

AN ABSTRACT OF THE DISSERTATION OF

Jerry Sun for the degree of Doctor of Philosophy in Physics presented on
May 12, 2025.

Title: Facilitating Searches For Nonlinear Gravitational Wave Memory in
Modern Pulsar Timing Arrays

Abstract approved: _____

Xavier Siemens

Pulsar timing arrays (PTAs) are GW detectors which can probe the nanohertz regime of the gravitational wave (GW) spectrum. Searches for various gravitational wave signals in this regime provide invaluable insights on populations of super massive black hole binary (SMBHB) systems and merger rates. These insights help elucidate the evolution of the universe over long cosmic timescales. In addition, as GW detectors, PTAs can be used to provide constraints on gravitational theory and test general relativity. One such prediction of general relativity is “nonlinear” GW memory, which is a permanent perturbation to the spacetime metric sourced by energy carried out of a system by gravitational waves. Robust Bayesian searches for this signal in PTAs can be expensive when using ignorance priors because of an insidious effect of the PTAs antenna response to the signal. In this work, I present two new techniques for searching for nonlinear GW memory in PTAs with long

timing baselines. The first technique uses precomputed likelihood lookup tables to expedite the traditional search for memory. I also give the results of a search for nonlinear memory in the NANOGrav 12.5-year data set using this new efficient Bayesian method. Overall, the sensitivity of this data set is similar to the previous NANOGrav 11-year data set, and, finding no compelling evidence of a signal, I continue to present updated upper limits. The second technique maximizes the likelihood ratio of a model including a nonlinear memory signal compared to the null-signal model. I show that this frequentist estimator follows the expected χ^2 -distribution, and that it is possible to use it to compute maximum likelihood estimators for the quantities that parameterize the amplitude of the nonlinear memory signal. This work represents novel developments in the search for nonlinear memory as a link to future generations of PTA data sets. These techniques and searches are the first to be carried out in pulsar timing data sets that are dominated by a stochastic gravitational wave background, and set a foundation for future searches for nonlinear memory.

©Copyright by Jerry Sun
May 12, 2025
All Rights Reserved

Facilitating Searches For Nonlinear Gravitational Wave Memory in
Modern Pulsar Timing Arrays

by

Jerry Sun

A DISSERTATION

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented May 12, 2025
Commencement June 2026

Doctor of Philosophy dissertation of Jerry Sun presented on May 12, 2025.

APPROVED:

Major Professor, representing Physics

Head of the Department of Physics

Dean of Graduate Education

I understand that my dissertation will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my dissertation to any reader upon request.

Jerry Sun, Author

ACKNOWLEDGEMENTS

First and foremost, I'd like to warmly and sincerely thank my doctoral advisor Dr. Xavier Siemens for his mentorship, patience, and support. There have been many moments in graduate school where I have felt burned out and inadequate. At each of these moments, Xavi reminded me to reflect on past achievements and to prioritize taking care of myself. His academic guidance throughout the course of my doctoral program has also been invaluable. Xavi's advice is often "Start with a simple thing and understand it well. Then, work on understanding the complications." This is good advice and I found it very grounding and sensible while working on difficult problems and fixing mysterious bugs. I expect to carry this wisdom with me through every future endeavor. Thank you, Xavi.

I also want to express my gratitude to Dr. Paul Baker, Dr. Dustin Madison, and Dr. Aaron Johnson. When I first began studying gravitational wave memory, they pointed me towards many helpful references and shared their knowledge, experience, and code. The work presented here would not have been possible without their contributions and guidance.

Furthermore, I'd like to thank Dr. Jeffrey Hazboun. Since the very beginning, Dr. Hazboun has been a valuable source of knowledge. My first attempts to install analysis software were rescued by him. Since then, Jeff has explained many things about pulsar timing arrays, noise, and data analysis to me.

To my doctoral committee, Dr. Craig, Dr. Lazzati, Dr. Sun, thank you for your support and enlightening questions. I have learned a great deal of physics

while thinking about these questions.

I'd like to thank my cohort and my friends, new and old. To my cohort, thank you for all the community, laughter, and most importantly, discussions about homework problems. I am glad we discovered the off-axis magnetic field of a solenoid together. To my friends Christian, Rand, Nick, Noah, Katie, Spencer, and Carter I will fondly remember the movie nights, laughs, and conversations we have had over these many years. To my office mates, Dallas, Jacob, and Robin, thank you for putting up with the permanent stream of grumbling from my corner of the office. Your camaraderie and collaboration over the last several years has been a constant source of joy. To Elliot, Sam, David, Adam, Alan, Fan, and Nolan, thank you for your friendship. The time and fun we share is something I look forward to every day. To Vic, thank you for your unwavering support and care.

Finally, to my siblings Joyce, Jennifer, Jason, and Jacy, thank you for all the kind wishes and support throughout this journey. To my rabbit Coco, thank you for sitting by me and reminding me that it's okay to take breaks to give you a carrot. To Mom and Dad, thank you for your support and wisdom. You remind me every day what it means to work hard to make the world better. I always look forward to the time we can spend together.

CONTRIBUTION OF AUTHORS

Drs. Xavier Siemens, Aaron D. Johnson, Dustin R. Madison, and Paul T. Baker were involved in advising and coordinating the writing of the articles for the projects in Chapters 3 and 4. Additionally, NANOGrav is a large collaboration, and its work would not be possible without the contributions of many people. As such, in adherence with the NANOGrav collaboration publication policies, in Chapter 4, I include many additional members of NANOGrav who have contributed in many different ways, including (but not limited to), code development, pulsar timing, data analysis, cyber-infrastructure management, and many other critical activities. Dr. Dustin R. Madison also assisted in editing the paper presented in Chapter 5. Finally, Dr. Xavier Siemens contributed many edits to all three of the manuscripts.

TABLE OF CONTENTS

	<u>Page</u>
1 Introduction	1
1.1 Pulsar Timing Arrays	1
1.2 Pulsar Timing Residuals from GWs	2
1.3 GW Signal Models	7
1.3.1 Nonlinear GW Memory	7
1.3.2 Stochastic GW Background	11
1.4 Pulsar Timing Models and the PTA likelihood	15
1.5 Statistical Inference in PTAs	22
2 Literature Review	28
2.1 Prediction	28
2.2 Search Methodologies and Detection Prospects for Memory in PTA Data	31
2.3 Current Limits on Nonlinear Memory	34
3 Implementation of an Efficient Bayesian Search for Gravitational-wave Bursts with Memory in Pulsar Timing Array Data	36
3.1 Abstract	37
3.2 Introduction	37
3.3 Signal and Data Model	39
3.4 Methodology	42
3.4.1 Bayesian Approach	42
3.4.2 Accelerated Bayesian Search	47
3.5 Results	53
3.5.1 Pulsar-Term Comparisons	56
3.5.2 Earth-Term Comparison	57
3.5.3 Computational Improvement	58
3.6 Conclusion	62
4 The NANOGrav 12.5 yr Data Set: Search for Gravitational Wave Memory 65	
4.1 Abstract	67

TABLE OF CONTENTS (Continued)

	<u>Page</u>
4.2 Introduction	68
4.3 Data	71
4.4 Signal and Noise Model	72
4.5 Methods	76
4.6 Results	80
4.6.1 Earth-term Memory Search	80
4.6.2 Pulsar-term Upper Limits	81
4.6.3 Earth-term Upper Limits	84
4.7 Discussion and Conclusion	89
5 Optimal strategies for gravitational wave memory detection in pulsar timing arrays	92
5.1 Abstract	93
5.2 Introduction	93
5.3 Background	96
5.3.1 Signal Model	96
5.3.2 Likelihood	99
5.3.3 General \mathcal{F} -statistic	101
5.3.4 Coherent \mathcal{F} -statistic	105
5.3.5 Incoherent \mathcal{F} -statistic	108
5.4 Methodology and Results	111
5.4.1 Simulated Data Sets	111
5.4.2 Results and Discussion	113
5.5 Conclusion	118
5.6 Acknowledgments	122
6 Conclusion	123
Bibliography	127
Appendices	158

LIST OF FIGURES

Figure	Page
1.1 An example of the expected residual signature in a pulsar whose rotational frequency is red-shifted by the GW memory effect. In this plot, the memory wavefront passes by the Earth at some burst epoch $t_0 = 2000$ MJD.	10
1.2 Spatial correlation pattern for the residuals induced by the SGWB from an ensemble of SMBHBs as a function of pulsar angular separation $\theta_{\alpha\beta}$	14
3.1 Percent difference in the pulsar term amplitude upper limits (ULs) for each pulsar used in [5]. These percent differences are computed by comparing the 95% ULs from Bayesian MCMC runs and lookup-table-based methods for positive- 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 upper limits of positively-signed memory, and the blue points are the differences for negatively-signed memory, and the error bars show the 95% confidence intervals computed from bootstrapping.	54

LIST OF FIGURES (Continued)

<u>Figure</u>	<u>Page</u>
<p>3.2 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 lookup tables are in red, and the marginal posteriors computed using the MCMC sampler are in red. 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 which has enough grid points to accurately characterize the features in the burst epoch posterior also gives very good agreement of the amplitude posteriors.</p>	55
<p>3.3 The 95% BWM amplitude upper limits (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 Aggarwal et al. [5]. 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 upper limits at these trial burst epochs.</p>	59
<p>3.4 Left: 95% BWM amplitude upper limits (ULs) as a function of sky location. The stars mark the locations of the pulsars in NANOGrav’s 11-year data release. As expected, the PTA is most sensitive to BWMs in sky locations where many pulsars are being timed.</p>	60
<p>4.1 A corner plot showing 1D and 2D marginalized posteriors for three key model parameters: burst epoch t_B, burst strain amplitude $\log_{10} h_0$, and CURN amplitude $\log_{10} A_{\text{CURN}}$. The good localization $\log_{10} A_{\text{CURN}}$ shows that the CURN is still detected in the presence of the memory model. Furthermore, the tail of $\log_{10} h_0$ extends to very low amplitude, which indicates that $h_0 \sim 0$ is still supported by the model.</p>	82

LIST OF FIGURES (Continued)

<u>Figure</u>	<u>Page</u>
<p>4.2 A plot of the pulsar-term upper limits on memory strain amplitude. The pulsars are listed in order of shortest to longest timing baseline. To find these upper limits, we compute amplitude posteriors from the pulsar-term lookup tables marginalized over the burst epoch, pulsar intrinsic red noise and a fixed spectral index common uncorrelated red noise (CURN) process. Overall, we do not find much difference in pulsar-term upper limits when comparing the results using a fixed CURN spectral index of $\gamma_{\text{SMBHB}} = 4.33$ and $\gamma_{\text{MAP}} = 5.5$.</p>	83
<p>4.3 Top left: The upper limits on memory strain amplitude as a function of skypixel including a CURN law process using a fixed spectral index of $\gamma_{\text{SMBHB}} = 4.33$, as expected for a stochastic gravitational wave background originating from an ensemble of uniformly, isotropically distributed SMBHBs. Top right: The upper limit on memory strain amplitude as a function of sky pixel including a CURN power law process using a fixed spectral index of $\gamma_{\text{MAP}} = 5.5$, the maximum a posteriori value for the detected CURN in Arzoumanian et al. [15]. Bottom: The difference between the top two panels. A positive value indicates that the upper limits on strain using a spectral index of 5.5 are higher. We see that overall, the upper limits change slightly when the red noise model uses the preferred red noise spectral index. However, these changes are localized to a small part of the sky. It can be shown that these upper limit differences can be largely attributed to PSR J0613–0200.</p>	87
<p>4.4 The memory strain amplitude upper limit as a function of burst epoch. The black curve shows the upper limits using a model which does not include CURN. The blue and orange curves show the upper limits using a model which does include a fixed spectral index CURN power law process with $\gamma_{\text{MAP}} = 5.5$ and $\gamma_{\text{SMBHB}} = 4.33$, respectively. We see that the sensitivity of the NANOGrav 12.5-year data set does not give significantly improved upper limits for the first 11 years. However, the additional pulsars and timing data give continuously improving upper limits later in the data set.</p>	88

LIST OF FIGURES (Continued)

<u>Figure</u>	<u>Page</u>	
4.5	<p>A plot of the predicted rate of SMBHB mergers (red) against the rate upper limits derived from the strain amplitude upper limit as a function of burst epoch. Because the strain upper limits do not improve significantly, the event rate upper limits are correspondingly similar to those presented in Aggarwal et al. [5].</p>	88
5.1	<p>This figure shows the best-fit χ^2-distributions to the \mathcal{F}_C and \mathcal{F}_I statistics in multiple scenarios containing different combinations of gaussian white noise, a common spatially uncorrelated red noise (CURN), and a nonlinear memory signal. We see that in these cases, the theoretically predicted distributions and the theoretical non-centrality parameter match almost exactly with the distributions of the \mathcal{F}-statistics from 5000 simulated realizations. Notably, the SNR decreases significantly when comparing the statistics in the presence of red noise.</p>	114
5.2	<p>This figure shows the results of maximum-likelihood parameter recovery using Eqs. 5.37 and 5.38 in multiple signal regimes. The top row show the recovered parameters for the data sets without injected pulsar-intrinsic red noise, and the bottom row shows the results for the data sets that contain additional pulsar-intrinsic red noise. The summary of the injection and recoveries may be found in Table 5.2. As expected, the maximum likelihood estimators for the strain amplitude h_0 and ψ do very poorly at small SNR (leftmost column, SNR = 1.74 and SNR = 1.46). However, with a moderately large signal (middle column, SNR = 5.03 and SNR = 4.22), the strain amplitude estimator becomes more accurate, while the polarization is not well measured in some realizations. Only when the signal is very loud relative to the noise (rightmost column, SNR = 15.67 and SNR = 13.34) does the maximum likelihood estimator for the ψ become accurate. The weak sensitivity to polarization angle is largely due to the total size (40 pulsars) and distribution (uniform) of the pulsars in this PTA. A PTA with more pulsars distributed more densely near the sky location of the signal source would better constrain the polarization of the signal.</p>	115

LIST OF TABLES

Table	Page
3.1	48
3.2 Priors used for each of the model parameters in the Bayesian search for global Earth-term GW BWMs using the full PTA. 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 subsection 3.4.1. The prior on t_{BWM} also varies depending on the particular upper limit calculation. For upper limits (ULs) as a function of sky location, we use priors between MJD 56000 and MJD 57000. For upper limits 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 subsection 3.4.2.	48
5.1 This table summarizes the values of the injected noise and signal parameters used to create 5000 realizations of simulated data sets. These realizations were used to compute the \mathcal{F} -statistics shown in Figure 5.1. For data sets that include intrinsic red noise, the red noise parameters were fixed to be constant for all realizations. . . .	112

LIST OF TABLES (Continued)

<u>Table</u>	<u>Page</u>
5.2	<p>This table shows the signal-to-noise ratio, the injected signal parameters, and maximum-likelihood estimators for the signal parameters recovered using \mathcal{F}_C in multiple signal regimes over 5000 realizations of simulations for data sets both including and excluding intrinsic pulsar red noise. The values in parentheses are the results for the data sets including intrinsic pulsar red noise. The signal regimes are determined based on the SNR of the signal compared to the injected noise, which was kept constant (see 5.1 for a summary of the injected noise and signals). The maximum-likelihood estimators for the injected memory strain-amplitude are very imprecise for low-SNR signals, but become much more accurate at higher SNR. This is because GW memory signal has a similar power spectrum to a red noise with $\gamma = 13/3$. However, the background is loud compared to the GW memory signal expected from a modest SMBHB merger. The sensitivity of the PTA to memory thus has a lower limit determined by the amplitude of the red noise.</p>
	113

Chapter 1: Introduction

In this chapter, we will introduce all the physics and mathematics necessary for a fundamental understanding of how pulsar timing arrays (PTAs) can be used to search for gravitational waves (GWs). First, in §1.1, we outline (very briefly) why PTAs can be used to detect GWs. Next, in §1.2, we will derive the pulsar timing residuals (see §1.1 for the definition of a “timing residual”) induced by a GW following the pedagogical derivation in Ref. [85]. Then, in §1.3 we will present the residuals induced by several different types of GW signals. Next, in §1.4, we will discuss the pulsar timing model in full and present a short derivation of the likelihood function used in PTA data analysis. Finally, in §1.5, we will discuss how Bayesian inference techniques are used in PTA data analysis.

1.1 Pulsar Timing Arrays

Pulsars are neutron stars that emit beams of electromagnetic radiation from their magnetic poles that can be detected as periodic pulses in telescopes on Earth. Millisecond pulsars (MSPs) fast-spinning pulsars that have incredibly stable rotational frequencies, and their rotational phases can be precisely tracked over very long timescales (\sim decades)[44, 79]. The electromagnetic radiation from a pulsar can be detected as periodic pulse in a telescope on Earth.

PTAs use a network of millisecond pulsars (MSPs) to detect changes in the space-time metric. As gravitational waves (GWs) pass over the line of sight from the Earth to an MSP, the geodesics of the emitted photons are lengthened or shortened, and the pulse times of arrival (TOAs) are delayed or advanced, respectively. The difference between the observed TOA and the expected TOA is defined as the “timing residual”. The rotational stability of an MSP makes it possible to detect the incredibly small timing residuals induced by a GW. As a result, an array of MSPs can be used to form a GW detector [44, 51, 57, 103].

1.2 Pulsar Timing Residuals from GWs

We can understand the effect of a GW on a timing residual by studying the effect that a GW has on the geodesic of photons emitted from a pulsar as they travel to Earth. The steps of this derivation closely follow those shown in Ref. [85]. Consider a system with the Earth, a single pulsar, and a GW. Without loss of generality, let us orient the system so that the GW is traveling along the z -axis and place the pulsar at the origin with the Earth a distance d_a away along the unit vector $\hat{\mathbf{n}} = (n^x, n^y, n^z)$ with $|\hat{\mathbf{n}}| = 1$. Finally, let T_{psr} be the rotational period of the pulsar. The spacetime interval ds^2 with a metric perturbation originating from a GW traveling along the z -axis in the transverse-traceless (TT) frame

$$ds^2 = -dt^2 + dx^i dx^j [\delta_{ij} + h_{ij}^{\text{TT}}(t, \mathbf{x}(t))], \quad (1.1)$$

where i and j index the three spatial dimensions, h_{ij}^{TT} is the time-dependent GW strain, $\mathbf{x}(t)$ gives the path of the photon, and we have used a unit system in which the speed of light $c = 1$. Then, we can parameterize the unperturbed trajectory of the photon with an affine parameter λ

$$x^i(\lambda) = n^i \lambda, \quad (1.2)$$

so that

$$dx^i = n^i d\lambda. \quad (1.3)$$

Substituting into Equation 1.1, and using the fact that photons travel along null geodesics ($ds^2 = 0$)

$$dt^2 = d\lambda^2 [1 + n^i n^j h_{ij}(t, \mathbf{x}(t))]. \quad (1.4)$$

Assuming the perturbation h_{ij} is small, we can approximate

$$dt \approx d\lambda [1 + \frac{n^i n^j}{2} h_{ij}(t, \mathbf{x}(t))]. \quad (1.5)$$

Rearranging and approximating again yields

$$d\lambda = dt [1 - \frac{n^i n^j}{2} h_{ij}(t, \mathbf{x}(t))]. \quad (1.6)$$

We can now integrate this from the photon emission time t_{em} to some observation time t_{obs}

$$\begin{aligned}\Delta\lambda = d_{\text{Earth}} &= \int_{t_{\text{em}}}^{t_{\text{obs}}} \left[1 - \frac{n^i n^j}{2} h_{ij}(t, \mathbf{x}(t))\right] dt \\ &= t_{\text{obs}} - t_{\text{em}} - \frac{n^i n^j}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + d_{\text{Earth}}} h_{ij}(t, \mathbf{x}(t)) dt,\end{aligned}\quad (1.7)$$

where d_{Earth} is the distance from the pulsar at the origin to the Earth. We have also used the unperturbed travel time of the photon in $h_{ij}(t, \mathbf{x}(t))$ since the integral on the right-hand side is already to first order in the small quantity h_{ij} . Solving for the observation time yields

$$t_{\text{obs}} = t_{\text{em}} + d_{\text{Earth}} + \frac{n^i n^j}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + d_{\text{Earth}}} h_{ij}(t, \mathbf{x}(t)) dt. \quad (1.8)$$

Repeating this for the subsequent pulse that later emitted one rotation of the pulsar (T_{psr}) and introducing a small variable change $t_{\text{em}} \rightarrow t_{\text{em}} + T_{\text{psr}}$, we obtain

$$t'_{\text{obs}} = t_{\text{em}} + T_{\text{psr}} + d_{\text{Earth}} + \frac{n^i n^j}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + d_{\text{Earth}}} h_{ij}(t + T_{\text{psr}}, \mathbf{x}(t)) dt. \quad (1.9)$$

Taking the difference between these two times yields

$$\begin{aligned}t'_{\text{obs}} - t_{\text{obs}} &= T_{\text{psr}} + \frac{n^i n^j}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + d_{\text{Earth}}} \frac{\partial h_{ij}(t, \mathbf{x}(t))}{\partial t} dt \\ &= T_{\text{psr}} + \Delta T,\end{aligned}\quad (1.10)$$

where the quantity ΔT is defined to be the perturbation to the rotational period of the pulsar induced by the GW (the integral term). Again, we use the unperturbed time and trajectory

$$\mathbf{x}(t) = -(t_{\text{em}} + d_a - t)\hat{\mathbf{n}}, \quad (1.11)$$

since the perturbations introduce higher order effects to the already-small h_{ij} . We also make the remark that this expression comes from Taylor approximation in the first argument of $h_{ij}(t, \mathbf{x}(t))$. As such, the partial derivative only acts on the explicit t -dependence, not the dependence of the path $\mathbf{x}(t)$ on the time. If we assume the functional form of the gravitational wave to be

$$h_{ij}(t, \mathbf{x}(t)) = A_{ij}(\hat{\mathbf{n}}_{\text{GW}})u(t - \hat{\mathbf{n}}_{\text{GW}} \cdot \mathbf{x}), \quad (1.12)$$

where the tensor A_{ij} is a linear combination of +- and \times -polarization tensors, e_{ij}^+ and e_{ij}^\times , and $u(t - \hat{\mathbf{n}}_{\text{GW}} \cdot \mathbf{x})$ carries the time-dependent propagation of the GW. For a GW travelling along the z -axis, these basis tensors are

$$e_{ij}^+ = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad e_{ij}^\times = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.13)$$

The general polarization vectors for any GW may be simply computed using standard rotation matrices. then, condensing $t_{\text{em}} + d_a = t_{\text{obs}}$ and using a substitution,

we can evaluate the integral

$$\int_{t_{\text{em}}}^{t_{\text{obs}}} \frac{\partial h_{ij}(t, \mathbf{x}(t))}{\partial t} dt = \frac{1}{(1 - \hat{\mathbf{n}}_{\text{GW}} \cdot \hat{\mathbf{n}})} [h_{ij}(t_{\text{obs}}, \mathbf{x}(t_{\text{obs}})) - h_{ij}(t_{\text{em}}, \mathbf{x}(t_{\text{em}}))]. \quad (1.14)$$

We find the surprising result that apparent fractional change in the pulsar rotational period $\Delta T/T$ depends only on the difference between the metric at the observation point and the deferred time at the emission point. We define this fractional change in the period (or equivalently rotational frequency) to be the redshift $z(t)$ induced by some perturbation to the metric $h_{ij}(t, \mathbf{x}(t))$

$$z(t) \equiv \frac{\nu_0 - \nu(t)}{\nu_0} = \frac{\Delta T}{T_{\text{psr}}}, \quad (1.15)$$

where ν_0 is the actual rotational frequency and $\nu(t)$ is the observed rotational frequency after accounting for the effect of the GW. Substituting our previous result to compute the redshift gives

$$z(t) = \frac{n^i n^j}{2(1 - \hat{\mathbf{n}}_{\text{GW}} \cdot \hat{\mathbf{n}})} [h_{ij}(t, \mathbf{x}_{\text{Earth}}) - h_{ij}(t - d_a, \mathbf{x}_{\text{psr}})]. \quad (1.16)$$

The timing residuals with respect to some reference time $t_0 = 0$ are computed by integrating the redshift until the observation time of the photon t_{obs}

$$R(t) = \int_0^{t_{\text{obs}}} dt' z(t'). \quad (1.17)$$

This quantity then gives the total advance/delay induced by perturbations to the

spacetime metric by some GW. We can use this formula to predict the residual signature of different kinds of GWs.

1.3 GW Signal Models

1.3.1 Nonlinear GW Memory

Linear GW memory was first predicted in Zel’dovich and Polnarev [130] to arise from unbound scattering events of compact objects like neutron stars. They find that over the course of a single scattering event, there is a change to the stress-energy tensor resulting from the change in the “Coulomb-like” gravitational potential. This causes a permanent change in the spacetime metric after the scattering event. This effect is called “linear” memory. Later, it was discovered that the energy carried away by GWs in any GW-producing event also sources a permanent change to the stress-energy tensor and thus the spacetime metric [19, 32, 115]. Although this effect comes from higher order terms that are usually neglected in linearized gravity, the resulting permanent changes in the metric from this type of memory are the same order of magnitude as the metric perturbations in linearized theory. Thus, the memory originating from emitted GWs is called “nonlinear” memory. While any GW-emitting event can produce this type of memory, it is expected that the only sources which will yield *detectable* memory are those which emit huge amounts of energy via GWs [115, 124].

For a circularized orbit of two black holes with reduced mass-ratio η , total mass

M_{tot} , at a comoving distance R , whose orbital plane has an inclination angle of ι , the difference in GW strain Δh_+ before and after the coalescence is [53]

$$\Delta h_+ = \frac{\eta M_{\text{tot}}}{384\pi R} \sin^2 \iota (17 + \cos^2 \iota) h_0 \quad (1.18)$$

There is no cross-polarized component from circularized binaries [10, 124]. The intrinsic strain h_0 is determined by the evolution of the mass quadrupole moment of the source [53]

$$h_0 = \frac{1}{\eta M_{\text{tot}}} \int_{-\infty}^{T_f} dt \left| \frac{d^3 I_{22}}{dt^3} \right|^2 \quad (1.19)$$

Pulsar timing experiments have observation cadences which are infrequent enough to be insensitive to the rise-time of the memory signal [35, 118]. Thus, we consider only the DC-component of the signal. In other words, as a memory wavefront crosses over either the Earth or over the α -th pulsar, there will be a nonzero difference in the metric at the two objects, resulting in an observed redshift in the rotational frequency of the observed pulsar. This can be quantitatively expressed using Equation 1.16

$$z_\alpha(t) = \Delta h_+ B_\alpha(\theta, \phi) (\Theta(t - t_e) - \Theta(t - t_p)) \quad (1.20)$$

where we have adopted a more condensed notation $B(\theta, \phi)$ for the antenna pattern more commonly used in literature [35, 82, 118]. $B(\theta, \phi)$ captures the geometric antenna response pattern of the pulsar to the GW, Θ is the Heaviside function representing the DC-effect of the memory signal, t_e is the time the burst passes

over the Earth, and t_p is the time at which the same burst passes (or passed) over a pulsar. The new parameters θ and ϕ give the angle between the propagation direction of the wave $\hat{\mathbf{n}}_{\text{GW}}$ and the pulsar direction $\hat{\mathbf{n}}_\alpha$, and ϕ is the angle between the principal polarization vector of the wave and the projection of the pulsar direction onto a plane perpendicular to $\hat{\mathbf{n}}_{\text{GW}}$, respectively.

This signal is completely deterministic. We can compute the integral of this redshift in the α -th pulsar analytically, since the only time dependence is in the Heaviside function. We also remark that the inter-pulsar distances and Earth-pulsar distances in a PTA experiment are long enough (\sim kpc) that pulsar timing experiments (with observation times lines of \sim decades) are not able to see the same memory wave front travel over both the Earth and a pulsar. Additionally, a memory wavefront that crosses over only one pulsar will not produce compelling evidence for a detection of memory. This is because pulsars are known to spontaneously “glitch”. A pulsar glitch is a rapid spin-up event resulting in a sudden increase in the rotational frequency [62]. It is believed that glitches are the result of While MSPs glitch less often than canonical pulsars, there have been observations of MSP glitches [34, 91]. There has also been a report of the first “anti-glitch”, or sudden decrease in rotational frequency, observed in a pulsar [116]. However, this phenomenon seems to be relatively rare.

With this in mind, we will consider the memory signal resulting from a memory wave front passing over the Earth. The integral of the Earth-term in Equation 1.20

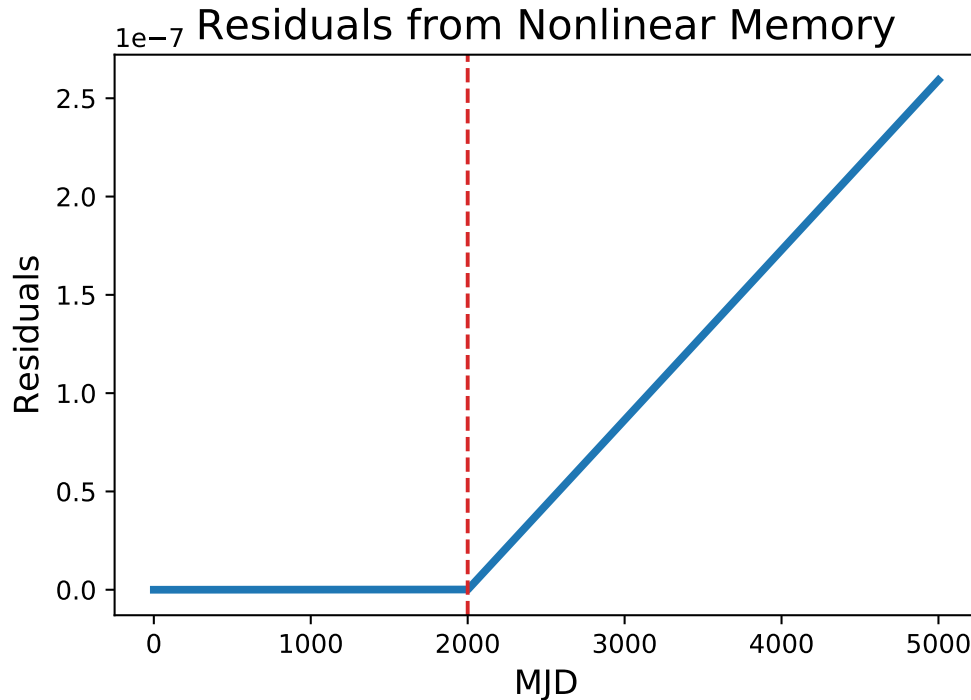


Figure 1.1: An example of the expected residual signature in a pulsar whose rotational frequency is red-shifted by the GW memory effect. In this plot, the memory wavefront passes by the Earth at some burst epoch $t_0 = 2000$ MJD.

gives the expected timing residuals for the *alphath* pulsar

$$R_\alpha^{\text{mem}}(t) = \Delta h_+ B_\alpha(\theta, \phi) \Theta(t - t_e)(t - t_e) \quad (1.21)$$

A cartoon of this residual signature is shown in Figure 1.1. When searching for this signal in PTAs, it is possible to use the deterministic antenna response to coherently search for this signal. In other words, the addition of data from multiple pulsars will compound the detectability of the signal.

We will conclude this brief discussion of the nonlinear GW memory signal here. However, there will be more details on the physics of nonlinear memory in section 2.1.

1.3.2 Stochastic GW Background

In this section, we will discuss the expected residual signature from a stochastic GW background (SGWB). While the focus of this work is on the search for nonlinear GW memory, it has been shown that there is strong evidence for a SGWB [1]. Additionally, several works have shown that the expected signature of the SGWB has a strong covariance with nonlinear GW memory and adversely affects detection prospects [35, 82, 111]. As a result, it is important to understand the effect that the SGWB has on pulsar timing residuals and the detectability of nonlinear memory.

The SGWB is assumed to be a direct result of the combination of many monochromatic signals from an ensemble of uniformly distributed SMBHB systems. This background is assumed to be isotropic, unpolarized, and stationary. The “smoking gun” signature of a SGWB comes in the form of correlations between the GW-induced redshifts of different pairs of pulsars. It can be shown that for two pulsars α and β , the correlation in redshifts $\langle z_\beta^*(t)z_\alpha(t) \rangle$ resulting from a stochastic background [64]

$$\langle z_\beta^*(t)z_\alpha(t) \rangle = \frac{1}{2} \sum_{A=+, \times} \int_{-\infty}^{\infty} df \int \frac{d^2 \hat{\mathbf{n}}_{\text{GW}}}{4\pi} S_h(f) F_\alpha^A(\hat{\mathbf{n}}_{\text{GW}}) F_\beta^A(\hat{\mathbf{n}}_{\text{GW}}), \quad (1.22)$$

where the quantity $F_\alpha^A(\hat{\mathbf{n}}_{\text{gw}})$ is the antenna response function of the α -th pulsar to the GW given by

$$F_\alpha^A(\hat{\mathbf{n}}_{\text{gw}}) = \frac{n_\alpha^i n_\alpha^j e_{ij}^A(\hat{\mathbf{n}}_{\text{gw}})}{2(1 + \hat{\mathbf{n}}_{\text{gw}} \cdot \hat{\mathbf{n}}_\alpha)}, \quad (1.23)$$

and $S_h(f)$ is the one-sided power spectral density of the SGWB defined by the correlator

$$\langle \tilde{h}_A^*(f, \hat{\mathbf{n}}) \tilde{h}_{A'}(f', \hat{\mathbf{n}}') \rangle = \frac{\delta^2(\hat{\mathbf{n}}_{\text{GW}} - \hat{\mathbf{n}}'_{\text{GW}})}{4\pi} \delta(f - f') \delta_{AA'} \frac{1}{2} S_h(f). \quad (1.24)$$

The integral over solid angle can be analytically computed after some tedious algebra [7]. We skip the details involved and report the final result. We define the function

$$C(\theta_{\alpha\beta}) \equiv \int \frac{d^2 \hat{\mathbf{n}}_{\text{GW}}}{4\pi} \sum_{A=+,\times} F_\alpha^A(\hat{\mathbf{n}}_{\text{GW}}) F_\beta^A(\hat{\mathbf{n}}_{\text{GW}}) \quad (1.25)$$

$$= x_{\alpha\beta} \log x_{\alpha\beta} - \frac{1}{6} x_{\alpha\beta} + \frac{1}{3}, \quad (1.26)$$

where $x_{\alpha\beta}$ is a parameter determined solely by the angular separation of the α -th and β -th pulsars

$$x_{\alpha\beta} \equiv \frac{1}{2}(1 - \cos \theta_{\alpha\beta}) \quad (1.27)$$

We can also express the correlations in the residuals instead of the redshifts

$$\langle R_\alpha(t) R_\beta(t) \rangle \simeq \zeta(\theta_{\alpha\beta}) \int_{f_i}^{f_h} df P_g(f), \quad (1.28)$$

where f_l and f_h are low- and high-frequency cutoffs determined by the total timing baseline and pulsar observation cadence (see §1.4 for how these are defined in PTA experiments), $P_g(f)$ is defined to be

$$P_g(f) = \frac{h_c^2(f)}{12\pi^2 f^{-13/3}}, \quad (1.29)$$

where $h_c(f)$ is called the “characteristic strain” of the background and $\zeta(\theta_{\alpha\beta})$ is related to the above correlator $C(\theta_{\alpha\beta})$ by

$$\zeta(\theta_{\alpha\beta}) \equiv \frac{3}{2}C(\theta_{\alpha\beta}) = \frac{3}{2}x_{\alpha\beta} - \frac{1}{4}x_{\alpha\beta} + \frac{1}{2}. \quad (1.30)$$

This correlation pattern is shown in Figure 1.2, and is the smoking gun for detecting a SGWB. This correlation pattern is nearly quadrupolar, resulting from the quadrupolar shape of the $+$ - and \times -polarization tensors. However, the direction of photon travel breaks the perfectly symmetric quadrupolar correlations by introducing a preferred direction. Some photons will “surf” along the same phase of any gravitational waves that are parallel to their trajectories. This results in the “broken” quadrupole shape for pulsars with separations of about 180° .

The characteristic strain is defined to be

$$h_c(f) \equiv \sqrt{f S_h(f)}, \quad (1.31)$$

and describes the strain of the SGWB at different frequencies weighted by the number of cycles that different frequency components spend in-band. For a SGWB

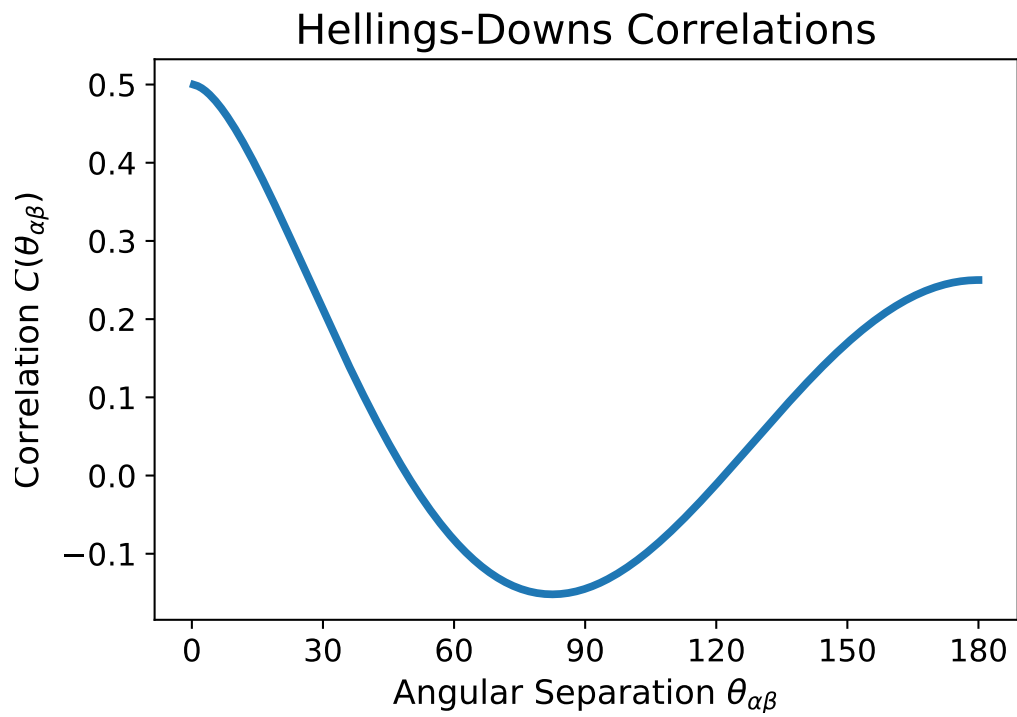


Figure 1.2: Spatial correlation pattern for the residuals induced by the SGWB from an ensemble of SMBHBs as a function of pulsar angular separation $\theta_{\alpha\beta}$.

of inspiralling SMBHB systems, the characteristic strain is predicted to be [99]

$$h_c(f) = A_{\text{SGWB},1\text{yr}} \left(\frac{f}{1\text{yr}^{-1}} \right)^{-2/3}, \quad (1.32)$$

where we have used the traditional expression that uses a reference frequency of $f_{\text{ref}} = 1\text{yr}^{-1}$ and a reference strain amplitude at this frequency $A_{\text{SGWB},1\text{yr}}$. Substituting this characteristic strain into Equation 1.29 yields the expected power spectrum of the SGWB

$$P_g(f) = \frac{A_{\text{SGWB},1\text{yr}}^2}{12\pi^2} \left(\frac{f}{1\text{yr}^{-1}} \right)^{-13/3}. \quad (1.33)$$

This signal is stochastic because it is the incoherent combination of many sources distributed throughout the universe. Furthermore, the GW strains at the Earth and at each pulsar are effectively decoupled, since they are separated by distances far greater than the decades-long operational timelines of pulsar timing experiments. Thus, the evidence for a SGWB comes from detection of significantly HD-correlated noise that has the power spectrum in Equation 1.33.

We conclude the description of the signal here, and discuss further how to detect this signal quantitatively in pulsar timing data in the next section.

1.4 Pulsar Timing Models and the PTA likelihood

In this section, we will explain the pulsar residual models as used in PTA GW data analysis. The full expression of the residual model for a single pulsar that

we will use is a sum of time series contributions from different physical processes. This can be written [15] as

$$\vec{r}_\alpha = \vec{n} + F\vec{a} + F_{\text{GW}}\vec{a}_{\text{GW}} + M\vec{\epsilon} + \vec{\delta t}_{\text{det}}, \quad (1.34)$$

where \vec{n} are contributions from Gaussian white noise (WN), F and \vec{a} are the Fourier basis and coefficients, respectively, which model Gaussian noise at different frequencies, F_{GW} and \vec{a}_{GW} are the Fourier basis and coefficients, respectively, which model the power spectrum of the background, M and $\vec{\epsilon}$ are the so-called “design matrix” and errors in the linearized pulsar timing model, and $\vec{\delta t}_{\text{det}}$ is a timeseries of deterministically computable residuals from a GW signal (such as nonlinear GW memory). We will discuss each of these terms in more detail.

The white noise term accounts for contributions to the residuals which are assumed to be Gaussian-distributed. In practice, this noise process is parameterized by a few different quantities that account for contributions to the noise from different parts of the data acquisition pipeline. However, it is still assumed that the overall statistics of the noise follow a normal distribution. Arzoumanian et al. [12] contains more details about each of the white noise parameters.

The components $F_{\text{GW}}\vec{a}_{\text{GW}}$ and $F\vec{a}$ both use a Fourier series to model noise at

different frequencies. Quantitatively:

$$F = \begin{bmatrix} \cos(2\pi f_1 t_0) & \sin(2\pi f_1 t_0) & \cdots & \cos(2\pi f_{N_f} t_0) & \sin(2\pi f_{N_f} t_0) \\ \cos(2\pi f_1 t_1) & \sin(2\pi f_1 t_1) & \cdots & \cos(2\pi f_{N_f} t_1) & \sin(2\pi f_{N_f} t_1) \\ \vdots & \vdots & & \vdots & \vdots \\ \cos(2\pi f_1 t_{N_{\text{toas}}}) & \sin(2\pi f_1 t_{N_{\text{toas}}}) & \cdots & \cos(2\pi f_{N_f} t_{N_{\text{toas}}}) & \sin(2\pi f_{N_f} t_{N_{\text{toas}}}) \end{bmatrix}, \quad (1.35)$$

and

$$a = \begin{bmatrix} a_1 \\ b_1 \\ a_2 \\ b_2 \\ \vdots \\ a_{N_f} \\ b_{N_f} \end{bmatrix}, \quad (1.36)$$

Where in F , each frequency f_n is defined to be

$$f_n = \frac{n}{T}. \quad (1.37)$$

where T is the total observing time span of either the pulsar (if we are considering intrinsic pulsar noise) or the entire PTA (when considering the SGWB). The column vector a contains all the Fourier coefficients that parameterize the Gaussian noise process. Following Nyquist's theorem, the high-frequency cut off is bounded by the cadence of pulsar timing observations, and the low-frequency cutoff is lim-

ited by the total timing baseline such that $f_1 = 1/T$. However, in practice, the high frequency cutoff is selected based on how much of the power spectrum will be contained within the frequency cutoffs for both the pulsar-intrinsic red noise and global GWB red noise.

Recall that the SGWB is expected to have a power spectrum consistent with red noise (as in Equation 1.29 and Equation 1.32). The intrinsic noise exhibited by MSPs has been shown to have very similar properties that can be similar to the SGWB power spectrum [105]. As a result, it is necessary to model both noise processes independently. It is common to model the intrinsic spin-noise of MSPs as a red noise process. Unlike the SGWB, there is no expected value for the amplitudes and spectral indices of these noise processes in different pulsars. Quantitatively, we use a similar red noise power law for the spin-noise as we use for the SGWB defined by some amplitude and spectral index

$$P_{\text{psr}}(f) = A_{\text{psr,1yr}}^2 \left(\frac{f}{f_{1\text{yr}}} \right)^{-\gamma}. \quad (1.38)$$

The second-to-last term on the right-hand side of Equation 1.34 ($M\vec{\epsilon}$) accounts for errors in the timing model for each pulsar. When the full TOA series for each pulsar is created, the timeseries is fit with a timing solution which estimates many different parameters that account for variations in pulse TOAs such as changes to the rotational period and period-derivative, parallax, proper motion, binary parameters (for pulsars in binary systems), and more. The parameters of these complex timing solutions have some uncertainties. We assume that the timing

parameters are close to the “true” values and that we can Taylor expand linearly around the parameter values chosen by the timing solution. The vector $\vec{\epsilon}$ represent the variations of each timing parameter, and the matrix M is called the “design matrix” and gives the basis which determines how small variations in timing parameters affect the full TOA series. For example, in a model in which we allow for variations in initial pulsar rotational phase, frequency and frequency-derivative, the design matrix M and parameter errors $\vec{\epsilon}$ schematically look like:

$$M = \begin{bmatrix} 1 & t_0 & t_0^2 \\ 1 & t_1 & t_1^2 \\ \vdots & & \\ 1 & t_{N_{\text{TOA}}} & t_{N_{\text{TOA}}}^2 \end{bmatrix}, \quad (1.39)$$

and

$$\vec{\epsilon} = \begin{bmatrix} \omega_0 \\ \omega_1 \\ \omega_2 \end{bmatrix}. \quad (1.40)$$

This corresponds to fitting out constant, linear, and quadratic contributions which arise from having incorrect values for the three aforementioned parameters in the timing solution.

The final term in the time series represents any deterministic contributions induced by a GW signal, such as nonlinear GW memory.

Using this residual model, it is possible to construct a likelihood for all model parameters described above. Let us call the collection of parameters that describe

the full residual model $\vec{\theta}$. Then, we construct the likelihood by computing the probability that the residuals minus the modeled processes gives a Gaussian white noise time series. In other words, from Equation 1.38, given a set of model parameters $\vec{\theta}$, we compute the corresponding contributions (denoted by primed quantities) to the total residual model and subtract them from the full residual model to form a guess for the underlying white noise process

$$\vec{\mathbf{r}} - F\vec{\mathbf{a}}' - F_{\text{GW}}\vec{\mathbf{a}}'_{\text{GW}} - M\vec{\boldsymbol{\epsilon}}' - \vec{\boldsymbol{\delta}}'_{\text{det}} = \vec{\mathbf{n}}'. \quad (1.41)$$

Thus, the likelihood that these $\vec{\theta}$, denoted by $p(\vec{\mathbf{r}}|\vec{\theta})$ below, accurately describe the residuals is

$$p(\vec{\mathbf{r}}|\vec{\theta}) = \left(\frac{1}{\sqrt{\det(2\pi N)}} \right) \exp\left(\vec{\mathbf{n}}'^T N^{-1} \vec{\mathbf{n}}' \right), \quad (1.42)$$

where above, the matrix N is a white-noise covariance matrix that describes the underlying white noise process. In practice, this turns out to be computationally impractical. The high dimensionality of total residual model means that it is impossible in practice to explore the entire parameter space when performing parameter estimation or model selection. Fortunately, it is possible to simplify the parameter space by marginalizing over some nuisance parameters. First, we make the remark that the Fourier coefficients used to model the SGWB and intrinsic pulsar spin noise are less interesting to GW astronomers than the parameters of the power spectra which describe them. This is because the dominant belief is that the background is completely stochastic, and the details of one particular realization

of the background (in the form of the Fourier coefficients), and more interested in the statistical characteristics of the background encoded by the hyperparameters. As a result, we can parameterize the red noise processes using the hyperparameters A and γ from Equation 1.38 and Equation 1.32. This makes it possible to marginalize over the independent Fourier coefficients using hyperpriors informed by these hyperparameters. Additionally, it has been shown that it is possible to analytically marginalize over timing model errors (carried by $\vec{\epsilon}$), further reducing the dimensionality of the parameter space. We will skip over the mathematical details, but more detail on these two elegant “tricks” are discussed in detail in Arzoumanian et al. [12], Lentati et al. [76]. The resulting marginalized likelihood used in practice is

$$L(\vec{\theta}) = p(\vec{r}|\vec{\theta}) = \left(\frac{1}{\sqrt{\det(2\pi C)}} \right) \exp \left((\vec{r}' - \vec{\delta t}'_{\text{det}})^T C^{-1} (\vec{r}' - \vec{\delta t}'_{\text{det}}) \right) \quad (1.43)$$

where the new covariance matrix is defined

$$C = N + TBT^T \quad (1.44)$$

with N being the white noise covariance matrix, T defined by

$$T = \begin{bmatrix} M & F & F_{\text{GW}} \end{bmatrix}, \quad (1.45)$$

and B defined to be a matrix which gives the covariances for the parameters which

we marginalize over and encode with hyperparameters

$$B = \begin{bmatrix} \infty & 0 & 0 \\ 0 & \phi & 0 \\ 0 & 0 & \phi_{\text{GW}} \end{bmatrix}, \quad (1.46)$$

where the ∞ gives the priors on the timing model errors, and ϕ and ϕ_{GW} are the Gaussian covariance prior matrices on the Fourier coefficients which describe pulsar spin noise and the SGWB-induced red noise informed by the amplitudes and spectral indices of the two processes.

In this discussion, we have provided a very brief description of the full marginalized likelihood used in PTA data analysis. The key point is that this form of the likelihood eliminates consideration of the timing model parameters and encodes the intrinsic pulsar spin noise and SGWB using only the hyperparameters required to characterize a power spectral density in the form of a red noise power law. Statistical inferences made about pulsar spin noise and the SGWB are mostly in the form of constraints on these hyperparameters.

1.5 Statistical Inference in PTAs

In this work, we often use Bayesian inference to perform two different kinds of analyses: model selection and parameter estimation. Bayesian inference relies on the use of Bayes' theorem to evaluate the posterior probability of model parameters conditioned on observed data. Consider a model \mathcal{M} with model parameters $\vec{\theta}_{\mathcal{M}}$

and some observed data $\vec{\mathbf{d}}$. Bayes' theorem gives the posterior probability for a set of model parameters

$$p(\vec{\boldsymbol{\theta}}_{\mathcal{M}}|\vec{\mathbf{d}}) = \frac{p(\vec{\mathbf{d}}|\vec{\boldsymbol{\theta}}_{\mathcal{M}})p(\vec{\boldsymbol{\theta}}_{\mathcal{M}})}{p(\vec{\mathbf{d}})}, \quad (1.47)$$

where the left-hand-side term is called the “posterior” probability of $\vec{\boldsymbol{\theta}}_{\mathcal{M}}$, $p(\vec{\mathbf{d}}|\vec{\boldsymbol{\theta}}_{\mathcal{M}})$ is called the “likelihood” (which we have shown an example of in Equation 1.43 above), $p(\vec{\boldsymbol{\theta}}_{\mathcal{M}})$ is called the “prior”, and $p(\vec{\mathbf{d}})$ is called the “evidence”. The posterior probability gives the probability distribution for the model parameters conditioned on the observed data. The likelihood gives the probability that the observed data resulted from the model parameters being evaluated. The prior encodes any previous knowledge about the model parameters; often in PTA data analysis, the priors used are ignorance priors in the form of uniform or log-uniform priors which are uniform over the parameter space or uniform over each order of magnitude of the parameter space, respectively. Finally, the evidence is the probability of having observed this data integrated over all possible values of the model parameters. In Bayesian parameter estimation problems, however, this term is constant and has the function of normalizing the posterior probability distribution. For **parameter estimation problems**, these posterior distributions can be estimated using a Monte-Carlo sampling technique. The resulting posteriors then represent the possible parameter values conditioned on observed data.

These posterior probability distributions can also be used to compute odds ratios for different models in **model selection problems** (e.g. selecting between

a model with a global HD-correlated SGWB versus a model with only intrinsic pulsar red noise, or selecting between models including or excluding a nonlinear GW memory event at some epoch). We can compute the posterior for some i -th model $p(\mathcal{M}_i|\vec{d})$ using Bayes' theorem.

$$p(\mathcal{M}_i|\vec{d}) = \frac{p(\vec{d}|\mathcal{M}_i)p(\mathcal{M}_i)}{p(\vec{d})}, \quad (1.48)$$

where the likelihood $p(\vec{d}|\mathcal{M}_i)$ is

$$p(\vec{d}|\mathcal{M}_i) = \int d\vec{\theta}_i p(\mathcal{M}_i|\vec{d}, \vec{\theta}_i)p(\vec{\theta}_i|\mathcal{M}). \quad (1.49)$$

Note that in this likelihood, we have integrated over all possible values of the model parameters $\vec{\theta}_i$ for the model \mathcal{M}_i so that this is now a global likelihood for this model. The odds ratio O_{ij} in favor of \mathcal{M}_i over \mathcal{M}_j is then the ratio of the posterior probabilities for the two models

$$O_{ij} = \frac{p(\mathcal{M}_i|d)}{p(\mathcal{M}_j|d)} = \frac{p(\vec{d}|\mathcal{M}_j)p(\mathcal{M}_i)}{p(\vec{d}|\mathcal{M}_j)p(\mathcal{M}_j)} \equiv B_{ij} \frac{p(\mathcal{M}_i)}{p(\mathcal{M}_j)}. \quad (1.50)$$

The likelihood ratio B_{ij} is called the ‘‘Bayes factor’’. In model selection problems in which each model is equally preferred by the prior, the Bayes factor is exactly equal to the odds ratio.

For both parameter estimation and model selection in PTA searches for GWs, the posterior probabilities are computed using Markov-Chain Monte Carlo (MCMC) sampling methods. In general, MCMC samplers can obtain samples from a target

probability distribution that may be otherwise infeasible to sample from. In this case, the target distribution is either the multivariate posterior probability distribution of model parameters for parameter estimation problems, or the posterior distribution over signal-model space. It has been shown that MCMCs are capable of directly computing odds ratios of different models through product-space sampling [77].

The common algorithm used for MCMC sampling in PTA data analysis is the Metropolis-Hastings algorithm. This algorithm uses the following steps to compute some target distribution $P_{\text{target}}(\vec{\theta})$:

1. Pick some random initial point in parameter space $\vec{\theta}$.
2. Use a proposal distribution $\pi(\vec{\theta}'|\vec{\theta})$ to choose a new point in parameter space
3. Accept or reject the new point $\vec{\theta}'$ with an acceptance probability defined by

$$p_{\text{acc}}(\vec{\theta}', \vec{\theta}) = \min \left(1, \frac{p_{\text{target}}(\vec{\theta}')\pi(\vec{\theta}|\vec{\theta}')}{p_{\text{target}}(\vec{\theta})\pi(\vec{\theta}'|\vec{\theta})} \right)$$

4. If the step is accepted, record the proposed point $\vec{\theta}'$ and replace the current point with the accepted new point $\vec{\theta} \rightarrow \vec{\theta}'$; otherwise, record the original point $\vec{\theta}$.
5. Repeat creating and accepting/rejecting proposals (steps 2 and 3) until there are enough recorded points in parameter space to resolve the target distribution $p_{\text{target}}(\vec{\theta})$

This algorithm is guaranteed to asymptotically approach the target distribution. If the acceptance ratio shown above is chosen to be the posterior probability ratio, then the target distribution is simply the full posterior probability distribution. Additionally, this algorithm is very efficient for sampling high-dimensional parameter spaces. For PTA analyses, which can require order 100-dimensional parameter spaces, this efficiency is key. There are additional features which can be used to further increase the efficiency, such as adaptive jump proposals and parallel tempering.

While MCMC sampling using an adaptive Metropolis-Hastings algorithm is generally popular, there are also alternative sampling algorithms, such as Hamiltonian Monte Carlo [58] and Gibbs sampling [72]. Each sampling algorithm has its own advantages and disadvantages, and the choice of sampling technique is best left determined by the problems involved in each data analysis problem. Additionally, for some GW analyses, it is possible to define elegant frequentist statistics which are generally faster to compute (see Refs. [7] and [46] for examples of frequentist statistics for the SGWB and continuous GWs). Useful frequentist statistics can take some work to construct. However, they are often useful as additional cross-validation tests against Bayesian results. They are also often fast to compute. We will present a frequentist method for nonlinear memory in Chapter 5.

In the remainder of this work, we will summarize the state of the literature in nonlinear GW memory in Chapter 2. Then, we will present the we will discuss a new expedited methodol for searching pulsar timing data for nonlinear GW memory in Chapter 3. We will then discuss results of the first search for nonlinear

memory in a data set which has been shown to include evidence for a global red noise process that is expected to be the SGWB that uses the aforementioned method in Chapter 4. Next, we will discuss the adaptation of a fast frequentist statistical technique designed for continuous GW searches to searches for nonlinear GW memory in Chapter 5. Finally, we conclude with a summative discussion of the future of searches for nonlinear memory in pulsar timing data.

Chapter 2: Literature Review

Since nonlinear memory can result from any GW-emitting event, it can serve as a catch-all for many exotic types of GW-emitting physical processes [38]. Additionally, it showcases a very interesting higher-order nonlinear effect of general relativity. Nonlinear memory does not appear in linearized theories. However, despite originating from a higher order expansion term (2.5 post-Newtonian), it miraculously contributes to the same order as effects from linearized gravity. Thus, any detection of nonlinear memory would serve as a very powerful validation of gravitational theory.

2.1 Prediction

Nonlinear memory was first predicted by Blanchet and Damour [19], Christodoulou [32] independently. Blanchet and Damour [19] and Thorne [115] understood this final permanent DC effect to be a hereditary result of all prior GW emission. Thorne [115] presents the metric perturbation formulated in terms of the energy emitted through GWs

$$\Delta h^{ij} = \frac{4}{r} \int d\Omega \frac{dE}{d\Omega} \frac{\Omega^i \Omega^j}{1 - \cos \theta_{\Omega, \text{det}}}, \quad (2.1)$$

where Ω is the solid angle around the source, $\hat{\Omega}$ is the unit vector in the direction of $d\hat{\Omega}$, r is the comoving distance from the source to the detector, and $\theta_{\Omega,\text{det}}$ is the angle between the detector and $\hat{\Omega}$. This is exactly the result found by Blanchet and Damour [19], Christodoulou [32] with the connection between emitted GWs and memory made explicit. Thorne [115] estimates that the order of magnitude of this effect is approximately

$$\Delta h \sim 0.1 \frac{E_{\text{em}}}{r}, \quad (2.2)$$

where E_{em} is the total energy emitted through GWs. This means that any strong transient GW event such as the coalescence of compact objects may come with an additional observable memory signal. Additionally, even strong events which are out of a GW detector’s band width may still produce a memory signal with no obvious source [93]. Such a signal is called “orphan memory”.

Kennefick [71], Wiseman and Will [124] evaluate Equation 2.1 for the inspiral phase of a circularized binary black hole system. Favata [53] extends this result using an effective one-body approach (Buonanno and Damour [24, 25]) calibrated to results of numerical relativity simulations to compute the mass quadrupole past the inspiral phase to show that there are significant contributions to the memory from the merger phase. Additionally, Favata [53] uses an analytic minimum-waveform model which continuously transitions the time-dependent mass quadrupole from the inspiral phase into the merger phase to compare with the numerical result from the effective one-body approach and finds that the memory saturates at similar values. However, there is a so-called “fudge factor” difference of 0.77 between the

two approaches, with the minimum-waveform model over-estimating compared to the numerical results (see Figure 1 in Favata [53]).

Early discussions of nonlinear memory were contextualized in ground-based and space-based GW detectors like LIGO and LISA. Then, following the reinvigoration of interest in nonlinear GW memory starting from Favata [53], Pshirkov et al. [100], Seto [104], van Haasteren and Levin [118] simultaneously release work regarding the prospects of finding nonlinear memory in PTA data. Seto [104] assumes that detection of nonlinear memory will occur in one of two cases: 1) PTA data is dominated by white noise and the most sensitive frequency bin is defined by $f_{\text{opt}} = 1/T_{\text{obs}}$; 2) PTA data is dominated by the red-noise SGWB, in which case the optimal frequency bin is the frequency at which the red noise and white noise spectra intersect (see Fig. 1 of Seto [104]). Using the Parkes PTA and planned Square Kilometer Array (SKA) sensitivities as representative of contemporary and next-generation projects, the expected number of detections of memory events is approximately 10^{-6} for each project (See Fig. 2 in Seto [104]). A very notable result from this work shows that if the PTA data set is SGWB-dominated, then the number of detectable events decreases.

This key fact is something that some contemporary papers (such as Pshirkov et al. [100], van Haasteren and Levin [118]) ignore in preliminary calculations of signal-to-noise ratios of memory in pulsar timing data. However, other work which does account for extraneous red noise corroborate the results of Seto [104]. Cordes and Jenet [35] show that “red noise has a strong deleterious effect on burst detection”. Using two different sets of simulations, one of which has only $100ns$ of

RMS white noise, and one which has a total of $100ns$ of RMS noise split equally between red and white noise, their changepoint analysis becomes much less reliable at estimating the epoch and amplitude of an injected memory signal (see Figure 1 of Cordes and Jenet [35]). Madison et al. [82] use simulated data to verify the theoretical predictions made in van Haasteren and Levin [118] regarding the sensitivity of PTAs to memory in white-noise dominated timing data. While van Haasteren and Levin [118] do not include effects of red noise, Madison et al. [82] extend their results to data sets including both types of noise and show that, all else held constant, the simple inclusion of red noise increases the uncertainty of any estimates of the memory amplitude by a factor of ~ 5 . This work shows that early estimates of the detectability of nonlinear memory in PTA data were too optimistic. Future searches for nonlinear memory will have limited sensitivity unless they can be made robust to red noise, especially considering recent results from the PTA community, which show strong evidence for a global red-noise signal in PTA data [1, 128].

2.2 Search Methodologies and Detection Prospects for Memory in PTA Data

Earlier searches for memory in PTAs used frequentist techniques to fit for memory signals. Wang et al. [123] define the memory source parameters $A_1 = h_{\text{mem}} \cos(2\psi)$ and $A_2 = h_{\text{mem}} \sin(2\psi)$ in addition to burst epoch and source sky position. They include these parameters in the `Tempo2` [65] timing model and use the built-in

linear least squares fit to minimize the Gaussian white noise. The amplitude parameters A_1 and A_2 are correlated, and so they compute a decorrelator matrix U^{-1} to obtain statistically uncorrelated amplitude estimates A_{1w} and A_{2w} . Then, the statistic $D = A_{1w}^2 + A_{2w}^2$ follows a χ^2 distribution. Then, the maximum value of D computed over the entire sky at each trial epoch gives the maximally loud memory signal. Note that this is a coherent fit because it fits one global signal to residuals in each pulsar simultaneously using the antenna responses determined by the pulsar locations relative to the source orientation.

Arzoumanian et al. [11] similarly uses **Tempo2** [65] to fit for a memory signal in the pulsar timing models. Additionally, they perform two searches; one is an incoherent search for memory in each individual pulsar’s residual time series (the “pulsar-term search”), and the other looks for a global signal coherently using the deterministic antenna response of the PTA to a memory signal (the “Earth-term search” similar to Wang et al. [123]). For the pulsar-term search, rather than maximizing the sum of the squared amplitude parameters $A_{1w}^2 + A_{2w}^2$, Arzoumanian et al. [11] compute the ratio of likelihoods for the residuals with the additional memory signal in the timing model to the residuals without the memory signal (see Equation (12) in Arzoumanian et al. [11]). This likelihood is χ^2 distributed, and may be used to compute the maximally likely memory event in each pulsar individually. For the Earth-term search, Arzoumanian et al. [11] recognize that the methods used in Wang et al. [123] work well for smaller data sets, but it may become too computationally intensive for larger PTA data sets. As a result, Arzoumanian et al. [11] propose an accelerated Earth-term likelihood calculation

relying on the factorizability of the PTA likelihood. This method allowed them to quickly compute a global Earth-term likelihood surface over the entire memory signal parameter space.

The PTA community has since moved towards using Bayesian inference to estimate GW parameters in the form of Markov-Chain Monte Carlo (MCMC) sampling. The most recent search for nonlinear memory in the NANOGrav 11-year data set was performed by using MCMC methods to compare various signal models and compute upper limits on GW signal parameters Aggarwal et al. [5]. MCMC samplers use a sampling scheme that allow them to numerically compute target probability distributions (in the case of PTAs, the posterior probability distribution for different noise and signal parameters) in very high-dimensional spaces very efficiently [94]. As such, for PTA data analysis, in which many noise and signal model parameters must be fit simultaneously, these methods allow for more efficient parameter estimation. These samplers can also include an index for different signal models. These additional indices allow for efficient model comparisons through product space sampling in addition to parameter estimation [28]. Notably, however, some GW searches, such as searches for continuous GWs and nonlinear memory, have parameter spaces that are large enough that MCMC methods remain onerous on a more absolute scale.

Finally, there have also been some recent efforts to develop a method to search for memory using astrometric deflection. Over long timescales, GWs can cause permanent deflections of the apparent locations of visible light sources in the universe [21]. These results have been extended to gravitational wave memory. Preliminary

calculations show that it may be possible to detect nonlinear memory by observing changes to the proper motions of distant light sources [22, 84]. This method may be promising, as there are many more distant sources of light than there are pulsars. The increased number of pairs means that it will be much easier to reduce excess noise when searching for cross-correlated power. However, the astrometric deflections are very small, and this technique will still benefit from improvements to current telescope sensitivity.

2.3 Current Limits on Nonlinear Memory

Since the discovery that nonlinear memory is potentially observable by PTAs, PTAs have published many results of searches for this signal. Thus far, there have been no detections. However, two PTAs have published upper limits on memory signals. Wang et al. [123] apply the D -statistic described above in §2.2 to the initial Parkes Pulsar Timing Array (PPTA) data release (Manchester et al. [90]) and find a maximum D -statistic of $D_{\max} = 12.4$, which corresponds to a false-alarm probability of 15%. They also compute rate event upper limits and find that the upper limits on event rates set in this search are still orders of magnitude greater than estimated event rates. In other words, the sensitivity of the PPTA was shown to be low enough to be unlikely to make a detection of a memory event. They find a global upper limit of $h_{\text{mem}} \leq 2.4 \times 10^{-13}$, with an event rate upper limit of 0.75 events yr^{-1} [123].

Similarly, Aggarwal et al. [5] find no compelling evidence for gravitational wave

memory in the NANOGrav 11-yr data set [14]. They also find a global maximum upper limit of $h_{\text{mem}} \lesssim 1.5 \times 10^{-13}$ (estimate based on the color map in Figure 4 of Aggarwal et al. [5]). These upper limits are still very far away from astrophysical estimates of merger rates. Islo et al. [67] finds that mergers which produce memory events with $\Delta h > 5 \times 10^{-16}$ are extremely rare, ranging from $\sim 1 \times 10^{-3}$ mergers per century to fewer than one merger per Hubble time, depending on the choice of merger model. Cordes and Jenet [35] estimate that the rate of mergers with memory signals $\Delta h > 10^{-15}$ to be up to 0.02 events per year. Overall, the estimates for merger rates vary because of uncertainties about the physics involved in SMBHB mergers, such as how SMBHBs evolve into a GW-dominated regime and dark matter halo interactions. However, it suffices to say that even optimistic estimates of merger rates predict that it will be unlikely for PTAs to make many detections of nonlinear memory events.

Chapter 3: Implementation of an Efficient Bayesian Search for
Gravitational-wave Bursts with Memory in Pulsar Timing Array
Data

Jerry P. Sun, Paul T. Baker, Aaron D. Johnson, Dustin R. Madison, Xavier
Siemens

The Astrophysics Journal
1667 K Street NW, Suite 800
Washington, DC 20006 USA
Volume 951, Issue 2

3.1 Abstract

The standard Bayesian technique for searching pulsar timing data for gravitational wave (GW) 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.

3.2 Introduction

Millisecond pulsars (MSPs) have very stable rotations. Because the rotational period is so stable, it is possible to detect small deviations in 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 [57, 64, 89, 92]. Pulsar timing arrays (PTAs) are expected to be able to use TOA data from many MSPs to either detect or provide constraints on GWs [44, 103].

One signal of interest is a gravitational wave 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 [32, 115]. In particular, it is expected that mergers of super massive black hole binaries (SMBHBs) 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. 67]. Additionally, because all GW events leave behind GW memory, detections of BWMs could lead to discoveries of new sources of GWs [38].

A GW BWM passing over an Earth-pulsar pair will shift the pulsar’s observed rotational frequency [e.g., 118]. 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 [35, 67, 82, 100, 104, 118]. The observed rotational frequency may change to be either faster or slower depending on the orientation and polarization of the memory wavefront, 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 [11, hereafter NG5-bwm] to expedite the Bayesian methods used in the NANOGrav 11 yr search for BWMs [5, hereafter NG11-bwm]. This search was performed on the NANOGrav 11 yr data set [14]. 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 3.3, we describe the effect of a BWM on the TOA residuals of a pulsar. In section 3.4, 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 3.5, 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.

3.3 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 [55, 82, 118]. 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 (θ, ϕ) . The event wavefront propagates in the direction $\hat{\mathbf{k}}$, with strain h_0 , passing over the Earth from which we observe a pulsar located at position \mathbf{p} . The memory wavefront has the principal polarization vector $\hat{\boldsymbol{\psi}}$ described by an angle ψ , which gives the principal polarization direction of the wavefront relative to an orthonormal basis $(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\beta}})$. In particular, the principal polarization vector $\hat{\boldsymbol{\psi}}$ is defined:

$$\hat{\boldsymbol{\psi}} = \hat{\boldsymbol{\delta}} \cos \psi + \hat{\boldsymbol{\beta}} \sin \psi \quad (3.1)$$

Following NG5-bwm and NG11-bwm, the perturbation to pulse times of arrival from this pulsar, δt_{bwm} , may be modeled as:

$$\delta t_{\text{bwm}}(t) = B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) h_{\text{mem}}(t), \quad (3.2)$$

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 [51, 64]. The geometric projection factor is:

$$B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{1}{2} \cos(2\psi_{\hat{\mathbf{p}}})(1 - \cos \alpha), \quad (3.3)$$

where α 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 $(\hat{\boldsymbol{\delta}}, \hat{\boldsymbol{\beta}})$ plane:

$$\alpha = \cos^{-1}(\hat{\mathbf{p}} \cdot \hat{\mathbf{k}}) \quad (3.4)$$

$$\psi_{\hat{\mathbf{p}}} = \tan^{-1} \left(\frac{\hat{\mathbf{p}} \cdot \hat{\boldsymbol{\beta}}}{\hat{\mathbf{p}} \cdot \hat{\boldsymbol{\delta}}} \right) - \psi \quad (3.5)$$

We include this description for completeness, and a diagram with more details may be found in Madison et al. [82]. For this analysis, we simply characterize the polarization using the principal polarization angle ψ . The time-dependent strain

term is [100, 118]:

$$h_{\text{mem}}(t) = h_0[(t - t_0)\Theta(t - t_0) - (t - t_p)\Theta(t - t_p)] \quad (3.6)$$

where h_0 is the strain of the memory; t_0 is the time that the memory wave front passed over the Earth; $t_p = t_0 + (|\mathbf{p}|/c) [1 + \cos(\alpha)]$, the time at which the memory wave front passed over the pulsar, with \mathbf{p} and α still defined to be the position of the pulsar and angle between pulsar and BWM propagation direction, as in Equation 3.3. Additionally, $\Theta(t)$ is the Heaviside function. Because each pulsar in NANOGrav’s 11-year data release (and more generally, current PTA datasets) is on the order of thousands of light-years away from the Earth, and total observation times are order tens of years, we only expect that one of the two terms in Equation 3.6 will be nonzero. The first term in Equation 3.6 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`¹ [113]

¹https://github.com/nanograv/enterprise_extensions

3.4 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-year data set for GW BWMs. Then, we will discuss the the adaptation of the techniques used in NG5-bwm that expedites both the pulsar- and Earth-term searches.

3.4.1 Bayesian Approach

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

$$\delta\mathbf{t} = \mathbf{s} + T\mathbf{b} + \mathbf{n}, \quad (3.7)$$

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 [14]. 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. $T\mathbf{b}$ are the contributions to the residuals from any Gaussian processes. In this paper, we consider two different Gaussian processes, and so our T -matrix and \mathbf{b} vector may be broken down into:

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

where above, M is the design matrix for the linearized timing model that accounts for uncertainty in the residuals from an imperfect timing model fit ϵ . 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 [99]:

$$P(f_j) = A_j^2 \left(\frac{f_j}{\text{yr}^{-1}} \right)^{-\gamma}, \quad (3.8)$$

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 yr^{-1} , and γ is the spectral index of the power law.

From Equation 3.7, we can construct an approximation of the Gaussian white noise given an estimation of the model parameters:

$$\mathbf{n} = \delta\mathbf{t} - \mathbf{s} - T\mathbf{b}. \quad (3.9)$$

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 residuals as:

$$p(\mathbf{n}) = \frac{\exp\left(-\frac{1}{2}\mathbf{n}^T N^{-1}\mathbf{n}\right)}{\sqrt{2\pi \det N}} \quad (3.10)$$

where \mathbf{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 is Gaussian white noise. In other words:

$$\begin{aligned} p(\delta\mathbf{t}|\mathbf{b}, \mathbf{s}) &= p(\delta\mathbf{t}|\mathbf{b}, h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) \\ &= \frac{\exp[-\frac{1}{2}(\delta\mathbf{t} - \mathbf{s} - T\mathbf{b})^T N^{-1}(\delta\mathbf{t} - \mathbf{s} - T\mathbf{b})]}{\sqrt{2\pi \det N}}, \end{aligned} \quad (3.11)$$

where we have explicitly written out the parameters that determine \mathbf{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 3.11 over the parameters that describe the Gaussian processes and reduce the dimensionality of the parameter space [76, 119, 120]. The reduced likelihood is:

$$p(\delta\mathbf{t}|h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{\exp(-\frac{1}{2}\mathbf{q}^T C^{-1}\mathbf{q})}{\sqrt{2\pi \det C}} \quad (3.12)$$

where

$$\mathbf{q} = \delta\mathbf{t} - \mathbf{s} \quad (3.13)$$

where C is defined as

$$C = N + TDT^T \quad (3.14)$$

and D is defined as

$$D = \begin{bmatrix} \infty & 0 \\ 0 & \phi \end{bmatrix}, \quad (3.15)$$

where ∞ is a diagonal matrix of infinities, which effectively give unconstrained priors on the timing model parameters, and ϕ is a diagonal matrix containing the red noise power at each frequency bin in Equation 3.8. The Woodbury [125] 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 ENTERPRISE² [49] 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 PTMCMCSAMPLER³ package [47].

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

²<https://github.com/nanograv/enterprise>

³<https://github.com/jellis18/PTMCMCSampler>

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 towards 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 which exactly cancels the selection effect, resulting in a posterior that is uniform in source-orientation. This is related to the technique used in Malmquist [88].

More specifically, to place an amplitude UL as a function of burst epochs, NG11-bwm created 48 HEALPIX⁴ [61] sky bins using `healpy`⁵ [132], with 8 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,

⁴<https://healpix.jpl.nasa.gov/>

⁵<https://github.com/healpy/healpy>

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 upper limit may be computed from the marginalized amplitude posterior for each sky position.

For a summary of the priors, see Table 3.2.

3.4.2 Accelerated Bayesian Search

For the accelerated Bayesian search, we mimic the Bayesian approach described in subsection 3.4.1 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:

$$\begin{aligned} p(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B) &= \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B) \\ &= \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|h_i, t_B), \end{aligned} \tag{3.16}$$

and:

$$h_i = B(\hat{\mathbf{k}}, \hat{\mathbf{p}}_i, \psi) \times h_{\text{mem}}, \tag{3.17}$$

Table 3.1.

Parameter	Prior	Description
$\log_{10} A_{\text{rm}}$	LinearExp(-17, -11)	Amplitude of intrinsic pulsar red noise
γ_{rn}	Uniform(0, 7)	Spectral index of intrinsic pulsar red noise
$\log_{10} A_{\text{BWM}}$	LinearExp(-17, -10)	Amplitude of global BWM
ψ_{BWM}	Uniform(0, π)	Polarization of BWM
θ_{BWM}	Uniform(0, π)	Polar angle of BWM source
ϕ_{BWM}	Uniform(0, 2π)	Azimuthal angle of BWM source
t_{BWM}	Uniform(MJD 56000, MJD 57000) Uniform(MJD 53216, MJD 57387)	Earth-term BWM epoch

Table 3.2. Priors used for each of the model parameters in the Bayesian search for global Earth-term GW BWMs using the full PTA. 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 subsection 3.4.1. The prior on t_{BWM} also varies depending on the particular upper limit calculation. For upper limits (ULs) as a function of sky location, we use priors between MJD 56000 and MJD 57000. For upper limits 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 subsection 3.4.2.

where $p(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_B)$ is the global likelihood of a burst propagating in the direction $\hat{\mathbf{k}}$, polarization ψ , 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 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_{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 pre-compute the individual pulsar-term likelihoods over a grid of only post-projection BWM amplitude and burst epoch without losing any information. Then at runtime, the geometric projection factor, Equation 3.3, 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 pre-computed lookup tables and combine them using Equation 3.16.

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 3.16). In addition to the BWM amplitude $|h_i|$, epoch t_B , and the sign of h_i , we include the amplitude A_{rn} and spectral index γ of a red-noise process described in Equation 3.8. 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 degrees 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 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 which contain the Bayesian likelihoods of finding an Earth-term trial burst with fixed sky position (θ, ϕ) , polarization ψ , 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 3.16. We compute one two-dimensional lookup table varying over trial bursts characterized by (h_0, t_0) for each set of trial parameters (θ, ϕ, ψ) .

We can then construct global amplitude posteriors as a function of sky position and epoch by integrating out any nuisance parameters against their prior distribu-

tions. We do so in the same way as in NG11-bwm (described in subsection 3.4.1); 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 samples from each source-orientation bin to compute the full-sky, polarization-marginalized 95% upper limits as a function of burst epoch.

To compute the ULs as a function of sky location, we use 768 HEALPix skypixels 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:

1. Compute pulsar-term BWM likelihoods on a grid of

$$\{\log_{10} |h_i|, \text{sign}(h_i), t_0, \log_{10} A_{\text{rn}}, \gamma_{\text{rn}}\}$$

2. Marginalize pulsar-term BWM likelihoods over red noise parameters.
3. Use pulsar-term likelihoods, Equation 3.16 and Equation 3.17 to compute Earth-term BWM likelihoods on a grid of $\{\log_{10} h_0, t_B\}$ for each set of trial burst parameters θ, ϕ, ψ
4. Marginalize over:
 - (a) Burst epoch and polarization to compute amplitude posterior over sky location
 - (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:

$$\begin{aligned} p_{\log\text{-uni}}(A_{\text{bwm}}|d) &\propto p(d|A_{\text{bwm}}) \pi_{\log\text{-uni}}(A_{\text{bwm}}) \\ &\propto p(d|A_{\text{bwm}}) \frac{1}{A_{\text{bwm}}} \end{aligned} \tag{3.18}$$

where $\pi(A_{\text{bwm}}) \propto \frac{1}{A_{\text{bwm}}}$ is the prior distribution on A_{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.

3.5 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.




Figure 3.1: Percent difference in the pulsar term amplitude upper limits (ULs) for each pulsar used in [5]. These percent differences are computed by comparing the 95% ULs from Bayesian MCMC runs and lookup-table-based methods for positive- 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 upper limits of positively-signed memory, and the blue points are the differences for negatively-signed memory, and the error bars show the 95% confidence intervals computed from bootstrapping.

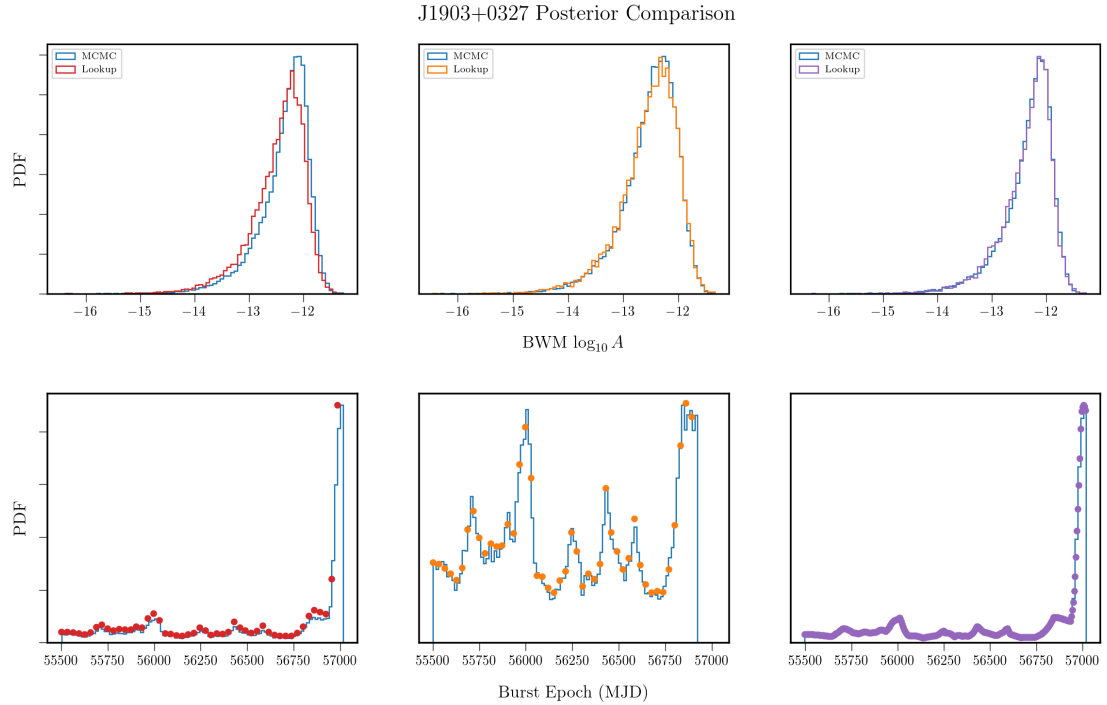


Figure 3.2: Marginal amplitude and burst epoch posteriors for three different sets of lookup tables. The leftmost pair show the marginal amplitude and burst epoch posteriors with 180 days excluded from the beginning and end of the epoch prior. The marginal posteriors computed using lookup tables are in red, and the marginal posteriors computed using the MCMC sampler are in red. 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 data set 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 which has enough grid points to accurately characterize the features in the burst epoch posterior also gives very good agreement of the amplitude posteriors.

3.5.1 Pulsar-Term Comparisons

Figure 3.1 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 upper limits separately in order to fully compare the two techniques.

Additionally, we narrowed the priors on 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 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 3.2 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 posterior (blue). We can see that without any special

adjustments made, the last two points of the burst epoch grid do not sufficiently characterize the posterior, and the resulting marginal amplitude posterior is biased towards high amplitudes. However, both the exclusion of early and late trial burst epochs (orange, center) and using a denser grid (purple, right) give agreement between the marginal amplitude posteriors.

3.5.2 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 Figure 3.3 and Figure 3.4, respectively.

In Figure 3.3, 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 non-constraining UL at these early epochs. More importantly, as more pulsars and more data are added to the PTA, the ULs become nearly identical.

Figure 3.4 shows the ULs on the BWM amplitude as a function of sky location using the method described in subsection 3.4.2. 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 NANOGrav 11 yr dataset, 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-year 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.

3.5.3 Computational Improvement

The computational complexity of computing the pulsar-term lookup tables is dominated by the cost of inverting the covariance matrix in Equation 3.12. This is an $N_{\text{gp}} \times N_{\text{gp}}$ matrix, where N_{gp} is the number of Gaussian process parameters needed for a single pulsar-term BWM signal model [47, 49]. One inversion has computational complexity $O(N_{\text{gp}}^3)$. To compute a full pulsar-term lookup




Figure 3.3: The 95% BWM amplitude upper limits (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 Aggarwal et al. [5]. 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 upper limits at these trial burst epochs.

table, we evaluate the likelihood once for each point on a five-dimensional grid $(A_{\text{rn}}, \gamma_{\text{rn}}, |h_0|, \text{sign}(h_0), t_0)$. Thus, the total cost of the inversions we must perform for one lookup table is $N_{A_{\text{rn}}} N_{\gamma} N_{A_{\text{bwm}}} N_{\text{sign}} N_{t_0} N_{\text{gp}}^3$, where each of these terms represent 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 which has 10 years of data, the total number of grid points is approximately 32×10^6 . 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_{psr} is the number of pulsars in the PTA.

The computational complexity of one pulsar-term search for BWMs using an




Figure 3.4: **Left:** 95% BWM amplitude upper limits (ULs) as a function of sky location. The stars mark the locations of the pulsars in NANOGrav’s 11-year data release. As expected, the PTA is most sensitive to BWMs in sky locations where many pulsars are being timed.

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_{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 upper limits as a function of trial burst epoch (the results shown in Figure 3.3), 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 upper limits as a function of burst epoch is approximately $N_{\text{iter}}N_\theta N_\phi N_\psi N_t N_{\text{psr}}N_{\text{gp}}^3$. 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 upper limits, it is far more efficient when computing certain full-PTA searches. On an Intel i9-9900K CPU with 8 physical cores operating at 3.60GHz, 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 upper limits would have taken approximately three years.

3.6 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 upper limits as a function of burst epochs (see Figure 3.3) match well with previously published results, with almost

no difference in the most sensitive regions of the data set (although the upper limits 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. [15], 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. [15], 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 is a very efficient way to perform search for BWMs as PTAs continue growing, 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 numbers 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.

Chapter 4: The NANOGrav 12.5 yr Data Set: Search for Gravitational Wave Memory

Gabriella Agazie, Zaven Arzoumanian, Paul T. Baker, Bence Bécsy, Laura Blecha, Harsha Blumer, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Rand Burnette, Robin Case, J. Andrew Casey-Clyde, Maria Charisi, Shami Chatterjee, Tyler Cohen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Dallas DeGan, Paul B. Demorest, Timothy Dolch, Brendan Drachler, Justin A. Ellis, Robert D. Ferdman, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Nate Garver-Daniels, Peter A. Gentile, Joseph Glaser, Deborah C. Good, Kayhan Gültekin, Jeffrey S. Hazboun, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey S. Key, Nima Laal, Michael T. Lam, William G. Lamb, T. Joseph W. Lazio, Natalia Lewandowska, Tingting Liu, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Chung-Pei Ma, Dustin R. Madison, Alexander McEwen, James W. McKee, Maura A. McLaughlin, Patrick M. Meyers, Chiara M. F. Mingarelli, Andrea Mitridate, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Nihan S. Pol, Scott M. Ransom, Paul

S. Ray, Joseph D. Romano, Shashwat C. Sardesai, Kai Schmitz, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Sophia V. Sosa Fiscella, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Joseph K. Swiggum, Jacob Taylor, Stephen R. Taylor, Jacob E. Turner, Caner Unal, Michele Vallisneri, Sarah J. Vigeland, Haley M. Wahl, Caitlin A. Witt, Olivia Young, and The NANOGrav Collaboration

The Astrophysics Journal
1667 K Street NW, Suite 800
Washington, DC 20006 USA
Volume 951, Issue 2

4.1 Abstract

We present the results of a Bayesian search for gravitational wave (GW) memory in the NANOGrav 12.5-yr data set. We find no convincing evidence for any gravitational wave memory signals in this data set. We find a Bayes factor of 2.8 in favor of a model which includes a memory signal and common spatially uncorrelated red noise (CURN) compared to a model including only a CURN. However, further investigation shows that a disproportionate amount of support for the memory signal comes from three dubious pulsars. Using a more flexible red noise model in these pulsars reduces the Bayes factor to 1.3. Having found no compelling evidence, we go on to place upper limits on the strain amplitude of GW memory events as a function of sky location and event epoch. These upper limits are computed using a signal model that assumes the existence of a common, spatially uncorrelated red noise in addition to a GW memory signal. The median strain upper limit as a function of sky position is approximately 3.3×10^{-14} . We also find that there are some differences in the upper limits as a function of sky position centered around PSR J0613–0200. This suggests that this pulsar has some excess noise which can be confounded with GW memory. Finally, the upper limits as a function of burst epoch continue to improve at later epochs. This improvement is attributable to the continued growth of the pulsar timing array.

4.2 Introduction

Any system which radiates gravitational waves (GWs) will also cause a permanent change in the spacetime metric. This effect, which was first derived in Zel'dovich and Polnarev [130], is called gravitational wave memory. Later on, it was discovered that the gravitational waves emitted by a system are themselves a source of memory [19, 32, 115, 124]. This effect is known as nonlinear memory, since it arises from nonlinearities in the Einstein field equations. Much work has been done to estimate the size of the effects of nonlinear GW memory, and it has been shown that there is a reasonable chance that the GW memory effect is significant enough to be observed using modern GW detectors [10, 53–55, 124].

One such GW detector is a pulsar timing array (PTA). A PTA is a collection of millisecond pulsars (MSPs) which have extremely stable rotational periods [79]. Because of their stability, it is expected that by carefully observing times-of-arrival (TOAs) of radio pulses from these MSPs, it is possible to observe timing residuals induced by the passage of GWs [44, 64, 103]. The combination of multiple MSPs into a PTA also offers boosted sensitivity when trying to detect GW signals that produce predictable correlations amongst multiple pulsars [57, 78]. Currently, there are several PTA collaborations in operation, including the North American Nanohertz Observatory for Gravitational Waves [NANOGrav, 92], the European Pulsar Timing Array [EPTA, 43], the Parkes Pulsar Timing Array [PPTA, 90], and the Indian Pulsar Timing Array [InPTA, 98]. Together, these collaborations form the International Pulsar Timing Array [IPTA, 121]. In addition, the Chi-

nese Pulsar Timing Array [CPTA, 127] and the MeerTime Pulsar Timing Array [MPTA, 96] have recently released their first analyses.

In the absence of any exotic physics, PTAs are expected to first detect a gravitational wave background originating from an ensemble of supermassive black hole binary (SMBHB) systems, followed by continuous waves from particularly bright SMBHBs [102]. Much work has already been done to characterize the GW background and place limits on continuous GWs [e.g., 1, 8, 12, 13, 15, 17, 30, 52]. During the final inspiral and merger of a SMBHB system, the SMBHB strongly emits GWs which are outside the frequency band detectable by PTAs. However, the accumulated memory from these GWs may be significant enough to be detected by PTAs [35, 82, 100, 104, 118]. In addition to these signals, it is also possible to detect or constrain more exotic GW sources using PTA data. For example, cosmic strings are expected to emit strong bursts of GWs [39, 106]. Yonemaru et al. [129] have placed limits on GW bursts from cosmic strings based on the second PPTA data release.

Several studies have already been done using PTA data to constrain GW memory. NANOGrav has published constraints on GW memory using their 5-year and 11-year data sets [5, 11, hereafter NG5-bwm and NG11-bwm, respectively]. The PPTA has also published constraints in Wang et al. [123]. Madison et al. [83] have published the results of a search for GW memory from five galaxy clusters using PPTA data.

Additionally, several studies have considered GW memory using ground- and space-based detectors like LIGO-Virgo-KAGRA and LISA. Lasky et al. [75] suggest

a method to detect accumulated memory of many individual mergers, each of which is too weak to see by itself. Hübner et al. [66] find no evidence of memory using data from 50 detections of GWs made during the third observation run of LIGO and Virgo. Boersma et al. [20] forecasted that GW memory could be detected with total $\text{SNR} = 3$ after approximately five years of aLIGO operation at design sensitivity. Favata [53] estimated that memory from SMBHB mergers may be detectable ($\text{SNR} \sim 5$) out to $z \lesssim 2$ with LISA, and Islo et al. [67] estimated that LISA may see between 1 and 10 memory events in its lifetime.

In this paper, we present our analysis of the NANOGrav 12.5-year data set [6] for GW memory. We find that there is no significant evidence for GW memory in the data set. The model including a memory signal and a common spatially uncorrelated red noise (CURN) process is only very marginally favored, with a Bayes factor of 2.8, when compared to a model including only a CURN process. The posteriors from a full PTA analysis show that there is a very weak hint for GW memory at three different epochs: MJDs 54000, 55400, and 57300. However, a more detailed analysis shows that these three features are spurious. Each event is only supported by one or two pulsars, and one even lies inside a data gap in which there are no TOAs. When using more flexible red noise models for three of the pulsars in the data set, the Bayes factor drops to $BF = 1.3$.

Thus, finding no GW memory events, we present upper limits constraining the amplitudes of any GW memory as functions of trial burst epoch and sky location. In addition, we use the constraints as a function of burst epoch to set constraints on rates of all astrophysical events which produce GW memory.

In section 4.3, we will describe the NANOGrav 12.5-year data set. Then, in section 4.4, we will discuss how a GW memory wavefront affects TOAs from a PTA. Next, in section 4.5, we will summarize the mathematical techniques and software used in this search. Finally, we will discuss the results in section 4.6.

4.3 Data

In this paper, we analyze the NANOGrav 12.5-year narrowband data set [6]. We will briefly summarize some key points about this data set, but more details may be found in Alam et al. [6] (hereafter NG12).

This data set contains TOAs from observations of 47 pulsars made between July 2004 and June 2017. However, in this analysis, we used only the 45 pulsars with at least three years of data. These observations were performed using the Arecibo Observatory (AO) and Green Bank Telescope (GBT). All pulsars in the declination range $0^\circ < \delta < +39^\circ$ were observed at Arecibo, with the remaining pulsars observed at GBT. In addition, PSRs B1937+21 and J1713+0747, which lie within the aforementioned declination range, were also observed at GBT. Based on work by Burt et al. [26] and Christy et al. [33], six pulsars were also observed in a high-cadence program: PSRs J0030+0451, J1640+2224, J1713+0747, J1909–3744, J2043+1711, and J2317+1439. These six pulsars were observed weekly starting from 2013 at GBT and 2015 at AO. The remaining pulsars were observed monthly.

Each pulsar was, where possible, observed using two different receivers at differ-

ent frequency ranges to help understand and model out interstellar medium (ISM) and dispersion measure (DM) effects. At Arecibo, observations were performed with the 1.4 GHz receiver, and one of either the 430 MHz or 2.1 GHz receivers depending on the noise characteristics of the observed pulsar. At the GBT, monthly observations used the 1.4 GHz and 820 MHz receivers. Weekly observations, however, only used the 1.4 GHz receiver. The observations were initially recorded using the ASP/GASP backends at AO and GBT, respectively [40]. Later, between 2010 and 2012, these backends were replaced by the wideband backends PUPPI/GUPPI at AO and GBT, respectively [45, 56].

Each pulsar’s timing model was fitted using TEMPO¹[97] and checked for consistency using TEMPO2²[65] and PINT³[80].

4.4 Signal and Noise Model

In this section, we will discuss the effects of GW memory on TOA residuals in pulsar timing data and summarize all of the components of the signal model and noise model used in this search.

Qualitatively, a memory wavefront passing over a single pulsar will cause the observed rotational frequency of the pulsar to suddenly increase or decrease by a constant amount. A memory wavefront passing over the Earth will cause the observed rotational frequencies of *all pulsars* in the PTA to suddenly increase

¹<https://tempo.sourceforge.net/>

²<https://bitbucket.org/psrsoft/tempo2>

³<https://github.com/nanograv/PINT>

or decrease by a constant amount. In either case, this sudden change to the observed rotational frequency will introduce timing residuals because of the difference between the pulsar's timing-model fitted rotational frequency and the observed rotational frequency. Because the difference between the expected frequency and observed frequency is a constant, this will cause residuals to accumulate linearly over the course of the observation. In the case of a memory wavefront passing over just one pulsar, we will only see residuals in the TOAs of that pulsar. If a wavefront passes over the earth, we will see the residuals begin accumulating in the TOAs of every single observed pulsar.

For a GW memory wavefront propagating in the direction $\hat{\mathbf{k}}$ with polarization angle ψ passing over the line-of-sight to a pulsar at sky position \mathbf{p} , the residuals induced may be calculated [51, 64]:

$$\delta t_{\text{mem}}(t) = B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) h_{\text{mem}}(t). \quad (4.1)$$

The projection factor $B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi)$ accounts for the fact that the GW memory has a quadrupolar antenna pattern. The effect of the GW memory on the TOAs from a pulsar depends on where that pulsar lies inside the antenna pattern. For a pulsar at \mathbf{p} and a wavefront propagating in the direction $\hat{\mathbf{k}}$ separated by an angle α , we can write the projection factor as:

$$B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{1}{2} \cos(2\psi_{\hat{\mathbf{p}}})(1 - \cos \alpha), \quad (4.2)$$

where α is the angle between $\hat{\mathbf{p}}$, and $\hat{\mathbf{k}}$, and $\psi_{\hat{\mathbf{p}}}$ is the angle between the principal

polarization vector (defined by ψ) and the pulsar line-of-sight projected onto a plane perpendicular to $\hat{\mathbf{k}}$.

The second factor in Equation 4.1 carries the strength and time dependence of the burst. For a wavefront with a characteristic strain of h_0 , we may write the time-dependence as

$$h_{\text{mem}}(t) = h_0[(t - t_0)\Theta(t - t_0) - (t - t_i)\Theta(t - t_i)] \quad (4.3)$$

where t_0 is the time at which the memory wavefront passes over the Earth, $t_i = t_0 + (|\vec{\mathbf{p}}_i|/c) [1 + \cos(\theta_i)]$ is the retarded time at which the same wavefront passed over the pulsar, and Θ is the Heaviside function. The left- and right-hand terms in Equation 4.3 are called the ‘‘Earth term’’ and ‘‘pulsar term’’, respectively. In reality, because the distance to each pulsar in our PTA is on the order of thousands of light-years, and the observation baselines of ongoing PTA experiments is tens of years, we expect only one nonzero term in Equation 4.3.

The characteristic strain of GW memory h_0 depends on the amount of energy radiated as GWs, ΔE_{rad} , and the orientation and distance of the source relative to the observer. For a binary merger

$$h_0 = \frac{\Delta E_{\text{rad}}}{24r} \sin^2 \iota (17 + \cos^2 \iota), \quad (4.4)$$

where r is the comoving distance to the source, ι is the orbital inclination angle of the binary, and ΔE_{rad} is a function of the individual masses and spins. As our

signal model includes only the memory portion, we are agnostic to the particulars of the signal’s origin and parameterize our model with h_0 directly.

Because of the significance of the detection of a CURN process in NG12gwb, we also include a CURN with a fixed spectral index as part of our model. The CURN is modeled as a power law, with a power spectrum characterized by two hyperparameters (A, γ) [99]:

$$P(f) = A^2 \left(\frac{f}{\text{yr}^{-1}} \right)^{-\gamma}, \quad (4.5)$$

where f is the frequency of the spectral component, A is the characteristic amplitude of the red-noise process at the reference frequency yr^{-1} , and γ is the spectral index of the process. A stochastic gravitational wave background generated by an ensemble of SMBHB is expected to have a spectral index of $\gamma_{\text{SMBHB}} = 13/3 \approx 4.33$. However, the maximum a posteriori value found for the spectral index of the CURN in NG12gwb was approximately $\gamma_{\text{MAP}} = 5.5$. In this paper, we present two sets of results using both of these fixed spectral indices for the CURN.

In addition to these signals, we include Gaussian white noise and Gaussian red noise on a per-pulsar basis. The Gaussian red noise accounts for long-timescale changes in the pulsar’s rotational frequency. Some examples of processes which can cause these changes include spin noise [74, 105], stochastic variations in dispersion measure [41, 69, 70], and mode changing [81, 95].

The Gaussian white noise is parameterized by three parameters known as EQUAD, EFAC, and ECORR. EQUAD and EFAC modify the measured TOA

uncertainty: EQUAD adds additional white noise in quadrature, and EFAC multiplies the total TOA uncertainty after including EQUAD. ECORR describes white noise that is correlated between TOAs gathered in the same observation epoch but uncorrelated between different observations. This term nominally accounts for pulse jitter noise. For this analysis, the white noise parameters are fixed to their median values as determined by single pulsar noise analyses for the sake of computational efficiency.

4.5 Methods

The techniques used in this search are documented in Sun et al. [110]. As such, in this section, we will give only a brief overview of the techniques. The residuals in a single pulsar’s TOAs may be written as the sum of multiple stochastic and deterministic processes:

$$\delta\mathbf{t} = \delta\mathbf{t}_{\text{mem}} + M\boldsymbol{\epsilon} + F\mathbf{a} + F_{\text{gw}}\mathbf{a}_{\text{gw}} + \mathbf{n}. \quad (4.6)$$

Above, $\delta\mathbf{t}$ are the residual timeseries for the pulsar. The term $\delta\mathbf{t}_{\text{mem}}$ are the residuals induced by GW memory; M is the design matrix accounting for small errors in the linearized pulsar timing model $\boldsymbol{\epsilon}$; F is the design matrix for a pulsar-intrinsic Gaussian red-noise process modeled as a Fourier series with coefficients \mathbf{a} ; Similarly, F_{gw} and \mathbf{a}_{gw} , are the design matrix and Fourier coefficients for the CURN; finally, \mathbf{n} are the uncertainties in the TOAs from Gaussian white noise.

Given estimations of the timing model parameters, GW memory signal, and

Gaussian process parameters, we can construct residuals \mathbf{r} :

$$\mathbf{r} = \delta\mathbf{t} - \delta\mathbf{t}_{\text{mem}} - M\boldsymbol{\epsilon} - F\mathbf{a} - F_{\text{gw}}\mathbf{a}_{\text{gw}}. \quad (4.7)$$

Since the residuals \mathbf{r} are expected to arise only from Gaussian white noise (having subtracted out all other effects), we can compute the likelihood of any set of model parameters as:

$$\begin{aligned} L(\boldsymbol{\epsilon}, \mathbf{a}, \mathbf{a}_{\text{gw}}, \delta\mathbf{t}_{\text{mem}}) &= p(\delta\mathbf{t}|\boldsymbol{\epsilon}, \mathbf{a}, \mathbf{a}_{\text{gw}}, h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) \\ &= \frac{\exp\left(-\frac{1}{2}\mathbf{r}^T N^{-1}\mathbf{r}\right)}{\sqrt{2\pi \det N}}. \end{aligned} \quad (4.8)$$

It is also possible to analytically marginalize this likelihood over the timing model and red noise parameters. For a full description, see Lentati et al. [76], van Haasteren and Vallisneri [119], and van Haasteren and Vallisneri [120]. The final marginalized likelihood is:

$$p(\delta\mathbf{t}|h_{\text{mem}}, t_0, \hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi) = \frac{\exp\left(-\frac{1}{2}\mathbf{q}^T C^{-1}\mathbf{q}\right)}{\sqrt{2\pi \det C}} \quad (4.9)$$

where we have the definitions:

$$\mathbf{q} = \delta\mathbf{t} - \delta\mathbf{t}_{\text{mem}}, \quad (4.10)$$

$$C = N + TDT^T, \quad (4.11)$$

$$T = \begin{bmatrix} M & F \end{bmatrix}, \quad (4.12)$$

$$D = \begin{bmatrix} \infty & 0 \\ 0 & \phi \end{bmatrix}, \quad (4.13)$$

where ∞ above is a diagonal matrix of infinities (which we can understand as unconstrained priors on timing model parameters), and ϕ is a covariance matrix for the individual red noise and CURN Fourier coefficients. Because we are using a CURN, these ϕ matrices are also diagonal, and simply contain the red noise power at each frequency bin given by Equation 4.5.

Following the procedure of Sun et al. [110], which is based on methods in §6 of NG5-bwm, we then compute the pulsar-term likelihoods (marginalized over intrinsic pulsar red noise and fixed spectral index CURN) on a grid of trial parameters $\{h_i, t_B\}$ where h_i is the post-projection strain of the memory signal in the i -th pulsar, and t_B is the burst epoch. The post-projection strain is given by the product of the projection factor (Equation 4.2) and the intrinsic strain of the memory signal h_0 . We can only use the post-projection strain since the residuals of one pulsar cannot break the degeneracy between the location of the signal's origin and the intrinsic strain. Additionally, Sun et al. [110] do not include a CURN, but we choose to include this additional noise process because of the results of NG12gwb, in which it was shown that there is significant evidence for a CURN in the NANOGrav 12.5-year data set [NG12]. These pulsar-term likelihood tables

may be used to set upper limits on pulsar-term GW memory.

Then, we can combine the pulsar-term likelihoods to compute Earth-term likelihoods by making use of the factorizability of the signal model:

$$\begin{aligned}
 p(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_0) &= \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|\hat{\mathbf{k}}, \psi, t_B, h_0) \\
 &= \prod_{i=1}^{N_{\text{psr}}} p_i(\delta\mathbf{t}|h_i, t_B),
 \end{aligned}
 \tag{4.14}$$

where above we have implicitly used Equation 3.3 to combine the burst parameters $\hat{\mathbf{k}}$, ψ , and h_0 into the post-projection, pulsar-term GW memory strain h_i . In this way, it becomes very computationally inexpensive to compute the red-noise-marginalized Earth-term likelihoods of any GW memory events on a full grid of trial parameters $\{\log_{10} h_0, t_B, \theta_B, \phi_B, \psi_B\}$ where h_0 is the *intrinsic* strain of the memory event, t_B is the event epoch, (θ_B, ϕ_B) are the polar and azimuthal angles of the sky location of the event source, and ψ_B is the polarization angle of the memory wavefront. Once we have the likelihoods on a grid of trial parameters, we can simply numerically marginalize over any of the trial parameters to obtain marginalized likelihoods or posterior probability distributions.

These signal models and this likelihood calculation are implemented in `enterprise`⁴ [49] and `enterprise_extensions`⁵ [113]

⁴<https://github.com/nanograv/enterprise>

⁵https://github.com/nanograv/enterprise_extensions

4.6 Results

4.6.1 Earth-term Memory Search

We began by performing an Earth-term Bayesian search for GW memory using MCMC sampling. We compared two models: 1) a noise-only model and 2) a noise and GW memory model. The noise-only model included intrinsic pulsar red noise, white noise, and a common red noise process. The signal model included the same noise processes with an additional GW memory signal. The two models were simultaneously sampled using the product-space sampling method [28, 59], allowing us to determine the posterior probability for the memory signal and compute the Bayes factor for the signal model compared to noise only. The resulting Bayes factor of 2.8, shows the GW memory model is marginally favored over noise only. However, this Bayes factor is too small to be considered a detection. The posterior probability distributions for the memory signal and global spatially uncorrelated red-noise process are shown in Figure 4.1. Based on the posterior probability of the burst epoch, we can identify three “hot spots” near MJDs 54000, 55400 and 57300.

The features near MJD 54000 and 55400 were both present in the analysis of NG11-bwm, where they were the most significant GW memory false alarm events in NG11 and NG9, respectively. The feature near MJD 54000 lies near the start of our observations and at a time where there were large data gaps for several pulsars. At early times in our dataset there were fewer pulsars being observed, and the observations were less regular. The sparsity of data makes it harder to

constrain any signal in these times. Events that occur early in the data are also more degenerate with the quadratic pulsar timing model fit to the pulsar rotational frequency and frequency derivative. This means that the signal model can be consistent with a high amplitude memory event that is effectively removed by the marginalization of the timing model.

Using a dropout analysis, we identified three pulsars in particular which supported each of the three aforementioned features: PSRs J0030+0451, J1744–1134 and J2043+1711. We performed another Bayesian search using a free spectral noise model, which treats the power in each frequency bin in the power spectral density as an independent parameter, rather than requiring a power-law red noise power spectral density for these three pulsars. We found that using a more flexible noise model in these pulsars completely removes the features at MJD 55400 and MJD 57300. Because each of these features has no support from any other pulsars, we conclude that these events are related to noise in individual pulsars and not actual GW memory events. The analysis using the free spectral noise model in three pulsars results in a Bayes factor for GW memory of 1.3. In general, more complex noise models like those used in Simon et al. [108, in prep] should help prevent noise features from contaminating searches for GW memory in the future.

4.6.2 Pulsar-term Upper Limits

Because we make no detection, we report upper limits on GW memory strain amplitude.

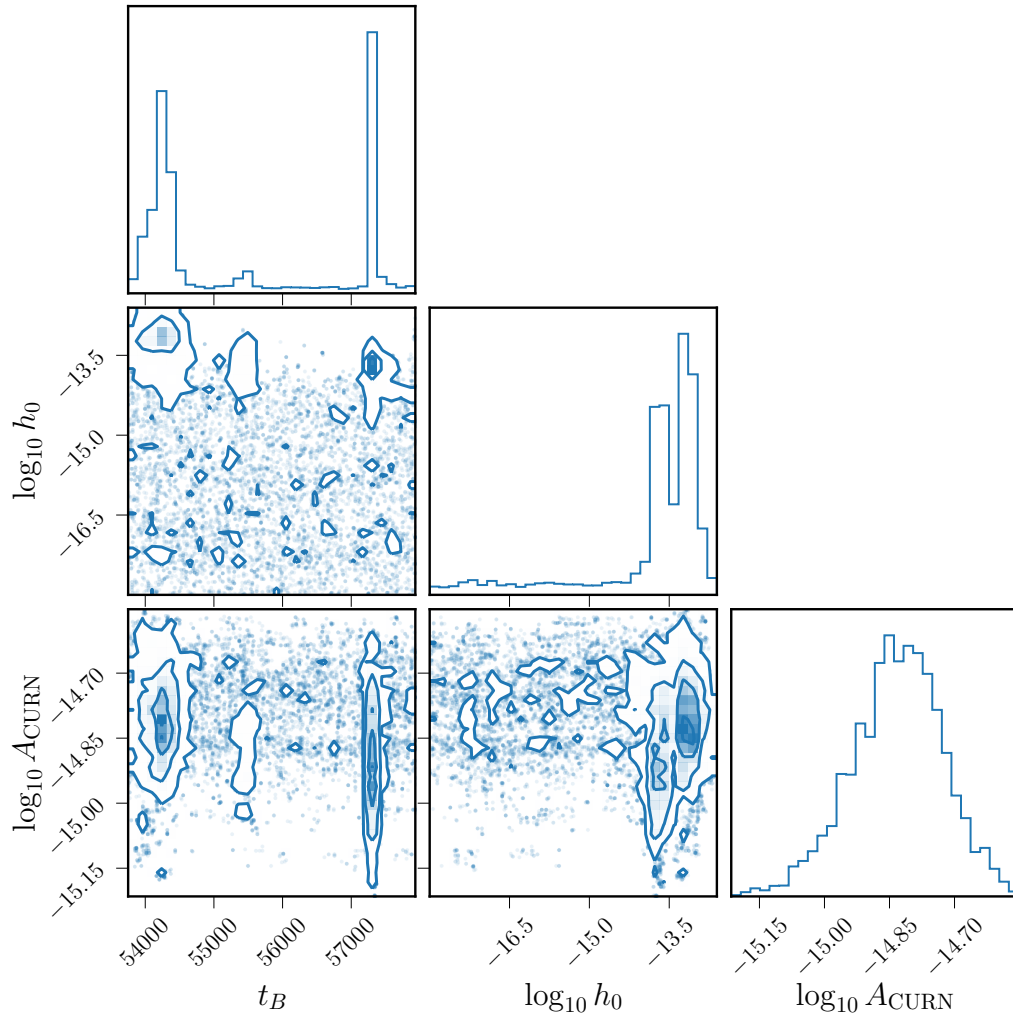


Figure 4.1: A corner plot showing 1D and 2D marginalized posteriors for three key model parameters: burst epoch t_B , burst strain amplitude $\log_{10} h_0$, and CURN amplitude $\log_{10} A_{\text{CURN}}$. The good localization $\log_{10} A_{\text{CURN}}$ shows that the CURN is still detected in the presence of the memory model. Furthermore, the tail of $\log_{10} h_0$ extends to very low amplitude, which indicates that $h_0 \sim 0$ is still supported by the model.

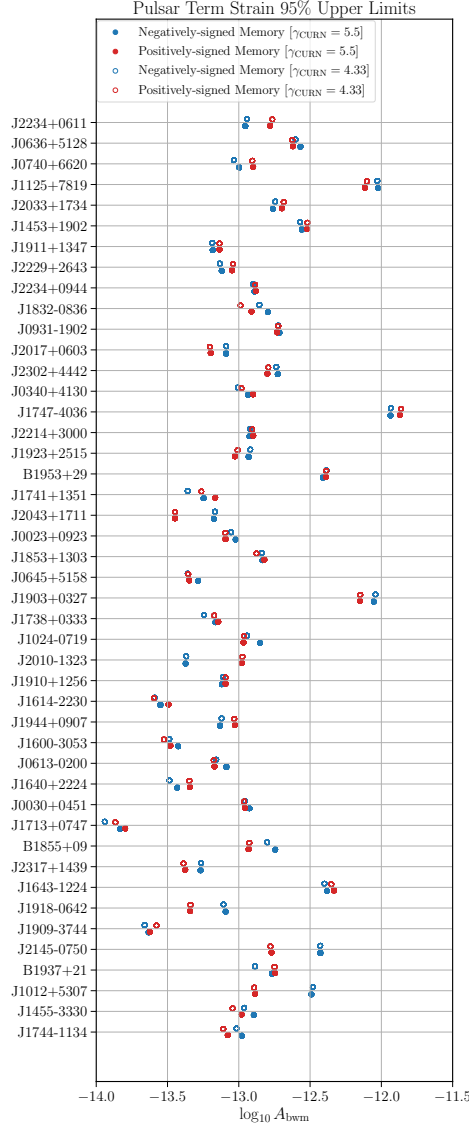


Figure 4.2: A plot of the pulsar-term upper limits on memory strain amplitude. The pulsars are listed in order of shortest to longest timing baseline. To find these upper limits, we compute amplitude posteriors from the pulsar-term lookup tables marginalized over the burst epoch, pulsar intrinsic red noise and a fixed spectral index common uncorrelated red noise (CURN) process. Overall, we do not find much difference in pulsar-term upper limits when comparing the results using a fixed CURN spectral index of $\gamma_{\text{SMBHB}} = 4.33$ and $\gamma_{\text{MAP}} = 5.5$.

Figure 4.2 shows the pulsar-term upper limits on GW memory using both γ_{MAP} and γ_{SMBHB} . Because the pulsar-term upper limits are computed one pulsar at a time, we lose all information relating to the sky location of the signal. This amplitude upper limit is a constraint on the product of $B(\hat{\mathbf{k}}, \hat{\mathbf{p}}, \psi)$ and h_0 , since these two terms are fully degenerate in the pulsar-term search. In other words, it is impossible to differentiate between a weak memory event or one that originated in the sky such that the antenna pattern is weak at the pulsar’s sky location. We see that the choice of spectral index does not affect the pulsar-term upper limits very much in most cases. Some pulsars (e.g., PSRs B1937+21, J0613–0200, J0645+5158, J1713+0747) show small, but significant differences.

4.6.3 Earth-term Upper Limits

Figure 4.3 and Figure 4.4 show the upper limits on GW memory strain amplitude in the NANOGrav 12.5-year data set as a function of burst epoch and sky location, respectively.

To compute the upper limits as a function of burst epoch (Figure 4.4), we must compute amplitude posteriors which have uniform priors over the sky and polarization. Thus, we started by splitting up the sky into 48 HEALPix⁶[61] sky pixels using healpy⁷ (`nside=2`) [132]. Then, for each sky pixel, we computed likelihood tables for global GW memory events using Equation 4.14 and the pulsar-term likelihood tables. Finally, for each trial burst epoch, we took an equal number

⁶<https://healpix.sourceforge.io/>

⁷<https://github.com/healpy/healpy>

of samples from the amplitude posteriors from each source-orientation bin at that trial epoch. We then concatenated the samples taken from each of these amplitude posteriors together to form a sky-averaged strain amplitude posterior.

We must sample each source-orientation bin independently to construct our sky-averaged posteriors because of the nature of the memory signal. Our PTA does not have uniformly-distributed pulsars, and as such, there are parts of the sky in which we have little to no sensitivity. If we are not careful about sampling, and instead search over the entire sky simultaneously, our amplitude would be dominated by samples taken from source orientations to which our PTA is completely insensitive. Furthermore, because there is much more prior volume at high amplitudes, these samples would all heavily bias our amplitude posteriors towards higher amplitudes which our PTA has no way of ruling out. This sampling scheme, in which we concatenate samples from different source orientation bins, guarantees that our posterior is marginalized uniformly over the prior [88].

Figure 4.3 shows the upper limits on memory strain as a function of sky location and fixed common red-noise spectral index. To obtain these upper limits, we first computed Earth-term lookup tables for 768 HEALPix sky pixels (`nside = 8`) marginalized over the polarization of the memory wavefront (c.f., the upper limits as a function of burst epoch). For these Earth-term lookup tables, we limit the prior on the burst epoch to the last three years because some pulsars do not have more than three years of data. We can see from this comparison that the upper limits differ slightly depending on the choice of fixed spectral index for the common red noise. While we do not believe these differences are significant, there

is a difference pattern that is very similar to the antenna response of a GW memory event around PSR J0613–0200. After repeating the same analysis, but omitting this pulsar, the nearby differences largely disappear. This suggests that this pulsar contains a noise feature which is difficult to model accurately using only a red noise power law and white noise. When mismodeled, the excess noise is conflated with a GW memory signal, thus causing the upper limit differences in Figure 4.3.

Figure 4.4 shows the upper limits on GW memory in the NANOGrav 12.5-year data set plotted as solid curves. We also show the results of the NANOGrav 11-year search for GW memory plotted as a dashed green curve. The upper limits computed from the NANOGrav 12.5-year data set do not improve significantly upon those computed from NG11-bwm in the overlapping epochs. However, the increased volume of timing data and number of pulsars do clearly result in improvements on the upper limits of approximately half an order of magnitude when compared to early upper limits. Continued observation and growth of this PTA will cause the upper limits in the future to be even lower, and give much more stringent limits on GW memory.

Figure 4.5 shows the upper limits on the rate of SMBHB mergers that produce GW memory computed using the results shown in Figure 4.4. We do this by counting the number of epochs which have lower strain upper limits than a given fixed strain. From this, we can then constrain the rate of events that have strains at or below this fixed strain. In addition, this figure shows the predicted rate in Islo et al. [67]. From the right-hand-side plot, we can see that our rate upper limits do not improve much when compared to the NANOGrav 11-year results. We also

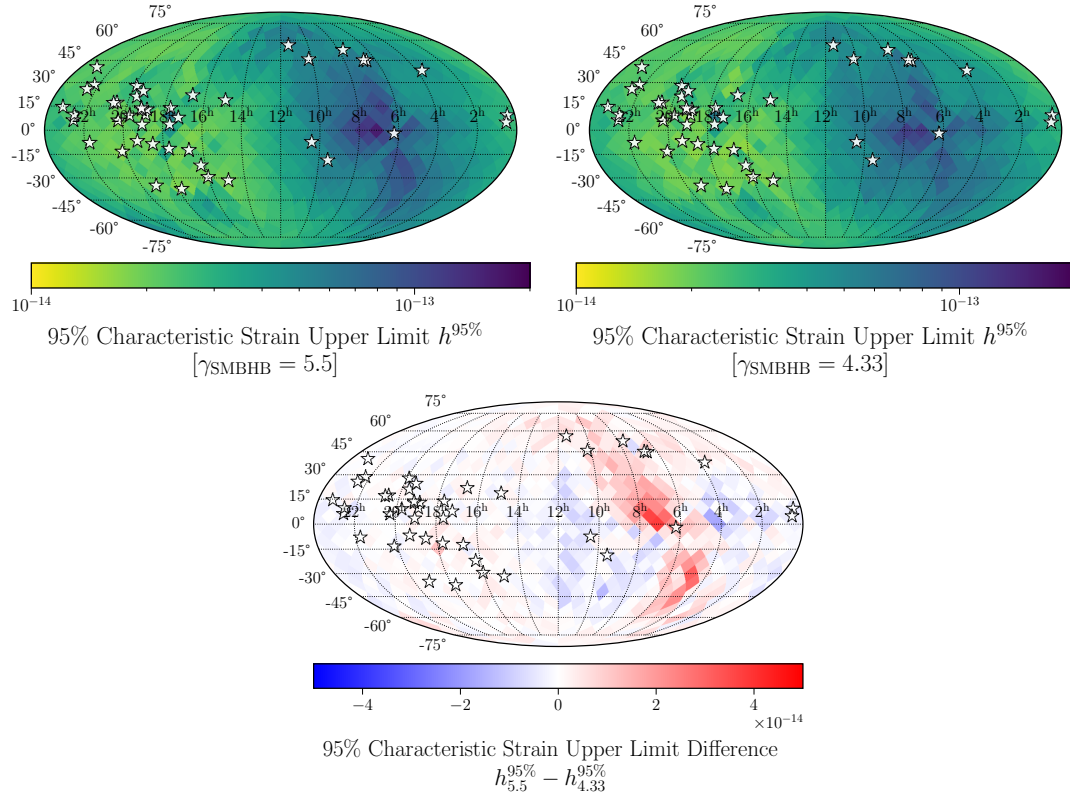


Figure 4.3: Top left: The upper limits on memory strain amplitude as a function of skypixel including a CURN law process using a fixed spectral index of $\gamma_{\text{SMBHB}} = 4.33$, as expected for a stochastic gravitational wave background originating from an ensemble of uniformly, isotropically distributed SMBHBs. Top right: The upper limit on memory strain amplitude as a function of sky pixel including a CURN power law process using a fixed spectral index of $\gamma_{\text{MAP}} = 5.5$, the maximum a posteriori value for the detected CURN in Arzoumanian et al. [15]. Bottom: The difference between the top two panels. A positive value indicates that the upper limits on strain using a spectral index of 5.5 are higher. We see that overall, the upper limits change slightly when the red noise model uses the preferred red noise spectral index. However, these changes are localized to a small part of the sky. It can be shown that these upper limit differences can be largely attributed to PSR J0613–0200.




Figure 4.4: The memory strain amplitude upper limit as a function of burst epoch. The black curve shows the upper limits using a model which does not include CURN. The blue and orange curves show the upper limits using a model which does include a fixed spectral index CURN power law process with $\gamma_{\text{MAP}} = 5.5$ and $\gamma_{\text{SMBHB}} = 4.33$, respectively. We see that the sensitivity of the NANOGrav 12.5-year data set does not give significantly improved upper limits for the first 11 years. However, the additional pulsars and timing data give continuously improving upper limits later in the data set.




Figure 4.5: A plot of the predicted rate of SMBHB mergers (red) against the rate upper limits derived from the strain amplitude upper limit as a function of burst epoch. Because the strain upper limits do not improve significantly, the event rate upper limits are correspondingly similar to those presented in Aggarwal et al. [5].

include the sky-marginalized pulsar term upper limits. Notably, in this analysis, the Earth term rate upper limits are more constraining than the combined pulsar term upper limits. This indicates that the PTA contains enough pulsars that the sensitivity at low strain amplitudes is no longer dominated by a few pulsars. In addition, the NANOGrav 12.5-year pulsar-term rate upper limits are worse than the 11-year rate upper limits. This is due to the additional red noise model used in the analysis of the 12.5-year data set.

4.7 Discussion and Conclusion

In this paper, we have shown that there is no significant detection of GW memory in the NANOGrav 12.5-year data set. We have therefore set upper limits on the strain amplitude of any GW memory events in the NANOGrav 12.5-year data set in the presence of the CURN detected in NG12gwb. The addition of a CURN to the noise model does not significantly affect the upper limits, but does have some covariance with the GW memory signal. We also see from Figure 4.3 that PSR J0613–0200 gives significantly different strain upper limits for sources in its vicinity depending on the choice of spectral index for a CURN process. Furthermore, these differences have a quadrupolar shape, similar to the antenna response of a GW memory signal. This indicates the presence of some excess low-frequency noise in this pulsar which can be conflated with GW memory. From Figure 4.4 and Figure 4.5, we see that the additional data in the 12.5-year data set continue to increase our sensitivity to GW memory, especially later in the data set. This in turn will allow us to continue

placing more stringent limits on the rates of memory-producing events.

It is important to remember that the predicted event rate upper limits shown in Figure 4.5 are only for SMBHB mergers. While the prospects for detecting GW memory from SMBHB mergers are low, there are many exotic sources which may be expected to emit GWs and produce GW memory [38]. Furthermore, it has been recently discovered that there is some evidence of a GW background in the NANOGrav 15-year data set [1]. The spectrum detected in Agazie et al. [1] may be used to update estimations of SMBHB systems and recompute merger event rates. This work is critical for continued development of future searches for GW memory. Additionally, pulsar glitches, which are instantaneous changes in the rotational frequencies, produce an almost identical signal as a pulsar-term GW burst with memory [35]. Pulsar glitches have been observed in two millisecond pulsars, PSRs B1821–24 and J0613–0200 thus far [34, 91]. Of these two, J0613–0200 is in NANOGrav’s timing data set. However, this glitch occurred before NANOGrav began timing the pulsar, and this glitch should therefore have no effect on this pulsar’s timing model or residuals. The pulsar-term upper limits presented in this analysis may be used to set upper limits on glitches in every other pulsar as well. In general, this analysis may be used to cross-validate any detection of any loud GW-producing event.

Finally, the search for GW memory can reveal interesting noise features of a PTA’s constituent pulsars. For example, the analysis presented in NG11-bwm shows that PSRs J1909–3744 and J0030+0451 had some excess, unmodeled noise. This analysis shows that there is some excess noise in PSRs J1744–1134 and

J2043+1711 that conspire to give support for a memory event at MJD 57300. This makes them good candidates for any future studies of noise characteristics, like those presented in Lam et al. [73] and Hazboun et al. [63]. In addition, Figure 4.3 suggests that PSR J0613–0200 may have a noise transient which is highly covariant with red noise. Previous work has shown that scattering variations may result in excess correlated noise in pulsar timing data sets [29, 60, 70]. Main et al. [86, 87] have shown that in particular, data from PSR J0613–0200 shows significant evidence of scattering variations. These scattering variations may be the source of the differences in GW memory upper limits in the vicinity of this pulsar when using different CURN spectral indices. As pulsar timing baselines become longer and PTA sensitivity to red noise increases, it will be critically important to explore how strong red noise and components of each pulsar’s timing models affect detection prospects of GW memory.

Overall, the search for GW memory remains a critical part of the GW analysis pipeline because of its use in cross-validation of any potential detections of loud GWs, its ability to reveal unmodeled noise, and the possibility that a GW memory event may reveal exotic GW sources. Continued methods development, as applied to both GW memory and intrinsic pulsar noise, will allow us to perform more robust searches for these sources using future data sets.

Chapter 5: Optimal strategies for gravitational wave memory
detection in pulsar timing arrays

Jerry P. Sun, Dustin R. Madison, Xavier Siemens

In-press in Physical Review D

1 Physics Ellipse

College Park, MD 20740

5.1 Abstract

In this work we derive two computationally efficient frequentist detection statistics that can be used in searches for gravitational-wave bursts with memory in pulsar timing data. By maximizing the likelihood ratio in two different ways we construct a coherent statistic and an incoherent statistic, which are analogs of the \mathcal{F}_e and \mathcal{F}_p statistics commonly used for continuous-wave searches in pulsar timing data. We show that both statistics are χ^2 -distributed with varying degrees of freedom and non-centrality parameters given by the signal-to-noise (SNR) ratio of the signal present in our data. The statistics can also be used to compute the maximum likelihood estimators of amplitude parameters of a possible gravitational-wave memory signal in pulsar timing data. We find that in the low-signal regime (SNR $\lesssim 5$), the estimators are inaccurate. However, in intermediate- to high-signal regimes, we show that these estimators can accurately determine the correct signal parameters. We also study the effects of red noise on these statistics, and find that the inclusion of red noise reduces the signal-to-noise ratio for injected memory signals.

5.2 Introduction

Pulsar timing arrays (PTAs) have been very successful at probing the low-frequency part of the gravitational wave (GW) spectrum. There are many ongoing efforts to search for a diverse set of GW signals including continuous GWs from super massive black hole binary (SMBHB) systems [2, 9, 52, 131], a stochastic background of GWs from ensembles of SMBHBs [1, 8, 50, 101, 128], nonlinear GW memory [3, 123],

non-Einsteinian GW polarization modes [16, 31, 37, 126], and short-duration GW bursts from various sources [27, 42]. Pulsar timing experiments use the fact that millisecond pulsars (MSPs) are extremely stable rotators [79]. The regularity of the radio pulses emitted by these pulsars make it possible to use changes in the pulse times of arrival (TOAs) to detect or set limits on various GW phenomena and perform tests of Einsteinian relativity [44, 64, 103]

One such test is the search for GW bursts with memory (BWMs). As a SMBHB system completes the final phase of a merger, the burst of gravitational waves radiated through the event act as a source of so-called “nonlinear” GW memory [19, 32, 115, 124, 130]. This phenomenon is called “nonlinear” because it originates from the nonlinearity of Einstein’s field equations. When this memory wave front crosses the line of sight between the Earth and an MSP, it has the effect of changing the observed rotational frequency of this MSP. Currently, there are efforts by PTAs and ground-based GW observatories to detect nonlinear memory [2, 75]. The work in Refs. [35, 67, 82, 118] also shows the detection prospects for nonlinear memory in space-based, and PTA GW experiments.

Several methods have been developed for searches for nonlinear GW memory across all GW-frequency regimes. PTA experiments largely use Markov Chain Monte Carlo sampling methods to compute posterior probabilities of memory model parameters. Because these methods are computationally intensive, we previously [110] proposed an efficient method using pre-computed likelihood tables to numerically compute posteriors on nonlinear memory model parameters.

In this work, we present two frequentist statistics which can be used in searches

for nonlinear GW memory by PTAs. These statistics are optimal in the sense that they are derived by maximizing the ratio of the likelihoods of a nonlinear GW signal event to the null signal model, following the optimal frequentist strategies laid out for the stochastic GW background [7], and for continuous GWs [46]. These statistics can also be used to compute estimators for the parameters that determine the amplitude of the signal. This work builds on [46], in which analogous statistics are derived for searches for continuous GWs. In turn, the results in [46] built on previous work [18, 36, 68] on frequentist statistics analytically maximized over waveform parameters for continuous GWs in LIGO, LISA, and PTA data. Finally, [11, 123] derived the maximum likelihood estimator for the strain amplitude of a GW memory event in PTA data. In this work we re-parameterize the signal differently allowing us to compute maximum likelihood estimators for the polarization, the amplitude, and, for the incoherent statistic, the sky location of the memory event.

In previous work for continuous GWs [36, 46, 68], the statistic derived by maximizing the likelihood ratio was referred to as the \mathcal{F} -statistic. Ref. [46] further split this into a coherent statistic which they call the Earth-term \mathcal{F} -statistic (\mathcal{F}_e), and an incoherent statistic called the pulsar-term \mathcal{F} -statistic (\mathcal{F}_p) depending on how the likelihood ratio was maximized. In this work, for GW memory events, we will refer to the two different statistics that result from the two maximizations as the coherent \mathcal{F} -statistic (\mathcal{F}_C), and the incoherent \mathcal{F} -statistic (\mathcal{F}_I). Despite progress in [110] to increase the efficiency of full Bayesian searches for nonlinear GW memory, such searches remain very computationally expensive. These new

\mathcal{F} -statistics are useful additions to the standard set of tools used in pulsar timing data analysis. They are very fast to compute and can be used to independently cross-validate a full Bayesian search without significant additional computational costs. We expect a noise-marginalized approach analogous to that used for GW stochastic background searches Vigeland et al. [122] to be necessary for robust GW memory searches using our frequentist techniques. We will develop these techniques in future work.

5.3 Background

In this section, we begin by reviewing the signal model for nonlinear GW memory in pulsar-timing data. We then apply a similar framework to that presented in Ellis et al. [46] for CWs to nonlinear GW memory and develop both coherent and incoherent frequentist statistics by maximizing the likelihood ratio in two different ways.

5.3.1 Signal Model

As a GW memory front passes through the Earth or a pulsar, the apparent rotational frequency of the pulsar changes by a fractional amount proportional to the strain-amplitude of the memory. This sudden mismatch between the pulsar's modeled rotational frequency and apparent rotational frequency induces a linear drift in the residuals of the pulsar's TOAs. This occurs because the change in

rotational frequency is a constant, and we accrue a constant timing residual with every subsequent rotation of the pulsar. In this paper we consider the effect of a GW memory front passing over only the Earth, the so-called Earth-term. The rationale for this assumption is explained in detail below. The residuals induced by the Earth-term GW memory signal in the a -th pulsar of a PTA may be written as

$$r_a(t; t_0, \hat{\Omega}, h_0, \psi) = h_0 \Theta(t - t_0) (t - t_0) (F_+^a(\hat{\Omega}) \cos(2\psi) + F_\times^a(\hat{\Omega}) \sin(2\psi)), \quad (5.1)$$

where h_0 is the intrinsic strain of the memory signal, t_0 is the time at which the memory wavefront passes over the Earth, $\hat{\Omega}$ is the location of the source of the GW memory, and ψ is the polarization angle, the angle between the principal polarization vector and pulsar line of sight projected onto the plane perpendicular to the propagation direction of the wave. The antenna pattern functions $F_+(\hat{\Omega})$ and $F_\times(\hat{\Omega})$ for plus- and cross-polarized GWs are given by

$$F_{a,+}(\hat{\Omega}) = \frac{1}{2} \frac{(\hat{\mathbf{m}} \cdot \hat{\mathbf{p}}_a)^2 - (\hat{\mathbf{n}} \cdot \hat{\mathbf{p}}_a)^2}{1 + \hat{\Omega} \cdot \hat{\mathbf{p}}_a}, \quad (5.2)$$

$$F_{a,\times}(\hat{\Omega}) = \frac{1}{2} \frac{(\hat{\mathbf{m}} \cdot \hat{\mathbf{p}}_a)(\hat{\mathbf{n}} \cdot \hat{\mathbf{p}}_a)}{1 + \hat{\Omega} \cdot \hat{\mathbf{p}}_a}, \quad (5.3)$$

where $\hat{\mathbf{p}}_a$ is the unit vector that points to the a -th pulsar. $\hat{\mathbf{m}}$ and $\hat{\mathbf{n}}$ are two orthogonal vectors that define the plane perpendicular to the propagation direction

of the memory wavefront and are given by

$$\hat{\mathbf{m}} = \sin \phi \hat{\mathbf{x}} - \cos \phi \hat{\mathbf{y}} \quad (5.4)$$

$$\hat{\mathbf{n}} = -(\cos \theta \cos \phi) \hat{\mathbf{x}} - (\cos \theta \sin \phi) \hat{\mathbf{y}} + (\sin \theta) \hat{\mathbf{z}}, \quad (5.5)$$

where θ and ϕ are the polar and azimuthal angles of the source.

It is worth pointing out that the choice of only including the Earth-term GW memory signal is well motivated. The distances from the Earth to pulsars in a PTA, and their distances from one another, are on the order of hundreds to thousands of light-years. Therefore, the time it takes a GW memory front to pass over both one pulsar and the Earth, or over two different pulsars in a PTA, is hundreds to thousands of years. This timescale is much longer than typical pulsar timing experiment durations which are on the order of a decade. This means we do not expect to observe the same GW memory event passing the Earth and a pulsar, or two different pulsars in a PTA, during the course of typical pulsar timing experiments.

It is difficult to make a compelling case for detection using single-pulsar GW memory measurements [35, 118]. This is because we observe sudden changes, called “glitches” [62], in the rotational frequencies of some pulsars which produce a signal in the timing residuals identical to a GW memory burst. These glitches have also been observed in MSPs, albeit very rarely [91].

Therefore only when a GW memory front passes over the Earth, affecting the residuals of all pulsars in the PTA in a correlated way, are we in a position to make

a compelling case for detection. Individual pulsar GW memory searches, however, can still be used to place upper limits.

5.3.2 Likelihood

The total signal and noise model for a pulsar's residuals $\delta\mathbf{t}$ may be written as

$$\delta\mathbf{t} = \delta\mathbf{t}_{\text{bwm}} + \mathbf{n}, \quad (5.6)$$

where $\delta\mathbf{t}_{\text{bwm}}$ are the residuals induced by the memory signal, and \mathbf{n} is a timeseries containing red and white Gaussian noise (see Ref. [119] for descriptions of Gaussian process noise). It is important to include red noise in \mathbf{n} for our analyses since recent PTA data sets have been found to contain a strong GW background red noise process in addition to pulsar-intrinsic red and white noises [1, 15]. Both the intrinsic and common red noise processes are parameterized using the amplitude A_{RN} and spectral index γ via [99]

$$P_{\text{RN}}(f) = A_{\text{RN}}^2 \left(\frac{f}{f_{1\text{yr}}} \right)^{-\gamma}, \quad (5.7)$$

where A_{RN} is the dimensionless noise amplitude at a reference frequency of $f_{1\text{yr}} = 1\text{yr}^{-1}$ and γ is the spectral index. For a sense of the expected magnitude of these parameters, current results from Agazie et al. [1] give estimates for the amplitude and spectral index of $A_{\text{RN}} = 6.4_{-2.7}^{+4.2} \times 10^{-15}$ and $\gamma = 3.2_{-0.6}^{+0.6}$ for the gravitational wave background. Ref. [12] contains a detailed explanation of the parameterization

of red and white noises in PTA data analysis.

The residuals are obtained from the TOAs by subtracting out a timing model. To account for the effects of this subtraction we use the R-matrix formalism described in refs. [41, 48], including only the quadratic terms which originate from fitting for the rotational phase, frequency, and frequency-derivative. The R matrix that fits out the least-squares-minimized quadratic is

$$R = I - M(M^T \Sigma^{-1} M)^{-1} M^T \Sigma^{-1}, \quad (5.8)$$

where I is the identity matrix and M is the following design matrix

$$M = \begin{bmatrix} 1 & t_1 & t_1^2 \\ 1 & t_2 & t_2^2 \\ \vdots & \vdots & \vdots \\ 1 & t_{N_{\text{toas}}} & t_{N_{\text{toas}}}^2 \end{bmatrix}. \quad (5.9)$$

The columns of M use all TOAs for the pulsar. Finally, Σ is the pre-fit noise variance matrix

$$\Sigma = \langle \mathbf{n} \mathbf{n}^T \rangle. \quad (5.10)$$

The post-fit residuals $\tilde{\delta \mathbf{t}}$ and post-fit noise covariance matrix $\tilde{\Sigma}$ then become

$$\tilde{\delta \mathbf{t}} = R \delta \mathbf{t} = R(\delta \mathbf{t}_{\text{bwm}} + \mathbf{n}) \equiv \tilde{\delta \mathbf{t}}_{\text{bwm}} + \tilde{\mathbf{n}}, \quad (5.11)$$

with

$$\tilde{\Sigma} = R\Sigma R^T. \quad (5.12)$$

We can then write the likelihood for the signal as the probability that the residuals with the signal subtracted out follow the expected Gaussian noise distribution

$$p(\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) = \frac{\exp\left[-\frac{1}{2}(\tilde{\boldsymbol{\delta}}\mathbf{t} - \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}})^T\tilde{\Sigma}^{-1}(\tilde{\boldsymbol{\delta}}\mathbf{t} - \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}})\right]}{\sqrt{\det(2\pi\tilde{\Sigma})}}. \quad (5.13)$$

5.3.3 General \mathcal{F} -statistic

The derivation presented in Ellis et al. [46] can also be applied to the case of nonlinear gravitational-wave memory. For completeness, we reproduce the steps involved in this derivation. We begin by defining the inner product between two vectors

$$(\mathbf{x}|\mathbf{y}) \equiv \mathbf{x}^T\tilde{\Sigma}^{-1}\mathbf{y}. \quad (5.14)$$

This allows us to rewrite the log-likelihood (Eq. 5.13) as

$$\log p(\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) = -\frac{1}{2}(\tilde{\boldsymbol{\delta}}\mathbf{t} - \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}|\tilde{\boldsymbol{\delta}}\mathbf{t} - \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) - \frac{1}{2}\log\det(2\pi\tilde{\Sigma}), \quad (5.15)$$

where we have suppressed the explicit dependence on the red noise parameters. For simplicity, in this work we will keep the red noise parameters fixed, and thus the resulting covariance matrix (Eq. 5.12) will be constant. Although the previous expressions for the likelihood (Eq. 5.13) are only for a single pulsar, we can easily

generalize this to the full PTA case by writing

$$\tilde{\boldsymbol{\delta t}} = \begin{bmatrix} \tilde{\boldsymbol{\delta t}}_1 \\ \tilde{\boldsymbol{\delta t}}_2 \\ \vdots \\ \tilde{\boldsymbol{\delta t}}_{N_{\text{psr}}} \end{bmatrix}, \quad (5.16)$$

and

$$\tilde{\boldsymbol{\delta t}}_{\text{bwm}} = \begin{bmatrix} \tilde{\boldsymbol{\delta t}}_{\text{bwm},1} \\ \tilde{\boldsymbol{\delta t}}_{\text{bwm},2} \\ \dots \\ \tilde{\boldsymbol{\delta t}}_{\text{bwm},N_{\text{psr}}} \end{bmatrix}. \quad (5.17)$$

In other words, the full residual time series is a vector containing each of the pulsar's timing residual time series. The full R matrix can be written in a block diagonal form

$$R = \begin{bmatrix} R_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & R_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \cdots & R_{N_{\text{psr}}} \end{bmatrix}. \quad (5.18)$$

Finally, for this work, we ignore the small off-diagonal Hellings and Downs (HD) spatial correlations that result from the gravitational-wave background and assume the full PTA noise covariance is block-diagonal (see § 5.5 for further discussion on

this assumption).

$$\tilde{\Sigma} = \begin{bmatrix} \tilde{\Sigma}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \tilde{\Sigma}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \cdots & \tilde{\Sigma}_{N_{\text{psr}}} \end{bmatrix}. \quad (5.19)$$

To construct the full PTA noise covariance for simulations including intrinsic red noise, we can simply add the intrinsic red noise covariance matrices for each pulsar in the corresponding block along the diagonal. We then write the log-likelihood ratio as

$$\log \Lambda = \log \frac{p(\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}})}{p(\tilde{\boldsymbol{\delta}}\mathbf{t}|0)} = (\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) - \frac{1}{2}(\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}), \quad (5.20)$$

where $p(\tilde{\boldsymbol{\delta}}\mathbf{t})$ is the likelihood that these residuals arise from only noise.

We can write the pre-fit signal template $\boldsymbol{\delta}\mathbf{t}_{\text{bwm}}$ as a sum of amplitudes a_i multiplied with template vectors \mathbf{A}^i

$$\boldsymbol{\delta}\mathbf{t}_{\text{bwm}} = \sum_i a_i \mathbf{A}^i, \quad (5.21)$$

Just like for the continuous-wave case [46], there is more than one way to decompose the signal into the product of an amplitude and template vector. In general, the template vectors \mathbf{A}^i carry the “ramp”-shaped template, and therefore the time dependence. In this work, Eqs. 5.26 and 5.42 show two different possible choices for template vectors. For now, we will not specify the decomposition used in Eq. 5.21 but will assume it gives us the full BWM-induced residual time series defined by Eq. 5.1 for each pulsar in the full PTA. This will allow us derive general results

now that we can apply to the coherent (see § 5.3.4) and incoherent (see § 5.3.5) decompositions later.

Using Eq. 5.21 for the BWM-induced residuals, we can write the log-likelihood ratio (Eq 5.20) as

$$\log \Lambda = a_i(\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\mathbf{A}}^i) - \frac{1}{2}a_i a_j(\tilde{\mathbf{A}}^i|\tilde{\mathbf{A}}^j), \quad (5.22)$$

where $\tilde{\mathbf{A}}^i = R\mathbf{A}^i$ are the post-fit time-dependent signal templates. We can then define the matrix elements $\mathbf{N}^i = (\tilde{\boldsymbol{\delta}}\mathbf{t}|\tilde{\mathbf{A}}^i)$ and $\mathbf{M}^{ij} = (\tilde{\mathbf{A}}^i|\tilde{\mathbf{A}}^j)$ and maximize this ratio over the amplitude parameters

$$\frac{\partial \log \Lambda}{\partial a_k} = 0 = \mathbf{N}^k - \mathbf{M}^{ik} a_i. \quad (5.23)$$

This gives the maximum likelihood estimator for the amplitude parameters

$$\hat{a}_i = \mathbf{M}_{ik} \mathbf{N}^k, \quad (5.24)$$

with $M_{ik} = [M^{-1}]^{ik}$, and the maximized log-likelihood ratio

$$\log \Lambda = \frac{1}{2} \mathbf{N}^i \mathbf{M}_{ij} \mathbf{N}^j \equiv \mathcal{F} \quad (5.25)$$

The statistic $2\mathcal{F}$ then follows a χ^2 distribution. As we will show below, depending on our choice of decomposition of the signal, this statistic may be coherent or incoherent. In subsequent sections, we show the choices that give the coherent and incoherent \mathcal{F} -statistics, \mathcal{F}_C and \mathcal{F}_I , respectively. We note that the dimensions of

the \mathbf{M} and \mathbf{N} matrices are determined solely by the number of template vectors. We also show that the coherent statistic \mathcal{F}_C can be computed using two template vectors, with \mathbf{M} a 2×2 matrix, and \mathbf{N} a vector of length 2, whereas the incoherent statistic \mathcal{F}_I only requires one template vector and both \mathbf{M} and \mathbf{N} are scalars.

5.3.4 Coherent \mathcal{F} -statistic

For the coherent statistic, we include the sky-location dependence that comes from the antenna-pattern functions F_+^a and F_\times^a in Eq. 5.1 as part of the time-dependent template \mathbf{A}^i when we decompose the signal using Eq. 5.21. We write this sky-location dependent time-domain template as

$$\mathbf{C}^{\alpha,m}(\hat{\Omega}, t; t_0) = F_{\alpha,m}(\hat{\Omega})\Theta(t - t_0)(t - t_0), \quad (5.26)$$

where $m = +, \times$ specifies the polarization, and α specifies the pulsar. The corresponding amplitudes are

$$c_+(h_0, \psi) = h_0 \cos(2\psi), \quad (5.27)$$

$$c_\times(h_0, \psi) = h_0 \sin(2\psi). \quad (5.28)$$

This allows us to rewrite the pre-fit signal template of Eq. 5.21 as follows

$$\delta \mathbf{t}_{\text{bwm}}(h_0, \psi, \hat{\Omega}, t; t_0) = \sum_{m=+,\times} c_m(h_0, \psi) \mathbf{C}^m(\hat{\Omega}, t; t_0), \quad (5.29)$$

where

$$\mathbf{C}^m(\hat{\Omega}, t; t_0) = \begin{bmatrix} \mathbf{C}^{1,m}(\hat{\Omega}, t; t_0) \\ \mathbf{C}^{2,m}(\hat{\Omega}, t; t_0) \\ \dots \\ \mathbf{C}^{N_{\text{psr}},m}(\hat{\Omega}, t; t_0) \end{bmatrix}. \quad (5.30)$$

and $\tilde{\mathbf{C}}^{+,\times} = R\mathbf{C}^{+,\times}$ are the post-fit time-dependent signal templates. Following the above derivation, the log-likelihood ratio is maximized for the amplitude parameters

$$\frac{\partial \log \Lambda}{\partial c_k} = 0 = (\tilde{\delta \mathbf{t}} | \tilde{\mathbf{C}}^i) - c_i (\tilde{\mathbf{C}}^i | \tilde{\mathbf{C}}^k) \quad (5.31)$$

$$= \mathbf{N}^k - c_i \mathbf{M}^{ik}, \quad (5.32)$$

where the \mathbf{M} and \mathbf{N} matrices are defined by

$$[\mathbf{M}^{ij}] = \begin{bmatrix} (\tilde{\mathbf{C}}^+ | \tilde{\mathbf{C}}^+) & (\tilde{\mathbf{C}}^+ | \tilde{\mathbf{C}}^\times) \\ (\tilde{\mathbf{C}}^\times | \tilde{\mathbf{C}}^+) & (\tilde{\mathbf{C}}^\times | \tilde{\mathbf{C}}^\times) \end{bmatrix} \quad (5.33)$$

$$[\mathbf{N}^i] = \begin{bmatrix} (\tilde{\delta \mathbf{t}} | \tilde{\mathbf{C}}^+) \\ (\tilde{\delta \mathbf{t}} | \tilde{\mathbf{C}}^\times) \end{bmatrix}. \quad (5.34)$$

Then, the maximum log-likelihood ratio is

$$2\mathcal{F}_C = \mathbf{N}^i \mathbf{M}_{ij} \mathbf{N}^j, \quad (5.35)$$

where \mathbf{M}_{ij} is the inverse of \mathbf{M}^{ij} . The full expression of $2\mathcal{F}_C$ includes the sum of two correlated random Gaussian variables with unit variance. It is possible to perform a linear transformation to uncorrelate them (§3 of Ref. [68] shows this in detail). As such, because $2\mathcal{F}_C$ is computed as the sum of two random Gaussian variables, it follows a χ^2 distribution with two degrees of freedom. In the presence of a signal, it becomes a non-central χ^2 -distribution with a non-centrality parameter $\bar{\rho}^2 = (\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}})$ [46, 68]. Note that this non-centrality parameter is also exactly the optimal signal-to-noise ratio (SNR),

$$\langle 2\mathcal{F}_C \rangle = 2 + \bar{\rho}^2 = 2 + (\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}|\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}). \quad (5.36)$$

In addition, for the coherent statistic, we can use the maximum likelihood amplitude estimators to easily find the maximum likelihood amplitude parameters for h_0 and ψ yielding

$$\hat{\psi} = \frac{1}{2} \tan^{-1} \left(\frac{\hat{c}_x}{\hat{c}_+} \right), \quad (5.37)$$

and

$$\hat{h}_0 = \sqrt{(\hat{c}_+)^2 + (\hat{c}_x)^2}. \quad (5.38)$$

We note that this calculation can only be done with two or more pulsars. In the case that the PTA consists of only a single pulsar, the matrix \mathbf{M}_{ij} is non-invertible. This is expected, since a single pulsar cannot simultaneously constrain the GW

strain h_0 and polarization ψ .

5.3.5 Incoherent \mathcal{F} -statistic

The other possibility is to decompose the signal using the time-domain signal template \mathbf{D}

$$\mathbf{D} = \begin{bmatrix} \mathbf{D}^1(t; t_0) \\ \mathbf{D}^2(t; t_0) \\ \vdots \\ \mathbf{D}^{N_{\text{psr}}}(t; t_0) \end{bmatrix}, \quad (5.39)$$

where each element of the column vector is the time-dependent template for a single pulsar

$$\mathbf{D}^\alpha(t, t_0) = \Theta(t - t_0)(t - t_0), \quad (5.40)$$

where again $\Theta(t - t_0)$ is the Heaviside function. The corresponding amplitude for a pulsar indexed by α is

$$d_\alpha(h_0, \psi, \hat{\Omega}) = h_0(\cos(2\psi)F_\alpha^+(\hat{\Omega}) + \sin(2\psi)F_\alpha^\times(\hat{\Omega})). \quad (5.41)$$

The post-fit signal for the entire PTA can then be expressed as

$$\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}(h_0, \psi, \hat{\Omega}, t; t_0) = \begin{bmatrix} d_1(h_0, \psi, \hat{\Omega})\tilde{\mathbf{D}}^1(t; t_0) \\ d_2(h_0, \psi, \hat{\Omega})\tilde{\mathbf{D}}^2(t; t_0) \\ \vdots \\ d_{N_{\text{psr}}}(h_0, \psi, \hat{\Omega})\tilde{\mathbf{D}}^{N_{\text{psr}}}(t; t_0) \end{bmatrix}. \quad (5.42)$$

where again we use the post-fit time-dependent template $\tilde{\mathbf{D}} = R\mathbf{D}$. The log-likelihood ratio for a memory signal to the null-signal model is

$$\log \Lambda = (\tilde{\boldsymbol{\delta}}\mathbf{t} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) - \frac{1}{2}(\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}) \quad (5.43)$$

$$= \sum_{\alpha}^{N_{\text{psr}}} \left[(\tilde{\boldsymbol{\delta}}\mathbf{t}^{\alpha} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}^{\alpha}) - \frac{1}{2}(\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}^{\alpha} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}^{\alpha}) \right] \quad (5.44)$$

$$= \sum_{\alpha}^{N_{\text{psr}}} \left[d_{\alpha}(\tilde{\boldsymbol{\delta}}\mathbf{t}^{\alpha} | \tilde{\mathbf{D}}^{\alpha}) - \frac{1}{2}d_{\alpha}^2(\tilde{\mathbf{D}}^{\alpha} | \tilde{\mathbf{D}}^{\alpha}) \right]. \quad (5.45)$$

Maximizing over the β -th amplitude gives

$$\frac{\partial \log \Lambda}{\partial d_{\beta}} = 0 = (\tilde{\boldsymbol{\delta}}\mathbf{t}^{\beta} | \tilde{\mathbf{D}}^{\beta}) - d_{\beta}(\tilde{\mathbf{D}}^{\beta} | \tilde{\mathbf{D}}^{\beta}) \quad (5.46)$$

Solving yields the maximal likelihood estimator for the β -th amplitude parameter

$$\hat{d}_{\beta} = \frac{(\tilde{\boldsymbol{\delta}}\mathbf{t}^{\beta} | \tilde{\mathbf{D}}^{\beta})}{(\tilde{\mathbf{D}}^{\beta} | \tilde{\mathbf{D}}^{\beta})} = \mathbf{N}_{\beta} \mathbf{M}_{\beta}^{-1} \quad (5.47)$$

with

$$\mathbf{M}_\beta = (\tilde{\mathbf{D}}^\beta | \tilde{\mathbf{D}}^\beta) \quad (5.48)$$

$$\mathbf{N}_\beta = (\tilde{\boldsymbol{\delta}}\mathbf{t} | \tilde{\mathbf{D}}^\beta), \quad (5.49)$$

where \mathbf{M}_β and \mathbf{N}_β are both scalars since there is only one amplitude and one template in this factorization of the signal.

The maximal log-likelihood (which we define as \mathcal{F}_I) is then Eq. 5.43 after substituting the maximum likelihood estimates for the amplitudes

$$\mathcal{F}_I = \sum_{\alpha}^{N_{\text{psr}}} \left[\hat{d}_\alpha(\tilde{\boldsymbol{\delta}}\mathbf{t}^\alpha | \tilde{\mathbf{D}}^\alpha) - \frac{1}{2} \hat{d}_\alpha^2(\tilde{\mathbf{D}}^\alpha | \tilde{\mathbf{D}}^\alpha) \right] \quad (5.50)$$

$$= \sum_{\alpha}^{N_{\text{psr}}} \left[\frac{1}{2} \frac{\mathbf{N}_\alpha^2}{\mathbf{M}_\alpha} \right]. \quad (5.51)$$

Finally,

$$2\mathcal{F}_I = \sum_{\alpha}^{N_{\text{psr}}} \frac{\mathbf{N}_\alpha^2}{\mathbf{M}_\alpha} \quad (5.52)$$

Because $2\mathcal{F}_I$ is a sum of N_{psr} independent Gaussian random variables with unit variance, it follows a χ^2 distribution with N_{psr} degrees of freedom. However, the non-centrality parameter $\bar{\rho}^2 = (\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}})$ remains the same as the non-centrality of the $2\mathcal{F}_C$ statistic

$$\langle 2\mathcal{F}_I \rangle = N_{\text{psr}} + \bar{\rho}^2 = N_{\text{psr}} + (\tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}} | \tilde{\boldsymbol{\delta}}\mathbf{t}_{\text{bwm}}). \quad (5.53)$$

As in the case for the \mathcal{F}_C statistic it is possible to solve for the signal parameters using the maximum likelihood estimators for the amplitude. Since there are four independent parameters in the amplitudes \hat{d}_β (see Eq. 5.41), it is possible to numerically solve for the maximum likelihood estimators by using a minimum of four pulsars.

5.4 Methodology and Results

5.4.1 Simulated Data Sets

To test our statistics we perform two different sets of simulations. First, we simulate PTAs with 40 pulsars placed in random, uniformly distributed sky locations, with time baselines of 10 years. Each pulsar timing residual data set is composed of Gaussian white noise, power-law red noise, and a BWM signal. The signal source is placed at the average of the random right-ascensions and declinations of all the pulsars in the PTA. The injected Gaussian white noise has an amplitude of $\sigma_{\text{WN}} = 100\text{ns}$. The red noise (following Eq. 5.7) has an amplitude of $A = 3.0 \times 10^{-15}$ and a spectral index $\gamma = \frac{13}{3}$. Finally, the memory signal is constructed with a log-strain $\log h_0 = -13.5$, burst epoch $t_0 = \frac{2}{5}T_{\text{PTA}}$, and a polarization angle $\psi = 0$. These injection parameters are summarized in Table 5.1. While red noise from a stochastic gravitational wave background has quadrupolar spatial correlations, we exclude these in this work for simplicity. The red noise injected in these data have the same power spectrum in each pulsar and are spatially

Parameter Name	Variable	Injected Value
Gaussian White Noise Amplitude	σ_{WN}	100 ns
Common Red Noise Amplitude	A_{RN}	3×10^{-15}
Common Red Noise Spectral Index	γ	13/3
Intrinsic Pulsar Red Noise Amplitude	A_{IRN}	log-uniform($-17, -14$)
Intrinsic Pulsar Red Noise Spectral Index	γ_{IRN}	uniform($2, 6$)
Memory Strain Amplitude	h_0	3.16×10^{-14}
Memory Epoch	t_0	4 yr
Memory Polarization	ψ	0 rad

Table 5.1: This table summarizes the values of the injected noise and signal parameters used to create 5000 realizations of simulated data sets. These realizations were used to compute the \mathcal{F} -statistics shown in Figure 5.1. For data sets that include intrinsic red noise, the red noise parameters were fixed to be constant for all realizations.

uncorrelated. This common uncorrelated red noise is often called CURN. We then create 5000 realizations for each combination of noises and signal, and compute both \mathcal{F}_I and \mathcal{F}_C . The results are shown in Figure 5.1.

We also show the parameter estimation capabilities of the \mathcal{F}_C statistic using Eqs. 5.37 and 5.38.

To do this, we simulate data sets with the same white noise ($\sigma_{\text{wn}} = 100\text{ns}$) and red noise ($A_{\text{RN}} = 3.0 \times 10^{-15}, \gamma = 13/3$), and inject signals at three different strengths $A_{\text{bwm}} = \{5 \times 10^{-15}, 8.5 \times 10^{-15}, 1.5 \times 10^{-14}\}$ at a polarization angle of $\psi = \pi/8 \sim 0.39$ rad. We also created data sets which include intrinsic pulsar red noise. The intrinsic red noise, like the common uncorrelated red noise, is parameterized using an amplitude and spectral index. These amplitudes and spectral indices were taken from uniform distributions with $\log_{10} A_{\text{IRN}} \sim \text{uniform}(-17, -14)$ and $\gamma_{\text{IRN}} \sim \text{uniform}(2, 6)$, and fixed throughout all simulations which include pulsar-

Signal Regime	SNR	Injected Parameters		Recovered Parameters	
		$\log_{10} h_0$	ψ	$\log_{10} \hat{h}_0$	$\hat{\psi}$
Weak	1.74 (1.46)	-14.30	0.39	$-14.19^{+0.20}_{-0.28}$ ($-14.17^{+0.20}_{-0.27}$)	$0.29^{+0.31}_{-0.67}$ ($0.29^{+0.32}_{-0.69}$)
Intermediate	5.03 (4.22)	-14.07	0.39	$-14.03^{+0.16}_{-0.21}$ ($-14.02^{+0.16}_{-0.22}$)	$0.38^{+0.20}_{-0.33}$ ($0.37^{+0.21}_{-0.33}$)
Strong	15.67 (13.34)	-13.82	0.39	$-13.81^{+0.11}_{-0.13}$ ($-13.81^{+0.11}_{-0.13}$)	$0.39^{+0.12}_{-0.15}$ ($0.39^{+0.13}_{-0.16}$)

Table 5.2: This table shows the signal-to-noise ratio, the injected signal parameters, and maximum-likelihood estimators for the signal parameters recovered using \mathcal{F}_C in multiple signal regimes over 5000 realizations of simulations for data sets both including and excluding intrinsic pulsar red noise. The values in parentheses are the results for the data sets including intrinsic pulsar red noise. The signal regimes are determined based on the SNR of the signal compared to the injected noise, which was kept constant (see 5.1 for a summary of the injected noise and signals). The maximum-likelihood estimators for the injected memory strain-amplitude are very imprecise for low-SNR signals, but become much more accurate at higher SNR. This is because GW memory signal has a similar power spectrum to a red noise with $\gamma = 13/3$. However, the background is loud compared to the GW memory signal expected from a modest SMBHB merger. The sensitivity of the PTA to memory thus has a lower limit determined by the amplitude of the red noise.

intrinsic red noise. These parameters are also included in Table 5.1.

Table 5.2 shows the injected signal parameters as well as the results of these parameter estimations. We also present 2-dimensional histograms of the recovered signal parameters in Figure 5.2.

5.4.2 Results and Discussion

We compute both the \mathcal{F}_C and \mathcal{F}_I statistics for each combination of injected noise and signals over 5000 realizations of the noise processes. To compare the results with analytic expectations we calculate the non-centrality parameters $\bar{\rho}^2$ as defined in Eqs. 5.36 and 5.53, and we use them to show the consistency between the best-fit χ^2 -distributions and the theoretically predicted χ^2 -distributions in Figure 5.1.

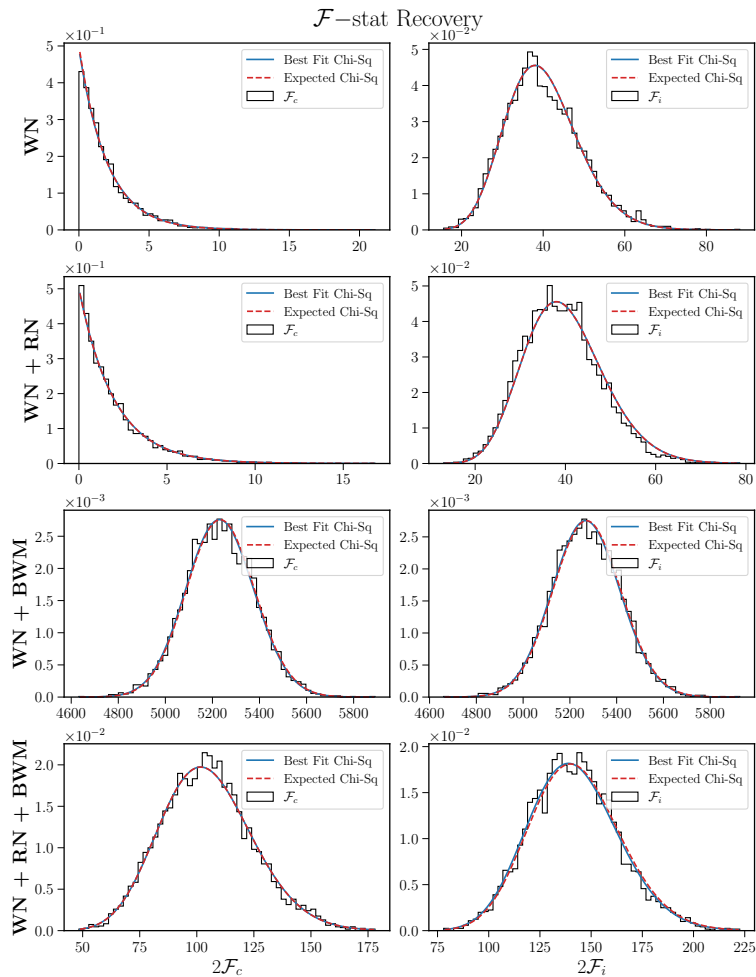


Figure 5.1: This figure shows the best-fit χ^2 -distributions to the \mathcal{F}_C and \mathcal{F}_I statistics in multiple scenarios containing different combinations of gaussian white noise, a common spatially uncorrelated red noise (CURN), and a nonlinear memory signal. We see that in these cases, the theoretically predicted distributions and the theoretical non-centrality parameter match almost exactly with the distributions of the \mathcal{F} -statistics from 5000 simulated realizations. Notably, the SNR decreases significantly when comparing the statistics in the presence of red noise.

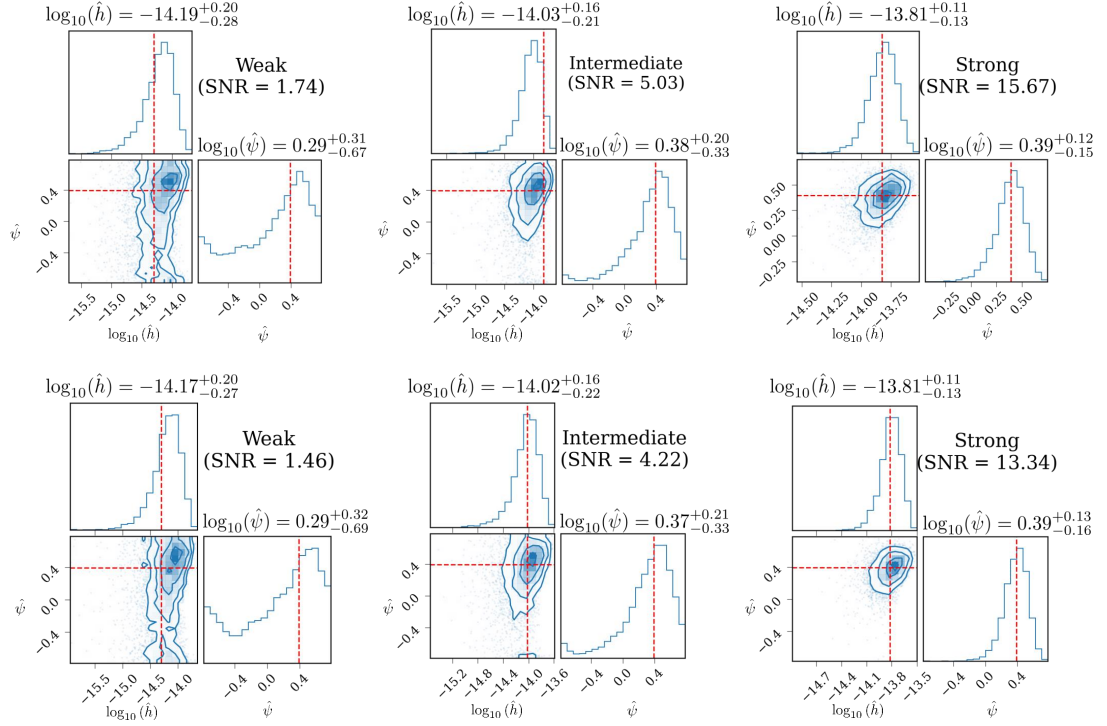


Figure 5.2: This figure shows the results of maximum-likelihood parameter recovery using Eqs. 5.37 and 5.38 in multiple signal regimes. The top row show the recovered parameters for the data sets without injected pulsar-intrinsic red noise, and the bottom row shows the results for the data sets that contain additional pulsar-intrinsic red noise. The summary of the injection and recoveries may be found in Table 5.2. As expected, the maximum likelihood estimators for the strain amplitude h_0 and ψ do very poorly at small SNR (leftmost column, SNR = 1.74 and SNR = 1.46). However, with a moderately large signal (middle column, SNR = 5.03 and SNR = 4.22), the strain amplitude estimator becomes more accurate, while the polarization is not well measured in some realizations. Only when the signal is very loud relative to the noise (rightmost column, SNR = 15.67 and SNR = 13.34) does the maximum likelihood estimator for the ψ become accurate. The weak sensitivity to polarization angle is largely due to the total size (40 pulsars) and distribution (uniform) of the pulsars in this PTA. A PTA with more pulsars distributed more densely near the sky location of the signal source would better constrain the polarization of the signal.

In all cases where there is no signal present, the top four plots in Fig. 5.1, the post-fit covariance matrix properly accounts for both red and white noise. In these cases, the best-fit χ^2 -distribution is nearly identical to the expected null-signal χ^2 -distribution. When a signal is present in our data, bottom four plots in Fig. 5.1, the best-fit non-centrality parameter is in excellent agreement with analytic predictions.

Notably, the addition of red noise into a data set with a signal significantly decreases the non-centrality parameter, which is just the optimal SNR. This is consistent with our expectations, since our BWM signal (which is just a ramp in the time domain) is partially covariant with powerlaw red noise processes. We find that the non-centrality parameter is reduced from approximately ~ 5000 in the case of signal and white noise to ~ 150 in the case of signal, white noise, and red noise. This reduction in the non-centrality parameter means that — even using the exact values of the injection parameters to compute the \mathcal{F} -statistics — our ability to detect a BWM signal is hampered significantly by the presence of red noise. This corroborates the results from refs. [35, 82] regarding the effects of red noise on the detectability of GW memory in PTAs.

We can also see this degeneracy between red noise and nonlinear GW memory in Figure 5.2. In these parameter estimation runs, we attempt to recover the injected signal parameters using the maximum-likelihood estimators in Eqs. 5.37 and 5.38 for simulated data sets both including and excluding additional pulsar-intrinsic red noise. The injected and recovered strain-amplitudes and polarizations are summarized in Table 5.2. The 2-dimensional histograms of the estimated

maximum-likelihood parameters are also shown in each constituent figure.

As seen in the figures, as the injected signal amplitude becomes weaker, parameter recovery becomes less accurate. In the weak-signal case (leftmost column of Fig. 5.2), the amplitude of the memory is small compared to the red noise amplitude. We see that the median recovered maximum likelihood amplitudes in both cases are larger than the injected value ($\log_{10} h_{0,\text{inj}} = -15$), and the polarization angle is not well recovered in many realizations. This is because the memory signal is completely hidden by the red noise. As a result, the coherent statistic, which makes use of the correlations between pulsars in the antenna response to the signal, cannot correctly recover the injected signal parameters. Thus, red noise sets a high noise “floor” for GW memory detection.

In the intermediate case (middle column of Fig. 5.2), we see that the maximum likelihood estimator for the amplitude of the memory signal estimates the injected strain-amplitude accurately ($\log_{10} h_0 = -14.07$) for many realizations. The estimator for the polarization is improved, but still inaccurate in some realizations.

In the case of very strong signals (rightmost column of Fig. 5.2), we see that the median maximum-likelihood estimators over all realizations are more accurate. To obtain an order-of-magnitude estimate of the size of the merging objects and provide a sense of scale, we can use Eq. 1 in [82]

$$h_0 = \frac{1 - \sqrt{8}/3}{24} \frac{G\mu}{c^2 r} \sin^2 \iota (17 + \cos^2 \iota), \quad (5.54)$$

where G is the gravitational constant, μ is the reduced mass of the SMBHB merger

system, c is the speed of light, r is the comoving distance to the SMBHB, and ι is the inclination angle of the binary. At a fiducial comoving distance of $r = 1$ Gpc, the strain amplitudes in the intermediate- and strong- regimes, $h_{\text{intermediate}} = 8.5 \times 10^{-15}$, and $h_{\text{strong}} = 1.5 \times 10^{-14}$, correspond to the merger of two equal-mass SMBHs of $2.7 \times 10^9 M_{\odot}$ and $4.8 \times 10^9 M_{\odot}$, respectively.

Finally, we can see that additional pulsar-intrinsic red noise does not significantly affect the recovery of the parameters when using the \mathcal{F} -statistic. However, the introduction of additional noise clearly reduces the SNR. In these simulations, the red noise amplitude and spectral index were fixed for each pulsar, with only new realizations generated for each data set. Because the memory signal is localized in the sky and only induces significant residuals in nearby pulsars, the relative reduction in SNR is mostly dependent on the red noise in pulsars which are close to the memory source location. As such, these SNR reductions will be sensitive to the source sky location, and it is difficult to draw any more generalized conclusions about effects pulsar-intrinsic red noise.

5.5 Conclusion

We have presented the derivation of two statistics analogous to the continuous GW \mathcal{F} -statistic for use in the detection of nonlinear GW memory in PTA data. These statistics are computed by analytically maximizing the likelihood over the amplitude parameters of the nonlinear GW memory signal model. This may be done in two different ways: coherently, in which case we fix a source sky location

and use the resulting antenna response of the PTA in the construction of the time-domain templates; or incoherently, in which case we do not use the antenna response of the PTA as part of the template and maximize over sky-locations along with the strain and polarization of the GW. The coherent statistic \mathcal{F}_C is the analog of the continuous wave \mathcal{F}_e , and the incoherent statistic \mathcal{F}_I is the analog of \mathcal{F}_p [46].

We have shown that in simple data sets including both Gaussian white and red noise, both statistics follow the expected non-central χ^2 -distributions with the non-centrality parameters given by the optimal SNR $\bar{\rho}^2 = (\tilde{\delta}\mathbf{t}_{\text{bwm}}|\tilde{\delta}\mathbf{t}_{\text{bwm}})$. We have also demonstrated the parameter-estimation capabilities of the coherent statistic. As expected, at low SNR ($\text{SNR} \lesssim 5$), the maximum likelihood estimators for both the strain amplitude and polarization of the nonlinear memory signal are inaccurate. In the intermediate SNR regime, only the amplitude estimator becomes reliably accurate. The polarization estimator only becomes accurate at very high SNR. This can be mitigated by increasing the number of pulsars in the PTA, since increased sky coverage allows the PTA to rule out more source orientations. We also further corroborate the result shown in many other works [e.g 35, 82, 104] that red noise heavily reduces the detectability of nonlinear memory. In the era of background-dominated PTA data sets, it is paramount to look for alternative ways to detect nonlinear memory, such as searching for persistent astrometric deflections as a result of memory [22, 84].

In our work, we have presented the results of computing the coherent and incoherent \mathcal{F} -statistics on data sets that do not contain HD-correlated red noise. In light of the discovery of GW-induced correlations in pulsar timing data sets Agazie

et al. [1], EPTA Collaboration and InPTA Collaboration: et al. [50], Reardon et al. [101], Xu et al. [128], we discuss some of the limitations of these techniques, along with expectations for the case of HD-correlated data, below.

For the coherent \mathcal{F} -statistic, this amounts to including off-diagonal components in the covariance matrix defined in Eq. 5.19. The \mathcal{F}_c -statistic is still χ^2 -distributed with two degrees of freedom, but has a different non-centrality parameter. The non-centrality parameter (the optimal SNR) changes because the covariance matrix used in the evaluation of the inner product in Eq. 5.36 has additional off-diagonal components. Preliminary checks with 5000 realizations of HD-correlated data, similar to those used in section 5.4, indicate that the inclusion of the off-diagonal components gives comparable SNR to the common (uncorrelated) red noise case. We believe this is due to the fact that the off-diagonal contributions from the correlations are small, since the HD coefficients are never greater than 0.5, and taking the covariance to be block diagonal is a reasonable approximation. The scope of these preliminary checks is limited, and more detailed work is required to fully understand the effects of correlated red noise on the \mathcal{F}_c -statistic, optimal SNR, and parameter estimation. It is worth pointing out that because the covariance matrix is no longer block-diagonal, its inversion becomes significantly ($O(N_{\text{psr}}^2)$) more computationally expensive.

On the other hand, the incoherent statistic relies on the likelihood being factorizable, which is a direct result of the block-diagonal structure of the covariance matrix (Eqs. 5.43 and 5.44). As such, the incoherent statistic as presented in this work cannot be adapted to use a covariance matrix with off-diagonal components.

We are currently working to incorporate the the \mathcal{F}_C and \mathcal{F}_I statistics into the standard pulsar timing analysis software extension `enterprise_extensions` package ¹[112]. This will allow the statistics to be used in future searches for nonlinear GW memory in PTA data sets. Compared with a full Bayesian search, the \mathcal{F}_C and \mathcal{F}_I statistics are efficient because we maximize, rather than marginalize, over signal parameters. However, we anticipate that as with other frequentist methods it will be necessary to develop noise-marginalized versions of these statistics to properly account for the broad posterior distributions of the pulsars' intrinsic noise parameters and avoid potential biases (such as those found in the optimal statistic for the GWB [122]). Despite the need for noise marginalization we still expect this procedure to be significantly more efficient than a full Bayesian search, similar to the case of the noise-marginalized optimal statistic for the stochastic GW background [122].

In future studies we will also use these statistics on more realistic data sets to better understand the impacts of different intrinsic red and white noises for each pulsar, different time baselines for different pulsars, searching over source sky locations, and correlations of the common red noise process that will more accurately model the stochastic gravitational-wave background.

¹https://github.com/nanograv/enterprise_extensions

5.6 Acknowledgments

We would like to thank Jeremy Baier and members of the NANOGrav Detection Working group for useful comments on the manuscript. We also extend much gratitude to the anonymous referee, whose questions and suggestions led to a more in-depth understanding of the statistics derived in this paper. JPS, XS, and DRM are members of the NANOGrav Collaboration. The NANOGrav collaboration receives support from National Science Foundation (NSF) Physics Frontiers Center award No. 2020265. This work is also partly supported by the George and Hannah Bolinger Memorial Fund in the College of Science at Oregon State University.

All the codes that support the findings of this article are openly available at [109]. The exact simulated data sets in this work are not published. However, the codes in the cited repository fully specify the random number generation seeds used in each simulation. The exact data sets used in the simulation are available upon reasonable request from the authors.

Chapter 6: Conclusion

The work presented in this dissertation represents significant advances in searches for nonlinear memory. First, we have developed techniques which greatly reduce the computational costs of Bayesian searches for nonlinear memory. Without the use of the efficient factorized likelihood tables, these searches would be computationally impossible to run without access to large clusters. However, using this new method, it is possible to carry these searches out on a single desktop computer in a shorter time with no demonstrably significant sacrifices. As PTA data sets grow and PTA collaborations produce larger and larger combined data sets [8], efficient data analysis techniques will be critical to intensive GW searches like the search for nonlinear memory. We have also demonstrated that this method remains viable in the coming era of background-dominated PTA data sets. However, the method shown in this dissertation assumed that the SGWB induces spatially uncorrelated residuals. This is a simplification, and it should be expected that as PTAs become better at probing the lower frequencies in the GW spectrum, these correlations will become more significant. As such, it may be possible to obtain even better sensitivity to nonlinear memory by using a more accurate model for the SGWB. This will require further innovation; as is, the method presented in this dissertation is incompatible with the inclusion of Hellings-Downs spatial correlations.

The fact that future data sets may include a spatially correlated background

poses a general problem for PTA data analysis. These spatial correlations introduce off-diagonal components into the noise covariance matrix, which causes the calculation of the inverse covariance matrix to be significantly more difficult. This problem compounds as pulsar timing data sets become larger, since matrix inversions scale nonlinearly with the dimensions of the matrix. As such, timing new pulsars and extending older timing baselines can quickly overwhelm available computational resources.

In this work, we have also introduced the frequentist \mathcal{F}_c and \mathcal{F}_i statistics. These statistics are similar to the frequentist statistics used earlier in searches for nonlinear memory. However, they can also be used to maximize polarization in addition to strain. This allows for better localization of a potential signal in the data set, and avoids the need to recompute the statistics over binned polarization angles. It may be the case that future searches for nonlinear memory will have to rely on these frequentist techniques to find preliminary high-SNR “hot spots” on the sky before resorting to the slower Bayesian parameter estimation techniques.

There are also much larger ongoing efforts by the PTA community to speed up GW data analysis in light of the burgeoning problem of larger data sets with more complex noise models. For example, work by Vallisneri et al. [117] shows that it can be possible to use a neural network to approximate a Bayesian posterior distribution by minimizing the Kullback-Leibler divergence of the two distributions. This allows for extremely fast parameter estimation. Additionally, there has been an effort to reimplement the standard GW analysis code (Ellis et al. [49], Taylor et al. [113]) to run on JAX (Bradbury et al. [23]), a Python library optimized for accelerated

computation and parallelization. Such work is crucial to the future of GW data analysis, and will allow for more complex and accurate models of ever-growing pulsar timing data sets. Needless to say, the PTA community is keenly aware of the necessity for improvements in computational efficiency to push GW data analysis smoothly towards the next generation. If these tools become standard, they will very easily facilitate searches for nonlinear memory.

Finally, as we have seen throughout this work, nonlinear memory is extremely covariant with red noise. This may pose a problem for future searches for nonlinear memory. The search for memory in the NANOGrav 12.5-yr data set shows that upper limits cannot be expected to become monotonically lower (in other words, the data sets become uniformly more sensitive to memory) [3]. This corroborates many results from past literature about the detectability of nonlinear memory in the presence of red noise [35, 82, 104]. Simple simulations in the efforts to search for nonlinear memory in the NANOGrav 15-yr data release (Agazie et al. [4], Figure 5) show that in the presence of red noise, the SNR of a nonlinear memory signal does not always increase with longer timing baselines. Instead, the signal eventually saturates when the timing baseline is long enough to see a strong background-dominated global red noise. The main result of this paper further corroborates the idea that memory is very easily contaminated by global red noise, and mimics results as the study of the NANOGrav 12.5-year search for memory. Neither data set offers a strict improvement in sensitivity to memory relative to their respective prior data sets. This result still needs to be verified with formal scaling law calculations (such as those for a stochastic background in Siemens et al.

[107]). The results of such a study will be paramount in strategic decisions about observation campaigns.

It is also worth noting that the results of the 12.5-year search for memory may induce pessimism about future searches. This is understandable since the upper limits set by PTAs are still very far away from predicted memory signals from SMBHB mergers (see Figure 4.5). However, it is important to remember that all GW events produce GW memory. As such, searches for nonlinear memory remain valuable as more generic probes of GWs. Furthermore, it is hard to overstate how monumental a detection of nonlinear memory would be as a confirmation of general relativity theory.

Through this work, we have successfully bridged the gap to background-dominated data sets by making advances in data analysis techniques to expedite traditionally computationally intensive analysis methods. We have also performed the first search for nonlinear memory in a global-red-noise-dominated data set and demonstrated the effects of red noise on upper limits on nonlinear memory. Finally, we have developed a fast, red-noise-robust frequentist technique for searching for and estimating parameters of a nonlinear memory signal. These advancements, along with contemporary work on even faster Bayesian inference methodologies, will be critical in the coming era of GW astronomy.

Bibliography

- [1] Gabriella Agazie, Akash Anumarpudi, Anne M. Archibald, Zaven Arzoumanian, Paul T. Baker, Bence Bécsy, Laura Blecha, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Rand Burnette, Robin Case, Maria Charisi, Shami Chatterjee, Katerina Chatziioannou, Belinda D. Cheeseboro, Siyuan Chen, Tyler Cohen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Kathryn Crowter, Curt J. Cutler, Megan E. DeCesar, Dallas DeGan, Paul B. Demorest, Heling Deng, Timothy Dolch, Brendan Drachler, Justin A. Ellis, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Nate Garver-Daniels, Peter A. Gentile, Kyle A. Gersbach, Joseph Glaser, Deborah C. Good, Kayhan Gültekin, Jeffrey S. Hazboun, Sophie Hourihane, Kristina Islo, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Matthew Kerr, Joey S. Key, Tonia C. Klein, Nima Laal, Michael T. Lam, William G. Lamb, T. Joseph W. Lazio, Natalia Lewandowska, Tyson B. Littenberg, Tingting Liu, Andrea Lommen, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Chung-Pei Ma, Dustin R. Madison, Margaret A. Mattson, Alexander McEwen, James W. McKee, Maura A. McLaughlin, Natasha McMann, Bradley W. Meyers, Patrick M. Meyers, Chiara M. F. Mingarelli, Andrea Mitridate, Priyamvada Natarajan, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Benetge B. P. Perera, Polina Petrov, Nihan S. Pol, Henri A. Radovan, Scott M. Ransom, Paul S. Ray, Joseph D. Romano, Shashwat C. Sardesai, Ann Schmiedekamp, Carl Schmiedekamp, Kai Schmitz, Levi Schult, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Abhimanyu Susobhanan, Joseph K. Swiggum, Jacob Taylor, Stephen R. Taylor, Jacob E. Turner, Caner Unal, Michele Vallisneri, Rutger Van Haasteren, Sarah J. Vigeland, Haley M. Wahl, Qiaohong Wang, Caitlin A. Witt, Olivia Young, and The NANOGrav Collaboration. The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background. *The Astrophysical Journal Letters*, 951(1):L8, July 2023. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/acdac6. URL <https://iopscience.iop.org/article/10.3847/2041-8213/acdac6>.

- [2] Gabriella Agazie, Akash Anumarlapudi, Anne M. Archibald, Zaven Arzoumanian, Paul T. Baker, Bence Bécsy, Laura Blecha, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Robin Case, J. Andrew Casey-Clyde, Maria Charisi, Shami Chatterjee, Tyler Cohen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Kathryn Crowter, Megan E. DeCesar, Paul B. Demorest, Matthew C. Digman, Timothy Dolch, Brendan Drachler, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Nate Garver-Daniels, Peter A. Gentile, Joseph Glaser, Deborah C. Good, Kayhan Gültekin, Jeffrey S. Hazboun, Sophie Hourihane, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Matthew Kerr, Joey S. Key, Nima Laal, Michael T. Lam, William G. Lamb, T. Joseph W. Lazio, Natalia Lewandowska, Tingting Liu, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Chung-Pei Ma, Dustin R. Madison, Alexander McEwen, James W. McKee, Maura A. McLaughlin, Natasha McMann, Bradley W. Meyers, Patrick M. Meyers, Chiara M. F. Mingarelli, Andrea Mitridate, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Benetge B. P. Perera, Polina Petrov, Nihan S. Pol, Henri A. Radovan, Scott M. Ransom, Paul S. Ray, Joseph D. Romano, Shashwat C. Sardesai, Ann Schmiedekamp, Carl Schmiedekamp, Kai Schmitz, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Abhimanyu Susobhanan, Joseph K. Swiggum, Jacob Taylor, Stephen R. Taylor, Jacob E. Turner, Caner Unal, Michele Vallisneri, Rutger Van Haasteren, Sarah J. Vigeland, Haley M. Wahl, Caitlin A. Witt, Olivia Young, and The NANOGrav Collaboration. The NANOGrav 15 yr Data Set: Bayesian Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries. *The Astrophysical Journal Letters*, 951(2):L50, July 2023. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/ace18a. URL <https://iopscience.iop.org/article/10.3847/2041-8213/ace18a>.
- [3] Gabriella Agazie, Zaven Arzoumanian, Paul T. Baker, Bence Bécsy, Laura Blecha, Harsha Blumer, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Rand Burnette, Robin Case, J. Andrew Casey-Clyde, Maria Charisi, Shami Chatterjee, Tyler Cohen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Dallas DeGan, Paul B. Demorest, Timothy Dolch, Brendan Drachler, Justin A. Ellis, Robert D. Ferdman, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Nate Garver-Daniels, Peter A. Gentile, Joseph Glaser, Debo-

rah C. Good, Kayhan Gültekin, Jeffrey S. Hazboun, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey S. Key, Nima Laal, Michael T. Lam, William G. Lamb, T. Joseph W. Lazio, Natalia Lewandowska, Tingting Liu, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Chung-Pei Ma, Dustin R. Madison, Alexander McEwen, James W. McKee, Maura A. McLaughlin, Patrick M. Meyers, Chiara M. F. Mingarelli, Andrea Mitridate, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Nihan S. Pol, Scott M. Ransom, Paul S. Ray, Joseph D. Romano, Shashwat C. Sardesai, Kai Schmitz, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Sophia V. Sosa Fiscella, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Joseph K. Swiggum, Jacob Taylor, Stephen R. Taylor, Jacob E. Turner, Caner Unal, Michele Vallisneri, Sarah J. Vigeland, Haley M. Wahl, Caitlin A. Witt, Olivia Young, and The NANOGrav Collaboration. The NANOGrav 12.5 yr Data Set: Search for Gravitational Wave Memory. *The Astrophysical Journal*, 963(1):61, March 2024. ISSN 0004-637X, 1538-4357. doi: 10.3847/1538-4357/ad0726. URL <https://iopscience.iop.org/article/10.3847/1538-4357/ad0726>.

- [4] Gabriella Agazie, Akash Anumalapudi, Anne M. Archibald, Zaven Arzoumanian, Jeremy G. Baier, Paul T. Baker, Bence Becsy, Laura Blecha, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Rand Burnette, J. Andrew Casey-Clyde, Maria Charisi, Shami Chatterjee, Tyler Cohen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Kathryn Crowter, Megan E. DeCesar, Paul B. Demorest, Heling Deng, Lankeswar Dey, Timothy Dolch, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Emiko C. Gardiner, Nate Garver-Daniels, Peter A. Gentile, Kyle A. Gersbach, Joseph Glaser, Deborah C. Good, Kayhan Gultekin, Jeffrey S. Hazboun, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, David L. Kaplan, Luke Zoltan Kelley, Matthew Kerr, Joey S. Key, Nima Laal, Michael T. Lam, William G. Lamb, Bjorn Larsen, T. Joseph W. Lazio, Natalia Lewandowska, Tingting Liu, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Chung-Pei Ma, Dustin R. Madison, Alexander McEwen, James W. McKee, Maura A. McLaughlin, Natasha McMan, Bradley W. Meyers, Patrick M. Meyers, Chiara M. F. Mingarelli, Andrea Mitridate, Priyamvada Natarajan, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Benetge B. P. Perera, Polina Petrov, Nihan S. Pol, Henri A. Radovan, Scott M. Ransom,

- Paul S. Ray, Jessie C. Runnoe, Alexander Saffer, Shashwat C. Sardesai, Ann Schmiedekamp, Carl Schmiedekamp, Kai Schmitz, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Sophia V. Sosa Fiscella, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Abhimanyu Susobhanan, Joseph K. Swiggum, Jacob Taylor, Stephen R. Taylor, Jacob E. Turner, Caner Unal, Michele Vallisneri, Rutger van Haasteren, Sarah J. Vigeland, Haley M. Wahl, Caitlin A. Witt, David Wright, and Olivia Young. The NANOGrav 15-year Data Set: Search for Gravitational Wave Memory, February 2025. URL <http://arxiv.org/abs/2502.18599>. arXiv:2502.18599 [gr-qc].
- [5] K. Aggarwal, Z. Arzoumanian, P. T. Baker, A. Brazier, P. R. Brook, S. Burke-Spolaor, S. Chatterjee, J. M. Cordes, N. J. Cornish, F. Crawford, H. T. Cromartie, K. Crowter, M. DeCesar, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. C. Ferrara, E. Fonseca, N. Garver-Daniels, P. Gentile, D. Good, J. S. Hazboun, A. M. Holgado, E. A. Huerta, K. Islo, R. Jennings, G. Jones, M. L. Jones, D. L. Kaplan, L. Z. Kelley, J. S. Key, M. T. Lam, T. J. W. Lazio, L. Levin, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, C. M. F. Mingarelli, C. Ng, D. J. Nice, T. T. Pennucci, N. S. Pol, S. M. Ransom, P. S. Ray, X. Siemens, J. Simon, R. Spiewak, I. H. Stairs, D. R. Stinebring, K. Stovall, J. K. Swiggum, S. R. Taylor, M. Vallisneri, R. van Haasteren, S. J. Vigeland, C. A. Witt, W. W. Zhu, and The NANOGrav Collaboration. The NANOGrav 11 yr Data Set: Limits on Gravitational Wave Memory. *The Astrophysical Journal*, 889(1):38, January 2020. ISSN 1538-4357. doi:10.3847/1538-4357/ab6083. URL <https://iopscience.iop.org/article/10.3847/1538-4357/ab6083>.
- [6] Md F. Alam, Zaven Arzoumanian, Paul T. Baker, Harsha Blumer, Keith E. Bohler, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Keeisi Caballero, Richard S. Camuccio, Rachel L. Chamberlain, Shami Chatterjee, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Paul B. Demorest, Timothy Dolch, Justin A. Ellis, Robert D. Ferdman, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Yhamil Garcia, Nathan Garver-Daniels, Peter A. Gentile, Deborah C. Good, Jordan A. Gusdorff, Daniel Halmrast, Jeffrey S. Hazboun, Kristina Islo, Ross J. Jennings, Cody Jessup, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey Shapiro Key, Michael T.

- Lam, T. Joseph W. Lazio, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin R. Madison, Kaleb Maraccini, Maura A. McLaughlin, Chiara M. F. Mingarelli, Cherry Ng, Benjamin M. X. Nguyen, David J. Nice, Timothy T. Pennucci, Nihan S. Pol, Joshua Ramette, Scott M. Ransom, Paul S. Ray, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Joseph K. Swiggum, Stephen R. Taylor, Michael Tripepi, Michele Vallisneri, Sarah J. Vigeland, Caitlin A. Witt, Weiwei Zhu, and The NANOGrav Collaboration. The NANOGrav 12.5 yr Data Set: Observations and Narrowband Timing of 47 Millisecond Pulsars. *The Astrophysical Journal Supplement Series*, 252(1): 4, December 2020. ISSN 1538-4365. doi: 10.3847/1538-4365/abc6a0. URL <https://iopscience.iop.org/article/10.3847/1538-4365/abc6a0>.
- [7] Melissa Anholm, Stefan Ballmer, Jolien D.E. Creighton, Larry R. Price, and Xavier Siemens. Optimal strategies for gravitational wave stochastic background searches in pulsar timing data. *Physical Review D*, 79(8):084030, April 2009. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.79.084030. URL <https://link.aps.org/doi/10.1103/PhysRevD.79.084030>.
- [8] J Antoniadis, Z Arzoumanian, S Babak, M Bailes, A-S Bak Nielsen, P T Baker, C G Bassa, B Bécsy, A Berthreau, M Bonetti, A Brazier, P R Brook, M Burgay, S Burke-Spolaor, R N Caballero, J A Casey-Clyde, A Chalumeau, D J Champion, M Charisi, S Chatterjee, S Chen, I Cognard, J M Cordes, N J Cornish, F Crawford, H T Cromartie, K Crowter, S Dai, M E DeCesar, P B Demorest, G Desvignes, T Dolch, B Drachler, M Falxa, E C Ferrara, W Fiore, E Fonseca, J R Gair, N Garver-Daniels, B Goncharov, D C Good, E Graikou, L Guillemot, Y J Guo, J S Hazboun, G Hobbs, H Hu, K Islo, G H Janssen, R J Jennings, A D Johnson, M L Jones, A R Kaiser, D L Kaplan, R Karuppusamy, M J Keith, L Z Kelley, M Kerr, J S Key, M Kramer, M T Lam, W G Lamb, T J W Lazio, K J Lee, L Lentati, K Liu, J Luo, R S Lynch, A G Lyne, D R Madison, R A Main, R N Manchester, A McEwen, J W McKee, M A McLaughlin, M B Mickaliger, C M F Mingarelli, C Ng, D J Nice, S Osłowski, A Parthasarathy, T T Pennucci, B B P Perera, D Perrodin, A Petiteau, N S Pol, N K Porayko, A Possenti, S M Ransom, P S Ray, D J Reardon, C J Russell, A Samajdar, L M Sampson, S Sanidas, J M Sarkissian, K Schmitz, L Schult, A Sesana, G Shaifullah, R M Shannon, B J Shapiro-Albert, X Siemens, J Simon, T L Smith, L Speri, R Spiewak, I H Stairs, B W Stappers, D R Stinebring, J K Swiggum, S R Taylor, G Theureau,

- C Tiburzi, M Vallisneri, E van der Wateren, A Vecchio, J P W Verbiest, S J Vigeland, H Wahl, J B Wang, J Wang, L Wang, C A Witt, S Zhang, and X J Zhu. The International Pulsar Timing Array second data release: Search for an isotropic gravitational wave background. *Monthly Notices of the Royal Astronomical Society*, 510(4):4873–4887, January 2022. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stab3418. URL <https://academic.oup.com/mnras/article/510/4/4873/6503453>.
- [9] J. Antoniadis, P. Arumugam, S. Arumugam, S. Babak, M. Bagchi, A. S. Bak Nielsen, C. G. Bassa, A. Bathula, A. Berthureau, M. Bonetti, E. Bortolas, P. R. Brook, M. Burgay, R. N. Caballero, A. Chalumeau, D. J. Champion, S. Chanlaridis, S. Chen, I. Cognard, S. Dandapat, D. Deb, S. Desai, G. Desvignes, N. Dhanda-Batra, C. Dwivedi, M. Falxa, I. Ferranti, R. D. Ferdman, A. Franchini, J. R. Gair, B. Goncharov, A. Gopakumar, E. Graikou, J. M. Grießmeier, L. Guillemot, Y. J. Guo, Y. Gupta, S. Hisano, H. Hu, F. Iraci, D. Izquierdo-Villalba, J. Jang, J. Jawor, G. H. Janssen, A. Jessner, B. C. Joshi, F. Kareem, R. Karuppusamy, E. F. Keane, M. J. Keith, D. Kharbanda, T. Kikunaga, N. Kolhe, M. Kramer, M. A. Krishnakumar, K. Lackeos, K. J. Lee, K. Liu, Y. Liu, A. G. Lyne, J. W. McKee, Y. Maan, R. A. Main, S. Manzini, M. B. Mickaliger, I. C. Nitu, K. Nobleson, A. K. Paladi, A. Parthasarathy, B. B. P. Perera, D. Perrodin, A. Petiteau, N. K. Porayko, A. Possenti, T. Prabu, H. Quelquejay Leclere, P. Rana, A. Samajdar, S. A. Sanidas, A. Sesana, G. Shaifullah, J. Singha, L. Speri, R. Spiewak, A. Srivastava, B. W. Stappers, M. Surnis, S. C. Susarla, A. Subbhanan, K. Takahashi, P. Tarafdar, G. Theureau, C. Tiburzi, E. van der Wateren, A. Vecchio, V. Venkatraman Krishnan, J. P. W. Verbiest, J. Wang, L. Wang, and Z. Wu. The second data release from the European Pulsar Timing Array V. Search for continuous gravitational wave signals, May 2024. URL <http://arxiv.org/abs/2306.16226>. arXiv:2306.16226 [astro-ph, physics:gr-qc].
- [10] K G Arun, Luc Blanchet, Bala R Iyer, and Moh'd S S Qusailah. The 2.5PN gravitational wave polarizations from inspiralling compact binaries in circular orbits. *Classical and Quantum Gravity*, 21(15):3771–3801, August 2004. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/21/15/010. URL <https://iopscience.iop.org/article/10.1088/0264-9381/21/15/010>.
- [11] Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, P. B. Demorest, X. Deng,

- T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, N. Garver-Daniels, F. Jenet, G. Jones, V. M. Kaspi, M. Koop, M. T. Lam, T. J. W. Lazio, L. Levin, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, D. J. Nice, N. Palliyaguru, T. T. Pennucci, S. M. Ransom, X. Siemens, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, M. Vallisneri, R. van Haasteren, Y. Wang, W. W. Zhu, and NANOGrav Collaboration. NANOGrav CONSTRAINTS ON GRAVITATIONAL WAVE BURSTS WITH MEMORY. *The Astrophysical Journal*, 810(2):150, September 2015. ISSN 1538-4357. doi: 10.1088/0004-637X/810/2/150. URL <https://iopscience.iop.org/article/10.1088/0004-637X/810/2/150>.
- [12] Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, K. Crowter, P. B. Demorest, X. Deng, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, N. Garver-Daniels, M. E. Gonzalez, F. Jenet, G. Jones, M. L. Jones, V. M. Kaspi, M. Koop, M. T. Lam, T. J. W. Lazio, L. Levin, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, C. M. F. Mingarelli, D. J. Nice, N. Palliyaguru, T. T. Pennucci, S. M. Ransom, L. Sampson, S. A. Sanidas, A. Sesana, X. Siemens, J. Simon, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, S. R. Taylor, M. Vallisneri, R. van Haasteren, Y. Wang, W. W. Zhu, and The NANOGrav Collaboration. THE NANOGrav NINE-YEAR DATA SET: LIMITS ON THE ISOTROPIC STOCHASTIC GRAVITATIONAL WAVE BACKGROUND. *The Astrophysical Journal*, 821(1):13, April 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/821/1/13. URL <https://iopscience.iop.org/article/10.3847/0004-637X/821/1/13>.
- [13] Z. Arzoumanian, P. T. Baker, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, F. Crawford, H. Thankful Cromartie, K. Crowter, M. DeCesar, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Ferrara, W. M. Folkner, E. Fonseca, N. Garver-Daniels, P. A. Gentile, R. Haas, J. S. Hazboun, E. A. Huerta, K. Islo, G. Jones, M. L. Jones, D. L. Kaplan, V. M. Kaspi, M. T. Lam, T. J. W. Lazio, L. Levin, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, C. M. F. Mingarelli, C. Ng, D. J. Nice, R. S. Park, T. T. Pennucci, N. S. Pol, S. M. Ransom, P. S. Ray, A. Rasskazov, X. Siemens, J. Simon, R. Spiewak, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, S. R. Taylor, M. Vallisneri,

- R. van Haasteren, S. Vigeland, W. W. Zhu, and The NANOGrav Collaboration. The NANOGrav 11 Year Data Set: Pulsar-timing Constraints on the Stochastic Gravitational-wave Background. *The Astrophysical Journal*, 859 (1):47, May 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aabd3b. URL <https://iopscience.iop.org/article/10.3847/1538-4357/aabd3b>.
- [14] Zaven Arzoumanian, Adam Brazier, Sarah Burke-Spolaor, Sydney Chamberlin, Shami Chatterjee, Brian Christy, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Kathryn Crowter, Megan E. DeCesar, Paul B. Demorest, Timothy Dolch, Justin A. Ellis, Robert D. Ferdman, Elizabeth C. Ferrara, Emmanuel Fonseca, Nathan Garver-Daniels, Peter A. Gentile, Daniel Halmrast, E. A. Huerta, Fredrick A. Jenet, Cody Jessup, Glenn Jones, Megan L. Jones, David L. Kaplan, Michael T. Lam, T. Joseph W. Lazio, Lina Levin, Andrea Lommen, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin Madison, Allison M. Matthews, Maura A. McLaughlin, Sean T. McWilliams, Chiara Mingarelli, Cherry Ng, David J. Nice, Timothy T. Pennucci, Scott M. Ransom, Paul S. Ray, Xavier Siemens, Joseph Simon, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Joseph K. Swiggum, Stephen R. Taylor, Michele Vallisneri, Rutger van Haasteren, Sarah J. Vigeland, Weiwei Zhu, and The NANOGrav Collaboration. The NANOGrav 11-year Data Set: High-precision Timing of 45 Millisecond Pulsars. *The Astrophysical Journal Supplement Series*, 235 (2):37, April 2018. ISSN 1538-4365. doi: 10.3847/1538-4365/aab5b0. URL <https://iopscience.iop.org/article/10.3847/1538-4365/aab5b0>.
- [15] Zaven Arzoumanian, Paul T. Baker, Harsha Blumer, Bence Bécsy, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Shami Chatterjee, Siyuan Chen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Paul B. Demorest, Timothy Dolch, Justin A. Ellis, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Nathan Garver-Daniels, Peter A. Gentile, Deborah C. Good, Jeffrey S. Hazboun, A. Miguel Holgado, Kristina Iso, Ross J. Jennings, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey Shapiro Key, Nima Laal, Michael T. Lam, T. Joseph W. Lazio, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin R. Madison, Maura A. McLaughlin, Chiara M. F. Mingarelli, Cherry Ng, David J. Nice, Timothy T. Pennucci, Nihan S. Pol, Scott M. Ransom, Paul S. Ray, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Renée Spiewak, Ingrid H. Stairs,

- Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Joseph K. Swiggum, Stephen R. Taylor, Jacob E. Turner, Michele Vallisneri, Sarah J. Vigeland, and Caitlin A. Witt. The NANOGrav 12.5 yr Data Set: Search for an Isotropic Stochastic Gravitational-wave Background. *The Astrophysical Journal Letters*, 905(2):L34, December 2020. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/abd401. URL <https://iopscience.iop.org/article/10.3847/2041-8213/abd401>.
- [16] Zaven Arzoumanian, Paul T. Baker, Harsha Blumer, Bence Bécsy, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Maria Charisi, Shami Chatterjee, Siyuan Chen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Dallas M. DeGan, Paul B. Demorest, Timothy Dolch, Brendan Drachler, Justin A. Ellis, Elizabeth C. Ferrara, William Fiore, Emmanuel Fonseca, Nathan Garver-Daniels, Peter A. Gentile, Deborah C. Good, Jeffrey S. Hazboun, A. Miguel Holgado, Kristina Islo, Ross J. Jennings, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey Shapiro Key, Nima Laal, Michael T. Lam, T. Joseph W. Lazio, Duncan R. Lorimer, Tingting Liu, Jing Luo, Ryan S. Lynch, Dustin R. Madison, Alexander McEwen, Maura A. McLaughlin, Chiara M. F. Mingarelli, Cherry Ng, David J. Nice, Ken D. Olum, Timothy T. Pennucci, Nihan S. Pol, Scott M. Ransom, Paul S. Ray, Joseph D. Romano, Shashwat C. Sardesai, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Magdalena S. Siwek, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Jerry P. Sun, Joseph K. Swiggum, Stephen R. Taylor, Jacob E. Turner, Michele Vallisneri, Sarah J. Vigeland, Haley M. Wahl, Caitlin A. Witt, and The NANOGrav Collaboration. The NANOGrav 12.5-year Data Set: Search for Non-Einsteinian Polarization Modes in the Gravitational-wave Background. *The Astrophysical Journal Letters*, 923(2):L22, December 2021. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/ac401c. URL <https://iopscience.iop.org/article/10.3847/2041-8213/ac401c>.
- [17] Zaven Arzoumanian, Paul T. Baker, Laura Blecha, Harsha Blumer, Adam Brazier, Paul R. Brook, Sarah Burke-Spolaor, Bence Bécsy, J. Andrew Casey-Clyde, Maria Charisi, Shami Chatterjee, Siyuan Chen, James M. Cordes, Neil J. Cornish, Fronefield Crawford, H. Thankful Cromartie, Megan E. DeCesar, Paul B. Demorest, Timothy Dolch, Brendan Drachler, Justin A. Ellis, E. C. Ferrara, William Fiore, Emmanuel Fonseca, Gabriel E. Freedman, Nathan Garver-Daniels, Peter A. Gentile, Joseph

- Glaser, Deborah C. Good, Kayhan Gültekin, Jeffrey S. Hazboun, Ross J. Jennings, Aaron D. Johnson, Megan L. Jones, Andrew R. Kaiser, David L. Kaplan, Luke Zoltan Kelley, Joey Shapiro Key, Nima Laal, Michael T. Lam, William G Lamb, T. Joseph W. Lazio, Natalia Lewandowska, Tingting Liu, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin R. Madison, Alexander McEwen, Maura A. McLaughlin, Chiara M. F. Mingarelli, Cherry Ng, David J. Nice, Stella Koch Ocker, Ken D. Olum, Timothy T. Pennucci, Nihan S. Pol, Scott M. Ransom, Paul S. Ray, Joseph D. Romano, Brent J. Shapiro-Albert, Xavier Siemens, Joseph Simon, Magdalena Siwek, Renée Spiewak, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Joseph K. Swiggum, Jessica Sydnor, Stephen R. Taylor, Jacob E. Turner, Michele Vallisneri, Sarah J. Vigeland, Haley M. Wahl, Gregory Walsh, Caitlin A. Witt, and Olivia Young. The NANOGrav 12.5 yr Data Set: Bayesian Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries. *The Astrophysical Journal Letters*, 951(2):L28, July 2023. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/acdbc7. URL <https://iopscience.iop.org/article/10.3847/2041-8213/acdbc7>.
- [18] Stanislav Babak and Alberto Sesana. Resolving multiple supermassive black hole binaries with pulsar timing arrays. *Physical Review D*, 85(4):044034, February 2012. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.85.044034. URL <https://link.aps.org/doi/10.1103/PhysRevD.85.044034>.
- [19] Luc Blanchet and Thibault Damour. Hereditary effects in gravitational radiation. *Physical Review D*, 46(10):4304–4319, November 1992. ISSN 0556-2821. doi: 10.1103/PhysRevD.46.4304. URL <https://link.aps.org/doi/10.1103/PhysRevD.46.4304>.
- [20] Oliver M. Boersma, David A. Nichols, and Patricia Schmidt. Forecasts for detecting the gravitational-wave memory effect with Advanced LIGO and Virgo. *Physical Review D*, 101(8):083026, April 2020. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.101.083026. URL <https://link.aps.org/doi/10.1103/PhysRevD.101.083026>.
- [21] Laura G. Book and Éanna É. Flanagan. Astrometric effects of a stochastic gravitational wave background. *Physical Review D*, 83(2):024024, January 2011. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.83.024024. URL <https://link.aps.org/doi/10.1103/PhysRevD.83.024024>.

- [22] Töre Boybeyi, Vuk Mandic, and Alexandros Papageorgiou. Astrometric deflections from gravitational wave memory accumulation over cosmological scales. *Physical Review D*, 110(4):043047, August 2024. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.110.043047. URL <https://link.aps.org/doi/10.1103/PhysRevD.110.043047>.
- [23] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.
- [24] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Physical Review D*, 59(8):084006, March 1999. ISSN 0556-2821, 1089-4918. doi: 10.1103/PhysRevD.59.084006. URL <https://link.aps.org/doi/10.1103/PhysRevD.59.084006>.
- [25] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D*, 62(6):064015, August 2000. ISSN 0556-2821, 1089-4918. doi: 10.1103/PhysRevD.62.064015. URL <https://link.aps.org/doi/10.1103/PhysRevD.62.064015>.
- [26] Brian J. Burt, Andrea N. Lommen, and Lee S. Finn. OPTIMIZING PULSAR TIMING ARRAYS TO MAXIMIZE GRAVITATIONAL WAVE SINGLE-SOURCE DETECTION: A FIRST CUT. *The Astrophysical Journal*, 730(1):17, March 2011. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/730/1/17. URL <https://iopscience.iop.org/article/10.1088/0004-637X/730/1/17>.
- [27] Bence Bécsy and Neil J Cornish. Bayesian search for gravitational wave bursts in pulsar timing array data. *Classical and Quantum Gravity*, 38(9):095012, May 2021. ISSN 0264-9381, 1361-6382. doi: 10.1088/1361-6382/abf1c6. URL <https://iopscience.iop.org/article/10.1088/1361-6382/abf1c6>.
- [28] Bradley P. Carlin and Siddhartha Chib. Bayesian Model Choice Via Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(3):473–484, September 1995. ISSN 00359246. doi: 10.1111/j.2517-6161.1995.tb02042.x. URL <https://onlinelibrary.wiley.com/doi/10.1111/j.2517-6161.1995.tb02042.x>.

- [29] A Chalumeau, S Babak, A Petiteau, S Chen, A Samajdar, R N Caballero, G Theureau, L Guillemot, G Desvignes, A Parthasarathy, K Liu, G Shaifullah, H Hu, E van der Wateren, J Antoniadis, A-S Bak Nielsen, C G Bassa, A Berthereau, M Burgay, D J Champion, I Cognard, M Falxa, R D Ferdman, P C C Freire, J R Gair, E Graikou, Y J Guo, J Jang, G H Janssen, R Karuppusamy, M J Keith, M Kramer, K J Lee, X J Liu, A G Lyne, R A Main, J W McKee, M B Mickaliger, B B P Perera, D Perrodin, N K Porayko, A Possenti, S A Sanidas, A Sesana, L Speri, B W Stappers, C Tiburzi, A Vecchio, J P W Verbiest, J Wang, L Wang, and H Xu. Noise analysis in the European Pulsar Timing Array data release 2 and its implications on the gravitational-wave background search. *Monthly Notices of the Royal Astronomical Society*, 509(4):5538–5558, 11 2021. ISSN 0035-8711. doi: 10.1093/mnras/stab3283. URL <https://doi.org/10.1093/mnras/stab3283>.
- [30] S Chen, R N Caballero, Y J Guo, A Chalumeau, K Liu, G Shaifullah, K J Lee, S Babak, G Desvignes, A Parthasarathy, H Hu, E van der Wateren, J Antoniadis, A-S Bak Nielsen, C G Bassa, A Berthereau, M Burgay, D J Champion, I Cognard, M Falxa, R D Ferdman, P C C Freire, J R Gair, E Graikou, L Guillemot, J Jang, G H Janssen, R Karuppusamy, M J Keith, M Kramer, X J Liu, A G Lyne, R A Main, J W McKee, M B Mickaliger, B B P Perera, D Perrodin, A Petiteau, N K Porayko, A Possenti, A Samajdar, S A Sanidas, A Sesana, L Speri, B W Stappers, G Theureau, C Tiburzi, A Vecchio, J P W Verbiest, J Wang, L Wang, and H Xu. Common-red-signal analysis with 24-yr high-precision timing of the European Pulsar Timing Array: inferences in the stochastic gravitational-wave background search. *Monthly Notices of the Royal Astronomical Society*, 508(4):4970–4993, October 2021. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stab2833. URL <https://academic.oup.com/mnras/article/508/4/4970/6410749>.
- [31] Zu-Cheng Chen, Chen Yuan, and Qing-Guo Huang. Non-tensorial gravitational wave background in NANOGrav 12.5-year data set. *Sci. China Phys. Mech. Astron.*, 64(12):120412, December 2021. ISSN 1674-7348, 1869-1927. doi: 10.1007/s11433-021-1797-y. URL <https://link.springer.com/10.1007/s11433-021-1797-y>.
- [32] Demetrios Christodoulou. Nonlinear nature of gravitation and gravitational-wave experiments. *Physical Review Letters*, 67(12):1486–1489, September 1991. ISSN 0031-9007. doi: 10.1103/PhysRevLett.67.1486. URL <https://link.aps.org/doi/10.1103/PhysRevLett.67.1486>.

- [33] Brian Christy, Ryan Anella, Andrea Lommen, Lee Samuel Finn, Richard Camuccio, and Emma Handzo. OPTIMIZATION OF NANOGRAV'S TIME ALLOCATION FOR MAXIMUM SENSITIVITY TO SINGLE SOURCES. *The Astrophysical Journal*, 794(2):163, October 2014. ISSN 1538-4357. doi: 10.1088/0004-637X/794/2/163. URL <https://iopscience.iop.org/article/10.1088/0004-637X/794/2/163>.
- [34] Ismaël Cognard and Donald C. Backer. A Microglitch in the Millisecond Pulsar PSR B1821-24 in M28. *The Astrophysical Journal*, 612(2):L125–L127, September 2004. ISSN 0004-637X, 1538-4357. doi: 10.1086/424692. URL <https://iopscience.iop.org/article/10.1086/424692>.
- [35] J. M. Cordes and F. A. Jenet. DETECTING GRAVITATIONAL WAVE MEMORY WITH PULSAR TIMING. *The Astrophysical Journal*, 752(1):54, June 2012. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/752/1/54. URL <https://iopscience.iop.org/article/10.1088/0004-637X/752/1/54>.
- [36] Neil J Cornish and Edward K Porter. The search for massive black hole binaries with LISA. *Classical and Quantum Gravity*, 24(23):5729–5755, December 2007. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/24/23/001. URL <https://iopscience.iop.org/article/10.1088/0264-9381/24/23/001>.
- [37] Neil J. Cornish, Logan O'Beirne, Stephen R. Taylor, and Nicolás Yunes. Constraining Alternative Theories of Gravity Using Pulsar Timing Arrays. *Phys. Rev. Lett.*, 120(18):181101, May 2018. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.120.181101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.181101>.
- [38] Curt Cutler, Sarah Burke-Spolaor, Michele Vallisneri, Joseph Lazio, and Walid Majid. The gravitational-wave discovery space of pulsar timing arrays. *Physical Review D*, 89(4):042003, February 2014. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.89.042003. URL <https://link.aps.org/doi/10.1103/PhysRevD.89.042003>.
- [39] Thibault Damour and Alexander Vilenkin. Gravitational Wave Bursts from Cosmic Strings. *Physical Review Letters*, 85(18):3761–3764, October 2000. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.85.3761. URL <https://link.aps.org/doi/10.1103/PhysRevLett.85.3761>.

- [40] P. B. Demorest. *Measuring the gravitational wave background using precision pulsar timing*. PhD thesis, University of California, Berkeley, August 2007.
- [41] P. B. Demorest, R. D. Ferdman, M. E. Gonzalez, D. Nice, S. Ransom, I. H. Stairs, Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlain, J. M. Cordes, J. Ellis, L. S. Finn, P. Freire, S. Giampanis, F. Jenet, V. M. Kaspi, J. Lazio, A. N. Lommen, M. McLaughlin, N. Palliyaguru, D. Perrodin, R. M. Shannon, X. Siemens, D. Stinebring, J. Swiggum, and W. W. Zhu. LIMITS ON THE STOCHASTIC GRAVITATIONAL WAVE BACKGROUND FROM THE NORTH AMERICAN NANOHERTZ OBSERVATORY FOR GRAVITATIONAL WAVES. *The Astrophysical Journal*, 762(2):94, January 2013. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/762/2/94. URL <https://iopscience.iop.org/article/10.1088/0004-637X/762/2/94>.
- [42] Heling Deng, Bence Bécsy, Xavier Siemens, Neil J. Cornish, and Dustin R. Madison. Searching for gravitational wave burst in pulsar-timing-array data with piecewise linear functions. *Physical Review D*, 108(10):102007, November 2023. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.108.102007. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.102007>.
- [43] G. Desvignes, R. N. Caballero, L. Lentati, J. P. W. Verbiest, D. J. Champion, B. W. Stappers, G. H. Janssen, P. Lazarus, S. Osłowski, S. Babak, C. G. Bassa, P. Brem, M. Burgay, I. Cognard, J. R. Gair, E. Graikou, L. Guillemot, J. W. T. Hessels, A. Jessner, C. Jordan, R. Karuppusamy, M. Kramer, A. Lassus, K. Lazaridis, K. J. Lee, K. Liu, A. G. Lyne, J. McKee, C. M. F. Mingarelli, D. Perrodin, A. Petiteau, A. Possenti, M. B. Purver, P. A. Rosado, S. Sanidas, A. Sesana, G. Shaifullah, R. Smits, S. R. Taylor, G. Theureau, C. Tiburzi, R. van Haasteren, and A. Vecchio. High-precision timing of 42 millisecond pulsars with the European Pulsar Timing Array. *Monthly Notices of the Royal Astronomical Society*, 458(3):3341–3380, May 2016. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw483. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stw483>.
- [44] S. Detweiler. Pulsar timing measurements and the search for gravitational waves. *The Astrophysical Journal*, 234:1100–1104, December 1979. ISSN 0004-637X. doi: 10.1086/157593. URL <https://ui.adsabs.harvard.edu/abs/1979ApJ...234.1100D>. ADS Bibcode: 1979ApJ...234.1100D.

- [45] Ron DuPlain, John Benson, and Eric Sessoms. Data Vault: providing simple web access to NRAO data archives. In Alan Bridger and Nicole M. Radziwill, editors, *Advanced Software and Control for Astronomy II*, volume 7019 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 70191A, August 2008. doi: 10.1117/12.789402.
- [46] J. A. Ellis, X. Siemens, and J. D. E. Creighton. OPTIMAL STRATEGIES FOR CONTINUOUS GRAVITATIONAL WAVE DETECTION IN PULSAR TIMING ARRAYS. *The Astrophysical Journal*, 756(2):175, September 2012. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/756/2/175. URL <https://iopscience.iop.org/article/10.1088/0004-637X/756/2/175>.
- [47] Justin Ellis and Rutger Van Haasteren. Jellis18/Ptmcmcsampler: Official Release, October 2017. URL <https://zenodo.org/record/1037579>.
- [48] Justin A. Ellis, Xavier Siemens, and Rutger van Haasteren. An Efficient Approximation to the Likelihood for Gravitational Wave Stochastic Background Detection Using Pulsar Timing Data. *Astrophys. J.*, 769:63, 2013. doi: 10.1088/0004-637X/769/1/63.
- [49] Justin A. Ellis, Michele Vallisneri, Stephen R. Taylor, and Paul T. Baker. ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, September 2020. URL <https://zenodo.org/record/4059815>.
- [50] EPTA Collaboration and InPTA Collaboration:, J. Antoniadis, P. Arumugam, S. Arumugam, S. Babak, M. Bagchi, A.-S. Bak Nielsen, C. G. Bassa, A. Bathula, A. Berthereau, M. Bonetti, E. Bortolas, P. R. Brook, M. Burgay, R. N. Caballero, A. Chalumeau, D. J. Champion, S. Chanlaridis, S. Chen, I. Cognard, S. Dandapat, D. Deb, S. Desai, G. Desvignes, N. Dhanda-Batra, C. Dwivedi, M. Falxa, R. D. Ferdman, A. Franchini, J. R. Gair, B. Goncharov, A. Gopakumar, E. Graikou, J.-M. Grießmeier, L. Guillemot, Y. J. Guo, Y. Gupta, S. Hisano, H. Hu, F. Iraci, D. Izquierdo-Villalba, J. Jang, J. Jawor, G. H. Janssen, A. Jessner, B. C. Joshi, F. Kareem, R. Karuppusamy, E. F. Keane, M. J. Keith, D. Kharbanda, T. Kikunaga, N. Kolhe, M. Kramer, M. A. Krishnakumar, K. Lackeos, K. J. Lee, K. Liu, Y. Liu, A. G. Lyne, J. W. McKee, Y. Maan, R. A. Main, M. B. Mickaliger, I. C. Nițu, K. Nobleson, A. K. Paladi, A. Parthasarathy, B. B. P. Perera, D. Perrodin, A. Petiteau, N. K. Porayko, A. Possenti, T. Prabu,

- H. Quelquejay Leclere, P. Rana, A. Samajdar, S. A. Sanidas, A. Sesana, G. Shaifullah, J. Singha, L. Speri, R. Spiewak, A. Srivastava, B. W. Stappers, M. Surnis, S. C. Susarla, A. Susobhanan, K. Takahashi, P. Tarafdar, G. Theureau, C. Tiburzi, E. Van Der Wateren, A. Vecchio, V. Venkatraman Krishnan, J. P. W. Verbiest, J. Wang, L. Wang, and Z. Wu. The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals. *Astronomy & Astrophysics*, 678:A50, October 2023. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202346844. URL <https://www.aanda.org/10.1051/0004-6361/202346844>.
- [51] Frank B. Estabrook and Hugo D. Wahlquist. Response of Doppler spacecraft tracking to gravitational radiation. *General Relativity and Gravitation*, 6(5):439–447, October 1975. ISSN 0001-7701, 1572-9532. doi: 10.1007/BF00762449. URL <http://link.springer.com/10.1007/BF00762449>.
- [52] M Falxa, S Babak, P T Baker, B Bécsy, A Chalumeau, S Chen, Z Chen, N J Cornish, L Guillemot, J S Hazboun, C M F Mingarelli, A Parthasarathy, A Petiteau, N S Pol, A Sesana, S B Spolaor, S R Taylor, G Theureau, M Vallisneri, S J Vigeland, C A Witt, X Zhu, J Antoniadis, Z Arzoumanian, M Bailes, N D R Bhat, L Blecha, A Brazier, P R Brook, N Caballero, A D Cameron, J A Casey-Clyde, D Champion, M Charisi, S Chatterjee, I Cognard, J M Cordes, F Crawford, H T Cromartie, K Crowter, S Dai, M E DeCesar, P B Demorest, G Desvignes, T Dolch, B Drachler, Y Feng, E C Ferrara, W Fiore, E Fonseca, N Garver-Daniels, J Glaser, B Goncharov, D C Good, J Griessmeier, Y J Guo, K Gültekin, G Hobbs, H Hu, K Islo, J Jang, R J Jennings, A D Johnson, M L Jones, J Kaczmarek, A R Kaiser, D L Kaplan, M Keith, L Z Kelley, M Kerr, J S Key, N Laal, M T Lam, W G Lamb, T J W Lazio, K Liu, T Liu, J Luo, R S Lynch, D R Madison, R Main, R Manchester, A McEwen, J McKee, M A McLaughlin, C Ng, D J Nice, S Ocker, K D Olum, S Osłowski, T T Pennucci, B B P Perera, D Perrodin, N Porayko, A Possenti, H Quelquejay-Leclere, S M Ransom, P S Ray, D J Reardon, C J Russell, A Samajdar, J Sarkissian, L Schult, G Shaifullah, R M Shannon, B J Shapiro-Albert, X Siemens, J J Simon, M Siwek, T L Smith, L Speri, R Spiewak, I H Stairs, B Stappers, D R Stinebring, J K Swiggum, C Tiburzi, J Turner, A Vecchio, J P W Verbiest, H Wahl, S Q Wang, J Wang, J Wang, Z Wu, L Zhang, S Zhang, and IPTA Collaboration. Searching for continuous Gravitational Waves in the second data release of the International Pulsar Timing Array. *Monthly Notices of*

- the Royal Astronomical Society*, 521(4):5077–5086, March 2023. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stad812. URL <https://academic.oup.com/mnras/article/521/4/5077/7083157>.
- [53] Marc Favata. Nonlinear Gravitational-Wave Memory from Binary Black Hole Mergers. *The Astrophysical Journal*, 696:L159–L162, May 2009. ISSN 0004-637X. doi: 10.1088/0004-637X/696/2/L159. URL <https://ui.adsabs.harvard.edu/abs/2009ApJ...696L.159F>. ADS Bibcode: 2009ApJ...696L.159F.
- [54] Marc Favata. Post-Newtonian corrections to the gravitational-wave memory for quasicircular, inspiralling compact binaries. *Physical Review D*, 80:024002, July 2009. ISSN 1550-7998/0556-2821. doi: 10.1103/PhysRevD.80.024002. URL <https://ui.adsabs.harvard.edu/abs/2009PhRvD..80b4002F>. ADS Bibcode: 2009PhRvD..80b4002F.
- [55] Marc Favata. The gravitational-wave memory effect. *Classical and Quantum Gravity*, 27(8):084036, April 2010. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/27/8/084036. URL <https://iopscience.iop.org/article/10.1088/0264-9381/27/8/084036>.
- [56] John M. Ford, Paul Demorest, and Scott Ransom. Heterogeneous real-time computing in radio astronomy. In Nicole M. Radziwill and Alan Bridger, editors, *Software and Cyberinfrastructure for Astronomy*, volume 7740 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77400A, July 2010. doi: 10.1117/12.857666.
- [57] R. S. Foster and D. C. Backer. Constructing a pulsar timing array. *The Astrophysical Journal*, 361:300, September 1990. ISSN 0004-637X, 1538-4357. doi: 10.1086/169195. URL <http://adsabs.harvard.edu/doi/10.1086/169195>.
- [58] Gabriel E. Freedman, Aaron D. Johnson, Rutger Van Haasteren, and Sarah J. Vigeland. Efficient gravitational wave searches with pulsar timing arrays using Hamiltonian Monte Carlo. *Phys. Rev. D*, 107(4):043013, February 2023. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.107.043013. URL <https://link.aps.org/doi/10.1103/PhysRevD.107.043013>.
- [59] Simon J Godsill. On the Relationship Between Markov chain Monte Carlo Methods for Model Uncertainty. *Journal of Computational and Graphical*

- Statistics*, 10(2):230–248, June 2001. ISSN 1061-8600, 1537-2715. doi: 10.1198/10618600152627924. URL <http://www.tandfonline.com/doi/abs/10.1198/10618600152627924>.
- [60] Boris Goncharov, D J Reardon, R M Shannon, Xing-Jiang Zhu, Eric Thrane, M Bailes, N D R Bhat, S Dai, G Hobbs, M Kerr, R N Manchester, S Osłowski, A Parthasarathy, C J Russell, R Spiewak, N Thyagarajan, and J B Wang. Identifying and mitigating noise sources in precision pulsar timing data sets. *Monthly Notices of the Royal Astronomical Society*, 502(1):478–493, January 2021. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/staa3411. URL <https://academic.oup.com/mnras/article/502/1/478/5957533>.
- [61] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759–771, April 2005. ISSN 0004-637X, 1538-4357. doi: 10.1086/427976. URL <https://iopscience.iop.org/article/10.1086/427976>.
- [62] Brynmor Haskell and Andrew Melatos. Models of pulsar glitches. *International Journal of Modern Physics D*, 24(03):1530008, March 2015. ISSN 0218-2718, 1793-6594. doi: 10.1142/S0218271815300086. URL <https://www.worldscientific.com/doi/abs/10.1142/S0218271815300086>.
- [63] J. S. Hazboun, J. Simon, S. R. Taylor, M. T. Lam, S. J. Vigeland, K. Islo, J. S. Key, Z. Arzoumanian, P. T. Baker, A. Brazier, P. R. Brook, S. Burke-Spolaor, S. Chatterjee, J. M. Cordes, N. J. Cornish, F. Crawford, K. Crowter, H. T. Cromartie, M. DeCesar, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Ferrara, E. Fonseca, N. Garver-Daniels, P. Gentile, D. Good, A. M. Holgado, E. A. Huerta, R. Jennings, G. Jones, M. L. Jones, A. R. Kaiser, D. L. Kaplan, L. Z. Kelley, T. J. W. Lazio, L. Levin, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, C. M. F. Mingarelli, C. Ng, D. J. Nice, T. T. Pennucci, N. S. Pol, S. M. Ransom, P. S. Ray, X. Siemens, R. Spiewak, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, J. E. Turner, M. Vallisneri, R. van Haasteren, C. A. Witt, W. W. Zhu, and (The NANOGrav Collaboration). The NANOGrav 11 yr Data Set: Evolution of Gravitational-wave Background Statistics. *The Astrophysical Journal*, 890(2):108, Febru-

- ary 2020. ISSN 1538-4357. doi: 10.3847/1538-4357/ab68db. URL <https://iopscience.iop.org/article/10.3847/1538-4357/ab68db>.
- [64] R. W. Hellings and G. S. Downs. Upper limits on the isotropic gravitational radiation background from pulsar timing analysis. *The Astrophysical Journal*, 265:L39, February 1983. ISSN 0004-637X, 1538-4357. doi: 10.1086/183954. URL <http://adsabs.harvard.edu/doi/10.1086/183954>.
- [65] G. B. Hobbs, R. T. Edwards, and R. N. Manchester. tempo2, a new pulsar-timing package - I. An overview: tempo2, a new pulsar-timing package - I. Overview. *Monthly Notices of the Royal Astronomical Society*, 369(2):655–672, June 2006. ISSN 00358711. doi: 10.1111/j.1365-2966.2006.10302.x. URL <http://mnras.oxfordjournals.org/cgi/doi/10.1111/j.1365-2966.2006.10302.x>.
- [66] Moritz Hübner, Paul Lasky, and Eric Thrane. Memory remains undetected: Updates from the second ligo/virgo gravitational-wave transient catalog. *Phys. Rev. D*, 104:023004, Jul 2021. doi: 10.1103/PhysRevD.104.023004. URL <https://link.aps.org/doi/10.1103/PhysRevD.104.023004>.
- [67] Kristina Islo, Joseph Simon, Sarah Burke-Spolaor, and Xavier Siemens. Prospects for Memory Detection with Low-Frequency Gravitational Wave Detectors. *arXiv:1906.11936 [astro-ph]*, June 2019. URL <http://arxiv.org/abs/1906.11936>. arXiv: 1906.11936.
- [68] Piotr Jaranowski, Andrzej Królak, and Bernard F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Physical Review D*, 58(6):063001, August 1998. ISSN 0556-2821, 1089-4918. doi: 10.1103/PhysRevD.58.063001. URL <https://link.aps.org/doi/10.1103/PhysRevD.58.063001>.
- [69] M. L. Jones, M. A. McLaughlin, M. T. Lam, J. M. Cordes, L. Levin, S. Chatterjee, Z. Arzoumanian, K. Crowter, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, M. E. Gonzalez, G. Jones, T. J. W. Lazio, D. J. Nice, T. T. Pennucci, S. M. Ransom, D. R. Stinebring, I. H. Stairs, K. Stovall, J. K. Swiggum, and W. W. Zhu. The NANOGrav Nine-year Data Set: Measurement and Analysis of Variations in Dispersion Measures. *The Astrophysical Journal*, 841(2):125, June 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/aa73df. URL <https://iopscience.iop.org/article/10.3847/1538-4357/aa73df>.

- [70] M. J. Keith, W. Coles, R. M. Shannon, G. B. Hobbs, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. J. Champion, A. Chaudhary, A. W. Hotan, J. Khoo, J. Kocz, S. Osłowski, V. Ravi, J. E. Reynolds, J. Sarkissian, W. van Straten, and D. R. B. Yardley. Measurement and correction of variations in interstellar dispersion in high-precision pulsar timing. *Monthly Notices of the Royal Astronomical Society*, 429(3):2161–2174, March 2013. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/sts486. URL <http://academic.oup.com/mnras/article/429/3/2161/1002440/Measurement-and-correction-of-variations-in>.
- [71] Daniel Kennefick. Prospects for detecting the Christodoulou memory of gravitational waves from a coalescing compact binary and using it to measure neutron-star radii. *Phys. Rev. D*, 50(6):3587–3595, September 1994. ISSN 0556-2821. doi: 10.1103/PhysRevD.50.3587. URL <https://link.aps.org/doi/10.1103/PhysRevD.50.3587>.
- [72] Nima Laal, William G. Lamb, Joseph D. Romano, Xavier Siemens, Stephen R. Taylor, and Rutger Van Haasteren. Exploring the capabilities of Gibbs sampling in pulsar timing arrays. *Phys. Rev. D*, 108(6):063008, September 2023. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.108.063008. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.063008>.
- [73] M. T. Lam, J. M. Cordes, S. Chatterjee, Z. Arzoumanian, K. Crowter, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, M. E. Gonzalez, G. Jones, M. L. Jones, L. Levin, D. R. Madison, M. A. McLaughlin, D. J. Nice, T. T. Pennucci, S. M. Ransom, R. M. Shannon, X. Siemens, I. H. Stairs, K. Stovall, J. K. Swiggum, and W. W. Zhu. THE NANOGRAV NINE-YEAR DATA SET: EXCESS NOISE IN MILLISECOND PULSAR ARRIVAL TIMES. *The Astrophysical Journal*, 834(1):35, December 2016. ISSN 1538-4357. doi: 10.3847/1538-4357/834/1/35. URL <https://iopscience.iop.org/article/10.3847/1538-4357/834/1/35>.
- [74] M. T. Lam, J. M. Cordes, S. Chatterjee, Z. Arzoumanian, K. Crowter, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, M. E. Gonzalez, G. Jones, M. L. Jones, L. Levin, D. R. Madison, M. A. McLaughlin, D. J. Nice, T. T. Pennucci, S. M. Ransom, R. M. Shannon, X. Siemens, I. H. Stairs, K. Stovall, J. K. Swiggum, and W. W. Zhu. THE NANOGRAV NINE-YEAR DATA SET: EXCESS NOISE IN MILLISECOND PULSAR AR-

- RIVAL TIMES. *The Astrophysical Journal*, 834(1):35, December 2016. ISSN 1538-4357. doi: 10.3847/1538-4357/834/1/35. URL <https://iopscience.iop.org/article/10.3847/1538-4357/834/1/35>.
- [75] Paul D. Lasky, Eric Thrane, Yuri Levin, Jonathan Blackman, and Yanbei Chen. Detecting Gravitational-Wave Memory with LIGO: Implications of GW150914. *Physical Review Letters*, 117:061102, August 2016. ISSN 0031-9007. doi: 10.1103/PhysRevLett.117.061102. URL <https://ui.adsabs.harvard.edu/abs/2016PhRvL.117f1102L>. ADS Bibcode: 2016PhRvL.117f1102L.
- [76] Lindley Lentati, P. Alexander, M. P. Hobson, S. Taylor, J. Gair, S. T. Balan, and R. van Haasteren. Hyper-efficient model-independent Bayesian method for the analysis of pulsar timing data. *Physical Review D*, 87(10):104021, May 2013. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.87.104021. URL <https://link.aps.org/doi/10.1103/PhysRevD.87.104021>.
- [77] Tom Lodewyckx, Woojae Kim, Michael D. Lee, Francis Tuerlinckx, Peter Kuppens, and Eric-Jan Wagenmakers. A tutorial on Bayes factor estimation with the product space method. *Journal of Mathematical Psychology*, 55(5):331–347, October 2011. ISSN 00222496. doi: 10.1016/j.jmp.2011.06.001. URL <https://linkinghub.elsevier.com/retrieve/pii/S0022249611000423>.
- [78] Andrea N Lommen. Pulsar timing arrays: the promise of gravitational wave detection. *Reports on Progress in Physics*, 78(12):124901, December 2015. ISSN 0034-4885, 1361-6633. doi: 10.1088/0034-4885/78/12/124901. URL <https://iopscience.iop.org/article/10.1088/0034-4885/78/12/124901>.
- [79] Duncan R. Lorimer. Binary and Millisecond Pulsars. *Living Reviews in Relativity*, 11(1):8, December 2008. ISSN 2367-3613, 1433-8351. doi: 10.12942/lrr-2008-8. URL <http://link.springer.com/10.12942/lrr-2008-8>.
- [80] Jing Luo, Scott Ransom, Paul Demorest, Rutger van Haasteren, Paul Ray, Kevin Stovall, Matteo Bachetti, Anne Archibald, Matthew Kerr, Jonathan Colen, and Fredrick Jenet. PINT: High-precision pulsar timing analysis package. *Astrophysics Source Code Library*, page ascl:1902.007, February 2019. URL <https://ui.adsabs.harvard.edu/abs/2019ascl.soft02007L>. ADS Bibcode: 2019ascl.soft02007L.

- [81] Andrew Lyne, George Hobbs, Michael Kramer, Ingrid Stairs, and Ben Stappers. Switched Magnetospheric Regulation of Pulsar Spin-Down. *Science*, 329(5990):408–412, July 2010. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1186683. URL <https://www.science.org/doi/10.1126/science.1186683>.
- [82] D. R. Madison, J. M. Cordes, and S. Chatterjee. ASSESSING PULSAR TIMING ARRAY SENSITIVITY TO GRAVITATIONAL WAVE BURSTS WITH MEMORY. *The Astrophysical Journal*, 788(2):141, June 2014. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/788/2/141. URL <https://iopscience.iop.org/article/10.1088/0004-637X/788/2/141>.
- [83] D. R. Madison, X. J. Zhu, G. Hobbs, W. Coles, R. M. Shannon, J. B. Wang, C. Tiburzi, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, J. Dempsey, M. Keith, M. Kerr, P. Lasky, Y. Levin, S. Osłowski, V. Ravi, D. Reardon, P. Rosado, R. Spiewak, W. van Straten, L. Toomey, L. Wen, and X. You. Versatile directional searches for gravitational waves with Pulsar Timing Arrays. *Monthly Notices of the Royal Astronomical Society*, 455:3662–3673, February 2016. ISSN 0035-8711. doi: 10.1093/mnras/stv2534. URL <https://ui.adsabs.harvard.edu/abs/2016MNRAS.455.3662M>. ADS Bibcode: 2016MNRAS.455.3662M.
- [84] Dustin R. Madison. Persistent Astrometric Deflections from Gravitational-Wave Memory. *Physical Review Letters*, 125(4):041101, July 2020. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.125.041101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.125.041101>.
- [85] Michele Maggiore. *Gravitational Waves: Volume 2: Astrophysics and Cosmology*. Oxford University Press Oxford, 1 edition, March 2018. ISBN 978-0-19-857089-9 978-0-19-181718-2. doi: 10.1093/oso/9780198570899.001.0001. URL <https://academic.oup.com/book/43950>.
- [86] R A Main, S A Sanidas, J Antoniadis, C Bassa, S Chen, I Cognard, M Gaikwad, H Hu, G H Janssen, R Karuppusamy, M Kramer, K J Lee, K Liu, G Mall, J W McKee, M B Mickaliger, D Perrodin, B W Stappers, C Tiburzi, O Wucknitz, L Wang, and W W Zhu. Measuring interstellar delays of PSR J0613-0200 over 7 yr, using the Large European Array for Pulsars. *Monthly Notices of the Royal Astronomical Society*, 499(1): 1468–1479, 09 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa2955. URL <https://doi.org/10.1093/mnras/staa2955>.

- [87] R A Main, J Antoniadis, S Chen, I Cognard, H Hu, J Jang, R Karuppusamy, M Kramer, K Liu, Y Liu, G Mall, J W McKee, M B Mickaliger, D Perrodin, S A Sanidas, B W Stappers, T Sprenger, O Wucknitz, C G Bassa, M Burgay, R Concu, M Gaikwad, G H Janssen, K J Lee, A Melis, M Pilia, A Possenti, L Wang, and W W Zhu. Variable scintillation arcs of millisecond pulsars observed with the large european array for pulsars. *Monthly Notices of the Royal Astronomical Society*, page stad1980, July 2023. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stad1980. URL <https://academic.oup.com/mnras/advance-article/doi/10.1093/mnras/stad1980/7220724>.
- [88] K. G. Malmquist. On some relations in stellar statistics. *Meddelanden fran Lunds Astronomiska Observatorium*, 100:1–52, March 1922.
- [89] R N Manchester. The International Pulsar Timing Array. *Classical and Quantum Gravity*, 30(22):224010, November 2013. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/30/22/224010. URL <https://iopscience.iop.org/article/10.1088/0264-9381/30/22/224010>.
- [90] R. N. Manchester, G. Hobbs, M. Bailes, W. A. Coles, W. Van Straten, M. J. Keith, R. M. Shannon, N. D. R. Bhat, A. Brown, S. G. Burke-Spolaor, D. J. Champion, A. Chaudhary, R. T. Edwards, G. Hampson, A. W. Hotan, A. Jameson, F. A. Jenet, M. J. Kesteven, J. Khoo, J. Kocz, K. Maciesiak, S. Osłowski, V. Ravi, J. R. Reynolds, J. M. Sarkissian, J. P. W. Verbiest, Z. L. Wen, W. E. Wilson, D. Yardley, W. M. Yan, and X. P. You. The Parkes Pulsar Timing Array Project. *Publications of the Astronomical Society of Australia*, 30:e017, 2013. ISSN 1323-3580, 1448-6083. doi: 10.1017/pasa.2012.017. URL https://www.cambridge.org/core/product/identifier/S1323358012000173/type/journal_article.
- [91] J. W. McKee, G. H. Janssen, B. W. Stappers, A. G. Lyne, R. N. Caballero, L. Lentati, G. Desvignes, A. Jessner, C. A. Jordan, R. Karuppusamy, M. Kramer, I. Cognard, D. J. Champion, E. Graikou, P. Lazarus, S. Osłowski, D. Perrodin, G. Shaifullah, C. Tiburzi, and J. P. W. Verbiest. A glitch in the millisecond pulsar J0613-0200. *Monthly Notices of the Royal Astronomical Society*, 461(3):2809–2817, September 2016. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw1442. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stw1442>.
- [92] M A McLaughlin. The North American Nanohertz Observatory for Gravitational Waves. *Classical and Quantum Gravity*, 30(22):224008,

November 2013. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/30/22/224008. URL <https://iopscience.iop.org/article/10.1088/0264-9381/30/22/224008>.

- [93] Lucy O. McNeill, Eric Thrane, and Paul D. Lasky. Detecting Gravitational Wave Memory without Parent Signals. *Physical Review Letters*, 118:181103, May 2017. ISSN 0031-9007. doi: 10.1103/PhysRevLett.118.181103. URL <https://ui.adsabs.harvard.edu/abs/2017PhRvL.118r1103M>. ADS Bibcode: 2017PhRvL.118r1103M.
- [94] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6): 1087–1092, June 1953. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.1699114. URL <https://pubs.aip.org/jcp/article/21/6/1087/202680/Equation-of-State-Calculations-by-Fast-Computing>.
- [95] M T Miles, R M Shannon, M Bailes, D J Reardon, S Buchner, H Middleton, and R Spiewak. Mode changing in J1909 - 3744: the most precisely timed pulsar. *Monthly Notices of the Royal Astronomical Society*, 510(4):5908–5915, January 2022. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stab3549. URL <https://academic.oup.com/mnras/article/510/4/5908/6456380>.
- [96] M T Miles, R M Shannon, M Bailes, D J Reardon, M J Keith, A D Cameron, A Parthasarathy, M Shamohammadi, R Spiewak, W van Straten, S Buchner, F Camilo, M Geyer, A Karastergiou, M Kramer, M Serylak, G Theureau, and V Venkatraman Krishnan. The MeerKAT Pulsar Timing Array: first data release. *Monthly Notices of the Royal Astronomical Society*, 519(3): 3976–3991, 12 2022. ISSN 0035-8711. doi: 10.1093/mnras/stac3644. URL <https://doi.org/10.1093/mnras/stac3644>.
- [97] D. Nice, P. Demorest, I. Stairs, R. Manchester, J. Taylor, W. Peters, J. Weisberg, A. Irwin, N. Wex, and Y. Huang. Tempo: Pulsar timing data analysis. *Astrophysics Source Code Library*, page ascl:1509.002, September 2015. URL <https://ui.adsabs.harvard.edu/abs/2015ascl.soft09002N>. ADS Bibcode: 2015ascl.soft09002N.
- [98] Ashis Paul, Krishnakumar M.A., Lankeswar Dey, Manjari Bagchi, Mayuresh Surnis, Neelam Dhanda, P. K. Manoharan, Prakash Arumugasamy, Sai Chaitanya Susarla, Siraprapa Sanpaarsa, Suryarao Bethapudi, Abhimanyu Su-

- sobhanan, Shantanu Desai, Yashwant Gupta, Yogesh Maan, Achamveedu Gopakumar, Arpita Choudhary, Avishek Basu, Arun Naidu, Bhal Chandra Joshi, Dhruv Pathak, and Kishalay De. The Indian Pulsar Timing Array (InPTA). In *2019 URSI Asia-Pacific Radio Science Conference (AP-RASC)*, pages 1–1, New Delhi, India, March 2019. IEEE. ISBN 978-90-825987-5-9. doi: 10.23919/URSIAP-RASC.2019.8738505. URL <https://ieeexplore.ieee.org/document/8738505/>.
- [99] E. S. Phinney. A Practical Theorem on Gravitational Wave Backgrounds, August 2001. URL <http://arxiv.org/abs/astro-ph/0108028>. Number: arXiv:astro-ph/0108028 arXiv:astro-ph/0108028.
- [100] M. S. Pshirkov, D. Baskaran, and K. A. Postnov. Observing gravitational wave bursts in pulsar timing measurements: Observing gravitational wave bursts. *Monthly Notices of the Royal Astronomical Society*, 402(1):417–423, February 2010. ISSN 00358711. doi: 10.1111/j.1365-2966.2009.15887.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2009.15887.x>.
- [101] Daniel J. Reardon, Andrew Zic, Ryan M. Shannon, George B. Hobbs, Matthew Bailes, Valentina Di Marco, Agastya Kapur, Axl F. Rogers, Eric Thrane, Jacob Askew, N. D. Ramesh Bhat, Andrew Cameron, Małgorzata Curyło, William A. Coles, Shi Dai, Boris Goncharov, Matthew Kerr, Atharva Kulkarni, Yuri Levin, Marcus E. Lower, Richard N. Manchester, Rami Mandow, Matthew T. Miles, Rowina S. Nathan, Stefan Osłowski, Christopher J. Russell, Renée Spiewak, Songbo Zhang, and Xing-Jiang Zhu. Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array. *The Astrophysical Journal Letters*, 951(1):L6, July 2023. ISSN 2041-8205, 2041-8213. doi: 10.3847/2041-8213/acdd02. URL <https://iopscience.iop.org/article/10.3847/2041-8213/acdd02>.
- [102] P. A. Rosado, A. Sesana, and J. Gair. Expected properties of the first gravitational wave signal detected with pulsar timing arrays. *Monthly Notices of the Royal Astronomical Society*, 451(3):2417–2433, June 2015. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stv1098. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stv1098>.
- [103] M. V. Sazhin. Opportunities for detecting ultralong gravitational waves. *Soviet Astronomy*, 22:36–38, February 1978. ISSN 0038-5301. URL

<https://ui.adsabs.harvard.edu/abs/1978SvA....22...36S>. ADS Bibcode: 1978SvA....22...36S.

- [104] Naoki Seto. Search for Memory and Inspiral Gravitational Waves from Super-Massive Binary Black Holes with Pulsar Timing Arrays. *Monthly Notices of the Royal Astronomical Society: Letters*, 400(1):L38–L42, November 2009. ISSN 17453925, 17453933. doi: 10.1111/j.1745-3933.2009.00758.x. URL <http://arxiv.org/abs/0909.1379>. arXiv: 0909.1379.
- [105] Ryan M. Shannon and James M. Cordes. ASSESSING THE ROLE OF SPIN NOISE IN THE PRECISION TIMING OF MILLISECOND PULSARS. *The Astrophysical Journal*, 725(2):1607–1619, December 2010. ISSN 0004-637X, 1538-4357. doi: 10.1088/0004-637X/725/2/1607. URL <https://iopscience.iop.org/article/10.1088/0004-637X/725/2/1607>.
- [106] Xavier Siemens, Jolien Creighton, Irit Maor, Saikat Ray Majumder, Kipp Cannon, and Jocelyn Read. Gravitational wave bursts from cosmic (super)strings: Quantitative analysis and constraints. *Physical Review D*, 73(10):105001, May 2006. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.73.105001. URL <https://link.aps.org/doi/10.1103/PhysRevD.73.105001>.
- [107] Xavier Siemens, Justin Ellis, Fredrick Jenet, and Joseph D Romano. The stochastic background: scaling laws and time to detection for pulsar timing arrays. *Class. Quantum Grav.*, 30(22):224015, November 2013. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/30/22/224015. URL <https://iopscience.iop.org/article/10.1088/0264-9381/30/22/224015>.
- [108] J. Simon et al. *in preparation*, 2023.
- [109] Jerry Sun. `bwm_fpfe`, 2024. URL https://github.com/jerry-sun1/bwm_fpfe.
- [110] Jerry Sun, Paul T. Baker, Aaron D. Johnson, Dustin R. Madison, and Xavier Siemens. Implementation of an Efficient Bayesian Search for Gravitational-wave Bursts with Memory in Pulsar Timing Array Data. *The Astrophysical Journal*, 951(2):121, July 2023. ISSN 0004-637X, 1538-4357. doi: 10.3847/1538-4357/acd2cc. URL <https://iopscience.iop.org/article/10.3847/1538-4357/acd2cc>.

- [111] Jerry P. Sun, Xavier Siemens, and Dustin R. Madison. Optimal Strategies for Gravitational Wave Memory Detection in Pulsar Timing Arrays. 2024. doi: 10.48550/ARXIV.2410.19977. URL <https://arxiv.org/abs/2410.19977>. Version Number: 1.
- [112] Stephen R. Taylor, Paul T. Baker, Jeffrey S. Hazboun, Joseph Simon, and Sarah J. Vigeland. `enterprise_extensions`, 2021. URL https://github.com/nanograv/enterprise_extensions. v3.4.2.
- [113] Stephen R. Taylor, Paul T. Baker, Jeffrey S. Hazboun, Joseph Simon, and Sarah J. Vigeland. `enterprise_extensions`, 2021. URL https://github.com/nanograv/enterprise_extensions. v2.3.3.
- [114] The NANOGrav Collaboration, Zaven Arzoumanian, Adam Brazier, Sarah Burke-Spolaor, Sydney Chamberlin, Shami Chatterjee, Brian Christy, James M. Cordes, Neil Cornish, Kathryn Crowter, Paul B. Demorest, Timothy Dolch, Justin A. Ellis, Robert D. Ferdman, Emmanuel Fonseca, Nathan Garver-Daniels, Marjorie E. Gonzalez, Fredrick A. Jenet, Glenn Jones, Megan L. Jones, Victoria M. Kaspi, Michael Koop, Michael T. Lam, T. Joseph W. Lazio, Lina Levin, Andrea N. Lommen, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin Madison, Maura A. McLaughlin, Sean T. McWilliams, David J. Nice, Nipuni Palliyaguru, Timothy T. Pennucci, Scott M. Ransom, Xavier Siemens, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Joseph K. Swiggum, Michele Vallisneri, Rutger van Haasteren, Yan Wang, and Weiwei Zhu. THE NANOGrav NINE-YEAR DATA SET: OBSERVATIONS, ARRIVAL TIME MEASUREMENTS, AND ANALYSIS OF 37 MILLISECOND PULSARS. *The Astrophysical Journal*, 813(1):65, October 2015. ISSN 1538-4357. doi: 10.1088/0004-637X/813/1/65. URL <https://iopscience.iop.org/article/10.1088/0004-637X/813/1/65>.
- [115] Kip S. Thorne. Gravitational-wave bursts with memory: The Christodoulou effect. *Physical Review D*, 45(2):520–524, January 1992. ISSN 0556-2821. doi: 10.1103/PhysRevD.45.520. URL <https://link.aps.org/doi/10.1103/PhysRevD.45.520>.
- [116] Youli Tuo, Muhammed Miraç Serim, Marco Antonelli, Lorenzo Ducci, Armin Vahdat, Mingyu Ge, Andrea Santangelo, and Fei Xie. Discovery of the First Antiglitch Event in the Rotation-powered Pulsar PSR B0540-69. *ApJL*, 967(1):L13, May 2024. ISSN 2041-8205, 2041-8213. doi:

- 10.3847/2041-8213/ad4488. URL <https://iopscience.iop.org/article/10.3847/2041-8213/ad4488>.
- [117] Michele Vallisneri, Marco Crisostomi, Aaron D. Johnson, and Patrick M. Meyers. Rapid parameter estimation for pulsar-timing-array datasets with variational inference and normalizing flows, May 2024. URL <http://arxiv.org/abs/2405.08857>. arXiv:2405.08857 [gr-qc].
- [118] Rutger van Haasteren and Yuri Levin. Gravitational-wave memory and pulsar timing arrays. *Monthly Notices of the Royal Astronomical Society*, 401(4):2372–2378, February 2010. ISSN 00358711, 13652966. doi: 10.1111/j.1365-2966.2009.15885.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2009.15885.x>.
- [119] Rutger van Haasteren and Michele Vallisneri. New advances in the Gaussian-process approach to pulsar-timing data analysis. *Physical Review D*, 90(10):104012, November 2014. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.90.104012. URL <https://link.aps.org/doi/10.1103/PhysRevD.90.104012>.
- [120] Rutger van Haasteren and Michele Vallisneri. Low-rank approximations for large stationary covariance matrices, as used in the Bayesian and generalized-least-squares analysis of pulsar-timing data. *Monthly Notices of the Royal Astronomical Society*, 446(2):1170–1174, January 2015. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu2157. URL <http://academic.oup.com/mnras/article/446/2/1170/2892265/Lowrank-approximations-for-large-stationary>.
- [121] J. P. W. Verbiest, L. Lentati, G. Hobbs, R. Van Haasteren, P. B. Demorest, G. H. Janssen, J.-B. Wang, G. Desvignes, R. N. Caballero, M. J. Keith, D. J. Champion, Z. Arzoumanian, S. Babak, C. G. Bassa, N. D. R. Bhat, A. Brazier, P. Brem, M. Burgay, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, I. Cognard, J. M. Cordes, S. Dai, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, J. R. Gair, N. E. Garver-Daniels, P. Gentile, M. E. Gonzalez, E. Graikou, L. Guillemot, J. W. T. Hessels, G. Jones, R. Karuppusamy, M. Kerr, M. Kramer, M. T. Lam, P. D. Lasky, A. Lassus, P. Lazarus, T. J. W. Lazio, K. J. Lee, L. Levin, K. Liu, R. S. Lynch, A. G. Lyne, J. McKee, M. A. McLaughlin, S. T. McWilliams, D. R. Madison, R. N. Manchester, C. M. F. Mingarelli, D. J. Nice, S. Osłowski, N. T. Palliyaguru, T. T.

- Pennucci, B. B. P. Perera, D. Perrodin, A. Possenti, A. Petiteau, S. M. Ransom, D. Reardon, P. A. Rosado, S. A. Sanidas, A. Sesana, G. Shaifullah, R. M. Shannon, X. Siemens, J. Simon, R. Smits, R. Spiewak, I. H. Stairs, B. W. Stappers, D. R. Stinebring, K. Stovall, J. K. Swiggum, S. R. Taylor, G. Theureau, C. Tiburzi, L. Toomey, M. Vallisneri, W. Van Straten, A. Vecchio, Y. Wang, L. Wen, X. P. You, W. W. Zhu, and X.-J. Zhu. The International Pulsar Timing Array: First data release. *Monthly Notices of the Royal Astronomical Society*, 458(2):1267–1288, May 2016. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw347. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stw347>.
- [122] Sarah J. Vigeland, Kristina Islo, Stephen R. Taylor, and Justin A. Ellis. Noise-marginalized optimal statistic: A robust hybrid frequentist-Bayesian statistic for the stochastic gravitational-wave background in pulsar timing arrays. *Physical Review D*, 98(4):044003, August 2018. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.98.044003. URL <https://link.aps.org/doi/10.1103/PhysRevD.98.044003>.
- [123] J. B. Wang, G. Hobbs, W. Coles, R. M. Shannon, X. J. Zhu, D. R. Madison, M. Kerr, V. Ravi, M. J. Keith, R. N. Manchester, Y. Levin, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, S. Osłowski, W. van Straten, L. Toomey, N. Wang, and L. Wen. Searching for gravitational wave memory bursts with the Parkes Pulsar Timing Array. *Monthly Notices of the Royal Astronomical Society*, 446(2):1657–1671, January 2015. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu2137. URL <http://academic.oup.com/mnras/article/446/2/1657/2892038/Searching-for-gravitational-wave-memory-bursts>.
- [124] Alan G. Wiseman and Clifford M. Will. Christodoulou’s nonlinear gravitational-wave memory: Evaluation in the quadrupole approximation. *Physical Review D*, 44:R2945–R2949, November 1991. ISSN 1550-79980556-2821. doi: 10.1103/PhysRevD.44.R2945. URL <https://ui.adsabs.harvard.edu/abs/1991PhRvD..44.2945W>. ADS Bibcode: 1991PhRvD..44.2945W.
- [125] Max A. Woodbury. *Inverting modified matrices*. Number 42 in SRG Memorandum report. Department of Statistics, Princeton University, Princeton, NJ, 1950.

- [126] Yu-Mei Wu, Zu-Cheng Chen, and Qing-Guo Huang. Constraining the Polarization of Gravitational Waves with the Parkes Pulsar Timing Array Second Data Release. *ApJ*, 925(1):37, January 2022. ISSN 0004-637X, 1538-4357. doi: 10.3847/1538-4357/ac35cc. URL <https://iopscience.iop.org/article/10.3847/1538-4357/ac35cc>.
- [127] Heng Xu, Siyuan Chen, Yanjun Guo, Jinchen Jiang, Bojun Wang, Jiangwei Xu, Zihan Xue, R. Nicolas Caballero, Jianping Yuan, Yonghua Xu, Jingbo Wang, Longfei Hao, Jingtao Luo, Kejia Lee, Jinlin Han, Peng Jiang, Zhiqiang Shen, Min Wang, Na Wang, Renxin Xu, Xiangping Wu, Richard Manchester, Lei Qian, Xin Guan, Menglin Huang, Chun Sun, and Yan Zhu. Searching for the nano-hertz stochastic gravitational wave background with the chinese pulsar timing array data release i. *Research in Astronomy and Astrophysics*, 23(7):075024, jun 2023. doi: 10.1088/1674-4527/acdfa5. URL <https://dx.doi.org/10.1088/1674-4527/acdfa5>.
- [128] Heng Xu, Siyuan Chen, Yanjun Guo, Jinchen Jiang, Bojun Wang, Jiangwei Xu, Zihan Xue, R. Nicolas Caballero, Jianping Yuan, Yonghua Xu, Jingbo Wang, Longfei Hao, Jingtao Luo, Kejia Lee, Jinlin Han, Peng Jiang, Zhiqiang Shen, Min Wang, Na Wang, Renxin Xu, Xiangping Wu, Richard Manchester, Lei Qian, Xin Guan, Menglin Huang, Chun Sun, and Yan Zhu. Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I. *Research in Astronomy and Astrophysics*, 23(7):075024, July 2023. ISSN 1674-4527. doi: 10.1088/1674-4527/acdfa5. URL <https://iopscience.iop.org/article/10.1088/1674-4527/acdfa5>.
- [129] N Yonemaru, S Kuroyanagi, G Hobbs, K Takahashi, X-J Zhu, W A Coles, S Dai, E Howard, R Manchester, D Reardon, C Russell, R M Shannon, N Thyagarajan, R Spiewak, and J-B Wang. Searching for gravitational-wave bursts from cosmic string cusps with the Parkes Pulsar Timing Array. *Monthly Notices of the Royal Astronomical Society*, 501(1):701–712, December 2020. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/staa3721. URL <https://academic.oup.com/mnras/article/501/1/701/6017303>.
- [130] Ya. B. Zel’dovich and A. G. Polnarev. Radiation of gravitational waves by a cluster of superdense stars. *Soviet Astronomy*, 18:17, August 1974. ISSN 0038-5301. URL <https://ui.adsabs.harvard.edu/abs/1974SvA....18..17Z>. ADS Bibcode: 1974SvA....18...17Z.

- [131] X.-J. Zhu, G. Hobbs, L. Wen, W. A. Coles, J.-B. Wang, R. M. Shannon, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. J. Keith, M. Kerr, Y. Levin, D. R. Madison, S. Osłowski, V. Ravi, L. Toomey, and W. Van Straten. An all-sky search for continuous gravitational waves in the Parkes Pulsar Timing Array data set. *Monthly Notices of the Royal Astronomical Society*, 444(4):3709–3720, November 2014. ISSN 1365-2966, 0035-8711. doi: 10.1093/mnras/stu1717. URL <http://academic.oup.com/mnras/article/444/4/3709/1029897/An-allsky-search-for-continuous-gravitational>.
- [132] Andrea Zonca, Leo Singer, Daniel Lenz, Martin Reinecke, Cyrille Rosset, Eric Hivon, and Krzysztof Gorski. healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python. *Journal of Open Source Software*, 4(35):1298, March 2019. ISSN 2475-9066. doi: 10.21105/joss.01298. URL <http://joss.theoj.org/papers/10.21105/joss.01298>.

APPENDICES

