

Wide-band profile domain pulsar timing analysis

L. Lentati,¹★ M. Kerr,² S. Dai,² M. P. Hobson,¹ R. M. Shannon,^{2,3} G. Hobbs,²
M. Bailes,⁴ N. D. Ramesh Bhat,³ S. Burke-Spolaor,⁵ W. Coles,⁶ J. Dempsey,⁷
P. D. Lasky,⁸ Y. Levin,⁸ R. N. Manchester,² S. Osłowski,^{4,9,10} V. Ravi,¹¹
D. J. Reardon,⁸ P. A. Rosado,⁴ R. Spiewak,^{4,12} W. van Straten,⁴
L. Toomey,² J. Wang,¹³ L. Wen,¹⁴ X. You¹⁵ and X. Zhu¹⁴

¹*Astrophysics Group, Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, UK*

²*Australia Telescope National Facility, CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia*

³*International Centre for Radio Astronomy Research, Curtin University, Bentley, Western Australia 6102, Australia*

⁴*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, Victoria 3122, Australia*

⁵*National Radio Astronomy Observatory Array Operations Center, PO Box O, Socorro, NM 87801, USA*

⁶*Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093, USA*

⁷*CSIRO Information Management and Technology, PO Box 225, Dickson ACT 2602, Australia*

⁸*Monash Centre for Astrophysics, School of Physics and Astronomy, Monash University, VIC 3800, Australia*

⁹*Fakultät für Physik, Universität Bielefeld, Postfach 100131, D-33501 Bielefeld, Germany*

¹⁰*Max Planck Institute for Radio Astronomy, Auf dem Hügel 69, D-53121 Bonn, Germany*

¹¹*Cahill Center for Astronomy and Astrophysics, MC 249-17, California Institute of Technology, Pasadena, CA 91125, USA*

¹²*Department of Physics, University of Wisconsin-Milwaukee, PO Box 413, Milwaukee, WI 53201, USA*

¹³*Xinjiang Astronomical Observatory, Chinese Academy of Sciences, 150 Science 1-Street, Urumqi, Xinjiang 830011, China*

¹⁴*School of Physics, University of Western Australia, Crawley, WA 6009, Australia*

¹⁵*School of Physical Science and Technology, Southwest University, Chongqing 400715, China*

Accepted 2016 December 22. Received 2016 December 22; in original form 2016 May 23

ABSTRACT

We extend profile domain pulsar timing to incorporate wide-band effects such as frequency-dependent profile evolution and broad-band shape variation in the pulse profile. We also incorporate models for temporal variations in both pulse width and in the separation in phase of the main pulse and interpulse. We perform the analysis with both nested sampling and Hamiltonian Monte Carlo methods. In the latter case, we introduce a new parametrization of the posterior that is extremely efficient in the low signal-to-noise regime and can be readily applied to a wide range of scientific problems. We apply this methodology to a series of simulations, and to between seven and nine years of observations for PSRs J1713+0747, J1744–1134 and J1909–3744 with frequency coverage that spans 700–3600 Mhz. We use a smooth model for profile evolution across the full frequency range, and compare smooth and piecewise models for the temporal variations in dispersion measure (DM). We find that the profile domain framework consistently results in improved timing precision compared to the standard analysis paradigm by as much as 40 per cent for timing parameters. Incorporating smoothness in the DM variations into the model further improves timing precision by as much as 30 per cent. For PSR J1713+0747, we also detect pulse shape variation uncorrelated between epochs, which we attribute to variation intrinsic to the pulsar at a level consistent with previously published analyses. Not accounting for this shape variation biases the measured arrival times at the level of ~ 30 ns, the same order of magnitude as the expected shift due to gravitational waves in the pulsar timing band.

Key words: methods: data analysis – pulsars: general – pulsars: individual.

* E-mail: ltl21@cam.ac.uk

1 INTRODUCTION

When the first pulsar was discovered in 1967 November (Hewish et al. 1968), the receivers used had a central frequency of 81.5 MHz and a bandwidth of 1 MHz. Today, observations with fractional bandwidths of $\gtrsim 1/3$ are common place (e.g. Ransom et al. 2009; Stappers et al. 2011; Cognard et al. 2013; Manchester et al. 2013). The development of new instrumentation for pulsar timing such as the ‘Ultra-broadband receiver’¹ (Karuppusamy et al., in preparation) and the ‘Parkes Ultra-Wideband Receiver’ (Manchester 2015) will result in observations with instantaneous frequency coverage from 600 MHz up to 4 GHz. This improvement in technology, however, has significantly increased the complexity of pulsar timing analysis.

In a typical observation of a pulsar, an average pulse profile for that epoch is formed using the best available estimate of the timing model and using that to ‘fold’ the individual pulses. A model of the average profile, the ‘template’ (see e.g. Manchester et al. 2013 for an example of forming templates) is then used to estimate both the pulse time of arrival (ToA) and its uncertainty. This is most commonly done via the ‘Fourier phase-gradient method’ (Taylor 1992) in which the phase offset between the folded profile data and the template is computed by taking the Fourier transform of both and performing a cross-correlation between the two.

When considering large bandwidths, however, additional considerations are needed. For example, as our line of sight to the pulsar varies with time, so too can the dispersion measure (DM) along that line of sight. If an incorrect DM is assumed when determining the ToA, this can cause a loss of precision by smearing the pulse and potentially bias the arrival times. Profile evolution as a function of frequency has also been readily observed in many pulsars. This can be a result of scattering in the ionized interstellar medium (IISM), instrumental effects or intrinsic variation of the profile. If the pulsar also scintillates as a result of scattering in the IISM (e.g. Narayan 1992), then this will cause a change in the observed flux density of the pulse as a function of frequency. When combined with profile evolution and DM variations, scintillation can further degrade the precision of the ToAs formed from the cross-correlation process if the profile data are averaged over frequency, or if a template is used that does not incorporate profile evolution.

These factors will be complicated further by any shape variation in the observed pulse at each epoch that results from the averaging of a finite number of individual pulses. At the single pulse level, profiles are known to exhibit significant variability (e.g. Hankins & Cordes 1981). As such, as the sensitivity of observations improves and the instrumental noise decreases, this intrinsic stochasticity in the pulse shape will unavoidably become more significant. In particular, if this process is a wide-band phenomenon then the significance of the shape variation will add coherently as the bandwidth increases.

Recently, several approaches have been proposed to mitigate some of the issues present in wide-band pulsar timing. In both Pennucci, Demorest & Ransom (2014) and Liu et al. (2014), two-dimensional extensions to the standard ToA forming process are introduced that fit for the DM at each observational epoch when estimating the ToA, and include profile evolution in the template. Fitting for both the ToA and the DM in this way, however, introduces a significant number of potentially unneeded free parameters. Ideally, one would want to leverage any smoothness present in the DM variations over many years of observations in order to better

constrain the model for DM at each epoch (e.g. Lee et al. 2014). This, however, requires a joint analysis of all the profile data for a particular pulsar.

In this paper, we extend the profile domain analysis framework described in Lentati, Alexander & Hobson (2015) and Lentati & Shannon (2015) (henceforth L15a and L15b, respectively, L15a,b collectively) to include effects such as profile evolution and wide-band shape variation. This framework allows for pulsar timing analysis to be carried out directly on the profile data, rather than estimating ToAs as in the standard pulsar timing paradigm. This means that simultaneous estimation of, for example, DM variations, pulse shape variation, profile evolution and the pulsar timing model can be carried out across the whole data set to ensure the greatest possible constraint. This approach has already been used to model secular profile variation in the MSP J1643–1224 (Shannon et al. 2016), where the global analysis of both shape variation and timing instabilities was shown to improve the sensitivity of the data set to gravitational waves by an order of magnitude.

In L15a,b, all analysis was performed using the MULTINEST (Feroz & Hobson 2008; Feroz, Hobson & Bridges 2009) and POLYCHORD (Handley, Hobson & Lasenby 2015) sampling algorithms. Both these samplers, however, are limited in the dimensionality of the problems that they can solve. While POLYCHORD can efficiently sample from ~ 200 dimensions, as our data sets become more complex, we will need to be able to include many thousands of parameters. As such in this work, we also make use of Hamiltonian Monte Carlo (HMC) sampling methods using the Guided Hamiltonian Sampler (GHS; Balan, Hobson & Ashdown, in preparation). In order to exploit this sampler over a wide range of problems, in Section 2, we first introduce a new parametrization of the stochastic parameters used in L15a,b, which is significantly more efficient in the low signal-to-noise (S/N) regime. While we will only use this parametrization here in the context of profile domain pulsar timing analysis, its potential applications are extremely wide ranging.

In Section 3, we then describe our extensions to the profile domain formalism that incorporates wide-band effects, including optimizations that speed the likelihood evaluation time up by approximately two orders of magnitude relative to L15a,b. In Section 4, we apply this framework to a series of simulations of increasing complexity that include scintillation, profile evolution and DM variations. We use these simulations to compare the timing precision obtained between the standard ToA-based pulsar timing paradigm and our profile domain analyses. In Sections 5 and 6, we then apply this approach to data sets for PSRs J1713+0747, J1744–1134 and PSR J1909–3744. These data sets span a range of observing frequencies from ~ 700 to 3600 MHz with observations taken over 7–9 yr. We both analyse the full data sets using the GHS and perform Bayesian model selection on the 2600–3600 MHz data using POLYCHORD to compare different descriptions of the evolution of the profile with frequency, and models that include stochastic, temporal pulse shape variation. Finally, in Section 7, we offer some concluding remarks.

2 A LOW S/N MODEL

Given a data vector \mathbf{d} of length N_d subject to Gaussian noise, one can write down the probability that the data are described by a model vector \mathbf{s} , which can be considered a function of some parameters θ as

$$\Pr(\mathbf{d}|\theta) = \frac{1}{\sqrt{(2\pi)^{N_d} \det(\mathbf{N})}} \exp\left(-\frac{1}{2}(\mathbf{d} - \mathbf{s})^T \mathbf{N}^{-1}(\mathbf{d} - \mathbf{s})\right), \quad (1)$$

¹ <http://www.mpifr-bonn.mpg.de/research/fundamental/ubb>

where without loss of generality, we will consider \mathbf{N} to be a diagonal matrix with elements $N_{ij} = \sigma_i^2 \delta_{ij}$, with σ_i the standard deviation of the uncorrelated noise in data point d_i . In the toy problems that follow, σ_i is assumed to be known, and the matrix \mathbf{N} is therefore a fixed input to the analysis.

In particular, we will parametrize the model as the product of a vector of N_m amplitude parameters \mathbf{a} , and an $N_d \times N_m$ matrix of basis vectors \mathbf{M} . These basis vectors could be Fourier modes as in Lentati et al. (2013) that describe a stochastic gravitational wave background in pulsar timing data, or spherical harmonics as in Taylor, Ashdown & Hobson (2008), where the model describes the fluctuations in the cosmic microwave background. In either case, if we consider a model vector, $\mathbf{s}(\mathbf{a}) = \mathbf{M}\mathbf{a}$, that describes a zero-mean stochastic process, one can then include a Gaussian prior on the amplitude parameters such that the probability that the model amplitudes, \mathbf{a} , are described by a set of hyper-parameters $\boldsymbol{\varphi}$ is given by

$$\Pr(\mathbf{a}|\boldsymbol{\varphi}) = \frac{1}{\sqrt{(2\pi)^{N_m} \det(\Psi)}} \exp\left(-\frac{1}{2}\mathbf{a}^T \Psi^{-1} \mathbf{a}\right), \quad (2)$$

where Ψ is a diagonal matrix with elements $\Psi_{ij} = \varphi_i^2 \delta_{ij}$, where φ_i describes the standard deviation of the i th amplitude parameter a_i . Note that there need not be a one-to-one correspondence between the amplitude parameters and the hyper-parameters. For example, in our analysis of the toy problems in this section, we will use a single hyper-parameter to describe the standard deviation of all the amplitude parameters included in the model.

One can then combine these terms to give the posterior probability of both the model amplitudes \mathbf{a} and the hyper-parameters $\boldsymbol{\varphi}$, given the data:

$$\Pr(\mathbf{a}, \boldsymbol{\varphi}|\mathbf{d}) = \Pr(\mathbf{d}|\mathbf{a})\Pr(\mathbf{a}|\boldsymbol{\varphi})\Pr(\boldsymbol{\varphi}), \quad (3)$$

where $\Pr(\boldsymbol{\varphi})$ is the prior probability distribution for $\boldsymbol{\varphi}$.

We now consider two toy problems using the posterior in equation (3) that henceforth we refer to as ‘Model Parametrization 1’ (MP1).

In the first, which we refer to as problem T1, we consider a zero S/N scenario. Here we take our data vector \mathbf{d} to be of length 10, and include uncorrelated Gaussian noise with unit variance. The matrix \mathbf{N} in equation (1) will thus be the identity matrix and, as stated previously, will be held constant in the analysis. We also take our matrix of basis vectors \mathbf{M} to be the identity matrix, such that the model vector \mathbf{s} describes an additional uncorrelated stochastic term in the data. We include a single hyper-parameter φ that describes the standard deviation of the model vector. In equation (2), this corresponds to having $\Psi_{ij} = \varphi^2 \delta_{ij}$. In total, we are therefore sampling from an 11-dimensional parameter space. Note that the amplitude parameters are not intended to model the uncorrelated Gaussian noise that is already factored into the matrix \mathbf{N} . Instead, the model amplitudes are describing an additional uncorrelated process in excess of the unit variance noise. As such, we would expect it to be consistent with zero in our analysis.

Rather than sample directly from the parameter φ , we instead sample from the parameter $\rho = \log_{10}(\varphi)$. We perform this analysis using both a prior that is uniform in the amplitude of φ and a prior that is uniform in the log of φ . The posterior probability distributions obtained using the MULTINEST algorithm for the first model amplitude, a_1 , and ρ using the log-uniform prior and uniform prior are shown in the top-left and middle-left panels of Fig. 1, respectively. One can immediately see that these are highly non-trivial parameter spaces to search over in the zero S/N case. When using a prior that is uniform in the log of φ , all scales are a priori

equally probable. Thus, when no signal is present, any amplitude less than the noise level will have equal weight in the posterior. The sampler must therefore explore progressively smaller values of a_1 up to the lower bound of the prior, which here we have set to be -8 . This covers approximately eight orders of magnitude in the model amplitude a_1 . When using a prior that is uniform in the amplitude of φ , this issue is somewhat alleviated, as the prior distribution favours larger values of ρ . However, it is still challenging to sample from, especially when using sampling methods that use a single step size for each parameter, which is not appropriate in this case, in order to fully explore the parameter space.

In order to solve these issues, we now describe a re-parametrization of equations (1) and (2). Here, rather than sample from the parameters \mathbf{a} , we instead sample from the related parameters \mathbf{u} , where, for the i th model amplitude, we will have

$$a_i = u_i \varphi, \quad (4)$$

whereas before φ is the model parameter describing the standard deviation amplitude parameters \mathbf{a} . In order to still sample uniformly in the original parameters, \mathbf{a} , we then include an additional term, the determinant of the Jacobian describing the transformation from a_i to u_i . The Jacobian in this case has elements:

$$J_{i,j} = \varphi \delta_{i,j}, \quad (5)$$

with $\delta_{i,j}$ the Kronecker delta. The determinant is therefore

$$\det(\mathbf{J}) = \prod_{i=0}^m \varphi, \quad (6)$$

which acts to cancel exactly with the determinant of the matrix Ψ in equation (2). Henceforth, we refer to this as ‘Model parametrization 2’ (MP2).

The top-right and middle-right panels of Fig. 1 show the one- and two-dimensional posterior probability distributions for the first amplitude parameter, u_1 , and ρ for the same toy problem T1, using log-uniform and uniform priors in φ for the two panels, respectively. The difference between the old and new likelihoods is clear, with a single step size in \mathbf{u} now being appropriate across the range of ρ sampled in both cases. When performing the sampling using MULTINEST, we find that this parametrization results in approximately a factor of 5–10 decrease in the number of samples required compared to the original parametrization.

In the bottom-left and bottom-right panels of Fig. 1, we then show the same parameters obtained from a second, high S/N, toy problem in which our data vector has an additional uncorrelated stochastic component with a standard deviation of 10. As in the first toy problem, we still include the unit-variance noise component that is accounted for in the fixed noise matrix \mathbf{N} . In this problem, each amplitude parameter will therefore be of the order of 10, and will take values in the approximate range from -30 to 30 . The hyper-parameter ρ will be approximately 1. In this case, the new parametrization is less effective, as the model amplitudes and power spectrum coefficients become correlated, whereas in the original parametrization, they are completely uncorrelated.

In the following work, when using the GHS, we will use both parametrizations dependent upon whether the model in question is in the high or low S/N regime.

3 A PROFILE DOMAIN MODEL

The methods used in this work are drawn from those presented in L15a,b. The key difference here is that while in these previous

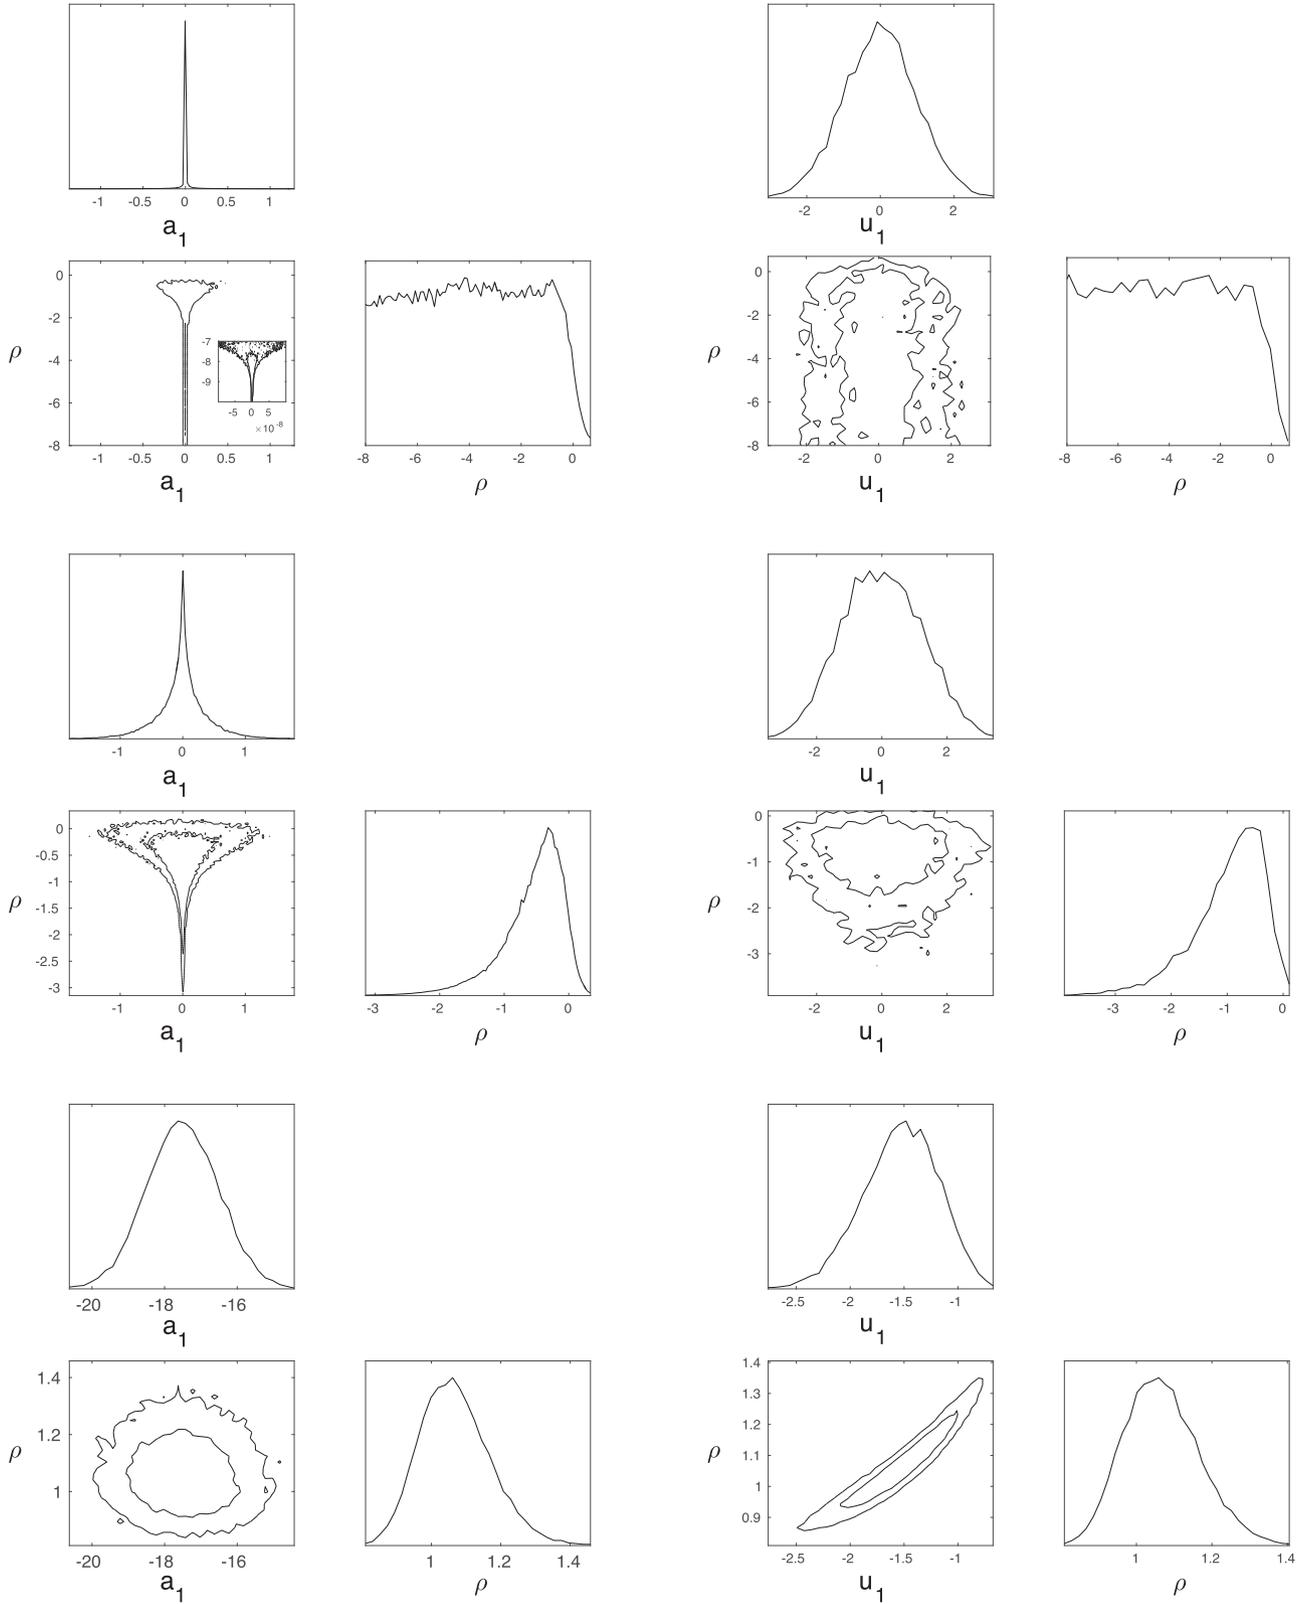


Figure 1. One- and two-dimensional posterior probability distributions for MP1 (left-hand panels) and MP2 (right-hand panels) when applied to two-, 11-dimensional toy problems. The top and middle panels are from a zero S/N example where we have used priors that are uniform in the log of φ (top) and uniform in φ (middle). The bottom panels are from a high S/N problem where each amplitude parameter is $\sim 10\sigma$. We show only the first amplitude parameter u_1 or a_1 for MP2 and MP1, respectively.

works, frequency-averaged profiles were used for each observational epoch, here we will be using multichannel profile data. As a result, in the following section, we will extend the existing formalism to incorporate band-wide variation and profile evolution. In addition, the use of multichannel data significantly increases the number of profiles that must be dealt with in the analysis. In the data sets used in Section 6, there are between 12 000 and 21 000 profiles, compared to 300 used in the analysis of PSR J1909–0747 in L15b. We therefore incorporate a shapelet interpolation scheme into our analysis that speeds up the likelihood evaluation by approximately two orders of magnitude for these larger data sets, with no detectable loss of timing precision. We describe this process in Section 3.2.

For a full description of the general profile domain framework, we refer the reader to L15a,b. Below we will give details of how the methodology has been changed to accommodate wide-band observations, and any differences that arise when performing the analysis with the GHS.

3.1 A profile model

As in L15a,b, we construct our profile model using the shapelet basis (Refregier 2003). A shapelet profile is described by a position t , a scalefactor Λ and a set of n_{\max} amplitude parameters with which we can construct the set of basis functions:

$$B_n(t; \Lambda) \equiv \Lambda^{-1/2} \Phi_n(\Lambda^{-1}t), \quad (7)$$

with $\Phi_n(t)$ given by

$$\Phi_n(t) \equiv [2^n n! \sqrt{\pi}]^{-1/2} H_n(t) \exp\left(-\frac{t^2}{2}\right), \quad (8)$$

where H_n is the n th Hermite polynomial. We can then write our profile model, $s(t, \zeta, \Lambda)$, as the sum:

$$s(t, \zeta, \Lambda) = \sum_{n=0}^{n_{\max}} \zeta_n(\nu) B_n(t; \Lambda), \quad (9)$$

where ζ_n are the shapelet amplitudes, which we have explicitly written as a function of the observing frequency ν , and where n_{\max} is the number of shapelet basis vectors included in the model.

In our analysis, as in L15a,b, we use the shapelet basis to describe the overall profile shape and then include an amplitude parameter for each profile in the data set that scales this average model. One component of the shapelet model therefore acts as the reference for the rest. In the work that follows, we use the zeroth-order term, leaving only $n_{\max} - 1$ free parameters ζ_n that are the amplitudes for the shapelet components with $n > 0$ and we take $\zeta_0 = 1$. Each profile will also have an arbitrary baseline offset, which we denote as γ . Written in this way, equation (9) becomes

$$s(t, A, \zeta, \Lambda, \gamma) = A \sum_{n=0}^{n_{\max}} \zeta_n(\nu) B_n(t; \Lambda) + \gamma, \quad (10)$$

with A the overall scaling factor for a particular profile.

Note that one is free to use multiple shapelets to model a single profile. We do this in Section 6 in our analysis of PSR J1744–1134, for which there are two well-separated profile components corresponding to the main pulse and the interpulse. In this case, we parametrize the amplitudes of the interpulse components relative to the zeroth-order Gaussian term from the main pulse, and include a parameter $\delta\phi$ corresponding to their separation in units of time. For simplicity, we henceforth write the set of parameters that describe

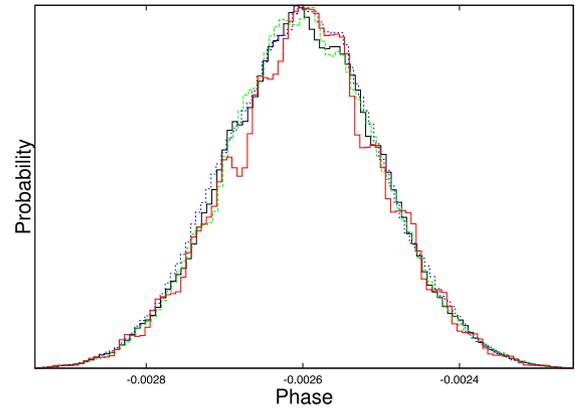


Figure 2. One-dimensional posterior probability distribution for the ToA in phase for a single epoch for PSR J1909–3744 using different interpolation intervals. Colours represent: no interpolation (black), 1 (green), 32 (blue) and 128 ns (red). The effect of interpolation only begins to impact parameter estimates beyond the largest interval.

the profile as $\theta \equiv (A, \zeta, \Lambda, \delta\phi, \gamma)$. Our model in this case will then be

$$s(t, \theta) = A \left(\sum_{n=0}^{n_{\max}} \zeta_n(\nu) B_n(t; \Lambda) + \sum_{m=0}^{m_{\max}} \zeta_m(\nu) B_m(t + \delta\phi; \Lambda) \right) + \gamma, \quad (11)$$

where the index n refers, as before, to the main pulse, and the index m refers to the interpulse, which includes m_{\max} amplitude parameters in the model. While PSR J1909–3744 is also known to have an interpulse (e.g. Dai et al. 2015), as it is extremely weak, we do not consider it in the profile model.

3.2 Shapelet interpolation

Evaluating the shapelet model for a large number of profiles rapidly becomes extremely computationally intensive for even a modest numbers of shapelet components ($n_{\max} \sim 10$). In L15b, the profile model for PSR J1909–3744 required approximately 30 components, and evaluating this model dominated the likelihood evaluation time. For large numbers of profiles, this approach is not computationally tractable. We therefore adopt an interpolation scheme, where the shapelet basis vectors are pre-computed on a grid from $t = 0$ up to the duration of a single phase bin for the maximum-likelihood value of the scalefactor Λ , determined using the fully folded profile data. When performing the sampling, we then linearize the width parameter and include this in our model simultaneously with the rest of the profile parameters, a process we describe in Section 3.4. We only need to compute the grid up to the duration of a single phase bin, as for shifts of greater than one bin, the interpolated model can be rotated by the relevant integer number of bins. The interpolation interval of this grid can be chosen based on the precision of the data set being analysed. When performing the sampling, rather than evaluate the shapelet basis vectors directly, we instead use the set of gridded basis functions that are closest to the ToA predicted by our model.

In Fig. 2, we show the estimated ToA in phase for a single profile from our PSR J1909–3744 data set when using no interpolation, and when choosing the closest possible interpolation interval to 1 (2879 interpolation bins), 32 (90 interpolation bins) and 128 ns (23 interpolation bins). For the 1 and 32 ns interpolation interval, no significant departure from the uninterpolated case can be seen, and

the parameter estimates and uncertainties derived for the ToA are almost identical. For an interpolation interval of 128 ns, the effects of quantization become visible and the parameter estimates and uncertainties begin to diverge.

For all the analysis that follows, we use an interpolation interval of 1 ns, chosen to be sufficiently small so that no bias enters our analysis as a result of the interpolation process.

3.3 Evaluating the timing model

In order to extend the timing framework described in L15a,b to incorporate multichannel profile data, we first consider the data in terms of a set of N_e epochs. Each epoch i then has $N_{c,i}$ channels, such that the total number of profiles $N_p = \sum_{i=1}^{N_e} N_{c,i}$. The profile in the j th channel of the i th epoch then consists of a set of $N_{i,j}$ values representing the flux density of the profile as measured at a set of times $t_{i,j}$. As in L15a,b, we write our profile model, $s_{i,j}$, as a function of

- (i) the timing model parameters for the pulsar in question ϵ ,
- (ii) the overall phase offset ϕ ,
- (iii) the profile parameters θ .

We will therefore have

$$s_{i,j}(t, \theta, \epsilon, \phi) = s(t - \tau(\epsilon)_{i,j} - \phi, \theta), \quad (12)$$

where $\tau(\epsilon)_{i,j}$ is the ToA predicted by the set of timing parameters ϵ . As in L15b, we compute the correction to the Solar system barycenter for each profile after including the phase offset and any instrumental offsets.

We can then write the likelihood that the data are described only by the timing model parameters, the phase offset, the shapelet parameters and baseline offsets as

$$\begin{aligned} \Pr(\mathbf{d}|\epsilon, \phi, \theta) &\propto \prod_{i=1}^{N_e} \prod_{j=1}^{N_{c,i}} \frac{1}{\sqrt{\det \mathbf{N}_{i,j}}} \\ &\times \exp \left[-\frac{1}{2} (\mathbf{d}_{i,j} - s_{i,j})^T \mathbf{N}_{i,j}^{-1} (\mathbf{d}_{i,j} - s_{i,j}) \right], \end{aligned} \quad (13)$$

where $\mathbf{N}_{i,j}$ is the white noise covariance matrix for the profile corresponding to the j th channel in the i th epoch, with elements $(N_{i,j})_{mn} = \sigma_{i,j} \delta_{mn}$, with $\sigma_{i,j}$ the root mean square (rms) deviation of the uncorrelated radiometer noise in the profile.

When using POLYCHORD, in order to decrease the dimensionality of the problem, we obtain an estimate of $\sigma_{i,j}$ from the off-model region of the profile data, where the fractional amplitude of the model profile is less than 0.1 per cent. As in L15b, this value will then be modified by a global scaling factor, referred to as PFAC, which will be an additional free parameter in the analysis. We note that in traditional pulsar timing, a parameter called EFAC is used to scale the ToA uncertainties. We refer to our scaling parameter as PFAC in the profile domain case because, in principle, it should have an equivalent effect. When performing the analysis with the GHS, however, we do not include a PFAC, as $\sigma_{i,j}$ is a free parameter in our analysis for every profile.

3.4 Models for profile evolution

A priori, little is known about the pulsar emission mechanism or in what way the pulse shape will vary with frequency. In the ‘radius-to-frequency mapping’ model, emission is considered to be narrow

band at a given altitude in the pulsars magnetosphere. In this framework, the emission frequency decreases with altitude, which can lead to a change in both the separation in phase, and in the width, of different components in the average pulse profile (e.g. Cordes 1978).

Given such a theoretical framework, one could potentially build a frequency-dependent profile model by searching over a basis that consists of a number of independent components where the evolution of the width and separation in phase for each component are determined by free parameters in the analysis. Such a model, however, will be highly non-linear and contain significant degeneracies making the parameter space difficult to explore.

In our analysis, we therefore consider two models for the evolution of the profile with frequency. The first is a general polynomial expansion of the shapelet amplitudes with frequency. While this model is not physically motivated, it is capable of describing any potential smooth profile evolution, and as with the model for the average profile, is linear in the amplitude parameters, thus making it simple to sample from. The model is defined such that for the p terms in the polynomial, we can write the i th shapelet amplitude $\zeta_i(v)$ as

$$\zeta_i(v) = \sum_{k=0}^p (v - v_c)^k \zeta_{i,k}, \quad (14)$$

where v_c is an arbitrary reference frequency, and $\zeta_{i,k}$ is the amplitude parameter for the k th polynomial of the i th term in the shapelet model. The second model is a simple linear change in the width of the profile with frequency. We can incorporate this model into our analysis by making the scalefactor, Λ , a function of frequency, such that equation (10) becomes

$$s(t, v, A, \zeta, \Lambda_c, \lambda) = A \sum_{n=0}^{n_{\max}} \zeta_n(v) B_n(t; \Lambda(v, \lambda, \Lambda_c)), \quad (15)$$

where the profile width as function of frequency is given by

$$\Lambda(v, \lambda, \Lambda_c) = \Lambda_c (1 + \lambda(v - v_c)), \quad (16)$$

with Λ_c the width at the reference frequency v_c . In order to avoid recomputing the interpolated shapelet basis vectors for different widths, we can expand our shapelet model for small, fractional changes in the width, which we denote λ . We can define a ‘width profile’ as

$$\begin{aligned} w_{i,j}(t, \theta) = \frac{A}{\Lambda_c} \sum_{n=0}^{n_{\max}} \zeta_n(v) &\left[\left(t_{\beta}^2 - \frac{1}{2} \right) B_n(t; \Lambda_c) \right. \\ &\left. - \sqrt{2n} t_{\beta} B_{n-1}(t; \Lambda_c) \right], \end{aligned} \quad (17)$$

with

$$t_{\beta} = \frac{t}{\Lambda_c}, \quad (18)$$

such that our altered shapelet model will be given by

$$s'_{i,j}(t, v, \theta, \lambda) = s_{i,j}(t, \theta) + \lambda(v - v_c) w_{i,j}(t, \theta). \quad (19)$$

3.5 Pulse jitter

As in L15b, we define pulse jitter as a shift in the arrival time of the mean profile relative to that predicted by the pulsar’s timing model, where the shifts are uncorrelated between epochs. Note that the term ‘jitter’ is often used to denote any stochastic shape variation, such as ‘stochastic wide-band impulse modulated self-noise’ (Osłowski

et al. 2011). In Section 3.7, we will discuss a more general model for epoch-to-epoch variation in the profile shape to model such processes. While stochastic shape variation that is uncorrelated between observational epochs would result in a white noise process in the ToAs, uncorrelated shifts in the arrival times could also result from the high-frequency tail of flat spectrum timing noise or small systematic offsets between observations. This approach makes the assumption that the jitter amplitudes at each epoch are Gaussian distributed, where the standard deviation of that distribution is a free parameter in the analysis.

When sampling with POLYCHORD, we take the same approach as in L15b where this shift is linearized. In this case, for the j th channel in the i th epoch, with profile model $s_{i,j}(t, \theta)$, we can write down a ‘jitter profile’ $j_{i,j}(t, \theta)$ given by

$$j_{i,j}(t, \theta) = \frac{1}{\sqrt{2}\Lambda} \sum_{n=0}^{n_{\max}} \zeta_n(v_j) (\sqrt{n} B_{n-1}(t; \Lambda) - \sqrt{n+1} B_{n+1}(t; \Lambda)), \quad (20)$$

such that our shifted shapelet model will be given by

$$s'_{i,j}(t, \theta, \delta t_i) = s_{i,j}(t, \theta) + \delta t_i j_{i,j}(t, \theta). \quad (21)$$

Note that the only change that was required to make this model for pulse jitter band wide was to define one shift, δt_i , for each epoch i . For a narrow-band description of pulse jitter, where the shifts are uncorrelated between different frequency channels, one would simply define a separate shift per channel j as $\delta t_{i,j}$.

When performing the analysis with the GHS, however, we use the non-linear model, and so include the shift parameter δt directly in our expression for the profile model:

$$s'_{i,j}(t, \theta, \Lambda, \delta t_i) = A_i \sum_{n=0}^{n_{\max}} \zeta_n B_n(t - \delta t_i; \Lambda). \quad (22)$$

Regardless of which sampler is used, the standard deviation of the N_e shift parameters δt is then incorporated into the analysis by including a Gaussian prior on the amplitude parameters. As in L15b, this is achieved by defining the covariance matrix \mathbf{J} of the jitter amplitudes as:

$$J_{ij} = \langle \delta t_i \delta t_j \rangle = \mathcal{J}_i \delta_{ij}, \quad (23)$$

where the hyper-parameter \mathcal{J}_i is the standard deviation of the arrival time shifts due to the jitter model at epoch i , and the angle brackets denote the expectation value. When using the linearized model, we then normalize the amplitudes as in L15b such that the shift amplitudes can be described by a single standard deviation \mathcal{J}_i for all observations in units of seconds.

We note that one of the advantages of sampling with the GHS is that it is not necessary to marginalize analytically over any of the model parameters. As such, it is much simpler to include a non-Gaussian prior for the amplitude parameters describing the pulse jitter by following the same process as in Lentati, Hobson & Alexander (2014). This will be investigated further in subsequent work.

In Section 3.1, when considering a pulse profile that consisted of a main pulse and an interpulse, we explicitly wrote the shapelet model as the sum of these two separate components. In this case, the total jitter profile will similarly be given by the sum of the two separate jitter profiles. In Section 6, in our analysis of PSR J1744–1134, this allows us to compare two models for pulse jitter. First, where both the main pulse and the interpulse are shifted by the same amount at each epoch, and secondly, where the main pulse is able to shift independently of the interpulse. Henceforth,

we refer to this second case as ‘Interpulse Jitter’, which is therefore equivalent to modelling epoch-to-epoch variation in the parameter $\delta\phi$ in equation (11), which described the separation of the main pulse and interpulse in the profile model.

3.6 Width jitter

In addition to pulse jitter, we can use the expression for small changes in the profile width given in equation (17) to model uncorrelated epoch-to-epoch variation in the profile width. As before, we can define the i th altered epoch given a change in width λ_i as

$$s'_{i,j}(t, \theta, \lambda_i) = s_{i,j}(t, \theta) + \lambda_i w_{i,j}(t, \theta). \quad (24)$$

As for the pulse jitter, we then include a constraint on the amplitudes of the width jitter parameters λ by fitting for the standard deviation of the distribution. The covariance matrix \mathbf{W} of the width jitter amplitudes is written as

$$W_{ij} = \langle \lambda_i \lambda_j \rangle = \mathcal{W}_i \delta_{ij}, \quad (25)$$

where the hyper-parameters \mathcal{W} represent the standard deviation of the changes in the profile width at epoch i .

3.7 Profile stochasticity

In addition to a model for pulse jitter, L15b introduced models for variation in the shape of the profile that were uncorrelated between epochs. Here we will refer to this profile stochasticity in terms of uncorrelated (UC) and phase-correlated (PC) variations in the pulse shape, where in both cases, the correlations we refer to are in pulse phase.

3.7.1 Uncorrelated stochasticity

In L15b, a ‘stochastic envelope’ was defined that represented an increase in the uncorrelated noise in the on-pulse region. This modelled ‘self-noise’ in the profile data (e.g. Lentati et al. 2014), representing shape variation on scales smaller than the width of a phase bin. In L15b, the stochastic envelope was defined to have the same shape as the mean profile, and was proportional to the amplitude of the profile.

In the analysis presented here, we will consider a two-component model for the stochastic envelope:

- (i) As in L15b, a term that is proportional to the profile, and
- (ii) a second term that is constant across the on-pulse region and is proportional to the amplitude of the pulse.

We therefore define two scaling parameters, α and β , which represent this increase in the variance, which add in quadrature to our instrument noise $\sigma_{i,j}$ such that the new variance in a bin k for epoch i and channel j is given by

$$\hat{\sigma}_{ijk}^2 = \sigma_{ij}^2 + \beta^2 s_{i,j}(t, \theta)_k^2 + \alpha^2 A_{i,j}^2 \delta_p, \quad (26)$$

where the delta function δ_p is 1 for the on-pulse region, and 0 for the off-pulse region, and $A_{i,j}$ is the amplitude of the model profile for epoch i and channel j .

When sampling with POLYCHORD, we marginalize analytically over the individual profile amplitudes. As such, as in L15b, we use a maximum-likelihood estimate of the amplitudes, $A_{i,j}$, in order to scale both terms in the stochastic envelope as in equation (26). This maximum-likelihood estimate is given by

$$\hat{A}_{i,j} = \frac{F_d}{F_m}, \quad (27)$$

with F_d the flux in the on-pulse region of the profile data, and F_m the flux in the profile model setting $A_{i,j} = 1$. With the GHS, however, we sample from the profile amplitude numerically, and thus simply include the amplitude parameters directly in equation (26).

3.7.2 Phase-correlated stochasticity

To model low-frequency stochastic shape changes in the profile, we use the same shapelet basis as for the template. Our goal is then to robustly determine the power spectrum of the shape changes as a function of the scale (or order) of the component in the model. We note that this is a very general model for epoch-to-epoch shape variation in the profile data. As such, the shape variation induced by the presence of either width jitter or pulse jitter could be modelled using this approach. However, when performing Bayesian analysis, if the data really are well described by a simple model for pulse jitter, then this will be reflected in the posterior probability distributions for the stochastic hyper-parameters. Qualitatively, this is simply because it is more probable that a single hyper-parameter describing the standard deviation of the pulse jitter is the appropriate amplitude, than it is that the many parameters that describe the low-frequency stochasticity all take appropriate values.

We define the set of shape variation power spectrum coefficients \mathcal{S} , such that for a given scale i the covariance matrix at a particular epoch k will be given by

$$(\mathcal{S}_i)_{jk} = \frac{1}{P_{d,k}} \langle (\zeta_i - \bar{\zeta}_i)(\zeta_j - \bar{\zeta}_j) \rangle = \frac{1}{P_{d,k}} \mathcal{S}_i \delta_{ij}, \quad (28)$$

where $\bar{\zeta}_i$ is the amplitude of the i th shapelet coefficient from the average profile and, as before, the angle brackets denote the expectation value. As each profile in the data set can have an arbitrary overall normalization for its amplitude, we need to scale the variance for each shapelet coefficient on a per-epoch basis. We therefore define the parameters \mathcal{S}_i for a profile with total power equal to one, and include the factor $P_{d,k}$ to scale these parameters for each epoch k . When sampling with POLYCHORD, we estimate this scaling parameter for each profile from the data, taking $P_{d,k} = F_{d,k}^2$, whereas when sampling with the GHS, we obtain $P_{d,k}$ directly using the model amplitudes being sampled. In either case, this gives us the power spectrum coefficients \mathcal{S}_i in units of the fraction of the total power in the profile.

Note that for all the models we have discussed that describe stochastic processes, the covariance matrix of the model amplitudes is defined to be diagonal. However, this does not mean that we are making the assumption that the modes in the signal are orthogonal to one another in the profile data where they are sampled. This is identical to the implementation of stochastic models for timing noise used elsewhere in pulsar timing analysis (see e.g. Lentati et al. 2013), where any covariances either between the signal parameters, or the hyper-parameters, are correctly accounted for in the likelihood.

3.8 Marginalizing analytically over the linear parameters

As in L15a,b, when performing our analysis with POLYCHORD, we will marginalize analytically over the linear amplitude parameters in order to decrease the parameter space. We will now describe this process for our wide-band models, for which the implementation is sufficiently different to the narrow-band models described in L15b to warrant detailing below.

In total, we will be marginalizing over the following:

- (i) the arbitrary offset for each profile $\gamma_{i,j}$,

- (ii) the overall amplitude for each profile $A_{i,j}$,
- (iii) the band-wide pulse jitter amplitudes δt_i ,
- (iv) the band-wide width jitter amplitudes w_i ,
- (v) the low-frequency profile stochasticity amplitudes ζ_i .

Given N_c channels in a particular epoch i , where the profiles in each channel j have N_b phase bins, we will have $N_c N_b$ phase bins in total for the entire epoch.

As the number of linear parameters we wish to marginalize over is N_m , we therefore define the $(2N_c + N_m) \times N_c N_b$ matrix \mathbf{P}_i , for which the first $2N_c$ rows will have a block diagonal format given by

$$\begin{aligned} (P_i)_{2(j-1)+1, N_b(j-1)+k} &= 1 \\ (P_i)_{2(j-1)+2, N_b(j-1)+k} &= (s_{i,j})_k, \end{aligned} \quad (29)$$

for each channel j and phase bin k , and will be zero otherwise.

The remaining N_m rows are filled with the basis vectors describing pulse jitter, width jitter and stochastic profile variation. For example, we will have

$$(P_i)_{2N_c+1, N_b(j-1)+k} = j_{i,j}(t, \theta)_k \quad (30)$$

and similarly for the other stochastic parameters.

We then define the diagonal, $(2N_c + N_m) \times (2N_c + N_m)$ matrix Ψ_i , which is zero for the elements corresponding to the overall amplitude and baseline offsets, and is equal to \mathcal{J}_i , \mathcal{W}_i and \mathcal{S}_i for the pulse jitter, width jitter and profile stochasticity parameters.

From here the marginalization process occurs as in L15b. Adopting the notation of that paper, we define the matrix Σ_i for each epoch as $\mathbf{P}_i^T \mathbf{N}_i^{-1} \mathbf{P}_i + \Psi_i^{-1}$ and define $\mathbf{P}_i^T \mathbf{N}_i^{-1} \mathbf{d}_i$ as $\bar{\mathbf{d}}_i$.

Our final probability distribution for the non-linear parameters is then given by

$$\begin{aligned} \Pr(\theta, \epsilon, \mathcal{S}, \mathbf{J} | \mathbf{d}) &\propto \prod_{i=1}^{N_e} \frac{\det(\Sigma_i)^{-1/2}}{\sqrt{\det(\mathbf{N}_i)}} \\ &\times \exp \left[-\frac{1}{2} \left(\mathbf{d}_i^T \mathbf{N}_i^{-1} \mathbf{d}_i - \bar{\mathbf{d}}_i^T \Sigma_i^{-1} \bar{\mathbf{d}}_i \right) \right]. \end{aligned} \quad (31)$$

3.9 Sampling with the GHS

HMC sampling methods have been successfully applied to a wide range of scientific problems with extremely high dimensionality (e.g. over 10^6 in Taylor et al. 2008). HMC uses local gradient information about the likelihood surface to draw on the mathematical framework used to describe the motion of particles in potential wells. In doing so, the random walk behaviour exhibited by conventional MCMC methods is suppressed and the sampler remains efficient even in high-dimensional problems.

One of the main drawbacks with the HMC method that has limited its acceptance within the scientific community is the large number of tuning parameters required in order to effectively explore the parameter space. In particular, the step size for each parameter and the number of steps n_s in the trajectory must be selected, typically requiring tuning runs. If the step size is too small, computational time is wasted taking many small steps, while if it is too large, the acceptance rate will decrease. Similarly, if the number of steps is too small, successive samples will be too close to one another in parameter space leading to high correlation within the chain, while taking too many steps will lead the HMC to follow trajectories that loop around the parameter space.

The GHS is designed to eliminate much of this tuning aspect. It makes use of the Hessian of the problem probability distribution calculated at its peak to set the step size for each parameter. The number of steps is then drawn from a uniform distribution, $U(1, n_{\max})$, with n_{\max} of 10 found to be suitable for all tested problems. A single global scaling parameter for the step size is then the only tunable parameter, chosen such that the acceptance rate for the GHS is ~ 68 per cent. In order to perform the sampling with the GHS, we therefore need the following:

- (i) the gradient of negative log likelihood for each parameter,
- (ii) the peak of the joint probability distribution,
- (iii) the Hessian calculated at that peak.

While, in principle, one might be tempted to include as much information in the Hessian as possible, in practice, we do not follow this approach for two reasons.

First, many of the parameters are almost completely uncorrelated. Storing the full Hessian in such cases therefore wastes both system memory and computing time. The latter effect is a result of the matrix–vector multiplications required by the GHS in each likelihood calculation that are of the size of the Hessian. We therefore divide our Hessian into $N_p + 1$ blocks. We have one 3×3 block for each profile, which includes the elements of the Hessian for the overall amplitude, baseline offset and system noise for that profile. The Hessian for all remaining parameters, which includes the timing model, profile model and any stochastic parameters, is then stored as a final block.

Secondly, the Hessian can become numerically unstable when including the cross-terms for certain parameters. For example, we have found that when using the MP2 likelihood, including the full cross-correlation between the amplitude parameters that define the model vector s and the timing model parameters can significantly reduce the sampling efficiency as a result of this instability.

3.9.1 Obtaining the maximum-likelihood solution

In order to obtain a suitable point at which to calculate the Hessian, we use a two-stage process. We first use a Nelder–Mead optimization algorithm (Nelder & Mead 1965) to find the maximum-likelihood parameter estimates using the likelihood in equation (31). Using these parameters, we then analytically solve for the remaining linear parameters that were marginalized over in Section 3.8.

3.9.2 Calculating gradients and the Hessian

While we will not list the gradients and elements of the Hessian for all parameters, we will provide details for some of the key parameters below. In particular, the gradients and elements of the Hessian for the non-linear timing model and all parameters that result in shifts of the profile warrant some consideration.

We find the basis vectors used by TEMPO2² to evaluate the linear timing model to be sufficient to compute the gradient and Hessian of the non-linear timing model parameters.

In order to calculate the gradient for shifts in the arrival time of the profile, we can use the jitter profile, \mathbf{j} , given in equation (20). For example, denoting the profile residuals after subtracting the model at a particular point in parameter space as \mathbf{r} , and the basis vector for the i th timing model parameter, ϵ_i , as \mathbf{M}_i , we can approximate

the log-likelihood for small changes in this parameter as

$$L(d|\epsilon_i) = \sum_{j=0}^{N_p} \frac{1}{2} (\mathbf{r}_j - A_j M_{ij} \epsilon_i \mathbf{j}_j)^T \mathbf{N}_j^{-1} (\mathbf{r}_j - A_j M_{ij} \epsilon_i \mathbf{j}_j), \quad (32)$$

with A_j the amplitude and M_{ij} the value of the basis vector \mathbf{M}_i for the j th profile. We can then compute the gradient of equation (32) with respect to ϵ_i :

$$\frac{dL(d|\epsilon_i)}{d\epsilon_i} = \sum_{j=0}^{N_p} -A_j M_{ij} \mathbf{j}_j^T \mathbf{N}_j^{-1} (\mathbf{r}_j - A_j M_{ij} \epsilon_i \mathbf{j}_j), \quad (33)$$

and the second derivative:

$$\frac{d^2L(d|\epsilon_i)}{d\epsilon_i^2} = \sum_{j=0}^{N_p} (A_j M_{ij})^2 \mathbf{j}_j^T \mathbf{N}_j^{-1} \mathbf{j}_j. \quad (34)$$

4 SIMULATIONS

In order to test the efficacy of the analysis method described in the preceding sections, we apply it to three simulations of increasing complexity. Details of these simulations are given below.

(i) Simulation 1

Simulation 1 consists of 100 observational epochs covering a total timespan of 6 yr. Each epoch assumes 256 Mhz of bandwidth with a central frequency of 1369 Mhz, which is split into 32, 8-MHz channels. The model profile used is shown in Fig. 3 (top left) and does not evolve over this frequency range. Each epoch includes 64 one-minute sub-integrations, with a total integrated S/N ~ 1000 for each epoch. We simulate a simple timing model including only position, period, period derivative and DM with values for all parameters chosen to be consistent with PSR J1713+0747. We list the simulated timing model parameters in Table 1.

(ii) Simulation 2

As Simulation 1, however, we include both profile evolution across the observed band and scintillation. The simulated profile at the top and bottom of the band is shown in Fig. 3 (top right), while an example of the simulated scintillation for one epoch is shown in Fig. 3 (bottom left). We simulate the dynamic spectrum of interstellar scintillation according to Dai et al. (2016). Our simulations are in the regime of strong scintillation, and are valid for a thin scattering layer of homogeneous isotropic turbulence with a Kolmogorov spectrum.

(iii) Simulation 3

As Simulation 2, however, we also include significant DM variations, consistent with those observed in PSR J1045–4509 (Reardon et al. 2016). The measured DM at each epoch is shown in Fig. 3 (bottom right, black points with error bars).

For each simulation, we obtain fully frequency-averaged ToAs using the timing model in Table 1 and also ToAs for the 32 channels separately. We then perform a standard timing analysis using the ToAs for these two data sets and a profile domain analysis on the profiles used to form the 8-MHz ToAs.

We list the relative timing precision, given as the mean ratio of the uncertainties in the timing model parameters, obtained from these different analyses for each of the three simulations in Table 2. In Simulation 1, we find no evidence for any stochastic parameters in either the ToA or profile domain, and obtain completely consistent results from each of the analysis methods. This is to be expected, as

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

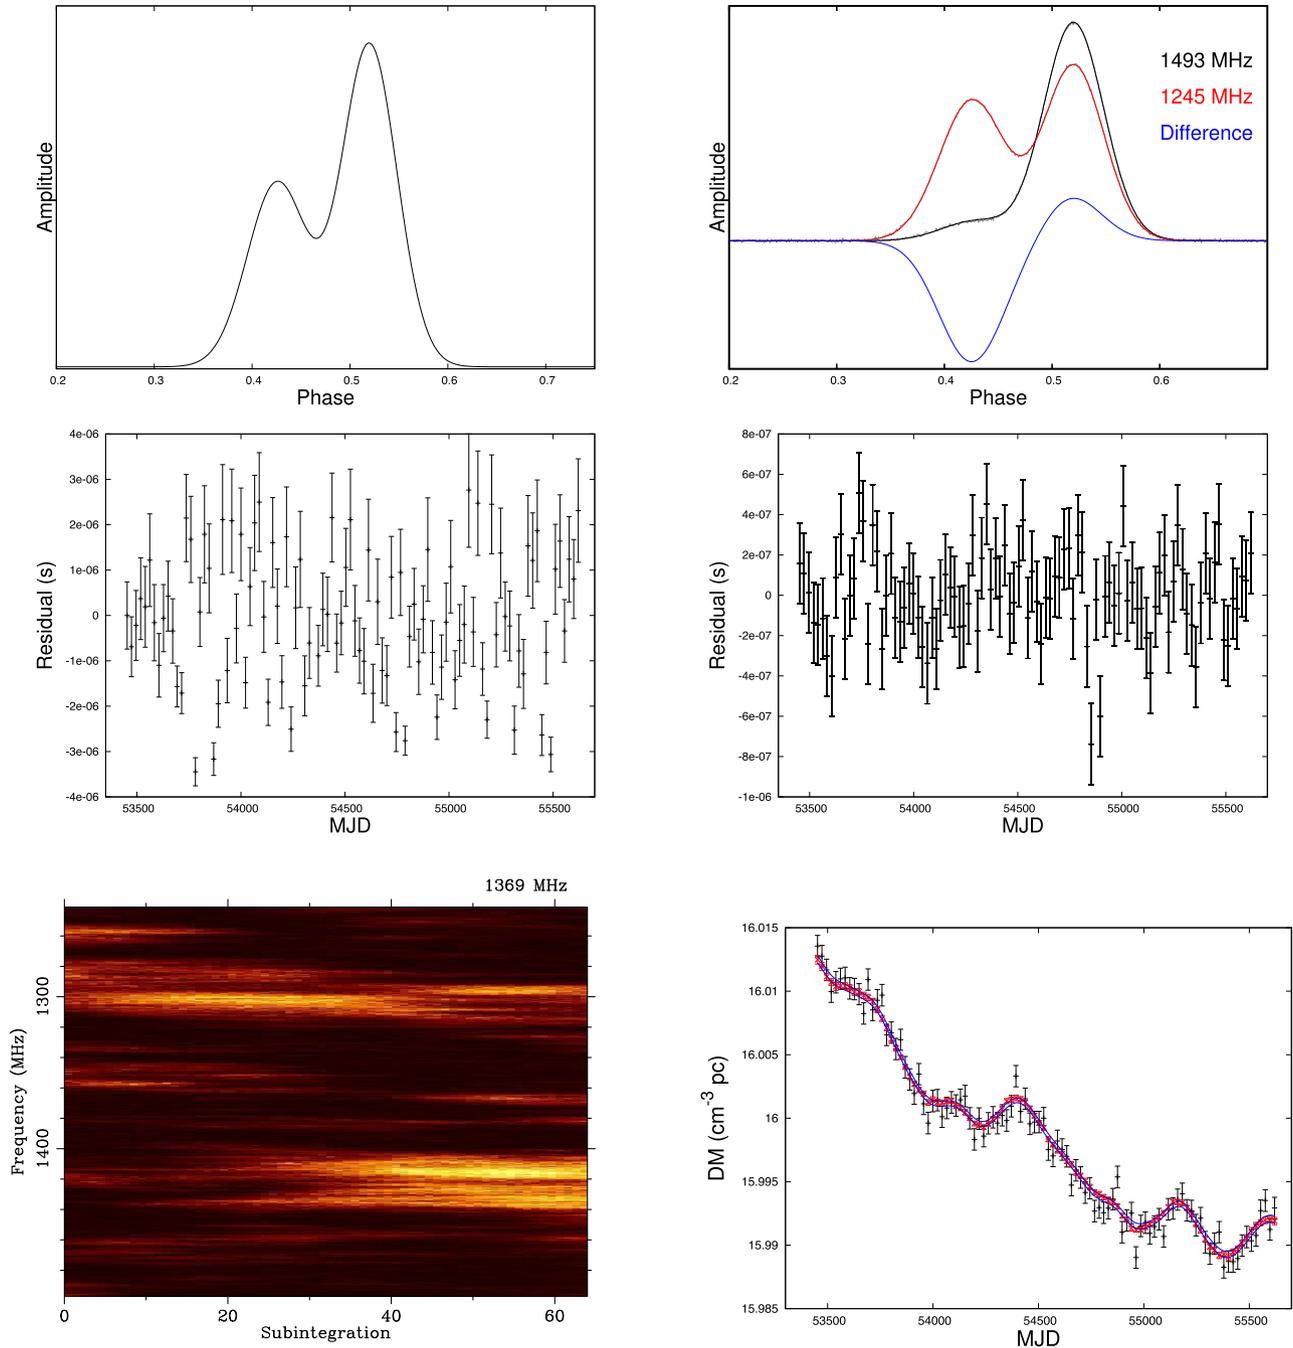


Figure 3. Top left: model profile used in Simulation 1. Top right: evolving profile model used in Simulations 2 and 3. The profile evolves linearly across the band from the red to the black curve. (Middle) Timing residuals from Simulation 2 when fully frequency averaging the profile data before forming the ToAs (left-hand panel), and when forming ToAs from 8-MHz channels, and then averaging the residuals for each epoch after fitting for the timing model (right-hand panel). In the latter case, we include ‘FD’ parameters in the timing model to act as a proxy to profile evolution in the ToA analysis. Bottom left: example of the scintillation for one epoch in Simulation 2. Bottom right: DM variations in Simulation 3 measured independently at each epoch (black points with error bars), measured using the DMX parametrization in a profile domain analysis (red points with error bars), and using a smooth model for the DM variations in the profile domain (blue line representing the 1σ confidence interval).

the assumptions made when forming the ToAs (no profile evolution, stationary DM) were correct.

This is not the case in either of the subsequent simulations. In Simulation 2, the combination of profile evolution and scintillation results in significant loss of precision in the fully frequency-averaged ToAs. This can be seen in Fig. 3 (middle-left panel), where we show the large amount of scatter present in the timing

residuals. When modelling this scatter in the frequency-averaged ToA analysis, we use an EQUAD parameter (an additional white noise term that adds in quadrature with the ToA uncertainty) with a value of 1.34×10^{-6} s. Given the typical uncertainty on the frequency-averaged ToAs is 0.8×10^{-6} s, this is a significant decrease in the precision of the data. The impact of the profile evolution and scintillation is much less significant in the 8-MHz ToAs. In

Table 1. Simulated timing model parameters.

Parameter	Value
Right ascension, α (hh: mm: ss)	17:13:49.532 5456
Declination, δ (dd: mm: ss)	+07:47:37.499 94
Pulse frequency, ν (s^{-1})	218.811 840 441 437 121 35
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-4.084 12 \times 10^{-16}$
Dispersion measure, DM ($cm^{-3}pc$)	16.0

Table 2. Measured timing precision in simulations relative to the 32 channel ToA analysis. For Simulation 3, we do not include the parameters describing the DM variations.

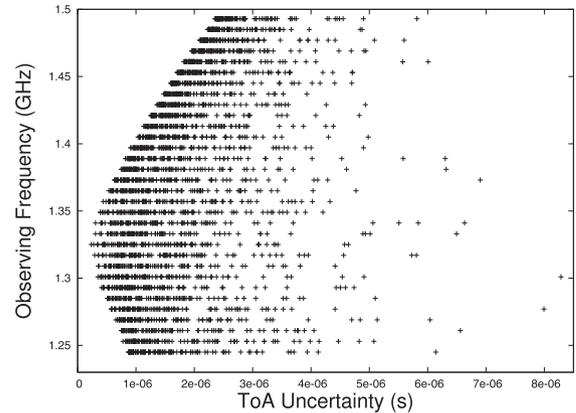
Analysis	Relative timing precision
Simulation 1	
Frequency averaged ToAs	1.0
32 channel ToAs	1.0
Profile domain	1.0
Simulation 2	
Frequency-averaged ToAs (EQUAD)	8.0
32 channel ToAs (FD)	1.0
Profile domain (PE)	0.75
Simulation 3	
Frequency-averaged ToAs (EQUAD)	9.0
32 channel ToAs (DMX, FD)	1.0
Profile domain (DMX, PE)	0.73
Profile domain – (smooth DM, PE)	0.43

Fig. 3 (middle-right panel), we show the epoch-averaged residuals when modelling the profile evolution using the ‘FD’ parametrization (Arzoumanian et al. 2015). The FD parameters model profile evolution as a shift in the arrival time given by

$$\Delta_{FD} = \sum_{i=1}^n c_i \log \left(\frac{\nu}{1 \text{ GHz}} \right)^i, \quad (35)$$

with the c_i free parameters to be fit for. In our analysis, we find only the first term is required to model the profile evolution in our simulations in the ToA domain with a value of $(33.4 \pm 1.0) \times 10^{-5}$. Higher order terms are highly correlated, possibly due to the narrow overall bandwidth of the simulation, and we find including them does not improve the fit.

We find no evidence for an EQUAD using the ToAs formed from the 8-MHz channels when including the first FD parameter, and obtain timing precision that is a factor of 8 better than the fully frequency-averaged ToAs. In our profile domain analysis, we model the profile evolution directly as a linear change in the profile parameters (cf. Section 3.4), as opposed to using a proxy as with the FD parametrization. In this case, we find an additional 25 per cent improvement in the timing precision compared to the 8-MHz ToAs. This improvement comes because, in the ToA domain, we have formed ToAs using a template that does not incorporate profile evolution, and are modelling that evolution purely as a shift in the arrival times. In Fig. 4, we show the ToA uncertainties obtained from using a single template across the observing band. There is a clear trend towards larger uncertainties towards the edge of the band, as the mismatch between the template and the profile increases. In our ToA domain analysis, while we include EFAC and EQUAD parameters that scale and add in quadrature to the error bars, no

**Figure 4.** Formal ToA uncertainties for the 8-MHz ToAs from Simulation 2. A clear trend towards larger error bars can be seen towards the highest and lowest frequencies as a result of a mismatch between the average template, and the evolving profile.

models currently include smooth, frequency-dependent scaling of the ToA uncertainties. This leads to an overall decrease in sensitivity compared to the profile domain analysis, where we are correctly modelling the evolution as a change in shape.

Finally, in Simulation 3, we see a further decrease in the timing precision obtained with the fully frequency-averaged ToAs relative to both the 8-MHz ToAs and the profile domain analysis. In this simulation, we have included significant DM variations, and thus the effect of profile evolution and scintillation is compounded by the fact that we are no longer dedispersing at the correct DM for each epoch. This leads to further perturbation of the frequency-averaged profile, and thus increases the scatter in the residuals. We first model the DM variations using the DMX parametrization (Demorest et al. 2013), in which a piecewise-constant DM(t) model is included in the analysis, with a separate value for each epoch in the data set. In this case, we find a similar improvement in timing precision when going from the 8-MHz ToAs to the profile domain as in Simulation 2. We find a further 60 per cent improvement in precision when using a smooth model for the DM variations in our profile domain analysis. We implement a model for DM variations using the same approach as described in L15a and find that both MP1 and MP2 give consistent results.

In Fig. 3 (bottom right), we show the parameter estimates from the DMX model (red points with errors) and the 1σ confidence interval on the signal from the smooth DM model (blue lines) obtained in the profile domain. We also compare these with the result of performing a two-dimensional analysis on each epoch independently where we fit for both a phase offset and the DM at that epoch (black points with errors). In Fig. 5 (top panel), we show the mean parameter estimates and 1σ uncertainties on the DM power spectrum for those frequencies in the model that are detected with high significance. In total, we find that only 12 frequencies from $1/T$ to $12/T$ are necessary in the model to describe the higher order DM variations, with T the length of the data set. The bottom panel of Fig. 5 shows the mean reduced χ^2 , $\chi_r^2 = \chi^2/\sqrt{N_b}$ for the on-pulse region of the profile data from the 32 channels at each epoch using this smooth model for the DM. We find the results are consistent with a value of 1.0 within the uncertainties. It is this significantly reduced parameter space compared to the 100 DMX parameters that results in the improvement in timing precision between the two models. We note that this improvement will depend upon the complexity of the DM variations present in the data set.

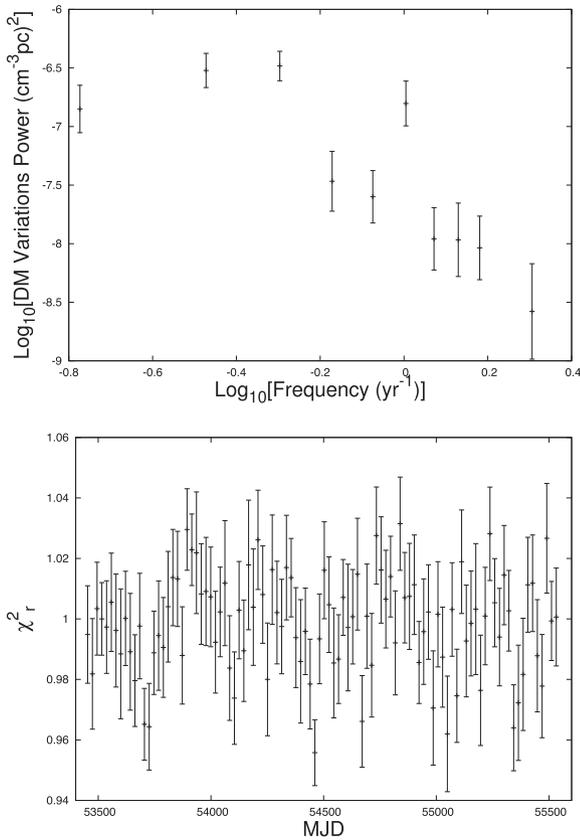


Figure 5. Top: mean parameter estimates and 1σ uncertainties on the DM power spectrum coefficients for those frequencies in the model that are detected with high significance. In total, we find that only 12 frequencies from $1/T$ to $12/T$ are necessary in the model to describe the higher order DM variations, with T the length of the data set. Bottom: mean value and standard deviation of the reduced χ^2 for the on-pulse region of the profile data at each epoch in Simulation 3.

If there are significant non-stationary fluctuations in the DM (e.g. Keith et al. 2013; Coles et al. 2015) that require more complex models (e.g. Lentati et al. 2016), then the difference between the two models will decrease.

Both the DMX and the smooth DM model provide significantly better constraints than modelling the DM independently at each epoch. In Pennucci et al. (2014) and Liu et al. (2014), the DM is estimated in this way in order to obtain ToA estimates that are not biased by dedispersing with an incorrect DM. The degree of improvement in modelling the DM coherently across the entire data set will naturally depend upon the bandwidth and the S/N of that data set. However, as the bandwidth and quality of observations improve, so too will the number of parameters required to model the data. For example, lower frequency observations will require phase and both time-variable DM and scattering terms to be modelled simultaneously for each epoch. If a pulsar shows detectable, band-wide shape variation, a statistical description of that variation should be included in the model in order to robustly estimate the remaining parameters of interest. Other factors, such as frequency-dependent DM across a wide bandwidth (Cordes, Shannon & Stinebrink 2016) will also eventually need to be considered. However, by estimating these parameters simultaneously over the whole data set as in our profile domain framework, one can always ensure that the optimal result is achieved.

Table 3. Details of the individual pulsar data sets. N_E is the total number of observing epochs, N_P is the total number of profiles in the data set and σ_w is the weighted rms of the ToAs formed from the profile data.

Data set PSR	Timespan (yr)	N_E	N_P	σ_w (μs)
J1713+0747	6.74	482	13 643	0.55
J1744–1134	6.74	496	13 610	1.4
J1909–3744	9.02	695	19 194	0.43

5 DATA SETS

We perform our analysis using observations of PSRs J1713+0747, J1744–1134 and J1909–3744 made with the 64-m Parkes radio telescope. Data were collected using two receiver packages, a co-axial system at 10 and 40 cm, and the centre pixel of a multi-beam 20 cm system. Data at 3100 (co-axial, hereafter ‘10 cm’) and 1369 MHz (multibeam, hereafter ‘20 cm’) were recorded using a digital polyphase filterbank (PDFB4) with a typical resolution of 1024 channels and 1024 bins and respective bandwidths of 1024 and 256 MHz. Data from the lower co-axial band, centred at 732 MHz (hereafter ‘40 cm’), were recorded over a 64 MHz bandwidth with a similar polyphase system, DFB3, until its demise in 2014 April. Subsequent data employed CASPSR, a coherently dedispersing system producing 512-channel, 1024-bin output. Data from the DFBs were averaged over 1-min sub-integrations, while CASPSR data maintain an 8-s resolution. Specific details of the three data sets are given in Table 3 and more details about the observing system and data reduction are given in Manchester et al. (2013). For completeness, we outline the data reduction below.

We carry out data reduction using the `PSRCHIVE` package (Hotan et al. 2004). We remove channels within 5 per cent of the band edge as they have low gain and may contain aliased signals. Narrow-band radio-frequency interference (RFI) is identified and removed with a median bandpass filter, and impulsive broad-band RFI is identified by eye and removed by deleting affected sub-integrations. We note that the 10-cm band is largely free of both varieties of RFI. The 20-cm band is affected primarily by the passage of global positioning system satellites through the telescope side lobes and by occasional strong, impulsive RFI from aircraft. The 40-cm band was largely free from RFI until 2015 April when a new mobile phone base station went online, necessitating a shift of observing band frequency and leaving strong contamination.

Observations are carried out with a cadence of approximately 14 ± 10 d for the co-axial and multibeam systems separately. While the 10 and 40-cm observations occur simultaneously for the majority of epochs, the 20-cm observations are interleaved between them, so that the overall cadence is higher, at 6 ± 7 d; however, we note that the true distribution is non-Gaussian. Before each pulsar observation, we record modulated noise injected at the receiver front end, allowing estimation and correction of differential gain and phase between the voltage probes through the amplification and down-conversion system. At each observing epoch, we observe the radio galaxy Hydra A, allowing bandpass and flux density calibration. We find that all quantities vary only modestly with time. We apply these corrections to the raw data to produce flux- and polarization-calibrated pulse profiles. For the 20-cm data, there is evidence for elliptical polarization/cross-coupling of the receiver feed, causing the observed pulse profile to depend slightly on the parallactic angle of the source. We correct these variations using a model of the parallactic-angle dependence obtained from long observations of PSR J0437–4715; see Manchester et al. (2013) for more details.

After calibration, the sub-integrations were time-averaged for each epoch.

There are several discrete time offsets (known as jumps) in the data sets. In particular, a firmware upgrade for the DFB systems at MJD 55319 resulted in an offset for each spectrometer mode that we include as free parameters in our analysis.

6 RESULTS

In the following three sections, we describe the results obtained when using our wide-band profile domain timing technique on the three data sets described in Section 5. In our analysis, we have considered both the 10-cm data on its own, for which we used POLYCHORD to perform evidence comparisons with the analytic likelihood, and also the full data sets covering 3 GHz of bandwidth where the sampling was performed with the GHS using the numerical likelihood. In the latter case, we perform only parameter estimation and do not obtain evidence values for model comparison.

In Section 6.1, we compare the two models for profile evolution described in Section 3.4 that we denote as following models.

(E1) A polynomial expansion of the shapelet amplitudes that describe the mean profile with frequency, where the change in each amplitude is independent, and

(E2) a one-dimensional change in the width of the profile as a function of frequency.

In Section 6.2, we discuss the results for our models for profile stochasticity, for which we consider the following:

(UC) Profile stochasticity uncorrelated in phase, and

(PC) Phase-correlated profile stochasticity.

In Section 6.3, we describe the results for the different jitter models we consider in our analysis. We denote these as follows:

(PJ) Pulse jitter as described in Section 3.5,

(WJ) Width jitter, and

(IPJ) Interpulse jitter.

In Section 6.4, we compare the piecewise DMX and smooth power-law models for DM variations in both our profile and ToA domain analysis, and finally in Section 6.5, we compare the timing precision from the different approaches.

For ease of notation when discussing evidence comparisons, we denote our most basic model, which includes only the pulsar timing model, a model for the mean profile and a PFAC parameter, as model (M0).

6.1 Profile evolution

6.1.1 10-cm data

We find all three data sets show extremely significant support for profile evolution across the 10-cm band. We find an increase in the log evidence, which we denote $\Delta\mathcal{Z}$, of over 100 relative to any model that does not incorporate profile evolution constructed from the elements listed in Section 6.

We find that for model (E2), the fractional change in the profile width is 0.01396 ± 0.0007 , -0.039 ± 0.003 , $-0.0592 \pm 0.0013 \text{ GHz}^{-1}$ for PSR J1713+0747, J1744–1134 and J1909–3744, respectively. Here, a negative value reflects the fact that the profile narrows as the frequency increases. In all three data sets, however, we find that the evidence supports the more complex model for profile evolution, with $\Delta\mathcal{Z} = 387$, 260 and 4 in

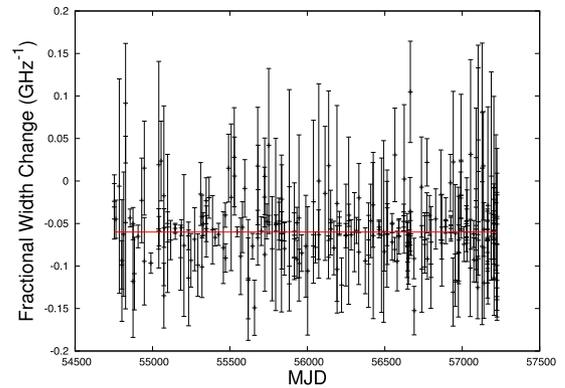


Figure 6. The mean parameter estimate and 1σ uncertainties for the evolution of the profile width as a function of frequency for PSR J1909–3744. Parameter estimates are obtained separately for each epoch from a three-dimensional analysis where the arrival time, DM and width change is evaluated simultaneously, with other timing parameters fixed to the maximum-likelihood estimates obtained from the analysis of the full data set. The horizontal line is set to the mean value of the evolution determined from the full analysis of -0.0592 .

favour of model (E1) for the three data sets, respectively. Apart from PSR J1909–3744, the profile evolution is thus dominated by very general shape change with frequency, as opposed to width changes, implying that little physical interpretation can be drawn from the exact value of the simple one-parameter model. Even in PSR J1909–3744, however, deviations from the simple model are still significant.

One might expect a change in the width as a function of frequency if the DM has not been modelled appropriately and the profile has been smeared across the band. We therefore perform a consistency check on the PSR J1909–3744 data set in which we perform a three-dimensional analysis on each epoch separately. In each case, we fit for the ToA, the DM and the change in width as a function of frequency. In Fig. 6, we plot the mean parameter estimate and 1σ uncertainties for the evolution of the profile width that results from this analysis. We find that the evolution measured at each epoch is consistent with the estimate determined from our global analysis, and obtain a mean and standard deviation from our per epoch analysis of $-0.06 \pm 0.02 \text{ GHz}^{-1}$. That the uncertainty in this case is an order of magnitude larger than the global analysis is simply the result of performing an incoherent analysis across multiple epochs, as opposed to a fully coherent analysis with the timing model.

6.1.2 700–3600 Mhz data

In Fig. 7, we show the mean model pulse profiles for PSRs J1713+0747 (top panel), J1744–1134 (middle panel) and J1909–3744 (bottom panel) at a frequency of 700 (black lines), 1400 (red lines) and 2800 Mhz (blue lines). We find that a cubic expansion of the shapelet coefficients with frequency is sufficient to model the profile evolution for both PSRs J1744–1134 and J1909–3744; however, PSR J1713+0747 required an additional quartic term. This is commensurate with the number of FD parameters required for each of the pulsars when performing a ToA domain analysis, with PSR J1713+0747 warranting the first three terms in the FD model, compared to two for PSR J1744–1134 and one for PSR J1909–3744.

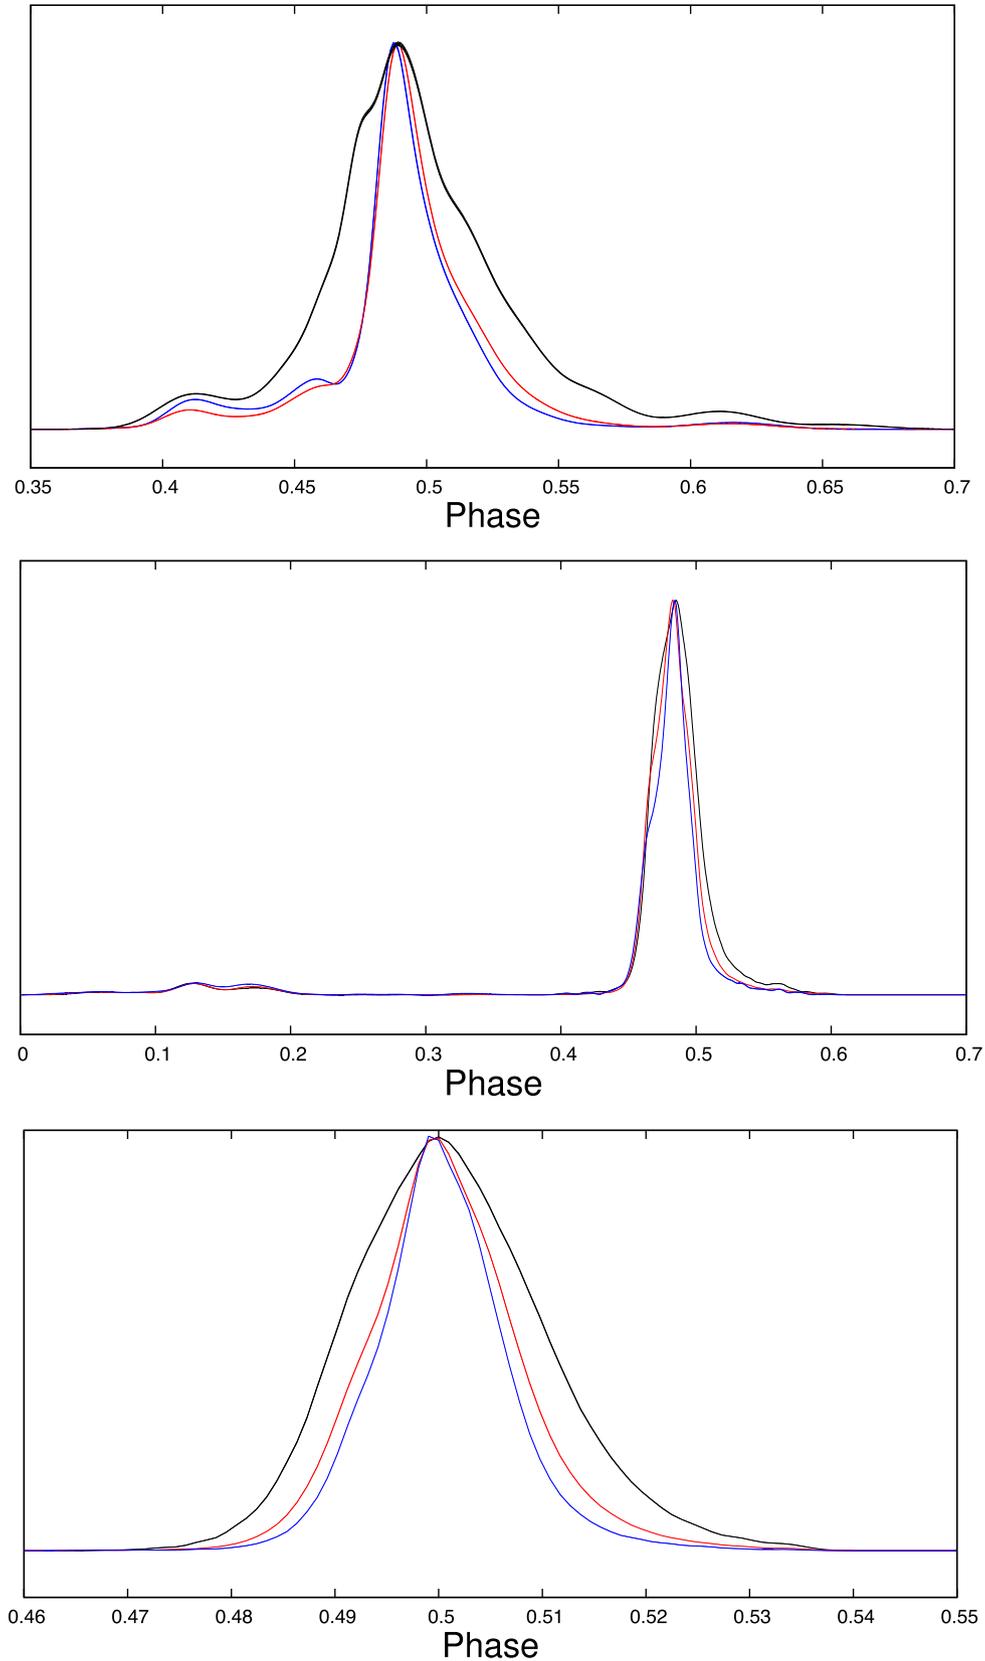


Figure 7. Model pulse profiles for PSR J1713+0747 (top), J1744–1134 (middle) and J1909–3744 (bottom), evaluated at frequencies of 700 (black lines), 1400 (red lines) and 2800 MHz (blue lines). These frequencies are approximately the centres of the 40- and 20-cm bands, and the bottom of the 10-cm band, respectively. In each case, the overall phase parameter is set to $\phi = 0.5$ (cf. equation 12).

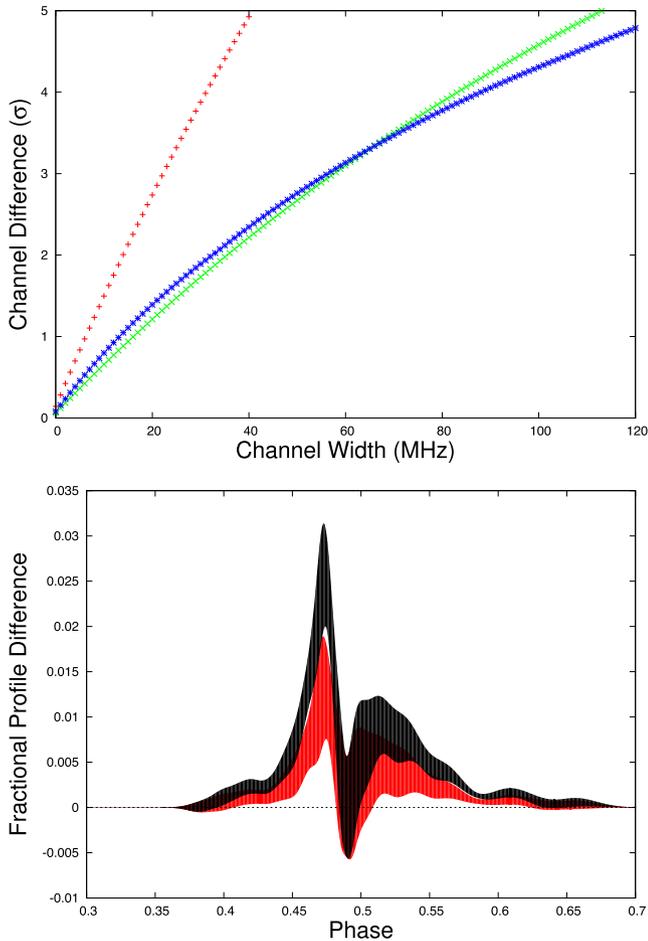


Figure 8. Top panel: the maximum change in the profile model due to frequency evolution as a function of channel bandwidth for PSRs J1713+0747 (red +), J1744–1134 (green ×) and J1909–3744 (blue *) for the 10-cm band. The difference is measured in terms of the confidence intervals on the profile model returned by our profile domain analysis. The curvature is a result of the uncertainty in the model for profile evolution, which increases as a function of frequency, gradually dominating over the uncertainty in the mean profile, which is constant as a function of frequency. Bottom panel: difference between the model for profile evolution for PSR J1713+0747 between frequencies of 724 and 728 Mhz (red interval), and between frequencies of 724 and 732 Mhz (black interval).

We can use the confidence intervals for our models of profile evolution in each data set to determine the channel width such that neighbouring channels are still consistent to some statistical level. We illustrate this in two ways in Fig. 8. First (top panel), using the 10-cm band, we show the maximum change in the profile as a function of observing frequency in terms of the uncertainties in the profile model. We find that for PSRs J1744–1134 and J1909–3744, a channel width of approximately 40 MHz is required such that the maximum difference in the profile model between neighbouring channels at any point in phase is less than 2σ . That these values are similar despite the difference in timing precision between the two data sets is simply a result of the magnitude of the profile evolution being much greater in the PSR J1744–1134 data set, while the S/N of the data set is much larger for PSR J1909–3744. For PSR J1713+0747, however, we find that a narrower channel width of 20 MHz is required to maintain a similar level of statistical consistency. Secondly (bottom panel), we show the difference in

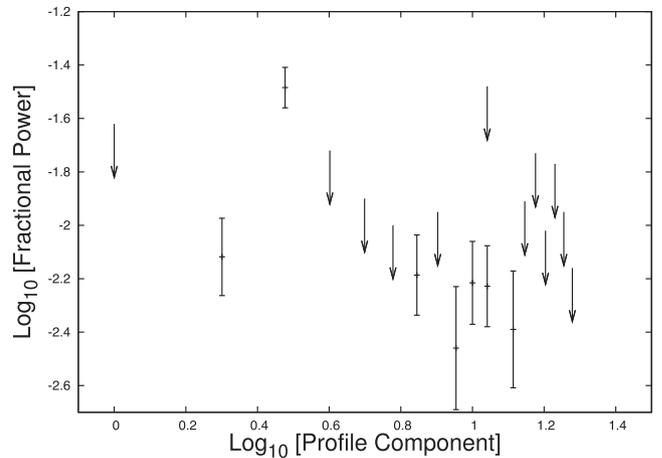


Figure 9. Power spectrum of the shape variations in PSR J1713+0747 from the 10-cm data obtained using POLYCHORD. Arrows indicate 95 per cent upper limits, points with error bars indicate significant detections. Increasing profile component reflects smaller scales in phase.

the profile model for PSR J1713+0747 between frequencies of 724 and 728 Mhz (red interval), and between frequencies of 724 and 732 Mhz (black interval). One can clearly see that even over small fractional changes in observing frequency ($\sim 5.5 \times 10^{-3}$), there is still significant detectable profile evolution. We note that as we were sampling simultaneously for our profile model and our model for DM, these uncertainties account for the covariances that exist between our models for these two processes. Thus, the observed evolution in the profile is not attributable to incorrect estimation of the DM as a function of time.

6.2 Profile stochasticity

6.2.1 10-cm data

When analysing the 10-cm data with POLYCHORD, we detect significant profile stochasticity only in the PSR J1713+0747 data set for which we obtain an increase in the log evidence for an (M0, E1, PC) model compared to an (M0, E1) model of approximately 200. For PSRs J1744–1134 and J1909–3744, we find no increase in the log evidence for either the (PC) or (UC) model parameters. In Fig. 9, we show the power spectrum for the shape variation as a function of the component in the shapelet model from this POLYCHORD analysis. From left to right, the increasing order of the shapelet coefficient can be thought of as representing progressively smaller scales in phase. We find the profile stochasticity to be at the level of approximately 1 per cent in the individual components, and find that the shape variation is detectable in individual epochs with over 3σ significance in approximately 15 per cent of observations.

If the origin of this shape variation were not intrinsic to the pulsar, we might expect that the signal would not be coherent across the full frequency band. As such, we use the 10-cm data to compare models for which the profile stochasticity is a coherent process across the band, with a model where it is incoherent (i.e. where the realization of the change in shape is allowed to vary from channel to channel). We find the evidence is significantly in favour of the coherent model, with an increase in the log evidence of approximately 60 compared to the incoherent model.

Whether this variation is intrinsic to the pulsar, or is related to systematic effects in the data, is impossible to determine with only a single telescope. However, as a consistency check, we calculate

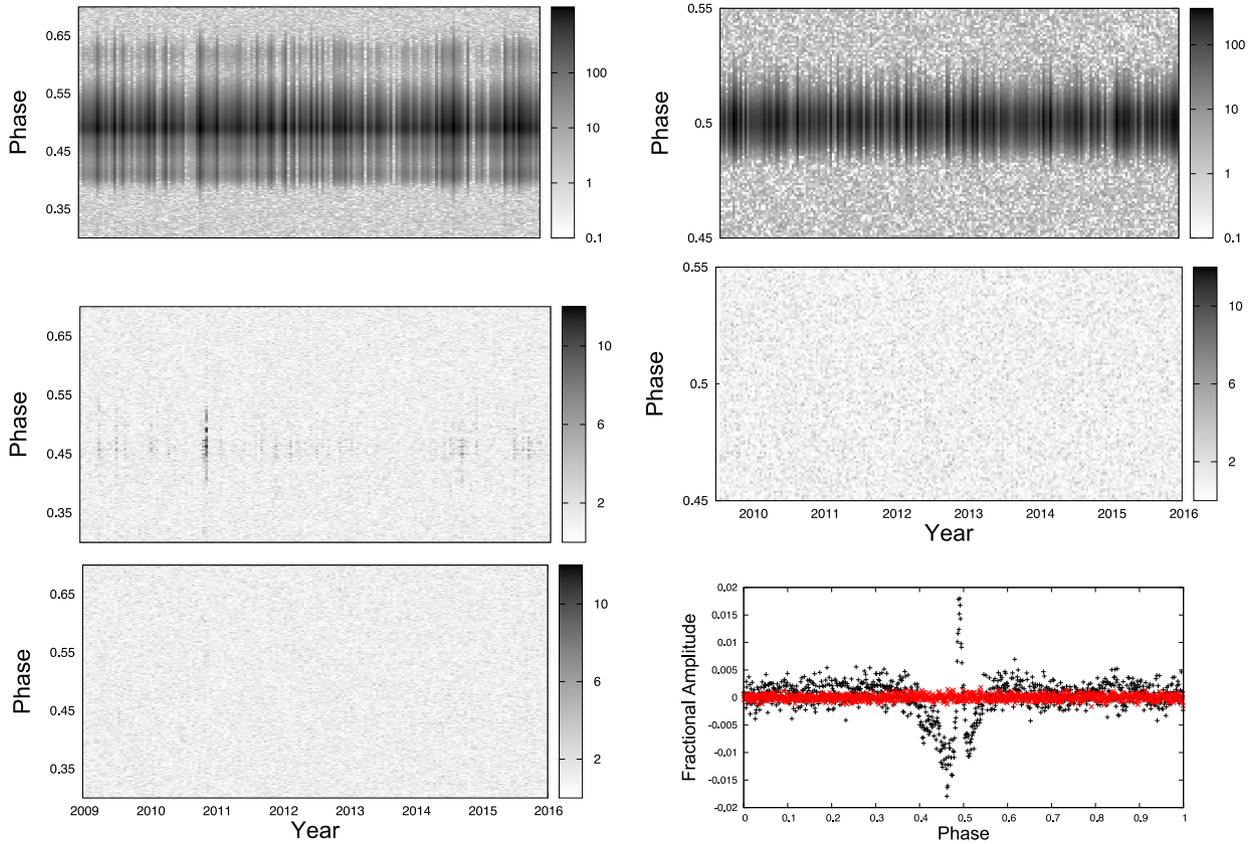


Figure 10. Log of absolute intensity (top panels) and absolute value of the profile residuals after subtracting out the mean profile model (middle panels) for the 20 cm profile data for PSR J1713+0747 (left) and for the 10-cm data for PSR J1909–3744 (right) in units of S/N. For the 20 cm PSR J1713+0747 profile data, we then additionally show the absolute value of the residuals after subtracting out our model for stochasticity in the pulse profile (bottom-left panel), and the residuals from the fully time-averaged 20-cm data (bottom-right panel, red points) with the residuals from the epoch for which the most significant profile stochasticity was detected (black points).

the impact the maximum-likelihood model for the variation has on the ToA of the pulse. We find this to be at the 10–30 ns level for the most significant detections, consistent with previously published estimates for this pulsar (e.g. Shannon et al. 2014; Arzoumanian et al. 2015).

While significant, we do not yet find that this profile stochasticity has a significant impact on the timing analysis, with consistent parameter estimates and uncertainties with or without the (PC) model. This consistency can be explained by estimating the ToA for the epoch with the most significant shape variation either assuming that the profile noise is white, or including the maximum-likelihood parameter estimates for the PC power spectrum in the model. We find that the log evidence increases by 9, and the mean arrival time shifts by 35 ns when including the (PC) model; however, the ToA uncertainty in both cases is 58 ns. We therefore see that even for the most significant example, the data set is not yet sensitive enough for the bias in arrival time when ignoring profile stochasticity to impact the timing results.

The most sensitive 95 per cent upper limit for an isotropic stochastic gravitational-wave background is currently $A < 10^{-15}$ at a reference frequency of 1 yr^{-1} (Shannon et al. 2015). A background of this amplitude would induce a shift in the ToA of a pulse at the level of ~ 100 ns. The change in the arrival time when accounting, or not, for this shape variation in the PSR J1713+0747 data set is already a significant fraction of this level. Observations from the

largest radio telescopes in the world (e.g. Arzoumanian et al. 2015; Desvignes et al. 2016; Reardon et al. 2016) have therefore already reached the stage where correctly modelling such shape variation has become of significant importance in order to reach the greatest levels of sensitivity to gravitational waves.

6.2.2 700–3600 Mhz data

We obtain consistent results using the full data set and performing the sampling with the GHS. In Fig 10 (top-left panel), we show log-intensity for the profile data over the MJD range included in the data set for the PSR J1713+0747 20-cm data in units of S/N. In the middle-left and bottom-left panels, we then show the absolute value of the residuals in units of S/N after subtracting out the mean profile model and additionally after subtracting out our model for stochasticity in the pulse profile. We find that the model for profile stochasticity has successfully modelled the shape variation, with the residuals that remain being noise-like for the duration of the data set. For comparison, in Fig. 10, we also show the log-intensity (top-right panel) and absolute value of the residuals (middle-right panel) for the 10 cm PSR J1909–3744 data set for which no significant profile stochasticity was detected.

We stress that because we are using the same basis to determine the profile stochasticity as for the mean profile model, and because it is being evaluated simultaneously with the mean profile model, this

Table 4. Parameter estimates for pulse jitter models. Upper limits are quoted at the 95 per cent level.

Jitter model	J1713+0747 (ns)	J1744–1134 (ns)	J1909–3744 (ns)
PJ _{10cm}	<140	<290	<80
WJ _{10cm}	<90	<260	<75
IPJ _{10cm}	–	<300	–
PJ _{20cm}	<87	240 ± 35	100 ± 10
PJ _{50cm}	<390	<300	<100

variation is not simply the result of mismodelling the mean profile. Instead, it reflects genuine variation in the shape of the profile between epochs. To demonstrate this, we show in the bottom-right panel of Fig. 10 the residuals from the fully time-averaged 20-cm data (red points) with the residuals from the epoch for which the most significant profile stochasticity was detected (black points). The average residuals are an order of magnitude smaller than the shape variation in the individual epoch.

6.3 Jitter

In the 10-cm data, we find no evidence for jitter in any of the three data sets, with changes in the log evidence of less than two when including jitter parameters compared to model M0, along with our model for evolution (E1). In Table 4, we list the 95 per cent upper limits on the different jitter models considered for each data set.

We find that the limit on the (IPJ) jitter model in PSR J1744–1134 is completely consistent with the (PJ) model. This can be understood by considering the S/N ratio of the interpulse at any epoch. Even in the brightest observations, the interpulse is only just visible. As such, the data are unable to discriminate between a model where both components are shifted by the same amount, as in the (PJ) model, or where it is the separation between the two components that is allowed to vary, as in the (IPJ) model.

When performing the analysis with the GHS, we include a separate parameter \mathcal{J} for each of the three bands. We find the upper limit on the pulse jitter in the 10-cm band is consistent with the analysis performed using POLUCHORD for all three pulsars. As we do not compute the evidence for different models, when using the GHS, we consider the posterior for a particular \mathcal{J} to represent an upper limit when there is greater than 5 per cent probability that the parameter takes a value of less than 1 ns. Given this definition, we only detect significant pulse jitter in the 20-cm bands for PSRs J1744–1134 and J1909–3744.

6.4 DM variations

All three pulsars in our data set are known to exhibit detectable variations in DM (see, e.g. Arzoumanian et al. 2015; Lentati et al. 2016). In Fig. 11, we show the estimates we obtain on the DM variations for three different approaches for each pulsar. We use the DMX parametrization with a 30-d window both with the ToA data (black points with 1σ uncertainties), and in the profile domain (red points with 1σ uncertainties). We compare this with a smooth power-law model in the profile domain (blue lines representing the 1σ confidence interval on the signal), where we include frequencies in our model from $1/T$ with T the length of the data set up to and including

30 d, and use a quadratic in DM to act as a proxy to the lower frequency fluctuations.

When using the DMX model, in order to minimize the covariance between the mean DM and the variations in DM, we take the epoch that has the best individual constraints on DM and keep that fixed as a reference. We then include the mean DM in our fit as a free parameter.

We find that all three models are consistent within Gaussian statistics, with no significant ($>3\sigma$) discrepancies between the DMX model and the power-law model. We note that the non-stationary DM event observed in previous analysis of PSR J1713+0747 occurs before the start of the data set at MJD 54780 (e.g. Lentati et al. 2016).

In Fig. 12, we plot the one- and two-dimensional marginalized posterior probability distributions for the spectral index and log amplitude of the DM power-law model for all three pulsars. We find that the results for PSR J1713+0747 and J1909–3744 are consistent with those from an analysis of the first International Pulsar Timing Array data release (Lentati et al. 2016); however, no significant detection of DM variations in PSR J1744–1134 was made in that analysis for a comparison to be made with the results obtained here.

6.5 Timing precision

We find that the profile domain analysis consistently results in higher precision compared to the standard ToA analysis. Additionally, we find that the measurement precision obtained for the timing model parameters when using the power-law model for DM variations is also superior for all three data sets. We list the mean and 1σ confidence intervals for the timing model parameters obtained from our profile domain analysis of PSRs J1713+0747, J1744–1134 and J1909–3744 using the power-law model for DM in Table 5.

Naturally, the level of improvement when using the power-law model depends both on the pulsar in question and the time-scale over which the timing model parameter occurs. For PSR J1909–3744, the binary period is ~ 1.5 d, and so is significantly shorter than the 30-d window used in the DMX model, or the shortest period in the power-law model. As such, the binary parameters are not highly correlated with either model for DM variations. We therefore find only an ~ 3 per cent improvement in precision when using the power-law model compared to DMX for the binary parameters. In contrast, the binary period of PSR J1713+0747 is approximately 67 d, and so in this case, the binary parameters will be more correlated with the unconstrained DMX model. For this pulsar, we therefore find an average improvement in the measured precision of the binary parameters of 20 per cent for the power-law model compared to DMX.

For all three data sets, the astrometric parameters, which have time-scales of six months or more, benefit more significantly, with an increase in the measured precision of up to 40 per cent. We show the ratios for all timing model parameters for the DMX model compared to the power-law model in the top panel of Fig. 13.

In the centre panel of Fig. 13, we show an equivalent plot comparing the measured timing precision obtained from the standard ToA analysis to the profile domain analysis. In both cases, we use the DMX model for DM variations. The improvement in the timing precision is more uniform across the different timing parameters in this case. We find the average improvement is 25 per cent for the three pulsars analysed. Given each data set is at least 7 yr long, for parameters such as parallax for which the measurement precision scales as the square-root of the total observing time, this

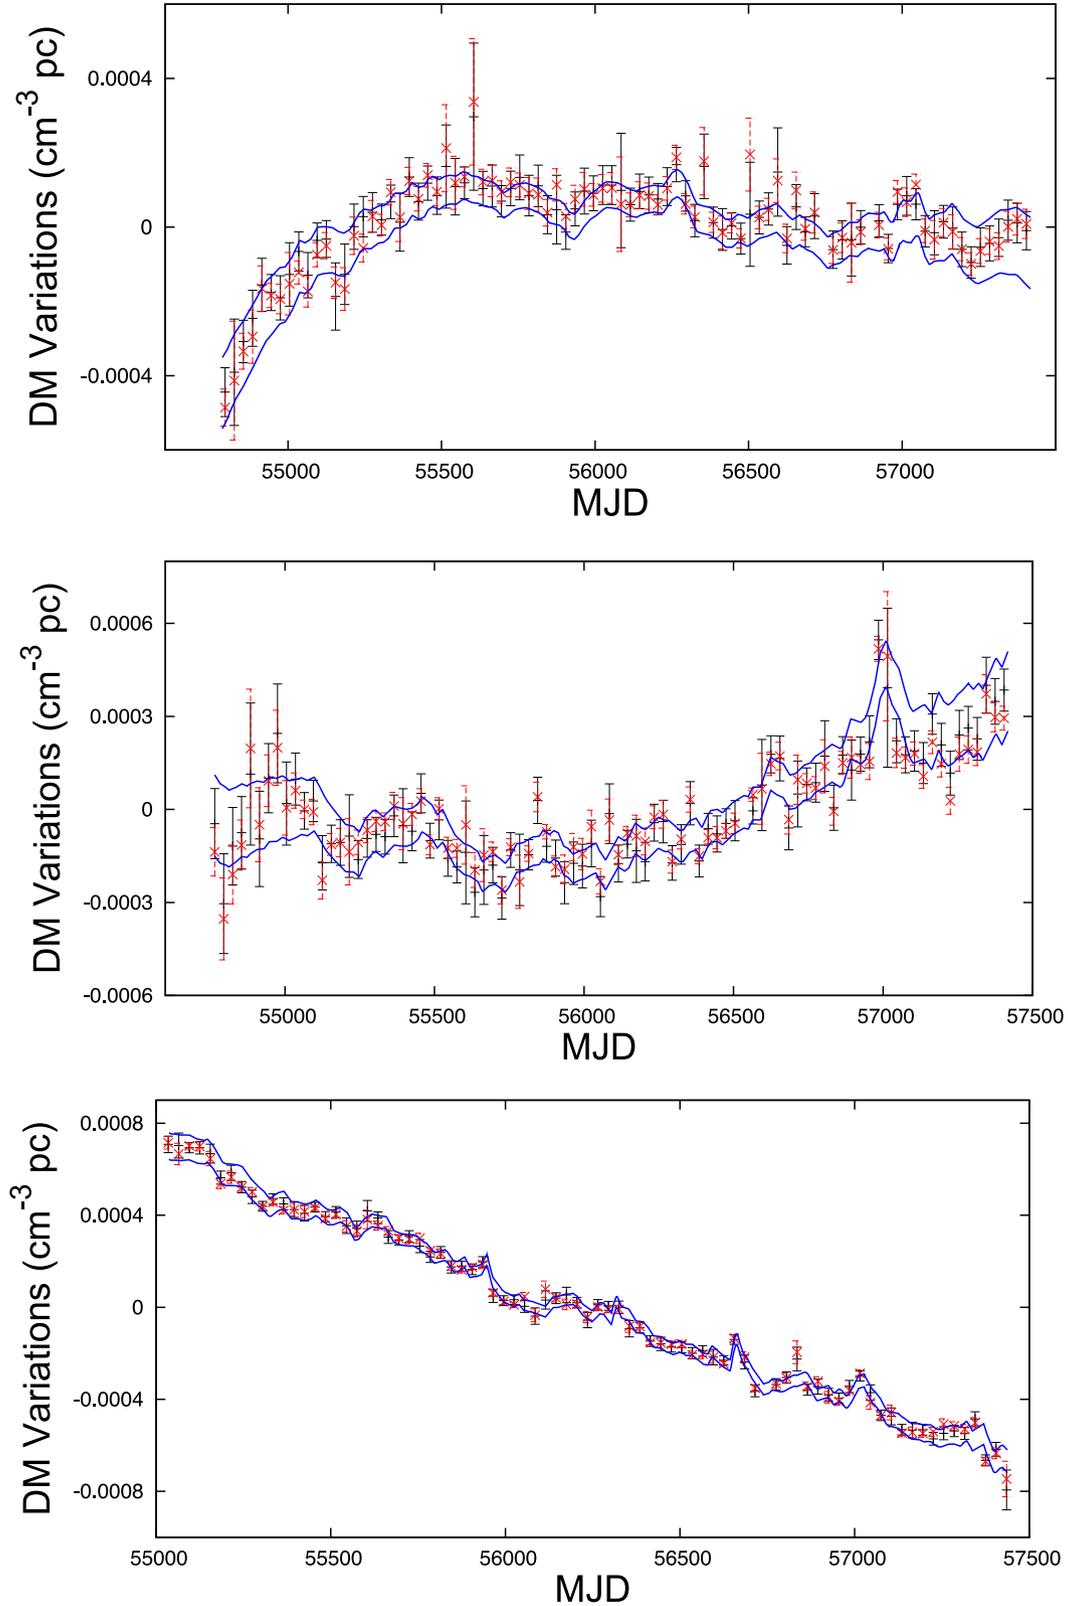


Figure 11. A comparison of different models for the DM variations in PSRs J1713+0747 (top panel), J1744-1134 (middle panel) and J1909-3744 (bottom panel) data sets. We compare the DMX parametrization with a 30 d window from a ToA analysis (black points with 1σ uncertainties), and from a profile domain analysis (red points with 1σ uncertainties), and find the profile domain analysis consistently results in improved constraints on DM by an average of 40 per cent. We also compare the result of using a time-stationary smooth power-law model (blue lines representing the 1σ confidence interval) and find that it is consistent with the DMX model.

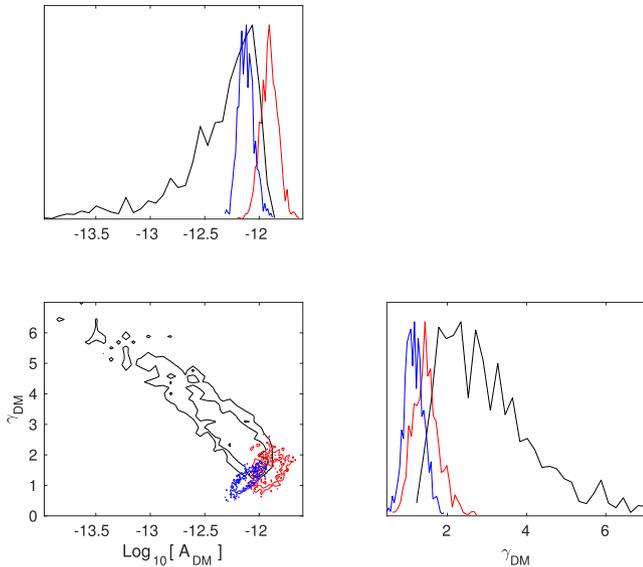


Figure 12. One- and two-dimensional marginalized posterior probability distributions for log amplitude and spectral index of the model for power-law DM variations in PSRs J1713+0747 (black lines), J1744–1134 (red lines) and J1909–3744 (blue lines).

corresponds to having an additional four years of data of equivalent quality.

As with the simulations, we also perform an analysis using the ToAs formed from the fully frequency-averaged profile data. In this case, we use a separate template for each band that does not evolve across the band. For PSR J1909–3744 and J1744–1134, we find

a decrease in the measured precision of the timing parameters of up to 20 per cent, with a mean decrease of 15 per cent, compared to using the sub-band ToAs.

For PSR J1713+0747, we find a much greater loss of precision, of approximately a factor of 2, when using the fully frequency-averaged ToAs. In this case, there is significantly more scatter present in the residuals, and we find the evidence supports an additional EQUAD term with an amplitude of approximately 200 ns, consistent with the properties of the data set used in Shannon et al. (2015). This is consistent with the effect of combining profile evolution and scintillation seen in Simulation 2 in Section 4.

Finally, we compare the parameter estimates obtained from our profile domain analysis both with an analysis of the corresponding ToAs and with other independent published results.

In the bottom panel of Fig. 13, we show the difference between the mean timing model parameter estimates for the ToA and profile domain analyses in terms of their uncertainties. For clarity, we exclude DM from this plot, as its value is highly dependent on the model used for the DM variations, and whether profile evolution has been incorporated. We find that the remaining parameters are all consistent within their uncertainties. In the following two sections, we also compare our profile domain results with parameters published in Reardon et al. (2016, hereafter R16), Desvignes et al. (2016, hereafter D16), Fonseca et al. (2016, hereafter F16) and Matthews et al. (2016, hereafter M16).

6.5.1 Comparison of astrometric parameters

We first consider the astrometric parameters, parallax and proper motion. Comparing parallax measurements, we find that we are

Table 5. Timing model parameter estimates.

	J1713+0747	J1744–1134	J1909–3744
Measured quantities			
Right ascension, α (hh:mm:ss)	17:13:49.532 7239(11)	17:44:29.405 787(3)	19:09:47.434 6730(12)
Declination, δ (dd:mm:ss)	+07:47:37.497 94(4)	–11:34:54.681 36(19)	–37:44:14.466 74(5)
Pulse frequency, ν (s^{-1})	218.811 840 378 344 90(13)	245.426 119 713 0575(11)	339.315 687 288 2423(5)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-4.083\ 91(3) \times 10^{-16}$	$-5.381\ 57(7) \times 10^{-16}$	$-1.614\ 806(3) \times 10^{-15}$
Proper motion in right ascension, $\mu_\alpha \cos \delta$ (mas yr $^{-1}$)	4.919(4)	18.801(8)	–9.519(3)
Proper motion in declination, μ_δ (mas yr $^{-1}$)	–3.926(7)	–9.43(4)	–35.770(9)
Parallax, π (mas)	0.77(4)	2.39(7)	0.870(18)
Orbital period, P_b (d)	67.825 130 973(3)	–	1.533 449 474 28(3)
Epoch of periastron, T_0 (MJD)	51 997.5797(4)	–	–
Longitude of periastron, ω_0 ($^\circ$)	176.1975(17)	–	–
Orbital eccentricity, e	$7.494\ 09(7) \times 10^{-5}$	–	–
First derivative of orbital period, \dot{P}_b	–	–	$5.18(7) \times 10^{-13}$
Companion mass, M_c (M_\odot)	0.277(14)	–	0.2074(20)
Longitude of ascending node, Ω ($^\circ$)	87(3)	–	–
Orbital inclination angle, i ($^\circ$)	72.4(8)	–	–
Sine of inclination angle, $\sin i$	–	–	0.998 18(9)
Projected semimajor axis of orbit, x (lt-s)	32.342 421 45(16)	–	1.897 991 09(4)
TASC (MJD)	–	–	53 113.950 742 05(3)
EPS1	–	–	$2.3(16) \times 10^{-8}$
EPS2	–	–	$-1.04(9) \times 10^{-7}$
Set quantities			
Epoch of position determination (MJD)	54 500	54 500	54 500
Epoch of frequency determination (MJD)	56 100	54 500	54 500
Derived quantities			
Distance from parallax (kpc)	1.30(6)	0.419(11)	1.16(2)
Distance from \dot{P}_b (kpc)	–	–	1.174 (16)

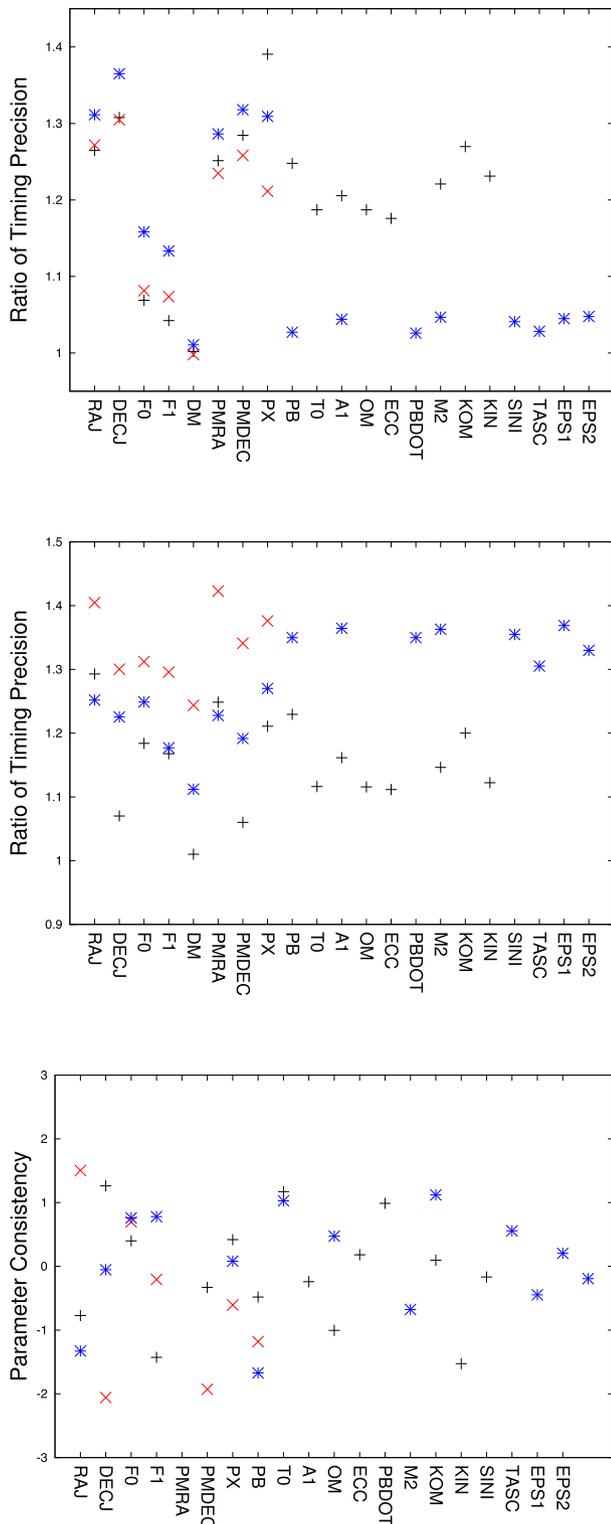


Figure 13. Ratio of the measured precision for the timing model parameters for PSRs J1713+0747 (black +), J1744–1134 (red x) and J1909–3744 (blue *) for the DMX and power-law models for DM variations (top panel) and for the ToA and profile domain analysis (middle panel). Larger numbers correspond to smaller uncertainties in the power-law model (top panel) and in the profile domain analysis (middle panel). Bottom panel: difference between the mean timing model parameter estimates for the ToA and profile domain analyses in terms of their uncertainties.

consistent with previously published values within 2σ confidence intervals for both PSRs J1744–1134 and J1909–3744;³ however, small discrepancies between the measurements can be seen for PSR J1713+0747.

For PSR J1713+0747, we find our parallax measurement lies 1σ , 1.6σ and 2.6σ below the values reported in R16 ($\pi = 0.86(9)$ mas), M16 ($\pi = 0.85(3)$ mas) and D16 ($\pi = 0.90(3)$ mas), respectively. In this case, a further independent check is provided by VLBA measurements (Chatterjee et al. 2009), which result in a value of $\pi = 0.95(6)$ mas, which lies 2.5σ above our value. While none of these offsets are extremely significant, in M16, inconsistency with previously reported values for parallax was attributed to insufficient modelling of DM variations in those earlier publications. In order to check the impact of possible ISM effects, we split our data set up into a 10 cm only, and a combined 20- and 40-cm data set, and measure the parallax in each of these separately. We obtain values of $\pi = 0.80(8)$ mas for the combined 20- and 40-cm data set, which is consistent with both the results from pulsar timing and the VLBA measurement at the 1.5σ level, and a value of $\pi = 0.71(5)$ mas for the 10-cm data, which is consistent at the 3σ level with VLBA. Given that the 10-cm data are unlikely to suffer significantly from mismodelling of the ISM, and has little RFI or other known sources of systemic noise, we conclude that this marginally lower value is simply the result of random variation, rather than any systematic offset.

We next compare measurements of proper motion for these three pulsars. Overall, proper motions are consistent between M16, D16, R16 and the values we report from our profile domain analysis. We are only inconsistent at greater than 2σ for measurements of the proper motion in declination for PSR J1744–1134 obtained in M16 ($\mu_\delta = -9.20(8)$ mas yr⁻¹, 2.6σ) and proper motion in right ascension for PSR J1713+0747 compared to the value in D16 ($\mu_\alpha = -3.888(14)$ mas yr⁻¹, 2.4σ). As for parallax, VLBA provides measurements of proper motion for PSR J1713+0747 of $\mu_\delta = -3.67(18)$ mas yr⁻¹ and $\mu_\alpha = 4.75_{-7}^{+17}$ mas yr⁻¹ that we are consistent with at the 1.0 , and 1.6σ level.

6.5.2 Comparison of binary parameters

We find the majority of binary parameters for PSRs J1909–3744 and J1713+0747 are consistent with those published in R16, D16 and F16. For example, measurements of the companion mass for PSR J1909–3744 were given as $M_c = 0.2067(19)$, $0.213(3)$ and $0.214(3) M_\odot$ in R16, D16 and F16, respectively, which are consistent with the value of $M_c = 0.2074(20) M_\odot$ given in Table 5 to within 2σ . Small differences are seen only in the inclination angle for PSR J1909–3744 compared to the value reported in D16 ($\sin i = 0.99771(13)$, 3σ).

In Lentati et al. (2016), it was shown that unmodelled systematic or frequency-dependent effects can significantly bias the parameter estimates for the stochastic signals in pulsar timing data sets. Neither in this paper, nor in M16, D16 or R16 have such models been incorporated into the analysis, and so it is possible that inconsistencies at the $\sim 3\sigma$ level could be due to systematic or frequency-dependent effects that are currently unmodelled in these different analyses.

³ The uncertainty on the parallax in R16 has a typographical error. The corrected value is $0.81(3)$, which is consistent with our measurement at the 2σ level (Reardon, private communication).

One of the key advantages of the profile domain framework, however, is that eventually we will be able to integrate more physical models for these effects into the analysis.

7 CONCLUSIONS

We have extended the profile domain analysis framework ‘Generative Pulsar Timing Analysis’ to incorporate wide-band effects such as profile evolution and broad-band shape variation in the pulse profile. We also incorporate models for epoch-to-epoch variation in both pulse width and in the separation in phase of the main pulse and interpulse. This framework allows for a full timing analysis to be performed using the profile data, rather than forming estimates of the pulse times of arrivals as in the standard pulsar timing paradigm, thereby enabling simultaneous estimation of, for example, DM variations, pulse shape variation, profile evolution and the pulsar timing model.

In order to handle the large number of profiles ($\sim 20\,000$), and the large number of parameters (> 1000) that must be included in an analysis of wide-band data sets, we used HMC sampling methods and introduced a new parametrization that is significantly more efficient in the low S/N regime for many of the stochastic parameters included in the model. While we only used this likelihood here in the context of pulsar timing analysis, its potential applications are extremely wide ranging.

We applied this analysis framework first to a series of simulations and then to over 7 yr of observations for PSRs J1713+0747, J1744–1134 and J1909–3744 made between 700 and 3600 Mhz. We found that in all but the most trivial simulation, the profile domain analysis resulted in higher precision compared to a standard timing analysis, with an average improvement of 25 per cent. For timing parameters that scale as the square-root of observing time, such as parallax, one would need an additional 4 yr of data with equivalent timing precision for a 7-yr data set to achieve the same improvement.

We compared a non-time-stationary piecewise model for the variations in the DM with a time-stationary power-law model for each data set and found that they were consistent in each case. By incorporating the inherent smoothness in the DM variations into the model, however, the measurement precision obtained for the astrometric timing parameters was found to improve by as much as 40 per cent, and the precision of the binary parameters for PSR J1713+0747 was found to improve by 20 per cent.

We also detected significant wide-band pulse shape variation in the PSR J1713+0747 data set at the level expected for hour-long integrations given previously published estimates of the intrinsic variation in the pulsar (e.g. Arzoumanian et al. 2015). Not accounting for this shape variation changes the measured arrival times at the level of ~ 30 ns, the same order of magnitude as the expected shift due to gravitational waves in the pulsar timing band.

As new instrumentation that continues to increase the simultaneous bandwidth of observations comes online, the ability to deal with effects that combine coherently over the band, such as intrinsic pulse shape variation and profile evolution, will continue to become more important. The profile domain framework developed in this work and preceding papers provides a statistically robust approach to simultaneously incorporating these effects into pulsar timing analysis. By directly modelling the effect of shape change, one can improve the sensitivity of the analysis to any process that results in a shift of the arrival time. Such processes include post-

Keplerian effects in highly relativistic systems (Kramer et al. 2006), changes in the binary period, which provide tests of the potential evolution in the gravitational constant G (Zhu et al. 2015) and the passage of nHz gravitational waves. As both models and sampling efficiency improve, and the analysis of progressively more complex data sets becomes tractable, profile domain analysis will ultimately provide the optimal approach to performing high-precision pulsar timing analysis.

ACKNOWLEDGEMENTS

The Parkes radio telescope is part of the Australia Telescope National Facility that is funded by the Commonwealth of Australia for operation as a National Facility managed by Commonwealth Science and Industrial Research Organization (CSIRO). We thank all of the observers, engineers and Parkes observatory staff members who have assisted with the observations reported in this paper.

LL was supported by a Junior Research Fellowship at Trinity Hall College, Cambridge University.

REFERENCES

- Arzoumanian Z. et al., 2015, *ApJ*, 813, 65
 Chatterjee S. et al., 2009, *ApJ*, 698, 250
 Cognard I., Theureau G., Guillemot L., Liu K., Lassus A., Desvignes G., 2013, in Cambresy L., Martins F., Nuss E., Palacios A., eds, SF2A-2013: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, p. 327
 Coles W. A. et al., 2015, *ApJ*, 808, 113
 Cordes J. M., 1978, *ApJ*, 222, 1006
 Cordes J. M., Shannon R. M., Stinebring D. R., 2016, *ApJ*, 817, 16
 Dai S. et al., 2015, *MNRAS*, 449, 3223
 Dai S., Johnston S., Bell M. E., Coles W. A., Hobbs G., Ekers R. D., Lenc E., 2016, *MNRAS*, 462, 3115
 Demorest P. B. et al., 2013, *ApJ*, 762, 94
 Desvignes G. et al., 2016, *MNRAS*, 458, 3341 (D16)
 Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
 Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
 Fonseca E. et al., 2016, *ApJ*, 832, 167 (F16)
 Handley W. J., Hobson M. P., Lasenby A. N., 2015, *MNRAS*, 453, 4384
 Hankins T. H., Cordes J. M., 1981, *ApJ*, 249, 241
 Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, *Nature*, 217, 709
 Keith M. J. et al., 2013, *MNRAS*, 429, 2161
 Kramer M. et al., 2006, *Science*, 314, 97
 Lee K. J. et al., 2014, *MNRAS*, 441, 2831
 Lentati L., Alexander P., Hobson M. P., Taylor S., Gair J., Balan S. T., van Haasteren R., 2013, *Phys. Rev. D*, 87, 104021
 Lentati L., Hobson M. P., Alexander P., 2014, *MNRAS*, 444, 3863
 Lentati L., Shannon R. M., 2015, *MNRAS*, 454, 1058 (L15b)
 Lentati L., Alexander P., Hobson M. P., 2015, *MNRAS*, 447, 2159 (L15a)
 Lentati L. et al., 2016, *MNRAS*, 458, 2161
 Liu K. et al., 2014, *MNRAS*, 443, 3752
 Manchester R. N., 2015, *IAU General Assembly*, 22, 2256190
 Manchester R. N. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, 17
 Matthews A. M. et al., 2016, *ApJ*, 818, 92 (M16)
 Narayan R., 1992, *Phil. Trans. R. Soc. A*, 341, 151
 Nelder J. A., Mead R., 1965, *Comput. J.*, 7, 308
 Osłowski S., van Straten W., Hobbs G. B., Bailes M., Demorest P., 2011, *MNRAS*, 418, 1258
 Pennucci T. T., Demorest P. B., Ransom S. M., 2014, *ApJ*, 790, 93

Ransom S. M., Demorest P., Ford J., McCullough R., Ray J., DuPlain R.,
Brandt P., 2009, Am Astron. Soc. Meeting Abstr., 214, 605.08
Reardon D. J. et al., 2016, MNRAS, 455, 1751 (R16)
Refregier A., 2003, MNRAS, 338, 35
Shannon R. M. et al., 2014, MNRAS, 443, 1463
Shannon R. M. et al., 2015, Science, 349, 1522
Shannon R. M. et al., 2016, ApJ, 828, L1

Stappers B. W. et al., 2011, A&A, 530, A80
Taylor J. H., 1992, Phil. Trans. R. Soc. A, 341, 117
Taylor J. F., Ashdown M. A. J., Hobson M. P., 2008, MNRAS, 389, 1284
Zhu W. W. et al., 2015, ApJ, 809, 41

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