THE NANOGRAV 11-YEAR DATA SET: LIMITS ON GRAVITATIONAL WAVE MEMORY

K. Aggarwal[1,2], Z. Arzoumanian[3], P. T. Baker[§4,2], A. Brazier[5,6], P. R. Brook[1,2], S. Burke-Spolaor[1,2], S. Chatterjee[5], J. M. Cordes[5], N. J. Cornish[7], F. Crawford[8], H. T. Cromartie[9], K. Crowter[10], M. DeCesar[†11], P. B. Demorest[12], T. Dolch[13], J. A. Ellis[†14,1,2], R. D. Ferdman[15], E. C. Ferrara[16], E. Fonseca[17], N. Garver-Daniels[1,2], P. Gentile[1,2], D. Good[10], J. S. Hazboun[†18], A. M. Holgado[19], E. A. Huerta[19], K. Islo[20], R. Jennings[5], G. Jones[21], M. L. Jones[20], D. L. Kaplan[20], L. Z. Kelley[22], J. S. Key[18], M. T. Lam[23,24,1,2], T. J. W. Lazio[25,26], L. Levin[27], D. R. Lorimer[1,2], J. Luo[17], R. S. Lynch[28], D. R. Madison[†1,2], M. A. McLaughlin[1,2], S. T. McWilliams[1,2], C. M. F. Mingarelli[29,30], C. Ng[31], D. J. Nice[11], T. T. Pennucci[†32], N. S. Pol[1,2], S. M. Ransom[9,33], P. S. Ray[34], X. Siemens[35,20], J. Simon[25,26], R. Spiewak[20,36], I. H. Stairs[10], D. R. Stinebring[37], K. Stovall[†12], J. K. Swiggum[†20], S. R. Taylor[38,25,26], M. Vallisneri[25,26], R. van Haasteren[‡26], S. J. Vigeland[20], C. A. Witt[1,2], W. W. Zhu[39]

(The NANOGrav Collaboration)[⋆]

⋆Author order alphabetical by surname

[1]Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA
[2]Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA
[3]X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA
[4]Department of Physics and Astronomy, Widener University, One University Place, Chester, PA 19013, USA
[5]Department of Astronomy, Cornell University, Ithaca, NY 14853, USA
[6]Cornell Center for Advanced Computing, Ithaca, NY 14853, USA
[7]Department of Physics, Montana State University, Bozeman, MT 59717, USA
[8]Department of Physics and Astronomy, Franklin & Marshall College, P.O. Box 3003, Lancaster, PA 17604, USA
[9]University of Virginia, Department of Astronomy, P.O. Box 400325, Charlottesville, VA 22904, USA
[10]Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada
[11]Department of Physics, Lafayette College, Easton, PA 18042, USA
[12]National Radio Astronomy Observatory, 1003 Lopezville Rd., Socorro, NM 87801, USA
[13]Department of Physics, Hillsdale College, 33 E. College Street, Hillsdale, Michigan 49242, USA
[14]Infinia ML, 202 Rigsbee Avenue, Durham, NC 27701, USA
[15]Department of Physics, University of East Anglia, Norwich, UK
[16]NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[17]Department of Physics, McGill University, 3600 University St., Montreal, QC H3A 2T8, Canada
[18]University of Washington Bothell, 18115 Campus Way NE, Bothell, WA 98011, USA
[19]NCSA and Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA
[20]Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA
[21]Department of Physics, Columbia University, New York, NY 10027, USA
[22]Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208
[23]School of Physics and Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA
[24]Laboratory for Multiwavelength Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA
[25]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
[26]Theoretical AstroPhysics Including Relativity (TAPIR), MC 350-17, California Institute of Technology, Pasadena, California 91125, USA
[27]Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, M13 9PL, United Kingdom
[28]Green Bank Observatory, P.O. Box 2, Green Bank, WV 24944, USA
[29]Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA
[30]Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA
[31]Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George St., Toronto, ON M5S 3H4, Canada
[32]Hungarian Academy of Sciences MTA-ELTE "Extragalatic Astrophysics Research Group", Institute of Physics, Eötvös Loránd University, Pázmány P. s. 1/A, 1117 Budapest, Hungary
[33]National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA
[34]Naval Research Laboratory, Washington DC 20375, USA
[35]Department of Physics, Oregon State University, Corvallis, OR 97331, USA
[36]Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia
[37]Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA
[38]Department of Physics and Astronomy, Vanderbilt University, 2301 Vanderbilt Place, Nashville, TN 37235, USA
[39]CAS Key Laboratory of FAST, Chinese Academy of Science, Beijing 100101, China
† NANOGrav Physics Frontiers Center Postdoctoral Fellow and
‡ currently employed at Microsoft Corporation
Draft version December 9, 2019

ABSTRACT

The mergers of supermassive black hole binaries (SMBHB) promise to be incredible sources of gravitational waves (GW). While the oscillatory part of the merger gravitational waveform will be outside the frequency sensitivity range of pulsar timing arrays (PTA), the non-oscillatory GW memory effect is detectable. Further, any burst of gravitational waves will produce GW memory, making memory a useful probe of unmodeled exotic sources and new physics. We searched the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) 11-year data set for GW memory. This dataset is sensitive to very low frequency GWs of

$\sim$ 3 to 400 nHz (periods of $\sim$ 11 yr $-$ 1 mon). Finding no evidence for GWs, we placed limits on the strain amplitude of GW memory events during the observation period. We then used the strain upper limits to place limits on the rate of GW memory causing events. At a strain of 2.5×10^{-14}, corresponding to the median upper limit as a function of source sky position, we set a limit on the rate of GW memory events at < 0.4 yr^{-1}. That strain corresponds to a SMBHB merger with reduced mass of $\eta M \sim 2 \times 10^{10} M_\odot$ and inclination of $\iota = \pi/3$ at a distance of 1 Gpc.

As a test of our analysis, we analyzed the NANOGrav 9-year data set as well. This analysis found an anomolous signal, which does not appear in the 11-year data set. This signal is not a GW, and its origin remains unknown.

Keywords: Gravitational waves – Methods: data analysis – Pulsars: general

1. INTRODUCTION

Non-oscillatory gravitational wave (GW) effects have been known since the 1970s (Zel'dovich & Polnarev 1974; Braginsky & Grishchuk 1985; Braginsky & Thorne 1987). For cases like supernovae explosions (Burrows & Hayes 1996) or hyperbolic passages of massive bodies (Turner & Will 1978), the non-oscillatory motion of the sources on unbound trajectories is encoded in the GWs as linear GW memory. For systems with purely oscillatory source motion, like binary black hole (BBH) inspirals, the GWs themselves follow unbound trajectories, generating non-linear GW memory (Christodoulou 1991; Wiseman & Will 1991; Blanchet & Damour 1992; Thorne 1992).

GW memory is a permanent change to the spacetime metric, contributing a DC component to the GW waveform. Non-linear memory builds throughout the whole history of a system's evolution, with the largest accumulation of memory occurring during periods of maximal GW emission. For compact binary sources, a burst of GW memory is produced during the highly relativistic merger. While non-linear memory is sourced at the 2.5 post-Newtonian order, it enters the GW amplitude evolution at the leading, Newtonian order (Arun et al. 2004), making the memory effect tantalizingly detectable. While the GW memory accumulated during a binary inspiral can be calculated using the post-Newtonian formalism (Wiseman & Will 1991; Arun et al. 2004), it was not until Favata (2009a,b, 2011) estimated the memory effect all the way through binary merger using an effective-one-body approach that interest in the subject was revitalized.

Pulsars act as highly stable galactic clocks (Lorimer 2008, and references therein). Their stability allows one to detect small changes in the arrival times of pulses caused by the passage of GWs between the pulsar and observer (Sazhin 1978; Detweiler 1979). Long term pulsar timing campaigns provide sensitivity to low frequency gravitational waves with periods of months to years. A pulsar timing array (PTA) combines observational data from multiple pulsars boosting sensitivity to common effects like GWs (Foster & Backer 1990; Lommen 2015). Supermassive black hole binaries (SMBHB) are the most promising sources of GWs for PTAs. Inspiraling SMBHB could be detected as individual, resolvable sources or as a stochastic background of many overlapping sources (Haehnelt 1994; Jaffe & Backer 2003; Sesana et al. 2009).

While the inspiral phase of a SMBHB emits GWs detectable by PTAs, the final merger phase emits GWs that are too high frequency for PTAs to detect (order days^{-1} for $\sim 10^9 M_\odot$ systems). Despite this, the non-linear GW memory associated with the merger could potentially be resolved independently of the oscillatory component (Seto 2009; van Haasteren & Levin 2010; Pshirkov et al. 2010; Cordes & Jenet 2012; Madison et al. 2014). Further, every GW producing system will produce non-linear memory, so searches for GW memory could uncover exotic sources of GWs or even new physics. Past studies by the North American Nanohertz Observatory for Gravitational Waves (NANOGrav, McLaughlin 2013) and the Parkes Pulsar Timing Array (PPTA, Hobbs 2013) have searched for and placed limits on GW memory in Arzoumanian et al. (2015, hereafter NG5mem) and Wang et al. (2015). Additionally, Madison et al. (2016) used PPTA data to constrain GW memory signals originating in five nearby galaxy clusters.

As a DC effect, detection prospects for memory have been considered for experiments spanning the GW spectrum. Favata (2009a,b, 2011) discussed detection prospects for both LIGO (LIGO Scientific Collaboration et al. 2015) and LISA (Amaro-Seoane et al. 2017). As mentioned above, the PTA community quickly realized the memory detection potential of their low-frequency GW experiments. More recently, Lasky et al. (2016), McNeill et al. (2017), and Talbot et al. (2018) have considered an approach to detecting GW memory with LIGO by stacking data from several BBH detections. Madison et al. (2017) have considered the detection prospects of GW memory originating in globular clusters. Cutler et al. (2014) point out that GW memory accompanying bursts of GWs at high redshift could be detectible. The ubiquitous nature of GW memory production makes it an excellent discovery tool capable of probing new and exotic physics.

NANOGrav recently published its 11-year data release (Arzoumanian et al. 2018a, hereafter NG11). Using this data set the NANOGrav collaboration has placed limits on a stochastic background of GWs (Arzoumanian et al. 2018b, hereafter NG11gwb) and on continuous GWs from individual inspiraling SMBH binaries (Aggarwal et al. 2019, hereafter NG11cw). In this work, we search the NANOGrav 9-year (Arzoumanian et al. 2015, hereafter NG9) and 11-year (NG11) data sets for GW memory. Our primary reported results use NG11.

2. THE NANOGRAV 9-YEAR AND 11-YEAR DATA SETS

For this analysis we used the NANOGrav 9-year (NG9) and 11-year data sets (NG11). The NANOGrav 9-year data set contains the times of arrival (TOAs) for 37 pulsars observed between 2004 and 2013. The NANOGrav 11-year data set extends the 9-year data set, including TOAs for 45 pulsars with observations extending to 2015. Several pulsars have been added to the array since regular observations began. Despite their names, not all pulsars in the 9-year and 11-year datasets have observations covering the whole timespan. For our analysis of the 11-year data set we used only the 34 pulsars with a minimum of 3 years of observations. For more details on NANGrav observations and data reduction see NG9

§ Corresponding author email: paul.baker@nanograv.org

and NG11.

NANOGrav observations use two radio telescopes: the 100-m Robert C. Byrd Green Bank Telescope (GBT) of the Green Bank Observatory in Green Bank, West Virginia; the 305-m William E. Gordon Telescope (Arecibo) of Arecibo Observatory in Arecibo, Puerto Rico. Prioritizing Arecibo's better sensitivity, all pulsars visible to Arecibo ($0° < δ < 39°$) were observed with it. Those outside Arecibo's declination range were observed with GBT. Two pulsars, PSRs J1713+0747 and B1937+21, were observed with both. We observed most pulsars once a month. In 2013 we began a high-cadence observing campaign aimed to increase our sensitivity to individual SMBH binary sources (Burt et al. 2011; Christy et al. 2014). Seven pulsars were observed weekly: PSRs J1713+0747 and J1909−3744 with GBT; PSRs J0030+0451, J1640+2224, J1713+0747, J2043+1711, and J2317+1439 with Arecibo.

In order to measure pulse dispersion due to the interstellar medium (ISM) we observed each pulsar at multiple radio frequencies. At GBT each pulsar was observed with both the 820 MHz and 1.4 GHz receivers. These two observations were typically separated by a few days due to mechanical and scheduling constraints. At Arecibo each pulsar was observed with the 1.4 GHz receiver and one of 430 MHz or 2.3 GHz receiver depending on the properties of the individual pulsars. The Arecibo observations are made back-to-back with the second frequency observation beginning minutes after the first completes. The telescopes' backend instrumentation systems were upgraded between 2010 and 2012. Earlier data were recorded using the 64 MHz bandwidth ASP (Arecibo) and GASP (GBT) systems. Newer data were recorded with the wideband PUPPI (Arecibo) and GUPPI (GBT) systems. During the transition data were simultaneously recorded with both systems for verification; however, only the data from the newer system is included in the release.

We fit a timing model to each pulsar's observed TOAs using `tempo`[1] and `tempo2`[2] (Hobbs et al. 2006; Edwards et al. 2006). The timing models for all pulsars include the spin period, spin period derivative, sky location, distance, and proper motion. Pulsars in binaries have additional Keplerian and post-Keplerian parameters describing the binary motion. The timing models also account for dispersion measure variations, as a piece-wise offset from the mean for each observing epoch, as first discussed in Demorest et al. (2013, hereafter NG5).

3. DATA ANALYSIS METHODS

We present the first fully Bayesian search for GW memory with a PTA. This work represents a leap forward in analysis sophistication over NANOGrav's previous search for GW memory (NG5mem).

3.1. *Model Overview*

We modeled the residual pulse time of arrival, $δ\mathbf{t}$, for a particular pulsar as the sum of stochastic and deterministic components

$$δ\mathbf{t} = \mathbf{s} + T\mathbf{b} + \mathbf{n}. \tag{1}$$

In this framework $\mathbf{s}$ represents deterministic effects such as those from gravitational waves or a solar system ephemeris (SSE) model. $T\mathbf{b}$ are stochastic processes described by a

Gaussian process model: T is the design matrix of basis functions for the models, and $\mathbf{b}$ are the model coefficients. This Gaussian process model was used for low-frequency (red), intrinsic pulsar noise and to account for uncertainty in the pulsar timing model. White noise sources are given by $\mathbf{n}$, including template fitting uncertainty and radio frequency correlated pulse jitter noise (see Cordes & Shannon 2010, and references therein).

3.2. *Gravitational Wave Model*

Non-linear GW memory is believed to accompany the oscillatory GWs produced by compact binaries. Favata (2009a) computed the total accumulated memory during BBH inspiral and merger, finding[3]

$$h_{\mathrm{mem}} = \frac{1}{24} \frac{ηM}{R} \sin^2 ι \left(17 + \cos^2 ι\right) \left[\frac{ΔE_{\mathrm{rad}}}{ηM}\right], \tag{2}$$

$$\frac{ΔE_{\mathrm{rad}}}{ηM} \sim 1 - \sqrt{8}/3 \sim 0.06, \tag{3}$$

where $M = m_1 + m_2$ is the binary total mass; $η = m_1 m_2/M^2$ is the reduced mass ratio; R is the co-moving distance to the source; $ι$ is the binary inclination; and $ΔE_{\mathrm{rad}}$ is the radiated energy, which is approximated in Equation 3 following Lousto et al. (2010). Assuming a modest inclination $ι = π/3$,

$$h_{\mathrm{mem}} \sim 1.5 \times 10^{-15} \left(\frac{ηM}{10^9 M_\odot}\right) \left(\frac{\mathrm{Gpc}}{R}\right). \tag{4}$$

For a non-precessing source, the GW memory signal is purely linearly polarized. Following the usual coordinate conventions for BBH inspiral and merger, the memory is "+" polarized in the source frame. The observed polarization angle will depend on the specific source-detector geometry.

We implemented the same memory model used by van Haasteren & Levin (2010), Pshirkov et al. (2010), Madison et al. (2014), Wang et al. (2015), and NG5mem. This model treats the GW memory as a step function that turns on (and off) as the wavefront passes by the Earth at time t_0 (and the pulsar at time t_i). This step in the spacetime metric causes a change in the distance between the pulsar and the Earth. Each radio pulse will arrive progressively more late (or early depending on the sign of the step) compared to the expected TOA. The response to GW memory in the timing residuals will therefore be a linear increase (or decrease).

For the amplitude of the step we use h_{mem} directly, bypassing the source specific amplitude dependencies (M, $η$, R, $ι$). This model ignores the details of memory accumulation. We assume that the bulk of the GW memory arises from a transient burst of GWs at a timescale shorter than our ~monthly pulsar observations. While PTA searches for GW memory were originally motivated by BBH systems, this generic model is agnostic to the source.

We choose to write the GW memory model in a form slightly different than previous work to better illuminate our search parameters. The GW's effect on the time of arrival of pulses from the ith pulsar is given by

$$s_i(t) = h_{\mathrm{mem}} B_i(\hat{k}, ψ; \hat{n}_i) \times$$
$$[(t - t_0) Θ(t - t_0) - (t - t_i) Θ(t - t_i)], \tag{5}$$

where h_{mem} is the GW strain of the memory; B_i is the angular response of the pulsar, which depends on its sky position $\hat{n}_i$

and the sky position $\hat{k}(\theta, \phi)$ and polarization (ψ) of the source (Estabrook & Wahlquist 1975); t_0 is the time that the GW wavefront passes the Earth, t_i is the retarded time for the GW wavefront passing the pulsar (GW passage time corrected for signal travel time from pulsar to Earth), and Θ is the Heaviside step function. This model assumes a plane-fronted GW with $R \gg \ell_i$, i.e., the distance to the source is much greater than the distance to any pulsar. The sky positions of the pulsars are very well constrained by the timing model, and we took these to be known exactly.

The final factor of Equation 5 contains the so called *Earth term* and *pulsar term*, describing the state of the GWs at each location. With the typical pulsar distance being ∼kpc the Earth and pulsar terms will not both fall within our 11-year observation window unless the propagation direction of the GW is nearly perpendicular to the Earth-pulsar separation. The source positions which result in both the Earth and pulsar term being active in the observation time represent a small but non-negligible fraction of the sky (Pitkin 2012). Further, a search including both the Earth and pulsar term would potentially add a small amount of signal-to-noise for these sources (Cordes & Jenet 2012; NG5mem). However, including the pulsar term would greatly complicate the analysis owing to the poorly constrained Earth-pulsar distances. For this reason, we did not do a simultaneous search for both the Earth and pulsar terms. Instead, we undertook two separate analyses:

- a search for the Earth term only

$$s_i(t) = h_{\mathrm{mem}} B_i(\hat{k}, \psi; \hat{n}_i)(t - t_0) \Theta(t - t_0). \qquad (6)$$

  This search combined data from all pulsars and had five free parameters in the GW model ($h_{\mathrm{mem}}, t_0, \theta, \phi, \psi$).

- a search for the pulsar term only

$$s_i(t) = s\, h_{\mathrm{opt}}(t - t_i) \Theta(t - t_i), \qquad (7)$$

  where h_{opt} is the strain amplitude assuming an optimally oriented source and s is the sign ($+/-$) of the memory effect. The pulsar term search was conducted on each pulsar individually and had three free parameters in the GW model (h_{opt}, t_i, s). The signal model was modified for the pulsar term search, because the angular response B is completely covariant with the GW amplitude h_{mem}. When using a single pulsar, the extrinsic parameters of the source ($\hat{k}, \psi$) are not constrained.

It is nearly impossible to make a confident detection with a single pulsar term search, because the GW memory signal is nearly indistinguishable from an intrinsic pulsar glitch (van Haasteren & Levin 2010). Even though it would be hard to trust a GW detection from a single pulsar search, we can still use non-detection in the pulsar term search to set upper limits. The individual pulsar term searches cover many independent time periods, making pulsar term limits especially useful for constraining the rate of GW memory producing events.

There are a small fraction of source locations that would result in the GW passing multiple pulsars but not the Earth during our observation. Analyzing multiple pulsar terms without the Earth term to cover these cases, could result in a small boost to signal-to-noise. As in the Earth term search, the poorly constrained Earth-pulsar distances would greatly complicate any multiple pulsar term search. For this reason, we performed only single pulsar term searches. We did not

search over cases with multiple simultaneous pulsar terms for simplicity.

3.3. *White Noise Model*

We used the pulsar noise model described in Lam et al. (2016) and also used in NG11 and NG11gwb. The white noise is parametrized in each pulsar per each observing system, k (each unique combination of telescope frontend and backend hardware, e.g., L-band GUPPI). This noise $\mathbf{n}$ is defined by a covariance matrix with each TOA specified by its observation time t and radio frequency ν

$$N_{\nu\nu' tt' k} = \delta_{tt'} \left[\delta_{\nu\nu'} \left(\mathcal{F}_k{}^2 \sigma^2 + \mathcal{Q}_k{}^2 \right) + \mathcal{J}_k{}^2 \right], \qquad (8)$$

where σ is the pulse template fitting uncertainty; $\mathcal{F}_k$ is 'EFAC', an additional scaling factor (in practice ∼1 for all); $\mathcal{Q}_k$ is 'EQUAD', an additional variance added in quadrature, 'EQUAD'; $\mathcal{J}_k$ is 'ECORR', a component that is correlated between different radio frequency channels for a given observation, but not correlated from one observation to the next. $\mathcal{J}$ includes pulse jitter noise. Finally, δ is the Kronecker delta. The parameters $\mathcal{F}$, $\mathcal{Q}$, and $\mathcal{J}$ account for additional noise which is empirically observed, but not accounted by pulse template fitting uncertainty σ alone (Lam et al. 2016). The 'E' names refer to the parameter names given in the `tempo` and `tempo2` pulsar timing software.

The white noise covariance matrix N is block diagonal. The σs, $\mathcal{F}$s, and $\mathcal{Q}$s run down the diagonal, and the $\mathcal{J}$s form blocks connecting each frequency channel from the same observation.

3.4. *Gaussian Process Models*

We modeled the remaining stochastic processes as Gaussian processes (see Appendix C of NG9 and section 3 of Arzoumanian et al. 2016 for more discussion of this methodology). In this framework each process is defined by T, a design matrix of basis functions, $\mathbf{b}$, a vector of basis coefficients, and B, a covariance matrix defining the Gaussian priors on $\mathbf{b}$.

We defined the intrinsic pulsar red noise as a Fourier-basis Gaussian process. The design matrix T_{red} contains the sine and cosine basis functions, and the coefficients $\mathbf{b}_{\mathrm{red}} = (a, b)_j$ are the Fourier coefficients such that

$$T_{\mathrm{red}} \mathbf{b}_{\mathrm{red}} = \sum_{j=1}^{N} \left[a_j \sin\left(2\pi f_j t\right) + b_j \cos\left(2\pi f_j t\right) \right] \qquad (9)$$

We restricted the sum to the first $N = 30$ Fourier components (starting at the inverse observation time), as this model is for *low-frequency* noise unaccounted elsewhere.

The red noise is modeled as a power-law spectrum

$$P(f_j) = A^2 \left(\frac{f_j}{\mathrm{yr}^{-1}} \right)^{-\gamma} \mathrm{yr}^3, \qquad (10)$$

where A is the characteristic amplitude at the reference frequency of yr^{-1}, and γ is the spectral slope. The spectral shape defines the priors on the Fourier coefficients. The individual Fourier components are approximated as uncorrelated, so the resulting covariance is diagonal:

$$(B_{\mathrm{red}})_{ii} = P(f_i). \qquad (11)$$

We also defined the timing model uncertainty in the Gaussian process framework. The design matrix, T_{tm}, has columns

that are the timing model linearized around the best fit parameters, and the basis coefficients, $\mathbf{b}_{\mathrm{tm}}$, are small offsets from the best fit timing parameters. We placed an unconstrained, uniform prior on the timing offsets by setting the covariance to a diagonal matrix of infinities, $(B_{\mathrm{tm}})_{ii} = \infty$.

We then concatenated the various Gaussian processes into the form of Equation 1.

$$T = \begin{bmatrix} T_{\mathrm{tm}} & T_{\mathrm{red}} \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} \mathbf{b}_{\mathrm{tm}} \\ \mathbf{b}_{\mathrm{red}} \end{bmatrix}, \quad B = \begin{bmatrix} B_{\mathrm{tm}} & \\ & B_{\mathrm{red}} \end{bmatrix}. \quad (12)$$

Additional Gaussian processes can be cleanly added in the same way. For instance, one could include a stochastic GW background as a correlated noise source in all pulsars by adding an additional Gaussian process model. We do not include it in this work, but this data analysis framework allows for a stochastic GW background to be added in a straightforward manner in the future.

3.5. *Solar System Ephemeris Model*

NG11gwb showed that the NANOGrav 11-year dataset is sensitive to uncertainty in the solar system ephemeris (SSE). SSE errors can appear as a spatially correlated stochastic process. As the GW background is also a spatially correlated stochastic effect, there is a natural covariance between the two (Tiburzi et al. 2015). Because GW memory appears as a transient, deterministic effect, there is little confusion between SSE errors and GW memory.

Despite this, we chose to use the same *BayesEphem* model described in NG11gwb to mitigate SSE uncertainty in this analysis. *BayesEphem* implements perturbations to a given SSE by varying 11 parameters: the masses of the gas giants (4), the rotation rate about the ecliptic pole (1), and Juptier's orbital elements (6). We repeated most analyses using two recent JPL SSEs: DE430 and DE436 (Folkner et al. 2014; Folkner & Park 2016). For each SSE we conducted two analyses: holding the SSE fixed and using the *BayesEphem* model. Results reported as "*BayesEphem*" used DE436 as the input SSE before perturbations. There was no measurable difference between using DE430 and DE436 as the *BayesEphem* input.

3.6. *Bayesian analysis*
3.6.1. *Likelihood*

We constructed a Gaussian likelihood based on the white noise covariance. The model residuals, $\mathbf{r}$, should follow the same distribution as the white noise, $\mathbf{n}$:

$$\mathbf{r} = \delta \mathbf{t} - T\mathbf{b} - \mathbf{s}$$
$$p\left(\delta \mathbf{t} \mid \mathbf{b}, \lambda\right) = \frac{\exp(-\frac{1}{2}\mathbf{r} \cdot N^{-1} \cdot \mathbf{r})}{\sqrt{\det(2\pi N)}}, \quad (13)$$

where λ are the model parameters (red and white noise, deterministic signals), and N is defined by Equation 8.

Following the scheme of Lentati et al. (2013) and van Haasteren & Vallisneri (2015), we can analytically marginalize over the Gaussian process coefficients $\mathbf{b}$, leaving us with

$$\mathbf{q} = \delta \mathbf{t} - \mathbf{s}$$
$$C = N + TBT^T$$
$$p\left(\delta \mathbf{t} \mid \lambda\right) = \frac{\exp(-\frac{1}{2}\mathbf{q} \cdot C^{-1} \cdot \mathbf{q})}{\sqrt{\det(2\pi C)}}. \quad (14)$$

The white $(\mathcal{F}, \mathcal{Q}, \mathcal{J})$ and red (A, γ) per-pulsar noise parameters contribute to the final covariance matrix, appearing in N and B, respectively. The Woodbury (1950) matrix identity can be used to evaluate Equation 14 efficiently. Sparse matrix algebra can provide an additional speedup.

The noise parameters that appear in C of Equation 14 were first fit with individual noise analyses for each pulsar. For computational efficiency in the GW analyses, we held the white noise parameters fixed to their median values. The per-pulsar red noise parameters were simultaneously searched with the global GW parameters and SSE parameters, owing to their covariance. With 34 pulsars the search space for the 11-year Earth term analysis contained 84 dimensions, $(34 \times 2_{\mathrm{RN}}) + 5_{\mathrm{GW}} + 11_{\mathrm{SSE}}$.

3.6.2. *Priors*

We prefered ignorance priors for our model parameters, implementing uniform or log-uniform priors for all. We used the same priors for noise parameters as NG11gwb and NG11cw.

For detection analysis we follow the philosophy of NG11gwb, setting a log-uniform prior on h_{mem}. This prior is improper for upper limit analysis: in order to set upper limits on the amplitude of GW memory we must integrate our posterior from a lower bound of $h_{\mathrm{mem}} = 0$. Following the detection analysis, we reran our pipeline using a uniform prior on h_{mem} for the purpose of setting upper limits.

3.6.3. *Inference*

We used the Bayes factor for the GW model compared to a noise only model, $\mathcal{B}_{\mathrm{gw}}$, as our detection statistic. We calculated Bayes factors using the Savage-Dickey approximation (Dickey 1971),

$$\mathcal{B}_{\mathrm{gw}} = \frac{\mathcal{E}_{\mathrm{gw}}}{\mathcal{E}_{\mathrm{noise}}} \approx \lim_{h_{\mathrm{mem}} \to 0} \frac{p(h_{\mathrm{mem}})}{p\left(h_{\mathrm{mem}} \mid \delta \mathbf{t}\right)}, \quad (15)$$

which approximates the evidence ratio $(\mathcal{E}_{\mathrm{gw}}/\mathcal{E}_{\mathrm{noise}})$ for the GW and noise only models as the ratio of the prior to posterior probability in the limit that GW amplitude goes to zero. This is calculation uses posterior samples near the low amplitude prior boundary and is much more computationally efficient than a full evidence integral. The GW model is favored when there is low posterior probability for small h_{mem}, i.e., the posterior for h_{mem} is peaked away from zero.

3.7. *Software*

Our analysis, like NG11gwb and NG11cw, used NANOGrav's core data analysis software, `enterprise`[4] (Ellis et al. 2019) to compute the posterior probability for our models. Likelihood evaluations were sped up using sparse matrix algebra with the `scikit-sparse`[5] Python package and the `SuiteSparse`[6] C library (Chen et al. 2008). We sampled our posterior distribution for Bayesian inference with `PTMCMCSampler`[7] (Ellis & van Haasteren 2017). We used `healpy`[8] (Zonca et al. 2019) and `HEALPix`[9] (Górski et al. 2005) to grid the sky for some analyses.

---

[4] github.com/nanograv/enterprise
[5] github.com/scikit-sparse/scikit-sparse
[6] faculty.cse.tamu.edu/davis/suitesparse.html
[7] github.com/jellis18/PTMCMCSampler
[8] https://github.com/healpy/healpy
[9] https://healpix.jpl.nasa.gov/

Figure 1. Bayes factor distribution for single pulsar term searches for GW memory. Bayes factors less than 1 imply the model without a GW memory burst is preferred. Three of the 34 pulsars have a Bayes factor slightly larger than one, meaning the GW memory model is slightly preferred over the noise model. None are significant.

4. RESULTS

4.1. *Detection Statistics*

We find no evidence for GW memory in the NANOGrav 11-year dataset. For the Earth term search we analyzed the 34 pulsars with a minimum observation baseline of three years (we did not analyze the 11 additional pulsars which had been observed for less than this). We set a log-uniform prior on GW memory strain amplitude and uniform priors on the other search parameters. We searched over burst epochs between 2005.7 and 2015.4. We did not search for GWs in the first 10% and last 5% of the dataset, owing to biases brought about by the loss of sensitivity near the edges of the dataset (this effect is clearly seen in our amplitude upper limits, e.g., Figure 2). The Earth term search results in a Bayes factor of $\mathcal{B}_{gw} \approx 0.7$ for the GW memory model compared to the noise only model.

In the pulsar term search we analyzed data from the same 34 pulsars, individually, calculating the Bayes factor for each. The results are shown in Figure 1. Most pulsars prefer the model without a GW memory burst, having $\mathcal{B}_{gw} < 1$. PSR J1744−1134 is the pulsar with the largest Bayes factor, $\mathcal{B}_{gw} \approx$ 3.5. This Bayes factor is on the threshold of worth mentioning according to the Jeffreys (1961) scale. The preferred burst time, t_0, occurs near the beginning of J1744−1134's observation period, when data were of lower quality (sparser sampling, narrow radio band).

We also searched the NANOGrav 9-year dataset, using the same methods as the 11-year. In this case we found an anomalous GW memory-like signal at MJD 55422±46, about 2010.6. If this were truly a GW signal, it would have appeared in the 11-year data analysis, as well. Seeing as it did not, we can confidently say it must not be a GW memory burst. Additionally, nearly any modification to the individual pulsar noise modeling reduces the significance of this event considerably. For further discussion of this anomalous event see Appendix A.

4.2. *Upper limits*

Finding no evidence for GW memory in our data we place upper limits on the strain of GW memory events during our observations. To compute strain upper limits we sampled the log-amplitude of GW strain, placing an exponential prior on this parameter. This is equivalent to sampling strain amplitude directly with a uniform prior. Our prior choice biases the pos-

Figure 2. 95% upper limit on gravitational wave memory strain amplitude as a function of burst epoch, marginalized over source orientation. The very small semi-transparent regions contain the 90% sampling uncertainty on upper limits. *Top*: Comparison of 9-year and 11-year data sets using *BayesEphem*. The elevated upper limit from the 9-year data set during 2010-2012 is a result of the anomalous signal. See subsection 4.1 and Appendix A for further discussion. *Bottom*: Comparison of 11-year data set under different SSEs. Note that the black curve, 11-year with *BayesEphem* is the same in both.

terior toward higher amplitude. This well known effect results in conservative limits and is not usually a problem. In our case the non-uniform sensitivity of our PTA combines with the amplitude prior in such a way that the most insensitive times and sky positions dominate the posterior.

For GW amplitudes below the sensitivity of our PTA, the likelihood is flat: changes in signal parameters do not change the likelihood. In this case the posterior is dominated by the amplitude prior. The highest probability regions of the posterior will correspond to insensitive source orientations (sky position, polarization) and times, where the GW amplitude can be made largest without affecting the likelihood. The end result is a posterior that is peaked at our PTA's blind spots. If we were to naively perform an all sky, all time search and use the 1D marginal posterior for strain amplitude to compute an upper limit, the limit would be dominated by the most insensitive times and source orientations. This limit would not be a fair representation of our search.

For physical reasons we do not expect GWs to originate from any particular direction. The non-uniformity of our posterior distribution in source orientation is caused by the amplitude prior. We can fix this by implementing a non-uniform prior on source orientation which exactly cancels the bias from the amplitude prior. This type of prior scheme is sometimes called a Malmquist prior, as it corrects for selection effects as Malmquist (1922) did for stellar absolute magnitude. We describe our method to unbias our rate upper limits below.

To place upper limits on the rate of GW memory bursts, we first need strain amplitude upper limits as a function of time. These are computed by determining a source orientation averaged upper limit for each of 40 time bins. We drew MCMC samples for each time bin resulting in a posterior biased to insensitive source orientations, as described above. We then resampled the biased posterior, effectively implementing a *post hoc* prior on source orientation to ensure uniform distributions. In our resampling scheme, we binned the source sky location using an `HEALPix` grid with 48 bins, `nside = 2`.

Figure 3. 95% upper limit on gravitational wave memory strain amplitude as a function of burst epoch. The three curves show the sky averaged Earth term upper limit (same as Figure 2); the Earth term upper limit for the most sensitive sky position for each epoch; and the pulsar term upper limit for an optimally oriented source using the most sensitive single pulsar, PSR J1713+0747.

We binned the polarization angle into 8 bins from 0 to π. We then drew samples from the biased posterior ensuring an equal number of samples lie in each of these 384 source orientation bins. By construction the resulting samples are uniformly distributed in source orientation. The final 1D marginal distribution for strain amplitude is uniformly averaged over source orientation and not biased toward insensitive locations. The results of this effort are shown in Figure 2. We find that our limits are not drastically affected by choice of SSE.

To produce a pulsar term limit we conducted the upper limit versus time analysis for each individual pulsar in the array. First, we computed the upper limit as a function of time for a optimally oriented source. The upper limit for the most sensitive pulsar, PSR J1713+0747, is shown in Figure 3. It is compared to the sky averaged Earth term limit, and the limit using the most sensitive sky position from the Earth term search. We see that for optimally oriented sources the Earth term which combines information from many pulsars is more limiting than the single best pulsar.

An individual pulsar cannot distinguish source orientation as discussed in subsection 3.2. In order to set a limit for all sources, we applied a correction factor to account for a source orientation average. The correction factor comes from analytically marginalizing over source orientation, assuming uniform priors on sky position and polarization angle, and is shown in Appendix B. The choice of prior for the pulsar term sky averaging matches the resampled posterior, ensuring the two can be compared directly.

To explicitly show how the sky position of the source affects the strain upper limit, we conducted a second analysis. In this case we set a limit for each pixel of an `nside = 8` `HEALPix` grid (768 sky locations) marginalizing over polarization angle. Because the sky sensitivity of the PTA changes as new pulsars are added to the array, we focused this analysis to more recent times, ~2012-2015. We used all of the observed TOAs in this analysis, but only searched for GW memory in that time span. The results of this analysis are shown in Figure 4. Our PTA is up to an order of magnitude more sensitive to sources originating from the most sensitive sky positions compared to the least.

Note that the optimal source upper limit shown in Figure 3 dips below 10^{-14}. This optimal source includes the optimal polarization angle, while the analysis shown in Figure 4 marginalizes over polarization angle.

Figure 4. 95% upper limit on gravitational wave memory strain amplitude as a function of sky location of source, using *BayesEphem*. We placed a prior on burst epoch to constrain the analysis to the more recent time span of the data ~2012-2015. The low density of pulsars in RA 0-12h makes us much less sensitive to GW memory originating in that hemisphere. Stars mark the locations of the 34 pulsars used in this work. This map is in equatorial coordinates

5. DISCUSSION AND CONCLUSIONS

From the strain upper limits as a function of burst epoch, we can construct a limit on the rate of GW memory events. Figure 5 shows these limits. The left panel compares the rate limits from this work to those predicted by Islo et al. (2019), who predicted the rate of memory events from SMBHB mergers for LISA and PTAs. Based on their analysis using a simulated galaxy stellar mass function from Sesana et al. (2014) and the $M_{\rm BH} - M_{\rm bulge}$ relation of McConnell & Ma (2013), we should expect a very small number of SMBHB mergers producing GW memory detectable by PTAs. The right panel of Figure 5 compares the limits from this work to the previous published NANOGrav limits of NG5mem.

The median sky position bin from Figure 4 has an upper limit of $\sim 2.5 \times 10^{-14}$. Taking this as representative of our Earth term strain sensitivity, this work improves the rate constraints at that strain by more than 2 orders of magnitude, a factor of ~ 160, relative to NG5mem. For larger amplitudes, $\gtrsim 3 \times 10^{-13}$ this work improves the NG5mem Earth term limits by a factor of 10. Finally, this work improves the large amplitude pulsar term limit by a factor of ~ 20 compared to NG5mem.

Our non-detection of GW memory from SMBHB merger should come as no surprise. While this work was motivated by the prospects of detecting SMBHB merger, GW memory is a generic feature of all GW producing events. Our rate limits are presented as agnostic limits on events that produce GW memory of a particular amplitude.

Future searches for GW memory with PTAs are unlikely to detect SMBHB mergers, but other sources can produce GW memory detectable by PTAs. Because the amplitude of GW memory is proportional to the source mass and inversely proportional to distance, GW memory from much smaller sources could be detectible much closer to home. For instance, Madison et al. (2017) discussed the prospects of detecting stellar mass compact binary mergers in globular clusters in the Milky Way. If memory sources are located in the Milky Way, the data analysis methods used will need to be altered. Nearby sources will violate the plane-wave assumption, so Madison et al. (2017) considered spherically fronted

Figure 5. 95% upper limit on the rate of memory causing events as a function of strain amplitude. *left*: Comparison of the 11-year data set rate upper limit to the memory rate from SMBHB mergers predicted by Islo et al. (2019). *right*: Comparison of the 11-year data set upper limit to the limits previously published in NG5mem. Note that the blue curves are the same in both panels.

waves. Depending on the location of the GW memory producing event relative to the pulsars in the array, an event could activate multiple pulsar terms and/or the Earth term. Searches for these sources will face many of the same challenges that have affected searches for continuous GWs from individual SMBHB sources. Effective methods to simultaneously determine the poorly constrained Earth-pulsar separations and the GW parameters is foremost among these problems (Corbin & Cornish 2010). The solution to this problem for GW memory searches will differ from those implemented in continuous GW searches (e.g., Ellis 2013; Taylor et al. 2014) owing to the transient nature of the GW memory signal. Like in continuous GW searches, incorporating good prior information from pulsar distance measurments will play a crucial role. Pulsar timing provides distance measurements via parallax for some pulsars (e.g., Matthews et al. 2016; NG11). For pulsars with no timing parallax, an independent distance measurement should be incorporated, if available. These could come from Very Long Baseline Interferometry (VLBI, e.g., Deller et al. 2009) or other astrometric experiments (e.g., Mingarelli et al. 2018; Jennings et al. 2018).

There remains a rich discovery space for exotic sources of GW memory (Cutler et al. 2014). As any burst of GWs will produce GW memory, exotic GW producing events such as cosmic strings (Damour & Vilenkin 2001) are possible PTA sources. Even some non-GW effects, such as a cosmic string crossing between the line of sight from the Earth to a pulsar, exhibit a similar response in pulsar timing data (Pshirkov & Tuntsov 2010). Pulsar glitches also produce a signal very similar to GW memory. While glitches are much more common in canonical pulsars, there are some observations of glitches in millisecond pulsars (Cognard & Backer 2004; McKee et al. 2016). It is possible to use our limits on GW memory to place limits on glitches in the 34 millisecond pulsars studied.

More generally, searches for transient GWs in PTAs can reveal transient features in the noise that are not currently modeled. This analysis discovered strange noise features in two of NANOGrav's longest timed pulsars, PSRs J1909−3744 and J0030+0451, in relation to the anomalous signal detected in NG9. Both of these pulsars were found to have unmodeled excess noise by Lam et al. (2017); however, neither stands out as extraordinary in its noise features in that study from other pulsars in the array. These newly discovered noise features, along with surprising noise features uncovered in concurent

NANOGrav data analysis (NG11cw; Hazboun et al. 2019), are driving the development of new PTA data analysis techniques.

Looking into the future, searches for GW memory should remain an integral part of the PTA data analysis regime. These analyses should be implemented on new datasets, like the second data release from the International Pulsar Timing Array (IPTA, Perera et al. 2019).

Figure 6. Results of dropout analysis for the anomalous event in the 9-year dataset. Bayes factors for the signal to be *on* in each pulsar. The event is dominated by two pulsars, PSRs J1909−3744 and J0030+0451. This event was not found in the analysis of the full 11-year dataset indicating it is not a real GW event.

APPENDIX

AN ANOMALOUS EVENT IN THE 9-YEAR ANALYSIS

As stated in subsection 4.1, when analyzing the data from NG9 we find an anomalous GW memory detection. The best fit sky position for this anomalous event is a region of low sensitivity (J2000, R.A. 4^h23^m, Dec. $5°44'$). This sky location combined with the best fit polarization angle conspire to hide this event from nearly all of our pulsars. To determine which pulsars are problematic we can perform a "dropout analysis" introduced in NG11cw. For a dropout analysis we conduct the standard search with a few modifications. First, we fix the parameters of the GW memory signal to match their best fit values from the previous analysis. Next, we introduce a new parameter for each pulsar which acts as a switch. Pulsars that are switched *on* have the GW memory signal included in the likelihood calculation. Pulsars that are switched *off* do not. From the posteriors of these parameters we are able to assess whether each pulsar prefers a noise only model or a noise plus GW memory model. If a pulsar prefers to be in the *off* state, then the noise only model is a better fit to its data.

Bayes factors for *on* versus *off* for each pulsar are shown in Figure 6. PSRs J1909−3744 and J0030+0451 are the only two that show significant preference for the GW memory signal being turned *on* with Bayes factors ∼100. This tells us that the anomalous signal is isolated to these two pulsars and therefore unlikely to be a true GW memory signal. We expect any real GW signal to appear significantly in several pulsars.

Based on the 11-year single pulsar limits, the four most sensitive pulsars in the array to GW memory are PSRs J1713+0747, J1909−3744, J2317+1439, and J1600−3053. The first three were observed for the whole length of the data set, while J1600−3053 was added to the array in 2008. It is worth noting that the three "most *off*" pulsars are included in this list. The presence of any *off* pulsars should be a red flag for validating a detection. If a pulsar is insensitive to a signal it should have no preference for *on* or *off*, therefore a Bayes factor ∼ 1. The anomolous signal has a source orientation that already minimizes the response to most pulsars in the array. The dropout analysis shows us that even with this supressed amplitude our most sensitive pulsars should still be able to see it, yet they do not.

While we are certain that this signal is not a GW, its origin remains a mystery. Deeper investigation into the noise properties of the NANOGrav PTA is ongoing. As the PTA continues to become more sensitive, new noise sources emerge which must be characterized and modeled.

AVERAGING OVER SOURCE ORIENTATION

In order to have a fair comparison between the upper limits found from the Earth term, which were marginalized over source orientation, and the optimal oriented upper limits found from the pulsar term, we must rescale the pulsar term limits accounting for the varying sensitivity depending on source orientation. Most previous work followed van Haasteren & Levin (2010) and used the RMS average of the pulsar angular response function $\langle B^2 \rangle$ to account for source orientation. Because the Earth term upper limit *marginalizes* over source orientation, the fair comparison would do the same. Here we analytically marginalize a pulsar's angular response $B_i(\hat{k}, \psi; \hat{n}_i)$, introduced in Equation 5 over the GW source orientation.

A pulsar's angular response to GW memory is given by Estabrook & Wahlquist (1975) as (dropping the subscript)

$$B(\hat{k}, \psi; \hat{n}) = B(\alpha, \beta) = \frac{1}{2} \cos(2\beta)(1 - \cos(\alpha)) \qquad (B1)$$

$$\cos(\alpha) = \hat{n} \cdot \hat{k},$$

where α is the angle between the line of sight to the pulsar $\hat{n}$ and the source $\hat{k}$, and β is the projected azimuthal angle between the source's principle polarization vector (defined by ψ) and $\hat{n}$ in the plane perpendicular to $\hat{k}$.

For fixed $\hat{n}$ and $\hat{k}$ we can marginalize over the projected source polarization. We integrate only the positive half cycle of β using a uniform prior.

$$\frac{2}{\pi} \int_{-\pi/4}^{\pi/4} d\beta \cos(2\beta) = \frac{2}{\pi} \qquad (B2)$$

Without loss of generality we can align the line of sight to the source with the $\hat{z}$-axis of our coordinate system. Using spherical polar coordinates, $\hat{n} \cdot \hat{k} = \cos\theta$, where θ is co-latitude, and (θ, ϕ) is the source position on the sky. We marginal-

ize over source position using a uniform prior on the whole sphere of solid angle.

$$\frac{1}{4\pi}\int_0^\pi d\theta \int_0^{2\pi} \sin\theta\, d\phi\,(1-\cos\theta) = 1 \qquad (B3)$$

Finally, we put it all together:

$$\frac{1}{2\pi^2}\int_{-\pi/4}^{\pi/4} d\beta \int_0^\pi d\theta \int_0^{2\pi} \sin\theta\, d\phi \left[\frac{1}{2}(1-\cos\theta)\cos(2\beta)\right] = \frac{1}{\pi} \qquad (B4)$$

We can rescale our single pulsar, optimally oriented upper limits by a factor of π to account for non-uniform sensitivity.

<div align="center">REFERENCES</div>

Aggarwal, K., Arzoumanian, Z., Baker, P. T., et al. 2019, ApJ, 880, 116 [1, 3.6.2, 3.7, 5, A]
Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv e-prints, arXiv:1702.00786 [astro-ph.IM] [1]
Arun, K. G., Blanchet, L., Iyer, B. R., & Qusailah, M. S. S. 2004, Class. Quantum Grav., 21, 3771 [1]
Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, ApJ, 810, 150 [1, 3, 3.2, 3.2, 5]
Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, ApJ, 813, 65 [1, 2, 3.4, 5, A]
Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, ApJ, 821, 13 [3.4]
Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018a, ApJS, 235, 37 [1, 2, 3.3, 5]
Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018b, ApJ, 859, 47 [1, 3.3, 3.5, 3.6.2, 3.7]
Blanchet, L., & Damour, T. 1992, Phys. Rev. D, 46, 4304 [1]
Braginsky, V. B., & Grishchuk, L. P. 1985, Sov. Phys. JETP, 62, 427, [Zh. Eksp. Teor. Fiz.89,744(1985)] [1]
Braginsky, V. B., & Thorne, K. S. 1987, Nature, 327, 123 [1]
Burrows, A., & Hayes, J. 1996, Phys. Rev. Lett., 76, 352 [1]
Burt, B. J., Lommen, A. N., & Finn, L. S. 2011, ApJ, 730, 17 [2]
Chen, Y., Davis, T. A., Hager, W. W., & Rajamanickam, S. 2008, ACM Trans. Math. Softw., 35, 22:1 [3.7]
Christodoulou, D. 1991, Phys. Rev. Lett., 67, 1486 [1]
Christy, B., Anella, R., Lommen, A., et al. 2014, ApJ, 794, 163 [2]
Cognard, I., & Backer, D. C. 2004, ApJ, 612, L125 [5]
Corbin, V., & Cornish, N. J. 2010, ArXiv e-prints, arXiv:1008.1782 [astro-ph.HE] [5]
Cordes, J. M., & Jenet, F. A. 2012, ApJ, 752, 54 [1, 3.2]
Cordes, J. M., & Shannon, R. M. 2010, arXiv e-prints, arXiv:1010.3785 [3.1]
Cutler, C., Burke-Spolaor, S., Vallisneri, M., Lazio, J., & Majid, W. 2014, Phys. Rev. D, 89, 042003 [1, 5]
Damour, T., & Vilenkin, A. 2001, Phys. Rev. D, 64, 064008 [5]
Deller, A. T., Tingay, S. J., Bailes, M., & Reynolds, J. E. 2009, ApJ, 701, 1243 [5]
Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, ApJ, 762, 94 [2]
Detweiler, S. 1979, ApJ, 234, 1100 [1]
Dickey, J. M. 1971, The Annals of Mathematical Statistics, 42, 204 [3.6.3]
Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549 [2]
Ellis, J., & van Haasteren, R. 2017, jellis18/PTMCMCSampler: Official Release [3.7]
Ellis, J. A. 2013, Class. Quantum Grav., 30, 224004 [5]
Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, https://github.com/nanograv/enterprise: enterprise [3.7]
Estabrook, F. B., & Wahlquist, H. D. 1975, General Relativity and Gravitation, 6, 439 [3.2, B]
Favata, M. 2009a, ApJ, 696, L159 [1, 3.2]
—. 2009b, Phys. Rev. D, 80, 024002 [1]
—. 2011, Phys. Rev. D, 84, 124013 [1]

Folkner, W. M., & Park, R. S. 2016, JPL planetary and Lunar ephemeris DE436, Tech. rep., Jet Propulsion Laboratory, Pasadena, CA, https://naif.jpl.nasa.gov/pub/naif/JUNO/kernels/spk/de436s.bsp.lbl [3.5]
Folkner, W. M., Williams, J. G., Boggs, D. H., Park, R. S., & Kuchynka, P. 2014, Interplanetary Network Progress Report, 196, 1 [3.5]
Foster, R. S., & Backer, D. C. 1990, ApJ, 361, 300 [1]
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759 [3.7]
Haehnelt, M. G. 1994, MNRAS, 269, 199 [1]
Hazboun, J. S., Simon, J., Taylor, S. R., et al. 2019, arXiv e-prints, arXiv:1909.08644 [5]
Hobbs, G. 2013, Class. Quantum Grav., 30, 224007 [1]
Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, MNRAS, 369, 655 [2]
Islo, K., Simon, J., Burke-Spolaor, S., & Siemens, X. 2019, arXiv e-prints, arXiv:1906.11936 [astro-ph.HE] [5]
Jaffe, A. H., & Backer, D. C. 2003, ApJ, 583, 616 [1]
Jeffreys, H. 1961, Theory of Probability, 3rd edn., Oxford Classic Texts in the Physical Sciences (Oxford University Press) [4.1]
Jennings, R. J., Kaplan, D. L., Chatterjee, S., Cordes, J. M., & Deller, A. T. 2018, ApJ, 864, 26 [5]
Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, ApJ, 819, 155 [3.3, 3.3]
—. 2017, ApJ, 834, 35 [5]
Lasky, P. D., Thrane, E., Levin, Y., Blackman, J., & Chen, Y. 2016, Phys. Rev. Lett., 117, 061102 [1]
Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, Phys. Rev. D, 87, 104021 [3.6.1]
LIGO Scientific Collaboration, Aasi, J., Abbott, B. P., et al. 2015, Class. Quantum Grav., 32, 074001 [1]
Lommen, A. N. 2015, Reports on Progress in Physics, 78, 124901 [1]
Lorimer, D. R. 2008, Living Reviews in Relativity, 11, 8 [1]
Lousto, C. O., Campanelli, M., Zlochower, Y., & Nakano, H. 2010, Class. Quantum Grav., 27, 114006 [3.2]
Madison, D. R., Chernoff, D. F., & Cordes, J. M. 2017, Phys. Rev. D, 96, 123016 [1, 5]
Madison, D. R., Cordes, J. M., & Chatterjee, S. 2014, ApJ, 788, 141 [1, 3.2]
Madison, D. R., Zhu, X.-J., Hobbs, G., et al. 2016, MNRAS, 455, 3662 [1]
Malmquist, K. G. 1922, Meddelanden fran Lunds Astronomiska Observatorium Serie I, 100, 1 [4.2]
Matthews, A. M., Nice, D. J., Fonseca, E., et al. 2016, ApJ, 818, 92 [5]
McConnell, N. J., & Ma, C.-P. 2013, ApJ, 764, 184 [5]
McKee, J. W., Janssen, G. H., Stappers, B. W., et al. 2016, MNRAS, 461, 2809 [5]
McLaughlin, M. A. 2013, Class. Quantum Grav., 30, 224008 [1]
McNeill, L. O., Thrane, E., & Lasky, P. D. 2017, Phys. Rev. Lett., 118, 181103 [1]
Mingarelli, C. M. F., Anderson, L., Bedell, M., & Spergel, D. N. 2018, arXiv e-prints, arXiv:1812.06262 [5]
Perera, B. B. P., DeCesar, M. E., Demorest, P. B., et al. 2019, arXiv e-prints, arXiv:1909.04534 [5]
Pitkin, M. 2012, MNRAS, 425, 2688 [3.2]
Pshirkov, M. S., Baskaran, D., & Postnov, K. A. 2010, MNRAS, 402, 417 [1, 3.2]
Pshirkov, M. S., & Tuntsov, A. V. 2010, Phys. Rev. D, 81, 083519 [5]
Sazhin, M. V. 1978, Soviet Ast., 22, 36 [1]
Sesana, A., Barausse, E., Dotti, M., & Rossi, E. M. 2014, ApJ, 794, 104 [5]
Sesana, A., Vecchio, A., & Volonteri, M. 2009, MNRAS, 394, 2255 [1]
Seto, N. 2009, MNRAS, 400, L38 [1]
Talbot, C., Thrane, E., Lasky, P. D., & Lin, F. 2018, Phys. Rev. D, 98, 064031 [1]
Taylor, S., Ellis, J., & Gair, J. 2014, Phys. Rev. D, 90, 104028 [5]
Thorne, K. S. 1992, Phys. Rev. D, 45, 520 [1]
Tiburzi, C., Hobbs, G., Kerr, M., et al. 2015, MNRAS, 455, 4339 [3.5]
Turner, M., & Will, C. M. 1978, ApJ, 220, 1107 [1]
van Haasteren, R., & Levin, Y. 2010, MNRAS, 401, 2372 [1, 3.2, 3.2, B]
van Haasteren, R., & Vallisneri, M. 2015, MNRAS, 446, 1170 [3.6.1]
Wang, J. B., Hobbs, G., Coles, W., et al. 2015, MNRAS, 446, 1657 [1, 3.2]
Wiseman, A. G., & Will, C. M. 1991, Phys. Rev. D, 44, R2945 [1]
Woodbury, M. A. 1950, Memorandum report, 42, 106 [3.6.1]
Zel'dovich, Y. B., & Polnarev, A. G. 1974, Soviet Ast., 18, 17 [1]
Zonca, A., Singer, L., Lenz, D., et al. 2019, The Journal of Open Source Software, 4, 1298 [3.7]