
Precision radio-frequency pulsar timing & interstellar scintillometry

By

DANIEL JOHN REARDON



MONASH University

School of Physics and Astronomy
MONASH UNIVERSITY

A dissertation submitted to MONASH UNIVERSITY in
accordance with the requirements of the degree of
DOCTOR OF PHILOSOPHY in the FACULTY OF SCIENCE.

MAY 2018

Advisors: Prof. Yuri Levin & Dr. George Hobbs

COPYRIGHT NOTICE

©The author (2018)

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

ABSTRACT

Pulsars are natural clocks that can be used as laboratories for strong-field gravity, and as sensitive probes of the interstellar plasma along our line-of-sight. Pulsar timing arrays (PTAs) monitor a set of the most stable millisecond pulsars (MSPs) over many years to function as a Galactic-scale gravitational wave (GW) detector, sensitive to nanohertz-frequency waves. The eventual detection of these GWs will rely on a detailed understanding of the pulsars and of the various sources of noise, including the interstellar plasma. In this thesis we improve the ephemerides for a set of these MSPs, and a relativistic binary pulsar, using the techniques of pulsar timing and scintillometry.

For each of the radio MSPs in the Parkes PTA, we have analysed their intrinsic spin noise and modelled small changes to the electron density along our line-of-sight. These are strong sources of non-stationary noise for pulsar timing, which we accounted for by using a new algorithm. We produced new timing models for each pulsar, resulting in the first parallax measurements for five pulsars and improvements to other parameters, such as the Shapiro delay to give new pulsar mass measurements. Our new distance measurement for PSR J0437–4715, $D = 156.79 \pm 0.25$ pc, is the most precise for any pulsar. For the first time, the uncertainty on this measurement is less than the typical wavelengths of GWs that PTAs are sensitive to, which will increase sensitivity to continuous GW sources.

Scintillation of radio pulsars is results from scattering by the turbulent interstellar plasma. The scattering leads to a time-variable interference pattern that is sampled by the observatory moving through space along the Earth’s orbit. We use measurements of this scintillation to model the motion of the line-of-sight to the pulsar, which depends on the relative transverse motions of the Earth, pulsar, and scattering plasma. We have measured long-term changes in the diffractive scintillation pattern of a relativistic binary pulsar, PSR J1141–6545, and used it to determine several previously unknown important parameters for the system. In particular, we have measured the orientation of the orbit in celestial coordinates, $\Omega = 23 \pm 3^\circ$ (N→E). We use our scattering model in a new way to estimate the pulsar distance $D = 9_{-3}^{+5}$ kpc, and to give the first estimate of its proper motion in right ascension $\mu_\alpha = 2.5 \pm 1.2$ mas yr^{−1} and in declination $\mu_\delta = 1.5 \pm 0.7$ mas yr^{−1}. With these measurements, we constrain the kinematic contamination to the relativistic orbital period-derivative for this system.

Finally, we have capitalised on the amazing long-term data collected by the PPTA collaboration for PSR J0437–4715. The intensity measurements reveal parabolic "scintillation arcs" on ~1300 observing epochs, which vary in curvature with the changing line-of-sight velocity. By modelling these arcs, we have determined some orbital parameters of the pulsar, and properties of the interstellar scattering medium, with unprecedented precision. This method is a promising alternative to the more common diffractive scintillation modelling for pulsars that have well-defined scintillation arcs or that are observed in the regime of weak scattering.

AUTHOR'S DECLARATION

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes one original paper published in peer reviewed journals and two unpublished publications. The core theme of the thesis is the understanding of noise and interstellar medium effects in pulsar timing residuals, and the use of scintillation in conjunction with pulsar timing to improve models of pulsars. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the student, working within the School of Physics and Astronomy under the supervision of Professor Yuri Levin.

The inclusion of co-authors reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research. In the case of Chapters 2, 3, and 4 my contribution to the work involved the following, shown in Table 1

NAME:

SIGNED:

DATE:

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work. In instances where I am not the responsible author I have consulted with the responsible author to agree on the respective contributions of the authors.

NAME:

SIGNED:

DATE:

Table 1: Declaration of published works contained in this thesis and acknowledgement of co-author contributions. No co-authors are students of Monash University.

Thesis chapter	Publication title	Status	Nature and % of student contribution	Nature and % of Co-author contributions
2	Timing analysis for 20 millisecond pulsars in the Parkes Pulsar Timing Array	Accepted	65%, tested split-Cholesky algorithm with monte carlo simulations, found noise models (for timing noise, white noise, and DM variations) and new timing models for all PPTA pulsars. Prepared draft manuscript.	1) George Hobbs, implemented split-Cholesky algorithm in TEMPO2 and input into manuscript, 10% 2) William Coles, co-develop split-Cholesky algorithm and input into manuscript and DM modelling procedure, 10% 3) Yuri Levin, input into manuscript, 6% 4) Michael Keith, co-develop split-Cholesky algorithm, 4% 5) Parkes Pulsar Timing Array (PPTA) collaboration, observing and manuscript input, 5%
3	Modelling the relativistic and annual variations in the scintillation of PSR J1141–6545 for precise astrometry	Submitted	60%, co-development of codes for measuring and modelling diffractive scintillations from dynamic spectra. Tested models and fit scintillations of PSR J1141–6545. Prepared draft manuscript	1) William Coles, co-development of scintillation codes and manuscript input, 20% 2) George Hobbs, manuscript input, 5% 3) Stephen Ord, observing and manuscript input, 3% 4) Matthew Kerr, assistance using PPTA data processing pipeline for dynamic spectra, and manuscript input, 3% 5) Matthew Bailes, Ramesh Bhat, and Vivek Venkatraman Krishnan, observations and manuscript input, 2%
Unpublished traditional thesis Chapter				
Thesis chapter	Chapter title	Nature and % of student contribution		Nature and % of other contributions
4	Precise scintillometry with annual and orbital variations in scintillation arcs for PSR J0437–4715	80%, Developed code for selecting, stacking, and editing dynamic spectra, and for measuring the scintillation arcs from secondary spectra. Sorted through and edited observations to improve the sample. Measured the scintillation arcs, then identified and modelled the arc curvature variations following these measurements. Wrote complete chapter draft.		1) William Coles, provided a framework of code for calculating secondary spectra, 5%, 2) Artem Tuntsov & Mark Walker used code from the work of Walker et al. (2017) to search stellar catalogues for possible associations and to calculate their probability of association for the given set of arc parameters. They also provided useful discussion and some comments on a draft manuscript, 10% 3) George Hobbs & Yuri Levin, discussion and comments on chapter draft, 5%

ACKNOWLEDGEMENTS

Thanks to my supervisors Yuri Levin and George Hobbs, and to my unofficial advisor and collaborator Bill Coles, for the amazing support and advice during my candidature. I owe my completion of this thesis to your stellar professional and personal guidance.

I am grateful to Alma and Bill Coles for their kindness and hospitality during my visit to San Diego, where I not only learned some incredible, scintillating science, but also discovered a love of cycling and mountain biking.

I am lucky to have been able to attend fascinating conferences in Banff, Canada; Stellenbosch, South Africa; the Blue Mountains, New South Wales; and at the Parkes radio telescope. This was possible thanks to the support of Monash University, CSIRO Astronomy and Space Science (CASS), and of course Yuri and George. Out of all the places I have been fortunate enough to travel to, there was something particularly awe-inspiring about observing with Parkes from inside the control tower. Experiences like this, and the exciting moments of discovery, are why I have a passion for studying astronomy.

I thoroughly enjoyed my many visits to CASS in Sydney to work with George and other members of the Parkes Pulsar Timing Array, again thanks to the support of Yuri and CASS. A definite highlight of my candidature has been sharpening my table tennis skills on these trips.

Thanks also to Barney Rickett for valuable discussions on the diffractive scintillation work during my San Diego visit, Matthew Bailes for reading and providing comments on my thesis introduction, and Mark Walker and Artem Tuntsov for helpful discussions on scintillation arcs and hot stars.

I am especially grateful for all of the wonderful and supportive people in my life who have encouraged me and always been interested in my work. In particular, my amazing partner Raluca who helped me through stressful moments and is always eager to learn about astronomy, my friend Ross Turner who is seemingly available to discuss problems in my work at any hour, and my parents and other family members for always being excited to hear about my research.

Thanks to everyone for the years of encouragement!

Science doesn't always go forwards. It's a bit like doing a Rubik's cube. You sometimes have to make more of a mess with a Rubik's cube before you can get it to go right.

– Jocelyn Bell Burnell

Twinkle, twinkle, little star

– Jane Taylor

TABLE OF CONTENTS

	Page
List of Tables	xv
List of Figures	xvii
1 Introduction	1
1.1 Discovery and Properties of Pulsars	1
1.1.1 Pulsar formation and evolution	2
1.1.2 Millisecond pulsars and binaries	3
1.2 Pulsar Timing	6
1.2.1 Radio observations	6
1.2.2 Data processing pipeline	8
1.2.2.1 Noise sources	9
1.2.2.2 Parameter estimation in the presence of noise	11
1.2.3 Pulsar ephemeris	13
1.2.3.1 Keplerian binary parameters	14
1.2.3.2 Post-Keplerian parameters	16
1.2.4 Other applications	17
1.2.4.1 Ultra-low frequency gravitational wave detection	17
1.2.4.2 Tests of general relativity and other theories of gravity	18
1.3 Effects of the Ionised Interstellar Medium	19
1.3.1 Frequency Dispersion	20
1.3.2 Structure of the IISM	21
1.3.3 Scattering	22
1.3.4 Scintillation	26
1.4 Thesis overview	27
2 Timing analysis for 20 millisecond pulsars in the Parkes Pulsar Timing Array	29
2.1 Introduction	29
2.2 Observations	31
2.3 The timing model	33

TABLE OF CONTENTS

2.3.1	Choosing parameters to include in the model	34
2.4	Parameter measurement in the presence of non-stationary red noise	35
2.4.1	Modelling the Dispersion Measure variations	36
2.4.2	Modelling the red timing noise	38
2.4.3	How do we know when our models are optimal?	39
2.5	Results	41
2.5.1	PSR J0437–4715	44
2.5.2	PSR J0613–0200	45
2.5.3	PSR J0711–6830	45
2.5.4	PSR J1022+1001	45
2.5.5	PSR J1024–0719	47
2.5.6	PSR J1045–4509	47
2.5.7	PSR J1600–3053	48
2.5.8	PSR J1603–7202	48
2.5.9	PSR J1643–1224	48
2.5.10	PSR J1713+0747	48
2.5.11	PSR J1730–2304	49
2.5.12	PSR J1732–5049	49
2.5.13	PSR J1744–1134	49
2.5.14	PSR J1824–2452A	50
2.5.15	PSR J1857+0943	50
2.5.16	PSR J1909–3744	50
2.5.17	PSR J1939+2134	51
2.5.18	PSR J2124–3358	51
2.5.19	PSR J2129–5721	51
2.5.20	PSR J2145–0750	51
2.6	Discussion	51
2.6.1	Advantages of using the split-Cholesky algorithm	51
2.6.2	PSR J0437–4715 Kinematic distance measurement from $\dot{\mathbf{P}}_{\mathbf{b}}$	52
2.7	Conclusion	55
3	Modelling relativistic and annual variations in the scintillation of PSR J1141–6545 for precise astrometry	57
3.1	Introduction	58
3.2	Dataset	60
3.2.1	Observations and dynamic spectra	60
3.2.2	Measurement of τ_d and Δv_d	62
3.2.3	The effects of inhomogeneity in the IISM	64
3.3	The models	66

3.3.1	Effective velocity and IISM anisotropy	69
3.3.2	Comparison to the earlier models	71
3.3.2.1	Harmonic coefficients	71
3.3.2.2	Physical model for individual epochs	72
3.4	Fitting and comparing the models	74
3.5	Results and model comparison	75
3.5.1	Isotropic scattering	75
3.5.1.1	Individual epochs	75
3.5.1.2	Long-term models	78
3.5.2	Anisotropic Scattering: Long-term models	80
3.5.2.1	Uniform medium	80
3.5.2.2	Thin screen	81
3.6	Discussion	84
3.6.1	Pulsar distance and proper motion estimates	84
3.6.2	Implications for timing and tests of general relativity	86
3.7	Conclusion	87
4	Precise scintillometry with annual and orbital variations in scintillation arcs for PSR J0437–4715	89
4.1	Introduction	90
4.2	Observations and data	93
4.2.1	Computing dynamic and secondary spectra	94
4.3	Interpreting and fitting the secondary spectra	96
4.3.1	Fitting arc curvature	98
4.4	Modelling arc curvature variations	99
4.5	Results	101
4.5.1	Primary arc	102
4.5.2	Secondary arc	104
4.6	Discussion	106
4.6.1	Screen distances and IISM velocity	106
4.6.2	Object candidates	107
4.6.2.1	Stellar associations	108
4.7	Conclusion	110
5	Conclusion	113
5.1	Suggestions for future work	117
	Bibliography	121
	Appendix	139

TABLE OF CONTENTS

Timing model parameter tables from Chapter 2	139
Reproducing our results from Chapter 3	144

LIST OF TABLES

1	Declaration of published works contained in this thesis and acknowledgement of co-author contributions. No co-authors are students of Monash University.	vi
	TABLE	Page
2.1	Parameters describing the DM model used for each pulsar. DM noise in the earliest residuals is described by the DM covariance function parameters (a and b) in Equation 7, which are calculated after the removal (if necessary) of a linear trend, $d\text{DM}/dt$, and annual variations described by the sine (A) and cosine (B) amplitudes. Δt_{DM} is the separation of ΔDM measurements in the multi-frequency section of the dataset for each pulsar.	39
2.2	Parameters for the red-noise model for each pulsar. The parameters are used to describe the frequency-independent noise in the data with a power law (Equation 4), where α is the spectral index and P_0 is the power at corner frequency, f_c	40
2.3	Parallax values and distance measurements for pulsars in our sample. Parallax and parallax-derived distance values are corrected for the Lutz-Kelker bias using the method of Verbiest et al. (2012). TC93 and NE2001 DM distances have approximate uncertainties of 25% and 20% respectively.	44
3.1	Measured parameters for each of the isotropic scintillation models. The first two columns are parameters from single-epoch fits of a physical model. κ_u and i are the weighted mean values from the fit, while Ω , s , $v_{\mu,\parallel}$, and $v_{\mu,\perp}$ are derived from a simple model of the annual variations in a time series of the measured $v_{C,\parallel}$, and $v_{C,\perp}$ (Section 3.5.1.1). Two degenerate solutions of equal quality are given for these single-epoch fits, and the equivalent solutions are also given in the next four columns of long-term isotropic models. Columns three and four are the parameters from a uniform medium while columns five and six are for a thin screen; these are discussed in Section 3.5.1.2. The chi-squared χ^2 and reduced chi-squared χ_r^2 values for each model are given, and these are used for model comparison. Numbers in brackets are the 1σ standard errors on the last quoted digit.	79

3.2	Measured parameters for anisotropic long-term models for model comparison. Columns one and two are the best-fit solutions for an anisotropic uniform medium and thin screen respectively. The uniform medium model is ruled out based on the low inclination angle and high space velocity (Section 3.5.2). Bounding the inclination angle to $1\text{ rad} < i < \frac{\pi}{2}\text{ rad}$ does not reveal a valid solution and instead returns the boundary condition $i = 1\text{ rad}$. We therefore only considered the thin screen model with inclination angle fixed at values corresponding to the $\sin i$ constraint of pulsar timing. The two best solutions for each of $i = 73^\circ$ and $i = 107^\circ$ are given in the right four columns for the case of an anisotropic thin screen (and these are discussed in Section 3.5.2.2). The chi-squared χ^2 and reduced chi-squared χ_r^2 values for each model are given, which we use for model comparison. Numbers in brackets are the 1σ standard errors on the last quoted digit.	81
3.3	Two best-fit long-term models for an anisotropic thin scattering screen, with derived parameters. For the pulsar distance D , we assume a 20% uncertainty on the A_{ISS} coefficient from (Cordes and Rickett 1998). The parameters are described and discussed in Section 3.6.1.	86
4.1	Parameters for isotropic and anisotropic scattering models for the curvature variations in the primary and secondary arcs for PSR J0437–4715. Parameters were measured with a weighted nonlinear least-squares analysis.	102
1	Parameters for the solitary pulsars J0711–6830, J1024–0719, J1730–2304, J1744–1134, J1824–2452A, J1939+2134, and J2124–3358. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.	140
2	Parameters for the binary pulsars described by the small-eccentricity ELL1 binary model, PSRs J0613–0200, J1045–4509, J1603–7202, J1732–5049, J1857+0943, J2129–5721, and J2145–0750. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.	141
3	Parameters for the binary pulsars described by DD binary model, PSRs J1022+1001, J1600–3053, and J1643–1224. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.	142
4	Parameters for the binary pulsars described by T2 binary model, PSRs J0437–4715, J1713+0747, and J1909–3744. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis. In each case where a companion mass is measured from the Shapiro delay, the corresponding $\sin i$ parameter is linked to the Kopeikin parameter, i	143

LIST OF FIGURES

FIGURE	Page
1.1 " $P-\dot{P}$ diagram", displaying spin period P and spin-down rate \dot{P} for all known pulsars with measurements taken from version 1.57 of the Australia Telescope National Facility (ATNF) pulsar catalogue (Manchester et al. 2005). Pulsars with radio emission are shown with blue dots and radio-quiet pulsars (with pulsed emission of higher-energy radiation) are shown with maroon triangles. Pulsars with binary companions have a red circle surrounding the pulsar type marker, and those with associated supernova remnants (SNRs) are surrounded by a yellow star. Pulsars with high-energy (x-ray and gamma-ray) emissions without pulsations, such as Anomalous X-ray Pulsars (AXP) or Soft Gamma-ray Repeater (SGR), are shown with green triangles. The dashed lines are lines of constant characteristic age $\tau_{\text{age}} \propto P/\dot{P}$, and dashed-dotted lines are of constant magnetic field strength $B \propto \sqrt{P\dot{P}}$. The solid line is the pulsar "death line" that encloses the pulsar "graveyard", beyond which electron-positron pair production in the pulsar's inner magnetosphere is expected to cease, which terminates the radio emission.	4
1.2 Schematic of the major steps involved in the observing and data processing pipeline for pulsar timing, from emission at the pulsar to measurements of timing model parameters. The steps illustrated here are summarised in Section 1.2.2.	7
1.3 Schematic definition of angles in a Keplerian binary orbit. Lines in the plane of the sky are grey and lines in the plane of the orbit are black; the intersection of these planes is the line of nodes (grey dashed line). The orbit is tilted from the plane by the inclination angle i . Periastron occurs at a distance of $a(1-e)$, for semi-major axis a and eccentricity e , and the longitude of periastron ω is defined from the line connecting the barycentre to the ascending node. The orientation of the orbit from the perspective of the observer is given by the longitude of the ascending node Ω in the plane of the sky (defined East of North). This figure was adapted from Figure 8.3 of Lorimer and Kramer (2004).	15

- 1.4 A schematic of interstellar medium effects on the emission from radio pulsars. The pulsar initially emits radiation that is spatially coherent. The radiation interacts with the ionised interstellar medium (IISM) during its propagation and becomes scattered and frequency-dispersed (the rotation measure is also altered by the Faraday effect since the IISM is magnetised). The IISM here is assumed to be thin compared to the total distance between the observer and pulsar D , and it is located at a distance D_s from the pulsar. The wavefronts are distorted from scattering and interfere to produce a frequency-dependent interference pattern that is sampled by the observatory. The observatory, IISM, and pulsar have transverse velocities (with respect to the local standard of rest) V_E , V_{IISM} , and V_P respectively (the direction chosen for the figure was arbitrary, and just symbolises the transverse velocity). The effective velocity V_{eff} (defined in schematic) is the transverse velocity of a point along the line-of-sight with respect to the medium (which in this case is defined at the location of the thin screen at D_s). The velocity of the line-of-sight through the diffraction pattern at the observer is V_{los} . The distance to the scattering screen from the observer (used in-text) is $D_e = D - D_s$ 25
- 2.1 (a) Whitenened and normalised post-fit residuals for PSR J1713+0747. (b) Power spectra of whitenened and normalised post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and dashed lines indicate expected mean and variance for the spectrum. The pseudo-frequency is determined by converting the whitenened components to a pseudo-time-series using the ToAs of the unwhitenened residuals. (c) Cumulative distribution of whitenened and normalised post-fit residuals (solid line) with expected distribution based on normal distribution with zero mean and unit variance (dashed line). Modified AD statistic for this distribution with the expected distribution is labelled. (d) Power spectra of post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and solid flat line is an estimate of the white-noise level. Dashed line is the power law model of the frequency-independent timing noise. 41
- 2.2 Final post-fit residuals for each of the pulsars in our sample. The vertical range of each subplot is given below the pulsar name. 43

- 2.3 (a) Whitenened and normalised post-fit residuals for PSR J0437–4715. (b) Power spectra of whitenened and normalised post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and dashed lines indicate expected mean and variance for the spectrum. The pseudo-frequency is determined by converting the whitenened components to a pseudo-time-series using the ToAs of the unwhitenened residuals. (c) Cumulative distribution of whitenened and normalised post-fit residuals (solid line) with expected distribution based on normal distribution with zero mean and unit variance (dashed line). Modified AD statistic for this distribution with the expected distribution is labelled. (d) Power spectra of post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and solid flat line is an estimate of the white-noise level. The timing noise model applies over entire data set in combined noise model, while the Kolmogorov DM noise model only applies to residuals prior to MJD 53430. 46
- 2.4 Distribution of parallax values from timing model fits to 500 realisations of simulated PSR J0437–4715 data. σ is the number of standard deviations from the true value for each of the N realisations. We used three different noise treatments: (a) No red-noise model, (b) red-noise model from DR1 data, extrapolated to apply over the entire dataset was used with the Cholesky algorithm, and (c) two-component red-noise model with Kolmogorov DM model is used in split-Cholesky algorithm. For each panel, the black line is a normal distribution fit to the distribution and the red, dashed line is a normal distribution with zero mean and standard deviation equal to the average of the standard TEMPO2 uncertainties for the 500 realisations, scaled to the same area as the black-line distribution. 53
- 2.5 Distributions of parameter values from timing model fits with the split-Cholesky algorithm to 500 realisations of simulated PSR J0437–4715 data. σ is the number of standard deviations from the true value for each realisation. Data were simulated with statistically identical red noise properties to those of PSR J0437–4715. 54
- 3.1 A dynamic spectrum of four PSR J1141–6545 observations made on MJD 56391 with $\sim 300 \text{ MHz}$ of used bandwidth about a centre frequency of 1366 MHz . The four observations cover a combined ~ 2.2 pulsar orbits. Stretching of the scintles (in black) in time is evident (at e.g. 100 and 400 minutes) when the pulsar orbital velocity reaches a minimum. The vertical white bars are periods between the individual observations, while horizontal white bars and patches were removed because of radio-frequency interference. The greyscale shows the normalised flux with the colour limits chosen to optimise the visualisation of scintles. 59

- 3.2 Decorrelation bandwidth ($\Delta\nu_d$) as a function of time for all archival observations in a ~ 6 year span with a total observing time of at least 20 minutes at ~ 1400 MHz for PSR J1141–6545. The data are split into 23 epochs separated by vertical, dashed lines. Each epoch contains multiple individual observations, and the dynamic spectra for these were cut into ~ 11 minute segments. Black crosses are the measured $\Delta\nu_d$ for each segment where $B_c < \Delta\nu_d$ for channel bandwidth B_c . Grey circles are a flux-based estimate of $\Delta\nu_d$ from the procedure described in Section 3.2.3 for segments where $B_c > \Delta\nu_d$. The measurements and estimates were used to calculate a weighted mean value of $\Delta\nu_d$ for each observing epoch. The anomalous region just before MJD 55500 is primarily over-estimated due to poor dynamic spectrum quality because of terrestrial radio interference in some observations at this time. This does not significantly affect the weighted mean for this epoch. 62
- 3.3 Scintillation velocity, V_{ISS} as a function of true anomaly for the 23 epochs of observations shown in Figure 3.2. The title of each panel gives the approximate starting date for the first observation in the group. V_{ISS} is defined as the scintillation velocity observed at the Earth for a uniform, Kolmogorov medium along the line-of-sight (Equation 3.3). The solid line for each panel is the best fit physical model, which is described in Section 3.3.2.2. 67
- 3.4 The normalised harmonic coefficients derived from a fit to $V_{\text{ISS}}(\phi)$ at each observing epoch, as described in Section 3.3.2.1. 73
- 3.5 The longitude of periastron, ω as a function of time, measured independently at each of the 23 epochs shown in Figure 3.3. The solid line is a weighted best fit, where the gradient, $\dot{\omega} = 5.6 \pm 0.3^\circ \text{ yr}^{-1}$, is the advance of periastron and is consistent with the measurement from pulsar timing (red dashed line). 77
- 3.6 Components of the constant (with orbital phase) transverse velocity parallel ($v_{C,\parallel}$, top panel) and perpendicular ($v_{C,\perp}$, bottom panel) to the line of nodes for each observing epoch. The measured velocity is a scaled combination of the Earth and IISM velocities and the pulsar proper motion. The Earth’s contribution is apparent from the clear annual modulation. The grey line is a weighted best fit annual sine wave to each of the time series. The red diamonds are the calculated Earth’s velocity at each epoch, scaled down by a screen distance $s = 0.36 \pm 0.07$, and rotated from celestial coordinates by $\Omega = 25 \pm 10^\circ$ (i.e. “Solution 1” from Table 3.1). 78

- 3.7 Scintillation velocity, V_{ISS} as a function of orbital phase, ϕ for ~ 6 years of PSR J1141–6545 observations at a frequency of ~ 1400 MHz. The orbital phase was calculated using the ω and $\dot{\omega}$ values measured in the pulsar timing model. V_{ISS} in this case is defined at the scattering screen (IISM frame), because the distance to the screen was a parameter in the model. The best-fit anisotropic-IISM model is shown as the red line, and the apparent white-noise in the model is due to the out-of-phase Earth’s velocity. The data in each epoch have been scaled to the mean V_{ISS} and the model is scaled by $\bar{\kappa}$ to match. 83
- 4.1 Dynamic spectra (left panels) and corresponding secondary spectra (right panels) for observations of PSR J0437–4715 from the Parkes 64 m radio telescope from the “20 cm” observing band (top panels) and “40 cm” observing band (bottom panels). The 20 cm dynamic spectrum is a combination of six consecutive observations, which have been concatenated with linear-interpolation applied to the gaps (as described in Section 4.2.1), while the 40 cm dynamic spectrum is from a single observation. The colour in the dynamic spectra shows the normalised flux after subtracting the mean, while the colour in the secondary spectrum shows the log power for these dynamic spectra, which saturates at 2.5 dex above the mean to optimise visualisation of the arcs. 91
- 4.2 Cropped secondary spectrum (left) and corresponding $P_{\text{arc}}(\eta)$ curve (in $\log_{10}(P_{\text{arc}})$ and $\sqrt{\eta}$) for a 20 cm observation on MJD 56946. The colour scale for the secondary spectrum is as described for Figure 4.1 and the white regions were discarded from the analysis since they contain scattered power. The red solid vertical line and dashed lines on the $P_{\text{arc}}(\eta)$ curve show the primary arc curvature measurement and error region, while the black vertical lines are for the secondary arc. The error was calculated from the mean and standard deviation power of the noise in the secondary spectrum, which is shown with the horizontal solid and dotted lines respectively. The measured curvature for the primary and secondary arcs are shown in secondary spectrum with red and black dashed lines respectively. 99
- 4.3 Top-panel: Equivalent curvature η^* (Equation 4.6) for each measurement of the primary scintillation arc across ~ 13 years of observations for PSR J0437–4715. Bottom panels: Equivalent curvatures with best-fit anisotropic model (solid line) as a function of orbital true anomaly θ (left; with fitted annual variation subtracted) and day-of-year (right; with fitted pulsar orbital variation subtracted). The parameter uncertainties have been scaled down by a factor of 1.4 because of the small χ_r^2 value (see Section 4.5.1). 103
- 4.4 Posterior probability distributions for the parameters of the anisotropic model for the primary arc using the MCMC sampler of Veitch and Vecchio (2010), implemented in MATLAB by Pitkin and Romano (2013). 103

- 4.5 Top-panel: Equivalent curvature η^* (Equation 4.6) for each measurement of the secondary scintillation arc across ~ 13 years of observations for PSR J0437–4715. Bottom panels: Equivalent curvatures with best-fit anisotropic model (solid line) as a function of orbital true anomaly θ (left; with fitted annual variation subtracted) and day-of-year (right; with fitted pulsar orbital variation subtracted). The parameter uncertainties have been scaled up by a factor of 1.8 because of the large χ_r^2 value. . . . 105
- 4.6 Posterior probability distributions for the parameters of the anisotropic model for the secondary arc using the MCMC sampler of Veitch and Vecchio (2010), implemented in MATLAB by Pitkin and Romano (2013). 105

INTRODUCTION

In modern astronomy, radio emissions from celestial objects are known to originate from an exceptionally diverse collection of sources. These range from abundant neutral hydrogen gas, to the violent cores of distant galaxies with supermassive black holes that power near-light-speed jets of plasma. The astrophysical insights enabled by discoveries in radio astronomy are equally diverse. One such discovery, of twinkling radio stars called pulsars, revolutionised numerous fields of astronomy and physics including studies of gravity, stellar evolution, the interstellar medium (ISM), dense-matter, magnetohydrodynamics, and high-energy astrophysics. These pulsars are natural clocks and laboratories of extreme physics with strong gravitational fields, and they are distributed throughout the Galaxy. Applications range from precise tests of gravitational theories and searches for gravitational waves, to the development of pulsar-based time-standards and eventually perhaps interstellar spacecraft navigation. This chapter gives an overview of these pulsars and some of their remarkable applications with a particular focus on the technique of pulsar timing and the phenomenon of interstellar scintillation, which are both used in the subsequent chapters to develop new methods for studying properties of the pulsars themselves and the ISM.

1.1 Discovery and Properties of Pulsars

Observations of a radio star in 1951 showed that compact sources can vary in intensity over time (Hewish 1951), analogous to the twinkling of stars in the night sky due to turbulence in the Earth's atmosphere. In this case the phenomenon, called scintillation, was caused by the scattering and subsequent interference of radio waves passing through the Earth's ionosphere. The combined effect of interference patterns from incident waves along the many lines-of-sight from an extended source diminishes the observed scintillation. Thus scintillation is only observed for compact (point-

like) sources with spatially-coherent radiation, which is the physical mechanism underlying the adage ‘stars twinkle and planets do not’. Further observations of compact radio sources showed that scattering by density irregularities in the solar wind could be another cause of scintillation (Hewish 1955). This was called "interplanetary scintillation", and could be used to study the physics of plasma in the solar wind (Hewish et al. 1964).

To investigate the interplanetary scintillation further, Antony Hewish constructed a radio telescope, the "Interplanetary Scintillation Array", and employed a novel observing strategy of sampling the intensity of radio waves with a high time resolution of 0.1 s in initial surveys. This fast sampling would allow for the first detailed studies of interplanetary scintillation in bright and compact radio sources, such as the energetic cores of distant galaxies (quasars), which had recently been discovered (Schmidt 1963). In 1967, while studying quasars with this telescope, Hewish’s student Jocelyn Bell discovered a signal that appeared to be tracking with the sky and regularly pulsating every 1.3 s (Hewish et al. 1968). She had discovered the first pulsar; a rapidly rotating "neutron star" with beamed radio emission that appears to pulse with each rotation as the beam sweeps over the observatory. Pulsars are therefore natural clocks with a lighthouse-like behaviour.

Neutron stars were first theoretically predicted to be the compact stellar remnant of a core-collapse supernova of a massive star by Baade and Zwicky (1934). Soon after Bell’s discovery of the first pulsating radio source, Gold (1968) proposed the rapidly rotating neutron star model for their origin. This model for pulsars, and their supernova origins were strongly supported by the discovery of pulsars with < 0.1 s periods in the Vela (Large et al. 1968) and Crab (Staelin and Reifenstein 1968) supernova remnants (the evolution of pulsars is discussed further in the following section). Early pulsar observations also showed strong modulations to the observed intensity (Lyne and Rickett 1968), caused by interstellar scintillation; scattering and interference by density irregularities in the plasma of the ISM. The observation of intensity scintillations and the first scattering model for the ISM (Scheuer 1968) appeared soon after the discovery of pulsars, because they are near-perfect point sources and emit spatially coherent radiation. Their fortuitous discovery during a study of scintillation is therefore especially apt, because pulsars themselves are ideal tools for studying scintillation. The physics of this interstellar scintillation, and its applications from pulsar observations will be discussed further in Section 1.3.

1.1.1 Pulsar formation and evolution

Pulsars (neutron stars) are the dense remnants of the core-collapse supernova explosions of stars with masses $M \gtrsim 8 M_{\odot}$. They are born highly-energetic, with powerful magnetic fields and a large angular momentum left over from the parent star. Since pulsars typically have a radius of about 10 km, this angular momentum gives the rapid rotation with typical spin periods $P \lesssim 0.1$ s (at birth). Much of the parent star’s mass is expelled in the supernova, but the remnant pulsar retains a mass of approximately $M_p \sim 1.3 M_{\odot}$ (Özel et al. 2012), meaning they are immensely

dense objects. They can also emit high-energy radiation such as x-rays and gamma-rays, in addition to the usual radio beams from the poles of their magnetic field. This energy expenditure is taken from the rotational kinetic energy of the pulsar, which can be observed as a spin-down rate \dot{P} . The spin evolution of a pulsar can be visualised in a " P – \dot{P} diagram" such as Figure 1.1, which plots the spin period and its derivative for all pulsars that are currently in the Australia Telescope National Facility (ATNF) pulsar catalogue (Manchester et al. 2005).

Pulsars are born in the upper-left region of the P – \dot{P} diagram, and rapidly drift towards the cluster of "normal" radio pulsars (with $P \approx 0.5$ s) as they age and their spin and spin-down rates both decrease. A characteristic age (dashed lines in Figure 1.1) can therefore be defined from these quantities alone, $\tau_{\text{age}} = P/(2\dot{P})$. Many of the younger pulsars in Figure 1.1 have known, or suspected associations with supernova remnants (SNRs), which supports the theory of their formation from the death of massive stars. However since supernovae are energetic explosions, the pulsar is expected to receive a high "kick velocity" at birth (often exceeding 100 km/s) and quickly escape the remnant. In such cases, a measurement of a young pulsar's distance, proper motion, and characteristic age can sometimes be used to trace its spatial trajectory backwards in time to an OB association (group of short-lived, hot stars) where the supernova likely took place. For older pulsars ($\tau_{\text{age}} \gtrsim 10^5$ yr), any supernova remnant would have merged with the general turbulent flow of the ISM and be unobservable.

The majority of pulsars occupy a large cluster in the P – \dot{P} diagram populated by solitary, radio-loud pulsars (blue dots in Figure 1.1) with $0.1 \text{ s} \lesssim P \lesssim 3 \text{ s}$ and $10^{-17} \lesssim \dot{P} \lesssim 10^{-13}$, and corresponding characteristic ages of $10^5 \text{ yr} \lesssim \tau_{\text{age}} \lesssim 10^8 \text{ yr}$. The high fraction of solitary "normal" pulsars is again consistent with a supernova birth scenario, since it would be common for the pulsar to be ejected from the system during supernova. The maximum age of solitary pulsars indicates that the decaying spin-down energy of the pulsar relates to the energy in its radio beams. Beyond a certain age, pulsars enter a "graveyard" where they no longer emit radio emission and become radio-quiet neutron stars (grey shaded region in Figure 1.1). Since the beams are powered by the magnetic field and spin of the pulsar, the spin evolution also relates to the characteristic magnetic field strength through $B = 3.2 \times 10^{19} \sqrt{P\dot{P}}$ G (dashed-dotted lines in Figure 1.1 Lorimer and Kramer 2004), which is an approximation to the surface magnetic field strength.

A small population of young, particularly energetic pulsars with enormous magnetic field strengths ($B \gtrsim 10^{14}$ G) occupy the upper-right region of the P – \dot{P} diagram, and these are called "magnetars". However, the main focus for the remainder of this thesis will be on the pulsars at the opposite end of the diagram. These are the "millisecond pulsars", which evolve from the normal pulsars with the help of a binary companion, as discussed in the following section.

1.1.2 Millisecond pulsars and binaries

A separate population of pulsars with fast spin periods ($P \lesssim 30$ ms), slow spin-down rates ($\dot{P} \approx 10^{-20}$), high characteristic ages ($\tau_{\text{age}} \gtrsim 10^8$ yr), and relatively weak magnetic fields ($B \lesssim 10^{10}$ G) occupy the lower-left region of the P – \dot{P} diagram.




Figure 1.1: " P – \dot{P} diagram", displaying spin period P and spin-down rate \dot{P} for all known pulsars with measurements taken from version 1.57 of the Australia Telescope National Facility (ATNF) pulsar catalogue (Manchester et al. 2005). Pulsars with radio emission are shown with blue dots and radio-quiet pulsars (with pulsed emission of higher-energy radiation) are shown with maroon triangles. Pulsars with binary companions have a red circle surrounding the pulsar type marker, and those with associated supernova remnants (SNRs) are surrounded by a yellow star. Pulsars with high-energy (x-ray and gamma-ray) emissions without pulsations, such as Anomalous X-ray Pulsars (AXP) or Soft Gamma-ray Repeater (SGR), are shown with green triangles. The dashed lines are lines of constant characteristic age $\tau_{\text{age}} \propto P/\dot{P}$, and dashed-dotted lines are of constant magnetic field strength $B \propto \sqrt{P\dot{P}}$. The solid line is the pulsar "death line" that encloses the pulsar "graveyard", beyond which electron-positron pair production in the pulsar's inner magnetosphere is expected to cease, which terminates the radio emission.

occupy the lower-left portion of the $P-\dot{P}$ diagram (Figure 1.1). These millisecond pulsars (MSPs), are thought to evolve from normal pulsars through accretion of material from a binary companion star (Bhattacharya and van den Heuvel 1991).

If, during the supernova of a massive star in a binary system, the explosion does not eject the remnant, a pulsar may be born with a main-sequence stellar companion. Since stars commonly spend $\gtrsim 10^9$ years on the main sequence, the pulsar may evolve through the $P-\dot{P}$ diagram to the graveyard during this time. However, at some point the companion star will enter the red giant phase of its evolution and swell to a point where the red giant may orbit within the Roche lobe (gravitationally bound region) of the neutron star and begin to transfer matter via Roche lobe overflow. Some of the angular momentum of infalling plasma is transferred to the neutron star; spinning it up to power the radio beam once again but also reducing its magnetic field strength in the process (Bisnovatyi-Kogan and Komberg 1974). Pulsars reborn in this way are called "recycled pulsars" and have masses in the range of $1.2 \lesssim M_p \lesssim 2 M_\odot$ (Özel et al. 2012). MSPs are recycled pulsars that have generally remained in the accretion phase for the longest times; until the companion sheds enough matter to become a white dwarf, or is evaporated completely (Bhattacharya and van den Heuvel 1991). As a consequence, $\sim 80\%$ of MSPs are observed to be in binary systems, and the majority of these have white dwarf companions (including all of the binary MSPs that are studied in this thesis). The binary orbit also becomes circularised in this process. For one of the pulsars studied in Chapter 2, PSR J1909–3744, we measure the orbital eccentricity to be $e = 1.14 \times 10^{-7}$, which corresponds to a difference between its semi-major and semi-minor axis of just $\sim 4 \mu\text{m}$ despite having a radius of $\sim 570,000 \text{ km}$.

The first millisecond pulsar to be discovered, PSR J1939+2134 (Backer et al. 1982) remains among the fastest known, with a spin frequency of 642 Hz. MSPs exhibit remarkable stability in their spin periods, since they are essentially massive flywheels. The stability of MSP rotation can rival that of the best atomic clocks on decadal timescales (Hobbs et al. 2012). Long-term monitoring of their spin stability, through models of the pulsar that can predict the arrival times of pulses, can therefore be used to create a pulsar-based time standard. The technique of pulsar timing has numerous other applications owing to the complex models required for precise predictions of the pulse arrival time, and these will be discussed in the following section.

The first pulsar to be discovered in a binary system, PSR J1915+1606 (also referred to as the Hulse-Taylor binary after its discoverers), was found to be in a $\sim 7.75 \text{ hr}$ eccentric orbit with another neutron star (Hulse and Taylor 1975). This small orbital period, combined with the large masses of the pulsar and its companion, makes this system highly relativistic and in fact it was the first system discovered that could allow tests of general relativity in the strong-field regime. Subsequent observations of this pulsar (Taylor and Weisberg 1982) showed that the orbital period was decreasing with time in a way that agreed with the prediction of general relativity for energy loss through the emission of gravitational waves (ripples of spacetime curvature that propagate at the speed of light). This was the first evidence for the existence of gravitational waves and

paved the way for continued tests of general relativity and other theories of gravity with future discoveries of relativistic binary pulsars (see Sections 1.2.3.2 and 1.2.4.2).

1.2 Pulsar Timing

The MSPs are incredibly stable natural clocks. Their beamed emission is observed as a train of pulses, spaced equally by the rotation period of the pulsar, which can be measured with high precision. Individual pulses show variation in their shape and amplitude, but when averaged over a ~ 30 min observation of the pulsar, a stable mean "pulse profile" is observed and an average pulse time-of-arrival (ToA) can be measured from this. These ToAs are the fundamental data of pulsar timing, in which they are compared with the prediction of a detailed "timing model" to give timing residuals. The timing model includes a precise measure of time with atomic clocks and a description of the astrometry (position and motion) and orbital dynamics of the pulsar and the Earth, as well as other bodies in the Solar System and the ISM along the line-of-sight. The most precisely-timed MSPs have root mean square (rms) timing residuals of $\lesssim 150$ ns over decades of monitoring (e.g. Shannon et al. 2015), meaning that the timing model is able to accurately track the relative distance between the Earth and these pulsars to a precision of $\lesssim 50$ m despite typically being of order 1 kpc ($\sim 10^{19}$ m) away. With this kind of precision and long-term monitoring of many pulsars, it should be possible to eventually measure slow changes to the distances to pulsars caused by a stretching of space-time by nanohertz-frequency gravitational waves (e.g. Sazhin 1978). Indeed this is the primary goal of a "pulsar timing array" (PTA), in which a set of millisecond pulsars are monitored over many years to attempt to measure a correlated signal in their timing residuals, which will be the signature of these gravitational waves. Throughout this thesis we use data collected for the Parkes Pulsar Timing Array (PPTA) project, which uses the Parkes 64 m radio telescope. Further details of the data processing pipeline for pulsar timing is given in the following section (1.2.1), before an explanation of some of the parameters in the timing model (Section 1.2.3) and sources of noise in timing residuals (Section 1.2.2.1), which will be used extensively in Chapter 2. The search for gravitational waves and other applications of PTAs are then briefly discussed in Section 1.2.4.

1.2.1 Radio observations

Pulsar timing is most commonly conducted with radio-frequency observations collected with a single-dish radio telescope such as the Parkes 64 m radio telescope that is used for the work in this thesis. The technique of pulsar timing is summarised in Figure 1.2, which gives a schematic of the observing and data processing pipeline. A train of pulses emitted from the pulsar first travels through the ISM to become frequency-dispersed and scattered (see Section 1.3) before arriving at the telescope, where it is focused at the feed of the receiver by a parabolic surface. The signal-to-noise ratio of the observation depends on several properties of the pulsar and observing

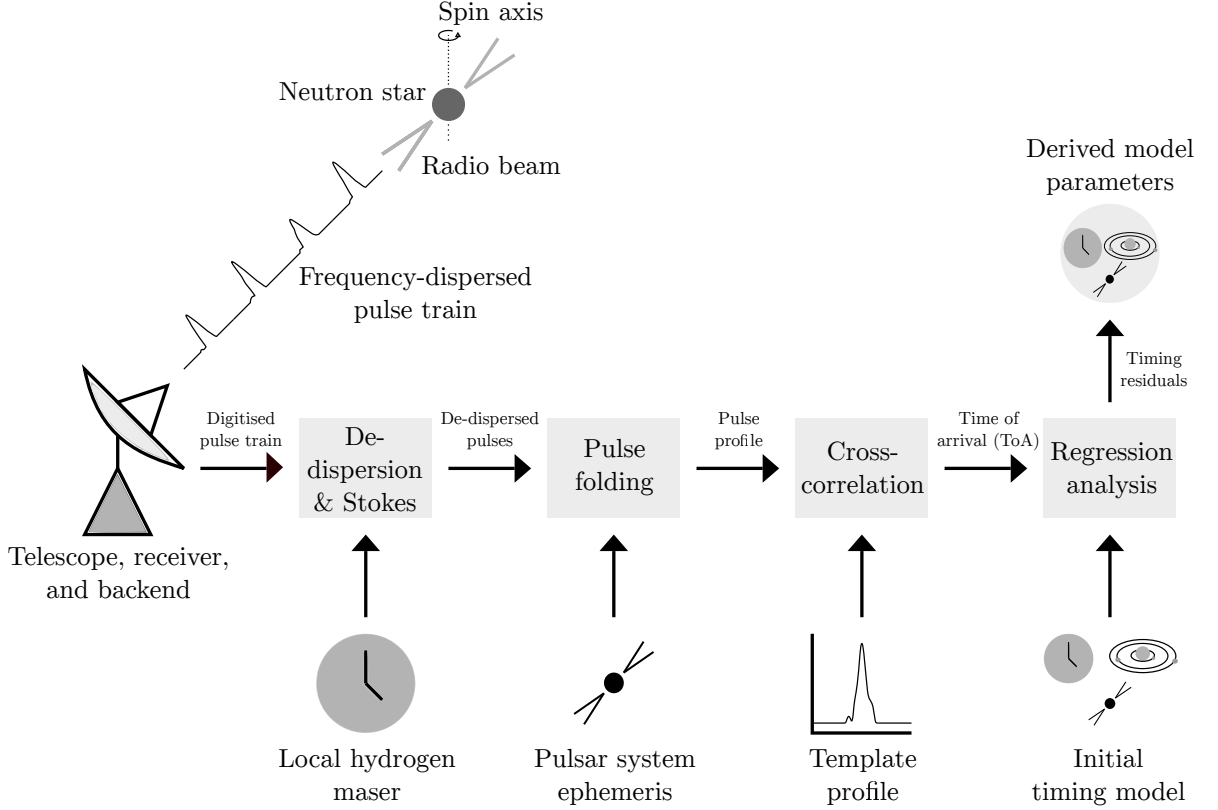


Figure 1.2: Schematic of the major steps involved in the observing and data processing pipeline for pulsar timing, from emission at the pulsar to measurements of timing model parameters. The steps illustrated here are summarised in Section 1.2.2.

system, and is given by the pulsar timing radiometer equation (Lorimer and Kramer 2004)

$$(1.1) \quad \frac{S}{N} = \frac{\hat{S}_f G \sqrt{n_p t_{\text{obs}} \Delta f}}{\beta T_{\text{sys}}} \sqrt{\frac{P - W}{W}}$$

where \hat{S}_f is the pulsar's mean flux density (mJy) at an observing frequency f , $G = A_e/2k$ is the telescope's gain (K Jy^{-1}) for effective collecting area A_e and Boltzmann constant k , t_{obs} is the observing integration time (s), Δf is the observing bandwidth (MHz) around f , T_{sys} is the system observing temperature (K; including sky noise), $\beta \approx 1$ is a correction factor for digitisation losses, and W (s) and P (s) are the pulse width and pulse period respectively. Essentially, Equation 1.1 quantitatively explains the unsurprising result that the ideal observatory has a large collecting area (large G) and a cooled receiver (small T_{sys}) with wide bandwidth (large Δf). Pulsars typically have steep flux density spectra, increasing towards lower frequencies with $\hat{S} \propto f^{-1.60 \pm 0.54}$ (Jankowski et al. 2018). However at higher frequencies, ISM effects are suppressed and the pulse width is typically narrower, giving a smaller $(P - W)/W$ owing to the emission originating from a region of the magnetic field closer to the surface of the pulsar (i.e. lower in the light cone; Radhakrishnan and Cooke 1969; Cordes 1978). Recently Shannon et al. (2015) suggested

that more pulsar timing observations should be conducted at higher frequencies to avoid the significant scatter associated with ISM effects that may limit sensitivity to gravitational waves. However higher frequency observations require a smoother telescope surface (since the rms deviations from a parabolic shape need to be much smaller than the observing wavelength), which is more costly to construct with a large area, while radio frequency interference is more common at lower frequencies. As a consequence of this balance between high and low frequencies, radio telescopes currently used for pulsar timing achieve best results in the range of ~ 300 MHz to ~ 4 GHz (Verbiest et al. 2016), depending largely on the nature of the pulsar itself.

1.2.2 Data processing pipeline

Once received by an instrument at the telescope's focus (for "single dish" observatories), the signal, in the form of a time series of voltages is sent to a "backend" digitising, dedispersion, and folding system before being processed further at a later date. The data processing pipeline for pulsar timing, as shown in Figure 1.2 is as follows:

- The digitised observation is time-tagged by a local clock that is maintained by inputs from a Hydrogen maser frequency standard, and from periodic steering to follow coordinated universal time UTC(GPS) as monitored by a network of atomic clocks used for the global positioning system (GPS). The pulses are de-dispersed to correct for the frequency-dependent time delay introduced by propagation through the ISM (see Section 1.3.1), and Stokes parameters for polarisation are formed (though for the work in this thesis we use only total intensity, Stokes I).
- The dedispersed pulse train is then folded by the known pulsar period (accounting for any delays caused by significant binary motion), using a model (ephemeris) of the pulsar system. The result is a high signal-to-noise ratio pulse profile, which is usually highly stable in shape with time (particularly for MSPs).
- The pulse profile is then cross-correlated with a template profile to find the pulse ToA. The template profile is often an analytical model that is fitted to a particularly high signal-to-noise ratio profile formed by stacking together many individual observations.
- A detailed timing model provides a prediction for the ToAs, and a regression analysis can be used to update parameters of the timing model and results in timing residuals.

The timing model used in the above procedure takes into account all significant sources for time delays to the pulse from its emission at the pulsar to its measurement at the observatory. To accurately track the relative motions of the pulsar and observatory, the ToA is first referenced to an inertial reference frame, which in this case is the Solar System barycentre (SSB; which is actually quasi-inertial). This time transformation takes the topocentric arrival time (t_{topo} ; in UTC as measured by the local clock) and applies corrections for light travel time, frequency-dispersion,

and relativistic effects. If the pulsar happens to be in a binary system, its emission time is also referenced to the barycentre of the pulsar's orbit. The total time transformation in the timing model (initially assuming a solitary pulsar) is given by (Edwards et al. 2006)

$$(1.2) \quad t_{\text{SSB}} = t_{\text{topo}} + \Delta t_{\text{clk}} - \Delta t_{\text{DM}} + \Delta t_{\text{R}\odot} + \Delta t_{\text{S}\odot} + \Delta t_{\text{E}\odot}$$

where Δt_{clk} is the total clock correction to convert the locally measured UTC to a chosen realisation of Terrestrial Time (TT) that is maintained by the *Bureau International des Poids et Mesures* (BIPM) (for example, in Chapter 2 we use a retroactive revision of TT that was computed until to the year 2013; TT(BIPM13)), Δt_{DM} is the total delay due to propagation through the ISM, solar wind, and ionosphere (see Section 1.3.1 for derivation), $\Delta t_{\text{R}\odot}$ is the Römer delay or classical light-travel time from the observatory to the SSB, $\Delta t_{\text{S}\odot}$ is the Shapiro Delay (Shapiro 1964), which is a relativistic correction to the light-travel time for the curvature of space due to massive bodies in the Solar System, and $\Delta t_{\text{E}\odot}$ is the Einstein delay that includes relativistic corrections for time dilation (from the motion and changing gravitational potential at the surface of the Earth) and gravitational redshift (by other bodies in the Solar System). The $\Delta t_{\text{R}\odot}$, $\Delta t_{\text{S}\odot}$, and $\Delta t_{\text{E}\odot}$ terms all require an accurate and precise model of the positions and motions of major bodies in the Solar System (a Solar System ephemeris). The timing model for a binary pulsar will require additional terms for the Römer, Einstein, and Shapiro delays to reference the emission to the barycentre of the pulsar's orbit (another quasi-inertial reference frame). These pulsar binary corrections are computed with parameters within the pulsar ephemeris, which is determined by modelling the ToAs.

Timing residuals will show all deviations from this timing model, including errors in the assumed parameters (which can often be refit with the addition of more ToAs to improve the model) or missing effects such as that of gravitational waves. In addition there are many sources of noise in the residuals originating from the pulsar itself, instrumentation, the ISM, drifting of the time standard, and errors in the Solar System ephemeris. These are discussed in the following section.

1.2.2.1 Noise sources

Noise in timing residuals can be characterised by their colour (e.g. white or red), dependence on observing frequency, and/or spatial correlation (i.e correlation between pulsars). The most fundamental noise process is that of white radiometer noise that originates from the sky background noise and the electronics of the observing system itself, and has a strength proportional to T_{sys} (Equation 1.1). This is the primary noise process used to estimate the ToA uncertainty for a given observation, which is given by cross-correlation with an analytical template pulse profile. However since the pulse shape/phase changes from pulse-to-pulse, small random fluctuations to the average profile are expected and this produces an additional white noise process referred to as "jitter noise" (Shannon et al. 2014). Jitter noise is also a white noise process, but can be

reduced by averaging over a larger number of pulses (i.e. having a longer integration time on the source t_{obs}). Other potential sources for white noise include unknown instrumental effects and errors in calibration. Since these processes are not accounted for to derive the initial ToA uncertainty, scaling factors are commonly required to produce normally-distributed residuals (see Chapter 2).

Many pulsars show significant low-frequency noise ("red noise") in their timing residuals. So-called "timing noise" that is intrinsic to the pulsar commonly originates from random variations to the pulsar's spin period (e.g. Shannon and Cordes 2010). These fluctuations may originate for example from variable coupling between the pulsar crust and a superfluid core (e.g. Jones 1990), spin-down fluctuations linked to changes in magnetosphere torque (e.g. Cheng 1987), or maybe even asteroid belts (Shannon et al. 2013a). Timing noise is therefore frequency-independent (unless magnetospheric changes affect the emission region as in Shannon et al. (2016)), uncorrelated between pulsars, and can be significantly reduced in MSPs owing to their significantly larger spin angular momentum and reduced magnetic fields. Measuring and understanding timing noise is not only essential for other applications of pulsar timing (such as gravitational wave detection), but is also a way to study the neutron star magnetosphere and interior structure.

The timing model (and therefore timing residuals) relies heavily on an accurate local time standard and Solar System ephemeris. However while highly stable on short timescales, atomic clocks drift randomly on longer timescales; as a red stochastic process. Since this is a local phenomenon intrinsic to atomic clocks, the induced red noise in timing residuals from time standard errors are identical (fully correlated) and frequency-independent. By combining the residuals of many stable MSPs to measure the fully correlated component, corrections can be made to the assumed time standard, thus producing a new, pulsar-based realisation of terrestrial time (e.g. Hobbs et al. 2010).

Accurately tracking the motion of the Earth with a Solar System ephemeris is a complex problem of n-body interaction, and therefore small errors in the mass or position of major bodies in the Solar System can induce long-timescale errors in the position of the SSB (and thus red timing noise). Ephemeris errors are frequency-independent but are correlated between pulsars with a dipole signature depending on the angular separation between pulsars. If this correlation is measured, it can be used to determine masses of major bodies in the Solar System (e.g. Champion et al. 2010).

Finally, there is significant frequency-dependent red noise in the timing residuals of many MSPs, owing to random changes in the column density of electrons along the line-of-sight through the ISM to a given pulsar (thus it is uncorrelated between pulsars). This ISM noise can often dominate timing noise in magnitude, but since it is frequency-dependent with a known dependence (which increases with f^{-2}), it can be corrected if there are near-simultaneous measurements of ToAs at multiple observing frequencies. However this correction may not be perfect and residual noise could be significant in the timing residuals of some of the best MSPs.

We discuss this ISM noise and other effects of the ISM in detail in Section 1.3.

1.2.2.2 Parameter estimation in the presence of noise

These noise processes (in particular the red noise), complicate the parameter fitting procedure used to improve the timing model. The noise must be accurately characterised to achieve unbiased parameter measurements and uncertainties from a simple linear least-squares fit. In this basic fitting procedure (which is summarised below, and in Hobbs et al. (2006); Coles et al. (2011)), the timing model is linearised about an initial guess of the parameters (many of the parameters are linear initially; nonlinearity is encountered with some "post-Keplerian" binary parameters; Section 1.2.3.2), and the multivariate regression model is applied

$$(1.3) \quad R = MP + E$$

where R is a column vector of the pre-fit timing residuals of length n (number of ToAs), M is the design matrix of size $n \times m$ for m parameters in the timing model (it contains values of each of the variables at each of the observations), P is a column vector of the fitted parameters with length m , and E is a column vector of "errors", or in this case the post-fit timing residuals for each of the n observations (which also have a mean of zero). The parameters are estimated by minimising the "squared error" $E^T E$ to give the ordinary least squares (OLS) solution

$$(1.4) \quad P_{\text{est}} = (M^T M)^{-1} M^T R$$

where the resulting post-fit timing residuals are assumed for now to be "homoscedastic", meaning they are uncorrelated (i.e. "white") and have equal variance (σ^2 for ToA uncertainty σ). The normalised squared error is the chi-squared value $\chi^2 = E^T E / \sigma^2$, which is used to quantify the "goodness-of-fit". The reduced chi-squared value $\chi_r^2 = \chi^2 / (n - m)$ (where $n - m$ is the number of degrees of freedom) is unity for Gaussian-distributed post-fit residuals with standard deviation equal to the ToA uncertainties σ . If $\chi_r^2 < 1$ then there may be too many parameters in the model, or the uncertainties may be over-estimated, while if $\chi_r^2 > 1$ then more parameters may be required, or the uncertainties are under-estimated (which is extremely common in ToA error estimation, as discussed in the previous section). To account for this deviation, it is common practice to simply scale the covariance matrix of the estimated parameters $\text{Cov}(P_{\text{est}}) = \sigma^2 (M^T M)^{-1}$ by the measured χ_r^2 , although this may result in over-optimistic error estimates.

In practice, the ToA errors are unequal so we instead find the weighted least squares (WLS) solution using the $n \times n$ covariance matrix of the residuals C , where the diagonal elements are the variances σ^2 of each ToA and the remaining matrix is zero for uncorrelated (white) noise. The weighted squared error is $E^T C^{-1} E$, which again is minimised to find a solution. For under-estimated ToA errors, giving $\chi_r^2 > 1$, the covariances in this case are either all scaled by χ_r^2 as in OLS, or have a constant added (this process is described further in Chapter 2). If the residuals are also correlated in time (e.g. if a source of red noise is present) then the off-diagonal elements

in the covariance matrix C are non-zero and the generalised least squares (GLS; Aitken 1936) solution is

$$(1.5) \quad P_{\text{est}} = (M^T C^{-1} M)^{-1} M^T C^{-1} R$$

if C is known. In this GLS problem, the covariance matrix is used to normalise and whiten the residuals and the timing model so it then resembles an OLS problem. However estimating the covariance matrix is non-trivial and if the red, correlated noise is not properly accounted for, the solution in P_{est} and the errors given by $\text{Cov}(P_{\text{est}})$ will be inaccurate.

Currently there are two main approaches for dealing with red noise successfully, which can often be characterised as stationary with a power-law spectral model (Shannon and Cordes 2010). The first is a "Frequentist" approach that uses a linear transformation to whiten the residuals and the timing model, following on from the GLS method described above (and is the primary focus here and in Chapter 2). One possible linear transformation to use in this approach, which was adopted for pulsar timing by Coles et al. (2011), is the Cholesky decomposition of C . In this process, C is factorised into $C = U^T U$ by "Cholesky lower triangle factorisation", and U^{-1} is used as the whitening and normalising transformation for the residuals and timing model in the following way. Applying the U^{-1} whitening transformation gives $R_w = U^{-1} R$, $M_w = U^{-1} M$, and $E_w = U^{-1} E$, and then the GLS problem is reduced to OLS (Equation 1.6) $R_w = M_w P + E_w$ with solution (Coles et al. 2011)

$$(1.6) \quad P_{\text{est}} = (M_w^T M_w)^{-1} M_w^T R_w.$$

This procedure is only optimal if C is known, but can still be used to whiten the residuals and timing model accurately. To do this, the red noise is first estimated through a spectral analysis. In most cases the noise is stationary and the power spectrum $P(f)$ is well-modelled with a single power-law of the form $P_m(f) = A f^{-\alpha}$, where A is the amplitude, f is the frequency of the signal, and α is the spectral index. Since the ToA samples are not spaced regularly in time, $P(f)$ cannot be computed directly with a fast Fourier transform (FFT), and is instead estimated with the "periodogram", which is the squared magnitude of the FFT, scaled to be an estimator of the spectral density. Methods for computing the periodogram include the Lomb-Scargle method (Lomb 1976; Scargle 1982), which we use in Chapter 2 for checking the whitened residuals, and a weighted version by Zechmeister and Kürster (2009), which is used for the original spectral analysis of pre- and post-fit timing residuals (with first-difference pre-whitening before computation for strong red noise with $\alpha > 2$, followed by post-darkening of the result; Coles et al. 2011).

The periodogram of residuals can be modelled with a function of the form

$$(1.7) \quad P_m(f) = \frac{A}{\left[1 + \left(\frac{f}{f_c}\right)^2\right]^{\alpha/2}}$$

which includes a "corner frequency" f_c to account for either a physical turn-over in the pulsar's intrinsic timing noise (e.g. Lasky et al. 2015), or simply the low-frequency power subtracted from the residuals during a fit of the spin period and its derivative (which is a quadratic). Importantly, this model does not diverge at low frequencies. The covariance matrix of the residuals is then estimated from this power spectrum, since $P_m(f)$ is the Fourier transform of the model covariance function $c(\tau)$ (where τ is the time lag between pairs of ToAs) for the red noise. This covariance function is computed for each pair of residuals and the uncorrelated measurement error is added. This method is computationally efficient (usually the most significant computation involved is the matrix inversion to find the whitening transformation U^{-1}), and has been shown to give unbiased parameter measurements and uncertainties if the noise is correctly characterised (Coles et al. 2011, see also Chapter 2). This method is extended, tested, and used for parameter measurement in Chapter 2, where we consider timing residuals with non-stationary noise characteristics (i.e. noise that changes at some known position in the dataset).

The second commonly-used method for dealing with red noise is a "Bayesian" approach which aims to find the "maximum likelihood" values for parameters (including those in a red noise model) as given by the posterior probability distribution (or just "posterior") calculated with Bayes' theorem. If the residuals from OLS have a Gaussian distribution, the solution is also the maximum-likelihood, so the two methods give identical results if the timing model is complete, and the covariance matrix is estimated accurately. Several algorithms have been developed for this Bayesian approach to fitting for the timing model (e.g. van Haasteren et al. 2009; van Haasteren and Levin 2013; Lentati et al. 2014a), each with a different method for sampling the posterior. This approach gives the full probability distribution for fitted parameters, so produces robust parameter uncertainties for a complete model (i.e. with no parameters missing). However unlike GLS, the method of fully sampling the posterior is computationally expensive, and is sensitive to the prior probabilities that are assumed for the parameters.

In Chapters 3 and 4 we use WLS for parameter estimation when fitting models to measured properties of scintillation (since there is not significant correlated noise). The models are non-linear, but a set of initial guesses for P are used, following by iterative WLS fitting until the process converges to a potential solution. In a nonlinear parameter space there may be multiple local $E^T E$ minima, and we use the χ^2 to compare solutions to find the "best-fit", and the parameter errors are derived from $\text{Cov}(P_{\text{est}})$.

1.2.3 Pulsar ephemeris

The component of the timing model that describes the spin, astrometry, and orbital dynamics of the pulsar is referred to as the pulsar ephemeris. It contains parameters describing the motion of the pulsar relative to the SSB, as well as the IISM along the line-of-sight (with a "dispersion measure" model described in Section 1.3.1). These parameters can be measured with GLS from a time series of ToAs for a given pulsar. In this way the pulsars themselves can be studied and

used as tools for experiments such as tests of general relativity and probes of the IISM.

Most fundamentally, each pulsar ephemeris describes the pulsar's spin period P and its derivative \dot{P} as well as its position on the sky in celestial coordinates; right ascension α and declination δ . The transverse motion on the sky for solitary pulsars and the barycentres of binary pulsars, relative to the SSB, is described by the proper motion in α and δ , μ_α and μ_δ respectively (usually expressed in units of milliarcseconds per year; mas/yr). Any constant radial motion simply changes the apparent spin period P with a Doppler shift. For some nearby, and/or precisely-timed pulsars, a timing parallax π (mas) may also be measured, which is a measurement of the curvature of emitted wavefronts. The effect of this parallax is often small, and depends on the pulsar's ecliptic latitude since its measurement depends on the variation in arrival time for a curved wavefront relative to a plane wave at different positions for the Earth's orbit (and perpendicular to the ecliptic plane there is no time dependence). If a parallax can be measured, it gives the distance to the system directly, which is useful for various applications including electron density models of the Galaxy (when combined with the dispersion measure, e.g. Taylor and Cordes 1993; Cordes and Lazio 2002; Yao et al. 2017). Proper motion measurements are useful for understanding the spatial velocities of pulsars, which is related to their "kick velocities" at birth and dynamical evolution in the gravitational potential of the Galaxy (e.g. Hobbs et al. 2005).

Errors in these spin and astrometric parameters each leave a unique signature in timing residuals. Spin parameters P and \dot{P} produce a linear and quadratic term respectively, while position errors produce an annual sinusoid (with a phase that depends on whether the error is in α or δ), proper motion errors produce an annual sinusoid that grows with time from the reference epoch at which the position was accurate, and a parallax error induces a six-month sinusoid.

1.2.3.1 Keplerian binary parameters

A simple non-relativistic binary system can be described by Kepler's laws of motion (neglecting for now any relative motions between the SSB and the barycentre of the pulsar's orbit, e.g. proper motion). In general the orbital motion in the radial direction (which is most relevant to pulsar timing) of such a system is well modelled with these laws using just five parameters; the orbital period P_b , the orbital eccentricity e , the projected semi-major axis $x = a \sin i$ (for semi-major axis a and orbital inclination angle i), the longitude of periastron ω (the angle between the pulsar's ascending passage through the plane of the sky that cuts the centre of mass of the orbit, and its closest approach to this centre of mass at periastron), and the epoch of periastron T_0 . A sixth parameter, the longitude of the ascending node Ω , describes the orientation of the orbit in celestial coordinates (typically defined as the position angle East of North, of the ascending node relative to the centre of mass). Together, these six parameters are the "Keplerian parameters", and are shown schematically in Figure 1.3. However, Ω is difficult to measure through pulsar timing alone since it is not directly related to the radial motion that ToAs are sensitive to. Instead,




Figure 1.3: Schematic definition of angles in a Keplerian binary orbit. Lines in the plane of the sky are grey and lines in the plane of the orbit are black; the intersection of these planes is the line of nodes (grey dashed line). The orbit is tilted from the plane by the inclination angle i . Periastron occurs at a distance of $a(1-e)$, for semi-major axis a and eccentricity e , and the longitude of periastron ω is defined from the line connecting the barycentre to the ascending node. The orientation of the orbit from the perspective of the observer is given by the longitude of the ascending node Ω in the plane of the sky (defined East of North). This figure was adapted from Figure 8.3 of Lorimer and Kramer (2004).

its small effect on radial motion is apparent when relative motions between the SSB and pulsar system are considered, since the projection of the orbit on the sky changes slightly with time (this and other kinematic effects are described in the following section).

Following from the third law of Kepler and Newton's law of gravitation, the binary mass function can be defined, which relates some Keplerian parameters to the masses of the pulsar

m_p and its companion m_c

$$(1.8) \quad f(m_p, m_c) = \frac{4\pi^2 x^3}{GP_b^2} = \frac{(m_c \sin i)^3}{(m_p + m_c)^2}$$

where G is Newton's gravitational constant.

The mass function $f(m_p, m_c)$ itself can be accurately determined for all solved binary orbits because x and P_b are easy to measure to very high precision. However for the masses to be determined separately, observations of relativistic effects are required.

1.2.3.2 Post-Keplerian parameters

There are numerous corrections to the basic Keplerian orbital model that must be made in order to accurately describe the motions of precisely-timed MSPs (where small kinematic effects are detectable) and relativistic systems (where general relativistic parameters can be measured). The parameters used to describe the general relativistic corrections are referred to as "post-Keplerian" (PK) parameters.

Approximating the pulsar and its compact companion (white dwarf or neutron star; no pulsars with black hole companions have yet been discovered) as point masses with no spin angular momentum, the general relativistic PK parameters are a function of only the Keplerian parameters and m_p and m_c . The parameters are required in the model if the radio beam of the pulsar passes through a strong gravitational field, or if the orbit is close enough for the orbital velocity to reach relativistic speeds. Among the binary MSPs in the PPTA, the smallest orbital periods are of order $P_b \sim 1$ day around a white dwarf, which is not close enough to show relativistic effects due to high orbital velocity at the current timing precision (Chapter 2). However, if the orbit is highly inclined (i.e. viewed near to edge-on) from our perspective, the radio emission experiences a detectable Shapiro delay from the strong gravitational field of the companion. The Shapiro delay is described by the "range" r and "shape" s parameters

$$(1.9) \quad r = T_\odot m_c$$

$$(1.10) \quad s = \sin i = T_\odot^{1/3} \left(\frac{P_b}{2\pi} \right)^{-2/3} x \frac{(m_p + m_c)^{2/3}}{m_c},$$

where i is the inclination angle, and $T_\odot = GM_\odot/c^3 = 4.925490947 \mu\text{s}$ is the mass of the Sun in units of time. These parameters are measured in several of the PPTA pulsars, and are the most common method for measuring the masses of the pulsar and companion in combination with the mass function. For more compact or eccentric orbits, parameters including the relativistic advance of periastron $\dot{\omega}$ and the orbital period decay due to gravitational radiation \dot{P}_b can be measured, and these are related to the masses through

$$(1.11) \quad \dot{\omega} = 3T_\odot^{2/3} \left(\frac{P_b}{2\pi} \right)^{-5/3} \frac{1}{1-e^2} (m_p + m_c)^{2/3}$$

$$(1.12) \quad \dot{P}_b = -\frac{192\pi}{5} T_\odot^{5/3} \left(\frac{P_b}{2\pi}\right)^{-5/3} f(e) \frac{m_p m_c}{(m_p + m_c)^{1/3}},$$

where $f(e) = (1 + (73/24)e^2 + (37/96)e^4)/(1 - e^2)^{7/2}$ (Damour and Deruelle 1985, 1986; Peters and Mathews 1963).

Periastron advance is famous for first being observed for the orbit of Mercury prior to general relativity, where the relativistic component is $\dot{\omega} = 43$ arcseconds per century. By contrast, the most relativistic binary pulsars have $\dot{\omega}$ on the order of a few degrees to >25 degrees (e.g. Stovall et al. 2018) per year (including PSR J1141–6545, which is studied in Chapter 3).

Additional geometric timing parameters from various kinematic effects will be introduced in Chapter 2 where they are measured for some MSPs in the PPTA. In brief, these parameters originate from the pulsar’s proper motion, which can alter the projection of the orbit on the sky to give the apparent time derivatives \dot{x} and $\dot{\omega}$ that depend on i and Ω (Kopeikin 1995, 1996), or induce a changing Doppler shift from the small radial acceleration associated with a transverse velocity (the proper motion is tangent to the surface of a sphere centred on the SSB with radius equal to the distance to the pulsar, meaning that a radial acceleration is observed; Shklovskii 1970).

1.2.4 Other applications

We have shown above in Section 1.2.2.1 that some noise processes have direct applications, including studies of neutron stars themselves, development of pulsar-based time standards, and measurements of masses in the Solar System. IISM applications are given in detail in Section 1.3 since this is the focus of Chapters 3 and 4. Highlighted in the following sections below is gravitational wave detection because this is the primary goal of PTAs (e.g. the PPTA), and tests of theories of gravity, which applies to work in Chapters 2 and 3.

1.2.4.1 Ultra-low frequency gravitational wave detection

The era of gravitational wave astronomy began recently, and suddenly, with the direct detection of gravitational waves from a binary black hole inspiral and merger event (Abbott et al. 2016). These short-lived ripples in spacetime were measured by the Laser Interferometer Gravitational Wave Observatory (LIGO Abbott et al. 2009). LIGO is a ground-based experiment involving powerful underground lasers that are pointed along two perpendicular 4 km arms before being reflected by suspended mirrors multiple times and eventually recombined, with any interference being measured by a detector. The mirrors are arranged such that the laser should cancel perfectly through destructive interference (disregarding noise sources) if the arms are identical lengths. As a gravitational wave (with frequency in the range 10 Hz to 10 kHz) passes through the detector, the proper length of the interferometer arms oscillate in a way that depends on the direction,

strength, and polarisation of the wave. A non-zero interference pattern is then detected and properties of the gravitational wave can be reconstructed from the data.

PTAs are sensitive to much lower frequency gravitational waves in a similar way to LIGO, except on a Galactic scale. A set of pulsars in the Galaxy, each emitting radio beams, is analogous to the LIGO test masses (mirrors) that are monitored with lasers. Instead of using interferometry however, changes to the relative distance to each pulsar is tracked precisely using a timing model. Gravitational waves transiting the line-of-sight to the pulsar will again oscillate the proper length and induce a change to the propagation time of the pulses, which can be measured in timing residuals (Sazhin 1978). Since the stability of MSPs is strongest on decadal timescales and longer, PTAs are most sensitive to nanohertz-frequency gravitational waves, emitted for example by the ensemble of supermassive black hole binary systems in the centres of galaxies throughout the Universe, producing a stochastic "gravitational wave background" (GWB). This GWB is currently predicted to be greater in amplitude than any individual sources of these low-frequency gravitational waves (e.g. Zhu et al. 2015). Because the GWB is characteristically stronger at low-frequencies and is random in nature, it is expected to induce red noise in timing residuals. This noise, with a characteristic spectral index, should be first noticed in the most precisely-timed pulsars. Eventually, the expected signal will show as a correlation in the timing residuals between multiple pulsars that depends on the separation angle on the sky for a given pair of pulsars (the exact dependency is called the "Hellings and Downs curve", Hellings and Downs 1983).

Successful isolation of the red noise due to gravitational waves, and measurement of the Hellings and Downs curve relies heavily on precise timing models for the pulsars, and on a thorough understanding of all possible sources of noise in the timing residuals. Since gravitational wave detection is the primary goal of all existing PTAs, the study of pulsar timing models, noise, and interstellar medium effects presented in later chapters is important for this purpose.

1.2.4.2 Tests of general relativity and other theories of gravity

Many pulsars in compact orbits are relativistic enough for measurements of the PK binary parameters described in Section 1.2.3.2. In such systems, the parameters can be used to test general relativity and other theories of gravity. Since each of the parameters is a different function of m_p and m_c , two measurements of any PK parameters can be used to determine measurements of these masses, under the assumption that general relativity is correct. If a third or more PK parameters are measured, the relativistic system can be used as a laboratory for testing general relativity, since additional parameters must be consistent with the implied masses. Relativistic binaries have been used in this way to precisely constrain general relativity in the strong field regime (Kramer et al. 2006).

The first famous example of this kind of system was the Hulse-Taylor binary, for which the measurement of \dot{P}_b gave the first indirect evidence for the existence of gravitational waves. Tests

of general relativity and theories of gravity have since improved with more relativistic systems, such as the first double pulsar system, PSR J0737–3039 (Lyne et al. 2004), which has provided tighter constraints to deviations general relativity in the strong-field than any other pulsar system (Kramer et al. 2006). The system is highly inclined, meaning that $\sin i$ was measured from the Shapiro delay with high precision. As we will see in the following sections, the rapidly changing orbital velocity is ideal for measuring changes to the scintillation of the pulsar, and this was used to independently measure the inclination angle i (uniquely) for this system (Rickett et al. 2014).

Independent measurements of i from scintillation can be used to provide mass measurements or further constrain theories of gravity. These measurements will be particularly valuable for systems where the Shapiro delay is not detectable through pulsar timing. PSR J1141–6545, a relativistic binary pulsar with white dwarf companion, was the first relativistic binary to be modelled in this way (Ord et al. 2002a). In the following section we discuss this method further and revisit the modelling of PSR J1141–6545 in Chapter 3.

1.3 Effects of the Ionised Interstellar Medium

The interstellar medium (ISM) is a broad term to collectively refer to all of the material between the stars in the Galaxy. It is the ordinary matter (i.e. excluding dark matter and dark energy), neutral or ionised, that fills interstellar space, and includes: atomic and molecular gas, dust, plasma, and cosmic rays. The ISM at a given location in the Galaxy can be classified into one of several "phases" depending on the temperature, density, and phase of its local constituents. The three most fundamental phases of the ISM are: the cool neutral medium (CNM) of neutral atomic and molecular gas with temperature $T < 300$ K, the warm ionised medium (WIM) of primarily plasma with $T \sim 10^4$ K, and the hot ionized medium (HIM) of plasma that has been shock-heated to $T \sim 10^6$ K, by supernovae and often forms galaxy outflows (McKee and Ostriker 1977).

Observations of compact radio sources such as pulsars, active galactic nuclei (AGN), and masers are affected in a number of ways by free electrons and magnetic fields in the WIM. This magnetised plasma, with typical electron densities of $n_e \approx 0.1 \text{ cm}^{-3}$ in the spiral arms of the Milky Way (Yao et al. 2017), is referred to as the ionised interstellar medium (IISM). This IISM interacts with all radio waves of centimetre to metre wavelength to produce a number of observational effects such as Faraday rotation, frequency dispersion, scintillation, and image distortion. The magnetised component of the IISM is primarily responsible for the Faraday rotation, which is a change to the position angle of the polarisation of a wave that is proportional to the magnetic field strength in the direction of propagation, n_e , and on the path length. For the remaining observational effects, which are described in more detail in the following sections, the magnetic field of the IISM can largely be ignored. These effects depend most strongly on the density, distribution, and turbulence of the plasma. However, it is assumed that magnetic fields may be

important for the confinement of some of the discussed structures in the IISM (e.g. Romani et al. 1987).

The interaction between the IISM and pulsar radiation can be a significant source of noise for precise experiments involving pulsar timing, such as the search for gravitational waves. But it can also be a useful tool for studying the astrometric and orbital parameters of some pulsars, as well as structures in the IISM itself. In the following sections we introduce observational effects of the IISM and discuss their consequences for pulsar timing and physics.

1.3.1 Frequency Dispersion

One major effect of the IISM on all pulsar observations is the frequency-dependence of the delay on pulse arrival times. It was noted even in the discovery of pulsars (Hewish et al. 1968) that higher frequencies arrived earlier than lower frequencies. This arises because the IISM is a plasma with free electrons in which a small perturbation will produce a restoring Coulomb force, and consequently an oscillation in n_e . The frequency of this oscillation is the "plasma frequency", $f_p \approx 8.5 \text{ kHz} \sqrt{n_e / \text{cm}^{-3}}$ (Lorimer and Kramer 2004). Radio waves at a frequency $f > f_p$ propagating through this oscillating plasma (and its accompanying electric field) will experience a frequency-dependent index of refraction $\mu(f) = \sqrt{1 - (f_p/f)^2}$. The reduced group velocity $v_g(f) = c\mu(f)$ of the propagating wave, therefore produces a frequency-dependent time delay relative to a wave of infinite frequency, which is given by

$$(1.13) \quad \Delta t_{\text{DM}} = \left(D \oint \frac{ds}{v_g(f)} \right) - \frac{D}{c},$$

where the path integral is taken over the fractional distance s from the pulsar at $s = 0$ to the observer at $s = 1$. The observed spread in time of different frequencies emitted simultaneously is referred to as frequency dispersion. For typical radio observations at MHz to GHz frequency, where $f \gg f_p$, the refractive index can be approximated by Taylor series with $\mu(f) \approx 1 - (f_p^2)/(f^2 + 2f^2)$ and the delay in arrival time between two pulses with frequencies f_1 and f_2 is then

$$(1.14) \quad \Delta t_{\text{DM}} = \text{DM} \left(\frac{\mathcal{D}}{f_1^2} - \frac{\mathcal{D}}{f_2^2} \right),$$

where $f_2 > f_1$, $\mathcal{D} = e^2/(2\pi m_e c)$, and DM, or "dispersion measure" is the total column density of electrons along the line-of-sight

$$(1.15) \quad \text{DM} = D \oint n_e ds,$$

and is expressed in units of $\text{cm}^{-3} \text{pc}$. If the distance to the pulsar is known (e.g. from a measurement of its parallax), a measurement of the DM from the curve of Δt with observing frequency f gives a good measure of the mean electron density along the line-of-sight to the pulsar. Many measurements of DM for pulsars with known distances have been used to construct a model of Galactic electron density of the IISM (e.g. Taylor and Cordes 1993; Cordes and Lazio 2002; Yao

et al. 2017). For pulsars where the distance is not independently measured, this problem can be reversed to estimate the pulsar distance using the model of Galactic electron density along the line-of-sight of interest and a measurement of the DM.

1.3.2 Structure of the IISM

The electron density models, derived from DM measurements, describe the large-scale structures in the distribution of electrons in the IISM, e.g. components of the disk, bulge, and halo of the Galaxy, H II regions, supernova remnants, and other known large regions of over- and under-density. However the structure of n_e is far from smooth on smaller scales, because the IISM is highly turbulent. This turbulence is expected to generate a power-law distribution of density irregularities, which can be modelled by the three-dimensional wavenumber k . The general form of this power spectrum is then

$$(1.16) \quad P_{n_e}(q) = C_{n_e}^2 k^{-\beta},$$

where C_{n_e} is the strength of fluctuations, and $\beta = 11/3$ for the canonical spectrum of Kolmogorov turbulence in which energy cascades from larger to smaller scales (as with other fluids). Many pulsar observations show IISM effects that follow this Kolmogorov spectrum closely (e.g. Armstrong et al. 1995; Keith et al. 2013) over many orders of magnitude in k . However the power spectrum is also observed to be truncated at some "inner" and "outer" scales, $l_i = 1/k_i \approx 10^8$ m and $l_o = 1/k_o \approx 10^{18}$ m respectively (Armstrong et al. 1995; Rickett 1990).

Other deviations from this Kolmogorov power-law are also observed, including variations to the exponent β , and AU-scale regions of significant over-density and turbulence that are theorised to be confined by magnetic fields in the Galaxy or ram pressure (Fiedler et al. 1994; Walker 2007). These intensely turbulent structures, which are like storms in the IISM, are referred to as "extreme scattering events" (ESEs) because their scattering efficiency dominates the combined scattering from the entire remaining line-of-sight (scattering and its observational effects will be described further in the following sections). ESEs were first discovered through observations of a quasar, where the compact and dense region of plasma acted as a diverging lens that refracted the flux to its outer edges (Fiedler et al. 1987; Romani et al. 1987). The flux of this quasar initially increased significantly as the ESE began to pass across the line-of-sight, before dropping below the mean for ~ 60 days because of the divergent flux, and it then increased at the trailing edge of the ESE and returned to the mean once it had passed. Since this original observation, many more ESEs have been discovered in quasars and pulsars through similar effects on the flux density, and in general they are found to be approximately AU-scale (e.g. Walker 2007).

The power-law distribution of density irregularities in the IISM means that the DM along a given line-of-sight changes with time as a low-frequency (red) stochastic process. The time-dependence of DM is a combination of the motion of the plasma itself that results in a time-dependent electron density $n_e(t)$, a changing pulsar distance from radial motion $D(t)$, and

transverse motions of the line-of-sight from Solar System and pulsar binary velocities. $D(t)$ is expected to contribute a linear term to $DM(t)$, while annual terms are also observed in some pulsars with transverse spatial variations in n_e at a location near to the Earth (Keith et al. 2013, Chapter 2).

Precise long-term timing of MSPs often reveals a frequency-dependent red noise, which is the time delay induced by the stochastic component of fractional changes to the DM, from $n_e(t)$ (e.g. You et al. 2007; Keith et al. 2013). These "DM variations", $DM(t)$, are a major source of noise in the timing residuals used by PTAs. They are often quantified with the "structure function" statistic, which is the mean squared difference in DM over a time lag τ

$$(1.17) \quad D_{DM}(\tau) = \langle [DM(t + \tau) - DM(t)]^2 \rangle$$

and for Kolmogorov turbulence with $\beta = 11/3$, $D_{DM}(\tau) \propto \tau^{5/3}$ (the exact relationship depends on the diffractive scintillation timescale introduced in the following sections; Foster and Cordes 1990). The delay associated with frequency dispersion, proportional to DM, also has a corresponding phase perturbation to the propagating photons $\phi = 2\pi f \Delta t_{DM} = e^2 DM / (m_e c f)$ relative to a photon at infinite frequency. The structure functions for $DM(t)$, Δt_{DM} and ϕ are related through

$$(1.18) \quad D_\phi(\tau) = (2\pi f)^2 D_{\Delta t}(\tau) = (e^2 / (m_e c f))^2 D_{DM}(\tau),$$

and in this way, measurements of DM variations can be used to study several properties of the IISM. The phase-perturbing nature of the IISM is responsible for the scintillation phenomenon as well, through scattering of incident wavefronts, which is discussed in the following section.

Using measurements of the TOAs at multiple observing frequencies, it is possible to measure the $DM(t)$ with accuracies on the order of one part in 10^5 for precisely-timed pulsars because $\Delta t_{DM} \propto f^{-2}$ (Equation 1.14). The $DM(t)$ red noise can largely be removed in this way, but the procedure is imperfect and residual red noise, which is stronger at lower-frequencies may affect the sensitivity of a PTA to gravitational waves (e.g. Shannon et al. 2015). Measurements of $DM(t)$ can also be used to discover phenomena such as ESEs (e.g. Coles et al. 2015), through the sudden increase in DM (paired with changes in observables of scattering; see the following sections) caused by the compact and dense plasma drifting across the line-of-sight. Similarly, other lensing events (e.g. Lam et al. 2018) or "holes" in the IISM (e.g. Coles et al. 2015) can be discovered through a sudden decrease in DM.

1.3.3 Scattering

It is often appropriate to approximate the IISM for a given line-of-sight as a single, thin (compared with the distance to the source from the observer), phase-changing screen, which is the "scattering screen". For the purpose of illustrating the physics of scattering, this thin screen model is initially used below with the assumption of isotropic scattering (following earlier reviews, e.g. Rickett 1990; Cordes and Lazio 1991; Narayan 1992; Cordes 2002; Lorimer and Kramer 2004). However

the basic mechanisms are the same for more complex scattering geometries, and in Chapters 3 and 4, models for extended (along the line-of-sight) and anisotropic scattering are used.

A two-dimensional wavefront with unit amplitude passing through the thin scattering screen will encounter a phase perturbation $\phi(x, y)$ at transverse position (x, y) due to the power-spectrum of density irregularities $P_{n_e}(q)$ in the screen (Rumsey 1975). Considering a source at infinite distance (for simplicity), the amplitude of the wavefront is $\exp[i\phi(x, y)]$ immediately after crossing the screen, and the amplitude at position (X, Y) on the plane of the observer is given by the Fresnel-Kirchhoff integral (Born et al. 1999; Narayan 1992) assuming small-angle scattering

$$(1.19) \quad \psi(X, Y) = \frac{e^{-i\pi/2}}{2\pi r_F^2} \iint \exp\left(i\phi(x, y) + i\frac{(x-X)^2 + (y-Y)^2}{2r_F^2}\right) dx dy$$

where $r_F = \sqrt{D_e/k}$ is the Fresnel scale for wavenumber k and distance to scattering screen D_e from the observer. The Fresnel scale is the approximate transverse scale at which irregularities in the medium transition from scattering via refraction (for irregularities larger than r_F) to diffraction (for irregularities smaller than r_F). Between the inner and outer scales of isotropic turbulence, the spatial phase structure function simply follows

$$(1.20) \quad D_\phi(z) = \left(\frac{z}{s_d}\right)^{5/3},$$

where $z = x^2 + y^2$ and s_d is the spatial scale of diffractive scattering. s_d is also referred to as the "coherence scale" because it is the transverse separation that corresponds to an rms phase difference of 1 radian (this coherence scale is usually written as s_0 in the literature, but since $s_0 \approx s_d$ and diffractive scattering is the focus for this thesis, the coherence scale is hereafter s_d ; e.g. Rickett 1990).

Scattering is the effect of wavefronts becoming deflected from their original paths, and these phase perturbations can be understood as originating from the scattering of the wavefronts by an angle θ_0 . The coherence scale s_d therefore gives the width of a scatter-broadened image of the source, which appears as a diffuse halo with angular size

$$(1.21) \quad \theta_d = \left(\frac{D_s}{D}\right) \theta_0 = \frac{1}{k s_d},$$

for pulsar at distance D and scattering screen at distance D_s from the pulsar. The additional geometric path length from this scattering means that a narrow pulse from a pulsar is scattered into a longer, quasi-exponential pulse, with $1/e$ "scattering timescale"

$$(1.22) \quad \tau_s = \frac{\theta_d^2 D_e}{c}.$$

If we consider the phase change due to this geometric delay across a finite frequency range (bandwidth) $\Delta\nu_d$, we have the following condition for interference (for a maximum phase difference between interfering waves of ~ 1 radian; Rickett 1977)

$$(1.23) \quad \delta\phi = 2\pi\Delta\nu_d\tau_s \sim 1,$$

and therefore the interference pattern formed by scattered waves is frequency-dependent and changes over the "decorrelation bandwidth" $\Delta\nu_d$.

Two distinct scattering regimes can be defined from this basic model of scattering and phase structure function, with the "scattering strength" parameter $u = r_F/s_d$:

- "Weak scattering" with $u < 1$ implies that phase perturbations within the first Fresnel zone are small ($D_\phi(r_F) < 1$), leading to slow variations in the interference pattern due to refraction, and decorrelation over a wide bandwidth ν_d .
- "Strong scattering" with $u > 1$ has the diffractive scale s_d as the dominant length scale and there are strong phase perturbations within the first Fresnel zone ($D_\phi(r_F) > 1$) that produce an interference pattern with narrow ν_d .

Pulsars are commonly observed in the regime of strong scattering at frequencies $f \gtrsim 1$ GHz, where the diffractive mechanism is dominant, but the "transition frequency" between strong and weak scattering varies with the pulsar and scattering medium. The result of the large phase perturbations in strong scattering is a strong intensity modulation in the interference pattern at the plane of the observer, which drifts across the line-of-sight to produce the phenomenon of scintillation (discussed in the following section). In strong scattering, the large phase perturbations mean there are multiple points of stationary phase on the scattering screen for each point at the plane of the observer (this is multi-path propagation). The size of the coherent regions around the points of stationary phase define the diffractive scale s_d , but all of these regions within a larger scattering region (the "scattering disk") of size $s_r = D_e\theta_d = r_F^2/s_d$ will contribute to the interference pattern at the observer. s_r is the refractive scale and irregularities on this scale will contribute slow intensity variations that modulate the rapid and narrow-bandwidth diffractive variations in this regime (Sieber 1982; Rickett et al. 1984).

A schematic of scattering is given in Figure 1.4, where θ_d and θ_0 are highlighted. The model of scattering has been generalised with the pulsar at finite distance D from the observer and screen at distance D_s from the pulsar, but the basic mechanism described with a simplified model above (Equation 1.19) is the same (e.g. Goodman and Narayan 1985). In Figure 1.4, the transverse velocities of the pulsar, observer, and IISM are shown, and are related to the "effective velocity" V_{eff} of the line-of-sight through the scattering medium at a given point (in this case at point D_s). This V_{eff} determines the rate of scintillation because it is related to the speed at which the interference pattern drifts across the observatory V_{los} (see following section), and therefore how rapidly the patches of correlated intensity, or "scintles", are sampled by the line-of-sight. This picture of scintillation and its observational properties are briefly introduced in the next section, with further introduction given in Chapters 3 and 4, where scintillations are used to study the IISM and properties of pulsars.

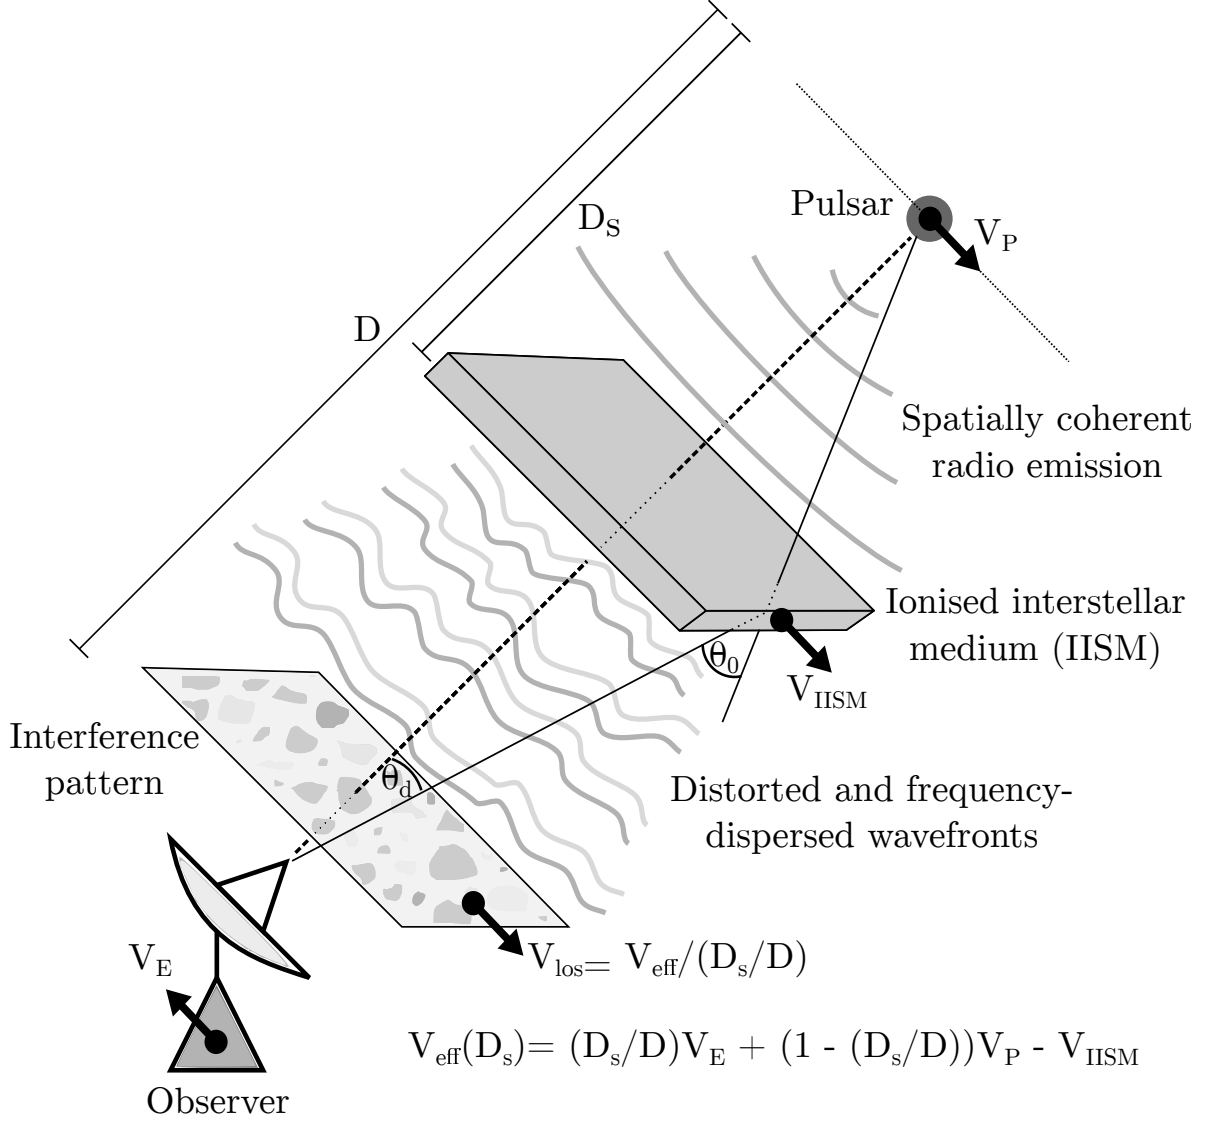


Figure 1.4: A schematic of interstellar medium effects on the emission from radio pulsars. The pulsar initially emits radiation that is spatially coherent. The radiation interacts with the ionised interstellar medium (IISM) during its propagation and becomes scattered and frequency-dispersed (the rotation measure is also altered by the Faraday effect since the IISM is magnetised). The IISM here is assumed to be thin compared to the total distance between the observer and pulsar D , and it is located at a distance D_s from the pulsar. The wavefronts are distorted from scattering and interfere to produce a frequency-dependent interference pattern that is sampled by the observatory. The observatory, IISM, and pulsar have transverse velocities (with respect to the local standard of rest) V_E , V_{IISM} , and V_P respectively (the direction chosen for the figure was arbitrary, and just symbolises the transverse velocity). The effective velocity V_{eff} (defined in schematic) is the transverse velocity of a point along the line-of-sight with respect to the medium (which in this case is defined at the location of the thin screen at D_s). The velocity of the line-of-sight through the diffraction pattern at the observer is V_{los} . The distance to the scattering screen from the observer (used in-text) is $D_e = D - D_s$.

1.3.4 Scintillation

Scintillation, known colloquially as "twinkling", is an observational manifestation of the scattering of light by irregularities in a turbulent medium. The widely-known optical twinkling of stars in the night sky is due to scattering in the Earth's turbulent atmosphere. Interference between scattered wavefronts produces a spatial pattern of intensity variations. These are observed as temporal scintillations if there are relative transverse motions of the source, observer, and/or scattering medium. Compact radio sources scintillate in the same way (though with longer timescales), but the turbulent scattering medium is a plasma, in which the scattering is dominated by the free electrons as described in the previous sections. The first observation of scintillation in radio sources came from observations of flux variation in a radio star, and was found to be caused by scattering from turbulence in the Earth's ionosphere (Hewish 1951). Later, pulsars were discovered serendipitously during a study of "interplanetary scintillation" of the bright and compact cores of distant active galaxies. This scintillation occurs typically on a timescale of a few seconds and is caused by scattering from electron density variations of the Solar wind (Hewish 1955, see also Section 1.1).

The strength of scattering (introduced in the previous section), and correspondingly the strength of scintillation, can be observationally characterised with a measurement of the decorrelation bandwidth Δv_d (also called the "scintillation bandwidth") through $u = \sqrt{f/\Delta v_d}$. The early examples of scintillation given above correspond to the regime of "weak scattering", with $u < 1$. The discovery of pulsars was shortly followed by the observation of "strong scattering", with $u > 1$ in most of the observations (Scheuer 1968; Rickett 1969). Pulsars have since been key tools for the study of small-scale structures in the IISM through measurements of the properties of scintillation.

The two-dimensional interference pattern at the observer produced by wavefronts scattered by a phase-changing screen shows two distinct length scales s_d , s_r . The shorter length scale s_d is associated with diffractive angular scattering and results in fast (of order minutes) strong scintillations for most pulsars. These are modulated by the weaker scintillations on longer timescales (of order days) that correspond to the longer refractive length scales s_r . As the line-of-sight to the source cuts through the interference pattern at the observer's plane with velocity V_{los} , it samples the spatial intensity distribution, leading to scintillations in time, with a characteristic scintillation timescale τ_d . We can therefore define a "scintillation velocity" from the spatial and temporal diffractive scales

$$(1.24) \quad V_{\text{ISS}} = \frac{s_d}{\tau_d},$$

which is equivalent to V_{los} but can be observationally derived from a measurement of τ_d and an estimate of the spatial scale s_d from a measurement of Δv_d . From Equations 1.21, 1.22, and 1.23, it is clear that $s_d \propto \sqrt{\Delta v_d D}/f$, but the exact relationship will depend on the observational definitions of Δv_d and τ_d , the condition for interference (Equation 1.23), and the geometry of the

scattering (e.g. a thin screen at distance D_e from the observer was assumed for Equation 1.22). These factors are included into a constant A_{ISS} , and the observationally-derived scintillation velocity is then

$$(1.25) \quad V_{\text{ISS}} = A_{\text{ISS}} \frac{\sqrt{D \Delta v_d}}{f \tau_d},$$

which can be modelled with an appropriate V_{los} and model for A_{ISS} (e.g. Gupta 1995; Cordes and Rickett 1998).

The "effective velocity" of the line-of-sight through the medium at any position x ranging from $x = 0$ at the pulsar to $x = D$ at the observer, is a linear combination of the transverse components of velocities for the Earth V_E , pulsar V_P , and the IISM itself V_{IISM}

$$(1.26) \quad V_{\text{eff}}(x) = \frac{x}{D} V_E + \left(1 - \frac{x}{D}\right) V_P - V_{\text{IISM}},$$

For a thin screen at $x = D_s$ (such as Figure 1.4), $V_{\text{los}} = V_{\text{eff}}(D_s)/s$, where s is the fractional distance along the line-of-sight $s = D_s/D$. For an extended scattering medium, the effective velocity is integrated along the line-of-sight (weighted by the distribution of scattering material, e.g. Cordes and Rickett 1998) to give V_{los} , which can be used to model the observations of V_{ISS} .

The interference pattern at the observatory is sampled as a function of time and frequency across the receiver bandwidth. Observations of compact sources, such as pulsars, can be presented in a "dynamic spectrum" of the source flux variations with frequency and time (e.g. see Chapter 3, Figure 3.1). It is from this dynamic spectrum that the scintillation bandwidth and timescale are typically measured (from a characteristic scintle, obtained for example by the autocovariance function of the dynamic spectrum, such as in Chapter 3). Since V_{ISS} is sensitive to the transverse velocity of the pulsar, measurements of τ_d and Δv_d can be used to model the components of pulsar motion that are least sensitive to pulsar timing (which is sensitive to radial motion through the pulse ToA). In this way, models of pulsars can be improved with careful measurements of their scintillation, independent to the primary method of pulsar timing. Determining properties of the interstellar plasma, scattering geometry, and the pulsars in this way is referred to as "scintillometry". This thesis demonstrates the complementary approaches of pulsar timing and scintillometry, and an overview of the contents and techniques is given in the following section.

1.4 Thesis overview

In this chapter, we have shown that the IISM interferes with pulsar observations in a way that can be destructive for precision pulsar timing (e.g. DM variations), while valuable for other applications (e.g. studying extreme scattering events, and transverse velocity modelling). In this thesis we investigate properties of pulsars using pulsar timing and scintillation observations in particular. In Chapter 2, we use a new method for measuring parameters of the pulsar timing model in the presence of non-stationary red noise (with WLS; see Section 1.2.2.1), to update

the models for the PPTA pulsars using the latest dataset. For each pulsar we measure and model its noise properties, including the timing noise that is intrinsic to the pulsar, and the frequency-dependent noise from the DM variations, $DM(t)$. We then fit for the parameters in the pulsar timing model, which were in most cases the most precise measurements to date. We find several new measurements of parallax and some post-Keplerian orbital parameters, as well as the most precise distance measurement to any star for PSR J0437–4715.

We then develop a procedure for measuring and modelling the diffractive interstellar scintillation of pulsars from a dynamic spectrum, and we use this to study the long-term scintillation of the relativistic binary pulsar PSR J1141–6545 in Chapter 3. As in Chapter 2, we focus on precision measurements of astrometric and orbital parameters, however this time we model the transverse velocity through the scintillation, rather than the radial velocity measured with pulsar timing (see Section 1.3.4). Modelling the transverse velocity allows us to measure the orbital inclination angle and the proper motion of the pulsar, while these are difficult or impossible to measure with pulsar timing at the current timing precision for this pulsar. We also start to explore the IISM itself by modelling anisotropic scattering and determining the dominant location of the scattering along the line-of-sight (the distance to the scattering screen).

In Chapter 4, we consider a different approach to modelling the long-term scintillation of pulsars. We instead study the Fourier transform of the dynamic spectrum (the secondary spectrum) for PSR J0437–4715, which is the brightest pulsar in the PPTA. Properties of this secondary spectrum are introduced in Chapter 4, but the primary feature is that of parabolic arcs that often appear for high signal-to-noise ratio observations. They are clearest when the scattering is localised, such as in a thin screen, and for PSR J0437–4715 we observe multiple arcs (suggesting multiple discrete screens). Similarly to the diffractive scintillation timescale, the curvature of parabolic arcs depends on the velocity of the line-of-sight through the scattering medium, and in Chapter 4 we use this fact to model the transverse orbital velocity of PSR J0437–4715 as well as other properties of the scattering medium. Remarkably, despite being one of the most precisely timed millisecond pulsars, we are able to measure the longitude of the ascending node with comparable precision to the timing model. The curvature in the arcs is also seen to vary with the Earth’s velocity and we use the relative amplitudes of the Earth and pulsar velocities to precisely constrain the distance and velocity for the thin scattering screens. We also consider whether the scattering is anisotropic, and discuss some possible explanations for the physical structures in the IISM that cause the scattering.

Overall, we demonstrate the close connection between studies with pulsar timing and effects of the IISM. Understanding of the IISM is essential for pulsar timing experiments including the future detection of low-frequency gravitation waves, particularly as timing precision improves in the near future with new telescopes (e.g. MeerKAT, ASKAP, and FAST) and observing instruments (e.g. ultra-wideband receivers for the Parkes 64 m radio telescope). These results are concluded in Chapter 5, including suggestions of future work stemming from the results presented.

TIMING ANALYSIS FOR 20 MILLISECOND PULSARS IN THE PARKES PULSAR TIMING ARRAY

We present timing models for 20 millisecond pulsars in the Parkes Pulsar Timing Array. The precision of the parameter measurements in these models has been improved over earlier results by using longer data sets and modelling the non-stationary noise. We describe a new noise modelling procedure and demonstrate its effectiveness using simulated data. Our methodology includes the addition of annual dispersion measure (DM) variations to the timing models of some pulsars. We present the first significant parallax measurements for J1024–0719, J1045–4509, J1600–3053, J1603–7202, and J1730–2304, as well as the first significant measurements of some post-Keplerian orbital parameters in six binary pulsars, caused by kinematic effects. Improved Shapiro delay measurements have resulted in much improved pulsar mass measurements, particularly for PSRs J0437–4715 and J1909–3744 with $M_p = 1.44 \pm 0.07 M_\odot$ and $M_p = 1.47 \pm 0.03 M_\odot$ respectively. The improved orbital period-derivative measurement for PSR J0437–4715 results in a derived distance measurement at the 0.16% level of precision, $D = 156.79 \pm 0.25$ pc, one of the most fractionally precise distance measurements of any star to date.

2.1 Introduction

The Parkes Pulsar Timing Array (PPTA; Manchester et al. 2013), like its North American (Demorest et al. 2013) and European (Kramer and Champion 2013) counterparts, is a program in which an array of millisecond pulsars (MSPs) is observed regularly over many years. The times of arrival (ToAs) of pulses from MSPs are highly predictable using timing models that describe the spin evolution and astrometric properties of the pulsar and any companions, as well as

taking into account the motion of Earth and pulse propagation through curved spacetime and the interstellar medium (ISM). The parameters of the timing model are determined, or improved, by a least squares-fit of the model to the ToAs. The differences between the measured and predicted ToAs (the “timing residuals”) after this fit contain the measurement error, stochastic fluctuations in the apparent pulsar rotation rate (known as timing noise; Shannon and Cordes 2010), and other unmodelled effects such as those of gravitational-waves (GWs), errors in the assumed time standard, and errors in the Solar System ephemeris. Most of these effects are stronger at lower frequencies, i.e., they have a “red” power spectrum.

The main goal of a pulsar timing array (PTA) is to search for and eventually detect and study nanohertz frequency GWs (e.g. van Haasteren et al. 2011; Demorest et al. 2013; Zhu et al. 2014; Wang et al. 2015), but there are many secondary objectives such as testing general relativity (GR; Freire et al. 2012; Zhu et al. 2015), constraining common models of supermassive black hole and galaxy formation (e.g. Shannon et al. 2015), measuring planetary masses (e.g. Champion et al. 2010), studying the ISM (e.g. You et al. 2007; Keith et al. 2013, hereafter K13), developing pulsar-based time standards (e.g. Hobbs et al. 2012), and precise measurements of properties of the pulsars themselves (e.g. Verbiest et al. 2008, hereafter V08). The latter includes for example the much improved distance and mass measurements for PSR J0437–4715 presented in this Chapter, which will be important for future limits to changes of Newton’s gravitational constant, and the Neutron Star Interior Composition Explorer (NICER) mission that will attempt to measure its radius (Gendreau et al. 2012). The new and improved distance measurements for pulsars presented in this Chapter are also useful for future Galactic electron density models (cf. Cordes and Lazio 2002).

It has long been known that least-squares fitting for the parameters of the timing model can be biased and can underestimate the uncertainties on the parameters when the residuals contain significant red noise. To account for this effect, V08 used Monte Carlo simulations for PSR J0437–4715 to determine parameter uncertainties. Verbiest et al. (2009, hereafter V09), however, prewhitened the residuals for the three pulsars with the most red noise in their sample by fitting harmonically-related sine/cosine pairs (Hobbs et al. 2006). Coles et al. (2011, hereafter C11) demonstrated that this method can result in biased parameter measurements and underestimated uncertainties. For all pulsars, V09 then doubled the formal uncertainties obtained from the fit.

For our work, we use an extension of the “Cholesky” algorithm developed by C11 and implemented in the timing software package TEMPO2 (Edwards et al. 2006; Hobbs et al. 2006). For the results presented in C11 the red noise was modelled as wide-sense stationary, i.e., having a single power spectrum. However, the algorithm only requires that the red noise be described by a covariance matrix. In our data set the red noise is not stationary because the earlier data contain uncorrected fluctuations in the dispersion measure (DM; the column density of electrons along the line of sight to the pulsar) of the ISM, while the later data do not. Accordingly, we use, and

describe below, a modification of the algorithm in C11 that we refer to as the “split-Cholesky” algorithm, which allows for two different red-noise models in the data set.

In Section 2.2, we describe the observations and methodology for determining white-noise parameters. In Section 2.3 we describe the parameters of the timing model. In Section 2.4 we describe the new split-Cholesky algorithm, modelling of the DM variations, and present parameters describing the red noise and DM noise models. In Section 2.5, for each pulsar, we present the new timing model parameter values. In Section 2.6 we present simulations of the split-Cholesky algorithm and compare the method to alternate Bayesian pulsar timing analysis algorithms (e.g. van Haasteren et al. 2011; van Haasteren and Levin 2013; Lentati et al. 2014a), and derive a precise pulsar distance for PSR J0437–4715.

2.2 Observations

The observations used here were published as the extended first Parkes Pulsar Timing Array (PPTA) data release (DR1E) by Manchester et al. (2013). All observations were taken using the Parkes 64m radio telescope. The data set includes observations at three observing bands (with approximate centre wavelengths of 10 cm, 20 cm, and 50 cm) from the PPTA project that commenced in 2005 (these observations alone are referred to as data release one; DR1), along with observations prior to 2005 (in the 20 cm band only) from previous observing programmes. The earliest data were obtained from a timing programme that commenced during the Parkes 70 cm survey (Bailes et al. 1994) and were published by Bell et al. (1997) and Toscano et al. (1999). The sample of pulsars was increased by MSPs discovered during the Swinburne intermediate-latitude survey (Edwards et al. 2001) and elsewhere. Updated timing solutions were published by Hotan et al. (2006) and Ord et al. (2006). Throughout this Chapter, we refer to the archival 20 cm observations taken prior to the PPTA as the “early” data, and the multi-band PPTA observations as “recent” data.

An intensive observing campaign was used to study PSR J0437–4715 in detail. Results were published in van Straten et al. (2001). V08 (for PSR J0437–4715) and V09 (for the other 19 pulsars in the PPTA) combined the earlier data with the initial PPTA data to determine timing ephemerides. Here we use the extra ~ 3 years of data provided by Manchester et al. (2013) to improve on the results of V08 and V09. Along with the extra data span, our new data set provides significantly improved observing cadence and, for the recent data, the ability to remove the effects of DM variations (ΔDM) more precisely than previously possible. Most of the raw observation files used in this analysis are available from the Parkes pulsar data archive (Hobbs et al. 2011).

Throughout this Chapter we make use of the TEMPO2 software package to analyse the pulse arrival times (Hobbs et al. 2006). Our analysis method, described below, relies on knowledge of the noise affecting the residuals. Radiometer noise affects all pulsars and is well-modelled by the ToA uncertainty that is obtained from the template-matching procedure carried out when determining

the ToA. However, in almost all cases the observed scatter in the residuals is greater than that expected from radiometer noise alone. This is not unexpected. Such excess can arise from intrinsic pulse jitter¹ (e.g. Osłowski et al. 2011; Shannon et al. 2014), calibration errors, instrumental effects, or a poor selection of templates used in the template matching process. TEMPO2 currently only has two methods for correcting the measured ToA uncertainties: 1) uncertainties for a set of observations can be multiplied by a scaling factor (this is termed an “EFAC”) or 2) adding a specified amount of extra noise in quadrature with the original uncertainties (termed an “EQUAD”). If both methods are implemented then the resulting uncertainty is:

$$(2.1) \quad \sigma'_i = \text{efac} \times \sqrt{\sigma_i^2 + \text{equad}^2}$$

where σ_i is the original uncertainty for the i ’th observation. Determining the EFAC and/or EQUAD is non-trivial as any low-frequency noise in the residuals must be accounted for and the EFAC and EQUAD parameters are covariant. We follow the procedure below for each data set using the EFACEQUAD plugin for TEMPO2 (Wang et al. 2015):

- Estimate the red noise by fitting a smooth model to the residuals. The default smooth model is a linear interpolation through a set of samples on 100-day intervals. The smooth model (red noise estimate) is then subtracted from the residuals, leaving only the white noise.
- Divide the data into groups based on observing systems that are expected to have the same EFAC and EQUAD values. For example, our data set includes data taken using different “backend” instruments, some of which have identical firmware, bandwidths etc., and are therefore expected to share the same noise properties.
- For a given group selected from the whitened data set, the reduced- χ^2 value (χ_r^2) is calculated. If $\chi_r^2 < 1$ then $\text{efac} = \sqrt{\chi_r^2}$ and $\text{equad} = 0$. Note that this is the only means to reduce a ToA uncertainty as an EQUAD will always increase the uncertainty.
- If $\chi_r^2 > 1$ then the normalised residuals (r_i/σ_i) are determined for a particular grid of EFAC and EQUAD values.
- For each grid position, we determine the probability that the normalised residuals are drawn from a Gaussian distribution (using a Kolmogorov-Smirnov test) to determine optimal EFAC and/or EQUAD.
- By default, we do not include the EFACs and EQUADs for a group that consists of less than 10 ToAs. They are added if it is necessary to produce normally-distributed residuals.

¹For this work, we do not use the jitter parameters introduced by Shannon et al. (2014) because we do not have all of the observation lengths for each ToA in our current dataset.

2.3 The timing model

The timing model describes the spin, astrometric, and orbital properties of a pulsar, the ISM along the line of sight, and requires the use of a terrestrial time standard, Solar System ephemeris, and solar wind model. For this work we use the DE421 Jet Propulsion Laboratory Solar System ephemeris and as a time reference use TT(BIPM2013). We use the default model within TEMPO2 to account for dispersion measure variations caused by the solar wind (wind density at 1 A.U.: 4 cm^{-3} ; Edwards et al. 2006). The pulsars have been observed over many years with various different “backend” instruments. These instruments have different time offsets which we also measure as part of the usual timing fit (with the split-Cholesky method).

The thirteen pulsars in our sample that are in binary systems each have a white dwarf companion. For such systems, we do not expect any time dependencies for the orbital parameters caused by mass loss or spin-orbit coupling. Because of the relatively low mass of the companion stars and relatively long orbital periods, relativistic effects are small in such systems. However, V08 do report a detection of the advance of the longitude of periastron, $\dot{\omega}$, for PSR J0437–4715 that is consistent with that predicted from GR, where the component masses were derived from the Shapiro delay measurement and mass function (Thorsett and Chakrabarty 1999). van Straten (2013) also measures $\dot{\omega}$ for PSR J1022–1001 that is consistent with GR.

For binary pulsars in orbits with small eccentricities, the longitude and epoch of periastron are not well defined and are highly correlated. The ELL1 model (Lange et al. 2001) is used to describe such systems, since it uses a small-eccentricity approximation to avoid this high correlation. For PSRs J1022+1001, J1600–3053, and J1643–1224 this approximation is not valid and so we instead use the DD (Damour and Deruelle 1986) model to describe the binary orbit.

Observed changes in the orbital parameters can be caused by kinematic effects. Changes in the apparent viewing geometry of the orbit caused by proper motion can lead to an apparent time derivative of the projected semi-major axis of the orbit, \dot{x} , and/or $\dot{\omega}$ (Kopeikin 1996). For some pulsars, this kinematic \dot{x} or $\dot{\omega}$ may be detected individually. However, if both are well determined we instead parametrise the effect with Ω and i , which are the longitude of the ascending node and inclination angle of the orbit respectively (van Straten and Bailes 2003). This parametrisation of the orbit describes the annual orbital parallax and is implemented through the T2 model (Edwards et al. 2006), which we use for PSRs J0437–4715, J1713+0747, and J1909–3744. For PSRs J1022+1001, J1600–3053, J1603–7202, J1643–1224, and J2145–0750 we use the \dot{x} measurement to place an upper limit on the inclination angle of the orbit (using $\tan i \leq x\mu/\dot{x}$, where μ is the total proper motion), as was done by Sandhu et al. (1997).

A pulsar can have an apparent spin frequency derivative ($\dot{\nu}$) or orbital period-derivative (\dot{P}_b) as a result of the Shklovskii effect (Shklovskii 1970); an apparent radial acceleration of the system caused by proper motion perpendicular to the line of sight. As we do not include this explicitly in the timing model, we expect non-zero values for these parameters. The $\dot{\nu}$ value caused by the Shklovskii effect is simply absorbed into the intrinsic spin-down rate for the pulsar.

However, if the expected intrinsic $\dot{P}_b = 0$ then as shown by Bell and Bailes (1996), the observed value can be used to determine the distance to the pulsar:

$$(2.2) \quad D = \frac{c}{\mu^2} \frac{\dot{P}_b^{\text{obs}}}{P_b}$$

where D is the pulsar's distance, c the vacuum speed of light, and μ the total proper motion. This equation neglects the differential acceleration of the pulsar and the Earth in the gravitational potential of the Galaxy. The pulsar's distance can therefore be determined from \dot{P}_b^{obs} or from the annual parallax when these effects are taken into account. We choose to decouple these parameters (i.e., fit separately for the two distances) in order to determine whether the distance measurements are consistent. The pulsar timing model requires a distance estimate when determining the orbital parallax and the annual orbital parallax (Kopeikin 1995). In all cases here, we use the default parallax distance. In Section 2.6.2 we discuss the significance of this \dot{P}_b distance measurement for PSR J0437–4715.

2.3.1 Choosing parameters to include in the model

For all pulsars we start with the timing models presented by V09 (and, for PSR J0437–4715, the model from V08). For solitary pulsars, model parameters include the spin (ν , $\dot{\nu}$), astrometric (position, proper motion, and parallax), and ISM parameters (Section 2.4.1). In all cases, we fit for the spin, position (right ascension α , and declination δ), and DM variation parameters (from Manchester et al. 2013). These parameters are all measured at a reference epoch of MJD 54500. In a few cases (in particular where the pulsar is either close to or almost perpendicular to the ecliptic plane respectively) we cannot obtain a significant measurement of the proper motion in declination (μ_δ) or the parallax (π).

To determine whether parameters beyond this base model are required by the data, we make use of the Akaike information criterion (AIC; Akaike 1998), which states that a model is a better fit to the data if

$$(2.3) \quad \Delta\chi^2 > 2k$$

where $\Delta\chi^2$ is the difference of the χ^2 value before and after a fit that includes k new parameters. To determine which parameters to fit for using the AIC, we use the following procedure with solitary pulsars:

- Remove non-essential parameters (proper motion in right ascension μ_α and declination μ_δ , and parallax π) from the timing model (if present) and note the χ^2 value of residuals.
- Fit for each parameter separately and note the new χ^2 value in each case.
- Include the parameter that results in the lowest χ^2 value permanently into the timing model if this parameter also satisfies the AIC.

- Repeat with remaining parameters until either all parameters are included in the timing model, or all remaining parameters fail to improve the timing solution, as determined by the AIC.

This procedure is also applied to pulsars in binary systems. All non-essential binary parameters (post-Keplerian) are initially removed and the AIC is used to find which should be included. If a Shapiro delay (Shapiro 1964) is detectable for the binary system, we parametrise this with the companion mass M_c , and sine of inclination angle of the orbit, $\sin i$, except in the cases of pulsars described with the T2 binary model, for which we link $\sin i$ with the measured inclination angle from the Kopeikin terms (Kopeikin 1996).

2.4 Parameter measurement in the presence of non-stationary red noise

In the original implementation of the C11 algorithm, the red noise was characterized by a power spectrum, from which a covariance function was estimated and finally a covariance matrix was constructed. This was satisfactory for analysis of the DR1 data set. However, DR1E data contain uncorrected DM variations in the 20 cm residuals prior to multi-band observations and thus have additional red noise. We cannot produce a single power-law model for the entire data set because of this extra red noise which may dominate the total red noise in a subset of the data. Instead we produce two separate models to describe the two sources of red noise in the residuals. One model describes the frequency-independent noise present throughout the dataset using a power-law, while the other describes the additional DM noise present only in the early data (Section 2.4.1). We have updated the implementation of the algorithm to synthesize a covariance matrix for DR1E observations using these two red-noise models, which we refer to as the split-Cholesky algorithm.

The method requires two red noise covariance matrices, one for the frequency-independent timing noise and another for the DM noise. Both of these are estimated from the DR1 data alone because for this data set we can estimate and remove the DM noise (described in Section 2.4.1), allowing us to model the timing noise (Section 2.4.2) independently with an analytical model,

$$(2.4) \quad P(f) = \frac{P_0}{\left[1 + \left(\frac{f}{f_c}\right)^2\right]^{\alpha/2}}$$

where P_0 is the amplitude of the power in yr^3 at a corner frequency f_c , and α is the spectral index. A stationary covariance matrix is then computed from this spectral model.

To build the DM noise model that applies to only the early DM-uncorrected observations, we first estimate DM variations, $\Delta\text{DM}(t)$ from the DM-corrected DR1 data using a process described

in Section 2.4.1. We can then choose to create an analytical model of the power spectrum of $\Delta\text{DM}(t)$ from which we compute a covariance function $C_f(\tau)$ (Keith et al. 2013), or we can model the covariance function of $\Delta\text{DM}(t)$ directly (Section 2.4.1). To create the final covariance matrix for the entire data set we must account for the fact that we do not know the mean ΔDM for the DR1 DM-corrected data. It is adjusted to match the end of the earlier uncorrected data with no discontinuity. We then compute the non-stationary $\Delta\text{DM}(t)$ covariance matrix $C_m(t_i, t_j)$ as follows:

$$C_m(t_i, t_j) = C_f(|t_i - t_j|)$$

for $t_i, t_j < T_c$

$$C_m(t_i, t_j) = C_m(T_c, T_c) = C_f(0)$$

for $t_i, t_j > T_c$, and

$$C_m(t_i, t_j) = C_f(|t_i - T_c|)$$

for $t_i < T_c, t_j > T_c$, where T_c is MJD 53430, the time in the dataset beyond which the data is DM corrected. Finally we sum the stationary covariance matrix for the red timing noise (C11), the non-stationary covariance matrix for the $\Delta\text{DM}(t)$, and the diagonal matrix of the variances of the white noise at each sample and apply the Cholesky algorithm as originally formulated. In Section 2.6.1 we demonstrate the effectiveness of this algorithm through the use of simulated data.

2.4.1 Modelling the Dispersion Measure variations

With the advent of the PPTA in 2005, regular observations occurred in multiple observing bands (10cm, 20cm, and 50cm). For these data we are, in principle, able to obtain a measurement of the DM because the group delay is $\propto \lambda^2$. However, as the pulse profiles of MSPs in our sample evolve significantly with frequency (Dai et al. 2015), the effect of an absolute DM on the residuals is coupled with the frequency evolution of the pulse profile and is therefore difficult to determine. We are however able to measure changes in DM, which we refer to as $\Delta\text{DM}(t)$, using interband measurements with the required accuracy ($\approx 1 : 10^5$). Residual errors in these corrections are an important source of noise in the PPTA (and all other PTAs). Such time series (obtained from the same observations we analyse here but with different EFACs and EQUADs) were analysed by K13.

We assume that the $\Delta\text{DM}(t)$ are caused by the movement of the line of sight from the Earth to the pulsar through spatial variations in the ISM. If the velocity of the line of sight were constant, $\Delta\text{DM}(t)$ would simply represent a cut through the ISM in the direction of the velocity. If the fluctuations are due to homogeneous Kolmogorov turbulence then the power spectrum of $\Delta\text{DM}(t)$ would be (K13)

$$(2.5) \quad P_{\text{DM}}(f) \simeq 3.539 D(\tau) \tau^{-5/3} f^{-8/3}$$

where $D(\tau)$ is the structure function at time lag τ . Here $D(\tau)$ is measured in s^2 , τ in s, f in yr^{-1} , and $P_{\text{DM}}(f)$ in yr^3 .

Many of the pulsars in K13 showed a clear linear trend in $\Delta\text{DM}(t)$, possibly indicating a constant spatial gradient over the observing span. In such cases the Earth’s orbital motion causes an annual sinusoid in $\Delta\text{DM}(t)$ and this was also observed by K13. Although this gradient may be part of a stochastic process, for the purpose of analysing our observations it can be considered deterministic and included in the timing model. Accordingly, if it is statistically significant (determined using an AIC test), we fit and remove a linear gradient ($d\text{DM}/dt$) and an annual sinusoid that has been added to the parameters of the timing model in TEMPO2 with the equation

$$(2.6) \quad \text{DM}_{\text{yr}} = A \sin(2\pi \text{yr}^{-1}(t - T_0)) + B \cos(2\pi \text{yr}^{-1}(t - T_0))$$

where A and B are the parameters in the timing model that describe the amplitude and phase of the DM annual variations and T_0 is the reference epoch for the DM measurements. If $\Delta\text{DM}(t)$ has a linear trend and/or annual variation, which we include in the timing model using $d\text{DM}/dt$, and/or A and B , then we need to measure and model the covariance, $\text{Cov}(\tau)$, of the residual $\Delta\text{DM}(t)$.

We measure $\Delta\text{DM}(t)$ with a 5 yr^{-1} cadence for each pulsar using the DR1 multi-band data, and convert each measurement to a time delay in the 20 cm band using $t_{\text{DM}}(\nu) = \text{DM}/(K\nu^2)$, where ν is the observing frequency (1400 MHz in this case) and $K = 2.410 \times 10^{-4} \text{ MHz}^{-2} \text{ cm}^{-3} \text{ pc s}^{-1}$ (You et al. 2007). We then model the covariance functions of the detrended (if required) $\Delta\text{DM}(t)$ with a function of the form

$$(2.7) \quad \text{Cov}(\text{DM}(t)) = a \exp\left(-\left(\frac{\tau}{b}\right)^\alpha\right).$$

The covariance (Cov) is a function of lag τ (in days), a is the amplitude of the red noise (in s^2), b is the characteristic timescale (in days), and the exponent was chosen to be $\alpha = 2$ so that the covariance function will have a positive definite Fourier transform, which is the power spectrum.

The annual variation, linear trend, and covariance function parameters that we have used to construct a DM model for each pulsar are given in Table 2.1. In some cases, no $d\text{DM}/dt$ is apparent in the data, but the $\Delta\text{DM}(t)$ noise is nevertheless well modelled by the covariance function since it is small. For all pulsars with the exception of PSR J0437–4715 (which is very well modelled by a Kolmogorov power law), we find that these covariance function models successfully whiten the residuals and we therefore include them in our combined red-noise models. For PSR J0437–4715 we use the Kolmogorov power law presented in K13 to model the DM noise present in the early data.

For PSR J1603–7202 there is an extreme scattering event (ESE; Fiedler et al. 1987), lasting ~ 250 days, which was reported by K13 and is described in detail in Coles et al. (2015). The ESE dominates the shape of the non-DM-corrected 20 cm residuals in that region. We examined the early 20 cm data searching for ESEs comparable with the one reported by K13 and found none. We obtained the covariance model for $\Delta\text{DM}(t)$ by linearly interpolating across the ESE before computing the covariance function.

Maitia et al. (2003) reported on a 3-year-long ESE (centred on the year 1998) detected in the direction of PSR J1643–1224 by studying flux variability of the pulsar using observations undertaken at the Nancay observatory. Unfortunately we have poor data during this time and consequently do not find evidence for such an event. For PSR J1713+0747 we see a peculiar “drop out” in the $\Delta\text{DM}(t)$ that is possibly related to an ESE-like structure (Coles et al. 2015). We therefore computed the covariance of $\Delta\text{DM}(t)$ for PSR J1713+0747 by interpolating across this drop out.

The deterministic terms in the model for each pulsar include an absolute DM measurement (which we hold constant because it is covariant with pulse profile evolution; Dai et al. 2015), and if required, the additional terms $d\text{DM}/dt$, A and B , or $\Delta\text{DM}(t)$ measurements taken at a spacing of Δt_{DM} days to remove any residual red noise. The sampling interval Δt_{DM} is ideally the widest spacing required to remove the red noise since the measurements add white noise to the residuals. We start with the Δt_{DM} published with the dataset in Manchester et al. (2013). We sample more frequently if there is residual red noise because the published Δt_{DM} were selected to minimise the rms residual in the best band for each pulsar, and not to absorb the most red noise. However adding the deterministic DM parameters often allows us to reduce the $\Delta\text{DM}(t)$ measurement cadence relative to that published in Manchester et al. (2013).

2.4.2 Modelling the red timing noise

The analytical model (Equation 2.4) for the frequency-independent red noise (e.g. timing noise) must be estimated with data that have the frequency-dependent red noise contributions from DM variations already removed. We therefore use the DM-corrected DR1 data and fit the analytical model (Equation 2.4) to the power spectrum. In most cases we find that $f_c = 1/T_1$ fits the data reasonably well, where T_1 is the length of the DR1 data. For these cases we assume that the model will fit the entire DR1E with $f_c = 1/T_E$, where T_E is the length of the DR1E dataset, and we scale P_0 to this new f_c accordingly.

In many cases, extrapolating the timing noise model by adjusting f_c will underestimate the noise at $f \approx 1/T_E$ because fitting for the spin frequency and its derivative removes much of the low-frequency power at $f \approx 1/T_1$. For pulsars where the timing noise dominates the DM noise, we create a single spectral model for the entire DR1E span directly. If this model ultimately whitens the residuals adequately we do not need to estimate the covariance matrix of the DM variations. If however the timing noise model does not extrapolate well and the DM noise is too significant to ignore, we construct a red-noise model using the DR1E data, but reduce the amplitude, P_0 , to account for the known DM noise in the data. The amplitude need only be reduced such that the residuals are sufficiently whitened by the red-noise model.

For PSRs J0613–0200, J1600–3053, and J2145–0750 we find that the red-noise model from the DR1 data alone does not extrapolate well over the entire dataset since the power spectrum of the DR1E data is observed to turn over at $f_c > 1/T_E$. This was not obvious when analysing

2.4. PARAMETER MEASUREMENT IN THE PRESENCE OF NON-STATIONARY RED NOISE

Table 2.1: Parameters describing the DM model used for each pulsar. DM noise in the earliest residuals is described by the DM covariance function parameters (a and b) in Equation 7, which are calculated after the removal (if necessary) of a linear trend, dDM/dt , and annual variations described by the sine (A) and cosine (B) amplitudes. Δt_{DM} is the separation of ΔDM measurements in the multi-frequency section of the dataset for each pulsar.

Pulsar Name	DM (cm^{-3}pc)	dDM/dt ($\text{cm}^{-3}\text{pc yr}^{-1}$)	Timing model parameters		Δt_{DM} (days)	DM covariance	
			A ($10^{-4} \text{ cm}^{-3}\text{pc}$)	B ($10^{-4} \text{ cm}^{-3}\text{pc}$)		a (s^2)	b (days)
J0437–4715*	2.64498	–	–	–	60	–	–
J0613–0200	38.7756	–	-1.0 ± 0.2	-1.0 ± 0.2	365.25	5.5×10^{-14}	317
J0711–6830	18.4099	$(9 \pm 7) \times 10^{-5}$	–	–	200	5.8×10^{-13}	331
J1022+1001	10.2531	–	–	–	200	1.2×10^{-13}	153
J1024–0719	6.48803	$(2.2 \pm 0.6) \times 10^{-4}$	–	–	–	–	–
J1045–4509	58.1438	$(-3.66 \pm 0.13) \times 10^{-3}$	-8.1 ± 2.3	-0.9 ± 4	182.62	1.7×10^{-11}	179
J1600–3053	52.3249	$(-6.3 \pm 0.3) \times 10^{-4}$	–	–	125	5.2×10^{-13}	146
J1603–7202	38.0489	–	–	–	100	6.3×10^{-13}	64
J1643–1224	62.4143	$(-1.23 \pm 0.005) \times 10^{-3}$	-2.9 ± 0.7	-5.9 ± 0.7	365.25	1.5×10^{-12}	113
J1713+0747	15.9903	–	–	–	365.25	2.6×10^{-14}	171
J1730–2304	9.61634	$(5.6 \pm 0.5) \times 10^{-4}$	–	–	–	–	–
J1732–5049	56.8365	$(8.8 \pm 1.2) \times 10^{-4}$	–	–	–	–	–
J1744–1134	3.13695	$(-1.32 \pm 0.18) \times 10^{-4}$	–	–	–	–	–
J1824–2452A	119.892	$(1.15 \pm 0.08) \times 10^{-3}$	–	–	82.5	–	–
J1857+0943	13.2984	$(2.8 \pm 0.5) \times 10^{-4}$	–	–	–	–	–
J1909–3744	10.3932	$(-2.97 \pm 0.06) \times 10^{-4}$	–	–	105	–	–
J1939+2134	71.0227	$(-5.9 \pm 0.3) \times 10^{-4}$	2.4 ± 1.1	1.6 ± 1.2	50	3.2×10^{-13}	112
J2124–3358	4.60096	–	–	–	–	–	–
J2129–5721	31.8509	$(-1.6 \pm 0.4) \times 10^{-4}$	–	–	–	–	–
J2145–0750	8.99761	$(1.2 \pm 0.3) \times 10^{-4}$	–	–	–	–	–

* Kolmogorov model from K13 is used to model the DM noise instead of a covariance function.

the DR1 data alone. We observe a similar turn over for PSR J1909–3744, however the dataset used for this pulsar includes multi-band observations across the entire span and so is largely DM-corrected (see Section 2.5.16). For this pulsar, we therefore model the power spectrum of the entire data set and use this with the original Cholesky algorithm.

The parameters describing the red timing noise models for each pulsar are given in Table 2.2. The DM models that are used with these red-noise models are described in Section 2.4.1, with parameters in Table 2.1.

2.4.3 How do we know when our models are optimal?

The split-Cholesky algorithm produces the best linear unbiased estimators if, and only if, the residuals after whitening are white and normally distributed. If this is not the case, one or more of the red-noise model, EFACs, or EQUADs are incorrect. We utilise two tests to check that the final residuals are what we require to have confidence in the parameter measurements and uncertainties.

Table 2.2: Parameters for the red-noise model for each pulsar. The parameters are used to describe the frequency-independent noise in the data with a power law (Equation 4), where α is the spectral index and P_0 is the power at corner frequency, f_c .

Pulsar Name	α	P_0 (yr ³)	f_c (yr ⁻¹)
J0437-4715	3	1.14×10^{-27}	0.067
J0613-0200	5	5.5×10^{-28}	0.40
J0711-6830	—	—	—
J1022+1001	—	—	—
J1024-0719	6	1.8×10^{-23}	0.066
J1045-4509	3	2.0×10^{-24}	0.059
J1600-3053	2.5	3.0×10^{-28}	0.40
J1603-7202	2.5	1.2×10^{-25}	0.065
J1643-1224	4	1.5×10^{-25}	0.15
J1713+0747	2	3.0×10^{-27}	0.059
J1730-2304	—	—	—
J1732-5049	2	3.0×10^{-27}	0.25
J1744-1134	—	—	—
J1824-2452A	3.5	3.0×10^{-25}	0.17
J1857+0943	—	—	—
J1909-3744	2	1.2×10^{-29}	0.50
J1939+2134	4.5	1.5×10^{-24}	0.064
J2124-3358	3.5	2.0×10^{-25}	0.06
J2129-5721	1	1.0×10^{-27}	0.065
J2145-0750	4	3.0×10^{-26}	0.3

The first of these is an Anderson-Darling (AD) test for normality (Anderson and Darling 1954). We apply this to our whitened, normalised residuals to determine whether they are consistent with a normal distribution with $\mu = 0$ and $\sigma = 1$. The result is the modified AD statistic, A^{*2} . This is used to test the hypothesis that the residuals obey the described normal distribution. The hypothesis is rejected if $A^{*2} > 2.492$ with 5% significance (Stephens 1974), since the expected distribution function for the normalised residuals is known. For example, panel (a) of Figure 2.1 shows the post-fit whitened and normalised residuals for PSR J1713+0747 and panel (c) shows the cumulative distribution of these residuals. The modified AD statistic, $A^{*2} = 1.03$, indicates that the normalised residuals are consistent with a standard normal distribution.

We then test for “whiteness” by inspecting the power spectrum of the whitened, normalised residuals (see also C11). Panel (b) of Figure 2.1 shows the Lomb-Scargle periodogram of these residuals for PSR J1713+0747. We compute this power spectrum by first converting the whitened components in panel (a) to a time series using the ToAs of the unwhitened residuals. The frequency axis of the resulting power spectrum is not well-defined, and we term it the pseudo-frequency. However the whitening process provides a diagonally-dominant whitening matrix such that low frequencies in the unwhitened residuals translate to low frequencies in the pseudo-time-series of

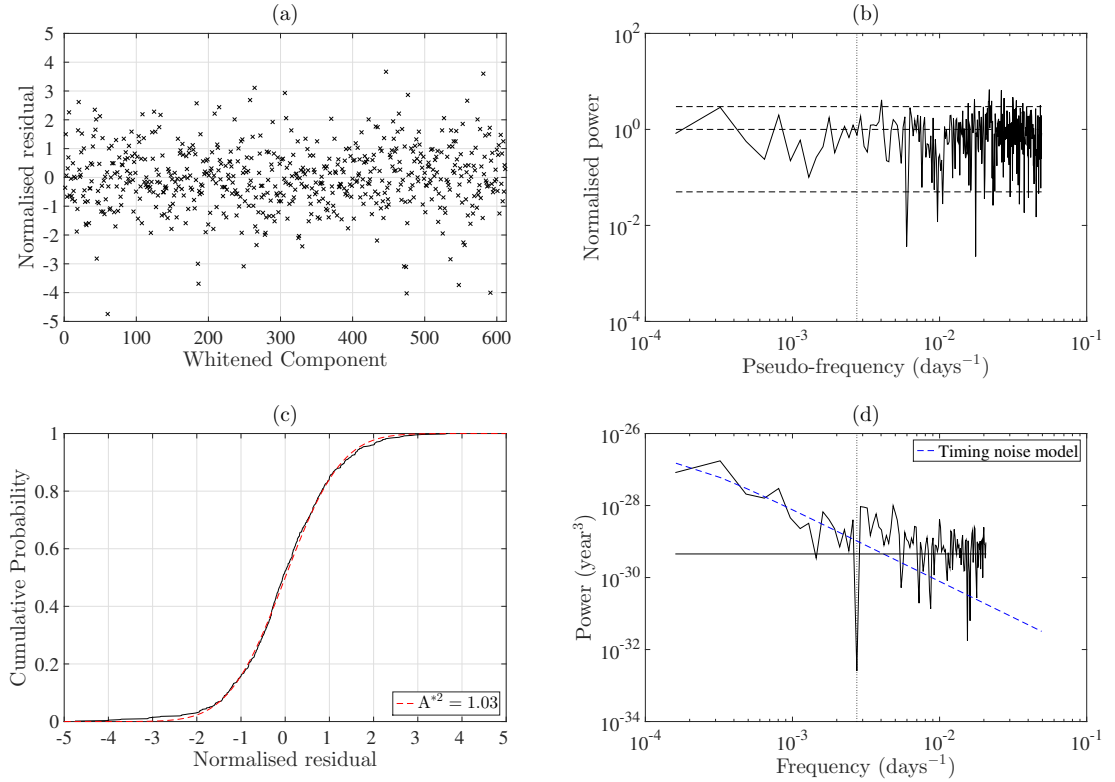


Figure 2.1: (a) Whitened and normalised post-fit residuals for PSR J1713+0747. (b) Power spectra of whitened and normalised post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and dashed lines indicate expected mean and variance for the spectrum. The pseudo-frequency is determined by converting the whitened components to a pseudo-time-series using the ToAs of the unwhitened residuals. (c) Cumulative distribution of whitened and normalised post-fit residuals (solid line) with expected distribution based on normal distribution with zero mean and unit variance (dashed line). Modified AD statistic for this distribution with the expected distribution is labelled. (d) Power spectra of post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and solid flat line is an estimate of the white-noise level. Dashed line is the power law model of the frequency-independent timing noise.

the whitened components. The pseudo-frequency power spectrum can therefore be a useful test of the red-noise model used in the whitening process, since it must be flat. The power spectrum of the unwhitened residuals is shown in panel (d) of Figure 2.1 with the timing noise model used for this pulsar. If the test fails for any pulsar, we update the red-noise model and re-fit the timing model until the residuals are successfully whitened as required.

2.5 Results

In this section, we present our final timing solutions for each pulsar, with the post-fit residuals given in Figure 2.2. We compare our measured parameters with those in the literature, including

VLBI measurements and distances derived from the pulsar DMs using the Taylor and Cordes (1993, hereafter TC93) and NE2001 (Cordes and Lazio 2002) Galactic free electron distribution models (giving distances accurate to approximately 25% and 20% respectively).

Much of the comparison in this section will be with V08 and V09, which used a subset of our dataset and is therefore not independent. However, V08 and V09 used a different Solar System ephemeris, DE405, and time standard, TT(TAI), to our analysis, which results in apparent changes to some parameters. Use of the DE405 ephemeris in particular induces significantly different position parameters compared to the newer DE421 ephemeris, while the use of TT(TAI) changes the apparent spin frequency and its derivative. For each pulsar we find that inconsistencies in the position parameters are explained by the use of this different ephemeris.

Although K13 measured DM variations for all PPTA pulsars and determined an optimal sampling time, Manchester et al. (2013) did not publish any Δ DM measurements for seven pulsars. This is because Δ DM values were only published by Manchester et al. (2013) if their inclusion improved the rms residuals. For our work, we must account for all red noise in the residuals since this is required for the split-Cholesky algorithm (this process is described in Section 2.4). The testing of red- and white-noise models for each pulsar is described in Section 2.4.3, where we present the analysis of PSR J1713+0747 as an example. We do not present the analysis of the other pulsars, with the exception of PSR J0437–4715, which shows some variation that we discuss below in Section 2.5.1.

For each pulsar we also compare the various distance measurements available in the literature to check that our new distance measurements from π or \dot{P}_b are consistent with these. For each parallax measurement, we calculate the Lutz-Kelker (Lutz and Kelker 1973) corrected parallax value and the corresponding corrected distance using the method of Verbiest et al. (2012) and with the mean flux of each pulsar given in Manchester et al. (2013). The parallax measurements for PSRs J1024–0719, J1045–4509, J1603–7202, and J1857+0943 have high fractional uncertainties (parallax measured with less than $3\text{-}\sigma$ confidence) and therefore the bias corrected values for these closely resemble the prior distributions used. The measurements however are important for placing an upper-bound on the parallax value, or similarly, a lower-bound on the distance measurement. A table of these values, and the distances from \dot{P}_b and DM are given in Table 2.3.

For each pulsar, we derive the characteristic age, surface magnetic field strength, and energy-loss rate from the spin-down. Where available, we use the pulsar mass function (Thorsett and Chakrabarty 1999) and a precise Shapiro delay measurement to calculate a pulsar mass, and use the kinematic \dot{x} measurements to place a limit on the inclination angle of the orbit. These derived parameters are given in the Appendix with the tables of parameters.

The red-noise models, ToA files, and parameter files for each pulsar are available from the The Commonwealth Scientific and Industrial Research Organisation (CSIRO) Data Access Portal².

The final parameters for the seven solitary pulsars J0711–6830, J1024–0719, J1730–2304,

²<http://dx.doi.org/10.4225/08/561EFD72D0409>




Figure 2.2: Final post-fit residuals for each of the pulsars in our sample. The vertical range of each subplot is given below the pulsar name. 43

Table 2.3: Parallax values and distance measurements for pulsars in our sample. Parallax and parallax-derived distance values are corrected for the Lutz-Kelker bias using the method of Verbiest et al. (2012). TC93 and NE2001 DM distances have approximate uncertainties of 25% and 20% respectively.

Pulsar Name	Measured values		L-K bias corrected values		DM distances (kpc)	
	Parallax π (mas)	P_b distance (kpc)	Parallax π (mas)	π distance (kpc)	TC93	NE2001
J0437–4715	6.37 ± 0.09	0.15679 ± 0.00025	6.37 ± 0.09	0.1569 ± 0.0022	0.14	0.189
J0613–0200	0.86 ± 0.13	–	0.81 ± 0.13	$1.09^{+0.18}_{-0.14}$	2.19	1.700
J0711–6830	–	–	–	–	1.04	0.854
J1022+1001	1.1 ± 0.3	1.2 ± 0.5	1.0 ± 0.3	$0.74^{+0.19}_{-0.13}$	0.60	0.443
J1024–0719	0.5 ± 0.3	–	$0.5^{+0.21}_{-0.16}$	$1.1^{+0.4}_{-0.3}$	0.35	0.381
J1045–4509	2.2 ± 1.1	–	$0.29^{+0.5}_{-0.15}$	$0.34^{+0.2}_{-0.10}$	3.24	1.945
J1600–3053	0.48 ± 0.11	–	0.43 ± 0.11	$1.8^{+0.5}_{-0.3}$	2.67	1.581
J1603–7202	1.1 ± 0.8	3.9 ± 1.8	$0.25^{+0.4}_{-0.12}$	$0.53^{+0.4}_{-0.16}$	1.64	1.159
J1643–1224	1.27 ± 0.19	–	1.18 ± 0.19	$0.74^{+0.12}_{-0.10}$	>4.86	2.320
J1713+0747	0.86 ± 0.09	3.1 ± 1.2	0.84 ± 0.09	$1.12^{+0.12}_{-0.11}$	0.89	0.889
J1730–2304	1.5 ± 0.3	–	1.2 ± 0.4	$0.62^{+0.15}_{-0.10}$	0.51	0.529
J1732–5049	–	–	–	–	1.81	1.392
J1744–1134	2.53 ± 0.07	–	2.52 ± 0.07	0.395 ± 0.011	0.17	0.415
J1824–2452A	–	–	–	–	3.64	3.042
J1857+0943	0.5 ± 0.3	–	$0.15^{+0.2}_{-0.07}$	$1.2^{+0.7}_{-0.4}$	0.70	1.168
J1909–3744	0.81 ± 0.03	1.140 ± 0.012	0.81 ± 0.03	1.23 ± 0.05	0.55	0.457
J1939+2134	0.52 ± 0.16	–	0.40 ± 0.16	$1.5^{+0.5}_{-0.3}$	3.58	3.550
J2124–3358	2.4 ± 0.4	–	2.15 ± 0.4	$1.39^{+0.08}_{-0.06}$	0.25	0.268
J2129–5721	–	3.2 ± 1.5	–	–	>2.55	1.686
J2145–0750	1.84 ± 0.17	–	1.80 ± 0.17	$0.53^{+0.06}_{-0.05}$	0.50	0.566

J1744–1134, J1824–2452A, J1939+2134, and J2124–3358 are given in Table 4. The binary pulsars are separated by the binary model used to describe their orbit. Parameters for the small-eccentricity pulsars described by the ELL1 model, PSRs J0613–0200, J1045–4509, J1603–7202, J1732–5049, J1857+0943, J2129–5721, and J2145–0750 are presented in Table 5; DD model pulsars J1022+1001, J1600–3053, and J1643–1224 are presented in Table 6; and T2 model pulsars J0437–4715, J1713+0747, and J1909–3744 are presented in Table 7.

2.5.1 PSR J0437–4715

PSR J0437–4715 is the closest MSP currently known and the brightest at radio wavelengths. van Straten et al. (2001) presented a timing solution including the full three-dimensional geometry of the binary orbit. An updated model was presented by V08. This included a precise distance estimate derived from an orbital period-derivative measurement.

In Figure 2.3, we show the two components of the red-noise model with the power spectrum of the post-fit residuals (panel d). As shown in panel (b), the model successfully whitens the residuals. Significant uncorrected DM noise is present in the early data. Therefore, as expected, the timing noise model underestimates the total noise, whereas the DM noise model alone

overestimates the total noise (since it does not apply to the entire dataset). There is excess noise at all frequencies (the mean of the power spectrum of normalised residuals is > 1). This could result, for example, by additional uncorrected short-timescale correlated noise in the residuals. The normalised residuals do not pass the AD test ($A^*2 = 6.75$). This may be because of instrumental effects, or because of pulse jitter (Shannon et al. 2014). Non-Gaussianity has been detected previously for PSR J0437–4715 by Lentati et al. (2014b). However it was shown that, at this level, the non-Gaussianity and high mean spectral power do not significantly affect the parameter measurements or uncertainties.

Using our new, precise measurement of \dot{P}_b , we can calculate an improved distance to PSR J0437–4715 of $D = 156.79 \pm 0.25$ pc. We discuss this measurement in detail in Section 2.6.2. As shown in Table 2.3, this measurement is consistent with independent distances, including our parallax distance measurement of $D = 157 \pm 2$ pc. V08 measured a high pulsar mass for this pulsar, of $M_p = 1.76 \pm 0.2 M_\odot$. Our improved measurement of the Shapiro delay reduces the uncertainty on the pulsar mass and we find $M_p = 1.44 \pm 0.07 M_\odot$, significantly smaller at the $1.5\text{-}\sigma$ level. This improved mass measurement will be important for the NICER mission which will attempt to measure the neutron-star radius, probing the neutron star equation of state.

Kinematic contributions to the measured $\dot{\omega}$ are included in the timing model through the measurement of Kopeikin terms. We therefore expect that our $\dot{\omega}$ measurement is solely due to the effects of GR, as was reported in V08. Under this assumption, the reported $\dot{\omega}$ corresponds to a combined pulsar and companion mass of $M_p + M_c = 1.44 \pm 0.2 M_\odot$, which is just consistent with the measured masses at the $1\text{-}\sigma$ level. In analysis of future data sets for this pulsar, the improving $\dot{\omega}$ measurement will be able to be used in combination with the Shapiro delay measurement to further constrain the pulsar mass.

2.5.2 PSR J0613–0200

The only new parameters included in our PSR J0613–0200 timing model are those describing the sinusoidal annual variations in DM that were presented in Section 2.4.1.

2.5.3 PSR J0711–6830

The pulsar timing model for this pulsar contains the same parameters as in V09. As expected, our measurements are more precise than previous work. We have not been able to obtain a parallax measurement for this pulsar because of its proximity to the ecliptic pole.

2.5.4 PSR J1022+1001

This pulsar has the smallest ecliptic latitude (-0.06 deg) of all pulsars in the sample. Observations when the line-of-sight to the pulsar passes close to the Sun have been used by You et al. (2007) to study the solar corona. However, these observations are generally removed for high-precision

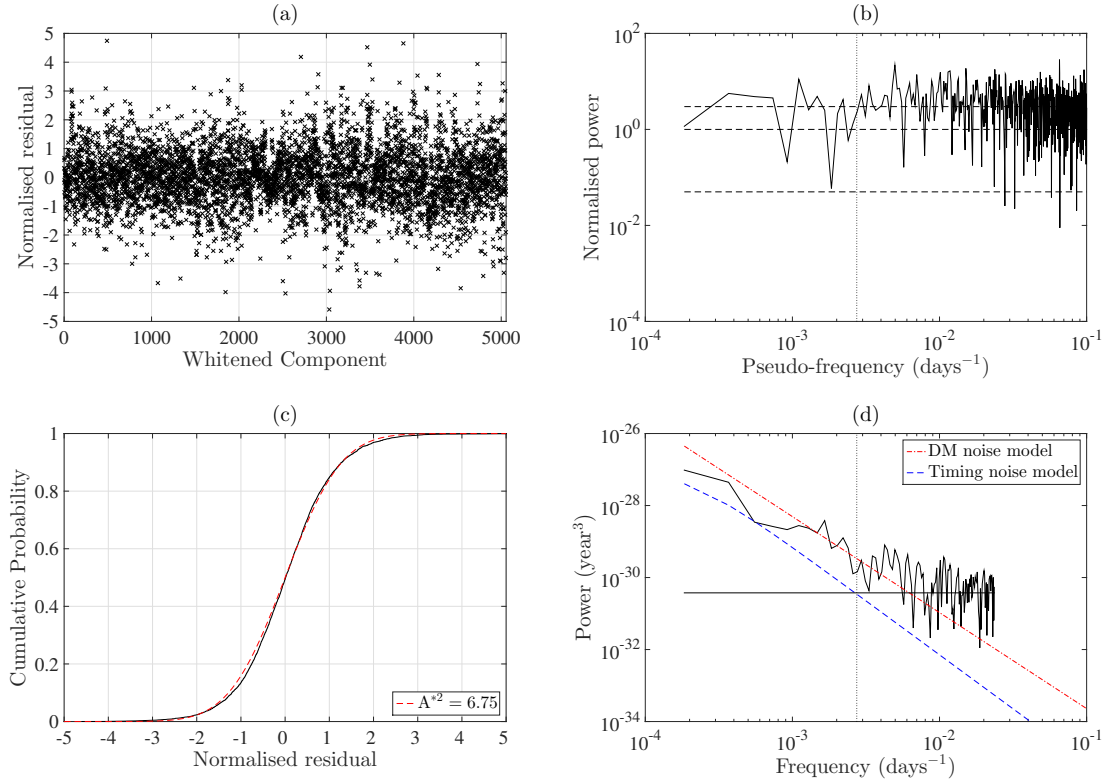


Figure 2.3: (a) Whitened and normalised post-fit residuals for PSR J0437–4715. (b) Power spectra of whitened and normalised post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and dashed lines indicate expected mean and variance for the spectrum. The pseudo-frequency is determined by converting the whitened components to a pseudo-time-series using the ToAs of the unwhitened residuals. (c) Cumulative distribution of whitened and normalised post-fit residuals (solid line) with expected distribution based on normal distribution with zero mean and unit variance (dashed line). Modified AD statistic for this distribution with the expected distribution is labelled. (d) Power spectra of post-fit residuals. Dotted line marks $f = 1 \text{ yr}^{-1}$ and solid flat line is an estimate of the white-noise level. The timing noise model applies over entire data set in combined noise model, while the Kolmogorov DM noise model only applies to residuals prior to MJD 53430.

timing applications. For our timing solutions, we removed the ToAs that were obtained when the line-of-sight to the pulsar passed within 5° of the Sun. The removal of these residuals produces normalised residuals that are consistent with zero mean and unit variance.

van Straten (2013) analysed 7.2 years of data for this pulsar from the Parkes radio telescope using a new method of polarimetric calibration to improve timing precision. The DR1 and DR1E datasets did not include this new calibration procedure. With our dataset we have detected a significant Shapiro delay using the traditional parameters M_c and $\sin i$. However while we found $\sin i = 0.69 \pm 0.18$, we did not measure a significant companion mass, $M_c = 2.2 \pm 2.4 M_\odot$. Our parameters are consistent with those of van Straten (2013).

We measured the orbital period-derivative, $\dot{P}_b = (5.5 \pm 2.3) \times 10^{-13}$, for the first time (this

parameter was not measured by van Straten 2013). Assuming a pulsar mass of $M_p = 1.4 M_\odot$ and companion mass of $M_c = 0.2 M_\odot$, the expected \dot{P}_b contribution from quadrupolar GW emission is $\dot{P}_b^{\text{GR}} = -1.7 \times 10^{-16}$, three orders of magnitude smaller than this measurement. We therefore expect that our measurement is an apparent orbital period increase caused by the Shklovskii effect. Given the proper motion of the pulsar, we can derive the distance to the pulsar, $D = 1.5 \pm 0.5$ kpc, which, as shown in Table 3, is consistent with other distance measurements.

The measured \dot{x} is expected to be from the Kopeikin kinematic effects discussed in Section 2.3. From this, we place an upper limit on the inclination angle of the orbit of $i \leq 84$ deg using the total proper motion presented in van Straten (2013), since we do not measure the proper motion in declination. The $\dot{\omega}$ measurement can have contributions from these same kinematic effects, but may also be consistent with the periastron advance expected from GR. Assuming that the observed $\dot{\omega}$ comes entirely from GR, we derive the combined mass of the system to be $M_p + M_c = 2.5 \pm 1.3 M_\odot$, which is consistent with a neutron star – white dwarf binary system.

2.5.5 PSR J1024–0719

The red timing noise for PSR J1024–0719 has a large spectral exponent (i.e., it is very steep)³. Our usual procedure, as described earlier, requires that we extrapolate the red-noise model obtained from the recent, multi-wavelength data into the earlier data. For PSR J1024–0719, we found that the red-noise model obtained from the recent data did not extrapolate well. Since DM noise is not detectable in the detrended data, we simply modelled the red noise in the entire dataset and did not use an additional DM noise model.

Our measurements are consistent with those of V09 and the measurement precision is improved in all cases. We measure a parallax of $\pi = 0.5 \pm 0.3$ mas (prior to Lutz-Kelker bias correction), which was undetected by V09. A parallax was previously measured by Hotan et al. (2006) to be $\pi = 1.9 \pm 0.8$ mas, which is also consistent with our measurements. However, since a parallax was undetected by V09, and PSR J1024–0719 has steep red noise that was unaccounted for by Hotan et al. (2006), the uncertainty for their original parallax measurement is likely to be severely underestimated.

2.5.6 PSR J1045–4509

A parallax measurement was presented by V09 of $\pi = 3 \pm 4$ mas, but this measurement may have been affected by uncorrected red noise. We now make the first significant parallax measurement of $\pi = 2.2 \pm 1.1$ mas.

³Shortly after the publication of this work in Reardon et al. (2016), it was shown simultaneously by Kaplan et al. (2016) and Bassa et al. (2016) that this steep red noise likely originates not from timing noise, but from a long-period orbit with a low-mass companion.

2.5.7 PSR J1600–3053

We present the first significant measurement of $\dot{x} = (-4.2 \pm 0.7) \times 10^{-15}$, which is attributed to the proper motion of the system (Kopeikin 1996), and gives us an upper limit on the inclination angle of $i \leq 67$ deg. We detect the first significant parallax of $\pi = 0.48 \pm 0.11$ mas (compared with $\pi = 0.2 \pm 0.3$ mas given by V09). From the Shapiro delay companion mass measurement $M_c = 0.34 \pm 0.15 M_\odot$ and the mass function, we can provide a constraint on the pulsar’s mass: $M_p = 2.4 \pm 1.7 M_\odot$.

2.5.8 PSR J1603–7202

For this pulsar, we have the first measurement of a parallax $\pi = 1.1 \pm 0.8$ mas (prior to Lutz-Kelker bias correction). We also present the measured first derivatives of the orbital period, $\dot{P}_b = (3.1 \pm 1.5) \times 10^{-13}$ and the projected semi-major axis, $\dot{x} = (1.36 \pm 0.16) \times 10^{-14}$. Since the GR contribution (for both \dot{P}_b and \dot{x}) is negligible for this system, we use the proper motion of the pulsar and our \dot{P}_b measurement to derive the distance to the pulsar, $D = 3.9 \pm 1.8$ kpc, which is marginally consistent with other distance measurements (Table 3). However all distance measurements for this pulsar are poor. Using the \dot{x} measurement, we place an upper limit on the orbital inclination angle of $i \leq 31$ deg.

2.5.9 PSR J1643–1224

We present the first measurement of $\dot{\omega} = -0.007 \pm 0.004$ deg. An improved measurement of $\dot{x} = (-5.25 \pm 0.16) \times 10^{-14}$ allows us to derive an upper limit on the inclination angle of $i \leq 28$ deg. For the measured $\dot{\omega}$ to be the result of GR effects, rather than the assumed kinematic effects, the combined mass of the system would need to be $M_p + M_c = 54 M_\odot$; an order of magnitude larger than expected. We therefore expect that this measurement is not contaminated by GR contributions and instead results from kinematic effects. However since these measurements are not well determined and we do not detect a Shapiro delay, we are unable to find a unique solution for the Kopeikin terms i and Ω and we therefore do not re-parametrise the orbit.

2.5.10 PSR J1713+0747

Splaver et al. (2005) reported on 12 years of timing observations of this pulsar from the Arecibo observatory. Their analysis was carried out with the JPL DE405 Solar System ephemeris and the time reference was TT(BIPM03). They obtained $\pi = 0.89 \pm 0.08$ mas, $\mu_\alpha \cos \delta = 4.917 \pm 0.004$ mas yr⁻¹, and $\mu_\delta = -3.933 \pm 0.01$ mas yr⁻¹. The orbital projection effects caused by this proper motion allowed them to determine $\Omega = (87 \pm 6)^\circ$. Their analysis was carried out by “whitening” the residuals using eight time derivatives of the pulse frequency. C11 showed that such whitening can lead to underestimated parameter uncertainties.

V09 obtained a less precise parallax determination of $\pi = 0.94 \pm 0.10$ mas and a proper motion components of $\mu_\alpha \cos \delta = 4.924 \pm 0.10$ mas yr⁻¹ and $\mu_\delta = -3.85 \pm 0.02$ mas yr⁻¹. They obtained $i = (78.6 \pm 1.7)^\circ$ and $\Omega = (67 \pm 17)^\circ$. V09 also included $\dot{P}_b = (41 \pm 20) \times 10^{-13}$.

The most recent VLBI observations of this pulsar (Chatterjee et al. 2009) give proper motion components of $\mu_\alpha = 4.75^{+0.16}_{-0.07}$ mas yr⁻¹ and $\mu_\delta = -3.67^{+0.06}_{-0.15}$ mas yr⁻¹ and parallax of $\pi = 0.95^{+0.06}_{-0.05}$ mas. These values are in fair agreement with our values in Table 7.

Zhu et al. (2015) have analysed 21 years of timing data from the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) for this pulsar to conduct tests of theories of gravity. In their analysis they measured parameters using a number of different noise models. Using TEMPO2 with a jitter-based white-noise model and a red-noise model, they measured $\pi = 0.87 \pm 0.03$ mas, $\mu_\alpha \cos \delta = 4.915 \pm 0.003$ mas yr⁻¹, and $\mu_\delta = -3.914 \pm 0.005$ mas yr⁻¹. For the binary model, they measured the Kopeikin terms $i = (71.9 \pm 0.7)^\circ$, and $\Omega = (88 \pm 2)^\circ$, as well as a companion mass of $M_c = 0.286 \pm 0.012 M_\odot$ and $\dot{P}_b = (0.36 \pm 0.17) \times 10^{-12}$. These parameters are consistent with, and more precise than our measurements below because of the longer data span and higher timing precision of the NANOGrav dataset for this pulsar.

With our analysis we obtain $\dot{P}_b = (1.7 \pm 0.7) \times 10^{-12}$. The intrinsic \dot{P}_b^{GR} from GW emission is negligible and so we expect that this result comes from the Shklovskii effect. This provides a pulsar distance of $D = 3.1 \pm 1.2$ kpc, which is marginally consistent with other distance measurements (Table 3). We also use the Shapiro delay companion mass measurement $M_c = 0.32 \pm 0.05 M_\odot$ and the mass function, to calculate the pulsar mass, $M_p = 1.7 \pm 0.4 M_\odot$.

2.5.11 PSR J1730–2304

We present the first measurement of a parallax for this pulsar of $\pi = 1.5 \pm 0.3$ mas, prior to Lutz-Kelker bias correction. All other parameters are consistent with the previous values from V09.

2.5.12 PSR J1732–5049

V09 was only able to determine $\mu_\delta = -9.3 \pm 0.7$ mas yr⁻¹. We now also present a measurement of the proper motion in right ascension, $\mu_\alpha \cos \delta = -0.41 \pm 0.09$ mas yr⁻¹, but parallax was not detected.

2.5.13 PSR J1744–1134

All parameters are consistent with V09 after accounting for apparent changes resulting from the different solar-system ephemeris and time standard used in the analysis.

2.5.14 PSR J1824–2452A

PSR J1824–2452A is a solitary pulsar located in the globular cluster M28. The timing residuals for this pulsar exhibit red noise, which may be caused by acceleration within the cluster potential, or timing noise. Our timing model includes dDM/dt but, as we do not include any single-frequency data for this pulsar, we use the original Cholesky routines with a single noise model. We did not measure a significant proper motion in declination nor a parallax even though V09 did publish a proper motion in declination.

2.5.15 PSR J1857+0943

This pulsar is in an orbit that is highly inclined to our line of sight, allowing for a precise measurement of the Shapiro delay. We do not improve parameter uncertainties for every parameter since V09 made use of publicly available data from the Arecibo observatory to extend the dataset, while we chose to use only the PPTA DR1E data set. Using our Shapiro delay companion mass measurement $M_c = 0.25 \pm 0.03 M_\odot$ and the mass function, we calculate the pulsar mass to be $M_p = 1.5 \pm 0.2 M_\odot$.

2.5.16 PSR J1909–3744

The narrow pulse width, particularly at 10 cm, allows us to achieve very low ToA uncertainties. Recent PPTA data for this pulsar, timed to sub-100 ns precision over more than ten years has led to the most stringent limit on the stochastic GW background to date (Shannon et al. 2015). For this reason, this pulsar is an important tool for testing models of galaxy and supermassive black hole formation.

For this pulsar, we use a corrected version of the DR1E dataset that is described in Shannon et al. (2013b). Previously undetected instrumental offsets were found and corrected, and additional archival 50cm observations were included to allow measurement of ΔDM over the entire dataset. We also include an additional jump corresponding to a software upgrade at MJD 55319.8 that was identified by Shannon et al. (2015). We include dDM/dt in the timing model, which removes the majority of the DM noise. While there is no evidence for red noise in the 10 cm residuals (Shannon et al. 2015), we identify some slight red noise originating from the 20 cm and 50 cm residuals. This could be the result of instrumental noise or residual interstellar dispersion noise. Because of this noise, we included a red-noise model that sufficiently whitens the residuals.

V09 did not include Kopeikin terms in the timing model, but instead fitted for \dot{x} ; $\dot{\omega}$ was not measured. We now include the Kopeikin terms, giving the inclination angle, $i = 93.52 \pm 0.09$ deg, and the longitude of ascending node, $\Omega = 39 \pm 10$ deg. We measure the orbital period-derivative to be $\dot{P}_b = (5.03 \pm 0.06) \times 10^{-13}$. The expected contribution from quadrupolar GW emission to this measurement is $\dot{P}_b^{\text{GR}} = -2.7 \times 10^{-15}$; two orders of magnitude smaller than this measurement. This expected value was calculated from the measured companion mass $M_c = 0.2067 \pm 0.0019 M_\odot$.

and the calculated pulsar mass from the Shapiro delay and mass function of $M_p = 1.47 \pm 0.03 M_\odot$. From this \dot{P}_b measurement and the proper motion, we derive a distance of $D = 1.140 \pm 0.012$ kpc, which is consistent with the parallax distance. The distances derived from the DM and Galactic electron density models are evidently under-estimated.

2.5.17 PSR J1939+2134

PSR J1939+2134 was the first MSP discovered (Backer et al. 1982), and it is currently the second fastest spinning pulsar known. The timing residuals for this pulsar are dominated by red noise.

2.5.18 PSR J2124–3358

All measured parameters for this pulsar are consistent with V09 and the uncertainties have decreased in all cases.

2.5.19 PSR J2129–5721

We have the first measurement of the orbital period-derivative, $\dot{P}_b = (7.9 \pm 3.6) \times 10^{-13}$. Using this measurement and the proper motion, we derive a distance of $D = 3.2 \pm 1.5$ kpc, which is consistent with the DM distances (Table 3), however all distance measurements for this pulsar are poor. We do not yet detect a parallax for this pulsar because of its proximity to the ecliptic pole.

2.5.20 PSR J2145–0750

We have the first measurement of $\dot{x} = (8.0 \pm 0.8) \times 10^{-15}$, resulting from the proper motion of the pulsar. Using this value, we place an upper limit on the inclination angle of the orbit of $i \leq 69$ degrees. All parameters are consistent with the previous results. V09 published \dot{x} , \dot{P}_b , and $\dot{\omega}$ values, but the measurements were not significant. Furthermore, the uncertainties were likely to be underestimated because of the red noise present in these observations.

2.6 Discussion

2.6.1 Advantages of using the split-Cholesky algorithm

The methodology that we have used in this Chapter is based on traditional, frequentist analysis of pulsar timing residuals. An alternate approach is through Bayesian algorithms, such as those described in van Haasteren et al. (2009), van Haasteren and Levin (2013), or the TEMPONEST algorithm developed by Lentati et al. (2014a). These algorithms have successfully been used by the NANOGrav and European Pulsar Timing Array (EPTA) groups (e.g. van Haasteren et al. 2011; Arzoumanian et al. 2014a). When these same algorithms are applied to PPTA datasets, uncertainties arise in the noise models because of the covariance between DM variations and

timing noise processes in the single-wavelength (20 cm) early data. At present, there is no way for current implementations of the Bayesian algorithms to model such non-stationary red noise in the way that the split-Cholesky algorithm allows ⁴. Instead the Bayesian algorithms assume that the noise is wide-sense stationary. Constructing separate red-noise models for the timing noise and DM variations allows us to better understand our noise model, by avoiding large uncertainties in our early data. We therefore chose to use a frequentist approach to analyse our dataset since it is less computationally expensive than the Bayesian alternatives and gives us greater control of our noise models.

To demonstrate the necessity for the split-Cholesky algorithm with our dataset, we created 500 realisations of PSR J0437–4715 data with red noise and DM noise at the level presented in Section 2.4. The parameters in the timing model were fitted (including jumps and DM variations) with three different noise treatments; no red-noise model, extrapolated DR1 red-noise model in the Cholesky algorithm only, and a two-component red-noise model with Kolmogorov DM noise in the split-Cholesky algorithm. In Figure 2.4 we show the distribution of post-fit parallax values represented by the number of standard deviations from the true value. When no noise model is used (panel a), the parameter uncertainties are clearly underestimated. When the single-component red-noise model is used (panel b), there is a significant improvement but the parameter uncertainties remain underestimated. Finally, when a two-component red-noise model is used with the split-Cholesky algorithm (panel c), we can accurately model the total red noise for the pulsar, and as a result we avoid underestimation of uncertainties. This is true for all parameters with the exception of ν and $\dot{\nu}$ (as was the case in the original Cholesky algorithm; see C11). The distributions of post-fit values for each parameter in the PSR J0437–4715 timing model (excluding ν and $\dot{\nu}$), using the split-Cholesky method, are given in Figure 2.5 with the Anderson-Darling statistic used to test the distribution. We see that the distributions are consistent with the expected zero mean, unit variance distribution for all parameters except for declination and proper motion, which have slightly overestimated uncertainties.

2.6.2 PSR J0437–4715 Kinematic distance measurement from \dot{P}_b

As described in Section 2.3, the measurement of \dot{P}_b for PSR J0437–4715 can be used to measure the pulsar’s distance, using Equation 2. Contributions to the observed \dot{P}_b value can come from changes intrinsic to the pulsar system, \dot{P}_b^{int} , from the kinematic or Shklovskii effect, \dot{P}_b^{kin} , and from differential acceleration of the Solar System and pulsar system caused by mass in the Galaxy, \dot{P}_b^{Gal} . Therefore,

$$(2.8) \quad \dot{P}_b^{\text{obs}} = \dot{P}_b^{\text{int}} + \dot{P}_b^{\text{Gal}} + \dot{P}_b^{\text{kin}}$$

where we have measured $\dot{P}_b^{\text{obs}} = (3.7276 \pm 0.0058) \times 10^{-12}$.

⁴They have since been updated to allow non-stationary noise models.




Figure 2.4: Distribution of parallax values from timing model fits to 500 realisations of simulated PSR J0437–4715 data. σ is the number of standard deviations from the true value for each of the N realisations. We used three different noise treatments: (a) No red-noise model, (b) red-noise model from DR1 data, extrapolated to apply over the entire dataset was used with the Cholesky algorithm, and (c) two-component red-noise model with Kolmogorov DM model is used in split-Cholesky algorithm. For each panel, the black line is a normal distribution fit to the distribution and the red, dashed line is a normal distribution with zero mean and standard deviation equal to the average of the standard TEMPO2 uncertainties for the 500 realisations, scaled to the same area as the black-line distribution.

The intrinsic orbital decay for neutron star-white dwarf systems such as PSR J0437–4715 is dominated by quadrupolar GW emission, which can be calculated using the relation:

$$(2.9) \quad \dot{P}_b^{\text{int}} \simeq \dot{P}_b^{\text{GR}} = - \frac{192\pi G^{5/3}}{5c^5} (P_b/2\pi)^{-5/3} (1-e^2)^{-7/2} \times \left(1 + \frac{73}{24}e^2 + \frac{37}{96}e^4 \right) m_p m_c (m_p + m_c)^{-1/3}$$

(Peters and Mathews 1963; Taylor and Weisberg 1982). For PSR J0437–4715 this GR contribution is $\dot{P}_b^{\text{GR}} = -3.2 \times 10^{-16}$, which is smaller than the value calculated by V08 because of their measurement of a high pulsar mass. This GR contribution is an order of magnitude smaller than the uncertainty in the observed value.

The Galactic acceleration component can be estimated by combining the accelerations from differential rotation and the Galactic potential. The differential rotation acceleration is found

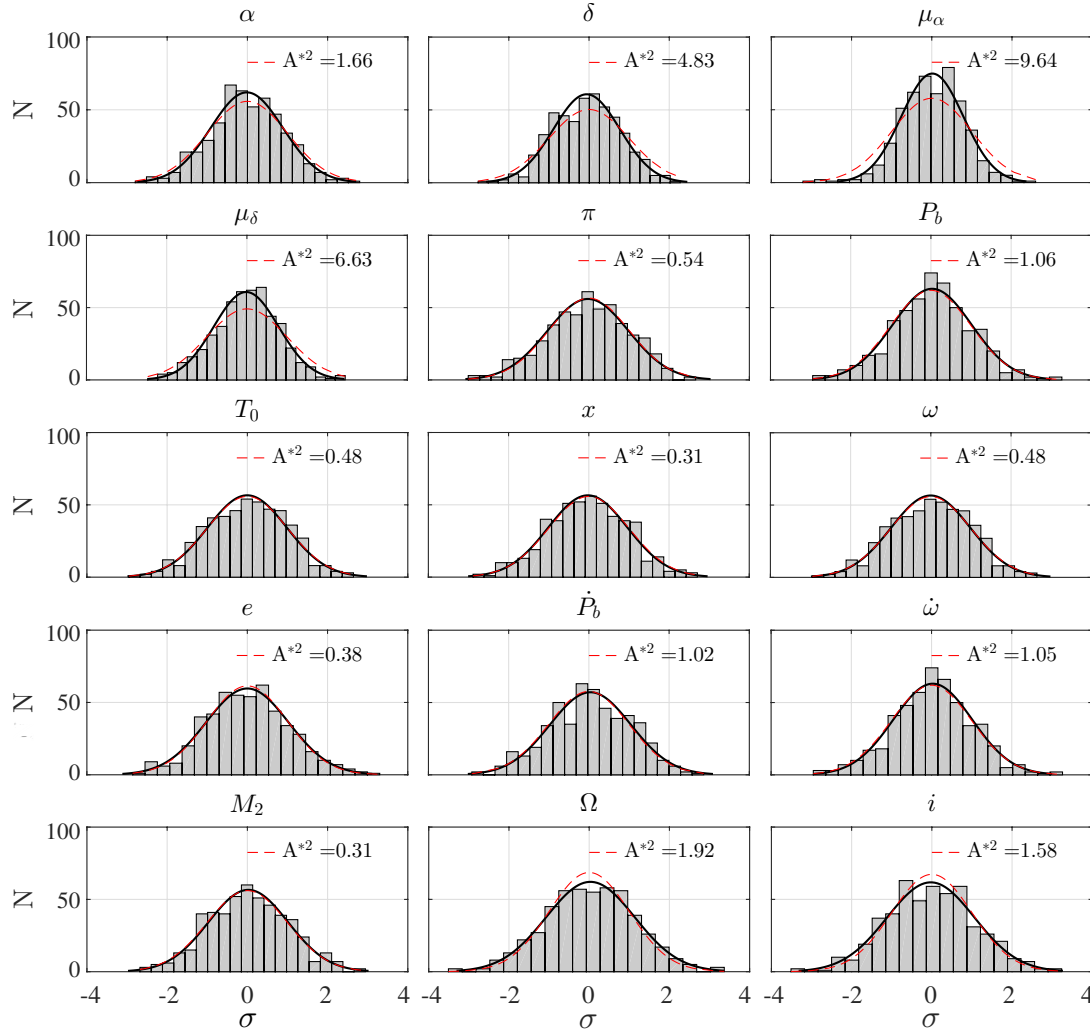


Figure 2.5: Distributions of parameter values from timing model fits with the split-Cholesky algorithm to 500 realisations of simulated PSR J0437–4715 data. σ is the number of standard deviations from the true value for each realisation. Data were simulated with statistically identical red noise properties to those of PSR J0437–4715.

using the galactic longitude and distance to the pulsar, and the galactocentric distance and circular velocity of the Sun. Acceleration in the Galactic potential varies as a function of height and can be computed from a model of the local surface density of the Galaxy given the Galactic latitude of the pulsar and its parallax distance. V08 used the Holmberg and Flynn (2004) and calculate the total Galactic contribution to be $\dot{P}_b^{\text{Gal}} = (-1.8 - 0.5) \times 10^{-14} = -2.3 \times 10^{-14}$. Since the value of the Galactic acceleration is larger than the uncertainty of our observed value, the uncertainty in this component may become important. Bovy et al. (2012) find that the uncertainty for the circular velocity of local sources is approximately 3% by using data from the Apache Point Observatory Galactic Evolution Experiment. From Holmberg and Flynn (2004), the uncertainty in the surface density resulting in the vertical component of acceleration is approximately 10%. Using these uncertainties, we have $\dot{P}_b^{\text{Gal}} = (-2.3 \pm 0.08) \times 10^{-14}$, which is small compared to the uncertainty in our measurement.

The kinematic contribution to \dot{P}_b^{obs} because of the Shklovskii effect and the distance to the pulsar can be found through

$$(2.10) \quad \dot{P}_b^{\text{obs}} - \dot{P}_b^{\text{GR}} - \dot{P}_b^{\text{Gal}} \simeq \dot{P}_b^{\text{kin}} = \frac{\mu^2 D}{c} P_b$$

where μ is the total proper motion of the pulsar and D is the distance. This gives $\dot{P}_b^{\text{kin}} = (3.7513 \pm 0.006) \times 10^{-12}$ and $D = 156.79 \pm 0.25$ pc, which is consistent with our independently measured parallax distance of $D = 156.9 \pm 2.2$ pc. This is the most precise distance measurement for all pulsars and one of the most fractionally precise distance measurements for any star. Our distance measurement is also consistent with the VLBI parallax distance of $D = 156.3 \pm 1.3$ pc, measured by Deller et al. (2008). Since \dot{P}_b produces an effect in the residuals that grows over time, we can expect the uncertainty in \dot{P}_b^{obs} to decrease significantly in future data sets, leaving the distance uncertainty to be dominated by the uncertainty in models used to calculate the contributions from Galactic acceleration. Precise distance measurements such as this will be important to PTAs in the future since it allows the use of the pulsar term in single-source GW detection, which is essential for accurately determining the source location (Lee et al. 2011).

2.7 Conclusion

We have presented new models for red noise caused by DM variations in single-frequency data sets, obtained by first including deterministic components in the timing model, and then modelling the covariance function of the remaining noise. For 14 of the pulsars in our sample, we detect a linear trend in the DM variations and include this in the timing model, and for four pulsars (PSRs J0613–0200, J1045–4509, J1643–1224, and J1939+2134) we also include new parameters that describe annual DM variations. The new DM noise models apply only to the early data where excess noise is present, and are used in conjunction with a red-noise model for the frequency-independent noise that is present in the entire dataset. These two-component models

were used in the new “split-Cholesky” algorithm to whiten the residuals to provide unbiased parameter measurements. We have described this algorithm and demonstrated its effectiveness on simulated data. Model parameters were shown to be unbiased and have accurate uncertainties through simulations based on PSR J0437–4715.

Determining new timing models for the 20 PPTA pulsars in the DR1E dataset required these new models and algorithm because of the non-stationary red noise for most pulsars. The models we present provide the best description of the noise currently possible with the PPTA data sets, and result in the most accurate and precise parameter measurements to date for most pulsars in our sample, as well as the detection of several new parameters. Most notably, we presented the first significant parallax measurements for PSRs J1024–0719, J1045–4509, J1600–3053, J1603–7202, and J1730–2304, and determined the distance to PSR J0437–4715 at the 0.16% level of precision. We also measured an improved pulsar mass for PSR J0437–4715, which at $M_p = 1.44 \pm 0.07 M_\odot$, is somewhat lower than the previous measurement.

The analysis described here can easily be applied to future PPTA data releases and to any PTA data with non-stationary noise processes. Longer datasets for PSR J0437–4715 will further improve the distance measurements based on both the parallax, and the orbital period-derivative from Schklovskii acceleration. If the parallax distance from pulsar timing becomes more precise than the current VLBI distance, these two independent measurements can be used for example to further improve current constraints on the change to Newton’s Gravitational constant (Freire et al. 2012). It may also be possible to measure, or place limits on, the acceleration of the pulsar system caused by mass in the Galaxy. With longer datasets, improving $\dot{\omega}$ measurements, particularly for PSR J0437–4715, can compliment the Shapiro delay to further improve measurements of the pulsar mass.

MODELLING RELATIVISTIC AND ANNUAL VARIATIONS IN THE SCINTILLATION OF PSR J1141–6545 FOR PRECISE ASTROMETRY

We have observed the relativistic binary pulsar PSR J1141–6545 over a period of ~ 6 years using the Parkes 64 m radio telescope, with a focus on precision measurements of the pulse times of arrival (TOAs), and diffractive intensity scintillations. Fitting models to these scintillations allows us to measure astrometric and orbital parameters that are difficult to measure with TOAs alone because TOA measurements are sensitive to the changing radial velocity of the pulsar while scintillations are sensitive to its transverse velocity. We have modelled the long-term scintillation, which shows relativistic and annual variations, to measure parameters that include: the inclination of the binary orbit i ; the longitude of the ascending node Ω ; the relativistic advance of periastron $\dot{\omega}$; and the pulsar system transverse velocity. We use the annual variations to resolve the previous ambiguity in the sense of the inclination angle. Using the correct sense, but the more precise value inferred from tests of general relativity ($i = 73 \pm 1.5^\circ$), we find $\Omega = 23 \pm 3^\circ$ and we estimate the pulsar distance to be $D = 9_{-3}^{+5}$ kpc, which then gives us an estimate of this pulsar's proper motion of $\mu_\alpha = 2.5 \pm 1.2 \text{ mas yr}^{-1}$ in right ascension and $\mu_\delta = 1.5 \pm 0.7 \text{ mas yr}^{-1}$ in declination. Finally, we obtain measurements of the spatial structure of the interstellar electron density fluctuations, including: the spatial scale and anisotropy of the diffraction pattern; the distribution of scattering material along the line of sight; and spatial variation in the strength of turbulence from epoch to epoch. We find that the scattering is dominated by a thin screen that is slightly anisotropic.

3.1 Introduction

Interstellar intensity scintillations are seen in all radio observations of highly compact sources at centimetre to metre wavelengths. They are caused by transverse fluctuations in the electron density of the turbulent ionised interstellar medium (IISM). The primary mechanism is diffractive angular scattering, where intensity variations arise from interference between the scattered waves (Rickett 1969). The diffractive scintillations are modulated by refractive scintillations on larger spatial scales (Rickett et al. 1984; Romani et al. 1986). The observed scintillations, which appear as variations in the source flux with time and observing frequency, are caused by a spatial interference pattern drifting across the line-of-sight. Thus the time scale of the observed scintillations τ_d (typically of order minutes) is inversely proportional to this drift velocity V_{los} , which is a linear combination of the transverse velocities of the observer, source, and IISM (Section 3.3.1). The diffraction pattern is also frequency-dependent and becomes decorrelated over a bandwidth $\Delta\nu_d$. The angular scattering broadens each pulse into a quasi-exponential pulse with timescale τ_s , which is related to the bandwidth by $2\pi\Delta\nu_d\tau_s \approx 1$ (Rickett 1977). Diffractive scintillations have a narrow bandwidth (typically of order MHz), which is a measure of the strength of scattering and can be used to estimate the diffractive spatial scale s_d (Section 3.2.3). Detailed overviews of pulsar scintillations are given by Rickett (1990) and Narayan (1992).

Measurements of a dynamic spectrum of intensity scintillation with time and frequency can therefore provide information on the spatial structure of the IISM, the transverse velocity of the pulsar, and the strength of scattering. For solitary pulsars V_{los} is often dominated by the pulsar proper motion and is relatively simple to model, depending only on the distance from the scattering region to the Earth (Lyne and Smith 1982). If the pulsar has a binary companion, orbital dynamics can also be studied from the transverse orbital motion. This was first used by Lyne (1984) to measure the orbital inclination angle of PSR B0655+64 for the first time. Then Ord et al. (2002a) analysed two consecutive orbits of PSR J1141–6545, a relativistic binary pulsar with a white dwarf companion in a ~ 4.7 hour eccentric orbit (discovered by Kaspi et al. 2000), with the assumption of isotropic scattering. The technique was extended to deal with anisotropic scattering of the double pulsar PSR J0737–3039A by Coles et al. (2005) and further extended to analyze multiple observations of PSR J0737–3039A including the effects of the variation in the Earth’s velocity by Rickett et al. (2014). We have been provided access to the latter data and used it to calibrate our analysis against that of Rickett et al. (2014).

For this work we use the dynamic spectrum as our basic observable and model the diffractive scintillations over a period of ~ 6 years for PSR J1141–6545. An example dynamic spectrum from some of our observations is shown in Figure 3.1. This spectrum shows the modulation of the scintillation timescale τ_d , caused by the orbital motion of the pulsar (see also Figure 1 of Ord et al. 2002a). We characterise the statistics of this spectrum by a two dimensional autocovariance function (ACF). Following convention (Cordes and Rickett 1998), the half-width at half maximum of this ACF in frequency is the decorrelation bandwidth $\Delta\nu_d$ and the half-width at $1/e$ in time




Figure 3.1: A dynamic spectrum of four PSR J1141–6545 observations made on MJD 56391 with ~ 300 MHz of used bandwidth about a centre frequency of 1366 MHz. The four observations cover a combined ~ 2.2 pulsar orbits. Stretching of the scintles (in black) in time is evident (at e.g. 100 and 400 minutes) when the pulsar orbital velocity reaches a minimum. The vertical white bars are periods between the individual observations, while horizontal white bars and patches were removed because of radio-frequency interference. The greyscale shows the normalised flux with the colour limits chosen to optimise the visualisation of scintles.

is τ_d . The scale of spatial variations in the diffraction pattern s_d is related to the scintillation timescale by $s_d = V_{\text{los}} \tau_d$. The spatial scale is determined by the strength of scattering, which can vary with time but is independent of V_{los} . If the spatial structure is anisotropic, τ_d will depend on the angle between the semi-major axis of the anisotropy and the velocity vector. Fortunately, since Δv_d is a direct measure of the strength of scattering, this can be used to correct for temporal variations in the spatial scale caused by changes in the strength of scattering, provided that the anisotropy does not also change with time. A formalism for using Δv_d to correct for changes in s_d has been provided by Cordes and Rickett (1998). The near-constant Δv_d for the dynamic spectrum in Figure 3.1 indicates that the strength of scattering does not change over the orbital period, because the projected size of the orbit is smaller than the scattering disk. However we do observe small strength of scattering changes from epoch to epoch (discussed in Section 3.2.3) and we use the scheme of Cordes and Rickett (1998) to account for this.

PSR J1141–6545 has already proved to be a valuable and unique laboratory for testing general relativity (Bhat et al. 2008; Manchester et al. 2010). The parameters measured from scintillation are important for these tests because poor timing precision limits the measurement of the Shapiro delay, which would otherwise provide a measurement of the inclination angle and companion mass to further constrain the tests of general relativity (Bhat et al. 2008). In

addition to this, the distance to the pulsar is poorly constrained at present, which means that the contamination from kinematic effects (in this case the Shklovskii effect; Shklovskii 1970) in the measured relativistic orbital parameters is unknown. A lower-bound distance estimate of 3.7 kpc was found by Ord et al. (2002b) using HI absorption spectra and was used to calculate the “scintillation velocity” V_{ISS} (observationally-derived V_{los}) in Ord et al. (2002a). Verbiest et al. (2012) then showed that distances derived in this way may be overestimated because of a luminosity bias, and presented a revised distance estimate of 3 ± 2 kpc (which is not significantly different, but we take this as an initial value for our models). At this large distance, the proper motion is too small to be measured with current timing precision. For this reason, both the measured transverse velocity (to estimate the Shklovskii effect) and inclination angle (to constrain the pulsar and companion masses) from scintillation modelling are useful for improving the tests of general relativity with this system. The system runaway velocity is also an important tool for investigating the formation of such white dwarf-neutron star systems (e.g. Tauris and Sennels 2000; Davies et al. 2002; Church et al. 2006).

In this paper we present new short-term and long-term scintillation models for two different scattering geometries. Our best long-term model is a thin scattering screen and includes anisotropy in the IISM. We show that the Earth’s velocity is detectable from scintillation timescale modulation over a year, and that the relativistic advance of periastron can be measured from scintillation alone. We use the annual variation to measure the orientation of the pulsar’s orbit in celestial coordinates, resolve the sense of the inclination angle, determine the distance to the scattering region, and calculate the proper motion of the pulsar. We compare these results with previous measurements from scintillation and pulsar timing, provide a revised distance estimate to the pulsar, and predict the contamination in the orbital period-derivative from the Shklovskii effect (Shklovskii 1970).

3.2 Dataset

3.2.1 Observations and dynamic spectra

For this work, we use a selection of archival PSR J1141–6545 observations from the Parkes 64 m radio telescope, spanning ~ 6 years from June 2009 to June 2015. The observations were part of the P361 observing project, which was a long-term campaign to monitor PSR J1141–6545 to improve the stringent tests of gravitational theories. The data were received with the central beam of the Parkes multibeam receiver, and recorded either with a digital polyphase filterbank system with a 256 MHz bandwidth and 0.25 MHz channel width, or a coherent dedispersion machine with a 400 MHz bandwidth and 0.78 MHz channel width. Each observation was from the “20 cm band” and had a centre frequency in the range of 1358 MHz to 1423 MHz.

We selected only the longest observations, which spanned at least 142 minutes, so that each observation covered at least half of an orbit (orbital period ~ 4.7 hrs). This yielded 126 individual

observations, which we organise into 23 distinct “epochs” separated by at least 60 days.

A dynamic spectrum for each observation was produced using the data processing pipeline designed for the upcoming data release 2 (DR2) of the Parkes Pulsar Timing Array (PPTA; Manchester et al. 2013) project. In brief, observations of a pulsed noise diode that excites both X and Y polarisations in phase are performed before each of the observations described above, to allow correction of the complex gain. The noise diode is itself calibrated to absolute flux density using on- and off-source observations of the bright radio galaxy Hydra A. Polarisation calibration is done using the noise diode observations combined with regular observations of the highly polarised pulsar PSR J0437–4715 (van Straten 2004). To compute dynamic spectra, we perform a least-squares fit of an analytic model of the pulse profile in total intensity (Stokes I) to the observed pulse profile for each sub-integration and frequency channel. This fit provides the amplitude and its uncertainty. Because the observed pulsar profiles are already absolutely calibrated, the amplitude measurement yields the pulsar flux density directly. In addition, using this analytic pulse profile simultaneously optimises the signal-to-noise ratio for both the pulse amplitude and time of arrival. The calibrations and measurements described above are performed with the PSRCHIVE (Hotan et al. 2004) package.

Figure 3.1 shows the dynamic spectra of four consecutive observations of PSR J1141–6545 on MJD 56391. We cut each observation into segments < 12 minutes in length (on average they are ~ 11 minutes) and measure the scintillation timescales and bandwidths for each segment (as described in the following section). In this way we can measure the modulation of τ_d across orbital phase for each observation. The decorrelation bandwidth $\Delta\nu_d$ measurements and the grouping of observations into the 23 epochs are shown in Figure 3.2. The measurements in this Figure are from all available archival observations (with $t_{\text{obs}} > 24$ mins) during our selected observing span, not just the 126 observations from P361 that we use for modelling. This is because many of the P361 observations have channel bandwidths $B_c = 0.78$ MHz, meaning that we overestimate the decorrelation bandwidth because $\Delta\nu_d < B_c$. We use all available data (which may not be suitable for τ_d measurements because of a short observing span for example) to measure $\Delta\nu_d$ from observations with $\Delta\nu_d > B_c$, and estimate $\Delta\nu_d$ using the variance of the flux in the dynamic spectrum when $\Delta\nu_d < B_c$. This method is described in Section 3.2.3, where we also show that it is sufficient to simply take the weighted mean of $\Delta\nu_d$ at each epoch and hold this fixed for each of the 126 long observations that we use for τ_d measurements and our modelling.

Since the binary orbital velocity can change rapidly for this system, our measured scintillation parameters are actually an average across the ~ 11 minute window. This has a small effect for our measurements near the minimum velocity, where the measured timescale is typically underestimated near the maximum. To account for this, we apply an 11 minute smoothing window to our models before fitting. Without taking this bias into account, the peak-to-peak difference in orbital velocity would be reduced, which would primarily result in a slightly over-estimated inclination angle.




Figure 3.2: Decorrelation bandwidth ($\Delta\nu_d$) as a function of time for all archival observations in a ~ 6 year span with a total observing time of at least 20 minutes at ~ 1400 MHz for PSR J1141–6545. The data are split into 23 epochs separated by vertical, dashed lines. Each epoch contains multiple individual observations, and the dynamic spectra for these were cut into ~ 11 minute segments. Black crosses are the measured $\Delta\nu_d$ for each segment where $B_c < \Delta\nu_d$ for channel bandwidth B_c . Grey circles are a flux-based estimate of $\Delta\nu_d$ from the procedure described in Section 3.2.3 for segments where $B_c > \Delta\nu_d$. The measurements and estimates were used to calculate a weighted mean value of $\Delta\nu_d$ for each observing epoch. The anomalous region just before MJD 55500 is primarily over-estimated due to poor dynamic spectrum quality because of terrestrial radio interference in some observations at this time. This does not significantly affect the weighted mean for this epoch.

3.2.2 Measurement of τ_d and $\Delta\nu_d$

We measure the scintillation parameters from ~ 11 minute segments of the dynamic spectrum because the scintillation timescale varies rapidly for the relativistic orbit. We characterise the statistics of these segments using the estimated autocovariance function (ACF), $C(\tau, \delta\nu)$. To calculate this we pad each segment with an equal length of zeroes in both dimensions, perform a 2-D FFT on the zero-padded segment, take the squared magnitude of the result, and perform an inverse 2-D FFT. We then perform a least squares fit of analytical models to $C(\tau, 0)$ and $C(0, \delta\nu)$ to obtain τ_d and $\Delta\nu_d$ respectively. First we fit $C(\tau, 0)$ with

$$(3.1) \quad C(\tau, 0) = A \exp\left(-\left|\frac{\tau}{\tau_d}\right|^{\frac{5}{3}}\right) \Lambda(\tau, T_{\text{obs}}), \text{ for } \tau > 0$$

$$C(0, 0) = N.$$

where T_{obs} is the length of the segment (~ 11 minutes) and $\Lambda(\tau, T_{\text{obs}})$ is the triangle function of length T_{obs} . This function is a slight modification to the previous standard of a Gaussian function (Cordes and Rickett 1998), where we have introduced the exponent of $5/3$ to be a better fit to the shape of Kolmogorov scintillations. After obtaining A , N , and τ_d , we keep A and N constant

because we typically have marginal resolution in frequency and it is easier to isolate the noise spike using the time cut, $C(\tau, 0)$. We then fit $C(0, \delta\nu)$ with

$$(3.2) \quad C(0, \delta\nu) = A \exp\left(-\left|\frac{\delta\nu}{\Delta\nu_d/\ln 2}\right|\right) \Lambda(\delta\nu, B), \text{ for } \delta\nu > 0$$

$$C(0, 0) = N$$

to obtain $\Delta\nu_d$, where B is the receiver bandwidth. Here $\Delta\nu_d$ is the half width at half power for the exponential function, which is standard in the field (Cordes and Rickett 1998).

This fitting provides straightforward estimators of τ_d and $\Delta\nu_d$, which is easily checked manually, but the least squares fit is not optimal because samples of the observed ACF are heavily correlated. We therefore repeat the fit using the same analytical models, but we perform the fit in the Fourier transform domain where we simply transform the autocovariance function and the models.

The 2-D FFT of $C(\tau, \delta\nu)$ is the power spectrum (or secondary spectrum) $P(f_{\text{dop}}, t_{\text{del}})$, where the dimensions are the differential time delay t_{del} and the differential Doppler shift f_{dop} of the interfering waves. To obtain τ_d from this data we first sum $P(f_{\text{dop}}, t_{\text{del}})$ over the t_{del} dimension and divide by the number of samples N_{del} and then we fit the transform of Equation 3.1. As before we hold A and N fixed from this fit before obtaining $\Delta\nu_d$ by summing $P(f_{\text{dop}}, t_{\text{del}})$ over the f_{dop} dimension, dividing by the number of samples N_{dop} , and fitting the resulting power spectrum with the transform of Equation 3.2.

The errors on these average power spectra are independent, but not equal. In fact they are proportional to the average power spectra itself, so we use a weighted least squares fit with the models providing the weights. This approach provides a second (but not independent) estimator for both τ_d and $\Delta\nu_d$ for which we believe the measurements and uncertainty estimates are more reliable. We check that the two methods agree within the uncertainty of the second method and review the dynamic spectra if they do not. For this work we use the τ_d and $\Delta\nu_d$ measurements from this Fourier-domain method.

As discussed in the previous section, our measurements of τ_d are actually the average across an ~ 11 minute segment, and we therefore must smooth our models to account for this. However, there is also a bias introduced by assuming that the shape of $C(\tau, 0)$ is given by Equation 3.1, even in the case of a rapidly changing velocity across the segment. The shape of the covariance function will change slightly since the data are a combination of multiple velocities, each contributing a covariance with the form of $C(\tau, 0)$, but with different widths. We have pre-computed this bias as a function of orbital-phase by averaging a range of $C(\tau, 0)$ corresponding to the range of observed scintillation velocities across the segment, and then measuring the change to τ_d relative to the scale of $C(\tau, 0)$ from the mean velocity. We find that the maximum bias is $\sim 3\%$ when the velocity change across the segment is a maximum. Our data are then corrected for this bias, which may otherwise have a small effect on the measured parameters, particularly the inclination angle because this is closely related to minimum transverse orbital velocity.

3.2.3 The effects of inhomogeneity in the IISM

The primary physical mechanism underlying scintillation is angular scattering by small scale irregularities in electron density that is diffractive in nature. Observations that are sensitive to the spectral exponent of the power-law of density fluctuations have shown a Kolmogorov spectrum truncated at an inner scale (which is at or near the ion inertial scale or ion Larmor radius) (e.g. Armstrong et al. 1995; Rickett et al. 2009; Spangler and Gwinn 1990). Thus it has been assumed that the microstructure in the IISM is turbulent in origin and by default they have been assumed homogeneous. However more recent observations of phenomena such as extreme scattering events (ESEs) (e.g. Fiedler et al. 1987; Coles et al. 2015) suggest that the turbulence is often inhomogeneous or that inhomogeneous structures that dominate the scattering are often present at some place on the line of sight from the source to the observer.

Furthermore the spatial structure is now often found to be localised (along the line-of-sight) and anisotropic, for example by the presence of enhanced parabolic arcs in the two-dimensional fourier transform of the dynamic spectrum that were first discovered by Stinebring et al. (2001). This has been called the “secondary spectrum” or the “delay-Doppler distribution” because its axes are the differential Doppler shift and the differential time delay of the interfering waves that cause the intensity variations (e.g. Walker et al. 2004; Cordes et al. 2006; Briskin et al. 2010). The parabolic arcs occur when the scattering is dominated by a compact local region on the line of sight and are enhanced when the scattering is anisotropic and is aligned with the drift velocity of the diffraction pattern. We have analysed the secondary spectra for PSR J1141–6545 and found no parabolic arcs, suggesting that the scattering may be isotropic and/or extended along the line of sight. PSR J1141–6545 is much further away than the pulsars observed by Stinebring et al. (2001), so it is not surprising that the scattering is not dominated by a single localised region. It is also likely that any anisotropy in the microstructure would be reduced by the long line of sight integration for a uniform distribution. For this reason we choose to fit a broad set of scattering models. However we may also be insensitive to the arcs if the scattering is localised, but extended (as in a “thick screen”), or if the curvature of the arc is not in the sensitivity range for our secondary spectra (the curvature could be large for a distant pulsar like PSR J1141–6545, which could hide the arcs along τ_{del} axis; Cordes et al. 2006).

The coherence spatial scale of the scattering medium (or the spatial scale of the diffractive scattering), s_d is defined as the transverse separation where incident waves have a 1 radian rms difference in phase. The width of the angular scattering is then $\theta_d \approx 1/(ks_d)$, for incident radiation with wavenumber k (Rickett 1990). The radiation received by the observer arrives from a scattering region of diameter $s_r = \theta_d D_e$, where D_e is the distance to the scattering screen from the observer. The intensity will also show variations from refraction, with a spatial scale equal to the diameter of this scattering region (the refractive scale). We estimate the size of the scattering region s_r from $s_r/r_F = \sqrt{v/\Delta v_d}$ (Rickett 1990), where $r_F = \sqrt{D_e(1 - D_e/D)/k}$ is the “Fresnel scale,” and find $s_r \approx 10^8$ km (AU-scale) for scintillation bandwidth $\Delta v_d = 0.36$ MHz

(mean from Figure 3.2). Because the projected size of the orbit is 5.6×10^5 km at this distance, the scattering disk is much larger than the orbital diameter regardless of the location of the scattering screen or a significant error in the pulsar distance. So for a single binary orbit, the line of sight to PSR J1141–6545 does not travel outside of the scattering region. Therefore we do not expect the strength of scattering or the anisotropy of the IISM to change during a single orbit (although the anisotropy could show small random variations even for isotropic scattering). In this case the variation of τ_d over an orbit can be exactly described by a sum of five harmonic terms (i.e. a constant plus sine and cosine of first and second harmonics; Rickett et al. 2014). In addition Δv_d should not vary over the orbit, which provides a 6th constraint on the model. However we know that the scattering may be inhomogeneous on a spatial scale of AU, so we must expect variations in the strength of scattering and perhaps other parameters of the IISM from epoch-to-epoch. Accordingly we fit $V_{\text{ISS}}(\phi)^2$ at each epoch with a five harmonic model, where $V_{\text{ISS}}(\phi)$ is the “scintillation velocity” as a function of orbital phase, ϕ and is defined in Equation 3.3 of Section 3.3. As expected we find the harmonic model fits very well at each epoch, as shown in Figure 3.3. The relationship between the harmonic coefficients and physical parameters of the scattering are given in Equation 3.17 of Section 3.3.2.1.

We compute the weighted mean and the rms of Δv_d at each epoch, but we find that the apparent value of Δv_d depends on the channel bandwidth B_c . In a few epochs we have observations with $B_c = 0.78$ MHz and also with $B_c = 0.25$ MHz. In these cases Δv_d should be the same, but the estimator determined from the ACF saturates near B_c and typically is close to 0.3 MHz. Unfortunately many of our observations are made with $B_c > \Delta v_d$. In this case fitting the ACF does not provide a useful estimate, but we know that the intensity variance V_I will be reduced by a factor of $\Delta v_d/B_c$. When $B_c < \Delta v_d$ we know that $V_I = M_I^2$ where M_I is the mean intensity. So when $B_c > \Delta v_d$ we can use $\Delta v_d = F_c B_c V_I / M_I^2$, where F_c is a calibration factor that is needed because the method depends on the actual shape of the ACF. We use the epochs for which we have observations with two different values of B_c to determine $F_c = 0.82$. This is shown in the lower panel of Figure 3.2, where the direct Δv_d measurements are in dark grey, the scaled estimates from V_I are in light grey, and the resulting weighted mean Δv_d for each epoch is in black. This technique was also used briefly by Kerr et al. (2018), but they had insufficient data to determine the calibration factor F_c . We show here for the first time that the method works well enough to detect small changes in bandwidth and generally agrees well with the direct measurement of Δv_d from the ACF.

From a single orbit one can hope to model six parameters which must include: s_d , the axial ratio of anisotropy A_r , the angle of anisotropy ψ , the distance to screen s , and two constant components of the velocity V_x , and V_y . If, as it happens, A_r is small (see Figure 3.4 and Section 3.3.2.1), then we can estimate ω and i at each epoch. We can then use the variation of the Earth’s velocity over the year, and the known relativistic advance of the longitude of periastron (from precise pulsar timing), to provide additional constraints. However one must re-estimate s_d from

Δv_d at each epoch as the strength of scattering changes on a spatial scale of AU; the calculated size of the scattering region.

In most cases the plasma turbulence causing the scattering has an outer scale which is much greater than the refractive scale. So it is not uncommon to see a spatial gradient in density which is essentially constant over the refractive scale. This will cause the apparent position of the pulsar to shift, and this will affect Δv_d , but it does not, to first order, affect τ_d . The spatial gradient can be estimated from the ACF, as can A_r , however our observations do not have sufficient sensitivity to do more than estimate the phase gradient. The phase gradient in the direction of the velocity causes a frequency-dependent spatial shift to the diffraction pattern, which is measured as a “tilt” in the characteristic scintle in the ACF (this is a chromatic aberration effect; Cordes et al. 1986; Rickett et al. 2014). We find that the amplitude and direction of these tilts vary across orbital phase for each observation, consistent with a constant phase gradient that is sampled by a rotating velocity vector. However, we do not find a strong correlation between these tilts and the measured Δv_d for each segment. The uncertainty on the weighted mean Δv_d , which we calculate as a constant for each epoch, is larger than any bias introduced by not accounting for this phase gradient. We therefore use the weighted mean Δv_d (Figure 3.2; black diamonds), measured at each epoch, to estimate the spatial scale and V_{ISS} as described in the following section.

3.3 The models

In this section we describe how we use the procedure of Cordes and Rickett (1998) to model the dynamic spectra. Their model is designed to account for variations in the strength of scattering, which will change the spatial scale and consequently the time scale. They define a “scintillation velocity” V_{ISS} , which is the ratio of the spatial scale of the diffraction pattern at the observer s_d , to the temporal scale τ_d . This V_{ISS} can be modelled if one knows the distribution and velocity of interstellar plasma along the line-of-sight. We choose two simple models for the plasma distribution:

Uniform medium: A continuous, uniform distribution of plasma with Kolmogorov turbulence along the line-of-sight.

Thin screen: A single compact “blob” of plasma with Kolmogorov turbulence, the “scattering screen”, at some position s between the pulsar $s = 0$ and the Earth $s = 1$.

For both models we assume that the density irregularities originate from turbulence and are described by a Kolmogorov spectrum. Cordes and Rickett (1998) derive the spatial scale s_d given the observed decorrelation bandwidth Δv_d for both of these models, assuming that the scattering is isotropic. In general they find that the scintillation velocity is then given by

$$(3.3) \quad V_{ISS} = A_{ISS} \frac{\sqrt{D \Delta v_d}}{f \tau_d},$$

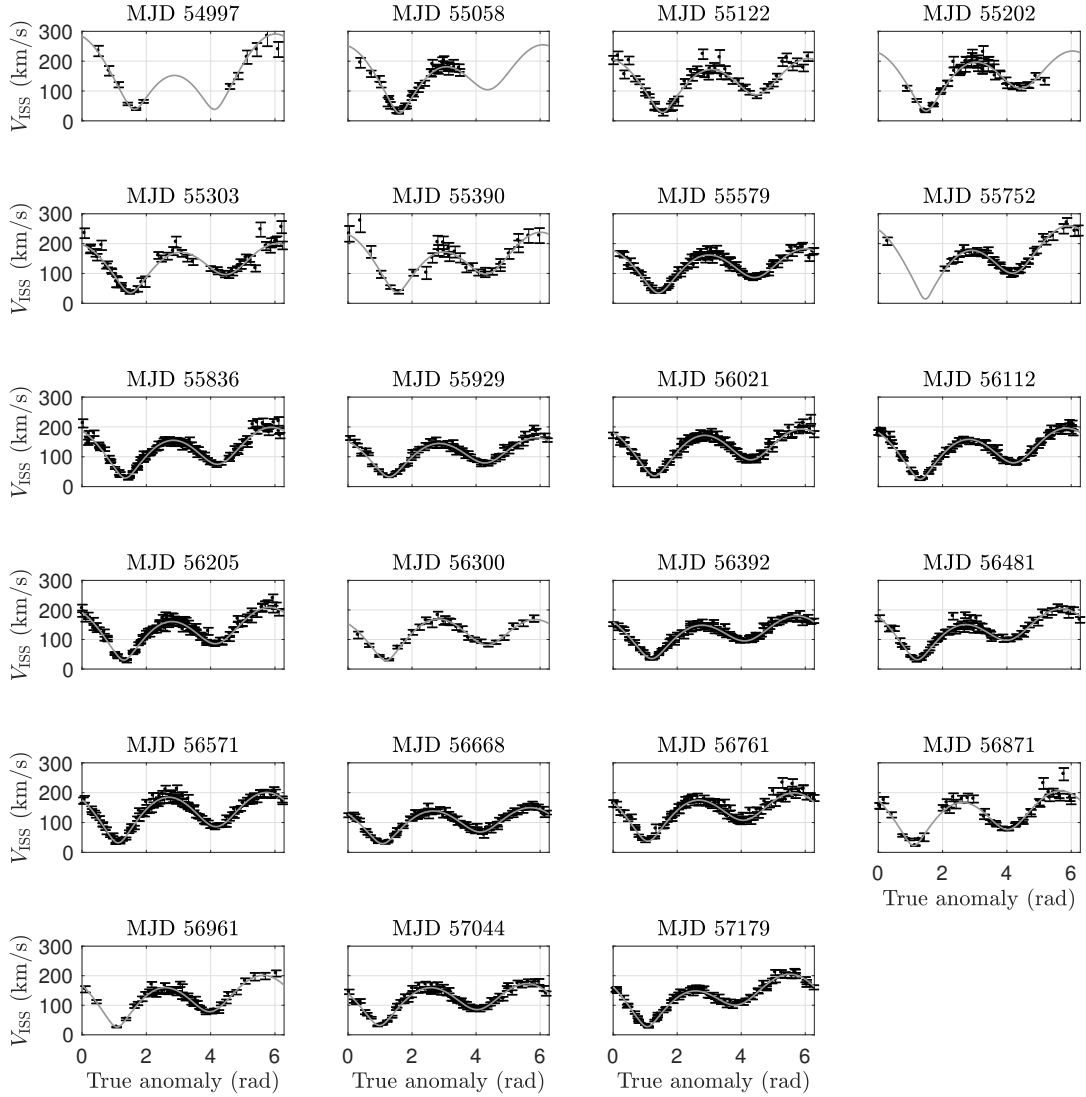


Figure 3.3: Scintillation velocity, V_{ISS} as a function of true anomaly for the 23 epochs of observations shown in Figure 3.2. The title of each panel gives the approximate starting date for the first observation in the group. V_{ISS} is defined as the scintillation velocity observed at the Earth for a uniform, Kolmogorov medium along the line-of-sight (Equation 3.3). The solid line for each panel is the best fit physical model, which is described in Section 3.3.2.2.

where D is the distance to the pulsar in kiloparsecs, f is the observing frequency in GHz, $\Delta\nu_d$ is in MHz, τ_d is in seconds, and the factor A_{ISS} depends on the assumed geometry of the scattering medium and on the exponent of the density spectrum. For the uniform medium model, $A_{\text{ISS}} = 2.53 \times 10^4 \text{ km s}^{-1}$, while for the thin screen model, $A_{\text{ISS}} = 2.78 \times 10^4 \sqrt{2(1-s)/s} \text{ km s}^{-1}$, where s is the fractional distance along the line-of-sight from the pulsar to the scattering screen (Cordes and Rickett 1998). We also include the extension to anisotropic scattering first presented by Coles et al. (2005) in both of our models, and we describe this extension in Section 3.3.1. In addition to these models that we use for the long-term scintillation, we fit the model of Ord et al. (2002a) to each of the 23 epochs of observations in Figure 3.3. We refer to this as the “physical model” for single epochs as opposed to the harmonic coefficient model of Rickett et al. (2014), both of which are described in Section 3.3.2 with results given in Section 3.5.1.1.

V_{ISS} is related to the effective transverse line-of-sight velocity $V_{\text{eff}}(s)$ through the scattering medium at position s , which is a weighted sum of the pulsar, Earth, and IISM velocities (Cordes and Rickett 1998):

$$(3.4) \quad V_{\text{eff}}(s) = (1-s)(V_p + V_\mu) + sV_E - V_{\text{IISM}}(s),$$

where V_p , V_μ , V_E , and V_{IISM} are the velocities from the pulsar’s orbit, the pulsar proper motion, the Earth, and the IISM respectively. Each is relative to the local standard of rest, and the IISM velocity can vary as a function of distance along the line-of-sight, s .

The proper motion of PSR J1141–6545 is not currently known from pulsar timing, however we expect that the transverse velocity of the pulsar system is larger than any IISM speed (Ord et al. 2002a), and accordingly we chose to set $V_{\text{IISM}} = 0$. In this case, our measurement of V_μ will include a contamination from any non-zero V_{IISM} . Therefore, when interpreting V_μ as the proper motion we implicitly make the assumption for the uniform medium model that the mean V_{IISM} along the line-of-sight is small compared with V_μ , and for the thin screen model our assumption is that the velocity of the scattering screen is small compared with V_μ .

V_{ISS} is found by integrating V_{eff} along the line-of-sight from $x = 0$ at the pulsar and $x = 1$ at the observer, with a weight that corresponds to the geometry of the scattering medium (Cordes and Rickett 1998)

$$(3.5) \quad V_{\text{ISS}} = \left[\frac{\int_0^1 \eta(x) |V_{\text{eff}}(x)|^\alpha dx}{\int_0^1 \eta(x) x^\alpha dx} \right]^{1/\alpha},$$

where $\alpha = 5/3$ for a Kolmogorov medium, and $\eta(x)$ is the mean-square scattering angle per unit distance and functions as the weight for the integral to describe different geometries. For the uniform medium model the weight is unity ($\eta(x) = 1$), and the integral simplifies to

$$(3.6) \quad V_{\text{ISS}} = \left[\frac{8}{3} \int_0^1 |V_{\text{eff}}(x)|^{5/3} dx \right]^{3/5}.$$

For the thin screen model, $\eta(x)$ is a delta function at $x = s$, and V_{ISS} reduces to the line-of-sight velocity V_{los} with respect to the diffraction pattern, defined at the location of the observer,

$$(3.7) \quad V_{\text{ISS}} = V_{\text{los}} = V_{\text{eff}}/s.$$

In the following section we describe the model for V_{eff} . This includes the definitions for the pulsar orbital velocity, the orientation of the orbit in celestial coordinates, and the extension to anisotropic scattering developed by Coles et al. (2005).

3.3.1 Effective velocity and IISM anisotropy

The scintillation velocity in Equation 3.3 is derived from the dynamic spectrum and can be modelled with a weighted sum of the pulsar, Earth, and IISM velocities (Equation 4.7) for a given scattering geometry. The transverse velocity of the pulsar V_p has components $v_{p,\parallel}$ along the line of nodes, and $v_{p,\perp}$ perpendicular to this in the plane of the sky. These velocities are defined as a function of orbital phase from the line of nodes $\phi = \theta + \omega$, where θ is the true anomaly and ω is the longitude of periastron,

$$(3.8) \quad \begin{pmatrix} v_{p,\parallel} \\ v_{p,\perp} \end{pmatrix} = \begin{pmatrix} -V_0 (e \sin \omega + \sin \phi) \\ V_0 \cos i (e \cos \omega + \cos \phi) \end{pmatrix} + \begin{pmatrix} v_{\mu,\parallel} \\ v_{\mu,\perp} \end{pmatrix}$$

where $V_0 = 2\pi xc/(\sin i P_b(1 - e^2)^{1/2})$ is the mean orbital velocity, x is the projected semi-major axis in seconds, P_b is the binary orbital period, e is the eccentricity, i is the inclination angle, and $v_{\mu,\parallel}$ and $v_{\mu,\perp}$ are the proper motion components of the pulsar velocity. These equations are in agreement with the previous work of Ord et al. (2002a) for PSR J1141–6545 and Rickett et al. (2014) for PSR J0737–3039A. The true anomaly was first calculated by numerically computing the eccentric anomaly, E from Kepler's equation $E - e \sin E = M$, with mean anomaly $M = (2\pi/P_b)(t - T_0)$, where T_0 is the epoch of periastron. The true anomaly is then given by

$$(3.9) \quad \theta = 2 \arctan \left[\sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2} \right].$$

To include the changing velocity of the Earth in our model, we calculate, for the position of PSR J1141–6545, the transverse component of the Earth's velocity V_E in right ascension $v_{E,\alpha}$, and declination $v_{E,\delta}$. We then rotate these into the coordinates of the pulsar velocity above, which is defined by the line of nodes and the plane of the sky. The angle of this rotation Ω , is the longitude of the ascending node, and the rotation is defined as

$$(3.10) \quad \begin{pmatrix} v_{E,\parallel} \\ v_{E,\perp} \end{pmatrix} = \begin{pmatrix} \sin \Omega & \cos \Omega \\ \cos \Omega & -\sin \Omega \end{pmatrix} \begin{pmatrix} v_{E,\alpha} \\ v_{E,\delta} \end{pmatrix}$$

where $v_{E,\parallel}$ and $v_{E,\perp}$ are the components of the Earth's velocity aligned with $v_{p,\parallel}$ and $v_{p,\perp}$ respectively. This was also used by Rickett et al. (2014) (but with the direction of the perpendicular axis reversed in their definition) to include the Earth's velocity and proper motion of PSR J0737–3039A into their scintillation model, which are both known in celestial coordinates. We then combine the pulsar and Earth velocities, scaling them appropriately by the distance to the scattering region s

$$(3.11) \quad \mathbf{V}(s) = \begin{pmatrix} v_{\parallel}(s) \\ v_{\perp}(s) \end{pmatrix} = s \begin{pmatrix} v_{E,\parallel} \\ v_{E,\perp} \end{pmatrix} + (1-s) \begin{pmatrix} v_{p,\parallel} \\ v_{p,\perp} \end{pmatrix}$$

In the case of isotropic scattering where the angular size of the pulsar orbit is compact enough to remain in the scattering disk, the spatial scale s_d is constant with orbital phase (but it may change on longer time scales if the IISM is inhomogeneous; see Section 3.2.3). Consequently, the decorrelation bandwidth Δv_d is also constant with orbital phase and the effective velocity is simply given by

$$(3.12) \quad V_{\text{eff}}(s) = |\mathbf{V}(s)| = \sqrt{v_{\parallel}(s)^2 + v_{\perp}(s)^2}.$$

However, for anisotropic scattering, s_d depends on the direction of V_{ISS} . To account for such scattering in PSR J0737–3039A, Coles et al. (2005) considered the spatial diffraction pattern as an ellipse. The pattern is then described by a quadratic form

$$(3.13) \quad Q(\mathbf{s}_d) = a s_{d,\parallel}^2 + b s_{d,\perp}^2 + c s_{d,\parallel} s_{d,\perp},$$

where the coefficients a , b , and c are parametrised by the axial ratio A_r of the ellipse and its orientation ψ with respect to the coordinates of the pulsar orbit as defined above. Rickett et al. (2014) used this anisotropy model for PSR J0737–3039A but parametrised the quadratic coefficients in terms of $R = (A_r^2 - 1)/(A_r^2 + 1)$, which is bound between 0 and 1. If the orientation angle ψ is defined clockwise from the line of nodes, then from Rickett et al. (2014) the coefficients are

$$(3.14) \quad \begin{aligned} a &= [1 - R \cos(2\psi)] / \sqrt{1 - R^2} \\ b &= [1 + R \cos(2\psi)] / \sqrt{1 - R^2} \\ c &= -2R \sin(2\psi) / \sqrt{1 - R^2}. \end{aligned}$$

Finally, we introduce a scaling factor κ to the model, which will account for any errors (for example an error in the pulsar distance D) in the calculation of V_{ISS} from the dynamic spectrum (Equation 3.3). Our final model for the effective velocity is then

$$(3.15) \quad V_{\text{eff}}(s) = \kappa \sqrt{Q(\mathbf{V}(s))},$$

which we use in Equation 3.6 for the uniform medium model and Equation 3.7 for the thin screen model.

The parameters in this model that are unknown from pulsar timing of PSR J1141–6545 are κ , s , i , Ω , $v_{\mu,\parallel}$, $v_{\mu,\perp}$, R , and ψ . As discussed in Section 3.2.3, the constraints from the known Earth’s velocity and $\dot{\omega}$ provide the additional degrees of freedom required to uniquely determine these parameters. For the uniform medium model we have one fewer parameter because we integrate over s (Equation 3.6). For the case of isotropic scattering, which we consider for both the uniform medium model and the thin screen model, we have $R = 0$ and the model reduces to Equation 3.12 scaled by the factor κ . However, we also found in Section 3.2.3 that it is necessary to re-estimate the spatial scale from Δv_d at each epoch because of AU-scale inhomogeneities in the IISM. The calculation of the spatial scale from these Δv_d measurements is imperfect, and Rickett et al. (2014) found some disagreement between Δv_d and the spatial scale (by analysing the harmonic coefficients described in the following section). To account for this, and some other correlated time-variability in properties such as the anisotropy and IISM velocity, we chose to use a scaling factor at each epoch. To do this, we scaled the data for each epoch to the mean V_{ISS} , then fitted for the mean scaling factor $\bar{\kappa}$ between the model and scaled data. Ideally $\bar{\kappa}$ would then represent the major systematic biases in the model, such as from an inaccurate assumed pulsar distance, while the scaling at each epoch accounts for time variations in s_d and perhaps some of the variation due to changing anisotropy and IISM velocity. Finally, as discussed in Section 3.2.1, because our τ_d are an average over an ~ 11 minute segment we also apply a smoothing window of this length to each of our models (including those described below) to avoid small biases in important parameters such as the inclination angle.

3.3.2 Comparison to the earlier models

3.3.2.1 Harmonic coefficients

Like PSR J1141–6545, which we analyse here, the double pulsar PSR J0737–3039A shows orbital modulation of the scintillation timescale, which was first measured and modelled by Ransom et al. (2004). The analysis was extended to include anisotropy in the IISM by Coles et al. (2005) and further extended to include the Earth’s velocity by Rickett et al. (2014), who also showed that the modulation of $1/\tau_d^2$ with orbital phase can be modelled as the sum of five harmonics

$$(3.16) \quad \frac{1}{\tau_d(\phi)^2} = K_0 + K_S \sin \phi + K_C \cos \phi + K_{S2} \sin(2\phi) + K_{C2} \cos(2\phi),$$

where the harmonic coefficients contain all of the information on the diffractive interstellar scintillation available in the data. The relationship between these coefficients and the physical parameters of the scattering and velocity are

$$\begin{aligned}
 (3.17) \quad K_0 &= [0.5V_0^2(a + b \cos^2 i) \\
 &\quad + a(v_{C,\parallel} - V_0 e \sin \omega)^2 \\
 &\quad + b(v_{C,\perp} + V_0 e \cos \omega \cos i)^2 \\
 &\quad + c(v_{C,\parallel} - V_0 e \sin \omega)(v_{C,\perp} + V_0 e \cos \omega \cos i)]/s_d^2 \\
 K_S &= -V_0[2a(V_{C,\parallel} - V_0 e \sin \omega) \\
 &\quad + c(v_{C,\perp} + V_0 e \cos i \cos \omega)]/s_d^2 \\
 K_C &= V_0 \cos i [c(v_{C,\parallel} - V_0 e \sin \omega) \\
 &\quad + 2b(v_{C,\perp} + V_0 e \cos i \cos \omega)]/s_d^2 \\
 K_{S2} &= -cV_0^2 \cos i / 2s_d^2 \\
 K_{C2} &= V_0^2(-1 + \cos^2 i) / 2s_d^2
 \end{aligned}$$

where $v_{C,\parallel}$ and $v_{C,\perp}$ are the constant (for a given epoch) components of the line-of-sight velocity, which will be dominated by the pulsar’s proper motion, but also includes the Earth’s velocity and any IISM velocity. Note that each of these coefficients are inversely proportional to s_d^2 , and that K_{S2} and K_{C2} are constant with time for a constant anisotropy. Rickett et al. (2014) used these facts to normalise the coefficients by K_{C2} , which corrects each one for changes to the spatial scale s_d with time and with observing frequency. This allowed them to model observations at multiple observing frequencies simultaneously.

We have measured these normalised harmonic coefficients ($k_0 = K_0/K_{C2}$, $k_S = K_S/K_{C2}$, $k_C = K_C/K_{C2}$, and $k_{S2} = K_{S2}/K_{C2}$) for each epoch of PSR J1141–6545 observations. The results are shown in Figure 3.4. It is worth noting that in the case of isotropic scattering (with $a = b = 1$ and $c = 0$), k_S only changes with time because of $V_{C,\parallel}$ (ignoring ω), k_C changes with $V_{C,\perp}$, and $k_{S2} = 0$. While we do find that k_{S2} is consistent with zero on average (Figure 3.4), there is some variation that is correlated with variations in the other normalised parameters. Since k_{S2} is constant with orbital phase (for a constant anisotropy over the orbit), this indicates epoch-to-epoch variation to the anisotropy, which will be an unmodelled source of noise in our data and is discussed further in Section 3.5.2.2. Because $k_{S2} = 0$ on average, we can also see that our long-term model must be consistent with $c = 0$, and thus should indicate that the scattering is isotropic, with $R = 0$, or anisotropic with orientation near $\psi = 0^\circ$ or $\psi = 90^\circ$.

3.3.2.2 Physical model for individual epochs

For the first scintillation analysis of PSR J1141–6545, Ord et al. (2002a) modelled the dynamic spectrum for a single 10 hr observation. From this dynamic spectrum they derived the V_{ISS} of Cordes and Rickett (1998) (Equation 3.3), and modelled this using only the pulsar velocities in




Figure 3.4: The normalised harmonic coefficients derived from a fit to $V_{\text{ISS}}(\phi)$ at each observing epoch, as described in Section 3.3.2.1.

Equation 4.9, with a scaling factor κ_u . They assumed that the scattering was isotropic and the velocity model was then given by

$$(3.18) \quad V_{\text{model}} = \kappa_u \sqrt{v_{\text{p},\parallel}^2 + v_{\text{p},\perp}^2},$$

with κ_u , i , ω , $v_{\text{C},\parallel}$, and $v_{\text{C},\perp}$ being the five fitted parameters. We use $v_{\text{C},\parallel}$ and $v_{\text{C},\perp}$ here instead of $v_{\mu,\parallel}$ and $v_{\mu,\perp}$ because they will include a contribution from the Earth’s velocity, and as a result will show annual variations with time. The parameter κ_u can be used to absorb any errors in the calculation of V_{ISS} , such as from an incorrect pulsar distance, which is what we intend for κ in Equation 3.15. However, since the pulsar velocity in this model is not scaled by a screen distance s because the changing Earth’s velocity is not included to constrain it, one could also interpret κ_u in this case as a scale factor for A_{ISS} . In this interpretation, κ_u is also related to the screen distance s and can account for different scattering geometries provided they are isotropic.

We repeat this analysis of Ord et al. (2002a) as described above for each of the 23 epochs of observations and we refer to this as the single-epoch “physical model”. We present the results from this model in Section 3.5.1.1 and show that this approach can also be used to measure the advance of periastron $\dot{\omega}$.

3.4 Fitting and comparing the models

The harmonic model for individual epochs was initially fitted with a weighted linear least-squares analysis to provide an initial solution. However this and all other models were also fitted with weighted nonlinear regression (using the LSQNONLIN function in MATLAB) after we accounted for the smoothing required for the models (since the data are smoothed by the ~ 11 minute window of the dynamic spectra segments). We use this method for all models because this was a convenient way of propagating the uncertainties to the physical parameters. For fitting the harmonic coefficients, the linear fit of the unsmoothed model provides the starting point for the fit, and the coefficients were unbounded. To fit the remaining models, we used a set of random initial guesses across the full parameter space, which were uniformly distributed in the physically-bound regions for each parameter (i.e. $0 < s < 1$, $0 \leq \Omega < 2\pi$, $0 \leq i < \pi$, $0 \leq \psi < \pi$ and $0 \leq R < 1$). For the pulsar's proper motion velocity components, we used normally distributed initial guesses with a mean and standard deviation of $\mu_v = 0 \text{ km s}^{-1}$ and $\sigma_v = 150 \text{ km s}^{-1}$ respectively, and for the scaling factor κ we used random guesses between $0 < \kappa \leq 2$. The nonlinear fit was used to find a local minimum from each initial guess, and we then selected the best-fit local minimum for each model based on the χ^2 value as a goodness-of-fit measure. In the cases where there are multiple interesting, or equally significant solutions, we present and discuss them in the following sections.

The uncertainties for each parameter in the resulting best-fit models are estimates drawn from the diagonal elements in the parameter covariance matrix \mathbf{C} , which was computed from the Jacobian of the model \mathbf{J} with $\mathbf{C} = (\mathbf{J}^T \mathbf{J})^{-1}$, where T denotes the transpose. Since there is excess noise in the data (which is likely due to time variability for scattering properties that are not accounted for in the model, as we have discussed) the parameter covariance matrix for each model fit is also scaled by the mean squared error (mse) to approximately account for this. The vector of parameter standard errors is then given by $\sigma = \sqrt{\text{diag}(\text{mse} \times \mathbf{C})}$.

For basic model comparison we use the Akaike Information Criterion (AIC; Akaike 1998), which compares the χ^2 statistics for two models differing by k parameters. The additional k parameters improved the model if $\Delta\chi^2 = \chi_2^2 - \chi_1^2 > 2k$. However there is excess noise in our data from variations in scattering properties from epoch-to-epoch (e.g. anisotropy, IISM velocity, and perhaps strength of scattering variations that are not fully accounted for by measurements of Δv_d), and as a result the reduced chi-squared values for the models ($\chi_r^2 = \chi^2/(n - m)$ for n observations and m parameters) will be large. To use the AIC for two models with large χ_r^2 , we scale the χ^2 of both models by the χ_r^2 of the model with additional parameters $\chi_{r,2}^2$, giving the new condition $\Delta\chi_c^2 = (\chi_2^2 - \chi_1^2)/\chi_{r,2}^2 > 2k$. This approximate solution is effectively a scaling of the measurement errors to account for the excess noise to set $\chi_{r,2}^2 = 1$, which is the requirement for the AIC to hold. We use this in the following sections to compare between uniform medium and thin screen models, as well as isotropic and anisotropic models.

3.5 Results and model comparison

Here we present and compare the various scintillation models for PSR J1141–6545. We first fitted each of the 23 epochs separately with a sum of five harmonics of orbital phase, which are described in Rickett et al. (2014) and in Section 3.3.2.1. We then fit a physical model for isotropic scattering to each epoch (Section 3.3.2.2), which provided measurements of ω and V_C over time to clearly show that the relativistic advance of periastron and the modulation from the Earth’s velocity can be recovered in the data, and to compare the results with the earlier analysis of Ord et al. (2002a). These models are presented below in Section 3.5.1.1.

We then held the periastron advance fixed at the precisely-measured value from pulsar timing, $\dot{\omega} = 5.3096^\circ \text{yr}^{-1}$ (Bhat et al. 2008), and calculated the components of the Earth’s velocity (V_E) transverse to the line-of-sight of PSR J1141–6545 for each observation. Then we combined the 23 epochs of data and fit several long-term scintillation models, using the additional degrees of freedom provided by $\dot{\omega}$ and V_E to help constrain additional parameters including the scattering anisotropy and the longitude of the ascending node Ω . We fit these long-term velocity models for two distributions of Kolmogorov turbulent plasma along the line-of sight: A uniform distribution, and a thin screen.

We first present and compare the isotropic models, for both individual epochs and long term data, in Section 3.5.1. For these models we compare the results of the long-term analysis to the single epochs, and discuss the multiple solutions that arise from a degeneracy between i and $v_{\mu,\perp}$. We show that the annual variations in the long-term analysis can be used to break the degeneracy and provides evidence for the sense of the inclination angle and distribution of plasma along the line-of-sight (though the effects of anisotropy should be considered before making strong conclusions). We then give the long-term models of anisotropic scattering for both scattering geometries in Section 3.5.2.

Finally, we have included in Appendix A the links to access the data (raw observations and processed dynamic spectra) used for this work, as well as the MATLAB code used for the analysis for the purpose of reproducibility, including the scripts used to measure scintillation parameters from the dynamic spectra and fit the scintillation velocity as described in the previous section.

3.5.1 Isotropic scattering

3.5.1.1 Individual epochs

For each epoch of observations shown in Figure 3.3, we fit the V_{ISS} of Equation 3.3 as a function of orbital phase ϕ , with a sum of five harmonics (Section 3.3.2.1). The resulting normalised harmonic coefficients are given in Figure 3.4 as a function of time. In the case of PSR J0737–3039A, the inclination of the orbit is essentially edge-on, which means that k_C and k_{S2} are almost zero (Rickett et al. 2014). This is not true for PSR J1141–6545, but we do still find that the mean of k_{S2} is consistent with zero. This suggests that the scattering may be nearly isotropic since k_{S2} is

independent of the changing pulsar and Earth velocities. There is however some time variability in k_{S2} , which is correlated with some other normalised harmonic coefficients. This may be due to time variability in the anisotropy, which one would expect even from random realisations of a truly isotropic medium. This time variability is not evidence for anisotropic scattering on average, but it is a source of noise for the long-term scintillation models. This may be the primary reason for a large reduced chi-squared value for the long-term models (Section 3.5.1.2 and Section 3.5.2).

The evidence for isotropic scattering is important because the single-epoch physical model (Section 3.3.2.2) depends on it. With only five degrees of freedom available from a single epoch of observations, accurate measurement of the pulsar proper motion and inclination angle can only be made for isotropic scattering (Coles et al. 2005). We therefore fit this physical model to each of the epochs under the assumption of isotropic scattering to obtain a time series for each of the physical parameters described in Section 3.3.2.2. However we find that there are four possible solutions for each epoch, arising from a known degeneracy between the proper motion and inclination angle (Lyne 1984). This produces a more edge-on solution with higher proper motion and a more face-on solution with lower proper motion, and their corresponding pairs with an opposite “sense” of inclination about $i = 90^\circ$. Ord et al. (2002a) considered only the $i < 90^\circ$ solutions and determined the more edge-on solution to be favorable using the implied pulsar mass. For each of our fits we also take the $i < 90^\circ$ solutions, and we show in Section 3.5.2.2 that this physical.

The time series of our fitted ω values is shown in Figure 3.5. The clear gradient is the advance of periastron, $\dot{\omega} = 5.6 \pm 0.3^\circ/\text{yr}$, and is close to the measurement of pulsar timing $\dot{\omega} = 5.3096 \pm 0.0004^\circ/\text{yr}$ (Bhat et al. 2008). As expected, κ_u and i are constant with time with the exception of random variations similar to those seen in the harmonic coefficients (Figure 3.4). The weighted mean and standard deviation for these parameters across all epochs are given in left two columns of Table 3.1, for both degenerate solutions. The inclination angles from these solutions $i = 80.1 \pm 1.1^\circ$ and $i = 71 \pm 3^\circ$ are both consistent with the previous scintillation measurement of 76 ± 2.5 , to within about 1σ of both measurements. However our more face-on solution is more consistent with the inferred $i = 73 \pm 1.5^\circ$ from the mass constraints derived from general relativity (Bhat et al. 2008). The reduced chi-squared value for these solutions (across all epochs, with total number of parameters $m = 115$; five per epoch), is $\sim \chi_r^2 = 1.5$, suggesting that either the measurement errors for scintillation parameters are underestimated, there is excess noise in the data, or the model is incomplete. The χ_r^2 for individual epochs with observations spanning several days is generally higher than the few epochs with a single day, which is usually close to unity. It is therefore likely that the cause is small random changes in the scattering on a timescale of \sim days, such as variations to the anisotropy. Such variations may be physical, or an observational effect since we are observing a finite number of scintles and would expect random variations from the ensemble average.

The more edge-on inclination angle is consistent with the equivalent from Ord et al. (2002a),




Figure 3.5: The longitude of periastron, ω as a function of time, measured independently at each of the 23 epochs shown in Figure 3.3. The solid line is a weighted best fit, where the gradient, $\dot{\omega} = 5.6 \pm 0.3^\circ \text{yr}^{-1}$, is the advance of periastron and is consistent with the measurement from pulsar timing (red dashed line).

but κ_u is significantly smaller. We note that the $V_{\text{ISS}}(\theta)$ shown in Figure 2 of Ord et al. (2002a) is larger by approximately a factor of two compared with our data in Figure 3.3. This explains the variation in κ_u , and is more likely the result of systematic errors than real changes in the IISM, given the stability with time seen in these later six years. We also note that Ord and Bailes (2005) attempted to model the long-term scintillation using data collected in 2002 and 2003, shortly after the work of Ord et al. (2002a), and their scintillation velocities are more consistent with our results.

In Figure 3.6 we show the time series for $v_{\text{C},\parallel}$ (top panel) and $v_{\text{C},\perp}$ (bottom panel). There is a clear annual modulation, demonstrating that the Earth’s velocity can be used to constrain additional parameters in a long-term model. We calculate the transverse components of the Earth’s velocity (in celestial coordinates) in the direction of PSR J1141–6545, and then scale it by the “screen distance” s , and rotate it by Ω to align with $v_{\text{C},\parallel}$ and $v_{\text{C},\perp}$. These parameters, for both solutions, are also given in the left two columns of Table 3.1 and are approximately equivalent to the parameters in the long-term models and we compare them in the next section. The scaled and rotated Earth velocity is then shown on Figure 3.6 with red diamonds.




Figure 3.6: Components of the constant (with orbital phase) transverse velocity parallel ($v_{C,||}$, top panel) and perpendicular ($v_{C,⊥}$, bottom panel) to the line of nodes for each observing epoch. The measured velocity is a scaled combination of the Earth and IISM velocities and the pulsar proper motion. The Earth’s contribution is apparent from the clear annual modulation. The grey line is a weighted best fit annual sine wave to each of the time series. The red diamonds are the calculated Earth’s velocity at each epoch, scaled down by a screen distance $s = 0.36 \pm 0.07$, and rotated from celestial coordinates by $\Omega = 25 \pm 10^\circ$ (i.e. “Solution 1” from Table 3.1).

3.5.1.2 Long-term models

Our first long-term model is that of a Kolmogorov turbulent plasma distributed uniformly along the line-of-sight (Equation 3.6). Such a situation may be rare in observations of nearby pulsars since scattering is often dominated by one or a few compact regions, or “thin screens”, (e.g. Stinebring et al. 2001; Coles et al. 2015). For more distant sources such as PSR J1141–6545 we might expect to find that the scattering is extended uniformly, or appears to be as a result of combining the effects of many scattering regions across the line of sight. We find that there are no clear parabolic scintillation arcs detected in secondary spectra for any observation over the six years we analyse. This might be because of insufficient data quality, or it may be suggestive of scattering that is nearly isotropic and/or extended along the line-of sight (Stinebring et al. 2001). Accordingly, we fit both isotropic and anisotropic models assuming a uniform medium. The parameters for each of the isotropic models are shown in Table 3.1, and the extension to anisotropy is discussed in the next section.

A thin screen model is often appropriate, but requires one additional parameter to describe,

Table 3.1: Measured parameters for each of the isotropic scintillation models. The first two columns are parameters from single-epoch fits of a physical model. κ_u and i are the weighted mean values from the fit, while Ω , s , $v_{\mu,\parallel}$, and $v_{\mu,\perp}$ are derived from a simple model of the annual variations in a time series of the measured $v_{C,\parallel}$, and $v_{C,\perp}$ (Section 3.5.1.1). Two degenerate solutions of equal quality are given for these single-epoch fits, and the equivalent solutions are also given in the next four columns of long-term isotropic models. Columns three and four are the parameters from a uniform medium while columns five and six are for a thin screen; these are discussed in Section 3.5.1.2. The chi-squared χ^2 and reduced chi-squared χ_r^2 values for each model are given, and these are used for model comparison. Numbers in brackets are the 1σ standard errors on the last quoted digit.

Model type	Measured parameters					
	Single epochs [†]		Long-term			
	Scattering geometry		Uniform medium		Thin screen	
Solution #	Solution 1	Solution 2	Solution 1	Solution 2	Solution 1	Solution 2
κ_u and $\bar{\kappa}$	0.78(7)	0.78(7)	0.794(4)	0.797(4)	0.626(19)	0.59(2)
s	0.36(7)	0.26(5)	—	—	0.203(10)	0.184(10)
Ω (°)	25(10)	39(12)	18.2(17)	33.8(18)	11(4)	38(4)
i (°) [†]	80.1(14)	71(3)	78.94(16)	72.5(3)	80.19(13)	70.32(17)
$v_{\mu,\parallel}$ (km s ⁻¹)	11.8(15)	14.2(15)	13.0(5)	14.1(5)	12.5(5)	14.6(5)
$v_{\mu,\perp}$ (km s ⁻¹)	75(2)	41.4(16)	69.1(9)	45.7(7)	77.0(7)	41.1(6)
χ^2	2211	2211	4307	4531	3790	4013
χ_r^2	1.499	1.499	2.718	2.858	2.390	2.530

[†] The previous measurements are $i = 76 \pm 2.5^\circ$ from scintillation (Ord et al. 2002a) and $i = 73 \pm 1.5^\circ$ from pulsar timing, assuming general relativity to be correct (Bhat et al. 2008).

compared with the uniform model. This is the fractional screen distance s , defined to be 0 at the pulsar and 1 at the observer. For both the uniform medium and thin screen models, we give the two solutions that correspond to the degenerate solutions from fits to individual epochs. The parameters are given in the right four columns of Table 3.1, and again we show only the $i < 90^\circ$ solutions because for these models with annual variations accounted for, the $i > 90^\circ$ solutions have a higher χ^2 (and this remains true when including anisotropy in the models; Section 3.5.2). Unlike the fits to individual epochs, the long-term models allow model comparison because of the additional degrees of freedom provided by the annual variations. We find that the more edge-on solutions are favoured for both the uniform medium solutions, with $\Delta\chi_c^2 \sim 82$, and the thin screen solutions, with $\Delta\chi_c^2 \sim 93$. This is consistent with Ord et al. (2002a), but these more edge-on solutions are higher than the inclination angle inferred from general relativity (Bhat et al. 2008), likely because of the assumption of isotropic scattering.

The addition of the extra parameter s for the thin screen model also comfortably passes the AIC, with $\Delta\chi_c^2 = 216$ (from Table 3.1, comparing the preferred “Solution 1” for both models). The best-fit model suggests that the scattering is located nearer to the pulsar, with $s = 0.203 \pm 0.010$, which is constrained simply by the relative amplitude of annual modulation to the orbital modulation.

3.5.2 Anisotropic Scattering: Long-term models

Diffraction scintillation is often dominated by a single compact and turbulent scattering region in the line-of-sight, and we found from the isotropic models in the previous section that the thin screen model provided a better fit than the uniform medium. However when accounting for anisotropy in the models, there is a potential for the results to be different because the model has the freedom to change the amplitude of the annual variations without the screen distance parameter s . We therefore consider again the uniform medium model as well as the thin screen for the case of anisotropic scattering.

From the results in the previous section, the Earth’s velocity is observed to be significantly modulating V_{ISS} for this pulsar. Including the known V_E in our long-term model provides important constraints and allows us to break model degeneracies. Ordinarily we would find a “reflected” solution with opposite sense of the inclination angle (and direction of proper motion) about $i = 90^\circ$. However these models also give different orbital orientations (parametrised in celestial coordinates by the longitude of the ascending node Ω), meaning that the Earth’s velocity can be used to distinguish them. This was first successfully applied by Rickett et al. (2014) for PSR J0737–3039A to determine a unique inclination angle. Although in our case the scattering appears to be fractionally further from the Earth, whether uniform (which is similar to a screen at $s \sim 0.5$) or a thin screen (at $s = 0.257$, from Section 3.5.2.2). Because of this, the effect of the Earth’s motion is diminished with respect to PSR J0737–3039, but we have more data than Rickett et al. (2014) had available. Using V_E with these long-term models, we are indeed able to resolve the inclination ambiguity and we find that the $i < 90^\circ$ solutions are physical (Section 3.5.2.2). For both models there is a unique best-fit solution, which is shown in the left two columns of Table 3.3, and discussed in the following sections.

3.5.2.1 Uniform medium

Our measurement of anisotropic scattering for a uniform medium is essentially the average across the entire line-of-sight. We found that the thin screen model was preferred in the case of isotropic scattering simply because of the relative amplitude of the annual variations in V_{ISS} . However this amplitude could also be reduced by anisotropic scattering in combination with the other parameters, so we have repeated the uniform medium fit with anisotropy. The result, in the left column of Table 3.3, is a unique solution with an inclination angle of $i = 38 \pm 3^\circ$, which is ruled out by the mass constraints from pulsar timing (Bhat et al. 2008). It also gives a very high proper motion velocity, with $v_{\mu,\perp} = 530 \pm 50 \text{ km s}^{-1}$ and an especially small scaling factor $\bar{\kappa} = 0.236 \pm 0.018$, which are both unlikely. We searched for other solutions by constraining the inclination angle in the range $1 \text{ rad} < i < \frac{\pi}{2} \text{ rad}$, but found no solutions with a χ^2 comparable to the anisotropic thin screen model presented in the following section. The conclusion is therefore that the dominant scattering is localised to some position on the line-of-sight and the thin screen model is appropriate.

Table 3.2: Measured parameters for anisotropic long-term models for model comparison. Columns one and two are the best-fit solutions for an anisotropic uniform medium and thin screen respectively. The uniform medium model is ruled out based on the low inclination angle and high space velocity (Section 3.5.2). Bounding the inclination angle to $1 \text{ rad} < i < \frac{\pi}{2} \text{ rad}$ does not reveal a valid solution and instead returns the boundary condition $i = 1 \text{ rad}$. We therefore only considered the thin screen model with inclination angle fixed at values corresponding to the $\sin i$ constraint of pulsar timing. The two best solutions for each of $i = 73^\circ$ and $i = 107^\circ$ are given in the right four columns for the case of an anisotropic thin screen (and these are discussed in Section 3.5.2.2). The chi-squared χ^2 and reduced chi-squared χ_r^2 values for each model are given, which we use for model comparison. Numbers in brackets are the 1σ standard errors on the last quoted digit.

Inclination angle Scattering geometry Solution #	Measured parameters					
	Free parameter		Fixed at $i = 73^\circ$		Fixed at $i = 107^\circ$	
	Uniform medium	Thin screen	Thin screen		Thin screen	
			Solution 1	Solution 2	Solution 1	Solution 2
$\bar{\kappa}$	0.236(18)	0.45(5)	0.552(18)	0.62(3)	0.44(3)	0.51(3)
s	–	0.295(18)	0.257(11)	0.175(11)	0.180(15)	0.126(12)
Ω ($^\circ$)	31(3)	27(3)	23(3)	40(5)	41(5)	4(6)
i ($^\circ$) [†]	38(3)	62(5)	73	73	107	107
$v_{\mu,\parallel}$ (km s^{-1})	29(3)	20.2(15)	18.4(12)	14.9(5)	17.7(14)	15.1(8)
$v_{\mu,\perp}$ (km s^{-1})	530(50)	221(45)	129(6)	37.1(9)	−129(4)	−36.8(12)
ψ ($^\circ$)	88.8(3)	87.7(7)	85.4(9)	10(3)	94.1(12)	169(6)
R	0.905(7)	0.75(7)	0.48(4)	0.112(14)	0.49(3)	0.12(2)
χ^2	3640	3570	3592	3929	4065	4214
χ_r^2	2.299	2.257	2.271	2.483	2.569	2.664

[†] The previous measurements are $i = 76 \pm 2.5^\circ$ from scintillation (Ord et al. 2002a) and $i = 73 \pm 1.5^\circ$ from pulsar timing, assuming general relativity to be correct (Bhat et al. 2008).

3.5.2.2 Thin screen

We showed in Section 3.5.1.2 that the isotropic thin screen model gave an inclination angle that was inconsistent with the current constraint from pulsar timing, which suggests either an underestimated parameter uncertainty or a bias due to an incomplete model, or both. The obvious source for this discrepancy would be the assumption of isotropic scattering, since Coles et al. (2005) showed that anisotropy in the IISM can produce a bias in proper motion velocity measurements (for PSR J0737–3039), and we find here that $v_{\mu,\perp}$ in particular is closely tied with i . Any anisotropy cannot be accounted for by measuring a single epoch of V_{ISS} alone because there are not enough degrees of freedom available for the additional two parameters. Instead, Coles et al. (2005) was able to use the correlated scintillations of both pulsars in the double pulsar system at their apparent closest approach (when the magnetosphere of pulsar B eclipses pulsar A) to produce a spatial correlation pattern that revealed the anisotropy. Later, Rickett et al. (2014) was able to use the annual variations in the harmonic coefficients for PSR J0737–3039 to measure the anisotropy and determine the sense of the inclination angle by comparing the two models with inclinations fixed by the $\sin i$ measurement from pulsar timing. They also showed that the anisotropy and IISM velocity (because the proper motion velocity was known from

timing) could be measured from a model of the individual ACFs alone, but that the measurements differed from the average value obtained with long-term fitting, suggesting some weakness in the model and difficulty in confidently determining the anisotropy. The annual variation approach is equivalent to the long-term physical models we use here, and we are unable to use the ACFs of PSR J1141–6545 for anything other than estimating the phase gradient for each observation (see Section 3.2.3).

The best anisotropic thin screen model, with inclination angle fitted, is shown in Figure 3.7 with parameters given in the second column of Table 3.3. Each observation epoch has been independently scaled to the mean V_{ISS} in Figure 3.7 as described in Section 3.3, and the apparent white noise in the model is the out-of-phase variation due to the Earth’s velocity. For this model we find $i = 62 \pm 5^\circ$, which is lower than the $73 \pm 1.5^\circ$ given by the general relativistic mass constraints, and the $80.19 \pm 0.13^\circ$ given by the isotropic thin screen model. The large uncertainty for i suggests that the isotropic model uncertainties were greatly underestimated, and that the anisotropy is important to consider. The anisotropy for this model is $R = 0.75 \pm 0.07$, which corresponds to an axial ratio $A_r = 2.8 \pm 0.5$.

If the scattering were truly isotropic, we would expect that each observation samples a single realisation of this, and thus would randomly appear slightly anisotropic, with an rms of $R_{\text{rms}} \sim 0.7(s_d/s_r)^{-1/6}$ (Romani et al. 1986; Rickett et al. 2014). In our case $R_{\text{rms}} \sim 0.18$, and from 23 epochs we would expect to observe $R_{\text{rms}}/\sqrt{23} \sim 0.04$. The measurement is significantly larger than this, which suggests that there is significant anisotropy in the scattering, which is supported by the additional parameters passing the AIC, with $\Delta\chi_c^2 \sim 97$ when compared to the isotropic thin screen model. We find $\psi = 87.7 \pm 0.7^\circ$, which is expected because we have $k_{S2} \approx 0$ on average (Figure 3.4), so any anisotropic solution must have $\sin(2\psi) = 0$ (from Equations 3.14 & 3.17) and thus $\psi \approx 0^\circ$ or $\psi \approx 90^\circ$. This close alignment with the pulsar orbit may be chance, or may be the result of the anisotropy absorbing some systematic errors in the data. Future observations and studies of PSR J1141–6545 could be designed to maximise the detail in the ACFs to allow for an independent measurements of the anisotropy and its time variability, which could then be used in these long-term V_{ISS} models to improve the inclination angle measurement.

To improve the measurements of the other parameters in the model, and to resolve the sense of the inclination angle, we repeat the fit of this anisotropic model with the inclination angle fixed at values corresponding to the $\sin i$ inferred from pulsar timing; $i = 73^\circ$ and $i = 107^\circ$ (Bhat et al. 2008). We find two solutions for each inclination angle; one significantly anisotropic solution (though less anisotropic than the model with i as a free parameter) with ψ near to 90° , and one slightly anisotropic model with a smaller $v_{\mu,\perp}$. These solutions are also shown in Table 3.3, and it is worth noting that the two inclination angles give opposite signs for $v_{\mu,\perp}$, meaning that resolving this ambiguity also provides a unique proper motion velocity for the first time. The solutions with higher anisotropy (that are approximately aligned with orbit) are preferred by the AIC, and the best solution with $i = 73^\circ$ is preferred over the equivalent with $i = 107^\circ$ with a high




Figure 3.7: Scintillation velocity, V_{ISS} as a function of orbital phase, ϕ for ~ 6 years of PSR J1141–6545 observations at a frequency of ~ 1400 MHz. The orbital phase was calculated using the ω and $\dot{\omega}$ values measured in the pulsar timing model. V_{ISS} in this case is defined at the scattering screen (IISM frame), because the distance to the screen was a parameter in the model. The best-fit anisotropic-IISM model is shown as the red line, and the apparent white-noise in the model is due to the out-of-phase Earth’s velocity. The data in each epoch have been scaled to the mean V_{ISS} and the model is scaled by $\bar{\kappa}$ to match.

significance $\Delta\chi_c^2 \sim 208$. The anisotropy for this model is $R = 0.48 \pm 0.04$, which corresponds to an axial ratio of $A_r = 1.70 \pm 0.08$.

Our model with i as a free parameter passes the AIC but gives only a marginal improvement over the fixed $i = 73^\circ$ model, with $\Delta\chi_c^2 \sim 10$. We therefore take the other parameters from the model with the more precise $i = 73^\circ$ measurement. This includes $\Omega = 23 \pm 3^\circ$, which we use in combination with the scaling factor $\bar{\kappa} = 0.552 \pm 0.18$ and unique velocities $v_{\mu,\parallel} = 18.4 \pm 1.2 \text{ km s}^{-1}$ and $v_{\mu,\perp} = 129 \pm 6 \text{ km s}^{-1}$ in Section 3.6.1 to estimate the pulsar proper motion in celestial coordinates for the first time.

The reduced chi-squared value for the best models is $\chi_r^2 \sim 2.3$, and (as mentioned for the isotropic model as well) this is likely high because of model errors, such as our assumption of a time-stationary anisotropy and proper motion. The proper motion of the pulsar is not expected to change, but the velocity of the IISM could be changing from epoch-to-epoch, as can the anisotropy. As discussed above, even for a truly isotropic screen we would see epoch-to-epoch variation in the apparent anisotropy. Indeed the normalised harmonic coefficients in Figure 3.4 show

some time variability that would not be explained by the Earth’s motion or by changes to the strength of scattering that would be taken into account by scaling of each epoch. This random time-variability is likely to be our largest source of noise and the major reason for a large χ_r^2 , but there may also be some contribution from measurement errors on the scintillation parameters for example. Obtaining independent constraints on the anisotropy and IISM velocity from modelling high-resolution ACFs directly (e.g. Rickett et al. 2014) is the best chance for future work to improve the quality of the fit and the precision of the inclination angle measurement.

3.6 Discussion

Our modelling of the long-term scintillation of PSR J1141–6545 shows that the dominant scattering region is centred at a fractional distance of $s = 0.257 \pm 0.011$ from the pulsar, and that the scattering is slightly anisotropic with an axial ratio of $A_r = 1.70 \pm 0.08$. We were able to provide an independent measurement of the orbital inclination angle $i = 62 \pm 5^\circ$ from the anisotropic thin screen model, however this was less precise than the value of $i = 73 \pm 1.5^\circ$ obtained from pulsar timing (Bhat et al. 2008). We repeated the fit with this more reliable inclination angle, to obtain improved measurements of the other scintillation parameters. The measured parameters from both of these models, the derived anisotropy axial ratio A_r , the mass of pulsar M_p , and the mass of the companion M_c are given in Table 3.6.1. We also derive a new estimate of the pulsar distance D and proper motion in celestial coordinates for the first time, as described below in section 3.6.1. We then discuss the implications of these derived measurements for future tests of general relativity in Section 3.6.2.

3.6.1 Pulsar distance and proper motion estimates

The distance to PSR J1141–6545 is currently poorly constrained, with the best estimate of $D = 3 \pm 2$ kpc from Verbiest et al. (2012) originating from a luminosity bias correction to an earlier lower-limit of 3.7 ± 1.7 kpc that was derived from neutral hydrogen absorption spectrum (Ord et al. 2002b). There is also a lower estimate of 1.7 kpc from the most recent Galactic electron-density model (Yao et al. 2017). This poor measurement precision is problematic for our data because our derivation of V_{ISS} from the scintillation parameters is proportional to \sqrt{D} . We accordingly fitted for a scaling factor $\bar{\kappa}$ (see Section 3.3.1), which would absorb any errors in the calculation of V_{ISS} , for example from an error in the assumed pulsar distance of $D = 3$ kpc. However, this scaling factor would also include systematic errors in the measurements of Δv_d and τ_d , as well as any error in the numerical relationship between Δv_d , the strength of scattering, and s_d derived by Cordes and Rickett (1998). This numerical relationship is summarised by the A_{ISS} coefficient, $A_{\text{ISS}} = 2.78 \times 10^4 \sqrt{2(1-s)/s} \text{ km s}^{-1}$ for the thin screen model.

We have also chosen to scale each epoch of observations to the mean V_{ISS} , which will primarily take into account any epoch-to-epoch errors in the relationship between the strength of scattering

and Δv_d . In Section 3.2.3 we showed that it was necessary to re-estimate Δv_d at each epoch due to inhomogeneities in the IISM on an AU scale. However we also know that the derivation of s_d from Δv_d is imperfect and can depend on the shape of the ACF among other factors. This was also noted by Rickett et al. (2014), who found discrepancies between their Δv_d measurements and the spatial scale inferred from measurements of the harmonic coefficient K_{C2} with time. For our analysis, we intend for our scaling of each epoch to take into account such discrepancies with time, which may also come partially from time-evolution in other parameters, such as the anisotropy. However, we assume that the scaling factor for the mean V_{ISS} , $\bar{\kappa}$, represents the largest systematic error in the calculation of V_{ISS} , and we assume that this is from an error in the pulsar distance D .

For the two best models, we show this scaling factor $\bar{\kappa}$ and its error in Table 3.6.1. For the model with fixed inclination angle, $\bar{\kappa} = 0.552 \pm 0.018$, suggesting a significant systematic error in the V_{ISS} . If this is indeed a measurement of the error in pulsar distance, we can use this to provide a new estimate. To do so, we first assume a 20% uncertainty on A_{ISS} , originating from a discrepancy of approximately this magnitude between the spatial scale s_d derived from the measured Δv_d , and that derived from K_{C2} in this work and that of Rickett et al. (2014). We then calculate the distribution for pulsar distance given that $D \propto 1/\bar{\kappa}^2$ and assume a Gaussian distribution for $\bar{\kappa}$ with mean and standard deviation given by the measurement in Table 3.6.1. We do this for both models (for comparison), and find that the distance is estimated to be larger than the assumed 3 kpc, but that it is poorly constrained by this method. The fixed- i model gives a large range of $D = 9_{-3}^{+5}$ kpc. This shows that in principle scintillation can be used to estimate pulsar distances, but in practice it is complicated by the $D \propto 1/\bar{\kappa}^2$ relationship and by the large uncertainty in A_{ISS} . With improved scintillation modelling and understanding of the IISM along the line-of-sight, we expect that the precision on distance estimates with this method can be improved in the future.

Using our own distance estimate (for self-consistency) with our measurements of the pulsar's transverse velocity in coordinates of the pulsar's orbit, $v_{\mu,\parallel}$ and $v_{\mu,\perp}$, and the longitude of the ascending node Ω we have been able to derive the pulsar proper motion in celestial coordinates (right ascension α and declination δ). We use the same distance distribution derived from $\bar{\kappa}$ above, and calculated the distribution of proper motions assuming Gaussian distributions for $v_{\mu,\parallel}$, $v_{\mu,\perp}$, and Ω given by the measurements and their uncertainties listed in Table 3.6.1. The derived proper motions in α and δ are shown in the table for both models, and are consistent due to the large uncertainty. For the fixed- i model, we have $\mu_\alpha = 2.5 \pm 1.2 \text{ mas yr}^{-1}$ and $\mu_\delta = 1.5 \pm 0.7 \text{ mas yr}^{-1}$. We believe that this is a fairly conservative estimate (because of the large uncertainty on D), but may be under-estimated if other existing distance estimates are correct. For example, taking the current lowest distance estimate of $D = 1.7$ kpc from Yao et al. (2017) and assuming a 20% uncertainty, we would have $\mu_\alpha = 15 \pm 4 \text{ mas yr}^{-1}$ and $\mu_\delta = 9 \pm 3 \text{ mas yr}^{-1}$. However a proper motion of this magnitude may soon be ruled out with improved pulsar timing sensitivity.

Table 3.3: Two best-fit long-term models for an anisotropic thin scattering screen, with derived parameters. For the pulsar distance D , we assume a 20% uncertainty on the A_{ISS} coefficient from (Cordes and Rickett 1998). The parameters are described and discussed in Section 3.6.1.

Measured parameters		
	Fitted i	Fixed $i = 73^\circ$
$\bar{\kappa}$	0.45(5)	0.552(18)
s	0.295(18)	0.257(11)
Ω ($^\circ$)	27(3)	23(3)
i ($^\circ$)	62(5)	73
$v_{\mu,\parallel}$ (km s^{-1})	20.2(15)	18.4(12)
$v_{\mu,\perp}$ (km s^{-1})	221(45)	129(6)
ψ ($^\circ$)	87.7(7)	85.4(9)
R	0.75(7)	0.48(4)
Derived parameters		
A_r	2.8(5)	1.70(8)
M_p (M_\odot)	1.18(6)	1.27
M_c (M_\odot)	1.11(5)	1.02
D (kpc)	13^{+9}_{-5}	9^{+5}_{-3}
μ_α (mas yr^{-1})	2.8(16)	2.5(12)
μ_δ (mas yr^{-1})	1.8(10)	1.5(7)
χ^2	3570	3592
χ_r^2	2.257	2.271

The estimated proper motion is highly uncertain, primarily because of uncertainty in the pulsar distance. For the following section we will assume the values derived from our model of $\mu_\alpha = 2.5 \pm 1.2 \text{ mas yr}^{-1}$ and $\mu_\delta = 1.5 \pm 0.7 \text{ mas yr}^{-1}$. However, if in the future the pulsar distance is constrained with higher confidence, the proper motion should be re-derived from our measurements of $v_{\mu,\parallel}$, $v_{\mu,\perp}$, and Ω . Alternatively, given a measurement of the pulsar proper motion in the near future from improved timing precision, our independent measurements of the velocity due to proper motion can be used to derive the distance.

3.6.2 Implications for timing and tests of general relativity

PSR J1141–6545 is a highly relativistic pulsar in an eccentric, asymmetrical mass system, which makes it an ideal laboratory for testing general relativity. Bhat et al. (2008) analysed the gravitational radiation losses from this system through pulsar timing and noted that with increasing precision of the orbital period-derivative, contamination from kinematic effects (e.g. the Shklovskii effect) and Galactic acceleration would start to dominate the uncertainty in the near future. The transverse velocity of the pulsar system is accompanied by a radial acceleration, which produces a time-dependent Doppler-shift to the pulsar spin frequency and orbital period. This is the Shklovskii effect (Shklovskii 1970), and it results in an apparent orbital period-derivative \dot{P}_b^{kin} that is considered a contamination to the orbital period measurement from gravitational

radiation losses, \dot{P}_b^{GR} . However, the exact level of this effect was unknown because the proper motion was not measured via pulsar timing. We are now able to determine the contribution from the Shklovskii effect for the first time because our scintillation work is sensitive to the transverse motion of the pulsar instead of the radial motion probed by pulsar timing.

This Shklovskii effect can be calculated from $\dot{P}_b^{kin} = DP_b\mu^2/c = P_bV_\mu^2/cD$ (Bell and Bailes 1996). With the distance and transverse velocity provided by our model, we calculate $\dot{P}_b^{kin} = 3.5 \times 10^{-15}$, which is $\sim 1\%$ of the \dot{P}_b measurement of Bhat et al. (2008) and well below the 6% measurement precision. However, if we again assume the distance from Yao et al. (2017), then $\dot{P}_b^{kin} = 1.9 \times 10^{-14}$. This is at the level of the current expected timing precision of $\sim 2\%$ for \dot{P}_b . It is therefore important to further constrain the system transverse velocity and/or distance from improved scintillation modelling and/or pulsar timing. In addition to the Shklovskii effect on \dot{P}_b , the proper motion changes the projected geometry of the binary orbit, resulting in apparent \dot{x} and $\dot{\omega}$. However, these kinematic contaminations are far below the measurement precision for this pulsar and are typically only observed in precisely-timed millisecond pulsars (Kopeikin 1996).

3.7 Conclusion

We have presented new scintillation models for PSR J1141–6545 using six years of data from the Parkes 64 m radio telescope. We found that like many pulsars, the scattering shows some anisotropy, and is dominated by a single scattering region centred at $s = 0.257 \pm 0.011$. By accounting for anisotropy in the scattering, we measured the system inclination angle of $i = 62 \pm 5^\circ$, which is less precise than previous measurements, suggesting that previous errors were underestimated. The value is also lower than the more precise constraint of $73 \pm 1.5^\circ$ from pulsar timing (inferred from the masses using general relativity), suggesting that a weakness in the model, such as time-stationary IISM velocity and anisotropy cause some systematic errors. However, using a model with the inclination angle fixed at 73° , we have been able to measure several astrometric parameters for the first time.

Using the significant annual and relativistic variations observed in the scintillation velocity to constrain our models, we have been able to resolve the “sense” of the inclination angle, and we find that $i < 90^\circ$. This in turn resolved the ambiguity in the direction of the proper motion velocity in pulsar coordinates. With our new measurement of the orientation of the orbit in celestial coordinates $\Omega = 23 \pm 3^\circ$, and estimate of the pulsar distance $D = 9_{-3}^{+5}$ kpc, we have been able to estimate the proper motion for the first time. We determine the proper motion in right ascension $\mu_\alpha = 2.5 \pm 1.2 \text{ mas yr}^{-1}$ and in declination $\mu_\delta = 1.5 \pm 0.7 \text{ mas yr}^{-1}$, and we use these numbers to calculate the contribution of the Shklovskii effect to \dot{P}_b . This effect is the most significant source of contamination for tests of general relativity, but our low proper motion suggests that it exists only at the $\sim 1\%$ level. Our improved accuracy and precision for the pulsar’s transverse velocity is also important for understanding the formation of this system.

Our distance measurement is model-dependent, with the accuracy determined partially by the relationship between the measured scintillation bandwidth and the spatial scale, for which we have used the method of Cordes and Rickett (1998). In the near future, wide-bandwidth observing systems, such as that of the MeerKAT radio telescope or the ultra-wideband low-frequency (UWL) receiver for the Parkes 64 m radio telescope, could provide an experimentally-derived relationship through analysis of the frequency-dependence on the scintillation bandwidth. In this way scintillation studies may be used to give distance measurements to pulsars with predictable modulation of the scintillation timescale (e.g. relativistic binaries). We may also soon see an independent proper motion measurement from pulsar timing with improved observation span and techniques, which would provide another method for determining the pulsar distance in combination with our velocity measurements. Future high-quality observations, with the sensitivity and resolution to estimate the scattering anisotropy from the autocovariance functions of individual dynamic spectra will be valuable for improving the accuracy and precision of future inclination angle measurements of scintillating binary pulsars.

PRECISE SCINTILLOMETRY WITH ANNUAL AND ORBITAL VARIATIONS IN SCINTILLATION ARCS FOR PSR J0437–4715

Intensity scintillations of radio pulsars are known to originate from interference between waves scattered by electron density irregularities in the ionised interstellar medium (IISM; see Section 1.3.3 of Chapter 1). Parabolic arcs are observed in a power spectrum of these scintillations when interference occurs between waves from scattered images of the source that extend beyond the root-mean-square scattering angle. The curvature of these arcs is related to the distance of the scattering plasma and its velocity with respect to the line-of-sight, as well as the orientation of the scattered image. In this Chapter we report the observation of annual and orbital variations in the curvature of scintillation arcs over a period of ~ 13 years for the bright PPTA millisecond pulsar, PSR J0437–4715. These variations are the signature of the relative transverse motions of the Earth, pulsar, and IISM, which we model to obtain precise measurements of parameters of the pulsar’s binary orbit and the scattering medium itself. We observe two clear scintillation arcs in most observations and we show that they originate from anisotropic scattering in thin screens located at distances $D_e = 90.6 \pm 0.7$ pc and $D_e = 122 \pm 3$ pc from the Earth, along the line-of-sight to the pulsar. From our velocity model of the "primary" (brightest) arc we measure the pulsar’s longitude of ascending node to be $\Omega = 207.2 \pm 0.7^\circ$, which is consistent with the best timing model (Chapter 2), but surpasses its precision despite this being one of the most precisely-timed pulsars. This is the highest-precision kinematic model achieved with scintillometry for any pulsar, demonstrating that scintillation arcs can potentially be a superior method to the traditional diffractive scintillation modelling for pulsars with sharp scintillation arcs, or pulsars in weak scattering. In such cases, this technique can be used in combination with pulsar timing to determine the full three-dimensional orbital geometries of binary pulsars.

4.1 Introduction

PSR J0437–4715 is the nearest and brightest millisecond radio pulsar. For this reason it is a key pulsar in the Parkes Pulsar Timing Array (PPTA; Manchester et al. 2013) and International Pulsar Timing Array (IPTA; Hobbs et al. 2010) projects that monitor the arrival times of a set of millisecond pulsars over many years. The main goal of these projects is the eventual detection and study of nanohertz-frequency gravitational waves. In order to achieve this detection, it is necessary to understand in detail the effects of the IISM in the pulsar observations, including dispersion measure (DM) variations (changes to the total column density of electrons along the line-of-sight) and scattering (see Section 1.3 of Chapter 1 for an introduction).

These effects are observed with all radio observations of compact sources at centimetre to metre wavelengths, but pulsars are a particularly sensitive tool for their study. Even small fractional changes to the DM with time, on the order of one part in 10^4 , are a significant source of noise in pulsar timing data, but can also be a tool for identifying compact and dense structures in the IISM such as extreme scattering events (ESEs) (e.g. Coles et al. 2015; Bannister et al. 2016).

These spatial fluctuations in the IISM electron density, which have a power-law distribution of sizes and densities originating from turbulence, also scatter incident wavefronts and produce an interference pattern of intensity variations at the observatory. As the interference pattern drifts across the telescope with velocity \mathbf{V}_{los} , the intensity varies as a function of time and frequency of the radio waves. These variations are insightfully captured in the dynamic spectrum (see left panels of Figure 4.1), and our observations of these are described in Section 4.2. Diffractive angular scattering is the dominant mechanism for the scintillations in pulsar observations (Rickett 1969), but the resulting diffractive scintillations with spatial scale s_d are modulated by refractive scintillations that occur on larger spatial and angular scales (Rickett et al. 1984). The timescale of the diffractive scintillations τ_d (of order minutes) is inversely proportional to the line-of-sight velocity $\mathbf{V}_{\text{los}} = s_d/\tau_d$, while the frequency scale $\Delta\nu_d$ (of order MHz for strong scattering) is the frequency change over which the interference pattern becomes decorrelated. This decorrelation occurs because the strength of scattering, and therefore the interference pattern, is frequency-dependent. In Chapter 3 we measured these diffractive scintillations to model the transverse velocity for PSR J1141–6545.

The multi-path propagation of a scattered pulse extends it into a quasi-exponential shape with a characteristic decay rate determined by the scattering timescale τ_s . The latter is inversely proportional to $\Delta\nu_d$ because the two quantities are related through the "uncertainty relation" $2\pi\Delta\nu_d\tau_s \approx 1$ (Rickett 1977). Since PSR J0437–4715 is near to the Earth at $D = 156.79$ pc (Chapter 2), it has a relatively small DM and τ_s , and thus the scintillation bandwidth $\Delta\nu_d$ is relatively large. It is often observed in the regime of weak scattering where $\Delta\nu_d > f$, for observing frequency f above a "transition frequency" of $f_t \sim 1$ GHz. Below the transition frequency, $f < f_t$, the pulsar is observed in strong scattering with $\Delta\nu_d < f$, which is common for most pulsar observations.

Pulsar observers often work with the secondary spectrum, which is the two-dimensional




Figure 4.1: Dynamic spectra (left panels) and corresponding secondary spectra (right panels) for observations of PSR J0437–4715 from the Parkes 64 m radio telescope from the “20 cm” observing band (top panels) and “40 cm” observing band (bottom panels). The 20 cm dynamic spectrum is a combination of six consecutive observations, which have been concatenated with linear-interpolation applied to the gaps (as described in Section 4.2.1), while the 40 cm dynamic spectrum is from a single observation. The colour in the dynamic spectra shows the normalised flux after subtracting the mean, while the colour in the secondary spectrum shows the log power for these dynamic spectra, which saturates at 2.5 dex above the mean to optimise visualisation of the arcs.

power spectrum of the dynamic spectrum (Figure 4.1, right panels). This may also be referred to as a “delay-Doppler” distribution because the variables on the plotted axes are the differential time delay τ_{del} and differential Doppler shift f_{dop} between the interfering waves from a scattered image of the pulsar and the main image (Cordes et al. 2006). The nature of τ_{del} is understood from the Fourier relationship between scattered pulse broadening and the decorrelation over frequency of the interference pattern, which leads to the uncertainty relation given above. The Fourier analysis of scintillations in time gives f_{dop} , which is proportional to the spatial wavenumber vector $k\theta$ (for arrival angle θ) and the effective velocity (see Section 4.3.1 for more detail). Parabolic arcs in a pulsar secondary spectrum were first observed by Stinebring et al. (2001), and the origin of their shape was explained by Walker et al. (2004) to be a result of the common dependence on scattering angle for the Fourier components, $f_{\text{dop}} \propto (\theta)_0$ and $\tau_{\text{del}} \propto (\theta)_0^2$, where $(\theta)_0$ is the angle between the interfering scattered waves.

The curvature of the arcs is related to several variables: the distance to the scattering screen, the “effective velocity” of the line-of-sight relative to the medium (which is a linear combination of the transverse motions of the Earth, pulsar, and IISM), and the angle between the effective velocity vector and the scattered image in the case of anisotropic scattering. Arcs are clearest when the scattering region is compact on the line-of-sight (e.g. a “thin screen”), and are enhanced when the scattered image is anisotropic and extended in the direction of the effective velocity vector. The secondary spectrum and arc curvature are explained further in Section 4.3, and a detailed explanation can be found in Walker et al. (2004) and Cordes et al. (2006).

With multiple measurements of parabolic arcs at different epochs, we would expect to observe cyclical variations corresponding to the orbital motions of the Earth and the pulsar. However earlier analyses of scintillation arcs have primarily involved solitary pulsars without binary motions (e.g. Stinebring et al. 2001), and individual epochs of observations (e.g. Briskin et al. 2010; Bhat et al. 2016). Only one previous example of annual and orbital velocity modulations to arc curvature has been reported, which was for an analysis of the double pulsar, PSR J0737–3039 (Stinebring et al. 2005). This arc curvature model was inferior to the earlier diffractive scintillation model of Ransom et al. (2004) for PSR J0737–3039 because the arcs were not as sharp as those we present here.

Scintillation arcs have previously been observed for PSR J0437–4715, in observations from the Parkes radio telescope and Murchison Wide-field Array (MWA; Bhat et al. 2016). However with only two observations of the brightest arc, Bhat et al. (2016) were only able to estimate the distance to one of the screens. We show that attempting to measure the scattering screen distance from individual observations in this way, with the necessarily restrictive assumptions of a stationary IISM and an isotropic scattered image, can result in significantly biased measurements. In this chapter, we show that the modelling of the long-term behavior of the scintillation arcs can provide a robust screen distance measurement, which is constrained by the relative amplitudes of the arc curvature modulation due to the pulsar and Earth velocities (as with the screen distance

for PSR J1141–6545 in Chapter 3). It can also provide precise measurements of the IISM velocity and anisotropy, and useful parameters of the pulsar ephemeris such as the orbital inclination angle i , and the longitude of the ascending node Ω . This is important for studies of relativistic binary pulsars because these parameters can be difficult to measure through pulsar timing alone, but they are useful in tests of general relativity (as discussed in Chapter 3). We found in Chapter 3 that modelling of the long-term changes to the diffractive scintillations can also be used to precisely measure these parameters (see also Rickett et al. 2014), but the work in this Chapter is an independent approach that can work well even in the regime of weak scattering.

PSR J0437–4715 was the first pulsar to have its full three-dimensional orbital motion determined (van Straten et al. 2001), and it also has the most precisely measured distance of any pulsar, $D = 156.79 \pm 0.25$ pc (Chapter 2). These facts make this system an ideal candidate for modelling of the transverse motion apparent in scintillations, and allow us to determine the distance to the scattering screens with unprecedented precision. In Section 4.4 we describe the model for arc curvature variations before presenting the results in Section 4.5. We identify two clear arcs corresponding to two discrete scattering screens, and we are able to model the long-term arc curvature modulations for both. The models show that the scattering is likely anisotropic, which is a common finding through scintillation phenomena (e.g. Brisken et al. 2010). The results are discussed in Section 4.6, and in Section 4.6.2 we give some suggestions for candidate structures in the IISM responsible for the scattering. Recently, Walker et al. (2017) found that the anisotropic plasma structures that cause intra-day variability in the flux of two radio quasars is associated with hot stars with small impact parameters for the line-of-sight to the quasars. In Section 4.6.2.1 we test this conclusion by searching for associations with stars that have small impact parameters, and measured distances and kinematics that are consistent with our precise scattering screen models.

4.2 Observations and data

Our observations of PSR J0437–4715 are from the Parkes 64 m radio telescope and span 12.7 years from MJD 52620 to 57670, with the majority being taken as part of the PPTA program (Manchester et al. 2013) that commenced regular observations in 2005 (MJD 53430). The PPTA observes a set of millisecond pulsars approximately every two weeks in three observing bands, 40/50 cm (at centre frequencies $f_c \sim 685$ MHz and $f_c \sim 732$ MHz respectively), 20 cm ($f_c \sim 1400$ MHz), and 10 cm ($f_c \sim 3100$ MHz). Here we use observations from the 40/50 cm and 20 cm bands because the 10 cm observations do not show clear scintillation arcs in the secondary spectra. This is likely because these observations are too far above the transition frequency and because the curvature of any arcs would be quite large and near to the noise and scattered power close to the $f_{\text{dop}} = 0$ axis of the secondary spectrum. The lower-frequency observations prior to mid-2009 were taken with the “10/50 cm” receiver in the 50 cm band and with an observing band-

width $B \sim 64$ MHz. After mid-2009, the observing band was shifted to the 40 cm band with the same receiver and bandwidth to avoid some significant terrestrial radio-frequency interference (RFI) in the 50 cm band. The 20 cm observations were captured either with the central beam of the “multibeam” receiver with a typical centre frequency of $f_c \sim 1369$ MHz and bandwidth $B \sim 300$ MHz, or the “H-OH” receiver with typical values $f_c \sim 1433$ MHz and $B \sim 600$ MHz.

PSR J0437–4715 is an important target for the PPTA and IPTA because it is not included in either of the Northern PTA counterparts, and its high flux density means observations yield extremely precise times of arrival (TOAs). It is out of the Galactic plane unlike most other pulsars, meaning that there are fewer other targets to observe when it is in the sky. Consequently, this pulsar is often selected for repeated observations during the PPTA observing program (archival data is found under the “P456” observing code through the CSIRO Data Access Portal: <https://data.csiro.au/dap/>). These additional observations give valuable TOAs that increases the sensitivity of the PPTA to gravitational waves, but simultaneously provides us with the opportunity to extend the dynamic spectrum in time by concatenating consecutive observations, thus improving our signal-to-noise ratio of any scintillation arcs in the secondary spectra. In addition, PSR J0437–4715 is highly linearly-polarized and for this reason it is the target of a separate observing campaign (observing code “P737”), which tracks the pulsar (a few times per year) for up to ~ 10 hours from rise to set for the purpose of polarization calibration and instrument commissioning.

These consecutive observations give particularly valuable dynamic spectra when concatenated, such as the 20 cm observation shown in the top-left panel of Figure 4.1, with a well-defined scintillation arc apparent in the secondary spectrum. However, while these observations give very clear scintillation arcs, and often show multiple arcs, we actually see the primary arc in all available observations in the 20 cm and 50 cm bands, provided they are not too contaminated with RFI. Below we describe the formation of the dynamic and secondary spectra from these observations.

4.2.1 Computing dynamic and secondary spectra

The dynamic spectra, $S(t, f)$, are computed using a procedure in the data processing pipeline that has been developed for the PPTA data release 2 (dr2). This uses the PSRCHIVE (Hotan et al. 2004) package and will be described in detail in Kerr et al. (in prep.). Dynamic spectra for the relativistic binary pulsar, PSR J1141–6545 that were computed with this same pipeline, were analysed in Chapter 3 for the purpose of modelling the long-term changes to the scintillation velocity derived from diffractive scintillations. In brief, observations of PSR J0437–4715 are preceded and followed by observations of a pulsed noise diode which excites the X and Y polarizations for measuring and correcting the complex gain. This noise diode is calibrated to absolute flux density using observations of the bright radio galaxy Hydra A, which are taken usually in the same observing session (i.e. within a day or two of the pulsar observation). As mentioned above, the

P737 observing campaign monitors PSR J0437–4715 for the purpose of polarisation calibration, and these observations are used in combination with the noise diode for polarization calibration of all other observations (e.g. from the P456 project). However for computing the dynamic spectra we are only concerned with the total intensity (Stokes I).

For each subintegration and frequency channel in an observation, an analytical model of the Stokes I pulse profile is fitted with linear regression to the observed pulse profile. The result of the fit is the pulsar flux density and its uncertainty because of the prior calibration. In this way, we obtain the dynamic spectrum of flux density as a function of frequency and time (left panels of Figure 4.1), with typical resolutions of order $B_c \sim 0.5$ MHz (channel bandwidth) and $t_{\text{sub}} \sim 30$ s (subintegration time) respectively. In the case of multiple consecutive observations of the pulsar, we concatenate the dynamic spectra to increase the effective observing time, which is particularly valuable for the 20 cm observations because of the larger spatial and temporal scales of the scintles. The gaps in time between individual observations (during which the noise diode is observed) is filled with linear interpolation. We find that this does not affect the curvature of the scintillation arcs, but it does reduce the scattered power near the axes of the secondary spectrum, thus improving our arc signal-to-noise ratio (we experimented with filling the gaps instead with the mean flux density, rather than linearly interpolating, and saw similar results). Each dynamic spectrum is also analysed for RFI by simply searching for channels and subintegrations that have outlying flux densities. Anything in the dynamic spectrum that is flagged in this way is re-filled with a linear interpolation from the neighboring clean channels/subintegrations. The edited dynamic spectra are then checked manually and any significantly contaminated spectra (e.g. with residual RFI or large regions that have been flagged as bad and require re-filling) are excluded from our analysis.

Once the consecutive dynamic spectra are concatenated and corrected for RFI, we correct the new dynamic spectrum by the measured bandpass of flux density, subtract the mean flux, and finally normalise before computing the secondary spectrum, $P(f_{\text{dop}}, \tau_{\text{del}})$, which is its two-dimensional Fourier transform. This is computed by first pre-whitening the dynamic spectrum (by means of a first-difference filter, e.g. Jenkins and Watts 1969; Coles et al. 2011) before it is Fourier transformed with zero padding. We take the squared magnitude of the transform, and shift and crop it to show only values for $t_{\text{del}} > 0$. The spectrum is then "post-darkened" (the reverse process of pre-whitening), normalised, and finally we take the logarithm of the result because the power in the noise is log-normally distributed. So in summary we have $P(f_{\text{dop}}, \tau_{\text{del}}) = \log(|\tilde{S}(t, f)|^2)$, where the tilde denotes the two-dimensional Fourier transform, f_{dop} and τ_{del} are the Fourier conjugates of t and f respectively, and $S(t, f)$ is the mean-subtracted and normalised dynamic spectrum. Two examples of the secondary spectra computed in this manner are shown in the right panels of Figure 4.1. The total observing time and receiver bandwidth of the dynamic spectrum define the range in the transformed secondary spectrum, but we crop the secondary spectrum for improved arc visualisation and curvature uncertainty estimation (see Section 4.3.1). In the

following section we describe some properties of the secondary spectrum and the method for measuring the curvature of the arcs.

4.3 Interpreting and fitting the secondary spectra

We have so far introduced the secondary spectra as being the power spectrum of intensity scintillations with axes corresponding to the differential geometric time delay τ_{del} and differential Doppler shift f_{dop} between interfering waves. Here we provide some basic theory to understand this spectrum and the parabolic arcs in more detail (see also Walker et al. 2004; Cordes et al. 2006). We then also describe the method we use for fitting the parabolic arcs and estimating the uncertainty in the curvature parameter (which follows Bhat et al. 2016).

A power-law distribution of density irregularities in the IISM scatters incident radiation, by means of diffraction, into an angular spectrum relative to the direct line-of-sight to the source. Consider the interference of waves arriving at the observatory from just two small angles in this spectrum, θ_1 and θ_2 . The result is a single frequency-dependent interference fringe pattern, which is sampled by the observatory in time and frequency to be observed as a sinusoid in the dynamic spectrum (Cordes et al. 2006). For a compact scattering region, or a so-called “thin screen,” at some fractional position s along the line of sight from the source (i.e. from $s = 0$ at the source to $s = 1$ at observatory), the axes of the secondary spectrum are related to these scattering angles with

$$(4.1) \quad \tau_{\text{del}} = \frac{D(1-s)}{2cs}(\theta_2^2 - \theta_1^2)$$

$$(4.2) \quad f_{\text{dop}} = \frac{f_c}{cs} \mathbf{V}_{\text{eff}} \cdot (\boldsymbol{\theta}_2 - \boldsymbol{\theta}_1),$$

where D is the distance to the source from the observatory, \mathbf{V}_{eff} is the velocity of the line-of-sight through the screen (a linear combination of Earth, IISM, and pulsar velocities which is described in Section 4.4), f_c is the centre frequency of the observation, and c is the speed of light. As mentioned previously, τ_{del} is the differential geometric time delay between the paths taken to arrive from the two angles, and f_{dop} is their differential Doppler shift, but is equally the “fringe rate” in the dynamic spectrum. Each Fourier component in the secondary spectrum corresponds to one of the sinusoidal fringe patterns and thus to the summation of all pairs of components of the angular spectrum with the appropriate τ_{del} and f_{dop} .

In the case of weak scattering (e.g. PSR J0437–4715 observations in the 20 cm band), the origin of a parabolic arc in the secondary spectrum can be understood by considering the interference between the unscattered main image of the pulsar (centred on the line-of-sight), and the scattered image. In this case we have Equations 4.1 and 4.2 with $\theta_1 = 0$ and $\theta_2 = \theta_0$, where θ_0 is now the position angle to a component of the scattered image. The arc comes from the parabolic

relationship between delay and Doppler shift through their dependence on θ_0 , and we define the curvature parameter η with $\tau_{\text{del}} = \eta f_{\text{dop}}^2$, and this is then given by

$$(4.3) \quad \eta = \frac{cDs(1-s)}{2f_c^2 \mathbf{V}_{\text{eff}}^2 \cos^2 \psi}$$

where ψ is the angle between \mathbf{V}_{eff} and the position vector along the anisotropy in the scattered image and comes from the dot product in Equation 4.2. For an isotropically-distributed image such as a ring or halo, the equation is the same but $\cos \psi$ has any value in the range $[-1, 1]$ for individual points in the image, giving a maximum extent of $\cos \psi = 1$ in the secondary spectrum (Cordes et al. 2006). In the case of strong scattering (such as in 40/50 cm observations of PSR J0437–4715), the argument is similar but instead a scattered image that extends beyond the root-mean-square scattering angle interferes with the main image.

Additional phenomena that can be observed in the secondary spectra include: arc asymmetries about $f_{\text{dop}} = 0$ formed by asymmetrical scattering about the line-of-sight in the direction of \mathbf{V}_{eff} (e.g. Cordes et al. 2006), and inverted “arclets” on or near to the main arc, which are formed by discrete multiple images that interfere with each other and the main image (e.g. Briskin et al. 2010). We do not observe any clearly asymmetrical arcs (from an asymmetrical image distribution) at any epoch or orientation of the \mathbf{V}_{eff} vector. This may indicate that the scattered image is not only symmetrical, but also continuous, such as a linear structure (for anisotropic scattering) or a circularly-symmetric halo (for isotropic scattering). Since \mathbf{V}_{eff} is dominated by the pulsar’s high proper motion, there are only small changes to the orientation of the vector with time, meaning that mostly symmetrical arcs could also be observed if any discontinuities in the scattered image were aligned perpendicularly to the mean direction of \mathbf{V}_{eff} . However in this case we would expect to see arcs that are discontinuous in τ_{del} , which we do not. We also do not observe any inverted arclets, which could also be expected from discontinuous/inhomogeneous scattering that appears as multiple images. Although this would also only be expected to appear in strong scattering such as in the 40/50 cm band, but these observations in general are less sensitive than the 20 cm observations. This is because when the scattered radiation is weaker than the main image, the scintillations show interference of waves on the direct path with itself and the scattered image, rather than multiple scattered images interfering with each other (Walker et al. 2004).

The power distribution in the secondary spectra itself can also give us some insight into the appropriate scattering model to apply. The relatively thin, symmetrical, and continuous arcs that we observe (Figure 4.1, right panels) are indicative of a scattered image that is continuous and symmetrical about a line perpendicular to \mathbf{V}_{eff} for an isotropic image, or about the direct line-of-sight for a one-dimensional anisotropic image. The decay of power in the arcs away from the $\tau_{\text{del}} = 0$ axis is simply due to the decline in intensity with increasing scattering angle. In addition, the thin arcs suggest a compact scattering region (thin screen) and the appearance of

multiple arcs in the highest signal-to-noise secondary spectra suggest multiple discrete scattering screens distributed along the line-of-sight, each with its own \mathbf{V}_{eff} . In Section 4.4 we describe the models for \mathbf{V}_{eff} and arc curvature, which includes models for both isotropic and anisotropic scattered images.

4.3.1 Fitting arc curvature

To measure the arc curvature and its uncertainty from the secondary spectrum of each observation, we follow the method of Bhat et al. (2016) in their analysis of Murchison Widefield Array (MWA) and Parkes observations of PSR J0437–4715. For a set of trial curvatures η_i , we find the power along the parabola $\tau_{\text{del}} = \eta_i f_{\text{dop}}^2$ by summing the linearly-interpolated power between pixels neighboring $f_{\text{dop},i} = \pm \sqrt{\tau_{\text{del}}/\eta_i}$ for each τ_{del} in the secondary spectrum $P(f_{\text{dop}}, \tau_{\text{del}})$. We then divide by twice the number of τ_{del} rows in the secondary spectrum N , to give the mean power per pixel as a function of curvature η

$$(4.4) \quad P_{\text{arc}}(\eta) = \frac{1}{2N} \sum_{i=1}^N P(f_{\text{dop},i}, \tau_{\text{del},i}).$$

The arcs in the secondary spectrum then correspond to different maxima in this function. Using this method, the secondary spectrum must be cropped (or truncated) at a τ_{del} value that is ideally just beyond the arc that is being measured, so that the mean power in the arc is not averaged down by including noise. This is difficult because the power in the arcs decays with increasing τ_{del} . We chose to crop each secondary spectrum at a fixed maximum time delay, $\tau_{\text{del,max}}$ beyond which most observations show no evidence of the primary arc. This delay depends on the observing frequency for the observation, and we have therefore defined $\tau_{\text{del,max}} = 0.25 \mu\text{s} \times (1400 \text{ MHz}/f)^2$ for the primary arc and $\tau_{\text{del,max}} = 0.1 \mu\text{s} \times (1400 \text{ MHz}/f)^2$ for fitting the secondary. In general, this is a conservative figure for the primary arc and our estimated uncertainties for the arc curvature are expected to be slightly overestimated because the measured power in the primary arc is averaged down as the power decays towards $\tau_{\text{del,max}}$, such as in Figure 4.2. We measured $P_{\text{arc}}(\eta)$ from these cropped secondary spectra with a set of trial curvatures spaced equally in $\sqrt{\eta}$ with a resolution of $0.001 \text{ s}^{3/2}$.

The curvature for the primary arc is found by using a three-point smoothing window on our measured $P_{\text{arc}}(\eta)$ (to reduce the noise), before finding the maximum. To estimate the uncertainty of this measurement (again following Bhat et al. 2016), we first determine the average and standard deviation of the noise far from the power in the arcs and the τ_{del} and f_{dop} axes, $\langle P_{\text{noise}} \rangle$ and σ_{noise} respectively. We use the region of the original (uncropped) secondary spectrum in the upper half of the τ_{del} range and upper half of $|f_{\text{dop}}|$ to measure $\langle P_{\text{noise}} \rangle$ and σ_{noise} for each observation. The uncertainty in the arc curvature measurement is then given by half of the domain around the measurement for which $P_{\text{arc}}(\eta)$ is within one σ_{noise} of the maximum. This curvature measurement for the primary and secondary arcs, the measured $\langle P_{\text{noise}} \rangle$ and σ_{noise} ,

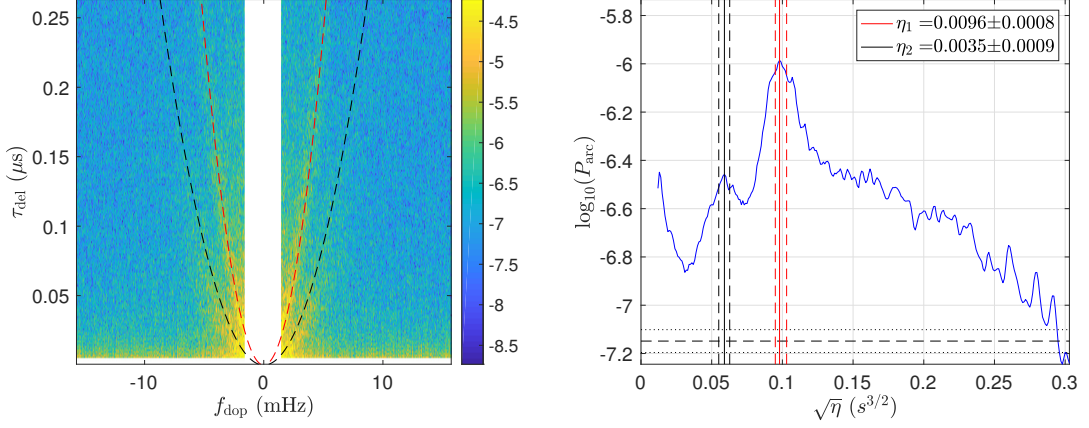


Figure 4.2: Cropped secondary spectrum (left) and corresponding $P_{\text{arc}}(\eta)$ curve (in $\log_{10}(P_{\text{arc}})$ and $\sqrt{\eta}$) for a 20 cm observation on MJD 56946. The colour scale for the secondary spectrum is as described for Figure 4.1 and the white regions were discarded from the analysis since they contain scattered power. The red solid vertical line and dashed lines on the $P_{\text{arc}}(\eta)$ curve show the primary arc curvature measurement and error region, while the black vertical lines are for the secondary arc. The error was calculated from the mean and standard deviation power of the noise in the secondary spectrum, which is shown with the horizontal solid and dotted lines respectively. The measured curvature for the primary and secondary arcs are shown in secondary spectrum with red and black dashed lines respectively.

and the corresponding uncertainty range for the curvature, are shown for one observation in Figure 4.2.

If $P_{\text{arc}}(\eta)$ clearly shows the second maximum corresponding to the secondary arc (as in Figure 4.2), we find the curvature corresponding to the local maximum power after using a five-point smoothing window on $P_{\text{arc}}(\eta)$. The secondary arc is defined as the highest curvature peak below the measured primary arc with a prominence of at least $1\sigma_{\text{noise}}$ (in the smoothed $P_{\text{arc}}(\eta)$ curve). The uncertainty for the curvature in this secondary arc is then estimated in the same manner as for the primary arc. However, the uncertainty may be underestimated because minor arcs are often difficult to separate from the scattered power near the axes and the power in the primary arc.

4.4 Modelling arc curvature variations

The curvature of arcs in the secondary spectrum depends on the distance to the scattering region s (assumed to be a thin screen), the velocity of the line-of-sight with respect to the medium at this distance \mathbf{V}_{eff} , and the orientation of the scattered image (if anisotropic) with respect to this velocity vector ψ , as given in Equation 4.3. We therefore expect the curvature to be time-dependent as \mathbf{V}_{eff} changes because of the changing transverse components of the Earth's velocity (\mathbf{V}_{E}) and the pulsar's binary velocity (\mathbf{V}_{p}), where both are relative to the Solar system

barycentre. The effective velocity is a linear combination of these velocities and the velocity of the medium itself (\mathbf{V}_{ISM}),

$$(4.5) \quad \mathbf{V}_{\text{eff}} = (1-s)(\mathbf{V}_p + \mathbf{V}_\mu) + s\mathbf{V}_E - \mathbf{V}_{\text{ISM}},$$

where s is defined as a fractional distance from $s = 0$ at the pulsar to $s = 1$ at the Earth, and \mathbf{V}_μ is the constant transverse velocity of the pulsar system (corresponding to its proper motion).

The curvature of the arcs is also radio frequency-dependent, $\eta \propto f_c^{-2}$. To simultaneously model arcs from observations at different frequencies, we define the “equivalent curvature”

$$(4.6) \quad \eta^* = \eta \left(\frac{f_c}{1400 \text{ MHz}} \right)^2$$

to be the arc curvature scaled to a reference frequency of 1400 MHz.

We model the variations in $\eta^*(t)$ with effective velocity components in right ascension (α) and declination (δ)

$$(4.7) \quad \begin{aligned} v_{\text{eff},\alpha} &= (1-s)(v_{p,\alpha} + v_{\mu,\alpha}) + sv_{E,\alpha} - v_{\text{ISM},\alpha} \\ v_{\text{eff},\delta} &= (1-s)(v_{p,\delta} + v_{\mu,\delta}) + sv_{E,\delta} - v_{\text{ISM},\delta} \\ \mathbf{V}_{\text{eff}} &= \sqrt{v_{\text{eff},\alpha}^2 + v_{\text{eff},\delta}^2}. \end{aligned}$$

The distance and proper motion for PSR J0437–4715 is known to high precision from pulsar timing, giving $v_{p,\alpha} = 90.26 \text{ km s}^{-1}$ and $v_{\mu,\delta} = -53.13 \text{ km s}^{-1}$ (Chapter 2). The precise timing model also allows us to derive the mean orbital velocity

$$(4.8) \quad \mathbf{V}_0 = \frac{2\pi xc}{\sin i P_b \sqrt{1-e^2}}$$

from the projected semi-major axis x (light-seconds), orbital period P_b , eccentricity e , and the inclination angle i . The transverse velocity is then defined in terms of the true orbital anomaly θ

$$(4.9) \quad \begin{aligned} v_{p,\parallel} &= -\mathbf{V}_0 (e \sin \omega + \sin(\theta + \omega)) \\ v_{p,\perp} &= \mathbf{V}_0 \cos i (e \cos \omega + \cos(\theta + \omega)), \end{aligned}$$

These are rotated into right ascension α and declination δ with the longitude of the ascending node Ω , defined East of North. This is the most uncertain parameter in the velocity model with $\Omega = 207 \pm 1.2^\circ$ because it is included into the timing model through a subtle kinematic effect caused by the pulsar’s proper motion changing the projection of the orbit. Remarkably, by modelling the annual and orbital modulation of diffractive scintillations (e.g. Rickett et al. 2014, and Chapter 3) or arc curvature, we are able to measure Ω and other parameters (such as i) often with higher precision than through timing model because they depend primarily on the transverse velocity. From Equations 4.7 and 4.9 we then see that the only parameters required for modelling the effective velocity are: s , $v_{\text{ISM},\alpha}$, and $v_{\text{ISM},\delta}$, although we also fit for Ω to improve the current measurement.

The variations in arc curvature follow from this velocity model through Equation 4.3. For an isotropic scattered image the parameters described above complete the model, while for a highly anisotropic model we must also consider the angle between \mathbf{V}_{eff} and the anisotropy ψ . For this work we consider both isotropic and anisotropic models and compare their goodness-of-fit to the arc curvature variations, and discuss the power distribution in the secondary spectra. Our anisotropic model is a one-dimensional approximation for highly-anisotropic scattering, originating for example from thin plasma filaments. In such a scenario, the data are only sensitive to the velocity parallel to the direction of the anisotropy ψ . If this anisotropy is caused by plasma filaments, then the arcs are only sensitive to the component of the effective velocity that is perpendicular to the filaments. Therefore instead of the two IISM velocity components in celestial coordinates, we instead fit for a single velocity component parallel to the anisotropy, $v_{\text{IISM},\psi}$. We then use $v_{\text{IISM},\alpha} = v_{\text{IISM},\psi} \sin \psi$ and $v_{\text{IISM},\delta} = v_{\text{IISM},\psi} \cos \psi$.

We fit both isotropic and anisotropic models to the arc curvature variations measured in the primary and secondary arcs. We use weighted nonlinear least squares regression and present the results in the following section. The models are fitted using the same procedure as described in Section 3.4 of Chapter 3 and we similarly follow the description therein of basic model comparison with the Akaike information criterion (AIC; Akaike 1998). However, for the anisotropic model of the primary arc, which we find is kinematically favoured, we also use a Markov Chain Monte Carlo (MCMC) sampler to analyse the full posterior probability distribution for the model parameters. This was done in order to confirm that the probability density functions (PDF) were near Gaussian, which is the assumption we make to simplify the calculation of chance associations with known stars (Section 4.6.2.1). The algorithm we use is described in Veitch and Vecchio (2010) and was implemented in MATLAB by Pitkin and Romano (2013). We use uniform priors on all parameters and find that the MCMC analysis gives results in agreement with the least squares fit, which is unsurprising because the residuals are normally distributed and the parameters are well-described by a multi-variate Gaussian function (Section 4.5.1).

4.5 Results

We have measured the curvature of scintillation arcs in a set of observations across two observing bands for PSR J0437–4715. We have found that all observations in the 20 cm and 40/50 cm bands show evidence for at least one arc, which is the strongest in all of the observations and we refer to as the "primary arc". The arc signal-to-noise ratio depends strongly on the pulsar flux, observation length t_{obs} , and observing bandwidth B . In observations with the highest signal-to-noise ratio, we see a fainter, "secondary arc" at a lower equivalent curvature η^* .

The time series of η^* for each measurement of the primary scintillation arc is shown in Figure 4.3. There is a clear annual modulation to the curvatures, as well as a ~ 5.7 day modulation corresponding to the pulsar's orbital period. We find similar annual and binary orbital modulation

Table 4.1: Parameters for isotropic and anisotropic scattering models for the curvature variations in the primary and secondary arcs for PSR J0437–4715. Parameters were measured with a weighted nonlinear least-squares analysis.

	Primary arc		Secondary arc	
	Isotropic	Anisotropic	Isotropic	Anisotropic
s	0.420(5)	0.422(4)	0.21(3)	0.218(14)
Ω ($^\circ$)	206.7(8)	207.2(7)	(207 $^\circ$) [†]	(207 $^\circ$) [†]
$v_{\text{IISM},\alpha}$ (km s $^{-1}$)	−10.9(8)	–	−4(9)	–
$v_{\text{IISM},\delta}$ (km s $^{-1}$)	31.7(7)	–	47(8)	–
$v_{\text{IISM},\psi}$ (km s $^{-1}$)	–	−31.2(5)	–	−41(5)
ψ ($^\circ$)	–	135.4(5)	–	141(2)
χ^2	734	705	2462	2440
χ_r^2	0.547	0.526	3.354	3.324
† Fixed value				

for the secondary arc (Figure 4.5), although the binary orbital modulation dominates for these because the scattering region is located closer to the pulsar on the line-of-sight. We have modelled the variations in both the primary and secondary arcs with both the isotropic and anisotropic scattering models described in Section 4.4. Table 4.1 gives the fitted parameters for each of these models, which are discussed further in the following sections.

We see weak evidence for additional faint arcs; one at a lower curvature and one at a higher curvature in a few observations. However we are unable to reliably measure the curvature for these arcs in multiple observations to find curvature modulations. This is because the arc with smaller curvature is generally faint, near to the $\tau_{\text{del}} = 0$ axis, and near the power on the leading-edge of secondary and/or primary arc. The potential arc with higher curvature is hidden mostly within the power inside the primary arc and the scattered power along the $f_{\text{dop}} = 0$ axis. Future observations can be optimised for fitting any arcs with higher curvature by using dynamic spectra with long and continuous integration times.

4.5.1 Primary arc

The primary arc is seen in all observations in the 20 cm and 40/50 cm bands that are not contaminated with RFI, giving us 1347 total measurements of η . This is the highest number for any single object owing to the frequent and long-term observations of the PPTA.

The best-fit model for the variations in η^* for the primary arcs is shown in the bottom panels of Figure 4.3. We have subtracted the annual modulation given by the model fit, in order to show the modulation as a function of orbital phase (bottom left of Figure 4.3), and similarly subtracted the orbital modulation to show the annual variation as a function of day of year (bottom right of Figure 4.3). The parameters for both the isotropic and anisotropic models are shown in Table 4.1. The anisotropic model is favoured by the chi-squared goodness-of-fit test, with $\Delta\chi_c^2 \sim 55$ (see Section 3.4 of Chapter 3).

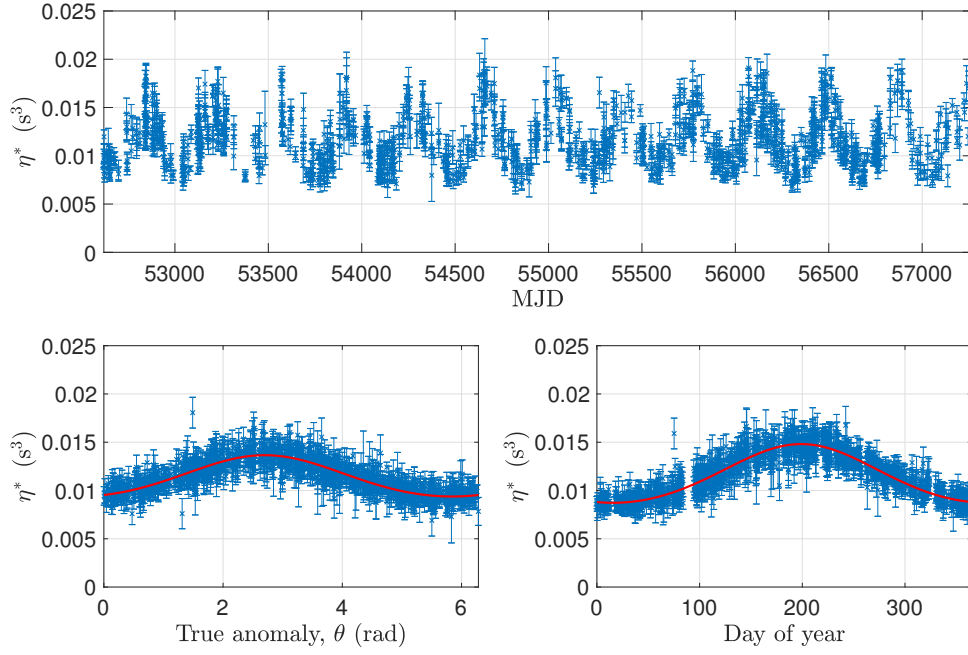


Figure 4.3: Top-panel: Equivalent curvature η^* (Equation 4.6) for each measurement of the primary scintillation arc across ~ 13 years of observations for PSR J0437–4715. Bottom panels: Equivalent curvatures with best-fit anisotropic model (solid line) as a function of orbital true anomaly θ (left; with fitted annual variation subtracted) and day-of-year (right; with fitted pulsar orbital variation subtracted). The parameter uncertainties have been scaled down by a factor of 1.4 because of the small χ_r^2 value (see Section 4.5.1).

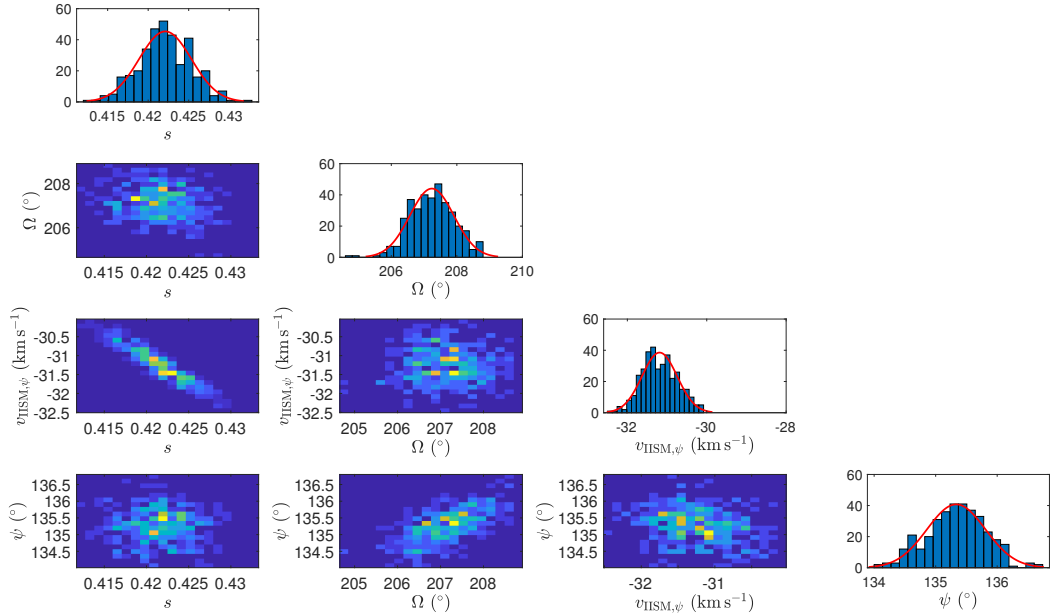


Figure 4.4: Posterior probability distributions for the parameters of the anisotropic model for the primary arc using the MCMC sampler of Veitch and Vecchio (2010), implemented in MATLAB by Pitkin and Romano (2013).

We also prefer the anisotropic model based on the power distribution across the arc. From Figure 4.2 it is clear that the power drops off just inside the arc and is somewhat symmetric about the peak. If the scattering were truly isotropic, we would expect to see power inside the arc "filled in" (e.g. Figure 10 of Cordes et al. 2006) such that the power across the arc in Figure 4.2 would be maximally asymmetric.

The χ_r^2 values for these models are low, suggesting that the uncertainties for the arc measurement were over-estimated by a factor of $\sim 1/\sqrt{\chi_r^2} = 1.4$ on average. We have scaled the measurement uncertainties shown in Figure 4.3 by this factor.

If we choose to also fit for the orbital inclination angle i with the anisotropic model, we measure $i = 136.1 \pm 0.5^\circ$, which differs from the timing solution of $i = 137.56 \pm 0.04^\circ$ (Chapter 2) by $\sim 3\sigma$, suggesting only a slight weakness in the model. The precision (and accuracy) of this measurement is impressive for scintillation studies, which can often be complicated by changes to properties of the scattering with time. The IISM velocity and anisotropy in the direction of PSR J0437–4715 remains incredibly stable over the ~ 13 years of our observations, meaning that the scattering properties of the screen are stable over a spatial scale of ~ 350 AU.

Finally, we have also repeated this fit with an MCMC analysis to sample the posterior probability distribution for the fitted parameters s , Ω , $v_{\text{IISM},\psi}$, and ψ . The result is shown in Figure 4.4 with a normal distribution fit to the PDF of each parameter. The mean and standard deviations are identical to the measurements from the least squares fit, with uncertainties taken from the parameter covariance matrix (Section 3.4, Chapter 3). We therefore use a normal distribution for the parameter probabilities in the calculation of chance stellar associations in Section 4.6.2.1.

4.5.2 Secondary arc

The secondary arc is observed in 739 of the highest signal-to-noise ratio observations (Figure 4.2). The measured curvature is less robust than that of the primary arc and as a consequence the data contains more noise. As with the primary arc, we have fitted isotropic and anisotropic models, which are given in Table 4.1. The anisotropic model is favoured, but only slightly, with $\Delta\chi_c^2 \sim 7$. This is not surprising because the two models are kinematically very similar, but once again the anisotropic model is also favoured because of the symmetric power distribution across the arc (Figure 4.2). The excess noise in the data from complications in fitting this secondary arc (see Section 4.3.1) is quantified by the large reduced chi-squared value, $\chi_r^2 = 3.324$, suggesting underestimated uncertainties by a factor of ~ 1.8 .

We have also repeated the fit with an MCMC analysis and find that the PDF for each parameter is well described by a normal distribution. The parameters are consistent with the least squares analysis, with only slightly larger uncertainties in two parameters: $s = 0.219 \pm 0.017$ and $v_{\parallel} = -41 \pm 6 \text{ km s}^{-1}$.

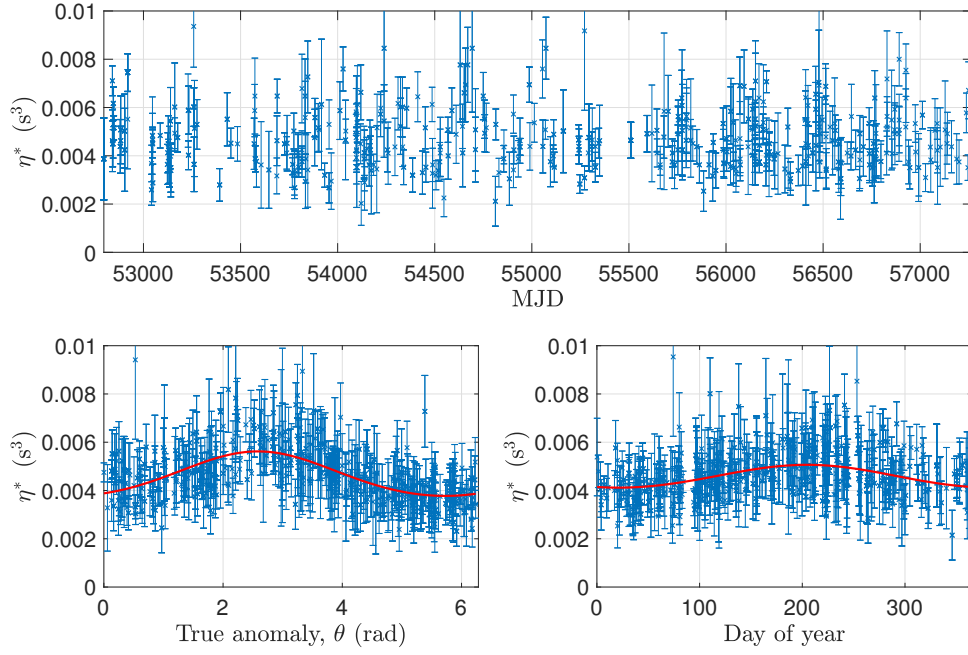


Figure 4.5: Top-panel: Equivalent curvature η^* (Equation 4.6) for each measurement of the secondary scintillation arc across ~ 13 years of observations for PSR J0437–4715. Bottom panels: Equivalent curvatures with best-fit anisotropic model (solid line) as a function of orbital true anomaly θ (left; with fitted annual variation subtracted) and day-of-year (right; with fitted pulsar orbital variation subtracted). The parameter uncertainties have been scaled up by a factor of 1.8 because of the large χ_r^2 value.

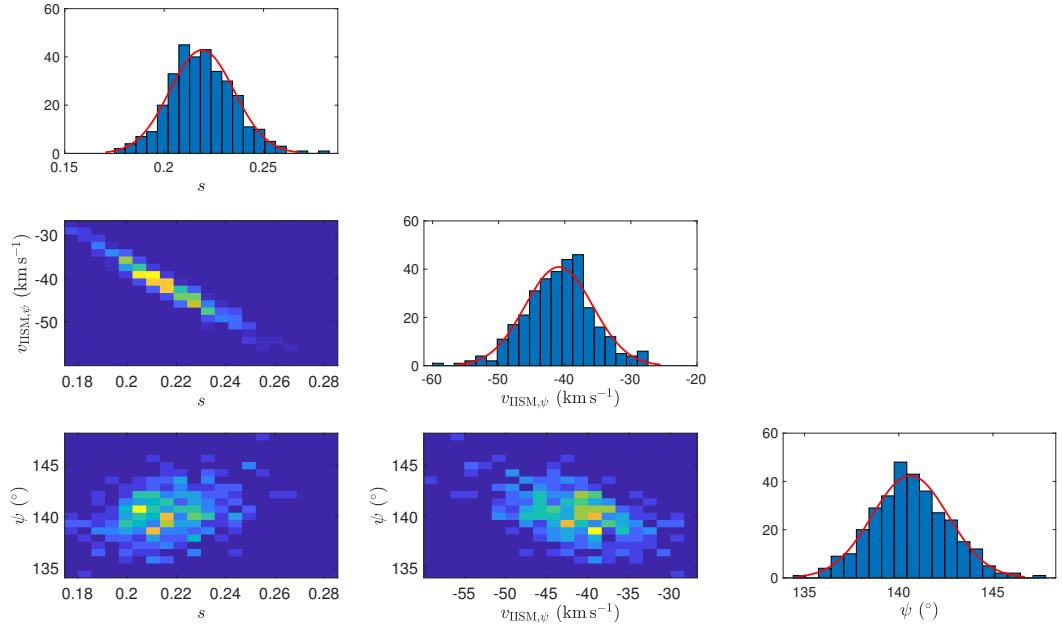


Figure 4.6: Posterior probability distributions for the parameters of the anisotropic model for the secondary arc using the MCMC sampler of Veitch and Vecchio (2010), implemented in MATLAB by Pitkin and Romano (2013).

4.6 Discussion

Scintillation studies are useful for studying the IISM and modelling the transverse velocity, which pulsar timing is not usually as sensitive to. This has important applications when parameters such as the inclination angle and longitude of ascending node can be measured, particularly for relativistic binary pulsars, where these parameters are important for testing theories of gravity in the strong-field. The scintillation of the relativistic binaries, PSRs J0737–3039 (Rickett et al. 2014) and J1141–6545 (Chapter 3), are well analysed with diffractive scintillations because there are many scintles in their dynamic spectra, meaning that one can measure a stable autocovariance function and it will remain stable over an orbit because the scattering region is larger than the projected size of the orbit. In these pulsars we see variations in the diffractive scintillation from orbit-to-orbit that complicate the processing because the level of turbulence varies and that changes both the timescale and the bandwidth (Chapter 3).

However, for PSR J0437–4715 we do not have enough scintles in an observation to get a stable scintillation bandwidth or timescale, and the strength of scattering changes from observation-to-observation. This makes it difficult to model the transverse motion through measurement of the diffractive scintillation parameters. Fortunately, the curvature of arcs in the secondary spectra depends on the screen velocity and distance, neither of which change rapidly with time. The arc curvature is also defined mostly by the very small scale structures in the dynamic spectrum, meaning that there may be many in a dynamic spectrum, even for weak scattering. From our observations, we see that the IISM velocity, distance, and anisotropy change slowly enough for the curvature to remain stable over many years, which gives us clean annual variations for precise transverse velocity modelling.

Our models for variations to the curvature of two scintillation arcs are given in the previous Section, and below we discuss these results. In Section 4.6.1 we give the absolute distance to each of the screens, and discuss the similarity of their measured velocities and anisotropy angles. In Section 4.6.2 we briefly discuss potential sources for the anisotropic plasma, and in 4.6.2.1 we search for stars with impact parameters that may be associated with the plasma.

4.6.1 Screen distances and IISM velocity

We find that the scattering is anisotropic for both arcs, based on the distribution of power across the arcs, and by the quality of the kinematic model fits. In both cases we have also measured a significant IISM velocity, which may be related to the Alfvén speed of the interstellar plasma (Goldreich and Sridhar 1995), or perhaps to the velocity of a comoving object. Taking these velocities into account, we were able to make robust screen distance measurements of $s = 0.422 \pm 0.004$ and $s = 0.219 \pm 0.017$ for the primary and secondary (taking the slightly larger uncertainty from the MCMC analysis) arcs respectively. Using the precise new measurement of the distance to PSR J0437–4715 from Chapter 2, $D = 156.79 \pm 0.25$ pc, the absolute distances to

these screens are $D_e = 90.6 \pm 0.7$ pc and $D_e = 122 \pm 3$ pc respectively.

The primary screen distance is significantly different from that reported in the earlier PSR J0437–4715 arc analysis of Bhat et al. (2016). This is because they had only two observations and so could not account for the IISM velocity or anisotropy. Our results show that screen distance estimates from individual arcs in single observations are unreliable, particularly because of a potentially significant IISM velocity.

We also note that the IISM velocity and anisotropy measurements for the two screens are very similar (Table 4.1), despite being separated by ~ 31 pc. Rather than being a coincidence or suggestive of an association between these two screens, this similarity is likely a due to a selection effect imposed by the sampling characteristics of our dynamic spectra (Mark Walker, private communication). If the IISM for these screens had zero velocity, then the pulsar’s high proper motion would be the only constant component of the effective velocity. Since we measure the screen to be moving approximately anti-parallel to the direction of proper motion (which is 120.5° East of North; Chapter 2), the effective velocity is significantly increased. A stationary IISM would therefore have a relatively much larger curvature and would be hidden in the scattered power on the $f_{\text{dop}} = 0$ axis. High-velocity screens moving parallel or perpendicular to the proper motion will not reduce the effective velocity in order for it to be in our range.

By the same selection effect, our data are most sensitive to anisotropy near to the direction of proper motion. If the line-of-sight contains a population of screens with velocity dispersion smaller than the proper motion velocity of the pulsar (which is large at 105 km s^{-1} ; Chapter 2), then those aligned with the proper motion with an anti-parallel velocity will be favoured. We therefore predict that there may be many more scintillation arcs observable for PSR J0437–4715 if future observations are planned to maximise the uninterrupted integration time for the dynamic spectra, and therefore maximise the f_{dop} resolution in the secondary spectra in order to measure arcs with higher curvature.

4.6.2 Object candidates

Structures in the IISM are generally poorly understood because they are difficult to study; appearing only indirectly through scintillation of compact radio-frequency sources. Compact, turbulent, and over-dense regions of electron density in the IISM are known to have a high scattering efficiency, that can dominate the scattering of the entire line-of-sight, meaning that the scintillation can often be described by a single thin screen scattering model. However the origin of such compact regions, including extreme scattering events (ESEs; e.g. Fiedler et al. 1987, 1994; Coles et al. 2015; Bannister et al. 2016), largely remains a mystery. Often plasma confinement by magnetic fields in the IISM is invoked to explain the observed scattering phenomena (e.g. Goldreich and Sridhar 1995).

Anisotropic scattering, resulting in pulsar scintillation arcs, has previously been suggested to originate from inclined, corrugated "reconnection current sheets" in the IISM (Pen and Levin

2014). These current sheets form at the boundaries between magnetic field configurations after they relax from an energetic disturbance such as a supernova. Other potential structures include the boundaries of local interstellar clouds (e.g. Linsky et al. 2008), or the ionised surfaces of small, self-gravitating molecular clouds (Walker and Wardle 1998). However, the latter suggests discrete AU-scale clouds, which is a more appropriate model for ESEs (Walker et al. 2017) than the sustained, stable scattering that we observe for this pulsar. Bhat et al. (2016) also suggested that the primary screen for PSR J0437–4715 may be associated with the edge of the Local Bubble (a local region of plasma under-density and high temperature), which is estimated to be at $\sim 100 - 120$ pc (e.g. Spangler 2009). Our updated distance for the primary screen is significantly less than the measurement of Bhat et al. (2016), but either of the screens (at 90.6 ± 0.7 pc and 122 ± 3 pc) could still be associated with the boundary of the Local Bubble, or other plasma structures surrounding it.

The observed power distribution in our secondary spectra makes studying individual structures difficult, since there are no arc asymmetries, inverted arclets, or other deviations from a parabola that passes through the origin, which would be associated with asymmetric scattering across the line-of-sight. Instead, the power distributions appear most consistent with an anisotropic Kolmogorov spectrum (e.g. Cordes et al. 2006), and we also observe that DM variations (which probe larger spatial scales in the IISM) agree well with a Kolmogorov spectrum (Keith et al. 2013, Chapter 2). The anisotropy is supported by our modelling of the curvature variations, for which an anisotropic scattered image is favoured based on the χ^2 value for the fit.

Recently, Walker et al. (2017) found that the scintillation of two intra-day variable quasars (IDVs) was associated with hot (O-, B-, and A-Type) stars with small impact parameters. They modelled the extreme scattering observed in these IDVs and found that the anisotropy vector was consistent with plasma filaments that were aligned radially to nearby hot stars and comoving with them. In addition to anisotropy position vectors for both screens and a component of the plasma velocity, we have precise measurements of the distance to the scattering region, giving another constraint on possible stellar associations. Accordingly we have searched the *Hipparcos* catalogue for stars with small impact parameters and we discuss the results in the following section.

4.6.2.1 Stellar associations

The precise measurements of IISM velocity, anisotropy angle, and screen distance that are provided by our velocity models (Section 4.5) are highly constraining for the hypothesis that the scattering plasma is in the form of radial filaments associated with nearby stars (Walker et al. 2017). We have searched for stars with small impact parameters in the *Hipparcos* (Perryman et al. 1997) catalogue that have measured distances within 3σ of our screen distance measurements. We take the distances from the recent parallax measurements of the *Gaia* data release 2 (dr2; Gaia Collaboration et al. 2018), which have much higher precision than those of *Hipparcos*. We

find two candidates for the primary arc, and none for the secondary arc. For each candidate we calculate the component of their proper motion velocity parallel to our measured anisotropy, $v_{\text{Star},\psi}$, which can then be compared directly to our measurements of $v_{\text{IISM},\psi}$.

The first candidate for the primary arc is an A8IV star, HIP 21849, at a position angle of $\phi = 126.9^\circ$ (East of North, relative to the pulsar), with impact parameter $b = 1.51$ pc, distance $D = 92.3 \pm 0.4$ pc, and velocity in direction of anisotropy $v_{\text{Star},\psi} = 3.5 \text{ km s}^{-1}$. The second is a G3V star, HIP 22073, with $\phi = 57.4^\circ$, $b = 2.35$ pc, $D = 91.43 \pm 0.18$ pc, and $v_{\text{Star},\psi} = -28.2 \text{ km s}^{-1}$.

The position angle of HIP 21849 is 8.5° from the measured direction of anisotropy, but since the scattered image is perpendicular to the plasma filaments, it is a poor candidate for a model of radial filaments. Despite this, it is interesting because of its close distance to the screen. We could speculate that the position angle instead supports a transverse filament model, with magnetic fields that confine the plasma and are aligned perpendicular to the radial direction from the star at a distance of 1.5 pc. However this model is not established with independent observations, unlike the radial filament model, and the discrepancy between our measured IISM velocity and the star's proper motion velocity in the direction of anisotropy would also imply a plasma velocity of $\sim 35 \text{ km s}^{-1}$ relative to the star. For this reason we simply calculate the probability of a chance association using only the distance, described below.

HIP 22073 is a cooler star than expected by the "hot star" hypothesis, and is relatively far from the line-of-sight at $b = 2.35$ pc. However it is still an interesting candidate for this radial filament model because the distance, velocity, and position angle are all relatively close to the measured screen properties (Table 4.1). The differences between these parameters are $(D - D_e) = 0.8$ pc, $(\phi - (\psi - 90^\circ)) = 12^\circ$, and $(v_{\text{Star},\psi} - v_{\text{IISM},\psi}) = 3 \text{ km s}^{-1}$.

The probability of chance associations, $P(\text{coincidence})$, is calculated using code developed for the earlier work of Walker et al. (2017). The local stellar population is characterised with estimates of local stellar density from the *Gaia* data release 1 (Gaia Collaboration et al. 2016; Bovy 2017), and the local kinematics from *Hipparcos* (Perryman et al. 1997; Dehnen and Binney 1998). The probability of getting a chance association at least as good as those observed is calculated as described in Walker et al. (2017).

Using only the distance, we find that the chance of finding a star as hot as HIP 21849 (or hotter) as close to the screen as we observe it (or closer), is $P(\text{coincidence}) \sim 1\%$. It is therefore unlikely. This probability estimate may be improved by using estimates of local stellar density and kinematics from the recent *Gaia* data release 2 (Gaia Collaboration et al. 2018). It may also be possible that the distribution of hot stars in the vicinity of the screen is not well characterised by the local (within 60 pc) sample.

Assuming the radial filament model for HIP 22073, we calculate the $P(\text{coincidence})$ for finding a G3V (or hotter) star with distance, velocity, and position angle that match at least as well. By integrating over the error ellipse for the measured parameters, assuming they are described by the normal distributions in Figure 4.4, we calculate $P(\text{coincidence}) \sim 2.5 \times 10^{-4}$.

This 1 in 4000 chance for an association is especially intriguing, but the conclusion that the primary screen is associated with HIP 22073 challenges the hypothesis that this phenomenon is linked with hot stars. The implication is that stars at least as cool as G3V can be associated with ionised filaments at a distance of 2.35 pc. This initially seems unlikely, but there are currently no theoretical models for relating FUV luminosity to scattering properties. As with the $P(\text{coincidence})$ calculation for HIP 21849, this probability estimate will be improved by considering a sample population characterised by the new *Gaia* dr2. Since this catalogue will contain more stars (although mainly late types) than *Hipparcos*, we will also extend the analysis to search again for associations with the secondary screen.

4.7 Conclusion

Measuring the curvature of scintillation arcs as they change with the velocity of the line-of-sight through the scattering medium can be an effective way to precisely measure properties of the scattering and orbit of a binary pulsar. We have measured annual and orbital modulations to the curvature of two separate arcs for the millisecond pulsar, PSR J0437–4715. The two arcs correspond to separate scattering screens, and we have precisely measured their distance, velocity, and anisotropy. We find that the primary (strongest) and secondary screens are located at $D_e = 90.6 \pm 0.7$ pc and $D_e = 122 \pm 3$ pc from the Earth respectively. We find that their scattering is anisotropic and that the interstellar plasma in the screen has a significant speed with respect to the local standard of rest at $v_{\text{IISM},\psi} = -31.2 \pm 0.5$ km s^{−1} and $v_{\text{IISM},\psi} = -41 \pm 5$ km s^{−1} respectively, in the direction of the measured anisotropy. Not accounting for this IISM velocity, and the scattering anisotropy, can lead to biased screen distance measurements.

We found that the kinematics of the scattering is very well modelled by just a single IISM velocity and anisotropy angle across the ~ 13 years of observations, meaning that the properties of the interstellar plasma remain stable over a spatial scale of at least ~ 350 AU. Our precise velocity model can provide a measurement of the longitude of the ascending node for PSR J0437–4715 independently of pulsar timing, and we find $\Omega = 207.2 \pm 0.7$, which is more precise than that obtained from the timing model of Chapter 2. The arc curvature is also sensitive to the orbital inclination angle, and we have measured $i = 136.1 \pm 0.5^\circ$. This is less precise than i measured in the timing model, but shows that modelling these variations can give precise inclination angle measurements for pulsars that do not have such precise timing. Our measurement of i is also relatively accurate compared with the modelling of diffractive scintillation in e.g. Chapter 3. This will be important for future relativistic binary pulsars for example.

If the proper motion of a pulsar is not known from pulsar timing, this method could be used to estimate it by assuming that its velocity is larger than any IISM velocity (as in Chapter 3). This method is also promising for pulsars observed in the weak scattering regime, where the scintillation bandwidth and timescale is too unstable from observation-to-observation to reliably

measure the properties of diffractive scintillation and their change with time.

We have discussed potential structures in the IISM that could generate the observed anisotropic scattering. Possibilities include inclined reconnection current sheets, boundaries between interstellar clouds, or the boundary of the Local Bubble. Since we have precise screen distance, velocity, and anisotropy information, we have been able to test a recent finding of plasma filaments oriented radially from hot stars with small impact parameters. We searched stellar catalogues for candidates, and calculated the probabilities of their association being a coincidence. For the primary screen, the best candidate is a G3V star located 2.35 pc off the line-of-sight, HIP 22073, with a very small chance of association by coincidence $P(\text{coincidence}) \sim 2.5 \times 10^{-4}$, owing to its distance, velocity, and position angle that all agree well with a radial filament model. Since this star is cooler and further from the line-of-sight than expected, this possible association is surprising.

This work will lead to a publication after an extension to include a targeted search for additional minor arcs. The dynamic spectra sampling will be optimised for this work with new observations. We will also improve the estimate of stellar association probabilities, using the new *Gaia* dr2 catalogue for our target search and for characterising the local stellar population.

CONCLUSION

High precision timing of millisecond pulsars (MSPs) is likely to produce the first evidence for nanohertz-frequency gravitational waves in the coming years. Advances such as this are accelerated by continuously improving data quality with next-generation observatories and instruments such as MeerKAT, FAST, and the ultra-wideband receiver for the Parkes radio telescope. Other applications of pulsar timing arrays will also see advances including improved tests of general relativity with better timing precision and new discoveries of relativistic binaries, and studies of IISM effects across wide observing bandwidths.

We are already discovering that uncorrected effects of the ionised interstellar medium (IISM) in pulsar timing data may in part be limiting our sensitivity to gravitational waves (Shannon et al. 2015). As data improves in the near future, it will become more important that we understand in detail the phenomena that cause small changes to scattering and frequency dispersion with time. With this comes the opportunity to use these effects to understand not only the IISM, but also the pulsars themselves. The primary result in this thesis was a demonstration of the complementary methods of pulsar timing and interstellar scintillation modelling for producing precise pulsar ephemerides. Since pulsar timing is most sensitive to radial components of orbital motion, while the pulsar scintillation is sensitive to the transverse motion, a scintillating binary pulsar can have its full three-dimensional orbit reconstructed without the need for the measurement of subtle kinematic orbital effects that are only seen in the most precisely-timed MSPs (e.g. Kopeikin 1995, 1996, see also Chapter 2). In addition, scintillometry can provide a new way for estimating distances to pulsars (Chapter 3), and can be useful in the regime weak scattering (Chapter 4) as well as the more common strong scattering.

In Chapter 2 of this thesis, the technique of precise pulsar timing in an era of constantly-improving data was demonstrated. New timing models were presented for the 20 MSPs of the

Parkes Pulsar Timing Array (PPTA) using the extended first data release (Manchester et al. 2013), which contains timing residuals with an abrupt change to their noise characteristics at a known epoch. The abrupt change in the noise characteristics corresponded with the start of the official PPTA project in 2005, beyond which regular observations at multiple frequencies were made for each pulsar. This allowed for measurement and correction of the significant dispersion measure (DM) variations in the data, which would otherwise produce strong red noise. However the legacy data prior to 2005, which is valuable for pulsar timing model improvements because of its long timespan, did not have observations at multiple frequencies and thus included this DM noise. To account for this, we presented an extension to an earlier method of accounting for red noise by choosing an appropriate whitening transformation for the data and model (the "Cholesky method" introduced in Chapter 1; Coles et al. 2011). The extension was simply to allow for multiple noise components to deal with non-stationary noise in the data. For the datasets in Chapter 2, we applied one noise model for the intrinsic "timing noise" across the entire dataset, and then characterised an additional noise component due to DM variations to apply to the legacy data only. We described this new "split-Cholesky" method and demonstrated that it gives unbiased parameter measurements and uncertainties when an appropriate noise model is used. This algorithm is currently the best alternative to the increasingly-popular Bayesian pulsar timing algorithms (e.g. Lentati et al. 2014a; van Haasteren and Levin 2013), which can now account for non-stationary noise processes, but are much more computationally expensive. The generalised least-squares solution from the split-Cholesky algorithm is equivalent to that of the maximum likelihood approach of these Bayesian algorithms if the noise model is accurate, and we therefore recommend its use in pulsar timing work for the purpose of computational efficiency, especially for problems that require multiple/repeated model fits.

For each of the PPTA MSPs, we presented new noise models that were used in this split-Cholesky algorithm, which included new timing noise models and novel DM models that included linear trends and annual DM variations. These parameters were included because the DM variations showed linear trends more frequently than one would expect for a purely Kolmogorov turbulent medium (e.g. Keith et al. 2013). The conclusion is that in some cases, the mean interstellar electron density along the line of sight changes steadily because of either radial motion of the pulsar through a relatively dense medium, or because of transverse motions of the line-of-sight through a density gradient. The latter is the likely cause of annual variations to the DM, where the density gradient may be located nearer to the Earth and thus is modulated as the line-of-sight cuts through it following the motions of the Earth. Through the same mechanism, there is likely to be DM variations following the binary orbital period for some pulsars. However in the PPTA sample, the majority have short orbital periods and are not observed with a high enough cadence to detect these changes. The DM variations are small and are usually measured as an average change across a time span of order ~ 100 days, but small orbital modulations caused by a transverse density gradient may be one of the sources for excess noise in the "corrected" DM

variations that affect the sensitivity of PTAs. Future scintillation studies can probe the IISM on much smaller time and length scales than the DM variations, and can potentially reveal these changes to the IISM. In the following section we summarise the prospects for future work to use measurements of the phase gradient across the line-of-sight as an estimator for DM variations; an approach that may be able to detect changes in DM over smaller time scales than pulsar timing, and with a comparable sensitivity (Bill Coles, private communication).

Using our new noise models for each pulsar, we then completed the timing model fit to improve their ephemerides and measure new parameters. The key results included new parallax measurements for five pulsars, improved masses for important pulsars such as PSRs J1909–3744 and J0437–4715, and also an exceptionally precise distance measurement for PSR J0437–4715. The mass measurements are important for constraining the neutron star equation of state, and will be used soon in the Neutron star Interior Composition Explorer (NICER; Arzoumanian et al. 2014b) mission, which aims to measure the radius of pulsars, including PSR J0437–4715, by modelling their x-ray emission. The distance measurement for PSR J0437–4715 of $D = 156.79 \pm 0.25$ pc is particularly promising for future science, since it is the first time that the distance to any millisecond pulsar has been measured with an uncertainty of less than 1 lightyear, which is comparable to the wavelength of gravitational waves that PTAs are sensitive to. For continuous gravitational waves originating from a single source, such as a supermassive black hole binary in a nearby galaxy, knowing the distance to this precision will allow us to determine the pulsar term and include it deterministically in the search template for the wave. This effectively increases the available signal because the pulsar term is often regarded as unknown noise during searches. We have also used this precise distance measurement in Chapter 4 to determine the distances to two thin scattering screens to high precision. Additional applications for these results are listed in the following section.

In Chapter 3 we continued the study of IISM effects by measuring and modelling the diffractive scintillations of the relativistic binary pulsar PSR J1141–6545. As in Chapter 2, the goal was to improve the ephemeris for this important pulsar. PSR J1141–6545 has previously been used to test general relativity (e.g. Bhat et al. 2008; Manchester et al. 2010) and it is unusual for being an asymmetrical-mass binary with an elliptical orbit, where the white dwarf companion formed first (Tauris and Sennels 2000; Davies et al. 2002). The model for the pulsar is therefore particularly interesting for its relativistic effects (to test theories of gravity) and its runaway velocity (to test models of its formation). The former can easily be measured through pulsar timing, but as discussed in Chapter 2 and Chapter 3, there are also kinematic contributions to these post-Keplerian parameters that are considered a contamination for the relativistic component. For PSR J1141–6545, the most significant contamination is from the Shklovskii effect (Shklovskii 1970), which depends on the distance and proper motion (and actually provided our precise new distance to PSR J0437–4715 in Chapter 2). However, for PSR J1141–6545 the distance is poorly constrained and the proper motion is unknown from pulsar timing. From our

modelling of scintillation timescale modulation in Chapter 3, we were able to determine the transverse velocity and inclination angle uniquely, which is currently impossible to do with timing of this pulsar. This was possible because of our improved long-term scintillation velocity models that account for anisotropy in the scattering. These models also included the Earth’s velocity, which allowed us to uniquely determine the sense of the inclination angle, and the relativistic advance of periastron, which we showed could be measured from scintillation alone.

The above method is most effective when the strength of scattering is such that the scintillation bandwidth can be measured easily. If the scattering is too weak, there may not be enough scintles in the dynamic spectra to estimate the bandwidth accurately. If the scattering is too strong, the scintillation bandwidth may be less than the channel bandwidth, as was the case for a subset of our observations. For these observations, we used the flux variance to estimate the number of scintles in the band, and thus their bandwidth, but this approach also required a small calibration factor (Chapter 3, Section 2). Measurement of the scintillation bandwidth enables corrections for changes to the strength of scattering with time using a numerical model (we used the model of Cordes and Rickett 1998). Following this correction, we were able to give a new estimate for the distance to PSR J1141–6545 of $D = 9_{-3}^{+5}$ kpc. This distance, in combination with our measured transverse velocity and longitude of ascending node, gave the first estimate of the proper motion. We then used this to estimate the contamination to the relativistic orbital period derivative from the Shklovskii effect, and showed that it exists only at the $\sim 1\%$ level. Our new approach for estimating the distance will be improved in the near future with improved numerical models for the scintillation velocity provided by long-term scintillation studies with observations across large frequency ranges. This is because it requires a strong understanding of the relationship between scintillation bandwidth and diffractive spatial scale, and of the distribution of interstellar plasma along the line-of-sight.

For pulsars in weak scattering and/or with long scintillation timescales, such as PSR J0437–4715, reliable measurement of the scintillation bandwidth and timescale may be almost impossible. However, in Chapter 4 we showed that the parabolic arc phenomenon observed in secondary spectra originates from finer-scale structures in dynamic spectra and can be used for precise velocity modelling. Using ~ 13 years of data for PSR J0437–4715 we measured the curvature of these arcs and found that they were modulated by the motion of the Earth and pulsar in the same way as the diffractive scintillations of Chapter 3. We identified two prominent parabolic arcs for many of the observations, each corresponding to a discrete, thin scattering screen on the line-of-sight. The annual and orbital modulation of the primary (highest signal-to-noise ratio) arc was so pronounced that we were able to measure parameters of the orbit to a precision that was comparable to the precise timing model. Our model used the timing model and distance measurement provided in Chapter 2, and used parts of the velocity model established in Chapter 3. We showed that like the well-known diffractive timescale modulation, changes to arc curvature can be used to measure useful astrometric and orbital parameters such as the inclination angle i

and longitude of the ascending node Ω . The latter was measured to a higher precision than the timing model of Chapter 2, and we found $\Omega = 207.2 \pm 0.7^\circ$. For relativistic binaries with suitable scintillation properties, this approach could give precise measurements of the inclination angle, and with it an additional constraint on the system masses for testing theories of gravity if the Shapiro delay cannot be measured through timing.

The scattering in the direction of PSR J0437–4715 appears incredibly stable, leading to arc curvatures that do not significantly deviate from our model over ~ 13 years, which corresponds to a spatial scale of ~ 350 AU. We favoured a model of anisotropic scattering for both of the scattering screens, and measured their distances to be $D_e = 90.6 \pm 7$ pc and $D_e = 123 \pm 4$ pc for the primary and secondary scattering screens respectively. Because the proper motion of the pulsar is known to high precision from the timing model (Chapter 2), we were able to include the velocity of the interstellar plasma as a parameter of the kinematic model (unlike in Chapter 3 where the IISM velocity could not be separated from the proper motion). However, for the preferred anisotropic model, the data were only sensitive to the component of IISM velocity parallel to the anisotropy. The velocity had magnitudes of $v_{\text{IISM},\psi} = -31.2 \pm 0.5$ km s $^{-1}$ and $v_{\text{IISM},\psi} = -41 \pm$ km s $^{-1}$ in the direction of $\psi = 135.4 \pm 0.5^\circ$ and $\psi = 141 \pm 2^\circ$ (East of North) for the primary and secondary screens respectively. The velocity measurements may represent the Alfvén speed of the interstellar plasma at these locations, or may be explained by an association with a comoving object. We searched for stars near to the line-of-sight with measured distances, velocities, and position angles that matched the measured screen properties, following a recent model of plasma filaments aligned radially to hot stars (Walker et al. 2017). Finally, we calculated the probability of a chance association for two candidate stars after characterising the statistics of a local stellar population with stellar catalogues, as described in Walker et al. (2017). We found a $\sim 2.5 \times 10^{-4}$ probability of a chance association (assuming this radial filament model) for the G3V star, HIP 22073, which is 2.35 pc off the line-of-sight. However this probability is for a chance association of a star with this spectral type or hotter. The probability will decrease if the search is broadened to include cooler stars, but an appropriate lower-bound temperature is not known. If this is a true association, it challenges the hypothesis that the plasma is associated with only the hottest (e.g. O-, B-, and A-type) stars. An absent association for the secondary scintillation arc may also suggest that the apparent primary arc association is spurious.

5.1 Suggestions for future work

The work in this thesis has demonstrated that scintillometry can be used in partnership with pulsar timing to improve pulsar ephemerides, and the results have opened up several pathways for further research.

The precise distance measurement to PSR J0437–4715 presented in Chapter 2, which was provided by the Shklovskii effect on the orbital period derivative, has several important applica-

tions. In that Chapter we discussed that the distance can be used in combination with another independent distance measurement (for example from an improved timing or VLBI parallax) to place limits on changes to Newton’s gravitational constant G . This is possible because the measured orbital period derivative \dot{P}_b is sensitive to any anomalous changes, which could be caused for example by a time derivative of G . An independent distance measurement is required because the contributions from the Shklovskii effect and differential Galactic acceleration must be accounted for accurately. By the same reasoning, a precise independent distance could also be used to directly measure the acceleration of the pulsar in the Galactic potential, under the assumption that G is constant. The distance will also be useful for future searches for single-source gravitational waves using the pulsar term. This will be especially useful as the precision of the distance measurement decreases with improved timing until the uncertainty is dominated by the relatively poorly understood contribution to \dot{P}_b from differential Galactic acceleration. The precise mass measurements, particularly for PSR J0437–4715 will also soon be used by the NICER mission and will then help to constrain the neutron star equation of state.

The diffractive scintillation analysis of PSR J1141–6545 in Chapter 3 showed that careful modelling of long-term annual and orbital variations can provide new measurements of: the inclination angle, longitude of ascending node, transverse velocity, screen distance, and scattering anisotropy. We were also able to estimate the pulsar distance in a new way and we used this to determine the proper motion for the first time. However, the model was limited by the assumptions of a stationary anisotropy and screen velocity. This problem was overcome in a long-term scintillation study of the double pulsar by Rickett et al. (2014). For that work, they used an analytical model of the two-dimensional autocovariance function (ACF) of the dynamic spectra to fit for the anisotropy and IISM velocity directly from the shape of the ACF. Unfortunately the scintillation bandwidth in our observations was too small to see any of the necessary structure in the ACF to fit for these additional parameters. Future observations can be processed with smaller channel bandwidths to improve the frequency resolution in the dynamic spectra so that the anisotropy, and its variation with time, can be measured independently of the scintillation velocity model. This will likely improve the accuracy and precision of the measured parameters; particularly the inclination angle, which is highly correlated with the axial ratio parameter for the anisotropy.

We have briefly analysed other pulsars in the PPTA with the methods presented in Chapter 3. We found annual variations in the scintillation timescale for each of the pulsars with annual DM variations measured in Chapter 2: PSRs J0613–0200, J1045–4509, J1643–1224, and J1939+2134. This supports the idea that the annual DM variations are caused by a region in the interstellar plasma near to the Earth with a transverse density gradient, because it also appears to affect the scattering. Further analysis of these pulsars should provide a measurement of the distance to the scattering regions, and may also allow measurements of astrometric and/or orbital parameters for these pulsars in the same way as PSR J1141–6545 in Chapter 3.

We also discussed in Chapter 3 the presence of a "phase gradient" across the line-of-sight, which was measured in the dynamic spectrum as a "tilt" to the scintles. A constant phase gradient causes a frequency-dependent angular displacement of the direct line-of-sight, which in turn causes a spatial displacement to the diffraction pattern. The frequency-dependent spatial displacement is the source of the observed tilts in the ACF, when sampled by the moving line-of-sight. For the work in Chapter 3 we did not have the resolution in the ACFs to measure the anisotropy accurately, but we were able to make simple measurements of this tilt. We simply reported that it varied with orbital phase but that it did not affect our measurements of the decorrelation bandwidth. Each observation is only sensitive to the component of the phase gradient in the direction of the effective velocity, but with an accurate model of this velocity, the true phase gradient can be recovered. If these individual phase gradient measurements are integrated over time, the total phase $\Phi(t)$ can be recovered, which is directly proportional to the DM variations, $DM(t)$ (Bill Coles, private communication). Regularly monitoring pulsars for precise measurements of the phase gradient can therefore potentially be used to measure $DM(t)$. Measurements and corrections of $DM(t)$ using the phase gradients can hopefully be done on shorter timescales than through timing, and with a comparable precision. This method requires only one relatively narrow observing band, rather than the wide frequency range required for $DM(t)$ measurement and correction in the traditional pulsar timing method (e.g. the 10 cm, 20 cm, and 40/50 cm bands in the PPTA). We have already investigated this technique for some pulsars in the PPTA, and found that it can potentially be used to identify extreme scattering events (ESEs). We observed a change to the tilts of scintles across the known ESE for PSR J1603–7202 (Keith et al. 2013; Coles et al. 2015), which was recovered despite having an incomplete model of the effective velocity. The method currently requires strong scattering for stable measurements of the diffractive scintles, a prior model of the line-of-sight velocity, and some calibration scaling factor to relate $\Phi(t)$ to the measured $DM(t)$. We are continuing to develop this method, which may provide a promising alternative to current $DM(t)$ measurement techniques and help to reduce the excess noise in "DM-corrected" timing residuals.

Our scintillation arc analysis in Chapter 4 shows that modelling arc curvature variations is a promising alternative to diffractive scintillation modelling for pulsars with sharply defined arcs and/or weak scattering. The characteristics of the dynamic spectra used in this chapter introduced a selection bias that favoured scattering screens with velocities aligned anti-parallel to the proper motion of the pulsar (to produce a higher effective velocity). We anticipate that many more screens may be detectable in the secondary spectra if the uninterrupted integration time of future observations is increased, which will increase the differential Doppler frequency resolution. It is possible that arcs of higher curvature can be measured with such observations, since we have already seen some evidence for structures inside the primary arcs.

All of the work presented in this thesis will benefit from the next generation of radio telescopes, observing instruments, analysis techniques, and computing. Wide-bandwidth observations with

the new Parkes ultra-wideband receiver will allow us to study IISM effects with unprecedented detail. We will be able to measure the frequency evolution of scattering properties to accurately model the relationship between scintillation bandwidth and diffractive spatial scale. This will be useful for the new distance estimation technique of Chapter 3. A wide bandwidth also allows for segments of a dynamic spectrum to be selected based on the optimal strength of scattering, to avoid radio frequency interference, or to optimise the clarity of scintillation arcs. Pulsar timing across a wide bandwidth will significantly improve the correction of DM variations for pulsar timing, both because of the improved sensitivity and because of increased observing efficiency, which may lead to a higher observing cadence. Improved observing efficiency will allow for regular monitoring of a larger set of pulsars for both pulsar timing and IISM applications, giving improved sensitivity to transient events in the pulsar such as glitches or magnetosphere disruptions, or in the IISM such as extreme scattering events. New techniques such as the use of cyclic spectroscopy (e.g. Demorest 2011) can significantly improve the frequency resolution in dynamic spectra, which will help with measurements of scintillation bandwidth in sources for which it is unresolved (e.g. some observations of PSR J1141–6545 in Chapter 3).

We hope that the work in this thesis will encourage other pulsar astronomers to utilise their pulsar timing observations for scintillation studies more frequently. This will be useful for better understanding of the pulsars in the array, and of the IISM noise in timing residuals.

BIBLIOGRAPHY

- B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, and et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901, July 2009. doi: 10.1088/0034-4885/72/7/076901.
- B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6):061102, February 2016. doi: 10.1103/PhysRevLett.116.061102.
- A. C. Aitken. Iv.—on least squares and linear combination of observations. *Proceedings of the Royal Society of Edinburgh*, 55:42–48, 1936. doi: 10.1017/S0370164600014346.
- Hirotougu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotougu Akaike*, pages 199–213. Springer, 1998.
- T. W. Anderson and D. A. Darling. A test of goodness of fit. *Journal of the American Statistical Association*, 49(268):765–769, 1954. doi: 10.1080/01621459.1954.10501232. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1954.10501232>.
- J. W. Armstrong, B. J. Rickett, and S. R. Spangler. Electron density power spectrum in the local interstellar medium. *ApJ*, 443:209–221, April 1995. doi: 10.1086/175515.
- Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, J. M. Cordes, P. B. Demorest, X. Deng, T. Dolch, J. A. Ellis, R. D. Ferdman, N. Garver-Daniels, F. Jenet, G. Jones, V. M. Kaspi, M. Koop, M. T. Lam, T. J. W. Lazio, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. A. McLaughlin, S. T. McWilliams, D. J. Nice, N. Palliyaguru, T. T. Pennucci, S. M. Ransom, A. Sesana, X. Siemens, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, M. Vallisneri, R. van Haasteren, Y. Wang, W. W. Zhu, and NANOGrav Collaboration. Gravitational Waves from Individual Supermassive Black Hole Binaries in Circular Orbits: Limits from the North American Nanohertz Observatory for Gravitational Waves. *ApJ*, 794:141, October 2014a. doi: 10.1088/0004-637X/794/2/141.
- Z. Arzoumanian, K. C. Gendreau, C. L. Baker, T. Cazeau, P. Hestnes, J. W. Kellogg, S. J. Kenyon, R. P. Kozon, K.-C. Liu, S. S. Manthripragada, C. B. Markwardt, A. L. Mitchell, J. W. Mitchell,

- C. A. Monroe, T. Okajima, S. E. Pollard, D. F. Powers, B. J. Savadkin, L. B. Winternitz, P. T. Chen, M. R. Wright, R. Foster, G. Prigozhin, R. Remillard, and J. Doty. The neutron star interior composition explorer (NICER): mission definition. In *Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*, volume 9144 of *Proc. SPIE*, page 914420, July 2014b. doi: 10.1117/12.2056811.
- W. Baade and F. Zwicky. Cosmic Rays from Super-novae. *Proceedings of the National Academy of Science*, 20:259–263, May 1934. doi: 10.1073/pnas.20.5.259.
- D. C. Backer, S. R. Kulkarni, C. Heiles, M. M. Davis, and W. M. Goss. A millisecond pulsar. *Nature*, 300:615–618, December 1982. doi: 10.1038/300615a0.
- M. Bailes, P. A. Harrison, D. R. Lorimer, S. Johnston, A. G. Lyne, R. N. Manchester, N. D’Amico, L. Nicastro, T. M. Tauris, and C. Robinson. Discovery of three binary millisecond pulsars. *ApJ*, 425:L41–L44, April 1994. doi: 10.1086/187306.
- K. W. Bannister, J. Stevens, A. V. Tuntsov, M. A. Walker, S. Johnston, C. Reynolds, and H. Bignall. Real-time detection of an extreme scattering event: Constraints on Galactic plasma lenses. *Science*, 351:354–356, January 2016. doi: 10.1126/science.aac7673.
- C. G. Bassa, G. H. Janssen, B. W. Stappers, T. M. Tauris, T. Wevers, P. G. Jonker, L. Lentati, J. P. W. Verbiest, G. Desvignes, E. Graikou, L. Guillemot, P. C. C. Freire, P. Lazarus, R. N. Caballero, D. J. Champion, I. Cognard, A. Jessner, C. Jordan, R. Karuppusamy, M. Kramer, K. Lazaridis, K. J. Lee, K. Liu, A. G. Lyne, J. McKee, S. Osłowski, D. Perrodin, S. Sanidas, G. Shaifullah, R. Smits, G. Theureau, C. Tiburzi, and W. W. Zhu. A millisecond pulsar in an extremely wide binary system. *MNRAS*, 460:2207–2222, August 2016. doi: 10.1093/mnras/stw1134.
- J. F. Bell and M. Bailes. New Method for Obtaining Binary Pulsar Distances and Its Implications for Tests of General Relativity. *ApJ*, 456:L33, January 1996. doi: 10.1086/309862.
- J. F. Bell, M. Bailes, R. N. Manchester, A. G. Lyne, F. Camilo, and J. S. Sandhu. Timing Measurements and Their Implications for Four Binary Millisecond Pulsars. *MNRAS*, 286: 463–469, April 1997. doi: 10.1093/mnras/286.2.463.
- N. D. R. Bhat, M. Bailes, and J. P. W. Verbiest. Gravitational-radiation losses from the pulsar white-dwarf binary PSR J1141 6545. *Phys. Rev. D*, 77(12):124017, June 2008. doi: 10.1103/PhysRevD.77.124017.
- N. D. R. Bhat, S. M. Ord, S. E. Tremblay, S. J. McSweeney, and S. J. Tingay. Scintillation Arcs in Low-frequency Observations of the Timing-array Millisecond Pulsar PSR J0437-4715. *ApJ*, 818:86, February 2016. doi: 10.3847/0004-637X/818/1/86.
- D. Bhattacharya and E. P. J. van den Heuvel. Formation and evolution of binary and millisecond radio pulsars. *Phys. Rep.*, 203:1–124, 1991. doi: 10.1016/0370-1573(91)90064-S.

- G. S. Bisnovatyi-Kogan and B. V. Komberg. Pulsars and close binary systems. *Soviet Ast.*, 18:217, October 1974.
- Max Born, Emil Wolf, A. B. Bhatia, P. C. Clemmow, D. Gabor, A. R. Stokes, A. M. Taylor, P. A. Wayman, and W. L. Wilcock. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 7 edition, 1999. doi: 10.1017/CBO9781139644181.
- J. Bovy. Stellar inventory of the solar neighbourhood using Gaia DR1. *MNRAS*, 470:1360–1387, September 2017. doi: 10.1093/mnras/stx1277.
- J. Bovy, C. Allende Prieto, T. C. Beers, D. Bizyaev, L. N. da Costa, K. Cunha, G. L. Ebelke, D. J. Eisenstein, P. M. Frinchaboy, A. E. García Pérez, L. Girardi, F. R. Hearty, D. W. Hogg, J. Holtzman, M. A. G. Maia, S. R. Majewski, E. Malanushenko, V. Malanushenko, S. Mészáros, D. L. Nidever, R. W. O’Connell, C. O’Donnell, A. Oravetz, K. Pan, H. J. Rocha-Pinto, R. P. Schiavon, D. P. Schneider, M. Schultheis, M. Skrutskie, V. V. Smith, D. H. Weinberg, J. C. Wilson, and G. Zasowski. The Milky Way’s Circular-velocity Curve between 4 and 14 kpc from APOGEE data. *ApJ*, 759:131, November 2012. doi: 10.1088/0004-637X/759/2/131.
- W. F. Briskin, J.-P. Macquart, J. J. Gao, B. J. Rickett, W. A. Coles, A. T. Deller, S. J. Tingay, and C. J. West. 100 μ as Resolution VLBI Imaging of Anisotropic Interstellar Scattering Toward Pulsar B0834+06. *ApJ*, 708:232–243, January 2010. doi: 10.1088/0004-637X/708/1/232.
- D. J. Champion, G. B. Hobbs, R. N. Manchester, R. T. Edwards, D. C. Backer, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, W. Coles, P. B. Demorest, R. D. Ferdman, W. M. Folkner, A. W. Hotan, M. Kramer, A. N. Lommen, D. J. Nice, M. B. Purver, J. M. Sarkissian, I. H. Stairs, W. van Straten, J. P. W. Verbiest, and D. R. B. Yardley. Measuring the Mass of Solar System Planets Using Pulsar Timing. *ApJ*, 720:L201–L205, September 2010. doi: 10.1088/2041-8205/720/2/L201.
- S. Chatterjee, W. F. Briskin, W. H. T. Vlemmings, W. M. Goss, T. J. W. Lazio, J. M. Cordes, S. E. Thorsett, E. B. Fomalont, A. G. Lyne, and M. Kramer. Precision Astrometry with the Very Long Baseline Array: Parallaxes and Proper Motions for 14 Pulsars. *ApJ*, 698:250–265, June 2009. doi: 10.1088/0004-637X/698/1/250.
- K. S. Cheng. Outer magnetospheric fluctuations and pulsar timing noise. *ApJ*, 321:799–804, October 1987. doi: 10.1086/165672.
- R. P. Church, S. J. Bush, C. A. Tout, and M. B. Davies. Detailed models of the binary pulsars J1141-6545 and B2303+46. *MNRAS*, 372:715–727, October 2006. doi: 10.1111/j.1365-2966.2006.10897.x.

- W. Coles, G. Hobbs, D. J. Champion, R. N. Manchester, and J. P. W. Verbiest. Pulsar timing analysis in the presence of correlated noise. *MNRAS*, 418:561–570, November 2011. doi: 10.1111/j.1365-2966.2011.19505.x.
- W. A. Coles, M. A. McLaughlin, B. J. Rickett, A. G. Lyne, and N. D. R. Bhat. Probing the Eclipse of J0737-3039A with Scintillation. *ApJ*, 623:392–397, April 2005. doi: 10.1086/428342.
- W. A. Coles, M. Kerr, R. M. Shannon, G. B. Hobbs, R. N. Manchester, X.-P. You, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. J. Keith, Y. Levin, S. Osłowski, V. Ravi, D. Reardon, L. Toomey, W. van Straten, J. B. Wang, L. Wen, and X. J. Zhu. Pulsar Observations of Extreme Scattering Events. *ApJ*, 808:113, August 2015. doi: 10.1088/0004-637X/808/2/113.
- J. M. Cordes. Observational limits on the location of pulsar emission regions. *ApJ*, 222:1006–1011, June 1978. doi: 10.1086/156218.
- J. M. Cordes. Pulsar Observations I. – Propagation Effects, Searching Distance Estimates, Scintillations and VLBI. In S. Stanimirovic, D. Altschuler, P. Goldsmith, and C. Salter, editors, *Single-Dish Radio Astronomy: Techniques and Applications*, volume 278 of *Astronomical Society of the Pacific Conference Series*, pages 227–250, December 2002.
- J. M. Cordes and T. J. Lazio. Interstellar scattering effects on the detection of narrow-band signals. *ApJ*, 376:123–133, July 1991. doi: 10.1086/170261.
- J. M. Cordes and T. J. W. Lazio. NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations. *ArXiv Astrophysics e-prints*, July 2002.
- J. M. Cordes and B. J. Rickett. Diffractive Interstellar Scintillation Timescales and Velocities. *ApJ*, 507:846–860, November 1998. doi: 10.1086/306358.
- J. M. Cordes, A. Pidwerbetsky, and R. V. E. Lovelace. Refractive and diffractive scattering in the interstellar medium. *ApJ*, 310:737–767, November 1986. doi: 10.1086/164728.
- J. M. Cordes, B. J. Rickett, D. R. Stinebring, and W. A. Coles. Theory of Parabolic Arcs in Interstellar Scintillation Spectra. *ApJ*, 637:346–365, January 2006. doi: 10.1086/498332.
- S. Dai, G. Hobbs, R. N. Manchester, M. Kerr, R. M. Shannon, W. van Straten, A. Mata, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, W. A. Coles, S. Johnston, M. J. Keith, Y. Levin, S. Osłowski, D. Reardon, V. Ravi, J. M. Sarkissian, C. Tiburzi, L. Toomey, H. G. Wang, J.-B. Wang, L. Wen, R. X. Xu, W. M. Yan, and X.-J. Zhu. A study of multifrequency polarization pulse profiles of millisecond pulsars. *MNRAS*, 449:3223–3262, May 2015. doi: 10.1093/mnras/stv508.
- T. Damour and N. Deruelle. General relativistic celestial mechanics of binary systems. I. The post-Newtonian motion. *Ann. Inst. Henri Poincaré Phys. Théor., Vol. 43, No. 1, p. 107 - 132*, 43: 107–132, 1985.

- T. Damour and N. Deruelle. General relativistic celestial mechanics of binary systems. II. The post-Newtonian timing formula. *Ann. Inst. Henri Poincaré Phys. Théor.*, Vol. 44, No. 3, p. 263 - 292, 44:263–292, 1986.
- M. B. Davies, H. Ritter, and A. King. Formation of the binary pulsars J1141-6545 and B2303+46. *MNRAS*, 335:369–376, September 2002. doi: 10.1046/j.1365-8711.2002.05594.x.
- W. Dehnen and J. J. Binney. Local stellar kinematics from HIPPARCOS data. *MNRAS*, 298: 387–394, August 1998. doi: 10.1046/j.1365-8711.1998.01600.x.
- A. T. Deller, J. P. W. Verbiest, S. J. Tingay, and M. Bailes. Extremely High Precision VLBI Astrometry of PSR J0437-4715 and Implications for Theories of Gravity. *ApJ*, 685:L67, September 2008. doi: 10.1086/592401.
- P. B. Demorest. Cyclic spectral analysis of radio pulsars. *MNRAS*, 416:2821–2826, October 2011. doi: 10.1111/j.1365-2966.2011.19230.x.
- P. B. Demorest, R. D. Ferdman, M. E. Gonzalez, D. Nice, S. Ransom, I. H. Stairs, Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, J. M. Cordes, J. Ellis, L. S. Finn, P. Freire, S. Giampanis, F. Jenet, V. M. Kaspi, J. Lazio, A. N. Lommen, M. McLaughlin, N. Palliyaguru, D. Perrodin, R. M. Shannon, X. Siemens, D. Stinebring, J. Swiggum, and W. W. Zhu. Limits on the Stochastic Gravitational Wave Background from the North American Nanohertz Observatory for Gravitational Waves. *ApJ*, 762:94, January 2013. doi: 10.1088/0004-637X/762/2/94.
- R. T. Edwards, M. Bailes, W. van Straten, and M. C. Britton. The Swinburne intermediate-latitude pulsar survey. *MNRAS*, 326:358–374, September 2001. doi: 10.1046/j.1365-8711.2001.04637.x.
- R. T. Edwards, G. B. Hobbs, and R. N. Manchester. TEMPO2, a new pulsar timing package - II. The timing model and precision estimates. *MNRAS*, 372:1549–1574, November 2006. doi: 10.1111/j.1365-2966.2006.10870.x.
- R. Fiedler, B. Dennison, K. J. Johnston, E. B. Waltman, and R. S. Simon. A summary of extreme scattering events and a descriptive model. *ApJ*, 430:581–594, August 1994. doi: 10.1086/174432.
- R. L. Fiedler, B. Dennison, K. J. Johnston, and A. Hewish. Extreme scattering events caused by compact structures in the interstellar medium. *Nature*, 326:675–678, April 1987. doi: 10.1038/326675a0.
- R. S. Foster and J. M. Cordes. Interstellar propagation effects and the precision of pulsar timing. *ApJ*, 364:123–135, November 1990. doi: 10.1086/169393.
- P. C. C. Freire, N. Wex, G. Esposito-Farèse, J. P. W. Verbiest, M. Bailes, B. A. Jacoby, M. Kramer, I. H. Stairs, J. Antoniadis, and G. H. Janssen. The relativistic pulsar-white dwarf binary PSR

BIBLIOGRAPHY

- J1738+0333 - II. The most stringent test of scalar-tensor gravity. *MNRAS*, 423:3328–3343, July 2012. doi: 10.1111/j.1365-2966.2012.21253.x.
- Gaia Collaboration, A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, F. Mignard, R. Drimmel, C. Babusiaux, C. A. L. Bailer-Jones, U. Bastian, and et al. Gaia Data Release 1. Summary of the astrometric, photometric, and survey properties. *A&A*, 595:A2, November 2016. doi: 10.1051/0004-6361/201629512.
- Gaia Collaboration, A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, C. Babusiaux, and C. A. L. Bailer-Jones. Gaia Data Release 2. Summary of the contents and survey properties. *ArXiv e-prints*, April 2018.
- K. C. Gendreau, Z. Arzoumanian, and T. Okajima. The Neutron star Interior Composition ExploreR (NICER): an Explorer mission of opportunity for soft x-ray timing spectroscopy. In *Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray*, volume 8443 of *Proc. SPIE*, page 844313, September 2012. doi: 10.1117/12.926396.
- T. Gold. Rotating Neutron Stars as the Origin of the Pulsating Radio Sources. *Nature*, 218: 731–732, May 1968. doi: 10.1038/218731a0.
- P. Goldreich and S. Sridhar. Toward a theory of interstellar turbulence. 2: Strong alfvénic turbulence. *ApJ*, 438:763–775, January 1995. doi: 10.1086/175121.
- J. Goodman and R. Narayan. Slow pulsar scintillation and the spectrum of interstellar electron density fluctuations. *MNRAS*, 214:519–537, June 1985. doi: 10.1093/mnras/214.4.519.
- Y. Gupta. On the Correlation between Proper Motion Velocities and Scintillation Velocities of Radio Pulsars. *ApJ*, 451:717, October 1995. doi: 10.1086/176258.
- R. W. Hellings and G. S. Downs. Upper limits on the isotropic gravitational radiation background from pulsar timing analysis. *ApJ*, 265:L39–L42, February 1983. doi: 10.1086/183954.
- A. Hewish. The diffraction of radio waves in passing through a phase-changing ionosphere. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 209(1096):81–96, 1951. ISSN 0080-4630. doi: 10.1098/rspa.1951.0189. URL <http://rspa.royalsocietypublishing.org/content/209/1096/81>.
- A. Hewish. The Irregular Structure of the Outer Regions of the Solar Corona. *Proceedings of the Royal Society of London Series A*, 228:238–251, February 1955. doi: 10.1098/rspa.1955.0046.
- A. Hewish, P. F. Scott, and D. Wills. Interplanetary Scintillation of Small Diameter Radio Sources. *Nature*, 203:1214–1217, September 1964. doi: 10.1038/2031214a0.
- A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a Rapidly Pulsating Radio Source. *Nature*, 217:709–713, February 1968. doi: 10.1038/217709a0.

- G. Hobbs, D. R. Lorimer, A. G. Lyne, and M. Kramer. A statistical study of 233 pulsar proper motions. *MNRAS*, 360:974–992, July 2005. doi: 10.1111/j.1365-2966.2005.09087.x.
- G. Hobbs, A. Archibald, Z. Arzoumanian, D. Backer, M. Bailes, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, D. Champion, I. Cognard, W. Coles, J. Cordes, P. Demorest, G. Desvignes, R. D. Ferdman, L. Finn, P. Freire, M. Gonzalez, J. Hessels, A. Hotan, G. Janssen, F. Jenet, A. Jessner, C. Jordan, V. Kaspi, M. Kramer, V. Kondratiev, J. Lazio, K. Lazaridis, K. J. Lee, Y. Levin, A. Lommen, D. Lorimer, R. Lynch, A. Lyne, R. Manchester, M. McLaughlin, D. Nice, S. Osłowski, M. Pilia, A. Possenti, M. Purver, S. Ransom, J. Reynolds, S. Sanidas, J. Sarkissian, A. Sesana, R. Shannon, X. Siemens, I. Stairs, B. Stappers, D. Stinebring, G. Theureau, R. van Haasteren, W. van Straten, J. P. W. Verbiest, D. R. B. Yardley, and X. P. You. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Classical and Quantum Gravity*, 27(8):084013, April 2010. doi: 10.1088/0264-9381/27/8/084013.
- G. Hobbs, D. Miller, R. N. Manchester, J. Dempsey, J. M. Chapman, J. Khoo, J. Applegate, M. Bailes, N. D. R. Bhat, R. Bridle, A. Borg, A. Brown, C. Burnett, F. Camilo, C. Cattalini, A. Chaudhary, R. Chen, N. D’Amico, L. Kedziora-Chudczer, T. Cornwell, R. George, G. Hampson, M. Hepburn, A. Jameson, M. Keith, T. Kelly, A. Kosmynin, E. Lenc, D. Lorimer, C. Love, A. Lyne, V. McIntyre, J. Morrissey, M. Pienaar, J. Reynolds, G. Ryder, J. Sarkissian, A. Stevenson, A. Treloar, W. van Straten, M. Whiting, and G. Wilson. The Parkes Observatory Pulsar Data Archive. *Publ. Astron. Soc. Australia*, 28:202–214, August 2011. doi: 10.1071/AS11016.
- G. Hobbs, W. Coles, R. N. Manchester, M. J. Keith, R. M. Shannon, D. Chen, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. Champion, A. Chaudhary, A. Hotan, J. Khoo, J. Kocz, Y. Levin, S. Osłowski, B. Preisig, V. Ravi, J. E. Reynolds, J. Sarkissian, W. van Straten, J. P. W. Verbiest, D. Yardley, and X. P. You. Development of a pulsar-based time-scale. *MNRAS*, 427:2780–2787, December 2012. doi: 10.1111/j.1365-2966.2012.21946.x.
- G. B. Hobbs, R. T. Edwards, and R. N. Manchester. TEMPO2, a new pulsar-timing package - I. An overview. *MNRAS*, 369:655–672, June 2006. doi: 10.1111/j.1365-2966.2006.10302.x.
- J. Holmberg and C. Flynn. The local surface density of disc matter mapped by Hipparcos. *MNRAS*, 352:440–446, August 2004. doi: 10.1111/j.1365-2966.2004.07931.x.
- A. W. Hotan, W. van Straten, and R. N. Manchester. PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis. *Publ. Astron. Soc. Australia*, 21:302–309, 2004. doi: 10.1071/AS04022.
- A. W. Hotan, M. Bailes, and S. M. Ord. High-precision baseband timing of 15 millisecond pulsars. *MNRAS*, 369:1502–1520, July 2006. doi: 10.1111/j.1365-2966.2006.10394.x.
- R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *ApJ*, 195:L51–L53, January 1975. doi: 10.1086/181708.

- F. Jankowski, W. van Straten, E. F. Keane, M. Bailes, E. D. Barr, S. Johnston, and M. Kerr. Spectral properties of 441 radio pulsars. *MNRAS*, 473:4436–4458, February 2018. doi: 10.1093/mnras/stx2476.
- G. M. Jenkins and D. G. Watts. *Spectral Analysis and Its Applications*. Holden Day, 1969.
- P. B. Jones. The Generation of Timing Noise by Superfluid Rotation in Pulsars. *MNRAS*, 246: 364, September 1990.
- D. L. Kaplan, T. Kupfer, D. J. Nice, A. Irrgang, U. Heber, Z. Arzoumanian, E. Beklen, K. Crowter, M. E. DeCesar, P. B. Demorest, T. Dolch, J. A. Ellis, R. D. Ferdman, E. C. Ferrara, E. Fonseca, P. A. Gentile, G. Jones, M. L. Jones, S. Kreuzer, M. T. Lam, L. Levin, D. R. Lorimer, R. S. Lynch, M. A. McLaughlin, A. A. Miller, C. Ng, T. T. Pennucci, T. A. Prince, S. M. Ransom, P. S. Ray, R. Spiewak, I. H. Stairs, K. Stovall, J. Swiggum, and W. Zhu. PSR J1024-0719: A Millisecond Pulsar in an Unusual Long-period Orbit. *ApJ*, 826:86, July 2016. doi: 10.3847/0004-637X/826/1/86.
- V. M. Kaspi, A. G. Lyne, R. N. Manchester, F. Crawford, F. Camilo, J. F. Bell, N. D’Amico, I. H. Stairs, N. P. F. McKay, D. J. Morris, and A. Possenti. Discovery of a Young Radio Pulsar in a Relativistic Binary Orbit. *ApJ*, 543:321–327, November 2000. doi: 10.1086/317103.
- M. J. Keith, W. Coles, R. M. Shannon, G. B. Hobbs, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. J. Champion, A. Chaudhary, A. W. Hotan, J. Khoo, J. Kocz, S. Osłowski, V. Ravi, J. E. Reynolds, J. Sarkissian, W. van Straten, and D. R. B. Yardley. Measurement and correction of variations in interstellar dispersion in high-precision pulsar timing. *MNRAS*, 429:2161–2174, March 2013. doi: 10.1093/mnras/sts486.
- M. Kerr, W. A. Coles, C. A. Ward, S. Johnston, A. V. Tuntsoy, and R. M. Shannon. Extreme scattering events towards two young pulsars. *MNRAS*, 474:4637–4647, March 2018. doi: 10.1093/mnras/stx3101.
- S. M. Kopeikin. On possible implications of orbital parallaxes of wide orbit binary pulsars and their measurability. *ApJ*, 439:L5–L8, January 1995. doi: 10.1086/187731.
- S. M. Kopeikin. Proper Motion of Binary Pulsars as a Source of Secular Variations of Orbital Parameters. *ApJ*, 467:L93, August 1996. doi: 10.1086/310201.
- M. Kramer and D. J. Champion. The European Pulsar Timing Array and the Large European Array for Pulsars. *Classical and Quantum Gravity*, 30(22):224009, November 2013. doi: 10.1088/0264-9381/30/22/224009.
- M. Kramer, I. H. Stairs, R. N. Manchester, M. A. McLaughlin, A. G. Lyne, R. D. Ferdman, M. Burgay, D. R. Lorimer, A. Possenti, N. D’Amico, J. M. Sarkissian, G. B. Hobbs, J. E.

- Reynolds, P. C. C. Freire, and F. Camilo. Tests of General Relativity from Timing the Double Pulsar. *Science*, 314:97–102, October 2006. doi: 10.1126/science.1132305.
- M. T. Lam, J. A. Ellis, G. Grillo, M. L. Jones, J. S. Hazboun, P. R. Brook, J. E. Turner, S. Chatterjee, J. M. Cordes, T. J. W. Lazio, M. E. DeCesar, Z. Arzoumanian, H. Blumer, H. T. Cromartie, P. B. Demorest, T. Dolch, R. D. Ferdman, E. C. Ferrara, E. Fonseca, N. Garver-Daniels, P. A. Gentile, V. Gupta, D. R. Lorimer, R. S. Lynch, D. R. Madison, M. A. McLaughlin, C. Ng, D. J. Nice, T. T. Pennucci, S. M. Ransom, R. Spiewak, I. H. Stairs, D. R. Stinebring, K. Stovall, J. K. Swiggum, S. J. Vigeland, and W. W. Zhu. A Second Chromatic Timing Event of Interstellar Origin toward PSR J1713+0747. *ApJ*, 861:132, July 2018. doi: 10.3847/1538-4357/aac770.
- C. Lange, F. Camilo, N. Wex, M. Kramer, D. C. Backer, A. G. Lyne, and O. Doroshenko. Precision timing measurements of PSR J1012+5307. *MNRAS*, 326:274–282, September 2001. doi: 10.1046/j.1365-8711.2001.04606.x.
- M. I. Large, A. E. Vaughan, and B. Y. Mills. A Pulsar Supernova Association? *Nature*, 220: 340–341, October 1968. doi: 10.1038/220340a0.
- P. D. Lasky, A. Melatos, V. Ravi, and G. Hobbs. Pulsar timing noise and the minimum observation time to detect gravitational waves with pulsar timing arrays. *MNRAS*, 449:3293–3300, May 2015. doi: 10.1093/mnras/stv540.
- K. J. Lee, N. Wex, M. Kramer, B. W. Stappers, C. G. Bassa, G. H. Janssen, R. Karuppusamy, and R. Smits. Gravitational wave astronomy of single sources with a pulsar timing array. *MNRAS*, 414:3251–3264, July 2011. doi: 10.1111/j.1365-2966.2011.18622.x.
- L. Lentati, P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon. TEMPONEST: a Bayesian approach to pulsar timing analysis. *MNRAS*, 437:3004–3023, January 2014a. doi: 10.1093/mnras/stt2122.
- L. Lentati, M. P. Hobson, and P. Alexander. Bayesian estimation of non-Gaussianity in pulsar timing analysis. *MNRAS*, 444:3863–3878, November 2014b. doi: 10.1093/mnras/stu1721.
- J. L. Linsky, B. J. Rickett, and S. Redfield. The Origin of Radio Scintillation in the Local Interstellar Medium. *ApJ*, 675:413–419, March 2008. doi: 10.1086/526420.
- N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Ap&SS*, 39:447–462, February 1976. doi: 10.1007/BF00648343.
- D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. December 2004.
- T. E. Lutz and D. H. Kelker. On the Use of Trigonometric Parallaxes for the Calibration of Luminosity Systems: Theory. *PASP*, 85:573, October 1973. doi: 10.1086/129506.

- A. G. Lyne. Orbital inclination and mass of the binary pulsar PSR0655 + 64. *Nature*, 310:300–302, July 1984. doi: 10.1038/310300a0.
- A. G. Lyne and B. J. Rickett. Radio Observations of Five Pulsars. *Nature*, 219:1339–1342, September 1968. doi: 10.1038/2191339a0.
- A. G. Lyne and F. G. Smith. Interstellar scintillation and pulsar velocities. *Nature*, 298:825–827, August 1982. doi: 10.1038/298825a0.
- A. G. Lyne, M. Burgay, M. Kramer, A. Possenti, R. N. Manchester, F. Camilo, M. A. McLaughlin, D. R. Lorimer, N. D’Amico, B. C. Joshi, J. Reynolds, and P. C. C. Freire. A Double-Pulsar System: A Rare Laboratory for Relativistic Gravity and Plasma Physics. *Science*, 303:1153–1157, February 2004. doi: 10.1126/science.1094645.
- V. Maitia, J.-F. Lestrade, and I. Cognard. A 3 Year Long Extreme Scattering Event in the Direction of the Millisecond Pulsar J1643-1224. *ApJ*, 582:972–977, January 2003. doi: 10.1086/344816.
- R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility Pulsar Catalogue. *AJ*, 129:1993–2006, April 2005. doi: 10.1086/428488.
- R. N. Manchester, M. Kramer, I. H. Stairs, M. Burgay, F. Camilo, G. B. Hobbs, D. R. Lorimer, A. G. Lyne, M. A. McLaughlin, C. A. McPhee, A. Possenti, J. E. Reynolds, and W. van Straten. Observations and Modeling of Relativistic Spin Precession in PSR J1141-6545. *ApJ*, 710: 1694–1709, February 2010. doi: 10.1088/0004-637X/710/2/1694.
- R. N. Manchester, G. Hobbs, M. Bailes, W. A. Coles, W. van Straten, M. J. Keith, R. M. Shannon, N. D. R. Bhat, A. Brown, S. G. Burke-Spolaor, D. J. Champion, A. Chaudhary, R. T. Edwards, G. Hampson, A. W. Hotan, A. Jameson, F. A. Jenet, M. J. Kesteven, J. Khoo, J. Kocz, K. Maciesiak, S. Osłowski, V. Ravi, J. R. Reynolds, J. M. Sarkissian, J. P. W. Verbiest, Z. L. Wen, W. E. Wilson, D. Yardley, W. M. Yan, and X. P. You. The Parkes Pulsar Timing Array Project. *Publ. Astron. Soc. Australia*, 30:e017, January 2013. doi: 10.1017/pasa.2012.017.
- C. F. McKee and J. P. Ostriker. A theory of the interstellar medium - Three components regulated by supernova explosions in an inhomogeneous substrate. *ApJ*, 218:148–169, November 1977. doi: 10.1086/155667.
- R. Narayan. The Physics of Pulsar Scintillation. *Philosophical Transactions of the Royal Society of London Series A*, 341:151–165, October 1992. doi: 10.1098/rsta.1992.0090.
- E. O. Ofek. MATLAB package for astronomy and astrophysics. Astrophysics Source Code Library, July 2014.
- S. M. Ord and M. Bailes. Long-Term Behaviour of the Scintillation Velocity of PSR J1141-6545. In F. A. Rasio and I. H. Stairs, editors, *Binary Radio Pulsars*, volume 328 of *Astronomical Society of the Pacific Conference Series*, page 345, July 2005.

- S. M. Ord, M. Bailes, and W. van Straten. The Scintillation Velocity of the Relativistic Binary Pulsar PSR J1141-6545. *ApJ*, 574:L75–L78, July 2002a. doi: 10.1086/342218.
- S. M. Ord, M. Bailes, and W. van Straten. A neutral hydrogen distance limit to the relativistic binary PSR J1141-6545. *MNRAS*, 337:409–412, December 2002b. doi: 10.1046/j.1365-8711.2002.05810.x.
- S. M. Ord, B. A. Jacoby, A. W. Hotan, and M. Bailes. High-precision timing of PSR J1600-3053. *MNRAS*, 371:337–342, September 2006. doi: 10.1111/j.1365-2966.2006.10646.x.
- S. Osłowski, W. van Straten, G. B. Hobbs, M. Bailes, and P. Demorest. High signal-to-noise ratio observations and the ultimate limits of precision pulsar timing. *MNRAS*, 418:1258–1271, December 2011. doi: 10.1111/j.1365-2966.2011.19578.x.
- F. Özel, D. Psaltis, R. Narayan, and A. Santos Villarreal. On the Mass Distribution and Birth Masses of Neutron Stars. *ApJ*, 757:55, September 2012. doi: 10.1088/0004-637X/757/1/55.
- U.-L. Pen and Y. Levin. Pulsar scintillations from corrugated reconnection sheets in the interstellar medium. *MNRAS*, 442:3338–3346, August 2014. doi: 10.1093/mnras/stu1020.
- M. A. C. Perryman, L. Lindegren, J. Kovalevsky, E. Hoeg, U. Bastian, P. L. Bernacca, M. Crézé, F. Donati, M. Grenon, M. Grewing, F. van Leeuwen, H. van der Marel, F. Mignard, C. A. Murray, R. S. Le Poole, H. Schrijver, C. Turon, F. Arenou, M. Froeschlé, and C. S. Petersen. The HIPPARCOS Catalogue. *A&A*, 323:L49–L52, July 1997.
- P. C. Peters and J. Mathews. Gravitational Radiation from Point Masses in a Keplerian Orbit. *Physical Review*, 131:435–440, July 1963. doi: 10.1103/PhysRev.131.435.
- M. Pitkin and J. Romano. *Matlabmultinest*, 2013.
- V. Radhakrishnan and D. J. Cooke. Magnetic Poles and the Polarization Structure of Pulsar Radiation. *Astrophys. Lett.*, 3:225, 1969.
- S. M. Ransom, V. M. Kaspi, R. Ramachandran, P. Demorest, D. C. Backer, E. D. Pfahl, F. D. Ghigo, and D. L. Kaplan. Green Bank Telescope Measurement of the Systemic Velocity of the Double Pulsar Binary J0737-3039 and Implications for Its Formation. *ApJ*, 609:L71–L74, July 2004. doi: 10.1086/422911.
- D. J. Reardon, G. Hobbs, W. Coles, Y. Levin, M. J. Keith, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. Kerr, P. D. Lasky, R. N. Manchester, S. Osłowski, V. Ravi, R. M. Shannon, W. van Straten, L. Toomey, J. Wang, L. Wen, X. P. You, and X.-J. Zhu. Timing analysis for 20 millisecond pulsars in the Parkes Pulsar Timing Array. *MNRAS*, 455:1751–1769, January 2016. doi: 10.1093/mnras/stv2395.

BIBLIOGRAPHY

- B. Rickett, S. Johnston, T. Tomlinson, and J. Reynolds. The inner scale of the plasma turbulence towards PSR J1644-4559. *MNRAS*, 395:1391–1402, May 2009. doi: 10.1111/j.1365-2966.2009.14471.x.
- B. J. Rickett. Frequency Structure of Pulsar Intensity Variations. *Nature*, 221:158–159, January 1969. doi: 10.1038/221158a0.
- B. J. Rickett. Interstellar scattering and scintillation of radio waves. *ARA&A*, 15:479–504, 1977. doi: 10.1146/annurev.aa.15.090177.002403.
- B. J. Rickett. Radio propagation through the turbulent interstellar plasma. *ARA&A*, 28:561–605, 1990. doi: 10.1146/annurev.aa.28.090190.003021.
- B. J. Rickett, W. A. Coles, and G. Bourgois. Slow scintillation in the interstellar medium. *A&A*, 134:390–395, May 1984.
- B. J. Rickett, W. A. Coles, C. F. Nava, M. A. McLaughlin, S. M. Ransom, F. Camilo, R. D. Ferdman, P. C. C. Freire, M. Kramer, A. G. Lyne, and I. H. Stairs. Interstellar Scintillation of the Double Pulsar J0737-3039. *ApJ*, 787:161, June 2014. doi: 10.1088/0004-637X/787/2/161.
- R. W. Romani, R. Narayan, and R. Blandford. Refractive effects in pulsar scintillation. *MNRAS*, 220:19–49, May 1986. doi: 10.1093/mnras/220.1.19.
- R. W. Romani, R. D. Blandford, and J. M. Cordes. Radio caustics from localized interstellar medium plasma structures. *Nature*, 328:324–326, July 1987. doi: 10.1038/328324a0.
- V. H. Rumsey. Scintillations due to a concentrated layer with a power-law turbulence spectrum. *Radio Science*, 10:107–114, January 1975. doi: 10.1029/RS010i001p00107.
- J. S. Sandhu, M. Bailes, R. N. Manchester, J. Navarro, S. R. Kulkarni, and S. B. Anderson. The Proper Motion and Parallax of PSR J0437-4715. *ApJ*, 478:L95–L98, April 1997. doi: 10.1086/310562.
- M. V. Sazhin. Opportunities for detecting ultralong gravitational waves. *Soviet Ast.*, 22:36–38, February 1978.
- J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, December 1982. doi: 10.1086/160554.
- P. A. G. Scheuer. Amplitude Variations in Pulsed Radio Sources. *Nature*, 218:920–922, June 1968. doi: 10.1038/218920a0.
- M. Schmidt. 3C 273 : A Star-Like Object with Large Red-Shift. *Nature*, 197:1040, March 1963. doi: 10.1038/1971040a0.

- R. M. Shannon and J. M. Cordes. Assessing the Role of Spin Noise in the Precision Timing of Millisecond Pulsars. *ApJ*, 725:1607–1619, December 2010. doi: 10.1088/0004-637X/725/2/1607.
- R. M. Shannon, J. M. Cordes, T. S. Metcalfe, T. J. W. Lazio, I. Cognard, G. Desvignes, G. H. Janssen, A. Jessner, M. Kramer, K. Lazaridis, M. B. Purver, B. W. Stappers, and G. Theureau. An Asteroid Belt Interpretation for the Timing Variations of the Millisecond Pulsar B1937+21. *ApJ*, 766:5, March 2013a. doi: 10.1088/0004-637X/766/1/5.
- R. M. Shannon, V. Ravi, W. A. Coles, G. Hobbs, M. J. Keith, R. N. Manchester, J. S. B. Wyithe, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, J. Khoo, Y. Levin, S. Osłowski, J. M. Sarkissian, W. van Straten, J. P. W. Verbiest, and J.-B. Wang. Gravitational-wave limits from pulsar timing constrain supermassive black hole evolution. *Science*, 342:334–337, October 2013b.
- R. M. Shannon, S. Osłowski, S. Dai, M. Bailes, G. Hobbs, R. N. Manchester, W. van Straten, C. A. Raithel, V. Ravi, L. Toomey, N. D. R. Bhat, S. Burke-Spolaor, W. A. Coles, M. J. Keith, M. Kerr, Y. Levin, J. M. Sarkissian, J.-B. Wang, L. Wen, and X.-J. Zhu. Limitations in timing precision due to single-pulse shape variability in millisecond pulsars. *MNRAS*, 443:1463–1481, September 2014. doi: 10.1093/mnras/stu1213.
- R. M. Shannon, V. Ravi, L. T. Lentati, P. D. Lasky, G. Hobbs, M. Kerr, R. N. Manchester, W. A. Coles, Y. Levin, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. J. Keith, S. Osłowski, D. J. Reardon, W. van Straten, L. Toomey, J.-B. Wang, L. Wen, J. S. B. Wyithe, and X.-J. Zhu. Gravitational waves from binary supermassive black holes missing in pulsar observations. *Science*, 349:1522–1525, September 2015. doi: 10.1126/science.aab1910.
- R. M. Shannon, L. T. Lentati, M. Kerr, M. Bailes, N. D. R. Bhat, W. A. Coles, S. Dai, J. Dempsey, G. Hobbs, M. J. Keith, P. D. Lasky, Y. Levin, R. N. Manchester, S. Osłowski, V. Ravi, D. J. Reardon, P. A. Rosado, R. Spiewak, W. van Straten, L. Toomey, J.-B. Wang, L. Wen, X.-P. You, and X.-J. Zhu. The Disturbance of a Millisecond Pulsar Magnetosphere. *ApJ*, 828:L1, September 2016. doi: 10.3847/2041-8205/828/1/L1.
- I. I. Shapiro. Fourth Test of General Relativity. *Physical Review Letters*, 13:789–791, December 1964. doi: 10.1103/PhysRevLett.13.789.
- I. S. Shklovskii. Possible Causes of the Secular Increase in Pulsar Periods. *Soviet Ast.*, 13:562, February 1970.
- W. Sieber. Causal relationship between pulsar long-term intensity variations and the interstellar medium. *A&A*, 113:311–313, September 1982.
- S. R. Spangler. Plasma Turbulence in the Local Bubble. *Space Sci. Rev.*, 143:277–290, March 2009. doi: 10.1007/s11214-008-9391-7.

- S. R. Spangler and C. R. Gwinn. Evidence for an inner scale to the density turbulence in the interstellar medium. *ApJ*, 353:L29–L32, April 1990. doi: 10.1086/185700.
- E. M. Splaver, D. J. Nice, I. H. Stairs, A. N. Lommen, and D. C. Backer. Masses, Parallax, and Relativistic Timing of the PSR J1713+0747 Binary System. *ApJ*, 620:405–415, February 2005. doi: 10.1086/426804.
- D. H. Staelin and E. C. Reifstein, III. Pulsating Radio Sources near the Crab Nebula. *Science*, 162:1481–1483, December 1968. doi: 10.1126/science.162.3861.1481.
- D. R. Stinebring, M. A. McLaughlin, J. M. Cordes, K. M. Becker, J. E. E. Goodman, M. A. Kramer, J. L. Sheckard, and C. T. Smith. Faint Scattering Around Pulsars: Probing the Interstellar Medium on Solar System Size Scales. *ApJ*, 549:L97–L100, March 2001. doi: 10.1086/319133.
- D. R. Stinebring, A. S. Hill, and S. M. Ransom. Scintillation Arcs and Binary Pulsars with an Application to PSR J0737-3039. In F. A. Rasio and I. H. Stairs, editors, *Binary Radio Pulsars*, volume 328 of *Astronomical Society of the Pacific Conference Series*, page 349, July 2005.
- K. Stovall, P. C. C. Freire, S. Chatterjee, P. B. Demorest, D. R. Lorimer, M. A. McLaughlin, N. Pol, J. van Leeuwen, R. S. Wharton, B. Allen, M. Boyce, A. Brazier, K. Caballero, F. Camilo, R. Camuccio, J. M. Cordes, F. Crawford, J. S. Deneva, R. D. Ferdman, J. W. T. Hessels, F. A. Jenet, V. M. Kaspi, B. Knispel, P. Lazarus, R. Lynch, E. Parent, C. Patel, Z. Pleunis, S. M. Ransom, P. Scholz, A. Seymour, X. Siemens, I. H. Stairs, J. Swiggum, and W. W. Zhu. PALFA Discovery of a Highly Relativistic Double Neutron Star Binary. *ApJ*, 854:L22, February 2018. doi: 10.3847/2041-8213/aaad06.
- T. M. Tauris and T. Sennels. Formation of the binary pulsars PSR B2303+46 and PSR J1141-6545. Young neutron stars with old white dwarf companions. *A&A*, 355:236–244, March 2000.
- J. H. Taylor and J. M. Cordes. Pulsar distances and the galactic distribution of free electrons. *ApJ*, 411:674–684, July 1993. doi: 10.1086/172870.
- J. H. Taylor and J. M. Weisberg. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *ApJ*, 253:908–920, February 1982. doi: 10.1086/159690.
- S. E. Thorsett and D. Chakrabarty. Neutron Star Mass Measurements. I. Radio Pulsars. *ApJ*, 512:288–299, February 1999. doi: 10.1086/306742.
- M. Toscano, J. S. Sandhu, M. Bailes, R. N. Manchester, M. C. Britton, S. R. Kulkarni, S. B. Anderson, and B. W. Stappers. Millisecond pulsar velocities. *MNRAS*, 307:925–933, August 1999. doi: 10.1046/j.1365-8711.1999.02685.x.
- R. van Haasteren and Y. Levin. Understanding and analysing time-correlated stochastic signals in pulsar timing. *MNRAS*, 428:1147–1159, January 2013. doi: 10.1093/mnras/sts097.

- R. van Haasteren, Y. Levin, P. McDonald, and T. Lu. On measuring the gravitational-wave background using Pulsar Timing Arrays. *MNRAS*, 395:1005–1014, May 2009. doi: 10.1111/j.1365-2966.2009.14590.x.
- R. van Haasteren, Y. Levin, G. H. Janssen, K. Lazaridis, M. Kramer, B. W. Stappers, G. Desvignes, M. B. Purver, A. G. Lyne, R. D. Ferdman, A. Jessner, I. Cognard, G. Theureau, N. D’Amico, A. Possenti, M. Burgay, A. Corongiu, J. W. T. Hessels, R. Smits, and J. P. W. Verbiest. Placing limits on the stochastic gravitational-wave background using European Pulsar Timing Array data. *MNRAS*, 414:3117–3128, July 2011. doi: 10.1111/j.1365-2966.2011.18613.x.
- W. van Straten. Radio Astronomical Polarimetry and Point-Source Calibration. *ApJS*, 152: 129–135, May 2004. doi: 10.1086/383187.
- W. van Straten. High-fidelity Radio Astronomical Polarimetry Using a Millisecond Pulsar as a Polarized Reference Source. *ApJS*, 204:13, January 2013. doi: 10.1088/0067-0049/204/1/13.
- W. van Straten and M. Bailes. Annual-orbital Parallax and Nearby Binary Pulsars. In M. Bailes, D. J. Nice, and S. E. Thorsett, editors, *Radio Pulsars*, volume 302 of *Astronomical Society of the Pacific Conference Series*, page 65, 2003.
- W. van Straten, M. Bailes, M. Britton, S. R. Kulkarni, S. B. Anderson, R. N. Manchester, and J. Sarkissian. A test of general relativity from the three-dimensional orbital geometry of a binary pulsar. *Nature*, 412:158–160, July 2001.
- J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys. Rev. D*, 81(6):062003, March 2010. doi: 10.1103/PhysRevD.81.062003.
- J. P. W. Verbiest, M. Bailes, W. van Straten, G. B. Hobbs, R. T. Edwards, R. N. Manchester, N. D. R. Bhat, J. M. Sarkissian, B. A. Jacoby, and S. R. Kulkarni. Precision Timing of PSR J0437-4715: An Accurate Pulsar Distance, a High Pulsar Mass, and a Limit on the Variation of Newton’s Gravitational Constant. *ApJ*, 679:675-680, May 2008. doi: 10.1086/529576.
- J. P. W. Verbiest, M. Bailes, W. A. Coles, G. B. Hobbs, W. van Straten, D. J. Champion, F. A. Jenet, R. N. Manchester, N. D. R. Bhat, J. M. Sarkissian, D. Yardley, S. Burke-Spolaor, A. W. Hotan, and X. P. You. Timing stability of millisecond pulsars and prospects for gravitational-wave detection. *MNRAS*, 400:951–968, December 2009. doi: 10.1111/j.1365-2966.2009.15508.x.
- J. P. W. Verbiest, J. M. Weisberg, A. A. Chael, K. J. Lee, and D. R. Lorimer. On Pulsar Distance Measurements and Their Uncertainties. *ApJ*, 755:39, August 2012. doi: 10.1088/0004-637X/755/1/39.
- J. P. W. Verbiest, L. Lentati, G. Hobbs, R. van Haasteren, P. B. Demorest, G. H. Janssen, J.-B. Wang, G. Desvignes, R. N. Caballero, M. J. Keith, D. J. Champion, Z. Arzoumanian, S. Babak,

- C. G. Bassa, N. D. R. Bhat, A. Brazier, P. Brem, M. Burgay, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, I. Cognard, J. M. Cordes, S. Dai, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, J. R. Gair, N. E. Garver-Daniels, P. Gentile, M. E. Gonzalez, E. Graikou, L. Guillemot, J. W. T. Hessels, G. Jones, R. Karuppusamy, M. Kerr, M. Kramer, M. T. Lam, P. D. Lasky, A. Lassus, P. Lazarus, T. J. W. Lazio, K. J. Lee, L. Levin, K. Liu, R. S. Lynch, A. G. Lyne, J. Mckee, M. A. McLaughlin, S. T. McWilliams, D. R. Madison, R. N. Manchester, C. M. F. Mingarelli, D. J. Nice, S. Osłowski, N. T. Palliyaguru, T. T. Pennucci, B. B. P. Perera, D. Perrodin, A. Possenti, A. Petiteau, S. M. Ransom, D. Reardon, P. A. Rosado, S. A. Sanidas, A. Sesana, G. Shaifullah, R. M. Shannon, X. Siemens, J. Simon, R. Smits, R. Spiewak, I. H. Stairs, B. W. Stappers, D. R. Stinebring, K. Stovall, J. K. Swiggum, S. R. Taylor, G. Theureau, C. Tiburzi, L. Toomey, M. Vallisneri, W. van Straten, A. Vecchio, Y. Wang, L. Wen, X. P. You, W. W. Zhu, and X.-J. Zhu. The International Pulsar Timing Array: First data release. *MNRAS*, 458:1267–1288, May 2016. doi: 10.1093/mnras/stw347.
- M. Walker and M. Wardle. Extreme Scattering Events and Galactic Dark Matter. *ApJ*, 498: L125–L128, May 1998. doi: 10.1086/311332.
- M. A. Walker. Extreme Scattering Events: Insights into the Interstellar Medium on AU-Scales. In M. Haverkorn and W. M. Goss, editors, *SINS - Small Ionized and Neutral Structures in the Diffuse Interstellar Medium*, volume 365 of *Astronomical Society of the Pacific Conference Series*, page 299, July 2007.
- M. A. Walker, D. B. Melrose, D. R. Stinebring, and C. M. Zhang. Interpretation of parabolic arcs in pulsar secondary spectra. *MNRAS*, 354:43–54, October 2004. doi: 10.1111/j.1365-2966.2004.08159.x.
- M. A. Walker, A. V. Tuntsov, H. Bignall, C. Reynolds, K. W. Bannister, S. Johnston, J. Stevens, and V. Ravi. Extreme Radio-wave Scattering Associated with Hot Stars. *ApJ*, 843:15, July 2017. doi: 10.3847/1538-4357/aa705c.
- J. B. Wang, G. Hobbs, W. Coles, R. M. Shannon, X. J. Zhu, D. R. Madison, M. Kerr, V. Ravi, M. J. Keith, R. N. Manchester, Y. Levin, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, S. Osłowski, W. van Straten, L. Toomey, N. Wang, and L. Wen. Searching for gravitational wave memory bursts with the Parkes Pulsar Timing Array. *MNRAS*, 446:1657–1671, January 2015. doi: 10.1093/mnras/stu2137.
- J. M. Yao, R. N. Manchester, and N. Wang. A New Electron-density Model for Estimation of Pulsar and FRB Distances. *ApJ*, 835:29, January 2017. doi: 10.3847/1538-4357/835/1/29.
- X. P. You, G. Hobbs, W. A. Coles, R. N. Manchester, R. Edwards, M. Bailes, J. Sarkissian, J. P. W. Verbiest, W. van Straten, A. Hotan, S. Ord, F. Jenet, N. D. R. Bhat, and A. Teoh. Dispersion

- measure variations and their effect on precision pulsar timing. *MNRAS*, 378:493–506, June 2007. doi: 10.1111/j.1365-2966.2007.11617.x.
- M. Zechmeister and M. Kürster. The generalised Lomb-Scargle periodogram. A new formalism for the floating-mean and Keplerian periodograms. *A&A*, 496:577–584, March 2009. doi: 10.1051/0004-6361:200811296.
- W. W. Zhu, I. H. Stairs, P. B. Demorest, D. J. Nice, J. A. Ellis, S. M. Ransom, Z. Arzoumanian, K. Crowter, T. Dolch, R. D. Ferdman, E. Fonseca, M. E. Gonzalez, G. Jones, M. L. Jones, M. T. Lam, L. Levin, M. A. McLaughlin, T. Pennucci, K. Stovall, and J. Swiggum. Testing Theories of Gravitation Using 21-Year Timing of Pulsar Binary J1713+0747. *ApJ*, 809:41, August 2015. doi: 10.1088/0004-637X/809/1/41.
- X.-J. Zhu, G. Hobbs, L. Wen, W. A. Coles, J.-B. Wang, R. M. Shannon, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. J. Keith, M. Kerr, Y. Levin, D. R. Madison, S. Osłowski, V. Ravi, L. Toomey, and W. van Straten. An all-sky search for continuous gravitational waves in the Parkes Pulsar Timing Array data set. *MNRAS*, 444:3709–3720, November 2014. doi: 10.1093/mnras/stu1717.

Timing model parameter tables from Chapter 2

The pulsar parameters from Chapter 2 are listed in the tables below, and a discussion for each pulsar is given in Section 2.5. The model parameter files, ToA files, and red-noise models for each pulsar are available from the CSIRO Data Access Portal: <http://dx.doi.org/10.4225/08/561EFD72D0409>.

Table 1: Parameters for the solitary pulsars J0711–6830, J1024–0719, J1730–2304, J1744–1134, J1824–2452A, J1939+2134, and J2124–3358. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.

Pulsar name.....	J0711–6830	J1024–0719	J1730–2304	J1744–1134	J1824–2452A	J1939+2134	J2124–3358
MJD range.....	49373.6–55619.2	50117.5–55619.5	49421.9–55598.8	49729.1–55619.0	53518.8–55619.1	49956.5–55619.0	49489.9–55618.2
Data span (yr).....	17.1	15.06	16.91	16.13	5.75	15.5	16.78
Number of TOAs.....	566	493	390	534	313	397	652
Rms timing residual (μ s).....	2.0	10.4	1.9	0.5	5.5	5.8	2.9
Measured Quantities							
Right ascension (RA), α (hh:mm:ss)...	07:11:54.189114(13)	10:24:38.678633(7)	17:30:21.66624(8)	17:44:29.4057891(11)	18:24:32.00788(3)	19:39:38.561213(3)	21:24:43.849372(10)
Declination (DEC), δ (dd:mm:ss).....	–68:30:47.41446(8)	–07:19:19.36778(19)	–23:04:31.19(2)	–11:34:54.68126(8)	–24:52:10.834(8)	+21:34:59.12628(6)	–33:58:44.8500(3)
Pulse frequency, ν (s^{-1}).....	182.1172346685786(10)	193.71568347859(4)	123.1102871605879(5)	245.4261197130557(4)	327.40559048006(3)	641.92822645342(4)	202.793893782574(6)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-4.94405(11) \times 10^{-16}$	$-6.9523(18) \times 10^{-16}$	$-3.05917(4) \times 10^{-16}$	$-5.38173(3) \times 10^{-16}$	$-1.735302(7) \times 10^{-13}$	$-4.33106(3) \times 10^{-14}$	$-8.4597(3) \times 10^{-16}$
Proper motion in RA, $\mu_\alpha \cos \delta$ (mas yr^{-1})	–15.57(3)	–35.33(4)	20.264(19)	18.790(6)	–0.69(13)	0.087(16)	–14.14(4)
Proper motion in DEC, μ_δ (mas yr^{-1})..	14.24(3)	–48.32(8)	–	–9.40(3)	–	–0.41(3)	–50.08(9)
Parallax, π (mas).....	–	0.5(3)	1.5(3)	2.53(7)	–	0.52(16)	2.4(4)
Set Quantities							
Dispersion measure, DM (cm^{-3}pc)....	18.4099	6.48803	9.61634	3.13695	119.892	71.0227	4.60096
Derived Quantities							
$\log_{10}(\text{Characteristic age, yr})$	9.77	9.64	9.8	9.86	7.48	8.37	9.58
$\log_{10}(\text{Surface magnetic field strength, G})$	8.46	8.5	8.61	8.29	9.35	8.61	8.51
$\log_{10}(\text{Edot, ergs/s})$	33.55	33.73	33.17	33.72	36.35	36.04	33.83

Table 2: Parameters for the binary pulsars described by the small-eccentricity ELL1 binary model, PSRs J0613–0200, J1045–4509, J1603–7202, J1732–5049, J1857+0943, J2129–5721, and J2145–0750. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.

Pulsar name	J0613–0200	J1045–4509	J1603–7202	J1732–5049	J1857+0943	J2129–5721	J2145–0750
MJD range	51526.6–55619.3	49405.5–55619.5	50026.1–55618.8	52647.1–55582.2	53086.9–55619.0	49987.4–55618.3	49517.8–55618.2
Data span (yr)	11.21	17.01	15.31	8.04	6.93	15.42	16.7
Number of TOAs	639	646	493	244	291	448	972
Rms timing residual (μ s)	1.0	11.9	2.7	3.1	1.1	1.4	1.7
Measured Quantities							
Right ascension (RA), α (hh:mm:ss) ...	06:13:43.975503(4)	10:45:50.18696(3)	16:03:35.67751(4)	17:32:47.766731(19)	18:57:36.390848(4)	21:29:22.766966(12)	21:45:50.46148(3)
Declination (DEC), δ (dd:mm:ss)	–02:00:47.21147(14)	–45:09:54.1223(4)	–72:02:32.72985(19)	–50:49:00.1917(4)	+09:43:17.21458(9)	–57:21:14.21183(12)	–07:50:18.4759(12)
Pulse frequency, ν (s^{-1})	326.6005620676858(16)	133.793149554823(14)	67.376581131844(3)	188.233512213289(4)	186.4940784047797(7)	268.3592273587338(16)	62.2958878423832(13)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-1.02293(3) \times 10^{-15}$	$-3.1616(8) \times 10^{-16}$	$-7.0956(12) \times 10^{-17}$	$-5.0296(6) \times 10^{-16}$	$-6.20417(14) \times 10^{-16}$	$-1.501784(14) \times 10^{-15}$	$-1.15599(7) \times 10^{-16}$
Proper motion in RA, $\mu_\alpha \cos \delta$ (mas yr^{-1})	1.811(16)	–6.07(9)	–2.46(4)	–0.41(9)	–2.69(3)	9.25(4)	–9.59(8)
Proper motion in DEC, μ_δ (mas yr^{-1}) ..	–10.36(4)	5.20(10)	–7.33(5)	–9.87(19)	–5.48(6)	–9.58(4)	–8.9(3)
Parallax, π (mas)	0.86(13)	2.2(11)	1.1(8)	–	0.5(3)	–	1.84(17)
Orbital period, P_b (d)	1.198512575218(18)	4.0835292548(3)	6.3086296691(5)	5.2629972182(5)	12.3271713817(5)	6.6254930923(13)	6.83890261536(5)
Projected semi-major axis, x (lt-s)	1.09144422(6)	3.0151313(3)	6.8806577(6)	3.9828703(4)	9.2307805(4)	3.50056678(14)	10.1641061(3)
Epoch of ascending node, T_{ASC} (MJD) .	50315.26949108(6)	50273.507005(3)	50426.28702402(13)	51396.3661225(3)	47520.4323457(3)	50442.6431238(4)	50802.29822944(3)
EPS1, $e \sin \omega_0$	$3.90(11) \times 10^{-6}$	$-2.096(17) \times 10^{-5}$	$1.60(5) \times 10^{-6}$	$2.08(16) \times 10^{-6}$	$-2.160(5) \times 10^{-5}$	$-3.58(8) \times 10^{-6}$	$-6.840(13) \times 10^{-6}$
EPS2, $e \cos \omega_0$	$3.40(11) \times 10^{-6}$	$-1.099(16) \times 10^{-5}$	$-9.20(4) \times 10^{-6}$	$-8.24(16) \times 10^{-6}$	$2.46(3) \times 10^{-6}$	$-1.165(8) \times 10^{-5}$	$-1.8059(14) \times 10^{-5}$
First derivative of orbital period, \dot{P}_b ...	–	–	$3.1(15) \times 10^{-13}$	–	–	$7.9(36) \times 10^{-13}$	–
First derivative of x , \dot{x}	–	–	$1.36(16) \times 10^{-14}$	–	–	–	$8.0(8) \times 10^{-15}$
Companion mass, M_c (M_\odot)	–	–	–	–	0.25(3)	–	–
Sine of inclination angle, $\sin i$	–	–	–	–	0.9988(8)	–	–
Set Quantities							
Dispersion measure, DM (cm^{-3}pc)	38.7756	58.1438	38.0489	56.8365	13.2984	31.8509	8.99761
Derived Quantities							
\log_{10} (Characteristic age, yr)	9.7	9.83	10.18	9.77	9.68	9.45	9.93
\log_{10} (Surface magnetic field strength, G)	8.24	8.57	8.69	8.44	8.5	8.45	8.84
\log_{10} (Edot, ergs/s)	34.12	33.22	32.28	33.57	33.66	34.2	32.45
Epoch of periastron, T_0 (MJD)	50315.432(4)	50276.256(5)	50429.268(5)	51398.790(16)	47529.900(3)	50446.270(7)	50806.1118(8)
Orbital eccentricity, e	$5.18(11) \times 10^{-6}$	$2.367(17) \times 10^{-5}$	$9.35(5) \times 10^{-6}$	$8.50(16) \times 10^{-6}$	$2.174(5) \times 10^{-5}$	$1.219(8) \times 10^{-5}$	$1.9311(14) \times 10^{-5}$
Longitude of periastron, ω_0 (deg)	48.9(12)	242.3(4)	170.1(3)	165.8(11)	276.49(7)	197.1(4)	200.75(4)
Pulsar mass, M_P (M_\odot)	–	–	–	–	1.5(0.2)	–	–
i limit from \dot{x} measurement (degrees) ..	–	–	≤ 31	–	–	–	≤ 69

Table 3: Parameters for the binary pulsars described by DD binary model, PSRs J1022+1001, J1600–3053, and J1643–1224. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis.

Pulsar name	J1022+1001	J1600–3053	J1643–1224
MJD range	52649.7–55618.6	52302.0–55618.8	49421.8–55618.9
Data span (yr)	8.13	9.08	16.97
Number of TOAs	615	715	488
Rms timing residual (μ s)	1.8	0.8	2.8
Measured Quantities			
Right ascension (RA), α (hh:mm:ss)...	10:22:58.0007(13)	16:00:51.903452(7)	16:43:38.160985(9)
Declination (DEC), δ (dd:mm:ss)	+10:01:52.77(5)	–30:53:49.3653(3)	–12:24:58.6783(6)
Pulse frequency, ν (s^{-1})	60.7794479636137(3)	277.9377070213120(12)	216.373337179973(9)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-1.60095(6) \times 10^{-16}$	$-7.3385(4) \times 10^{-16}$	$-8.6433(5) \times 10^{-16}$
Proper motion in RA, $\mu_\alpha \cos \delta$ (mas yr^{-1})	–17.09(3)	–0.99(4)	5.94(5)
Proper motion in DEC, μ_δ (mas yr^{-1}) ..	–	–7.22(15)	3.94(18)
Parallax, π (mas)	1.1(3)	0.48(11)	1.27(19)
Orbital period, P_b (d)	7.8051360(16)	14.3484577721(3)	147.01728(7)
Projected semi-major axis, x (lt-s)	16.765395(14)	8.8016536(13)	25.0726150(7)
Epoch of periastron, T_0 (MJD)	49778.4080(11)	53295.5390(7)	49577.972(3)
Orbital eccentricity, e	$9.683(17) \times 10^{-5}$	$1.73729(10) \times 10^{-4}$	$5.05753(9) \times 10^{-4}$
Longitude of periastron, ω_0 (deg)	97.64(5)	181.832(17)	321.857(6)
First derivative of orbital period, \dot{P}_b ...	$5.5(23) \times 10^{-13}$	–	–
First derivative of x , \dot{x}	$1.15(16) \times 10^{-14}$	$-4.2(7) \times 10^{-15}$	$-5.25(16) \times 10^{-14}$
Periastron advance, $\dot{\omega}$ (deg/yr)	0.012(4)	–	–0.0007(4)
Companion mass, M_c (M_\odot)	2.2(2.4)	0.34(15)	–
Sine of inclination angle, $\sin i$	0.69(18)	0.87(6)	–
Set Quantities			
Dispersion measure, DM (cm^{-3}pc)	10.2531	52.3249	62.4143
Derived Quantities			
\log_{10} (Characteristic age, yr)	9.78	9.78	9.6
\log_{10} (Surface magnetic field strength, G)	8.93	8.27	8.47
\log_{10} (Edot, ergs/s)	32.58	33.91	33.87
Pulsar mass, M_P (M_\odot)	–	2.4(1.7)	–
i limit from \dot{x} measurement (degrees) ..	≤ 84	≤ 67	≤ 28

Table 4: Parameters for the binary pulsars described by T2 binary model, PSRs J0437–4715, J1713+0747, and J1909–3744. Numbers in brackets are the TEMPO2 1-sigma uncertainties on the last quoted decimal place, including split-Cholesky analysis. In each case where a companion mass is measured from the Shapiro delay, the corresponding $\sin i$ parameter is linked to the Kopeikin parameter, i .

Pulsar name.....	J0437–4715	J1713+0747	J1909–3744
MJD range.....	50191.0—55619.2	49421.9—55618.9	52618.4—55619.1
Data span (yr).....	14.86	16.97	8.22
Number of TOAs.....	5065	622	1368
Rms timing residual (μ s)	0.3	0.4	0.2
Measured Quantities			
Right ascension (RA), α (hh:mm:ss)...	04:37:15.8961737(6)	17:13:49.5327220(19)	19:09:47.4346749(11)
Declination (DEC), δ (dd:mm:ss)....	–47:15:09.110714(7)	+07:47:37.49795(6)	–37:44:14.46674(5)
Pulse frequency, ν (s^{-1})	173.6879458121843(5)	218.8118404348011(11)	339.3156872882446(3)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-1.728361(5) \times 10^{-15}$	$-4.08380(6) \times 10^{-16}$	$-1.614817(5) \times 10^{-15}$
Proper motion in RA, $\mu_\alpha \cos \delta$ (mas yr^{-1})	121.4385(20)	4.912(7)	–9.517(5)
Proper motion in DEC, μ_δ (mas yr^{-1})..	–71.4754(20)	–3.888(14)	–35.797(17)
Parallax, π (mas).....	6.37(9)	0.86(9)	0.81(3)
Orbital period, P_b (d)	5.7410459(4)	67.825130978(4)	1.533449474406(13)
Projected semi-major axis, x (lt-s)....	3.36671444(5)	32.3424210(5)	1.89799118(4)
Epoch of periastron, T_0 (MJD).....	54501.4671(3)	51997.5804(9)	53631.39(4)
Orbital eccentricity, e	$1.91811(15) \times 10^{-5}$	$7.49373(17) \times 10^{-5}$	$1.14(10) \times 10^{-7}$
Longitude of periastron, ω_0 (deg)....	1.363(17)	176.201(5)	156(8)
First derivative of orbital period, \dot{P}_b ...	$3.728(6) \times 10^{-12}$	$1.7(7) \times 10^{-12}$	$5.03(6) \times 10^{-13}$
Periastron advance, $\dot{\omega}$ (deg/yr).....	0.0138(13)	–	–
Companion mass, M_c (M_\odot)	0.224(7)	0.34(5)	0.2067(19)
Inclination angle, i (degrees)	137.56(4)	69(3)	93.52(9)
Longitude of ascending node, Ω (degrees)	207.0(12)	99(4)	39(10)
Set Quantities			
Dispersion measure, DM (cm^{-3}pc)....	2.64498	15.9903	10.3932
Derived Quantities			
\log_{10} (Characteristic age, yr)	9.2	9.93	9.52
\log_{10} (Surface magnetic field strength, G)	8.76	8.3	8.31
\log_{10} (Edot, ergs/s)	34.07	33.55	34.34
Pulsar mass, M_P (M_\odot).....	1.44(7)	1.7(4)	1.47(3)

Reproducing our results from Chapter 3

The raw data from the Parkes radio telescope that were used for this work are available from the CSIRO data access portal (DAP; <https://data.csiro.au>), and were processed with a pipeline developed for the second PPTA data release. This pipeline is used to produce TOAs and dynamic spectra, and is briefly summarised in Section 3.2. The dynamic spectra files that were produced from this pipeline, and the collection of MATLAB codes used to analyse these dynamic spectra for the results presented in this paper, are available from the DAP at <https://doi.org/10.4225/08/5ae7b7b1b65c8>.

The code is presented for reproducibility and in general is not intended to be used for other applications, however a brief description of each script and function is included in README.TXT, and some may be useful for other scintillation studies. For example, we have included a function (GETDYNspecsPARAMS.M) that may be used to measure scintillation parameters for any dynamic spectrum given as a two-dimensional matrix of intensity versus observing time and frequency, using the methods described in Section 3.2.2. We have also included the scintillation velocity model (VISSMODEL.M), which is described in Section 3.3, and a script that numerically calculates the true anomaly from MJD and then uses nonlinear regression to fit the model to V_{ISS} derived from scintillation parameters (MODELdYNspecsPARAMS.M).

The MATLAB code was executed in version 2017b, and requires JD2DATE.M from the ASTRO-MATLAB library (<https://webhome.weizmann.ac.il/home/eofek/matlab/>; Ofek 2014), and the “Image Processing”, “Statistics and Machine Learning”, “Signal Processing”, “Optimization”, “Curve Fitting”, “Symbolic Math”, and “Parallel Computing” MATLAB toolboxes.