

Robust estimation of scattering in pulsar timing analysis

L. Lentati,¹★ M. Kerr,² S. Dai,² R. M. Shannon,^{2,3} G. Hobbs² and S. Osłowski^{4,5,6}

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

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

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

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

⁵*Fakultät für Physik, Universität Bielefeld, Postfach 100131, D-33501 Bielefeld, Germany*

⁶*Max Planck Institute for Radio Astronomy, Auf dem Hügel 69, D-53121 Bonn, Germany*

Accepted 2017 March 6. Received 2017 March 6; in original form 2016 December 20

ABSTRACT

We present a robust approach to incorporating models for the time-variable broadening of the pulse profile due to scattering in the ionized interstellar medium into profile-domain pulsar timing analysis. We use this approach to simultaneously estimate temporal variations in both the dispersion measure (DM) and scattering, together with a model for the pulse profile that includes smooth evolution as a function of frequency, and the pulsar’s timing model. We show that fixing the scattering time-scales when forming time-of-arrival estimates, as has been suggested in the context of traditional pulsar timing analysis, can significantly underestimate the uncertainties in both DM and the arrival time of the pulse, leading to bias in the timing parameters. We apply our method using a new, publicly available, GPU-accelerated code, both to simulations and observations of the millisecond pulsar PSR J1643–1224. This pulsar is known to exhibit significant scattering variability compared to typical millisecond pulsars, and we find including low-frequency (<1 GHz) data without a model for these scattering variations leads to significant periodic structure in the DM, and also biases the astrometric parameters at the 4σ level, for example, changing proper motion in right ascension by 0.50 ± 0.12 . If low-frequency observations are to be included when significant scattering variations are present, we conclude it is necessary to not just model those variations, but also to sample the parameters that describe the variations simultaneously with all other parameters in the model, a task for which profile domain pulsar timing is ideally suited.

Key words: methods: data analysis – pulsars: general – pulsars: individual: PSR J1643–1224 – ISM: general.

1 INTRODUCTION

The eventual detection of gravitational waves in the nanohertz window using a pulsar timing array (Foster & Backer 1990) will require a thorough understanding of the myriad mechanisms that can impact either the shape or the time of arrival (ToA) of the pulses of light from those pulsars. The ionized interstellar medium (IISM) is known to be the dominant such mechanism in pulsar timing experiments (e.g. Lam et al. 2016), and introduces *both* changes in the shape of the profile and delays in the arrival times.

The delays in the arrival times are primarily the result of interstellar dispersion. As the pulse propagates through the ionized plasma that makes up the IISM, it interacts with free electrons causing a frequency dependent delay. This delay is proportional the integrated column density of electrons along our line of sight to the pulsar,

called the dispersion measure (DM), and scales as ν^{-2} , with ν the observing frequency. As our line of sight to the pulsar changes with time, so too does the observed column density, introducing a time-variable delay in arrival times.

Over the last several years significant progress has been made in modelling these variations in DM, and in propagating our uncertainties in that model through to the timing parameters, with multiple different methods in use by groups around the world (e.g. Demorest et al. 2013; Keith et al. 2013; Lentati et al. 2014).

A more subtle effect due to the IISM is that of scattering, where inhomogeneities in the IISM cause both intensity variations (called diffractive scintillation, e.g. Narayan 1992) and also a broadening of the pulse profile. The impact of scattering, as for DM variations, will also change with time as our line of sight to the pulsar changes, and in this paper we will be concerned with how to robustly incorporate scattering variations into pulsar timing analysis.

Given the complexity of the IISM, and the steep frequency dependence of the impact it has on pulsar timing, one can rightly

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

ask why we should not just migrate to high-frequency observations alone. Indeed, this approach has already resulted in the most sensitive limit on an isotropic gravitational wave background using a pulsar timing array to date (Shannon et al. 2015). The simple answer is that pulsars are known to have steep spectral indices ($S_\nu \propto \nu^{-1.8}$ on average, Maron et al. 2000), and so are much brighter at lower observing frequencies. Given the strong dependence of the detection probability of gravitational waves using a pulsar timing array on the number of pulsars in the array (e.g. Taylor et al. 2016), being able to include more pulsars by virtue of observing at lower frequencies where they are brightest has clear benefits.

Previous attempts to model scattering have either made significant assumptions about a priori unknown quantities, such as the intrinsic pulse shape, or only consider scattering independent of other aspects of the model such as the intrinsic variability in the pulse shape, or DM variations. In both cases, the result is a failure to propagate the uncertainties from these other parameters into the scattering measurements.

For example, Löhmer et al. (2001) use profile data from 4.9 GHz observations to construct a model for the intrinsic pulse profile, and then assume this model when determining the scattering time-scales and scaling with frequency. However, significant profile evolution is known to occur across wide frequency ranges that can bias these parameter estimates. In this case, the uncertainties on the measured parameters were multiplied by a factor of 3; however, this is clearly unsatisfactory, and a more statistically robust approach would be preferred. In Bhat, Cordes & Chatterjee (2003), no assumptions are made about the intrinsic shape of the pulse profile; however, the DM is assumed fixed, and the CLEAN algorithm is sub-optimal when the scattering time-scale is small.

Further complications arise when wanting to include scattering variations in timing analysis. Existing approaches have suggested obtaining an estimate for the scattering time-scale, and then using that to ‘correct’ the ToAs (e.g. Levin et al. 2016). While cyclic spectroscopy has shown significant promise, the robustness of the method breaks down for low signal-to-noise ratio (S/N) profiles, where the S/N is less than ~ 100 (Palliyaguru et al. 2015). Even in the high-S/N simulations; however, the approach advocated is to ‘correct’ the ToAs using estimates of the scattering delay obtained from the cyclic spectroscopy. In the PSR J1643–1224 data set analysed in Section 4, the mean S/N per epoch is ~ 90 , with a minimum of only 15. We therefore require an approach that is robust across all observed profile S/N. Further, neither approach accounts for the significant covariances that exist between the scattering time-scales and other parameters of interest.

We illustrate two of the key challenges associated with incorporating scattering measurements into timing analysis in Fig. 1. We simulate a pulsar with a 4.6 ms period, using a Gaussian pulse shape with a full-width at half-maximum of 5 per cent pulse phase. The scattering time-scale is chosen to be $10^{-4.5}$ s at a reference frequency of 1 GHz, consistent with the values observed in our analysis of PSR J1643–1224 in Section 4. We use a frequency range of 1.25–1.5 GHz, separated into eight channels, again chosen to be consistent with a typical 20 cm observation used in Section 4. In the top left-hand panel, we show the simulated pulse profile at the lowest frequency channel (red lines, 1263 MHz), and the highest frequency channel (black lines, 1475 MHz). We have aligned the leading edge of the profiles (shown as a zoomed-in region in the bottom sub-panel), and can see that the trailing edge has been broadened in the low-frequency channel (shown as a zoomed-in region in the top sub-panel). In the top right-hand panel, we show

the same thing, however after adding noise to the simulation. By eye, the scattering is undetectable; however, we will now show that it is still sufficient to significantly bias parameter estimates when ignored in the analysis.

In the bottom left-hand panel, we show in black the one- and two-dimensional marginalized posterior distributions from a three-dimensional analysis of this simulated observation. The parameters included in the analysis are as follows:

- (i) the ToA, measured in units of phase;
- (ii) the DM in units of the TEMPO2 uncertainty;
- (iii) \log_{10} of the scattering time-scale, τ , measured in seconds at a reference frequency of 1 GHz.

One can clearly see that the scattering time-scale is correlated in a non-linear way with both the pulse ToA and the DM. The ToA in particular has a bi-modal distribution, with one peak associated with ‘large’ scattering time-scales ($\log_{10} \tau \gtrsim -5$), and one associated with small scattering time-scales ($\log_{10} \tau \lesssim -5$). Note that the peak of the one-dimensional probability distribution for the scattering time-scale is consistent with the simulated value, indicating that despite being practically invisible by eye, it is still marginally detected in our analysis.

In the same panel, we show the posterior distributions for the phase and DM when fixing the scattering time-scale at 10^{-10} (red lines) and $10^{-4.4}$ (blue lines). These correspond to models with no scattering, or using the maximum-likelihood scattering time-scale obtained from the three-dimensional analysis. In both cases, this results in significant bias in both the measured values and the uncertainties of the other two parameters. In particular, the measured ToA is completely inconsistent between the two models, with a difference of approximately 10 μ s, two to three orders of magnitude more than the shift expected from an isotropic gravitational wave background in pulsar timing observations given the most stringent current upper limits (Shannon et al. 2015). While the frequency dependence of scattering variations will help to decrease their covariance with gravitational waves, and so will limit how much this bias truly impacts our sensitivity, clearly leaving such variations unmodelled in the analysis must be sub-optimal.

In Fig. 1 (bottom right-hand panel), we then compare the posterior distribution for the ToA from the three-dimensional analysis (black line) to a Gaussian approximation of the posterior (red line). Typically when forming ToAs, a Gaussian approximation is made of the true probability density function (PDF) of the arrival time. While methods such as those presented in Liu et al. (2014) and Pennucci, Demorest & Ransom (2014), which perform a simultaneous fit to broad-band profile data in order to fit for both phase and DM can be extended to also fit for scattering, the fundamental approximation of Gaussianity implicit in the ToA forming process is invalid in regimes such as that shown in Fig. 1. While the Gaussian approximation is conservative, it is sub-optimal compared to incorporating the full PDF of the arrival time, and while in principle the full posterior distribution from this kind of analysis could be used in subsequent pulsar timing analysis, such an approach will always be an approximation to simply performing the analysis in the profile domain.

In this paper, we present a solution to these problems by extending the profile domain pulsar timing framework (Lentati et al. 2017, and references therein, henceforth L16) to incorporate scattering variations. This means that the parameters that describe the scattering as a function of time can be estimated simultaneously with all other parameters in the model. This includes DM variations, the

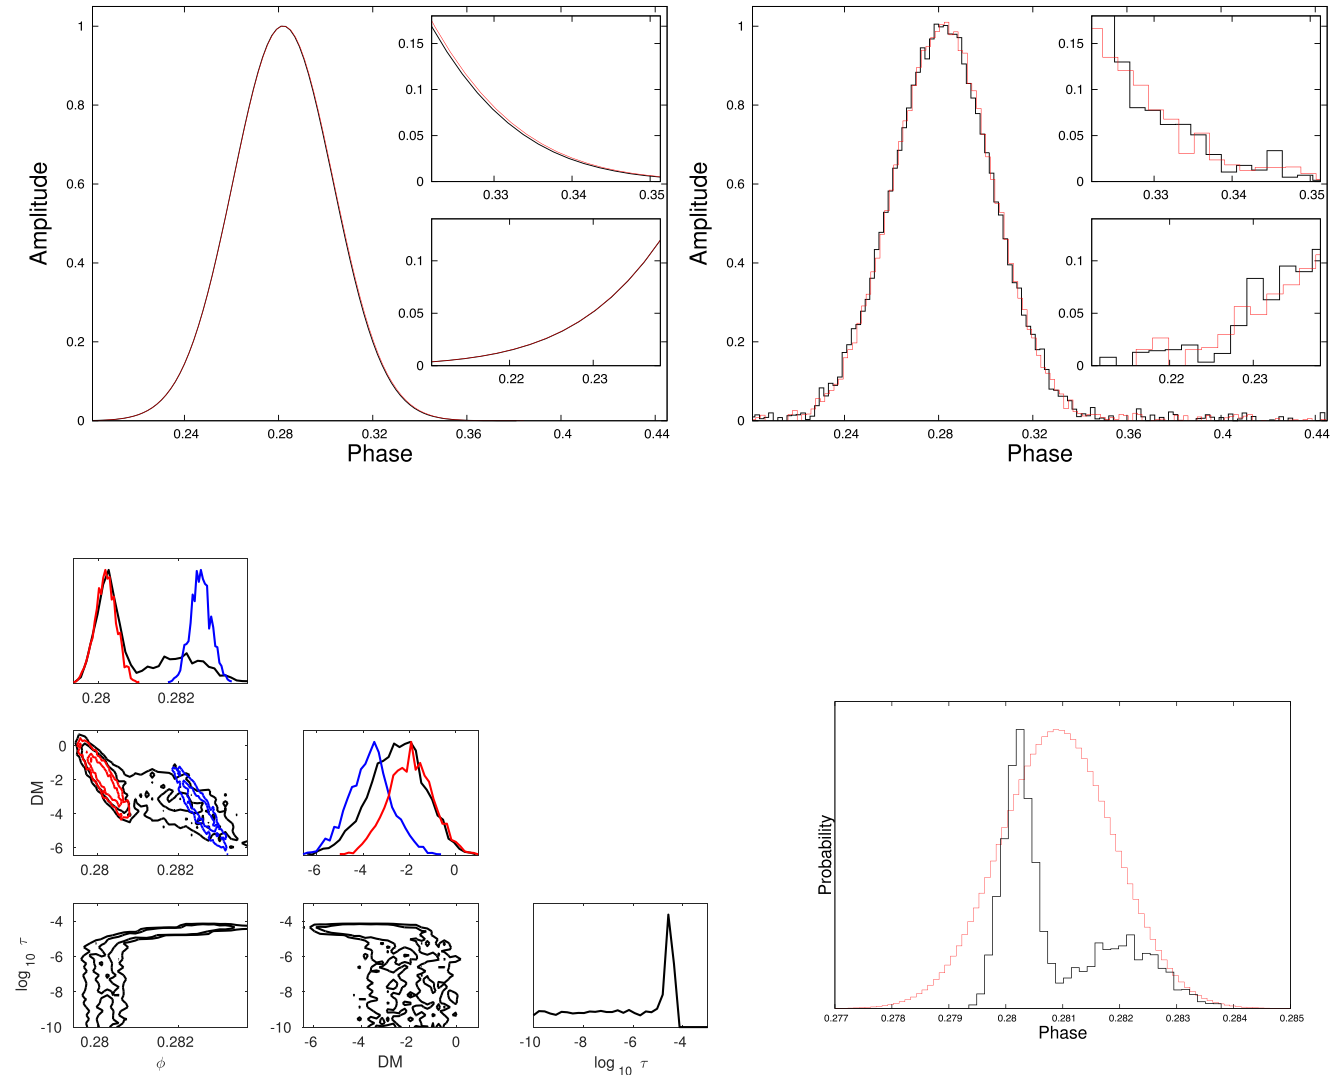


Figure 1. Top panels: simulated data before adding noise (left-hand panel) and after (right-hand panel) for a Gaussian pulse profile with a scattering time-scale of $10^{-4.5}$ s at a reference frequency of 1 GHz. Sub-panels show zoomed-in regions of the trailing (top) and leading edges (bottom). Bottom left-hand panel: one- and two-dimensional marginalized posterior distributions from a three-dimensional analysis of the simulated observation (black lines). We fit simultaneously for the ToA, the DM and \log_{10} of the scattering time-scale. We compare these parameter estimates with those obtained from a two-dimensional analysis where we fix the scattering time-scale at 10^{-10} (red lines) and $10^{-4.4}$ (blue lines) chosen to represent no scattering, and the maximum-likelihood time-scale from the three-dimensional analysis. Bottom right-hand panel: comparison of the true posterior probability distribution for the ToA from the three-dimensional analysis compared to a Gaussian approximation, as would be performed in traditional pulsar timing analysis.

pulsar timing model, timing noise, pulse jitter, and also any models for the profile, profile evolution as a function of frequency and pulse shape instability.

In order to incorporate scattering into the existing framework, we perform the full analysis in the Fourier domain, as opposed to the time domain as in L16. In Section 2, we describe this modified framework, before applying it to simulated, and then real data in Sections 3 and 4, before ending with some concluding remarks in Section 5. All the analysis presented subsequently is performed using a new, publicly available GPU-accelerated code,¹ and we have included the simulated data sets used in Section 3 in the repository with instructions for how to replicate the results presented in this paper.

¹ <https://github.com/LindleyLentati/TempoNest2>

2 A FOURIER-DOMAIN MODEL

In the following section, we will build on the methods presented in L16. In that previous work, the profile-domain analysis was performed in the time-domain. Here, we will define our models in the Fourier-domain which allows us to optimize the analysis in two ways. First, the Fourier representation makes it trivial to truncate our model for the profile at a particular harmonic, decreasing the size of the linear-algebra operations required to evaluate it. Secondly, when including scattering in the analysis, we only need to multiply the pulse broadening function (PBF) with our profile model in the Fourier domain, as opposed to computing a convolution in the time domain, and the gradients can likewise be obtained at less computational expense for all parameters.

For a full description of the general profile domain framework, we refer the reader to L16. Below we will give details of how the

methodology has been changed to allow analysis in the Fourier domain, and of our implementation of models for scattering.

2.1 The profile model

As in L16, we construct our profile model using the shapelet basis (Refregier 2003), for which the Fourier-transform can be obtained analytically. In the time domain, shapelets are described by a position t , a scale factor Λ and a set of n_{\max} amplitude parameters, with which we can construct the set of basis functions:

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

with $\Phi_n(t)$ given by

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

where H_n is the n th Hermite polynomial. In the Fourier domain, the dimensionless basis functions $\Phi_n(t)$ become

$$\tilde{\Phi}_n(f) = i^n \Phi_n(f) \quad (3)$$

and the basis functions $B_n(t; \Lambda)$ become

$$\tilde{B}_n(f; \Lambda) = i^n B_n(f; \Lambda^{-1}) \quad (4)$$

for Fourier frequencies f . We can then write our profile model in the Fourier domain, $\tilde{s}(f, \zeta, \Lambda)$, as the sum

$$\tilde{s}(f, \zeta, \Lambda) = \sum_{n=0}^{n_{\max}} \zeta_n(\nu) \tilde{B}_n(f; \Lambda^{-1}), \quad (5)$$

where ζ_n are the shapelet amplitudes and n_{\max} is the number of shapelet basis vectors included in the model.

As in L16, we have explicitly written the shapelet amplitudes as a function of the observing frequency ν , and use a general polynomial expansion of the shapelet amplitudes with frequency in order to model any potential smooth profile evolution. This model is defined such that for the p terms in the polynomial, we can write the n th shapelet amplitude $\zeta_n(\nu)$ as

$$\zeta_n(\nu) = \sum_{k=0}^p (\nu - \nu_c)^k \zeta_{n,k}(\nu), \quad (6)$$

where ν_c is an arbitrary reference frequency and $\zeta_{n,k}$ is the amplitude parameter for the k th polynomial of the n th term in the shapelet model.

Finally, as in L16, we use the shapelet basis to describe the overall profile shape, and then scale this in amplitude for each epoch. We therefore use the zeroth-order term as a reference, taking $\zeta_0 = 1$, leaving only $n_{\max} - 1$ free parameters ζ_n , which are the amplitudes for the shapelet components with $n > 0$. Written in this way equation (5) becomes

$$\tilde{s}(f, A, \zeta, \Lambda) = A \sum_{n=0}^{n_{\max}} \zeta_n(\nu) \tilde{B}_n(f; \Lambda), \quad (7)$$

with A the overall scaling factor for a particular epoch.

2.2 Shapelet interpolation

As in L16, we do not re-evaluate our shapelet model for every likelihood calculation, but adopt an interpolation scheme, where the shapelet basis vectors are pre-computed on a grid from $t = 0$ up to the duration of the longest phase bin in the data set. In principle, as we are performing our analysis in the Fourier domain, we can rotate

the shapelet model exactly by multiplying our complex shapelet model by a rotation vector:

$$R(f, \delta\phi) = \exp(2\pi i f \delta\phi), \quad (8)$$

where $\delta\phi$ is the amount of phase by which to rotate the model. However, when performing the sampling with Hamiltonian Monte Carlo (HMC) methods, we must compute the gradients for all the shapelet amplitudes. As such, if we rotate the profile model to match the data, we must also rotate all the basis vectors for each profile in the data set. We find it is more efficient to pre-compute sub-bin shifts in the profile model, and then to rotate the data by an integer number of bins, again using equation (8) so that it aligns with the nearest interpolated set of basis functions. In this way, we need to perform only a single rotation per profile, which is a significant computational saving.

As in L16, we use an interpolation interval of 1 ns, chosen to be sufficiently small that no bias enters our analysis as a result of the interpolation process.

2.3 Truncating the shapelet model

In addition to the interpolation scheme described in Section 2.2, the Fourier representation makes it straightforward to truncate the shapelet model at a particular harmonic. This means that we need to include only the subset of the harmonics in the profile model that contribute above some threshold in the analysis. While the interpolated shapelet model used in L16 significantly reduced the time taken to evaluate the mean profile compared to numerical evaluation, it still required an $N_b \times N_c$ matrix-vector product, where N_c is the number of amplitudes in the shapelet model and N_b is the number of bins in the profile data. By only including N_h harmonics in the profile model, we can reduce this to a $2N_h \times N_c$ product, providing an immediate decrease in the computation time of the profile model by a factor of $N_b/2N_h$.

In Fig. 2, we compare the time-domain (left-hand panel), and Fourier-domain (right-hand panel) representations of the time-averaged profile for PSR J1643–1224, using the 1200 to 1500 MHz data that we analyse in Section 4. While the time-domain profile extends across ~ 300 phase bins, the Fourier representation contains over 99 per cent of the total signal in only 40 harmonics, representing a significant reduction in the size of the matrix-vector multiplication required to evaluate the model. In Section 4, we include up to the 80th harmonic in our model. Beyond this, the relative contribution of higher harmonics in the profile model is less than one part in 10^{10} .

In principle, variations in the shape of the pulse profile may introduce fluctuations at harmonics higher than that required for the mean profile model. In the analysis presented here, such fluctuations will be absorbed by the white noise component of our model. In principle, however, one could compare models that incorporate these higher frequency terms and determine the optimal maximum harmonic to include in the analysis at the expense of increased computation time.

2.4 Scattering

The primary addition to the model used in our analysis in this work, as opposed to L16, is broadening of the pulse profile due to scattering as the pulse passes through the IISM. In the time-domain approach, pulse broadening would have to be included as a convolution of the pulse profile with a PBF. Typically, this is performed by performing a fast Fourier transform (FFT) of the model and the PBF, multiplying them in the Fourier domain, and




Figure 2. Time-averaged profile for PSR J1643–1224, using the 1200–1500 MHz data that we analyse in Section 4 in the time domain (left-hand panel) and Fourier domain (right-hand panel). While the profile extends across over 300 phase bins, over 99 per cent of the signal is contained in the first 40 Fourier bins.

then transforming back to the time domain. When calculating the gradient of the likelihood for each parameter in the model, these too require Fourier-transforms to be performed, and so the complexity of the likelihood calculation grows rapidly. In the Fourier domain, we need to perform only the multiplication of the Fourier transform of the PBF with our profile model, and the gradients can likewise be obtained trivially for all parameters.

In principle, any model for scattering that allows for the calculation of a gradient can be incorporated into our analysis. Even if the analytic Fourier transform of the PBF is not known, if the gradient can be computed in the time domain, it can then be transformed via FFT for use in our analysis framework, although this will be less computationally efficient. In our analysis in this work, we simply assume a thin screen model for the PBF (Williamson 1972), using a single parameter for the time-scale, τ , which in the time domain will be given by

$$\text{PBF}(t, \bar{\tau}, \nu, \alpha) = H(t) \exp\left(-\frac{t}{\nu^{-\alpha} 10^{\bar{\tau}}}\right), \quad (9)$$

where ν is the observing frequency, and $H(t)$ is the Heavyside step function. Note that as we always require the scattering time-scale to be positive, we fit for \log_{10} of the scattering time-scale, $\bar{\tau}$. The analytic Fourier transform of equation (9) is then given by

$$\text{PBF}(f, \bar{\tau}, \nu, \alpha) = \frac{1}{(2\pi f \nu^{-\alpha} 10^{\bar{\tau}})^2 + 1} + \frac{-2\pi f \nu^{-\alpha} 10^{\bar{\tau}}}{(2\pi f \nu^{-\alpha} 10^{\bar{\tau}})^2 + 1} i, \quad (10)$$

with gradients with respect to $\bar{\tau}$ and α :

$$\frac{d\text{PBF}}{d\bar{\tau}} = \frac{\log(10)}{((\omega \nu^{-\alpha} \bar{\tau})^2 + 1)^2} \times (-2\omega^2 \nu^{-2\alpha} 10^{2\bar{\tau}} + \omega \nu^{-\alpha} 10^{\bar{\tau}} ((10^{\bar{\tau}} \omega \nu^{-\alpha})^2 - 1) i), \quad (11)$$

$$\frac{d\text{PBF}}{d\alpha} = \frac{-\log(\nu)}{((\omega \nu^{-\alpha} \bar{\tau})^2 + 1)^2} \times (-2\omega^2 \nu^{-2\alpha} 10^{2\bar{\tau}} + \omega \nu^{-\alpha} 10^{\bar{\tau}} ((10^{\bar{\tau}} \omega \nu^{-\alpha})^2 - 1) i), \quad (12)$$

where as before ‘ i ’ indicates a complex number and $\omega = 2\pi f$. We can then multiply the Fourier representation of our shapelet model, \tilde{s} , by equation (10) to get the scattered profile model.

In our analysis of both simulated data and the PSR J1643 data set, we use a piecewise-constant $\tau(t)$ model for the scattering. This therefore makes the same assumptions made when modelling DM variations using the DMX parametrization (Demorest et al. 2013). In principle, one could also model the power spectrum of the scattering, under the assumption that it is a smooth, wide-sense stationary process in the same manner as the smooth model for DM variations in L16. Deviations from this simple model, such as those predicted from simulations (Coles et al. 2010), could then be encapsulated as additional, non-stationary features in the model, in the same manner as the extreme scattering events observed in PSR J1713+0747 (e.g. Desvignes et al. 2016; Lentati et al. 2016). However, for this demonstration of the profile domain approach, we do not yet consider such a model.

2.5 Evaluating the profile domain model

The remainder of the timing framework described in L16 requires only minimal changes to operate on the Fourier representation of the profile data. As in L16, we 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 amplitude of the profile as measured at a set of times $t_{i,j}$, which we denote $\mathbf{d}_{i,j}$. The Fourier representation of the profile data, $\tilde{\mathbf{d}}_{i,j}$, will therefore be a set of $N_{i,j}/2 + 1$ complex values. While in L16 a baseline offset for each profile was required as part of the model, in the Fourier representation, we simply discard both the DC offset term and the Nyquist term, so that our final data vector for each profile is of length $N_{i,j}/2 - 1$.

All further parameters, such as those describing pulse jitter or shape variation, can then be incorporated into our Fourier-domain shapelet model \tilde{s} in exactly the same manner as in L16. We can then write the final likelihood that the data are described by the parameters in our model, which we collectively refer to as θ , as

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

Table 1. Timing parameters for PSR J1643–1224.

Scenario	Measured quantities		
	S1	S2	S3
Right ascension (hh:mm:ss)	16:43:38.161512(21)	16:43:38.161572(11)	16:43:38.161510(12)
Declination (dd:mm:ss)	–12:24:58.6724(14)	–12:24:58.6725(8)	–12:24:58.6734(9)
Pulse frequency (s^{-1})	216.373 337 142 644(5)	216.373 337 142 639 (3)	216.373 337 142 639(4)
First derivative of pulse frequency (s^{-2})	$-8.6450(5) \times 10^{-16}$	$-8.6448(3) \times 10^{-16}$	$-8.6450(4) \times 10^{-16}$
Proper motion in right ascension (mas yr^{-1})	6.00(10)	5.50(6)	5.85(8)
Proper motion in declination (mas yr^{-1})	3.8(5)	3.5(3)	3.9(4)
Parallax, π (mas)	1.3(4)	1.3(3)	1.0(3)
Orbital period, P_b (d)	147.017 397 74(6)	147.017 397 75(4)	147.017 397 72(5)
Epoch of periastron, T_0 (MJD)	49283.9336(5)	49283.9337(4)	49283.9336(5)
Projected semi-major axis of orbit, x (lt-s)	25.072 6165(29)	25.072 6182(19)	25.072 6165(23)
Longitude of periastron, ω_0 (deg)	321.8487(13)	321.8489(10)	321.8485(11)
Orbital eccentricity, e	$5.05755(12) \times 10^{-4}$	$5.05757(10) \times 10^{-4}$	$5.05752(11) \times 10^{-4}$
First derivative of x , \dot{x} (10^{-12})	$-5.2(5) \times 10^{-14}$	$-5.5(3) \times 10^{-14}$	$-5.2(4) \times 10^{-14}$
	Set quantities		
Epoch of frequency determination (MJD)	55000	55000	55000
Epoch of position determination (MJD)	55000	55000	55000
Epoch of DM determination (MJD)	55000	55000	55000

where $\mathbf{N}_{i,j}$ is the white noise covariance matrix for the Fourier domain 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 performing the analysis with the guided Hamiltonian sampler (GHS), $\sigma_{i,j}$ is a free parameter in our analysis for every profile.

3 APPLICATION TO SIMULATED DATA

In order to test the efficacy of the analysis method described in the preceding sections, we first apply it to a simulated data set. In particular, we simulate 100 observational epochs, each assuming 256 MHz of bandwidth split into eight channels, with a central frequency of 1369 MHz and covering a total time span of 4.5 yr. Each epoch has an integrated S/N across the observing band of ~ 500 . Both the MJD and the bandwidth of each simulated epoch are taken from the 1400 MHz data used in Section 4. We use a non-evolving, Gaussian profile and a timing model consistent with that given in Table 1 for PSR J1643–1224.

In addition to this timing model, each epoch is subject to a random change in the DM with a standard deviation for the variations of $0.007 \text{ cm}^{-3} \text{ pc}$ and changes in the scattering of $0.81 \pm 0.16 \text{ ms}$ at a reference frequency of 1.4 GHz. The scattering is simulated using the thin screen model described in Section 2.4, with the frequency scaling parameter, α , equal to 4. We then perform a simultaneous analysis of all epochs, fitting for the following:

- (i) the DM and scattering time-scale at each epoch;
- (ii) the frequency scaling of the scattering, α ;
- (iii) the timing model parameters;
- (iv) the pulse amplitude and instrumental noise level in each profile.

In Fig. 3, we show both the injected values (red crosses) and mean parameter estimates with 1σ confidence intervals obtained from this analysis (black points with error bars) for the scattering (top left-hand panel) and DM variations (top right-hand panel). We find that the parameter estimates are consistent in both cases, and that the histogram of the residuals after subtracting the maximum-likelihood values is consistent with a Gaussian distribution with

unit variance (middle panels, red and black lines respectively). The Pearson product–moment correlation coefficient between the simulated and recovered values is 0.999 ± 0.005 . In Fig. 3 (bottom left-hand panel), we then show the one-dimensional marginalized posterior distribution for the frequency dependence of the scattering model, α . We indicate the simulated value with a vertical red line, and find that our inferred values are consistent with the simulation.

We also use this simulated data set to check the impact of not modelling the scattering variations on the parameter estimates for the DM variations. In Fig. 3 (bottom right-hand panel), we plot the difference between the measured DM variations from the simulated data set when including scattering variations in the model, or when assuming only a mean scattering time-scale (red points). Not modelling the variations in the scattering leads to significant bias in the measured DM variations. Note the different scales on the y-axis of the bottom right-hand panel and top right-hand panel, the bias introduced in this case is a factor of 2 larger than the DM variations themselves. We compare this bias with the structure present in the scattering variations (a shifted and scaled version of the scattering variations is overplotted in black). The Pearson product–moment correlation coefficient between the two sets of measurements is 0.92 ± 0.04 , indicating a significant correlation between the unmodelled scattering, and the resultant bias in the DM variations.

4 APPLICATION TO REAL DATA

We perform our analysis using observations of PSR J1643–1224 made with the 64-m Parkes radio telescope. We choose to analyse this particular MSP because it is known to display significant scattering variability, which has been attributed to its location behind an H II region (Gaustad et al. 2001; Villamariz & Herrero 2005). As such, we stress that the results presented will likely not be typical for a standard MSP. However, the benefit for more typical systems will naturally depend on combinations of factors such as the observing frequency of the data set and the brightness of the pulsar at low observing frequencies. Further application of the techniques we have discussed in the preceding sections to a broader and more typical population of pulsars will occur in subsequent work. Data were collected using two receiver packages, a co-axial system at

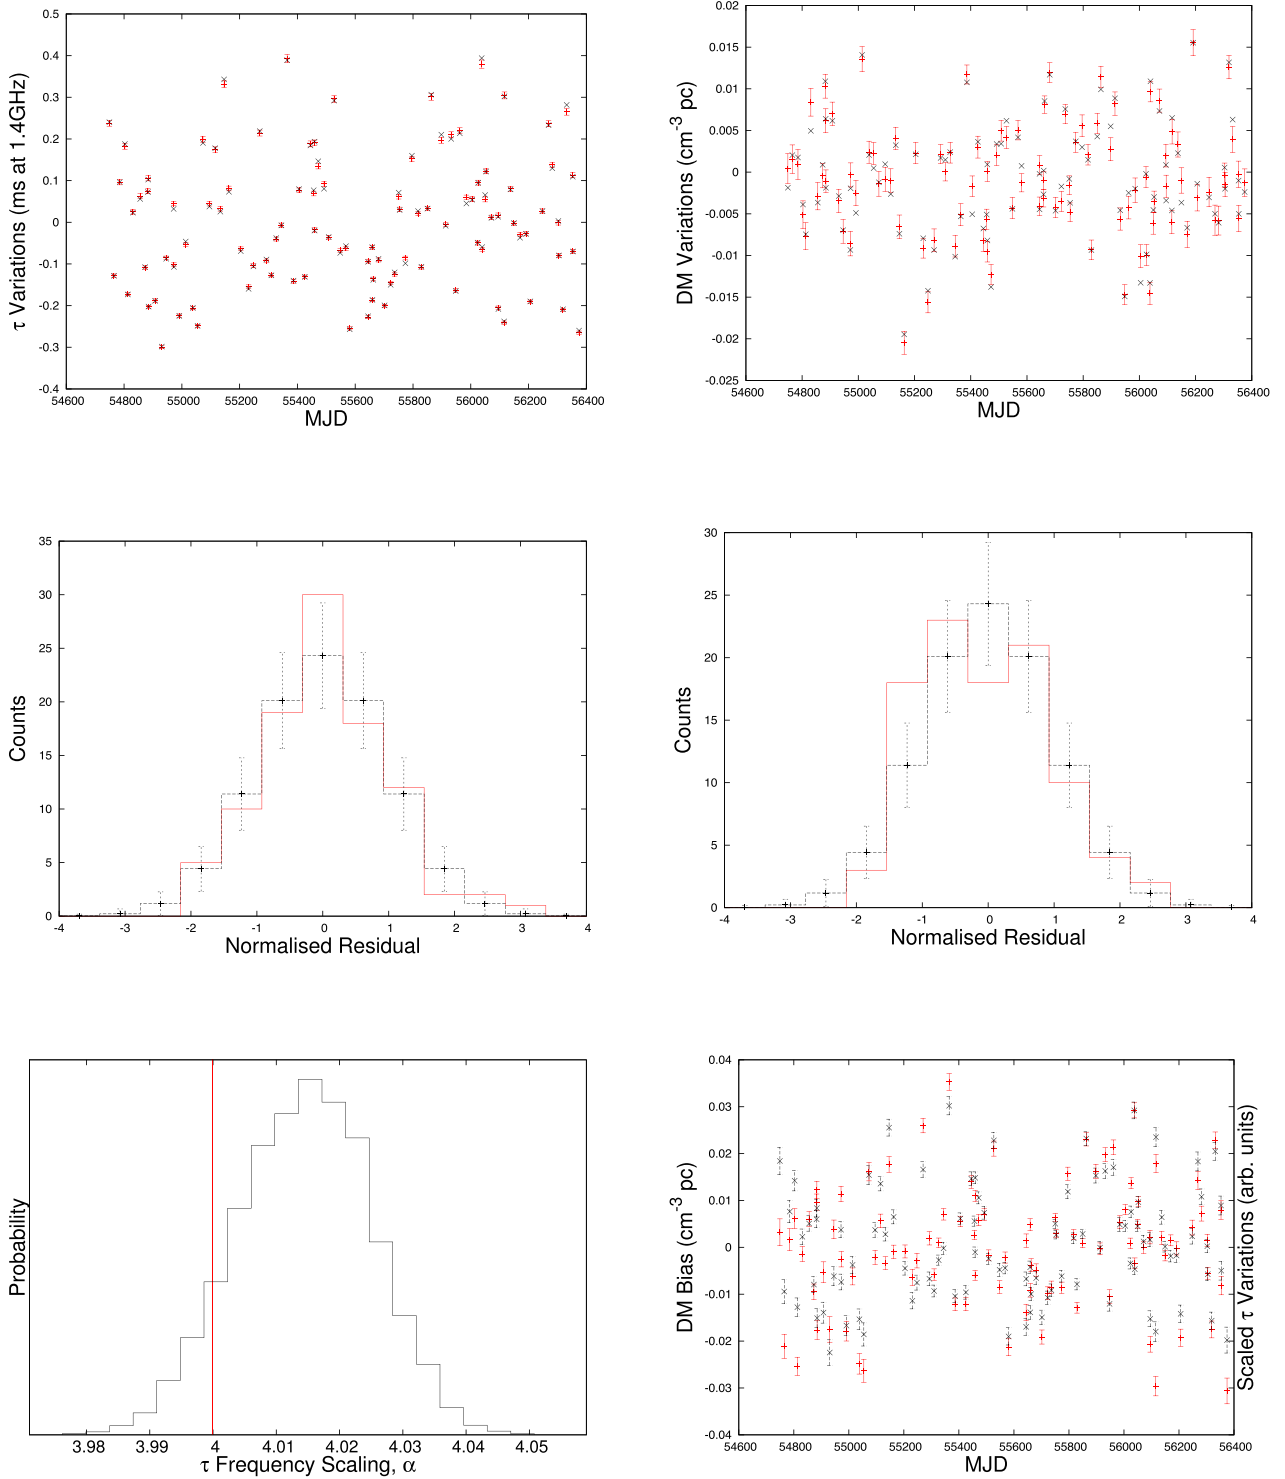


Figure 3. Top panels: injected values (red crosses) for scattering (left-hand panel) and DM variations (right-hand panel) used in the simulated data set analysed in Section 3, and the recovered parameter estimates and 1σ confidence intervals from the global analysis (black points with uncertainties). Middle panels: histograms of the residuals (red lines) after subtracting the maximum-likelihood parameter estimates from the injected values, compared to a Gaussian distribution with unit variance (black lines) for the scattering (left-hand panel) and DM variations (right-hand panel). Error bars represent \sqrt{N} uncertainties on the expected Gaussian distribution. Bottom left-hand panel: One-dimensional marginalized posterior probability distribution for the frequency scaling parameter, α (cf. equation 10) from our analysis of the simulated data set described in Section 3. The simulated value is indicated by the vertical red line. Bottom right-hand panel: Difference between the measured DM variations from the simulated data set when including scattering variations in the model, or when assuming only a mean scattering parameter (red points, 1σ uncertainties). We compare this with a shifted and scaled measurement of the scattering variations (black points, 1σ uncertainties). The Pearson product–moment correlation coefficient between the red and black points is 0.92 ± 0.04 .

10 and 40 cm and the centre pixel of a multibeam 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, PDFB3. More details about the observing systems and data reduction process are given in Manchester et al. (2013). For all three systems, we average the profile data into eight channels per band.

As there will be discrete time offsets (known as jumps) between the different observing systems, we include these as free parameters in our analysis together with the rest of the timing model parameters given in Table 1. In all the analysis that follows, we include a model for the profile and smooth profile evolution as a function of frequency using a quadratic polynomial in the shapelet amplitudes, which we found was sufficient to describe the observed variation. In addition, we then include DM variations using the DMX parametrization (Demorest et al. 2013) in which the variations in DM are modelled as a piecewise constant function, and we use an epoch length of 30 d, chosen so that the median number of 40 cm observations per epoch is one. Below we summarize three scenarios used to investigate the impact scattering variations have on the parameter estimates and uncertainties for the DM variations and timing model parameters.

(S1) We include only the 10 and 20 cm data. No scattering variations are included in the model.

(S2) It is the same as (S1); however, we also include the 40 cm data.

(S3) It is the same as (S2); however, we include scattering variations with the model described in Section 2.4, with the same cadence as the DMX parameters.

As in our analysis of simulated data in Section 3, we also attempt to recover the scaling of the scattering time-scale with frequency. However, we find that the relatively narrow 40 cm band, combined with the fact that the scattering variations are already sub-dominant to the DM variations in the 20 cm band, means we are unable to constrain the quantity in our analysis. We find that the only quantity that varies significantly when changing the scaling factor across a range of 3–5 is the mean scattering ; however, the variations in scattering and their impact on the other parameter estimates are consistent throughout. We have therefore set $\alpha = 4$ in all the results discussed below, and will refer only to the magnitude of variations in the scattering time-scale, rather than the absolute value.

4.1 DM variations

In Fig. 4 (top panel), we compare parameter estimates for the DM variations in our PSR J1643–1224 data set for scenarios S1–S3 (red, blue and black points, respectively).

Excess time-correlated noise has been observed in this pulsar previously at low observing frequencies (<1 GHz; e.g. Keith et al. 2013; Lentati et al. 2016, henceforth L16). In particular, L16a analysed the first International Pulsar Timing Array data set for PSR J1643–1224 and found that this ‘band noise’ had a steep dependence on observing frequency, such that above 1 GHz the timing fluctuations were consistent with DM variations only. We can therefore expect that in scenario (S1) we will also be dominated by DM variations, rather than scattering variations.

We find that the changes in the DM are consistent with the smooth variations observed in L16a when including band-noise terms.

When including the 40 cm data in our analysis in scenario (S2), significant additional structure can be seen in the DM that is inconsistent with the variations observed from scenario (S1). This structure has been reported previously in the literature as being the result of yearly variations in the DM (e.g. Arzoumanian et al. 2015; Jones et al. 2016)

In scenario (S3), the additional structure seen in scenario (S2) is no longer present, and we find the parameter estimates are consistent with those obtained from scenario (S1). That the parameter estimates for the DM variations are inconsistent between scenarios (S1) and (S2) is highly suggestive that unmodelled effects present in the 40 cm data are introducing bias into our analysis, and are subsequently mitigated when including scattering variations into the model.

4.2 Scattering variations

In the middle panel of Fig. 4, we show the parameter estimates for the variations in the scattering time-scale, τ , from scenario (S3). In the bottom panel of Fig. 4, we compare this structure with the additional structure seen in the DM variations in scenario (S2) compared to scenario (S1). In particular, we overplot the difference in the DM variations from scenarios (S1) and (S2) (red points) with a shifted and scaled version of the scattering variations (black points). The two clearly mirror one another, implying that the additional DM variations observed when adding low-frequency data in scenario (S2) are the result of unmodelled scattering variations, consistent with the result of not modelling the scattering variations observed in the analysis of simulated data in Section 3.

It is clear, then, that simply adding low-frequency data without considering scattering can significantly bias estimates of the DM variations. With LOFAR (Stappers et al. 2011) already performing pulsar timing at frequencies of 100–200 MHz, and upcoming telescopes such as the Square Kilometre Array providing unprecedented sensitivity in the frequency range of 500–800 MHz, this bias introduced into the analysis when failing to appropriately model scattering variations will become extremely important, especially in the context of gravitational wave astronomy using a pulsar timing array. Looking forward, this will become increasingly true even if we consider only more ‘typical’ pulsars compared to PSR J1643–1224. This is both because in the SKA era we can expect to add many distant pulsars to PTAs, for which the presence of detectable variations in scattering will be increasingly probable, and also simply because by observing for longer the impact of such effects will grow.

4.3 Timing bias

Existing methods for incorporating scattering variations into pulsar timing have suggested using maximum-likelihood estimates of scattering delays and ‘correcting’ the measured ToA by fixing that delay in the subsequent timing analysis. This can be done either by obtaining estimates of the delay from measurements of scintillation, as in Levin et al. (2016), or using cyclic spectroscopy (e.g. Palliyaguru et al. 2015). With the profile domain framework, we can directly compare the impact of fixing the scattering time-scales in the analysis at their maximum-likelihood values, which is equivalent to simply fixing the time delay induced by the scattering when

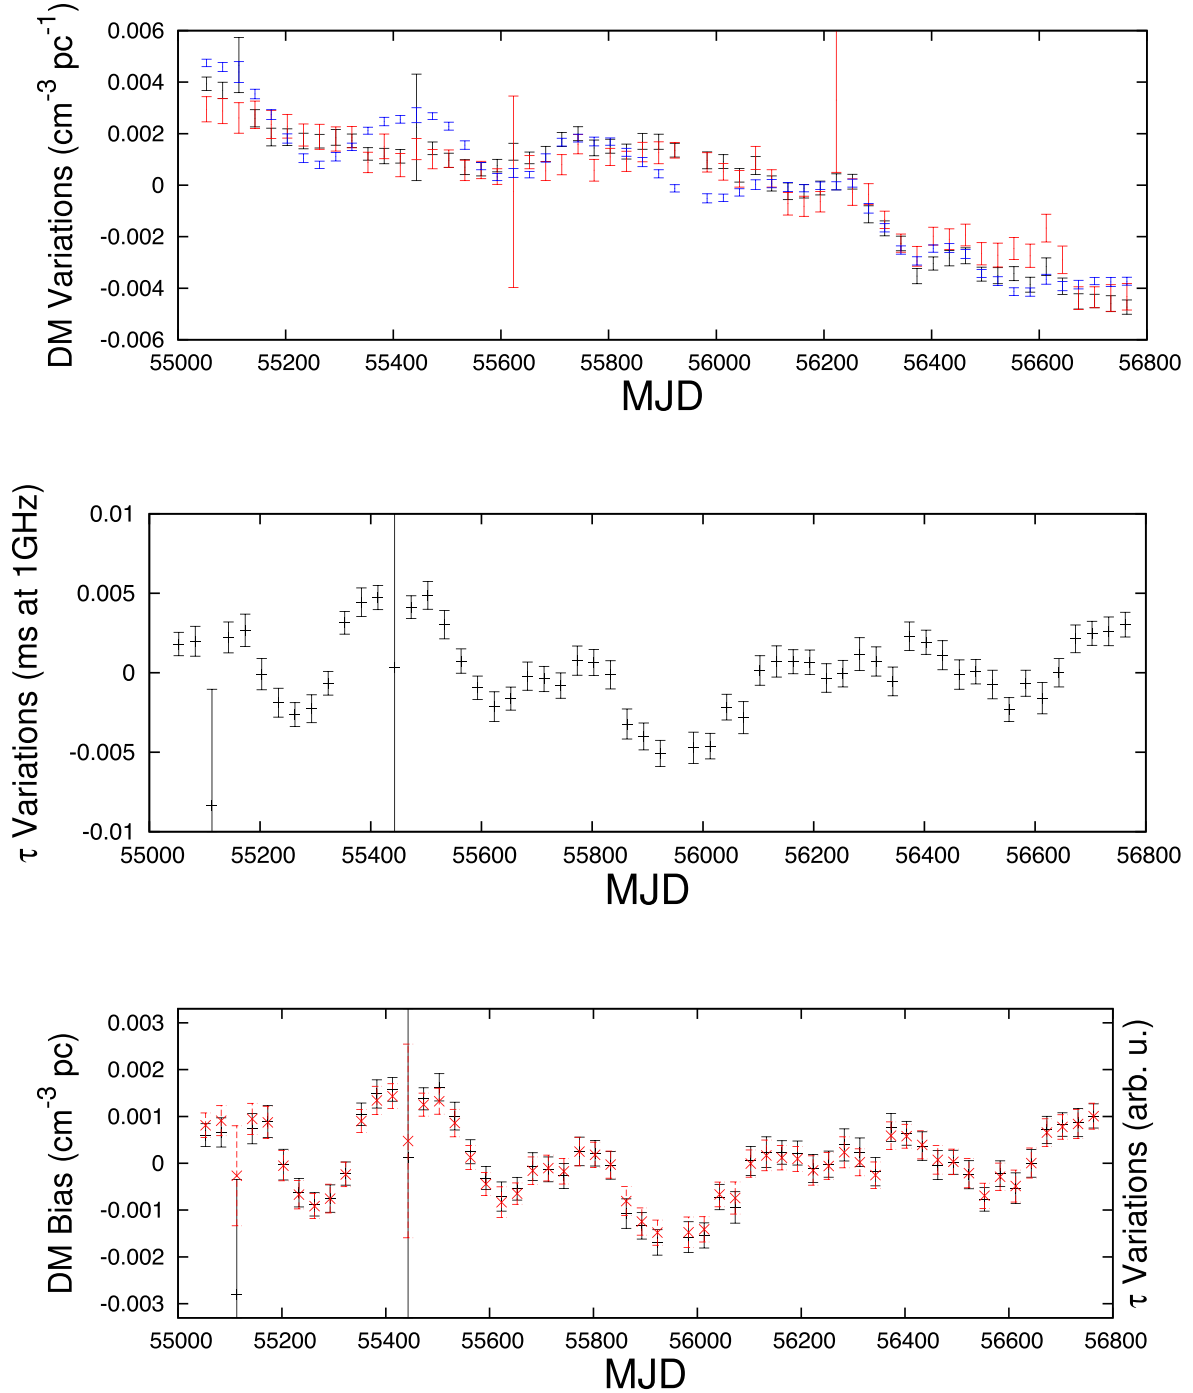


Figure 4. Top panel: mean parameter estimates and 1σ confidence intervals for the DM variations in the PSR J1643–1224 data set for three different scenarios: (1) including only the 10 and 20 cm data, and including no scattering variations in the model (red points), (2) including 10, 20 and 40 cm data, and including no scattering variations in the model (blue points), and (3) Including 10, 20 and 40 cm data, and including scattering variations in the model (black points). Middle panel: the variations in the scattering time-scale, τ and 1σ confidence intervals from our analysis of the full PSR J1643–1224 data set (black points). Bottom panel: Difference between the measured DM variations from scenarios (S2) and (S3) (red points), compared to a shifted and scaled version of the scattering variations from scenario (S3) (black points).

performing the analysis using ToAs. We find that fixing the scattering time-scales results in significant bias in the uncertainties for the DM variations of a factor of ~ 1.8 , with some epochs being much larger when less multifrequency information was available.

In Table 1, we list the timing model parameter estimates for all three scenarios, and in Fig. 5, we then compare the parameter estimates between scenarios (S1) and (S2), and scenarios (S3) and (S2). The bias in the parameter estimates introduced as a result of




Figure 5. Difference in the measured timing parameters in terms of their uncertainties between scenarios (S1) and (S2) (red points), and scenarios (S3) and (S2) (black points).

not modelling the scattering variations is substantial in several of the astrometric parameters. From the middle panel of Fig. 4, there is a clear annual term in the scattering signal that we would expect to correlate strongly with both position and proper motion. This is indeed the case, with both position in RA and proper motion in RA suffering $\sim 4\sigma$ changes in the parameter estimates. For example, the proper motion in RA changes from 6.0 ± 0.1 to 5.50 ± 0.06 when first adding in the 40 cm data without a model for scattering variations, and then changes to 5.85 ± 0.08 when including scattering variations in the model.

A natural question to ask is whether it is worth including low-frequency pulsar timing data into the analysis, if one is interested only in the timing model parameters and not the behaviour of the ISM. From Table 1, we can see that the uncertainties on the parameters from scenario (3) are still improved compared to scenario (1), in which only the 10 and 20 cm data were included in the analysis. Therefore, if the scattering variations are modelled appropriately such that no bias is introduced in the parameter estimates, the addition of the 40 cm data does improve the precision of the results.

However, this implies we must also ask whether the model used for scattering is sufficient. If the thin screen assumption is not valid, then bias will still enter into our parameter estimates. In Fig. 6, we show the timing residuals after subtracting the timing model given in Table 1 obtained from our analysis of scenario (S3). In the middle panel, we then show the timing residuals obtained from the maximum-likelihood model profile data that includes variation in the pulse width due to scattering. In both cases, significant additional structure in the 40 cm data (magenta points) can clearly be seen as a result of the scattering. In the bottom panel of Fig. 6, we then show the difference between the first two panels. The residuals, after subtracting out delays that result from scattering variations, are consistent with white noise. We stress that this process is performed purely for visualization of the profile domain model. As discussed, simply subtracting the delay induced by scattering is not an appropriate means of correcting for scattering variations.

5 CONCLUSIONS

Over the last several years, different groups worldwide have sought to improve the precision with which pulsar timing can be performed

by incorporating data across wide-band widths, extending to low radio frequencies (< 1 GHz). The principal aim was to better correct for DM variations, which are known to dominate the noise budget for the majority of millisecond pulsars being used as part of an ongoing effort to detect gravitational waves in the nHz band.

As better models for these DM variations were developed, excess noise in this low-frequency data has been observed in an increasing number of pulsars. In Lentati et al. (2016), four pulsars in the International Pulsar Timing Array data release were found to have ‘band noise’, including PSR J1643–1224. Even the most stable millisecond pulsar known to date, PSR J1909–3744, has been observed to suffer from band noise in the low-frequency data, which, unless modelled appropriately, decreased the sensitivity of the data set to gravitational waves by a factor of 2 (Shannon et al. 2015).

These band-noise models described the excess noise as time-correlated shifts in the arrival time of the pulses. Ultimately, however, if the origin of this noise is the result of time-variable broadening of the pulse profile due to scattering in the ISM, describing it in terms of shifts in the ToAs will be sub-optimal compared to appropriately modelling the changes in the pulse shape.

We have described a statistically robust approach to incorporating models for scattering into pulsar timing analysis by extending the existing profile-domain timing framework. This approach makes it possible to simultaneously estimate temporal variations in both the DM and scattering, together with a model for the pulse profile that includes smooth evolution as a function of frequency, the pulsars timing model and models for pulse jitter.

We have shown that unless both the DM variations and scattering are sampled together in the timing analysis, one can significantly underestimate the uncertainties in both the DM and the arrival time of the pulse, leading to bias in the timing parameters. This means that fixing the scattering time-scales and simply ‘correcting’ the ToAs, as has been suggested as a means of handling scattering variations in the literature, is not a viable approach to performing a robust analysis of low-frequency pulsar timing data. We also showed, however, that the true PDF of the arrival time of the pulse can be highly non-Gaussian when scattering is only marginally detected. In this case, the typical procedure for forming a ToA, which assumes a Gaussian approximation to the true PDF, is inappropriate. However, by performing a profile domain analysis the non-Gaussianity of the PDF is automatically incorporated into the final parameter estimates.

We have applied our method using a new, publicly available, GPU-accelerated code both to simulations and observations of the millisecond pulsar PSR J1643–1224 that cover a frequency range from 700 up to 3800 MHz over a period of 4.7 yr. This particular MSP was chosen because it is known to display significant scattering variability compared to typical MSPs. The techniques we have introduced will be applied to a broader and more typical population of pulsars in subsequent work. When including the low-frequency data in the PSR J1643–1224 analysis, the parameter estimates for both the DM variations and timing model are significantly biased by the variations in the pulse width due to scattering. Further, simply fixing the scattering time-scales to maximum-likelihood values – equivalent to fixing the time offset when making ToAs – results in almost a factor 2 underestimation in the uncertainties of the DM variations compared to the simultaneous analysis.

If low-frequency observations from telescopes such as LOFAR or the Square Kilometre Array are to be incorporated into pulsar timing analysis, then the bias that results from failing to appropriately model scattering variations will become extremely important. The increased sensitivity afforded by these instruments will likely




Figure 6. Top panel: timing residuals after subtracting the maximum-likelihood timing model, including DM variations obtained from the profile domain analysis performed in scenario (S3). Colours in this, and low panels, indicate the observing frequency of the data: 10 (cyan points), 20 (black points) and 40 cm (magenta points). Middle panel: model residuals obtained by forming ToAs from the maximum-likelihood model profiles from scenario (S3), and subtracting the same timing-model as in the top panel. The time-variable broadening introduced by the scattering model results in completely consistent structure as the observed ToAs. (Bottom panel) The difference between the residuals in the top and middle panels. The ‘corrected’ residuals are now consistent with white-noise; however, we stress this is only to be used figuratively. The residuals obtained by subtracting the maximum-likelihood timing offsets introduced by scattering should not be used for analysis as they do not propagate the uncertainties in the scattering parameters through to the timing uncertainties.

require more complex models than simple thin-screen approximations for scattering. Any differences between the true PBF and the models used will manifest as additional band-dependent noise, downweighting the low-frequency data. Given finite observing time, one must ask how much of that time to spend observing at low frequencies. If we are able to completely model the variations in pulse shape, then low-frequency observations will naturally be of great

benefit. The actual improvement, however, will vary from pulsar to pulsar depending on the strength and complexity of the scattering variations.

For pulsars that are bright at high frequencies ($\gtrsim 3$ GHz), the optimal approach may well be to spend all available observing time in a regime where IISM effects are negligible. Such an approach has already resulted in the most sensitive limit on an isotropic

gravitational wave background using a pulsar timing array to date (Shannon et al. 2015). However, pulsars are known to have steep spectral indices ($S_\nu \propto \nu^{-1.8}$ on average, Maron et al. 2000), and so in many cases observations at lower frequencies are necessary. In such instances, our ability to propagate our uncertainties in the scattering variability through to all other parameters will be critical, especially in the context of gravitational wave astronomy. This is all handled automatically as part of a profile domain analysis, which will inevitably form the basis of modern, broad-band pulsar timing analysis moving forward.

ACKNOWLEDGEMENTS

The Parkes radio telescope is part of the Australia Telescope National Facility, which 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. SO was supported by the Alexander von Humboldt Foundation.

We thank Stephen Taylor for his assistance in building the `PYTHON` software used in this paper.

REFERENCES

- Arzoumanian Z. et al., 2015, *ApJ*, 813, 65
 Bhat N. D. R., Cordes J. M., Chatterjee S., 2003, *ApJ*, 584, 782
 Coles W. A., Rickett B. J., Gao J. J., Hobbs G., Verbiest J. P. W., 2010, *ApJ*, 717, 1206
 Demorest P. B. et al., 2013, *ApJ*, 762, 94
 Desvignes G. et al., 2016, *MNRAS*, 458, 3341
 Foster R. S., Backer D. C., 1990, *ApJ*, 361, 300
 Gaustad J. E., McCullough P. R., Rosing W., Van Buren D., 2001, *PASP*, 113, 1326
 Jones M. L. et al., 2016, preprint ([arXiv:e-prints](https://arxiv.org/abs/1608.07692))
 Keith M. J. et al., 2013, *MNRAS*, 429, 2161
 Lam M. T. et al., 2016, *ApJ*, 819, 155
 Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, *MNRAS*, 437, 3004
 Lentati L. et al., 2017, *MNRAS*, 466, 3706
 Lentati L. et al., 2016, *MNRAS*, 458, 2 (L16)
 Levin L. et al., 2016, *ApJ*, 818, 166
 Liu K. et al., 2014, *MNRAS*, 443, 3752
 Löhmer O., Kramer M., Mitra D., Lorimer D. R., Lyne A. G., 2001, *ApJ*, 562, L157
 Manchester R. N. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, 17
 Maron O., Kijak J., Kramer M., Wielebinski R., 2000, *A&AS*, 147, 195
 Narayan R., 1992, *Phil. Trans. R. Soc. A*, 341, 151
 Palliyaguru N., Stinebring D., McLaughlin M., Demorest P., Jones G., 2015, *ApJ*, 815, 89
 Pennucci T. T., Demorest P. B., Ransom S. M., 2014, *ApJ*, 790, 93
 Refregier A., 2003, *MNRAS*, 338, 35
 Shannon R. M. et al., 2015, *Science*, 349, 1522
 Stappers B. W. et al., 2011, *A&A*, 530, A80
 Taylor S. R., Vallisneri M., Ellis J. A., Mingarelli C. M. F., Lazio T. J. W., van Haasteren R., 2016, *ApJ*, 819, L6
 Villamariz M. R., Herrero A., 2005, *A&A*, 442, 263
 Williamson I. P., 1972, *MNRAS*, 157, 55

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.