University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Physics and Astronomy Faculty Publications
and Presentations

College of Sciences

9-10-2015

NANOGrav CONSTRAINTS on GRAVITATIONAL WAVE BURSTS with MEMORY

Z. Arzoumanian

A. Brazier

S. Burke-Spolaor

S. J. Chamberlin

S. Chatterjee

See next page for additional authors

Follow this and additional works at: https://scholarworks.utrgv.edu/pa_fac

Part of the [Astrophysics and Astronomy Commons](#)

Recommended Citation

Z. Arzoumanian, et. al., (2015) NANOGrav CONSTRAINTS on GRAVITATIONAL WAVE BURSTS with MEMORY. *Astrophysical Journal* 810:2. DOI: <http://doi.org/10.1088/0004-637X/810/2/150>

This Article is brought to you for free and open access by the College of Sciences at ScholarWorks @ UTRGV. It has been accepted for inclusion in Physics and Astronomy Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.

Authors

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, and D. J. Nice

NANOGrav CONSTRAINTS ON GRAVITATIONAL WAVE BURSTS WITH MEMORY

Z. ARZOUMANIAN¹, A. BRAZIER², S. BURKE-SPOLAOR^{3,21}, S. J. CHAMBERLIN⁴, S. CHATTERJEE², B. CHRISTY⁵, J. M. CORDES², N. J. CORNISH⁶, P. B. DEMOREST³, X. DENG⁷, T. DOLCH², J. A. ELLIS^{8,22}, R. D. FERDMAN⁹, E. FONSECA¹⁰, N. GARVER-DANIELS¹¹, F. JENET¹², G. JONES¹³, V. M. KASPI⁹, M. KOOP⁷, M. T. LAM², T. J. W. LAZIO⁸, L. LEVIN¹¹, A. N. LOMMEN⁵, D. R. LORIMER¹¹, J. LUO¹², R. S. LYNCH¹⁴, D. R. MADISON², M. A. MCLAUGHLIN¹¹, S. T. MCWILLIAMS¹¹, D. J. NICE¹⁵, N. PALLIYAGURU¹¹, T. T. PENNUCCI¹⁶, S. M. RANSOM¹⁷, X. SIEMENS⁴, I. H. STAIRS¹⁰, D. R. STINEBRING¹⁸, K. STOVALL¹⁹, J. SWIGGUM¹¹, M. VALLISNERI⁸, R. VAN HAASTEREN^{8,22}, Y. WANG¹², AND W. W. ZHU^{10,20}

(NANOGrav COLLABORATION)

¹ Center for Research and Exploration in Space Science and Technology and X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

² Department of Astronomy, Cornell University, Ithaca, NY 14853, USA; drm252@cornell.edu

³ National Radio Astronomy Observatory, 1003 Lopezville Rd., Socorro, NM 87801, USA

⁴ Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA

⁵ Department of Physics and Astronomy, Franklin & Marshall College, P.O. Box 3003, Lancaster, PA 17604, USA

⁶ Department of Physics, Montana State University, Bozeman, MT 59717, USA

⁷ Department of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, USA

⁸ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91106, USA

⁹ Department of Physics, McGill University, 3600 University St., Montreal, QC H3A 2T8, Canada

¹⁰ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

¹¹ Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

¹² Center for Gravitational Wave Astronomy, University of Texas at Brownsville, Brownsville, TX 78520, USA

¹³ Department of Physics, Columbia University, New York, NY 10027, USA

¹⁴ National Radio Astronomy Observatory, P.O. Box 2, Green Bank, WV 24944, USA

¹⁵ Department of Physics, Lafayette College, Easton, PA 18042, USA

¹⁶ University of Virginia, Department of Astronomy, P.O. Box 400325, Charlottesville, VA 22904, USA

¹⁷ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

¹⁸ Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA

¹⁹ Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131, USA

²⁰ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

Received 2015 February 5; accepted 2015 July 20; published 2015 September 9

ABSTRACT

Among efforts to detect gravitational radiation, pulsar timing arrays are uniquely poised to detect “memory” signatures, permanent perturbations in spacetime from highly energetic astrophysical events such as mergers of supermassive black hole binaries. The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) observes dozens of the most stable millisecond pulsars using the Arecibo and Green Bank radio telescopes in an effort to study, among other things, gravitational wave memory. We herein present the results of a search for gravitational wave bursts with memory (BWMs) using the first five years of NANOGrav observations. We develop original methods for dramatically speeding up searches for BWM signals. In the directions of the sky where our sensitivity to BWMs is best, we would detect mergers of binaries with reduced masses of $10^9 M_\odot$ out to distances of 30 Mpc; such massive mergers in the Virgo cluster would be marginally detectable. We find no evidence for BWMs. However, with our non-detection, we set upper limits on the rate at which BWMs of various amplitudes could have occurred during the time spanned by our data—e.g., BWMs with amplitudes greater than 10^{-13} must encounter the Earth at a rate less than 1.5 yr^{-1} .

Key words: gravitational waves – pulsars: general

1. INTRODUCTION

Due to the intrinsically nonlinear nature of Einstein’s equations, all systems that radiate gravitational waves (GWs) are anticipated to produce “memory.” Once a GW propagates through a region of space, a memory component of the wave leaves the space permanently modified (Smarr 1977; Bontz & Price 1979; Braginskii & Thorne 1987; Christodoulou 1991; Blanchet & Damour 1992). Supermassive black hole binaries (SMBHBs), during the final stages of their merger, on a timescale approximately equal to the light-crossing time of the post-merger black hole’s event horizon, are expected to generate GW bursts with memory (BWMs) with sufficiently

large amplitudes to make them potentially detectable with pulsar timing arrays (PTAs; Favata 2009; Seto 2009; Pshirkov et al. 2010; van Haasteren & Levin 2010; Cordes & Jenet 2012; Madison et al. 2014). Cutler et al. (2014) recently singled out memory as a key detection target for PTAs as a probe of exotic and unexpected GW sources like phase transitions in the early universe.

To facilitate GW detection, several international consortia are currently using sensitive radio telescopes paired with pulsar-optimized hardware to realize the PTA concept (Hellings & Downs 1983; Foster & Backer 1990). The European Pulsar Timing Array (EPTA; Kramer & Champion 2013), the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; McLaughlin 2013), and the Parkes Pulsar Timing Array (PPTA; Hobbs 2013) are pushing precision

²¹ Jansky Fellow.

²² Einstein Fellow.

pulsar timing to its limits and developing new data analysis techniques to usher in the era of PTA GW astronomy. By pooling their data and expertise, these consortia have formed the International Pulsar Timing Array (IPTA; Manchester & IPTA 2013), which is poised to become the most sensitive of all PTAs.

In recent years, the various PTAs have begun to place astrophysically meaningful upper limits on continuous GWs from individually resolvable SMBHBs (Arzoumanian et al. 2014; Zhu et al. 2014) and a stochastic background of GWs (van Haasteren et al. 2011; Demorest et al. 2013; Shannon et al. 2013). Wang et al. (2015) searched for BWMs in the first approximately six years of PPTA data; they detected nothing, but determined with 95% confidence that BWMs with amplitudes greater than 10^{-13} occur at a rate less than 0.8 yr^{-1} .

In this paper, we search for BWMs in the first approximately five years of NANOGrav data using new techniques that lead to considerable computational speedups over the search methods described by Madison et al. (2014) and those used by Wang et al. (2015). In Section 2, we describe the data used in our BWM search. In Section 3, we describe the BWM signal model we use for our analysis. In Section 4, we discuss models of noise and how our sensitivity to BWMs is influenced by them. In Sections 5 and 6, we describe our search techniques, differentiating between searches for so-called pulsar-term and Earth-term events. In Sections 7 and 8, we present the results of our pulsar-term and Earth-term searches, respectively. In Section 9, we place upper bounds on BWM rates and amplitudes. In Section 10 we summarize our key results and offer concluding remarks. We provide a summary of our important notational elements in Table 2 in the Appendix.

2. PULSAR TIMING DATA SET

In this section, we discuss several aspects of the data that are relevant to our analysis. For a more thorough description of the data, see Demorest (2007) and Demorest et al. (2013).

The five-year data set consists of pulse times-of-arrival (TOAs) collected at approximately monthly intervals for each of 17 millisecond pulsars (MSPs). All observations were done with the Arecibo radio telescope and the Green Bank Telescope (GBT). Of these 17 MSPs, those visible to Arecibo were observed with Arecibo; all others were observed with the GBT. One pulsar, J1713+0747, was observed with both telescopes. All observations were done using one of two identical backend systems: the Astronomical Signal Processor (ASP) and the Green Bank Astronomical Signal Processor (GASP). These backends performed real-time coherent dedispersion over bands up to 64 MHz wide and recorded the results averaged over channels of width 4 MHz each.

For each observing epoch, several TOAs are reported from various frequency channels of the 64 MHz band. At Arecibo, observations were typically conducted at two widely separated frequencies (usually 430 and 1400 MHz) within one day; at the GBT, observations at two frequencies (usually 820 and 1400 MHz) were conducted within several days of each other. Approximately contemporaneous observations at multiple frequencies allow epoch-to-epoch timing fluctuations caused by changes in dispersion measure (DM) to be accounted for in the timing model fit (Lam et al. 2015).

Three of the pulsars comprising the five-year NANOGrav data set, J1853+1308, J1910+1256, and B1953+29, do not

have sufficient dual-frequency coverage to correct for timing errors from variations in DM over time; we exclude these pulsars from our analysis for this reason. We also exclude the data set for J1600–3053 from our analysis because it is comparatively very short, spanning just two years. Searching for BWMs in such short data sets is feasible in principle, but we avoid it for two reasons. First, the minimum detectable BWM amplitude in a particular data set scales approximately as $T^{-3/2}$ where T is the span of the data set (van Haasteren & Levin 2010), so we do not anticipate that this data set will greatly improve our sensitivity to BWMs. Second, in such short data sets, many timing model parameters are highly covariant with each other (e.g., spin and astrometric parameters; Madison et al. 2013) and with any BWM signal present.

Finally we exclude the data for J1643–1224 from our analysis because we believe it contains chromatic timing biases from unaccounted-for phenomenology in the interstellar medium (ISM; see Figure 2). The DM of J1643–1224 is approximately 62 pc cm^{-3} , nearly a factor of two greater than any other pulsar in our sample. With high DM pulsars, chromatic timing errors that deviate from the ν^{-2} scaling expected from cold plasma dispersion alone become more significant (Cordes & Shannon 2010). Furthermore, this pulsar is directly behind a complex region of H II associated with ζ -Ophiuchi, a massive, runaway O-type star spinning very near breakup (Gaustad et al. 2001; Villamariz & Herrero 2005). This intervening H II region could conceivably contribute to non-trivial and currently unaccounted-for chromatic effects on the timing behavior of J1643–1224.

3. SIGNAL MODEL

For a given pulsar, the timing perturbation from a BWM of amplitude h_B is well-modeled as

$$\Delta t(t) = h_B B(\theta, \phi) [(t - t_0)\Theta(t - t_0) - (t - t_1)\Theta(t - t_1)]. \quad (1)$$

See Pshirkov et al. (2010), van Haasteren & Levin (2010), Cordes & Jenet (2012), or Madison et al. (2014) for discussion of this signal model. The function $B(\theta, \phi) = (1/2)(1 - \cos \theta)\cos(2\phi)$ ranges between -1 and 1 and is common to all pulsar timing efforts to detect point-like sources of GWs (Estabrook & Wahlquist 1975; Hellings & Downs 1983; Lee et al. 2011). The angle between the direction the burst propagates and the line of sight from Earth to the pulsar is θ ; ϕ is the angle between the principal polarization vector of the wave and the projection of the line of sight from the Earth to the pulsar onto the plane normal to the wave propagation direction. The BWM encounters the Earth at a time t_0 and is observed from Earth to encounter the pulsar at a time $t_1 = t_0 + (l/c)(1 + \cos \theta)$ where l is the distance from the Earth to the pulsar (van Haasteren & Levin 2010; Cordes & Jenet 2012; Madison et al. 2014). The function Θ is the Heaviside step function. The amplitude of a BWM coming from a SMBHB merger of reduced mass $\mu \equiv M_1 M_2 / (M_1 + M_2)$ (M_1 and M_2 are the masses of the black holes in the binary) with a typical inclination angle of $\mathcal{I} = \pi/3$

at a luminosity distance D_L from Earth is (Madison et al. 2014)

$$h_B \approx 1.5 \times 10^{-12} \left(\frac{\mu}{10^9 M_\odot} \right) \left(\frac{1 \text{ Mpc}}{D_L} \right). \quad (2)$$

Equation (1) describes a timing perturbation that begins to grow linearly with time instantaneously at t_0 and stops growing instantaneously at t_1 ; the instantaneity of this signal model is an approximate description of an actual BWM from a SMBHB merger, but it is a well justified one. Formally speaking, the memory component of a GW builds up over the entire pre-merger history of a SMBHB, but it grows most precipitously during the final orbits of the binary just prior to merger, requiring a time to build up that is approximately equal to the light-crossing time of the post-merger event horizon—approximately one day for a $10^9 M_\odot$ merger (Cordes & Jenet 2012; Madison et al. 2014). The timescale necessary for the memory component of a GW to build up is substantially smaller than typical pulsar timing measurement cadences of several weeks to a month. The detailed turning-on of the memory signal will not be sufficiently resolved with these relatively low measurement cadences making Equation (1) sufficient to describe the anticipated memory signal.

The distances to the pulsars in the NANOGrav array are on the order of kiloparsecs. Our timing baseline is approximately five years. Unless θ differs from π by less than $\sim 3^\circ$ for a particular pulsar, we expect that t_0 and t_1 will not both fall within our observing window. Since we consider only 12 pulsars in our analysis, less than 1% of the sky is within 3° of one of our pulsars. Assuming BWMs occur isotropically, there is a less than 1% chance of both t_0 and t_1 occurring within our five-year observing span if a BWM occurs at all. And if a BWM occurred with a small enough angular separation from a pulsar in our array that we could see the timing perturbation both turn on at t_0 and turn off at t_1 , it would only be observed to turn off in that one pulsar. So, in each of our pulsar timing data sets, we need only to look for evidence of timing perturbations of the form

$$\Delta t(t) = h_p(t - t_B)\Theta(t - t_B), \quad (3)$$

where $h_p = \pm h_B B(\theta, \phi)$, what we call the projected burst amplitude, and t_B is either t_0 or t_1 , what we call the burst epoch. Bursts arriving at an individual pulsar only influence the timing behavior of that pulsar; we refer to these as pulsar-term bursts (see Section 5). Bursts arriving at the Earth will simultaneously begin to influence the timing behavior of all pulsars in our sample; we refer to these as Earth-term bursts (see Section 6).

4. NOISE MODEL

In this section, we describe several well-known sources of noise we anticipate to be present in our data and the correlations they induce in our measured TOAs (see Equation (7)). We then introduce the empirical parameterized noise model used in our analysis that can accommodate the types of noise we anticipate as well as potentially unexpected noise-like contributions (see Equation (8)). For one pulsar, J1713 + 0747, the noise parameters in Equation (7) are well-measured allowing a direct assessment of the empirical noise model used in our analysis (see Figure 1). Finally, we discuss how to use our noise models to make so-called “epoch-averaged” timing residuals and we display our data set in Figure 2.

Figure 1. Uncertainty on the projected amplitude of a BWM at various trial burst epochs with the NANOGrav five-year data set for PSR J1713+0747 using three different noise models. The lowest curve assumes that only radiometer noise is present in the data and is identical to the results of Madison et al. (2014). The middle curve assumes that only radiometer noise and jitter noise are present in the data (as in Equation (7)); the scale of the jitter contribution is consistent with Shannon & Cordes (2012). The most conservative curve is based on the fixed noise model used in Arzoumanian et al. (2014) (as in Equation (8)) sans a red noise component.

4.1. Physically Motivated Noise Models

Consider TOAs measured using pulse profiles obtained by synchronously averaging $\mathcal{N}$ pulses for each of several radio-frequency channels. The pulse profiles are functions of pulse phase, φ , channel frequency, ν , and observing epoch, τ . A TOA from a particular frequency channel and observing epoch can be written as

$$t_{\nu,\tau} = t_{\infty,\tau} + t_{\text{DM},\tau} + t_{\text{C},\tau} + \epsilon_{\text{S/N},\tau} + \epsilon_{\text{J},\tau} + \epsilon_{\text{DISS},\tau}, \quad (4)$$

where $t_{\infty,\tau}$ is the TOA at infinite frequency, t_{DM} is the dispersive delay from propagation through ionized interstellar plasma, t_{C} is an additional, non-dispersive chromatic perturbation, such as intrinsic profile evolution with frequency, pulse broadening from multipath propagation, and interstellar refraction (Cordes & Shannon 2010). If unaccounted for, t_{DM} and t_{C} can lead to systematic errors in TOA estimates. With multi-frequency observations at each observing epoch, effects from t_{DM} are mitigated in the NANOGrav data set. Potential effects from t_{C} are combatted by fitting for constant inter-channel offsets unique to each pulsar. Unlike t_{DM} and t_{C} , the “ ϵ ” terms in Equation (4) are random errors from a variety of noise sources.

The term $\epsilon_{\text{S/N},\tau}$ is uncorrelated between frequency channels and between observing epochs, i.e.,

$$\langle \epsilon_{\text{S/N},\tau} \epsilon_{\text{S/N},\tau'} \rangle = \sigma_{\text{S/N}}^2(\mathcal{N}) \delta_{\nu\nu'} \delta_{\tau\tau'}, \quad (5)$$

where $\sigma_{\text{S/N}} \propto (\text{S/N})^{-1} \propto \mathcal{N}^{-1/2}$. The quantity $\sigma_{\text{S/N}}$ is the TOA uncertainty from radiometer noise assuming that the radiometer noise adds to a fixed pulse shape under the assumptions of matched filtering against a very high signal-to-noise ratio (S/N) template pulse profile.

The second random contribution to $t_{\nu,\tau}$, $\epsilon_{\text{J},\tau}$, is from phase jitter in single pulses. Jitter is highly correlated between frequency channels, but is known to decorrelate over widely

Figure 2. Epoch-averaged timing residuals for the 12 pulsars used in our analysis (and the excluded pulsar J1643–1224). The various colors indicate different observing frequencies: 327 MHz is green, 430 MHz is blue, 820 MHz is red, 1400 MHz is black, and 2300 MHz is cyan. We show the residuals for J1643–1224 to illustrate the apparent chromatic problems with the data: the residuals from 820 and 1400 MHz are consistently anti-correlated and are often extreme outliers.

separated frequencies; the probability distribution function of the phase of single-pulse centroids evolves with frequencies similar to pulse profiles (Shannon et al. 2014). Jitter noise can be modeled as

$$\langle \epsilon_{J,\nu,\tau} \epsilon_{J,\nu',\tau'} \rangle = \sigma_J^2(\mathcal{N}) \delta_{\tau\tau'} \rho_{J,\nu,\nu'} \approx \sigma_J^2(\mathcal{N}) \delta_{\tau\tau'}, \quad (6)$$

where $\sigma_J \propto \mathcal{N}^{-1/2}$. The quantity $\rho_{J,\nu,\nu'}$ is approximately unity unless ν and ν' are very widely separated.

Usually, $\sigma_J \lesssim \sigma_{S/N}$. Jitter noise begins to dominate radiometer noise only when the S/N of single pulses begins to exceed unity. For our data, collected with ASP and GASP, jitter noise should be significantly subdominant to radiometer noise for all pulsars. Both radiometer and jitter noise are significantly larger than contributions from $\epsilon_{\text{DISS},\nu,\tau}$, a random error associated with diffractive interstellar scintillation (DISS; Cordes 1990; Cordes & Shannon 2010). In a detailed study of J1713+0747, Shannon & Cordes (2012) found that for single pulses, $\sigma_J \approx 27 \pm 1 \mu\text{s}$, while timing errors from DISS

were only a few nanoseconds. A recent 24 hr study of J1713+0747 (Dolch et al. 2014) confirmed the jitter measurement of Shannon & Cordes (2012). Errors from DISS can be substantially more important in the timing of high DM pulsars like B1937+21. We expect DISS to be a subdominant component of our timing error budget for the MSPs we include in our sample, so we ignore it in our analysis. In summary, the noise covariance matrix for a pulsar in our array, with these anticipated sources of noise, can be approximated as

$$\mathbf{C}_{\nu\nu',\tau\tau'} = \langle \epsilon_{\nu,\tau} \epsilon_{\nu',\tau'} \rangle \approx \delta_{\tau\tau'} \left[\delta_{\nu\nu'} \sigma_{S/N}^2(\mathcal{N}) + \sigma_J^2(\mathcal{N}) \right], \quad (7)$$

where $\epsilon_{\nu,\tau}$ is the sum of all noise influencing the TOA from observing epoch τ and frequency channel ν .

4.2. Empirical Noise Models

In a previous analysis of the first five years of NANOGrav data, Arzoumanian et al. (2014) searched for continuous GWs using a different noise model:

$$\mathbf{C}_{\nu\nu',\tau\tau'} = \delta_{\tau\tau'} \left[\delta_{\nu\nu'} (\mathcal{Q}^2 + E^2 \sigma_{S/N}^2(\mathcal{N})) + \mathcal{J}^2 \right] + \mathbf{C}_{\text{red}}(A, \gamma). \quad (8)$$

This noise model expands on the physically motivated model of Equation (7), allowing for additional, otherwise-unmodeled contributions to the noise. The terms $\mathcal{Q}$ and E are commonly referred to as EQUAD (a source of Gaussian white noise with time units added in quadrature to radiometer noise) and EFAC (a dimensionless constant multiplier on the anticipated amount of radiometer noise), respectively. The term $\mathcal{J}$ mimics the correlations induced between TOAs from different frequency channels of the same observing epoch by jitter, but qualitatively differs from jitter in that its value does not change to reflect the number of pulses, $\mathcal{N}$, averaged together to yield a TOA. Most observations in the five-year NANOGrav data set we considered had similar integration times, so $\mathcal{N}$ does not fluctuate substantially between epochs and $\mathcal{J}$ is thus very much like a jitter-induced correlation. Different values of $\mathcal{Q}$, E , and $\mathcal{J}$ were used for each widely separated observing frequency of each pulsar, e.g., one value of $\mathcal{Q}$ is used for observations of a pulsar at frequencies centered on 820 MHz while another value is used for observations of that pulsar centered on 1400 MHz. We regard these frequency bands as widely separated since their spacing is much larger than the maximum 64 MHz bandwidth used in our observations. The quantities $\mathcal{Q}$, E , and $\mathcal{J}$ describe inter-channel correlations between TOAs from within a single frequency band; TOAs from different frequency bands measured at different times are uncorrelated by the types of noise we are considering.

The Arzoumanian et al. (2014) noise model also included red noise with a power-law power spectrum $P(f) = A[f/(1 \text{ yr}^{-1})]^\gamma$. This component of the noise model is not to be thought of as part of the TOA error budget; TOAs measured with extreme accuracy may still differ from the expectations of a timing model because of spin noise intrinsic to a pulsar, which is well modeled as a stochastic process with a power-law spectrum (Shannon & Cordes 2010). In the Arzoumanian et al. (2014) analysis, under the assumption that no GW signal was present in the data set, the noise model parameters $\Xi = [A, \gamma, E, \mathcal{Q}, \mathcal{J}]$ were determined by finding

the maximum of the likelihood function

$$\mathcal{L}(\mathbf{R}|\delta p, \Xi) = \frac{\exp\left[-\frac{1}{2}(\mathbf{R} - \mathbf{M}\delta p)^T \mathbf{C}^{-1}(\mathbf{R} - \mathbf{M}\delta p)\right]}{\sqrt{(2\pi)^{N_{\text{TOA}}} \det \mathbf{C}}}, \quad (9)$$

where $\mathbf{R}$ are the timing residuals from an initial timing model, δp are deviations of the timing model parameters from the initial timing model, $\mathbf{M}$ is the timing model design matrix, and N_{TOA} is the total number of TOAs in the data set.

We applied the noise model assessment done by Arzoumanian et al. (2014) to two simulations of the five-year NANOGrav data set for PSR J1909–3744, one that contained only simulated radiometer noise and one that contained radiometer noise and a bright BWM occurring at the midpoint of the data set. The resulting best-fit noise models were identical except for in the red noise parameters A and γ . The BWM signal is heavily covariant with the red noise component of the noise model, resulting in a large value for A and a poorly constrained value of γ .

An independent analysis of the five-year NANOGrav data set by Perrodin et al. (2013) found that except for J1643–1224 and J1910+1256, pulsars we already excluded from our analysis, all pulsars were consistent with white noise alone. Because the demonstrated covariance between red noise and BWM signatures would complicate detection of a BWM (Cordes & Jenet 2012) and because of the lack of strong evidence for red noise, we have opted to consider the Arzoumanian et al. (2014) fixed noise models without the red noise component.

4.3. Comparing Noise Models

In Figure 1, we show how the uncertainty on the projected BWM amplitude, $\sigma_{h,p}$, varies in the NANOGrav five-year data set for PSR J1713+0747 over several trial burst epochs under three different noise models. The quantity $\sigma_{h,p}$ is a direct measure of how sensitive a particular data set is to BWMs and it is directly influenced by the noise model; it is discussed by, e.g., van Haasteren & Levin (2010) and Madison et al. (2014).

The bottom curve of Figure 1 is based on a noise model that only includes radiometer noise. The middle curve uses the noise model described by Equation (7) with the jitter measurement from Shannon & Cordes (2012). The top curve is based on the noise model used by Arzoumanian et al. (2014) and described in Equation (8) (without the red noise component). The discrepancy between the physically motivated noise model of Equation (7) and the empirical noise model of Equation (8) is not currently well understood. Detailed studies like Dolch et al. (2014) are being carried out to better understand the noise budget of pulsar timing experiments and any systematic effects that influence the timing procedure, but bridging the gap between the top and middle curves of Figure 1 is an ongoing area of research. The Arzoumanian et al. (2014) noise model is the most conservative of the three we consider and we use it for the remainder of our analysis.

4.4. Epoch Averaging

In Figure 2, we depict the epoch-averaged timing residuals from the 12 pulsars in our sample (and the excluded pulsar J1643–1224). Residuals are obtained individually for multiple frequency channels for each integration lasting 15–45 minutes.

The residuals from that integration are combined to form a single epoch-averaged residual. The noise model is central to this procedure. We define an operator,

$$\mathbf{A} = (\mathbf{U}^T \mathbf{C}^{-1} \mathbf{U})^{-1} \mathbf{U}^T \mathbf{C}^{-1}, \quad (10)$$

which maps the raw residuals, $\mathbf{R}$, to epoch-averaged residuals, $\mathbf{R}_E = \mathbf{A}\mathbf{R}$. The matrix $\mathbf{C}$ is the noise covariance matrix as described in Equation (8). The matrix $\mathbf{U}$ is the “exploder” matrix discussed in Arzoumanian et al. (2014) that maps epochs to the full set of TOAs. The uncertainties on the epoch-averaged residuals are the square roots of the diagonal entries of the matrix

$$\mathbf{C}_E = \langle \mathbf{R}_E \mathbf{R}_E^T \rangle = (\mathbf{U}^T \mathbf{C}^{-1} \mathbf{U})^{-1}. \quad (11)$$

5. SEARCHES FOR BWMS IN THE PULSAR TERM

Information in the pulsar terms of the 12 pulsars in our sample comes from causally distinct regions of spacetime (Cordes & Jenet 2012). A BWM is 12 times more likely to encounter a single pulsar in our array than it is to encounter the Earth. However, it is not true that BWMS encountering individual pulsars in the array are 12 times more likely to be confidently detected. Different pulsars are timed with varying degrees of precision and with differing observing cadences making them unequally sensitive in searches for BWMS. A particular BWM may be polarized in such a way or from such a part of the sky that the projection factor, $B(\theta, \phi)$, makes its influence on the timing behavior of a particular pulsar vanishingly small. Furthermore, non-BWM phenomena such as intrinsic pulsar spin noise (Shannon & Cordes 2010) and microglitches (Cognard & Backer 2004) can be confused as BWMS. Without the signal appearing concurrently in multiple pulsars as in an Earth-term BWM, it is difficult to rule out pulsar-specific phenomena. Nonetheless, we carry out a search for BWMS in each of our individual pulsars as non-detections in many pulsars allow us to place constraints on otherwise inaccessible regions of BWM parameter space.

Madison et al. (2014) describe techniques for searching for BWMS in individual pulsar timing data sets and for assessing the minimal projected amplitude detectable at a particular epoch. The timing perturbation from a BWM is deterministic and can be included as part of the timing model. The projected amplitude enters the timing model as a linear parameter and can be fit for using least-squares methods (Gregory 2010), yielding an estimate for the projected amplitude, $\hat{h}_p$, and its uncertainty, $\sigma_{h,p}$. We deal with the nonlinear parameter t_B by searching over a grid of trial burst epochs. For all of our timing model fits and calculation of timing residuals, we use the software package TEMPO2 (Edwards et al. 2006).

The modification to the timing solution caused by including a burst of projected amplitude h_p at time t_B can be assessed by computing the likelihood ratio for a model with and without a burst:

$$\Gamma(h_p, t_B) = \exp\left(-\frac{1}{2}[\chi^2(h_p, t_B) - \chi^2(0, t_B)]\right). \quad (12)$$

In this expression,

$$\chi^2(h_p, t_B) = \mathbf{R}^T(h_p, t_B) \mathbf{C}^{-1} \mathbf{R}(h_p, t_B), \quad (13)$$

where $\mathbf{R}(h_p, t_B)$ are the timing residuals when a burst of projected amplitude h_p at time t_B is included as part of the timing model. The least-squares estimator for the projected burst amplitude at a trial burst epoch t_B , $\hat{h}_p(t_B)$, maximizes Γ at that trial burst epoch; call this value $\hat{\Gamma}(t_B)$. For a fixed trial burst epoch, if the data are consistent with the noise model, $\hat{D} \equiv 2 \ln \hat{\Gamma}$, the reduction in the χ^2 of the residuals caused by introducing a projected burst amplitude to the timing model fit will be a random variable following a χ^2 distribution with one degree of freedom:

$$f_1(\hat{D}) = (2\pi\hat{D})^{-1/2} \exp(-\hat{D}/2). \quad (14)$$

Since the burst amplitude is a linear parameter in the timing model, we expect the χ^2 value of the residuals to respond quadratically to the burst amplitude, i.e.,

$$\chi^2(0, t_B) = \chi^2(\hat{h}_p(t_B), t_B) + [\hat{h}_p(t_B)/\sigma_{h,p}(t_B)]^2, \quad (15)$$

or $\hat{D}(t_B) = [\hat{h}_p(t_B)/\sigma_{h,p}(t_B)]^2$. The $\sigma_{h,p}$ values are the least-squares 1σ amplitude uncertainties on $\hat{h}_p$ and are the quantities plotted in Figure 1. For a trial burst time t_B , $\sigma_{h,p}(t_B)$ is the square root of the diagonal element of the parameter covariance matrix, $(\mathbf{M}^T \mathbf{C}^{-1} \mathbf{M})^{-1}$, associated with the amplitude of a BWM occurring at t_B ($\mathbf{M}$ is the fit design matrix; see Madison et al. 2014 for more discussion of $\mathbf{M}$ and $\sigma_{h,p}$).

If we compute $\hat{D}$ for two trial epochs that are very close together, the results will be correlated. Because of these correlations, when computing $\hat{D}$ along a densely sampled grid of N_t trial burst epochs, the number of effectively independent trial burst epochs tested is $N_l < N_t$. The probability distribution function for the maximum value of $\hat{D}$ along the grid, $\hat{D}_{\max}$, is

$$f_{N_l}(\hat{D}_{\max}) = N_l f_1(\hat{D}_{\max}) F_1(\hat{D}_{\max})^{(N_l-1)}, \quad (16)$$

where F_1 is the cumulative distribution of f_1 ,

$$F_1(\hat{D}) = \text{erf}\left(\sqrt{\hat{D}/2}\right). \quad (17)$$

The cumulative distribution for $\hat{D}_{\max}$ associated with f_{N_l} is then simply

$$F_{N_l}(\hat{D}_{\max}) = \text{erf}^{N_l}\left(\sqrt{\hat{D}_{\max}/2}\right). \quad (18)$$

The anticipated false alarm probability for noise alone to exceed a threshold value of $\hat{D}_{\max}$, $\hat{D}_{\text{thresh}}$, is just

$$\mathcal{F}_{N_l}(\hat{D}_{\text{thresh}}) = 1 - F_{N_l}(\hat{D}_{\text{thresh}}). \quad (19)$$

For a fixed allowable false alarm probability, $\hat{D}_{\text{thresh}}$ grows logarithmically with N_l if $N_l \gg 1$.

To estimate N_l , for each pulsar, we generated 1000 simulated sets of TOAs that matched the real data set in number of TOAs, observing schedule, and timing model, but yielded timing residuals consistent with our noise models. We then fit for $\hat{h}_p(t_B)$ along an equispaced grid of twenty trial burst epochs between MJD 53500 and 54900 and computed $\hat{D}(t_B)$ to get $\hat{D}_{\max}$.

In Figure 3, we show the cumulative distribution of the 1000 $\hat{D}_{\max}$ values from our simulations for three pulsars: J0030+0451 and J2145–0750, observed by Arecibo and the GBT,

Figure 3. Cumulative distributions of $\hat{D}_{\max}$ values from 1000 simulations of noise-like timing residuals for three of the pulsars in our sample. Overlaid are the anticipated cumulative distributions if N_l , the effective number of trial burst epochs tested, is one, five, or twenty. For all pulsars in our sample, the analytically anticipated cumulative distribution best fits the results of our simulations with N_l between 4 and 5.

respectively, with rms timing residuals of ~ 100 ns, and J1713+0747, observed by both Arecibo and the GBT with an rms timing residual of ~ 30 ns. We also plot the theoretically anticipated curves for N_l equal to one, five, and twenty. In fitting the anticipated cumulative distributions of $\hat{D}_{\max}$ with N_l as a free parameter to the results of our simulations, we find that for all pulsars, N_l is between four and five. As N_l increases, large values of $\hat{D}_{\max}$ occur more frequently in the presence of pure noise. We take the conservative approach and assume that there are five independent trial burst epochs to test. With N_l fixed at five, inverting Equation (19) allows us to compute $\hat{D}_{\text{thresh}}$ for any desired allowable false alarm probability; false alarm probabilities of approximately five and one percent are expected for $\hat{D}_{\text{thresh}}$ equal to 6.60 and 9.54, respectively.

6. SEARCHES FOR BWMS IN THE EARTH TERM

Searches for BWMs in the Earth term have several advantages over searches in pulsar terms. The timing perturbation from a BWM will turn on simultaneously for all pulsars in the PTA if it occurs in the Earth term, providing a powerful means by which pulsar-specific phenomena (e.g., glitches) can be ruled out. Variations in the projected BWM amplitude from pulsar to pulsar provide information about the location of the GW source and its polarization. As such, for a trial BWM source location, polarization, and epoch, the residuals of all pulsars in the PTA can be combined into a coherent simultaneous fit for all timing model parameters and the amplitude of a BWM, h_B . We refer to these simultaneous multi-pulsar fits as global fits. Such global fits, described in Madison et al. (2014), have better amplitude sensitivity than what is attainable with any one timing data set and can be carried out with a similar least-squares apparatus as is used in pulsar term searches.

Global least-squares fitting techniques were recently used by the PPTA to search for BWMs (Wang et al. 2015). The PPTA data set those authors analyzed contained many fewer TOAs than the NANOGrav data set we are considering here owing to

NANOGrav’s practice of reporting many TOAs from a single observing epoch but from different frequency channels. With these multifrequency TOAs, NANOGrav includes as part of its timing models numerous chromatic parameters to account for profile evolution across our wide-bandwidth receivers and time-variable DM, so, we fit for many more timing model parameters than the PPTA. The PPTA attempts to mitigate effects of time variable DM in a very different way that does not require the simultaneous fitting of many chromatic timing model parameters (see Keith et al. 2013); it is unclear which practice is better. Optimal methods for mitigating DM variations are an active area of pulsar timing research (e.g., Lam et al. 2015).

Suppose N is the number of timing model parameters being fit for each of M pulsars in a PTA (in reality, N varies from pulsar to pulsar). Then carrying out a global fit for the amplitude of a BWM requires the inversion of an $(NM + 1) \times (NM + 1)$ matrix, a procedure that requires computation time $\propto (NM + 1)^3$. This global fit has to be done over $N_\Omega N_\psi N_t$ trials in the four-dimensional phase space of sky position, burst epoch, and burst polarization (N_Ω , N_ψ , and N_t are, respectively, the number of sky positions, polarization angles, and burst epochs tested). Methods requiring these global fits are thus computationally onerous for NANOGrav because of the large number of timing model parameters for which we fit. More importantly, the total number of suitably stable MSPs being timed by NANOGrav is growing with time as projects like the Arecibo PALFA survey (e.g., Cordes et al. 2006; Swiggum et al. 2014), the Arecibo All-sky 327 MHz Drift Pulsar Survey (e.g., Deneva et al. 2013), and the GBT Northern Celestial Cap Pulsar Survey (e.g., Stovall et al. 2014) continue to find more pulsars. A search procedure based on global fits will not scale well to future PTA endeavors. We use a different technique that avoids having to carry out any global fits and is computationally much more efficient.

6.1. Accelerated Earth-term Searches

Over a five-dimensional grid of trial burst sky positions, Ω_{ij} , polarizations, ψ_k , epochs, $t_{B,l}$, and amplitudes, $h_{B,m}$, we search for the maximum in the global likelihood ratio surface,

$$\Gamma_G(\Omega_{ij}, \psi_k, t_{B,l}, h_{B,m}) = \prod_{K=1}^M \Gamma_K[B_K(\Omega_{ij}, \psi_k)h_{B,m}, t_{B,l}], \quad (20)$$

where B_K is the projection factor, $B(\theta, \phi)$, for the K th pulsar. The Γ_K quantities are equal to the likelihood ratios in Equation (12). The angles θ and ϕ depend on the location of the K th pulsar and the trial burst location, Ω_{ij} . The trial burst polarization angle, ψ_k , influences ϕ (see Equation (6) in Madison et al. 2014). The total number of trials in this five-dimensional search is $N_\Omega N_\psi N_t N_h$, where N_h is the number of trial burst amplitudes tested.

Ostensibly, the procedure appears to require that for each grid point in the five-dimensional space and for each pulsar in our sample, we incorporate the signature of a BWM of projected amplitude $B_K(\Omega_{ij}, \psi_k)h_{B,m}$ occurring at time $t_{B,l}$ into the timing model of the K th pulsar, refit the timing model, and compute Γ_K . This would take computation time $\propto N_\Omega N_\psi N_t N_h MN^3$, a speed up over the global fitting scheme so long as $N_h < (NM + 1)^3/MN^3$. The average N for our

sample of $M = 12$ pulsars is approximately 90, meaning N_h needs to be less than approximately 140 for this technique to be faster.

However, there is a greater speed-up to be had. Variations in Ω_{ij} , ψ_k , and $h_{B,m}$ only alter the projected BWM amplitude along different pulsar lines of sight. We can precompute the two-dimensional (projected burst amplitude and burst epoch) likelihood ratio surface for each pulsar in our sample, a procedure requiring computation time $\propto N_p N_t MN^3$, where N_p is the number of trial projected burst amplitudes considered. We then construct the full five-dimensional likelihood ratio surface from Equation (20) by figuring out what the projected burst amplitude in the K th pulsar would be for a particular choice of Ω_{ij} , ψ_k , and $h_{B,m}$, looking to the precomputed two-dimensional likelihood ratio surface for the K th pulsar, and interpolating between the nearest projected burst amplitudes tested and the projected amplitude of interest from the five-dimensional search; this step requires no additional timing model fits (the computationally costly step), is very fast, and is essentially an exercise in efficiently searching lookup tables.

Using precomputed two-dimensional likelihood ratio surfaces for a global Earth-term BWM search as we have just described will lead to a speed up over the global fitting scheme if $N_p < N_\Omega N_\psi (NM + 1)^3/MN^3$. For our search, we have tested $N_t = 40$ trial burst epochs evenly spaced between MJDs 53541 and 54995, $N_\psi = 17$ trial burst polarization angles evenly spaced between 0 and π , and $N_\Omega = 1598$ trial burst sky positions isotropically distributed on the sky. As a brief but important aside, we have chosen to search within the particular window of dates just mentioned because for each pulsar in our sample, there is at least one collection of multifrequency observations before and after this window, with the limiting pulsar on the early side being J0613–0200 and J2145–0750 on the late side. Since we are searching for changes in the observed rotational frequencies of the pulsars in our array occurring at some trial burst epoch, a pulsar data set for which there is no data before or after that epoch should not bear any weight in the search. With the grid parameters we have chosen, so long as N_p is less than approximately 4×10^6 , using precomputed two-dimensional likelihood ratio surfaces will be faster than global fitting schemes. In practice, we have set $N_p = 300$, testing 150 trial projected amplitudes logarithmically spaced between 5×10^{-17} and 10^{-12} and their negatives. Utilizing the precomputed two-dimensional likelihood ratio surfaces is thus faster than using global fitting by a factor of over thirteen thousand and is the only reason why searching such a densely sampled grid is feasible.

Once the global, five-dimensional likelihood ratio surface is computed, for fixed Ω_{ij} , ψ_k , and $t_{B,l}$, we isolate the value of $h_{B,m}$ that maximizes Γ_G ; we call it $\tilde{h}_B$. This four-dimensional surface, which we call $\tilde{\Gamma}_G$, is approximately equal to $\hat{\Gamma}_G$, what we would get if we had carried out the computationally costly global fits we have taken great care to avoid, differing from $\hat{\Gamma}_G$ only because of the finite resolution of our trial burst amplitude grid. We can then, as in the pulsar-term case, consider the false alarm statistics of the quantity $\tilde{D}_G = 2 \ln \tilde{\Gamma}_G$. For any fixed grid point in our four-dimensional parameter space, if the data are consistent with our noise models, $\tilde{D}_G$ again follows χ^2 statistics with one degree of freedom. With noise-like data, if we consider the probability distribution of $\tilde{D}_{G,\max}$, the maximum value of $\tilde{D}_G$ over the whole grid, also follows Equation (16), but with a larger number of effective independent trials

sampled than in the pulsar-term case owing to a search not just over time, but also over polarization angle and sky position. We call the effective number of independent trials in the global search N_G . We discuss N_G in more detail in Section 8.

6.2. Assessing BWM Amplitude Uncertainty

A similar relationship exists between $\tilde{D}_G$, $\tilde{h}_B$, and σ_h as exists between $\hat{D}$, $\hat{h}_p$ and $\sigma_{h,p}$ as discussed in Section 5, i.e., $\hat{D} = (\hat{h}_p/\sigma_{h,p})^2$. Since Γ_G is just the product of all the pulsar-specific likelihood ratio surfaces (with appropriate projection factors, as in Equation (20)),

$$\tilde{D}_G = \sum_{K=1}^M D_K (B_K \tilde{h}_B) = \tilde{h}_B^2 \sum_{K=1}^M (B_K/\sigma_{h,p,K})^2. \quad (21)$$

So, we have

$$\tilde{D}_G = (\tilde{h}_B/\sigma_h)^2, \quad (22)$$

where

$$\sigma_h = \left[\sum_{K=1}^M (B_K/\sigma_{h,p,K})^2 \right]^{-1/2}. \quad (23)$$

Equation (23) is equivalent to Equation (32) from van Haasteren & Levin (2010). While those authors analyzed idealized pulsar timing data sets with equispaced observing epochs, uniform TOA uncertainties, and very simple timing models, their result holds more generally and applies here.

6.3. Cross-checks with Bayesian Methods

As an independent check, we also carry out the Bayesian method of van Haasteren & Levin (2010), implemented in the software package *Piccard*²³, developed independently of the Madison et al. (2014) method. In brief, this Bayesian method takes the likelihood of Equation (9) as a function of the noise parameters Ξ , the timing model parameters δp , and the BWM parameters $(\Omega_{ij}, \psi_k, t_{B,i}, h_{B,m})$. Similar to what we do for our frequentist search, we keep the noise parameters fixed to the values obtained in Arzoumanian et al. (2014). Our prior distributions are flat in all stated BWM parameters (including the BWM amplitude h_p), where we note that the prior in Ω_{ij} is taken flat over the sphere. Given a significance level, an upper limit can then be set using the empirical cumulative marginalized posterior distribution as a function of h_p .

We have been able to confirm that the results we describe in this paper are identical for the frequentist and Bayesian methods; specifically, we have compared sky-slices of the detection statistic $\tilde{D}_G$ similar to what is depicted in Figure 5. At the time of this analysis, a full Bayesian analysis would have been significantly more computationally expensive than the frequentist analysis we primarily discuss in this work. However, the new techniques we have developed for accelerating Earth-term searches can be readily integrated into Bayesian searches and will yield comparable computational speedups as what we discussed in Section 6.1. We are working to implement these changes and plan to use the accelerated search techniques discussed here in analyses of future data sets.

²³ See <https://github.com/vhaasteren/piccard> for access to the publicly available code, a description of its capabilities, implementation details, and several illustrative examples.

Table 1
Summary of Results from Pulsar Term BWM Search

PSR	$\hat{h}_p^{\max}/10^{-13}$	$t_{B,i}^{\max}$ (MJD)	$\hat{D}_{\max}$	$\log_{10}(\mathcal{F}_5)$
J0030+0451	-0.87 ± 0.77	54100	1.2	-0.11
J0613-0200	0.68 ± 0.50	54584	1.8	-0.20
J1012+5307	1.47 ± 0.79	54062	3.4	-0.54
J1455-3330	-2.57 ± 3.35	54211	0.5	-0.02
J1640+2224	-2.65 ± 1.45	53801	3.3	-0.52
J1713+0747	0.08 ± 0.07	54808	1.4	-0.13
J1744-1134	-2.06 ± 1.19	53541	2.9	-0.45
B1855+09	1.00 ± 0.73	53578	1.8	-0.21
J1909-3744	-1.29 ± 0.74	54994	3.0	-0.45
J1918-0642	-8.29 ± 3.39	54994	5.9	-1.15
J2145-0750	26.60 ± 11.51	54994	5.3	-0.99
J2317+1439	-2.40 ± 1.20	53541	3.9	-0.67

Note. Description of most significant BWM-like signal detected in the pulsar term of each of 12 NANOGrav data sets. From left to right, the columns are the name of the pulsar, the most significant projected BWM amplitude calculated from least-squares fitting, the trial burst epoch at which the most significant amplitude was found, the corresponding $\hat{D}_{\max}$ value, and the logarithm of the anticipated false alarm probability for that value of $\hat{D}_{\max}$ assuming N_t , the number of effectively independent trial burst epochs tested, is five.

7. PULSAR TERM RESULTS

For all twelve of the pulsar data sets we analyze, our search for pulsar-term BWMs occurring between MJDs 53541 and 54994 yields results that are entirely consistent with our noise models. In Table 1, we summarize the key results in our search for pulsar-term BWMs. For each pulsar in our sample, we list the pulsar name, the most significant value of $\hat{h}_p$ (according to the $\hat{D}$ number) along with its 1σ uncertainty, the trial burst epoch at which this most significant amplitude was found, the $\hat{D}_{\max}$ value for that pulsar, and the logarithm of the false alarm probability anticipated from noise alone for that value of $\hat{D}_{\max}$ assuming $N_t = 5$. The most significant event we find is in the data for J1918-0642; it is consistent with approximately seven percent of noise realizations.

The epochs at which the most significant projected BWM amplitudes occur are distributed nearly uniformly throughout the range of trial epochs we tested with repeat values occurring only at the very first and last epochs tested. Clustering at the edges of the window of tested epochs is not surprising. If there are few timing residuals outside of the window of tested trial burst epochs, when testing the first or last trial epoch, the pre- or post-burst timing model is constrained by a small number of data points. If the residuals at the edge of the data set are slight outliers, allowing for an instantaneous change in the spin period of the pulsar near the beginning or end of the data set can bring them more in line with zero and lead to a modest reduction of the χ^2 value of the residuals.

The individual data set in our sample that most tightly constrains BWMs is for J1713+0747. This is in line with the expectations of Madison et al. (2014). In Figure 4, we show the detailed BWM search results for J1713+0747. The bottom panel is simply the epoch-averaged residuals as in Figure 2; we plot them again here to explicitly show how the interval of trial burst epochs tested overlaps with the TOA coverage. The second panel from the bottom shows $\hat{h}_p \pm \sigma_{h,p}$ at the 40 trial burst epochs we tested. The best-fit projected amplitude is completely consistent with zero at each trial burst epoch

considered. Further echoing this, in the third panel from the bottom, we show the value of $\hat{D}$ as it varies with trial burst epoch. Never does $\hat{D}$ approach the values necessary to be inconsistent with noise at the five- or one-percent level. In the remaining panels of Figure 4, we show a “heat map” of the two-dimensional likelihood ratio surface, as in Equation (12), that we use in our Earth-term search.

8. EARTH TERM RESULTS

In Figure 5, we show a two-dimensional slice of our Earth-term search results. For fixed trial burst polarization angle and epoch, we depict the quantity $\tilde{D}_G^{1/2}$ as it varies with trial burst source location. The quantity $\tilde{D}_G^{1/2}$ is the best-fit BWM amplitude, $\tilde{h}_B$, in units of the uncertainty on the amplitude, σ_h (see Equation (22)). The slice we show contains the smallest value of σ_h found in our entire search, i.e., the point in our parameter space where we are most sensitive to a BWM; its location on the sky is indicated by the black circle very near the location of J1713+0747.

The single largest value of $\tilde{D}_G$ we find is $\tilde{D}_{G,\max} = 6.03$, which corresponds to $\tilde{h}_G = 2.46\sigma_h$. We mentioned at the end of Section 6.1 that the effective number of trials tested in an Earth-term BWM search, N_G , will be larger than $N_1 = 5$ because of the search over not just many trial burst epochs, but also over many trial burst source locations and polarizations, implying that for a fixed allowable false-alarm probability we have to raise the threshold we impose on $\tilde{D}_G$ above the threshold used on the detection statistic in a pulsar-term search. However, $\tilde{D}_{G,\max}$ is small enough that even if N_G were only five, this result would be entirely consistent with more than 95% of realizations of our noise.

In the BWM analysis carried out by Wang et al. (2015), when faced with marginally high values of their test statistic, they conducted an extensive suite of simulations to assess N_G , comparable to what we have done to assess N_1 (and depicted in Figure 3), and were able to justifiably raise the detection threshold on their test statistic and rule out a detection. We find no comparably large values of $\tilde{D}_G$ that exceed our detection threshold even if we underestimate N_G as five.

We do still want an estimate of N_G as it will allow us to apply upper limits to the population of BWMs given our non-detection. Several recent papers have demonstrated that the response of a PTA to GWs of any waveform can be decomposed into a linear combination of a finite number of modes, or sky maps (Cornish & van Haasteren 2014; Gair et al. 2014). The number of modes required is equal to twice the number of pulsars in the array (the factor of two accounts for the two possible polarization modes of GWs). We use their result to estimate that the number of statistically independent samples in our Earth-term search over source-position and polarization space is 24, twice the number of pulsars used in our analysis. Given five independent samples in time, we thus adopt $N_G = 120$.

Adopting $N_G = 120$ is a conservative estimate as our sensitivity to BWMs is so strongly dominated by a single pulsar, so the effective number of pulsars in our array is fewer than 12. As mentioned following Equation (19), for a fixed allowable false alarm probability, if $N_G \gg 1$, $\tilde{D}_{G,\text{thresh}}$ only diverges logarithmically with N_G , so overestimates of N_G are not exceedingly deleterious for the purpose of setting upper limits. With $N_G = 120$, setting $\tilde{D}_{G,\text{thresh}} = 12.4$ assures a false alarm probability of less than five percent. This is equivalent to requiring that $\tilde{h}_B$ be greater

than $3.52\sigma_h$ in order for it to be inconsistent with 95% of realizations of noise. Again, in our Earth-term search, we find no signal that meets or exceeds this level of significance.

In Figure 6, we have averaged σ_h over trial burst polarization angles and epochs and shown the maximum luminosity distance at which a SMBHB merger with an inclination angle of $\pi/3$ and $\mu = 10^9 M_\odot$ (consistent with Equation (2)) would be detectable with our data set with 95% confidence, or where $\tilde{h}_B = 3.52\langle\sigma_h\rangle_{\psi,t}$. As per Equation (2), these luminosity distance limits scale linearly with the fiducial reduced mass which we have here set to $10^9 M_\odot$. The total volume in which we would be sensitive to such an event is approximately $25.6 \times 10^3 \text{ Mpc}^3$. Our sensitivity is worst near the position of PSR J0613–0200. This has little to do with J0613–0200, but is instead because these sky positions are antipodal to our greatest concentration of pulsars, especially our most sensitive pulsar, J1713+0747. The green triangle in Figure 6 indicates the position of the Virgo Cluster. Just 16.5 Mpc from Earth, the Virgo Cluster falls very near the edge of the volume over which we are sensitive to BWMs from $\mu = 10^9 M_\odot$, $\mathcal{I} = \pi/3$ binary mergers.

9. BWM RATE-AMPLITUDE CONSTRAINTS

To synthesize the results from both our Earth- and pulsar-term analyses, we derive constraints on $\Lambda(>h_B)$, the annual rate of BWMs from any part of the sky with any polarization having amplitudes greater than h_B , assuming they occur as a Poisson process. This quantity is readily relatable to astrophysical models of processes producing BWMs, i.e., Equation (15) of Cordes & Jenet (2012), Equation (21) of Madison et al. (2014), or Equation (17) of Wang et al. (2015).

Toward this end, define the quantities,

$$\tau_E(h_B, n) = \frac{\Delta t \Delta \Omega \Delta \psi}{4\pi^2} \sum_i^{N_t} \sum_j^{N_\Omega} \sum_k^{N_\psi} \Theta(h_B - n\sigma_{h,ijk}), \quad (24)$$

and

$$\tau_P(h_B, n) = \frac{\Delta t}{4\pi^2} \sum_K^M \sum_i^{N_t} \int \Theta\left(h_B - \frac{n\sigma_{h,p,K}}{B(\theta, \phi)}\right) d\psi d\Omega, \quad (25)$$

where Δt , $\Delta \Omega$, and $\Delta \psi$ describe the grid spacing in the Earth-term search, $\sigma_{h,p,K}$ is the uncertainty on the projected burst amplitude from the K th pulsar timing data set, and $\sigma_{h,ijk}$ is the uncertainty on the true (unprojected) burst amplitude evaluated at the i th trial burst time, j th trial burst sky location, and k th trial burst polarization angle. The quantities τ_E and τ_P are the total time that the PTA had n - σ sensitivity to a BWM of amplitude at least h_B in the Earth-term and pulsar-term, respectively, weighted by the fraction of the total source-location and polarization angle space over which that sensitivity was achieved. Our definitions for τ_E and τ_P differ from nearly identical definitions in Madison et al. (2014) by an overall factor of two. In that work, it was mistakenly assumed that only BWM polarization angles between 0 and $\pi/2$ must be considered. We here correctly carry out a search that tests polarization angles between 0 and π .

With the definitions for τ_E and τ_P in place, we can derive constraints on $\Lambda(>h_B)$ from our Earth-term and pulsar-term analyses:

$$\Lambda(>h_B) < -\frac{\ln(1 - \mathcal{Q})}{\tau_i}, \quad (26)$$

Figure 4. Results of a search for a BWM in the NANOGrav five-year data set for PSR J1713+0747, the single most sensitive data set in our sample for such searches. The bottom panel shows the epoch-averaged timing residuals for J1713+0747 as they are in Figure 2. The second panel from the bottom shows $\hat{h}_p \pm \sigma_{h,p}$ for the 40 trial burst epochs we tested. The third panel from the bottom shows the detection statistic $\hat{D}$ varying over the span of trial burst epochs tested. The remaining panels show a “heat map” of the two-dimensional likelihood ratio surface, as in Equation (12), that we use for our Earth-term analysis.

Figure 5. Time and polarization slice of our Earth-term search. The plotted quantity, $\hat{D}_G^{1/2}$, is equivalent to the best-fit BWM amplitude, $\hat{h}_B$, in units of the $1\text{-}\sigma$ uncertainty on the amplitude, σ_h . At the location indicated by the black circle, $\sigma_h = 6.07 \times 10^{-15}$, the smallest value of σ_h in our entire search. In our whole search, the maximum value of $\hat{D}_G^{1/2}$ we find is approximately 2.46, entirely consistent with our noise models at the 95% level even if N_G is no greater than 5 as we used in our pulsar-term searches. The largest diamonds indicate the positions of the 12 pulsars in our analysis. The largest diamond represents J1713+0747.

where Q is the probability that at least one BWM occurring at rate $\Lambda(>h_B)$ encounters the PTA during the time τ_i and i is a placeholder for either E or P .

In Figure 7, we show our constraints on $\Lambda(>h_B)$ from our Earth- and pulsar-term analyses. We have set $Q = 0.95$. We have set $n = 3.52$ for both our Earth- and pulsar-term constraints. This number comes from requiring a false alarm probability of less than five-percent in an Earth-term search with $N_G = 120$. Our estimate of N_G is likely larger than it needs to be. We can compare our choice of $N_G = 120$ to the work done in Wang et al. (2015). They determined, through extensive simulations, that the number of independent trials in their analysis was approximately 80. They used 20 pulsars in their analysis (compared to our 12) and avoided a search through burst polarization space by fitting for the polarization angle. If we had conducted our search as they did, based on the arguments we used to arrive at our chosen value of $N_G = 120$, we would increase N_G by a factor of $5/3$ to account for the use of 20 rather than 12 pulsars and then divide it by two to account for the lack of search through polarization space to arrive at 100 independent trials. This is more conservative than the value of 80 that they used, but for a fixed allowable false alarm probability, the amplitude a BWM must exceed to be inconsistent with the noise scales as the square root of the natural logarithm of N_G , so modest overestimates of N_G have a minimal impact on upper limits. With the chosen values for Q and n , we are setting a 95% confidence frequentist upper limit. We do this because it is a reasonably stringent and commonly used confidence threshold and to facilitate direct comparison with the results of Wang et al. (2015).

Since there are fewer effectively independent trials in a pulsar-term search, we should use a lower value of n for a 95% confidence upper limit. But, by treating the Earth- and pulsar-terms comparably, we are able to highlight an important fact about our data set. The near-total convergence of our Earth- and pulsar-term constraints for low values of h_B indicates that our upper limits are almost entirely dominated by a single pulsar data set: J1713+0747. This was anticipated by Madison et al. (2014) and holds true even with the more sophisticated noise modeling we have employed in this analysis. The pulsar-term constraint plunges to lower rates than the Earth-term constraint for high values of h_B because the pulsar terms can be treated as causally independent and the effective time baseline in the pulsar term search is greater than for the Earth term. Asymptotically, the two curves will differ by a factor of 12 for large h_B since our analysis uses equal amounts of observing time from 12 different pulsars.

10. CONCLUSION

In this paper, we have conducted a search for BWMs in the first five years of NANOGrav data. We did not detect any BWMs. Based on our current understanding of SMBHBs, the most conventional anticipated source of bright BWMs, it is unsurprising that we did not. Wang et al. (2015) predict that BWMs from SMBHB mergers exceeding amplitudes of 10^{-14} occur at a rate of just a few every 10^5 year. However, Cordes & Jenet (2012) conclude that because of large uncertainties in things such as the inspiral rate of SMBHBs, the actual rate of BWMs is essentially unconstrained. If, for example, the rate at which SMBHBs inspiral is enhanced through interactions with gas or stars in the innermost portions of a galaxy, the amount of time that individual binaries dwell at a single orbital frequency will decrease and low-frequency power in the stochastic GW background will be attenuated. This could adversely affect the prospects for detecting individual isolated inspiralling

Figure 6. Maximum luminosity distance of a SMBHB merger causing a BWM detectable with 95% confidence given our sensitivity averaged over burst epoch and polarization angle. We have assumed the binary had a typical inclination angle of $\pi/3$ and a reduced mass of $10^9 M_\odot$; the plotted distances scale linearly with this fiducial reduced mass. The diamonds indicate the positions of the 12 pulsars in our analysis. The largest diamond represents J1713+0747. The green triangle represents the position of the center of the Virgo Cluster. Just 16.5 Mpc from Earth, the Virgo Cluster is near the edge of the volume in which we can detect BWMs with these properties.

SMBHBs and could delay the detection of a stochastic background, but the rate of BWMs would be enhanced. Also, we stress that memory is a very general feature of bright GW events and large-amplitude BWMs may be produced by wholly unexpected phenomena occurring at unconstrained rates; ongoing searches for BWMs are crucial for utilizing the raw discovery potential of PTA observations (Cutler et al. 2014). The methods developed in this paper are readily generalizable to future, more sensitive data sets and the accelerated search techniques we have developed for Earth-term analyses will greatly expedite future BWM searches.

Our Figure 7 and Figure 10 from Wang et al. (2015) show constraints on the rate of BWMs from initial NANOGrav and PPTA data releases, respectively. Both probe to minimum BWM amplitudes of approximately 2×10^{-14} . The PPTA limits the rate of BWMs with amplitudes of 10^{-13} and 4×10^{-14} to less than approximately 0.8 and 3 yr^{-1} , respectively. Our upper limits on the rate of BWMs at these amplitudes are 1.5 and 6.2 yr^{-1} , respectively, a factor of approximately two less constraining. This difference can be in part attributed to the slightly longer timing baseline used in the PPTA’s analysis as the amplitude of the minimum detectable BWM scales as $T^{-3/2}$ where T is the length of the data set. Another contribution to the difference in our upper limits is that the PPTA and NANOGrav observe different sets of pulsars; the PPTA data set for PSR J0437–4715, a very bright and well-timed MSP that is too far south to be observed with Arecibo or the GBT, is a powerful tool in constraining BWMs.

Our Figure 6 and Figure 9 from Wang et al. (2015) illustrate the sensitivity of NANOGrav and the PPTA to BWMs as it varies over the sky. Combined, these figures help support the science case of the IPTA. The area of the sky where NANOGrav’s sensitivity to BWMs is worst is non-concentric with the area of the sky where the PPTA’s sensitivity is worst. Joint analysis of data from NANOGrav and the PPTA will lead to more uniform sensitivity to BWMs over the whole sky and more constraining upper limits. Though the IPTA has not yet conducted a search for BWMs with their data, the inclusion of their data in the joint IPTA data set will likely play a similarly important role in improving the uniformity of the IPTA’s BWM sensitivity.

NANOGrav has made great strides in improving its BWM sensitivity since the collection of this initial data set. The pulsar

Figure 7. 95% confidence Earth- and pulsar-term upper bound on the rate, $\Lambda(>h_B)$, of BWMs occurring with amplitudes at or above h_B . The two curves come from our Earth-term and pulsar-term analyses. The pulsar-term probes lower-rate events at high amplitudes because individual pulsar terms contain causally independent information. The near total convergence of the Earth- and pulsar-term curves at low BWM amplitudes demonstrates that a single pulsar, J1713+0747, is dominating our sensitivity.

timing backends at Arecibo and Green Bank, ASP and GASP, have been upgraded to the Puerto Rican Ultimate Pulsar Processing Instrument and the Green Bank Ultimate Pulsar Processing Instrument (DuPlain et al. 2008), backends with exceptionally wide bandwidths that have reduced the rms timing errors on most pulsars being timed by a factor of two to three. Furthermore, NANOGrav is now regularly timing more than 40 pulsars rather than just 17 as in the first five years. Finally, just having a longer timing baseline on the pulsars we have analyzed in this paper is a great boon to our BWM sensitivity. We reiterate that all else being equal, sensitivity to BWMs scales approximately as $T^{-3/2}$, where T is the span of the data set. Contrast this with the $T^{-1/2}$ scaling of a PTA’s sensitivity to individual inspiralling SMBHBs. If we focus on SMBHBs above a minimum reduced mass (as in Figure 6), the volume of space in which we are sensitive to memory from their mergers grows as $T^{9/2}$. The volume of space to which we are sensitive to individual inspiralling SMBHBs of a fixed chirp mass grows only as $T^{3/2}$. As the volume in which we are sensitive to billion solar mass mergers already encompasses parts of the Virgo cluster, this strong scaling of volume probed with time means that many more astrophysically interesting systems will enter our detection horizon with continued PTA observations. NANOGrav is preparing a collection of nine years worth of data that will be a significantly more sensitive probe of BWMs than any PTA data set that has yet been analyzed. We plan to apply the techniques used in this paper to the nine-year NANOGrav data set to produce unprecedented constraints on BWMs.

We thank our anonymous referee for many thoughtful and thorough comments that helped us to improve this manuscript. We thank T. Loredó for helping us to appropriately treat extreme-value statistics and G. Hobbs for his maintenance of and continual improvements to TEMPO2. The work of Z.A., A.B., S.B.-S., S.J.C., S.C., B.C., N.J.C.J.M.C., P.B.D., T.D., J.A.E., N.G.-D., F.J., G.J., M.T.L., T.J.W.L., L.L., A.N.L., D.R.L., J.L., R.S.L., D.R.M., M.A.M., S.T.M., D.J.N., N.P.,

T.T.P., S.M.R., X.S., D.R.S., K.S., J.S., M.V., R.vH. and Y.W. was partially supported through the National Science Foundation (NSF) PIRE program award number 0968296. NANOGrav research at UBC is supported by an NSERC Discovery Grant and Discovery Accelerator Supplement and by the Canadian Institute for Advanced Research. The work of the NANOGrav collaboration is partially supported by the NANOGrav Physics Frontiers Center NSF Award PHYS-1430284. D.R.M. acknowledges partial support through the New York Space Grant Consortium. J.A.E. and R.vH. acknowledge support by NASA through Einstein Fellowship grants PF4-150120 and PF3-140116, respectively. M.V. acknowledges support from the JPL RTD program. Portions of this research were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Parts of the analysis in this work

were carried out on the Nimrod cluster of S.M.R. Data for this project were collected using the facilities of the National Radio Astronomy Observatory and the Arecibo Observatory. The National Radio Astronomy Observatory is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc. The Arecibo Observatory is operated by SRI International under a cooperative agreement with the NSF (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana and the Universities Space Research Association.

APPENDIX

In Table 2 we list and briefly describe the key symbols used throughout this paper.

Table 2
Table of Key Symbols

Symbol	Description	Dimensions
$\hat{}$	... a “hat” on any quantity indicates that the BWM amplitude used to calculate it is the result of least-squares fitting	... dimensionless
$\sim$	... a “tilde” on any quantity indicates that the BWM amplitude used to calculate it maximizes the global likelihood (i.e., Equation (20)) among trial amplitudes in a grid search	... dimensionless
$B(\theta, \phi)$	... projection factor accounting for the geometric configuration of the Earth, pulsar, and burst source and the burst polarization angle	... dimensionless
$C_{\nu\nu', \tau\tau'}$	... covariance from noise between a TOA from epoch τ and frequency channel ν and a TOA from epoch τ' and frequency channel ν'	... time ²
$D(h_p, t_B)$	... $2 \ln \Gamma(h_p, t_B)$	... dimensionless
D_G	... $2 \ln \Gamma_G$	... dimensionless
D_L	... luminosity distance from Earth to BWM source	... distance
E	... EFAC, a constant multiplier on the rms perturbation attributed to radiometer noise	... dimensionless
f_k	... the p.d.f. for the maximum of k random numbers drawn from a χ^2 distribution with one degree of freedom	... dimensionless
F_k	... the c.d.f. associated with f_k	... dimensionless
$\mathcal{F}_k$	... one minus the c.d.f. associated with f_k	... dimensionless
h_B	... amplitude of a BWM	... dimensionless
h_p	... projected BWM amplitude, $\pm h_B B(\theta, \phi)$	... dimensionless
$\mathcal{J}$	... jitter-like covariance between TOAs from the same epoch but different frequency channels	... time
N_I	... the effective number of trial BWM epochs tested in a single-pulsar BWM search	... dimensionless
N_G	... the effective number of trials in the space of BWM epoch, sky position, and polarization angle in a global Earth-term BWM search	... dimensionless
$\mathcal{Q}$	... EQUAD, the rms of a Gaussian noise process added in quadrature to radiometer noise	... time
t_0	... time at which BWM wavefront encounters the Earth	... time
t_1	... time at which BWM wavefront is observed from Earth to encounter a pulsar	... time
t_B	... epoch BWM is observed to occur, either t_0 or t_1	... time
U	... the “exploder” matrix that maps observation epochs to the full set of TOAs, as discussed in Arzoumanian et al. (2014)	... dimensionless
$\Gamma(h_p, t_B)$	... likelihood of a pulsar’s TOAs assuming a BWM of projected amplitude h_p occurred at epoch t_B divided by the likelihood of the TOAs assuming no BWM occurred	... dimensionless
Γ_G	... the global likelihood ratio, or product of pulsar-wise likelihood ratios (i.e., Equation (20)) for a trial BWM of fixed sky location, polarization, epoch, and amplitude	... dimensionless
Δt	... perturbation to pulse times of arrival	... time
$\Lambda(>h_B)$	... the rate that BWMs with amplitudes greater than h_B encounter our PTA	... time ⁻¹
μ	... reduced mass of a SMBHB, $M_1 M_2 / (M_1 + M_2)$	... mass
σ_h	... uncertainty on the amplitude of a BWM	... dimensionless
$\sigma_{h,p}$	... uncertainty on the projected amplitude of a BWM	... dimensionless
σ_j	... rms timing perturbation from pulse phase jitter noise	... time
$\sigma_{S/N}$	... rms timing perturbation from radiometer noise	... time
$\tau_E(h_B, n)$	... total time that an Earth-term search could yield $n\text{-}\sigma_h$ detections of a BWM with an amplitude greater than h_B	... time
$\tau_P(h_B, n)$	... total time that a pulsar-term search could yield $n\text{-}\sigma_h$ detections of a BWM with an amplitude greater than h_B	... time
$\chi^2(h_p, t_B)$	... the square of the timing residuals when a BWM is included in the timing model weighted by the inverse noise covariance matrix (see Equation (13))	... dimensionless

Note. For reference, descriptions of important or frequently used symbols.

REFERENCES

- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2014, *ApJ*, **794**, 141
- Blanchet, L., & Damour, T. 1992, *PhRvD*, **46**, 4304
- Bontz, R. J., & Price, R. H. 1979, *ApJ*, **228**, 560
- Braginskii, V. B., & Thorne, K. S. 1987, *Natur*, **327**, 123
- Christodoulou, D. 1991, *PhRvL*, **67**, 1486
- Cognard, I., & Backer, D. C. 2004, *ApJL*, **612**, L125
- Cordes, J. M. 1990, in *Low Frequency Astrophysics from Space*, Vol. 362, ed. N. E. Kassim & K. W. Weiler (Berlin: Springer), 165
- Cordes, J. M., Freire, P. C. C., Lorimer, D. R., et al. 2006, *ApJ*, **637**, 446
- Cordes, J. M., & Jenet, F. A. 2012, *ApJ*, **752**, 54
- Cordes, J. M., & Shannon, R. M. 2010, arXiv:1010.3785
- Cornish, N. J., & van Haasteren, R. 2014, arXiv:1406.4511
- Cutler, C., Burke-Spolaor, S., Vallisneri, M., Lazio, J., & Majid, W. 2014, *PhRvD*, **89**, 042003
- Demorest, P. B. 2007, PhD thesis, Univ. California Berkeley
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, *ApJ*, **762**, 94
- Deneva, J. S., Stovall, K., McLaughlin, M. A., et al. 2013, *ApJ*, **775**, 51
- Dolch, T., Lam, M. T., Cordes, J., et al. 2014, *ApJ*, **794**, 21
- DuPlain, R., Ransom, S., Demorest, P., et al. 2008, *Proc SPIE*, **7019**, 1
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, **372**, 1549
- Estabrook, F. B., & Wahlquist, H. D. 1975, *GReGr*, **6**, 439
- Favata, M. 2009, *ApJL*, **696**, L159
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300
- Gair, J., Romano, J. D., Taylor, S., & Mingarelli, C. M. F. 2014, *PhRvD*, **90**, 082001
- Gaustad, J. E., McCullough, P. R., Rosing, W., & van Buren, D. 2001, *PASP*, **113**, 1326
- Gregory, P. 2010, *Bayesian Logical Data Analysis for the Physical Sciences* (Cambridge: Cambridge Univ. Press)
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39
- Hobbs, G. 2013, *CQGra*, **30**, 224007
- Keith, M. J., Coles, W., Shannon, R. M., et al. 2013, *MNRAS*, **429**, 2161
- Kramer, M., & Champion, D. J. 2013, *CQGra*, **30**, 224009
- Lam, M. T., Cordes, J. M., Chatterjee, S., & Dolch, T. 2015, *ApJ*, **801**, 130
- Lee, K. J., Wex, N., Kramer, M., et al. 2011, *MNRAS*, **414**, 3251
- Madison, D. R., Chatterjee, S., & Cordes, J. M. 2013, *ApJ*, **777**, 104
- Madison, D. R., Cordes, J. M., & Chatterjee, S. 2014, *ApJ*, **788**, 141
- Manchester, R. N. & IPTA 2013, *CQGra*, **30**, 224010
- McLaughlin, M. A. 2013, *CQGra*, **30**, 224008
- Perrodin, D., Jenet, F., Lommen, A., et al. 2013, arXiv:1311.3693
- Pshirkov, M. S., Baskaran, D., & Postnov, K. A. 2010, *MNRAS*, **402**, 417
- Seto, N. 2009, *MNRAS*, **400**, L38
- Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, **725**, 1607
- Shannon, R. M., & Cordes, J. M. 2012, *ApJ*, **761**, 64
- Shannon, R. M., Osłowski, S., Dai, S., et al. 2014, *MNRAS*, **443**, 1463
- Shannon, R. M., Ravi, V., Coles, W. A., et al. 2013, *Sci*, **342**, 334
- Smarr, L. 1977, *PhRvD*, **15**, 2069
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, **791**, 67
- Swiggum, J. K., Lorimer, D. R., McLaughlin, M. A., et al. 2014, *ApJ*, **787**, 137
- van Haasteren, R., & Levin, Y. 2010, *MNRAS*, **401**, 2372
- van Haasteren, R., Levin, Y., Janssen, G. H., et al. 2011, *MNRAS*, **414**, 3117
- Villamariz, M. R., & Herrero, A. 2005, *A&A*, **442**, 263
- Wang, J. B., Hobbs, G., Coles, W., et al. 2015, *MNRAS*, **446**, 1657
- Zhu, X.-J., Hobbs, G., Wen, L., et al. 2014, *MNRAS*, **444**, 3709