



# Implementation of an Efficient Bayesian Search for Gravitational-wave Bursts with Memory in Pulsar Timing Array Data

Jerry Sun<sup>1</sup>, Paul T. Baker<sup>2</sup>, Aaron D. Johnson<sup>3</sup>, Dustin R. Madison<sup>4</sup>, and Xavier Siemens<sup>1</sup>

<sup>1</sup>Department of Physics, Oregon State University, Corvallis, OR 97331, USA; [sun.jerry.1996@gmail.com](mailto:sun.jerry.1996@gmail.com)

<sup>2</sup>Department of Physics and Astronomy, Widener University, Chester, PA 19013, USA

<sup>3</sup>Center for Gravitation, Cosmology and Astrophysics, University of Wisconsin–Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA

<sup>4</sup>Department of Physics, University of the Pacific, 3601 Pacific Avenue, Stockton, CA 95211, USA

Received 2022 September 30; revised 2023 April 24; accepted 2023 May 3; published 2023 July 7

## Abstract

The standard Bayesian technique for searching pulsar timing data for gravitational-wave bursts with memory (BWMs) using Markov Chain Monte Carlo (MCMC) sampling is very computationally expensive to perform. In this paper, we explain the implementation of an efficient Bayesian technique for searching for BWMs. This technique makes use of the fact that the signal model for Earth-term BWMs (BWMs passing over the Earth) is fully factorizable. We estimate that this implementation reduces the computational complexity by a factor of 100. We also demonstrate that this technique gives upper limits consistent with published results using the standard Bayesian technique, and may be used to perform all of the same analyses of BWMs that standard MCMC techniques can perform.

*Unified Astronomy Thesaurus concepts:* [Gravitational waves \(678\)](#); [General relativity \(641\)](#); [Gravitational wave detectors \(676\)](#); [Gravitational wave astronomy \(675\)](#)

## 1. Introduction

Millisecond pulsars (MSPs) have very stable rotations. Because the rotational period is so stable, it is possible to detect small deviations in the times of arrival (TOAs) of radio pulses from an array of these pulsars caused by gravitational waves (GWs) passing between the pulsar and radio observatories on Earth (Hellings & Downs 1983; Foster & Backer 1990; Manchester 2013; McLaughlin 2013). Pulsar timing arrays (PTAs) are expected to be able to use TOA data from many MSPs to either detect or provide constraints on GWs (Sazhin 1978; Detweiler 1979).

One signal of interest is a GW burst with memory (GW BWM). “Memory” is a permanent change in the spacetime metric that remains after a GW passes through a region of space arising from the nonlinearity of Einstein’s field equations (Christodoulou 1991; Thorne 1992). In particular, it is expected that mergers of supermassive black hole binaries will leave behind detectable memory. Detections of (or constraints on) the rates of BWM events would allow for a better understanding of the rates at which these events occur in the universe (e.g., Islo et al. 2019). Additionally, because all GW events leave behind GW memory, detections of BWMs could lead to discoveries of new sources of GWs (Cutler et al. 2014).

A GW BWM passing over an Earth–pulsar pair will shift the pulsar’s observed rotational frequency (e.g., van Haasteren & Levin 2010). This shift causes a difference between the observed frequency of the pulsar and the frequency expected from the timing model fit, and will therefore contribute to a potentially detectable signal in the pulsar’s TOAs (Seto 2009; Pshirkov et al. 2010; van Haasteren & Levin 2010; Cordes & Jenet 2012; Madison et al. 2014; Islo et al. 2019). The observed rotational frequency may change to be either faster or slower,

depending on the orientation and polarization of the memory wave front, which determines the sign of the memory.

In this paper, we will discuss the adaptation of analysis techniques used in the NANOGrav 5 yr search for GW BWMs (Arzoumanian et al. 2015; hereafter, [NG5–bwm](#)) to expedite the Bayesian methods used in the NANOGrav 11 yr search for BWMs (Aggarwal et al. 2020; hereafter, [NG11–bwm](#)). This search was performed on the NANOGrav 11 yr data set (Arzoumanian et al. 2018). In [NG11–bwm](#), no detection of GW BWMs was reported. Thus, the authors presented Earth-term upper limits (ULs) as a function of burst epoch and sky location (among other results, but these will be our focus). Our goal in this paper is to show that the adapted techniques from [NG5–bwm](#) may be used to efficiently perform Bayesian analyses comparable to those in [NG11–bwm](#) with a similar degree of accuracy.

In Section 2, we describe the effect of a BWM on the TOA residuals of a pulsar. In Section 3, we discuss the current standard Bayesian approach to searching for GW BWMs in PTA data and the implementation of an efficient technique for speeding up this search. In Section 4, we compare the results of UL calculations using our more efficient technique against results previously published in the literature. We also discuss the improvements in computational efficiency that come from this technique.

## 2. Signal and Data Model

The rise time for the memory component of a GW BWM is much shorter than the typical observing cadence of PTAs; thus, we may ignore it and consider the frequency-shifting effect to be instantaneous (Favata 2010; van Haasteren & Levin 2010; Madison et al. 2014). This manifests as a linear “ramp” in the residuals, since a constant excess or deficit of pulse phase will accrue with each rotation of the pulsar. Consider a memory event from a source at  $(\theta, \phi)$ . The event wave front propagates in the direction  $\hat{k}$ , with strain  $h_0$ , passing over the Earth, from

which we observe a pulsar located at position  $\mathbf{p}$ . The memory wave front has the principal polarization vector  $\hat{\psi}$ , described by an angle  $\psi$ , which gives the principal polarization direction of the wave front relative to an orthonormal basis  $(\hat{\delta}, \hat{\beta})$ . In particular, the principal polarization vector  $\hat{\psi}$  is defined

$$\hat{\psi} = \hat{\delta} \cos \psi + \hat{\beta} \sin \psi. \quad (1)$$

Following `NG5-bwm` and `NG11-bwm`, the perturbation to pulse TOAs from this pulsar,  $\delta t_{\text{bwm}}$ , may be modeled as

$$\delta t_{\text{bwm}}(t) = B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) h_{\text{mem}}(t), \quad (2)$$

where  $h_{\text{mem}}(t)$  is the time-dependent strain of the memory wave front, and the geometric factor  $B$  accounts for the relative orientation of the source and pulsar (Estabrook & Wahlquist 1975; Hellings & Downs 1983). The geometric projection factor is

$$B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{1}{2} \cos(2\psi_{\hat{\mathbf{p}}}) (1 - \cos \alpha), \quad (3)$$

where  $\alpha$  is the angle between  $\hat{\mathbf{p}}$  and  $\hat{\mathbf{k}}$  (pulsar location and propagation direction, respectively) and the angle  $\psi_{\hat{\mathbf{p}}}$  is defined to be the angle between the principal polarization vector and the projection of the pulsar line of sight onto the  $(\hat{\delta}, \hat{\beta})$  plane:

$$\alpha = \cos^{-1}(\hat{\mathbf{p}} \cdot \hat{\mathbf{k}}), \quad (4)$$

$$\psi_{\hat{\mathbf{p}}} = \tan^{-1} \left( \frac{\hat{\mathbf{p}} \cdot \hat{\beta}}{\hat{\mathbf{p}} \cdot \hat{\delta}} \right) - \psi. \quad (5)$$

We include this description for completeness, and a diagram with more details may be found in Madison et al. (2014). For this analysis, we simply characterize the polarization using the principal polarization angle  $\psi$ . The time-dependent strain term is (Pshirkov et al. 2010; van Haasteren & Levin 2010)

$$h_{\text{mem}}(t) = h_0 [(t - t_0)\Theta(t - t_0) - (t - t_p)\Theta(t - t_p)], \quad (6)$$

where  $h_0$  is the strain of the memory;  $t_0$  is the time that the memory wave front passed over the Earth; and  $t_p = t_0 + (|\mathbf{p}|/c) [1 + \cos(\alpha)]$  is the time at which the memory wave front passed over the pulsar, with  $\mathbf{p}$  and  $\alpha$  still defined to be the position of the pulsar and the angle between the pulsar and BWM propagation direction, as in Equation (3). Additionally,  $\Theta(t)$  is the Heaviside function. Because each pulsar in NANOGrav’s 11 yr data release (and, more generally, current PTA data sets) is on the order of thousands of light-years away from the Earth, and total observation times are on the order of tens of years, we only expect that one of the two terms in Equation (6) will be nonzero. The first term in Equation (6) is called the “Earth term,” and the second is called the “pulsar term.” A BWM may be observed either when it passes over a single pulsar, or when it passes over the Earth. In the former case, we will see the frequency of a single pulsar spontaneously change. In the latter case, we expect to see the rotational frequency of each pulsar change simultaneously with a characteristic quadrupolar amplitude pattern. In either case, the time at which

the BWM wave front causes an apparent rotational frequency change is defined as the burst epoch.

This signal model is implemented in `enterprise_extensions`<sup>5</sup> (Taylor et al. 2021).

### 3. Methodology

We begin by discussing the standard Bayesian approach to searching for a GW BWM. This discussion will summarize the approach taken in `NG11-bwm`, which searched the NANOGrav 11 yr data set for GW BWMs. Then, we will discuss the adaptation of the techniques used in `NG5-bwm` that expedites both the pulsar- and Earth-term searches.

#### 3.1. Bayesian Approach

`NG11-bwm` modeled the timing residuals  $\delta \mathbf{t}$  of a pulsar as

$$\delta \mathbf{t} = \mathbf{s} + \mathbf{T}\mathbf{b} + \mathbf{n}, \quad (7)$$

where  $\delta \mathbf{t}$  are the remaining perturbations to the TOAs from a pulsar after fitting parameters in the pulsar’s timing model using a general least-squares fit (Arzoumanian et al. 2018). These remaining perturbations, the timing residuals, are expected to originate from a combination of noise processes, errors in the timing model fit, and GW signals.  $\mathbf{s}$  are the contributions to the timing residuals from a GW BWM.  $\mathbf{T}\mathbf{b}$  are the contributions to the residuals from any Gaussian processes. In this paper, we consider two different Gaussian processes, and so our  $\mathbf{T}$ -matrix and  $\mathbf{b}$  vector may be broken down into

$$\mathbf{T} = [\mathbf{M} \ \mathbf{F}], \quad \mathbf{b} = \begin{bmatrix} \boldsymbol{\epsilon} \\ \mathbf{a} \end{bmatrix},$$

where  $\mathbf{M}$  is the design matrix for the linearized timing model that accounts for uncertainty in the residuals from an imperfect timing model fit  $\boldsymbol{\epsilon}$ .  $\mathbf{F}$  is the design matrix for pulsar-intrinsic red noise, modeled as a Fourier series with coefficients  $\mathbf{a}$ . Finally, the elements of vector  $\mathbf{n}$  are Gaussian white-noise uncertainties in the observed TOAs.

The red-noise spectrum, for example from a stochastic background of GWs, is expected to behave as a power law (Phinney 2001):

$$P(f_j) = A_j^2 \left( \frac{f_j}{\text{yr}^{-1}} \right)^{-\gamma}, \quad (8)$$

where  $P(f_j)$  is the power spectral density of the red-noise process,  $A_j$  is the characteristic amplitude of the red-noise process in the  $j$ th frequency bin, using a reference frequency of  $\text{yr}^{-1}$ , and  $\gamma$  is the spectral index of the power law.

From Equation (7), we can construct an approximation of the Gaussian white noise given an estimation of the model parameters:

$$\mathbf{n} = \delta \mathbf{t} - \mathbf{s} - \mathbf{T}\mathbf{b}. \quad (9)$$

This is only an approximation of the white noise, since the terms on the right-hand side are estimations. However, if the white noise is expected to be Gaussian, we can write the probability of observing this particular series of white-noise

<sup>5</sup> [https://github.com/nanograv/enterprise\\_extensions](https://github.com/nanograv/enterprise_extensions)

residuals as:

$$p(\mathbf{n}) = \frac{\exp(-\frac{1}{2}\mathbf{n}^T N^{-1}\mathbf{n})}{\sqrt{2\pi \det N}}, \quad (10)$$

where  $N$  is a covariance matrix of white-noise uncertainties in each observed TOA, and  $\mathbf{n}^T$  is the transpose of  $\mathbf{n}$ .

Then, the likelihood of a BWM signal in the pulsar timing residuals is equivalent to the likelihood that the remaining residuals after subtracting out deterministic effects are Gaussian white noise. In other words,

$$p(\delta\mathbf{t}|\mathbf{b}, s) = p(\delta\mathbf{t}|\mathbf{b}, h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{\exp[-\frac{1}{2}(\delta\mathbf{t} - s - \mathbf{Tb})^T N^{-1}(\delta\mathbf{t} - s - \mathbf{Tb})]}{\sqrt{2\pi \det N}}, \quad (11)$$

where we have explicitly written out the parameters that determine  $s$ . This parameter space, when including each Fourier coefficient and timing model parameter, is very high-dimensional.

It is possible to analytically marginalize the likelihood in Equation (11) over the parameters that describe the Gaussian processes and reduce the dimensionality of the parameter space (Lentati et al. 2013; van Haasteren & Vallisneri 2014, 2015). The reduced likelihood is

$$p(\delta\mathbf{t}|h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{\exp(-\frac{1}{2}\mathbf{q}^T C^{-1}\mathbf{q})}{\sqrt{2\pi \det C}}, \quad (12)$$

where

$$\mathbf{q} = \delta\mathbf{t} - s, \quad (13)$$

where  $C$  is defined as

$$C = N + \mathbf{TDT}^T \quad (14)$$

and  $D$  is defined as

$$D = \begin{bmatrix} \infty & 0 \\ 0 & \phi \end{bmatrix}.$$

where  $\infty$  is a diagonal matrix of infinities, which effectively gives unconstrained priors on the timing model parameters, and  $\phi$  is a diagonal matrix containing the red noise power at each frequency bin in Equation (8). The Woodbury (Woodbury 1950) matrix identity is used to evaluate  $C^{-1}$  efficiently. The  $D$  matrix also only appears as an inverse in this identity. Thus, in practice, the diagonal matrix of infinities only appears as a matrix of zeros in the likelihood calculation. This likelihood is implemented in the ENTERPRISE<sup>6</sup> (Ellis et al. 2020) pulsar-timing GW analysis software package.

Now that the likelihood has been constructed, samples from the posterior distributions are drawn using the Markov Chain Monte Carlo (MCMC) sampler implemented in the PTMCMCSAMPLER<sup>7</sup> package (Ellis & Haasteren 2017).

Great care must be taken when computing ULs over the sky because of a strong selection bias. If there is no support for a signal in the data, then the maximum posterior probability will be determined largely by the prior. Because our amplitude prior

spans many orders of magnitude, there is much more prior volume at higher amplitudes. This means the posterior will be maximized for bursts with very large amplitudes at insensitive areas of the sky. Because the burst is placed at an insensitive area of the sky, the data cannot exclude this strong signal. This will cause the one-dimensional marginal posterior to be biased toward very high amplitude for combinations of burst epochs, sky locations, and polarization angles where the PTA has low sensitivity. This would not fairly represent the sensitivity of the Earth-term search (NG11-bwm).

To remedy this, NG11-bwm sampled individual ‘‘source-orientation’’ bins, in which the burst epoch, sky location, and polarization are all fixed. Then, a full Earth-term posterior is constructed by concatenating an equal number of samples from each source-orientation bin. This sampling scheme is the equivalent of implementing a prior that exactly cancels the selection effect, resulting in a posterior that is uniform in source-orientation. This is related to the technique used in Malmquist (1922).

More specifically, to place an amplitude UL as a function of burst epochs, NG11-bwm created 48 HEALPIX<sup>8</sup> (Gorski et al. 2005) sky bins using healpy<sup>9</sup> (Zonca et al. 2019), with eight polarizations in each sky bin. This gives a total of 384 source-orientation bins in each of the 40 burst epoch bins. An MCMC sampler is then used to sample the posterior probability distributions of the BWM amplitude. Then, to compute an amplitude posterior marginalized over source-orientations for a fixed burst epoch, equal numbers of samples are taken from each source-orientation bin and concatenated.

To place ULs as a function of sky position, NG11-bwm used 768 HEALPIX sky bins and directly sampled polarization (rather than sampling in fixed polarization bins). Then, the amplitude UL may be computed from the marginalized amplitude posterior for each sky position.

For a summary of the priors, see Table 1.

### 3.2. Accelerated Bayesian Search

For the accelerated Bayesian search, we mimic the Bayesian approach described in Section 3.1 as closely as possible. We found that the computational cost of the MCMC sampling required was prohibitively expensive to perform on machines we have access to. Thus, to expedite the Bayesian search, we leverage a fact from NG5-bwm: the Earth-term likelihood is able to be factorized into a product of pulsar-term likelihoods. In other words,

$$p(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B) = \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B) = \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|h_i, t_B), \quad (15)$$

and

$$h_i = B(\hat{\mathbf{k}}, \hat{\mathbf{p}}_i, \psi) \times h_{\text{mem}}, \quad (16)$$

where  $p(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B)$  is the global likelihood of a burst propagating in the direction  $\hat{\mathbf{k}}$ , with polarization  $\psi$ , an Earth-term epoch  $t_B$ , and strain  $h_B$ . Additionally,  $p_i$  is the pulsar-term likelihood of this burst in the  $i$ th pulsar, with  $h_i$  being the

<sup>6</sup> <https://github.com/nanograv/enterprise>

<sup>7</sup> <https://github.com/jellis18/PTMCMCSampler>

<sup>8</sup> <https://healpix.jpl.nasa.gov/>

<sup>9</sup> <https://github.com/healpy/healpy>

**Table 1**  
Priors Used for Each of the Model Parameters in the Bayesian Search for Global Earth-term GW BWMs Using the Full PTA

Parameter	Prior	Description
$\log_{10} A_m$	LinearExp( $-17, -11$ )	Amplitude of intrinsic pulsar red noise
$\gamma_{rn}$	Uniform( $0, 7$ )	Spectral index of intrinsic pulsar red noise
$\log_{10} A_{\text{BWM}}$	LinearExp( $-17, -10$ )	Amplitude of global BWM
$\psi_{\text{BWM}}$	Uniform( $0, \pi$ )	Polarization of BWM
$\theta_{\text{BWM}}$	Uniform( $0, \pi$ )	Polar angle of BWM source
$\phi_{\text{BWM}}$	Uniform( $0, 2\pi$ )	Azimuthal angle of BWM source
$t_{\text{BWM}}$	Uniform(MJD 56000, MJD 57000)	Earth-term BWM epoch
	Uniform(MJD 53216, MJD 57387)	

**Note.** There are a total of five global BWM parameters, as well as two parameters for each pulsar in the PTA. The priors on the logarithm of the amplitude are equivalent to setting uniform priors over the amplitude. Because of selection effects, it is nontrivial to implement uniform priors over the sky location of the burst. More details on this may be found in Section 3.1. The prior on  $t_{\text{BWM}}$  also varies depending on the particular UL calculation. For ULs as a function of sky location, we use priors between MJD 56000 and MJD 57000. For ULs as a function of burst epoch, we use priors that encompass all the timing data (approximately MJD 53216 to MJD 57387). There is more detail on the burst epoch prior in Section 3.2.

observed amplitude of the burst after accounting for the geometric projection,  $B(\hat{\mathbf{k}}, \hat{\mathbf{p}}_i, \psi)$ , of the burst onto the pulsar–Earth line of sight. Finally,  $N_{\text{psr}}$  is the number of pulsars. As pointed out in *NG5–bwm*, the pulsar’s TOAs have no information about the parameters of the burst, other than the apparent burst amplitude after being projected onto the pulsar–Earth line of sight. This allows us to precompute the individual pulsar-term likelihoods over a grid of only post-projection BWM amplitude and burst epoch without losing any information. Then, at run time, the geometric projection factor, Equation (3), may be applied to give the correct post-projection amplitude for any given global trial burst. This way, we may then look up the corresponding likelihoods of the global burst in the precomputed lookup tables and combine them using Equation (15).

With this in mind, we begin the accelerated Bayesian search by first generating five-dimensional lookup tables for the likelihood of each pulsar (the far right-hand side of Equation (15)). In addition to the BWM amplitude  $|h_i|$ , epoch  $t_B$ , and the sign of  $h_i$ , we include the amplitude  $A_m$  and spectral index  $\gamma$  of the red-noise process described in Equation (8). We emphasize that we must keep track of the sign separately, because any trial BWM may delay or advance the TOAs, depending on the relative orientation of the BWM polarization to the pulsar–Earth line of sight. Recall that a single BWM has a quadrupolar antenna pattern. Consider a pulsar that is in a part of the sky such that a particular BWM would cause the TOAs to be advanced by some amount. Rotating the trial burst by  $90^\circ$  would cause the TOAs to be delayed—rather than advanced—by the same amount. Thus, for global searches of BWMs with every possible orientation, we have to include the likelihoods for every amplitude of BWM with both positive and negative signs. We then numerically integrate over the red-noise parameters using the composite Simpson’s rule to obtain one red-noise-marginalized three-dimensional likelihood lookup table for each pulsar (with the remaining parameters being  $\{|h_i|, t_B, \text{sign}(h_i)\}$ ). Then, we may compute marginal amplitude likelihoods for any pulsar-term BWM by integrating over the burst epoch.

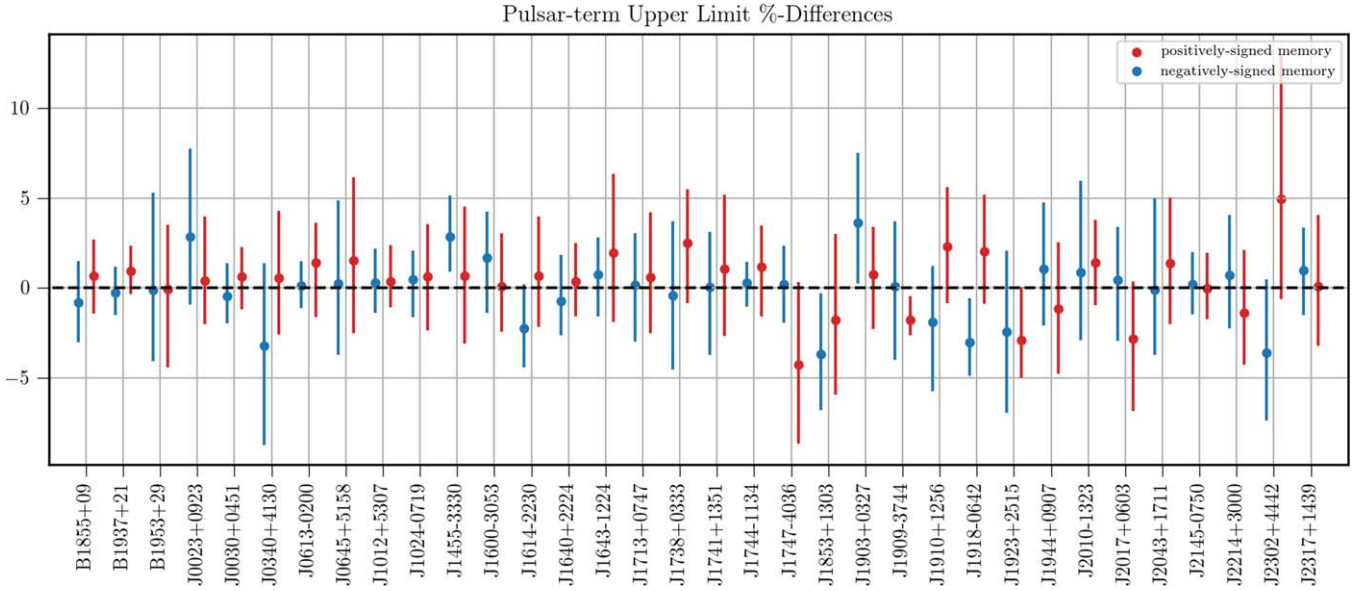
Next, we can combine the pulsar-term likelihood tables to construct global likelihood lookup tables that contain the Bayesian likelihoods of finding an Earth-term trial burst with fixed sky position  $(\theta, \phi)$ , polarization  $\psi$ , and strain  $h_0$  at some

fixed trial burst epoch  $t_0$ . To do so, we project a burst with these fixed global parameters onto each pulsar line of sight to find the amplitude and sign with which this burst will appear in the pulsar’s timing residuals. This allows us to compute the observed amplitude in each of the pulsar terms. We then simply look up the likelihood for each pulsar-specific projected amplitude in the single-pulsar lookup tables. Finally, the global likelihood of this trial burst is computed by multiplying the pulsar-term likelihoods, Equation (15). We compute one two-dimensional lookup table varying over trial bursts characterized by  $(h_0, t_0)$  for each set of trial parameters  $(\theta, \phi, \psi)$ .

We can then construct global amplitude posteriors as a function of sky position and epoch by integrating out any nuisance parameters against their prior distributions. We do so in the same way as in *NG11–bwm* (described in Section 3.1); whenever we marginalize over source-orientation, we are careful to do this by taking equal samples from each source-orientation bin to demand a posterior that is uniform in source-orientation.

Specifically, to compute the ULs as a function of burst epoch, we compute two-dimensional posterior distributions of BWM amplitude and epoch in each of 48 `HEALPIX` sky pixels with one of eight fixed polarizations (for a total of 384 total source-orientation bins). Then, we compute the marginal BWM amplitude posterior for each trial burst epoch. Finally, we concatenate an equal number of samples from each source-orientation bin to compute the full-sky, polarization-marginalized 95% ULs as a function of burst epoch.

To compute the ULs as a function of sky location, we use 768 `HEALPIX` sky pixels and eight polarization bins, with the prior for BWM epochs limited between MJD 56000 and MJD 57000. We use this limited prior because after MJD 56000, there are no new pulsars added to the PTA. It is challenging to come up with a scheme for determining representative BWM amplitude posteriors over a period in which new pulsars are continually added, so we limit our search only to the period in which we already have data for each pulsar. For each source-orientation bin, we then marginalize over burst epochs to obtain the marginal BWM amplitude posterior, and concatenate samples from all eight polarization bins. Finally, we marginalize over polarization by concatenating samples from each polarization bin to obtain the marginal amplitude posterior for each sky location. In very brief summary:



**Figure 1.** Percent difference in the pulsar-term amplitude ULs for each pulsar used in *NG11-bwm*. These percent differences are computed by comparing the 95% ULs from Bayesian MCMC runs and lookup table-based methods for positively and negatively signed memory. Overall, we see good agreement, with percent differences less than 5% and largely consistent with 0%. For this comparison, we limit the search by excluding the first 180 days and last 270 days from the data set. This is because many pulsars have very sparse observations early on. Furthermore, there will be little evidence for a BWM near the end of a data set, since there will not be enough observed TOAs after the trial epoch to accurately detect a BWM. This results in extremely large posterior probabilities for bursts at late times, which heavily reduced the accuracy of our numerical marginalization. The red points are the percent differences for amplitude ULs of positively signed memory, and the blue points are the differences for negatively signed memory, while the error bars show the 95% confidence intervals computed from bootstrapping.

1. Compute pulsar-term BWM likelihoods on a grid of  $\{\log_{10}|h_i|, \text{sign}(h_i), t_0, \log_{10}A_m, \gamma_m\}$ .
2. Marginalize pulsar-term BWM likelihoods over red-noise parameters.
3. Use pulsar-term likelihoods, Equations (15) and (16), to compute Earth-term BWM likelihoods on a grid of  $\{\log_{10}h_0, t_B\}$  for each set of trial burst parameters  $\theta, \phi, \psi$ .
4. Marginalize over:
  - (a) Burst epoch and polarization to compute amplitude posterior over sky location; and
  - (b) Sky location and polarization to compute amplitude posterior over burst epoch.

We computed these global BWM amplitude posteriors using a prior that is log-uniform in the burst amplitude. However, to compute ULs on the burst amplitude, we need to use a posterior with a prior that is *uniform* in the burst amplitude. Although our marginal posteriors have log-uniform priors built in, we can still readjust the prior. Under a log-uniform prior, the burst amplitude posterior is

$$p_{\log\text{-uni}}(A_{\text{bwm}}|d) \propto p(d|A_{\text{bwm}}) \pi_{\log\text{-uni}}(A_{\text{bwm}}) \propto p(d|A_{\text{bwm}}) \frac{1}{A_{\text{bwm}}}, \quad (17)$$

where  $\pi(A_{\text{bwm}}) \propto \frac{1}{A_{\text{bwm}}}$  is the prior distribution on  $A_{\text{bwm}}$ . We can see that multiplying the (log-uniform) posterior by the amplitude will then correctly adjust the prior to have equal volume at each burst amplitude, instead of equal volumes at each order of magnitude of burst amplitude. Once this posterior is recomputed with the correct prior, we can compute the 95% amplitude ULs by numerical integration or rejection sampling.

We would also like to emphasize a new, unique advantage of this accelerated search for BWMs. One challenge of using PTAs to detect GWs is the necessity of accurate, well-

understood pulsar noise models. Because our global likelihood is computed using individual pulsar-term likelihoods, we are free to experiment with different noise models for each pulsar individually. In contrast, using the traditional techniques would require a full recomputation of the Bayesian posteriors using MCMC, even when altering just one of the pulsars' noise models. There has been much work done to improve pulsar noise models, and this factorized approach very robustly allows for adjustments of noise models during analysis, while minimizing the computational cost. Furthermore, this would also allow us to use bespoke noise models for each pulsar, if necessary.

## 4. Results

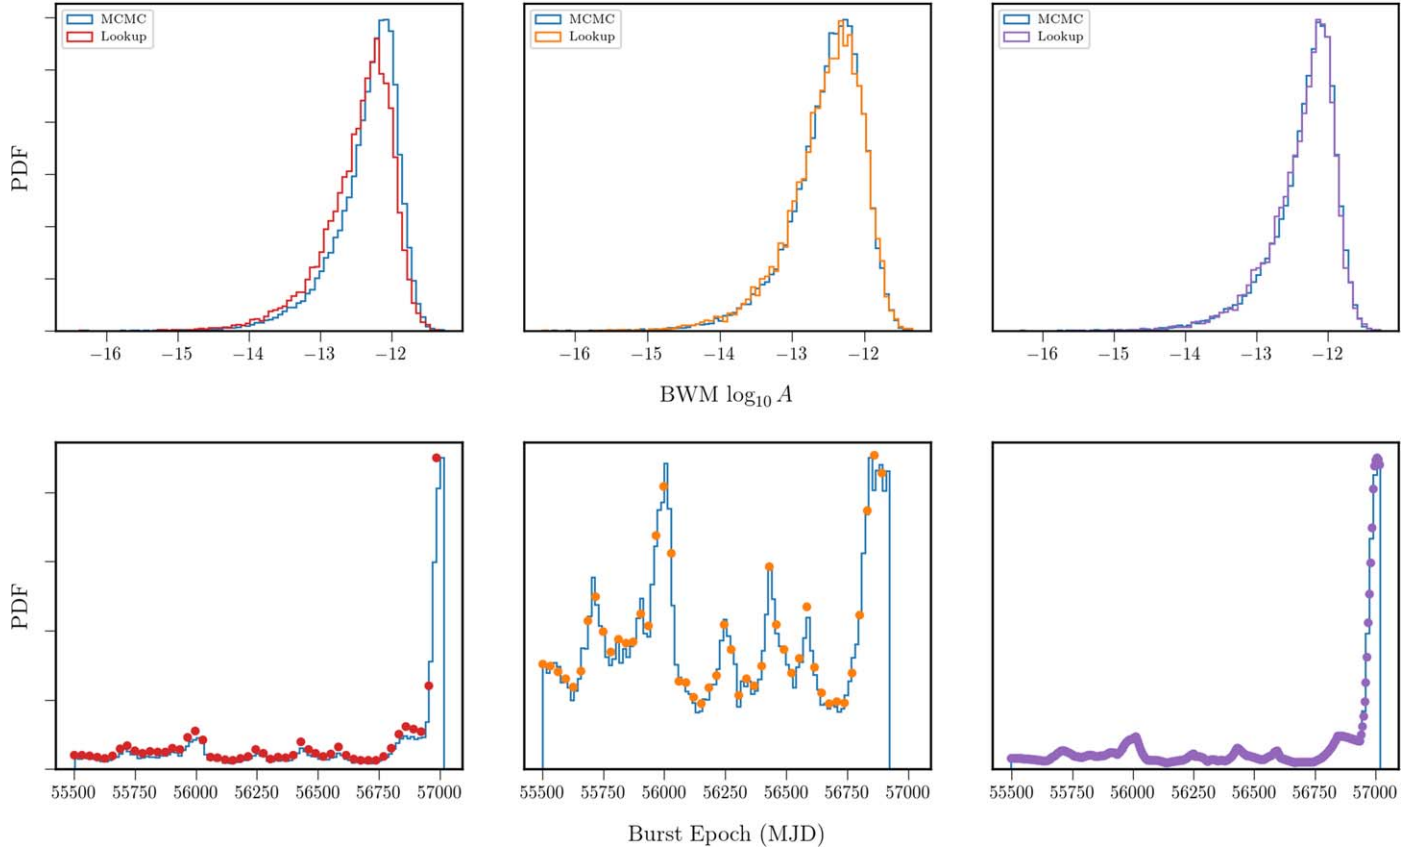
In this section, we will compare ULs on the amplitudes of GW BWMs computed using this more efficient search and the previously published ULs in *NG11-bwm*. Then, we discuss the improvements in efficiency.

### 4.1. Pulsar-term Comparisons

Figure 1 shows the percent difference between the pulsar-term BWM ULs for both positively and negatively signed memory computed using direct MCMC methods and our lookup table-based method. It also shows the 95% confidence intervals computed by bootstrap sampling the posteriors from the likelihood tables and the MCMC runs. We can see that these confidence intervals are essentially consistent with zero difference for the chosen grid density. We chose to compare the differently signed memory ULs separately in order to fully compare the two techniques.

Additionally, we narrowed the priors on the burst epoch to exclude the first 180 and last 270 days of each pulsar's TOAs. This choice is largely motivated by extremely large posterior probabilities for bursts at very early and late times in several

## J1903+0327 Posterior Comparison



**Figure 2.** Marginal amplitude and burst epoch posteriors for three different sets of lookup tables. The leftmost pair show the marginal amplitude and burst epoch posteriors with 180 days excluded from the beginning and end of the epoch prior. The marginal posteriors computed using the MCMC sampler are in blue. In this case, the grid density is too low, and the last two points in the burst epoch grid do not fully characterize the true distribution. This results in a biased marginal amplitude posterior. The center pair shows the posteriors if we exclude an additional 90 days from the end of the burst epoch prior (orange). We can see that removing the large feature at the end of the data set gives good agreement between the amplitude posteriors. Additionally, the rightmost pair shows that a higher grid density that has enough grid points to accurately characterize the features in the burst epoch posterior also gives very good agreement of the amplitude posteriors.

pulsars. These early- and late-time bursts are not credible, and only exist because they cannot be ruled out by data (there is not enough data before/after an early/late burst to constrain the amplitude). Furthermore, the posteriors in these cases vary on very short timescales, and therefore require more grid points to fully capture the feature. We can ameliorate this by either excluding early and late times from the pulsar-term search, or by using a much denser grid to characterize the posterior probability as a function of burst epoch. Figure 2 shows that both of these methods sufficiently address this problem. The left-hand side of the figure shows both the numerically marginalized posteriors (red) and the MCMC-computed posteriors (blue). We can see that without any special adjustments being made, the last two points of the burst epoch grid do not sufficiently characterize the posterior, and the resulting marginal amplitude posterior is biased toward high amplitudes. However, both the exclusion of early and late trial burst epochs (orange, center) and the use of a denser grid (purple, right) give agreement between the marginal amplitude posteriors.

#### 4.2. Earth-term Comparison


For the Earth-term ULs, we report two results: (1) the ULs as a function of burst epoch; and (2) the ULs as a function of

position in the sky. These results are shown in Figures 3 and 4, respectively.

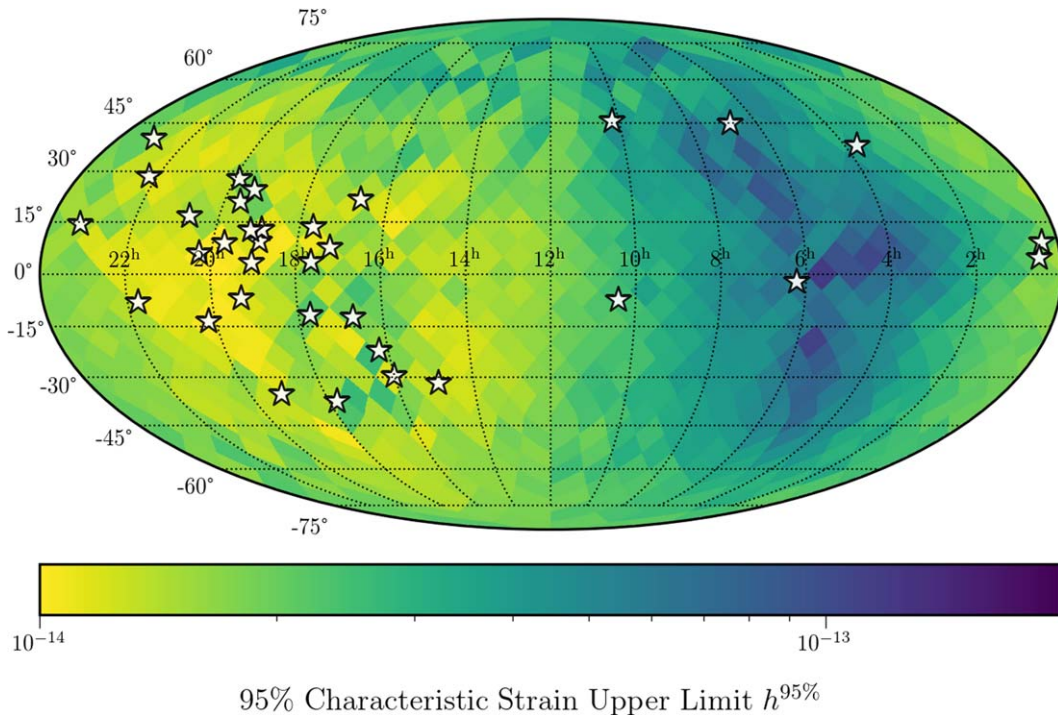
In Figure 3, we see that both methods return nearly identical ULs as a function of burst epoch. There are some significant differences, however, at early epochs. Although the ULs appear very discrepant, at these early epochs, there are very few recorded TOAs. As such, it is impossible to place very accurate limits on a BWM, since very large amplitude BWMs can be fit to the sparse data. Therefore, despite the apparent differences, we are not very concerned, since we expect a very nonconstraining UL at these early epochs. More importantly, as more pulsars and more data are added to the PTA, the ULs become nearly identical.

Figure 4 shows the ULs on the BWM amplitude as a function of sky location using the method described in Section 3.2. The resulting amplitude posterior is sampled to compute the 95% UL. We find that the amplitude ULs as a function of sky location are similar to those reported in *NG11-bwm*.

Although we can comment on general similarities between the results, we cannot directly compare them. *NG11-bwm* included an additional model, called *BayesEphem*, in their analysis. The *BayesEphem* model accounts for uncertainty in the solar system ephemeris. This is especially important in the



**Figure 3.** The 95% BWM amplitude ULs as a function of observation epoch. The original results are plotted in red, and the blue curve is used with permission from the authors of *NG11-bwm*. There is good agreement for the vast majority of the data set, with some discrepancy at early times. We believe that these discrepancies arise from the lack of data early in the data set, and expect uninformative, unconstraining ULs at these trial burst epochs.



**Figure 4.** Left: 95% BWM amplitude ULs as a function of sky location. The stars mark the locations of the pulsars in NANOGrav’s 11 yr data release. As expected, the PTA is most sensitive to BWMs in sky locations where many pulsars are being timed.

NANOGrav 11 yr data set, because the observation baseline is very close to Jupiter’s orbital period.

This model introduces 11 extra parameters, which is far too many to use on our parameter grid. It is therefore impossible to include *BayesEphem* using the techniques described in this work. Since there are no published results for ULs on BWMs in the NANOGrav 11 yr data set as a function of sky position that do not include *BayesEphem*, we report our results without a comparison.

This technique of using pulsar-term likelihood tables can be used to reproduce the same types of analyses and results that MCMC-based methods can. The fundamental Bayesian methodology is identical; both techniques compute marginalized posterior probabilities for model parameters. This

method simply takes advantage of the factorizable likelihood to more efficiently carry out the marginalization.

### 4.3. Computational Improvement

The computational complexity of computing the pulsar-term lookup tables is dominated by the cost of inverting the covariance matrix in Equation (12). This is an  $N_{\text{gp}} \times N_{\text{gp}}$  matrix, where  $N_{\text{gp}}$  is the number of Gaussian process parameters needed for a single-pulsar-term BWM signal model (Ellis & Haasteren 2017; Ellis et al. 2020). One inversion has computational complexity  $O(N_{\text{gp}}^3)$ . To compute a full pulsar-term lookup table, we evaluate the likelihood once for each point on a five-dimensional grid ( $A_m, \gamma_m, |h_0|, \text{sign}(h_0), t_0$ ).

Thus, the total cost of the inversions we must perform for one lookup table is  $N_{\text{Am}}N_{\gamma}N_{\text{Abwm}}N_{\text{sign}}N_{t_0}N_{\text{gp}}^3$ , where each of these terms represents the number of grid points in the lookup tables for red-noise amplitudes, red-noise spectral indices, BWM amplitudes, BWM signs, burst epochs, and Gaussian process parameters, respectively. For this paper, for a pulsar that has 10 yr of data, the total number of grid points is approximately  $32 \times 10^6$ . Then, if we compute one lookup table for each pulsar, for the purpose of convenient comparison, we can consider the complexity to be approximately  $10^7 N_{\text{psr}}N_{\text{gp}}^3$ , where  $N_{\text{psr}}$  is the number of pulsars in the PTA.

The computational complexity of one pulsar-term search for BWMs using an MCMC sampler may be approximated to be just the product of the complexity of one likelihood evaluation and the number of evaluations needed. This means the complexity of the pulsar-term search for BWMs is approximately  $N_{\text{iter}}N_{\text{psr}}N_{\text{gp}}^3$ , where  $N_{\text{iter}}$  is the number of iterations used per sampling run. Normally,  $N_{\text{iter}} \approx 10^6$  is sufficient for parameter estimates to converge, so we can consider the complexity to be  $10^6 N_{\text{psr}}N_{\text{gp}}^3$ .

It is very clear that the cost of producing one lookup table is significantly more expensive than performing one pulsar-term BWM search. However, once the pulsar-term likelihood tables are computed, it is very cheap to compute the global likelihoods in a full-PTA, Earth-term BWM search. For example, to compute the ULs as a function of trial burst epoch (the results shown in Figure 3), a full-PTA covariance matrix must be inverted. Because the signal model does not contain correlations between pulsar pairs, we may take advantage of the block diagonal structure of the covariance matrix and invert it in  $O(N_{\text{psr}}N_{\text{gp}}^3)$ . In other words, the matrix inversion itself is no less expensive. However, because of the sampling scheme, we must perform one MCMC sampling run for each set of  $(N_{\theta}, N_{\phi}, N_{\psi}, N_t)$ . Thus, the total complexity of computing ULs as a function of burst epoch is approximately  $N_{\text{iter}}N_{\theta}N_{\phi}N_{\psi}N_tN_{\text{psr}}N_{\text{gp}}^3$ . In *NG11-bwm*, this total cost is approximately  $8 \times 10^8 N_{\text{psr}}N_{\text{gp}}^3$ .

We see that although the search is less efficient for computing pulsar-term ULs, it is far more efficient when computing certain full-PTA searches. On an Intel i9-9900K CPU with eight physical cores operating at 3.60 GHz, it takes approximately two weeks to compute all the single-pulsar lookup tables. Once the lookup tables have been produced, each of the full-PTA searches may be completed in approximately two days. Using only MCMC sampling to compute full-PTA ULs would have taken approximately 3 yr.

## 5. Conclusion

In this paper, we have implemented a more efficient technique for performing a Bayesian search for GW BWMs by using precomputed lookup tables to circumvent repeated, expensive matrix inversions to compute a factorizable likelihood. This method is faster and gives very similar results to those given by MCMC sampling. In addition, because all deterministic signals necessarily factorize, this method is not limited only to GW BWM searches. However, the BWM signal lends itself very well to this method because both the pulsar-term and Earth-term signals have a very low-dimensional parameter space. This is not generally true of GW signals, and any extra parameter incurs significant costs in both computation and storage.

We believe that there are still improvements to be made. For example, a robust solution for any errors arising from our finite-density grid may be implementing a scheme for adaptive grid spacing, depending on the local variation of the likelihood surface. This way, we would spend less time over-characterizing regions of parameter space that do not vary much, while maintaining accuracy in quickly varying regions of parameter space.

Overall, we find that our sky-averaged ULs as a function of burst epochs (see Figure 3) match well with previously published results, with almost no difference in the most sensitive regions of the data set (although the ULs differ somewhat significantly at early trial epochs). This is somewhat unsurprising; there is very little timing data at early epochs, and we expect very weak constraints on any BWMs appearing this early in the data set.

Furthermore, we are able to perform the same full-PTA search for GWs over the entire sky. Although we cannot compare results with *NG11-bwm*, since they use an additional Bayesian ephemeris model, our results are still quite similar. For future data sets with more accurate ephemeris models, we expect these differences to become smaller. Specifically, when using the ephemeris model DE438 in Arzoumanian et al. (2020), the Bayesian ephemeris model, *BayesEphem*, no longer made a significant difference in common noise parameter estimation.

In the future, given the results in Arzoumanian et al. (2020), in which a detection of a common red-noise process was made, it will be important to include this common process in the signal model for future BWM searches. This additional signal requires introducing two new model parameters. While this would make this method take significantly longer, it may be possible to find improvements in computational costs by using Python vectorization or simply by reducing the resolution of the parameter grid. Preliminary testing shows that a reduction in grid resolution of approximately 20% still maintains a similar degree of accuracy to the results shown in this work. Even with the addition of two more signal parameters, we expect that this method will still be significantly faster than the traditional MCMC sampling method.

As pulsar timing baselines become longer and PTAs become populated with more pulsars, it will be difficult to use current MCMC sampling techniques to search for GW BWMs, and it will be important to find faster methods to do so. This method provides a very efficient way to perform searches for BWMs as PTAs continue to grow and data sets become too large for MCMC sampling to be tractable without significant computational resources.

## Acknowledgments

The NANOGrav project receives support from National Science Foundation (NSF) Physics Frontiers Center award Nos. 1430284 and 2020265. This work was also partly supported by the George and Hannah Bolinger Memorial Fund in the College of Science at Oregon State University. We also thank the anonymous reviewer for their suggestions and insightful questions, which led to a clearer manuscript and deeper understanding of the results of this paper.



## ORCID iDs

Jerry Sun  <https://orcid.org/0000-0002-7933-493X>  
 Paul T. Baker  <https://orcid.org/0000-0003-2745-753X>  
 Aaron D. Johnson  <https://orcid.org/0000-0002-7445-8423>  
 Dustin R. Madison  <https://orcid.org/0000-0003-2285-0404>  
 Xavier Siemens  <https://orcid.org/0000-0002-7778-2990>

## References

- Aggarwal, K., Arzoumanian, Z., Baker, P. T., et al. 2020, *ApJ*, 889, 38  
 Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, 905, L34  
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, *ApJ*, 810, 150  
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *ApJS*, 235, 37  
 Christodoulou, D. 1991, *PhRvL*, 67, 1486  
 Cordes, J. M., & Jenet, F. A. 2012, *ApJ*, 752, 54  
 Cutler, C., Burke-Spolaor, S., Vallisneri, M., Lazio, J., & Majid, W. 2014, *PhRvD*, 89, 042003  
 Detweiler, S. 1979, *ApJ*, 234, 1100  
 Ellis, J., & Haasteren, R. V. 2017, Jellis18/Ptmcmcsampler: Official Release, Zenodo, doi:10.5281/ZENODO.1037579  
 Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2020, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, Zenodo, doi:10.5281/ZENODO.4059815  
 Estabrook, F. B., & Wahlquist, H. D. 1975, *GRGr*, 6, 439  
 Favata, M. 2010, *CQGra*, 27, 084036  
 Foster, R. S., & Backer, D. C. 1990, *ApJ*, 361, 300  
 Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759  
 Hellings, R. W., & Downs, G. S. 1983, *ApJL*, 265, L39  
 Iso, K., Simon, J., Burke-Spolaor, S., & Siemens, X. 2019, arXiv:1906.11936  
 Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, *PhRvD*, 87, 104021  
 Madison, D. R., Cordes, J. M., & Chatterjee, S. 2014, *ApJ*, 788, 141  
 Malmquist, K. G. 1922, *MeLuF*, 100, 1  
 Manchester, R. N. 2013, *CQGra*, 30, 224010  
 McLaughlin, M. A. 2013, *CQGra*, 30, 224008  
 Phinney, E. S. 2001, arXiv:astro-ph/0108028  
 Pshirkov, M. S., Baskaran, D., & Postnov, K. A. 2010, *MNRAS*, 402, 417  
 Sazhin, M. V. 1978, *SvA*, 22, 36  
 Seto, N. 2009, *MNRAS Lett.*, 400, L38  
 Taylor, S. R., Baker, P. T., Hazboun, J. S., Simon, J., & Vigeland, S. J. 2022, enterprise\_extensions v2.4.3, GitHub, [https://github.com/nanograv/enterprise\\_extensions](https://github.com/nanograv/enterprise_extensions)  
 Thorne, K. S. 1992, *PhRvD*, 45, 520  
 van Haasteren, R., & Levin, Y. 2010, *MNRAS*, 401, 2372  
 van Haasteren, R., & Vallisneri, M. 2014, *PhRvD*, 90, 104012  
 van Haasteren, R., & Vallisneri, M. 2015, *MNRAS*, 446, 1170  
 Woodbury, M. A. 1950, *Inverting Modified Matrices* (Princeton, NJ: Department of Statistics, Princeton Univ.)  
 Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, 4, 1298