

Versatile directional searches for gravitational waves with Pulsar Timing Arrays

D. R. Madison,^{1,2★} X.-J. Zhu,^{3,4} G. Hobbs,⁴ W. Coles,⁵ R. M. Shannon,^{4,6}
J. B. Wang,⁷ C. Tiburzi,^{8,9,10,11} R. N. Manchester,⁴ M. Bailes,¹² N. D. R. Bhat,⁶
S. Burke-Spolaor,¹³ S. Dai,^{4,14} J. Dempsey,¹⁵ M. Keith,¹⁶ M. Kerr,⁴ P. Lasky,^{17,18}
Y. Levin,¹⁷ S. Osłowski,^{10,11} V. Ravi,¹⁹ D. Reardon,^{4,17} P. Rosado,¹² R. Spiewak,²⁰
W. van Straten,¹² L. Toomey,⁴ L. Wen³ and X. You²¹

¹Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

²National Radio Astronomy Observatory, 520 Edgemont Rd., Charlottesville, VA 22903, USA

³School of Physics, University of Western Australia, 35 Stirling Hwy, Crawley WA 6009, Australia

⁴CSIRO Astronomy and Space Science, Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia

⁵Electrical and Computer Engineering, University of California at San Diego, La Jolla, 92093 California, USA

⁶International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

⁷Xinjiang Astronomical Observatory, CAS, 150 Science 1-Street, Urumqi, Xinjiang 830011, China

⁸INAF - Osservatorio Astronomico di Cagliari, Via della Scienza, 09047 Selargius (CA), Italy

⁹Dipartimento di Fisica, Università di Cagliari, Cittadella Universitaria 09042 Monserrato (CA), Italy

¹⁰Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

¹¹Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany

¹²Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia

¹³National Radio Astronomy Observatory, 1003 Lopezville Rd., Socorro, NM 87801, USA

¹⁴Department of Astronomy, School of Physics, Peking University, Beijing 100871, China

¹⁵CSIRO Information Management & Technology (IM&T), PO Box 225, Dickson ACT 2602, Australia

¹⁶Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

¹⁷Monash Centre for Astrophysics (MoCA) School of Physics and Astronomy, Monash University, VIC 3800, Australia

¹⁸School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

¹⁹California Institute of Technology, Pasadena, CA 91125, USA

²⁰Department of Physics, University of Wisconsin-Milwaukee, PO Box 413, Milwaukee, WI 53201, USA

²¹School of Physical Science and Technology, Southwest University, Chongqing 400715, China

Accepted 2015 October 26. Received 2015 October 19

ABSTRACT

By regularly monitoring the most stable millisecond pulsars over many years, pulsar timing arrays (PTAs) are positioned to detect and study correlations in the timing behaviour of those pulsars. Gravitational waves (GWs) from supermassive black hole binaries (SMBHBs) are an exciting potentially detectable source of such correlations. We describe a straightforward technique by which a PTA can be ‘phased-up’ to form time series of the two polarization modes of GWs coming from a particular direction of the sky. Our technique requires no assumptions regarding the time-domain behaviour of a GW signal. This method has already been used to place stringent bounds on GWs from individual SMBHBs in circular orbits. Here, we describe the methodology and demonstrate the versatility of the technique in searches for a wide variety of GW signals including bursts with unmodelled waveforms. Using the first six years of data from the Parkes Pulsar Timing Array, we conduct an all-sky search for a detectable excess of GW power from any direction. For the lines of sight to several nearby massive galaxy clusters, we carry out a more detailed search for GW bursts with memory, which are distinct signatures of SMBHB mergers. In all cases, we find that the data are consistent with noise.

Key words: gravitational waves – methods: data analysis – pulsars: general.

1 INTRODUCTION

Pulsar timing arrays (PTAs) provide a unique means for detecting gravitational waves (GWs) with frequencies between approximately

*E-mail: drm252@cornell.edu

10^{-9} and 10^{-6} Hz. Sensitivity at such frequencies makes PTAs key for searching for and eventually studying GWs from supermassive black hole binaries (SMBHBs) with masses greater than $10^7 M_\odot$ (Sesana & Vecchio 2010). PTAs are therefore indispensable for understanding galaxy evolution over cosmological time-scales, for investigating the mechanisms by which the final parsec problem (Milosavljević & Merritt 2003) may be solved, and for probing strong gravitational fields.

The PTA concept was initially conceived decades ago (Detweiler 1979; Hellings & Downs 1983; Foster & Backer 1990), and is now being realized by several international collaborations. 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; Manchester et al. 2013) collaborations use sensitive radio telescopes to observe the most rotationally stable millisecond pulsars (MSPs) known and measure the highest-precision pulse times of arrival (ToAs) possible on a regular basis. Measured ToAs are compared to predictions of timing models that aim to account for all of the known physical effects that modulate the regularity with which pulses arrive at Earth-based observatories (Edwards, Hobbs & Manchester 2006). After sufficiently long periods of time, the differences between measured ToAs and model predictions, the timing residuals, may begin to show structure indicative of errors in terrestrial time standards (Hobbs et al. 2012), incorrect Solar system ephemerides (Champion et al. 2010), or the influence of GWs on the pulsar-Earth system. To analyse such effects, the various PTAs are also combining their data and expertise in the International Pulsar Timing Array (IPTA; Hobbs et al. 2010; Manchester 2013) which is poised to become the most sensitive tool for such investigations.

The GW signals potentially detectable by PTAs fall into two classes. First, individually resolvable sources such as SMBHBs in circular (Arzoumanian et al. 2014; Zhu et al. 2014; Babak et al. 2015) or eccentric orbits (Ravi et al. 2014; Huerta et al. 2015), and burst sources, especially so-called ‘bursts with memory’ (BWMs) from the final merger of SMBHBs (van Haasteren & Levin 2010; Cordes & Jenet 2012; Wang et al. 2015; Arzoumanian et al. 2015a) or potentially from exotic sources like phase transitions in the early Universe (Cutler et al. 2014). Secondly, an isotropic stochastic background (SB) of GWs created by the incoherent superposition of many unresolved SMBHBs scattered throughout the Universe (van Haasteren et al. 2011; Demorest et al. 2013; Shannon et al. 2013, 2015; Arzoumanian et al. 2015b; Lentati et al. 2015). An isotropic SB is an idealization and recent work has begun to place limits on anisotropic features of the background (Taylor et al. 2015). In this paper, we are concerned primarily with individually resolvable sources or small groups of bright sources clustered in a particular direction, though our techniques may also prove useful in studies of a SB.

When GWs from a single direction interact with the Earth, regardless of their detailed waveforms, they produce a distinctly quadrupolar correlation pattern in the timing residuals of the pulsars in a PTA. By exploiting this fact, pulsar timing data sets from many different pulsars can be coherently combined, or ‘phased-up’, so as to enhance a PTA’s sensitivity to GWs from particular directions. Here, we develop a formalism that enables this procedure. The strengths of this technique include the following.

(i) GW signals are included as part of the pulsar timing model. This means that all the issues relating to actual data sets (uneven observing cadence, different pulsars having different data spans and noise properties, the necessity to fit for pulsar parameters, etc.) are

modelled simultaneously with the GW search using standard pulsar timing techniques.

(ii) As part of the timing model, the covariances between the GW signal and the other timing model parameters are easily obtained.

(iii) The algorithms underlying our technique are fast, so they do not require large computing resources. The GW signal is condensed out of the raw observations into a much smaller data volume.

(iv) The technique makes no assumption about the actual form of the GW and so is useful for detecting unexpected GW signals as well as anticipated signals like BWMs or continuous waves (CWs) from SMBHBs in circular orbits. Inspection of the GW signal stream produced with our technique can help to determine what type of optimized signal detection schemes to implement.

In Section 2, we discuss how GWs manifest themselves in pulsar timing measurements and introduce our so-called \hat{A}_+ and \hat{A}_\times technique. In Section 3, we describe how \hat{A}_+ and \hat{A}_\times can be used in matched-filter searches for specific gravitational waveforms. In Section 4, we describe the PPTA Data Release 1 (DR1) and the simulated data sets we analyse for the remainder of the paper and indicate where these publicly available data can be accessed. In Section 5, using \hat{A}_+ and \hat{A}_\times , we conduct an all-sky search for detectable GW power from any direction with any time-domain behaviour. In Section 6, we study in greater detail the lines-of-sight to the Virgo, Fornax, Norma, Perseus, and Coma galaxy clusters, nearby massive clusters that are likely origins for the first detection by PTAs of an isolated GW source; for each of these special directions, we conduct a search for BWMs, clear indicators of the final merger of an SMBHB. In Section 7, we demonstrate that the correlations induced in PTA residuals from clock errors, inaccuracies in Solar system ephemerides, and GWs can be measured simultaneously without confusion. Finally, in Section 8, we summarize and conclude our work with some final remarks.

2 PLANAR GRAVITATIONAL WAVES AND PULSAR TIMING

A planar GW, $h_{ij}(t)$, coming from the direction \hat{n} , can be expressed as

$$h_{ij}(t) = a_+(t)\epsilon_{ij}^+ + a_\times(t)\epsilon_{ij}^\times, \quad (1)$$

where $a_+(t)$ and $a_\times(t)$ describe the time dependence of the two polarization modes of the GW and ϵ_{ij}^+ and ϵ_{ij}^\times are the relevant polarization tensors. Explicit representations of ϵ_{ij}^+ and ϵ_{ij}^\times in ecliptic longitude and latitude are given in Lee et al. (2011). For a PTA with N_p pulsars, the apparent pulsation frequency of the K^{th} pulsar in direction \hat{n}_K is influenced by the GW as

$$\frac{\Delta\nu_K}{\nu_K} = -\frac{1}{2} \frac{\hat{n}_K^i \hat{n}_K^j}{(1 - \cos\theta_K)} \times \left[h_{ij}(t)|_E - h_{ij} \left(t - \frac{D_K}{c} \right) |_{r_K} \right], \quad (2)$$

where $\cos\theta_K = \hat{n} \cdot \hat{n}_K$, the $|_E$ and $|_{r_K}$ notation indicates that the strain field is to be evaluated at the location of the Earth and the K^{th} pulsar, respectively, and D_K is the distance from the Earth to the pulsar (Estabrook & Wahlquist 1975; Lee et al. 2011). Despite the $(1 - \cos\theta_K)$ dependence in the denominator of equation (2), the fractional frequency shift does not diverge as θ_K approaches zero because the metric perturbation is transverse to the wave propagation direction, i.e. $\hat{n}_K^i h_{ij} \hat{n}_K^j \sim \sin^2\theta_K$. The perturbation to ToAs from the K^{th} pulsar caused by this GW, $\delta t_K^h(t)$, is given by the integral of

the fractional frequency change in equation (2):

$$\begin{aligned} \delta t_K^h(t) &= -\frac{1}{2} \frac{\hat{n}_K^i \hat{n}_K^j}{(1 - \cos \theta_K)} \times [\epsilon_{ij}^+ A_+(t) + \epsilon_{ij}^\times A_\times(t)], \\ &\equiv G_K^+ A_+(t) + G_K^\times A_\times(t), \end{aligned} \quad (3)$$

where

$$A_\star(t) = \int_0^t [a_\star(t')|_E - a_\star(t' - D_K/c)|_{r_K}] dt'. \quad (4)$$

The ‘ \star ’ subscript acts as a placeholder for either ‘+’ or ‘ \times ’. The factors G_K^\star depend on the relative angle between the GW source and the pulsar; they are maximized when the pulsar and GW source are in approximately the same direction of the sky and minimized when they are on opposite sides of the sky. This means that PTA sensitivity to point sources is best in the directions of the sky with the highest concentration of well-timed pulsars.

2.1 Timing model development and noise

We can express the n_K timing residuals for the K^{th} pulsar in vector form as

$$\delta \mathbf{t}_K = \mathbf{M}_K \delta \mathbf{p}_K + \delta \mathbf{t}_K^h + \delta \mathbf{t}_K^n. \quad (5)$$

The first term on the right side of equation (5) describes any structure in the residuals from inaccuracies in the timing model parameters. In this linearized approximation, $\delta \mathbf{p}_K$, a vector describing how much the m_K timing model parameters deviate from their true values, is assumed to be small. The design matrix, \mathbf{M}_K , is the $n_K \times m_K$ matrix describing how changes in the timing model parameters influence the residuals. The second term describes any structure in the residuals induced by the GW. The third term, $\delta \mathbf{t}_K^n$, describes any noise that influences the residuals. This may consist of white radiometer and pulse-phase-jitter noise, red spin noise, and a variety of additional sources (Cordes & Shannon 2010; Shannon & Cordes 2010). Additional correlated timing structure can be caused by things like inaccuracies in terrestrial time standards or errors in the position of the Solar system barycentre (SSB). We discuss these issues in Section 7. For our purposes here, these effects can be subsumed into the noise term of equation (5), $\delta \mathbf{t}_K^n$.

If we temporarily neglect the influence of GWs on the timing residuals, $\delta \mathbf{t}_K^h$, with the noise covariance matrix $\mathbf{C}_K = \langle (\delta \mathbf{t}_K^n) \cdot (\delta \mathbf{t}_K^n)^T \rangle$, we can estimate the maximum-likelihood corrections to the timing model parameters, $\delta \hat{\mathbf{p}}_K$, and the parameter covariance matrix, \mathbf{C}_K^P (Gregory 2010):

$$\mathbf{C}_K^P = (\mathbf{M}_K^T \mathbf{C}_K^{-1} \mathbf{M}_K)^{-1}, \quad \text{and} \quad (6)$$

$$\delta \hat{\mathbf{p}}_K = \mathbf{C}_K^P \mathbf{M}_K^T \mathbf{C}_K^{-1} \delta \mathbf{t}_K. \quad (7)$$

The use of a general noise model, first utilized by the pulsar timing community in Coles et al. (2011), is now commonly used for iteratively refining timing models as additional ToAs are acquired. Coles et al. (2011) demonstrated spectral methods for adequately estimating the noise covariance matrix and showed that properly modelling the noise, especially the highly temporally correlated red spin noise seen in some pulsars, is crucial for mitigating biases in timing model parameter estimation. However, if the residuals are being influenced by GWs, failing to account for them can lead to improper noise modelling and biased parameter estimation. With the many PTA data sets currently available, the influence of GWs in the vicinity of Earth, which causes correlated residual structure across all timing data sets with a distinct quadrupolar pattern, can

be disentangled from noise processes specific to single pulsars or errors in individual timing models.

As equation (4) indicates, A_\star is an integral of the difference in two terms: an ‘Earth term’ common to all pulsar timing data sets and a ‘pulsar term’ that differs between data sets owing to the different positions and distances of the pulsars. For burst GWs with durations much shorter than the light travel times between Earth and the pulsars in a PTA (thousands of years, typically), when the Earth term is active the various pulsar terms are all quiescent and the pulsar terms are negligible. However, the GWs from a SMBHB will be in the band of frequencies potentially detectable by PTAs for the order of millions of years. If we treat a SMBHB as a monochromatic source of sinusoidal GWs (i.e. ignore the secular frequency evolution of the source over thousand-year time-scales), then each pulsar term will be a sinusoid of the same frequency as the Earth term, a nearly identical amplitude, and an essentially random and uniformly distributed phase. The pulsar term can perfectly add to the Earth term, perfectly cancel it, or anything in-between. The pulsar terms are uncorrelated in different pulsar timing data sets. If N_P pulsars with approximately equal timing precision are analysed and their residuals coherently combined, the self-noise from these pulsar terms is suppressed by $N_P^{1/2}$ (Hellings & Downs 1983). For discussion of frequency evolution anticipated from SMBHBs and prospects for including pulsar terms in the signal model rather than treating them as sources of noise, see e.g. Arzoumanian et al. (2014) and Wang, Mohanty & Jenet (2014).

2.2 Building GWs into timing models

We model $A_+(t)$ and $A_\times(t)$ as linearly interpolated functions on a grid of N_τ times τ_μ . In principle, different grids could be used for the + and \times polarizations, but for simplicity, we will use an identical grid for each of them. The grid does not need to be evenly spaced; for some pulsar timing data sets where the ToA sampling is sparse in decades-old data but relatively uniform and dense in more modern data, variable spacing in the interpolation grid may be appropriate. An advantage to modelling $A_+(t)$ and $A_\times(t)$ as linear interpolants rather than a higher-order polynomial or Fourier series is that individual bad ToAs or small time-spans of bad ToAs, either with large uncertainties or apparent biases, mainly influence the estimates of $A_+(t)$ and $A_\times(t)$ locally rather than over larger spans of data or the whole data set.

These models for A_+ and A_\times , which we refer to as \mathcal{A}_+ and \mathcal{A}_\times , are incorporated into a global fit to the full set of ToAs that is carried out simultaneously with all the timing models for all the pulsars in the array. The $\mathcal{A}_+(t)$ and $\mathcal{A}_\times(t)$ time series must be constrained so that they are not covariant with the astrometric or spin parameters of the individual pulsar timing models. The products of this fitting procedure are $\hat{\mathcal{A}}_+$ and $\hat{\mathcal{A}}_\times$, the maximum-likelihood estimators of \mathcal{A}_+ and \mathcal{A}_\times , and a matrix, $\mathbf{C}_{+\times}$, describing covariances in $\hat{\mathcal{A}}_+$ and $\hat{\mathcal{A}}_\times$.

The $\hat{\mathcal{A}}_+$ and $\hat{\mathcal{A}}_\times$ time series are a complete representation of the GWs coming from a particular direction in that they carry all the information with respect to such GWs that is contained in the ToAs. One cannot obtain a better signal estimator using the ToAs than can be done with the auxiliary time series alone. Full details of the implementation of the global fit and the constraints that need to be applied to $\mathcal{A}_+(t)$ and $\mathcal{A}_\times(t)$ during the fitting procedure are given in Appendix A. These algorithms have been implemented in the TEMPO2 software package (Edwards et al. 2006, see appendix B for usage details).

3 SEARCHING \hat{A}_+ AND \hat{A}_\times FOR ANTICIPATED GW WAVEFORMS

Although one of the great strengths of the \hat{A}_+ and \hat{A}_\times time series is that it can be used to study GWs without an underlying signal model, we demonstrate here how searches for specific waveforms can be implemented directly in the \hat{A}_+ and \hat{A}_\times time series after they have been computed. Zhu et al. (2014) developed a procedure for conducting a matched-filter search for a specific waveform in \hat{A}_+ and \hat{A}_\times and applied it to DR1 from the PPTA in a search for CWs. The upper limit on the amplitude of CWs derived in that paper is among the most stringent achieved to date (for competitive limits, see Arzoumanian et al. 2014; Babak et al. 2015). Here, we will review the essential elements needed to carry out a matched-filter search in \hat{A}_+ and \hat{A}_\times and elaborate on certain points, but will focus on a different, simpler signal model: a BWM. We will apply these methods in a small targeted search for BWMs from nearby massive galaxy clusters using DR1 later in this paper. Wang et al. (2015) searched DR1 for BWMs but did not utilize \hat{A}_+ and \hat{A}_\times time series. Our search is different and complementary to the work described in that paper; \hat{A}_+ and \hat{A}_\times techniques produce consistent results with Wang et al. (2015) and are ultimately more computationally efficient.

In conducting the global fit that generates \hat{A}_+ and \hat{A}_\times , we must compute a global parameter covariance matrix (see Appendix A for a detailed discussion). A $2N_\tau \times 2N_\tau$ sub-block of this matrix describes the correlated uncertainties in \hat{A}_+ and \hat{A}_\times ; we call this sub-block $C_{+\times}$. As discussed in Zhu et al. (2014), $C_{+\times}$ is not a full-rank matrix because of the constraints applied to \hat{A}_+ and \hat{A}_\times and is thus non-invertible. We define $\tilde{C}_{+\times} = \mathbf{E}\mathbf{F}\mathbf{E}^T$ where \mathbf{E} is a matrix containing the eigenvectors of $C_{+\times}$ associated with non-null eigenvalues and \mathbf{F} is a square diagonal matrix containing those non-null eigenvalues. This generalized covariance matrix is the covariance matrix on the subspace orthogonal to the constraints and is manifestly invertible: $\tilde{C}_{+\times}^{-1} = \mathbf{E}\mathbf{F}^{-1}\mathbf{E}^T$. Also, we define the projection operator $\mathbf{P} = \mathbf{E}\mathbf{E}^T$.

We define a vector $\hat{A}^T = [\hat{A}_+^T, \hat{A}_\times^T]$, i.e. \hat{A}_+ and \hat{A}_\times stacked into a single vector. If we assume that the structure in \hat{A} is due to a BWM occurring at t_{inj} , if constraints had not been applied, \hat{A} could be expressed as a linear combination of two basis elements:

$$\beta_+ = \begin{bmatrix} (\boldsymbol{\tau} - t_{\text{inj}})\Theta(\boldsymbol{\tau} - t_{\text{inj}}) \\ 0 \cdot \boldsymbol{\tau} \end{bmatrix}, \quad (8)$$

$$\beta_\times = \begin{bmatrix} 0 \cdot \boldsymbol{\tau} \\ (\boldsymbol{\tau} - t_{\text{inj}})\Theta(\boldsymbol{\tau} - t_{\text{inj}}) \end{bmatrix}, \quad (9)$$

where $\boldsymbol{\tau}$ is the vector of N_τ times at which the \hat{A}_+ and \hat{A}_\times time series are sampled and Θ is the Heaviside step function. This signal model for a BWM is discussed by, e.g. Pshirkov, Baskaran & Postnov (2010), van Haasteren & Levin (2010), Cordes & Jenet (2012), and Madison, Cordes & Chatterjee (2014). The appropriately constrained version of \hat{A} can be expressed as a linear combination of these basis elements after they have been projected into the subspace orthogonal to the constraints, i.e. as a linear combination of $\tilde{\beta}_+ = \mathbf{P}\beta_+$ and $\tilde{\beta}_\times = \mathbf{P}\beta_\times$. We can say

$$\hat{A} = \alpha_+ \tilde{\beta}_+ + \alpha_\times \tilde{\beta}_\times, \quad (10)$$

and we can compute maximum-likelihood estimates of the coefficients, $\hat{\alpha}_+$ and $\hat{\alpha}_\times$, as follows:

$$\hat{\alpha} = \left(\mathbf{B}^T \tilde{C}_{+\times}^{-1} \mathbf{B} \right)^{-1} \mathbf{B}^T \tilde{C}_{+\times}^{-1} \hat{A}, \quad (11)$$

where $\alpha^T = [\alpha_+, \alpha_\times]$ and $\mathbf{B} = [\tilde{\beta}_+, \tilde{\beta}_\times]$. The uncertainties on the estimates $\hat{\alpha}_+$ and $\hat{\alpha}_\times$ are encoded in the 2×2 matrix

$$\boldsymbol{\Sigma} = \left(\mathbf{B}^T \tilde{C}_{+\times}^{-1} \mathbf{B} \right)^{-1}. \quad (12)$$

This signal parameter estimator, $\hat{\alpha}$, is biased. The bias is a consequence of the constraints applied to \hat{A}_+ and \hat{A}_\times . As an extreme example, if a GW signal was purely quadratic, that signal would be entirely filtered out of \hat{A}_+ and \hat{A}_\times by the constraints and the GW signal could never be reconstructed. Blandford, Narayan & Romani (1984) and Madison, Chatterjee & Cordes (2013) describe the procedure of fitting certain functional forms out of timing residuals as equivalent to applying a frequency-domain filter to the residual fluctuation spectrum. Analogously, the constraints applied to \hat{A}_+ and \hat{A}_\times destructively filter out some features of the GW signal. However, the spectral shape of constraint-induced filters depends upon the time-span of the data set. As more data are gathered, the constraints and the GW signal become less covariant, less of the GW signal is filtered out, and the bias in estimation of the signal parameters is reduced. Furthermore, through simulations, the bias can be quantified, allowing for robust and trustworthy parameter estimation.

4 REAL AND SIMULATED DATA SETS

For much of the remainder of this paper, we work with both simulated and real data sets. Our simulated data sets were produced using a newly developed software package called PTASIMULATE.¹ The simulated data sets we analyse are as follows.²

(i) SIM1: 20 synthetic pulsars are generated with random sky positions (statistically uniform on the sphere). We assume weekly ToA measurements with 50 ns timing precision. The timing models for each pulsar only include fits for three spin parameters (rotational phase, frequency, and frequency derivative). The timing residuals for each pulsar are consistent with 50 ns rms white Gaussian noise.

(ii) SIM2: as in SIM1, but a bright GW burst has been injected into both GW polarization channels of the simulation from a source at right ascension 0 h and declination 0°. The pulsar locations differ from those in SIM1, but are again drawn from a distribution that is uniform on the sphere. The white Gaussian noise added to each ToA again has an rms of 50 ns, but we use a different realization of noise from that in SIM1.

(iii) SIM3: we have generated 20 simulated pulsars with the same positions as the 20 pulsars comprising the PPTA DR1. We have assumed these pulsars are observed every two weeks and have ToA uncertainties consistent with 100 ns rms Gaussian white noise. In each pulsar timing model, we fit only for rotational phase, frequency, and frequency derivative. We have simulated a clock error by generating the simulated data with TAI and reprocessing them with BIPM2013, two different realizations of terrestrial time. We have simulated an ephemeris error by using Jet Propulsion Laboratory ephemeris DE414 (Standish 2006) to generate the simulated data and DE421 (Folkner, Williams & Boggs 2008) to reprocess it.

(iv) SIM4: as in SIM3, but a GW burst has also been injected into the simulated timing data. The realization of 100 ns rms Gaussian white noise in SIM4 differs from that in SIM3.

¹ PTASIMULATE and documentation for it can be accessed at (<https://bitbucket.org/psrsoft/ptasimulate>).

² All the data sets (real and simulated) used in this paper are available from <http://dx.doi.org/10.4225/08/560A00E2036F6>

We analyse a version of the PPTA DR1 (Manchester et al. 2013) that has been updated to include detailed noise models developed and used by Wang et al. (2015) and Zhu et al. (2014). The noise models account for additional white noise in the data beyond that anticipated from radiometer noise alone (jitter is a likely contribution to this) and, using the techniques developed in Coles et al. (2011), correlated red timing noise.

5 ALL-SKY SEARCHES FOR EXCESS GW POWER

Here, we make use of the \hat{A}_+ and \hat{A}_\times time series to search for GW power coming from any direction on the sky. This search would detect sufficiently strong signals of any waveform. As a detection statistic for this search, we define a dimensionless measure of the total GW power contained in \hat{A}_+ and \hat{A}_\times :

$$D = \hat{A}^T \tilde{C}_{+\times}^{-1} \hat{A}. \quad (13)$$

This is the mean squared error statistic for the null hypothesis. Using slightly different terminology, $D = \hat{A}_W^T \hat{A}_W$ where $\hat{A}_W = U^{-1} \hat{A}$ and U is a Cholesky whitening matrix, i.e. $\tilde{C}_{+\times} = U U^T$ (Coles et al. 2011). In the absence of signal, assuming the noise is correctly modelled by $\tilde{C}_{+\times}$ and thus U , the whitened data stream, \hat{A}_W , will consist of Gaussian noise, be unit variance, and white; D will follow χ^2 statistics with $N_{\text{dof}} = 2(N_\tau - N_c)$ degrees of freedom, where N_c is the number of constraints applied individually to the \hat{A}_+ and \hat{A}_\times time series and, again, N_τ is the number of grid points at which \hat{A}_+ and \hat{A}_\times are sampled. If \hat{A}_W is not unit variance and white, either GW signal is present in the data or the individual pulsar noise models are inaccurate. The noise models for the data we use here have been vetted by the analyses of Zhu et al. (2014) and Wang et al. (2015).

We compute \hat{A}_+ and \hat{A}_\times over a grid of trial source locations on the sky and compute D at each grid point. The number of statistically independent samples on this grid of sky positions, N_s , influences our anticipated false alarm probability. The false alarm probability for a given value of D is

$$f(D) = 1 - c(D; N_{\text{dof}})^{N_s}, \quad (14)$$

where $c(D; N_{\text{dof}})$ is the cumulative distribution for a χ^2 distribution with N_{dof} degrees of freedom evaluated at D . Cornish & van Haasteren (2014) recently demonstrated that the response of a PTA to any type of GW can be expressed as a linear combination of N_p orthogonal modes, or sky maps, where, again, N_p is the number of pulsars in the array. A similar result was derived by Gair et al. (2014). Since our actual data consist of observations of 20 pulsars, we will use the result from these authors and conservatively set $N_s = 20$.

In Fig. 1, we display the results of this search procedure when applied to SIM1, SIM2, and DR1 (listed in order from top to bottom). We have sampled \hat{A}_+ and \hat{A}_\times at 74 evenly spaced epochs, corresponding to a cadence of 30 d. The left panels indicate the value of D at each sky position. The black dot in these left panels indicates the position on the sky that yielded the greatest value of D . The right panels depict the \hat{A}_+ and \hat{A}_\times time series associated with the direction on the sky that yielded this greatest value of the detection statistic. For the analysis of SIM1, our null example, \hat{A}_+ and \hat{A}_\times appears by eye to be consistent with zero signal and, indeed, the maximum value of D that is realized is 174.8, consistent with approximately 18.7 per cent of realizations of noise.

Our analysis of SIM2 is shown in the middle row of Fig. 1. The position of the GW source was set to a right ascension of 0 h

and a declination of 0° . Notice the change in the colour scales between the top left and middle left panels. In the middle row, the detection statistic at nearly every trial sky position is inconsistent with noise, but D is peaked around the location of the simulated source demonstrating the efficacy of this procedure in localizing a bright GW source on the sky. The \hat{A}_+ and \hat{A}_\times time series produced at the position of the maximal detection statistic reproduces the injected GW signal,

$$b(t) = b_0(t - t_*)e^{-(t-t_*)^2/2w^2}, \quad (15)$$

where $b_0 = 3 \times 10^{-9}$, $t_* = \text{MJD } 54500$, and $w = 75$ d. For simplicity, we injected the same functional form into each polarization channel. This burst, at maximum amplitude, produces an approximately 136 ns ToA perturbation. In our analysis of SIM1 and SIM2, we have shown that with idealized data, our search method correctly returns a null result in the absence of a GW signal, but does correctly detect, reproduce, and localize on the sky, a bright signal that is present.

The bottom row of plots in Fig. 1 depicts the results of our search for directional GWs in the PPTA's DR1. The search is carried out in the same way as for the simulated data sets, but we have used a generalized least-squares fitting routine (as described in Coles et al. 2011) in order to account for the different red noise present in different pulsars.³ The white stars in the bottom left panel indicate the positions of the 20 pulsars comprising DR1. The distribution of pulsars on the sky does not approach statistical uniformity and the noise properties of the pulsars are not all equal as in SIM1 and SIM2. With $N_\tau = 74$, $N_c = 7$ (as discussed in Appendix A), and $N_s = 20$, D would have to exceed approximately 194 to claim an excess of GW power with 99 per cent confidence or 184 for 95 per cent confidence. The maximum value of D we find in our analysis of DR1 is 188.5, which is between the 95 per cent and 99 per cent confidence thresholds.

Even though this value of D is consistent with a null detection, it is close to values of interest. However, there is a reason to be sceptical of this marginally high value of D . The \hat{A}_+ time series associated with this maximal value of D from DR1 has a 4.2σ outlier at an epoch near September of 2008. This epoch is contemporaneous with the commissioning of a new pulsar timing backend at Parkes. Many ToAs from this brief era have been excluded from DR1 for displaying obvious systematic issues, but it is likely that systematic artefacts may still be present in the data. If this one outlier in the \hat{A}_+ time series is artificially forced towards zero until its 1σ error bar is consistent with zero, the value of D from this pointing is reduced to approximately 171, consistent with approximately 30 per cent of noise realizations.

This result with DR1 underscores the importance of the IPTA project in two ways. First, if the marginally significant result we have discussed is indeed just a systematic artefact in PPTA data, comparison with measurements from the EPTA and NANOGrav from around September of 2008 of a subset of pulsars that overlaps with the PPTA sample should be able to resolve the issue. Secondly, if the outlier in the PPTA data set is a very short duration, very bright GW burst, it will be present in the data sets of all three PTAs comprising the IPTA and a joint analysis of the combined data sets will bolster the significance of the result.

³ In contrast to searches for a GW background, the signals of interest here have significantly different spectral characteristics than the intrinsic red noise often observed in pulsar data sets.




Figure 1. Left: the total power detection statistic, D , as it varies over the sky. The black circle indicates the position on the sky that yielded the greatest value of D . Right: the \hat{A}_+ and \hat{A}_\times time series for the direction of the sky that yielded the greatest value of D . Top: analysis of SM1. The detection statistic never exceeds values that are inconsistent with noise. Middle: analysis of SM2. The detection statistic clearly indicates that the data are inconsistent with the noise, we are able to localize the location of the injected GW source, and we can reconstruct the waveform of the injected burst. Bottom: analysis of the six-year DR1 from the PPTA. The stars indicate the positions of the 20 pulsars included in DR1. The coloured diamonds indicate the positions of five nearby massive galaxy clusters: Virgo (blue), Fornax (green), Norma (magenta), Perseus (cyan), and Coma (yellow). There is no significant evidence for GW power in DR1.

6 TARGETED INVESTIGATIONS OF NEARBY MASSIVE GALAXY CLUSTERS

Simon et al. (2014) recently conducted a detailed analysis of surveys of local galaxies in order to identify potential directions on the sky from which an initial detection of resolvable GWs by PTAs is likely to originate, so-called ‘GW hotspots’. They singled out, in order of increasing distance from Earth, several massive galaxy clusters: Virgo, Fornax, Norma, Perseus, and Coma. The directions to these clusters are indicated in the bottom left panel of Fig. 1 with coloured diamonds. These clusters are all within 100 Mpc of Earth and all but Fornax, the smallest of the five, contain upwards of 500 galaxies. In Fig. 2, we plot the \hat{A}_+ and \hat{A}_\times time series produced when DR1 is phased up to the directions of each of these galaxy clusters. All are consistent with noise.

There are several notable features in Fig. 2. First, all five pointings have some marginal outliers near the 2008 September commissioning of the new pulsar timing backend at Parkes that was discussed earlier. We demonstrated in our analysis of SM2 (depicted in Fig. 2) that a GW burst can be localized on the sky and the five clusters we consider here are from widely separated directions; this gives

further weight to the argument that there are systematic issues with DR1 near September of 2008 and it is very unlikely that the excess power in \hat{A}_+ and \hat{A}_\times is from a spatially localized GW source. Secondly, the scatter in all pointings is larger earlier in the time-span; this is due largely to an increase in the observing cadence over time and thus a greater number of ToAs contributing to each \hat{A}_+ and \hat{A}_\times sample later in the data set. The change in the observing cadence has not led to any issues in our analysis of DR1, but this illustrates the possible need for unequal spacing in the \hat{A}_+ and \hat{A}_\times sample grid in other data sets. Finally, the scatter in the \hat{A}_+ and \hat{A}_\times time series is biggest in our pointing towards the Perseus cluster and smallest in our pointing towards the Norma cluster. The Perseus cluster is in the opposite direction of the sky from the peak concentration of PPTA pulsars while the Norma cluster is very nearly in the same direction as many of the PPTA pulsars. This is a reflection of the point we discussed in Section 2 that pulsar timing measurements are more sensitive to GWs coming from directions of the sky near the direction of the pulsar. It is well known that an anisotropic distribution of well-timed pulsars leads to anisotropic sensitivity to GWs; this again highlights the importance of the IPTA and close




Figure 2. The \hat{A}_+ (black) and \hat{A}_\times (red) time series produced when DR1 is phased up to the locations of the massive Virgo, Fornax, Norma, Perseus, and Coma galaxy clusters. All pointings are consistent with noise only (as indicated by the D value quoted in each panel). The diamonds in each panel are colour coded with the diamonds in the bottom left panel of Fig. 1 indicating the directions to these clusters.

collaboration and data sharing between pulsar astronomers in the Northern and Southern Hemispheres.

6.1 A BWM Search in \hat{A}_+ and \hat{A}_\times

Although we find no evidence for excess GW power in the directions of any of these five clusters, to search for a specific waveform, a matched-filter search of \hat{A}_+ and \hat{A}_\times is more sensitive than a total power search. Building on our discussion from Section 3, as a test case, we here conduct a search for BWMs in these five pointings.

We define a different detection statistic for this search: $D_B = \hat{\alpha}^T \Sigma^{-1} \hat{\alpha}$. This is nearly identical to the detection statistic used in the BWM search conducted by Wang et al. (2015), but here we do our search directly in \hat{A}_+ and \hat{A}_\times . In the absence of signal, D_B will follow χ^2 statistics with 2 degrees of freedom. We will search epochs within the innermost 80 per cent of each of the five \hat{A}_+ and \hat{A}_\times pointings in Fig. 2 for evidence of a BWM occurring. We restrict ourselves to this window because detecting a BWM requires the ability to accurately assess the pulsar timing behaviour both pre- and post-burst. Arzoumanian et al. (2015a) recently showed that in a BWM search over many trial burst epochs, there are approximately five statistically independent trials; this fact must be accounted for in assessing the false alarm probability in our search. In order for the data to be inconsistent with 95 per cent of realizations of noise, with $N_s = 5$, D_B must exceed approximately 9.2.

In the top panel of Fig. 3, we show the values of D_B derived from our search; we find nothing inconsistent with noise. Knowing Σ , however, allows us to compute the minimum value of $|\alpha|$ needed to exceed the 95 per cent confidence threshold in D_B . We display this quantity, which we call $h_{B,\min}$, in the middle panel of Fig. 3. The strain amplitude of a BWM is $h_B \approx 1.5 \times 10^{-13} (\mu/10^9 M_\odot)(d/10\text{Mpc})^{-1}$ where μ is the reduced mass of the binary and d is the luminosity distance between the binary and Earth; we have assumed the binary has a typical inclination angle of $\pi/3$ (Madison et al. 2014). Taking the luminosity distances to Virgo, Fornax, Norma, Perseus, and Coma, as 17 Mpc, 19 Mpc, 68 Mpc, 74 Mpc, and 99 Mpc, respectively (see Simon et al. 2014, and references therein), we have generated the bottom panel of

Fig. 3 displaying the minimal reduced mass, μ_{\min} , of a merging SMBHB that would have produced a BWM bright enough to exceed our 95 per cent confidence threshold on D_B . The likelihood of such a massive merger occurring in one of these galaxy clusters during our observing span is exceedingly small, but as the span of our data set grows, our sensitivity to BWMs will improve and we will become sensitive to less massive mergers.

The \hat{A}_+ and \hat{A}_\times time series associated with the directions towards these five galaxy clusters can be accessed along with DR1 and our simulated data sets following the link mentioned above. Instructions by which results for any other pointing can be quickly produced with TEMPO2 can be found in the usage details of Appendix B. Following the BWM example detailed here and the analysis of Zhu et al. (2014), analogous searches for any type of parametrized waveform can be easily carried out.

7 MULTIPOLE MENAGERIE: CLOCK ERRORS, INACCURATE EPHEMERIDES, AND GRAVITATIONAL WAVES

As we mentioned in Section 1, GWs are not the sole means by which timing residuals from all the pulsars in a PTA can become correlated. Faults in terrestrial time standards can induce monopolar correlations between pulsar timing data sets (Hobbs et al. 2012). Errors in our estimated position for the SSB from, for example, inaccuracies in the measured mass of Saturn can lead to dipolar correlations between pulsar timing residuals (Champion et al. 2010). In order for \hat{A}_+ and \hat{A}_\times time series to be useful in studies of GW-induced inter-pulsar timing correlations, they must be able to reliably differentiate the distinctly quadrupolar signature of a GW from the monopolar or dipolar signatures of these other effects; in Fig. 4, we demonstrate that our techniques can effectively do this.

The left panels of Fig. 4 depict our analysis of SIM3 in which clock errors and ephemeris errors have been simulated but there is no GW signal present. We simultaneously fit for the timing models of all the pulsars in our simulated array along with linear-interpolant models for clock errors (as in Hobbs et al. 2012), vectorial offsets between the true SSB and its assumed location (as in Champion




Figure 3. Top: the detection statistic, D_B , derived from a search for BWMs in the \hat{A}_+ and \hat{A}_\times time series associated with the directions to five nearby massive galaxy clusters. D_B needs to exceed approximately 9.2 in order to be inconsistent with more than 95 per cent of noise realizations. Middle: the minimum amplitude of a BWM that would have yielded a value of D_B exceeding 9.2. Bottom: assuming a statistically average inclination angle for a merging SMBHB, based on the luminosity distance to each of the five galaxy clusters we investigate, the minimum reduced mass of a merging SMBHB that would have produced a detectable BWM.

et al. 2010), and the \hat{A}_+ and \hat{A}_\times time series for the direction on the sky from which we eventually inject a GW burst in SIM4. From top to bottom, we display the clock signal (we call this ΔT), the three Cartesian components of the simulated vector displacement in the assumed SSB (called ΔX , ΔY , and ΔZ), and the \hat{A}_+ and \hat{A}_\times time series. The red lines indicate the injected signal with the best-fitting (in a least-squares sense) quadratic removed; the quadratic removal is necessary because of the quadratic fit carried out for each pulsar. The black dots indicate the recovered signal. We are able to accurately recover the clock and ephemeris-error signals present in the data and the \hat{A}_+ and \hat{A}_\times are consistent with a null result. Significant monopolar and dipolar correlations present in the data have not produced spurious signal in the \hat{A}_+ and \hat{A}_\times channels.

In the right panels of Fig. 4 we have carried out the same procedure described above but applied to SIM4 which contains a simulated GW burst from RA = 0.2 rad and Dec. = -0.3 rad, a place on the sky that is remote from the core concentration of PPTA pulsars. The waveforms injected into the two GW polarization channels are

$$F_*(t) = \mathcal{F}_* e^{-(t-t_*)^2/2w_*^2}, \quad (16)$$

where $\mathcal{F}_+ = 10^{-5}$, $\mathcal{F}_\times = -8 \times 10^{-6}$, $t_+ = \text{MJD } 50500$, $t_\times = \text{MJD } 50600$, $w_+ = 224$ d, and $w_\times = 387$ d. The injected clock and ephemeris error signals are identical between SIM3 and SIM4. The injected GW signal is successfully recovered, and even in the presence of a bright GW signal, the clock- and ephemeris-error signals are also successfully recovered.

The examples depicted in Fig. 4 illustrate the ability to successfully differentiate the correlations induced in pulsar timing data sets from clock errors, ephemeris errors, and GWs and to recover the relevant signals without substantial bias. However, suppose someone was only interested in searching for a GW signal. If a clock or ephemeris error were leading to correlations across multiple pulsar timing data sets and only \hat{A}_+ and \hat{A}_\times were fitted, the unaccounted-for correlations from the clock or ephemeris error would bias the GW signal-recovery and potentially lead to spurious claims of GW detection (for related discussion, see Tiburzi et al. 2015). We emphasize the importance of accounting for potential clock- or ephemeris-error correlations in any PTA effort to detect GWs.

8 CONCLUSIONS

In this paper, we have described a technique that allows the signal of a planar GW from a particular direction of the sky to be distilled from PTA data. This plane wave could be from a single distant source, such as an individual SMBHB, or from a collection of distant sources that are sufficiently close together on the sky like a massive galaxy cluster; how close together sources need to be depends on the angular resolution of the PTA which depends on the spatial distribution of well-timed pulsars. The angular resolution of PTAs will improve over time as timing techniques are improved and as new pulsars are discovered and incorporated into PTAs.

We have assumed that GWs behave according to the predictions of linearized general relativity. Some alternative theories of gravity allow for additional polarization modes of GWs such as longitudinal or breathing modes (Lee 2013); our formalism cannot currently address these alternative theories of gravity, but could with modification. With these caveats regarding the polarization of GWs in mind, since our methods require no assumptions about the time-domain behaviour of GWs, they provide a general and flexible tool that will help maximize the discovery potential of PTAs, potentially facilitating the discovery of wholly unanticipated sources.




Figure 4. Simultaneous extraction from a 20-pulsar simulated PTA of the monopolar signature of clock errors (top panel), the dipolar signature of an inaccurate Solar system ephemeris (the second, third, and fourth panels from the top), and the two polarization modes of a GW (the bottom two panels). The left panels, depicting an analysis of Sim3, contain no GW signal, while the right panels, depicting an analysis of Sim4, contain the bright GW burst described in equation (16).

Most modern efforts to detect GWs with PTAs utilize sophisticated Bayesian inference codes (van Haasteren et al. 2011; Vigeland & Vallisneri 2014; Arzoumanian et al. 2014; Lentati et al. 2015). Our total-power all-sky GW search in Section 5 and our targeted search for BWMs in Section 6 are frequentist in nature, but as \hat{A}_+ and \hat{A}_x are built into the PTA timing model, they can be straightforwardly incorporated into Bayesian analyses. We feel that the \hat{A}_+ and \hat{A}_x method will complement Bayesian GW searches by providing important diagnostic time series that will allow by-eye examination of GW signals. Furthermore, \hat{A}_+ and \hat{A}_x time series facilitate a significant condensation of the data volume necessary for GW investigations, and as such, can potentially reduce the computational resources required for even very advanced Bayesian search codes.

Finally, we point out that modern PTA data sets are constantly growing and evolving. As technology advances and allows for upgrades in radio instrumentation and as more nations and observatories become major contributors to PTA endeavours, the data come in a wider variety of forms and from more varied sources. PTA data are becoming more and more heterogeneous and unwieldy to work with. For the purposes of GW investigations, \hat{A}_+ and \hat{A}_x time series will aid in the creation of a compressed and homogenized auxiliary data set that isolates all of the important GW information.

ACKNOWLEDGEMENTS

For this work, we have used data from the Parkes radio telescope. The Parkes radio telescope is part of the Australia Telescope, which is funded by the Commonwealth of Australia for operation as a National Facility managed by the Commonwealth Scientific and Industrial Research Organization (CSIRO). DRM is a Jansky Fellow of the National Radio Astronomy Observatory (NRAO). NRAO is a facility of the National Science Foundation (NSF) operated under cooperative agreement by Associated Universities, Inc. DRM also acknowledges support from a sub-award to Cornell University from West Virginia University through NSF/PIRE grant 0968296, the NANOGrav Physics Frontiers Center NSF award PHY-1430284, and the New York NASA Space Grant Consortium. X-JZ and

LW acknowledge funding support from the Australian Research Council. GH is supported by Australian Research Council grant FT120100595. JW is supported by NSFC project No. 11403086 and the West Light Foundation CAS XBBS201322. SO is supported by the Alexander von Humboldt Foundation.

REFERENCES

- Arzoumanian Z. et al., 2014, *ApJ*, 794, 141
 Arzoumanian Z. et al., 2015a, *ApJ*, 810, 150
 Arzoumanian Z. et al., 2015b, preprint (arXiv:1508.03024)
 Babak S. et al., 2015, *MNRAS*, in press, preprint (arXiv:1509.02165)
 Blandford R., Narayan R., Romani R. W., 1984, *JA&A*, 5, 369
 Champion D. J. et al., 2010, *ApJ*, 720, L201
 Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, *MNRAS*, 418, 561
 Cordes J. M., Jenet F. A., 2012, *ApJ*, 752, 54
 Cordes J. M., Shannon R. M., 2010, preprint (arXiv:1010.3785)
 Cornish N. J., van Haasteren R., 2014, preprint (arXiv:1406.4511)
 Cutler C., Burke-Spolaor S., Vallisneri M., Lazio J., Majid W., 2014, *Phys. Rev. D*, 89, 042003
 Demorest P. B. et al., 2013, *ApJ*, 762, 94
 Detweiler S., 1979, *ApJ*, 234, 1100
 Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
 Estabrook F. B., Wahlquist H. D., 1975, *Gen. Relativ. Gravit.*, 6, 439
 Folkner W. M., Williams J. G., Boggs D. H., 2008, *The Planetary and Lunar Ephemeris DE421*. JPL Interoffice Memorandum IOM 343R-08-003, California
 Foster R. S., Backer D. C., 1990, *ApJ*, 361, 300
 Gair J., Romano J. D., Taylor S., Mingarelli C. M. F., 2014, *Phys. Rev. D*, 90, 082001
 Gregory P., 2010, *Bayesian Logical Data Analysis for the Physical Sciences*. Cambridge University Press, Cambridge
 Hellings R. W., Downs G. S., 1983, *ApJ*, 265, L39
 Hobbs G. et al., 2010, *Class. Quant. Grav.*, 27, 084013
 Hobbs G. et al., 2012, *MNRAS*, 427, 2780
 Huerta E. A., McWilliams S. T., Gair J. R., Taylor S. R., 2015, *Phys. Rev. D*, 92, 063010
 Keith M. J. et al., 2013, *MNRAS*, 429, 2161
 Kramer M., Champion D. J., 2013, *Class. Quant. Grav.*, 30, 224009

- Lee K. J., 2013, *Class. Quant. Grav.*, 30, 224016
 Lee K. J., Wex N., Kramer M., Stappers B. W., Bassa C. G., Janssen G. H., Karuppusamy R., Smits R., 2011, *MNRAS*, 414, 3251
 Lentati L. et al., 2015, *MNRAS*, 453, 2576
 McLaughlin M. A., 2013, *Class. Quant. Grav.*, 30, 224008
 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., 2013, *Class. Quant. Grav.*, 30, 224010
 Manchester R. N. et al., 2013, *PASA*, 30, 17
 Milosavljević M., Merritt D., 2003, *ApJ*, 596, 860
 Pshirkov M. S., Baskaran D., Postnov K. A., 2010, *MNRAS*, 402, 417
 Ravi V., Wyithe J. S. B., Shannon R. M., Hobbs G., Manchester R. N., 2014, *MNRAS*, 442, 56
 Sesana A., Vecchio A., 2010, *Class. Quant. Grav.*, 27, 084016
 Shannon R. M., Cordes J. M., 2010, *ApJ*, 725, 1607
 Shannon R. M. et al., 2013, *Science*, 342, 334
 Shannon R. M. et al., 2015, *Science*, 349, 1522
 Christy B., Simon J., Polin A., Lommen A., Stappers B., Finn L. S., Jenet F. A., 2014, *ApJ*, 784, 60
 Standish E. M., Jr, 2006, *JPL Planetary Ephemeris DE414*. JPL Interoffice Memorandum IOM 343.R-06-002, California
 Taylor S. R. et al., 2015, *Phys. Rev. Lett.*, 115, 041101
 Tiburzi C. et al., 2015, *MNRAS*, in press, preprint ([arXiv:1510.02363](https://arxiv.org/abs/1510.02363))
 van Haasteren R., Levin Y., 2010, *MNRAS*, 401, 2372
 van Haasteren R. et al., 2011, *MNRAS*, 414, 3117
 Vigeland S. J., Vallisneri M., 2014, *MNRAS*, 440, 1446
 Wang Y., Mohanty S. D., Jenet F. A., 2014, *ApJ*, 795, 96
 Wang J. B. et al., 2015, *MNRAS*, 446, 1657
 Zhu X.-J. et al., 2014, *MNRAS*, 444, 3709

APPENDIX A: APPROXIMATE REPRESENTATION OF A_+ AND A_\times

The interpolated approximations to A_+ and A_\times that we make use of can be written as

$$\mathcal{A}_*(t) = \sum_{\mu=1}^{N_\tau-1} l_{*,\mu}(t) w(t, \tau_\mu, \tau_{\mu+1}), \quad (\text{A1})$$

where

$$l_{*,\mu}(t) = \mathcal{A}_{*,\mu} + \frac{(\mathcal{A}_{*,\mu+1} - \mathcal{A}_{*,\mu})}{(\tau_{\mu+1} - \tau_\mu)} (t - \tau_\mu), \quad (\text{A2})$$

and

$$w(t, x, y) = \Theta(t - x) - \Theta(t - y). \quad (\text{A3})$$

Here, $\mathcal{A}_{*,\mu} \equiv \mathcal{A}_*(\tau_\mu)$. This piecewise linear interpolation scheme has been used in Hobbs et al. (2012) to search for monopolar timing residual correlations associated with terrestrial clock errors and Keith et al. (2013) in an effort to account for time-variable dispersion measures in individual pulsar timing data sets. As we wish to incorporate this linear interpolant into the timing model, we will need to know how the timing residuals change as the parameters of the interpolant are changed, i.e.

$$\begin{aligned} \frac{\partial \delta t_K(t)}{\partial \mathcal{A}_{*,\lambda}} &= G_K^* \left[1 - \frac{t - \tau_1}{\tau_2 - \tau_1} \right] w(t, \tau_1, \tau_2); \quad \text{if } \lambda = 1, \\ &= G_K^* \left\{ \frac{t - \tau_{\lambda-1}}{\tau_\lambda - \tau_{\lambda-1}} w(t, \tau_{\lambda-1}, \tau_\lambda) \right. \\ &\quad \left. + \left[1 - \frac{t - \tau_\lambda}{\tau_{\lambda+1} - \tau_\lambda} \right] w(t, \tau_\lambda, \tau_{\lambda+1}) \right\}; \quad \text{if } 1 < \lambda < N_\tau, \\ &= G_K^* \frac{t - \tau_{N_\tau-1}}{\tau_{N_\tau} - \tau_{N_\tau-1}} w(t, \tau_{N_\tau-1}, \tau_{N_\tau}); \quad \text{if } \lambda = N_\tau. \end{aligned} \quad (\text{A4})$$

To estimate $\mathcal{A}_{+, \mu}$ and $\mathcal{A}_{\times, \mu}$, they must be fit for simultaneously with the timing models of several pulsars. Without doing this, structure induced in the residuals of a single pulsar from a GW could be highly covariant with the structure anticipated from an inaccurate timing model. The timing model parameters would become biased away from their maximum-likelihood values and much of the power in the residuals from the GW would be absorbed. To execute a simultaneous fit of all the pulsars in the PTA that can accommodate global parameters that are shared by all the timing models, we can construct modified design and noise covariance matrices and use standard least-squares fitting techniques as in equations (6) and (7). A modified design matrix that allows us to carry out such a fit is \mathbf{M}_g , structured as follows:

$$\begin{bmatrix} \mathbf{M}_1 & \dots & \mathbf{0} & \frac{\partial \delta t_1(t_1)}{\partial \mathcal{A}_{+,1}} & \dots \\ \vdots & \ddots & \vdots & \vdots & \ddots \\ \mathbf{0} & \dots & \mathbf{M}_{N_p} & \frac{\partial \delta t_{N_p}(t_{N_p})}{\partial \mathcal{A}_{+,1}} & \dots \\ & & & \frac{\partial \delta t_1(t_1)}{\partial \mathcal{A}_{+,N_\tau}} & \frac{\partial \delta t_1(t_1)}{\partial \mathcal{A}_{\times,1}} & \dots & \frac{\partial \delta t_1(t_1)}{\partial \mathcal{A}_{\times,N_\tau}} \\ & & & \vdots & \vdots & \ddots & \vdots \\ & & & \frac{\partial \delta t_{N_p}(t_{N_p})}{\partial \mathcal{A}_{+,N_\tau}} & \frac{\partial \delta t_{N_p}(t_{N_p})}{\partial \mathcal{A}_{\times,1}} & \dots & \frac{\partial \delta t_{N_p}(t_{N_p})}{\partial \mathcal{A}_{\times,N_\tau}} \end{bmatrix}. \quad (\text{A5})$$

If the K^{th} pulsar has a $n_K \times m_K$ design matrix, \mathbf{M}_g will be a $(\sum_K n_K) \times (2N_\tau + \sum_K m_K)$ matrix that is block diagonal except for the rightmost $2N_\tau$ columns. The maximum possible value of N_τ is set by the requirement that \mathbf{M}_g have more rows than columns, or that there be more ToAs than model parameters; in practice we choose N_τ to be significantly smaller than this theoretical limit. The global noise covariance matrix will be $(\sum_K n_K) \times (\sum_K n_K)$ and block diagonal:

$$\mathbf{C}_g = \begin{bmatrix} \mathbf{C}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{C}_{N_p} \end{bmatrix}. \quad (\text{A6})$$

The matrices \mathbf{C}_g and \mathbf{M}_g can be used along with a vector of all of the timing residuals, $\delta t^T = [\delta t_1^T \dots \delta t_{N_p}^T]$, in equations (6) and (7) to carry out a global fit simultaneously for the timing models of all the pulsars in the array and the model for a quadrupolar GW signal coming from direction \hat{n} . However, without some additional conditioning, the fit will be ill-behaved because certain components of signals in A_+ and A_\times will induce structure in each pulsar's timing residuals that will be indistinguishable from structure caused by incorrect estimates of that pulsar's timing model parameters, resulting in a singular parameter covariance matrix.

For example, suppose that $A_+(t)$ and $A_\times(t)$ are simply quadratics, $Q_+(t)$ and $Q_\times(t)$. Then, according to equation (3), the timing perturbation from this GW in the K^{th} pulsar's residuals will be $\delta t_K^h(t) = G_K^+ Q_+(t) + G_K^\times Q_\times(t)$, which is itself a simple quadratic. The most fundamental parameters in any pulsar timing model are the spin parameters: a reference rotational phase, ϕ_K , the pulsar's spin frequency, f_K , and the pulsar's spin-down rate, \dot{f}_K . Errors in the timing model estimates of these parameters lead to quadratic structure in the timing residuals. In fact, the only means we have of measuring these quantities for a particular pulsar is by fitting a quadratic to that pulsar's timing residuals. If a GW were creating additional quadratic structure in a pulsar's timing residuals beyond what is expected from incorrect estimates of that pulsar's spin parameters, the two effects could never be meaningfully differentiated. For this reason, if we are to simultaneously fit the timing

models for all the pulsars in the PTA with models for \mathcal{A}_+ and \mathcal{A}_\times , we must constrain the models to not contain quadratics. This can be implemented by enforcing these six equality constraints:

$$\begin{aligned} \sum_{\mu=1}^{N_\tau} \mathcal{A}_{\star,\mu} &= 0, \\ \sum_{\mu=1}^{N_\tau} \tau_\mu \mathcal{A}_{\star,\mu} &= 0, \\ \sum_{\mu=1}^{N_\tau} \tau_\mu^2 \mathcal{A}_{\star,\mu} &= 0. \end{aligned} \quad (\text{A7})$$

We can approximately satisfy these six constraints while simultaneously conducting the global fit by appending six appropriately constructed columns to \mathbf{M}_g , six zero-value pseudo-residuals to δt , and six very small diagonal elements, ϵ , to \mathbf{C}_g . We find that this procedure is numerically stable when we make ϵ five to six orders of magnitude smaller than the largest diagonal element of \mathbf{C}_g . Modern PTA data sets include pulsars with rms ToA uncertainty ranging between several tens of nanoseconds and a few microseconds. Our condition on ϵ thus enforces that the linear equality constraints on \mathcal{A}_+ and \mathcal{A}_\times are satisfied with a precision consistent with or in excess of the timing precisions achievable with the best-timed known pulsars.

In addition to the spin parameters of a pulsar, for the purposes of precision pulsar timing, the celestial coordinates of the pulsar are always fit for, and if the precision with which the pulsar can be timed warrants it, two proper motion terms and a parallax are also fit for. The influence of these astrometric parameters on timing models principally concerns the calculation of the Roemer delay – the difference in light-travel-time for pulses arriving at Earth-based observatories and the SSB. Small estimation errors in the position parameters of a pulsar lead to annual sinusoidal fluctuations in the timing residuals of that pulsar. The sinusoid can be of any phase. Similarly, incorrect estimates for proper motion lead to annual sinusoidal fluctuations of any phase in a pulsar’s residuals, but the amplitude of the sinusoid grows linearly in time as the expected pulsar position deviates more and more from the true pulsar position.

Like a quadratic, if a GW were to induce sinusoidal structure with a one-year period (possibly with an amplitude changing linearly in time) in the residuals of any or all the pulsars in a PTA, that structure could not be differentiated from the signatures of incorrect estimates for the positions and proper motions of the affected pulsars. We must thus also constrain \mathcal{A}_+ and \mathcal{A}_\times to not contain such signatures:

$$\begin{aligned} \sum_{\mu=1}^{N_\tau} \sin(\omega_1 \tau_\mu) \mathcal{A}_{\star,\mu} &= 0, \\ \sum_{\mu=1}^{N_\tau} \cos(\omega_1 \tau_\mu) \mathcal{A}_{\star,\mu} &= 0, \\ \sum_{\mu=1}^{N_\tau} \tau_\mu \sin(\omega_1 \tau_\mu) \mathcal{A}_{\star,\mu} &= 0, \\ \sum_{\mu=1}^{N_\tau} \tau_\mu \cos(\omega_1 \tau_\mu) \mathcal{A}_{\star,\mu} &= 0, \end{aligned} \quad (\text{A8})$$

where $\omega_1 = 2\pi \text{ yr}^{-1}$.

In the context of fitting linear interpolants to individual pulsar timing data sets to model time-variable dispersion measure, Keith

et al. (2013) were the first to discuss and implement the constraints discussed above. Because pulsar timing models often fit for parallax, errors in which induce sinusoidal oscillations with a half-year period in timing residuals, Keith et al. (2013) also constrained their linear interpolants to not contain such biannual sinusoids. But, a parallax constraint is unnecessary for the global, multi-pulsar fits we are discussing. Fitting for parallax in all the pulsars in an array does not generate problems when globally fitting for \mathcal{A}_+ and \mathcal{A}_\times in the same way as fitting for the pulsars’ rotational parameters, positions, or proper motions. Estimation errors in the parallax induce sinusoidal structure in the timing residuals of a pulsar with a specific phase that depends only on the equatorial longitude of that pulsar. If a GW were to induce biannual sinusoidal structure in a pulsar’s residuals, it would have to have a very specific phase to potentially be confused as an error in the parallax estimate for that pulsar, but then, the residuals of other pulsars in the PTA would also show evidence of biannual sinusoidal fluctuations with the same phase in their timing residuals and with amplitudes that vary in accordance with the quadrupolar nature of a GW. If they do, these biannual sinusoidal fluctuations can be reliably ascribed to the activity of a GW in the vicinity of Earth. If they do not, the fluctuations can instead be ascribed to an estimation error in the parallax of one pulsar. For these reasons, no additional constraints regarding parallax must be enforced on \mathcal{A}_+ and \mathcal{A}_\times .

APPENDIX B: TEMPO2 USAGE

Here we provide details for how to make use of the routines described in this paper. In almost all cases the $\mathcal{A}_+(t)$ and $\mathcal{A}_\times(t)$ time series will be included as part of a global fit. The following parameters should therefore be provided in a global parameter file (which we assume is called ‘apac.par’):

```
# The following line sets (S) the
#quadrupolar (Q) interpolation function
#(IFUNC) for the plus (p) polarization.
#The 1 indicates the type of interpolation
#(linear) and the 2 indicates that this
#will be fitted globally to all pulsars
SQIFUNC_p 1 2
# The next line is the same for the
#cross polarization
SQIFUNC_c 1 2
# We now need to define the direction
#(in radians) of the quadrupole for the
#plus and cross functions
#(usually they are the same)
QIFUNC_POS_P <ra> <dec>
QIFUNC_POS_c <ra> <dec>

# We now define the actual grid points
#for the plus polarization
QIFUNC_p1 <mjd> <val> <err>
QIFUNC_p2 <mjd> <val> <err>

# and also for the cross polarization
QIFUNC_c1 <mjd> <val> <err>
QIFUNC_c2 <mjd> <val> <err>

# We then ensure that the constraints
#are correctly applied
```

```
CONSTRAIN QIFUNC_p
CONSTRAIN QIFUNC_c
```

Assuming that four pulsars have been observed, their parameters are in individual files (psr1.par, psr2.par, psr3.par and psr4.par) and their ToAs are in psr1.tim, psr2.tim etc., then the TEMPO2 fit can be carried out using:

```
tempo2 -f psr1.par psr1.tim -f psr2.par
psr2.tim -f psr3.par psr3.tim -f psr4.par
psr4.tim -global apac.par
```

To include red noise models (defined in model.dat), use:

```
tempo2 -f psr1.par psr1.tim -f psr2.par
psr2.tim -f psr3.par psr3.tim -f psr4.par
psr4.tim -global apac.par -dcf model.dat
```

The \hat{A}_+ and \hat{A}_\times time series are written to files called 'aplus_t2.dat' and 'across_t2.dat'. The covariance matrix $C_{+ \times}$ is written to a file called 'aplus_across.cvm'.

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.