nyquist frequency is given by

$$Ny_{new} = Ny_{old}/D \tag{16}$$

where D is the downsampling factor. A low pass filter eliminates the problems of aliasing, since after the filter is applied, the time series will be nyquist limited. Any downsampling technique will suffice for such a nyquist limited time series. A simple but effective technique is picking every D^{th} point.

3. Resampling

The heterodyned, band limited, downsampled time series will be a complex time series, which we can call $x(t_d^k)$, with

$$t_d^{k+1} - t_d^k = dt = \frac{1}{\Delta F} \tag{17}$$

where t_d^k is the k^{th} datum as measured in the detector's frame of reference and ΔF is the size of the band used for the analysis. Also

$$\frac{\Delta F}{2} = Ny_{new} \tag{18}$$

This time series contains all the information in the band of interest spanning the time of the analysis. Thus we can calculate the $\mathcal{F}$-statistic as laid out in 6 using discrete versions of equations 12 and 13. The discrete versions for a time series with N points are -

$$F_a(f) = \sum_{k=1}^{N} x(t_b^k) a(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} dt_b, \tag{19}$$

and

$$F_b(f) = \sum_{k=1}^{N} x(t_b^k) b(t_b^k) e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} dt_b. \tag{20}$$

Here t_b^k is the k^{th} datum in the time series as measured in the barycentric frame. Since we know the relationship between

$$t_b^k = G(t_d^k, \alpha, \delta, f^{(k)}) \tag{21}$$

where G is a function linear in t_d^k. This relationship between t_d^k and t_b^k can be used to calculate $x(t_b^k)$ from the time series, which is $x(t_d^k)$.

The way to accomplish this during computation, is to create an array of times which are linearly spaced in the barycentric time and the use interpolation to calculate the value of $x(t_b^k)$. The algorithm is as follows

1. Start with $x(t_d^k)$ with t_d^k linearly spaced in time, in the detector's frame of reference.

2. These $x(t_d^k)$ correspond to datums non-linearly spaced in the barycentric frame. Calculate $t_b^k(t_d^k)$.

3. Start at t_b^0 and using interpolation calculate $x(t_b^0 + k \cdot dt)$, which is the $x(t_b^k)$ used in equations 19 and 20.

4. Similarly from $a(t_d^k)$ and $b(t_d^k)$ calculate $a(t_b^k)$ and $b(t_b^k$ respectively.

5. Evaluate equations 19 and 20 to calcualte $F_a(f)$ and $F_b(f)$.

6. Use equation 6 to calculate the $\mathcal{F}$-statistic.

4. Heterodyne Correction

For any practical search for gravitational waves, heterodyning is an essential component of the analysis. Heterodyning is accomplished by multiplying the data stream with $e^{2*\pi i f_h t_d}$, which has an effect of shifting the frequency f_h to D.C or 0. All other frequencies follow this linear shift as is shown by 15. Since barycentric resampling relies on a wholesale change in the reference frame, any t_d functionality in the time series will interfere with the recovery of all the $\mathcal{F}$-statistic. Thus due to heterodyning, equations 19 and 20 become

$$F_a(f) = \sum_{k=1}^{N} x(t_b^k)a(t_b^k)e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \cdot e^{2\pi i f_h t_d^k}, \quad (22)$$

and

$$F_b(f) = \sum_{k=1}^{N} x(t_b^k)b(t_b^k)e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \cdot e^{2\pi i f_h t_d^k}. \quad (23)$$

The effect desired from the process of heterodyning is to shift the frequencies of interest in the barycentric frame, since that is the frame in which we carry out our calculations. But equations 22 and 23 will fail to achieve that goal. In order to correct this effect due to heterodyning, we need to multiply the data by $e^{2\pi i f_h \Phi_m}$, where Φ_m is as defined in equation 11. Thus this makes equations 22 and 23,

$$F_a(f) = \sum_{k=1}^{N} x(t_b^k)a(t_b^k)e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \cdot e^{2\pi i f_h t_d^k} e^{2\pi i f_h \Phi_m^k}, \quad (24)$$

and

$$F_b(f) = \sum_{k=1}^{N} x(t_b^k)b(t_b^k)e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} \cdot e^{2\pi i f_h t_d^k} e^{2\pi i f_h \Phi_m^k} \quad (25)$$

But -

$$e^{2\pi i f_h t_d^k} e^{2\pi i f_h \Phi_m} = e^{2\pi i f_h (t_d^k + \Phi_m^k)} = e^{2\pi i f_h t_b^k} \quad (26)$$

Thus -

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

and

$$F_b(f) = \sum_{k=1}^{N} x(t_b^k)b(t_b^k)e^{-2\pi i f t_b^k} e^{i\Phi_s(t_b^k)} . e^{2\pi i f_h t_b^k} = F_b(f - f_h) \quad (28)$$

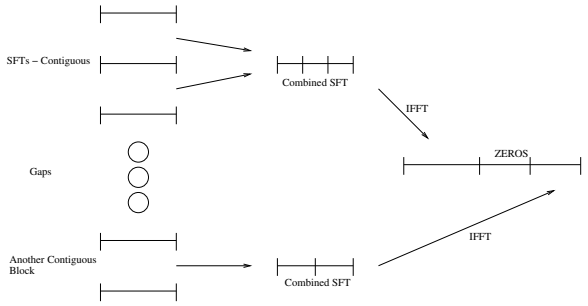

FIG. 1: Pictorial description of data pre-processing

B. Frequency Domain Analysis

A heterodyned, downsampled and low-pass filtered time series can also be obtained in the frequency domain. The frequency domain data can be whitened as a function of time if the data is pre-divided into small chunks each corresponding to a short interval in time. This has the advantage of dealing with the problem of non-stationarity of the detector. But there are subtle issues when dealing with data in the frequency domain. A simple algorithm to produce a time series equivalent to the one used for the time domain analysis is as follows -

1. Divide raw data into small chunks of time.

2. Take the fourier transform of each chunk, henceforth referred to as Short Fourier Transforms (SFTs).

3. Identify contiguous SFTs.

4. Combine each contiguous chunk of SFTs into one long time baseline SFT.

5. Pick required frequencies from the long time baseline SFT.

6. Heterodyne, Low-Pass filter and Down-Sample these frequencies.

7. Inverse Fourier transform to convert to time domain.

8. Stitch all these time domain chunks, filling gaps with zeros.

1. SFTs

Cite S2 paper. SFTs are described there.

2. Combining SFTs

The SFT data, after whitening, can be combined using the Dirichlet kernel to form longer time-baseline Fourier transforms. The procedure is as follows

3. Filtering and Down-Sampling

Since we limit our analysis to a certain band, all the information outside of that band of interest is irrelevant and should be removed in order to keep the minimum amount of information required. After heterodyning the data, the simplest thing one can do to get rid of all the unwanted frequencies is to just ignore them completely. But this means that there will be a sharp edge in the data, which can lead to some erratic behavior at the boundaries. An easy way to overcome this problem is to use a window on the data.

A window is just a series of weights applied to the data. To preserve all the relevant information and to avoid edge effects, a smoothly rising window applied outside the region of interest on either side with a flat top of 1.0 is used asa window. Such a window is called a Tukey window and is a Rectangular window with a split Hann window on either edge.

Multiplying by a window serves the same purpose as a low-pass filter does in the time domain. It supresses all the unwanted frequencies, which are now higher than the ones of interest. The act of ignoring all the remaining frequencies is equivalent to the procedure of down-sampling in the time domain. The down-sampling factor is given by $\frac{Numberoffrequencybinskept}{Totalnumberoffrequencybins}$.

Wonder if we need to show my matlab simulation to prove all this here.

4. Patching

A complex time series is obtained once an Inverse Fast Fourier Transform is taken on the heterodyned, filtered and down-sampled. This complex time series contains all the information in the frequencies of interest for the whole block of contiguous data. A Tukey window is applied to each such time series to bring ends to zero smoothly and to ensure that there is a smooth rise in the beginning. Then each of these time series is patched to the neighboring contiguous block with an appropriate number of zeros representing a gap in the data. This patched up time series is then used to calculate the $\mathcal{F}$-statistic.

C. Interpolation Issue

Ignoring this piece for a while. Need to discuss if before I write it.

IV. RESULTS AND CONCLUSIONS

A. Speed

The current scheme of calculating the $\mathcal{F}$-statistic involves using the Dirichlet kernel to patch up a series of short Fourier transforms (SFTs), which are calculated for 30 minutes of data taken at 16 kHz. The 30 minute window is set by the maximum Doppler shift due to the motion of the Earth. This algorithm uses a code called ComputeLALDemod to calculate the $\mathcal{F}$-statistic.

There are nested loops over the sky location , α and δ, the frequency f and the spin-downs $f^{(k)}$, a loop over the number of SFTs and a loop to calculate the Dirichlet kernel in the calculation of the $\mathcal{F}$-statistic in ComputeLALDemod . At this point in the analysis, we can ignore the loops over α, δ and $f^{(k)}$, since they will be similar in both the methods. 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 $N_{\rm ops}$. If the number of Dirichlet kernel points used are $N_{\rm D\_Ker}$, then the total number of operations used by ComputeLALDemod are -

$$N_{\rm Tot}^{\rm Demod} = N_{\rm ops} \cdot N_{\rm D\_Ker} \cdot N_{\rm SFTs} \cdot N \tag{29}$$

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

Comparing 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. Interpolate and calculate the beam patterns.
4. Take the Fourier Transform.

Each of these steps involves operations of order 10, but all of these steps are sequential, therefore they only add and we have say $N_{\rm ops}^{\rm resamp}$ number of operations, which is approximately 30 or so operations per data point. The last step is the Fourier transform, which is of order $N \log N$, therefore total number of steps are -

$$N_{\rm Tot}^{\rm Resamp} = (N_{\rm ops}^{\rm Resamp} + \log N) \cdot N \tag{30}$$

Therefore, the ratio of operations between the two methods is

$$\frac{N_{\rm Tot}^{\rm Demod}}{N_{\rm Tot}^{\rm Resamp}} = \frac{N_{\rm ops} \cdot N_{\rm D\_Ker} \cdot N_{\rm SFTs}}{N_{\rm ops}^{\rm Resamp} + \log N} \tag{31}$$

To first order,

$$\frac{N_{\rm Tot}^{\rm Demod}}{N_{\rm Tot}^{\rm Resamp}} \approx \frac{N_{\rm SFTs}}{\log N} \tag{32}$$

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.

[1] L. Bildsten, Gravitational radiation and rotation of accreting neutron stars. *Astrophys. J.* **501** L89, 1998.

[2] N. Andersson, K. D. Kokkotas, and N. Stergioulas, On the relevance of the r-mode instability for accreting neutron stars and white dwarfs. *Astrophys. J.* **516** 307, 1999.

[3] G. Ushomirsky, C. Cutler, and L. Bildsten, Deformations of accreting neutron star crusts and gravitational wave emission. *Mon. Not. Roy. Astron. Soc.* **319** 902, 2000.

[4] C. Cutler, Gravitational waves from neutron stars with large toroidal B-fields. *Phys. Rev.* **D66** 084025, 2002.

[5] A. Melatos and D. J. B. Payne, Gravitational radiation from an accreting millisecond pulsar with a magnetically confined mountain. *Astrophys. J.* **623** 1044, 2005.

[6] B. J. Owen, Maximum elastic deformations of compact stars with exotic equations of state. *Phys. Rev. Lett.* **95** 211101, 2005.

[7] T. Gold, Rotating neutron stars as the origin of the pulsating radio sources. *Nature* **218** 731, 1968.

[8] F. Pacini, Rotating neutron stars, pulsars, and supernova remnants. *Nature* **219** 145, 1968.

[9] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, Observation of a rapidly pulsating radio source. *Nature* **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] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The ATNF Pulsar Catalogue 2004.

[12] D. Lorimer, Binary and Millisecond Pulsars. http://relativity.livingreviews.org/Articles/lrr-2005-7/.

[13] P. Jaranowski, A. Krolak, and B. F. Schutz, Data analysis of gravitational-wave signals from spinning neutron stars. I: The signal and its detection. *Phys. Rev.* **D58** 063001, 1998.

[14] B. Abbott et al., Setting upper limits on the strength of periodic gravitational waves using the first science data from the GEO 600 and LIGO detectors. *Phys. Rev.* **D69** 082004, 2004.

[15] B. Abbott et al., All-sky search for periodic gravitational waves in LIGO S4 data 2007.